Veeam Backup & Replication

Version 7.0

User Guide
VMware Environments
September, 2013

CONTENTS

CONTACTING VEEAM SOFTWARE

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

Customer Support

Should you have a technical concern, suggestion or question, please visit our Customer Center Portal at cp.veeam.com to open a case, search our knowledge base, reference documentation, manage your license or obtain the latest product release.

Company Contacts

For the most up to date information about company contacts and offices location, please visit www.veeam.com/contacts.html.

Online Support

If you have any questions about Veeam Backup & Replication, you can use the following resources:

- Full documentation set: www.veeam.com/vmware-esx-backup/resources.html
- Community forum at www.veeam.com/forums

ABOUT THIS GUIDE

This user guide provides information about main features, installation and use of Veeam Backup & Replication in VMware environments. The document applies to version 7.0 and all subsequent versions until it is replaced with a new edition.

Intended Audience

The user guide is intended for anyone who wants to use Veeam Backup & Replication. It is primarily aimed at VMware administrators, consultants, analysts and any other IT professionals using the product.

Related Documentation

The complete set of Veeam Backup & Replication documentation can be found on the product resources web page at www.veeam.com/vmware-esx-backup/resources.html.

Document Revision History

Revision #	Date	Description of Changes
Revision 1	8/15/2013	Initial version of the document for Veeam Backup & Replication 7.0
Revision 2	8/19/2013	List of FS supported by the Multi-OS Restore wizard updated
Revision 3	9/23/2013	Minor changes in the Backup, Backup Copy, Tape Support and vCloud Director Support sections.

INTRODUCTION

Veeam Backup & Replication is a data protection and disaster recovery solution for VMware vSphere and Microsoft Hyper-V virtual environments of any size and complexity. Combining all of the necessary functions in one intuitive interface, Veeam Backup & Replication serves to solve the most critical problems of virtualized infrastructure management and protects mission-critical virtual machines (VMs) from both hardware and software failures.

Veeam Backup & Replication for VMware vSphere provides the following features and functionality:

VMware vSphere Support

Veeam Backup & Replication provides full support for VMware vSphere and VMware Infrastructure 3 (VI), including the newest version VMware vSphere 5.1. Veeam Backup & Replication capabilities are supported for both ESX and ESXi hosts in the same manner, providing top performance and reliability no matter of the hypervisor architecture.

Distributed Backup Architecture

Veeam Backup & Replication features a distributed backup architecture to take backup workload off the backup server and speed up backup, replication and restore over wide area networks (WANs) or slow links. Distributed architecture lets you move data processing to backup proxies that act as data movers, and backup repositories used as common backup locations. A single backup server acts as a 'point of control' managing multiple backup proxies and backup repositories across a large backup site or multiple branch offices. Installation and configuration of architecture components is fully automated, which helps streamline deployment and maintenance of remote sites and large installations.

2-in-1 Backup and Replication

To provide the most comprehensive protection of your virtual infrastructure,
Veeam Backup & Replication complements image-based backup with image−based replication. You can back up any VM, VM container or VM disk, as well as replicate VMs onsite for high availability (HA) or offsite for disaster recovery (DR), across local area and wide area networks.

By leveraging the latest technological advancements of the virtualization technology,
Veeam Backup & Replication delivers unprecedented replication speed. It provides near−continuous data protection (or near−CDP) at a fraction of cost of traditional CDP systems — you can capture changes and update VM images as often as every few minutes.

Veeam vPower

Veeam Backup & Replication offers vPower™ — Veeam's technology that allows you to:

- Immediately recover a failed VM, thus reducing downtime of production VMs to the minimum (see Instant VM Recovery).
- Verify recoverability of every backup and replica by starting and testing VMs directly from VM backups and replicas in an isolated environment (see SureBackup™ Recovery Verification).
- Restore items from any virtualized applications with U-AIR™ (see Universal Application-Item Recovery™).

Advanced Replication Options

Veeam Backup & Replication offers advanced replication options to cope with any disaster recovery scenario. If a VM goes down for some reason, you can fail over to its replica in several seconds. Once the production host is up again, you can use the failback option to get back to the original VM, or to a VM restored from the backup, in the original site or at a new location. During failback, Veeam Backup & Replication compares source and target replica states and transfers only differences, which dramatically decreases use of WAN traffic.

In addition to failover and failback capabilities, Veeam Backup & Replication provides enhanced operations with replicas — you can use replica seeding to minimize use of WAN traffic, map replicas to existing VMs in the DR site, and perform replica re−IP and network mapping in case the network settings of your production site does not match those of the DR site.

Native Tape Support

Veeam Backup & Replication allows you to archive backups of your VMs and Windows or Linux files to the tape media, and also to recover from tape records. Specially designed jobs can be run either manually or automatically on schedule, supporting customizable retention settings, hardware compression, post-backup export of media, and other options. You can specify tape media as a secondary target destination when configuring your standard backup jobs, fully automating the workflow and following '3-2-1' backup approach (3 copies, 2 types of media, 1 off-site location) considered as best practice for data protection and disaster recovery. Standalone drives, tape libraries and virtual tape libraries (VTL) are supported. Tracking of restore points facilitates restore of archived VMs.

Backup Copy and WAN Acceleration

To facilitate adoption of the 3-2-1 backup strategy, Veeam Backup & Replication offers backup copying capabilities. Backup copying allows you to create several instances of the same backup file in different locations, whether onsite or offsite. Copied backup files have the same format as those created by backup jobs and you can use any data recovery option for them. And with built-in WAN acceleration developed specifically for backup copy jobs you can copy data to offsite locations up to 50 times faster.

vCloud Director Support

With enhanced support for vCloud Director, you can protect your vApps and their contents. When Veeam Backup & Replication creates backup files for VMs, it captures vApp metadata along with the VM data. As a result, you can quickly restore vCloud Director VMs back to the vCloud Director hierarchy and do not need to perform any additional actions on import and VM configuration.

Multiple Data Recovery Options

Veeam Backup & Replication uses the same image−level backup to recover a full VM image, VM files, VM virtual drives, application items (U-AIR) and individual guest OS files. You can recover a full VM to any good-to-know point in time and place it to its original location with one click, or select a new location for a restored VM. You can also perform granular VM guest OS file- or folder-level recovery for FAT16, FAT32, NTFS and ReFS file systems without extracting a full VM image to the local drive. File-level restore for VMs running other file systems can be performed with the Multi-OS File-Level Restore wizard, or through utilizing Instant VM Recovery capabilities. Along with guest OS files recovery, Veeam Backup & Replication allows restoring specific VM files (VMDK, VMX and so on) and VM virtual drives, either to their latest state or to a specific point in time.

Forward Incremental and Reversed Incremental Backup

Depending on the type of backup storage you use, you can choose between two backup methods — incremental and reversed incremental. Incremental backup is recommended for disk–to–disk–to–tape and remote site backups — it reduces the time spent to move backups to tape or a remote site, and the amount of tape required. Reversed incremental backup is recommended for disk–to–disk backup, enabling you to keep the latest image of a VM in a ready–to–restore state on disk. With advanced options in Veeam Backup & Replication, you can select to perform incremental backup and schedule creation of synthetic full backups on specific days, which will let you combine advantages of incremental backup with those of reversed incremental.

Application-Aware Image Processing

Veeam Backup & Replication supports Windows Volume Shadow Copy (VSS) enabling backup and replication on live systems running Windows applications or working with databases (for example, Domain Controller, Exchange Server, SQL Server) without shutting them down. It also provides advanced options to control truncating of transaction logs so that you can ensure correct backup of applications that use transaction logs, and meet requirements of any backup scenario. You can select to truncate transaction logs after every backup job, every successful backup or not to truncate logs at all.

Indexing, Search and 1-Click Restore of VMs and Files

Veeam Backup & Replication provides possibilities for indexing guest OS files in Windows-based VMs, enabling you to perform quick and accurate search for files within backed up VM images without the need to restore them first. Using Veeam Backup Enterprise Manager, you can browse and search for files in a single selected VM backup or use the advanced search option to find necessary files in all VM backups within your backup infrastructure. You can use flexible delegation options to allow authorized users restore VMs or individual files with one click directly from Veeam Backup Enterprise Manager to the original location on the source VM disks, or download them to the local machine.

Data Deduplication and Compression

In order to decrease the size of created backups, Veeam Backup & Replication deduplicates identical blocks inside a backup file. Higher deduplication rates are achieved when backing up multiple VMs created from a single template or VMs with lots of free space. You can also decrease the backup file size by using compression.

Support for Storage Snapshots

Veeam Backup & Replication lets you leverage storage snapshots as a part of a comprehensive backup and recovery strategy, where SAN snapshots and image-level backups complement each other. The Backup from Storage Snapshots capability dramatically improves RPOs and reduce impact of backup activities on the production environment. And Veeam Explorer for SAN Snapshots allows you to recover VMs and VM data directly from HP LeftHand, HP StoreVirtual VSA and HP StoreServ snapshots. You can quickly restore an entire VM, VM guest OS files or Microsoft Exchange and Microsoft SharePoint items directly from SAN snapshots that can be taken throughout the day with very little impact on production systems.

Veeam Explorer for Exchange

Veeam Explorer for Exchange is a built-in tool that you can use to browse Microsoft Exchange mailbox stores inside Veeam backups. Veeam Explorer for Exchange features a familiar, easy-to-use interface and allows you to quickly locate mailboxes or items you need and restore them to the original location or to a new location.

Veeam Explorer for SharePoint

Veeam Explorer for SharePoint is a built-in tool that you can use to browse Microsoft SharePoint content databases inside Veeam backups. Veeam Explorer for SharePoint allows you to quickly locate documents, items, document libraries and lists and restore them to the original or new SharePoint site, send them by e-mail or save to the specified location.

Veeam Backup Enterprise Manager

Veeam Backup & Replication comes with Veeam Backup Enterprise Manager — a management and reporting component that allows you to manage multiple Veeam Backup & Replication installations from a single web console. In case of a distributed backup infrastructure, Veeam Backup Enterprise Manager acts as a single management point, enabling you to perform, edit and clone backup and replication jobs, and providing enhanced notification and reporting options.

Veeam Backup Enterprise Manager is also responsible for replicating and consolidating index files from backup servers to enable you to browse and search for files, and restore found files in one click. In addition to that, Veeam Backup Enterprise Manager acts as a license center, allowing you to centrally update licenses and get statistics on their usage.

vSphere Web Client Plug-in

The new plug-in for VMware's next-generation administrative interface for vSphere delivers at-a-glance and detailed views of your backup infrastructure directly from the vSphere web client. You can monitor job status and backup resources, as well as open Veeam Backup Management Suite reports on repository capacity, job statistics and identification of unprotected VMs.

Integration with Veeam ONE

Veeam Backup & Replication integrates with Veeam ONE allowing you to collect real-time statistics from your backup servers. You can use Veeam ONE to track the latest status of data protection operations in your virtual environment, receive immediate alarms whenever a potential problem can cause data loss, monitor performance of backup infrastructure components to optimize workloads and plan capacity of backup infrastructure resources.

OVERVIEW

Veeam Backup & Replication provides a set of features for building and maintaining a flexible backup infrastructure, performing data protection tasks (such as, regular backup and replication of VMs), and carrying out disaster recovery procedures. This section contains a high-level overview of Veeam Backup & Replication, its architecture and features, as well as data protection and disaster recovery concepts necessary to understand Veeam Backup & Replication background operations and processes.

Solution Architecture

Veeam Backup & Replication is a modular solution which allows flexible scalability for environments of different sizes and configuration. The installation package of Veeam Backup & Replication includes a set of components used to configure your backup infrastructure. Some of them are mandatory and provide core functionality; some of them are optional and can be installed to provide additional functionality for your business and deployment needs. You can consolidate Veeam Backup & Replication components on the same machine, either physical or virtual, or you can set them up separately for a more scalable approach.

Components

Veeam Backup & Replication comprises the following components. Some of the components are installed using a setup file; others are configured while working with the product.

Veeam Backup Server

The Veeam backup server is a Windows-based physical or virtual machine on which Veeam Backup & Replication is installed. It is the core component in the backup infrastructure that fills the role of the "configuration and control center". The Veeam backup server performs all types of administrative activities:

- Coordinates backup, replication, recovery verification and restore tasks
- Controls job scheduling and resource allocation
- Is used to set up and manage backup infrastructure components as well as specify global settings for the backup infrastructure

In addition to its primary functions, a newly deployed Veeam backup server also performs the roles of the default backup proxy and the backup repository (it manages data handling and data storing tasks).

The Veeam backup server uses the following services and components:

- **Veeam Backup Service** is a Windows service that coordinates all operations performed by Veeam Backup & Replication such as backup, replication, recovery verification and restore tasks. Veeam Backup Service runs under the administrator account with the *Log on as* right granted.
- **Veeam Backup Shell** provides the application user interface and allows user access to the application's functionality.
- **Veeam Backup Catalog Service** is a Windows service that manages guest OS file system indexing for VMs and replicates system index data files to enable search through guest OS files. Index data is stored in the Veeam Backup Catalog — a folder on the Veeam backup server. The Veeam Backup Catalog Service running on the Veeam backup server works in conjunction with search components installed on Veeam Backup Enterprise Manager and (optionally) a dedicated Microsoft Search Server.

- **Veeam Backup SQL Database** is used by Veeam Backup Service, Veeam Backup Shell and Veeam Backup Catalog Service to store data about the backup infrastructure, jobs, sessions and so on. The database instance can be located on a SQL Server installed either locally (on the same machine where the Veeam backup server is running) or remotely.
- **Veeam Backup PowerShell Snap-In** is an extension for Microsoft Windows PowerShell 2.0. Veeam Backup PowerShell adds a set of cmdlets to allow users to perform backup, replication and recovery tasks through the command-line interface of PowerShell or run custom scripts to fully automate operation of Veeam Backup & Replication.
- **Backup Proxy Services**. In addition to dedicated services, the Veeam backup server runs a set of data mover services (for details, see Backup Proxy).

Backup Proxy

When Veeam Backup & Replication is initially installed, the Veeam backup server coordinates all job activities and handles data traffic itself. That is, when you run a backup, replication, VM copy, VM migration job or perform restore operations, VM data is moved from source to target through the Veeam backup server. This scenario is acceptable for virtual environments where few backup jobs are performed; in large-scale environments, however, the workload on the Veeam backup server will be significant.

To take the workload off the Veeam backup server, Veeam Backup & Replication uses backup proxies. A backup proxy is an architecture component that sits between data source and target and is used to process jobs and deliver backup traffic. In particular, the backup proxy tasks include retrieving VM data from the production storage, compressing and sending it to the backup repository (for example, if you run a backup job) or another backup proxy (for example, if you run a replication job). As the data handling task is assigned to the backup proxy, the Veeam backup server becomes the "point of control" for dispatching jobs to proxy servers.

The role of a backup proxy can be assigned to a dedicated Windows server (physical or virtual) in your virtual environment. You can deploy backup proxies both in the primary site and in remote sites. To optimize performance of several concurrent jobs, you can use a number of backup proxies. In this case, Veeam Backup & Replication will distribute the backup workload between available backup proxies.

Use of backup proxies lets you easily scale your backup infrastructure up and down based on your demands. Backup proxies run light-weight services that take a few seconds to deploy. Deployment is fully automated — Veeam Backup & Replication installs the necessary components on a Windows-based server when you add it to the product console. As soon as you assign the role of a backup proxy to the added server, Veeam Backup & Replication starts the required services on it.

The primary role of the backup proxy is to provide an optimal route for backup traffic and enable efficient data transfer. Therefore, when deploying a backup proxy, you need to analyze the connection between the backup proxy and storage with which it is working. Depending on the type of connection, the backup proxy can be configured in one of the following ways (starting from the most efficient):

- A machine used as a backup proxy should have direct access to the storage on which VMs reside or the storage where VM data is written. This way, the backup proxy will retrieve data directly from the datastore, bypassing LAN.
- The backup proxy can be a VM with HotAdd access to VM disks on the datastore. This type of proxy also enables LAN-free data transfer.
- If neither of the above scenarios is possible, you can assign the role of the backup proxy to a machine on the network closer to the source or the target storage with which the proxy will be working. In this case, VM data will be transported over LAN using NBD protocol.

Depending on the type of backup proxy and your backup architecture, the backup proxy can use one of the following data transport modes: *Direct SAN Access*, *Virtual Appliance* or *Network*. If the VM disks are located on the SAN storage and the SAN storage is added to the Veeam Backup & Replication console, the backup proxy can also use the *Backup from Storage Snapshots* mode. You can explicitly select the transport mode or let Veeam Backup & Replication replication automatically choose the mode. For details, see Transport Modes and Backup from Storage Snapshots.

The backup proxy uses the following services and components:

- **Veeam Installer Service** is an auxiliary service that is installed and started on any Windows server once it is added to the list of managed servers in the Veeam Backup & Replication console. This service analyses the system, installs and upgrades necessary components and services depending on the role selected for the server.
- **Veeam Transport** is responsible for deploying and coordinating executable modules that act as "data movers" and perform main job activities on behalf of Veeam Backup & Replication, such as communicating with VMware Tools, copying VM files, performing data deduplication and compression and so on.

Backup Repository

A backup repository is a location used by Veeam Backup & Replication jobs to store backup files, copies of VMs and metadata for replicated VMs. Technically, a backup repository is a folder on the backup storage. By assigning different repositories to jobs and limiting the number of parallel jobs for each one, you can balance the load across your backup infrastructure.

In the Veeam backup infrastructure, you can use one of the following repository types:

- **Windows server with local or directly attached storage.** The storage can be a local disk, directly attached disk−based storage (such as a USB hard drive), or iSCSI/FC SAN LUN in case the server is connected into the SAN fabric.
  On a Windows repository, Veeam Backup & Replication deploys Veeam transport service (when you add a Windows−based server to the product console, Veeam Backup & Replication installs a set of components including the Veeam transport service on that server). When any job addresses the repository, the transport service on the repository establishes a connection with the source-side transport service on the backup proxy, enabling efficient data transfer over LAN or WAN.
  Windows repositories can be configured to function as **vPower NFS Servers**. In this case, Veeam Backup & Replication will run the Veeam vPower NFS Service directly on the backup repository (namely, on the managing Windows server to which storage is attached) and provide ESX(i) hosts with transparent access to backed up VM images stored on the repository. For details, see Veeam vPower NFS Service.
- **Linux server with local, directly attached storage or mounted NFS.** The storage can be a local disk, directly attached disk-based storage (such as a USB hard drive), NFS share, or iSCSI/FC SAN LUN in case the server is connected into the SAN fabric.
  On the Linux repository, Veeam Backup & Replication deploys and starts the Veeam transport service when a job addressing this repository is launched. This transport service establishes a connection with the source−side transport service on the backup proxy, enabling efficient data transfer over LAN or WAN.
- **CIFS (SMB) share.** SMB share does not support Veeam transport services, therefore data to the SMB share is written from a Windows−based proxy server. By default, this role performs a backup proxy that is utilized by the job for data transport.
  However, if you plan to move VM data to an offsite SMB repository over a WAN link, it is recommended that you deploy an additional Windows proxy server in the remote site, closer to the SMB repository. Veeam Backup & Replication will deploy a Veeam transport service on that proxy server, which will improve data transfer performance.

Veeam Backup Enterprise Manager

Veeam Backup Enterprise Manager is an optional component intended for distributed enterprise environments with multiple backup servers. Veeam Backup Enterprise Manager federates Veeam backup servers and offers a consolidated view of these servers through a web browser interface. You can centrally control and manage all jobs through a single "pane of glass", edit and clone jobs, monitor job 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 one click.

Veeam Backup Enterprise Manager can be installed on a physical or virtual machine. You can deploy it on the Veeam backup server or use a dedicated machine.

Veeam Backup Enterprise Manager uses the following services and components:

- **Veeam Backup Enterprise Manager** coordinates all operations of Veeam Backup Enterprise Manager, aggregates data from multiple Veeam backup servers and provides control over these servers.
- **Veeam Enterprise Manager Enterprise SQL Database** is used by Veeam Backup Enterprise Manager for storing data. The database instance can be located on a SQL Server installed either locally (on the same machine as Veeam Backup Enterprise Manager Server) or remotely.
- **Veeam Backup Catalog Service** replicates and consolidates guest OS file system indexing data from Veeam backup servers added to Veeam Backup Enterprise Manager. Index data is stored in Veeam Backup Enterprise Manager Catalog (a folder on the Veeam Backup Enterprise Manager Server) and is used to search for Windows guest OS files in backups created by Veeam Backup & Replication.

Veeam Backup Search

In Veeam Backup & Replication, search for guest OS files in backups is performed with Veeam Backup Enterprise Manager. However, if you frequently need to search through a great number of backups, it is recommended to install Veeam Backup Search from the installation package on a machine running Microsoft Search Server. Veeam Backup Search is an optional component in the backup infrastructure that is used for the purpose of search performance optimization.

The Veeam Backup Search server runs the **MOSS Integration Service** that invokes updates of index databases on Microsoft Search Server. The service also sends search queries to Microsoft Search Server which processes them and returns necessary search results to Veeam Backup Enterprise Manager.

U-AIR Wizards

Universal Application−Item Recovery (U−AIR), enabled by the Veeam vPower technology, allows you to recover individual items from any virtualized application. For such applications as Active Directory, Microsoft SQL and Microsoft Exchange, U−AIR is a wizard−driven process (that is, you can recover necessary items from applications using application−specific wizards). For other applications, U−AIR is user−driven (that is, Veeam Backup & Replication starts the application and all components required for its proper work in a virtual lab so that users can connect to that application and recover items themselves). For details, see Virtual Lab.

U−AIR wizards are not tied to the installation — they are standalone components that can be downloaded, installed and updated independent of the product release. You can install U−AIR wizards on any machine in your production environment from which you plan to perform the restore process.

Deployment Scenarios

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

Before installing Veeam Backup & Replication, it is strongly advised to familiarize yourself with common deployment scenarios and carefully plan your backup infrastructure layout.

Simple Deployment

In a simple deployment scenario, one instance of Veeam Backup & Replication is installed on a physical or virtual Windows−based machine. This installation is referred to as a Veeam backup server.

Simple deployment implies that the Veeam backup server fills three major roles:

- It functions as a management point, coordinates all jobs, controls their scheduling and performs other administrative activities.
- It acts as the default backup proxy for handling job processing and transferring backup traffic. All s necessary for the backup proxy functionality are installed on the Veeam backup server locally.
- It is used as the default backup repository. By default, backup files are stored to the `C:\backup` folder on the Veeam backup server.

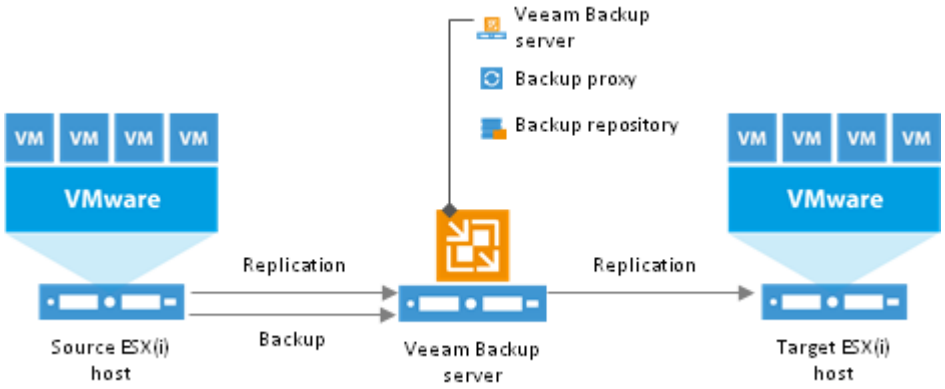

If you plan to back up and replicate only a small number of VMs or evaluate Veeam Backup & Replication, this configuration is enough to get you started.
Veeam Backup & Replication is ready for use right out of the box — as soon as it is installed, you can start using the solution to perform backup and replication operations. To balance the load of backing up and replicating your VMs, you can schedule jobs at different times.

Note If you decide to use simple deployment scenario, it is recommended that you install Veeam Backup & Replication on a VM, which will enable you to use the Virtual Appliance transport mode, allowing for LAN−free data transfer. For details, see Transport Modes.

The drawback of a simple deployment scenario is that all data is handled and stored on the Veeam backup server locally. For medium-size or large-scale environments, the capacity of a single Veeam backup server may not be enough. To take the load off the Veeam backup server and balance it throughout your backup infrastructure, we recommend that you use the advanced deployment scenario. For details, see Advanced Deployment.

Advanced Deployment

In large−scale virtual environments with a large number of jobs, the load on the Veeam Backup Server is heavy. In this case, it is recommended to use the advanced deployment scenario which moves the backup workload to dedicated backup proxies and backup repositories.

The essence of the advanced deployment is that the backup proxy takes off a part of Veeam backup server activities (namely, it collects and processes data and moves backup traffic from source to target). In addition, the Veeam backup server no longer acts as a storage location — the backup proxy transports VM data to a dedicated backup repository which is the location for keeping backup files, VM copies, metadata and so on. The Veeam backup server in this scenario functions as a "manager" for deploying and maintaining backup proxies and repositories.

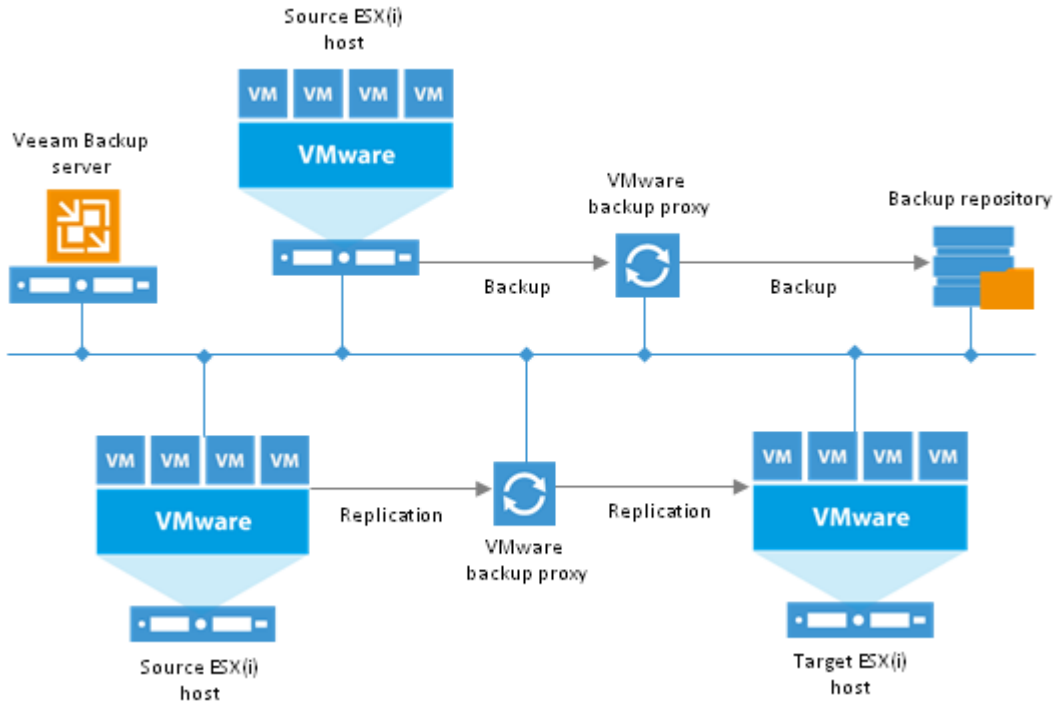

To deploy a backup proxy and/or a backup repository, you should add a server to Veeam Backup & Replication and assign a proxy and/or repository role to it. Veeam Backup & Replication will automatically install light-weight components and services onto these servers. A Backup proxy does not require a separate SQL database — all settings are stored centrally, within the SQL database used by the Veeam backup server.

With the advanced deployment scenario, you can easily meet your current and future data protection requirements. You can expand your backup infrastructure horizontally in a matter of minutes to match the amount of data you want to process and available network throughput. Instead of growing the number of backup servers or constantly tuning job scheduling, you can install multiple backup proxies and repositories and distribute the backup workload among them. The installation process is fully automated, which simplifies deployment and maintenance of the backup infrastructure in your virtual environment.

In virtual environments with several proxies, Veeam Backup & Replication dynamically distributes backup traffic among these proxies. A job can be explicitly mapped to a specific proxy. Alternatively, you can let Veeam Backup & Replication choose the most suitable proxy. In this case, Veeam Backup & Replication will check settings of available proxies and select the most appropriate one for the job. The proxy server to be used should have access to the source and target hosts as well as to the backup repository to which files will be written.

The advanced deployment scenario can be a good choice for backing up and replicating offsite. You can deploy a backup proxy in the production site and another one in the DR site, closer to the backup repository. When a job is performed, backup proxies on both sides establish a stable connection, so this architecture also allows for efficient transport of data over a slow network connection or WAN.

To regulate backup load, you can specify the maximum number of concurrent tasks per proxy and set up throttling rules to limit proxy bandwidth. The maximum number of concurrent tasks can also be specified for a backup repository in addition to the value of the combined ingestion rate for it.

Another advantage of the advanced deployment scenario is that it contributes to high availability — jobs can migrate between proxies if one of them becomes overloaded or unavailable.

Distributed Deployment

The distributed deployment scenario is recommended for large geographically dispersed virtual environments with multiple Veeam backup servers installed across different sites. These backup servers are federated under Veeam Backup Enterprise Manager — an optional component that provides centralized management and reporting for these servers through a web interface.

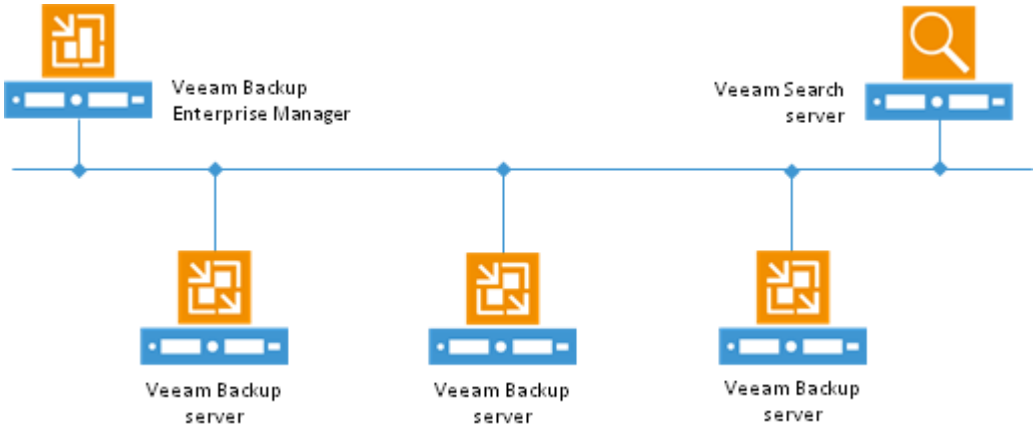

Veeam Backup Enterprise Manager collects data from Veeam backup servers and enables you to run backup and replication jobs across the entire backup infrastructure through a single "pane of glass", edit them and clone jobs using a single job as a template. It also provides reporting data for various areas (for example, all jobs performed within the last 24 hours or 7 days, all VMs engaged in these jobs and so on). Using indexing data consolidated on one server, Veeam Backup Enterprise Manager provides advanced capabilities to search for guest OS files of Windows–based VM backups created on all Veeam backup servers (even if they are stored in repositories on different sites), and recover them in a single click. Search for guest OS files is enabled through Veeam Backup Enterprise Manager itself; to streamline the search process, you can optionally deploy a Veeam Backup Search server in your backup infrastructure.

With flexible delegation options and security roles, IT administrators can delegate the necessary file restore or VM restore rights to authorized personnel in the organization – for example, allow database administrators to restore Oracle or SQL server VMs.

If you use Veeam Backup Enterprise Manager in your backup infrastructure, you do not need to install licenses on every Veeam backup server you deploy. Instead, you can install one license on the Veeam Backup Enterprise Manager server and it will be applied to all servers across your backup infrastructure. This approach simplifies tracking license usage and license updates across multiple Veeam backup servers.

In addition, VMware administrators will benefit from Veeam plug-in for vSphere Web Client that can be installed using Veeam Backup Enterprise Manager. They can analyze cumulative information on used and available storage space view and statistics on processed VMs, review success, warning, failure counts for all jobs, easily identify unprotected VMs and perform capacity planning for repositories, all directly from vSphere.

Resource Scheduling

With its built−in mechanism of resource scheduling, Veeam Backup & Replication is capable to automatically select and use optimal resources to run configured jobs. Resource scheduling is performed by the Veeam Backup Service running on the Veeam backup server. When a job starts, it communicates with the service to inform it about the resources it needs. The service analyzes job settings, parameters specified for backup infrastructure components, current load on the components, and automatically allocates optimal resources to the job.

For resource scheduling, Veeam Backup Service uses a number of settings and features:

Network Traffic Throttling and Multithreaded Data Transfer

To limit the impact of Veeam Backup & Replication jobs on network performance, you can throttle network traffic for jobs. Network traffic throttling prevents jobs from utilizing the entire bandwidth available in your environment and makes sure that enough traffic is provided for other critical network operations. It is especially recommended that you throttle network traffic if you perform offsite backup or replicate VMs to a DR site over slow WAN links.

In Veeam Backup & Replication, network traffic throttling is implemented through rules that apply to backup proxies, so you do not have to make any changes to your network infrastructure.

Network traffic throttling rules are enforced globally, at the level of the Veeam backup server. Every throttling rule limits the maximum throughput of traffic going between two backup proxies (in case of replication), between a backup proxy and a repository (in case of backup to a repository), or between a backup proxy and a proxying Windows server (in case of backup to an SMB share).

Rules are set for a pair of IP address ranges (the range can include a single IP address) and are applied to the source backup proxy and the target server (backup proxy, repository or proxying server) between which data is transferred over the network.

When a new job starts, Veeam Backup & Replication checks network traffic throttling rules against the pair of backup proxies assigned for the job — the backup proxy on the source side and the server (backup proxy, backup repository or proxying server) on the target side. If the source and target IP addresses fall into specified IP ranges, the rule will be applied. For example, if for a network traffic throttling rule you specify 192.168.0.1 – 192.168.0.255 as the source range and 172.16.0.1 – 172.16.0.255 as the target range, and the backup proxy on the source side has IP address 192.168.0.12, while the target backup proxy has IP address 172.16.0.31, the rule will be applied. The network traffic going from source to target will be throttled.

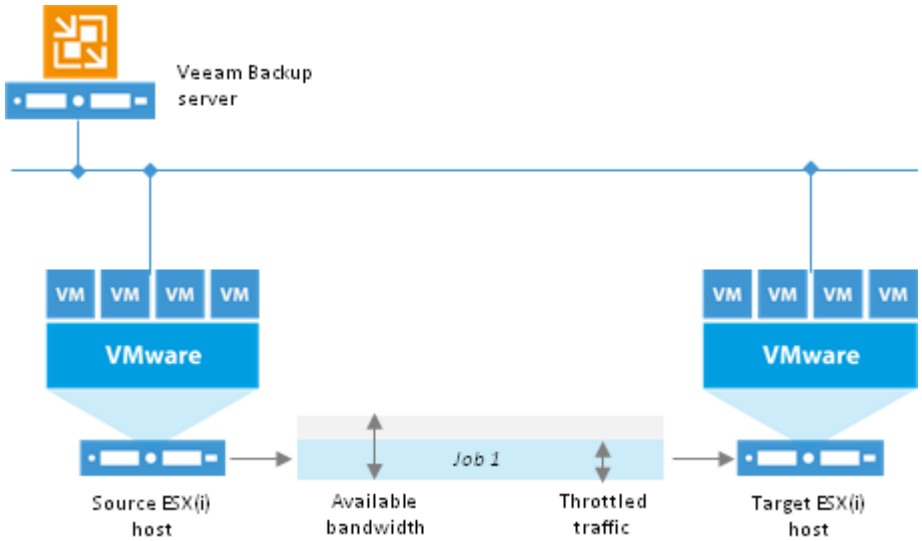

Note Throttling rules are reversible — they function in two directions. If the IP address of the server on the source side falls into the target IP range, and the IP address of the server on the target side falls into the source IP range, the rule will be applied in any case.

The Veeam backup server equally splits available bandwidth between all jobs that use backup proxies to which a network throttling rule applies. For example, if you run one job that uses a pair of proxies to which the rule applies, the job will get the entire bandwidth allowed by the rule. If you run two jobs at a time, the allowed bandwidth will be equally split between them. As soon as one of the jobs completes, the bandwidth assigned to it will be freed, and the remaining job will use the entire bandwidth allowed by the rule.

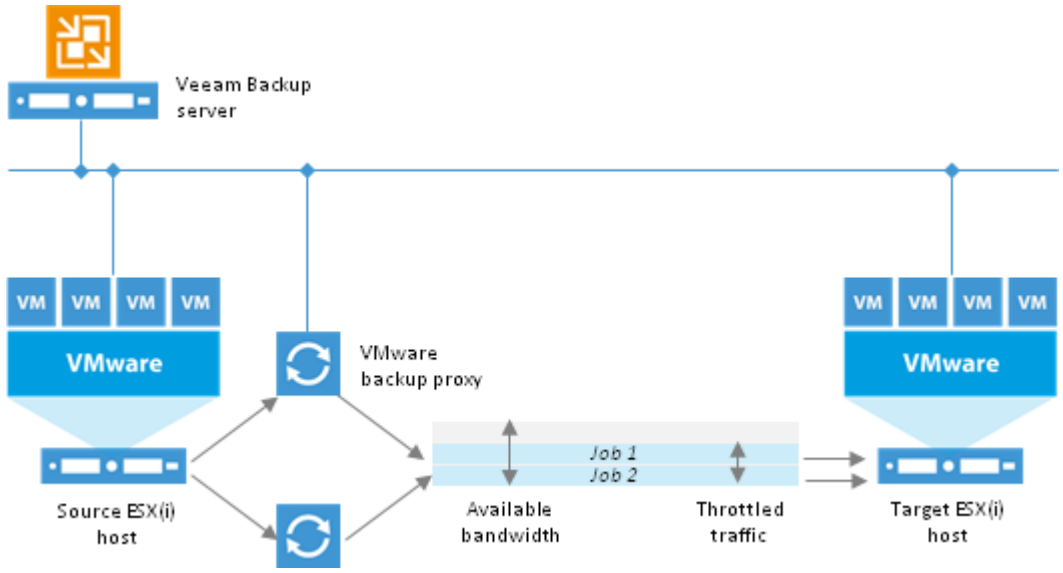

Throttling rules can be scheduled to only be active during specific time intervals (for example, during business hours). This way, you will minimize the impact of job performance spikes on the production network. Alternatively, you can select to apply throttling rules regardless of the time.

In addition to traffic throttling, Veeam Backup & Replication offers another possibility for network traffic management — management of data transfer connections. Normally, within one backup session Veeam Backup & Replication opens five parallel TCP/IP connections to transfer data from source to target. Multithreaded data transfer increases the transfer speed but can place additional load on the network. If required, you can disable multithreaded data transfer and limit the number of connections per session to one.

Note Veeam Backup & Replication performs a CRC check for the TCP traffic going between the source and the target. When you perform backup, replication or VM copy operations, Veeam Backup & Replication calculates checksums for data blocks going from the source. On the target, it re-calculates checksums for received data blocks and compares them to the checksums created on the source. If the CRC check fails, Veeam Backup & Replication automatically re-sends data blocks without any impact on the job.

Limiting the Number of Concurrent Tasks

To avoid overload of backup proxies and backup repositories, Veeam Backup & Replication allows you to limit the number of concurrent tasks performed on a backup proxy or targeted at a backup repository.

Before processing a new task, Veeam Backup & Replication detects what backup infrastructure components (backup proxies and repositories) will be involved. When a new job starts, Veeam Backup & Replication analyzes the list of VMs that the job includes, and creates a separate task for each disk of every VM to be processed. With Veeam Backup & Replication 7.0, tasks in the job can be processed in parallel (that is, VMs and VM disks within a single job can be processed simultaneously), optimizing your backup infrastructure performance and increasing the efficiency of resource usage.

Note To use this capability, proxy server(s) should meet system requirements – each task requires a single CPU core, so for two tasks to be processed in parallel, 2 cores is the recommended minimum. Also, parallel VM processing should be enabled in Veeam Backup & Replication options.

Task limiting is performed by the Veeam Backup Service that is aware of all backup proxies and backup repositories connected to it, and settings specified for these backup proxies and repositories (namely, the number of allowed concurrent tasks). When a job starts, it informs the Veeam Backup service about its task list and polls the service about allocated resources to its tasks at a 10 second interval after that. Before a new task targeted at a specific backup proxy or repository starts, the Veeam Backup Service checks the current workload (the number of tasks currently working with the proxy or repository) and the number of allowed tasks per this component. If this number is exceeded, a new task will not start until one of the currently running tasks is finished.

Limiting Data Ingestion Rate for Backup Repositories

Veeam Backup & Replication can limit the data ingestion rate for backup repositories. The Veeam Backup Service is aware of data ingestion settings configured for all repositories in the backup infrastructure. When a job targeted at some backup repository is performed, the Veeam Backup Service informs the Veeam transport service running on the repository about the allowed write speed specified for this repository so that the Veeam transport service can limit the write speed to the specified value.

If the repository is used by a number of jobs, the allowed write speed is equally split between all of them.

Detecting Performance Bottlenecks

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

Veeam Backup & Replication processes VM data in cycles. Every cycle includes a number of stages:

- Reading VM data blocks from the source
- Processing VM data on the backup proxy
- Transmitting data over the network
- Writing data to the target

When one data processing cycle is over, the next cycle begins. VM data therefore goes over the “data pipe”.

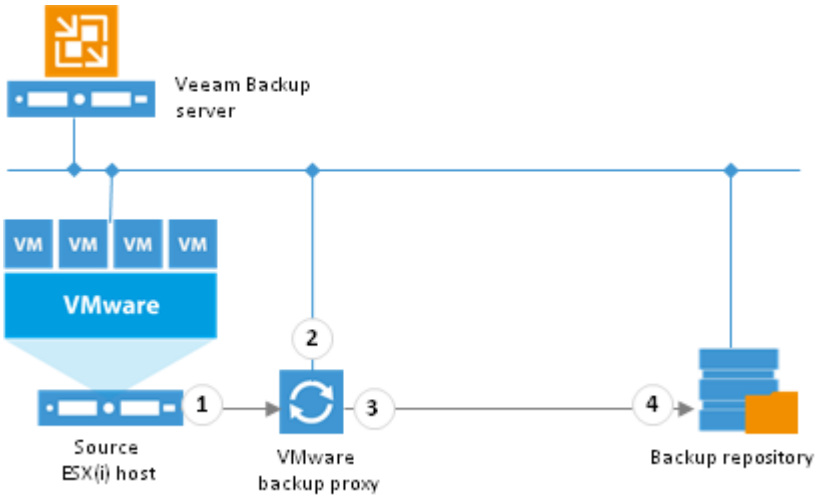

To evaluate the data pipe efficiency, Veeam Backup & Replication analyzes performance of all components in the data flow working as the cohesive system and evaluates key factors on the source and target sides. In the data pipe, the following points, or components, are considered:

(1) **Source** — the source disk reader component responsible for retrieving data from the source storage.

(2) **Proxy** — the backup proxy component responsible for processing VM data.

(3) **Network** — the network queue writer component responsible for getting processed VM data from the backup proxy and sending it over the network to the backup repository or another backup proxy.

(4) **Target** — the target disk writer component (backup storage or replica datastore).

The resource usage level for the four points is evaluated in percent. This percent rate defines the amount of time for which components are busy during the job. An efficient data flow assumes that there is no latency at any point of the data pipe and all its components work for approximately equal amount of time.

If any of the components operates inefficiently, there may appear a bottleneck in the data path. The insufficient component will work 100% of time while the others will be idling, waiting for data to be transferred. As a result, the whole data flow will slow down to the level of the slowest point in the data path and the overall time of data processing will increase.

To identify a bottleneck in the data path, Veeam Backup & Replication detects the component with the maximum workload: that is, the component that works for the most time of the job. For example, you use a low-speed storage device as the backup repository. Even if VM data is retrieved from SAN storage on the source side and transmitted over a high-speed link, VM data flow will still be impaired at the backup repository. The backup repository will be trying to consume transferred data at the rate that exceeds its capacity and the other components will stay idle. As a result, the backup repository will be working 100% of job time, while other components may be employed, for example, for 60% only. In terms of Veeam Backup & Replication, such data path will be considered insufficient.

The bottleneck statistics for a job is displayed in the job session data. The bottleneck statistics does not necessarily mean that you have a problem in your backup infrastructure; it simply informs you about the weakest component in the data path. However, if you feel that the job performance is low, you may try taking some measures to resolve the bottleneck. For instance, in the case described above, you can limit the number of concurrent tasks for the backup repository.

Backup

Unlike traditional backup tools designed to work with physical machines, Veeam Backup & Replication is built specifically for virtual environments. It operates at the virtualization layer and uses an image−based approach for VM backup. To retrieve VM data, no agent software needs to be installed inside the guest OS. Instead, Veeam Backup & Replication leverages ESX snapshot capabilities. When a new backup session starts, a snapshot is taken to create a cohesive point−in−time copy of a VM including its configuration, OS, applications, associated data, system state and so on.
Veeam Backup & Replication uses this point−in−time copy to retrieve VM data. Image−based backups can be used for different types of recovery, including full VM recovery, VM file recovery,Instant VM Recovery, file-level recovery and other.

Use of the image−based approach allows Veeam Backup & Replication to overcome shortfalls and limitations of traditional backup (such as, the necessity to provide guest OS credentials for every VM, significant resource overhead on the VM and hypervisor during the backup process, management overhead and so on). It also helps streamline recovery verification and the restore process — to recover a single VM, there is no need to perform multiple restore operations.
Veeam Backup & Replication uses a cohesive VM image from the backup to restore a VM to the required state without the necessity for manual reconfiguration and adjustment.

In Veeam Backup & Replication, backup is a job−driven process where one backup job can be used to process one or more VMs. A job is a configuration unit of the backup activity. Essentially, the job defines when, what, how and where to back up. It indicates what VMs should be processed, what components should be used for retrieving and processing VM data, what backup options should be enabled and where to save the resulting backup file. Jobs can be started manually by the user or scheduled to run automatically.

The resulting backup file stores compressed and deduplicated VM data. All backup files created by the job are located in a dedicated job folder on a backup repository. Veeam Backup & Replication creates and maintains the following types of backup files:

- Full backup (.vbk) to store copies of full VM images
- Backup increment (.vib or .vrb) to store incremental changes to VM images
- Backup metadata (.vbm) to provide information on the backup job, VMs in the backup, number and structure of backup files, restore points, and so on. The metadata file facilitates import of backups or mapping of backup jobs to existing backups.

To back up VMs, you can use one of two available methods: incremental backup or reversed incremental backup. Regardless of the method you use, the first run of a job creates a full backup of VM image. Subsequent job runs are incremental — Veeam Backup & Replication copies only those data blocks that have changed since the last backup job run. To keep track of changed data blocks, Veeam Backup & Replication uses different approaches, including VMware's Changed Block Tracking (CBT) technology.

Transport Modes

Efficiency of a backup job and time required for its completion in many respects depends on the transport mode. Transport mode is a method that is used by the Veeam transport service to retrieve VM data from the source host and write VM data to the target destination.

For data retrieval, Veeam Backup & Replication offers the following modes (starting from the most efficient):

- Direct SAN Access
- Virtual Appliance
- Network

The Veeam transport service responsible for data retrieval runs on a backup proxy server. Correspondingly, the transport mode can be defined in the settings of the backup proxy that performs the job. When configuring backup proxy settings, you can manually select a transport mode or let Veeam Backup & Replication select the most appropriate mode automatically. If you use automatic mode selection, Veeam Backup & Replication will scan backup proxy configuration and its connection to the VMware infrastructure to choose the optimal transport mode. If multiple transport modes are available for the same proxy, Veeam Backup & Replication will choose the mode in the following order: Direct SAN Access > Virtual Appliance > Network.

For writing data to the target destination, Veeam Backup & Replication normally uses the Network mode. In some cases, such as VM replication or full VM recovery, Veeam Backup & Replication also supports the Virtual Appliance mode. You cannot select the transport mode for writing data — Veeam Backup & Replication selects it automatically, based on the backup proxy configuration. Whenever possible, the Virtual Appliance mode is used.

For all transport modes, Veeam Backup & Replication leverages VMware vStorage APIs for Data Protection (VADP). VADP can be used with VMware vSphere 4 and 5 (including ESX/ESXi), vCenter Server 4, VMware Virtual Infrastructure (ESX/ESXi 3.5 and VirtualCenter 2.5). VADP is not supported for ESX 3.0.

Applicability and efficiency of each transport mode primarily depends on the type of datastore used by the source host (local or shared), and on the backup proxy server type (physical or virtual). The table below shows recommendations for installing the backup proxy, depending on the storage type and desired transport mode.

<table>
<tr><th>Production Storage Type</th><th>Direct SAN Access</th><th>Virtual Appliance</th><th>Network Mode</th></tr>
<tr><td>Fiber Channel (FC) SAN</td><td>Install the backup proxy on a physical server with direct FC access to the SAN.</td><td rowspan="3">Install the backup proxy on a VM running on an ESX(i) host connected to the storage device.</td><td rowspan="3">Not recommended</td></tr>
<tr><td>iSCSI SAN</td><td>Install the backup proxy on a physical or virtual server.</td></tr>
<tr><td>NFS Storage</td><td>Not supported</td></tr>
<tr><td>Local Storage</td><td>Not supported</td><td>Install backup proxies on VMs in every ESX(i) host.</td><td>Install the backup proxy on one physical or virtual server in the environment.</td></tr>
</table>

Direct SAN Access

This mode is recommended if the ESX(i) host uses shared storage. The backup proxy leverages VADP to retrieve VM data directly from FC or iSCSI storage in the Storage Area Network (SAN). To retrieve data blocks from SAN LUN, the backup proxy uses metadata about the layout of VM disks on the SAN. Since data blocks are not retrieved over the Local Area Network (LAN), this mode minimizes disruptions to your production network during backup.

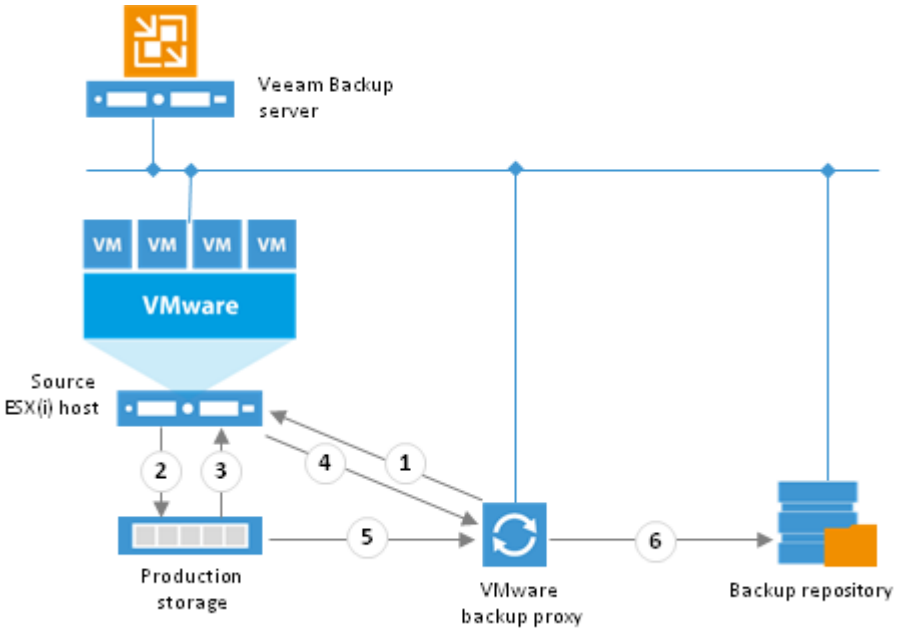

The process of data retrieval in Direct SAN Access mode includes the following steps:

1. The backup proxy sends a request to the ESX(i) host to locate the necessary VM on the datastore (1).
2. The host locates the VM and retrieves metadata about the layout of virtual disks on the SAN (that is, physical addresses of data blocks) and sends this metadata to the backup proxy (2, 3, 4).
3. The backup proxy uses the metadata to copy data blocks directly from the SAN and sends them to the target (5, 6).

Important! The backup proxy using Direct SAN Access mode must be connected directly into the SAN fabric. VM processing will fail if a direct SAN connection is not configured or not available when the job starts.

Veeam Backup & Replication allows you to switch to the Network mode and retrieve VM data through the ESX(i) host over the LAN if the SAN becomes inaccessible. This option is selected by default to ensure that backup jobs can be completed in any case. However, it puts additional load on your production network and thus may potentially affect performance if you are running backup and replication jobs during business hours.

If you install Veeam Backup & Replication in a VM and plan to use the Direct SAN access mode, you need to make sure that the Veeam backup server has direct access to the production storage. Refer to this post on Veeam's blog for details on additional configuration of the Veeam backup server.

Virtual Appliance

This mode is recommended and can only be used if the backup proxy is deployed on a VM. The Virtual Appliance mode uses the SCSI hot-add capability of ESX hosts to attach disks of the backed up VM to the backup proxy VM (or to the helper VM, depending on vCenter version you are using). In this mode, VM data is retrieved directly from storage through the ESX I/O stack, instead of going through the network stack, which improves performance. Note that the ESX(i) host on which the backup proxy VM resides must have access to the storage where disks of a backed up VM are located.

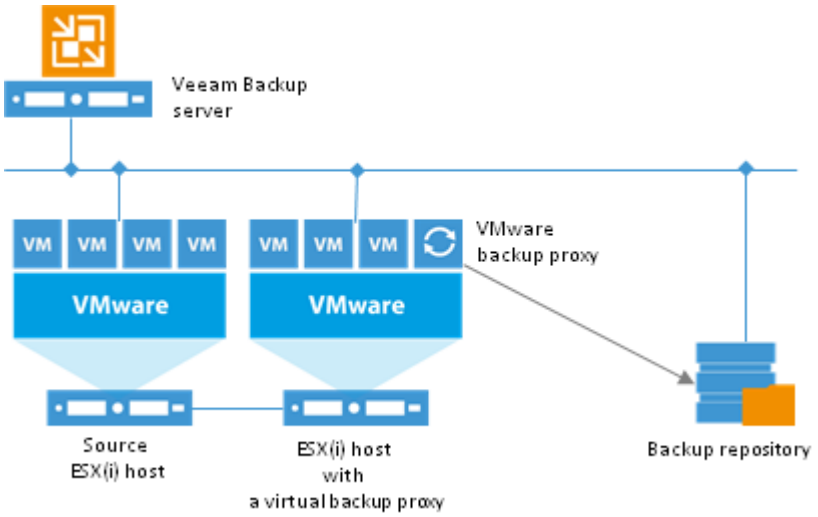

Important! If you are using vCenter Server earlier than version 4.0, a helper vM named *VeeamBackupVMName(VCBHELPER)* must also be created on the same ESX(i) host where the backup proxy is running. For example, if your backup proxy server is named *vbsrv01*, the helper appliance name must be *vbsrv01(VCB-HELPER)*.
The helper VM is a blank dummy VM without virtual disks or OS installed. This VM is only used to temporarily 'hot-add' disks of VMs that are being backed up.
To get a full list of requirements and limitations for the Virtual Appliance mode, see http://www.veeam.com/KB1054.

Network Mode

This mode can be used with any infrastructure configuration. However, when an alternative transport mode is applicable, the Network mode is not recommended because of the lowest data retrieval speed. It is the only applicable mode when the backup proxy is a physical machine and the host uses local storage. In this mode, data is retrieved via the ESX(i) host over the LAN using Network Block Device protocol (NBD).

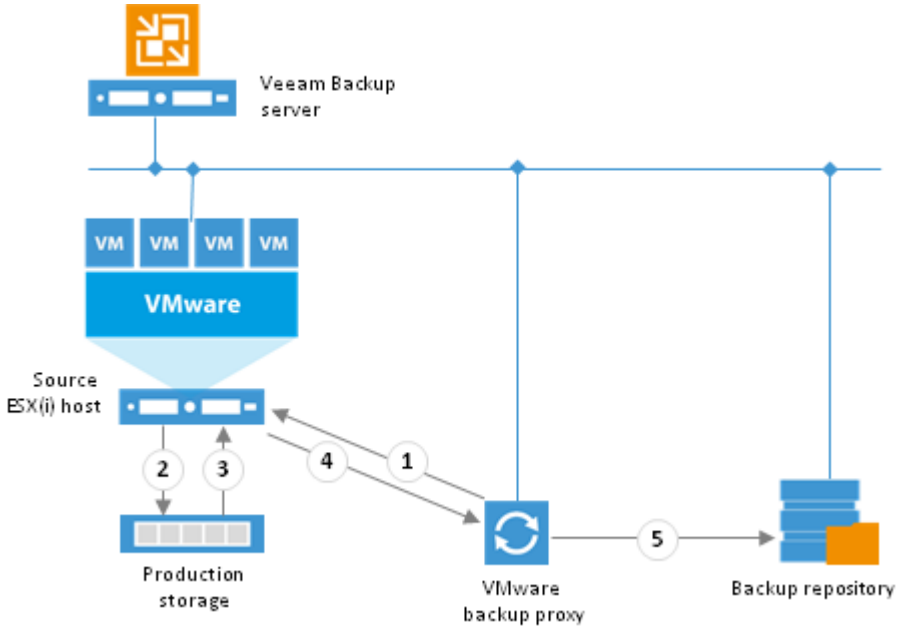

The process of data retrieval in Network mode includes the following steps:

1. The backup proxy sends a request to the ESX(i) host to locate the necessary VM on the datastore (1).
2. The host locates the VM, copies blocks of data and sends them to the backup proxy over the LAN (2, 3 and 4).
3. The backup proxy sends the data to target (5).

The Network mode is not recommended because of low traffic throughput via the LAN (the copy of the VM disk usually contains a lot of data). In order to take the load off the LAN, Veeam Backup & Replication provides two alternative modes: Direct SAN Access and Virtual Appliance.

Veeam Backup & Replication processes VM disks one by one or in parallel, depending on selected data processing settings. If VM disks are located on different storages (for example, on the SAN and local storage subsystem), Veeam Backup & Replication will use different transport modes to process VM disks. In such scenario, it is strongly recommended that you select the **Failover to network mode if primary transport modes fail or are unavailable** option when configuring the mode settings for the necessary backup proxy.

Backup Architecture

The backup infrastructure in a VMware vSphere environment comprises the following components:

- One or more source hosts with associated datastores
- One or more backup proxy servers
- One or more backup repositories

The source host and the repository produce two terminal points between which VM data is moved. Backup data is collected, transformed and transferred with the help of Veeam transport services. Veeam Backup & Replication uses two−service architecture — one Veeam transport service controls interaction with the source host during data transfer and the other one controls interaction with the backup repository. The Veeam transport services communicate with each other and maintain a stable connection. All backup infrastructure components engaged in the job make up a data pipe. VM data is moved over this data pipe block by block, so that processing of a single VM includes multiple processing cycles.

When a new backup session is started, the target−side Veeam transport service obtains the job instructions and communicates with the source−side Veeam transport service to begin data collection.

1. The **source-side Veeam transport service** copies VM data from the snapshot using one of VMware transport modes, as prescribed by the backup proxy server settings. While copying, the source−side Veeam transport service performs additional processing — it consolidates the content of virtual disks by filtering out overlapping snapshot blocks, zero−data blocks and blocks of swap files. During incremental job runs, the Veeam transport service retrieves only those data blocks that have changed since the previous job run. Copied blocks of data are compressed and moved from the source−side Veeam transport service to the target−side transport service.
2. The **target−side Veeam transport service** deduplicates similar blocks of data and writes the result to the backup file in the backup repository.

Veeam Backup & Replication supports a number of backup scenarios that depend on the type of repository you use and its location. For details, see Backup Repository.

Onsite Backup

To back up to an onsite Windows or Linux repository, you need to deploy a backup proxy on a server that has access to the source datastore, and point the backup job to this proxy. In this scenario, the source–side Veeam transport service is started on the proxy server, and the target–side Veeam transport service is started on the Windows or Linux repository server. Backup data is sent from the backup proxy to the repository over LAN.

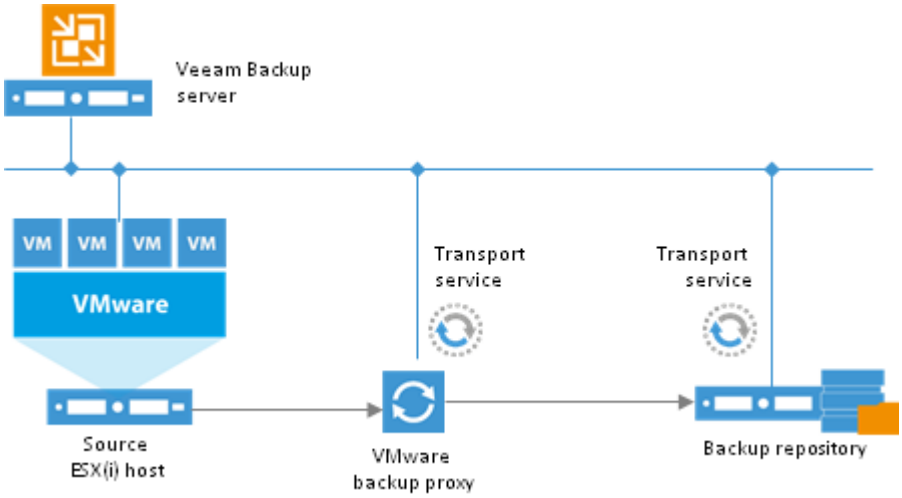

To back up to an onsite SMB share, you need a Windows–based proxying server that has access to the SMB share. This can be either the Veeam backup server or another Windows server added to the Veeam Backup & Replication console.

You can use one Windows server as the backup proxy and proxying server for SMB. In this scenario, Veeam Backup & Replication starts the source–side and target–side Veeam transport services on the same server. Backup data is sent from the proxy to the target SMB share over LAN.

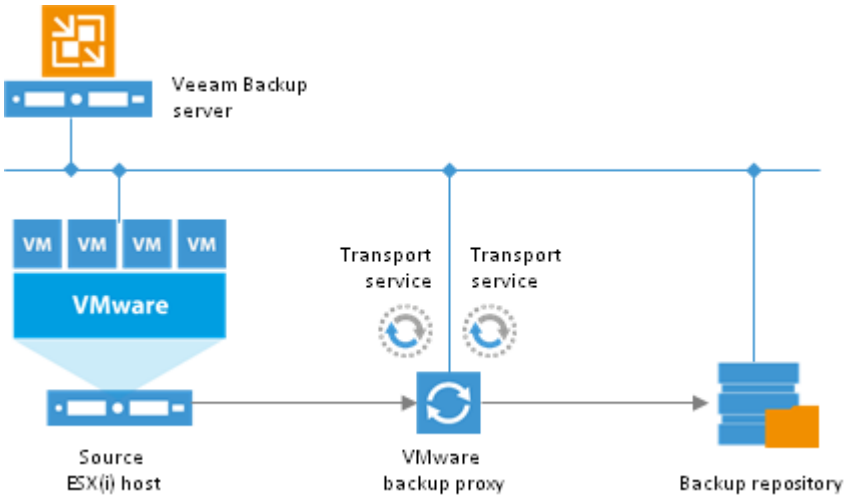

Offsite Backup

The common requirement for offsite backup is that one Veeam Transport Service runs in the production site (closer to the source datastore) and the other Veeam Transport Service runs in the remote target site (closer to the repository). During backup, Veeam Transport Services maintain a stable connection, which allows for uninterrupted operation over WAN or slow links.

To perform offsite backup to a Windows or Linux repository, you need to deploy a backup proxy in the production site (closer to the source datastore). In this scenario, the source–side Veeam Transport Service is started on the proxy server and the target–side Veeam Transport Service is started on the Windows or Linux repository server. Backup data is sent from the proxy to the repository over WAN.

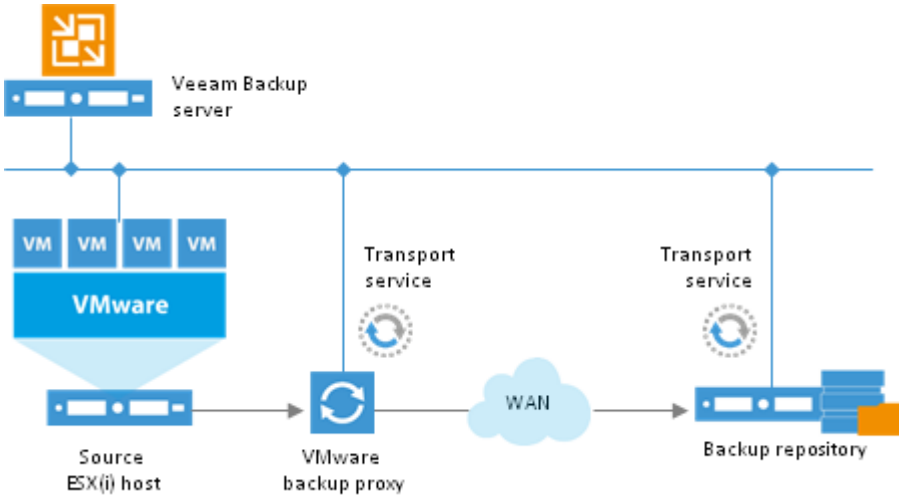

To back up VMs to an offsite SMB share, you should deploy a backup proxy in the source site and an additional Windows–based proxying server in the remote site. The SMB repository should be configured to point to the target–side proxying server. During backup the source–side Veeam Transport Service runs on the source proxy in the production site and the target–side Veeam Transport Service runs on the target proxying server in the remote site. Backup data is transferred between the backup proxy and the proxying server over WAN.

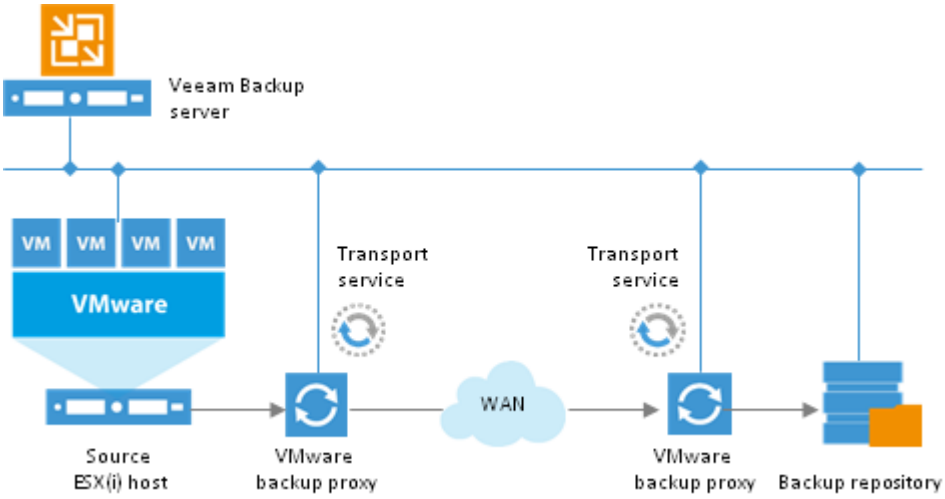

Backup Methods

Veeam Backup & Replication provides two methods for creating regular backup files:

- Reversed incremental backup (recommended for disk-based backup)
- Forward incremental backup (recommended for disk-to-disk-to-tape and remote site backups)

Additionally, it is possible to create periodic active full and synthetic full backups.

Note With Veeam Backup & Replication, you can easily switch between backup methods. Veeam Backup & Replication will not transform the previously created chain. Instead, it will create a new chain next to the existing one in the following manner:

- If you switch from the reversed incremental method to the forward incremental method, Veeam Backup & Replication will create a set of incremental backups next to the reversed incremental chain. The full backup in the reversed incremental chain will be used as a starting point for produced increments.
- If you switch from the forward incremental method to the reversed incremental method, Veeam Backup & Replication will first create a full backup next to the incremental backup chain. At every new job cycle, Veeam Backup & Replication will transform this full backup and add reversed incremental backups to the chain.

Reversed Incremental Backup

Reversed incremental backup implies that during the first run of a backup job a full backup of a VM is created. VM data is copied block by block, compressed at an appropriate compression level, and stored in a resulting full backup file (.vbk). All subsequent backups are incremental (that is, Veeam Backup & Replication copies only those data blocks that have changed since the last job run). During reverse incremental backup, Veeam Backup & Replication "injects" changes into the .vbk file to rebuild it to the most recent state of the VM. It also creates a reversed incremental backup file (.vrb) containing data blocks that are replaced when the full backup file is rebuilt. Therefore, the most recent restore point is always a full backup, and it gets updated after every backup cycle.

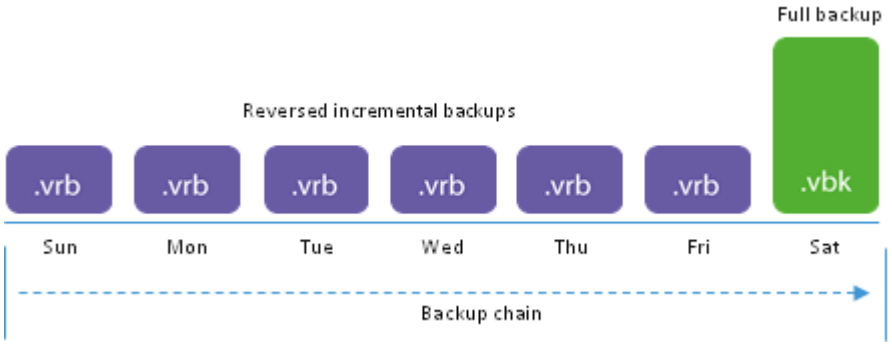

This backup method lets you perform forever-incremental backup and save disk space as you have to store only one full backup on the backup repository. If the number of restore points allowed by the retention policy is exceeded, Veeam Backup & Replication will simply delete the oldest reversed increment. For details, see Retention Policy.

Reversed incremental backup enables you to immediately restore a VM to the most recent state without extra processing, because the most recent restore point is a full backup file. If you need to restore a VM to a particular point in time, Veeam Backup & Replication will apply the required .vrb files to the .vbk file to get to the required restore point.

Forward Incremental Backup

During the first run of a forward incremental backup, or simply *incremental backup*, Veeam Backup & Replication creates a full backup file (.vbk). At subsequent backups, it only gets changes that have taken place since the last performed backup (whether full or incremental) and saves them as incremental backup files (.vib) next to the full backup.

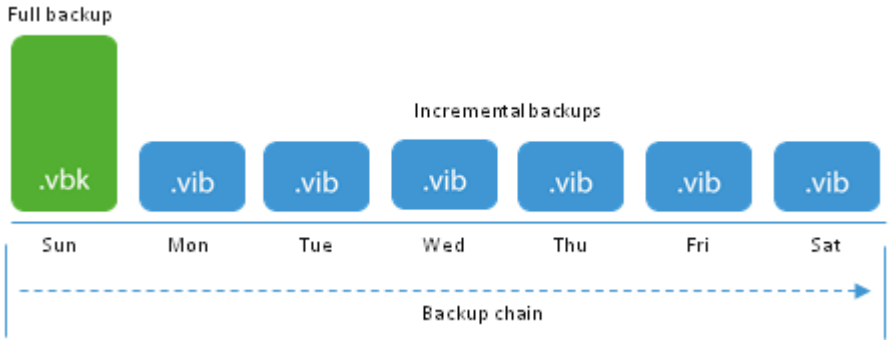

Incremental backup is the best choice if company regulation and policies require you to regularly move a created backup file to tape or a remote site. With incremental backup, you move only incremental changes, not the full backup file, which takes less time and requires less tape. You can initiate writing backups to tape or a remote site in Veeam Backup & Replication itself, by configuring post-backup activities.

Important! If you decide to use the forward incremental backup method, it is necessary to schedule the creation of periodic active full or synthetic full backups. This will help you avoid long chains of increments, ensure safety of backup data and allow you to meet the requirements of your retention policy. For details, see Retention Policy.

Active and Synthetic Full Backups

To let you get the most out of incremental backup, Veeam Backup & Replication enables you to create active full backups and schedule creation of synthetic full backups on specific days.

Active Full Backup

In some cases, you need to regularly create a full backup. For example, your corporate backup policy may require that you create a full backup on weekend and run incremental backup on work days. To let you conform to these requirements, Veeam Backup & Replication offers the ability to periodically perform active full backups.

The active full backup produces a full backup of a VM, just as if you run the backup job for the first time. Veeam Backup & Replication retrieves data for the whole VM from the source, compresses and deduplicates it and stores it to the full backup file — .vbk.

The active full backup resets the chain of increments: all subsequent increments use the latest active full backup as a new starting point. A previously used full backup file remains on disk until it is automatically deleted according to the backup retention policy.

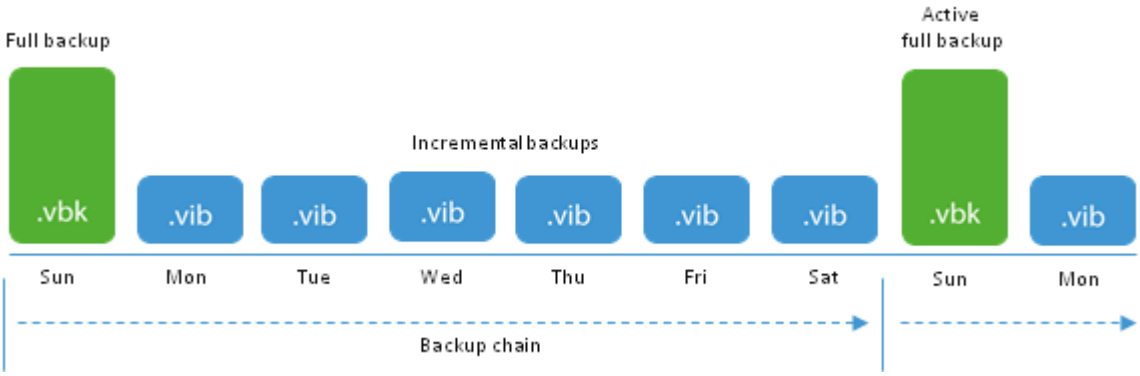

You can create active full backups manually or schedule a backup job to create active full backups with a certain periodicity.

- To create an active full backup manually, use the **Active Full** command from the shortcut menu of a corresponding backup job.
- To schedule active full backups, specify scheduling settings in the **Advanced** section of a corresponding backup job. You can schedule active full backups to run weekly, for example, every Saturday, or monthly, for example, every fourth Sunday of a month.

Synthetic Full Backup

In some situations, running active full backups periodically may not be an option. Full backups are very resource-intensive and consume considerable amount of network bandwidth. As an alternative, you can create synthetic full backups.

In terms of data, the synthetic full backup is identical to a regular full backup. The synthetic full backup is a .vbk file that contains data of the whole VM. The difference between these two backup types lies in the way how VM data is retrieved:

- When you perform full backup, Veeam Backup & Replication retrieves VM data from the source datastore (volume) where the VM resides, compresses and deduplicates it and writes it to the .vbk file on the backup repository.
- When you perform synthetic full backup, Veeam Backup & Replication does not retrieve VM data from the source datastore (volume). Instead, it "synthesizes" a full backup from data you already have on the backup repository. Veeam Backup & Replication accesses the previous full backup file and a chain of subsequent increments on the backup repository, consolidates VM data from these files and writes consolidated data into a new full backup file. As a result, the created synthetic full backup file contains the same data you would have if you created a full backup in a regular manner.

Veeam Backup & Replication treats a synthetic full backup as a regular full backup. As well as any other full backup, the synthetic full backup resets the chain of increments. All subsequent increments use the created synthetic full backup as a new starting point. A previously used full backup file remains on disk until it is automatically deleted according to the backup retention policy.

The synthetic full backup has a number of advantages:

- The synthetic full backup does not use network resources: it is created from backup files you already have on disk.
- The synthetic full backup imposes less load on the production environment: it is created right on the backup repository.

With Veeam Backup & Replication, you can schedule creation of synthetic full backups on specific days.

For example, you can configure a backup job to perform daily forward incremental backups and schedule synthetic fulls on Thursday. Veeam Backup & Replication will perform incremental backup Sunday through Wednesday as usual. On Thursday it will perform a synthetic full backup in the following way:

1. Veeam Backup & Replication will first perform incremental backup in the regular manner.
2. At the end of the backup job, the Veeam transport service on the backup repository will build a new synthetic full backup from the backup files you already have on disk: the full backup created on Sunday and a chain of increments, Monday through Wednesday plus the new increment created on Thursday.
3. The Veeam transport service will delete the increment created on Thursday.

As a result, you will have a backup chain consisting of the full backup created on Sunday, three increments created on Monday through Wednesday and a synthetic full backup created on Thursday.

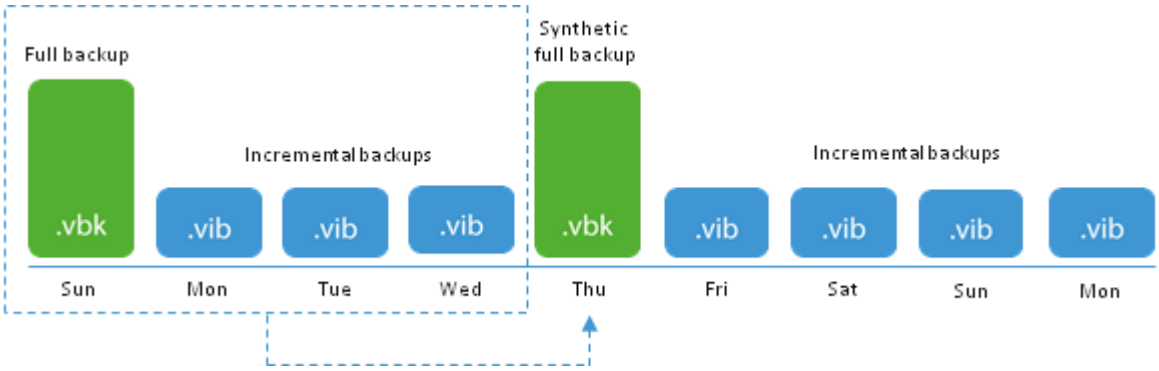

Every next run of the backup job will create an incremental backup starting from the synthetic full backup until next Thursday. On the next Thursday, Veeam Backup & Replication will create a new synthetic full backup.

Note Veeam Backup & Replication creates a synthetic full backup only once a day on which it is scheduled. If you run the backup job again on Thursday, Veeam Backup & Replication will perform incremental backup in the regular manner.

Transforming Incremental Backup Chains into Reversed Incremental Backup Chains

If you select to create synthetic full backups, you can additionally choose to transform a previous forward incremental backup chain into a reversed incremental backup chain. In this case, Veeam Backup & Replication will transform the latest backup chain consisting of the .vbk and .vib files into reversed incremental backups - .vrb files.

The transform option lets you dramatically reduce the amount of space required to store backups. Instead of two full backups — the regular full backup and the synthetic full backup — you will have only one synthetic full backup on disk. Note, however, that the transform process takes more time than simply creating a periodic synthetic full backup.

For example, you have configured a backup job to perform daily forward incremental backups and scheduled synthetic fulls on Thursday. Additionally, you have selected to transform the incremental backup chain into the reversed incremental backup sequence. The backup job starts on Sunday. In this case, Veeam Backup & Replication will perform backup in the following way:

1. On Sunday, Veeam Backup & Replication will create a full backup; Monday through Wednesday Veeam Backup & Replication will create incremental backups and add them to the backup chain.
2. On Thursday, Veeam Backup & Replication will first create an incremental backup in the regular manner and add it to the backup chain.
3. After that, Veeam Backup & Replication will transform the incremental backup chain into the reversed incremental chain. As a result, you have a full backup “created” on Thursday and a set of reversed increments Sunday through Wednesday.
4. When you run the backup job next time, Veeam Backup & Replication will add a new incremental backup to the chain; the synthetic full backup will be used as a starting point.

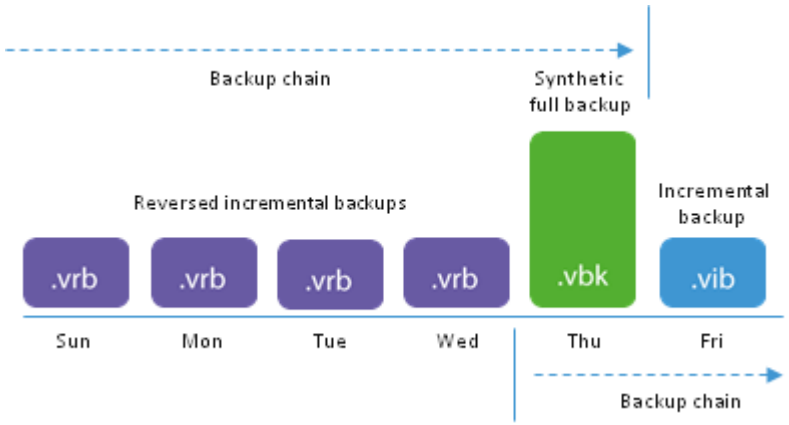

Veeam Backup & Replication always transforms only the latest incremental backup chain. For example, you have a backup chain that consists of one full backup file and set of increments. In the middle of the chain, you create an active full backup. When you run a transformation task, Veeam Backup & Replication will transform the most recent active full backup plus increments that follow it. All backups that precede the active full backup will stay intact.

Note The transform process is accounted for as an active backup repository task. Make sure you properly plan use of backup repository resources when you schedule backup jobs.

Retention Policy

Every successful run of a job creates a new restore point that lets you return your data to an earlier point in time. When you define retention policy, you specify how many restore points you want to keep and thus how 'far' you want to be able to roll back. Once the specified number is exceeded, the earliest restore point will be automatically removed. So if the retention policy is set to three and you already have three restore points, the fourth successful run of a job will delete the restore point created at the first job run.

Note When the allowed number of restore points in the backup chain is exceeded, Veeam Backup & Replication deletes the whole backup file, not separate VMs from it. To learn more, see Removing Restore Points from the Backup Chain.

Veeam Backup & Replication handles restore points in different ways for incremental and reversed incremental backups.

Retention for Reversed Incremental Backup

In case of reversed incremental backup, Veeam Backup & Replication immediately deletes the earliest reverse increment as soon as it meets the retention policy. For example, if the retention policy is set to three restore points, two latest reverse increments and a full backup will be retained.

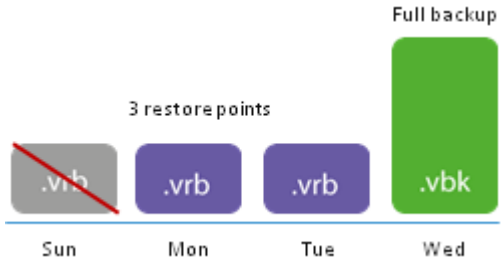

Retention for Incremental Backup

To be able to restore from a forward incremental backup, you need to have a full backup and a chain of subsequent increments on disk. If you delete a full backup, the whole chain of increments will become useless. In a similar manner, if you delete any increment before the restore point to which you want to roll back, you won't be able to restore your data (since later increments depend on earlier increments).

For this reason, if you select forward incremental backup, in some days there will be more restore points on disk than specified by your retention policy. Veeam Backup & Replication will remove the full backup chain only after the last increment in the chain meets your retention policy (which will happen once the retention policy reaches the next full backup).

For example, the retention policy is set to three restore points. A full backup is performed on Sunday, incremental backups are performed Monday through Saturday, and a synthetic full backup is scheduled on Thursday. Although the policy is already breached on Wednesday, the full backup is not deleted because without it the chain of increments would be useless, leaving you without any restore point at all. Thus, Veeam Backup & Replication will wait for the next full backup and two increments to be created, and only then delete the whole previous chain consisting of the full backup and increment, which will happen on Saturday.

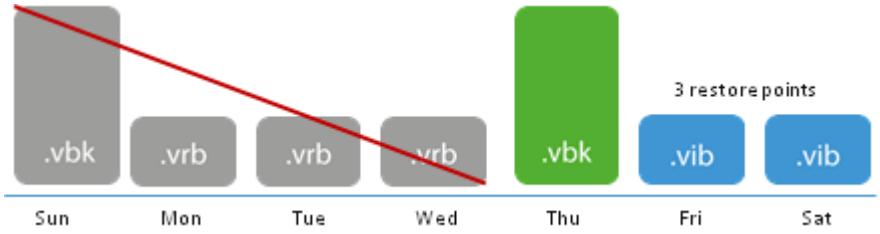

Retention Policy for Deleted VMs

In some situations, after configure and run backup jobs in Veeam Backup & Replication, you may want to change something in your virtual environment or even in your backup strategy. For example, you may remove some VMs from the virtual infrastructure or move them to some other location. You may also exclude some VMs from jobs that have already been run for some time.

By default, when you remove a VM protected by Veeam Backup & Replication from the virtual infrastructure or exclude it from a job, its backup files still remain on the backup repository. To avoid keeping redundant data on disk, you can select to control retention policy for deleted VMs.

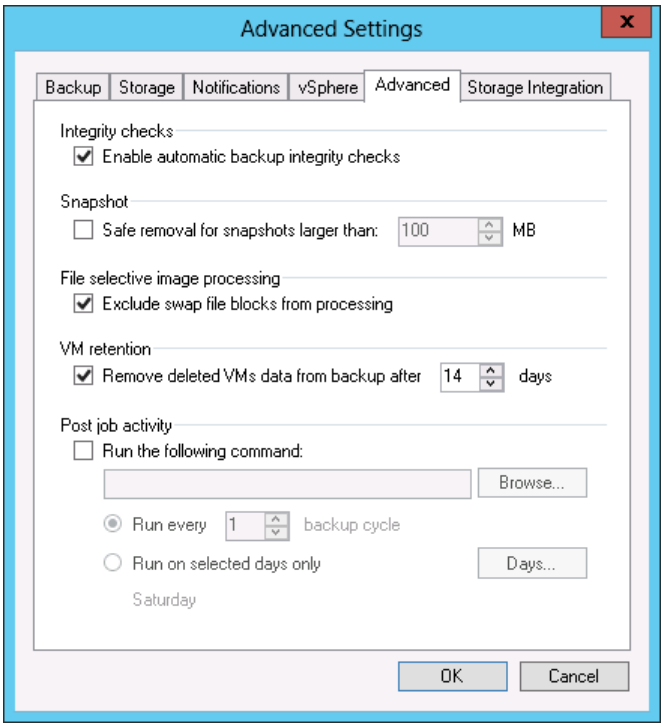

The retention policy for deleted VMs is an option in the backup job settings. By default, this option is disabled. To use this option, you need to select the **Remove deleted VMs data from backup after** check box specify the desired period of time for which data for deleted VMs must be retained on the backup repository.

With this option enabled, Veeam Backup & Replication will check the list of VMs included in the job when a job starts. If a VM is no longer available, for example, it was deleted or moved to another location, Veeam Backup & Replication will keep its data in the backup file for the specified period only. As soon as the specified retention period is over, data of the deleted VM will be removed from backup files on the backup repository.

Important! Retention policy for deleted VMs is applied only to reversed incremental backup chains and forward incremental backup chains for which synthetic full backups with subsequent transform is enabled.

Removing Restore Points from the Backup Chain

To keep up with the retention policy, Veeam Backup & Replication deletes the whole backup file, not separate VMs from it. For this reason, in some situation a certain VM may have more restore points in the backup chain than it is allowed by the retention policy settings. This can happen if a backup job processes a number of VMs or VM containers and some VMs or containers fail to be processed in some job sessions.

For example, a backup job processes two VMs: *VM1* and *VM2*. According to the retention policy settings, the backup chain must contain 7 restore points. The backup job has already had 7 job sessions and VMs have been processed in the following way:

- *VM1* has been successfully backed up 7 times and has 7 restore points
- *VM2* has failed to be processed in two job sessions and therefore has 5 restore points

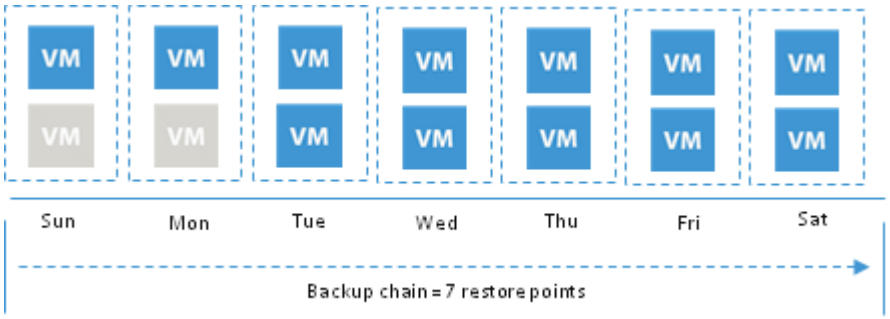

When a new restore point is added to the chain, Veeam Backup & Replication will not remove the earliest restore point because the number of restore points for *VM2* has not reached 7. Instead, Veeam Backup & Replication will wait for two more successful backup job sessions for VM2 and remove the earliest restore point only after that.

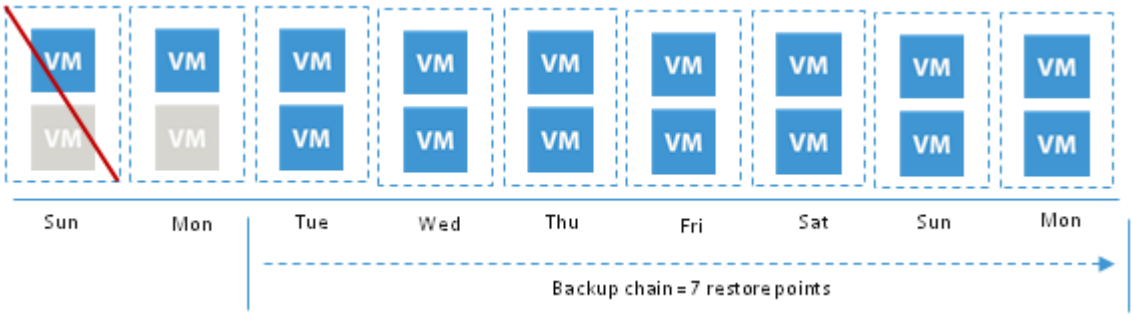

Scheduling

When you create a job, you can simply start it manually whenever it is necessary. However, as the number of backup and replication jobs increases, it may become hard to keep track of them. Veeam Backup & Replication provides a number of job scheduling options which enables you to set up automatic startup schedule for jobs, automatic retries for failed jobs, and a backup window to limit the time when jobs are performed.

Automatic Startup Schedule

To perform a job on a regular basis, you can schedule it to start automatically. The Veeam Backup Service running on the backup server continuously checks configuration settings of jobs and starts them in accordance with their schedules.

Jobs can also be scheduled to run continuously, that is, in a non-stop manner. Technically, a job running continuously is launched as soon as previous job processing is complete. With Veeam Backup & Replication, you can run jobs continuously or with an interval as low as one minute to implement near-continuous data protection (near-CDP) for the most critical applications and workloads.

Note Even if you have scheduling set up for a job, you can still start it manually at any moment.

Automatic Job Retry

Veeam Backup & Replication can be configured to retry a job for a certain number of times if the initial job pass fails. By default, Veeam Backup & Replication automatically retries a failed job for three times within one job session. If necessary, however, you can change the number of retries in the job settings.

Veeam Backup & Replication retries a job only if the previous job session has failed and one or several VMs in the job has not been processed. Veeam Backup & Replication does not perform a retry if a job session has finished with the *Success* or *Warning* status. During the job retry, Veeam Backup & Replication processes only those VMs that have failed.

Veeam Backup & Replication creates only one backup file within one job session. That is, if a job includes several VMs and some of them fail to be processed during the first job pass, Veeam Backup & Replication will create a backup file containing data for those VMs that have been successfully processed. At the job retry, Veeam Backup & Replication will attempt to process failed VMs; in case of success, Veeam Backup & Replication will write data of the processed VMs to the backup file that was created at the previous job pass.

In some situations, Veeam Backup & Replication may fail to process VMs during all job retries. In this case, failed VMs will be processed within the next job session; its data will be written to the backup file created within the current job session.

For example, you have configured a job for two VMs: *VM1* and *VM2*. The job uses the forward incremental method.

During the first job session, Veeam Backup & Replication has successfully processed *VM1* only and created a full backup file for it. *VM2* has failed to be processed during all three job retries. In this case, Veeam Backup & Replication will attempt to process the failed *VM2* within the next job session. Data for *VM2* will be written to the backup file created within this job session, which will be an incremental backup. As a result, at the end of the second backup job session you will have two files:

- Full backup file containing a full restore point for *VM1*
- Incremental backup file containing a full restore point for *VM2* and an incremental restore point for *VM1*

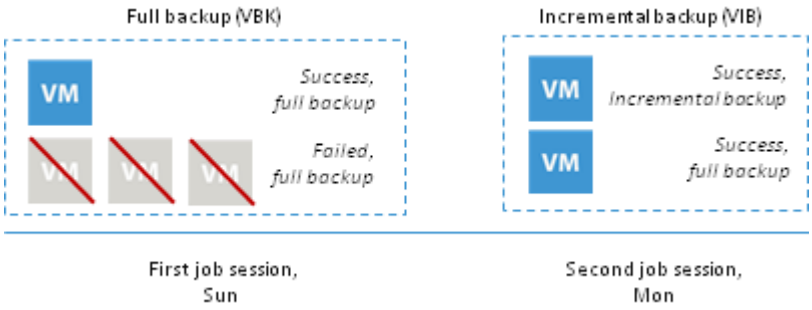

Backup Window

To prevent a backup or replication job from overlapping with production hours and ensure it does not provide unwanted overhead on your virtual environment, you can limit all jobs to a specific backup window. A backup window is a period of time on week days when backup and replication jobs are permitted to run. If the job exceeds the allowed window, it will be automatically terminated.

Backup Content

When creating a backup, replication or copy job, you can select to process separate VMs or VM containers.

Alongside with a general case of backing up a VM or VM container as a whole, Veeam Backup & Replication allows you to determine the content of the created backup by including or excluding specific VM disks and VM templates.

In some situations it may be necessary to back up only specific VM disks. For example, you may want to back up only the system disk instead of creating a full backup which would take much more space than you actually require. Veeam Backup & Replication provides the following options for disks selection:

- Back up all VM disks (selected by default)
- Back up the 0:0 disks (which are commonly the system disks of VMs)
- Back up custom disks at your discretion

Disk processing settings are specified granularly for every VM in the job.

When creating a job, you can select to include VM templates into the created backup. Backing up VM templates warranties supplementary safety of your production environment, though demands additional space. Veeam Backup & Replication allows you to include a VM template only in the full backup and omit it in all subsequent increments.

While processing VM data, Veeam Backup & Replication consolidates the content of virtual disks to present data in the same manner as it is seen by the guest OS. As part of this process, Veeam Backup & Replication filters out overlapping blocks of snapshots, blocks of swap files and zero-data blocks. To reduce the size of backups, Veeam Backup & Replication also excludes VM log files from processing.

Changed Block Tracking

To perform incremental backup, Veeam Backup & Replication needs to know which data blocks have changed since the previous job run.

For VMware VMs with hardware version 7 or later, Veeam Backup & Replication employs VMware vSphere Changed Block Tracking (CBT) — a native VMware feature. Instead of scanning VMFS, Veeam Backup & Replication queries CBT on vSphere through VADP and gets the list of blocks that have changed since the last run of this particular job. Use of CBT increases the speed and efficiency of block–level incremental backups. CBT is enabled by default; if necessary, you can disable it in the settings of a specific backup job.

In some situations, Veeam Backup & Replication cannot leverage VMware vSphere CBT due to VMware limitations. Whenever Veeam Backup & Replication cannot leverage VMware vSphere CBT (for example, if your VMs run an earlier version of virtual hardware or CBT is disabled at the ESX host level), it fails over to Veeam's proprietary filtering mechanism. Instead of tracking changed blocks of data, Veeam Backup & Replication filters out unchanged data blocks. During backup, Veeam Backup & Replication consolidates virtual disk content, scans through the VM image and calculates a checksum for every data block. Checksums are stored as metadata to backup files next to VM data. When incremental backup is run, Veeam Backup & Replication opens all backup files in the chain of previous full and incremental backups, reads metadata from these files and compares it with checksums calculated for a VM in its current state. If a match is found (which means the block already exists in the backup), the corresponding block is filtered out.

Compression and Deduplication

To decrease traffic and disk space required for storing backup files, Veeam Backup & Replication provides mechanisms of compression and deduplication.

Compression

Compression decreases the size of created backups but affects duration of the backup procedure. With the new data compression algorithm used in 7.0, Veeam Backup & Replication allows you to select one of the following compression levels:

- **None** compression level is recommended if you use storage devices with hardware compression and deduplication tools to store created backups.
- **Dedupe-friendly** is an optimized compression level for very low CPU usage. It is recommended if you want to decrease the proxy load.
- **Optimal** (default setting) is the recommended compression level providing the best ratio between the size of the backup file and time of the backup procedure.
- **High** compression level provides additional 10% compression ratio over **Optimal**, but at the cost of about 10x higher CPU usage.
- **Extreme** compression provides the smallest size of the backup file but reduces backup performance. We recommend that you run backup proxies on computers with modern multi-core CPUs (6 cores recommended) if you intend to use the extreme compression.

Note If you are upgrading to 7.0 from the previous version of Veeam Backup & Replication, consider that compression level you have previously set for your jobs will be preserved (compression level named **Optimal** in version 6.5 will appear as **High** in the user interface after the upgrade). However, all the newly created jobs will have the new compression level in effect.

Deduplication

You can apply deduplication when backing up multiple VMs that have similar data blocks (for example, if VMs were created from the same template) or great amount of free space on their logical disks. Veeam Backup & Replication does not store zero byte blocks or space that has been pre-allocated but not used. With deduplication, identical blocks or blocks of free space are eliminated, which decreases the size of the created backup file.

Depending on the type of storage you select as a backup target, Veeam Backup & Replication uses data blocks of different size to process VMs, which optimizes the size of a backup file and job performance. You can choose one of the following storage optimization options:

- The **Local target (16 TB + backup size)** option is recommended for backup jobs that can produce very large full backup files — larger than 16 TB.
  If you select to use data blocks of small size to dedupicate a large backup file, the backup file will be cut into a great number of data blocks. As a result, Veeam Backup & Replication will produce a very large deduplication metadata table which can potentially overgrow memory and CPU resources of your backup repository. For backup files over 16 TB, it is recommended to choose the **Local target (16 TB + backup size)** option. With this option selected, Veeam Backup & Replication will use data blocks of 8 MB. Large data blocks produce a smaller metadata table that requires less memory and CPU resources to process. Note, however, that this storage optimization option will provide the lowest deduplication ratio and the largest size of incremental backup files.
- The **Local target** option is recommended for backup to SAN, DAS or local storage. The SAN identifies larger blocks of data (1024 KB) and therefore can process large amounts of data at a time. This option provides the fastest backup job performance but reduces the deduplication ratio, because with larger data blocks it is less likely to find identical blocks.
- The **LAN target** option is recommended for backup to NAS and onsite backup. It provides a better deduplication ratio and reduces the size of a backup file because of reduced data block sizes (512 KB).
- The **WAN target** option is recommended if you are planning to use WAN for offsite backup. Veeam Backup & Replication will use small data blocks (256 KB), which will result in the maximum deduplication ratio and the smallest size of backup files, allowing you to reduce the amount of traffic over the WAN connection.

Note Changing the compression level and deduplication settings in an existing job will not have any effect on previously created backup files. It will affect only those backups that will be created after you set the new settings:

- New compression settings are applied at the next run of the job.
- New deduplication settings are applied only after a new active full backup is created.

Transaction Consistency

When you perform backup of a running VM, it is necessary to quiesce (or 'freeze') it to bring the file system and application data to a consistent state suitable for backup. Backing up a VM without quiescence produces a crash-consistent backup. Restoring a crash-consistent backup is essentially equivalent to rebooting a server after a hard reset. In contrast to it, restoring transactionally consistent backups (produced with VM data quiesced) ensures safety of data for applications running on VMs.

To create transactionally consistent backup images of VMware vSphere VMs,
Veeam Backup & Replication provides two options: application-aware image processing (utilizing Windows VSS) and VMware Tools quiescence.

Note When you select both options for a job at the same time, only the VSS option will be used for processing VMs. However, if you choose both of these two options and additionally select the **Ignore application processing failures** option for the job, Veeam Backup & Replication will attempt to process VMs with VSS first and if it fails or is not available (for example, in case of Linux VMs) it will use VMware Tools quiescence. This can be very useful when you have both Windows and Linux VMs in one job, so all VMs will be processed in a transactionally consistent manner by using either VSS or VMware Tools quiescence.

Application-Aware Image Processing

To create a transactionally consistent backup of a VM running VSS−aware applications (such as Active Directory, Microsoft SQL, Microsoft Exchange, Sharepoint) without shutting them down, Veeam Backup & Replication uses application−aware image processing. It is Veeam's proprietary technology that ensures successful VM recovery, as well as proper recovery of all applications installed on the VM without any data loss.

Veeam Backup & Replication does not deploy persistent agents inside VMs. Instead, it uses a runtime coordination process on every VM that is started once the backup operation is launched, and removed as soon as it is finished. This helps avoid agent−related drawbacks such as pre−installing, troubleshooting and updating.

To trigger VSS freeze, Veeam Backup & Replication finds out if there is a VSS−aware application running inside a VM. It then requests Microsoft VSS to create a consistent and reliable view of application data prior to taking a VM snapshot. Windows VSS interfaces with VSS−aware applications and Windows OS to quiesce all I/O at a specific point in time. This way, it ensures that there are no unfinished database transactions or incomplete application files during data copying operations.

Tip To find out which VSS−aware applications are installed on a VM, open the Windows command prompt and run the following command:

```
vssadmin list writers
```

The command will bring up a list of all VSS writers installed and registered in the Microsoft VSS framework.

As part of application−aware image processing, Veeam Backup & Replication also applies application−specific settings to prepare every application for VSS−aware restore at the next VM startup. And finally, if backup is successful, it performs transaction logs pruning for specific applications.

Microsoft Windows VSS integration is supported for the following OSes:

- Microsoft Windows XP (32-bit only)
- Microsoft Windows 2003
- Microsoft Windows Vista
- Microsoft Windows 2008
- Microsoft Windows 2008 R2
- Microsoft Windows 7
- Microsoft Windows 8
- Microsoft Windows 2012

Backup and replication with application-aware image processing enabled requires that your guest OS has VMware Tools and all the latest packs and patches installed.

Tip To learn more about Windows VSS and what it is used for, refer to the article VMware and Microsoft VSS: What You Need to Know by Greg Shields.

VMware Tools Quiescence

When taking snapshots of a running VM, VMware Tools will quiesce the VM file systems to ensure integrity of on-disk data. For ESX versions prior to 4.1 (actually, 3.5 U2), this was achieved by using a SYNC driver that held incoming I/O and flushed all dirty data to disk. However, under heavy I/O load, this delay in I/O could become too long. This could be an issue for highly-transactional applications (Exchange Server, SQL Server, and others), because if writes from these applications get delayed for too long, the applications may stop responding.

Since vSphere 4.1, for creating quiesced snapshots of the VM volumes, VMware Tools has supported VSS application consistency. In order to use the VSS component of VMware Tools, the virtual machines should have the following guest operating systems:

- Windows Server 2003 32-bit/64-bit
- Windows Vista 32-bit/64-bit
- Windows 7 32-bit/64-bit
- Windows Server 2008 32-bit/64-bit
- Windows Server 2012

Also, consider that supported quiescence features differ depending on the guest OS:

- For virtual machines running Windows Server 2003 as the guest OS, the VSS component uses application VSS writers, making sure that the VSS snapshots are application-consistent.
- For virtual machines running Windows Vista and Windows 7, the VSS component does not use application writers and, as a result, the snapshots are file-system consistent.
- For virtual machines running Windows Server 2008 or Windows Server 2012, the VSS component may or may not use application writers, based on specific criteria (as defined in the VMware referenced document below). Appropriately, the snapshots will be either file-system or application-consistent.

For details, see VMware documentation: http://pubs.vmware.com/vsphere-51/topic/com.vmware.vddk.pg.doc/vddkBkupVadp.9.6.html

The **Enable VMware Tools quiescence** job option enables freezing of the file system for proper snapshot creation, using VMware VSS component. It can be accessed on the **vSphere** tab of the **Advanced Settings** dialog after you click **Advanced** at the **Storage** step of the backup job wizard. By default, this option is disabled.

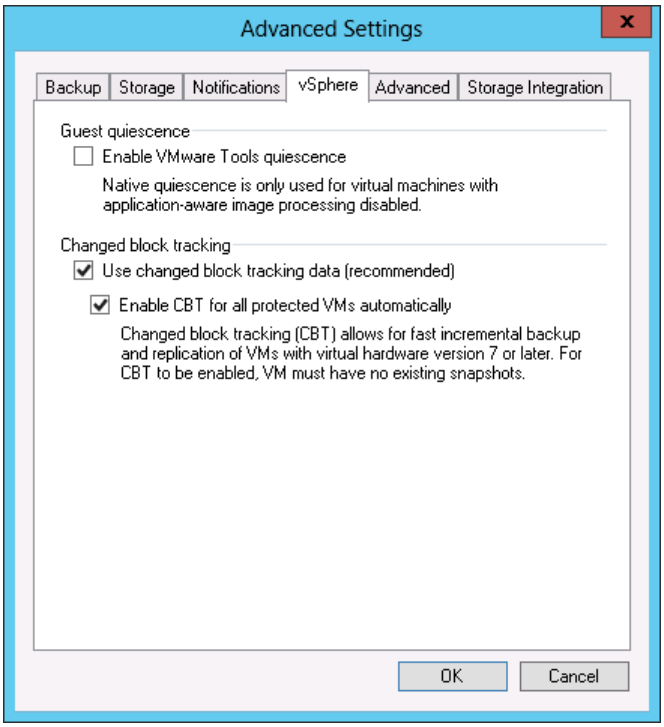

If you select the **Enable VMware Tools quiescence** option instead of the **Application-aware image processing**, VSS will perform application-level VSS quiescence (for example, for the Windows 2008 VMs) without any application-specific steps required for VSS backup and VSS restore.

Important! VMware Tools VSS does not support log truncation recommended for highly-transactional applications like Exchange Server or SQL Server.

That is why **Application-aware image processing** is the recommended option to use for backup and replication of Exchange Server, Active Directory and other VSS-aware applications. However, if for some reason, you cannot leverage this feature, select **Enable VMware Tools quiescence** to place the applications into a consistent state for the snapshot.

If you want to enable both features together at the same time within one job (in case of you need to back up different OS, Windows/Linux), you should do the following:

1. At the **Storage** step of the job wizard, click **Advanced** and select **Enable VMware Tools quiescence** on the **vSphere** tab.
2. At the **Guest Processing** step, select the **Enable Application-aware image processing** check box:

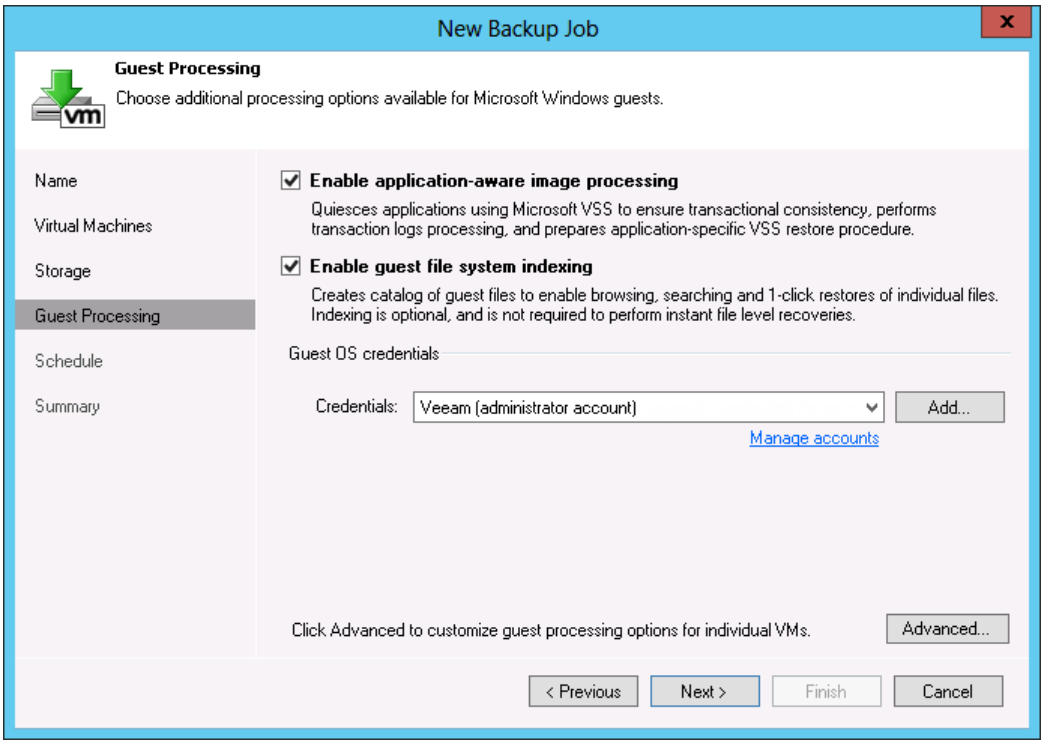

3. When configuring advanced option for individual VM, select **Ignore application processing failures**:

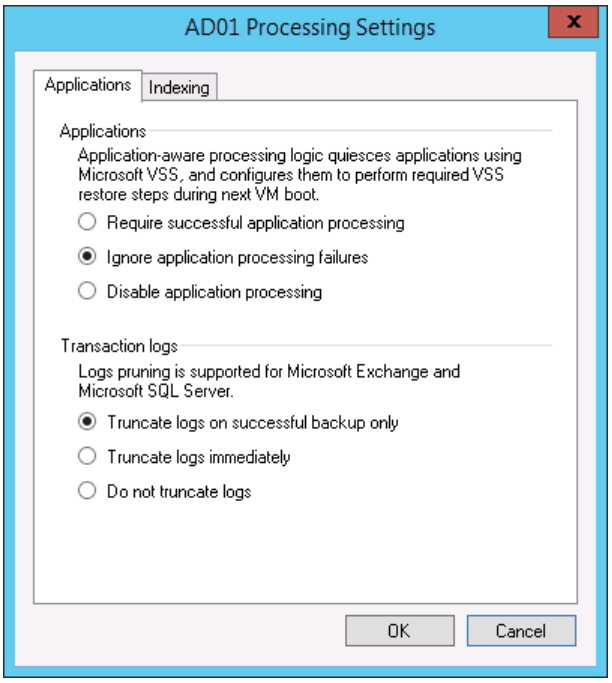

This combination of settings will enable VMware quiescence on all VMs that failed to use native Windows VSS (probably, it would be Linux VMs). So, all VMs will be processed in a transactionally consistent manner by using VSS or VMware Tools quiescence.

Note If you enable both Windows VSS and VMware Tools quiescence but do not select the **Ignore application processing failures** option, then Windows VSS will only be used for processing VMs.

Truncation of Transaction Logs

If you are performing backup or replication of database systems that use transaction logs (for example, Microsoft Exchange or Microsoft SQL), you can select to truncate transaction logs after the job so that they do not overflow the storage space. Veeam Backup & Replication provides advanced options of transaction logs handling for different backup scenarios.

- You can choose to truncate transaction logs after any VM backup to save disk on storage.
- You can choose to truncate logs after successful VM backup only. With this option selected, Veeam Backup & Replication will behave in the following way:
    - If the job fails, Veeam Backup & Replication will not truncate transaction logs. In case a problem occurs and you need to recover the database to some point in time in the past in this scenario, you can apply the transaction logs to the database and so get it to the necessary point in time.
    - If the job is successfully completed, Veeam Backup & Replication will truncate transaction logs. In case a problem occurs and you need to recover the database to some point in time in the past in this scenario, you can restore the database from the successful backup, get transaction logs from this backup, apply them to a restored database so get it to the necessary point in time.
- You can choose not to truncate transaction logs at all. This option is recommended if you are using another backup tool along with Veeam Backup & Replication to perform guest−level backup and this tool maintains consistency of the database state. In this case, truncation of logs with Veeam Backup & Replication will break the guest−level backup chain and cause it to fall out of sync.

VeeamZIP

With Veeam Backup & Replication, you can quickly perform backup of one or several VMs with VeeamZIP.

VeeamZIP is similar to full VM backup. The VeeamZIP job always produces a full backup file (.vbk) that acts as an independent restore point. You can store the backup file to a backup repository, to a local folder on the Veeam backup server or to a network share.

When you perform backup with VeeamZIP, you do not have to configure a backup job and schedule it. Instead, you can start the backup process for selected VMs immediately. This type of backup requires minimum settings — you should only select the backup destination, choose the necessary compression level and enable or disable application-aware image processing if necessary.

The VeeamZIP job is not registered in the database used by Veeam Backup & Replication and the backup file produced with it is not available under the **Backups** node in the **Backup & Replication** view. To be able to restore data from such file, you will need to import it to Veeam Backup & Replication. For import, you can simply double-click the necessary backup file on the machine where Veeam Backup & Replication is installed.

SureBackup Recovery Verification

To guarantee recoverability of your data, Veeam Backup & Replication offers the SureBackup technology.

SureBackup is intended to automate and simplify the backup verification process — one of the most crucial parts of data management and protection. SureBackup lets validate backups of your VMs without impacting the production infrastructure. You can automatically verify every created restore point of every VM and ensure that they will function as expected in case a disaster strikes.

Note SureBackup, or automatic recovery verification, is available in *Enterprise* and *Enterprise Plus Editions* of Veeam Backup & Replication. If you use the *Standard Edition*, you can manually verify VM backups with Instant VM Recovery. To learn more, see Manual Recovery Verification.

How It Works

SureBackup is Veeam's technology that lets you test a VM backup and ensure you will be able to recover data from it. To validate a VM backup, Veeam Backup & Replication performs its "live" verification: it automatically boots the VM from the backup in the isolated environment, performs tests against it, then powers the VM off and creates a report on the VM backup state. You can verify a VM from the latest backup or from any necessary restore point.

SureBackup recovery verification uses a regular image-based backup created with Veeam Backup & Replication. The procedure of the VM verification is the following:

1. Veeam Backup & Replication leverages the vPower technology to publish a VM in the isolated virtual environment. Veeam Backup & Replication runs VMs directly from backup files without restoring them to the production datastore. To achieve this, Veeam Backup & Replication utilizes the Veeam vPower NFS Service. This presents VM images on the backup storage as an NFS datastore to an ESX(i) host so a VM backup is seen as a regular VM image.
2. Veeam Backup & Replication performs a number of tests against the verified VM.
3. If a SureBackup job is configured to validate backup files, Veeam Backup & Replication performs backup file validation for verified VMs and optionally for VMs in the application group. Backup file validation is performed after all verification tests for all VMs in the SureBackup job are completed.
4. When the recovery verification process is over, Veeam Backup & Replication unpublishes the VM and creates a report on its state. The report is sent to the backup administrator by email.

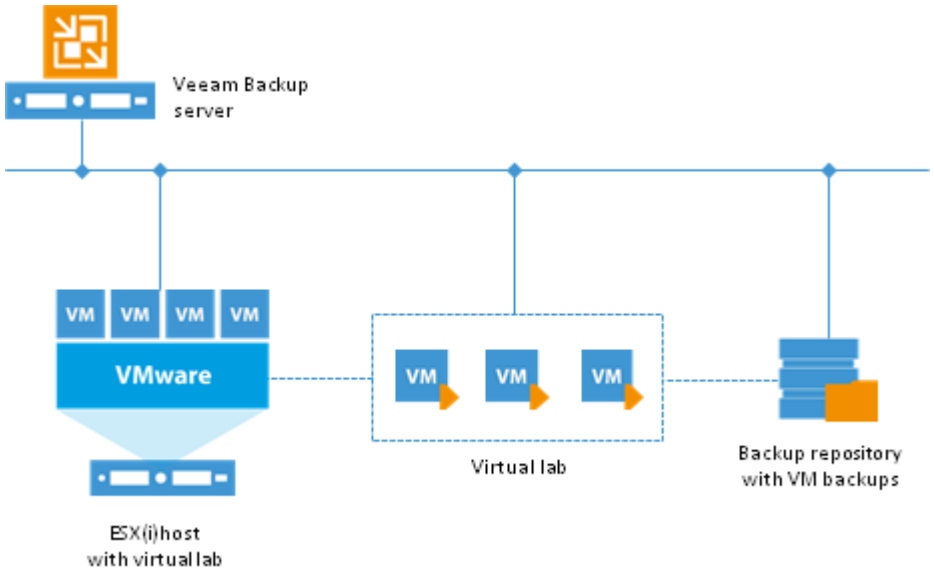

During verification, a backed up VM image remains in read−only state. All changes that take place when a VM is running are written to redo log files that are stored on a selected datastore in the production environment. Once the recovery verification process is complete, the redo logs are removed.

To perform VM verification, you need to create the following entities:

1. Application group. During recovery verification, the verified VM is not started alone: it is started together with VMs on which it is dependent. Starting the verified VM in conjunction with other VMs enables full functionality of applications running inside the VM and lets you run these applications just like in the production environment.
2. Virtual lab. SureBackup leverages the virtual lab technology to verify a VM backup. The virtual lab is the isolated virtual environment in which the verified VM and VMs from the application group are started and tested.
3. SureBackup job. The SureBackup job is a task to run the recovery verification process. You can run the SureBackup job manually or schedule it to run automatically according to some schedule.

Veeam vPower NFS Service

The Veeam vPower NFS Service is a Windows service that runs on a Windows backup repository server and enables it to act as an NFS server. vPower NFS allows Veeam Backup & Replication to mount a compressed and deduplicated backup file as a regular VMDK file directly to the ESX(i) host via NFS, so ESX(i) hosts get transparent access to backed up VM images.

If you store backups on a Windows repository, it is highly recommended to enable the vPower NFS Server on this repository. In this case, the vPower NFS Service will run on the managing Windows server.

Besides Windows−based backup repository servers, Veeam vPower NFS Service can run on any Windows server you choose, including the Veeam backup server itself. However, in this case performance may be much lower, because instead of a direct connection between the ESX(i) host and the backup repository, the connection will be split into two parts: ESX(i) host to NFS server and NFS server to backup repository.

The vPower technology is used to enable the following features:

- SureBackup: Recovery Verification
- Instant VM Recovery
- Multi-OS File-Level Recovery
- Universal Application-Item Recovery (U-AIR)

vPower-Specific Settings

To be able to successfully connect an ESX(i) host to the NFS server, you should make sure that the ESX(i) host has a proper network interface configuration and can access the server on which Veeam vPower NFS Service is running.

When connecting to the NFS server, the ESX(i) host uses a VMkernel interface. That is why the ESX(i) host you are using must have a VMkernel interface. Otherwise, vPower NFS mounting on ESX(i) host will fail.

By default, VMkernel interfaces are not available for non−ESXi versions, so you will have to add them on the ESX host to be able to connect to the NFS server.

- If the NFS server and ESX host are located in the same subnet, the ESX host should have a VMkernel interface in the same IP network as the NFS server.
- If the NFS server and ESX host are located in different subnets and use a router for network access, in addition to creating a new VMkernel interface, you will have to manually specify routing settings in the IP routing table on the ESX host.

Tip To check whether an ESX host can access the NFS server or not, you can use the vmkping utility on the ESX host. The vmkping utility is similar to the ping tool; the only difference is that ICMP packets are sent via the VMkernel interface rather than the console interface.

Recovery Verification Tests

To verify a VM started in the virtual lab, you can run Veeam's predefined tests or perform your own tests against VMs. The predefined tests include the following ones:

1. **Heartbeat test**. As soon as the VM is started, Veeam Backup & Replication performs a heartbeat test. It waits for a heartbeat signal from VMware Tools installed inside the VM to determine that the guest OS inside the VM is running. If the signal comes regularly at specific time intervals, the test is passed.
2. **Ping test**. During the ping test, Veeam Backup & Replication checks if the VM in the virtual lab can respond to the ping requests. If VM responds to ping requests from the Veeam backup server, the test is passed.
3. **Application test**. Veeam Backup & Replication waits for applications to start inside the VM and runs a script that checks application-specific network ports. For example, to verify a Domain Controller, Veeam Backup & Replication probes port 389 for a response. If the response is received, the test is passed.

Beside these predefined tests, you can use custom scripts to verify the VM backup.

Note To run the heartbeat and ping tests, you must have VMware Tools installed inside the VM you start from the backup. Otherwise these tests will be skipped; Veeam Backup & Replication will display a warning in the SureBackup job session results.

Backup File Validation

In addition to recovery verification tests, Veeam Backup & Replication allows you to perform backup file validation — a CRC check that runs for backup files of VMs verified by the SureBackup job. You can also optionally validate backup files for VMs from the application group.

For validation of a backup file, Veeam Backup & Replication uses the checksum algorithm. When Veeam Backup & Replication creates a backup file for a VM, it calculates a checksum for every data block in the backup file and stores this data together with VM data. During the backup file validation test, Veeam Backup & Replication de-compresses the backup file, re-calculates checksums for data blocks in the uncompressed backup file and compares them with initial checksum values. If the results match, the test is passed.

The backup file validation test is started after SureBackup recovery verification tests. As soon as Veeam Backup & Replication completes all "live" verification tests for all VMs in the SureBackup job, it powers off VMs in the virtual lab and starts the backup file validation test for these VMs and for VMs in the application group (provided you have chosen to validate backup files for the application group).

The result of the backup file validation test impacts the state of the SureBackup job session. If the validation tests are completed successfully but the backup validation is not passed, Veeam Backup & Replication marks the SureBackup job session with the *Warning* or *Error* status.

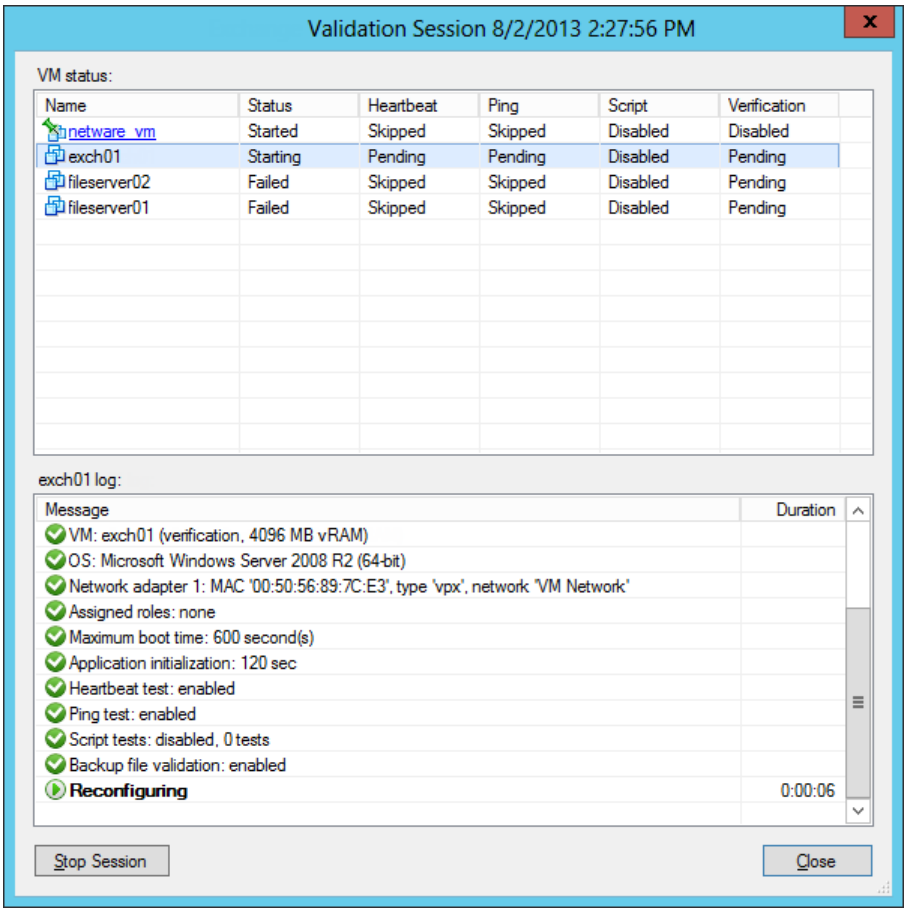

Application Group

In most cases, a VM works not alone but in cooperation with other services and components. To verify such VM, you first need to start all services and components on which the VM is dependent. To this aim, Veeam Backup & Replication uses the notion of application group.

The application group creates the “surroundings” for the verified VM. The application group contains one or several VMs on which the verified VM is dependent. These VMs run applications and services that must be started to enable fully functional work of the verified VM. Typically, the application group contains at least a domain controller, DNS server and DHCP server.

When you set up an application group, you specify a role of every VM, its boot priority and boot delay. Additionally, you specify what tests must be performed for VMs in the application group.

When a SureBackup job is launched, Veeam Backup & Replication first starts in the virtual lab VMs from the application group in the required order and performs necessary tests against them. This way, Veeam Backup & Replication creates the necessary environment to start the verified VM. Only after all VMs from the application group are started and tested, Veeam Backup & Replication starts the verified VM in the virtual lab.

For example, if you want to verify a Microsoft Exchange Server, you need to test its functionality in cooperation with other components: domain controller and DNS server. Subsequently, you must add to the application group a virtualized domain controller and DNS server. When Veeam Backup & Replication runs a SureBackup job, it will first start and verify the domain controller and DNS server in the virtual lab to make verification of the Exchange Server possible.

Note All VMs added to the application group must belong to the same platform — VMware or Hyper-V. Mixed application groups are not supported.

Virtual Lab

The virtual lab is an isolated virtual environment in which Veeam Backup & Replication verifies VMs. In the virtual lab, Veeam Backup & Replication starts a verified VM and VMs from the application group. The virtual lab is used not only for the SureBackup verification procedure, but also for U-AIR and On-Demand Sandbox processing.

A virtual lab does not require provisioning of additional resources. You can deploy it on the existing ESX(i) host in your virtual environment.

The virtual lab is fully fenced off from the production environment. The network configuration in the virtual lab mirrors the network configuration of the production environment. For example, if verified VMs are located in two logical networks in your production environment, the virtual lab will also have two networks. The networks in the virtual lab will be mapped to corresponding production networks.

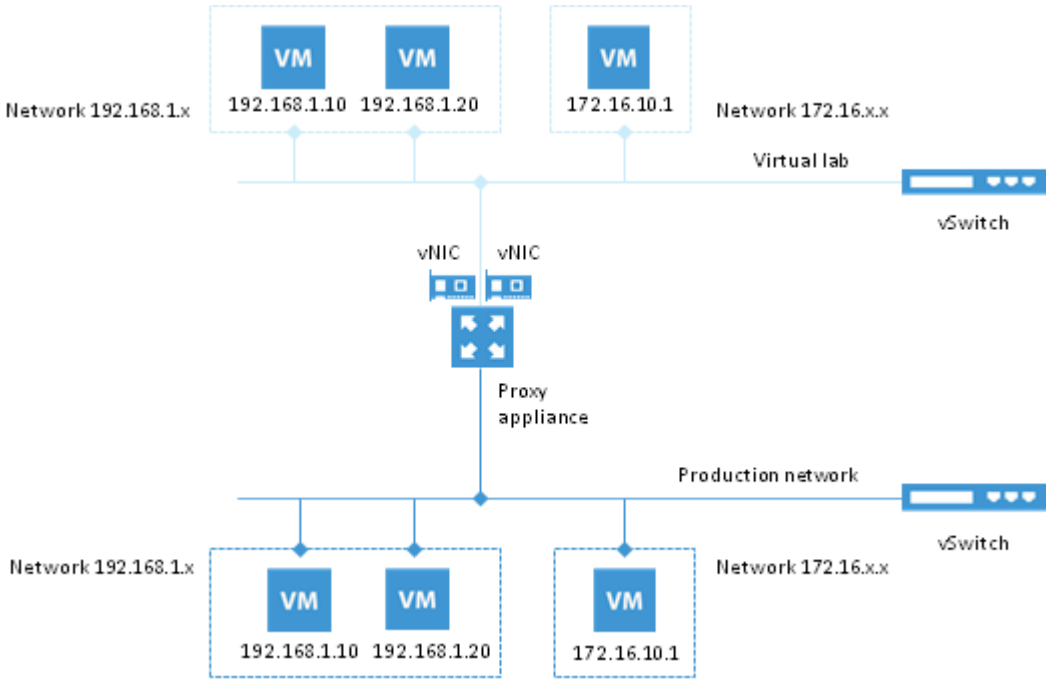

Tip You can optionally connect VMs to the same network in the virtual lab, even if corresponding VMs in the production environment are connected to different networks.

VMs in isolated networks have the same IP addresses as in the production network. This lets VMs in the virtual lab function just as if they would function in the production environment.

Proxy Appliance

To enable communication between the production environment and the isolated networks in the virtual lab, Veeam Backup & Replication uses a proxy appliance. The proxy appliance is a Linux-based auxiliary VM created on the ESX(i) host where the virtual lab is created. The proxy appliance VM is assigned an IP address from the production network and placed to the dedicated virtual lab folder and resource pool created on the ESX(i) host.

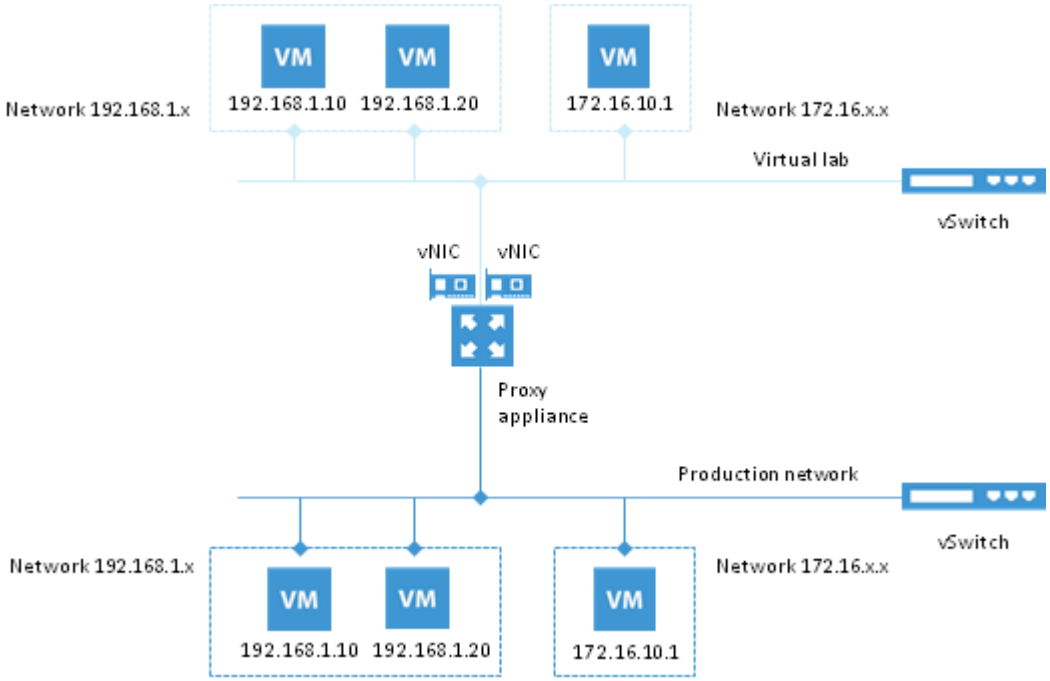

The proxy appliance is connected to the production network and to the isolated network and so has visibility of the production environment and the virtual lab. In essence, the proxy appliance acts as a gateway between the two networks, routing requests from the production environment to VM replicas in the virtual lab.

The proxy appliance connects to isolated networks using network adapters.
Veeam Backup & Replication adds to the proxy appliance one network adapter per each isolated network. For example, if there are two networks in the virtual lab, Veeam Backup & Replication will add two network adapters to the proxy appliance. The network adapter gets an IP address from the isolated network. Typically, this IP address is the same as the IP address of the default gateway in the corresponding production network.

Note The proxy appliance is an optional component. Technically, you can create a virtual lab without a proxy appliance. However, in this case, you will not be able to perform automatic recovery verification of VMs. VMs will be simply started from backups in the virtual lab; you will have to access them using the VM console and perform necessary tests manually.

IP Masquerading

To let the traffic into the virtual lab, Veeam Backup & Replication uses masquerade IP addressing.

Every VM in the virtual lab has a masquerade IP address, along with the IP address from the production network. The masquerade IP address resembles the IP address in the production network: for example, if the IP address of a VM is 172.16.1.13, the masquerade IP address may be 172.18.1.13.

The masquerade IP address can be thought of as an entry point to the VM in the virtual lab from the production environment. When you want to access a specific VM in the virtual lab, Veeam Backup & Replication addresses it by its masquerade IP address.

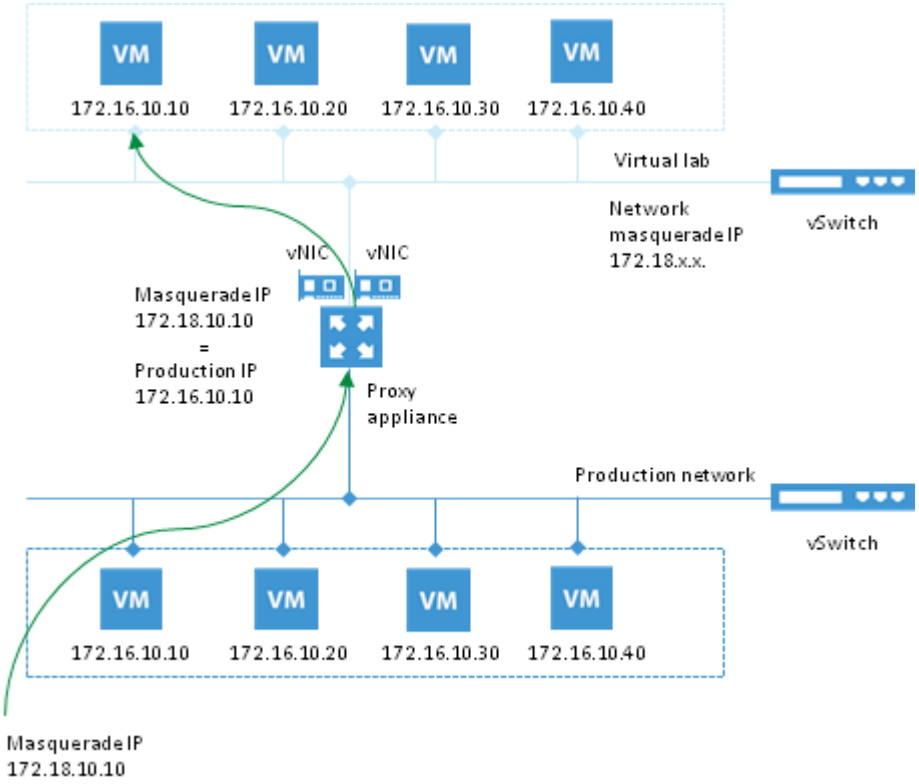

The rules routing requests to VMs in the virtual lab are specified in the routing table on the server from which you want to access VMs in the virtual lab. The routing table can be updated on the following servers:

- **Veeam backup server**. Veeam Backup & Replication automatically creates the necessary static route in the routing table on the Veeam backup server at the moment you launch a SureBackup job and Veeam Backup & Replication starts the virtual lab.
- **Client machine**. If you want to provide your users with access to VMs in the virtual lab, you need to manually update routing tables on their machines and add a new static route. See also: Static IP Mapping.

The added static route destines the masquerade network traffic to the proxy appliance. The proxy appliance here acts as a NAT device: it resolves the masquerade IP address, replaces it with "real" IP address of a VM from the production network and then directs the request to the necessary VM in the virtual lab. The static route is non-persistent: when you power off the virtual lab, the route is removed from the routing table on the Veeam backup server or client machine.

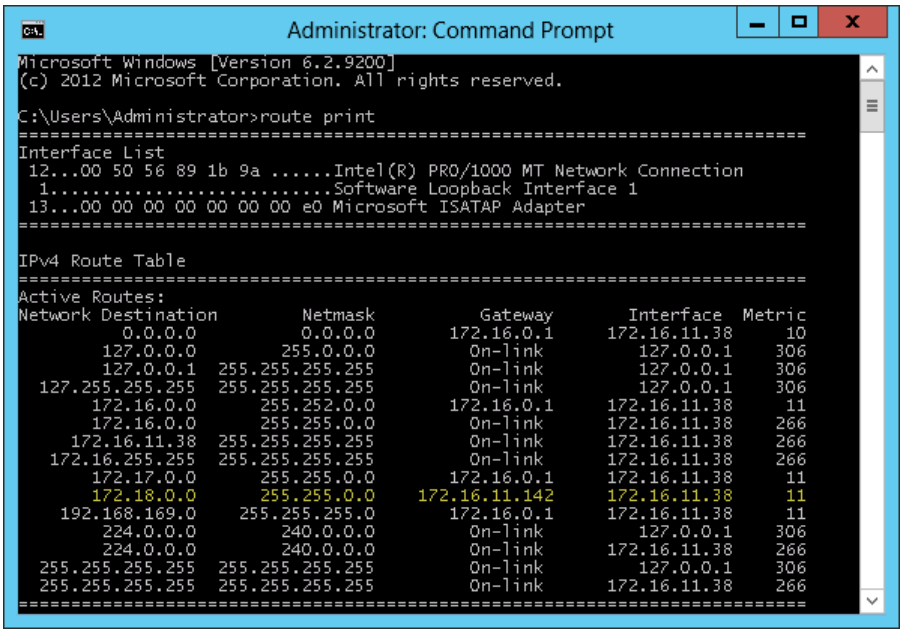

For example, when trying to access a VM with IP address 172.16.10.10 in the isolated network during recovery verification, Veeam Backup & Replication sends a request to the masquerade IP address 172.18.10.10. According to the routing rule added to the IP routing table, all requests are first sent to the next hop — the proxy appliance. The proxy appliance performs address translation, substitutes the masquerade IP address with the IP address in the isolated network and forwards the request to the necessary VM in the isolated network — in the given example, to 172.16.10.10.

Static IP Mapping

Sometimes it is necessary to provide many clients with access to a restored VM, which is especially the case for user-directed application item-level recovery. For example, you may want to provide your users with access to the Exchange Server started in the virtual lab using web−based access (like Outlook Web Access). Technically, you may update the routing table on every client machine; however, this will demand a lot of administrative effort.

For such situations, Veeam Backup & Replication enables you to get access to a VM in the virtual lab directly from the production environment. To be able to access to a VM in the virtual lab, you should reserve a static IP address in the pool of production IP addresses and map this IP address to the IP address of a VM in the virtual lab.

The static IP address is assigned to the proxy appliance network adapter connected to the production network. IP traffic directed to the specified static IP address is routed by the proxy appliance to the VM powered on in the isolated network.

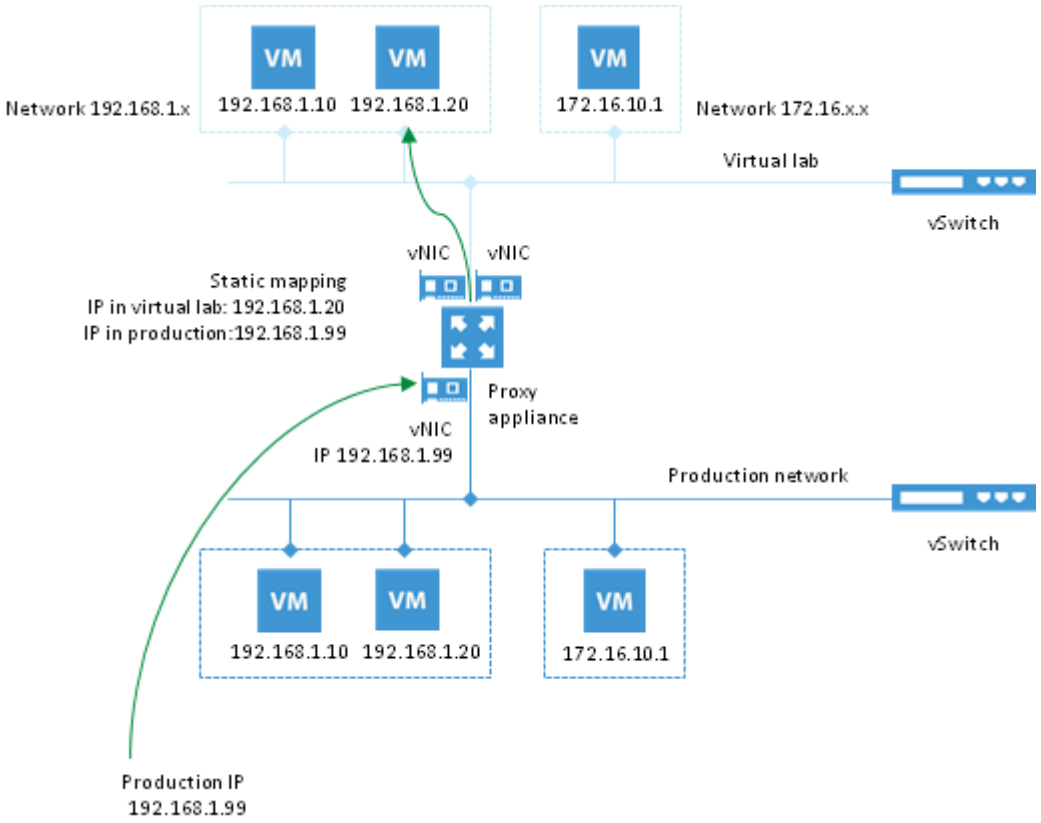

For example, for a VM with IP address 192.168.1.20 in the isolated network, you can reserve IP address 192.168.1.99 (a free IP address from the production network). As a result, you will be able to use IP address 192.168.1.99 to access the VM in the virtual lab from the production side.

You can also register an alias record in the production DNS server for the reserved IP address. For example, you can register backup.exchange.local as an alias for the IP address 192.168.1.99.

Virtual Lab Configuration

For SureBackup recovery verification, Veeam Backup & Replication offers two types of the virtual lab configuration:

- Basic single-host virtual lab
- Advanced single-host virtual lab

Basic Single-Host Virtual Labs

The basic single-host virtual lab (formerly known as the virtual lab with basic networking configuration) should be used if all VMs you want to verify, VMs from the application group and the Veeam backup server are connected to the same network.

For the basic single-host virtual lab, Veeam Backup & Replication creates one virtual network that is mapped to the corresponding production network. Additionally, Veeam Backup & Replication automatically adds a number of new instances on the ESX(i) host where the virtual lab is created:

- A new resource pool
- A new VM folder
- A new standard vSwitch

The vSwitch is only used by the VMs started in the virtual lab: there is no routing outside the virtual lab to other networks.

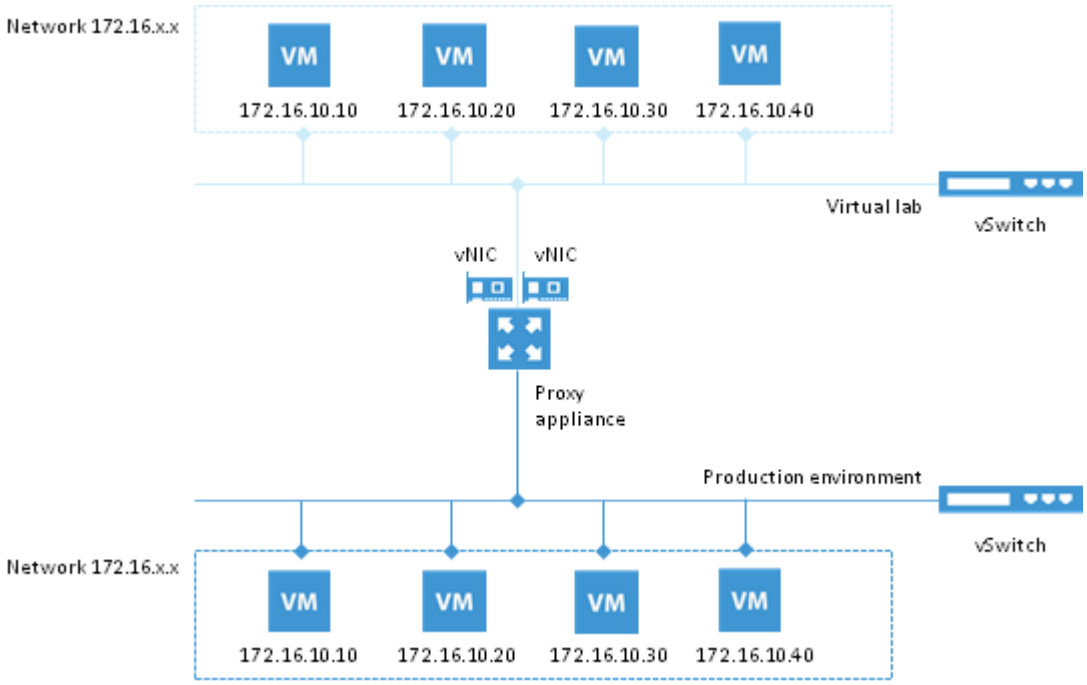

Advanced Single-Host Virtual Labs

The advanced single-host virtual lab (formerly known as the virtual lab with advanced networking configuration) should be used if VMs you want to verify and/or VMs from the application group are connected to different networks.

In the advanced single-host virtual lab, Veeam Backup & Replication creates several virtual networks for the virtual lab. The number of virtual networks corresponds to the number of production networks to which verified VMs are connected. Networks in the virtual lab are mapped to corresponding production networks.

Veeam Backup & Replication automatically adds a number of new VMware instances on the ESX(i) host where the virtual lab is created:

- A new resource pool
- A new VM folder
- A new standard vSwitch

The vSwitch is only used by the VMs started in the virtual lab: there is no routing outside the virtual lab to other networks.

When you create an advanced single-host virtual lab, Veeam Backup & Replication configures basic settings for networks that should be created in the virtual lab. You need to review these settings and manually adjust them if needed.

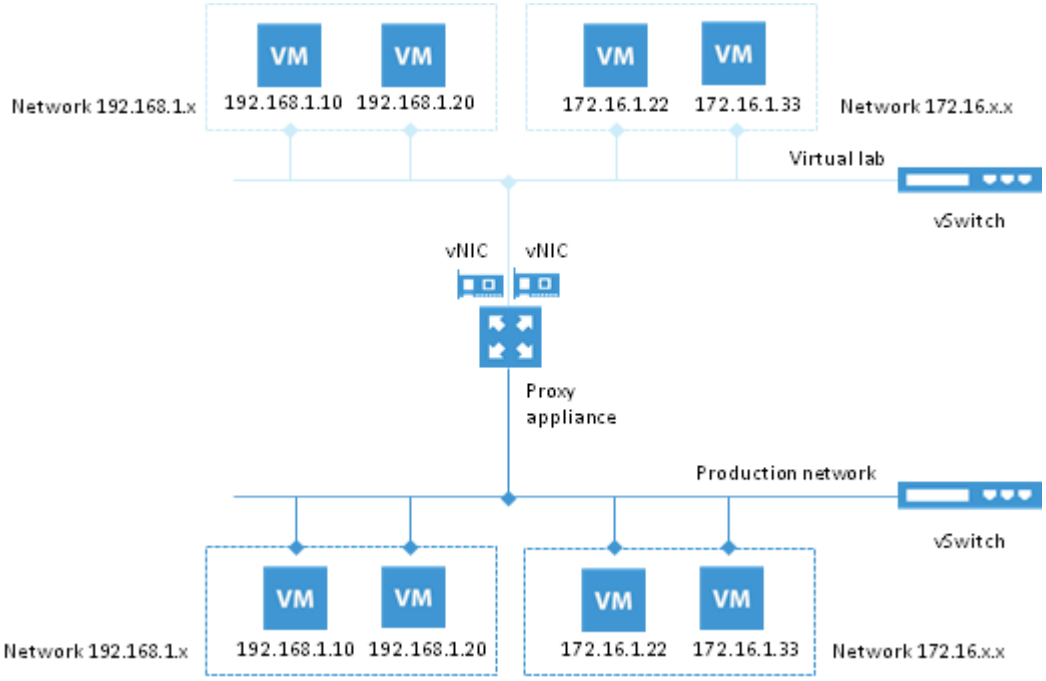

SureBackup Job

A SureBackup job is a task for VM backup recovery verification. The SureBackup job aggregates all settings and policies of a recovery verification task, such as application group and virtual lab to be used, VM backups that should be verified in the virtual lab and so on. The SureBackup job can be run manually or scheduled to be performed automatically.

When a SureBackup job runs, Veeam Backup & Replication first creates an environment for VM backups verification:

1. Veeam Backup & Replication starts the virtual lab.
2. In the virtual lab, it starts VMs from the application group in the required order. VMs from the application group remain running until the verified VMs are booted from backups and tested. If Veeam Backup & Replication does not find a valid restore point for any of VMs from the application group, the SureBackup job will fail.

Once the virtual lab is ready, Veeam Backup & Replication starts verified VMs from the necessary restore point, tests and verifies them one by one or, depending on the specified settings, creates several streams and tests a number of VMs simultaneously. If Veeam Backup & Replication does not find a valid restore point for any of verified VMs, verification of this VM fails, but the job continues to run.

By default, you can start and test up to three VMs at the same time. You can also increase the number of VMs to be started and tested simultaneously. Keep in mind that if these VMs are resource demanding, performance of the SureBackup job as well as performance of the ESX(i) host holding the virtual lab may decrease.

Once the verification process is complete, VMs from the application group are powered off. Optionally, you can leave the VMs from the application group running to perform manual testing or enable user-directed application item-level recovery.

In some cases, the SureBackup job schedule may overlap the schedule of the backup job linked to it. The backup file may be locked by the backup job and the SureBackup job will be unable to verify such backup. In this situation, Veeam Backup & Replication will not start the SureBackup job until the corresponding backup job is over.
To overcome the situation of job overlapping, you may chain the backup and SureBackup jobs or define the timeout period for the SureBackup job. To learn more, see Specifying Job Schedule.

Note VMs from the application group and verified VMs must belong to the same platform — VMware or Hyper-V. Mixed scenarios are not supported.

SureBackup Job Processing

The recovery verification process includes the following steps:

1. **Getting virtual lab configuration**. Veeam Backup & Replication gets information about configuration of the virtual lab where verified VMs should be started.
2. **Starting virtual lab routing engine**. Veeam Backup & Replication starts a proxy appliance used as a gateway to provide access to the virtual lab.
3. **Publishing**. Veeam Backup & Replication creates an NFS datastore with a VM backup and registers it on the selected ESX server. Veeam Backup & Replication does not deploy the whole VM from the backup file, it deploys VM configuration files only. Virtual disks are deployed per force and per required data blocks.
4. **Reconfiguring**. Veeam Backup & Replication updates configuration files for VMs that should be run in the isolated network.
5. **Registering**. Veeam Backup & Replication registers the verified VM on the selected ESX(i) host.
6. **Configuring DC**. If a verified VM has the Domain Controller or Global Catalog role, the VM is reconfigured.
7. **Powering on**. Veeam Backup & Replication powers on the verified VM in the isolated network.
8. **Heartbeat test**. Veeam Backup & Replication checks whether the VMware Tools heartbeat signal (green or yellow) is coming from the VM or not. If the VM has no VMware Tools, the test will not be performed, and a notification will be written to the session details.
9. **Running ping tests**. Veeam Backup & Replication checks if the VM responds to the ping requests or not. If the VM has no NICs and mapped networks for them and/or has no VMware Tools installed, the ping test will not be performed, and a notification will be written to the session details.
10. **Application initialization**. Veeam Backup & Replication waits for the applications installed in the VM (for example, SQL Server, web server, mail server) to start. The application initialization period is defined in the corresponding properties of the SureBackup job, and by default equals to 120 sec. However, depending on the software installed in a VM, the application initialization process may require more time than specified in the SureBackup job settings. If applications installed in a VM are not initialized within the specified period of time, test scripts can be completed with errors. If such an error situation occurs, you will need to increase the **Application initialization timeout** value and start the job once again.

11. **Running test scripts**. Veeam Backup & Replication runs scripts to test whether the application installed in the VM is working correctly or not. If the VM has no VMware Tools installed and/or there are no NICs and mapped networks for them, Veeam Backup & Replication will skip tests that use variables *%vm_ip%* and *%vm_fqdn%* as the IP address and fully qualified domain name of the VM cannot be determined. Test results are written to the session details. To define whether the script has completed successfully or not, Veeam Backup & Replication uses return codes. If the return code is equal to 0, the script is considered to complete successfully. Other values in the return code mean that the script has failed.
12. **Powering off**. After all tests have been performed, Veeam Backup & Replication powers off the verified VM.
13. **Unregistering**. Veeam Backup & Replication unregisters the verified VM on the selected ESX(i) host.
14. **Clearing redo logs**. Veeam Backup & Replication deletes redo logs that were created to store changes made to the VM while it was running from the backup file.
15. **Unpublishing**. Veeam Backup & Replication unpublishes the content of the backup file on the ESX(i) host.

Stabilization Algorithm

To be able to perform tests for a verified VM without errors, Veeam Backup & Replication needs to know that the VM is ready for testing. To determine this, Veeam Backup & Replication waits for the VM to reach a "stabilization point": that is, waits for the VM to boot completely and report it is ready for tests. After the stabilization point has been established, Veeam Backup & Replication can start performing heartbeat tests, ping tests and running test scripts against the VM.

Veeam Backup & Replication establishes the stabilization point with the help of VMware parameters that it gets from the VM. Depending on the VM configuration, it uses one of the three algorithms to do that:

- **Stabilization by IP**. This algorithm is used if the VM has VMware Tools installed, there are NICs and mapped networks for these NICs. In this case, Veeam Backup & Replication waits for an IP address of the VM for mapped networks, which is sent by VMware Tools running in the VM. The sent IP address should be valid and should not change for a specific period of time.
- **Stabilization by heartbeat**. This algorithm is used if the VM has VMware Tools installed but there are no vNICs and mapped networks for them. In this case Veeam Backup & Replication waits for a corresponding heartbeat signal (green or yellow) to come from the VM. As well as in the first case, the signal is sent by VMware Tools running in the VM.
- **Stabilization by Maximum allowed boot time**. This algorithm is used if the VM has neither VMware Tools installed, nor NICs and mapped networks for them. In this case, Veeam Backup & Replication will simply wait for the time specified in the **Maximum allowed boot time** field, which is considered to be a stabilization period for the VM. Once this time interval is exceeded, Veeam Backup & Replication will consider that the VM is successfully booted and is ready for testing.

The stabilization process cannot exceed the value specified in the **Maximum allowed boot** time field. If the stabilization point cannot be determined within the **Maximum allowed boot time**, the recovery verification process will be finished with the timeout error. For this reason, you should be careful when specifying this value — typically, the VM started within the frames of a SureBackup job requires more time to boot if compared to a regular VM startup. When such an error situation occurs, you will need to increase the **Maximum allowed boot time** value and start the job again.
Once the stabilization point has been established, Veeam Backup & Replication runs ping, heartbeat tests and performs test scripts against the verified VM.

Manual Recovery Verification

Beside automatic recovery verification, you can also perform manual verification of VM backups. Manual verification can be performed with all editions of Veeam Backup & Replication.

- To perform a VM boot test, you can go through the **Instant VM Recovery** wizard and power the VM on without connecting it to the production network.
- To perform the application recovery test, you should first create an isolated network. After that, you need to pass through the **Instant VM Recovery** wizard to restore a VM from the backup. At the **Network** step of the wizard, you should connect the VM to the created isolated network. The same procedure should be performed for all VMs that run applications on which a verified VM is dependent, such as domain controller and DNS. All VMs must be connected to the same isolated network and started in the correct order: for example, DNS > domain controller > verified VM.

SureReplica Recovery Verification

To guarantee recoverability of your data, Veeam Backup & Replication complements the SureBackup recovery verification technology with SureReplica.

SureReplica is in many respects similar to the SureBackup recovery verification. It lets you validate your DR environment without impacting the production infrastructure: you can automatically verify every created restore point of every VM replica and ensure that they are functioning as expected.

The SureReplica technology is not limited only to VM replica verification. Just like SureBackup, it provides the following capabilities:

- SureReplica: automated VM replica verification
- On-Demand Sandbox: an isolated environment for testing VM replicas, training and troubleshooting
- U-AIR: recovery of individual items from applications running on VM replicas

How It Works

SureReplica is Veeam's technology that lets you test a VM replica for recoverability. To ensure that the VM replica is functioning properly, Veeam Backup & Replication performs its "live" verification: it automatically boots the VM replica to the necessary restore point in the isolated environment, performs tests against it, powers the VM replica off and creates a report on the VM replica state.

The SureReplica technology does not require the vPower engine. A VM replica is essentially an exact copy of a VM with a set of restore points. The VM replica data is stored in the raw uncompressed format native to VMware. Therefore, to start a VM replica in the virtual lab, you do not need to translate its data via the vPower NFS datastore to the ESX(i) host. Veeam Backup & Replication re-configures the VM replica settings necessary for recovery verification, connects the VM replica to the isolated virtual lab and powers it on.

As there is no need to publish the VM from the backup file, the SureReplica processing is typically faster than SureBackup. Correspondingly, the U-AIR and On-Demand Sandbox operations are faster, too.

The procedure of the VM replica verification is the following:

1. Veeam Backup & Replication triggers a VMware snapshot for a VM replica. The snapshot helps protect the VM replica from changes when it is started and verified. All changes made to the VM replica are written to delta files.
2. Veeam Backup & Replication starts the VM replica in the isolated virtual environment.

3. Veeam Backup & Replication performs a number of tests against the verified VM replica.
4. When the verification process is over, Veeam Backup & Replication removes delta files of the VM replica snapshot, powers off the VM replica and creates a report on its state. The report is sent to the backup administrator by email.

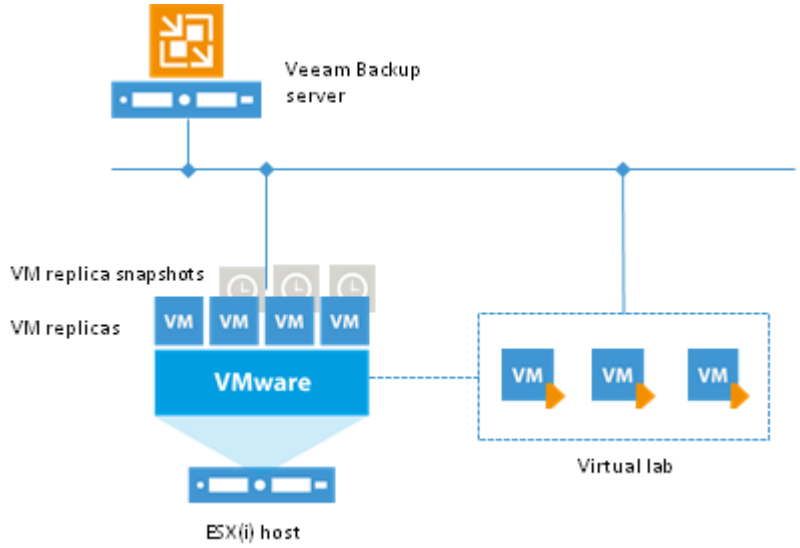

Note Veeam Backup & Replication verifies only VM replicas in the *Normal* state. If a VM replica is in the *Failover* or *Failback* state, the verification process will fail.
When Veeam Backup & Replication verifies the VM replica, it puts the VM replica to the *SureBackup* state. You cannot perform failback and failover operations for a VM replica in the *SureBackup* state until the recovery verification or U-AIR process is over and the VM replica returns to the *Normal* state.

To perform VM replica verification, you need to create the following entities:

1. Application group. During recovery verification, the VM replica is not started alone: it is started together with VMs on which the VM replica is dependent. Starting a VM replica in conjunction with other VMs enables full functionality of applications running inside the VM replica and lets you run these applications just like in the production environment.
2. Virtual lab. Just like SureBackup, SureReplica leverages the virtual lab technology to verify a VM replica. The virtual lab is the isolated virtual environment in which the VM replica and VMs from the application group are started and tested.
3. SureReplica job. The SureReplica job is a task to run the SureReplica verification process. You can run the SureReplica job manually or schedule it to run automatically according to some schedule.

Recovery Verification Tests

To verify a VM replica started in the virtual lab, you can run Veeam's predefined tests or perform your own tests against VMs. The predefined tests include the following ones:

- **Heartbeat test**. As soon as the VM replica is started, Veeam Backup & Replication performs a heartbeat test. It waits for a heartbeat signal from VMware Tools installed inside the VM to determine that the guest OS inside the VM replica is running. If the yellow and green signals come, the test is passed; if the red signal comes, the test is failed.
- **Ping test**. Veeam Backup & Replication sends ping requests to the VM replica started in the virtual lab. If VM replica can respond to ping requests from the Veeam backup server, the test is passed.
- **Application test**. Veeam Backup & Replication waits for applications to start inside the VM replica and runs a script that checks application-specific network ports. For example, to verify a Domain Controller, Veeam Backup & Replication probes port 389 for a response. If the response is received, the test is passed.

Beside these predefined tests, you can use custom scripts to verify the VM replica.

Note To run the heartbeat and ping tests, you must have VMware Tools installed inside the VM replica. Otherwise these tests will be skipped; Veeam Backup & Replication will display a warning in the SureReplica job session results.

Application Group

In most cases, a VM works not alone but in cooperation with other services and components. To verify a replica of such VM, you first need to start all services and components on which the VM replica is dependent. To this aim, Veeam Backup & Replication uses the notion of application group.

The application group creates the "surroundings" for the verified VM replica. The application group contains one or several VMs on which the verified VM replica is dependent. These VMs run applications and services that must be started to enable fully functional work of the verified VM replica. Typically, the application group contains at least a domain controller, DNS server and DHCP server.

When you set up an application group, you specify a role of every VM, its boot priority and boot delay. Additionally, you specify what tests must be performed for VMs in the application group.

When a SureReplica job is launched, Veeam Backup & Replication first starts in the virtual lab VMs from the application group in the required order and performs necessary tests against them. This way, Veeam Backup & Replication creates the necessary environment to start the verified VM replica. Only after all VMs from the application group are started and tested, Veeam Backup & Replication starts the verified VM replica in the virtual lab.

For example, if you want to verify a Microsoft Exchange Server, you need to test its functionality in cooperation with other components: domain controller and DNS server. Subsequently, you must add to the application group a virtualized domain controller and DNS server. When Veeam Backup & Replication runs a SureReplica job, it will first start and verify the domain controller and DNS server in the virtual lab to make verification of the Exchange Server possible.

Note Veeam Backup & Replication supports mixed application groups. You can add to the same application groups both VMs from backups and VMs from replicas. Keep in mind that all VMs from the application group must belong to the same platform — VMware or Hyper-V, and must have at least one valid restore point created by the time the SureReplica job starts.

Virtual Lab

The virtual lab is an isolated virtual environment in which Veeam Backup & Replication verifies VM replicas. In the virtual lab, Veeam Backup & Replication starts a verified VM replica and VMs from the application group. The virtual lab is used not only for the SureReplica verification procedure, but also for U-AIR and On-Demand Sandbox processing.

A virtual lab does not require provisioning of additional resources. Instead, you can deploy it using existing resources in your virtual environment. For example, you can create a virtual lab on an ESX(i) host in the DR site whose resources are typically under-utilized.

The virtual lab is fully fenced off from the production environment. The network configuration in the virtual lab mirrors the network configuration of the production environment. For example, if verified VM replicas are located in two logical networks in your production environment, the virtual lab will also have two networks. The networks in the virtual lab will be mapped to corresponding production networks.

VM replicas in isolated networks have with the same IP addresses as in the production network. This lets VM replicas in the virtual lab function just as if they would function in the production environment.

Note You can connect VMs that are connected to different networks in the production environment to the same network in the isolated virtual lab.

Proxy Appliance

To enable communication between the production environment and the isolated network in the virtual lab, Veeam Backup & Replication uses a proxy appliance. The proxy appliance is a Linux-based auxiliary VM created on the ESX(i) host where the virtual lab is created. The proxy appliance VM is assigned an IP address from the production network and placed to the virtual lab folder and resource pool created on the ESX(i) host.

The proxy appliance is connected to the production network and to the isolated network and so has visibility of the production environment and the virtual lab. In essence, the proxy appliance acts as a gateway between the two networks, routing requests from the production environment to VM replicas in the virtual lab.

The proxy appliance connects to isolated networks using vNIC adapters. Veeam Backup & Replication adds to the proxy appliance one vNIC adapter per each isolated network. For example, if there are two networks in the virtual lab, Veeam Backup & Replication will add two vNIC adapters to the proxy appliance. The vNIC adapter gets an IP address from the isolated network. Typically, this IP address is the same as the IP address of the default gateway in the corresponding production network.

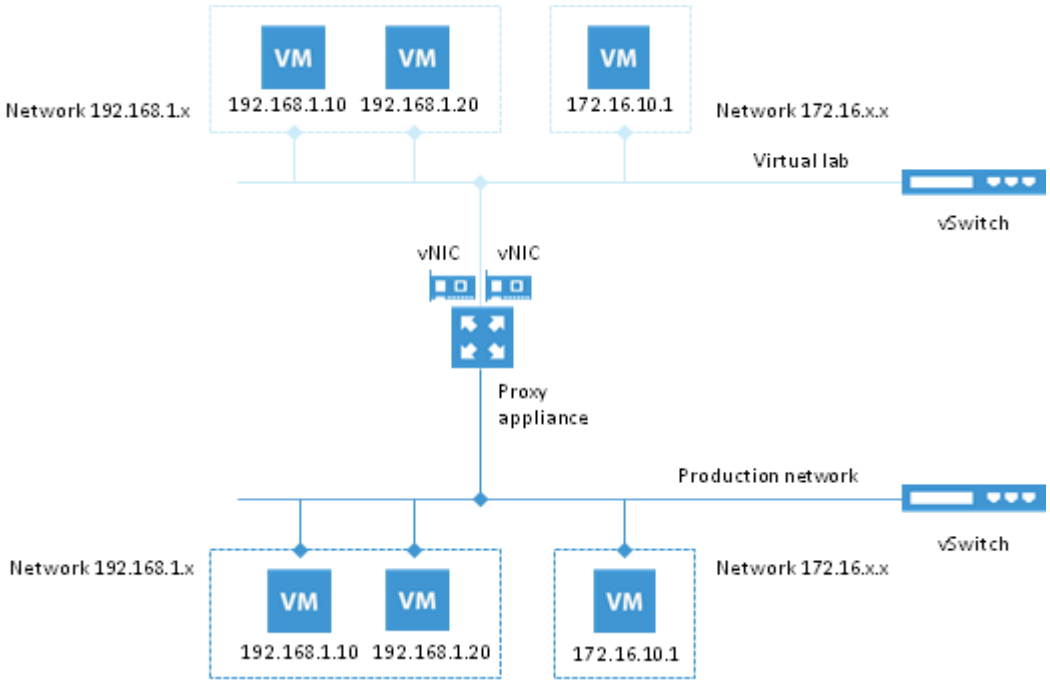

Note The proxy appliance is an optional component. Technically, you can create a virtual lab without a proxy appliance. However, in this case, you will not be able to perform automatic recovery verification of VM replicas. VM replicas will be simply started in the virtual lab; you will have to access them using the VM console and perform necessary tests manually.

IP Masquerading

To let the traffic into the virtual lab, Veeam Backup & Replication uses masquerade IP addressing.

Every VM replica in the virtual lab has a masquerade IP address, along with the IP address from the production network. The masquerade IP address resembles the IP address in the production network: for example, if the IP address of a VM replica is 172.16.1.13, the masquerade IP address may be 172.18.1.13.

The masquerade IP address can be thought of as an entry point to the VM replica in the virtual lab from the production environment. When you want to access a specific VM replica in the virtual lab, Veeam Backup & Replication addresses it by its masquerade IP address.

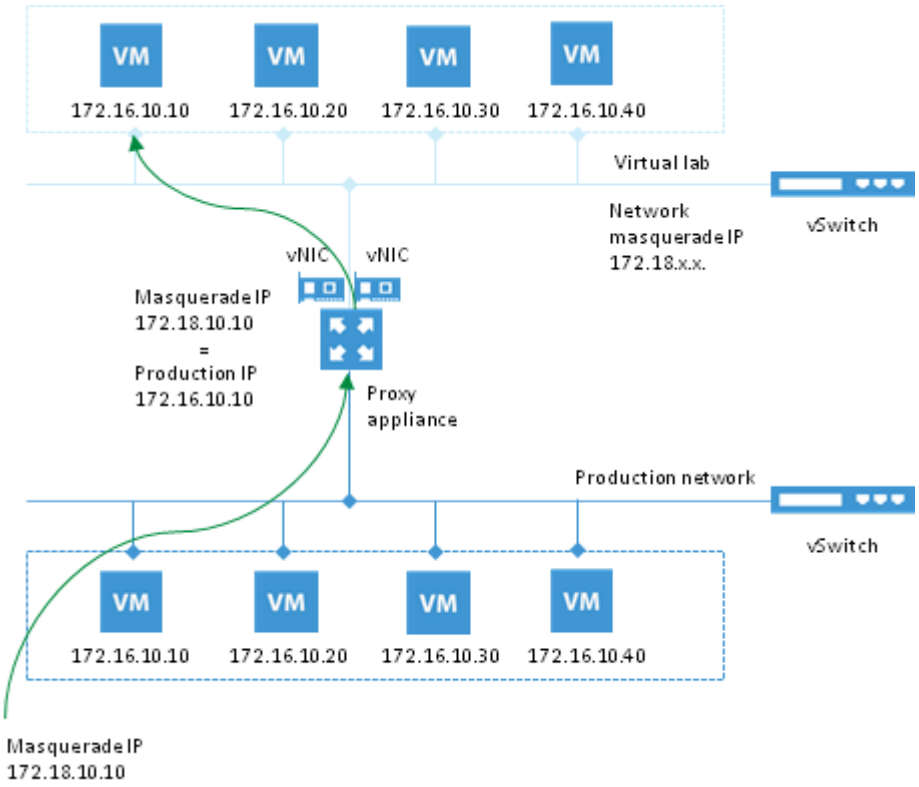

The rules routing requests to VMs in the virtual lab are specified in the routing table on the server from which you want to access VMs in the virtual lab. The routing table can be updated on the following servers:

- **Veeam backup server**. Veeam Backup & Replication automatically creates the necessary static route in the routing table on the Veeam backup server at the moment you launch a SureReplica job and Veeam Backup & Replication starts the virtual lab.
- **Client machine**. If you want to provide your users with access to VM replicas in the virtual lab, you need to manually update routing tables on their machines and add to them a new static route. See also: Static IP Mapping.

The added static route destines the masquerade network traffic to the proxy appliance. The proxy appliance here acts as a NAT device: it resolves the masquerade IP address, replaces it with "real" IP address of a VM from the production network and then directs the request to the necessary VM in the virtual lab. The static route is non-persistent: when you power off the virtual lab, the route is removed from the routing table on the Veeam backup server or the client machine.

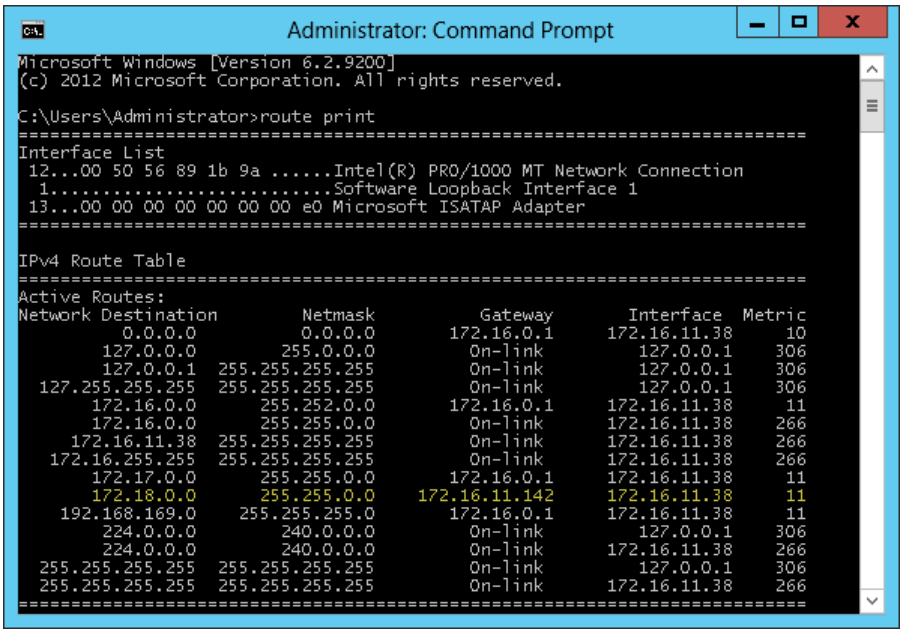

For example, when trying to access a VM with IP address 172.16.10.10 in the isolated network during recovery verification, Veeam Backup & Replication sends a request to the masquerade IP address 172.18.10.10. According to the routing rule added to the IP routing table, all requests are first sent to the next hop — the proxy appliance. The proxy appliance performs address translation, substitutes the masquerade IP address with the IP address in the isolated network and forwards the request to the necessary VM in the isolated network — in the given example, to 172.16.10.10.

Static IP Mapping

Sometimes it is necessary to provide many clients with access to a restored VM, which is especially the case for user−directed application item-level recovery. For example, you may want to provide your users with access to the Exchange Server replica using web−based access (like Outlook Web Access). Technically, you may update the routing table on every client machine; however, this will demand a lot of administrative work.

For such situations, Veeam Backup & Replication enables you to get access to a VM replica in the virtual lab directly from the production environment. To be able to access to a VM replica in the virtual lab, you should reserve a static IP address in the pool of production IP addresses and map this IP address to the IP address of a VM replica in the virtual lab.

The static IP address is assigned to the proxy appliance vNIC connected to the production network. IP traffic directed to the specified static IP address is routed by the proxy appliance to the VM powered on in the isolated network.

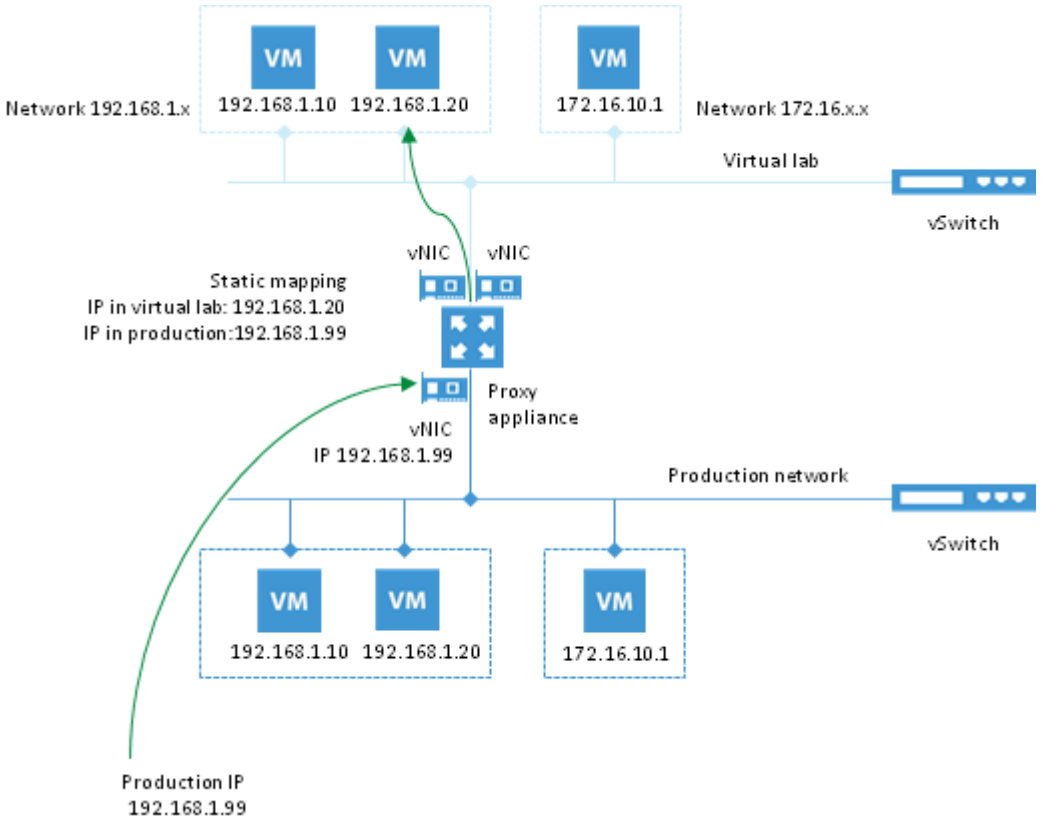

For example, for a VM replica with IP address 192.168.1.20 in the isolated network, you can reserve IP address 192.168.1.99 (a free IP address from the production network). As a result, you will be able to use IP address 192.168.1.99 to access the VM replica in the virtual lab from the production side.

You should also register an alias record in the production DNS server for the reserved IP address. For example, you can register *backup.exchange.local* as an alias for the IP address 192.168.1.99.

Virtual Lab Configuration

For SureReplica recovery verification, Veeam Backup & Replication offers three types of the virtual lab configuration:

- Basic single-host virtual lab
- Advanced single-host virtual lab
- Advanced multi-host virtual lab

See also:

Limitations of Single-Host Virtual Labs

Basic Single-Host Virtual Labs

The basic single-host virtual lab configuration (formerly known as basic networking mode) should be used if your DR site is configured in the following way:

- All VM replicas you want to verify are located on the same ESX(i) host.
- All VM replicas you want to verify are connected to the same network.

Important! For this configuration type, the virtual lab must be created on the same ESX(i) host where VMs replicas are located. If you create the virtual lab on some other ESX(i) host, the SureReplica job will fail.

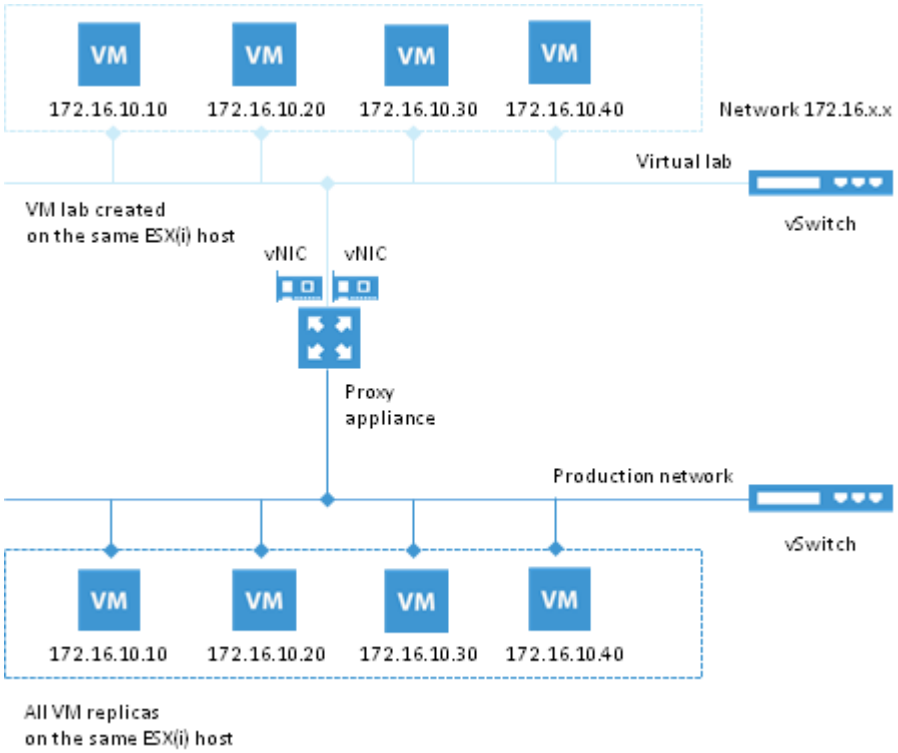

For the basic single-host virtual lab, Veeam Backup & Replication creates one virtual network that is mapped to the corresponding production network. Additionally, Veeam Backup & Replication automatically adds a number of new VMware instances on the ESX(i) host where the virtual lab is created:

- A new resource pool
- A new VM folder
- A new standard vSwitch

The vSwitch is only used by the VMs started in the virtual lab: there is no routing outside the virtual lab to other networks.

Veeam Backup & Replication automatically configures all settings for the basic single-host virtual lab. The proxy appliance is also created and configured automatically and placed to the virtual lab folder and resource pool created on the ESX(i) host.

Advanced Single-Host Virtual Labs

The advanced single-host virtual lab configuration (formerly known as advanced networking mode) should be used if your virtual environment is configured in the following way:

- All VM replicas you want to verify are located on the same ESX(i) host.
- VM replicas you want to verify are connected to different networks.

Important! For this configuration type, the virtual lab must be created on the same ESX(i) host where VMs replicas are located. If you create the virtual lab on some other ESX(i) host, the SureReplica job will fail.

In the advanced single-host virtual lab, Veeam Backup & Replication creates several virtual networks. The number of virtual networks corresponds to the number of production networks to which verified VM replicas are connected. Networks in the virtual lab are mapped to corresponding production networks.

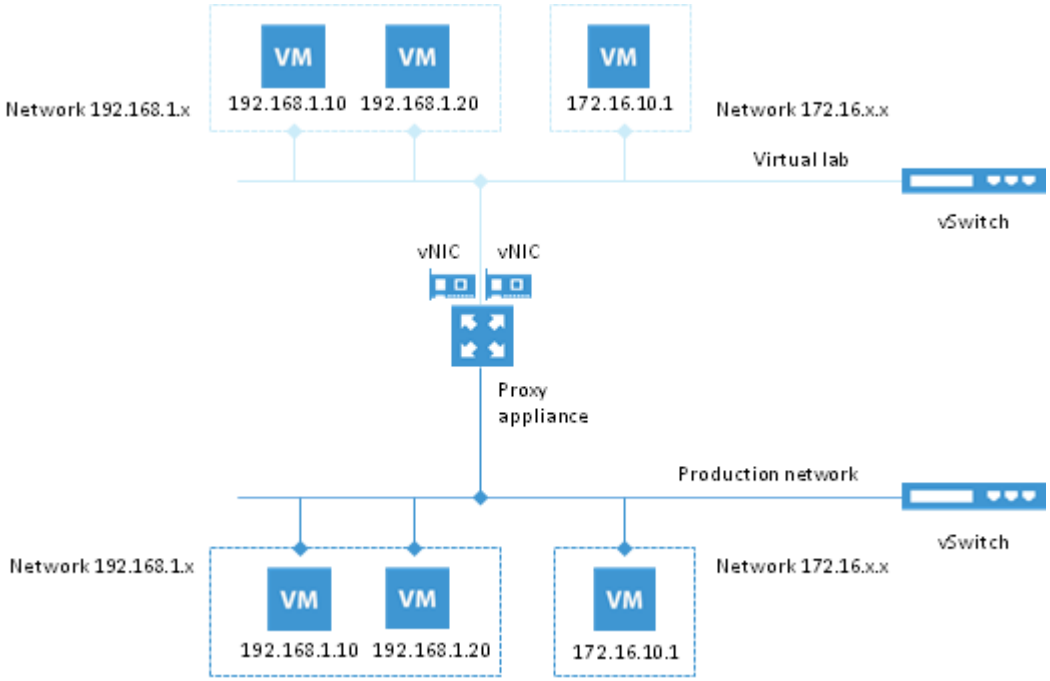

Veeam Backup & Replication automatically adds a number of new VMware instances on the ESX(i) host where the virtual lab is created:

- A new resource pool
- A new VM folder
- A new standard vSwitch

The vSwitch is only used by the VMs started in the virtual lab: there is no routing outside the virtual lab to other networks.

When you create an advanced single-host virtual lab, Veeam Backup & Replication configures basic settings for networks that should be created in the virtual lab. You need to review these settings and manually adjust them if needed.

Limitations of Single-Host Virtual Labs

In case VM replicas are located on different hosts, you cannot use the single-host virtual lab configuration (either basic or advanced). A single-host virtual lab, either basic or advanced, uses standard vSwitches (vSS) that have specific configuration limitations.

When you create or edit a virtual lab, Veeam Backup & Replication creates a new port group for each isolated network in the virtual lab. All VMs from the isolated network are added to this port group. Such configuration helps differentiate the traffic passing through the vSS to the isolated network in the virtual lab.

However, the vSS has a specific restriction: it is "limited" to a certain ESX(i) host. A vSS is configured on a specific ESX(i) host. The configuration of the vSS, such as information about port groups, resides on the ESX(i) host where it is configured. Other ESX(i) hosts in the virtual environment do not have access to this information.

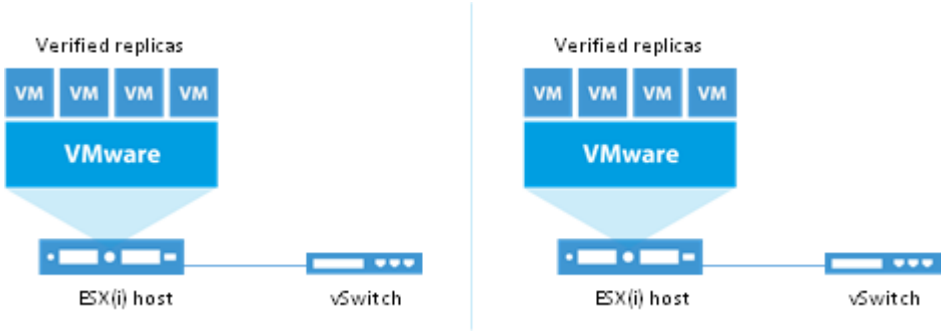

Therefore, the single-host configuration can only be used if all VM replicas are registered on the same ESX(i) host. If you start VM replicas registered on different ESX(i) hosts in the single-host virtual lab, VMs from different port groups will not be able to "see" each other and communicate with each other.

Advanced Multi-Host Virtual Labs

The advanced multi-host virtual lab configuration of virtual lab configuration should be used if your DR site is configured in the following way:

- All VM replicas you want to verify are located on the different ESX(i) hosts
- VM replicas you want to verify are connected to one or several networks

Important! DVS is limited to one datacenter. For this reason, all verified VM replicas and VM replicas from the application group that you plan to start in the virtual lab must belong to the same datacenter. If VM replicas belong to different datacenters, you will be able to start them in the virtual lab but Veeam Backup & Replication will not be able to automatically verify them with SureBackup.

To verify VM replicas registered on different ESX(i) hosts, you should use the advanced multi-host configuration of the virtual lab. The advanced multi-host virtual lab leverages the VMware Distributed vSwitch (DVS) technology. To learn more, see http://www.vmware.com/products/datacenter-virtualization/vsphere/distributed-switch.html.

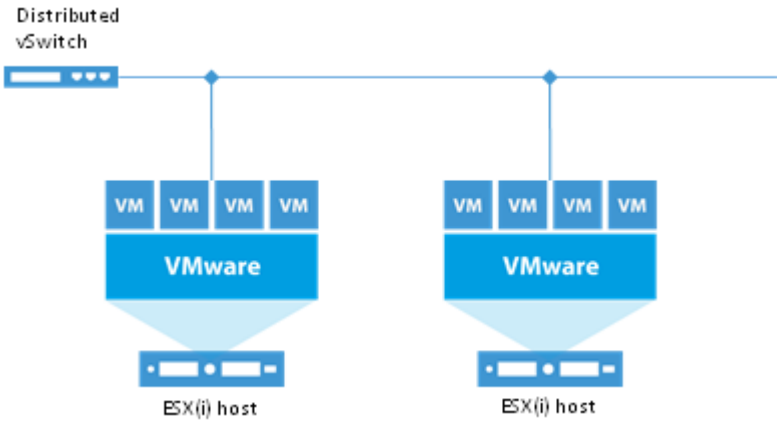

When you configure an advanced multi-host virtual lab, you should select an ESX(i) host on which the proxy server will be created and a DVS on which Veeam Backup & Replication will create isolated network(s). Veeam Backup & Replication does not offer an option to automatically configure the DVS. The DVS you plan to use must be pre-configured in your virtual environment.

With Veeam Backup & Replication, you can optionally connect VMs from different production networks to one network in the isolated virtual lab. In this case, all VM replicas in the virtual lab will be started in the same network.

Isolated Networks on DVS

For every isolated network in the virtual lab, Veeam Backup & Replication add a new DVS port group to the DVS. The added DVS port group is named after the isolated network.

The DVS port groups created on the DVS must be isolated from the production environment. To isolate port groups, you can use one of the following methods:

1. **Connect DVS uplinks to an isolated network**. You can link the DVS you plan to use for recovery verification to an external isolated network using uplink adapters. Note that these network configurations must be performed manually by the backup administrator.

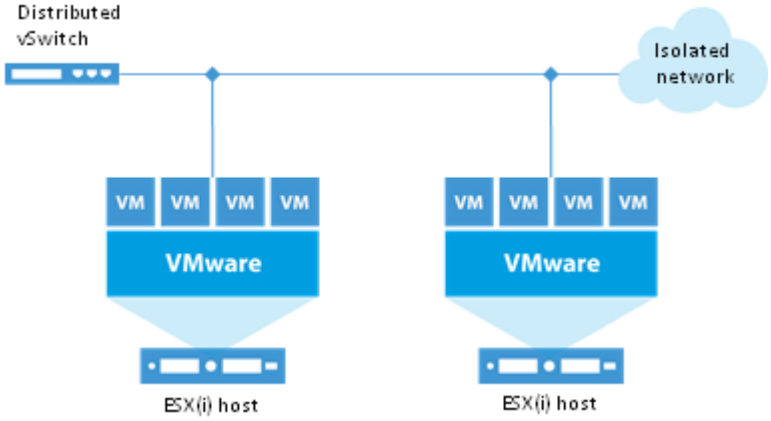

2. **Use VLAN tagging**. This method can be used only if your router supports VLAN ID tagging. When specifying settings for isolated networks in Veeam Backup & Replication, you can define different VLAN IDs for different isolated networks. Setting VLAN IDs restricts communication of VM replicas in the isolated network with the production environment.

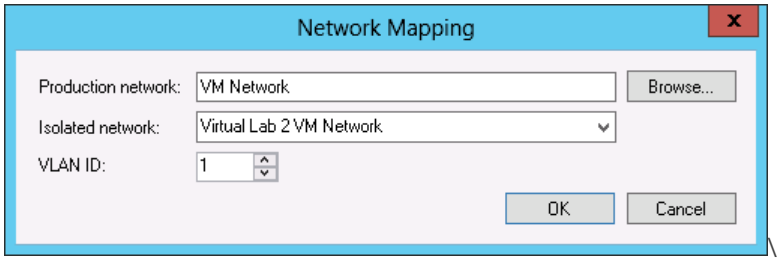

Important! If the router does not support VLAN ID tagging or the virtual lab has been incorrectly configured, VM replicas will be started in the virtual lab but Veeam Backup & Replication will not be able to automatically verify them with SureBackup.

Port Groups and VLAN IDs

When you configure the advanced multi-host virtual lab, you need to be extremely careful when specifying the port group and VLAN ID settings.

Port Groups in Advanced Multi-Host Virtual Labs

For the advanced multi-host virtual lab, Veeam Backup & Replication uses an existing DVS that was configured by the backup administrator beforehand. On the DVS, Veeam Backup & Replication creates a number of new port groups, one per isolated network created in the virtual lab.

When Veeam Backup & Replication creates a new port group, it performs a check of the DVS selected for the virtual lab:

- If a port group with the specified name already exists, Veeam Backup & Replication starts using it for the virtual lab. However, in this case, Veeam Backup & Replication will not be the owner of this port group.
- If a port group with the specified name does not exist, Veeam Backup & Replication creates it and becomes the owner of the created port group.

When a virtual lab is removed, Veeam Backup & Replication checks the ownership of the port group:

- If Veeam Backup & Replication is not the owner of the port group, the port group remains on the DVS; Veeam Backup & Replication simply stops using it.
- If Veeam Backup & Replication is the owner of the port group, it removes this port group from the DVS.

Several virtual labs can use the same port group. For this reason, you should be extremely careful when removing virtual labs. If Veeam Backup & Replication is the owner of the virtual lab and the port group is removed, other virtual labs using the removed port group may fail to start.

VLAN IDs in Advanced Multi-Host Virtual Labs

A DVS port group has VLAN ID settings. In case you select an existing port group to be used for the virtual lab, you should specify its VLAN ID in the virtual lab settings.

- If VLAN ID settings are specified correctly, Veeam Backup & Replication will be able to configure the virtual lab and verify VM replicas in it.
- If VLAN ID settings are specified not correctly, Veeam Backup & Replication will report an error informing that the selected port group exists but cannot be used due to incorrect VLAN ID settings.

SureReplica Job

A SureReplica job is a task for VM replica recovery verification. It aggregates all settings and policies of a recovery verification task, such as application group and virtual lab to be used, VM replicas that should be verified in the virtual lab and so on. The SureReplica job can be run manually or scheduled to be performed automatically.

When a SureReplica job runs, Veeam Backup & Replication first creates an environment for VM replica verification:

1. Veeam Backup & Replication starts the virtual lab.
2. In the virtual lab, it starts VMs from the application group in the required order. VMs from the application group remain running until the verified VM replicas are booted and tested. If Veeam Backup & Replication does not find a successful VM replica or backup for any of VMs from the application group, the SureReplica job will fail.

Once the virtual lab is ready, Veeam Backup & Replication starts the VM replicas to the necessary restore point, tests and verifies them one by one or, depending on the specified settings, creates several streams and tests a number of VM replicas simultaneously. If Veeam Backup & Replication does not find a successful restore point for any of verified VM replicas, verification of this VM replica fails, but the job continues to run.

By default, you can start and test up to three VM replicas at the same time. You can also increase the number of VMs to be started and tested simultaneously. Keep in mind that if these VMs are resource demanding, performance of the SureReplica job may decrease.

Once the verification process is complete, VMs from the application group are powered off. Optionally, you can leave the VMs from the application group running to perform manual testing or enable user-directed application item-level recovery.

In some cases, the SureReplica job schedule may overlap the schedule of the replication job linked to it. The VM replica files may be locked by the replication job and the SureReplica will be unable to verify such replica. In this situation, Veeam Backup & Replication will not start the SureReplica job until the replication job is over.
To overcome the situation of job overlapping, you may chain the replication and SureReplica jobs or define the timeout period for the SureReplica job. To learn more, see Specifying Job Schedule.

Note You can mix backups and replicas in the recovery verification job. For example, the application group may contain VMs that will be started from backup files and the job linked to the recovery verification job may be a replication job. Veeam Backup & Replication supports any type of a mixed scenario. Note that VM backups and VM replicas must belong to the same platform — VMware or Hyper-V.

SureReplica Job Processing

The recovery verification process for VM replicas includes the following steps:

1. **Getting virtual lab configuration**. Veeam Backup & Replication gets information about configuration of the virtual lab where verified VM replicas should be started.
2. **Starting virtual lab routing engine**. Veeam Backup & Replication starts a proxy appliance that is used as a gateway providing access to VM replicas the virtual lab.
3. **Publishing**. Veeam Backup & Replication triggers a protective VMware snapshot for the verified VM replica.
4. **Reconfiguring.** Veeam Backup & Replication updates configuration files to connect the VM replica to the isolated network in the virtual lab.
5. **Configuring DC**. If a verified VM has the Domain Controller or Global Catalog role, the VM is reconfigured.
6. **Powering on**. Veeam Backup & Replication powers on the verified VM replica in the isolated network.
7. **Heartbeat test**. Veeam Backup & Replication checks whether the VMware Tools heartbeat signal (green or yellow) is coming from the VM replica or not. If the VM replica has no VMware Tools, the test will not be performed and a notification will be written to the session details.
8. **Running ping tests**. Veeam Backup & Replication checks if the VM replica responds to the ping requests or not. If the VM replica has no NICs and mapped networks for them and/or has no VMware Tools installed, the ping test will not be performed and a notification will be written to the session details.
9. **Application initialization**. Veeam Backup & Replication waits for the applications installed in the VM (for example, SQL Server, web server, mail server) to start. The application initialization period is defined in the corresponding properties of the recovery verification job and by default is equal to 120 sec. However, depending on the software installed in a VM, the application initialization process may require more time than specified in the recovery verification job settings. If applications installed in a VM are not initialized within the specified period of time, test scripts can be completed with errors. If such error situation occurs, you will need to increase the **Application initialization timeout** value and start the job once again.
10. **Running test scripts**. Veeam Backup & Replication runs scripts to test whether the application installed in the VM replica is working correctly or not. If the VM replica has no VMware Tools installed and/or there are no NICs and mapped networks for them, Veeam Backup & Replication will skip tests that use variables *%vm_ip%* and *%vm_fqdn%*, as the IP address of the VM cannot be determined. Test results are written to the session details. To define whether the script has completed successfully or not, Veeam Backup & Replication uses return codes. If the return code is equal to 0, the script is considered to complete successfully. Other values in the return code mean that the script has failed.
11. **Powering off**. After all tests have been performed, Veeam Backup & Replication powers off the verified VM replica.
12. **Unpublishing**. Veeam Backup & Replication deletes the protective VMware snapshot and rollbacks all changes made to the VM replica while it was running in the virtual lab.
13. **Stopping virtual lab engine**. Veeam Backup & Replication powers off the proxy appliance in the virtual lab.

Stabilization Algorithm

To be able to perform tests for a verified VM replica without errors, Veeam Backup & Replication needs to know that the VM replica is ready for testing. To determine this, Veeam Backup & Replication waits for the VM replica to reach a "stabilization point": the VM replica has been booted and reports it is ready for tests. After the stabilization point has been established, Veeam Backup & Replication can start heartbeat tests, ping tests and test scripts against the VM replica.

Veeam Backup & Replication establishes the stabilization point with the help of VMware parameters that it gets from the VM replica. Depending on the VM replica configuration, it uses one of the three algorithms to do that:

- **Stabilization by IP**. This algorithm is used if the VM replica has VMware Tools installed, there are NICs and mapped networks for these NICs. In this case, Veeam Backup & Replication waits for an IP address of the VM replica for mapped networks that is sent by VMware Tools running in the VM replica. The sent IP address should be valid and should not change for a specific period of time.
- **Stabilization by heartbeat**. This algorithm is used if the VM replica has VMware Tools installed but there are no NICs and mapped networks for them. In this case, Veeam Backup & Replication waits for a corresponding heartbeat signal (green or yellow) to come from the VM replica. As well as in the first case, the signal is sent by VMware Tools running in the VM replica.
- **Stabilization by Maximum allowed boot time**. This algorithm is used if the VM replica has neither VMware Tools installed, nor NICs and mapped networks for them. In this case, Veeam Backup & Replication will simply wait for the time specified in the Maximum allowed boot time field, which is considered to be a stabilization period for the VM replica. Once this time interval is exceeded, Veeam Backup & Replication will consider that the VM replica is successfully booted and is ready for testing.

The stabilization process cannot exceed the value specified in the **Maximum allowed boot time** field. If the stabilization point cannot be determined within the **Maximum allowed boot time**, the recovery verification process will be finished with the timeout error. For this reason, you should be careful when specifying this value — typically, the VM started by a recovery verification job requires more time to boot that a VM replica started regularly. When such error situation occurs, you will need to increase the **Maximum allowed boot time** value and start the job again.
Once the stabilization point has been established, Veeam Backup & Replication runs ping, heartbeat tests and test scripts against the verified VM replica.

Data Recovery

Veeam Backup & Replication offers a number of recovery options for various disaster recovery scenarios:

- Instant VM Recovery enables you to instantly start a VM directly from a backup file
- Full VM recovery enables you to recover a VM from a backup file to its original or another location
- VM file recovery enables you to recover separate VM files (virtual disks, configuration files and so on)
- Virtual drive restore enables you to recover a specific hard drive of a VM from the backup file, and attach it to the original VM or to a new VM
- Windows file-level recovery enables you to recover individual Windows guest OS files (from FAT, NTFS and ReFS file systems)
- Multi-OS file-level recovery enables you to recover files from 15 different guest OS file systems
- Universal Application-Item Recovery (U-AIR) enables you to recover application objects (such as, AD entries, SQL database objects and so on) directly from backup files

Veeam Backup & Replication uses the same image−level backup for all data recovery operations. You can restore VMs, VM files and drives, application objects and individual guest OS files to the most recent state or to any available restore point.

Instant VM Recovery

With instant VM recovery, you can immediately restore a VM into your production environment by running it directly from the backup file. Instant VM recovery helps improve recovery time objectives (RTO), minimize disruption and downtime of production VMs. It is like having a "temporary spare" for a VM: users remain productive while you can troubleshoot an issue with the failed VM. Instant VM recovery supports bulk processing so you can immediately restore multiple VMs at once.

When instant VM recovery is performed, Veeam Backup & Replication uses the Veeam vPower technology to mount a VM image to an ESX(i) host directly from a compressed and deduplicated backup file. Since there is no need to extract the VM from the backup file and copy it to production storage, you can restart a VM from any restore point (incremental or full) in a matter of minutes.

The archived image of the VM remains in read−only state to avoid unexpected modifications. By default, all changes to virtual disks that take place while the VM is running, are logged to auxiliary redo logs residing on the NFS server (Veeam backup server or backup repository). These changes are discarded as soon as a restored VM is removed, or merged with the original VM data when VM recovery is finalized.

To improve I/O performance for a restored VM, you can redirect VM changes to a specific datastore. In this case, instead of using redo logs, Veeam Backup & Replication will trigger a snapshot and put it to the *Veeam IR* directory on the selected datastore, together with metadata files holding changes to the VM image. Keep in mind, though, that if you redirect VM changes, you will not be able to use Storage vMotion (a VMware feature for live migration of VMs within the environment) to complete VM recovery.

To finalize instant VM recovery, you can do one of the following:

- Use Storage vMotion to quickly migrate the restored VM to the production storage without any downtime. In this case, original VM data will be pulled from the NFS datastore to the production storage and consolidated with VM changes while the VM is still running. Storage vMotion, however, can only be used if you select to keep VM changes on the NFS datastore without redirecting them. Please note that Storage vMotion is only available with VMware Enterprise licenses and above.
- Use replication or VM copy functionality of Veeam Backup & Replication. In this case, you can create a copy of a VM and fail over to it during the next maintenance window. In contrast to Storage vMotion, this approach requires you to schedule some downtime while you clone or replicate the VM, power it off and then power the cloned copy or replica on.
- Use Quick Migration. In this case, Veeam Backup & Replication will perform a two−stage migration procedure — instead of pulling data from the vPower NFS datastore, it will restore the VM from the backup file on the production server, then move all changes and consolidate them with the VM data. For details, see Quick Migration.

In many respects, instant VM recovery gives results similar to failover of a VM replica. Both features can be used for tier−1 applications with little tolerance for business interruption and downtime. However, when you perform replica failover, you do not have dependencies on the Veeam backup server. And, unlike instant VM recovery that provides only limited I/O throughput, replication guarantees full I/O performance.

Beside disaster recovery matters, instant VM recovery can also be used for testing purposes. Instead of extracting VM images to production storage to perform regular DR testing, you can run a VM directly from the backup file, boot it and make sure the VM guest OS and applications are functioning properly.

Full VM Recovery

With Veeam Backup & Replication, you can restore an entire VM from a backup file to the latest state or to any good-to-know point in time if the primary VM fails.

In contrast to instant VM recovery, full VM restore requires you to fully extract the VM image to the production storage. Though full VM restore takes more resources and time to complete, you do not need to perform extra steps to finalize the recovery process. Veeam Backup & Replication pulls the VM data from the backup repository to the selected storage, registers the VM on the chosen ESX host and, if necessary, powers it on. Full VM recovery enables full disk I/O performance while Instant VM recovery provides a “temporary spare” for a VM as the vPower NFS throughput is limited.

Full VM recovery can be performed in one of the following modes (for details, see Transport Modes):

- If the backup proxy is virtualized and resides on the ESX host to which a VM should be restored, Veeam Backup & Replication will use the Virtual Appliance transport mode to deliver VM data. The Virtual Appliance mode utilizes VMware ESX capabilities of hot−adding disks to a VM and thus eliminates the need to transfer the backup data across the network. Veeam transport services deployed on the backup repository and a backup proxy retrieve VM data from the backup file and put it directly to the shared storage.
- If the Virtual Appliance transport mode cannot be utilized, VMs are restored with the Network transport mode.

Note You can restore a VM to an ESX host of the same or later version than the host on which the backup was created. For instance, if you created a backup of a VM running on ESX 3.0, you can restore this VM to ESX 3.0, 3.0.1, 3.0.2 and later or to ESXi.

A VM can be restored to its original location or to a new location. When you restore a VM to its original location, the primary VM is automatically turned off and deleted before the restore. This type of

restore ensures the quickest recovery and minimizes the number of mistakes which can be potentially caused by changes in VM settings.

When you restore a VM to a new location, you need to specify new VM settings such as the new VM name, the host and datastore where the VM will reside, disk format (thin or thick provisioned) and network properties. Veeam Backup & Replication will change the VM configuration file and store the VM data to the location of your choice.

VM File Recovery

Veeam Backup & Replication can help you to restore specific VM files (.vmdk, .vmx and others) if any of these files are deleted or the datastore is corrupted. This option provides a great alternative to full VM restore, for example, when your VM configuration file is missing and you need to restore it. Instead of restoring the whole VM image to the production storage, you can restore the specific VM file only.

When you perform VM file restore, VM files are restored from regular image−level backups. Veeam transport services deployed on the backup repository and the backup proxy retrieve VM data from the backup file and send it to the original VM location, or to a new location specified by the user.

Virtual Drive Recovery

Another data recovery option provided by Veeam Backup & Replication is restore of a specific hard drive of a VM. If a VM hard drive becomes corrupted for some reason (for example, with a virus), you can restore it from the image−based backup to any good−to−know point in time. The recovered hard drive can be attached to the original VM to replace a corrupted drive, or connected to any other VM. With the virtual drive restore, you can preserve the format of a recovered drive or convert the drive to the thin or thick format on the fly.

Guest OS File Recovery

With Veeam's Instant File-Level Recovery (IFLR), you can recover an individual file from a backup file or replica to the most recent state or to any point in time in just a few seconds. IFLR does not require you to extract VM image to the local drive or to start up the VM prior to restore — you can recover files directly from a regular image−level backup or replica.

IFLR is available for any virtualized file system, although, Veeam Backup & Replication provides different approaches for different file systems:

- For Windows-based VMs with NTFS, FAT and ReFS file systems, Veeam Backup & Replication uses built−in Windows file-level recovery
- For most commonly used file systems on Windows, Linux, Solaris, BSD, Novell Netware, Unix and Mac machines, Veeam Backup & Replication offers multi-OS file-level recovery
- For any other file system, Veeam Backup & Replication enables you to leverage Instant VM Recovery to perform manual file-level recovery

Windows File-Level Recovery

For FAT, NTFS and ReFS guest OS systems, Veeam Backup & Replication uses built−in file−level restore functionality.

When you perform file−level recovery, the VM image is not extracted from the backup file. The content of the backup file is mounted directly to the Veeam backup server and displayed in the built-in Veeam Backup Browser. For mounting file systems of VM guest OSs, Veeam Backup & Replication uses its proprietary driver. After that you can copy necessary files and folders to your local machine drive, save them anywhere within the network or simply point any applications to the files and use them normally. The backup file (or replica) remains in read−only state no matter what you do.

Multi-OS File-Level Recovery

Because Windows cannot read other file systems natively, Veeam Backup & Replication additionally provides multi−OS file−level recovery that allows reading data from 16 different file systems:

OS	Supported File Systems
Windows	FAT FAT32 NTFS ReFS
Linux	ext ext2 ext3 ext4 ReiserFS JFS XFS
BSD	UFS UFS2
Mac	HFS HFS+
Solaris	ZFS (up to pool version 23)
Novell Netware, Novell Open Enterprise Server	Novell Storage Services

Multi−OS file−level recovery understands not only basic disks, but also Linux LVM (Logical Volume Manager) and Windows LDM (Logical Disk Manager) partitions and ZFS pools.

Multi−OS file−level recovery is a wizard−driven process. To restore files from VM guest OS, Veeam Backup & Replication utilizes its patent−pending approach based on the use of a special FLR helper. The FLR helper is a virtual appliance running a stripped down Linux kernel that has a minimal set of components. The appliance is very small — around 20 MB and takes only 10 seconds to boot.

The FLR helper appliance is created directly on the selected ESX(i) host. Whenever you perform file−level restore, Veeam Backup & Replication automatically starts the appliance and mounts the VMDK files to the FLR appliance as virtual hard drives. VMDK files are mounted directly from backup files, without prior extraction of the backup content.

Once the restore process is complete, the wizard displays the file browser window providing you with direct access to the VM file system. You can then copy necessary files and folders to your local machine drive or save them anywhere within the network. Alternatively, you can allow users to restore files on their own through enabling an FTP server on the virtual appliance.

Tip When you perform recovery directly to a Linux host, you can recover files with correct permissions.

File-Level Recovery for Any File System

With the vPower technology, Veeam extends its IFLR to any file system, not just Windows FAT, NTFS, ReFS and those restored with the multi−OS file−level recovery wizard.

This type of file−level restore is not wizard−driven. You should leverage instant VM recovery to publish VMDK from a backup on the vPower NFS datastore (without actually starting the VM). With the VMDK files readily available, you can mount these VMDKs to any VM that can read the corresponding file system (including the original VM), and restore the required files using native OS file management tools. Alternatively, you can mount the VMDK to a Windows VM, and use a tool such as Portlock Explorer.

Universal Application-Item Recovery

Universal Application Item−Level Recovery (or U−AIR) addresses one of the most common IT problem — it enables you to restore individual objects from virtualized applications (for example, email messages, database records, directory objects and so on).

For recovery of application objects, U−AIR leverages the vPower technology. It starts the application and all components required for its proper work in an isolated virtual lab directly from compressed and deduplicated backup files. Once the VM is started, U−AIR provides transparent access to the backed up VM image through a proxy appliance that has visibility of both the virtual lab and production environment. Users can then extract the necessary application objects from the earlier VM images and bring them back to the production environment.

U−AIR does not require any special backups or additional tools — the application is started directly from the image−level backup file and users can restore application objects with the native management tools.

Technically, U−AIR is a set of wizards that guide you through the process of application objects recovery. For such applications as Active Directory, Microsoft SQL and Microsoft Exchange, U−AIR offers application−specific wizards (that is, you can restore necessary items from these applications using only Veeam's wizards). For other applications, U−AIR offers a universal wizard (that is, Veeam Backup & Replication starts the application and all required components in the virtual lab so that users can connect to that application with the native management tools and restore items manually).

U−AIR wizards are not tied to Veeam Backup & Replication — these are standalone components that can be downloaded, installed and updated independent of the product release. You can install U−AIR wizards on any machine in your production environment.

Replication

To ensure efficient and reliable data protection in your virtual environment, Veeam Backup & Replication complements image−based backup with image−based replication. Replication is the process of copying a VM from its primary location (source host) to a destination location (redundant target host). Veeam Backup & Replication creates an exact copy of the VM (replica), registers it on the target host and maintains it in sync with the original VM.

Replication provides the best recovery time objective (RTO) and recovery point objective (RPO) values, as you actually have a copy of your VM in a ready−to−start state. That is why replication is commonly recommended for the most critical VMs (which run tier 1 applications) that need minimum RTOs. Veeam Backup & Replication provides means to perform both onsite replication for high availability (HA) scenarios and remote (offsite) replication for disaster recovery (DR) scenarios. To facilitate replication over WAN or slow connections, Veeam Backup & Replication optimizes traffic transmission — it filters out unnecessary data blocks (such as, duplicate data blocks, zero data blocks or blocks of swap files) and compresses replica traffic. Veeam Backup & Replication also allows you to apply network throttling rules to prevent replication jobs from consuming the entire bandwidth available in your environment.

Replication is a job−driven process with one replication job used to process one or more VMs. You can start the job manually every time you need to copy VM data or, if you want to run replication unattended, create a schedule to start the job automatically. Scheduling options for replication jobs are similar to those for backup jobs. For details, see Scheduling.

In many respects, replication of VMware VMs works similarly to forward incremental backup. During the first run of a replication job, Veeam Backup & Replication copies the original VM running on the source host and creates its full replica on the target host. Unlike backup files, replica virtual disks are stored uncompressed in their native format. All subsequent replication job runs are incremental (that is, Veeam Backup & Replication copies only those data blocks that have changed since the last replication cycle).

For every replica, Veeam Backup & Replication creates and maintains a configurable number of restore points. If the original VM fails for any reason, you can temporary or permanently fail over to a replica and thus restore critical services with minimum downtime. If the latest state of a replica is not usable (for example, if corrupted data was replicated from source to target), you can select previous restore point to fail over to. Veeam Backup & Replication utilizes VMware ESX snapshot capabilities to create and manage replica restore points. A new incremental run of the replication job takes a regular snapshot of a replica. Blocks of data that have changed since the last job run are written to the snapshot delta file and the snapshot delta file acts as a restore point.

VMware replica restore points are stored in a native VMware format next to replica virtual disk files, which allows Veeam Backup & Replication to accelerate failover operations. To restore a replica in the required state, there is no need to apply rollback files. Instead, Veeam Backup & Replication uses native VMware mechanism of reverting to a snapshot.

As well as for backup jobs, for replication jobs you can define a retention period. Veeam Backup & Replication will keep only the specified number of points, removing any snapshots that breach the retention policy. For details, see Retention Policy.

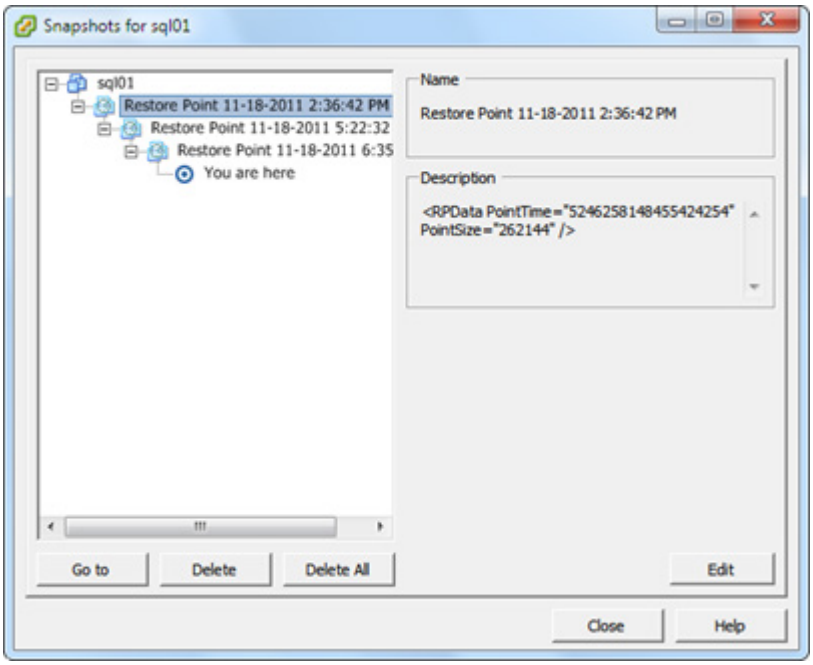

If the Veeam backup server becomes unavailable, you can fail over to any existing replica restore point without Veeam Backup & Replication (via vSphere Client or using a PowerCLI script).

You can define a retention period for replication jobs in the same way as backup jobs. Veeam Backup & Replication will keep only the specified number of points, removing any snapshots that breach the retention policy. For details, see Retention Policy.

Replicas include the following files:

- Full VM replica (a set of VM configuration files and virtual disks)
- Replica restore points (snapshot delta files)
- Replica metadata (.vbk files) used to store replica checksums. Veeam Backup & Replication uses this file to quickly detect changed blocks of data between two replica states. For details, see Changed Block Tracking.

The full VM replica along with its restore points is stored in a dedicated folder on the target datastore. Replica metadata files are located on a backup repository.

Replication Architecture

The replication infrastructure in VMware vSphere environment comprises the following components:

- Source host and target host with associated datastores
- One or two backup proxy servers
- Backup repository

The source host and the target host produce two terminal points between which replicated data is moved. The role of a target can be assigned to a single ESX(i) host or to an ESX(i) host cluster. Assigning a cluster as a target ensures uninterrupted replication in case one of the cluster hosts fails.

Replicated data is collected, transformed and transferred with the help of Veeam transport services. Veeam Backup & Replication uses three Veeam transport services for each replication job — a source–side Veeam transport service, a target–side Veeam transport service and a transport service hosted on the repository. These transport services communicate with each other and maintain a stable connection. During replication, the source–side Veeam transport service interacts with the source host and the target–side Veeam transport service interacts with the target host. The Veeam transport service hosted on the repository works with replica metadata files.

Note Although the most part of replica data is written to the target datastore, replica metadata files are located on the backup repository. This backup repository should be deployed closer to the source host, so that the source−side Veeam transport service can easily communicate with the Veeam transport service hosted on the repository to obtain metadata required for incremental job runs.

All replication infrastructure components engaged for the job make up a data pipe. VM data is moved over this data pipe block by block with multiple processing cycles for replication of each VM. When a new replication session is started, the target−side Veeam transport service obtains job instructions and communicates with the source−side Veeam transport service to begin data collection.

1. **The source−side transport service** accesses the VM image and copies VM data using one of VMware transport modes, as prescribed by the backup proxy server settings. While copying, the source−side Veeam transport service performs additional processing — it consolidates the content of virtual disks by filtering out overlapping snapshot blocks, zero data blocks and blocks of swap files. During incremental job runs, the Veeam transport service retrieves only those data blocks that have changed since the previous job run.
   In all cases when VMware CBT is not available, the source−side Veeam transport service interacts with the Veeam transport service hosted on the repository to obtain replica metadata. The source−side Veeam transport service uses this metadata to detect blocks that have changed since the previous job run.
   Copied blocks of data are compressed and moved from the source−side Veeam transport service to the target−side Veeam transport service.
2. **The target-side Veeam transport service** decompresses replica data and writes the result to the destination datastore.

Note In case of onsite replication, the source−side Veeam transport service and the target−side transport service may run on the same backup proxy server. In this case, no compression is performed.

Veeam Backup & Replication supports a number of replication scenarios that depend on the location of the target host.

To streamline the replication process, you can deploy the backup proxy on a VM. The virtual backup proxy must be registered on an ESX(i) host that has a direct connection to the target datastore. In this case, the backup proxy will be able to use the Virtual Appliance transport mode for writing replica data to target. For details, see Transport Modes.

During the first run of a replication job, Veeam Backup & Replication creates a replica with empty virtual disks on the target datastore. If the Virtual Appliance mode is applicable, replica virtual disks are mounted to the backup proxy and populated through the ESX host I/O stack. This results in increased writing speed and fail−safe replication to ESXi targets.

If the backup proxy is deployed on a physical server, or the Virtual Appliance mode is not available for other reasons, Veeam Backup & Replication will use the Network transport mode to populate replica disk files.

Onsite Replication

If the source and target hosts are located in the same site, you can deploy one backup proxy for data processing and a backup repository for storing replica metadata. This backup proxy must have access to the source host and to the target host at the same time. In this scenario, the source−side Veeam transport service and the target−side Veeam transport service will be started on the same backup proxy. Replication traffic will be transferred uncompressed between the two Veeam transport service.

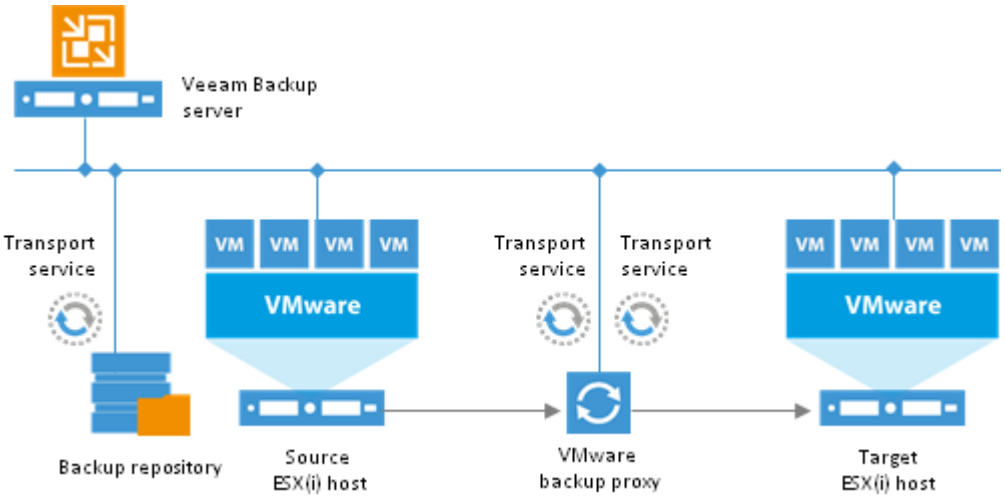

Offsite Replication

The common requirement for offsite replication is that one Veeam transport service runs in the production site (closer to the source host) and another Veeam transport service runs in the remote DR site (closer to the target host). During backup, the Veeam transport services maintain a stable connection, which allows for uninterrupted operation over WAN or slow links.

Thus, to replicate across remote sites, you should deploy at least one local backup proxy in each site — a source backup proxy in the production site, and a target backup proxy in the remote DR site. The backup repository should be deployed in the production site, closer to the source backup proxy.

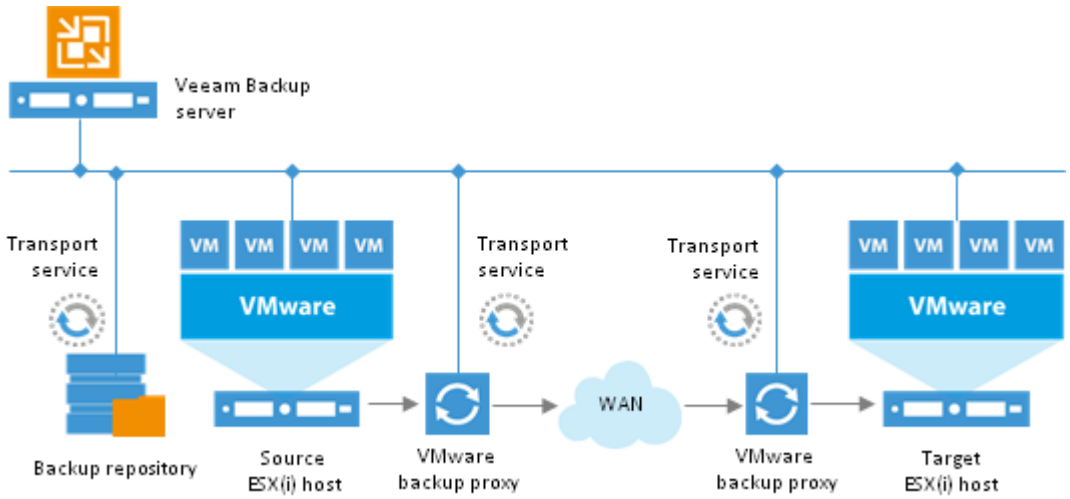

Note When planning for offsite replication, consider advanced possibilities to reduce the amount of replication traffic and streamline replica configuration. These include replica seeding, replica mapping, network mapping and re-IP.

Replica Seeding

If you replicate a VM to a remote DR site, you can use replica seeding. Replica seeding helps significantly minimize the amount of traffic going from the production site to the disaster recovery site over WAN or slow LAN links.

With replica seeding, you do not have to transfer all of VM data from the source host to the target host across the sites when you perform initial replication. Instead, you can use a VM backup created with

Veeam Backup & Replication as a replica "seed". When the replication job starts, Veeam Backup & Replication will use the seed to build a VM replica.

Replica seeding includes the following steps:

1. As a preparatory step for replica seeding, you need to create a backup of a VM that you plan to replicate.
2. The created backup should then be copied from the backup repository in the production site to the backup repository in the DR site.
3. When you create a replication job, you should point it to the backup repository in the DR site. During the first run of a replication job, Veeam Backup & Replication accesses the repository where the replica seed is located, and restores the VM from the backup. The restored VM is registered on the replication target host in the DR site. Files of the restored VM are placed to the location you specify as the replica destination datastore.
   Virtual disks of a replica restored from the backup preserve their format (that is, if the original VM used thin provisioned disks, virtual disks of the VM replica are restored as thin provisioned).
4. Next, Veeam Backup & Replication synchronizes the restored VM with the latest state of the original VM. After successful synchronization, in the **Backup & Replication** view of Veeam Backup & Replication, under **Replicas** node you will see a VM replica with two restore points. One point will contain the state of the VM from the backup file; the other point will contain the latest state of the original VM you want to replicate.
5. During all subsequent runs of the replication job, Veeam Backup & Replication transfers only incremental changes in a regular manner.

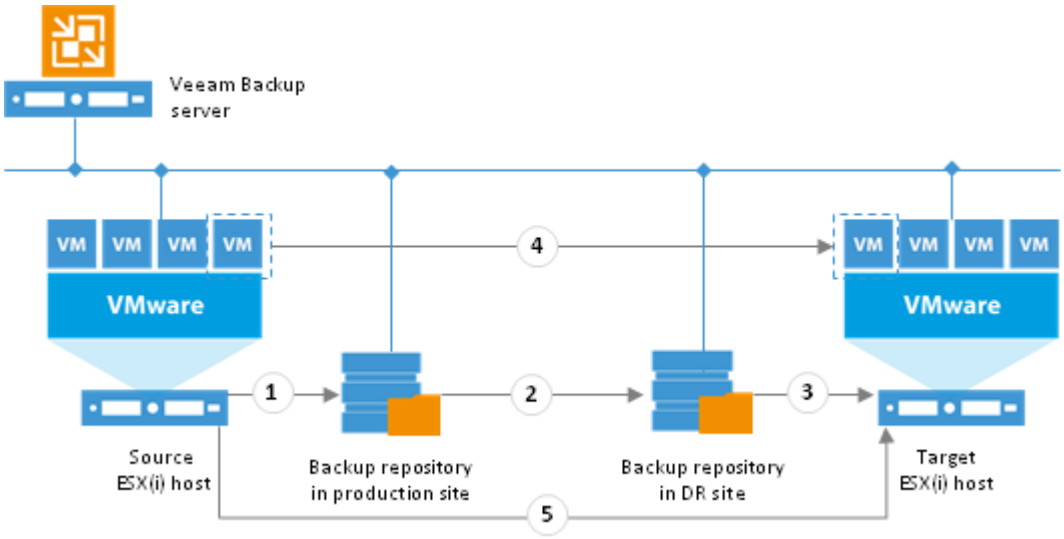

Replica seeding dramatically reduces traffic sent over WAN or slow connections because Veeam Backup & Replication does not send the full contents of the VM image. Instead, it transmits only differential data blocks.

Tip If you add new VMs to an already existing replication job, you can enable replica seeding settings for these VMs. In this case, the newly added VMs will be seeded from the selected backups at the next pass of the replication job. VMs that have already been processed by the job by the time you add new VMs will be processed in a regular manner.

Replica Mapping

To replicate VMs over WAN and slow connections, you can use replica mapping. Similar to replica seeding, replica mapping helps reduce traffic sent to the target host. Replica mapping can be a valuable option when you need to reconfigure or recreate replication jobs, for example, if you need to split one replication job into several jobs.

When configuring a new replication job, you can map an original VM in the production site to an already existing VM in the DR site. For example, this can be a replica VM created with a previous replication job or a VM restored from a backup on a DR target host.

Replication to a mapped VM is performed in the following way:

1. During the first run, the replication job will calculate the differences between the original and mapped VM. Instead of copying and transferring the whole of the original VM, the first replication job will transfer only increments to synchronize the state of the mapped VM with the state of the original VM. After successful synchronization, in the Backup & Replication view of Veeam Backup & Replication, under Replicas node you will see a VM replica with two restore points. One point will contain the latest state of the mapped VM (the VM located on the target host); the other point will contain the latest state of the original VM on the source host.
2. All subsequent runs of the replication job will transfer only increments as well.

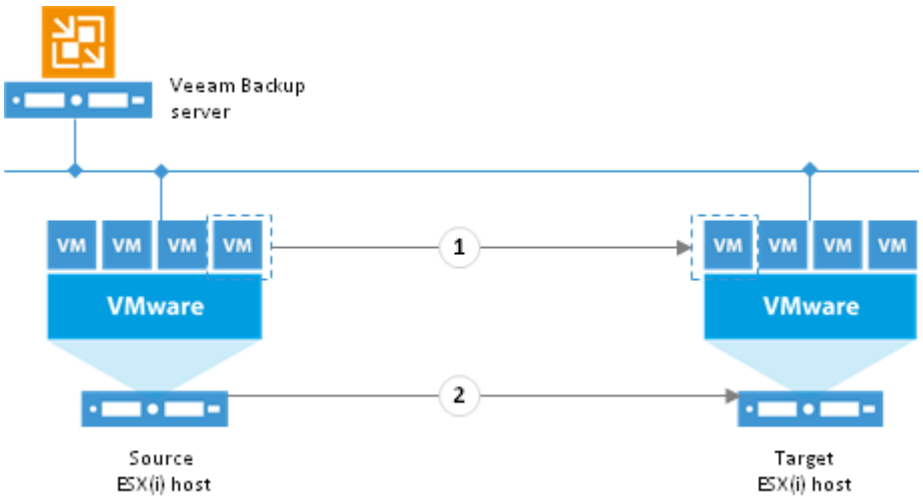

Note When you perform replica mapping, all snapshots of the mapped VM (the VM running on the target host) will be deleted. As a result, you will have only the most recent state of the mapped VM.

Network Mapping and Re-IP

If you use different network and IP schemes in the production and DR site, in the common case you would need to change the network configuration of a VM replica before starting it. To eliminate the need for manual replica reconfiguration and ensure minimum failover downtime, Veeam Backup & Replication offers possibilities of network mapping and automatic IP address transformation.

With Veeam Backup & Replication, a replicated VM uses the same network configuration as the original VM. If the network in your DR site does not match the production network, you can create a network mapping table for the replication job. The table maps source networks to target networks.

During every job run, Veeam Backup & Replication checks the network configuration of the original VM against the mapping table. If the original VM network matches a source network in the table, Veeam Backup & Replication updates the replica configuration file to replace the source network with the target one. The VM replica is then re−registered. Thus, network settings of a VM replica are always kept up to date with the DR site requirements. In case you choose to fail over to the VM replica, it will be connected to the correct network.

For Windows−based VMs, Veeam Backup & Replication also automates reconfiguration of VM IP addresses. If the IP addressing scheme in the production site differs from the DR site scheme, you can create a number of Re−IP rules for the replication job.

When you fail over to the replica, Veeam Backup & Replication checks if any of the specified Re−IP rules apply to the replica. If a rule applies, Veeam Backup & Replication console mounts image−based disks of the replica and changes its IP address configuration via the Windows registry. The whole operation takes less than a second. If failover is undone for any reason or if you fail back to the original location, replica IP address is changed back to the pre−failover state.

Replica Failover and Failback

In case of software or hardware malfunction, you can quickly recover a corrupted VM by failing over to its replica. When you perform failover, a replicated VM takes over the role of the original VM. You can fail over to the latest state of a replica or to any of its good known restore points.

In Veeam Backup & Replication, failover is a temporary intermediate step that should be further finalized. Veeam Backup & Replication offers the following options for different disaster recovery scenarios:

- You can perform **permanent failover** to leave the workload on the target host and let the replica VM act as the original VM. Permanent failover is suitable if the source and target hosts are nearly equal in terms of resources and are located on the same HA site.
- You can perform **failback** to recover the original VM on the source host or in a new location. Failback is used in case you failed over to a DR site that is not intended for continuous operations and would like to move the operations back to the production site when the consequences of a disaster are eliminated.

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

Failover

Failover is a process of switching over from the original VM on the source host to its VM replica on the target host.

During failover, Veeam Backup & Replication recovers a fully functional VM to the required restore point on the target 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 minimum disruption.

The failover operation is performed in the following way:

1. Veeam Backup & Replication rolls back the VM replica to the required restore point.
2. The VM replica is powered on.
3. All changes made to the VM replica while it runs in the failover state are written to the delta file of the snapshot, or restore point, to which you have selected to roll back.

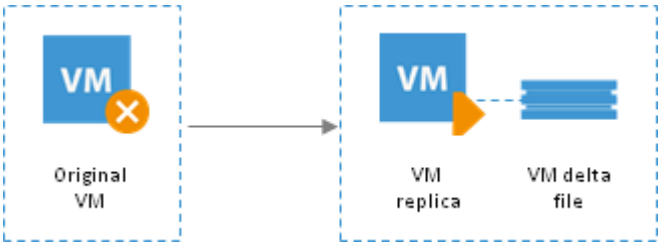

As a result of failover, the state of the replica is changed from *Normal* to *Failover*.
Veeam Backup & Replication temporarily puts replication activities for the original VM on hold until its replica is returned to the *Normal* state.

In Veeam Backup & Replication, the actual failover is considered a temporary stage that should be further finalized. While the replica is still in the *Failover* state, you have the option to undo failover, perform failback, or make failover permanent. In a disaster recovery scenario, after you test the VM replica and make sure the VM runs stable, you should take another step to perform permanent failover.

Important! If possible, avoid powering on a replica manually as it may disrupt further replication operations or cause loss of important data. It is strongly recommended to use Veeam Backup & Replication functionality to perform failover operations.

Permanent Failover

To confirm failover and finalize recovery of a VM replica on the target host, you need to perform permanent failover. As a result of permanent failover, the VM replica ceases to exist as a replica and takes on the role of the original VM.

The permanent failover operation is performed in the following way:

1. Veeam Backup & Replication removes replica restore points from the list of replicas in the Veeam Backup & Replication console and clears associated files from the datastore.
2. Changes written to the delta file of the snapshot, or restore point, are committed to the VM replica disk files to bring the VM replica to the most recent state.

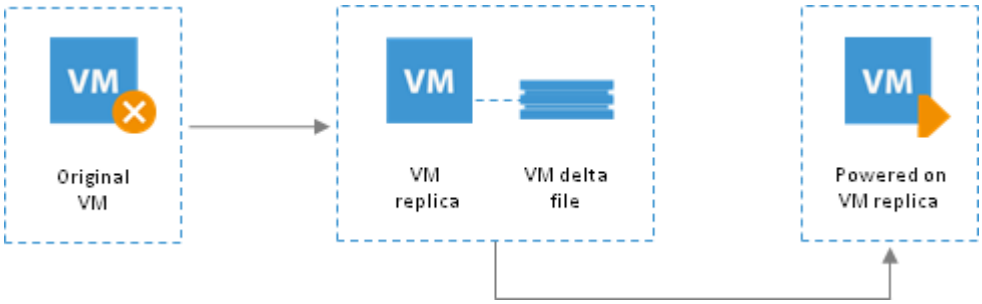

To protect the VM replica from corruption after performing a permanent failover, Veeam Backup & Replication reconfigures the replication job and adds the original VM to the list of exclusions. When the replication job that processes this VM starts, the VM will be skipped from processing and no data will be written to the working VM replica.

Undo Failover

To switch back to the original VM, revert replication operations and discard changes made to the working VM replica, you can undo failover.

When failover is undone, the VM replica reverts to its pre-failover state and all changes that have taken place since the VM replica was powered on are discarded.

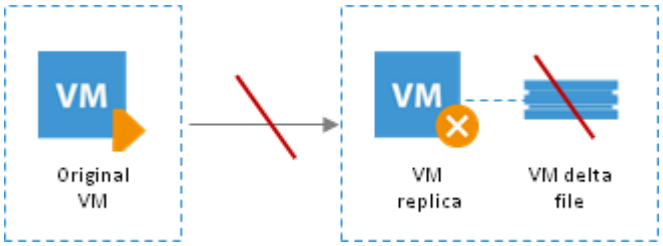

As a result of the undo failover operation, the state of a replica changes back to *Normal* – this means that during the next run, the replication job will process the original VM and create a new replica restore point.

Note During failover, the state of the original VM on the source host is not affected in any way. Basically, if you need to test the replica and its restore points for recoverability, you can perform actual failover as a background process, while the original VM is running. After all necessary tests, you can undo failover and go back to the normal mode of operation.

Failback

Veeam Backup & Replication streamlines and automates disaster recovery by providing replica failback capabilities. Failback is the process of switching from the VM replica to the production VM. During failback, Veeam Backup & Replication uses the working replica to recover the original VM and switch back to it.

If you managed to restore operation of the source host, you can switch back to the original VM on the source host. However, if the source host is not available, you can restore the original VM to a new location and switch back to it. Veeam Backup & Replication offers three failback options:

- Fail back to a VM in the original location on the source host
- Fail back to a VM that has been restored up-front from a backup in a new location
- Fail back to an entirely new location by transferring all replica files to the selected destination

The first two options help you decrease recovery time and use of the network traffic, as Veeam Backup & Replication will transfer only differences between the two VMs. The third option is used in cases when there is no way to use the original VM or restore the VM before performing failback.

During failback, Veeam Backup & Replication protects a running VM replica with a failback snapshot. The snapshot acts as a restore point and saves the pre-failback state of a replica to which you can return afterwards.

Veeam Backup & Replication uses the VM replica to restore the original VM in the selected location.

- When the VM replica is failed back to an existing VM (either the original VM on the source host or a VM restored from backup in a new location), Veeam Backup & Replication calculates the differences and synchronizes the original VM with the VM replica. The original VM is then powered on.
- When the VM replica is failed back to an entirely new location, all of its files are transferred to the target datastore.

After failover, the VM replica is running; changes stored in the snapshot differential file are locked. During failback, Veeam Backup & Replication transfers replica data to the target destination in two stages.

1. First, Veeam Backup & Replication updates the restored VM to the replica failover state. The VM replica is then stopped, and a failback protective snapshot is taken. The replica stays stopped until the moment when failback is committed or undone.
2. Next, Veeam Backup & Replication transfers all the changes made after failover (that is, changes made in the interval between the failover and failback protective snapshots).

The state of the replica is changed from *Failover* to *Failback*. Replication activities for the recovered VM are put on hold.

In Veeam Backup & Replication, failback is considered a temporary stage that should be further finalized. That is, after you test the recovered original VM and make sure it is running correctly, you should take another step to commit failback. However, while the replica is still in the failback state, you also have an option to undo failback and return the replica back to the failover state.

Commit Failback

To confirm failback and finalize recovery of the original VM, you need to commit failback.

As a result of failback commit, Veeam Backup & Replication removes the protective snapshots and unlocks replica disk files. The state of the replica is changed from *Failback* to *Normal*.

Further operations of Veeam Backup & Replication depend on the location to which the VM is failed back:

- If the VM replica is failed back to a new location, Veeam Backup & Replication additionally reconfigures the replication job and adds the former original VM to the list of exclusions. The VM restored in the new location takes the role of the original VM, and is included into the replication job instead of the excluded VM. When the replication job starts, Veeam Backup & Replication will skip the former original VM from processing, and will replicate the newly restored VM instead.
- If the VM replica is failed back the original location, the replication job is not reconfigured. When the replication job starts, Veeam Backup & Replication will process the original VM in the normal mode.

Undo Failback

If the VM to which you failed back from a replica is non-operational or corrupted, you can undo failback and switch the replica back to the failover state.

When failback is undone, the replica deletes the protective failback snapshot. Changes made while the VM replica was in the failback state are discarded. As a result of the undo failback operation, the state of a replica reverts from *Failback* to *Failover*.

VM Copy

With Veeam Backup & Replication, you can run a VM copy job to create an independent fully−functioning copy of a Vm or VM container on the selected storage. VM copying can be helpful if you want to move your datacenter, create a test lab and so on.

The produced copy of a VM is stored uncompressed, in a native VMware vSphere format, so it can be started right away. Although VM copy is similar to replication in many respects, there are several important differences.

- VM copy is a single−use process (that is, every run of a VM copy job mirrors a VM in its latest state). Due to their nature, VM copy jobs do not support incremental runs.
- Veeam Backup & Replication does not create and maintain restore points for VM copies. If you schedule to run a VM copy job periodically, every new run will overwrite the existing copy.
- With the VM copy job, all VM disks are copied as thick, while replication allows you to preserve the format of disks or convert the disk format on the fly.
- There are no failover or failback possibilities for a VM copy.

VM copy jobs use the same infrastructure components as backup jobs (for details, see Backup Architecture). In addition to available scenarios, you can also copy VMs to a target folder on any server or host connected to the Veeam backup server.

File Copy

Veeam Backup & Replication includes file copy possibilities, providing a natural way to deliver image files to hosts, make backup copies of existing VMs, exchange VMs and templates between servers or move backups across repositories. Using Veeam Backup & Replication, you can copy files and folders between and within servers connected to the Veeam Backup Server.

Note When file copy destination is located on a server managed by Veeam Backup & Replication, traffic compression can be used to minimize network bandwidth and improve performance of file copy activities.

Quick Migration

Veeam Quick Migration enables you to promptly migrate one or more VMs between ESX(i) hosts and datastores. Veeam Backup & Replication allows migration of VMs in any state with minimum disruption to business operations and end user access to services. You can use Quick Migration as a self−contained capability, solely for VM migration, or combine it with Instant VM Recovery.

Veeam Backup & Replication analyzes your virtual environment, its configuration, the state of VMs and selects the most appropriate relocation method. Whenever possible, Veeam Backup & Replication coordinates its operations with vCenter Server and uses native VMware vCenter migration mechanisms: vMotion and Storage vMotion. When VMware vCenter migration methods cannot be used (for example, if your VMware vSphere license does not provide support for vMotion and Storage vMotion, or you need to migrate VMs from one standalone ESX(i) host to another), Veeam Backup & Replication uses its proprietary SmartSwitch technology to relocate VMs.

Veeam Quick Migration provides means for fast background migration of VMs ensuring continuous uptime of your virtual environment. Quick Migration supports hot VM migration (with SmartSwitch) and cold VM migration (with cold switch).

Migration of a VM is performed in several stages:

1. Veeam Backup & Replication copies VM configuration (.vmx) to the target host and registers the VM.
2. Veeam Backup & Replication triggers a VM snapshot and copies VM disk content to the new destination.
3. VM state and changes made after snapshot creation are moved to a new location. Veeam Backup & Replication uses different approaches to move the VM state between hosts with compatible and non−compatible CPUs.
    - If you move a VM between two hosts with compatible CPUs, Veeam Backup & Replication uses SmartSwitch (that is, it suspends a VM to move its state file and changes made after snapshot creation). The VM is then resumed on the new host. This ensures minimum downtime, and completely eliminates any data loss during migration.
    - If you move a VM between two hosts with non−compatible CPUs, Veeam Backup & Replication stops the VM to move changes made after snapshot creation, and then starts the VM on the new host.

Integration with Instant VM Recovery

When you restore a VM using Instant VM Recovery, Veeam Backup & Replication starts the VM directly from a compressed and deduplicated backup file. To finalize recovery of a VM, you still need to move it to a new location. Moving the VM with VMware Storage vMotion or hot replication may require a lot of time and resources, or it may cause loss of valuable data.

Veeam Quick Migration was designed to complement Instant VM Recovery. Instead of pulling data from vPower NFS datastore, Quick Migration registers the VM on the target host, restores the VM contents from the backup file located in the backup repository and synchronizes the VM restored from backup with the running VM.

Quick Migration Architecture

Quick Migration architecture in a VMware vSphere environment comprises the following components:

- Source host and target host with associated datastores
- One or two backup proxy servers

Similar to backup, Quick Migration uses two-service architecture: the source–side Veeam transport service interacts with the source host, and the target–side Veeam transport service interacts with the target host. To perform onsite migration, you can deploy one backup proxy for data processing and transfer. This backup proxy must have access to the source host and to the target host at the same time. In this scenario, the source–side transport service and the target–side transport service are started on the same backup proxy.

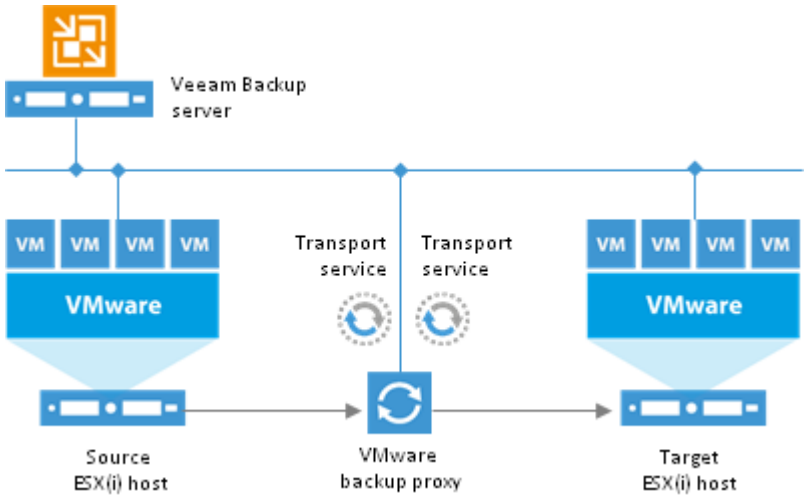

The common requirement for offsite migration is that one transport service runs in the production site (closer to the source host and datastore), and the other transport service runs in the remote target site (closer to the target host and datastore). During backup, the transport services maintain a stable connection, which allows for uninterrupted operation over WAN or slow links.

For offsite migration, you need to deploy at least one local backup proxy in each site: a source backup proxy in the production site, and a target backup proxy in the remote target site.

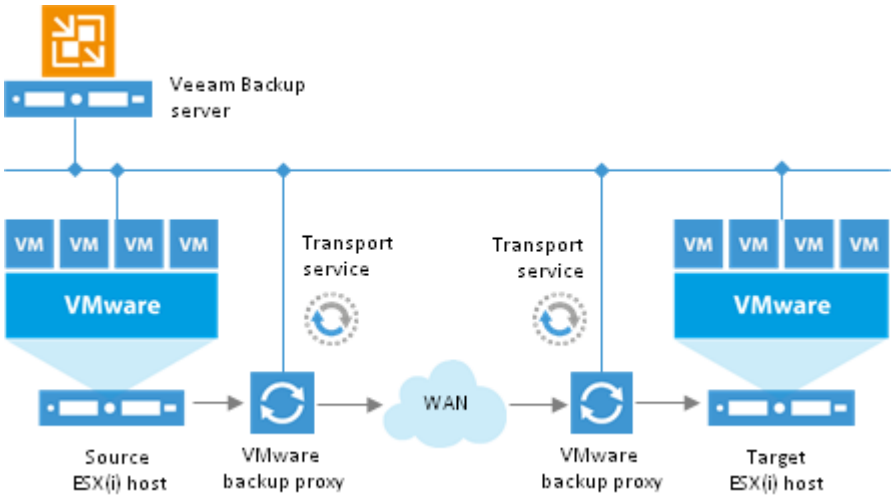
Veeam Backup server
VM
VM
VM
VM
VMware
Transport service
Transport service
VM
VM
VM
VM
VMware
WAN
Source ESX(i) host
VMware backup proxy
VMware backup proxy
Target ESX(i) host

HP SAN Support

Veeam Backup & Replication lets you leverage SAN snapshots as a part of a comprehensive backup and recovery strategy, where SAN snapshots and image-level backups complement each other. With Veeam Backup & Replication, you can:

- Perform backup from HP SAN storage snapshots
- Restore data directly from HP SAN storage snapshots

Veeam Backup & Replication supports the following HP SAN storage systems:

- HP StoreVirtual
- HP StoreVirtual VSA
- HP StoreServ

Backup from Storage Snapshots

Starting from version 7, Veeam Backup & Replication offers a new capability — Backup from Storage Snapshots — that lets you dramatically improve RPOs and reduce impact of backup activities on the production environment. Backup from Storage Snapshots is supported for HP StoreVirtual Storage and HP 3PAR StoreServ Storage systems and is available for data protection in the VMware vSphere environment.

Note The Backup from Storage Snapshots functionality is available only in the Enterprise Plus Edition of Veeam Backup & Replication.

VM Data Processing

In the regular processing course, when Veeam Backup & Replication backs up or replicates a VM, it triggers a VMware snapshot. The snapshot "freezes" the state and data of the VM at a specific point in time. This way, the VM data is brought to a consistent state suitable for backup.

The procedure of backup or replication is the following:

1. VM disks are put to the read-only state.
2. Every virtual disk receives a delta file named like *vmname-00001.vmdk*.
3. Veeam Backup & Replication starts copying VM data from read-only disks of the VM. All changes that the user makes to the VM are written to delta files.
4. When the backup or replication job is over, the snapshot is committed: the VM disks resume writes, data from the delta file is merged to the VM disks and the snapshot is removed.

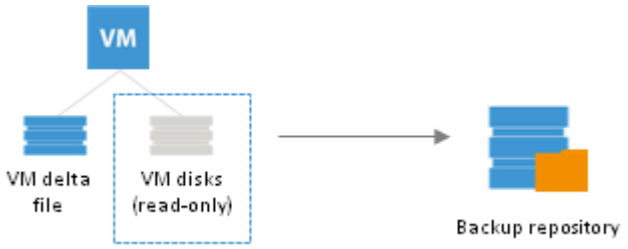

Sometimes VM data processing may take long: for example, if you back up a very large VM. And if backup or replication is performed for a VM running a highly transactional application, the delta file will grow very large, too. Consequently, the snapshot commit process will take much time and the VM may even freeze during this process.

To overcome this situation, Veeam Backup & Replication introduces Backup from Storage Snapshots.

VM Processing with Backup from Storage Snapshots

The Backup from Storage Snapshots technology lets you leverage HP SAN snapshots for VM processing and speed up backup and replication operations.

The procedure of backup or replication from storage snapshots is the following:

1. Veeam Backup & Replication first triggers a VMware snapshot for VMs whose disks are located on the HP SAN storage.
2. Veeam Backup & Replication triggers a HP SAN snapshot of the volume holding the VM and the VM snapshot.
3. The VMware snapshot on the original HP SAN volume is immediately deleted after that. To back up or replicate VM data, Veeam Backup & Replication accesses the 'cloned' VMware snapshot on the HP SAN snapshot.

As a result, the VMware snapshot exists for a very short time, namely for several seconds. Delta files do not grow large and the time of snapshot commit decreased up to 20 times.

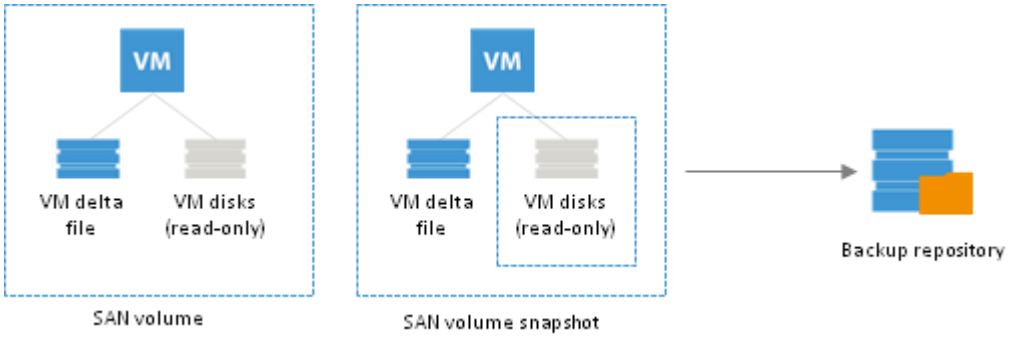

Important! If Backup from Storage Snapshots cannot be used due to incorrect infrastructure setup, Veeam Backup & Replication will not fail over to the regular processing mode and the backup or replication job will be finished with the *Failed* status.

How It Works

From the user's side, Backup from Storage Snapshots is an option that can be enabled for a specific backup or replication job. When you enable this option for the job, you define that VM data processing should be performed with use of HP SAN snapshots.

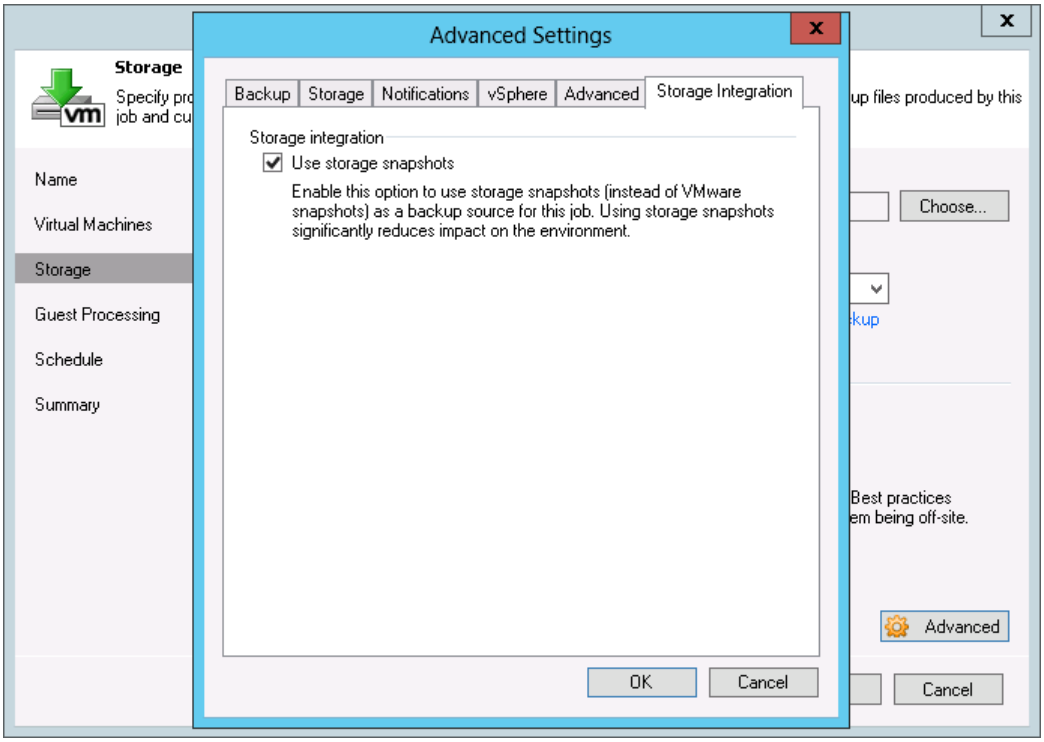

Veeam Backup & Replication can use the Backup from Storage Snapshots functionality only if you have met the following conditions:

- You have added the SAN storage system to the Veeam Backup & Replication console.
- You have the **Use storage snapshots** option enabled in the job settings.
- You have properly configured a VMware backup proxy. To learn more, see Configuring a VMware Backup Proxy for HP SAN Snapshots.

The Backup from Storage Snapshots functionality is used only for those VMs whose disks are located on the HP SAN storage.

- If the job includes a number of VMs whose disks are located on different types of storage, Veeam Backup & Replication will apply the Backup from Storage Snapshots option only to VMs with disks on the HP SAN storage.
- If a VM has several disks, some on the HP SAN and some on the other type of storage, Veeam Backup & Replication will not use Backup from Storage Snapshots for this VM and will process it in a regular manner.

As the backup or replication job typically contains a number of VMs, Veeam Backup & Replication processes such VMs in different ways. First, Veeam Backup & Replication triggers VMware and HP SAN snapshots for VMs on the HP SAN storage. Only after the HP SAN volume snapshot has been created, Veeam Backup & Replication triggers a VMware snapshot for other VMs. Therefore, the procedure of backup or replication from storage snapshots is the following:

1. Veeam Backup & Replication analyzes which VMs in the job have disks on the HP SAN storage and which VMs have disks on another type of storage, and groups VMs of two types.
2. Veeam Backup & Replication triggers a VMware snapshot for VMs whose disks are located on the HP SAN volume.

3. Veeam Backup & Replication triggers a snapshot of the HP SAN volume holding VM disks and VM snapshots.
4. Veeam Backup & Replication gets the Changed Block Tracking information about changed data blocks for VMs whose disks are located on the HP SAN storage.
5. Veeam Backup & Replication removes the VMware snapshot on the production HP SAN volume. The cloned VMware snapshot still remains on the HP SAN volume snapshot. After that, Veeam Backup & Replication triggers a VMware snapshot for VMs whose disks are located on another type of storage. These VMs are processed in the regular manner further on, in parallel with VMs whose disks are located on the HP SAN storage.
6. Veeam Backup & Replication detects if there is a VMware backup proxy having a direct connection to the HP SAN storage system in the backup infrastructure.
   - For HP StoreVirtual and HP StoreVirtual VSA, iSCSI connection is required
   - For HP StoreServ Storage, Fibre Channel connection is required.

   When such VMware backup proxy is detected, Veeam Backup & Replication mounts the HP SAN volume snapshot as a new volume to this backup proxy.
7. Veeam Backup & Replication reads and transports VM data blocks via the VMware backup proxy to the backup repository. In the incremental backup or replication course, Veeam Backup & Replication uses the CBT data to retrieve only changed data blocks.
8. When the processing is finished, Veeam Backup & Replication unmounts the HP SAN snapshot from the VMware backup proxy and issues a command to the HP SAN to remove the HP SAN volume snapshot.

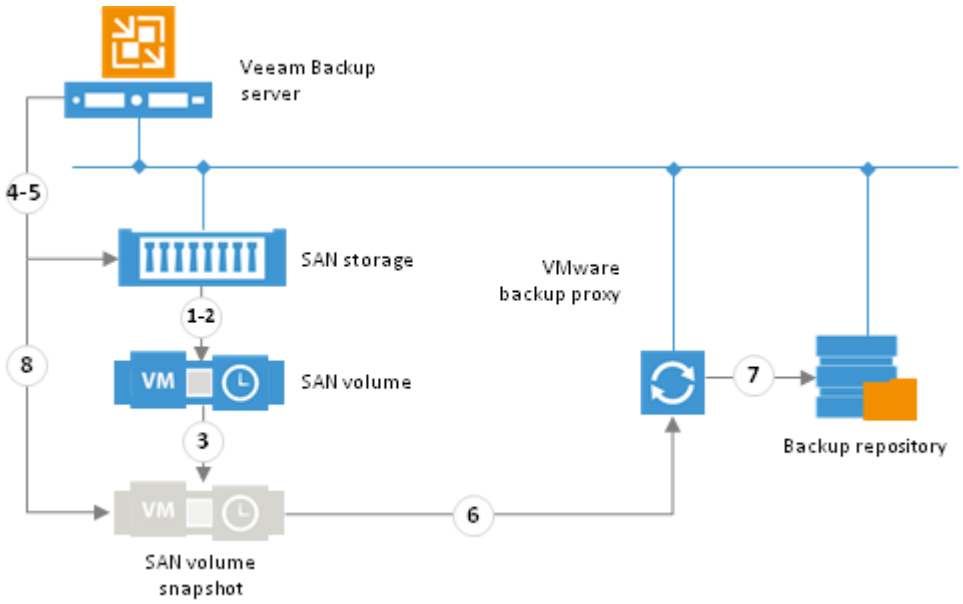

Veeam Explorer for SAN Snapshots

Many organizations use SAN snapshots for data protection. SAN snapshots allow for very low RPO: they have minimal impact on storage performance and can be created really fast. Administrators can take snapshots several times a day or even schedule them as often as every hour.

However, in the virtual environment, SAN snapshots add more difficulty to the restore process. SAN snapshots are created per-volume. A volume typically holds multiple VMs. For this reason, restore from SAN snapshots is not a simple rollback operation; it is a multi-task process. To restore a VM from the SAN snapshot manually, you need to perform the following actions:

1. Present the SAN snapshot to the ESX(i) host.
2. Perform an HBA rescan.
3. Mount the SAN snapshot to an ESX(i) host.
4. Browse the SAN snapshot to locate the VM files (VMDK).
5. Add the VM to the inventory or copy VM files to another VMFS datastore.
6. Power on the VM.
7. Perform restore operations.
8. Perform cleanup operations after the recovery is completed.

If you need to restore guest OS files and application objects from the VM on a SAN snapshot, the procedure will be even more complicated. As a result, the restore process takes much time.

To make VM recovery from SAN fast and easy, Veeam Backup & Replication offers Veeam Explorer™ for SAN Snapshots. Veeam Explorer for SAN Snapshots is a technology in Veeam Backup & Replication that lets you restore VMware VM data directly from SAN snapshots on HP LeftHand, HP StoreVirtual VSA and HP 3PAR StoreServ. Veeam Explorer for SAN has been designed and developed in collaboration with HP and it uses the native HP APIs.

Veeam Explorer for SAN Snapshots offers a variety of restore options:

- You can instantly restore an entire VM
- You can restore VM guest OS files (Windows, Linux, FreeBSD and other)
- You can restore Microsoft Exchange objects from HP SAN snapshots
- You can restore Microsoft SharePoint objects from HP SAN snapshots

Benefits of Veeam Explorer for SAN Snapshots

Veeam Explorer for SAN Snapshots lets you leverage the low overhead of SAN snapshots and use flexible restore options of Veeam Backup & Replication. While SAN snapshots provide good RPOs, Veeam Explorer for SAN Snapshots allows for very short RTOs. Veeam Explorer for SAN Snapshots:

- Reduces time to mount snapshots up to 10 times or more.
- Eliminates the need of staging and intermediate restores. You only need to select an ESX(i) host to which the HP SAN snapshot should be mounted and Veeam Backup & Replication will perform all other operations for you.
- Eliminates human errors that can potentially occur during the mount process.
- Leverages advantages of the storage system that you already have in place.

How It Works

Veeam Backup & Replication fully automates the operation of mounting the SAN snapshot to the ESX(i) host. You do not need to install additional agents or perform complex configuration actions. Veeam Backup & Replication does not convert HP SAN snapshots into backups. Instead, it uses them "as is" and lets you restore VM data directly from native HP SAN snapshots.

When you start the restore procedure, Veeam Backup & Replication performs the following actions "behind the scene":

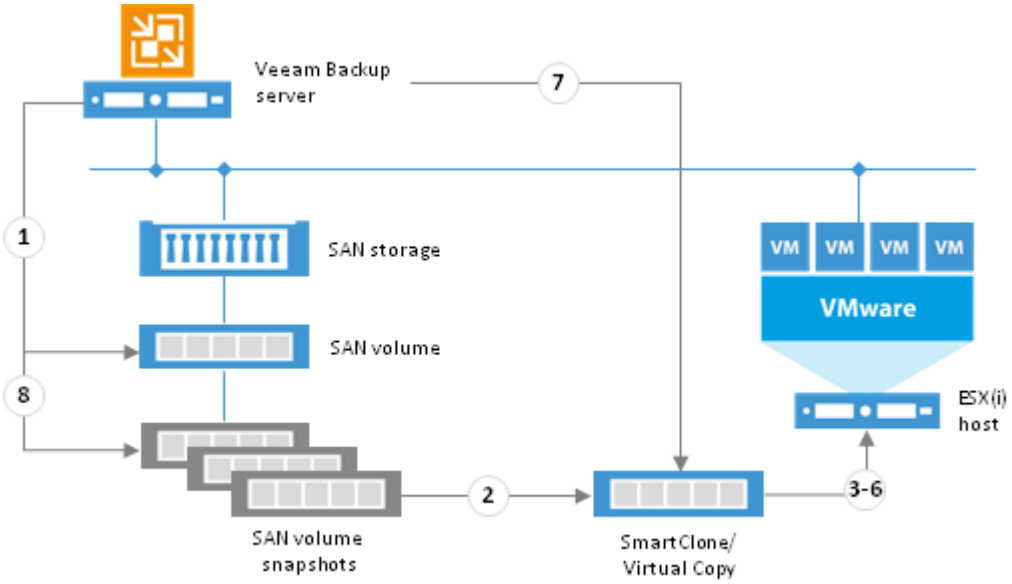

1. Veeam Backup & Replication issues a command to the HP SAN using the native HP APIs.
2. The HP SAN creates a SmartClone (for HP P4000) or a Virtual Copy (for HP 3PAR StoreServ) of the requested volume snapshot. The SmartClone/Virtual Copy is a read-write snapshot of the volume snapshot. For restore operations, Veeam Backup & Replication uses the created SmartClone/Virtual Copy, not the volume snapshot itself.
   The SmartClone/Virtual Copy is used to protect the VMFS volume metadata integrity on the LUN. During FLR and Instant VM Recovery operations, the ESX(i) host working with the datastore updates VMFS metadata on the LUN. Use of the SmartClone/Virtual Copy helps protect the volume snapshot from these changes.
3. The ESX(i) host to which the user wants to mount the SmartClone/Virtual Copy is added to the list of *Allowed Servers* for the SmartClone. As a result, the ESX(i) host has access to the SmartClone/Virtual Copy and can read/write data to/from it.
4. The HP SAN makes sure that the IP address of the HP SAN storage is in the list of static targets on the ESX(i) host. By default, the HP SAN uses the IP address of the HP cluster on which the SAN volume is allocated.
5. The HP SAN issues a HBA re-scan command to the vCenter Server. On re-scanning, the SmartClone/Virtual Copy appears in the discovered targets list on the ESX(i) host.
6. The HP SAN performs re-signature for LUN volumes.
7. Veeam Backup & Replication performs recovery operations.
8. After recovery is completed, Veeam Backup & Replication issues a command to the HP SAN. The HP SAN deletes the SmartClone/Virtual Copy volume from the ESX(i) host and performs cleanup operations.

From user's side, the restore process is very fast and easy. Restore takes a couple of clicks in the Veeam Backup & Replication interface and requires very little time. For example, to restore a VM or a guest OS file, you will need 2 minutes or less, compared to 20-30 minutes required for manual restore from SAN snapshots without Veeam Backup & Replication.

vCloud Director Support

Backup and restore of vCloud Director vApps and VMs has always been a hot topic. Up to now backup tools offered no option of backup in the vCloud Director environment. The only way was to perform backup at the level of the underlying vCenter Server. For restore, the administrators would first need to restore VMs to the vCenter Server level and then bring them to vCloud Director through import.

With version 7.0, Veeam Backup & Replication provides support for vCloud Director. It uses vCloud Director API to help you back up vApps and VMs and restore them directly to the vCloud Director hierarchy.

The main entity with which Veeam Backup & Replication works during backup is a vApp. A vApp is a virtual system that contains one or more individual VMs along with parameters that define operational details — vApp metadata. When Veeam Backup & Replication performs backup of VMs, it captures not only data of VMs being a part of vApps, but also vApp metadata. As a result, you can restore vCloud Director objects back to the vCloud Director hierarchy and do not need to perform any additional actions on import and VM configuration.

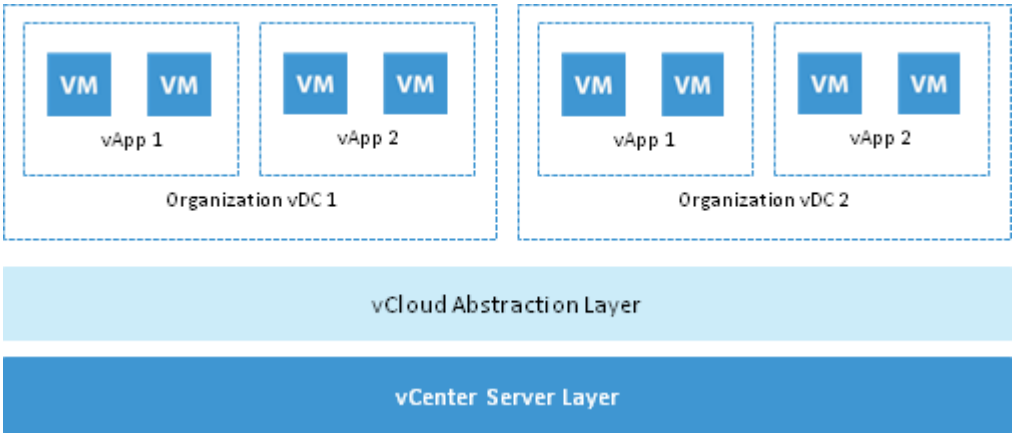

Backup and Restore of vApps

Veeam Backup & Replication provides you with an option to back up vCloud Director vApps and restore them back to the vCloud Director hierarchy.

In terms of vCloud Director, a vApp is a coherent system that contains one or more VMs. Every vApp is described with a set of operational details, or vApp metadata, that include the following ones:

- vApp owner settings
- Access rights settings
- vApp network settings: information about organization networks to which the vApp is connected
- Lease settings and so on

When Veeam Backup & Replication performs backup of a vApp, it backs up all VMs being a part of this vApp along with the vApp metadata. Backup of the vApp is performed with the vCD backup job. The vCD backup job may contain one or several vApps. If necessary, you can exclude specific VMs and VM disks from the backup when configuring a vCD backup job.

Veeam Backup & Replication offers the following restore options for backed up vApps:

- Restoring vApps to vCloud Director
- Restore of separate VMs being a part of the vApp to vCloud Director

Note Just like vCloud Director, Veeam Backup & Replication treats a vApp as a coherent system. For this reason, it is recommended that you add entire vApps, not separate VMs from the vApp, to the vCD backup job. If you do not want to back up specific VMs in the vApp, you can use exclusion settings in the vCD job.

Backup of vCloud Director VMs

Veeam Backup & Replication lets you perform backup for vApps and VMs, as well as VM containers in vCloud Director such as Organization vDC, Organization and even the vCloud Director instance.

When Veeam Backup & Replication performs backup of vApps and VMs, it additionally captures vApp metadata.

vApp metadata includes:

- General information about the vApp where VMs reside, such as: vApp name, description, VMs descriptions
- Information about vApp networks and organization networks to which the vApp is connected
- VMs startup options
- User information
- Lease
- Quota
- Storage template and so on

vApp metadata is stored together with the VM content. Capturing vApp metadata is extremely important for restore: without it, you will not be able to restore vApps and VMs back to vCloud Director.

Data to Back Up

With Veeam Backup & Replication, you can back up regular VMs and linked clone VMs.

Backup of Regular VMs

When you perform backup of regular VMs, Veeam Backup & Replication captures and stores to the backup file the following data:

- VM disk content
- vApp metadata
- VM metadata

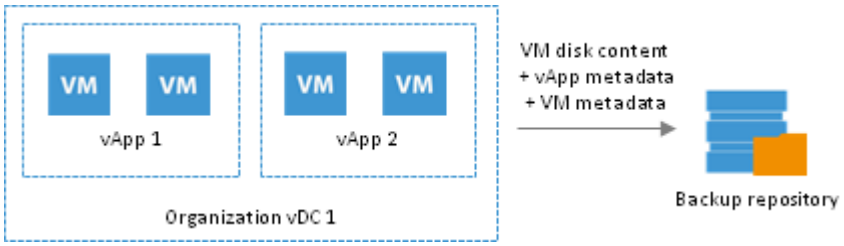

Backup of Linked Clone VMs

When you perform backup of linked clone VMs, Veeam Backup & Replication captures and stores to the backup file the following data:

- Content of the template to which the VM is linked
- Content of the VM user disk — delta disk
- vApp metadata
- VM metadata

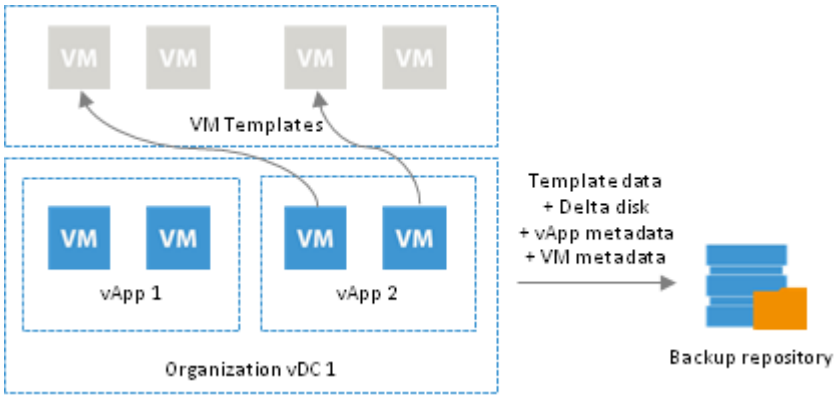

During full backup of linked clone VMs, Veeam Backup & Replication consolidates data of the VM template and delta disk and saves it as a regular VM disk in the backup file. Data merging guarantees proper VM restore: even if a VM template is lost by the time of recovery, you will still be able to restore the linked clone VM from the backup.

During incremental backup, Veeam Backup & Replication saves only changed data of the delta file.

vCD Backup Jobs

For VMs managed by vCloud Director, Veeam Backup & Replication offers a special type of the backup job — vCD backup job. vCD backup jobs have been specifically developed to process vCloud Director objects, ensure their proper restore and support of vCloud-specific features.

You should always use vCD backup jobs to back up VMs managed by vCloud Director. If you back up VMs managed by vCloud Director using a regular backup job, Veeam Backup & Replication will perform backup at the level of the underlying vCenter Server and will not capture vApp metadata. As a result, you will not be able to restore a fully-functioning VM to vCloud Director.

Restore of vCloud Director VMs

Veeam Backup & Replication enables full-fledged restore of VMs to vCloud Director. You can restore separate VMs to vApps, as well as VM data.

For restore, Veeam Backup & Replication uses VM metadata saved to a backup file and restores specific VM attributes. As a result, you get a fully-functioning VM in vCloud Director, do not need to import the restored VM to vCloud Director and adjust the settings manually.

Backed up objects can be restored to the same vCloud Director hierarchy or to a different vCloud Director environment. Restore options include:

- Instant VM recovery
- Full restore for vApps and VMs
- Restore of VM disks
- Restore of VM files
- Guest OS file-level restore for VMs

Restoring Regular VMs to vCloud Director

If you restore regular VMs back to the vCloud Director hierarchy, the restore process includes the following steps:

1. Veeam Backup & Replication uses the captured vApp metadata to define the vApp settings and VM initial location in the vCloud Director hierarchy.
2. Veeam Backup & Replication restores VMs from the backup file to their initial location or to other location. Additionally, Veeam Backup & Replication restores all VM settings.

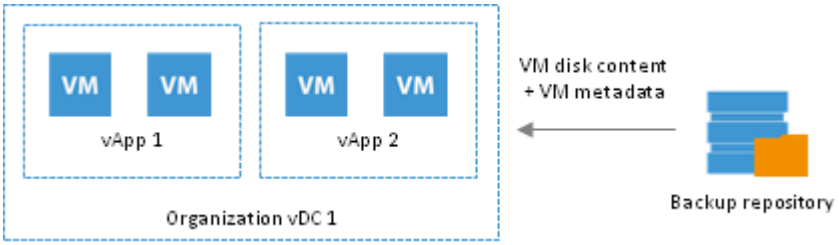

Restoring Linked Clone VMs to vCloud Director

Veeam Backup & Replication lets you restore linked clone VMs – VMs that were deployed from a VM template using the fast provisioning technology. There are several mechanisms for processing linked clone VMs.

Restore of Existing VMs

If you are restoring a vCD linked clone VM that exists in the vCloud Director hierarchy, the restore process includes the following steps:

1. Veeam Backup & Replication uses the captured vApp metadata to define the initial settings of the VM.
2. Veeam Backup & Replication calculates a signature for the consolidated VM disk in the backup file (containing the VM template data and data of the delta file) and the signature for the VM existing in vCloud Director. Veeam Backup & Replication then compares the disk signatures to define what data blocks have changed.
3. Veeam Backup & Replication restores only changed data blocks from the backup file and writes them to the user delta file.

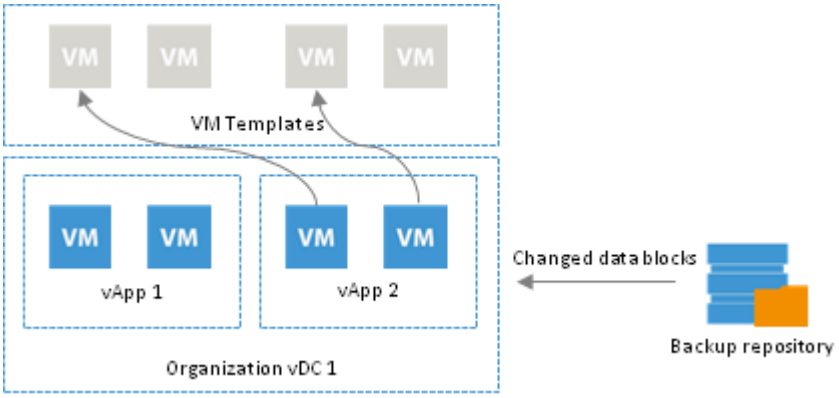

Restore of Deleted VMs

If you are restoring a VM that no longer exists in vCloud Director hierarchy, the restore process includes the following steps:

1. Veeam Backup & Replication uses vCloud Director to create a new linked clone VM from the VM template that the user selects. The new VM has a blank user delta file.
2. Veeam Backup & Replication calculates a signature for the consolidated VM disk in the backup file (containing the VM template data and data of the delta file) and the signature for the created VM in vCloud Director. Veeam Backup & Replication then compares the disk signatures to define what data blocks need to be restored.
3. Veeam Backup & Replication restores only those data blocks that need to be restored from the backup file and writes them to the blank user delta file.

By default, Veeam Backup & Replication links the VM to the same VM template that was used by the initial VM. During restore, Veeam Backup & Replication checks the settings of the VM template to which the restored VM is linked: verifies connection settings, makes sure the disk size coincide and so on.

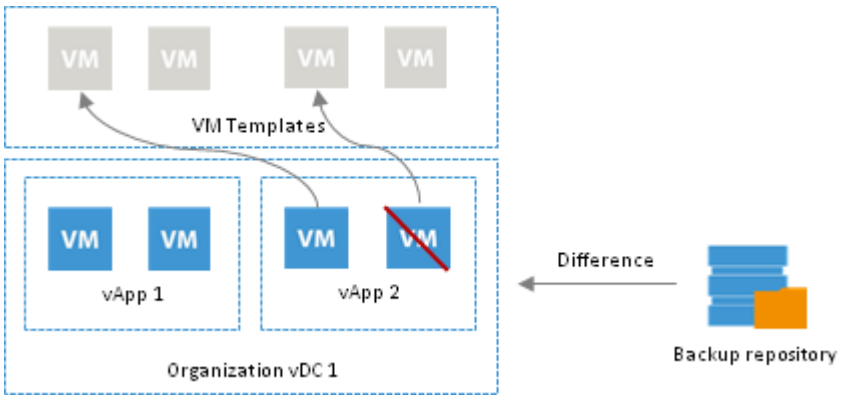

Restore of Linked Clone VMs as Regular VMs

In some cases, Veeam Backup & Replication can restore a VM from a backup file as a regular VM. This type of restore is accomplished in the following situations:

- You have intentionally chosen to restore a linked clone VM as a regular VM.
- You are restoring a VM to the Organization vDC which has the fast provisioning option disabled.
- A VM template to which the restored VM should be linked is not accessible in the location to which the VM is restored.

In this case, Veeam Backup & Replication uses the same algorithm as for restore of full VMs in the virtual environment. It retrieves the data of the consolidated VM disk from the backup file and restores the VM in the vCloud Director hierarchy.

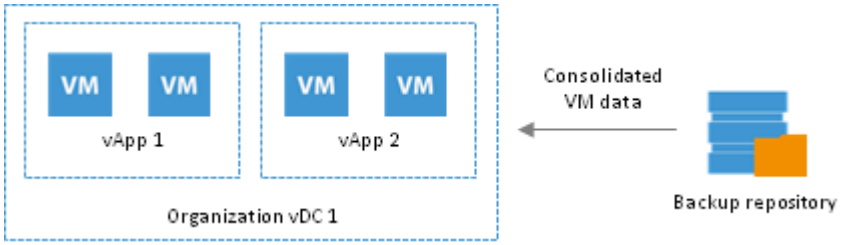

Backup Copy

The main backup purpose is to protect your data against disasters and VM failures. However, having one copy of a backup file does not provide the necessary level of safety. A backup file may get corrupted or lost, leaving you with not data to restore at all.

Backup experts claim that to build a successful data protection and disaster recovery plan, you must follow the 3-2-1 rule:

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

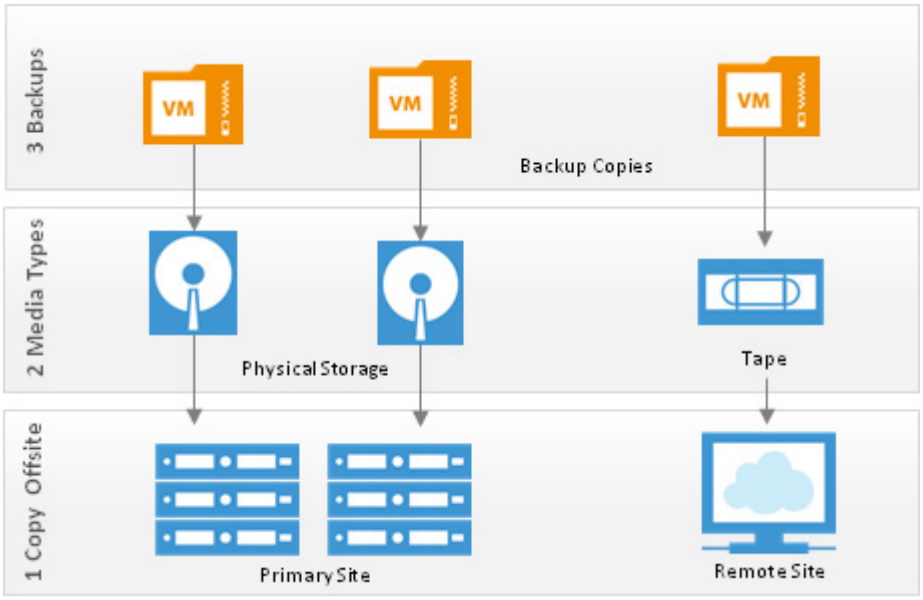

Thus, according to the first statement of the 3-2-1 backup strategy, you must have three different copies of a backup file in different locations. In case a disaster strikes, multiple backup copies increase your chances in data restore.

To let you adopt the 3-2-1 backup strategy, Veeam Backup & Replication offers backup copying capabilities. Backup copying allows you to create several instances of the same backup file in different locations, whether onsite or offsite. Copied backup files have the same format as those created by backup jobs and you can use any data recovery option for them.

Backup copy is a job-driven process. Veeam Backup & Replication fully automates the backup copying process and lets you specify retention policy settings to maintain the desired number of restore points for copied backups.

Backup Copying Process

Backup data is copied per VM at the block level. When the backup copying process starts, Veeam Backup & Replication accesses VM backup files in the source backup repository, retrieves data blocks for a specific VM from the backup file, copies them to the target backup repository and composes copied blocks into a backup file in the target backup repository. Therefore, the backup copying process does not affect virtual infrastructure resources, does not require an additional snapshot of a VM and does not produce any load on VMs whose backups are copied.

In the target backup repository, the backup copy job creates a chain of restore points using the incremental backup method. The target backup repository always contains only one active incremental backup chain. Restore points in the chain are rotated according to the specified retention policy. To learn more, see Retention Policy for Backup Copy Jobs.

The backup chain on the target backup repository is created in the following manner:

1. The first synchronization interval of the backup copy job always produces a full backup file. The full backup file is created in the following way:
   a. From the backup chain on the source backup repository, Veeam Backup & Replication copies data blocks that are necessary to build a full backup of a VM as of the most recent state. Data blocks can be copied from one or several backup files in the chain. If the backup chain on the source backup repository was created using the reversed incremental backup method, Veeam Backup & Replication simply copies data blocks of the latest full backup.

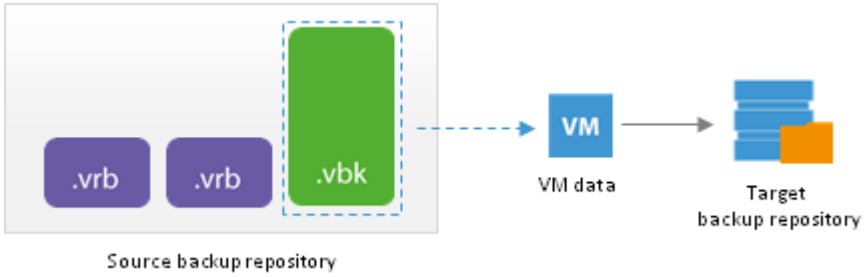

   If the backup chain on the source backup repository was created using the forward incremental backup method, Veeam Backup & Replication copies data blocks from the first full backup and a set of incremental backups to form a full backup of a VM as of the most recent state.

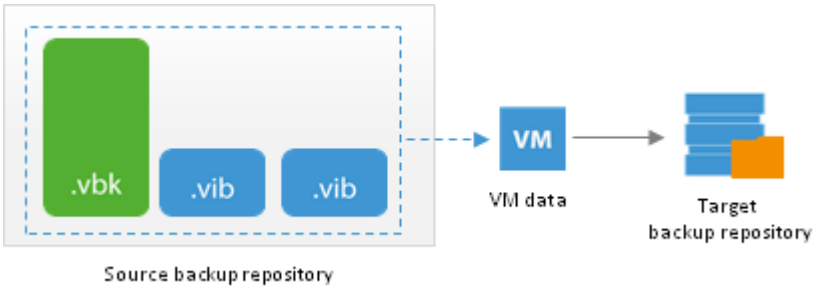

   b. On the target backup repository, Veeam Backup & Replication writes all copied data blocks to the same full backup file.

2. At every next synchronization interval, when a new restore point appears on the source backup repository, Veeam Backup & Replication copies incremental changes from this most recent restore point and transfers them to the target backup repository. On the target backup repository, Veeam Backup & Replication writes the copied data blocks to the incremental backup file.

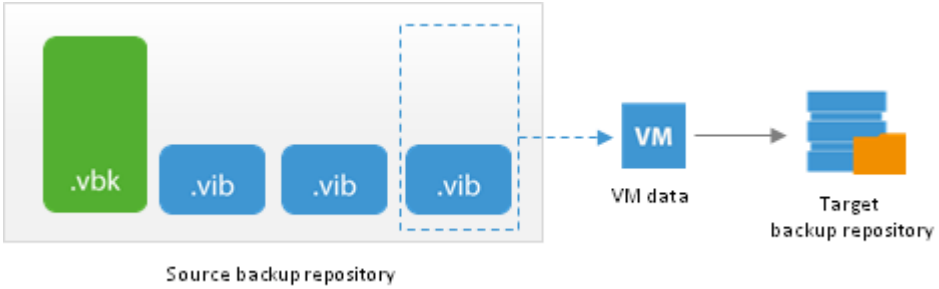

The backup copy job can be created for one VM or several VMs. If the backup copy job is created for several VMs, you can define the order in which the VMs should be processed.
Veeam Backup & Replication will subsequently process VMs one by one in the defined order. If any VM cannot be processed for some reason, for example, in case a new restore point for this VM is not yet available, Veeam Backup & Replication will pass to the next VM. Once this VM is processed, Veeam Backup & Replication will attempt to copy the unprocessed VM once again.

Even if a backup copy job processes several VMs, it creates one backup file on the target backup repository and stores to it data for all VMs processed by the job.

Note Backup copy jobs do not support parallel processing. Multiple VMs in the job are copied one by one, subsequently. Data for VM disks are also copied subsequently, not in parallel.

To minimize the amount of traffic going over the network, Veeam Backup & Replication uses the data compression and deduplication technologies. To learn more, see Compression and Deduplication.

Restore Point Selection

Veeam Backup & Replication does not necessarily use a backup created by one job and one backup repository as a source of data. It can copy VM data from backups created by different jobs and even from different backup repositories. When you set up a backup copy job, you only define what VM(s) you want to process. During the backup copy job, Veeam Backup & Replication searches for the most recent restore point in all available backup repositories, copies data blocks from it and saves them to a backup file on the target backup repository.

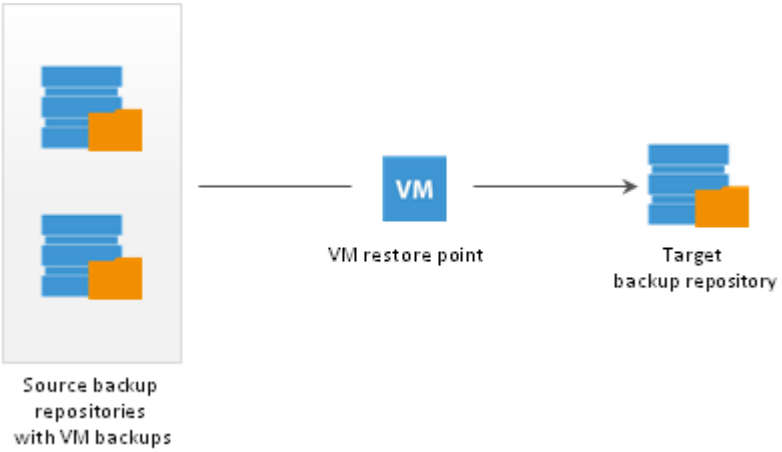

You can specify a search scope for the backup copy job: that is, define in which backup repositories Veeam Backup & Replication should search for restore points. In this case, Veeam Backup & Replication will skip all other backup repositories from searching.

Veeam Backup & Replication always copies the most recent restore point from the source backup repository. Even when backup copying is performed for the first time and the source backup repository already contains a chain of restore points, Veeam Backup & Replication will only copy a restore point containing data as of the most recent VM state. To learn more, see Backup Copying Process.

Veeam Backup & Replication identifies new restore points using the following rule:

Time of restore point creation >= current time – synchronization interval

For example, you have set the synchronization interval to 24 hours. Today's date and time are 7/1/2013, 12:00 PM and the restore point was created 23 hours ago, on 6/30/2013 at 1:00 PM. In this case, Veeam Backup & Replication will copy this new restore point, because:

6/30/2013, 1:00 PM >= 7/1/2013, 12:00 PM – 24 hours

The rule above is applied to all synchronization intervals, both the first one, copying a full backup file, and subsequent ones, copying incremental restore points. After you create a backup copy job and the first synchronization interval starts, Veeam Backup & Replication checks if there is some restore point falling into the necessary search scope on the source backup repository. If there is no restore point matching this condition, Veeam Backup & Replication will not copy data from the source backup repository. Instead, it will wait for the new restore point to appear on the source backup repository. Only after that Veeam Backup & Replication will copy the first, full restore point, to the target repository. This mechanism helps ensure that the backup chain produced by the backup copy job contains only the most recent VM data.

The backup copy job has the following limitations:

1. Veeam Backup & Replication does not copy restore points from the target backup repository.
2. Veeam Backup & Replication does not copy restore points from imported backups.
3. Veeam Backup & Replication does not copy restore points that have already been copied by the same backup copy job to the target backup repository.
4. Veeam Backup & Replication does not copy corrupted restore points.
5. Veeam Backup & Replication does not copy restore points that are locked by some tasks: for example, a backup job creating a backup chain with the reversed incremental method or a restore process.
6. Veeam Backup & Replication does not copy restore points if the block size of the restore point on the source backup repository differs from the block size of restore points on the target backup repository.
   The data block size for restore points on the target backup repository is set at the first synchronization cycle of the backup copy job. This size is taken from the corresponding settings of the primary backup job — the backup job that creates the backup chain on the source backup repository. If after the first synchronization cycle you add to the backup copy job new sources that use a different data block size, Veeam Backup & Replication will detect such restore points and display the *Restore point is located in backup file with different block size* message.
7. If you select a backup job as a source for the backup copy job, Veeam Backup & Replication will only copy restore points created by this very backup job. Veeam Backup & Replication will not perform search in other backup repositories.

Tip You can configure several backup copy jobs to copy one restore point from the source backup repository to different target locations.

Data Transport Path

To transport data from the source backup repository to the target backup repository, the backup copy job uses one of the following paths:

- **Direct transport path**: Veeam Backup & Replication transports data directly from the source backup repository to the target backup repository. This type of data transport is recommended for copying backups to onsite backup repositories or offsite backup repositories over fast connections.
  When Veeam Backup & Replication uses the direct transport path, it starts Veeam Transport Services on the following backup infrastructure components:
  - In case of Windows- and Linux based repositories: the source Veeam Transport Service is started on the source backup repository; the target Veeam Transport Service is started on the target backup repository.

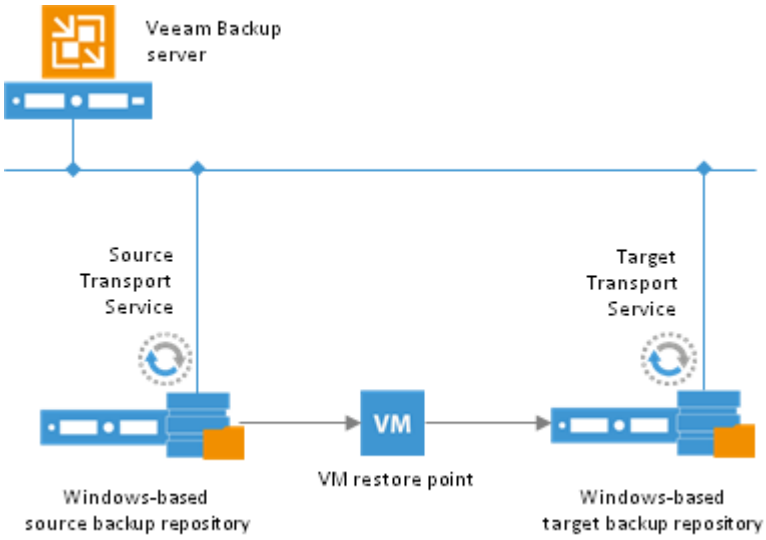

  - In case of CIFS share: the source Veeam Transport Service is started on the proxying server in the source site; the target Veeam Transport Service is started on the proxying server on the target site.

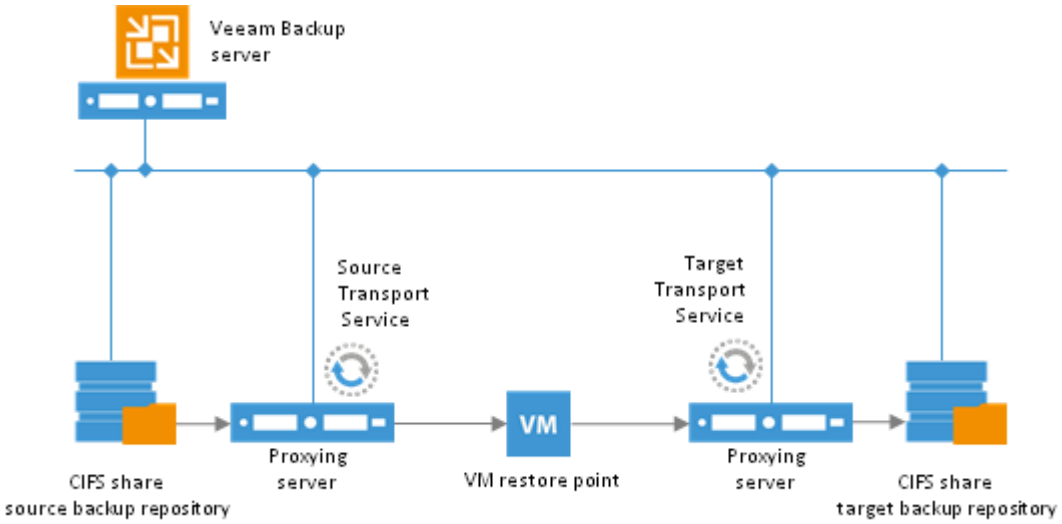

- **Through built-in WAN accelerators**: Veeam Backup & Replication transports data through a pair of WAN accelerators: one deployed on the source side and the other one deployed on the target side. WAN accelerators remove redundant blocks before transferring VM data and thus significantly reduce the amount of traffic going over the network. This type of data transport is recommended for copying backups offsite over slow connections or WAN.

When Veeam Backup & Replication uses the transport path via WAN accelerators, it starts the source Veeam Transport Service on the source backup repository (in case of Windows- and Linux based repositories) or on the proxying server in the source site (in case of a CIFS share). The target Veeam Transport Service is started on the target backup repository (in case of Windows- and Linux based repositories) or on the proxying server in the target site (in case of a CIFS share).

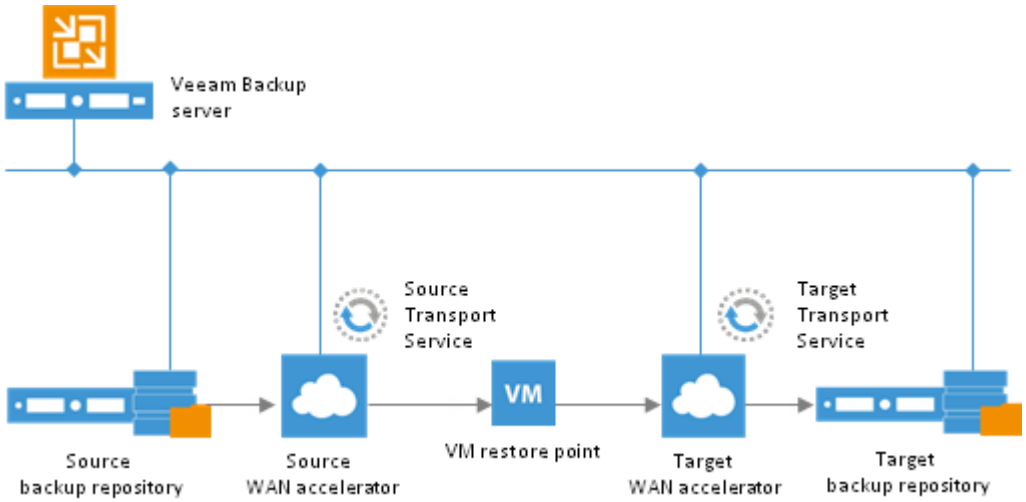

Important! The WAN acceleration technology is available only in the Enterprise Plus edition of Veeam Backup & Replication. To learn more, see WAN Acceleration.

Backup Copy Job

The backup copy job is a separate task that needs to be set apart from the backup job.

The aim of the backup copy job is to copy a VM restore point from the source backup repository to the target backup repository. Every backup copy job creates its own folder on the target backup repository and stores to it all copied restore points. The folder has the same name as the backup copy job.

The backup copy job runs continuously and has several phases:

- **Idle state**. For the most time, the backup copy job remains in the idle state, waiting for a new restore point to appear on the source backup repository.
- **Synchronization process**. The synchronization phase starts at a specific time interval. You can define any interval needed in minutes, in hours, in days.
  At the beginning of a new interval, Veeam Backup & Replication checks if a new restore point is available on the source backup repository:
  - If a new restore point is found, the backup copy job starts the synchronization process and copies the latest VM restore point from the source backup repository to the target backup repository.
  - If a new restore point is not found, the backup copy job is back to the idle state.
- **Transform operations**. After the backup copying task or at the end of the synchronization interval, Veeam Backup & Replication can perform a number of additional transform operations on the target backup repository. Transform operations include three tasks:
  - **Transforming a backup chain**. When a new VM restore point is copied to the target backup repository, Veeam Backup & Replication checks the retention policy settings for the backup copy job. If the limit in restore points is exceeded,
    Veeam Backup & Replication transforms the backup chain to make room for a new restore point. To learn more, see Retention Policy for Backup Copy Jobs.

After the transform process, Veeam Backup & Replication can perform additional operations: remove data for deleted VMs from the backup chain and compact a full backup file.

- **Removing deleted VMs from restore points**. In the backup copy job settings, you can select to maintain retention policy for deleted VMs. In this case, Veeam Backup & Replication will check the list of VMs included in the job and remove data for deleted VMs from the backup chain on the target backup repository. To learn more, see Specifying Advanced Settings.
- **Compacting a full backup file**. In the backup copy job settings, you can select to periodically compact a full backup file to reduce its size and increase the speed of read and write operations. To learn more, see Compacting Full Backup File.

- **Post-job activities**. In the properties of the backup copy job, you can select to perform post-job activities, such as execution of custom scripts or sending job results by email. Post-job activities are performed after all transform operations are completed.

The synchronization process and transform operations make up a separate session of the backup copy job.

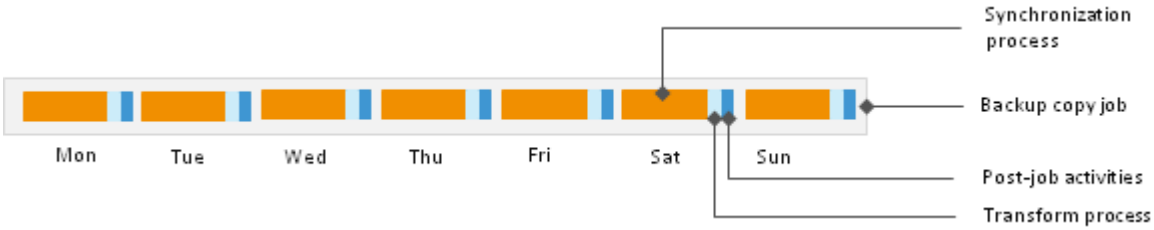

Disabling Backup Copy Job

A backup copy job can be disabled for some time. The disabled backup copy job does not monitor source backup repositories and does not copy restore points to the target backup repository.

The instance of the disabled backup copy job still remains in the Veeam Backup & Replication database and in the product console. You can enable the disable job at any time.

Synchronization Intervals

When creating a backup copy job, you should specify its synchronization interval.

The synchronization interval is a time span in which the backup copy job must copy a VM restore point from the source backup repository to the target backup repository. When a new synchronization interval starts, Veeam Backup & Replication checks if a new restore point is available on the source backup repository. In case a new restore point is found, Veeam Backup & Replication copies it from the source backup repository to the target backup repository. Note that the duration of the synchronization interval affects the restore point selection process. To learn more, see Restore Point Selection.

You can specify the synchronization interval in minutes, hours or days.

Minutely and Hourly Synchronization Intervals

If you set the synchronization interval in minutes or hours, Veeam Backup & Replication runs the backup copy process in cycles, one following another. When one synchronization interval is over, Veeam Backup & Replication starts a new synchronization interval.

For example, if you set the synchronization interval to 4 hours and start the backup copy job at 12 PM, Veeam Backup & Replication will create new synchronization intervals at 12 PM, 4 PM, 8 PM and so on.

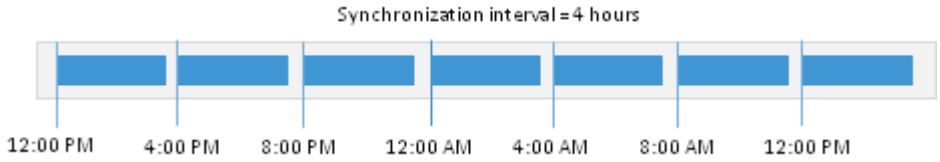

Daily Synchronization Intervals

If you set the synchronization interval to one or several days, Veeam Backup & Replication requires that you define the start time for the synchronization interval. This start time acts as a milestone, or control point for the backup copy process. When the specified point in time occurs, Veeam Backup & Replication starts a new synchronization interval.

For example, if you set the synchronization interval to 1 day and specify to start a new interval at 12 PM, Veeam Backup & Replication will force a new synchronization interval at 12 PM daily.

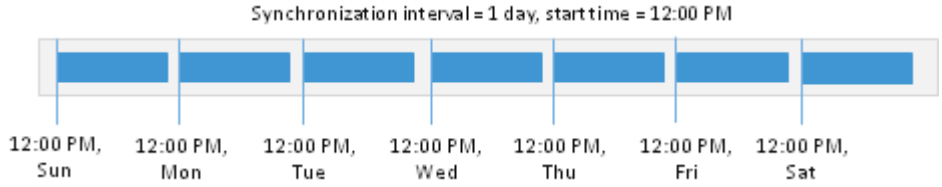

In some cases, the start time of the backup copy job and the start time of the synchronization interval start may not coincide. For example, when configuring a backup copy job, you may set the start time of the synchronization interval to 12 PM; the backup copy job itself may be launched at 12 AM. In this case, the first synchronization interval will be started immediately after you launch the job and will be run for a shorter period of time: in our example, for 12 hours only instead of one day. All subsequent synchronization intervals will be created and run as usual.

Backup Copy Window

If necessary, you can specify a window for the backup copy job. The backup copy window is a 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. In this case, you can define the time interval in which the job must not transfer data over the network.

The backup copy window affects only the data transport process. The backup copy job behavior during the 'prohibited' period of time depends on the length of the synchronization interval:

- If the synchronization interval is greater than the 'prohibited period', the backup copy job will simply put on hold the backup copying operations and wait for allowed hours. The backup copy job is put to the *Idle* state and remains in this state for the whole "prohibited period".
- If the synchronization interval is smaller than the 'prohibited period', Veeam Backup & Replication will finish all backup copy job sessions that must run during the 'prohibited period' with the *Failed* status. During the first synchronization interval on allowed hours, Veeam Backup & Replication will copy the restore point to the target backup repository. The copied restore point will contain all data for the 'prohibited period'. That is, it will aggregate all data that has changed between the latest restore point on the target backup repository and latest restore point on the source backup repository.

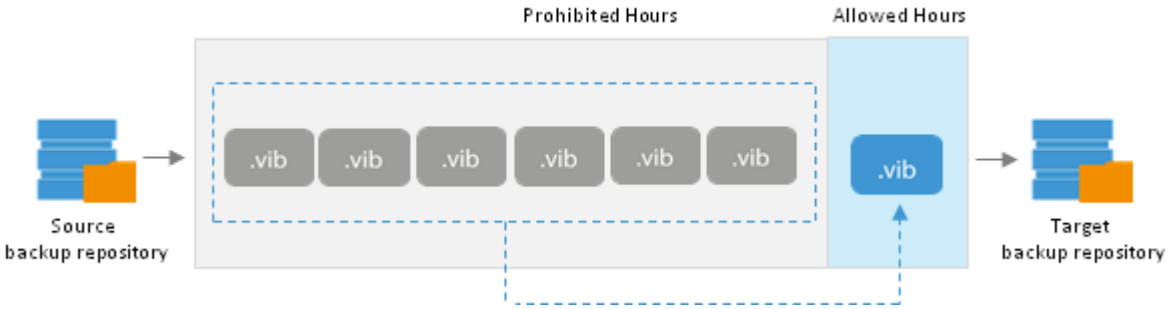

For example, you have set the synchronization interval to 2 hours and defined the backup copy window from 8 AM to 8 PM. Without the backup copy window, Veeam Backup & Replication would transport 6 restore points to the target backup repository between 8 AM and 8 PM. With the backup window, the backup copy job will not copy data from 8 AM to 8 PM. At 8 PM, however, a new synchronization interval will start. Veeam Backup & Replication will transport one restore points from the source backup repository. This restore point will contain VM data for those 6 restore points that might have been copied during the 'prohibited period' plus one that must be created within this new synchronization interval.

Automatic Job Retries

Veeam Backup & Replication automatically retries several operations that are performed within a backup copy job sessions.

Job Tasks Retry

By default, Veeam Backup & Replication automatically retries a failed backup copy task 5 times within one backup copy job session. A new task is started immediately after the previous one, without any interval.

The backup copy task is retried only if the previous task has failed and a restore point has not been copied to the target backup repository. Veeam Backup & Replication does not perform a retry if a task has finished with the *Success* status.

The backup copy task is retried during the same synchronization interval only. If a restore point fails to be copied during all retries in the current synchronization interval, Veeam Backup & Replication marks the synchronization task as failed and waits for the expiration of the synchronization interval. After that, Veeam Backup & Replication performs the necessary transform operations and starts a new synchronization interval.

A backup copy job can process several VMs. If only some VMs are successfully processed by the backup copy task, Veeam Backup & Replication creates a restore point holding data for these VMs on the target backup repository. Veeam Backup & Replication will attempt to process restore points for all VMs during the next synchronization cycle.

Note Some errors from WAN accelerators can block backup copy job retries. For example, if there is no space in the global cache on the target WAN accelerator, Veeam Backup & Replication put backup copying operations on hold and wait for the expiration of the synchronization interval.

Transform Retry

After the backup copying task, Veeam Backup & Replication performs a number of additional transform operations on the target backup repository if necessary. These operations include the backup chain transform, removing of deleted VMs from restore points and compacting a full backup file. To learn more, see Backup Copy Job.

Veeam Backup & Replication may fail to perform transform operations for some reason: for example, if the backup file on the target backup repository is locked by the file-level restore session. By default, Veeam Backup & Replication automatically retries transform operations for 5 times. The first interval between retries is 1 minute; the interval doubles with every new attempt. If Veeam Backup & Replication fails to perform transform operations during all retries in this synchronization interval, the job is put to the idle state, waiting for the new synchronization interval to begin.

Virtual Infrastructure Access Retry

At the beginning of every synchronization interval, Veeam Backup & Replication accesses the virtual infrastructure to make up a list of VMs processed by the job.

Veeam Backup & Replication may fail to access the virtual infrastructure for some reason: for example, in case the vCenter Server is not responding. By default, Veeam Backup & Replication automatically retries access operations for 5 times with a 5 minute interval.

Handling Backup Copy Job Issues

Being a scheduled activity, the backup copy job may fail to run as expected.
Veeam Backup & Replication automatically handles some issues that can occur with the backup copy job.

Short Synchronization Intervals

In some cases, Veeam Backup & Replication may fail to transport the restore point within the synchronization interval of the backup copy job. This can happen, for example, if the synchronization interval is too short and is not sufficient for the amount of data to be copied.

Veeam Backup & Replication handles this situation differently for the first and subsequent synchronization intervals.

- The first synchronization interval always produces a full backup file — the starting point in the backup chain. If Veeam Backup & Replication fails to copy data for the full backup file during the first synchronization interval, it marks the job session as finished with the *Warning* status. During the next synchronization interval, Veeam Backup & Replication attempts to copy data for the full backup file in the following manner:

    1. When a new synchronization interval begins, the restore point that was previously copied no longer corresponds to the restore point selection rules. That is, the time of the restore point creation falls out of the search scope. For this reason, Veeam Backup & Replication waits for a new restore point to appear on the source backup repository.
    2. When a new restore point appears on the source backup repository, Veeam Backup & Replication detects what data blocks still need to be copied to make up a full backup file on the target backup repository, and copies these data blocks.

    This process continues until there is a full backup file on the target backup repository.

- At subsequent synchronization intervals, Veeam Backup & Replication copies incremental restore points. If Veeam Backup & Replication fails to transport an incremental restore point, it marks the synchronization task as failed. Veeam Backup & Replication waits for the expiration of the synchronization interval; after that, Veeam Backup & Replication marks the job session as finished with the *Error* status.

Simultaneous Use of Backup Files

In some cases, restore points on the source backup repository may be locked by the backup job when a new synchronization interval starts. Such situation can occur if the backup job creating restore points on the source repository uses the reversed incremental mode. In this case, during every job session the backup job will lock the full backup file for some time to rebuild it to the most recent state.

If the backup job session and a backup copy job session overlap, Veeam Backup & Replication behaves in the following manner:

- If the synchronization process has started and Veeam Backup & Replication has already managed to copy the restore point to the target backup repository, all other backup copy job activities, such as transform and post-job operations, are put on hold. The backup copy job waits the backup file to be released and continues its operations only after the backup file is unlocked.
- If the synchronization process has started but Veeam Backup & Replication has not managed to copy the restore point to the target backup repository, the backup copy activity is forcibly terminated and the backup copy task is marked as *Failed*. Right after that, Veeam Backup & Replication starts a new backup copy task that remains in the *Idle* state until the backup file on the source backup repository is released and unlocked.

Change of the Synchronization Interval Start Time

If you have selected to run a backup copy job with a daily synchronization interval, you must define the start time of the synchronization interval. However, you may want to change the start time afterwards. After the start time change, Veeam Backup & Replication behaves in the following manner:

1. Veeam Backup & Replication finishes the current synchronization interval running according to the 'old' start time value as usual.
2. After the current synchronization interval is over, Veeam Backup & Replication immediately starts the synchronization interval, not waiting for the 'new' start time point to come. At that, Veeam Backup & Replication "stretches" the started interval: the interval lasts for the time remaining till the new start time plus the time of the synchronization interval itself.
3. All subsequent synchronization intervals are created and started in the regular manner by the new schedule.

For example, when you first created a backup copy job, you set a daily synchronization interval with the start time at 8 AM. After that, you changed the start time to 10 AM. In this case, Veeam Backup & Replication will first finish the synchronization interval that is currently running — that is, the synchronization interval that was started at 8 AM — as usual. After that, it will immediately start a new synchronization interval. This interval will run for 26 hours — from 8 AM of the current day until 10 AM of the next day. All subsequent synchronization intervals will be started at 10 AM every day.

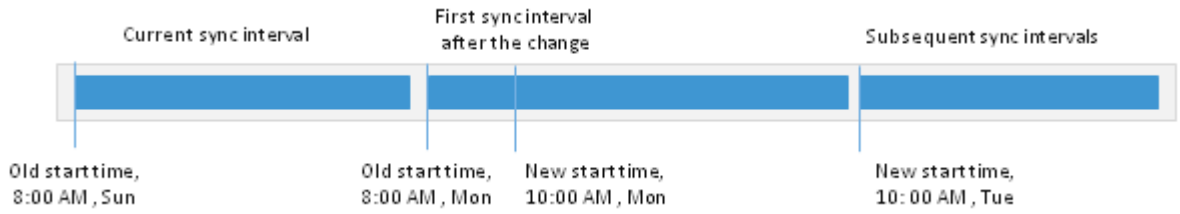

The first synchronization interval that is run after the start time change is typically longer than a regular one. This happens because of the synchronization interval "stretch" mentioned above. To start the synchronization process right away, you can use the **Sync Now** option after you change the start time value. In this case, Veeam Backup & Replication will behave in the following manner:

1. When you start the synchronization process manually, Veeam Backup & Replication forcibly finishes the current synchronization interval and begins a new synchronization interval according to the new start time value. This synchronization interval lasts until a new synchronization interval by the new schedule must be started.
2. All subsequent synchronization intervals are created and started in the regular manner.

As a result, the first synchronization interval after the start time change will begin immediately.

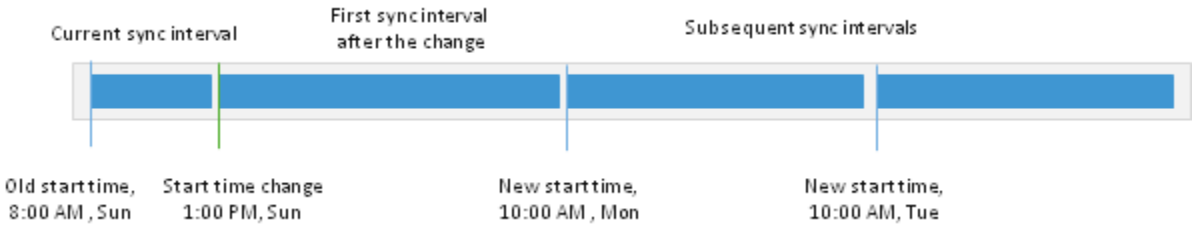

For example, when you first created a backup copy job, you set a daily synchronization interval with the start time at 8 AM. After that, you changed the start time to 10 AM. On the start time change, you started the manual synchronization process at 1 PM. In this case, Veeam Backup & Replication will finish the current synchronization interval — that is, the synchronization interval that was started at 8 AM — immediately at 1 PM. After that, it will start a new synchronization interval. This interval will run for 21 hours — from 1 PM of the current day until 10 AM of the next day. All subsequent synchronization intervals will be started at 10 AM every day.

Retention Policy for Backup Copy Jobs

The backup copy job has its own retention policy settings, independent of retention policy settings specified for a backup job. The retention policy of a backup copy job defines for how long Veeam Backup & Replication must retain copied restore points on the target backup repository.

Veeam Backup & Replication offers two retention policy schemes for backup copy jobs:

- Simple Retention Policy
- GFS Retention Policy

Simple Retention Policy

A simple retention policy scheme is intended for short-time archiving. When you specify retention policy settings for a simple scheme, you define how many restore points you want to retain on the target backup repository.

With this scheme, Veeam Backup & Replication creates a chain of restore points, subsequently following one another. The first restore point in the chain is always a full backup. All other restore points in the chain are incremental backups. By default, Veeam Backup & Replication keeps 7 restore points on the target backup repository.

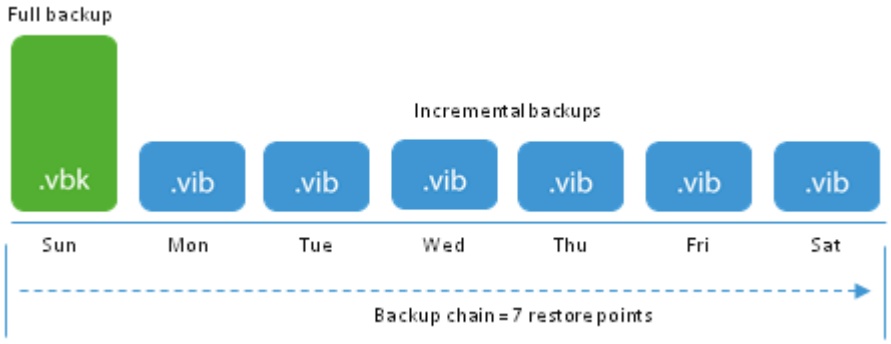

To maintain the desired number of restore points, Veeam Backup & Replication uses the following rotation scheme:

1. At the first synchronization interval, Veeam Backup & Replication copies the first restore point — full backup — to the target backup repository.
2. At every next synchronization interval, Veeam Backup & Replication adds a new restore point — incremental backup — to the chain on the target backup repository. This happens until the number of restore points in the backup chain reaches the number specified in the retention policy settings.
3. After the new restore point is added, the allowed number of restore point is exceeded. Veeam Backup & Replication transforms the backup chain to make room for the most recent restore point.

The backup chain transformation is performed in the following way:

1. Veeam Backup & Replication re-builds the full backup file to include changes of the incremental backup following the full backup. More specifically, it injects data blocks from the first incremental backup in the chain into the full backup. This way, a full backup 'moves' one step forward in the backup chain.
2. The first incremental backup is removed from the chain as redundant: its data has already been injected into the full backup and so the full backup file contains the same data as this incremental restore point.

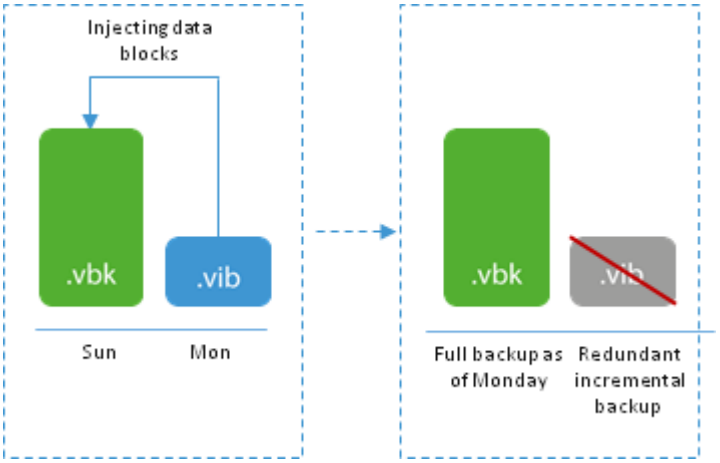

For example, you want to keep 7 restore points. The synchronization interval is 1 day; the backup copy job starts on Sunday.

During the first synchronization interval on Sunday, Veeam Backup & Replication creates the first restore point — a full backup. Monday through Saturday Veeam Backup & Replication adds six incremental backups to the chain.

The next Sunday, Veeam Backup & Replication adds a new incremental backup to the backup chain. The number of allowed restore point in the backup chain is therefore exceeded.

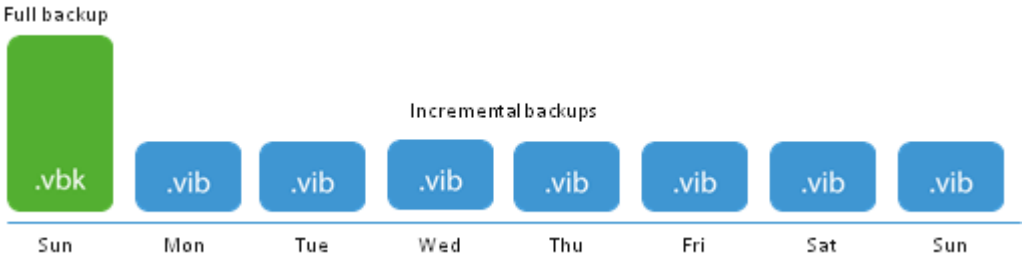

For this reason, Veeam Backup & Replication transforms the backup chain in the following way:

1. Veeam Backup & Replication merges data blocks from the incremental backup copied on Monday into the full backup copied on Sunday. This way, the full backup file 'moves' one step forward — from Sunday to Monday.

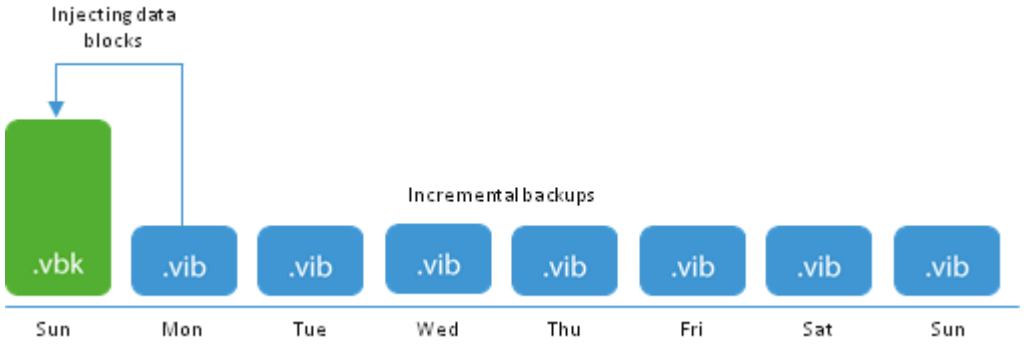

2. The incremental backup copied on Monday becomes redundant and is removed from the chain.

As a result, you have a chain of a full backup as of Monday and six incremental backups Tuesday through Sunday.

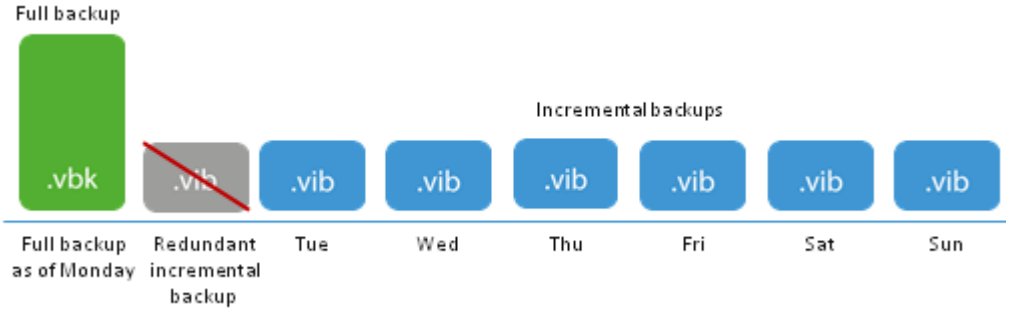

GFS Retention Policy

In most cases, simple backup retention policy is not enough. You cannot store an unlimited number of restore points on the target backup repository forever because it is not rational and is resource consuming. If you want to retain VM data for longer periods of time, it is recommended that you use the GFS retention policy scheme.

The GFS, or Grandfather-Father-Son retention policy is a backup rotation scheme intended for long-term archiving. It lets you keep backups of VMs for an entire year using minimum amount of storage space.

GFS is a tiered retention policy scheme. It uses a number of cycles to maintain backups at different tiers:

- Regular backup cycle performed according to the specified synchronization interval
- Weekly backup cycle
- Monthly backup cycle
- Quarterly backup cycle
- Yearly backup cycle

Backups created on the weekly basis are known as 'sons', monthly backups are known as 'fathers' and yearly backup are known as 'grandfathers'. Additionally, Veeam Backup & Replication maintains quarterly backups.

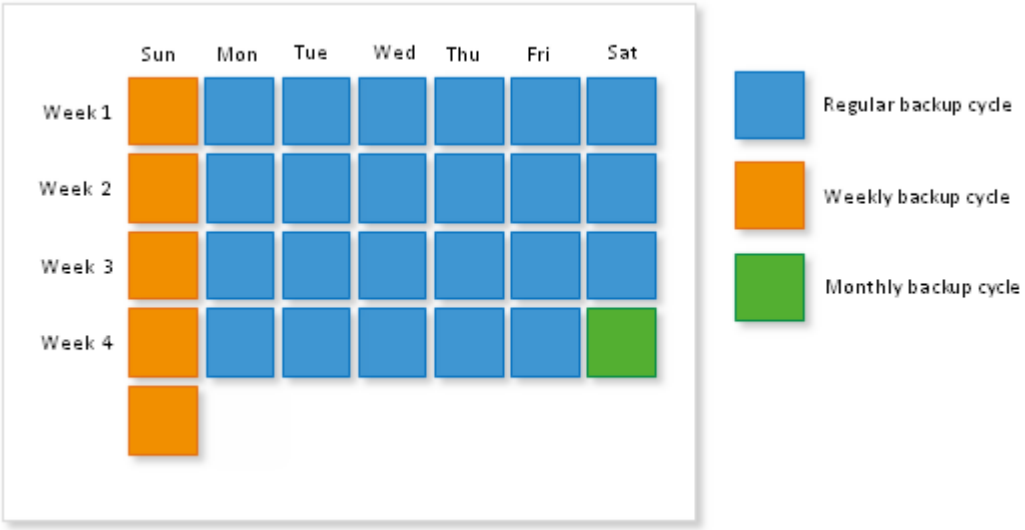

Regular Backup Cycle

The regular backup cycle is based on the simple retention policy scheme. When you specify retention policy settings, you define how many restore points you want to retain in the backup chain.

Veeam Backup & Replication runs the regular backup cycle in the following way:

1. During the first synchronization interval, Veeam Backup & Replication creates the first restore point — a full backup.
2. The next synchronization intervals add incremental backups to the backup chain.

As a result, the regular backup cycle produces a chain of a full backup and a set of incremental backups on the target backup repository.

For example, you have selected to keep 7 restore points. The synchronization interval is 1 day, the backup copy job starts on Sunday. Veeam Backup & Replication will create a full backup on Sunday and add 6 incremental backups Monday through Saturday.

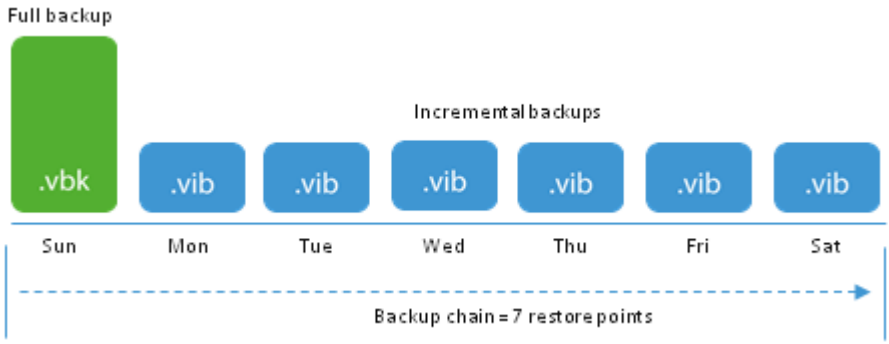

Weekly Backup Cycle

In the GFS scheme, the weekly backup is created during the weekly backup cycle.

Weekly backups are always full backups containing data of the whole VM image as of specific date. When you define retention policy settings for a weekly backup cycle, you specify how many weekly backups you want to retain per month and define the week day on which a full backup must be created.

Weekly backups are not created in a separate task. Veeam Backup & Replication re-uses a full backup created in the regular backup cycle and propagates it to the weekly tier.

The procedure of creating a weekly backup is performed in the following way:

1. Veeam Backup & Replication creates a chain of backups in the regular backup cycle. The chain consists of a full backup and a set of subsequent incremental backups.
   For example, you have selected to keep 7 restore points. The synchronization interval is 1 day, the backup copy job starts on Sunday. During the week, Veeam Backup & Replication creates a backup chain on the target backup repository. The backup chain consists of a full backup copied on Sunday and a set of incremental backups copied Monday through Saturday.

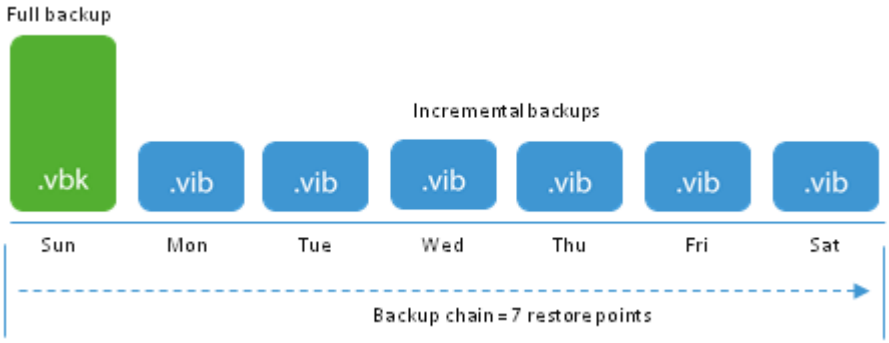

2. With every new synchronization interval, Veeam Backup & Replication transforms the backup chain moving the full backup forward. This procedure repeats until the full backup file reaches the day when the weekly backup is scheduled.

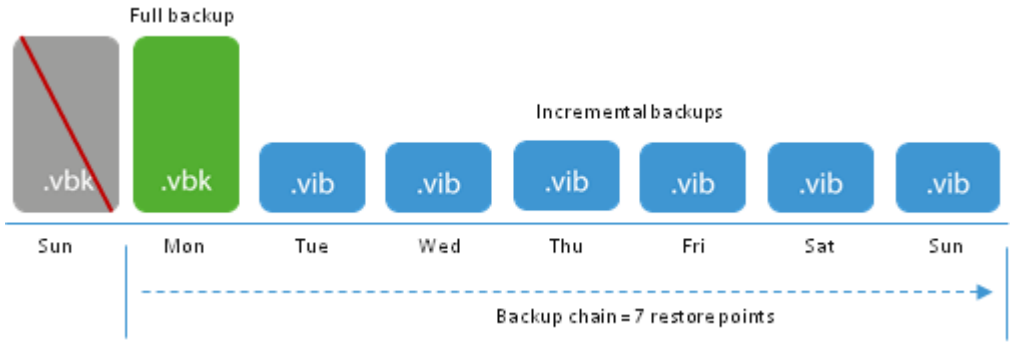

3. During the synchronization interval, Veeam Backup & Replication transforms the backup chain and creates a weekly backup at the same time. The procedure is the following:

   a. Veeam Backup & Replication adds a new restore point to the chain.

   b. As the allowed number of restore points is exceeded, Veeam Backup & Replication transforms the backup chain. The transformation process slightly differs from a regular one. Veeam Backup & Replication does not inject data from the incremental backup to the full backup. Instead, it copies data from full and incremental backups and stores them to a new full backup file, next to the primary backup file.

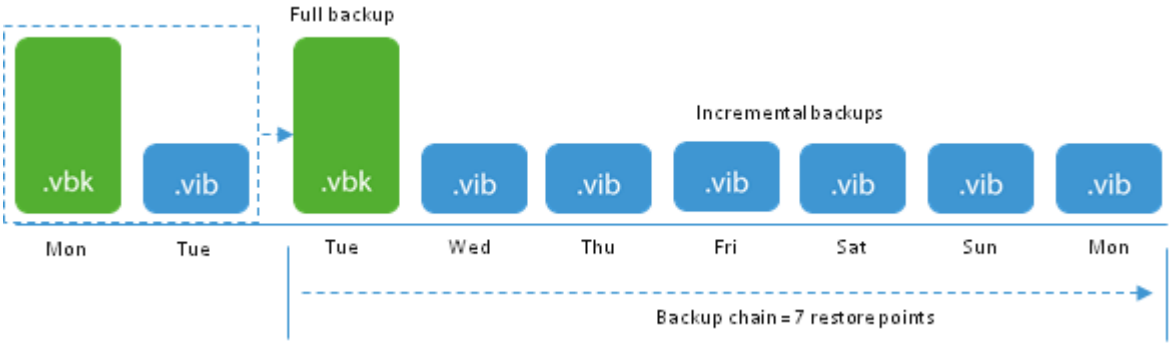

4. The incremental backup from which data was copied is removed as obsolete.

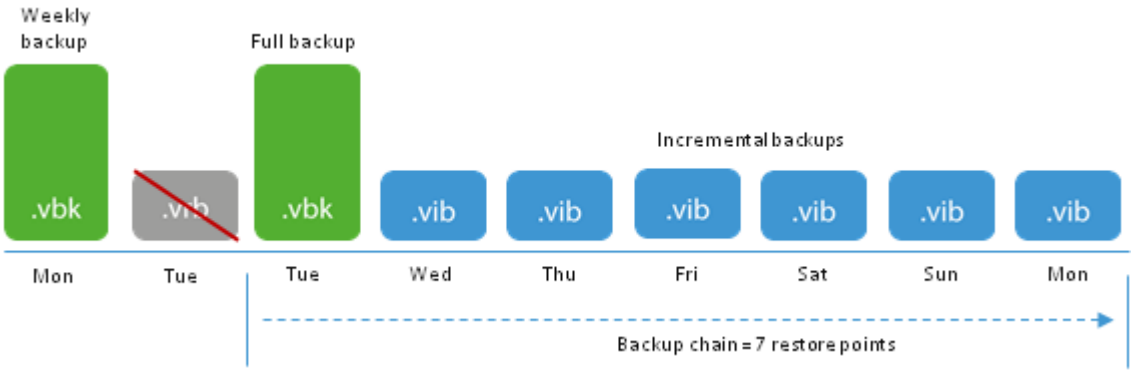

5. The primary full backup file remains on the target backup repository.
   Veeam Backup & Replication sets it aside and marks it as a weekly backup. The weekly backup is no longer used in the backup chain.

6. The newly created full backup file remains in the backup chain and is used as a starting point for incremental backups created by the regular backup cycle.

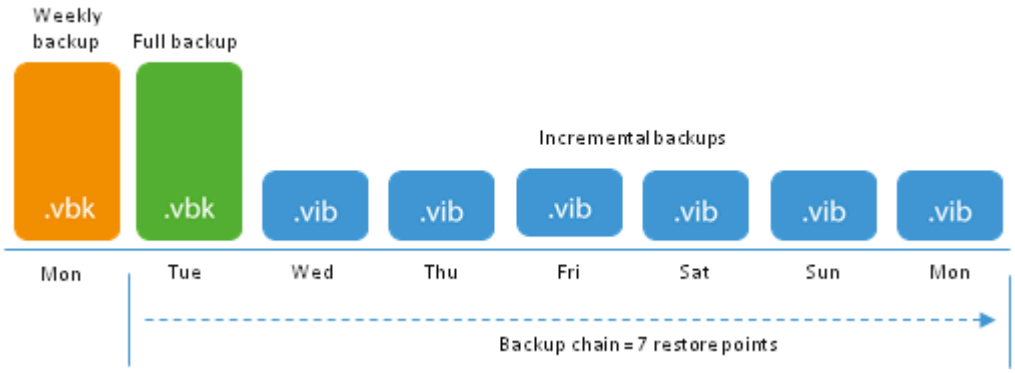

For example, weekly backup is scheduled on Monday. Veeam Backup & Replication will keep transforming the backup chain until the full backup file reaches Monday. During the next synchronization interval, Veeam Backup & Replication will transform the backup chain. To do that, it will copy data from the Monday full backup and Tuesday incremental backup to a new full backup file and store it next to the primary full backup file.

As a result, on the target backup repository you will have a full backup created on Monday and a backup chain that includes a full backup as of Tuesday and a chain of increments Wednesday through Monday. The full backup as of Monday will be marked as a weekly backup and set aside. The full backup as of Tuesday will be used as a new starting point in the backup chain.

Retention Policy for Weekly Backups

Veeam Backup & Replication repeats the weekly backup cycle until the number of weekly backups allowed by the retention policy is exceeded. After that, Veeam Backup & Replication removes the earliest weekly full backup from the target backup repository to make room for the most recent weekly full backup.

When deleting obsolete weekly backups, Veeam Backup & Replication considers weekly intervals, not separate backup files. For this reason, in some situations the target backup repository may contain a greater number of weekly backups than specified in the GFS retention policy scheme.

For example, you have selected to retain 4 weekly backups. The weekly backup is scheduled on Sunday. Right after the first weekly backup is created, you change the weekly backup schedule and schedule the weekly backup on Thursday. As a result, during the first week
Veeam Backup & Replication will create two weekly backups. At the end of the month you will have 5 weekly backups on the target backup repository:

- One weekly backups created on Sunday
- Four weekly backups created on Thursday

The next week, Veeam Backup & Replication will add a new weekly backup to the target backup repository. At the same time, it will remove all backups that were created on the first week – Sunday backup and Thursday backup.

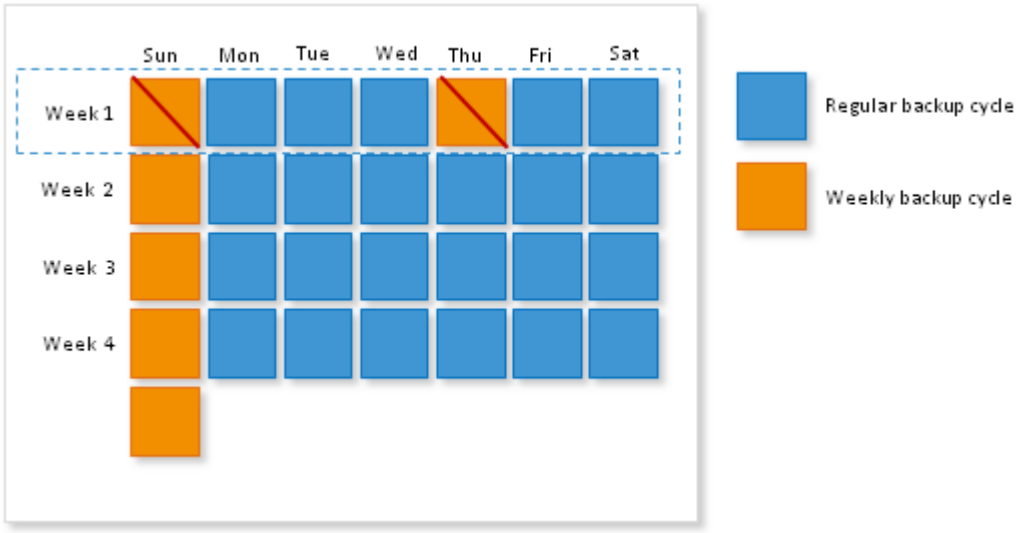

Important! One and the same full backup can be marked as weekly, monthly, quarterly and/or yearly. When transforming weekly, monthly, quarterly and yearly backup chains, Veeam Backup & Replication checks the flags set for the full backup file. If the full backup file belongs to some other retention policy tier and must be retained on the target backup repository, such backup file will not be removed.

Restore Point Selection for Weekly Backup

Typically, when the weekly backup is performed, Veeam Backup & Replication takes a full backup as of this day and marks it as a weekly backup. In some cases, however, Veeam Backup & Replication may fail to find a full backup on the day when the weekly backup is scheduled. In this situation, Veeam Backup & Replication will use the nearest full backup file created within the next synchronization interval.

For example, you have set the synchronization interval to 1 week and started the backup copy job on Sunday. As a result, every new restore point is created on Sunday. When Veeam Backup & Replication transforms the backup chain, the full backup will move from the previous Sunday to the next Sunday.

Imagine the weekly backup is schedule on Wednesday. As all backups are created on Sunday, Veeam Backup & Replication will not find a full backup as of Wednesday. For this reason, it will use the full backup from the next synchronization interval – a full backup as of Sunday.

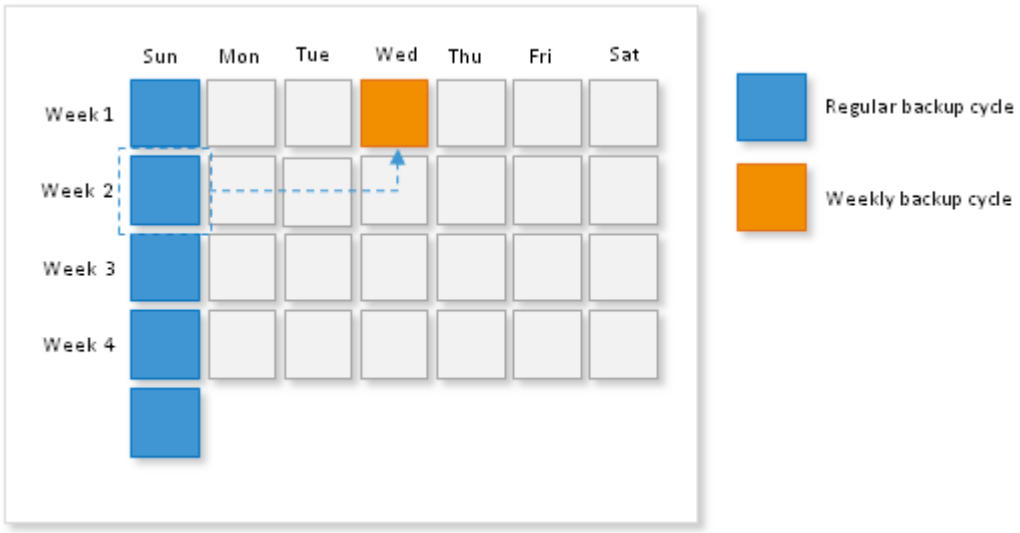

Monthly, Quarterly and Yearly Backup Cycles

The monthly, quarterly and yearly backup cycles use the same algorithms as the weekly backup cycle. When you define retention policy settings for these backup cycles, you specify how many backups you want to retain in the specified time interval and define the week day on which the monthly, quarterly or yearly backup should be created.

Veeam Backup & Replication repeats the monthly, quarterly or yearly backup cycle until the number of backups allowed by the retention policy is exceeded. After that, Veeam Backup & Replication removes the earliest full backup from the target backup repository to make room for the most recent monthly, quarterly or yearly backup.

Compacting a Full Backup File

The backup copy job constantly transforms the full backup file in the backup chain to meet the desired retention policy settings. The transformation process, however, has an opposite side. In the long run, the full backup file grows large and gets badly fragmented. The file data occurs to be written to non-contiguous clusters on the storage, which slows down reading from and writing to the backup file.

To resolve this issue, Veeam Backup & Replication offers an advanced option to compact a full backup file. Backup file compacting can be performed periodically – weekly or monthly on specified days.

During the file compact operation, Veeam Backup & Replication creates a new empty VBK file and copies to it all data blocks from the full backup file. As a result, the full backup file gets defragmented, its size reduces and the speed of writing and reading to/from the file increases.

Note The full backup compacting operation has the following limitations:

- The **Compact full backup** option can be enabled only for the simple retention policy scheme.
- The target backup repository must have enough space to hold a backup file of the full backup size. During the compacting process, Veeam Backup & Replication creates an auxiliary VBK file that exists on the backup repository till the end of the compacting operation.

Health Check for Copied Backups

Veeam Backup & Replication offers an option to perform periodic health check for restore points stored on the target backup repository. The health check is actually a CRC check performed for restore points stored on the target backup repository.

The health check is started at the beginning of the synchronization cycle before data transport. Veeam Backup & Replication always verifies only the most recent restore points for the VM(s) processed by the backup copy job. By default, the health check is performed on the last Sunday of every month. You can change the health check schedule and select to perform it weekly or monthly on specific days.

The health check is performed in the following way:

1. When a new restore point is copied from the source backup repository, Veeam Backup & Replication calculates checksums for data blocks in the restore point and saves them along with the backup data on the target backup repository.
2. During the health check, Veeam Backup & Replication calculates checksums for data blocks in the backup files stored on the target backup repository and compares them to the checksums that were previously stored to the backup files.

If a health check fails, Veeam Backup & Replication displays a warning in the job session report. During the next synchronization interval, Veeam Backup & Replication transports valid data blocks from the source backup repository and stores them to the newly copied restore point on the target backup repository.

Note In case the backup copy job uses WAN accelerators, Veeam Backup & Replication will attempt to find data blocks in the global cache not to transfer data over the network. To learn more, see WAN Acceleration.

Mapping Backup Copy Jobs

If you plan to copy VM restore points over the WAN and slow connections, you can use backup mapping.

Backup mapping can only be used if you already have a full backup file for the VM you plan to process with the backup copy job on the target backup repository. In this case, you can point the backup copy job to this backup file. This full backup from the backup chain will be used as a "seed" for the backup copy job and you will need to copy only small incremental changes over the network.

Important! The backup copy job can be mapped to the backup only if the backup chain you plan to use as a "seed" contains one restore point — a full backup file. If the chain contains a number of restore points, Veeam Backup & Replication will fail to map the backup copy job to the selected backup. To overcome this situation, you can create a backup "seed" by means of an auxiliary backup copy job on the target repository. To learn more, see Creating a Seed for the Backup Copy Job.

A backup copy job with a mapped backup is performed in the following way:

1. Veeam Backup & Replication accesses a full backup file to which you mapped the backup copy job. This backup file is used as a seed for further backup copying process.
2. Data processing during all subsequent synchronization intervals is performed in a regular manner. Veeam Backup & Replication copies only incremental changes and stores them as new restore points next to the full backup file.

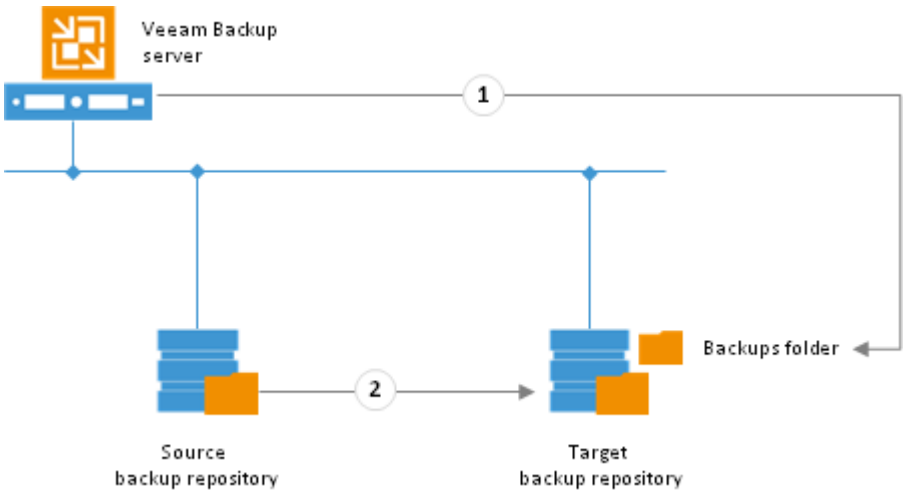

Unlike a regular backup copy job that uses a dedicated folder on the target backup repository, a mapped backup copy job stores copied restore points to the same folder where the backup file used as a seed resides.

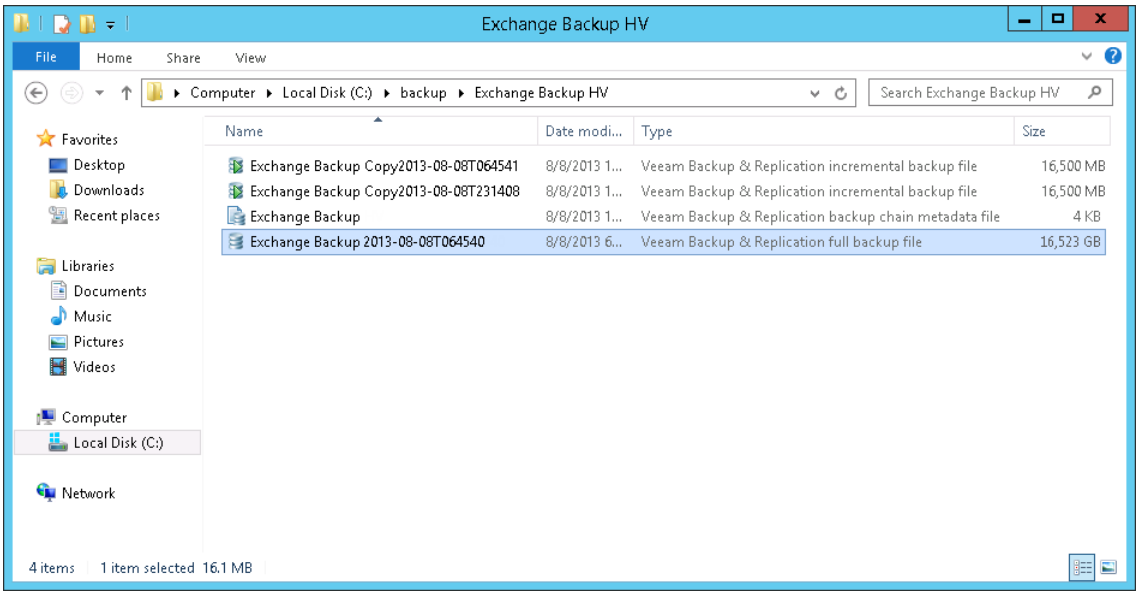

Creating a Seed for the Backup Copy Job

When you map a backup copy job to the backup file, you need to make sure that the backup chain used as a "seed" contains only one restore point. In the opposite case, the backup copy job will not be mapped to the selected backup.

If the backup chain contains a number of restore points, you can use a workaround scenario. To do that, you need to configure an auxiliary backup copy job in addition to the initial backup copy job. The auxiliary backup copy job will produce a full backup file and you will be able to use it as a "seed" for the initial backup copy job mapping.

To create a "seed", perform the following actions:

1. Create a new, auxiliary, backup copy job. The created backup copy job must copy backups of VMs processed by the initial backup copy job. Target the created auxiliary backup copy job to some backup repository on the source side. This backup repository will be used as an intermediate one.
2. Run the auxiliary backup copy job to create a full backup file (VBK) on the intermediate backup repository.
3. Move the created VBK file and VBM file from the intermediate backup repository to the backup repository on the target side.
4. Perform repository rescan to populate the backup repository on the target side.
5. Edit settings of the initial backup copy job: point the initial backup copy job to the full backup file that you have created and moved to the target backup repository.
6. Click **Sync Now** to start a new synchronization interval.

As a result, Veeam Backup & Replication will use the backup file that you have created and moved to the target backup repository as a "seed", or starting point, for the backup chain produced by the initial backup copy job. When a new restore point for the VM is available on the source backup repository, Veeam Backup & Replication will transfer only incremental changes and store them next to the "seed".

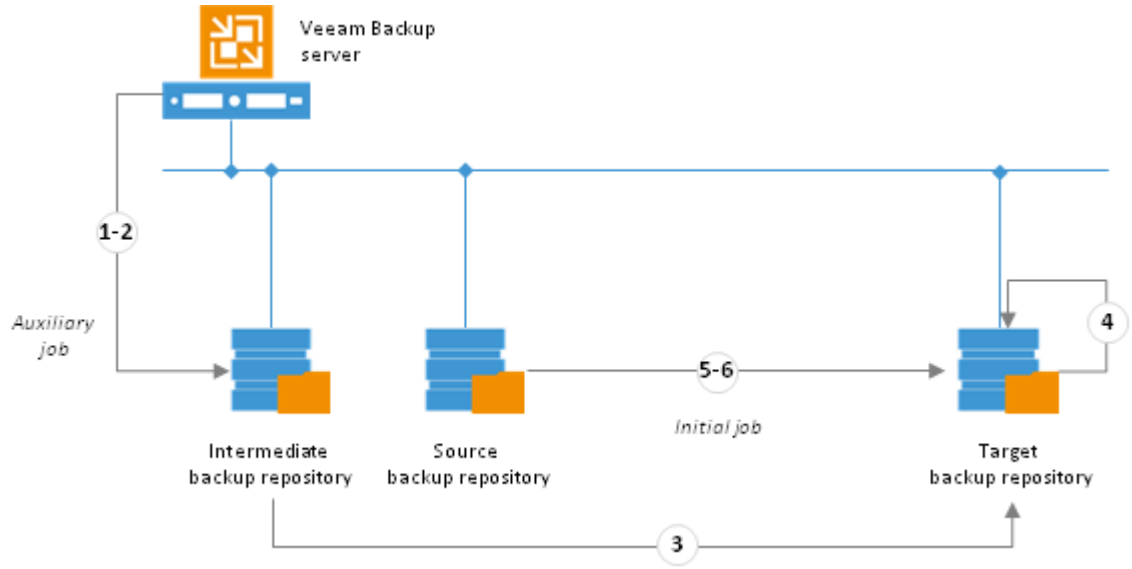

Important! When you configure an auxiliary backup copy job, make sure that its synchronization interval covers the whole chain of restore points on the backup repository from which you plan to copy backup files. The length of the synchronization interval has an impact on the algorithm of restore point selection. Veeam Backup & Replication copies only those restore points that match the following criterion:

Time of restore point creation >= current time – synchronization interval

That is, if you have a backup chain whose earliest restore point is one week old, you need to set the synchronization interval to one week. If you set the synchronization interval to a smaller time span, for example, 1 day, all restore points that are older than 1 day will fall out of the search scope and Veeam Backup & Replication will not transfer their data. To learn more, see Restore Point Selection.

WAN Acceleration

Storing backups offsite always involves moving large volumes of data between remote sites. The most common problems that backup administrators encounter during offsite backup are:

- Insufficient network bandwidth to support VM data traffic
- Transmission of redundant data

To solve these problems, Veeam Backup & Replication offers the WAN acceleration technology that helps optimize data transfer over the WAN.

The WAN acceleration technology is specific for backup copy jobs. Being a built-in feature, Veeam's WAN acceleration does not add complexity and cost to the backup infrastructure and does not require agents. The technology has been developed for copying backup files, with consideration of the VM backup file content.

Note WAN acceleration is available only in the Enterprise Plus Edition of Veeam Backup & Replication.

Global Data Deduplication

The goal of WAN acceleration is to send less data over the network. To reduce the amount of data going over the WAN, Veeam Backup & Replication uses the global data deduplication mechanism.

1. When you first run the backup copy job, Veeam Backup & Replication analyzes data blocks going over WAN.
2. With every synchronization interval of the backup copy job, Veeam Backup & Replication uses the data redundancy algorithm to find duplicate data blocks in the copied backup file. Veeam Backup & Replication analyzes data blocks in the backup file on the source side and compares them with those that have been previously transferred over the WAN. If an identical data block is found, Veeam Backup & Replication deduplicates it.

Veeam Backup & Replication uses three sources for data deduplication:

- **VM disks**. Veeam Backup & Replication analyses data blocks within the same VM disk. If identical blocks are found, duplicates are eliminated.
  For example, in case of a virtualized Microsoft Exchange server, the same email is typically stored in sender's Outbox folder of the sender and recipient's Inbox folder, which results in duplicate data blocks. During the backup copy job, Veeam Backup & Replication detects such VM data blocks and performs deduplication.
- **Previous restore points for the processed VM on the target repository**. Veeam Backup & Replication analyses data in the restore point that is about to be copied and the restore point(s) that are already stored on the target backup repository. If an identical block is found on the target repository, Veeam Backup & Replication eliminates the redundant data block in the copied restore point.
- **Global cache**. When a backup file is sent over the WAN, Veeam Backup & Replication creates a global cache holding data blocks that repeatedly go over the WAN. In a new synchronization interval, Veeam Backup & Replication analyzes data blocks to be sent and compares them with data blocks stored in the global cache. If an identical data block is already available in the global cache, its duplicate on the source side is eliminated and not sent over the WAN.

As a result, only unique data blocks go over the WAN. Data blocks that have already been sent are not sent. This way, Veeam Backup & Replication eliminates transfer of redundant data over the WAN.

Note Veeam Backup & Replication deduplicates data blocks within one VM disk and in restore points for one VM only. Deduplication between VM disks and restore points of different VMs is performed indirectly, via the global cache. To learn more, see WAN Global Cache.

WAN Accelerators

To enable WAN acceleration and data deduplication technologies, you must deploy a pair of WAN accelerators in your backup infrastructure.

- One WAN accelerator is deployed on the source site, closer to the source backup repository.
- The other one is deployed on the target site, closer to the target backup repository.

Technically, WAN accelerators add a new layer in the backup infrastructure — between the Veeam transport service on the source backup repository and the Veeam transport service on the target backup repository.

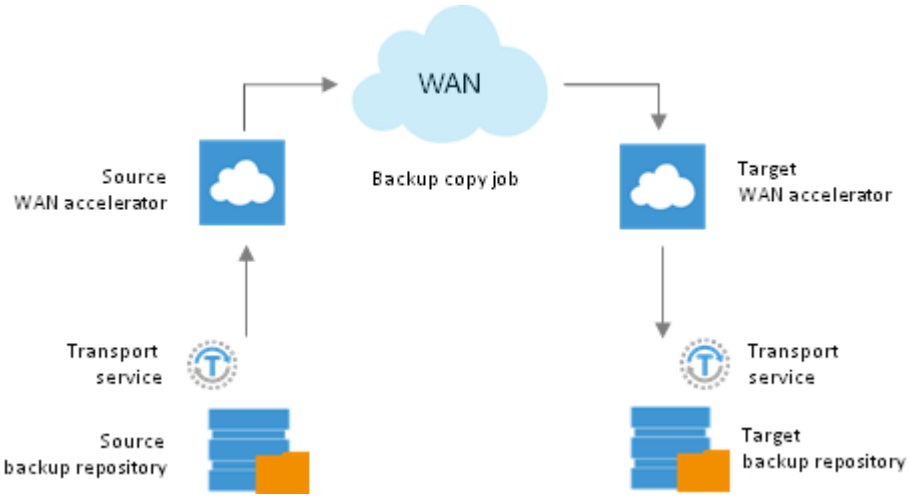

WAN accelerators are dedicated components responsible for global data caching and data deduplication. On each WAN accelerator, Veeam Backup & Replication installs the Veeam WAN Accelerator Service responsible for WAN acceleration tasks.

On each WAN accelerator Veeam Backup & Replication creates the *VeeamWAN* folder containing the following data:

- The *VeeamWAN* folder on the source WAN accelerator stores files with digests required for global deduplication. To learn more, see How It Works.
- The *VeeamWAN* folder on the target WAN accelerator stores global cache data.

To create a WAN accelerator, you need to assign the WAN accelerator role to a specific machine. You can use any 64-bit Windows-based machine in your environment, either physical or virtual. You can even assign the WAN accelerator role to the existing backup proxies and backup repositories. The machine that will perform the role of the target WAN accelerator must have enough free disk space to store the global cache data.

WAN Global Cache

From the technical point of view, the global cache is a folder on the target WAN accelerator. By default, global cache data is stored in the `VeeamWAN` folder on the disk with the most amount of space available. However, you can define any folder of your choice when you configure the target WAN accelerator.

By default, the size of the global cache is 100 GB. You can increase the size or decrease it if necessary. The more space you allocate, the more repeating data blocks will be written to the global cache and the more efficient WAN acceleration will be. It is recommended that you allocate at least 40 GB to the global cache storage.

The global cache size is specified per source WAN accelerator. That is, if you plan to use one target WAN accelerator with several source WAN accelerators, the specified amount of space will be allocated for every source WAN accelerator that will be working with the target WAN accelerator and the size of the global cache will increase proportionally.

The WAN global cache is actually a “library” that holds data blocks repeatedly going from the source repository to the target repository. The global cache is populated at the first synchronization interval of the backup copy job. The priority is given to data blocks of Windows-based OSes, other OSes like Linux/Unix and standard applications such as Microsoft Exchange Server.

Veeam Backup & Replication constantly maintains the global cache in the actual state. To do that, it continuously monitors data blocks going over the WAN and data blocks in the global cache.

- If some new data block is constantly sent over the WAN, it is added to the global cache.
- If some data block in the global cache is not sent over the WAN and re-used for some period of time, it is removed from the global cache to make room for new data blocks.

Veeam Backup & Replication also performs periodic consistency checks. If some data block in the global cache gets corrupted, Veeam Backup & Replication removes it from the global cache.

The efficiency of the WAN acceleration increases with new every synchronization interval in the backup copy job. During the first synchronization interval in the backup copy job, the WAN acceleration level is minimal. Veeam Backup & Replication populates the global cache. With every new synchronization interval, Veeam Backup & Replication updates the global cache to include the most "popular" data blocks and the WAN acceleration efficiency increases.

How It Works

When you create a backup copy job, you can select to use WAN acceleration in its properties.

The procedure of backup copying with WAN acceleration enabled is performed in the following way:

1. Before processing the backup file with the backup copy job, Veeam Backup & Replication uncompresses the backup file to analyze its content.
2. The Veeam WAN Accelerator Service on the source WAN accelerator analyzes data blocks of the uncompressed backup file and creates a file with digests for these data blocks. The created file with digests is stored to the *VeeamWAN* folder on the source WAN accelerator.
3. Veeam Backup & Replication compresses the backup file data and copies it to the target backup repository. At this point, Veeam Backup & Replication can perform deduplication within the VM itself — that is, deduplicate identical data blocks in every VM disk.
4. During the data transfer process, the Veeam WAN Accelerator Service on the target WAN accelerator populates the global cache storage with data blocks from the copied backup file.
5. During the next synchronization interval, the Veeam WAN Accelerator Service on the source WAN accelerator analyzes data blocks in the backup file that should be copied this time and creates digests for these data blocks.
6. The Veeam WAN Accelerator Service compares the created digests with the digests that have been previously stored to the VeeamWAN folder on the source WAN accelerator. If duplicate data blocks are found, the actual block data in the backup file is not copied over the WAN. Instead, it is taken from the global cache and written to the restore point in the backup copy folder.
7. Additionally, Veeam Backup & Replication analyzes restore points that have been previously copied to the target backup repository. If duplicates are found, Veeam Backup & Replication does not copy such blocks over the WAN but takes them from the global cache.

As a result, Veeam Backup & Replication copies only new data blocks to the target backup repository and uses data blocks that are already stored in the global cache.

Note If the target WAN accelerator is used by several backup copy jobs, the target backup repository may already contain data blocks of the necessary VM type. In this situation, Veeam Backup & Replication will copy the required data blocks to the global cache before the copying process starts and use these data blocks further on. To learn more, see Many to One WAN Acceleration.

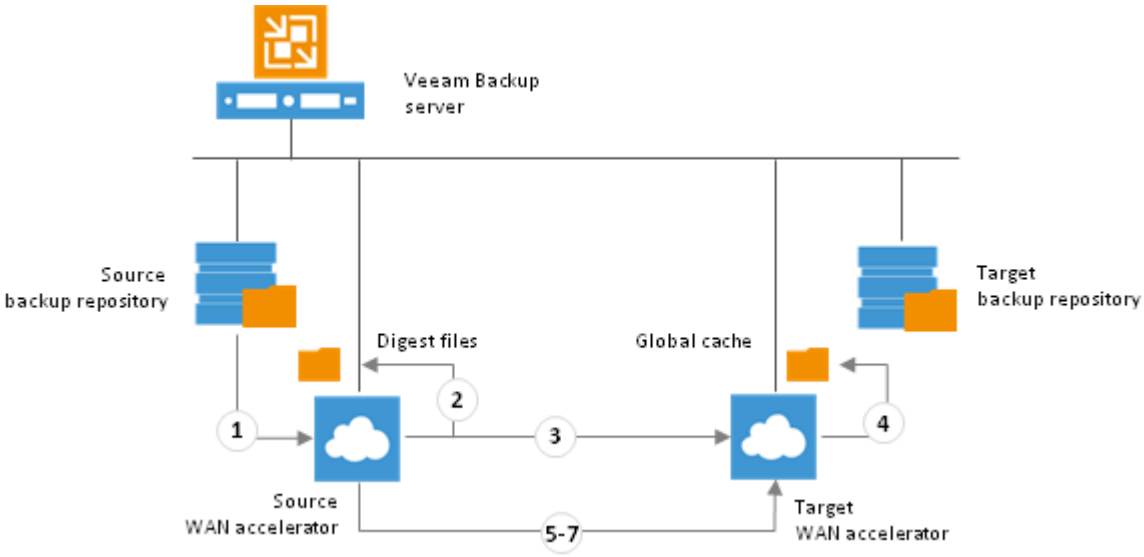

Many to One WAN Acceleration

The WAN global cache can be used by several source WAN accelerators simultaneously. For example, if you have several remote/branch offices, you can configure several source WAN accelerators in remote sites and one target WAN accelerator in the head office.

In this case, the global cache will hold cache data for separate source WAN accelerators. The cache data for every source WAN accelerator will be stored in a dedicated subfolder in the global cache folder.

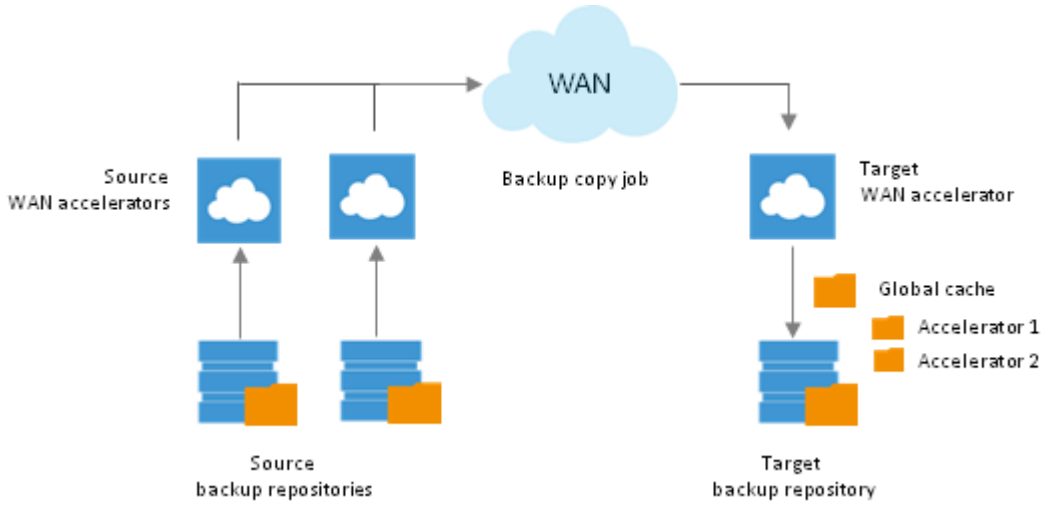

When one target WAN accelerator is used by several source WAN accelerators, Veeam Backup & Replication can copy data blocks between global cache subfolders created for them.

For example, you have two backup copy jobs: *Job 1* and *Job 2*. The Job 1 uses the source WAN accelerator *Source 1* and the target WAN accelerator *Target 3*. The *Job 2* uses the source WAN accelerator *Source 2* and the same target WAN accelerator *Target 3*. In the global cache folder, Veeam Backup & Replication will create two subfolders: *Source 1* and *Source 2*.

Imagine that the *Job 1* processes a VM running Microsoft Windows 2008 R2 and it has been running for some time. In the global cache, there is already data for this type of OS.

Now imagine that Job 2 should also process a VM running Microsoft Windows 2008 R2. When you start the *Job 2* for the first time, in its global cache subfolder there is no data for this type of OS. In such situation, Veeam Backup & Replication will simply copy the necessary data block from the *Source 1* cache folder to the *Source 2* cache folder and will not transport this data block over the WAN.

Note Beside using global cache of other WAN accelerator, Veeam Backup & Replication also utilizes backup files residing on the backup repository. For example, if a backup repository contains a backup file created with a backup job and the backup copy job starts copying a backup of a VM of the same type, Veeam Backup & Replication will copy data blocks from the backup file to the global cache folder not to transfer them over the WAN.

Data Block Verification

During the VM copy process, Veeam Backup & Replication performs a CRC check for the VM traffic going between the source and target WAN accelerators. The CRC check helps ensure that the correct VM data goes to the target side and no corrupted data blocks are written to the global cache or to backup files in the target backup repository.

The check is performed in the following way:

1. Before sending a data block to the target side, Veeam Backup & Replication calculates a checksum for the copied data block.
2. Once the data block is copied over the WAN and before it is written to the global cache or to the target backup repository, Veeam Backup & Replication re-calculates the checksum for this data block on the target side.
3. The source and target checksums are compared. If the checksums do not coincide, the target WAN accelerator sends a request to the source WAN accelerator for the correct data block. The source WAN accelerator re-sends the necessary data blocks to the target WAN accelerator as is and the re-sent data block is written to the global cache or to the backup file on the target backup repository on the fly.

Tape Device Support

Long-term archiving and compliance are listed as primary reasons for using tape that appears to be one of the most widely used media for off-site backup. Veeam Backup & Replication offers support for tape devices allowing you to archive your data to tape and restore it from tape whenever needed.

Veeam Backup & Replication allows working with tape devices directly attached to the Veeam backup server (for example, connected over Fibre Channel (FC), Serial Attached SCSI (SAS), SCSI) or remote devices connected to the Veeam backup server via iSCSI (you can use Microsoft iSCSI initiator to establish the connection).

Both physical and virtual tape libraries and standalone drives are supported.

Archiving to Tape

With Veeam Backup & Replication, you can archive to tape both backup files stored in backup repositories and regular files you might want to write to tape, such as Windows and Linux files. Veeam Backup & Replication supports file backup from any server which has been added as a managed server to the Veeam Backup console (that is, Windows or Linux server, including physical boxes). You can also archive files residing on NAS devices.

The archiving options include the following ones:

- **Backup to tape jobs**. This option allows you to archive to tape media backups created by Veeam Backup & Replication. Using backup to tape jobs, you can implement the '3-2-1' backup approach (3 copies, 2 types of media, 1 off-site location) considered as a best practice for data protection and disaster recovery. Veeam Backup & Replication provides flexible retention and scheduling settings that help automate backup archiving.
- **Files to tape jobs**. This option allows you to archive to tape media files from Windows and Linux servers connected to Veeam backup server. You can create both full and incremental backups of files on tape.

Veeam Backup & Replication uses the MTF (Microsoft Tape Format) industry format to write data to tape.

Restoring from Tape

Veeam Backup & Replication offers multiple options for restoring data from tape:

- **Restoring backup files from tape to disk**. You can also restore full backups or even backup chains to a repository or any location of your choice. The restored backup is registered in the Veeam Backup & Replication console so that you can work with it and use it for any restore scenario later on. To learn more, see Restoring Backups from Tape.
- **Restoring VMs from tape into the virtual infrastructure**. You can restore an entire VM from backup archive on tape. Veeam Backup & Replication supports all options available for regular full VM recovery, including selecting a restore point, choosing target location, changing VM configuration settings and so on. To learn more, see Restoring VMs from Tape to Virtual Infrastructure.
- **Restoring files and folders from tape to the original location or to another directory**. Flexible recovery options allow you to recover files or folders back to the original location or another server, preserving ownership and access permissions. To learn more, see Restoring Files from Tape.

How It Works

To back up data to the tape media and restore from tape archives, Veeam Backup & Replication uses several components:

1. **Source**. The source is the initial location where backup files or regular files you want to archive to tape reside. This can be a backup repository where Veeam backups are stored or a Windows or Linux server hosting files that should be archived.
2. **Data path**. During archiving and restore processes, data is transferred between two terminal points; the process is controlled by two Veeam Transport services:
   - *Source-side transport service* communicates with the source (backup repository or server hosting the files) and initiates reading data from that source.
   - *Target-side transport service* runs locally on the Veeam backup server machine to which a tape device is connected; it initiates writing data to the tape device.

   The transport services are responsible for the following operations:
   - Archiving files and folders from Windows and Linux servers or Veeam backup files from backup repository to the specified media in the tape library.
   - Restore of files and folders or Veeam backup files recorded on tape and placing them in the backup repository or folder. Virtual machines can be then restored through Veeam Backup & Replication recovery capabilities.
3. **Veeam backup server**. The core component in the backup infrastructure that manages all operations, schedules and executes the jobs.
4. **Tape device**. Physical tape library, a virtual tape library or a stand-alone tape drive. The tape device should be located in the same site with the Veeam backup server. Connection between them is established directly (using FC/SAS/SCSI) or over the network (using iSCSI).
5. **Veeam Backup database**. This component is used to store the following tape-related data:
   - *Tape Catalog* stores files/folders archived to tape media, as well as VBK and VIB backup files. It is updated during file-to-tape and VM-to-tape jobs. The content of the Tape Catalog can be examined under the Tape node in the Files view.
   - *Backup Catalog* stores information about VMs whose backups were archived to tape media. The content of the Backup Catalog is updated during VM-to-tape jobs and can be examined under the **Backups** > **Tape** node in **Backup & Replication** view.

VM Backup to Tape

When Veeam backup server executes a VM Backup to Tape job (started manually or on schedule), it performs the following operations:

1. Veeam Backup & Replication enumerates backup files using the Backup Catalog in the Veeam Backup & Replication database to detect if any data has been modified since the latest backup. Detected changes are queued for archiving.
2. Veeam Backup & Replication connects to transport services and starts the data transfer process.
3. Transport services take over from this point. The source transport gets data from the backup repository and target transport service controls recording to tape.
4. While tape recording is performed, Veeam Backup service updates data in the Backup Catalog and Tape Catalog in Veeam Backup database. The Veeam Backup console displays refreshed information about VM-to-tape backups and shows job statistics.

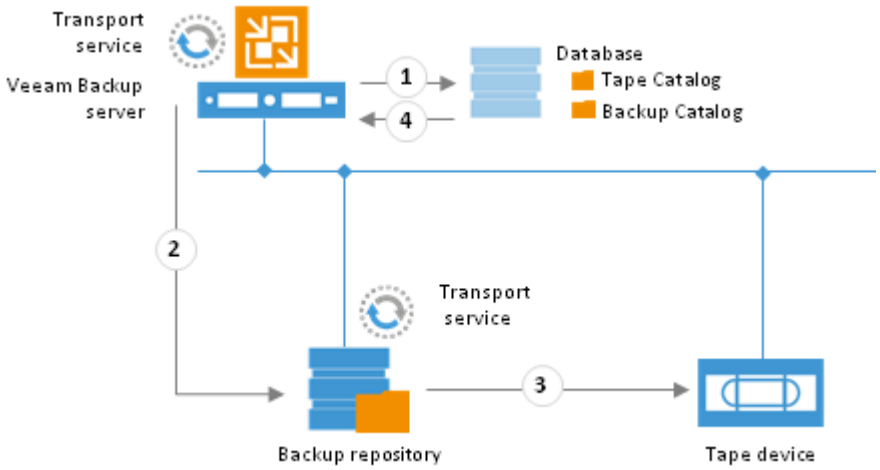

File Backup to Tape

With file backup to tape, you can archive files and folders from Windows or Linux machines added to your backup infrastructure. When a *File to Tape* job runs, the process goes in a similar way:

1. Veeam Backup & Replication scans the file system to detect file's latest modification date and compares it to the information stored in the database. If the file has been changed since last stored backup, this means it is necessary to archive the latest changes to tape.
2. Veeam Backup & Replication connects to transport services and starts data transfer process. Veeam transport services take over from that point, obtaining files from their host (Windows or Linux server) and recording them to tape.
3. While archiving to tape, Veeam Backup service updates Tape catalog data in the Veeam Backup database.
4. The Veeam Backup console displays refreshed information about file-to-tape backups and shows job statistics.

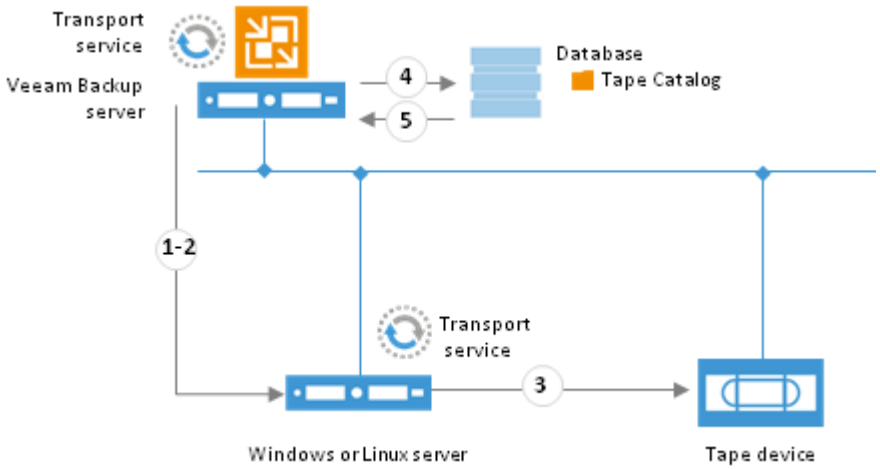

Data Cataloging

To facilitate and streamline recovery from tape, Veeam Backup & Replication catalogs information about all archived backup and file content and stores this data in the Veeam Backup & Replication database. The catalog includes up-to-date information on the backup sets on tape and is updated with every backup to tape or files to tape job session.

With data cataloging, Veeam Backup & Replication quickly detects location of the required items on tape, regardless of whether the tape is online or offline, which helps restore data from tape much quicker when necessary.

PLANNING AND PREPARATION

This section describes the background information that you should be aware of before building your data protection infrastructure using Veeam Backup & Replication.

Prerequisites

- Veeam Backup & Replication requires .NET Framework 4. If it is not available, the Veeam Backup & Replication setup will install it on your computer.
- Veeam Backup & Replication uses SQL Server instance installed either locally or remotely. In case it is not installed, the Veeam Backup & Replication setup will install Microsoft SQL Server 2008 R2 Express on your computer. If a Microsoft SQL Server instance has already been installed by the previous version, Veeam Backup & Replication will connect to the existing database, upgrade it (if necessary) and use it for work.
- Veeam Backup & Replication 7.0 setup performs configuration check to determine if all prerequisite software is available on the machine where you are installing Veeam backup server. If some of the required software components are missing, the setup wizard will offer you to install missing software automatically. You can accept automatic installation, or deploy the missing software manually and then re-check the configuration.

Requirements

This section covers the list of system requirements to the VMware vSphere environment, Veeam Backup & Replication console, virtual machines and backup targets, necessary rights and permissions, as well provides information on ports used by Veeam Backup & Replication.

Platform Support

Veeam Backup & Replication provides full support for the VMware vSphere virtualization platform.

Virtual Infrastructure

Specification	Requirement
Platform	VMware vSphere 5.0, 5.1 VMware vSphere 4.x VMware Infrastructure 3.5
Hypervisor	ESXi 5.0, 5.1 ESX(i) 4.x ESX(i) 3.5 Free ESXi is not supported
Management Server (optional)	vCenter Server 5.0, 5.1 vCenter Server 4.x Virtual Center 2.5

Virtual Machines

Specification	Requirement
Hardware	All types and versions of virtual hardware are supported, except raw disk mapping (RDM) in physical mode and independent disks. You can use disk exclusion functionality to exclude some of the unsupported disks from backup.
OS	Any OS supported by VMware vSphere. Application−aware image−level processing option is supported on Microsoft Windows XP x86, Microsoft Windows 2003, Microsoft Windows Vista, Microsoft Windows 2008, Microsoft Windows 2008 R2, Microsoft Windows 7, Microsoft Windows 8 and Microsoft Windows Server 2012. Windows file-level restore option is supported on NTFS, FAT, FAT32 and ReFS file systems (ReFS is supported only if Veeam Backup & Replication is installed on Windows 2012 Server). To restore files from non-Windows guests (Linux, Solaris, BSD, Novell Netware), use the Multi-OS File Level Restore wizard.
Software	VMware Tools all latest service packs and patches (required for application-aware image-level processing).

vCloud Director

Specification	Requirement
vCloud Director	VMware vCloud Director 5.1

System Requirements

To ensure successful usage of Veeam Backup & Replication, the following system requirements should be met:

Veeam Backup Server

Specification	Requirement
Hardware	CPU: any modern x64 processor (minimum 2 cores recommended) RAM: 4 GB HDD: 2 GB on the system disk (required to extract all components from the setup package during the product installation). 1 GB per 500 VM for guest file system catalog folder (persistent data). Sufficient free disk space for Instant VM Recovery cache (non-persistent data, at least 10GB recommended). Network: 1 Gbps recommended for onsite backup. 1 Mbps or faster WAN for offsite backup and replication. High latency links are supported, but TCP/IP connection must not drop.
OS	Only 64-bit versions of the following operating systems are supported: • Microsoft Windows Server 2008 SP2 • Microsoft Windows Server 2008 R2 SP1 • Microsoft Windows 7 SP1 • Microsoft Windows 8 • Microsoft Windows Server 2012
Software	During the setup, the installer will perform system configuration check to determine if all prerequisite software is available on the machine where you are installing Veeam backup server. If some of the required software components are missing, the setup wizard will offer you to install missing software automatically. This refers to: • Microsoft .NET Framework 4.0 (included in the setup) • Microsoft SQL Server Management Objects • Microsoft SQL Server System CLR Types • Microsoft Visual C++ 2010 Service Pack 1 redistributable package The following software needs to be installed manually: • Microsoft PowerShell 2.00 (required for PowerShell snap-in).
SQL Database	Local or remote installation of the following versions of Microsoft SQL Server are supported: • Microsoft SQL Server 2005 (Full and Express Edition) • Microsoft SQL Server 2008 (Full and Express Edition) • Microsoft SQL Server 2008 R2 (Full and Express Edition; Express Edition is included in the setup) • Microsoft SQL Server 2012 (Full and Express Edition)

Note Due to its limitations, Microsoft SQL Server Express Edition should only be used for evaluation purposes or in case of a small-scale production environment. For environments with a lot of VMs, it is necessary to install a fully functional commercial version of Microsoft SQL Server.

VMware Backup Proxy

Specification	Requirement
Hardware	CPU: any modern x86/x64 processor (minimum 2 cores recommended) RAM: 2 GB RAM + 500MB per each concurrent task. Using faster memory (DDR3) improves data processing performance. HDD: 200 MB free disk space Network: 1 Gbps recommended for onsite backup. 1 Mbps or faster WAN for offsite backup and replication. High latency links are supported, but TCP/IP connection must not drop.
OS	Both 32-bit and 64-bit versions of the following operating systems are supported: • Microsoft Windows XP SP3 • Microsoft Windows Server 2003 SP2 • Microsoft Windows Vista SP2 • Microsoft Windows Server 2008 SP2 • Microsoft Windows Server 2008 R2 SP1 • Microsoft Windows 7 SP1 • Microsoft Windows 8 • Microsoft Windows Server 2012

Backup Repository

Specification	Requirement
Hardware	CPU: any modern x86/x64 processor RAM: 4 GB HDD: 200 MB for Veeam Backup & Replication components and sufficient disk space to store backup files and replicas (high-RPM drives and RAID10 configuration recommended) Network: 1 Gbps recommended for onsite backup. 1 Mbps or faster WAN for offsite backup and replication. High latency links are supported, but TCP/IP connection must not drop.
OS	Both 32-bit and 64-bit versions of the following operating systems are supported: • Linux (SSH and Perl required) • Microsoft Windows XP SP3 • Microsoft Windows Server 2003 SP2 • Microsoft Windows Vista SP2 • Microsoft Windows Server 2008 SP2 • Microsoft Windows Server 2008 R2 SP1 • Microsoft Windows 7 SP1 • Microsoft Windows Server 2012 • Microsoft Windows 8

WAN Accelerator

Specification	Requirement
Hardware	CPU: any modern x64 processor (minimum 2 cores recommended) RAM: 8 GB HDD: minimum 50 GB for WAN accelerator cache Network: 1 Gbps recommended for onsite data transfer. 1 Mbps or faster WAN for offsite data transfer. High latency links are supported, but TCP/IP connection must not drop.
OS	Only 64-bit versions of the following operating systems are supported: • Microsoft Windows Server 2003 SP2 • Microsoft Windows Vista SP2 • Microsoft Windows Server 2008 SP2 • Microsoft Windows Server 2008 R2 SP1 • Microsoft Windows 7 SP1 • Microsoft Windows 8 • Microsoft Windows Server 2012

Veeam Backup Enterprise Manager Server

Specification	Requirement
Hardware	CPU: any modern x64 processor RAM: 4 GB HDD: 2 GB on the system disk (required to extract all components from the setup package during the product installation). Network: 1 Gbit/sec recommended
OS	Only 64-bit versions of the following operating systems are supported: • Microsoft Windows Server 2008 SP2 • Microsoft Windows Server 2008 R2 SP1 • Microsoft Windows 7 SP1 • Microsoft Windows 8 • Microsoft Windows Server 2012
Software	• Microsoft .NET Framework 4 (included in the setup), with update 4.0.3 (http://support.microsoft.com/kb/2600211) • Microsoft SQL Server Management Objects • Microsoft SQL Server System CLR Types • Microsoft Visual C++ 2010 Service Pack 1 redistributable package • Microsoft Internet Information Services 7.0 or later (during setup, the following components are installed automatically: Default Document Component, Directory Browsing Component, HTTP Errors Component, Static Content Component, Windows Authentication Component). **Note:** If the above components are not installed, they can be deployed automatically during the setup Internet Explorer 9.0 or later, latest versions of Mozilla Firefox and Google Chrome are supported; the browser needs to have JavaScript enabled. Microsoft Excel 2003 or later (to view report data exported from Veeam Backup Enterprise Manager).

Specification	Requirement
SQL Database	Local or remote installation of the following versions of Microsoft SQL Server are supported: • Microsoft SQL Server 2005 (Full and Express Edition) • Microsoft SQL Server 2008 (Full and Express Edition) • Microsoft SQL Server 2008 R2 (Full and Express Edition; Express Edition is included in the setup) • Microsoft SQL Server 2012 (Full and Express Edition)

Veeam Backup Search

Specification	Requirement
Hardware	Refer to corresponding Microsoft Search Server version system requirements
OS	Both 32-bit and 64-bit versions of the following operating systems are supported: • Microsoft Windows Server 2003 SP2 • Microsoft Windows Server 2008 SP2 • Microsoft Windows Server 2008 R2 SP1
Software	• Microsoft Search Server 2008 (including Express edition) • Microsoft Search Server 2010 (including Express edition)

Veeam Explorer for Exchange

Specification	Requirement
Microsoft Exchange	Veeam Explorer for Exchange only supports database files (.edb) created with Microsoft Exchange 2010 and Microsoft Exchange 2013. To open database files, Veeam Explorer for Exchange requires a service dynamic link library (ese.dll) which is installed together with Microsoft Exchange.
Software	Restore of folders and items to their original location is available to users of Veeam Backup & Replication Enterprise Edition only. If you plan to export folders and items as Personal Folder Files (.pst), you should have a 64-bit version of Microsoft Outlook 2010 installed on the Veeam Backup & Replication server.

Veeam Explorer for SharePoint

Specification	Requirement
Microsoft SharePoint	Veeam Explorer for SharePoint currently supports primary content database files (.mdf), secondary and log files (.ldf) and remote BLOB stores created with Microsoft SharePoint Server 2010 (Foundation, Server and Enterprise editions), virtualized on VMware or Hyper-V platform.
Staging SQL Server	The staging Microsoft SQL Server must run on the machine where Veeam Explorer for SharePoint is installed (that is, on the machine running Veeam backup server). The staging system must run the same or a later version of Microsoft SQL Server as the server that hosts restored Microsoft SharePoint content databases. **Note:** As the staging system, you can use the Microsoft SQL Server Express 2008 R2 SP1 that is shipped with the Veeam Backup & Replication setup. However, consider that content databases that exceed 10 GB cannot be attached to this SQL Server.

Required Permissions

The accounts used for installing and using Veeam Backup & Replication should have the following permissions:

Account	Required Permission
Setup Account	*Local Administrator* permissions on the Veeam backup server to install Veeam Backup & Replication
Target/Source Host Permissions	Root permissions on the source ESX(i) host. *Root* (or equivalent) permissions on Linux backup repository. *Write* permission on the target folder and share. If vCenter Server is used, administrator credentials are required.
SQL Server	The account used to run Veeam Backup Management Service must have **database owner** role for the *VeeamBackup* database on the SQL Server instance. The account used to run Veeam Backup Enterprise Manager must have **database owner** role for the *VeeamBackupReporting* database on the SQL Server instance.
Veeam Backup Enterprise Manager	*Local Administrator* permissions on the Veeam Backup Enterprise Manager server to install Veeam Backup Enterprise Manager. To be able to work with Veeam Backup Enterprise Manager, users should be assigned the *Portal Administrator*, *Restore Operator* or *Portal User* role.
Veeam Backup Search	*Local Administrator* permissions on the Microsoft Search Server to install Veeam Backup Search
Veeam Explorer for Exchange	Full access to Microsoft Exchange database and its log files for item recovery. The account you plan to use for recovery should have both read and write permissions to all files in the folder with the database. Access rights can be provided through impersonation, as described in the Configuring Exchange Impersonation article.
Veeam Explorer for SharePoint	The account used for working with Veeam Explorer for SharePoint requires membership in the **sysadmin** fixed server role on the staging Microsoft SQL Server. The account used for connection with target SharePoint server where document item(s)/list will be restored needs the following: • If permissions of the item being restored are inherited from the parent item (list) - **Full Control** for that list is required. • If permissions are not inherited, and restored item will replace an existing item - then **Contribute** for the item and **Full Control** for its parent list are required.

Used Ports

This section covers typical connection settings for the Veeam Backup & Replication components.

Veeam Backup Server Connections

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

From	To	Protocol	Port	Notes
Veeam backup server	vCenter Server	HTTPS	443	Default VMware web port that can be customized in vCenter settings
	ESX(i) server	HTTPS	443	Default VMware web port that can be customized in ESX host settings. Not required if vCenter connection is used.
		TCP	902	VMware data mover port.
		TCP	22	Default SSH port used as a control channel, only for jobs with a full ESX target with the console agent enabled.
	Linux server	TCP	22	Port used as a control channel from the console to the target Linux host.
	Windows server	TCP UDP	135, 137 to 139, 445	Ports required for deploying Veeam Backup & Replication components.
		TCP	6160	Default port used by the Veeam Installer Service.
		TCP	6162	Default port used by the Veeam Backup Proxy Service.
		TCP	6161	Default port used by the Veeam vPower NFS Service.
		TCP UDP	111, 2049+, 1058+	Standard NFS ports. Note: If ports 2049 and 1058 are occupied, the succeeding port numbers will be used).
Linux server	Veeam backup server	TCP	2500 to 5000	Default range of ports used as transmission channels for jobs writing to Linux target. For every TCP connection that a job uses, one port from this range is assigned.
Windows server	Veeam backup server	TCP	2500 to 5000	Default range of ports used as transmission channels for jobs writing to Windows target. For every TCP connection that a job uses, one port from this range is assigned.

Backup Proxy Connections

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

From	To	Protocol	Port	Notes
Communication with VMware Servers				
VMware Backup Proxy	vCenter Server	HTTPS	443	Default VMware web service port that can be customized in vCenter settings.
	ESX(i) server	TCP	902	VMware data mover port.
		HTTPS	443	Default VMware web service port that can be customized in ESX host settings. Not required if vCenter connection is used.
Communication with Backup Repositories				
VMware Backup Proxy	Linux server	TCP	22	Port used as a control channel from the backup proxy to the target Linux host.
	Shared folder CIFS (SMB) share	TCP UDP	135, 137 to 139, 445	Ports used as a transmission channel from the backup proxy to the target CIFS (SMB) share.
Linux server	VMware Backup Proxy	TCP	2500 to 5000	Default range of ports used as transmission channels for backup jobs. For every TCP connection that a job uses, one port from this range is assigned.
Windows server	VMware Backup Proxy	TCP	2500 to 5000	Default range of ports used as transmission channels for backup jobs. For every TCP connection that a job uses, one port from this range is assigned.
Communication with Backup Proxies				
VMware Backup Proxy	VMware Backup Proxy	TCP	2500 to 5000	Default range of ports used as transmission channels for replication jobs. For every TCP connection that a job uses, one port from this range is assigned.

WAN Accelerator Connections

The following table describes network ports that must be opened to ensure proper communication between WAN accelerators.

From	To	Protocol	Port	Notes
Communication with Veeam Backup Server				
Veeam backup server	WAN accelerator (source and target)	TCP	6160	Default port used by the Veeam Installer Service.
		TCP	6164	Controlling port for RPC calls.
Communication with Backup Repositories				
WAN Accelerator (source and target)	Backup repository (source and target)	TCP	2500 to 5000	Default range of ports used by data transport services for transferring files of a small size such as .nvram, .vmx, .vmxf, GuestIndexData.zip and others. A port from the range is selected dynamically.
Communication Between WAN Accelerators				
WAN accelerator	WAN accelerator	TCP	6164	Controlling port for RPC calls.
		TCP	6165	Default port used for data transfer between WAN accelerators. Ensure this port is open in firewall between sites where WAN accelerators are deployed.

VM Guest OS Connections

The following table describes network ports that must be opened to ensure proper communication of the Veeam backup server with the runtime coordination process deployed inside the VM guest OS for application-aware image processing and indexing.

From	To	Protocol	Port	Notes
Veeam backup server	Runtime coordination process inside the VM guest OS	TCP	1025 to 5000 (for Windows 2003) 49152-65535 (for Windows 2008)	Dynamic RPC port range used by the runtime coordination process deployed inside the VM guest OS for application-aware image processing (when working over network, not over VIX API).*

* If you use default Windows firewall settings, you do not need to configure dynamic RPC ports: during setup, Veeam Backup & Replication automatically creates a firewall rule for the runtime process. If you use firewall settings other than default ones or application-aware image processing fails with the "*RPC function call failed*" error, you need to configure dynamic RPC ports. For more information, see http://technet.microsoft.com/en-us/library/cc732839(WS.10).aspx.

Veeam Backup Enterprise Manager Connections

The following table describes network ports that must be opened to ensure proper communication of Veeam Backup Enterprise Manager with other components.

From	To	Protocol	Port	Notes
Veeam Backup Enterprise Manager	Veeam Backup Server	TCP	9392	Default port used by Veeam Backup Enterprise Manager for collecting data from Veeam backup servers. Can be customized during Veeam Backup & Replication installation.
			9393	Default port used by the Veeam Backup Catalog Service for catalog replication. Can be customized during Veeam Backup & Replication installation.
			2500 to 2600	Ports used by the Veeam Backup Catalog Service for replicating catalog data.
	Microsoft Search Server	TCP	9395	Default port used by the Veeam Backup Search service integration component. Can be customized during Veeam Backup Search installation.
IIS extension	Veeam Backup Enterprise Manager	TCP	9393	Default port used to enable file search. Can be customized during Veeam Backup Enterprise Manager installation.
			9394	Default port used by IIS extension to communicate with Veeam Backup Enterprise Manager. Can be customized during Veeam Backup Enterprise Manager installation.
Browser	IIS extension	HTTP	9080	Default ports used to communicate with the website. Can be customized during Veeam Backup Enterprise Manager installation.
		HTTPS	9443	
Enterprise Manager Web API Client	Enterprise Manager Web API	HTTP	9399	Default HTTP and HTTPS ports used to communicate with Veeam Backup Enterprise Manager Web API. Can be customized during Veeam Backup Enterprise Manager installation.
		HTTPS	9398	

Veeam U-AIR Wizards Connections

The following table describes network ports that must be opened to ensure proper communication of U-AIR wizards with other components.

From	To	Protocol	Port	Notes
U-AIR Wizards	Veeam Backup Enterprise Manager	TCP	9394	Default port used for communication with Veeam Backup Enterprise Manager. Can be customized during Veeam Backup Enterprise Manager installation.
	Local inter-process communication	TCP	9396	Default port used for interaction between the U-AIR wizard and Session Manager running on the machine where the U-AIR wizard is installed.

Note During installation, Veeam Backup & Replication automatically creates firewall rules for default ports to allow communication for the application components.

Hardware Recommendations

- Using faster processors configuration on the Veeam Backup & Replication console generally improves the backup performance. We recommend installing Veeam Backup & Replication on powerful computers with multi-core processors (Intel Core Duo/Quad, AMD Phenom X2/X4).
- You can additionally improve the backup speed by ensuring that a backup file is saved to the fast storage (high-RPM hard drives, RAID10 configurations).

DEPLOYMENT

This section describes the procedure for installing, upgrading and removing Veeam Backup & Replication. You will also find here information about Veeam Backup & Replication licensing, differences between product editions and functionality modes and Veeam Backup & Replication interface.

Installing Veeam Backup & Replication

Before you begin the installation process, take the following steps to prepare for deployment:

- **Check platform-specific and system requirements**. Make sure the computer on which Veeam Backup & Replication is to be installed meets the system requirements. For details, see System Requirements.
- **Check account permissions**. Make sure all accounts you will be using have sufficient permissions. You will not be able to use Veeam Backup & Replication successfully if the accounts do not have required permissions. For details, see Required Permissions.
- **Check ports**. Communication between components requires a number of ports to be open. Carefully plan your backup strategy and infrastructure layout. For details, see Used Ports.

To install Veeam Backup & Replication, follow the next steps:

Step 1. Start the Setup Wizard

Download the latest version of Veeam Backup & Replication installation image from www.veeam.com/downloads.

1. 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.
2. 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.

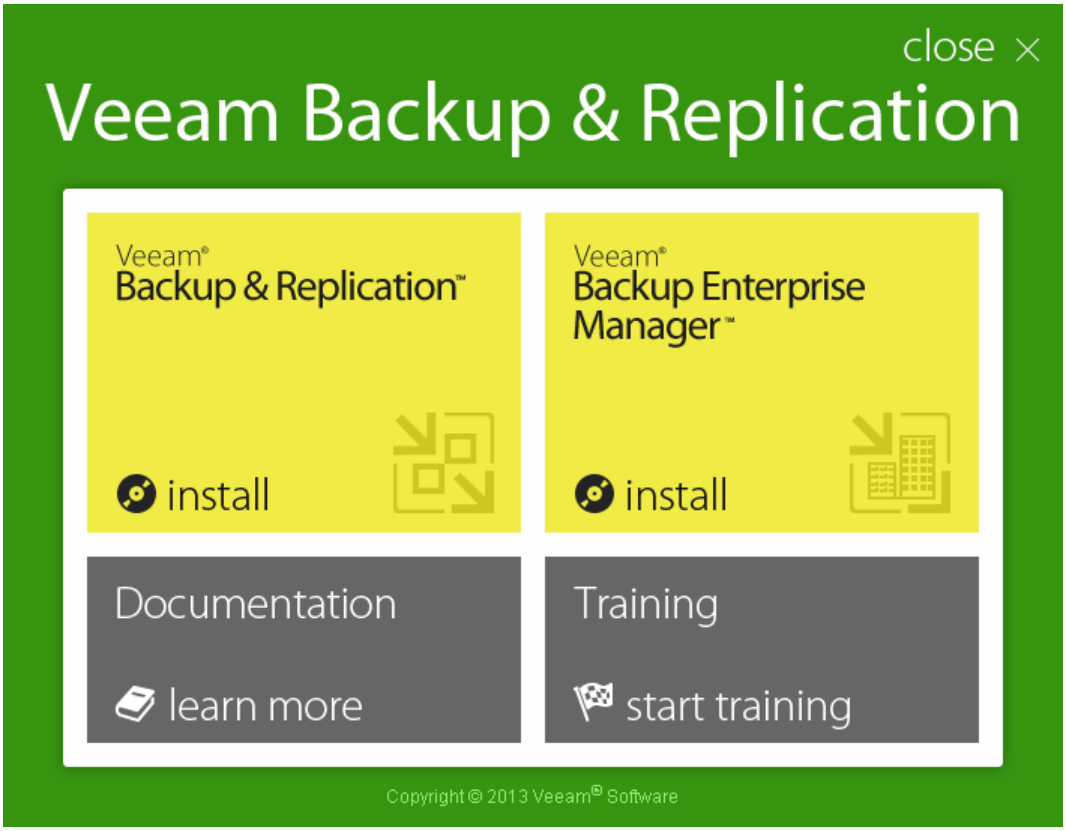

3. Click the **Install** link in the **Veeam Backup & Replication** section of the splash screen.
4. On the **Welcome** step of the wizard, click **Next** to start the installation.

Important! It is strongly recommended to run the setup using **Autorun**. Otherwise (if you use other installation files from CD/DVD folders), you may miss some components that need to be installed and Veeam Backup & Replication will not work as expected.

Step 2. Read and Accept License Agreement

To install Veeam Backup & Replication, you must accept the license agreement. Read the license agreement, select the **I accept the terms in the license agreement** option and click **Next**.

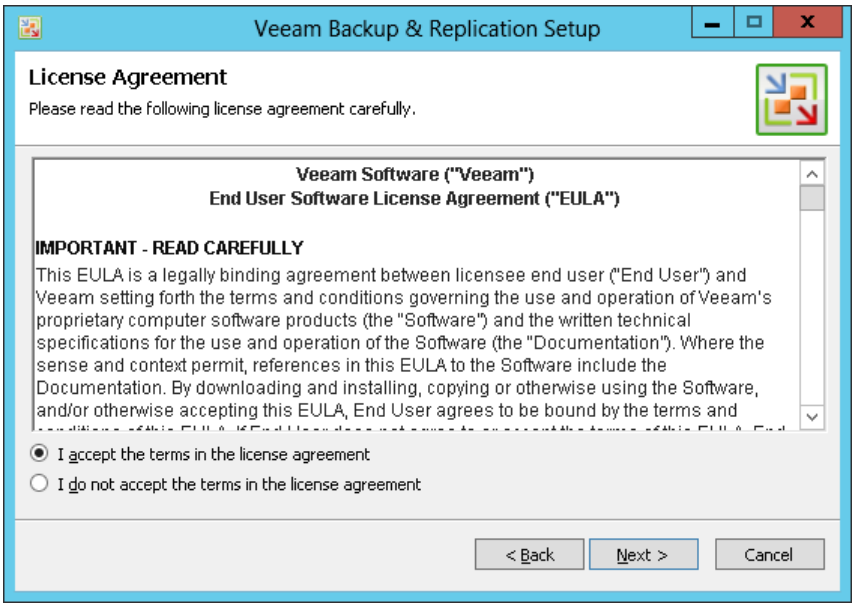

Step 3. Provide a License File

You can install Veeam Backup & Replication with a trial license that was sent to you after registration, a purchased full license or without any license at all. In the latter case, Veeam Backup & Replication will be run in the free functionality mode.

To install a license, click **Browse** and select a valid license file for Veeam Backup & Replication.

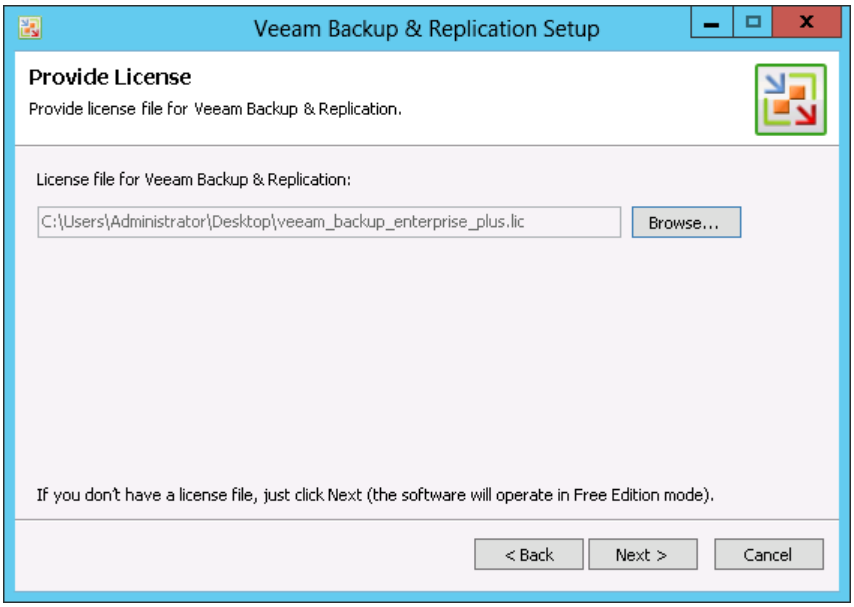

To learn about obtaining a license file, see Veeam Backup & Replication Licensing.

Note If a valid license file has been previously installed on the machine, the setup wizard will inform you about it. In this case, you can skip this step and move forward.

Step 4. Select Components

Select the components you want to install. The Veeam Backup & Replication setup includes the following components:

- Veeam Backup & Replication
- Veeam Backup Catalog responsible for indexing VM guest OS files
- Veeam Backup PowerShell snap-in for automating backup and replication activities via scripts. Note that the Veeam Backup PowerShell component is disabled by default. To be able to install it, you need to install the Windows Management Framework Core package first.

If necessary, you can change the installation folder.

Note Veeam Backup & Replication requires Microsoft .NET Framework 4 and Microsoft SQL Server (.NET Framework 4 and Microsoft SQL Server 2008 R2 Express are included in the setup). If you plan to install these components with the setup, you should have at least 2 GB of free space on the system disk.

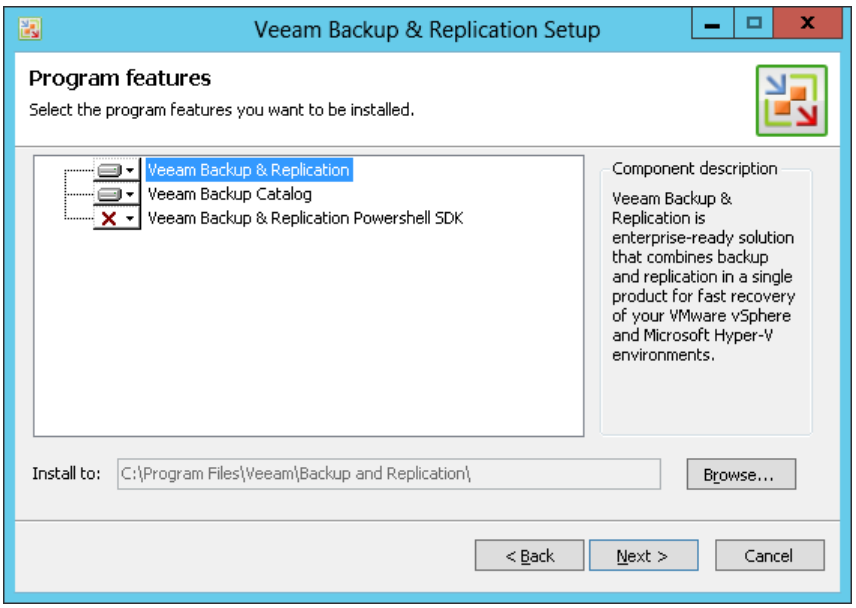

The setup will also install the following components in the silent mode:

- Veeam Explorer for Exchange
- Veeam Explorer for SharePoint
- HP StoreServ Plug-in
- HP StoreVirtual Plug-in

These components do not require additional licenses; they are integrated with Veeam Backup & Replication. Veeam Explorer for Exchange and Veeam Explorer for SharePoint can be launched either from the management console or from the **Start** menu. Refer to the product documentation for details.

Before proceeding with the installation, the setup wizard will perform a system configuration check to determine if all prerequisite software is available on the machine.

- To learn what software is required for the installation, see the System Requirements section.
- If some of the required software components are missing, the wizard will offer you to install missing software automatically.

- To install missing software components in the current session without interrupting the setup, click the **Install** button.
- If you cancel automatic installation, you should install and enable the missing software manually; otherwise, you will not be able to proceed to the next step. When all required software is installed, click **Re-run** to repeat verification.

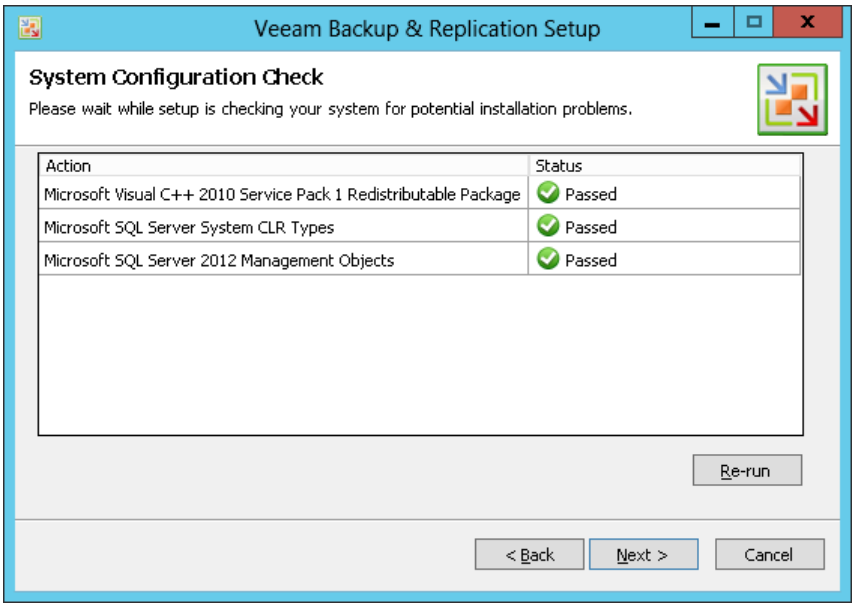

Step 6. Select a SQL Server Instance

At this step, you should select an SQL Server instance on which the Veeam Backup & Replication database should be created or choose to install a new SQL Server instance.

- If the SQL Server is not installed, select the **Install new instance of SQL Server** option.
- If the SQL Server is already installed, select the **Use existing instance of SQL Server** option. Enter the instance name in the *HOSTNAME\INSTANCE* format. In the **Database** field, specify the name of the database to be used.

The user account under which the installation is being performed should have sufficient rights to log on to the selected SQL Server instance using Windows integrated authentication and create a database on the selected instance.

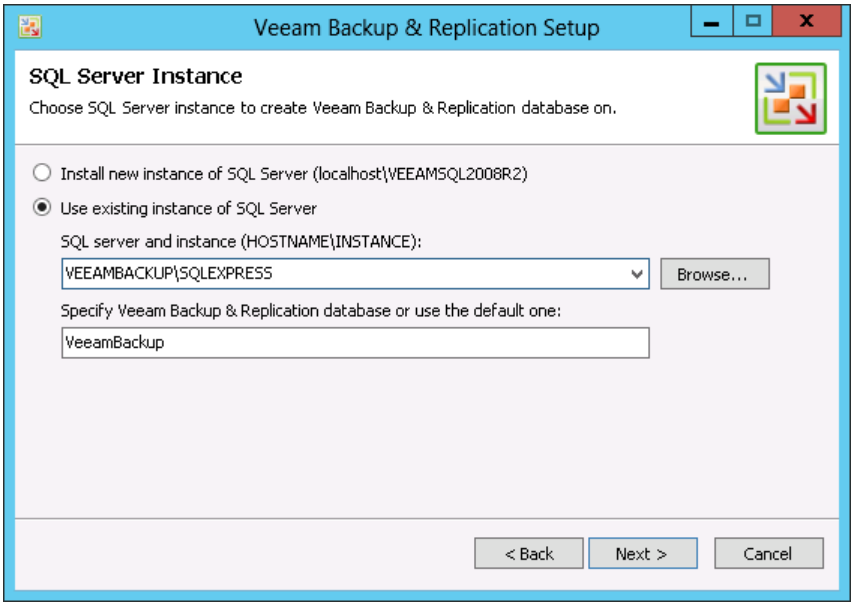

In case the Veeam Backup & Replication database already exists on the SQL Server instance (that is, it was created by the previous installations of Veeam Backup & Replication), the setup will display a warning notifying about it. Click **Yes** to connect to the detected database. If necessary, the existing database will be upgraded to the latest version.

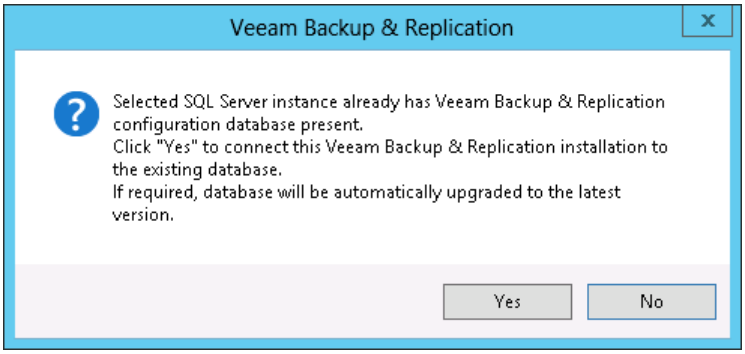

Step 5. Specify Service Account Credentials

Enter the administrative credentials of the account under which you want to run the Veeam Backup Service. The user name should be specified in the *DOMAIN\USERNAME* format. This user account must have:

- *Database owner* rights for the Veeam Backup & Replication database on the SQL Server instance
- *Full control* NTFS permissions on the `VBRCatalog` folder where index files are stored.

The *Log on as* right will be automatically granted to the specified user account.

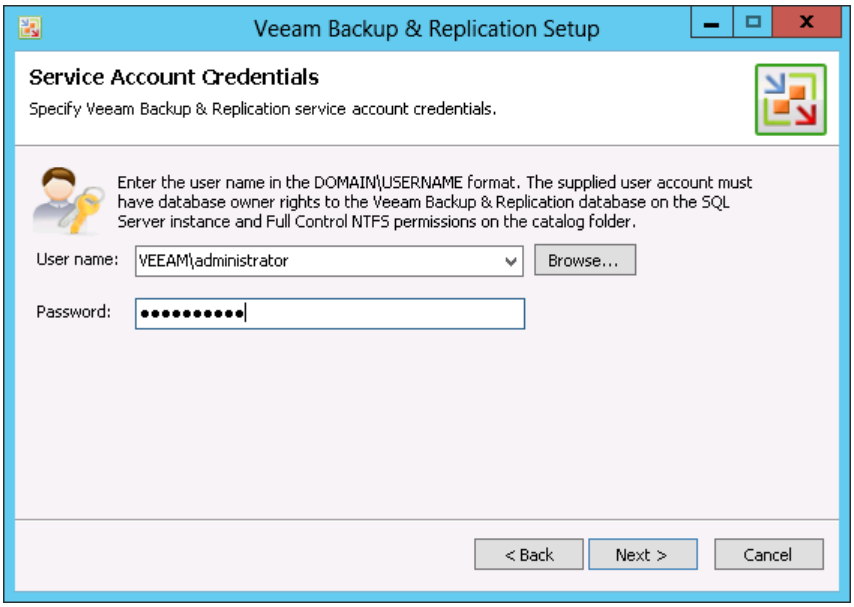

Step 7. Specify Service Ports

If required, update the port numbers to be used by the following components:

- Veeam Backup Service (default value is **9392**)
- Veeam Backup Catalog Service (default value is **9393**)

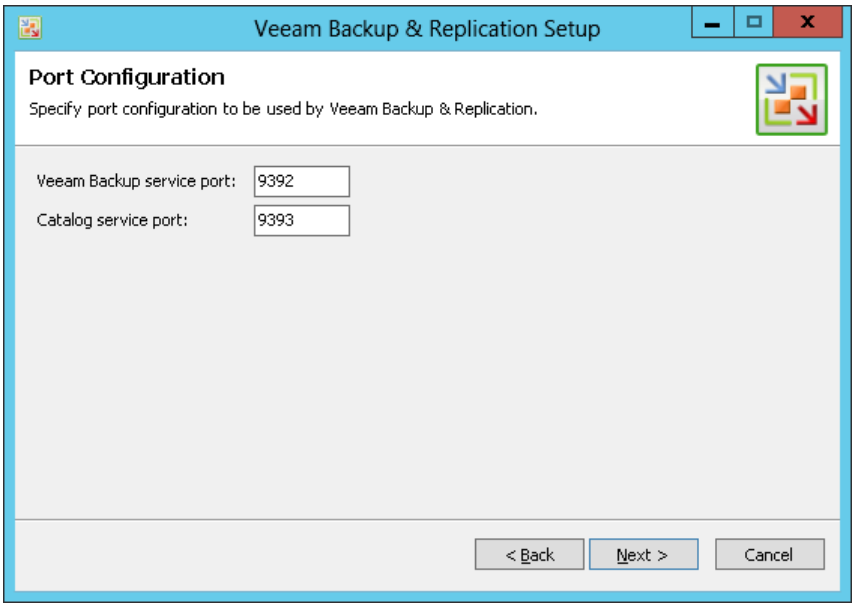

Step 8. Specify Directories

In the vPower NFS section, specify the folder to which instant VM recovery write cache will be stored. The default location is `C:\ProgramData\Veeam\Backup\NfsDatastore`. Please note that the selected volume should have at least 10 GB of free disk space.

Specify the name and destination for the catalog folder where index files should be stored. The default location is `C:\VBRCatalog`.

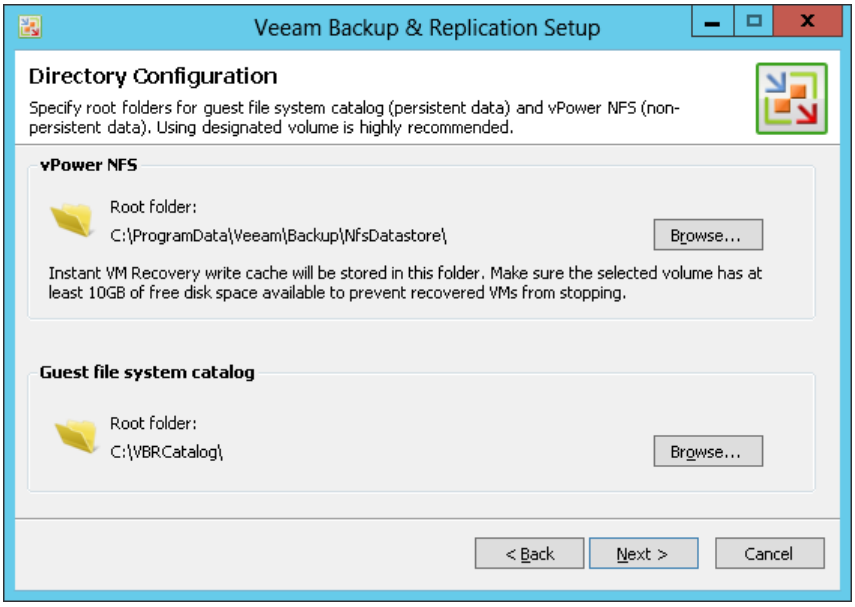

Step 9. Begin Installation

Finally, review the installation settings. You can go back, review and modify previous steps using the **Back** button. If you are sure that all settings are configured correctly, click **Install** to begin the installation.

When the installation completes, click **Finish** to exit the setup wizard. You can now start Veeam Backup & Replication.

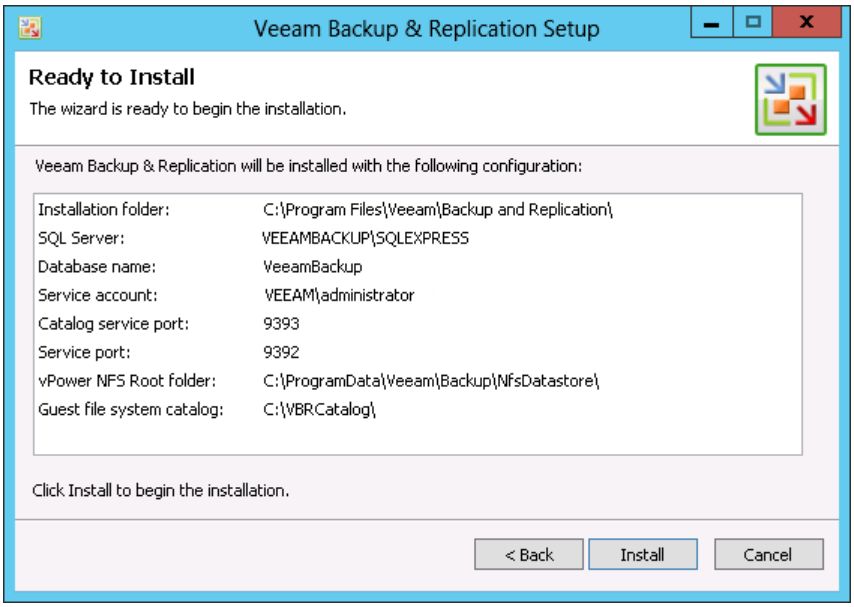

Upgrading Veeam Backup & Replication

Upgrade to version 7.0 is supported for Veeam Backup & Replication 6.0, 6.1, and 6.5. To perform upgrade, run the Veeam Backup & Replication setup file. For details on the upgrade procedure, refer to Veeam Backup & Replication 7.0 Release Notes on the resources page.

Before You Upgrade

- When you upgrade the product, the Veeam Backup & Replication database gets upgraded, too, and becomes available for usage with the newly installed version. It is strongly recommended that you perform backup for the Veeam Backup & Replication database before upgrading the product. In this case, you will be able to easily go back to a previous version in case of issues with upgrade. To perform database backup, you can use the native Microsoft SQL tools. If you are upgrading from version 6.5, make sure your existing Veeam Backup & Replication configuration was backed up, as described in product documentation. If you are upgrading from version 6.5, you can perform configuration backup.
- In Veeam Backup & Replication 7.0, support for legacy replica jobs is discontinued. Before upgrade, it is strongly recommended to remove legacy replica jobs or re-create and re-run them with Veeam Backup & Replication 6.x. Remove from inventory replica VMs that were created with legacy replica jobs. If legacy replica jobs or replica VMs created with such jobs are detected, upgrade will fail.
- In version 7.0, installation of Veeam Backup & Replication server component on a number of operating systems is no longer supported. If your current version of Veeam Backup & Replication server runs on an OS that is no longer supported, it is recommended to install Veeam Backup & Replication 7.0 on a machine with a supported OS. However, if you want to keep your current configuration, you can follow the instructions provided in section Upgrade for OSs with Discontinued Support. For details on supported platforms and configurations, see System Requirements.

After You Upgrade

When you upgrade the product, the Veeam Backup & Replication database gets upgraded, too, and becomes available for usage with the newly installed version. Once upgrade is completed, it is recommended to check all existing jobs, backup infrastructure components and make changes to the default upgrade configuration if necessary. Consider the following:

- During upgrade to Veeam Backup & Replication 7.0, backup files and VM replicas that were created with 6.x are not impacted in any way.
- Compression levels configured for 6.x jobs will be preserved. However, consider that compression level named **Optimal** in version 6.5 will be renamed and appear as **High** after the upgrade.
- The **Use storage snapshots** option will be disabled for all backup and replication jobs.
- After the upgrade, the **Enable parallel processing** option in Veeam Backup & Replication options is turned off.
- Credentials Manager will be populated with credentials that were used for connecting to virtual infrastructure servers, Windows servers, Linux servers and credentials used for guest processing.

Upgrade for OSs with Discontinued Support

In version 7.0, installation of Veeam Backup & Replication server component on a number of operating systems is no longer supported. If your current version of Veeam Backup & Replication server runs on an OS that is no longer supported by version 7.0, follow the instructions provided in this section.

- Upgrading from Veeam Backup & Replication 6.5
- Upgrading from Veeam Backup & Replication 6.0 or 6.1

Upgrading from Veeam Backup & Replication 6.5

To upgrade from Veeam Backup & Replication 6.5 running on an OS that is no longer supported, follow the next steps:

1. Wait for all jobs, restore and other data protection operations to complete.
2. Disable jobs that are scheduled to run automatically.
3. Perform configuration backup as described in section Running Configuration Backups Manually.
4. Install Veeam Backup & Replication 6.5 on a machine running OS that is supported by Veeam Backup & Replication 7.0. For details on supported platforms and configurations, see System Requirements. You can point Veeam Backup & Replication to a SQL Server installed locally or choose a remote SQL Server.

Note You can point the new instance of Veeam Backup & Replication 6.5 to the existing database. In this case, you will not need to restore configuration backup as described in step 11. However, after installation of Veeam Backup & Replication 6.5 is complete, you will need to re-define all credentials used for connecting to virtual infrastructure servers, Windows servers, Linux servers and credentials used for guest processing.

5. If you have previously applied hotfixes, re-apply these hotfixes to the new instance of Veeam Backup & Replication 6.5.
6. If you have previously changed registry values for Veeam Backup & Replication, move these changes to the new instance of Veeam Backup & Replication 6.5.

7. Move all backups and replica metadata files stored in the default backup repository and all local repositories to the machine running the new instance of Veeam Backup & Replication 6.5. It is required that you preserve exact paths to folders where backups or replica metadata filesare stored.
8. Move the configuration backup (created at step 3) to the machine running the new instance of Veeam Backup & Replication 6.5.
9. Move the contents of the `VBRCatalog` folder.
10. Start Veeam Backup & Replication console to upgrade components on connected servers.
11. Restore configuration from backup as described in section Restoring Configuration Data.
12. Check backup infrastructure components and all existing jobs to make sure configuration settings have been properly restored.
13. Remove all legacy replica jobs (if present). Remove from inventory and replica VMs that were created with legacy replica jobs.
14. Run upgrade to Veeam Backup & Replication 7.0.

Upgrading from Veeam Backup & Replication 6.0 or 6.1

To upgrade from Veeam Backup & Replication 6.0 or 6.1 running on an OS that is no longer supported, follow the next steps:

1. Upgrade Veeam Backup & Replication server component to version 6.5.
2. Run the upgrade procedure described in section Upgrading from Veeam Backup & Replication 6.5.

Uninstalling Veeam Backup & Replication

To uninstall Veeam Backup & Replication:

1. From the **Start** menu, select **Control Panel > Programs and Features**.
2. In the programs list, right-click **Veeam Backup & Replication** and select **Uninstall**. Note that if you have Veeam Backup Enterprise Manager installed on this machine, too, then this will uninstall both components. Wait for the process to complete.

The SQL database instance installed and used by Veeam Backup & Replication is not removed during the uninstall process. All jobs data stored in it remains as well.

Veeam Backup & Replication Licensing

Veeam Backup & Replication is licensed per CPU Socket ("CPU Sockets") for each Managed Server. For more information, see www.veeam.com/eula.html.

The trial license key is sent to you after registering the product with Veeam Software at: www.veeam.com/vmware-esx-backup/download.html. After registering the product you will receive a trial license key. The trial license is valid for 30 days from the moment of registration.

To obtain a full license key for the desired number of sockets, refer to www.veeam.com/buy-end-user.html. The full license includes a one-year maintenance plan. To renew your maintenance plan, please contact Veeam customer support at: renewals@veeam.com.

Installing a License

At installing Veeam Backup & Replication, you will be asked to specify the license file that was sent to you after registration. If you do not have a license, Veeam Backup & Replication will be run in the free functionality mode.

To view information on the currently installed license, select **Help** > **License** from the main menu. To change the license, click **Install License** and browse to the necessary .lic file.

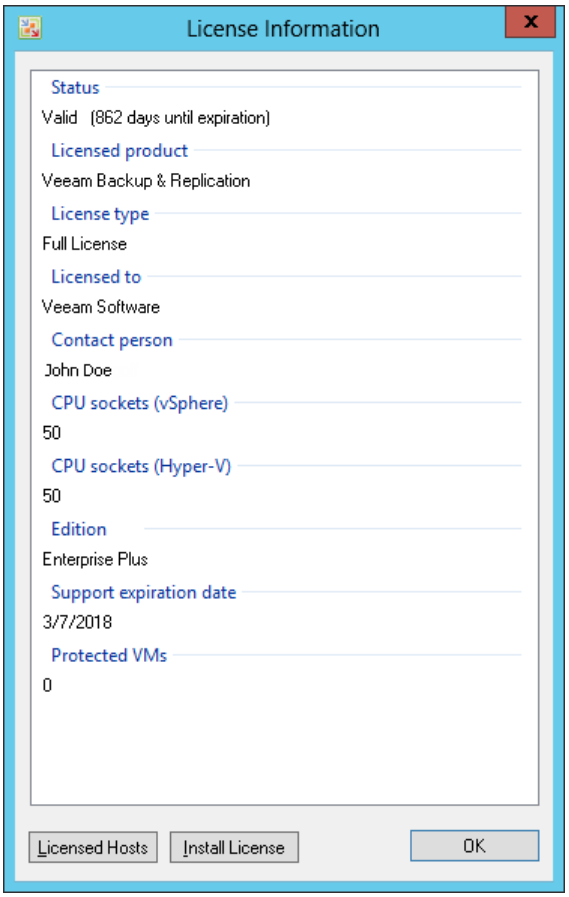

If Veeam Backup & Replication servers are connected to Veeam Backup Enterprise Manager, Veeam Backup Enterprise Manager collects information about all licenses installed on backup servers added to it. You can so manage and activate licenses for the whole of the backup infrastructure from Veeam Backup Enterprise Manager and thus reduce administration overhead.

When Veeam Enterprise Manager replicates databases from backup servers, it also synchronizes license data: checks if the license installed on the backup server coincides with the license installed on the Veeam Backup Enterprise Manager server. If the licenses do not coincide, the license on the backup server will be automatically replaced with that on Veeam Backup Enterprise Manager.

Revoking Servers from the License

Veeam Backup & Replication offers you a possibility to revoke unused ESX(i) hosts or Hyper−V hosts from the license.

When you run a job that uses a specific host, a license is applied to it. However, you may want to revoke the license applied by this host and re-use it for some other host. Revoking a host from the license may be required if the host to which the license is applied does not need backup or replication anymore, for example, in case it is no longer used.

To display the list of licensed hosts:

1. Select **Help** > **License** from the main menu.
2. In the displayed window, click **Licensed Hosts**. As a result, the list of hosts using the license will be displayed.

The **Licensed Hosts** list displays all hosts to which the license is applied. When you start Veeam Backup & Replication for the first time, the list will be empty. After you run a backup or replication job targeted at some objects, this section will display a list of servers that were engaged in the job, with the number of sockets per each.

To revoke a specific server, select it in the list and click **Revoke**. Licensed sockets used by it will be freed and will become available for use by other servers.

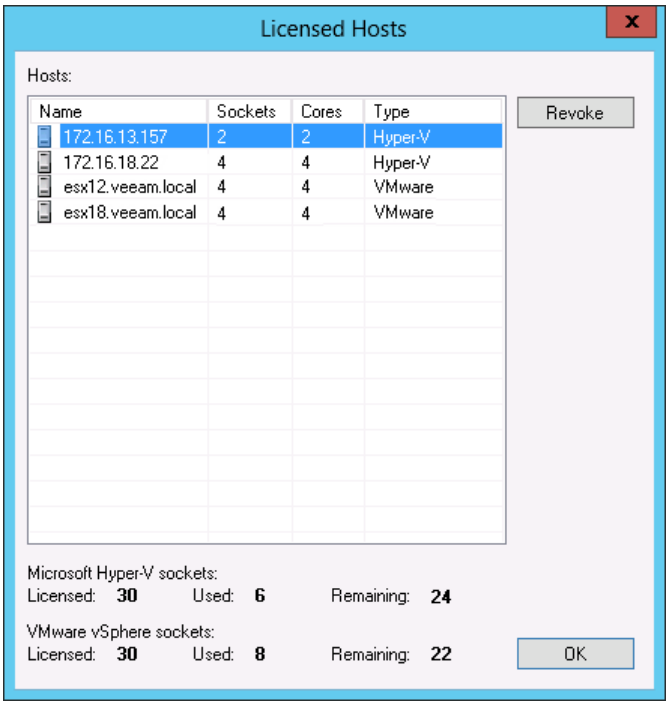

Product Editions

Veeam Backup & Replication is available in Standard, Enterprise and Enterprise Plus editions. The Enterprise and Enterprise Plus editions include additional features to accommodate the requirements of large enterprise environments. The differences in features for Standard, Enterprise and Enterprise Plus editions applicable to VMware vSphere environments are shown in the table below.

Feature	Standard Edition	Enterprise Edition	Enterprise Plus Edition
Backup Copy Job	Available Support for copying jobs to remote locations over WAN (direct operations only).	Available Direct operations only.	Available Direct operations and WAN acceleration.
Built-in WAN Acceleration	Not available	Not available	Available Support for caching, variable-length deduplication and optimizations for transferring Veeam backups across the WAN.
Backup from Storage Snapshots	Not available	Not available	Available Support for creating backups and replicas from SAN snapshots.
Native Tape Support	Limited Windows files archiving only.	Full support Includes archiving Veeam backups to tape, with full tracking of backups and restore points.	Full support
Support for vCloud Director	Limited Visibility of vCloud Director (vCD) infrastructure, backup via VeeamZIP (including backup of vApp and VM meta data and attributes) and direct restore to vCD.	Full support Also includes scheduled incremental backup jobs of vCD VMs.	Full support
Universal Application-Item Recovery	Not available	Available Includes specialized Active Directory, Microsoft Exchange and Microsoft SQL wizards as well as the universal wizard for any application.	Available
SureBackup Recovery Verification	Manual You can verify the recoverability of backup files by mounting VM disks from backup files using Instant VM Recovery and manually testing them.	Automated You can select to perform automatic recovery verification jobs after every backup and verify any restore point.	Available

Feature	Standard Edition	Enterprise Edition	Enterprise Plus Edition
SureReplica Recovery Verification	Not available	Available You can automatically verify every restore point in every replica.	Available
On-Demand Sandbox	Not available	Available You can run one or several VMs from backup in an isolated environment, providing a working copy of the production environment for troubleshooting, testing, training and so on.	Available
File system indexing	Restricted You can browse and search for files in backups which are currently on disk.	Not restricted You can browse and search for files in both current and archived backups (for example, backups which have been moved to tape storage).	Not restricted
1-Click Restore	Not available	Available Allows help desk administrators to restore VMs and guest files through the Enterprise Manager web UI.	Available
Delegation and self- recovery of VMs and guest files	Not available	Not available	Available Delegation options allow authorized users to restore VMs and guest files through the web UI.
Job cloning and editing via the web UI	Not available	Available You can clone existing jobs and edit their settings from the Veeam Backup Enterprise Manager web UI.	Available
Veeam Explorer for Microsoft Exchange	Limited support You can browse and restore mail items via save, send and export.	Full support Also includes restore to original location.	Full support
Veeam Explorer for Microsoft SharePoint	Limited support You can browse and restore documents and lists via save, send and export.	Full support Also includes restore to original location	Full support
Veeam Power-Shell snap-in	Available	Available	Available

Feature	Standard Edition	Enterprise Edition	Enterprise Plus Edition
Veeam Backup Enterprise Manager Web API	Not available	Not available	Available

All editions are installed with the same setup file, however, the extra functionality becomes available only after installing a full license for Veeam Backup & Replication Enterprise Edition or Enterprise Plus Edition. You can install the license for the necessary version during the setup process and change the license file later. You can also change the type of license used. For details, see Veeam Backup & Replication Licensing.

Full and Free Functionality Modes

Veeam Backup & Replication can operate in two functionality modes: full mode and free mode.

- When you run Veeam Backup & Replication in the full functionality mode, you get a commercial version of the product that provides access to all functions.
- When you run Veeam Backup & Replication in the free functionality mode, you get a free version of the product that offers limited capabilities: you can back up single VMs (create VeeamZIP files), recover VM data from backups, perform file copy operations, migrate VMs, restore VM data from HP SAN snapshots, archive files to tape, restore items from Exchange and Sharepoint backups locally and perform configuration backup and restore.

If you have a valid license installed, Veeam Backup & Replication operates in the full functionality mode. As soon as your license expires, you will be offered to install a new license or switch to the free functionality mode. To switch to the free mode, select **View** > **Free functionality only** from the main menu.

To switch back to the full mode, do either of the following:

- Install a valid license: select **Help** > **License** from the main menu. In the displayed window, click **Install License** and select the necessary license file.
- Select **View** > **Full functionality (advanced)** from the main menu. Note that if you do not have a valid license installed, you will not be able to use the functionality provided by the full mode.

Getting to Know User Interface

The user interface of Veeam Backup & Replication is designed to let you quickly find commands you need and perform necessary data protection and disaster recovery tasks. This section will familiarize you with elements of the application user interface.

Main Menu

The main menu in Veeam Backup & Replication contains commands related to general application settings. You can perform the following operations using the main menu:

- Update components installed on servers connected to Veeam Backup & Replication
- Start PuTTy
- Set up user roles
- Configure traffic throttling rules
- Manage credentials
- Enable and disable HP SAN plug-ins
- Perform configuration backup and restore
- Define general application options
- View program help, work with licenses and program logs
- Exit Veeam Backup & Replication

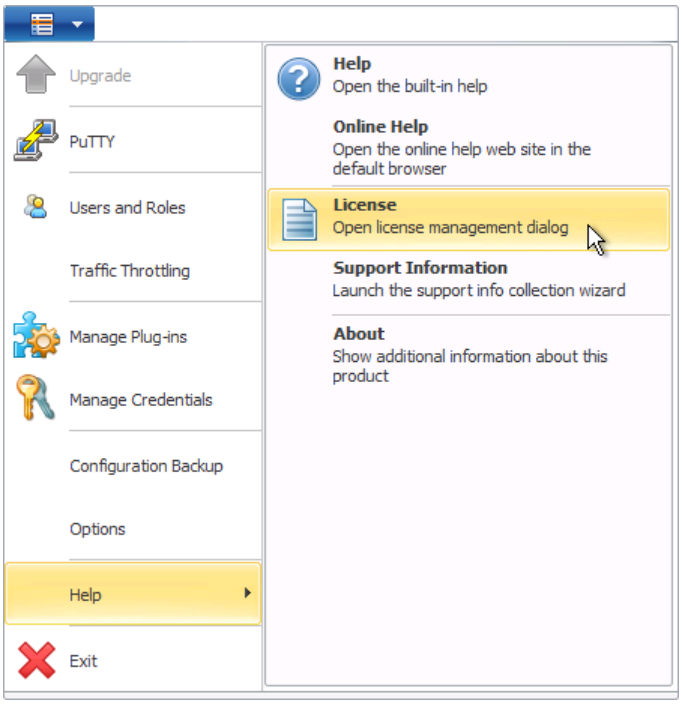

Navigation Pane

The navigation pane, located on the left side of the window, provides centralized navigation and enables you to easily access Veeam Backup & Replication items organized in views.

The navigation pane is made up of two areas:

- The upper, or the inventory pane, displays a hierarchy or a list of all items relevant for a specific view. The content of the inventory pane is different for different views. For example, in the **Infrastructure** view, the inventory pane displays a list of backup infrastructure components — virtual infrastructure servers, backup proxies and backup repositories. In the **Virtual Machines** view, the inventory pane displays a list of servers connected to Veeam Backup & Replication.
- The lower pane contains a set of buttons that enable you to switch between Veeam Backup & Replication views.

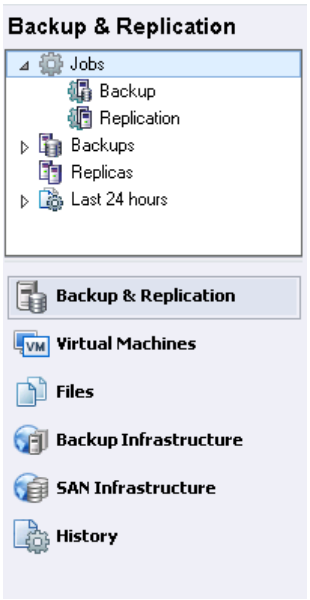

Ribbon and Tabs

Operation commands in Veeam Backup & Replication are organized into logical groups and collected together under tabs on the ribbon. The ribbon is displayed at the top of the main application window; it contains the **Home** tab that is always present, and context-sensitive tabs.

- The **Home** tab provides quick access to the most commonly performed operations. It enables you to create different types of jobs, perform restore and import operations. This tab is always displayed, no matter which view is currently opened.
- Context-sensitive tabs contain commands specific for certain items and appear when these items are selected. For example, if you open the **Backup & Replication** view and select a backup job in the working area, the **Jobs** tab containing buttons for operations with jobs will appear on the ribbon. In a similar manner, if you open the **Files** view and select a file or folder, the **File** tab containing buttons for operations with files will appear on the ribbon.

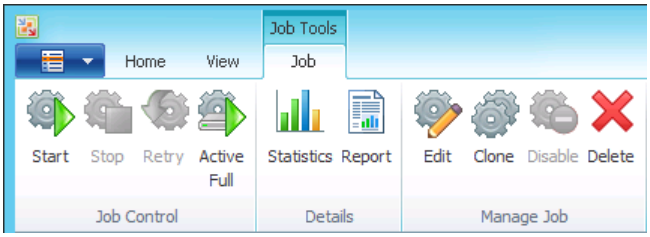

Note Commands for operations with items in Veeam Backup & Replication are also available from the shortcut menu.

Views

Veeam Backup & Replication displays its items in views. When you click the button of a specific view in the navigation pane, its content is displayed in the working area of Veeam Backup & Replication.

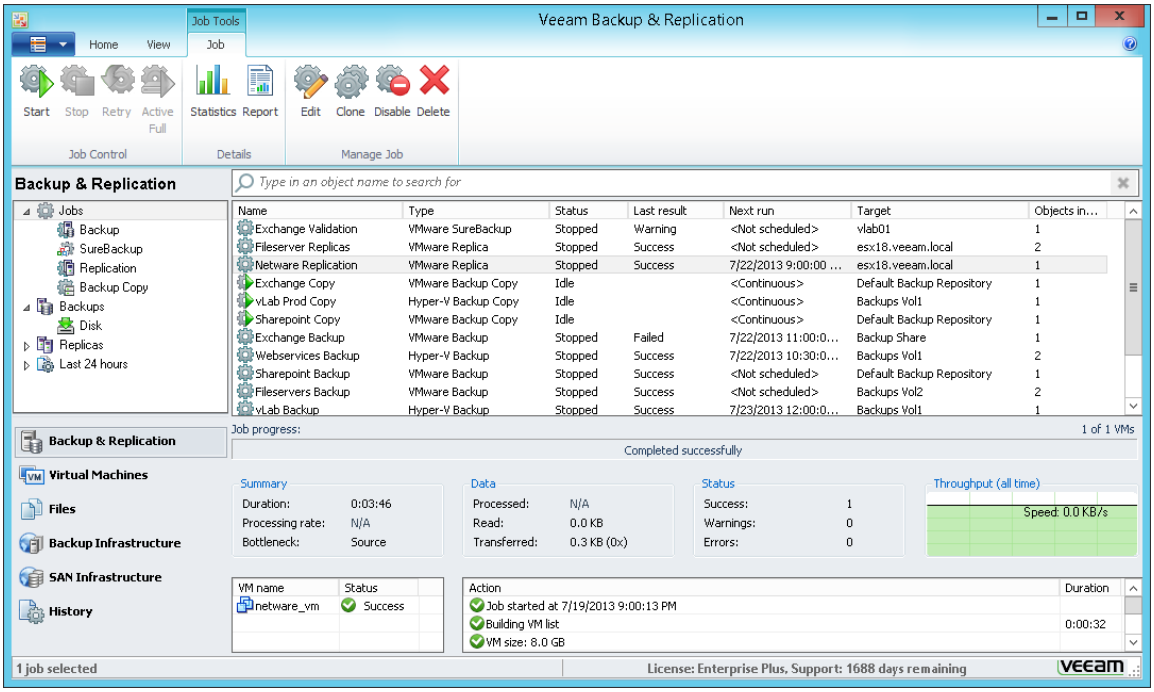

Veeam Backup & Replication offers the following views:

- The **Backup & Replication** view is used for work with all kind of jobs. It also displays a list of created backups and replicas that can be used for various restore operations, and provides statistics on recently performed jobs.
- The **Virtual Machines** view displays the inventory of your virtual infrastructure. The inventory can be presented from different perspectives — **Hosts**, **Categories** and **Storage**. You can use the **Virtual Machines** view to work with inventory objects and quickly add them to Veeam Backup & Replication jobs.
- The **Files** view displays a file tree for servers connected to Veeam Backup & Replication, and is primarily used for file copying operations.
- The **Backup Infrastructure** view displays a list of all backup infrastructure components — virtual infrastructure servers, backup proxies, backup repositories, tape devices and WAN accelerators. This view is used to set up the backup infrastructure that will be used for various data protection and disaster recovery tasks.
- The **SAN Infrastructure** view displays a list of HP storages, volumes and snapshots. This view is used to restore data from HP SAN snapshots (available only for VMware VMs).
- The **History** view displays statistics on operations performed with Veeam Backup & Replication. You can use this section for viewing statistics on performed tasks and reporting.

Working Area

The working area of Veeam Backup & Replication displays a list of items relevant to a specific view. The working area looks different depending on the view that is currently opened. For example, if you open the **History** view, the working area will display a list of job sessions and restore tasks performed with Veeam Backup & Replication. If you open the **Virtual Machines** view, the working area will display a list of virtual machines that reside on servers connected to Veeam Backup & Replication.

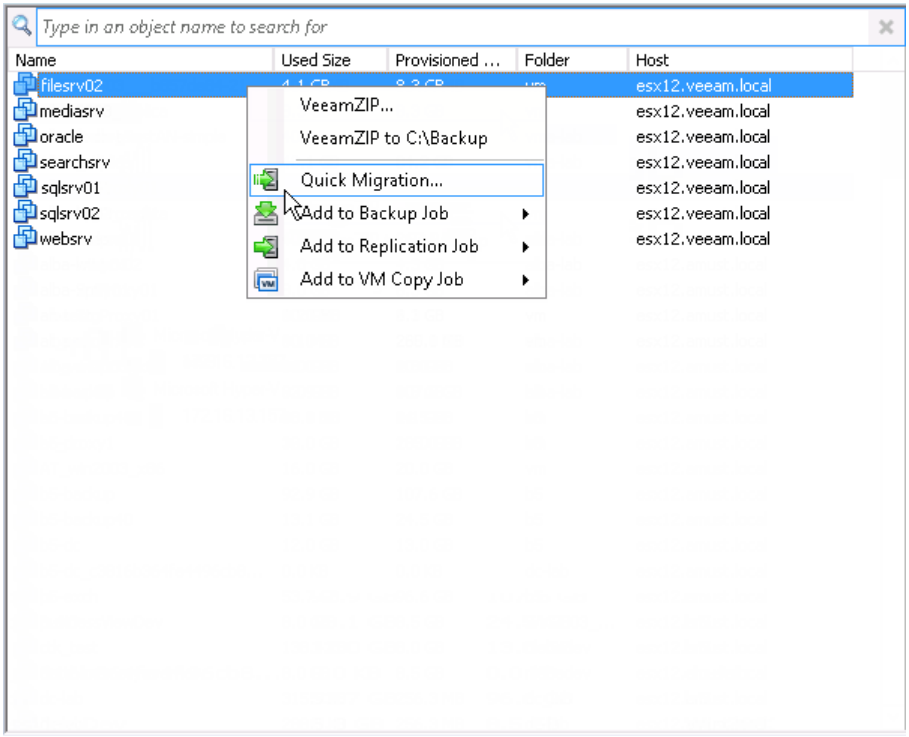

ADMINISTRATION

This section describes Veeam Backup & Replication administration tasks:

Setting Up Backup Infrastructure

Before creating backup, replication and other types of jobs, you need to plan and set up your backup infrastructure. Though the architecture may vary depending on the virtual environment and data protection requirements, a typical Veeam Backup deployment comprises the following components:

- Veeam backup server
- Backup proxies
- Backup repositories
- Virtual infrastructure servers – ESX(i) hosts used as source and target servers for backup, replication, VM copy and other types of jobs, as well as servers used for various types of restore operations.

To learn more about the purpose of each Veeam Backup & Replication infrastructure component, see Solution Architecture.

In general, the procedure of infrastructure setup includes the following steps:

1. **Adding servers.** First of all, you need to connect to the Veeam backup server all servers that you targeted as backup proxies, backup repositories, virtual infrastructure servers.
2. **Assigning proxy and repository roles.** After you have connected necessary servers, you need to assign the roles of backup proxies and backup repositories to the appropriate servers.

A newly deployed Veeam backup server performs the role of VMware backup proxy and backup repository in addition to its primary functions. That is why the Veeam backup server is also added to the list of managed servers, proxies and repositories. Essentially, this means that immediately after Veeam Backup & Replication installation, you can connect servers, configure and run the required jobs. The Veeam Backup & Replication server will be used as the backup server, proxy and repository at the same time.

Such scenario, however, is acceptable only if you plan to protect a small number of VMs or perform pilot testing. For a full-fledged backup infrastructure, you need to configure dedicated backup proxies and backup repositories. Components in such Veeam Backup deployment will be organized around the Veeam backup server which will function as the point of control for job processing. Data processing tasks will be offloaded to backup proxies and backup repositories.

Managing Credentials

To maintain a list of accounts authorized to perform the certain operation (for example, connection to server, guest OS access, and others), you can use the **Manage Credentials** command from **Main** menu.

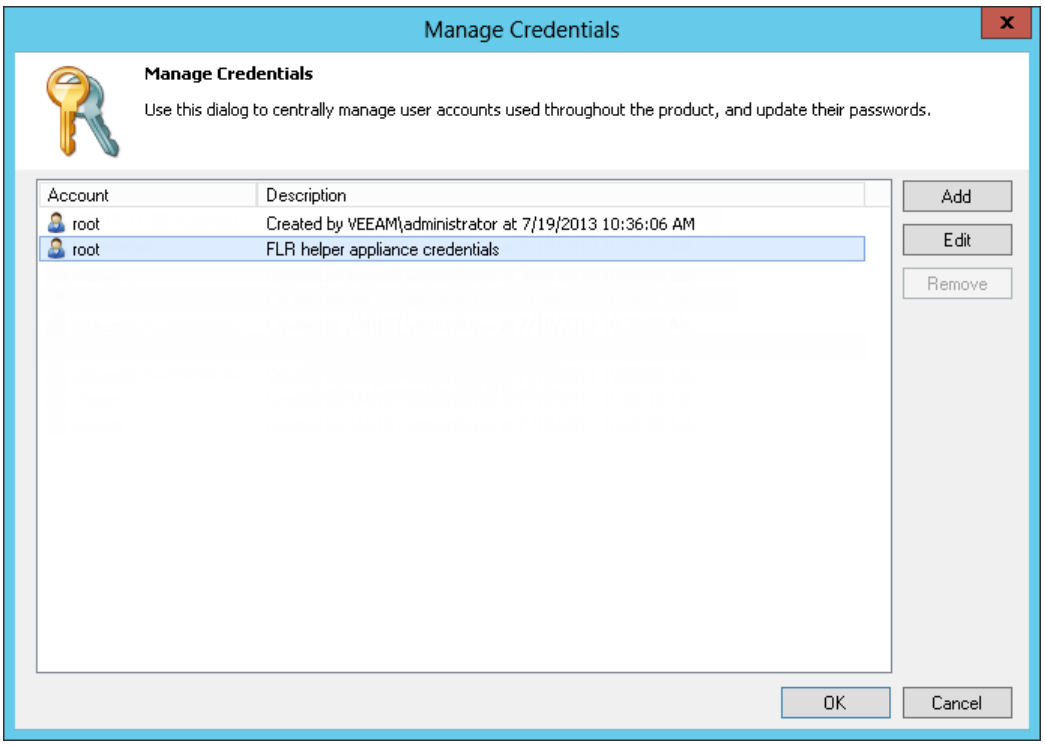

To enter a new user account:

1. Click **Add** and in the **Credentials** window.
2. Enter user name, password and description or browse for the user you need.

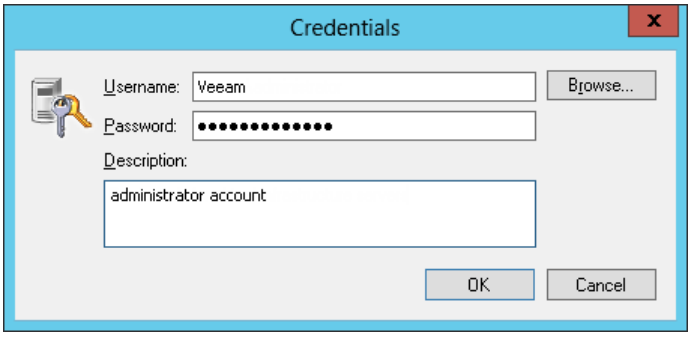

Tip As there can be a number of alike account names (for example, *Administrator*), it is highly recommended that you supply a meaningful unique description for the account name, so that you can distinguish these accounts when displayed in the list. The description will be shown in brackets, following the user name.

The **Manage Credentials** link is available on the **Credentials** step of the **Add Server** wizards. You can select the account from the list of available accounts, or use the **Add** button to add a different account (make sure it has sufficient access rights). Similarly, you can supply a user account when specifying guest OS access credentials at the **Guest processing** step of the **Backup Job** wizard.

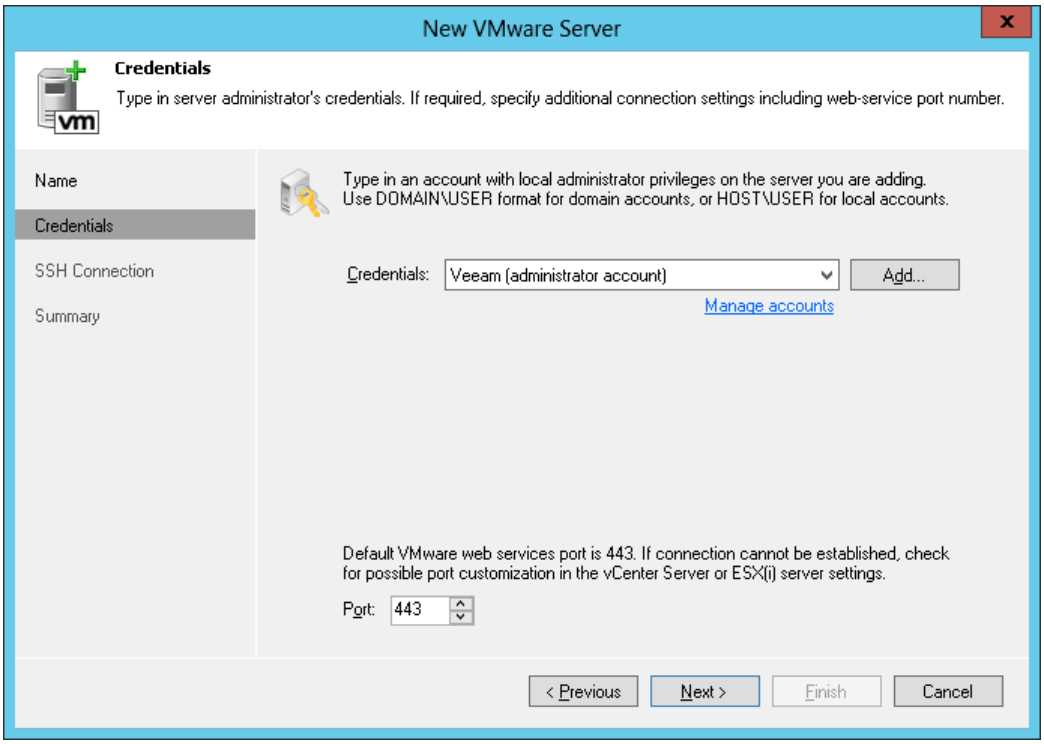

Adding Servers

For building your backup infrastructure in a VMware vSphere environment, Veeam Backup & Replication supports the following types of servers:

- VMware Server
- Windows Server
- Linux Server
- vCloud Director

Managed servers are physical or virtual machines used as source and target hosts, backup proxies, repositories and other servers included into the backup infrastructure. The table below shows which roles can be assigned to the different types of servers managed by Veeam Backup & Replication.

Server Type	Source	Replication Target	Backup Proxy	Backup Repository
VMware Server (standalone ESX(i) host or vCenter Server)	✓	✓	✗	✗
Windows Server	✗	✗	✓	✓
Linux Server	✗	✗	✗	✓
vCloud Director	✓	✗	✗	✗

Note	Any ESX host is essentially a Linux server. Thus, you can add ESX hosts both as virtualization servers and as standard file servers, depending on the role which you wish to assign for them. Note that if you plan to use the same host as an ESX host and Linux server, you should add it to Veeam Backup & Replication twice.

Adding a VMware Server

Veeam Backup & Replication allows you to connect both vCenter servers and standalone ESX(i) hosts. If possible, avoid adding ESX(i) hosts which are part of the vCenter Server hierarchy. Add the corresponding vCenter Server instead. Adding the vCenter Server facilitates management of the backup infrastructure and can be a recommended condition for certain types of operations (such as Quick Migration).

To add a VMware server, follow the next step:

Step 1. Launch the New VMware Server Wizard

To launch the **New VMware Server** wizard, do one of the following:

- Open the **Backup Infrastructure** view, select the **Managed servers** node in the inventory pane and click **Add Server** on the ribbon or right-click the **Managed servers** node and select **Add server**. In the **Add Server** window, select **VMware vSphere**.
- Open the **Virtual Machines** view, select the **VMware vSphere** node in the inventory pane and click **Add Server** on the ribbon or right-click the **VMware vSphere** node in the inventory tree and select **Add Server**.
- Open the **Virtual Machines** or **Files** view, right-click anywhere in the inventory pane and select **Add server**. In the **Add Server** window, select **VMware vSphere**.

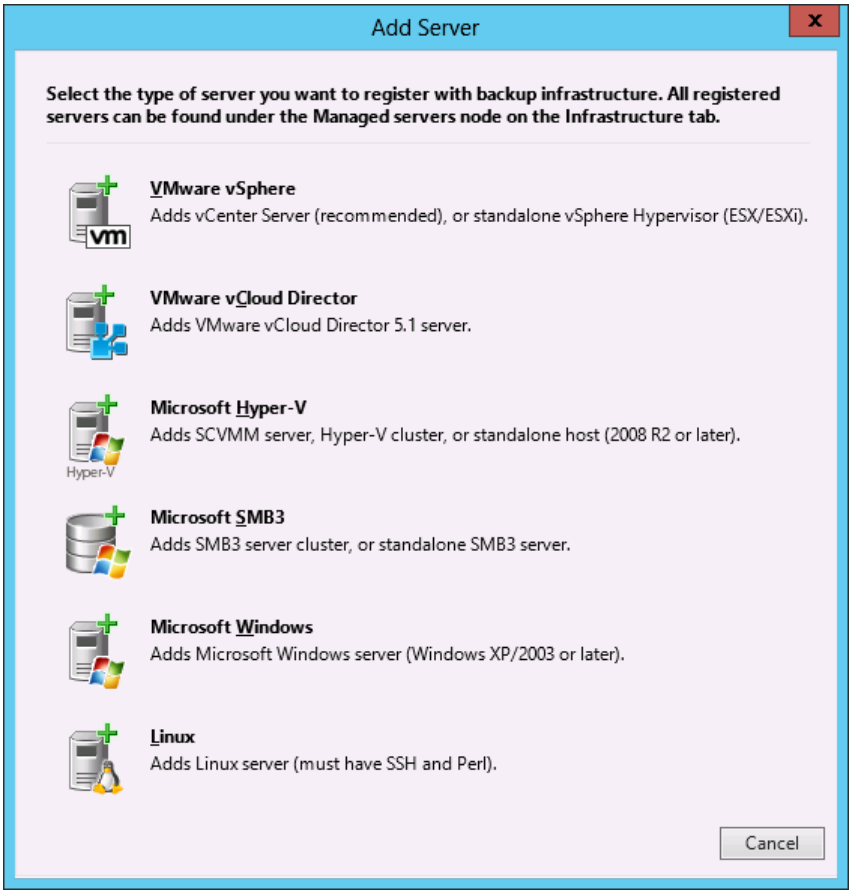

Step 2. Specify Server Name or Address

Enter a full DNS name or IP address of the vCenter Server or standalone ESX(i) host. Provide a description for future reference. The default description contains information about the user who added the server, as well as the date and time when the server was added.

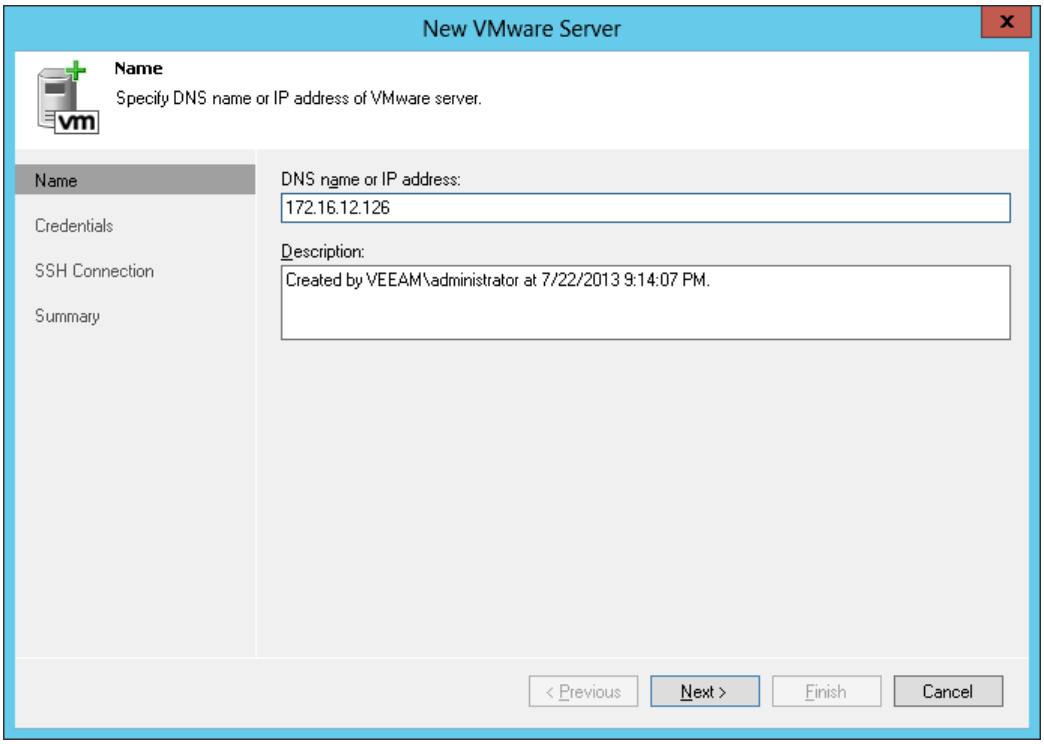

Step 3. Specify Credentials

At the **Credentials** step of the wizard, you should specify credentials and port settings for the added VMware Server.

1. From the **Credentials** list, select credentials for the account having administrator permissions on the added VMware Server. If you have not set up the necessary credentials beforehand, click the **Manage accounts** link at the bottom of the list or click **Add** on the right to add the necessary credentials. To learn more, see Managing Credentials.
2. The default port number for VMware web service communication is 443. If the connection with the vCenter Sever or ESX(i) host over this port cannot be established, you can customize the port number on the vCenter Server/ESX(i) host settings and specify the new port number in this field.

Note The user name of the account should be specified in the *DOMAIN\USERNAME* format.

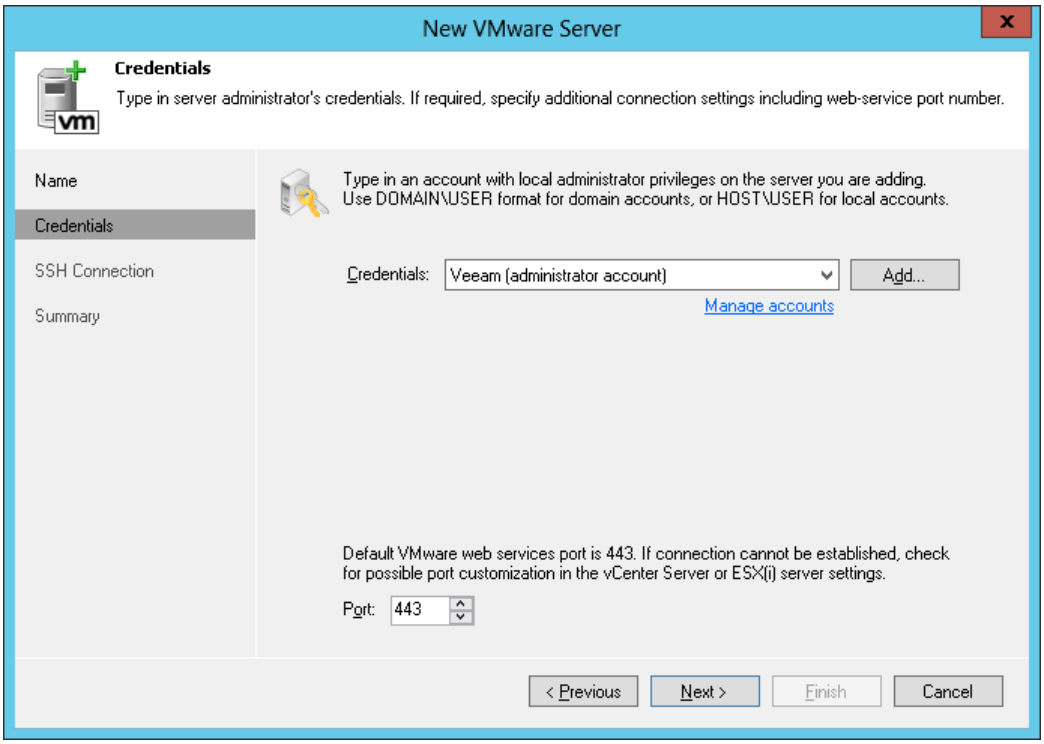

Step 4. Specify Service Console Connection Settings

This step is available only if you are adding an ESX host. When adding a vCenter Server or ESXi host, the wizard will skip this step and move on to the Summary step of the wizard.

If necessary, you can use a SSH connection for file copying operations and file copying jobs. These settings are optional. If you do not want to use SSH, clear the **Use service console connection to this server** check box. In this case, Veeam Backup & Replication will work with the server in the agentless mode. Agentless mode may be used for ESX 3.5 and later versions (ESX 3.0 and prior versions do not support agentless mode).

To use a SSH connection:

1. Select the **Use service console connection** to this server check box.
2. Enter the user name and password to connect to the service console of the server. If you want to use these credentials only during the current Veeam Backup & Replication session, clear the **Save password** check box. After you close and start the console next time, you will be asked to enter these credentials when a job addresses this ESX host.
3. If you choose to use non-root account that does not have root permissions on the ESX server, you can use the **Non-root account** section to grant sudo rights to this account. Select the **Elevate specified account to root** check box to provide a non-root user with access to the added server. You can add the account to sudoers file automatically by selecting the **Add account to the sudoers file automatically** check box. If you do not select this check box, you will have to manually add the user to the sudoers file.

Note Make sure that in the sudoers file the *NOPASSWD:ALL* option is enabled for the user account you want to elevate to root. Otherwise, jobs addressing the server will be failing, as sudo will request the password.

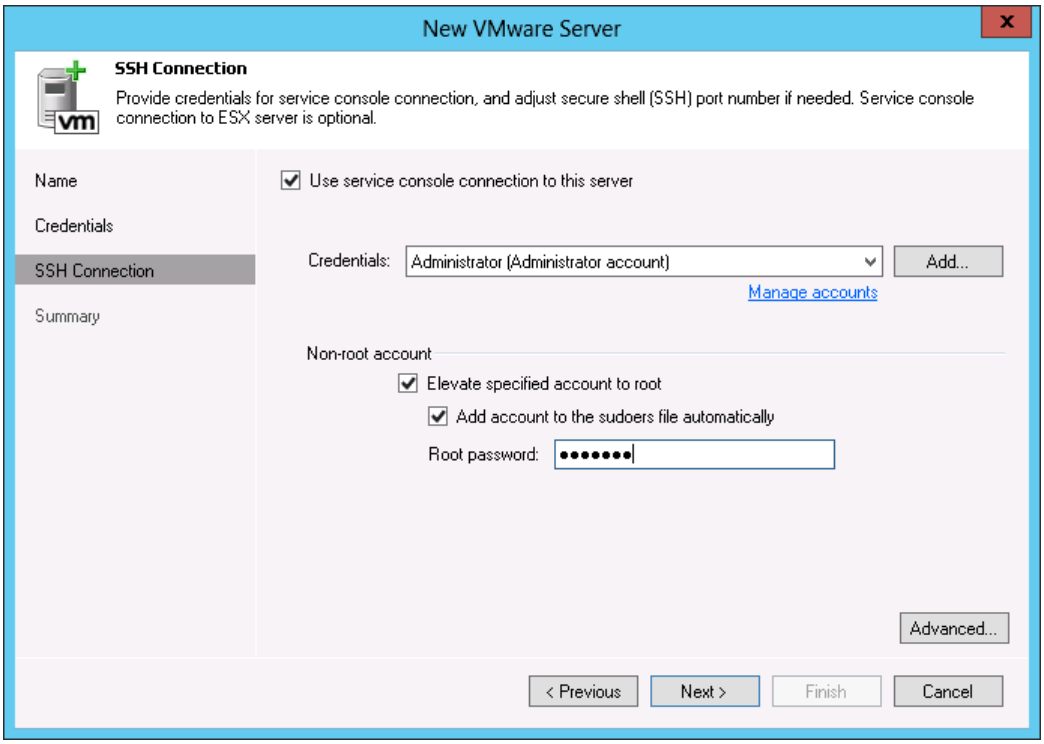

Click **Advanced** to configure advanced SSH settings.

1. In the **Service console connection** section, specify the SSH port to be used and SSH timeout. By default, SSH uses port number 22.
2. In the **Data transfer options** section of the SSH Settings window, specify a range of ports to be used as transmission channels between the source host and the target host (one port per task), and define the size of transmitted packets. By default, the port range is 2500 to 5000. However, depending on your environment, you can specify a smaller range of ports (for example, 2500 to 2510 will allow you to run 10 concurrent jobs at a time).
3. If the ESX host is deployed outside NAT, select the **Run the server on this side** check box in the **Preferred TCP connection role** section. In the NAT scenario, the outside client cannot initiate a connection with the server on the NAT network. Therefore, services that require the initiation of connection from outside can be disrupted. With this option selected, you will be able to overcome this limitation and initiate a 'server-client' connection (that is, a connection in the direction of the ESX server).

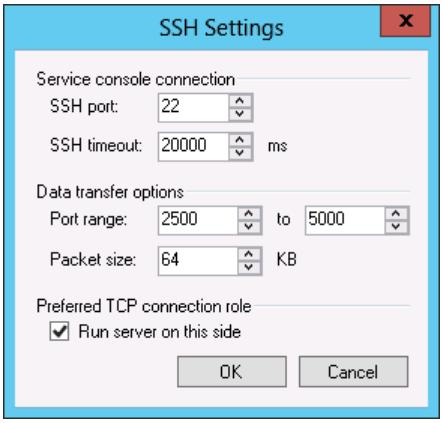

Tip	You can safely log on to a remote server directly from Veeam Backup & Replication using PuTTY — a popular SSH client. To open the **PuTTY Configuration** window, select **PuTTY** from the application tools menu. To learn about PuTTY, see http://www.chiark.greenend.org.uk/~sgtatham/putty/docs.html.

Step 5. Finish Working with the Wizard

Review the configuration information and click **Finish** to exit the wizard.

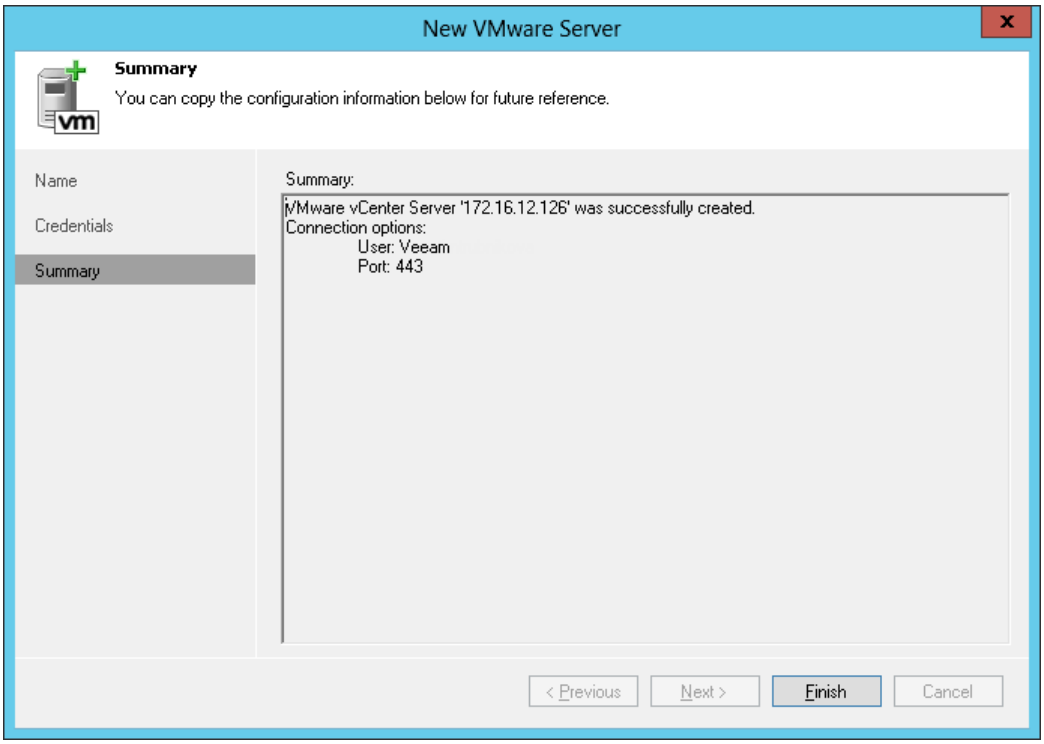

Adding a Windows Server

On every added Windows server, Veeam Backup & Replication deploys two components:

- Veeam Installer Service
- Veeam Transport

Make sure that *File and Printer Sharing* is enabled in the network connection settings of the server. Otherwise, the components will not be installed.

To add a Windows server, follow the next steps:

Step 1. Launch the New Windows Server Wizard

To launch the **New Windows Server** wizard, do one of the following:

- Open the **Backup Infrastructure** view, select the **Microsoft Windows** node in the inventory tree and click **Add Server** on the ribbon.
- Open the **Backup Infrastructure** or **Files** view, right-click the **Microsoft Windows** node in the inventory tree and select **Add Server**.
- Open the **Virtual Machines** or **Files** view, right-click anywhere in the inventory pane and select **Add server**. In the **Add Server** window, select **Microsoft Windows**.

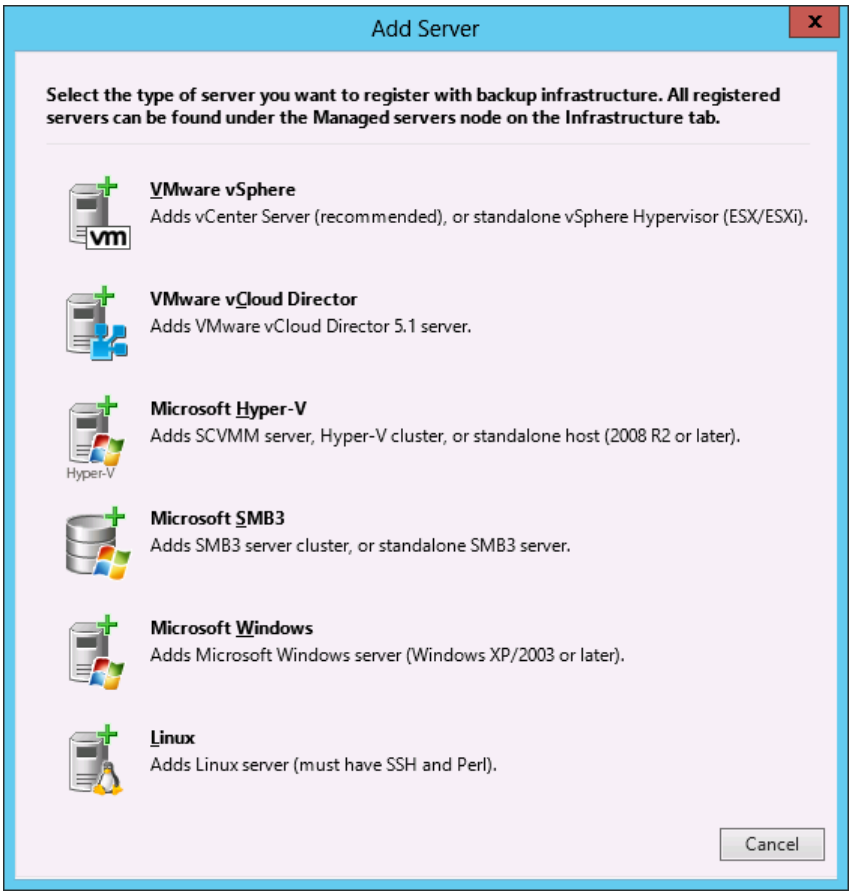

Step 2. Specify Server Name or Address

Enter a full DNS name or IP address of the Microsoft Windows server. Provide a description for future reference. The default description contains information about the user who added the server, as well as the date and time when the server was added.

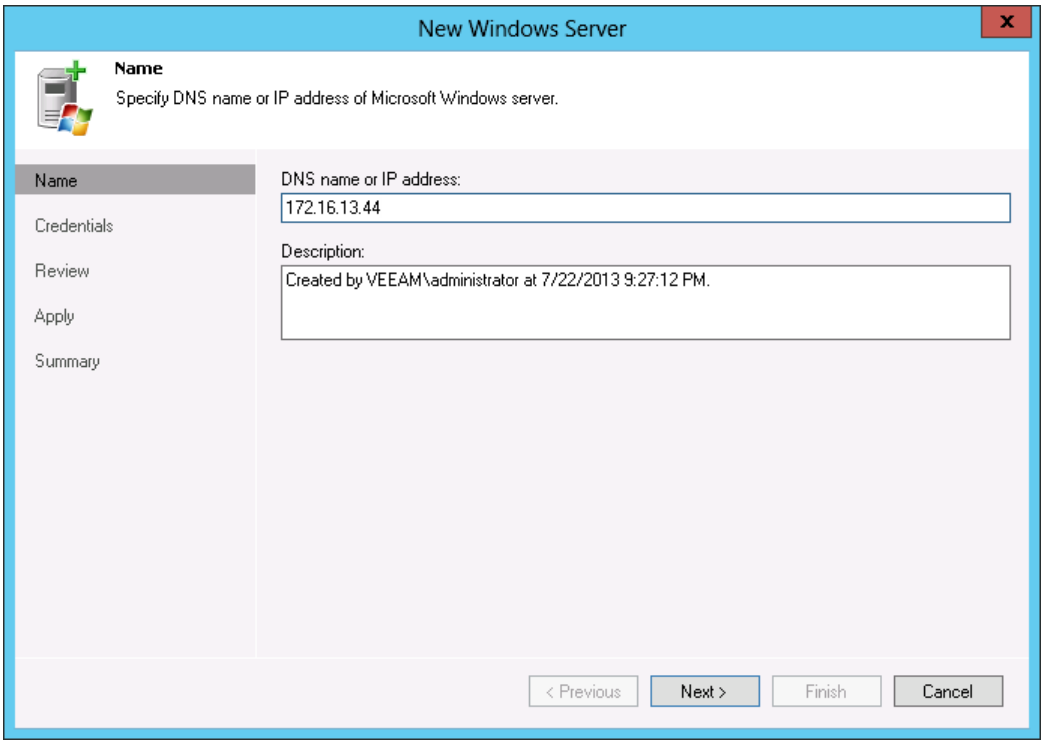

Step 3. Specify Credentials

At the **Credentials** step of the wizard, you should specify credentials for the Windows server you are adding.

From the **Credentials** list, select credentials for the account having local administrator permissions on the added Windows server. If you have not set up the necessary credentials beforehand, click the **Manage accounts** link at the bottom of the list or click **Add** on the right to add the necessary credentials. To learn more, see Managing Credentials.

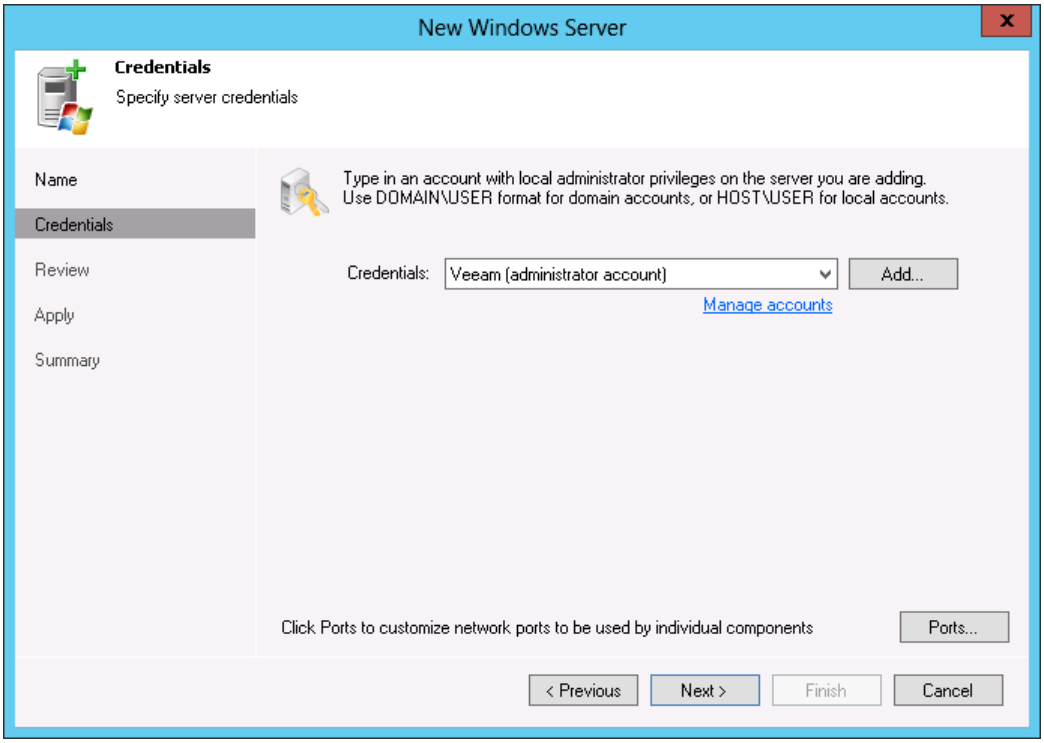

To customize network ports used by these components, click **Ports**.

- Veeam Installer Service is responsible for deploying the Veeam Transport on the Windows server. By default, the Veeam Installer Service uses port number 6160.
- Veeam Transport is responsible for deploying the corresponding during backup and replication processes. By default, the service uses port number 6162.

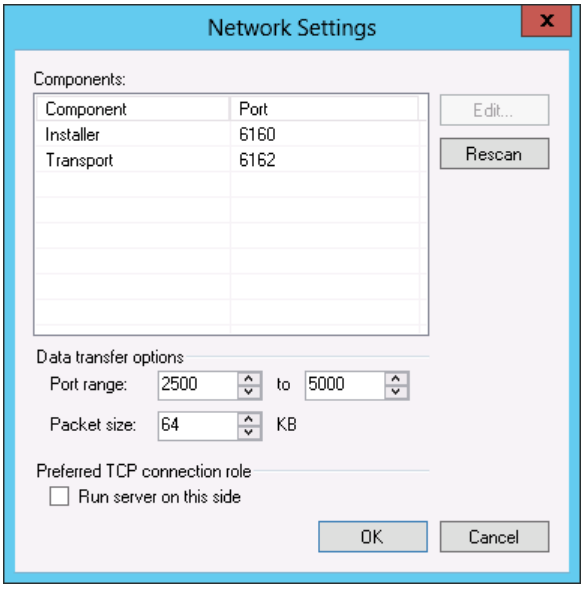

In the **Data transfer options** section of the **Network Settings** window, specify connection settings for FastSCP operations. Provide a range of ports to be used as transmission channels between the source host and the target host (one port per job), and define the size of transmitted packets. By default, the port range is 2500 to 5000. However, depending on your environment, you can specify a smaller range of ports (for example, 2500 to 2510 will allow you to run 10 concurrent jobs at a time).

If the Windows server is deployed outside NAT, select the **Run the server on this side** check box in the **Preferred TCP connection role** section. In the NAT scenario, the outside client cannot initiate a connection with the server on the NAT network. Therefore, services that require the initiation of connection from outside can be disrupted. With this option selected, you will be able to overcome this limitation and initiate a 'server-client' connection (that is, a connection in the direction of the Windows server).

Step 4. Review Components

At this step, you can review the list of components to be installed on the Windows server. If some of them are missing, Veeam Backup & Replication will automatically install them.

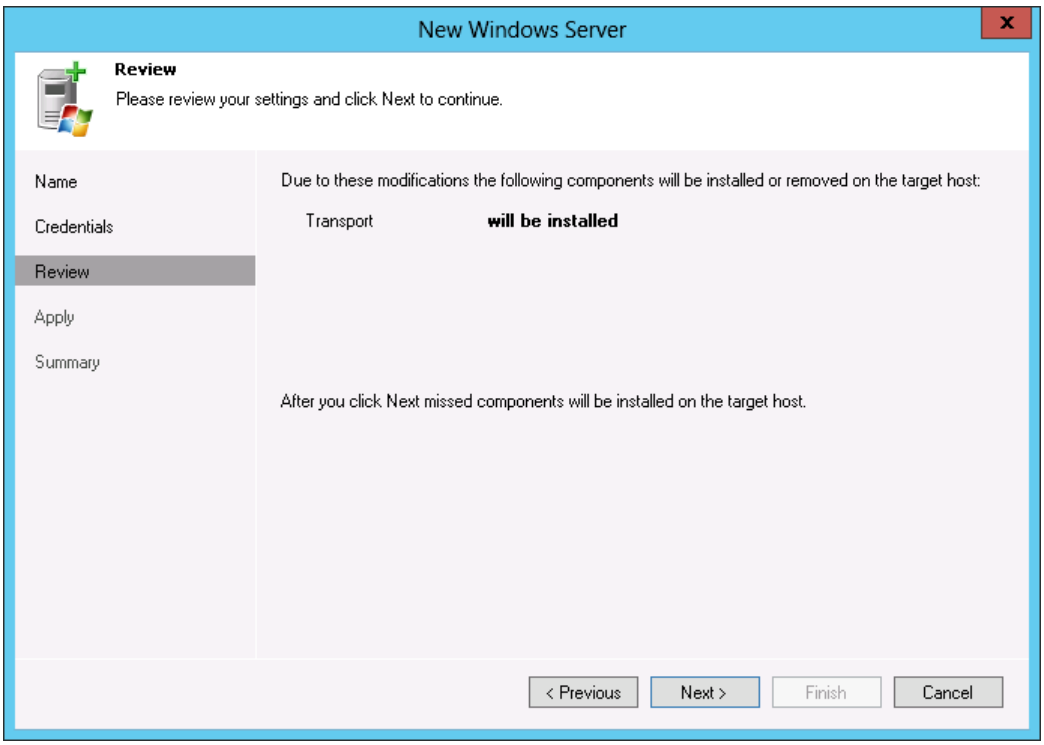

Step 5. Assess Results

At this step of the wizard, Veeam Backup & Replication will install the components on the added Windows server. You will see real time processing results for each server which is being added to the list of managed servers. Wait for the required operations to be performed. At the end, you can assess the results of processing and click **Next** to continue.

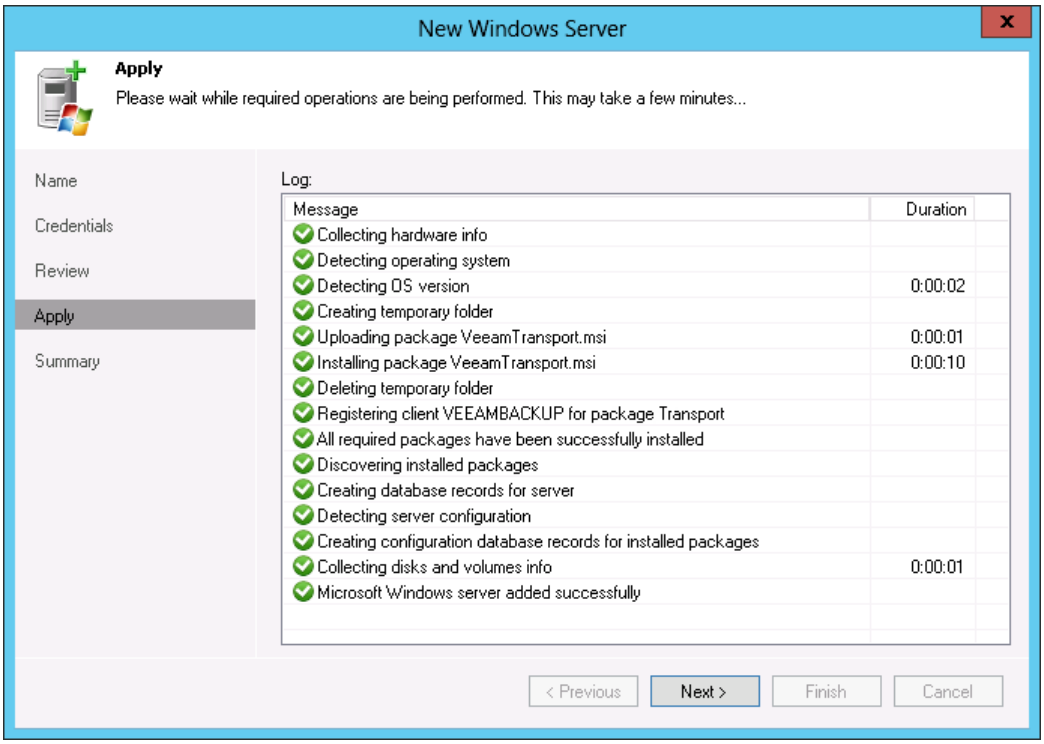

Step 6. Finish Working with the Wizard

Review the configuration information and click **Finish** to exit the wizard.

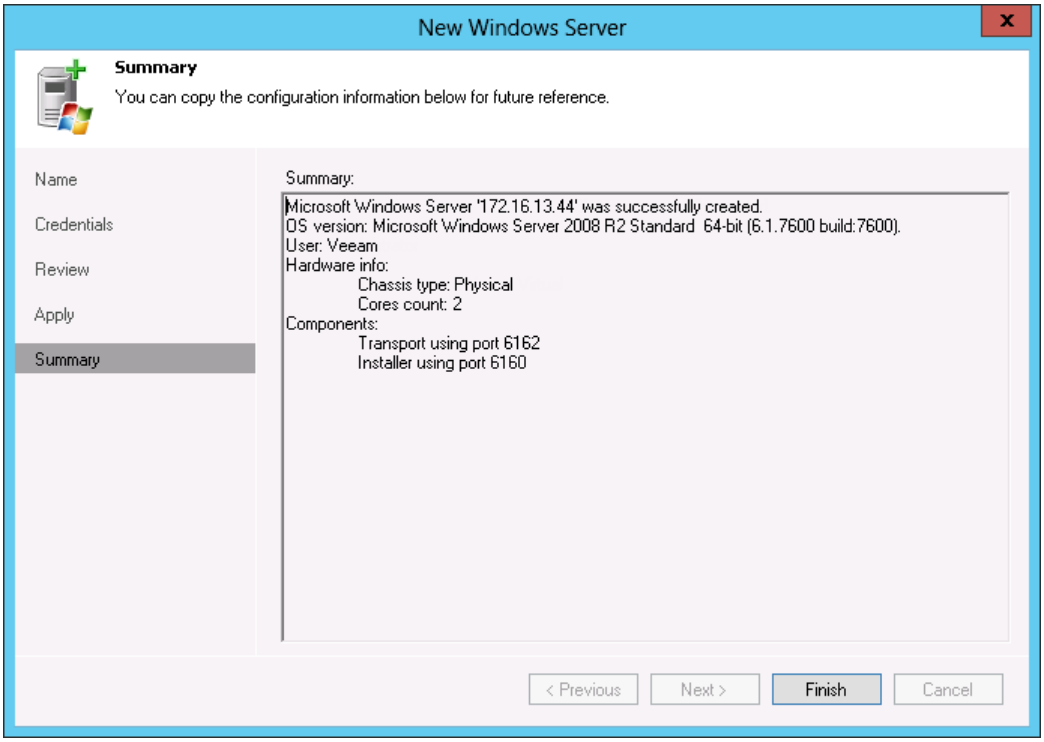

Adding a Linux Server

To add a Linux server, follow the next steps:

Step 1. Launch the New Linux Server Wizard

To launch the **New Linux Server** wizard, do one of the following:

- Open the **Backup Infrastructure** view, select the **Managed servers** node in the inventory tree and click **Add Server** on the ribbon or right-click the **Managed servers** node and select **Add server**. In the **Add Server** window, select **Linux**.
- Open the **Virtual Machines** or **Files** view, right-click anywhere in the inventory pane and select **Add server**. In the **Add Server** window, select **Linux**.

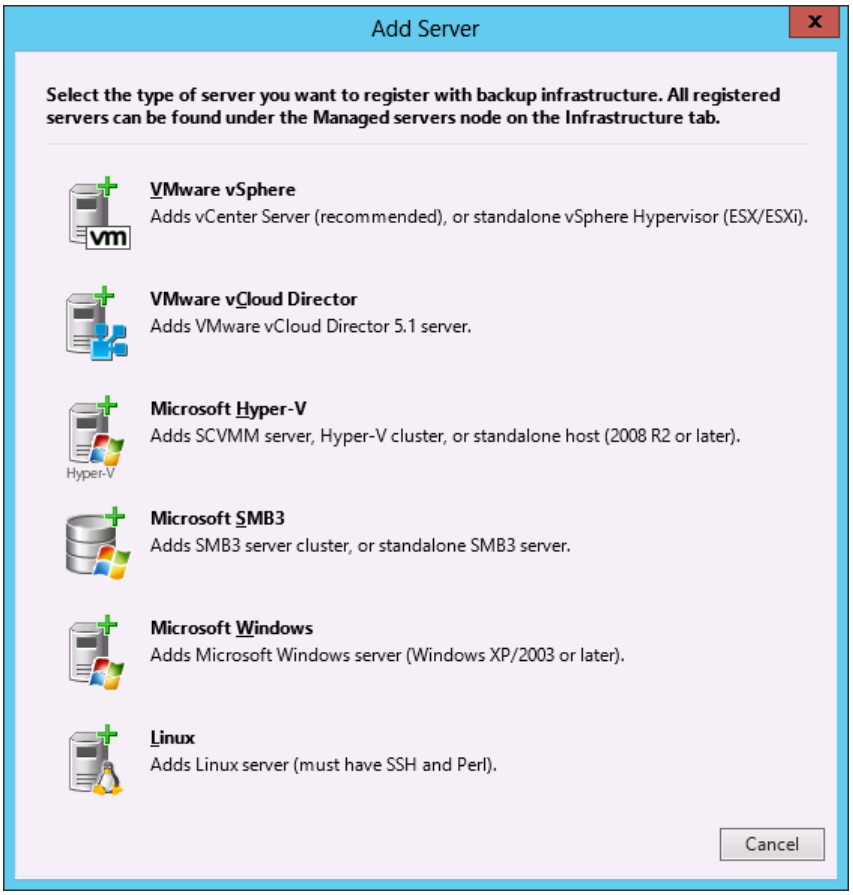

Step 2. Specify Server Name or Address

Enter a full DNS name or IP address of the Linux server. Provide a description for future reference. The default description contains information about the user who added the server, as well as the date and time when the server was added.

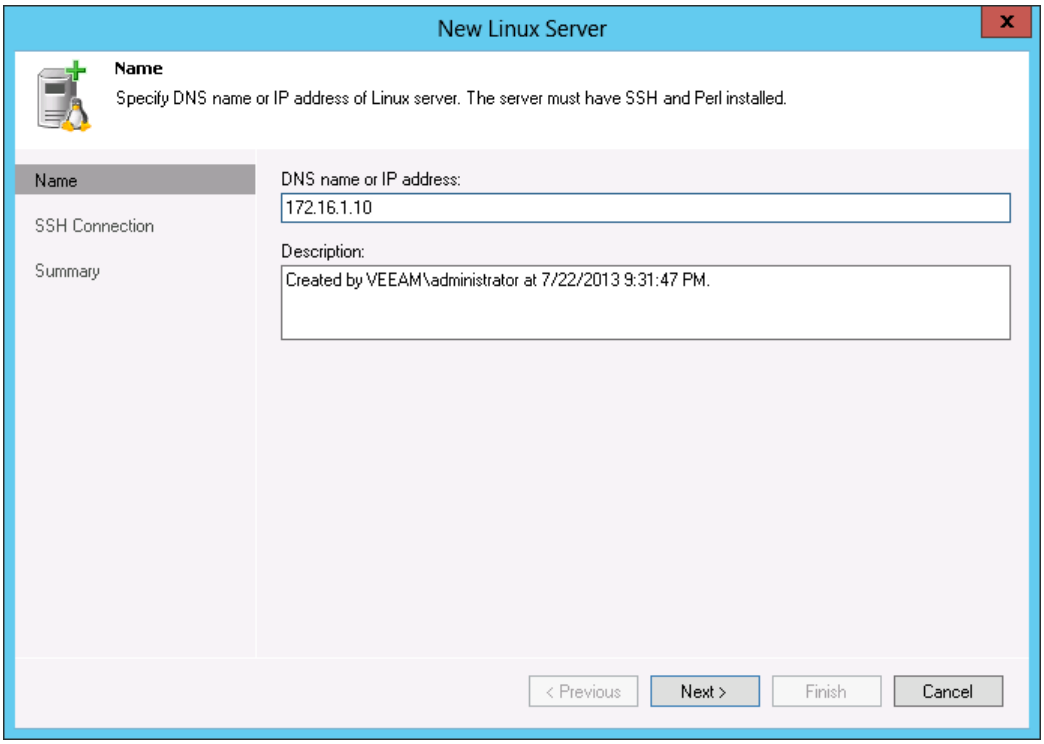

Step 3. Specify Credentials and SSH Port

At the **Credentials** step of the wizard, you should specify credentials for the Linux server you are adding.

1. From the **Credentials** list, select credentials for the account having local administrator permissions on the added Linux server. If you have not set up the necessary credentials beforehand, click the **Manage accounts** link at the bottom of the list or click **Add** on the right to add the necessary credentials. To learn more, see Managing Credentials.
2. Clear the **Save password** check box to use these credentials only during the current Veeam Backup & Replication session. After you close and start the console next time, you will be asked to enter these credentials when a job addresses this server.
3. If you choose to use non-root account that does not have root permissions on the Linux server, you can use the Non-root account section to grant sudo rights to this account. Select the **Elevate account to root** check box to provide a non-root user with access to the added server. You can add the account to sudoers file automatically by selecting the **Add account to the sudoers file automatically** check box. If you do not select this check box, you will have to manually add the user to the sudoers file.

Note Make sure that in the sudoers file the *NOPASSWD:ALL* option is enabled for the user account you want to elevate to root. Otherwise, jobs addressing the server will be failing as sudo will request the password.

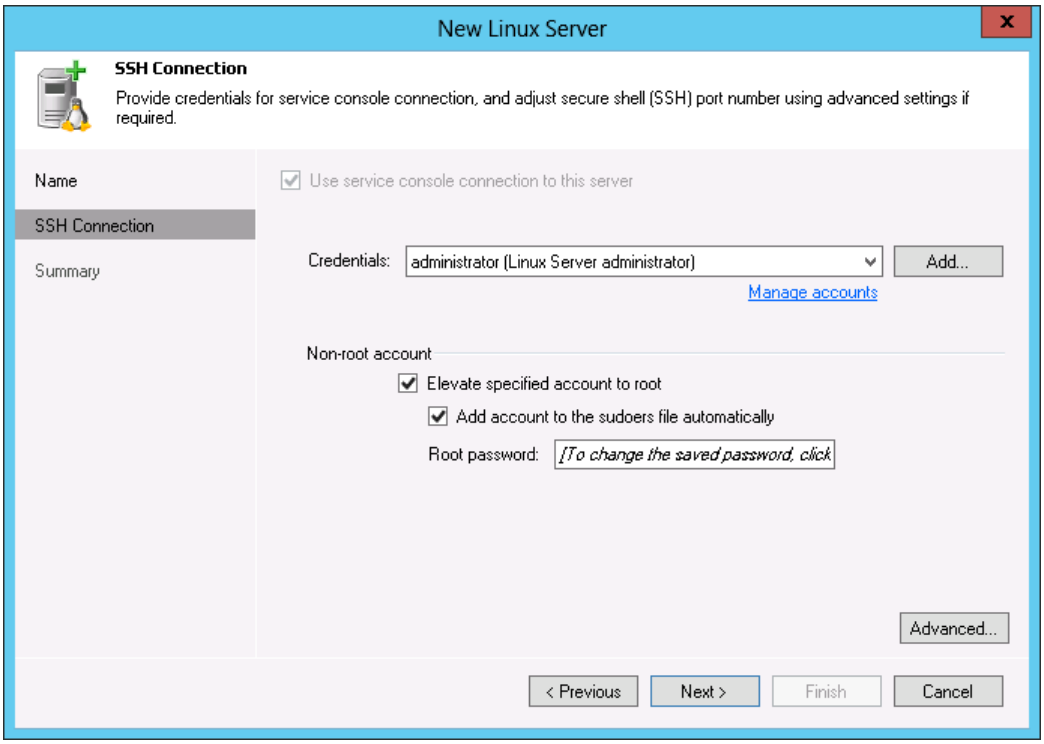

4. Click **Advanced** to configure advanced SSH settings.

    a. In the **console connection** section, specify the SSH port to be used and SSH timeout. By default, SSH uses port number 22.

    b. In the **Data transfer options** section, specify a range of ports to be used as transmission channels between the source host and the target host (one port per task) and define the size of transmitted packets. By default, the port range is 2500 to 5000. However, depending on your environment, you can specify a smaller range of ports (for example, 2500 to 2510 will allow you to run 10 concurrent jobs at a time).

    c. If the Linux server is deployed outside NAT, select the **Run the server on this side** check box. In the NAT scenario, the outside client (Linux server) cannot initiate a connection with the server (Veeam backup server) on the NAT network. For this reason, services that require the initiation of connection from outside can be disrupted. With this option selected, you will be able to overcome this limitation and initiate a 'server-client' connection, that is, a connection in the direction of the Linux server.

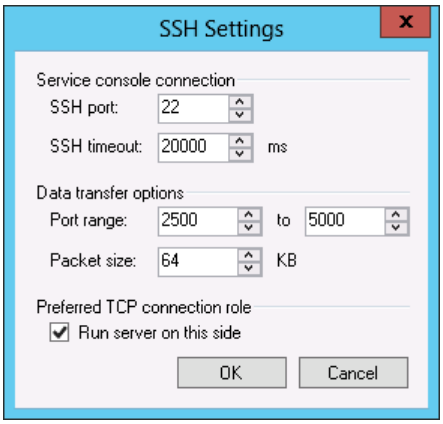

Step 4. Finish Working with the Wizard

Review the configuration information and click **Finish** to exit the wizard.

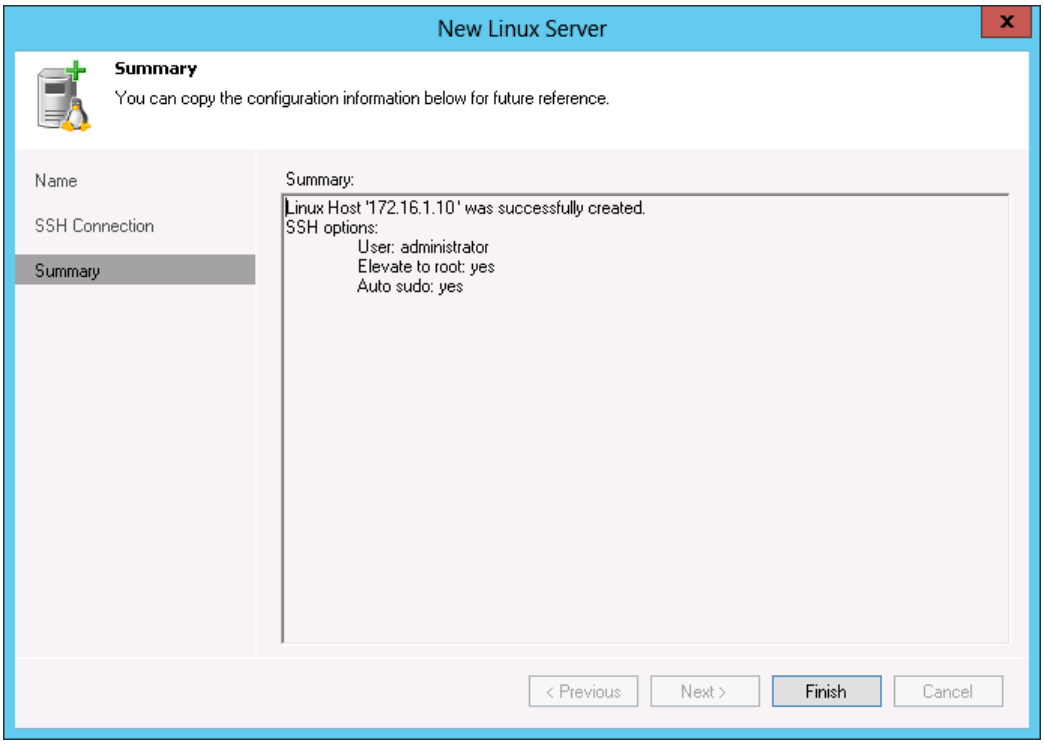

Adding vCloud Director

To work with VMs managed by vCloud Director, you need to add the vCloud Director server to Veeam Backup & Replication.

When you add vCloud Director, its hierarchy becomes available in the **Virtual Machines** > **vCloud Director** view. As a result, you can work with VMs managed by vCloud Director directly from the Veeam Backup & Replication console.

Note In case your vCloud Director infrastructure comprises several cells, you can specify connection settings for any cell in the vCloud Director hierarchy when adding vCloud Director to Veeam Backup & Replication.

To add the vCloud Director host, follow the next steps:

Step 1. Launch the New VMware vCloud Director Server Wizard

To launch the **New VMware vCloud Director** wizard, do either of the following:

- Open the **Backup Infrastructure** view, select the **Managed servers** node in the inventory tree and click **Add Server** on the ribbon or right-click the **Managed servers** node and select **Add server**.
- Open the **Virtual Machines** or **Files** view, right-click on the blank area in the inventory pane and select **Add server**.

In the **Add Server** window, select **VMware vCloud Director**.

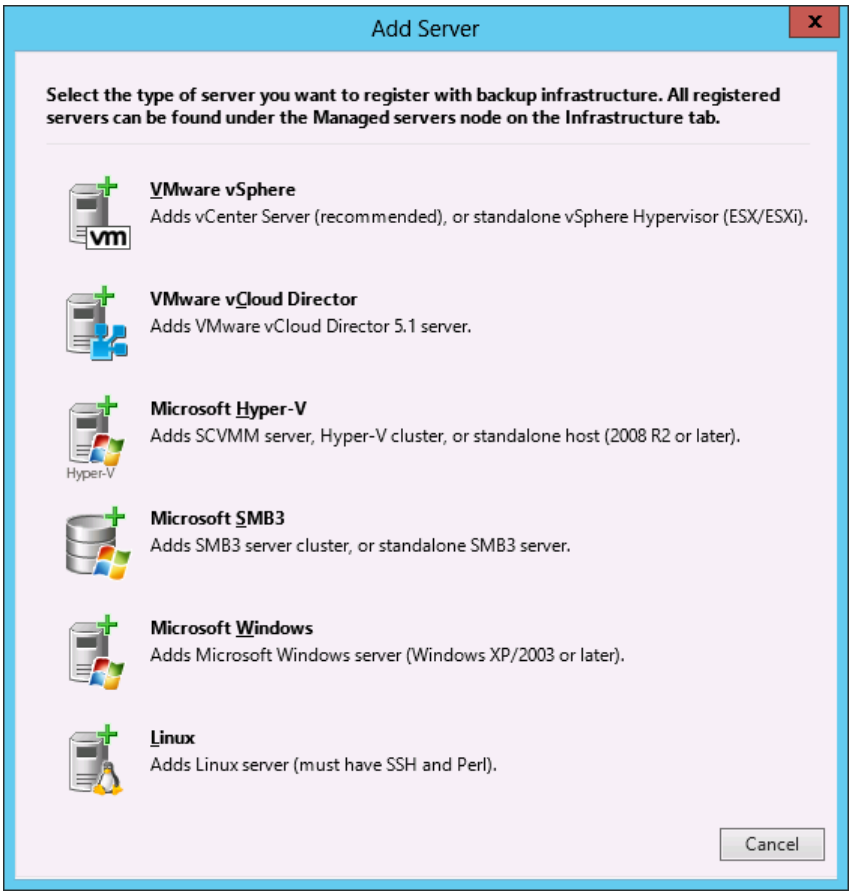

Step 2. Specify Server Name or Address

At the **Name** step of the wizard, specify connection settings for the added vCloud Director:

1. In the **DNS name or IP address** field, enter a full DNS name or IP address of the vCloud Director server or any cell in the vCloud Director infrastructure.
2. In the **URL** field, enter the URL of the vCloud Director server. By default, Veeam Backup & Replication uses the following URL: `https://<vcdservername>:443`, where `<vcdservername>` is the name or IP address of the vCloud Director server that you have specified in the field above and `443` is the default port for communication with vCloud Director.
3. In the **Description** field, provide a description for future reference. The default description contains information about the user who added the server, as well as the date and time when the server was added.

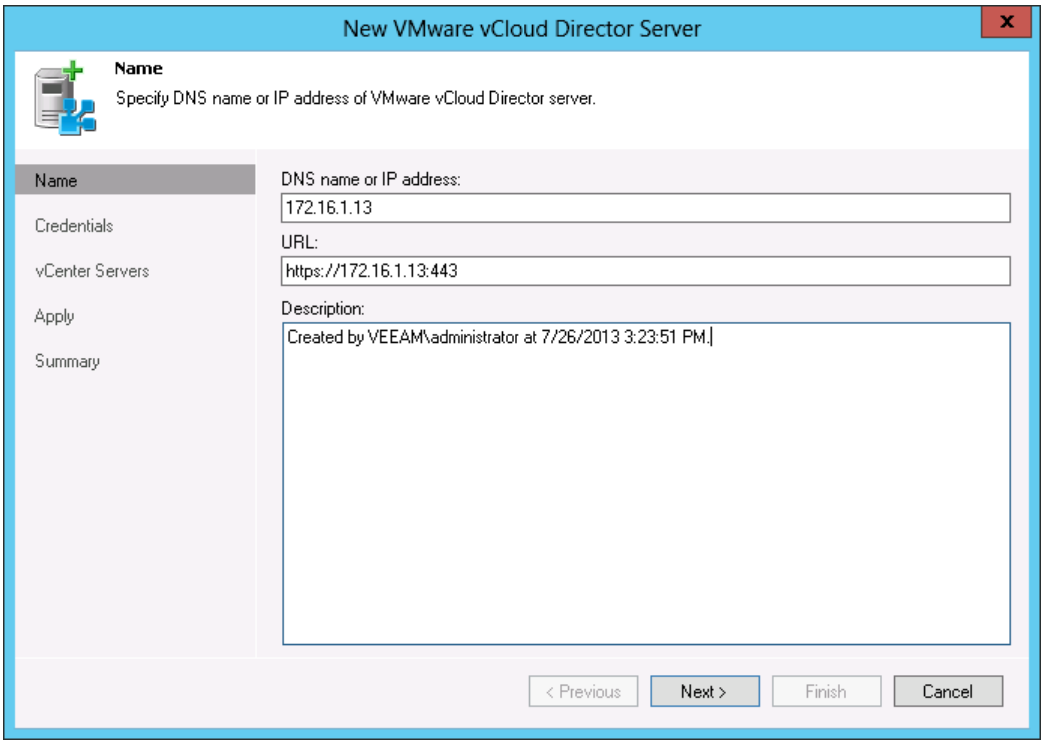

Step 3. Specify vCloud Director Credentials

At the **Credentials** step of the wizard, you should specify credentials to connect to the added vCloud Director.

From the **Credentials** list, select credentials for the account having administrative permissions on the added vCloud Director. If you have not set up the necessary credentials beforehand, click the **Manage accounts** link at the bottom of the list or click **Add** on the right to add the necessary credentials. To learn more, see Managing Credentials.

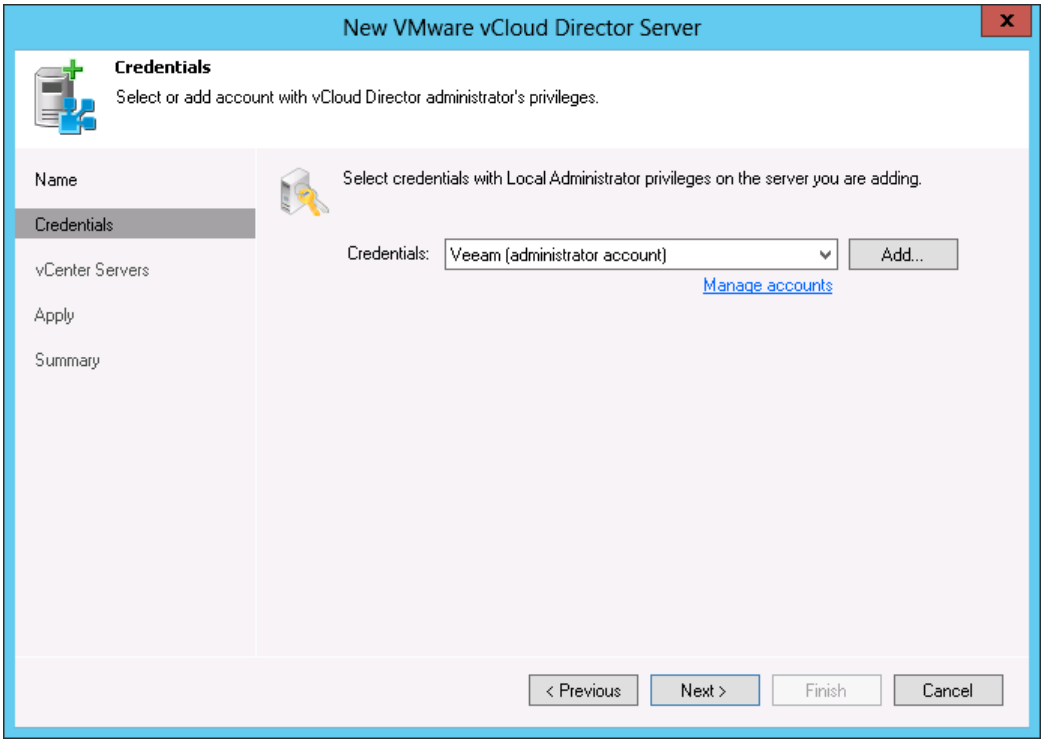

Step 4. Specify Credentials for Underlying vCenter Servers

At the **vCenter Servers** step of the wizard, you should specify credentials for every vCenter Server attached to the vCloud Director server:

1. From the **vCenter servers** list, select a vCenter Server.
2. Click **Account** on the right and select credentials to connect to the vCenter Server. By default, Veeam Backup & Replication uses the same credentials that you have specified for the vCloud Director at the previous step of the wizard.
   If you have not set up the necessary credentials beforehand, click the **Manage accounts** link at the bottom of the list or click **Add** on the right to add the necessary credentials. To learn more, see Managing Credentials.
3. By default, Veeam Backup & Replication automatically detects a port used to communicate with the vCenter Server. If necessary, you can change the connection port for the added vCenter Server. Click **vCenter** on the right and change the port number as required.
4. Repeat steps 1-3 for all vCenter Servers attached to vCloud Director.

Note If the vCenter Server attached to vCloud Director is already added to the Veeam Backup & Replication console, you do not need to enter credentials for it once again. Veeam Backup & Replication will automatically detect the credentials you provided when adding this vCenter Server and use them.

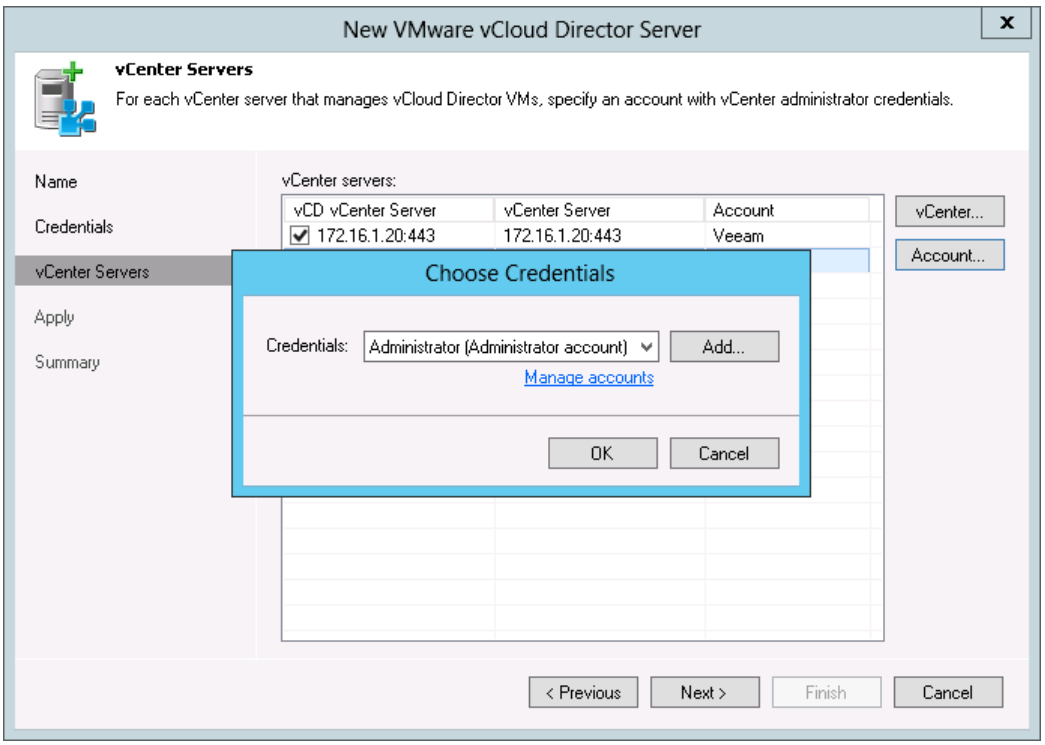

Step 5. Assess Results

At the **Apply** step of the wizard, Veeam Backup & Replication will add vCenter Servers attached to vCloud Director in the real time mode. Wait for the required operations to complete and assess results of the server processing.

Note If the vCenter Server attached to vCloud Director is already added to the Veeam Backup & Replication console, it will not be added for the second time: Veeam Backup & Replication will simply create an association with the added vCenter Server and display it in the vCloud Director hierarchy.

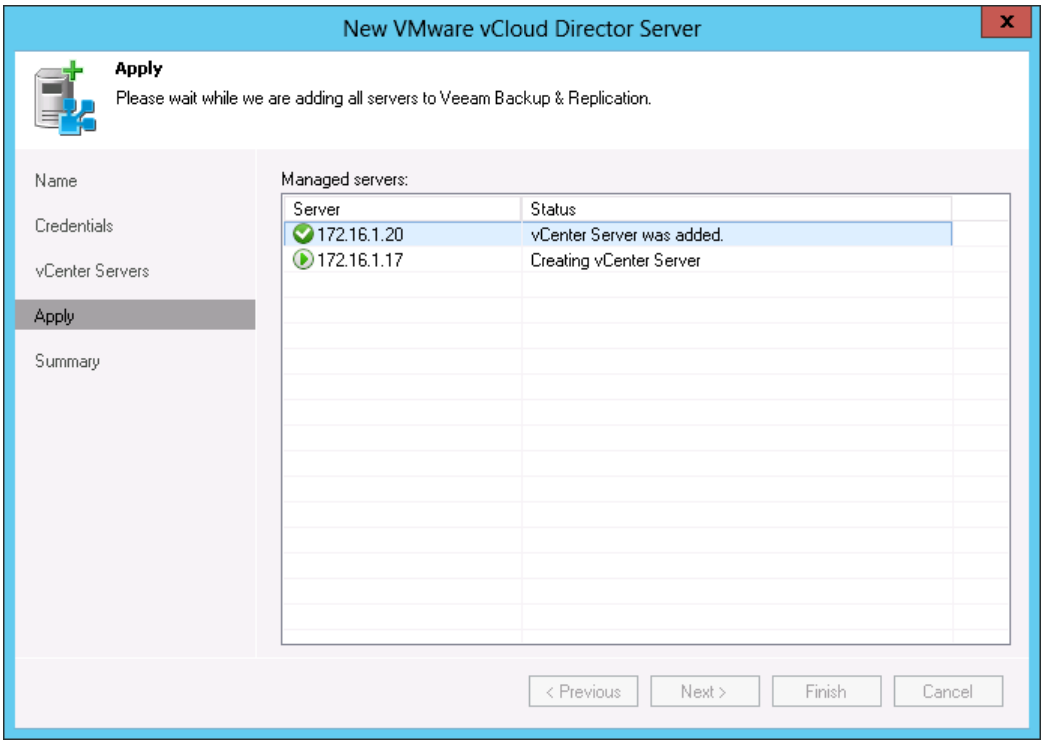

Step 6. Finish Working with the Wizard

At the **Finish** step of the wizard, review the configuration information and click **Finish** to exit the wizard.

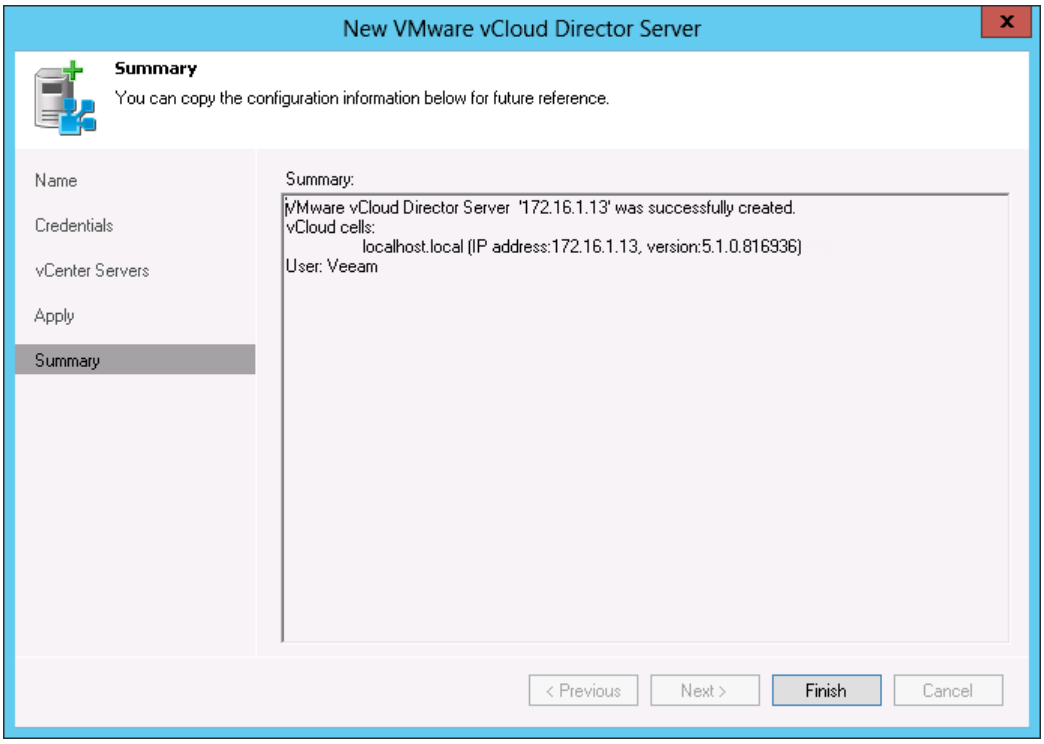

Managing Servers

You can edit settings of added servers, update components installed on servers, and remove servers from Veeam Backup & Replication.

Updating Server Components

Every time you launch Veeam Backup & Replication, it automatically checks if the components installed on managed servers are up to date. If there is a later version of a component available (usually, if you have upgraded Veeam Backup & Replication), the **Components Update** window will be displayed, prompting you to update components on managed servers.

You can also open the **Components Update** window if you select **Upgrade** from the main menu. If components on all managed servers are up to date, the menu item will be disabled.

The **Components Update** section lists servers that have outdated components deployed. To see the current and the latest available versions for deployed components, select a server in the list and click **Details**. Select check boxes next to servers for which you want to upgrade components and click **Next**.

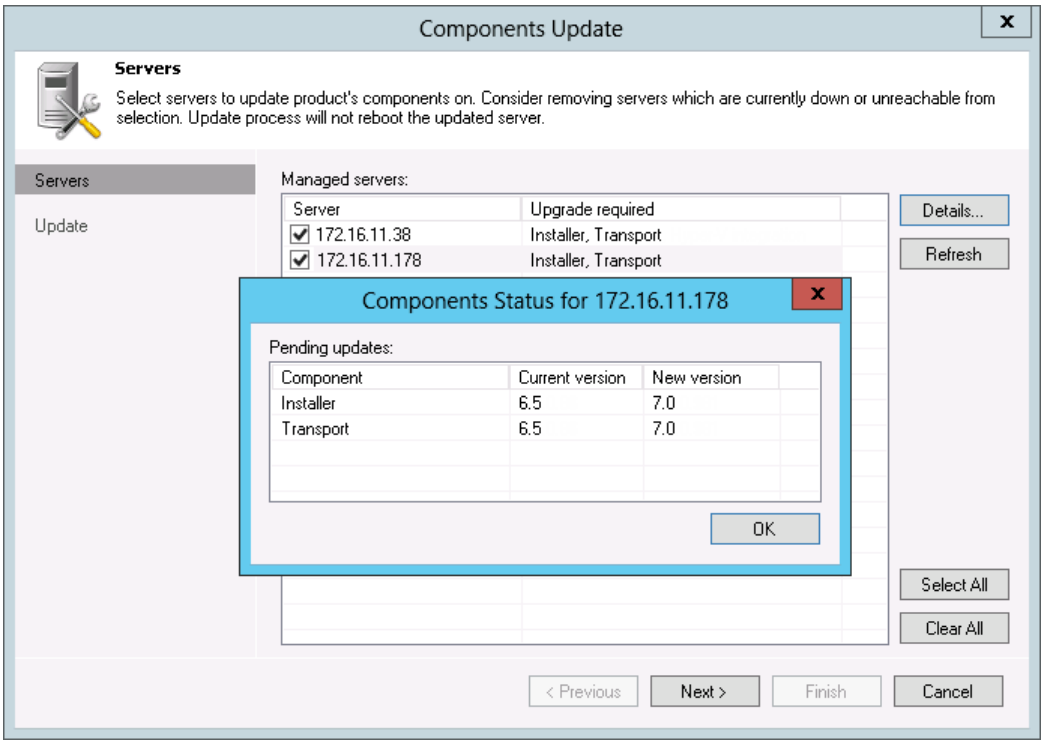

You can also update components on every managed server separately. Veeam Backup & Replication displays a warning next to Hyper-V and Windows server icons in the management tree to alert you when components on a server require updates. To update components, open the **Infrastructure** view, select the **Managed Servers** node in the inventory pane, select the necessary server in the working area and click **Upgrade** on the ribbon. Alternatively, you can open the **Infrastructure** view, select the **Managed Servers** node in the inventory pane, right-click the necessary server in the working area and select **Upgrade**.

Editing Server Settings

To edit settings of an added server:

1. Open the **Backup Infrastructure** view.
2. Click the **Managed Servers** node in the inventory tree.
3. Select the necessary server in the working area and click **Edit Server** on the ribbon or right-click the necessary server in the working area and select **Properties**.

You will follow the same steps as you have followed when adding the server. For details, see the description of the corresponding wizard for adding a new server under Adding Servers.

Removing Servers

To remove a server from the backup infrastructure:

1. Open the **Backup Infrastructure** view.
1. Click the **Managed Servers** node in the inventory tree.
2. Select the necessary server in the working area and click **Remove Server** on the ribbon or right-click the necessary server in the working area and select **Remove**.

A server that has any dependencies cannot be deleted. For example, you cannot delete a server that is referenced by a backup or replication job, or that is configured as a backup proxy or repository. To remove such server, you will need to delete all referencing jobs or objects first.

When you remove a server that was used as a target host or as a repository, actual backup files (.vbk, .vrb and .vib) and replica files (.vmdk and .vrb) are left on the server. You can easily import these files later to the Veeam Backup & Replication console for restore operations if needed.

Note When you remove the vCloud Director server from the Veeam Backup & Replication console, vCenter Servers attached to it are not removed. To remove such servers, expand the **vCenter Servers** node in the inventory pane, right-click the necessary server and select **Remove**.
You cannot remove vCenter Servers attached to the vCloud Director until the vCloud Director server is removed from the Veeam Backup & Replication console.

Adding a VMware Backup Proxy

In the backup infrastructure, a backup proxy acts as a "data mover". While the backup server fills the role of the job manager, the backup proxy actually performs main data handling – it retrieves data, processes it and transfers to the target destination. Use of backup proxies enables you to take the job processing off the backup server and allows for better scalability of your backup infrastructure.

Proxy deployment recommendations depend on a particular configuration of your virtual infrastructure. To learn about different deployment scenarios, see the Backup Architecture section.

To add a backup proxy to your backup infrastructure, you should assign this role to a Windows server that is already added to the list of managed servers.

By default, Veeam Backup & Replication adds the Veeam backup server to the list of backup proxies. Resources of such local backup proxy may be sufficient in simplest backup or replication scenarios. However, for larger VMware environments, you will need to deploy a number of additional VMware proxy servers to offload the Veeam backup server.

To add a VMware backup proxy, follow the next steps.

Step 1. Launch the New VMware Proxy Wizard

To launch the wizard, do one of the following:

- Open the **Backup Infrastructure** view, select the **Backup Proxies** node in the inventory pane, click **Add Proxy** on the ribbon and select **VMware**.
- Open the **Backup Infrastructure** view, right-click the **Backup Proxies** node in the inventory pane and select **Add VMware Backup Proxy**.

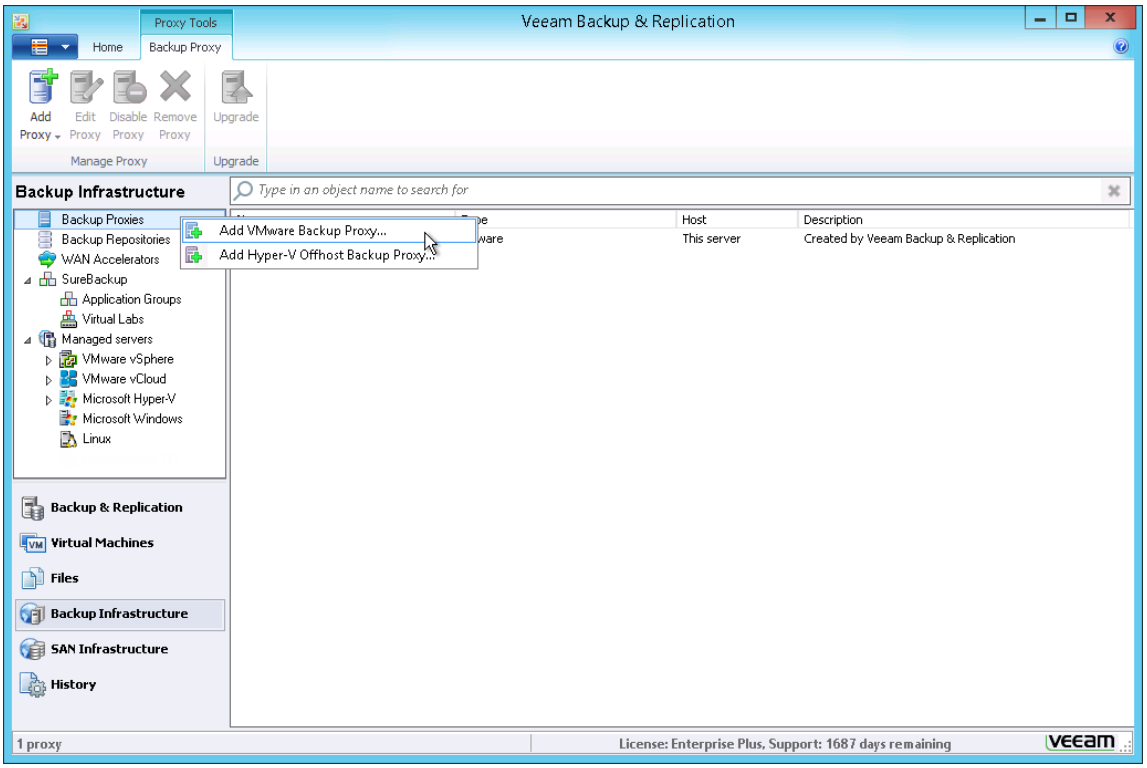

Step 2. Choose a Server

From the **Choose server** list, select a Windows server that has been previously added to Veeam Backup & Replication. If the server is not added yet, you can click **Add New** to open the **New Windows Server** wizard. In the **Proxy description** field, provide a description for future reference.

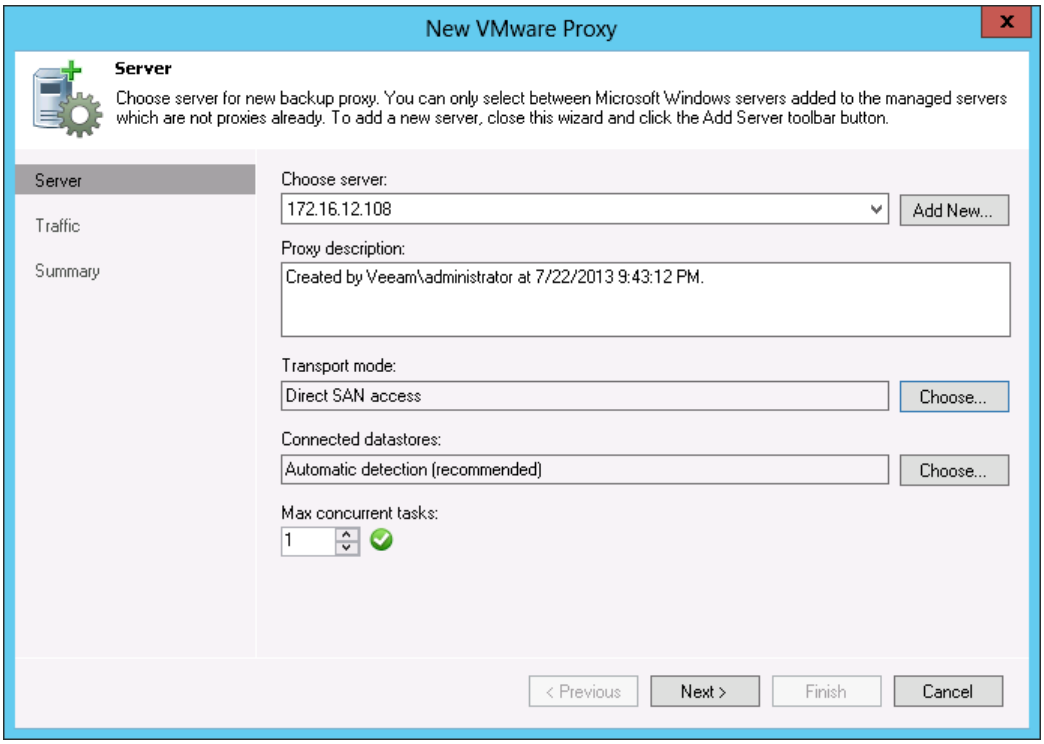

In the **Transport mode** field, specify the backup mode that the proxy will use to retrieve VM data. By default, Veeam Backup & Replication analyzes the backup proxy configuration, defines to which datastores it has access and automatically selects the best transport mode depending on the type of connection between the backup proxy and the source storage. However, you can manually select what mode you want to use for VM data retrieval. Click the **Choose** button on the right and select one of the following modes: **Direct SAN access**, **Virtual Appliance** or **Network**. For more information, see Transport Modes.

In the **Advanced** section, specify additional options for the selected mode:

- By default, if the **Direct SAN access** or **Virtual Appliance** mode is selected, Veeam Backup & Replication will automatically fail over to **Network** mode in case the primary selected backup mode fails during the job run. To disable failover, clear the **Failover to network mode if primary transport modes fail or are unavailable** check box.
- If the **Network** mode is selected, you can choose to transfer disks data over encrypted SSL connection. For this, select the **Encrypt LAN traffic in the network mode (SSL)** check box. Traffic encryption puts more stress on the CPU of an ESX server, providing, however, secure data transfer.

In the **Connected datastores** field, specify datastores to which this backup proxy has direct SAN connection. By default, Veeam Backup & Replication automatically detects all datastores that the backup proxy can access via the Direct SAN Access mode.

If Veeam Backup & Replication cannot detect accessible datastores for some reason, you can also define the list of datastores yourself. To assign datastores to the backup proxy manually, click the **Choose** button on the right, switch to the **Manual selection** mode and add datastores from which VM data should be retrieved in the Direct SAN Access mode.

In the **Max concurrent tasks** field, you can specify the number of tasks the backup proxy should handle in parallel. Recommended number of concurrent tasks is calculated automatically in accordance with the available resources. When configuring this parameter manually, consider that each data processing task requires one CPU core. For example, a 2-core CPU (minimum recommended for a proxy) can handle two concurrent tasks. If the specified number of tasks is exceeded, the backup proxy will not start a new task until one of the current tasks is finished. Also, when entering the number of concurrent tasks, you should keep in mind the network traffic throughput in your virtual infrastructure.

To learn more about limiting the number of concurrent tasks, see Limiting the Number of Concurrent Tasks.

Step 3. Configure Traffic Throttling Rules

At this step of the wizard, you can configure throttling rules to limit the outbound traffic rate for the backup proxy. Throttling rules will help you manage bandwidth usage and minimize the impact of backup jobs on network performance. For details, see Setting Network Traffic Throttling Rules.

The list of throttling rules contains only those rules that are applicable to the backup proxy you are adding. The rule is applied to the backup proxy if its IP address falls under the source IP range of the rule. To view the rule settings, select it in the list and click the **View** button on the right.

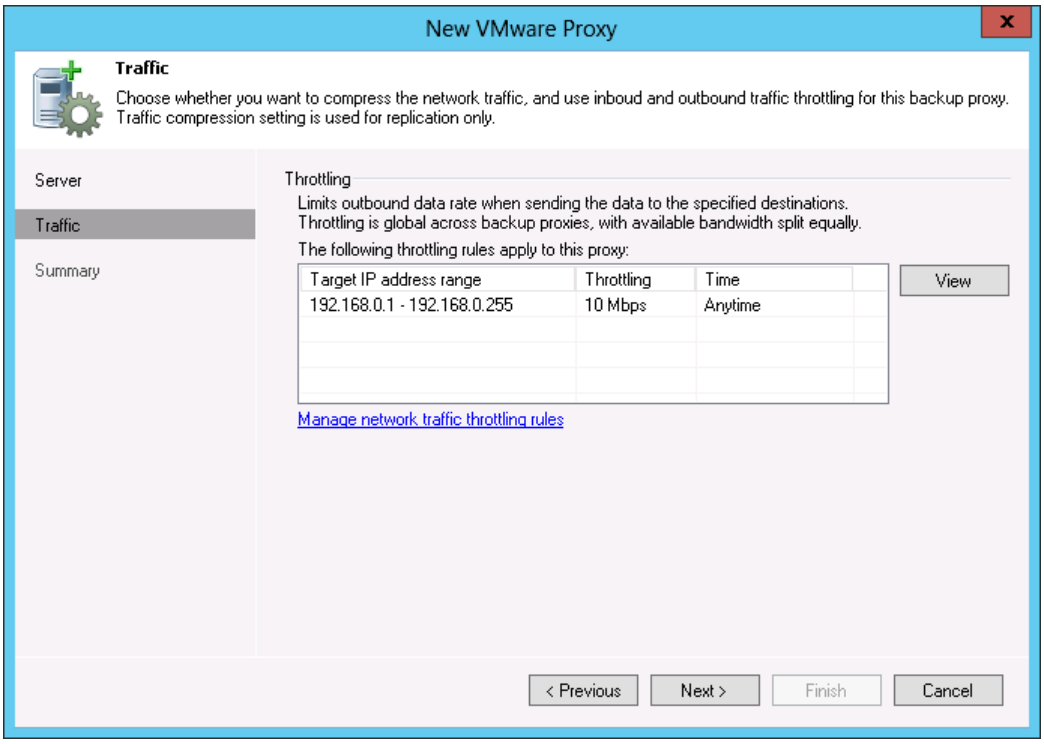

You can also open global throttling settings and modify them directly from the wizard by clicking **Manage network traffic throttling rules** at the bottom of the window.

Step 4. Finish Working with the Wizard

At this step of the wizard, Veeam Backup & Replication will add the backup proxy in the real-time mode. Once the backup proxy is added, click **Finish** to exit the wizard.

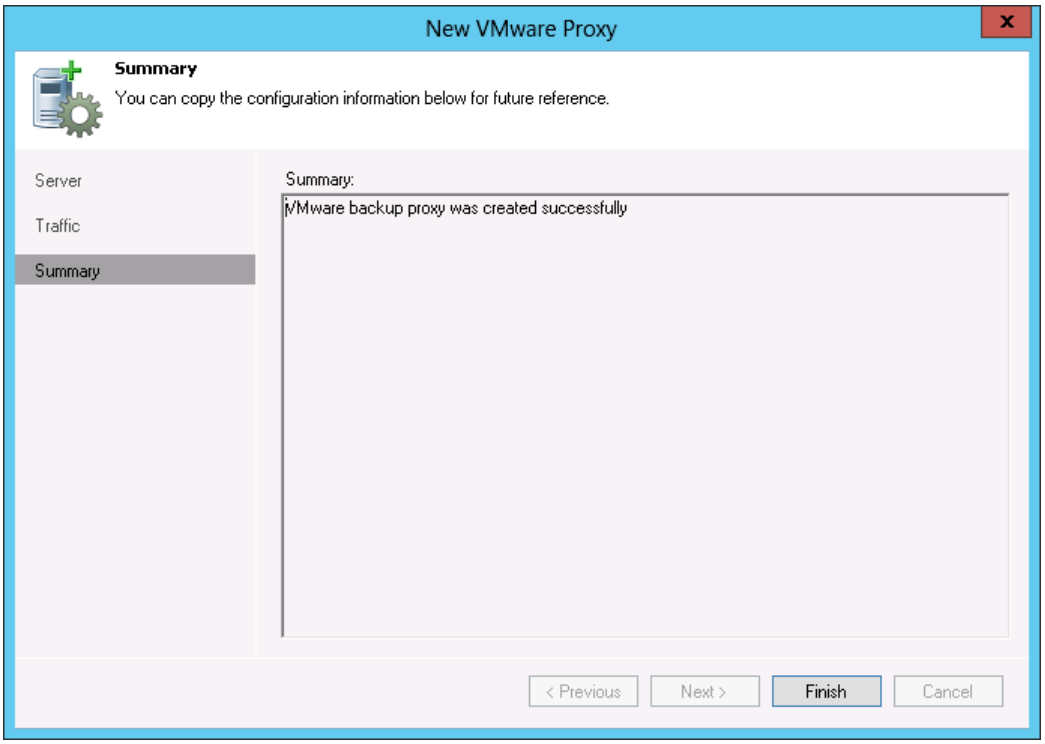

Managing Backup Proxies

To edit settings of a backup proxy:

1. Open the **Backup Infrastructure** view.
2. Select the **Backup Proxies** node in the inventory pane.
3. Select the necessary proxy in the working area and click **Edit Proxy** on the ribbon. You can also right-click the necessary proxy in working area and select **Properties**.

Then edit the backup proxy settings as required.

You can temporarily disable a backup proxy. In this case, it will not be used by any job.

1. Open the **Backup Infrastructure** view.
2. Select the **Backup Proxies** node in the inventory pane.
3. Select the proxy in the working area and click **Disable Proxy** on the ribbon. You can also right-click the necessary proxy in the working area and select **Disable proxy**. To enable the backup proxy, select it and click the **Disable Proxy** button on the ribbon or right-click the proxy in the working area and select **Disable proxy** once again.

To remove a backup proxy:

1. Open the **Backup Infrastructure** view.
2. Select the **Backup Proxies** node in the inventory pane.
3. Select the proxy in the working area and click **Remove Proxy** on the ribbon. You can also right-click the necessary proxy in the working area and select **Remove**.

When you remove a backup proxy, Veeam Backup & Replication unassigns the proxy role from the server, so it is no longer used as a backup proxy. The actual server remains connected to Veeam Backup & Replication.

Important! You cannot remove a backup proxy that is explicitly selected in any backup, replication or VM copy job. To remove such a proxy, you need to delete all job references to it first.

Adding Backup Repositories

Backup repositories are locations for storing backup data and auxiliary files. You can assign the role of a backup repository to any Windows or Linux server added to the list of managed servers in Veeam Backup & Replication, or to any shared CIFS folder to which the backup server has access. Windows-based backup repositories can also perform the role of the Veeam vPower NFS server enabling advantages of the vPower technology for multi-OS file-level restore, Instant VM Recovery, SureBackup and U-AIR capabilities. For more information, see Veeam vPower NFS Service.

To add a backup repository, follow the next steps:

Step 1. Launch the New Backup Repository Wizard

To launch the wizard, do either of the following:

- Open the **Backup Infrastructure** view, select the **Backup Repositories** node in the inventory pane and click **Add Repository** on the ribbon.
- Open the **Backup Infrastructure** view, right-click the **Backup Repositories** node in the inventory pane and select **Add Backup Repository**.

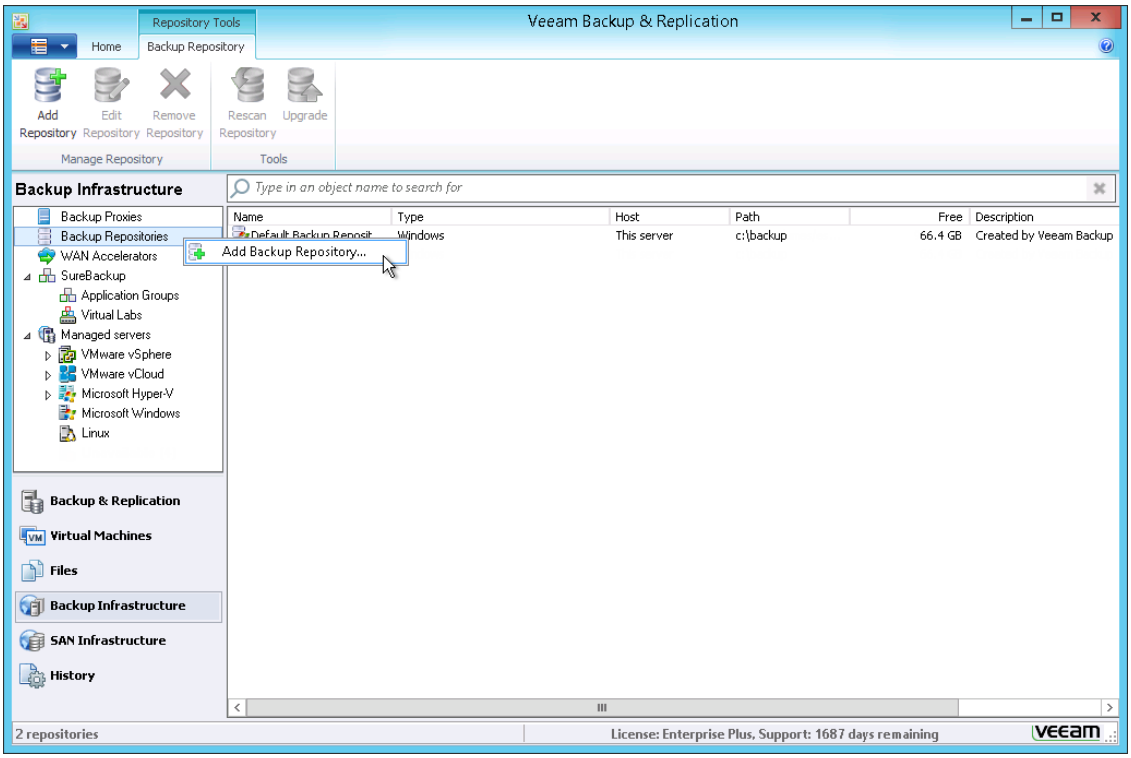

Step 2. Specify Name and Description

Specify a name for the repository and provide a description for future reference. By default, the description contains information about the user who created the backup repository, as well as the date and time when the repository was added.

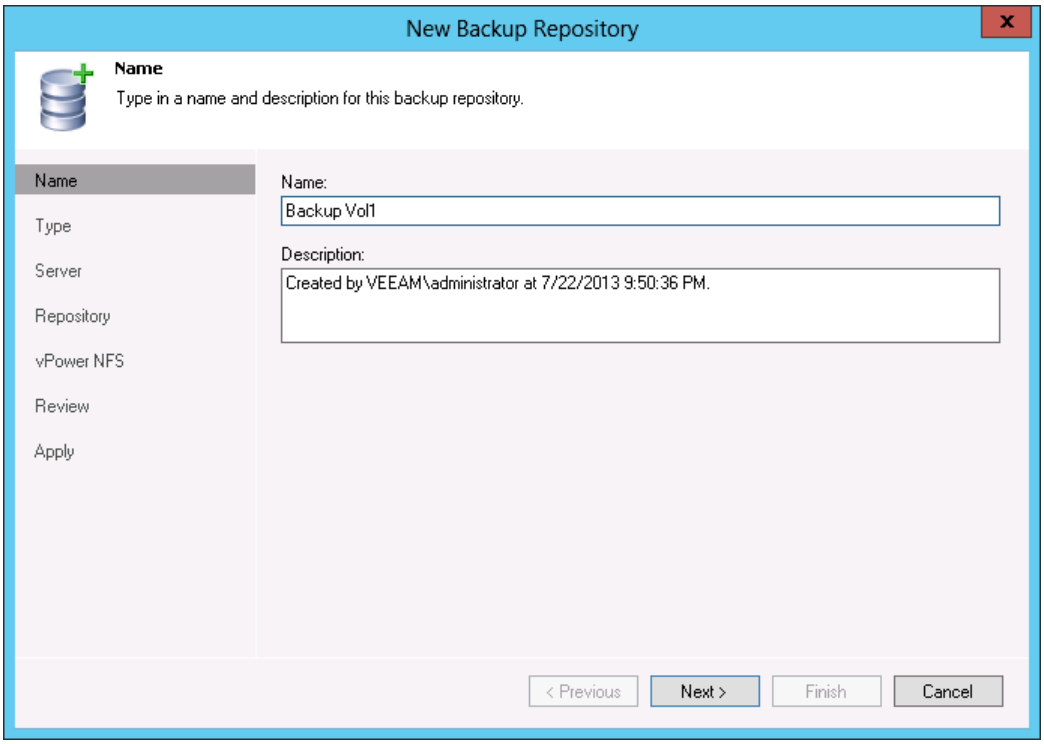

Step 3. Choose the Type of Repository

Select the type of repository you would like to add:

- **Microsoft Windows server** with local or directly attached storage. In this case, Veeam Backup & Replication will deploy the Veeam transport service on the Windows server connected to the storage system. The transport service is responsible for data processing tasks, enabling efficient backups over slow connections. This type of a backup repository can also act as the vPower NFS server (for this, remember to select corresponding settings at the next steps of the wizard). For more information, see Veeam vPower NFS Service.
- **Linux server** with local, directly attached or mounted NFS storage. In this case, Veeam Backup & Replication will deploy the Veeam transport service on the Linux server connected to the storage system. The transport service is responsible for data processing tasks, enabling efficient backups over slow connections.
- **Shared folder** using CIFS (SMB) protocol. This type of storage cannot run a Veeam transport service. If you select a shared folder as a backup repository in the offsite backup scenario (that is, the shared folder is located offsite) and your connections for sending VM data are slow, it is recommended to deploy a backup proxy closer to the backup repository.

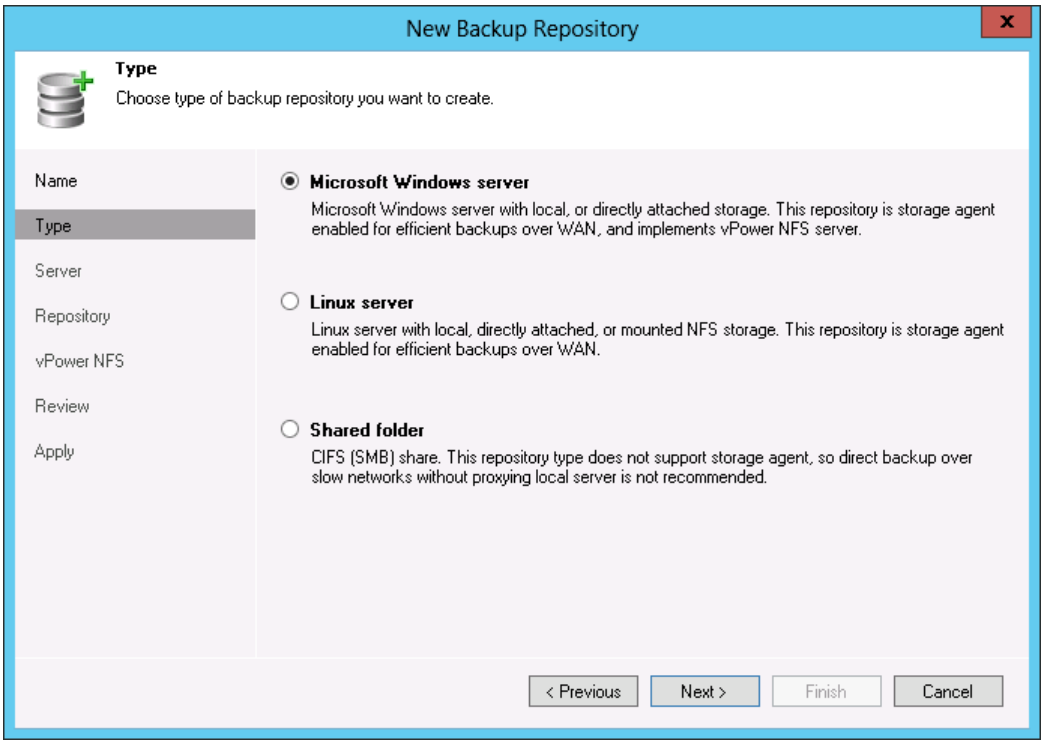

Step 4. Specify Server or Share

This step depends on the type of backup repository you selected.

Microsoft Windows Server or Linux Server

From the **Repository server** list, select the necessary Windows or Linux server to be used as a backup repository. The **Repository servers** list contains only those servers that have been added to Veeam Backup & Replication beforehand. You can also click **Add New** on the right to add a new server to be used as the repository.

Click **Populate** to see a list of datastores connected to the selected server, their capacity and free space.

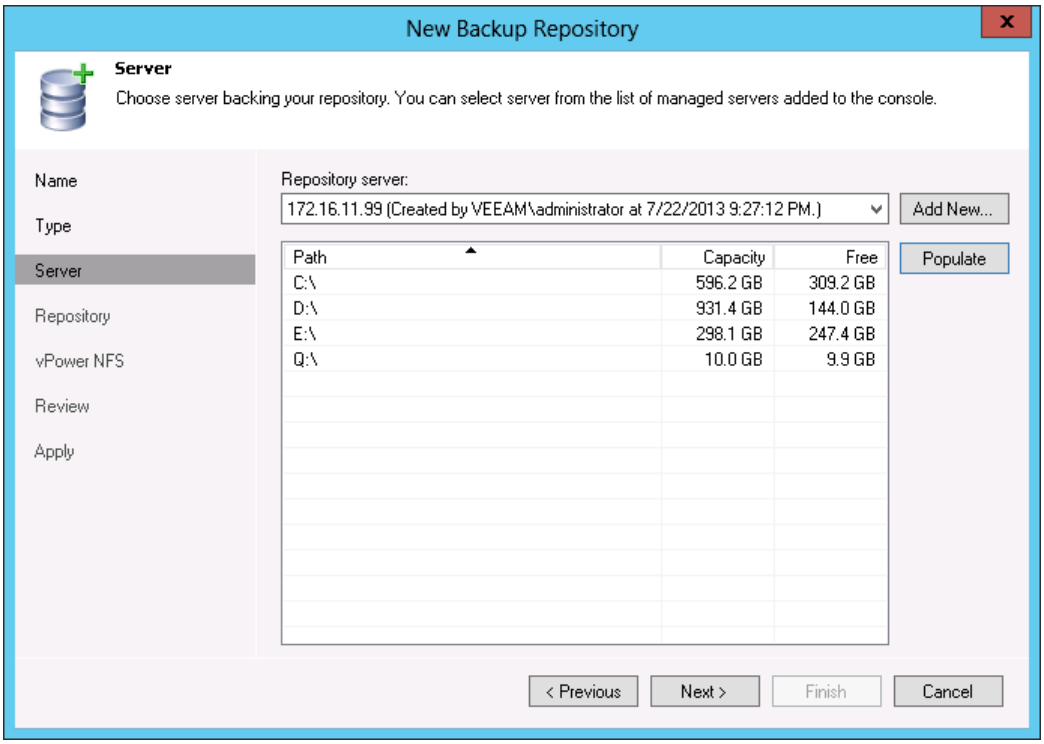

Shared folder

In the **Shared folder** field, specify the UNC path to the shared folder you want to use as a backup repository. Select credentials of an account with administrative privileges on the share. If you have not set up the necessary credentials beforehand, click the **Manage accounts** link at the bottom of the list or click **Add** on the right to add the necessary credentials. To learn more, see Managing Credentials.

Specify the way in which VM data should be written to the shared folder:

- If you are using fast connections, select the **Directly from backup proxy server** option to write VM data directly from the source-side backup proxy to the repository.
- If you are planning to perform offsite backup over WAN connections, select the **Through the following proxying server** option and specify an additional proxy server on the target side which will be used for moving data to the backup repository.

If you use a shared folder as a backup repository and do not specify the proxying server, Veeam Backup & Replication will deploy the target Veeam transport service on any available backup proxy having access to the shared folder. The backup proxy is picked at random: Veeam Backup & Replication may use one backup proxy for one job session and another backup proxy for another job session.

In some cases, however, such behavior may cause problems. For example, during one job session Veeam Backup & Replication may use a 64-bit backup proxy to create a backup file. If during the next job session Veeam Backup & Replication uses a 32-bit backup proxy, Veeam Backup & Replication will fail to open the created backup file on the backup repository. To overcome this situation, it is recommended to explicitly define the proxying server on which the target Veeam transport service will be deployed.

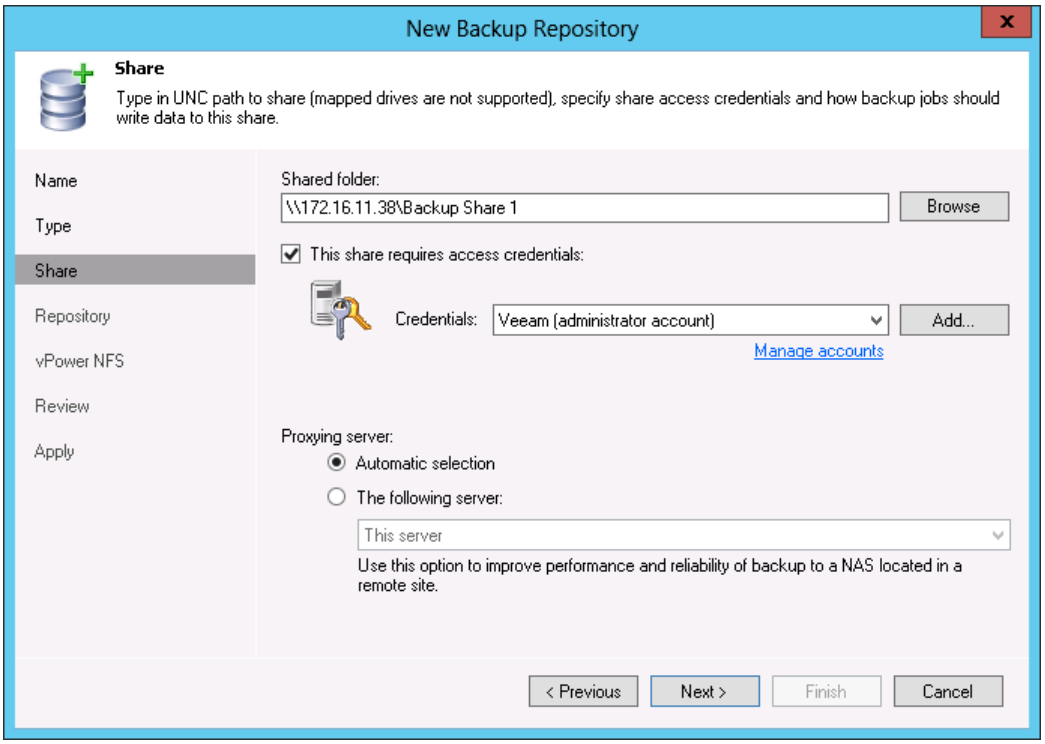

Step 5. Configure Path and Load Control Settings

In the **Location** section, specify the path to the folder to which backup files should be stored. Click **Populate** to see the capacity and available free space on the selected partition.

In the **Load control** section, set the necessary values to limit the number of concurrent jobs for the repository. If the specified threshold has been reached, a new job using this repository will not start. You can also limit data ingestion rate to restrict the total speed of writing data to the repository disk. Limiting the number of concurrent tasks and data ingestion rate will help you control the load on the repository and prevent possible timeout of storage I/O operations. For more information, see Resource Scheduling.

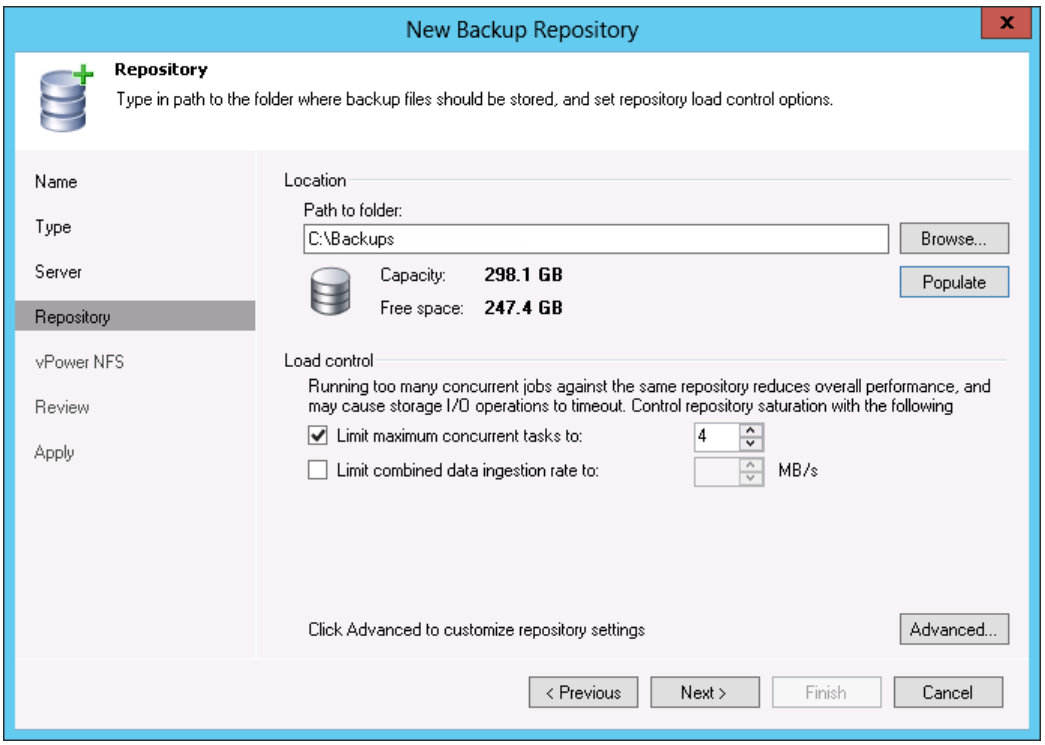

If you plan to use a deduplicating storage appliance, click **Advanced** to configure additional repository settings:

- For storage systems using fixed block size, select the **Align backup file data blocks** check box. Veeam Backup & Replication will align VM data saved to a backup file to a 4Kb block boundary. This option provides better deduplication across backup files, but can result in greater amount of unused space on the storage device and higher level of fragmentation.
- When you enable compression for a backup job, VM data is compressed at the source side before it is transmitted to the target. However, compressing data prior to writing it to deduplicating storage appliance results in poor deduplication ratios as the number of matching blocks decreases. To overcome this situation, you can select the **Decompress backup data blocks before storing** check box. If data compression is enabled for a job, Veeam Backup & Replication will compress VM data, transmit it over LAN, uncompress data on the target side and write raw VM data to the storage device to achieve a higher deduplication ratio.

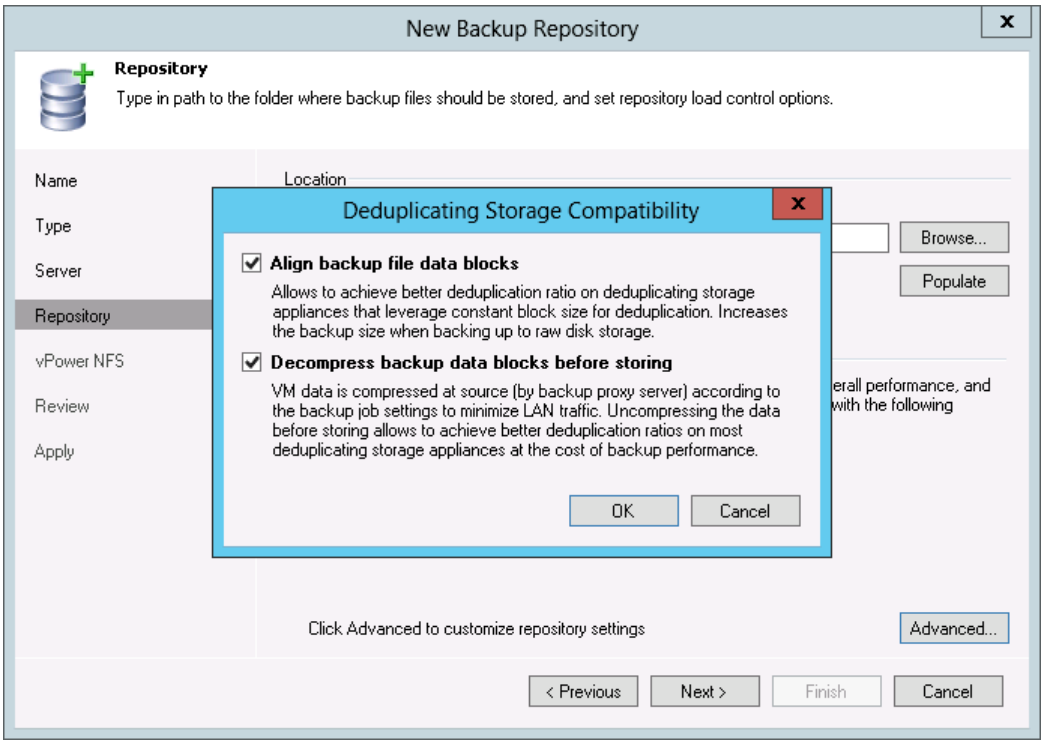

Step 6. Specify vPower NFS Settings

Select the **Enable vPower NFS server** check box to make the repository accessible by the vPower NFS Service. From the list below, choose a Windows server that will be used as the vPower NFS server. You can select any Windows server from the list of those added to Veeam Backup & Replication, or choose the **Add Server** option to assign this role to a Windows server not added to the application console. Veeam Backup & Replication will install the Veeam vPower NFS Service on the selected server, which will enable running VMs directly from backup files for advanced backup verification and recovery functionality. For more information, see Veeam vPower NFS .

In the **Folder** field, specify the folder where instant VM recovery write cache will be stored. Please note that the selected volume should have at least 10 GB of free disk space.

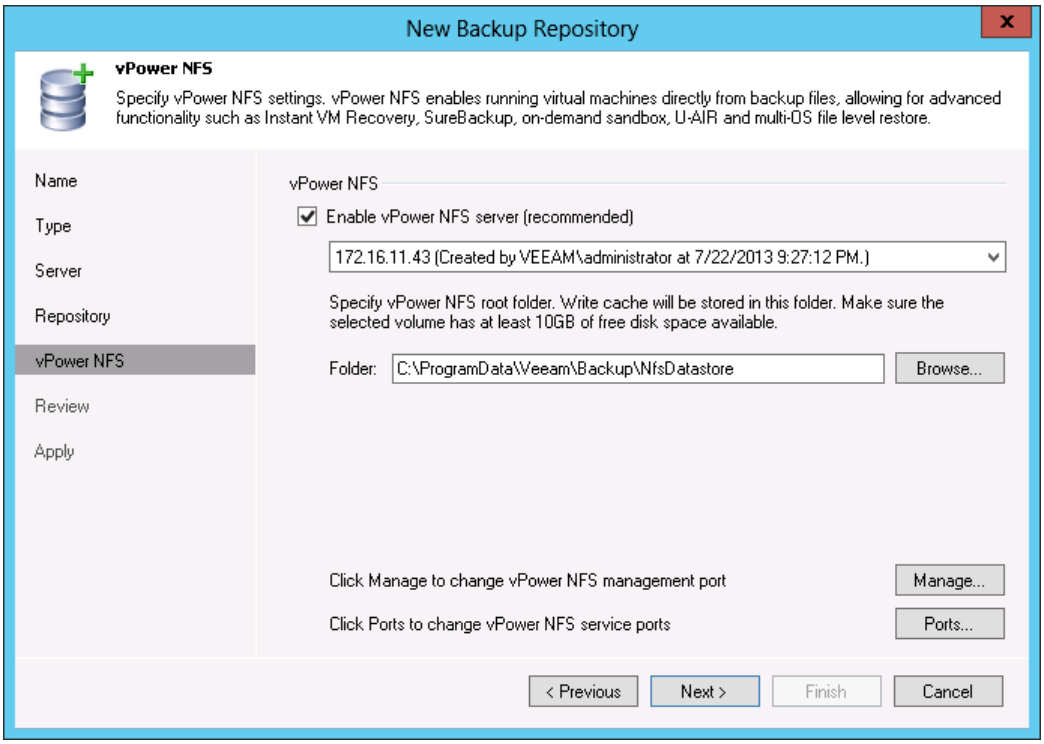

Click **Manage** to open the **Network Settings** window and customize network ports for individual components:

- Veeam Installer Service is responsible for deploying the *vPower NFS* and Veeam Transport on the Windows server. By default, the Veeam Installer Service uses port number 6160.
- vPower NFS Service is responsible for providing ESX(i) hosts with transparent access to backed up VM images. By default, the vPower NFS Service uses port number 6162.
- Veeam Transport is responsible for deploying the corresponding services during backup and replication processes. By default, the service uses port number 6162.

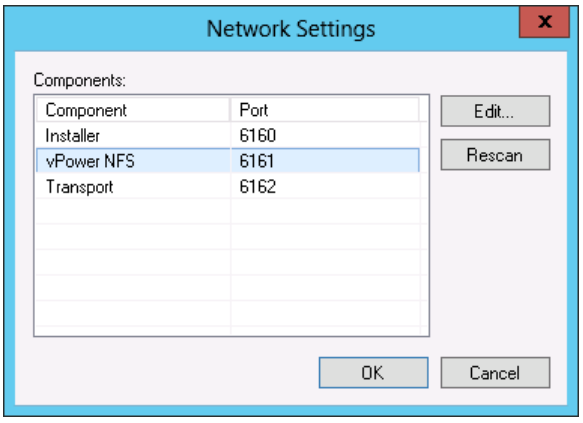

Click **Ports** to open the **Ports Settings** window and customize the following ports for vPower NFS Service:

- RPC port (by default, port number 6161 is used)
- Mount port (by default, port number 1058 is used)
- vPower NFS port (by default, port number 2049 is used)

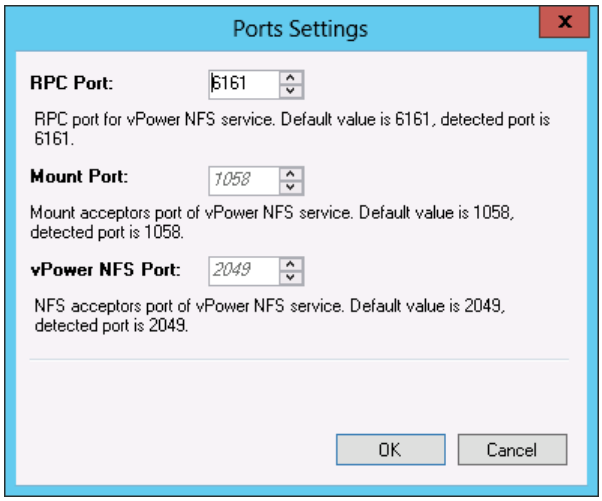

Step 7. Review Properties and Components

After the wizard checks for existing components, you can review the repository properties and installed components.

Select the **Import existing backups automatically** check box. Veeam Backup & Replication will scan the repository folder for existing backup files and automatically add them to the Veeam Backup & Replication console under the **Backups** node.

If the repository folder contains guest file system index files that were previously created by Veeam Backup & Replication, select the **Import guest file system index** check box. The index files will be imported along with backups, so you will be able to search for guest OS files inside the imported backups.

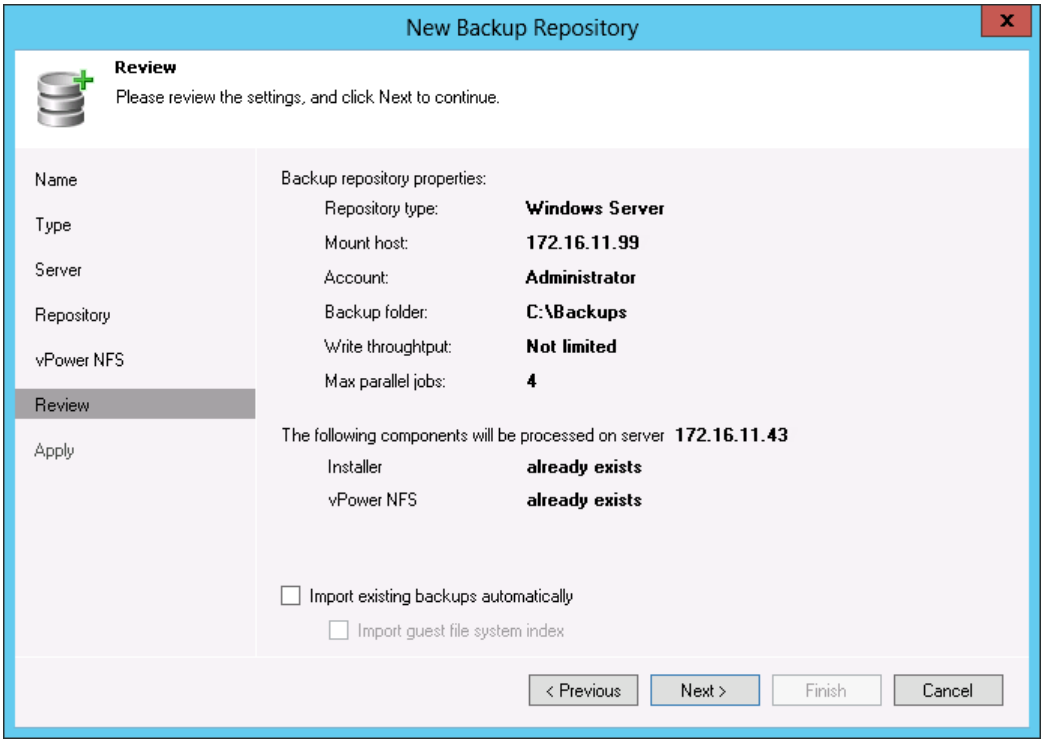

Step 8. Finish Working with the Wizard

You will see real time processing results in the log. Wait for the required operations to be performed. When the wizard completes adding the backup repository, you can review the log information. Click **Finish** to exit the wizard.

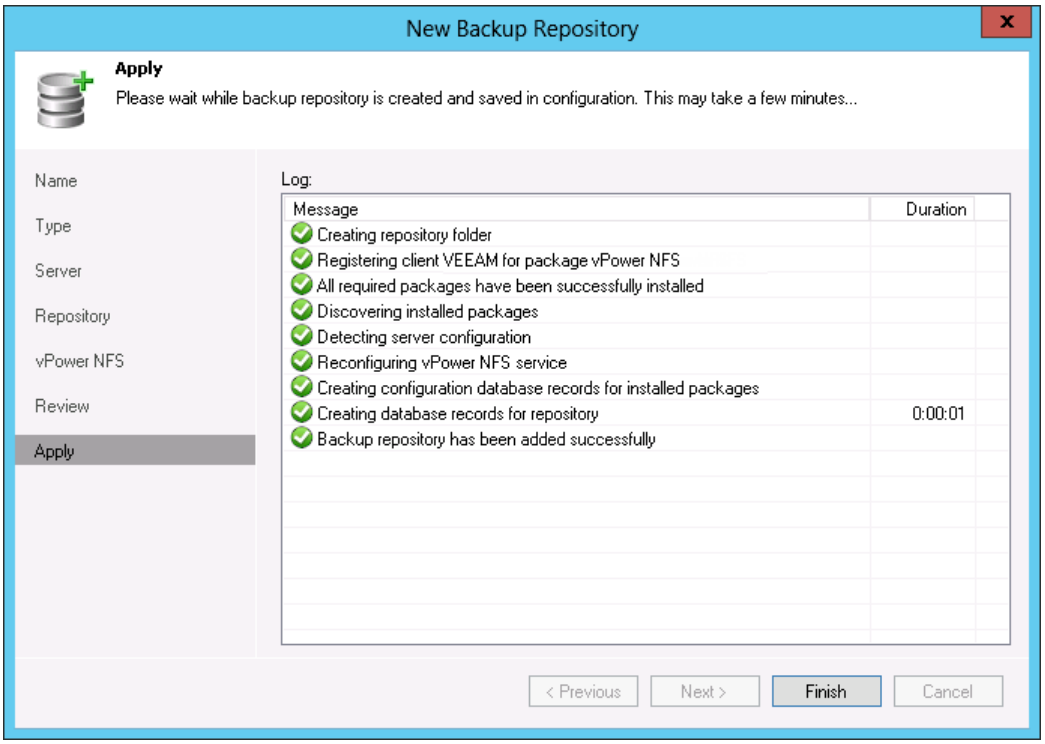

Managing Backup Repositories

To edit settings of an added backup repository:

1. Open the **Backup Infrastructure** view.
2. Select the **Backup Repositories** node in the inventory pane.
3. Select the necessary backup repository in the working area and click **Edit Repository** on the ribbon. You can also right-click the necessary repository in working area and select **Properties**.

Then edit the repository settings as required.

To update information on backups stored on a repository, you can perform repository rescan. Rescanning can be required if you have moved backups from repository to tape or if you have copied backups to the repository. To make sure that the Veeam Backup & Replication database stores up-to-date information about the backups that the repository hosts, open the **Infrastructure** view, select the **Backup Repositories** node in the inventory pane, right-click the necessary repository in the working area and select **Rescan repository**.

To remove a backup repository:

1. Open the **Backup Infrastructure** view.
2. Select the **Backup Repositories** node in the inventory pane.
3. Select the necessary backup repository in the working area and click **Remove Repository** on the ribbon. You can also right-click the necessary repository in working area and select **Remove**.

When you remove a backup repository, Veeam Backup & Replication unassigns the repository role from the server, so it is no longer used as a backup destination. The actual server remains connected to Veeam Backup & Replication.

Important! You cannot remove a backup repository that is selected in any backup or replication job. To remove such a repository, you need to delete all job references to it first.

Adding WAN Accelerators

To optimize traffic going over the WAN during backup copy jobs, you need to configure a pair of WAN accelerators.

- One WAN accelerator must be configured on the source side, closer to the source repository.
- The other WAN accelerator must be configured on the target side, closer to the target repository.

To deploy a WAN accelerator, you should assign this role to a Microsoft Windows machine added to the list of managed servers in Veeam Backup & Replication. The Microsoft Windows machine must meet the following requirements:

1. You can use either physical or virtual Microsoft Windows machine as a WAN accelerator. The role can be assigned to backup proxies and backup repositories existing in your backup infrastructure as well.
2. You can use only 64-bit Microsoft Windows machines as WAN accelerators. 32-bit versions of Microsoft Windows are not supported.
3. WAN acceleration operations are resource-consuming. When creating a WAN accelerator, mind available CPU and RAM resources on the Microsoft Windows machine that you plan to use as a WAN accelerator. It is recommended to assign this role to machines with 8 GB RAM and more. Otherwise the WAN acceleration process may fail.

To configure a WAN accelerator, follow the next steps:

Step 1. Launch the New WAN Accelerator Wizard

To launch the **New WAN Accelerator** wizard, do one of the following:

- Open the **Backup Infrastructure** view, select the **WAN Accelerators** node in the inventory pane and click **Add WAN Accelerator** on the ribbon.
- Open the **Backup Infrastructure** view, right-click the **WAN Accelerators** node in the inventory pane and select **Add WAN Accelerator**.

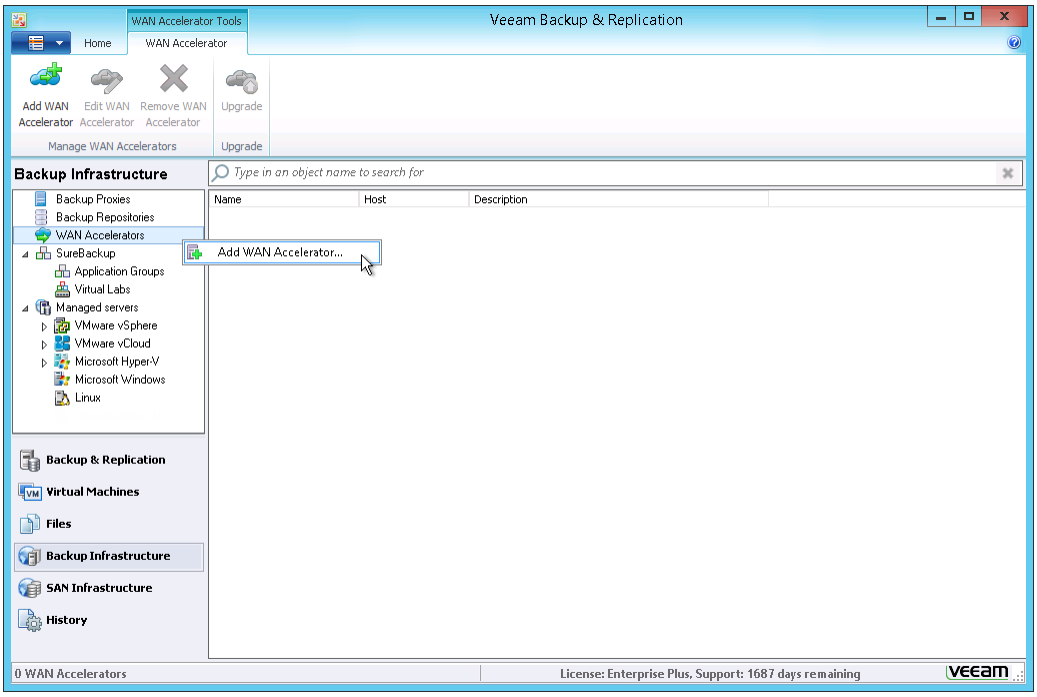

Step 2. Choose a Server

At the **Server** step of the wizard, you should select a Microsoft Windows server to which the WAN accelerator role will be assigned and define port and connection settings for the added WAN accelerator:

1. From the **Choose server** list, select a Windows server added to Veeam Backup & Replication. If the server is not added yet, you can click **Add New** to open the **New Windows Server** wizard.
2. In the **Description** field, provide a description for future reference. It is recommended that you describe the added WAN accelerator as the source or the target one. When you create a backup copy job, this hint will be displayed in brackets next to the WAN accelerator name, which will help you choose the necessary WAN accelerator to be used in the source and target sites.
3. In the **Traffic port** field, specify the number of the port over which WAN accelerators will communicate with each other. By default, port 6165 is used.
4. In the **Streams** field, specify the number of connections that should be used to transmit data between WAN accelerators. By default, 5 connections are used.

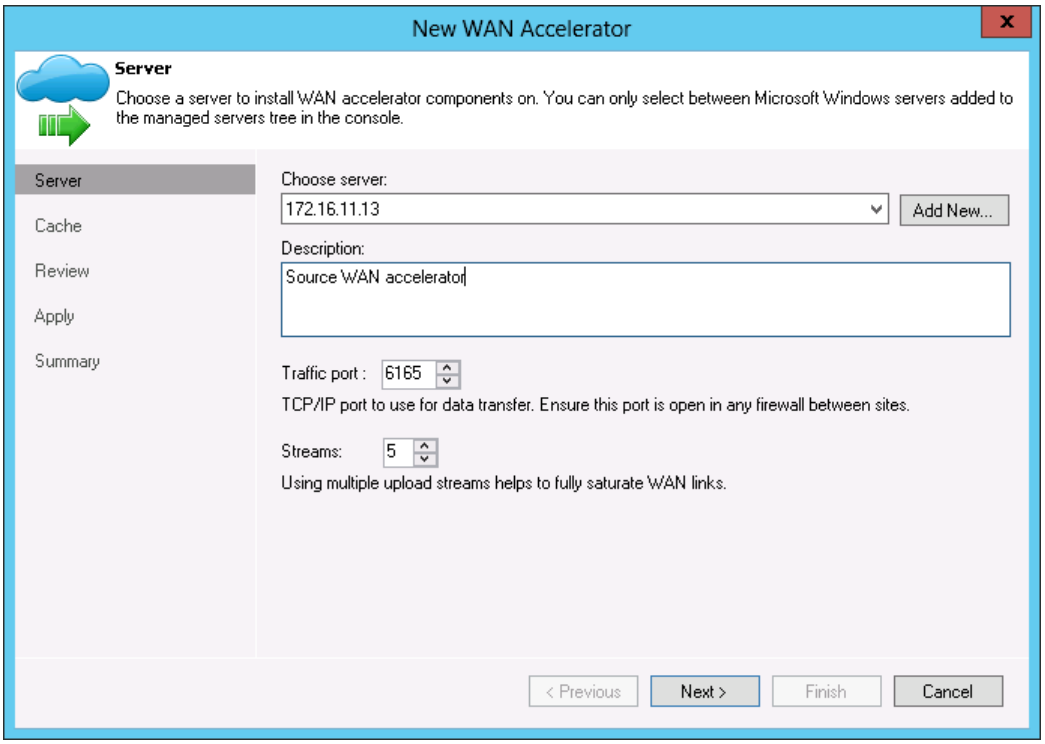

Step 3. Define Cache Location and Size

At the **Cache** step of the wizard, you should define settings for the *VeeamWAN* folder that will be created on the added WAN accelerator.

1. In the **Folder** field, specify a path to the folder in which global cache data and Veeam WAN service files should be stored. When selecting a folder for the target WAN accelerator, make sure that there is enough space for holding global cache data.
2. (For the target WAN accelerator) In the **Cache size** field, specify the size for the global cache. The global cache size is specified per source WAN accelerator. That is, if you plan to use one target WAN accelerator with several source WAN accelerators, the specified amount of space will be allocated to every source WAN accelerator that will be working with the target WAN accelerator and the size of the global cache will increase proportionally. To learn more, see Many to One WAN Acceleration.

Important! It is not recommended to specify a path of significant depth for the global cache folder. During WAN acceleration operations, Veeam Backup & Replication generates service files having long file names. Placing these files to a folder of significant depth may cause problems in the NTFS file system.

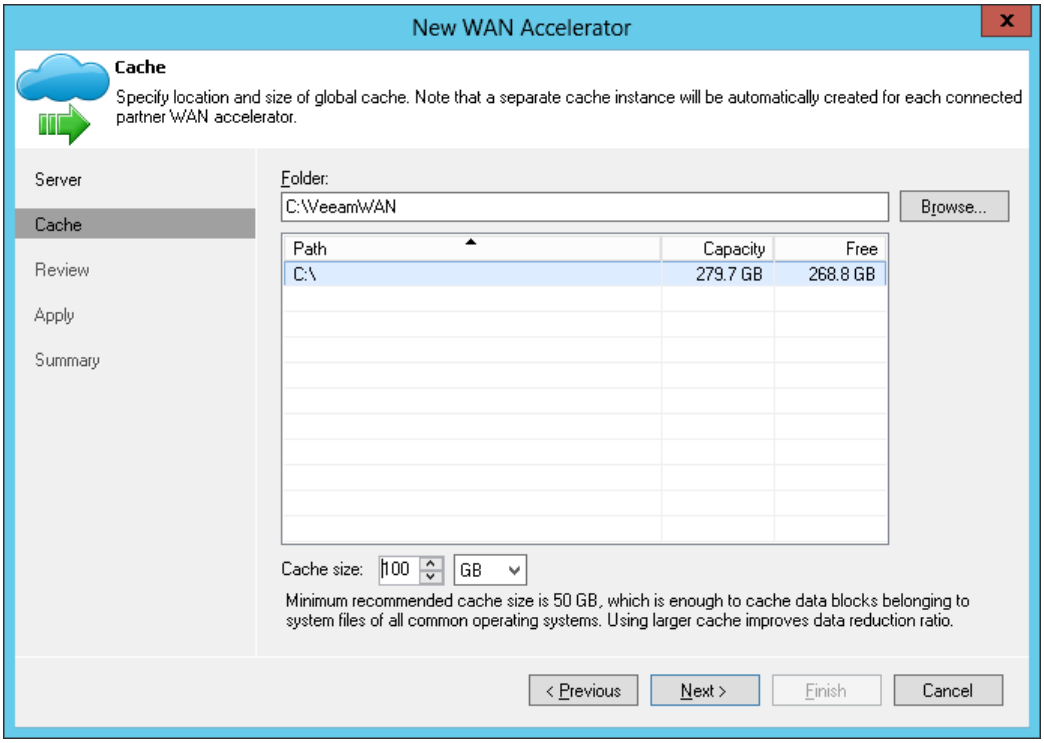

Step 4. Review Components

At the **Review** step of the wizard, Veeam Backup & Replication will display the list of components required for work of the WAN accelerator:

- Veeam Transport
- Veeam WAN Accelerator

If any of them is missing, Veeam Backup & Replication will automatically install them on the selected server.

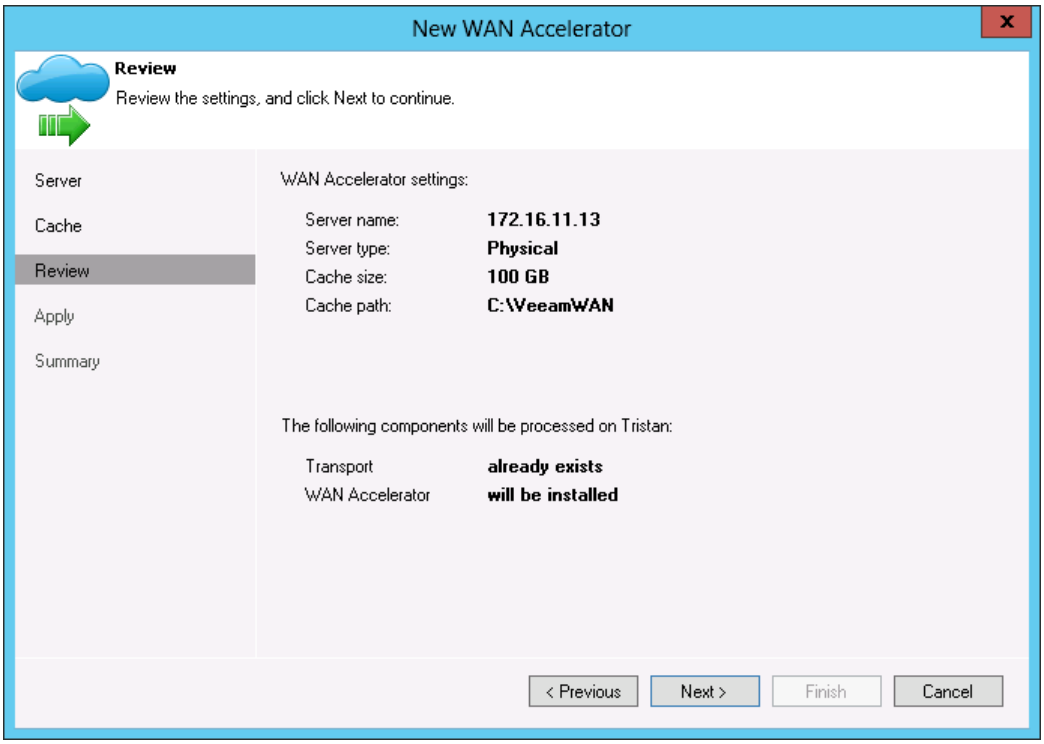

Step 5. Assess Results

At the **Apply** step of the wizard, Veeam Backup & Replication will add the WAN accelerator to the backup infrastructure in the real time mode.

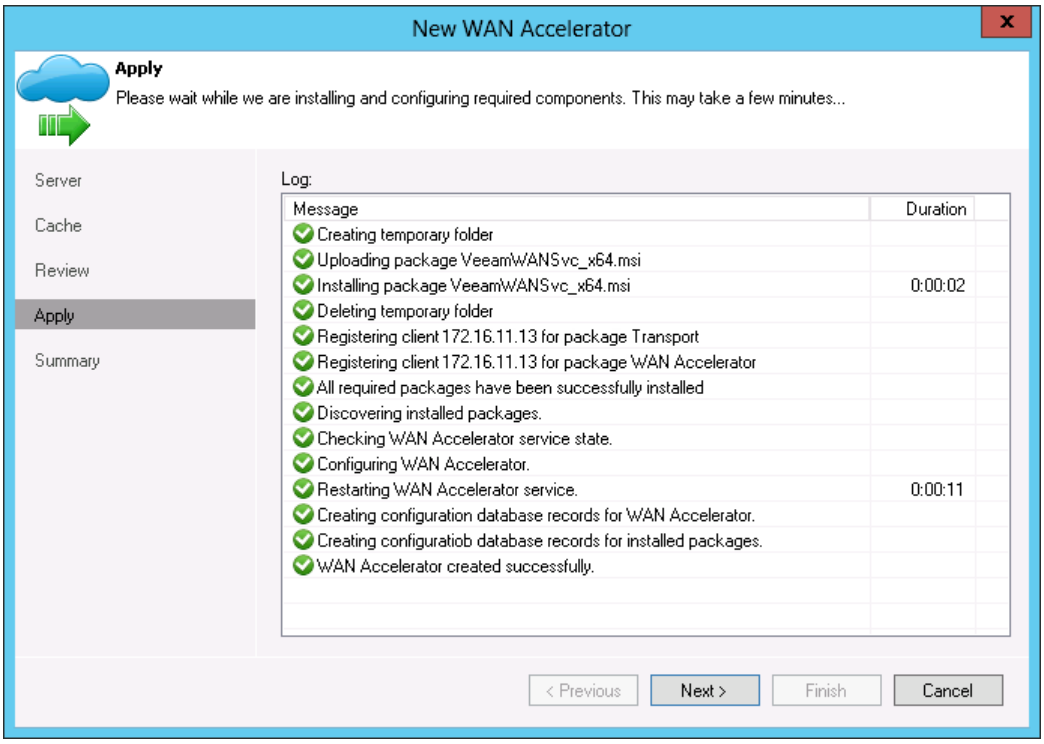

Step 6. Finish Working with the Wizard

Once the WAN accelerator is added, review the summary and click **Finish** to exit the wizard.

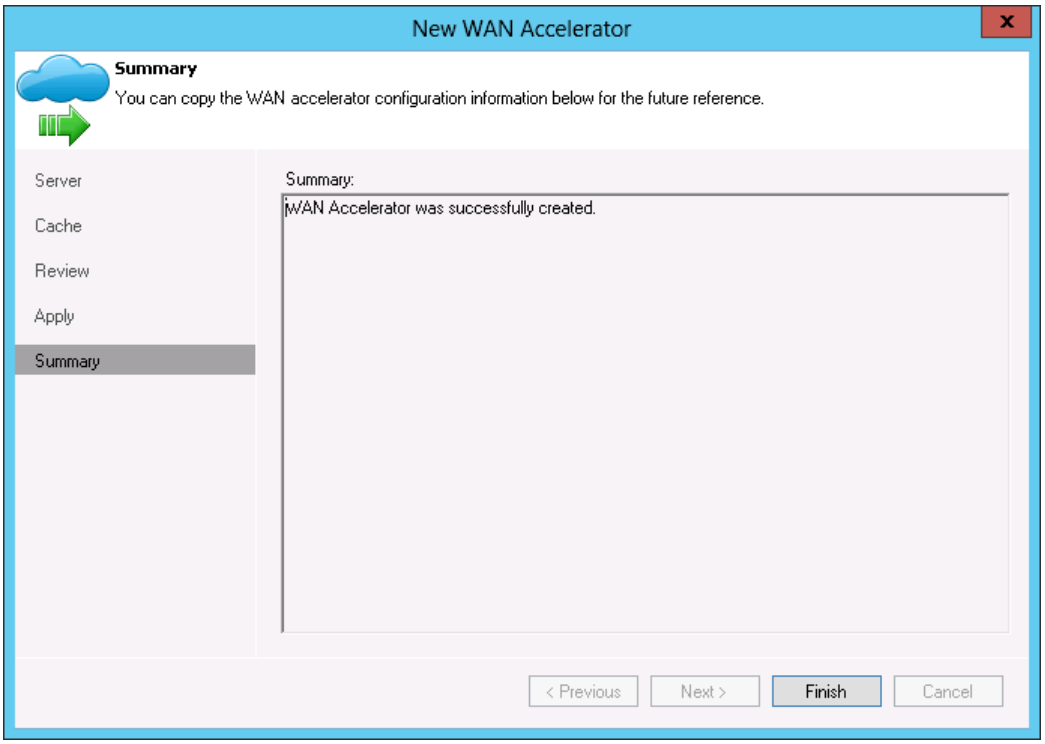

Clearing Global Cache on WAN Accelerator

In some cases, you may need to remove data from the global cache. Such situation can occur, for example, if data in the global cache gets corrupted.

Another use case is to clear the global cache to remove existing data and re-load new data to it. This situation may occur if you plan to copy VMs of a completely different type. For example, if you have copied VM running Microsoft Windows 2008R2 for some time, the global cache will contain data blocks of Microsoft Windows 2008R2 VMs. After that, you may move to another version of OS, for example, Microsoft Windows 2012 and may want to copy VMs running this OS. In this case, existing data blocks in the global cache will be inappropriate and of no use for the backup copy job. In such situation, it is recommended to clear the global cache before you start processing new types of VMs. Veeam Backup & Replication will not need to continuously remove "old" data blocks from the global cache and replace them with the "new" ones. During the first run of the new backup copy job, Veeam Backup & Replication will populate the global cache with data blocks of appropriate type and these blocks will be re-used further on.

To clear the global cache:

1. Open the **Backup Infrastructure** view.
2. In the inventory pane, click the **WAN Accelerators** node.
3. In the working area, right-click the necessary WAN accelerator and select **Clear cache**.

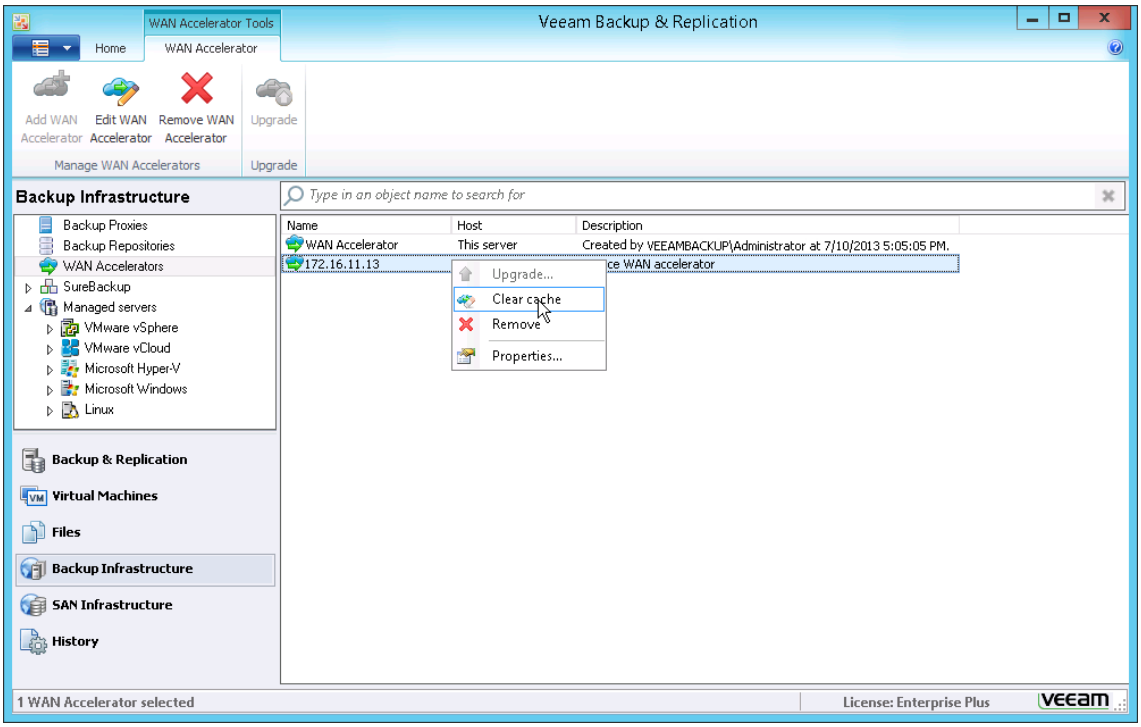

Managing Network Traffic

If you plan to perform offsite backup or replicate VMs to a remote DR site, you can manage network traffic by applying traffic throttling rules or limiting the number of data transfer connections. To learn more about network traffic management possibilities, see Network Traffic Throttling and Multithreaded Data Transfer.

Setting Network Traffic Throttling Rules

Managing Data Transfer Connections

Setting Network Traffic Throttling Rules

Network throttling rules are applied to limit the maximum throughput of traffic going from source to target. In Veeam Backup & Replication, network traffic throttling rules are created and enforced globally, at the level of the Veeam backup server. Rules are set not for a single IP address, but for a pair of IP address ranges, on the source side and on the target side.

Throttling rules are checked against backup infrastructure components between which VM data is transferred over the network: more specifically, on which Veeam Transport Services engaged in the job are started. In case of backup, replication or VM copy jobs, throttling rules are checked against the following components:

- On the source side: backup proxies, either dedicated or the default one — Veeam backup server.
- On the target side:
  - For backup jobs: Windows- or Linux-based backup repository or a proxying server (in case you use a CIFS share as a backup repository).
  - For replication jobs: backup proxy.

In case of backup copying jobs, throttling rules are checked against the following components:

- If you copy backup files over the direct path, the throttling rule is checked against the Windows- or Linux-based backup repository or a proxying server (in case you use a CIFS share as a backup repository).
- If you copy backup files via WAN accelerators, the throttling rule is checked against the source and target WAN accelerators.

If the IP address of the backup infrastructure component falls into the specified source IP range of a rule, the corresponding rule will be applied to it. For example, if you specify 192.168.0.1 – 192.168.0.30 as the source range for a network traffic throttling rule and the backup proxy on the source side has IP address 192.168.0.12, this rule will be applied to the proxy. The network traffic going from this backup proxy to the target side will be throttled.

To create a throttling rule:

1. Select **Traffic Throttling** from the main menu.
2. In the **Global Network Traffic Throttling** section click **Add**.
3. In the **Source IP address range** section, specify a range of IP addresses for backup infrastructure components from which VM data will be transferred over the network.
4. In the **Target IP address range** section, specify a range of IP addresses for backup infrastructure components to which transferred VM data will be targeted.
5. In the **Throttle network traffic to** section, specify the maximum speed that can be used to transfer traffic from source servers to target servers.
6. In the **Apply this rule** section, specify the period during which the rule should be enforced. You can select to use throttling rules all the time, or schedule traffic throttling for specific time intervals (for example, for business hours, to minimize the impact of job performance on the production network).

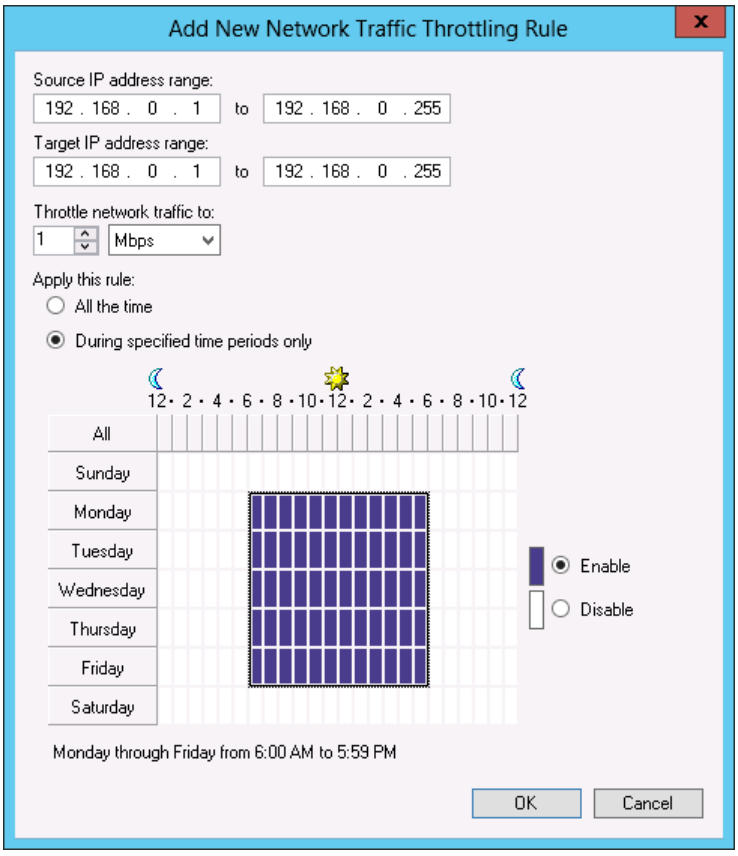

For example, to manage network traffic during business and non-business hours, you can create two throttling rules:

- Limit the speed to 1 Mbps Monday through Friday from 7 AM to 7 PM
- Limit the speed to 10 Mbps on weekends and from 7 PM to 7 AM on weekdays

In this case, Veeam Backup & Replication will limit the transfer speed to 1 Mbps during business hours, while during non-business hours the speed will be limited to 10 Mbps.

Note If you create several traffic throttling rules for the same range of IP addresses, make sure that time intervals when these rules are enabled do not overlap.

Network traffic throttling rules that apply to a specific proxy can be viewed at the **Traffic** step of the backup proxy wizard.

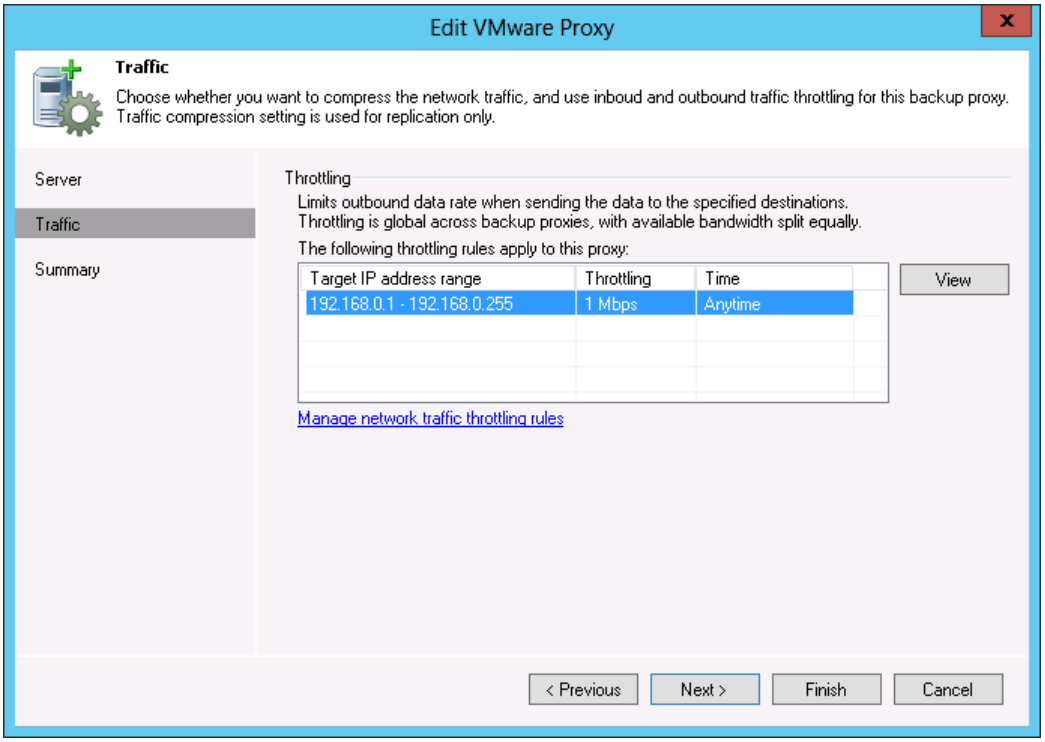

Note that several network traffic rules applied to the same backup infrastructure component may cover the same range of target IP addresses. In case two throttling rules use the same target IP address range, but have different speed limits, the rule with the lowest transfer speed will be used.

For example, there is a 4 Mbps throttling rule for a server with the 192.168.0.12 address and a 1 Mbps rule for the 192.168.0.1 – 192.168.0.30 range. While both rules apply to the server with the 192.168.0.12 address, Veeam Backup & Replication will use the lowest transfer speed for the 192.168.0.12 server: that is, 1 Mbps.

Managing Data Transfer Connections

By default, for every job session, Veeam Backup & Replication uses multithreaded data transfer. VM data from source to target is transferred over five TCP/IP connections. However, when several jobs are scheduled to run at the same time, the load on the network may be heavy. If the network capacity is not sufficient to support multiple data transfer connections, you can configure network traffic throttling rules or disable multithreaded data transfer.

To disable multithreaded data transfer:

1. Select **Traffic Throttling** from the main menu.
2. In the **Global Network Traffic Throttling** section, clear the **Use multiple upload streams per job** check box. With this option disabled, Veeam Backup & Replication will use only one TCP/IP transfer connection for every job session.

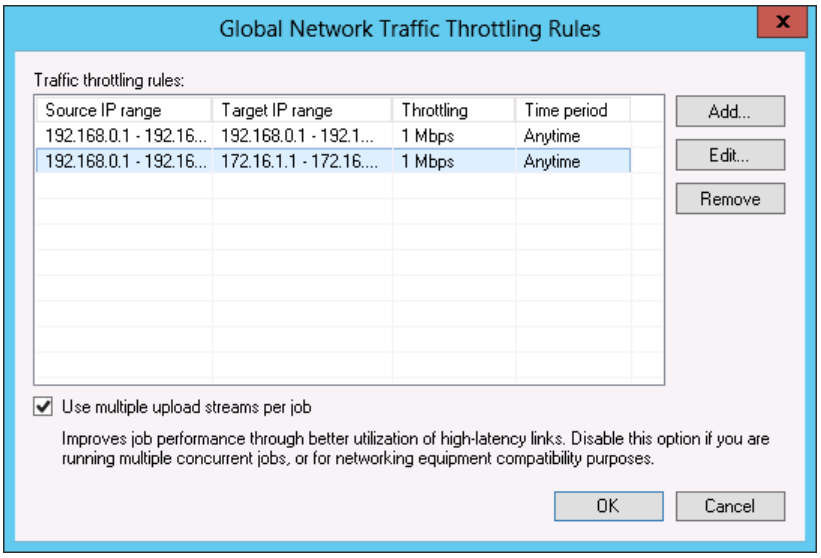

Managing Jobs

Any backup, replication or VM copy operation performed with Veeam Backup & Replication is a job-driven process. A job is a specific task that can be accomplished immediately after its creation, saved for future or scheduled to run automatically, at a specific time. To create a job, you should run a corresponding wizard and complete all wizard steps.

To view all created jobs, open the **Backup & Replication** view and select the **Jobs** node in the inventory pane. The list of available jobs will be displayed in the working area. You can edit job properties, start and stop jobs, restart failed jobs, view job statistics data and delete unnecessary jobs. Commands for all listed operations are available from the shortcut menu.

Creating Backup Jobs

To perform backup of VMs, you should create a backup job by means of the **New Backup Job** wizard. This section will guide you through all steps of the wizard and provide explanation on available options.

Before You Begin

- Prior to creating a backup job, make sure you have set up all necessary backup infrastructure components for the job. Open the **Backup Infrastructure** view and check if the source hosts are available under the **Managed servers** node in the inventory pane, backup proxies and backup repositories are available under the **Backup Proxies** and **Backup Repositories** nodes and properly configured. You will not be able to add backup infrastructure components or change their configuration once the **New Backup Job** wizard is launched; you will need to edit settings of the job after it is created to add these components. To learn more, see Setting Up Backup Infrastructure.
- During every job run, Veeam Backup & Replication checks disk space on the destination storage. If the disk space is below a specific threshold value, Veeam Backup & Replication will display a warning in the job session log. To specify the disk space threshold, select **Options** from the main menu. On the **Notifications** tab, specify the amount of free disk space required in percent.

To create a backup job, follow the next steps:

Step 1. Launch the New Backup Job Wizard

To run the **New Backup Job** wizard, do one of the following:

- On the **Home** tab, click **Backup Job** and select **VMware**.
- Open the **Backup & Replication** view, right-click the **Jobs** node and select **Backup** > **VMware**.
- Open the **Virtual Machines** view, select one or several VMs in the working area, click **Add to Backup** on the ribbon and select **New job**. Alternatively, you can right-click one or several VMs in the working area and select **Add to Backup Job** > **New job**. In this case, the selected VMs will be automatically included into the backup job. You can add other VMs to the job when passing through the wizard steps.

You can quickly include VMs to already existing jobs. To do that, open the **Virtual Machines** view, right-click necessary VMs in the working area and select **Add to Backup Job** > *name of a created job*.

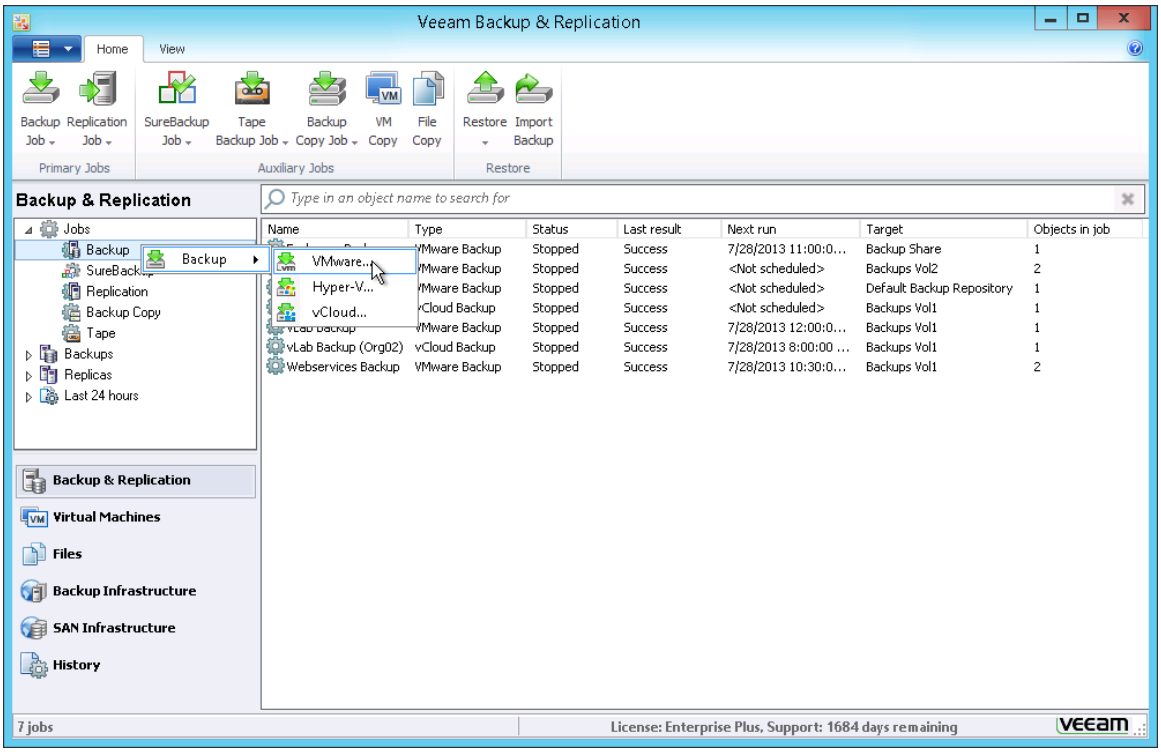

Step 2. Specify Job Name and Description

At the first step of the wizard, enter a name and description for the job. The default description contains information about the user who created the job, as well as the date and time when the job was created.

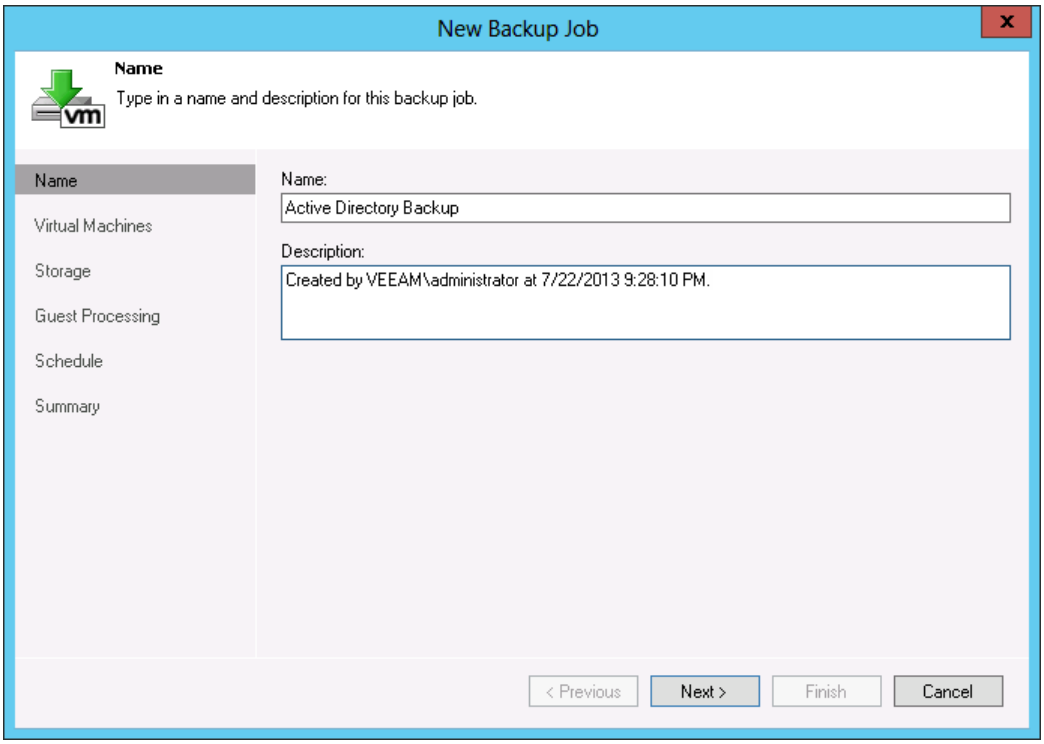

Step 3. Select Virtual Machines to Back Up

At this step, you should select an individual VM, multiple VMs or VM containers which you want to back up. Jobs with VM containers are dynamic in their nature: if a new VM is added to the container after a backup job is created, the job will be automatically updated to include the added VM.

Click **Add** to browse to VMs and VM containers that should be backed up. In the displayed tree, select the necessary object and click **Add**.

To facilitate objects selection, you can switch between views by clicking **Hosts and Clusters**, **VMs and Templates** or **Datastores and VMs** at the top of the tree. In addition, you can use the search field at the bottom of the **Add Objects** window: click the button to the left of the field and select the necessary type of object to search for (**Everything**, **Folder**, **Cluster**, **Host**, **Resource pool**, **VirtualApp** or **Virtual machine**), enter the object's name or a part of it and click the **Start search** button on the right or press **[ENTER]**.

Note	Depending on the view you select, some objects may not be available (for example, if you select the **VMs and Templates** view, no resource pools, hosts or clusters will be shown in search results).

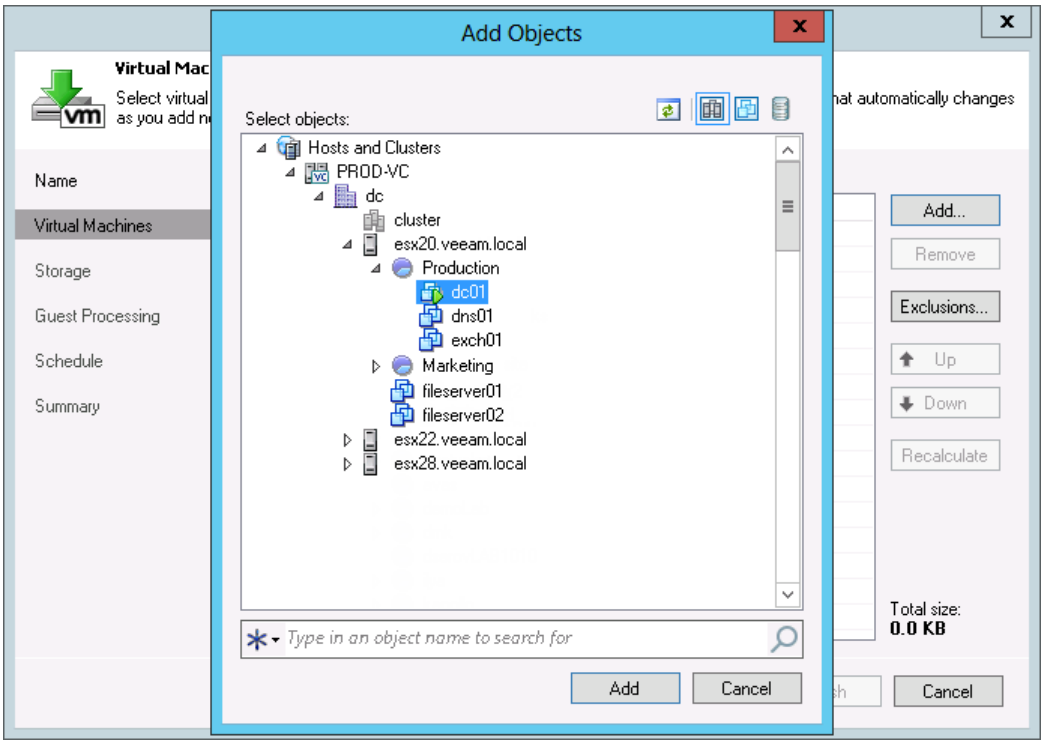

To remove an object from the list, select it and click **Remove** on the right.

The initial size of VMs and VM containers added to the backup job is displayed in the **Size** column in the list. The total size of backed up objects is displayed in the **Total size** field. Use the **Recalculate** button to refresh the total size value after you add a new object to the job.

Step 4. Exclude Objects from the Backup Job

After you have added VMs and VM containers to the list, you can specify which objects should be excluded from backup. Veeam Backup & Replication allows excluding the following types of objects: VMs and VM templates from VM containers, as well as specific VM disks.

To select which objects should be excluded, click **Exclusions**.

- To exclude VMs from a VM container (for example, if you need to back up the whole ESX(i) host excluding several VMs), click the **VMs** tab. Click **Add** on the right and select VMs that should be excluded. To facilitate objects selection, you can switch between the **Hosts and Clusters**, **VMs and Templates** and **Datastores and VMs** views, as well as use the search field to find necessary objects by their name.
- To exclude specific VM disks from backup, click the **Disks** tab, select the necessary VM in the list and click **Edit**. If you want to exclude disks of a VM added as part of a container, use the **Add** button to include the VM in the list as a standalone instance.
  You can choose to process all disks, 0:0 disks (typically, the system disks) or select custom disks. If you select the **Remove excluded disks from VM configuration** check box, Veeam Backup & Replication will modify the VMX file to remove excluded disks from VM configuration. If this option is used, you will be able to restore, replicate or copy VM from its backup file to a location where excluded disks are not accessible with the original paths. If you do not use this option, you will have to manually edit the VM configuration file to be able to power on the VM.
- When processing VM containers, Veeam Backup & Replication includes VM templates. To exclude VM templates from backup, open the **Templates** tab and clear the **Backup VM templates** check box. The **Exclude templates from incremental backup** option allows you to include VM templates into full backups only.

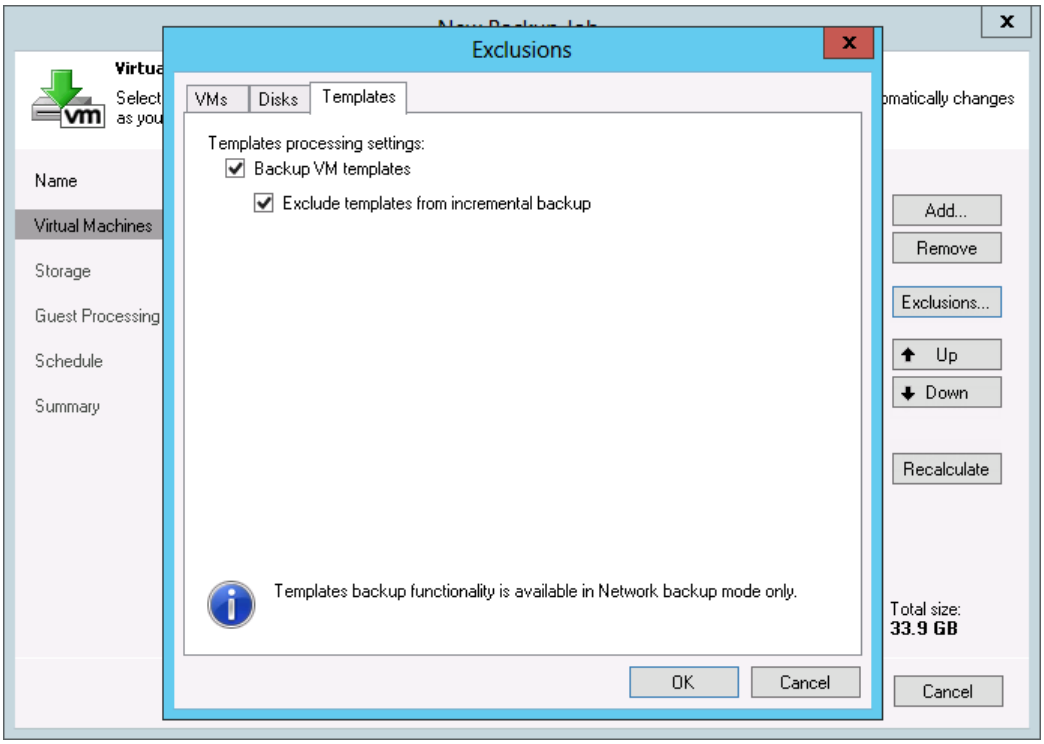

Note Veeam Backup & Replication automatically excludes VM log files from backup to make backup process faster and reduce the size of the backup file.

Step 5. Define VM Backup Order

If you want to back up certain VMs before others, you can define the order in which the backup job must process VMs. VM backup order can be helpful if you want to ensure that backup of a VM does not overlap with other scheduled activities, or that backup is completed before a certain time.

To define VM backup order, select the necessary VMs and move them up or down the list using the **Up** and **Down** buttons on the right. In the same manner, you can set the backup order for containers in the backup list. Note, however, that if you choose to back up a container, VMs inside the container will be processed at random. To ensure that VMs are processed in the defined order, you should add them as standalone VMs, not as part of a container.

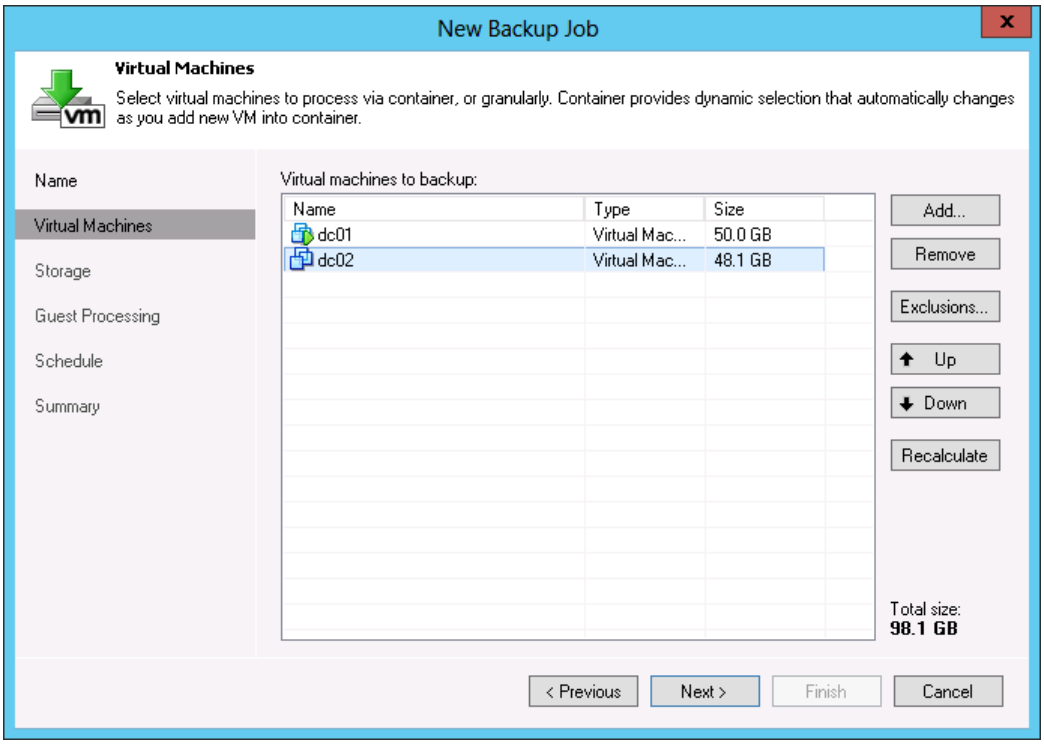

Step 6. Specify Backup Storage Settings

At this step of the wizard, you should select backup infrastructure components (backup proxy and backup repository) and define backup storage settings.

Click **Choose** next to the **Backup proxy** field to select a backup proxy for the job.

- If you choose **Automatic selection**, Veeam Backup & Replication will detect backup proxies that are connected to the source datastore and will automatically assign optimal proxy resources for processing VM data.
  Veeam Backup & Replication assigns resources to VMs included in the backup job one by one. Before processing a new VM in the VM list, Veeam Backup & Replication checks available backup proxies. If more than one proxy is available, Veeam Backup & Replication analyzes transport modes that the proxies can use for data retrieval and the current workload on the proxies to select the most appropriate resource for VM processing.
- If you choose **Use the backup proxy servers specified below**, you can explicitly select proxies that the job can use. It is recommended that you select at least two proxies to ensure that the backup job will be performed should one of the proxies fail or lose its connectivity to the source datastore.

From the **Backup repository** list, select a repository where the created backup should be stored. Make sure you have enough free space on the storage device. When you select a repository, Veeam Backup & Replication checks how much free space is available on the backup repository.

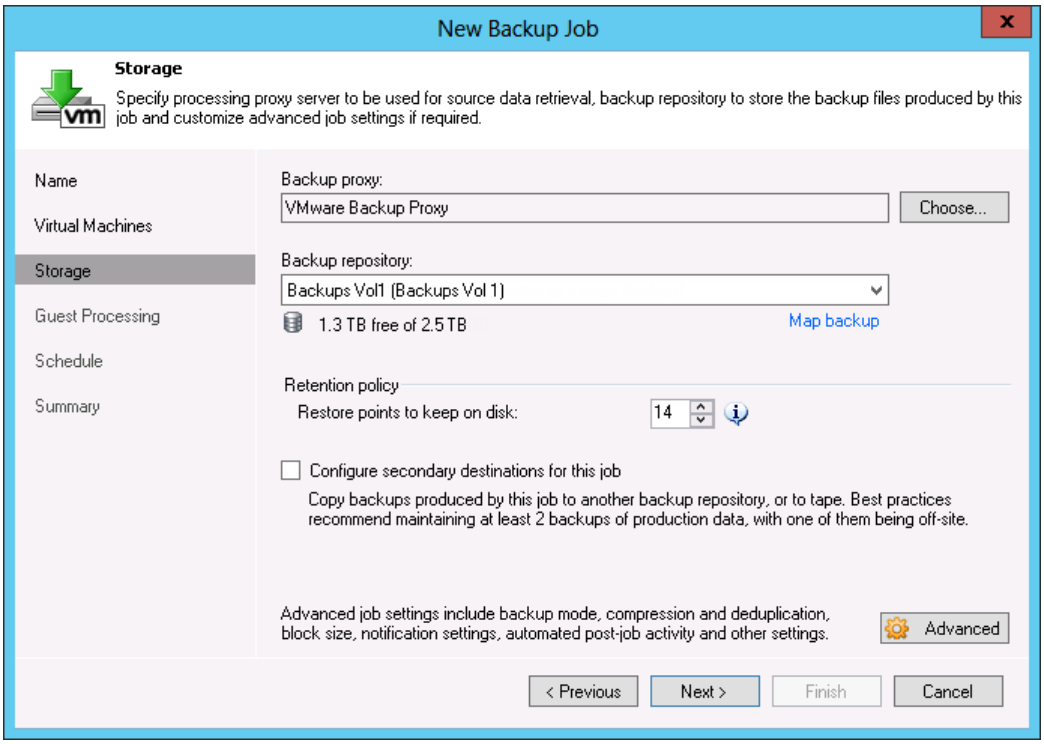

You can map the job to a specific backup stored in a repository. Mapping can be used if you moved backup files to a new repository and you want to point the job to an existing backup on a new repository. Note that before configuring mapping settings, you need to rescan the repository to which you moved backups. Otherwise, Veeam Backup & Replication will not be able to recognize the backups. For details on rescanning repositories, see Managing Backup Repositories.

Mapping can also be helpful if you need to reconfigure an existing backup job or if the Veeam Backup & Replication configuration database is corrupted, so it is necessary to recreate the jobs. To set up job mapping, click the **Map backup** link and point to the necessary backup in the repository. Backups stored in a repository can be easily identified by job names. To facilitate search, you can also use the search field at the bottom of the window.

Note To point a job to an existing backup, make sure the corresponding backup folder includes a complete set of backup files (.vbm, .vbk and .vib/.vrb).
Note that backups created with Veeam Backup & Replication v5 do not include .vbm files, and therefore, cannot be used for mapping "as is". To generate a .vbm file for a v5 backup, you need to perform at least one job pass after upgrading to Veeam Backup & Replication 7.0.

In the **Retention policy** section, specify the number of restore points that should be kept on disk. If this number is exceeded, the earliest restore point will be deleted. The number of restore points is a relative value and doesn't correspond to the number of days to store them. Please keep in mind that such retention policy mechanism works for reversed incremental backup; for incremental backup another mechanism is applied. To learn about the retention policy for incremental backup, see Retention Policy.

If you want to archive the backup file created by the backup job to tape or create a copy of the file in some other location, select the **Configure secondary destination for this job** check box. With this option enabled, the **New Backup Job** wizard will include an additional step, **Secondary Target**. At this step of the wizard, you can link a VM tape backup job or a backup copy job to the backup job you create. To learn more, see Step 8. Specify Secondary Target.

Step 7. Specify Advanced Backup Settings

Click **Advanced** to specify advanced options for the backup job.

Backup settings

Select the method you want to use to back up VMs: **Reversed incremental** or **Incremental**. For details, see Backup Methods.

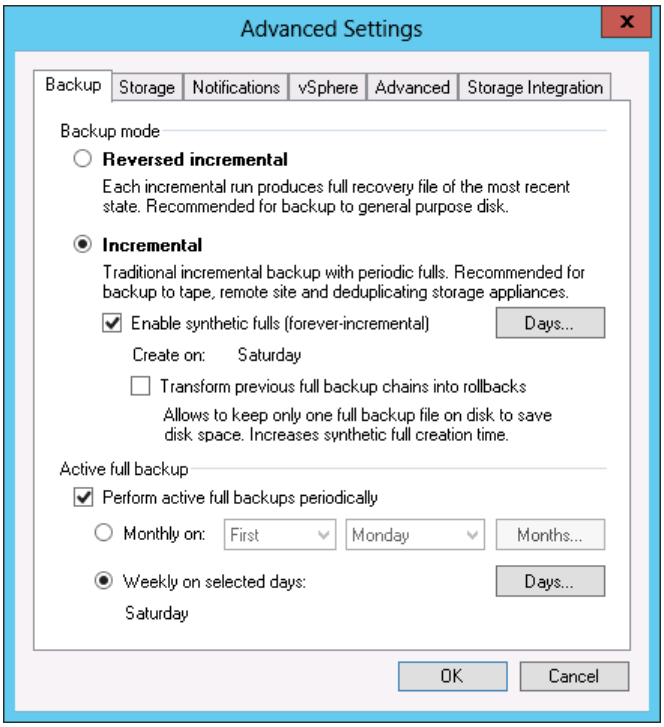

If you choose the incremental backup method, you must select to periodically create a synthetic full backup or perform active full backups.

- To create a synthetic full backup, select the **Enable synthetic fulls (forever-incremental)** check box and click **Days** to schedule synthetic fulls on the necessary days. The created synthetic full backup will be used as a starting point for subsequent increments.
  You can additionally choose to transform the previous full backup chain into the reversed incremental backup chain. To do so, select the **Transform previous full backup chains into rollbacks** check box. To learn more about the transform process, see Transforming Incremental Backup Chains into Reversed Incremental Backup Chains.
- To perform full backups regularly, select the **Perform active full backups periodically** check box and define scheduling settings. The created full backup will be used as a starting point for subsequent increments.

Note If you schedule the active full backup and the synthetic full backup with or without the transform task on the same day, Veeam Backup & Replication will perform only active full backup and skip the synthetic backup and the transform task.

When scheduling periodic full backups, you need to make sure you have enough free space on the backup repository. As an alternative, you can perform active full backup manually. To do so, right-click the ready backup job in the list and select **Active Full**.

Storage settings

On the **Storage** tab, specify deduplication, compression and optimization settings for backup files that the job will provision. For details, see Compression and Deduplication.

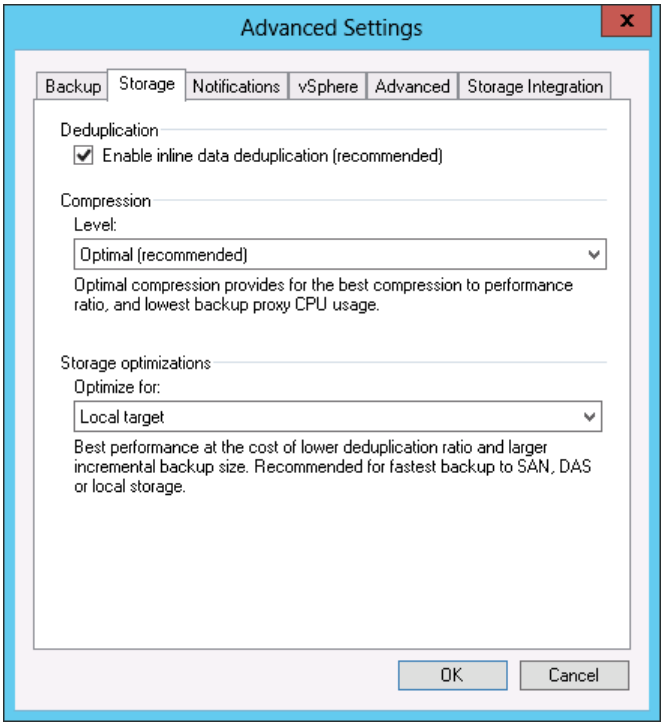

By default, Veeam Backup & Replication performs deduplication before storing VM data to a backup repository. Deduplication provides a smaller size of the resulting backup file but may reduce backup performance.

You can disable deduplication at all by clearing the **Enable inline data deduplication** check box. By disabling this option, you also change the mechanism of incremental backup. If Changed Block Tracking is enabled for the job, Veeam Backup & Replication will save all data block CBT has marked as new to the destination storage, without performing additional check or using Veeam's filtering mechanism. This will result in faster incremental backup. To learn more, see Changed Block Tracking.

In the **Compression** section, specify a compression level for the created backup: *None*, *Dedupe-friendly*, *Optimal*, *High* or *Extreme*.

In the **Storage optimizations** section, select the type of backup target you are planning to use. Depending on the chosen option, Veeam Backup & Replication will use data blocks of different size to optimize the size of backups and job performance:

- **Local target (16 TB + backup size)**. This option provides the lowest deduplication ratio and produces the largest incremental backup file. This option is recommended for backup jobs that can produce very large full backup files — larger than 16 TB.
- **Local target**. This option is recommended if you are planning to use SAN, DAS or local storage as a target. SAN identifies larger blocks of data and therefore can process larger quantities of data at a time. This option provides the fastest backup job performance but reduces the deduplication ratio — the larger a data block is, the lower is the chance to find an identical block.
- **LAN target**. This option is recommended for NAS and on-site replication. It provides a better deduplication ratio and reduces the size of an incremental backup file.
- **WAN target**. This option is recommended if you are planning to use WAN for offsite backup. Veeam Backup & Replication uses small data blocks, which involves significant processing overhead but results in the maximum deduplication ratio and the smallest size of increment files.

Notifications settings

Use the **Notifications** tab if you want to be notified when the backup job is completed.

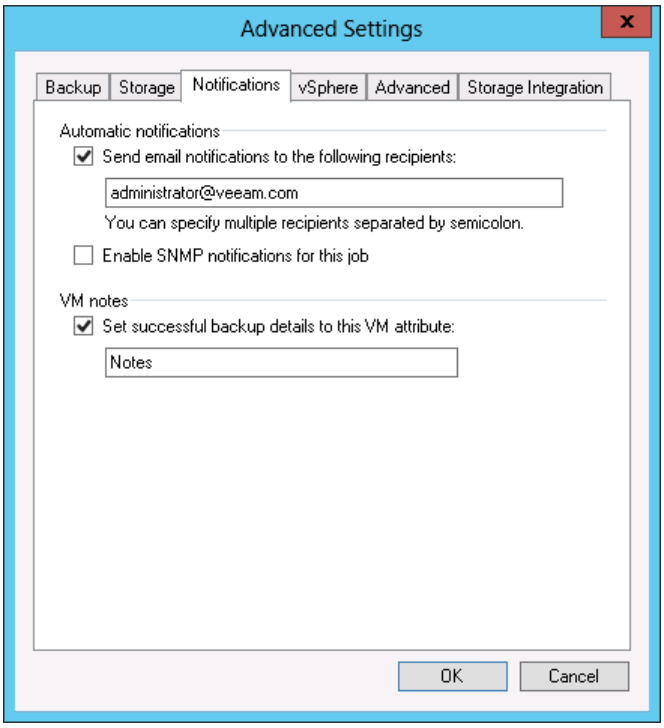

- Select the **Send email notifications to the following recipients** check box if you want to receive notifications by email in case of job failure or success. In the field below, specify a recipient's email address. You can enter several addresses separated by a semicolon. Email notifications will be sent only if you have selected the **Enable email notification** check box in the **Options** window and specified email notification settings (select **Tools** > **Options** from the main menu). For details, see Specifying Notification Settings.
- Select the **Enable SNMP notification for this job** check box if you want to receive SNMP traps when a job is completed successfully. SNMP traps will be sent if you configure SNMP settings in Veeam Backup & Replication and on the recipient's computer. For details, see Specifying SNMP Settings.

- In the **VM notes** section, select the **Set successful backup details to this VM attribute** check box to write to a VM custom attribute information about successfully performed backup and data on backup results (backup date and time, backup console name and path to the backup file). In the field below, enter the name of the necessary attribute. If the specified attribute does not exist, Veeam Backup & Replication will create it.

vSphere settings

On the **vSphere** tab, specify if file system freezing and changed block tracking should be used.

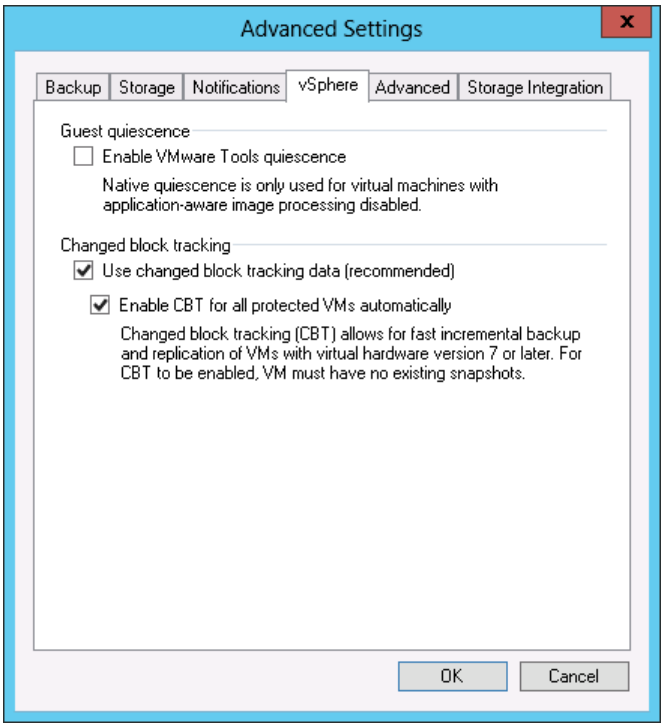

The **Enable VMware tools quiescence** option enables freezing of the file system for proper snapshot creation. With this option enabled, creation of a snapshot is performed with the help of the sync driver responsible for holding incoming I/O and flushing all dirty data to a disk, thus making the file systems consistent.

In the **Changed block tracking** section, specify if vSphere Changed Block Tracking (CBT) should be used. By default, this option is selected. If you want to force using changed block tracking even if CBT is disabled on the ESX(i) host, select the **Enable changed block tracking for all processed VMs** check box.

Important! You can use this option only for VMs using virtual hardware version 7 or later.

Advanced settings

On the **Advanced** tab, specify miscellaneous advanced settings for the job.

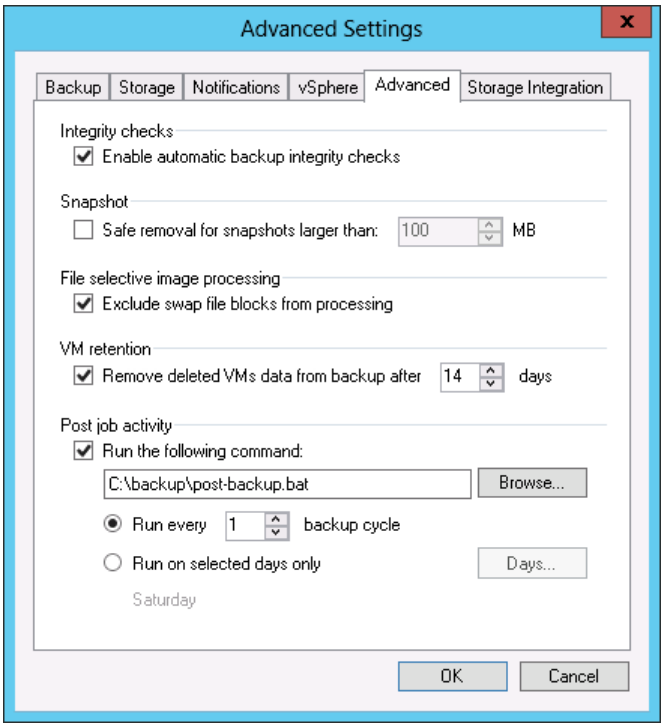

- Select the **Enable automatic backup integrity checks** check box if you want Veeam Backup & Replication to periodically check the full backup file. An automatic backup check allows you to verify integrity of the backup file and avoid a situation when a full backup is corrupted, making all further increments corrupted, too.
  A backup check is performed every time the job is started. During the backup check, Veeam Backup & Replication verifies the service data written to the backup file. If the check fails, Veeam Backup & Replication displays a notification message, prompting you to perform a new active full backup. During such full backup, no integrity check is performed.

Tip The integrity check verifies only service data in the full backup file. To perform a CRC check, you can create a SureBackup recovery verification job and instruct it to validate the verified backup file. To learn more, see Performing Recovery Verification.

- If you are running pre-ESX 3.5 Update 2 hosts, consider enabling the safe snapshot removal option. Because a full image-level backup can take long time depending on the VM size, the VM snapshot can grow very large. When a large snapshot is removed on a VM with heavy disk I/O, a consolidation helper snapshot may grow large too, and will then require a long time to be committed. While a helper snapshot is being committed into VM virtual disk files, the VM remains completely "frozen", and depending on the consolidation helper snapshot size, the freeze time may be so long that some applications running on a VM will time out. To prevent such situation, Veeam Backup & Replication offers a procedure of safe snapshot removal which includes creating an additional snapshot in cases when the "main" snapshot size is above the specified threshold. An additional snapshot is used to host writes while the "main" snapshot is being deleted. This ensures that a consolidation helper snapshot does not grow large. To use this option, select the **Safe removal for snapshots larger than ... MB** check box and specify a threshold for the size of a snapshot that should not be exceeded.

- In the **File selective image processing** section, define whether you want to exclude blocks of Windows page files from the backup. During backup, Veeam Backup & Replication checks the NTFS MFT file on Windows-based VMs to identify blocks of the Windows pagefile, and excludes these blocks from processing. Windows page files are dynamic in their nature and change intensively between backup job runs, even if VMs do not change much. Therefore, page file processing results in reduced backup performance and increased size of backup increments.
  Clear the **Exclude swap file blocks from processing** check box if Windows page files should be processed during backup.
- In the **VM retention** section, specify the number of days to keep backup data for deleted VMs. If a VM is no longer available (for example, it was deleted or moved to another location), Veeam Backup & Replication will keep its data in the backup for the period you specify in the **Remove deleted VMs data from backup after** field. When this retention period is over, data of the deleted VM will be removed from backup files. To learn more, see Retention Policy for Deleted VMs.
  The retention period for deleted VMs is particularly useful if the job is configured to create synthetic full backups, and you want to make sure that the full backup does not include redundant data.
- Select the **Run the following command** check box if you want to execute post-backup actions (for example, to launch a script recording the resulting backup file to tape). Use the **Browse** button to select an executable file.
  You can select to execute post-backup actions after a number of backup cycles or on specific week days. If you select the **Run every... backup cycle** option, specify the number of the backup job run after which the file should be executed. If you select the **Run on selected days only** option, click the **Days** button and specify week days when actions should be performed.

Storage integration settings

On the **Storage Integration** tab, define whether you want to use the Backup from Storage Snapshots technology or not. Backup from Storage Snapshots lets you leverage HP SAN snapshots for VM data processing. The technology dramatically improves RPOs and reduces impact of backup activities on the production environment.

By default, the **Use storage snapshots** option is enabled. If you do not want to use the Backup from Storage Snapshots technology, clear the check box. To learn more, see Performing Backup from Storage Snapshots.

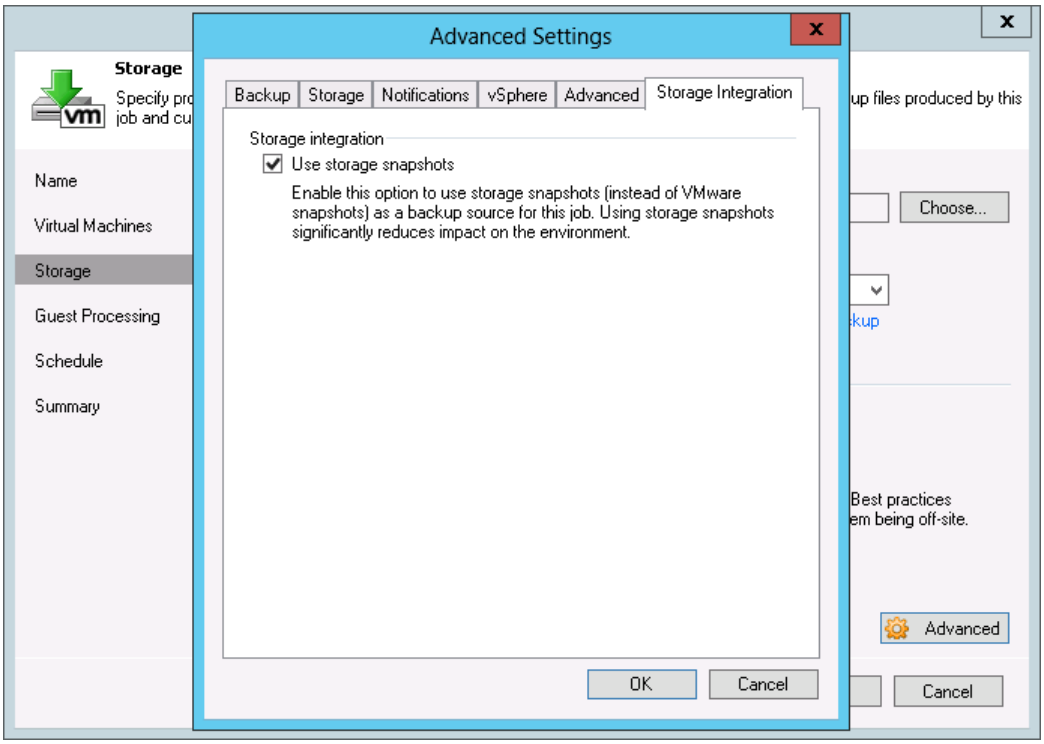

Step 8. Specify Secondary Target

The **Secondary Target** step is available if you have selected the **Configure secondary destination for this job** check box at the **Storage** step of the wizard. Using this step of the wizard, you can link a VM tape backup job or a backup copy job to the created backup job. As a result, the backup file created by the backup job you configure will be automatically archived to tape or copied to another location.

To link a job, click **Add** on the right and select the necessary job from the list. Note that the VM tape backup job or the backup copy job you want to link must be pre-configured on the Veeam backup server.

To learn more, see Linking Backup Jobs to Backup Copy Jobs and Linking Backup Jobs to Backup to Tape Jobs.

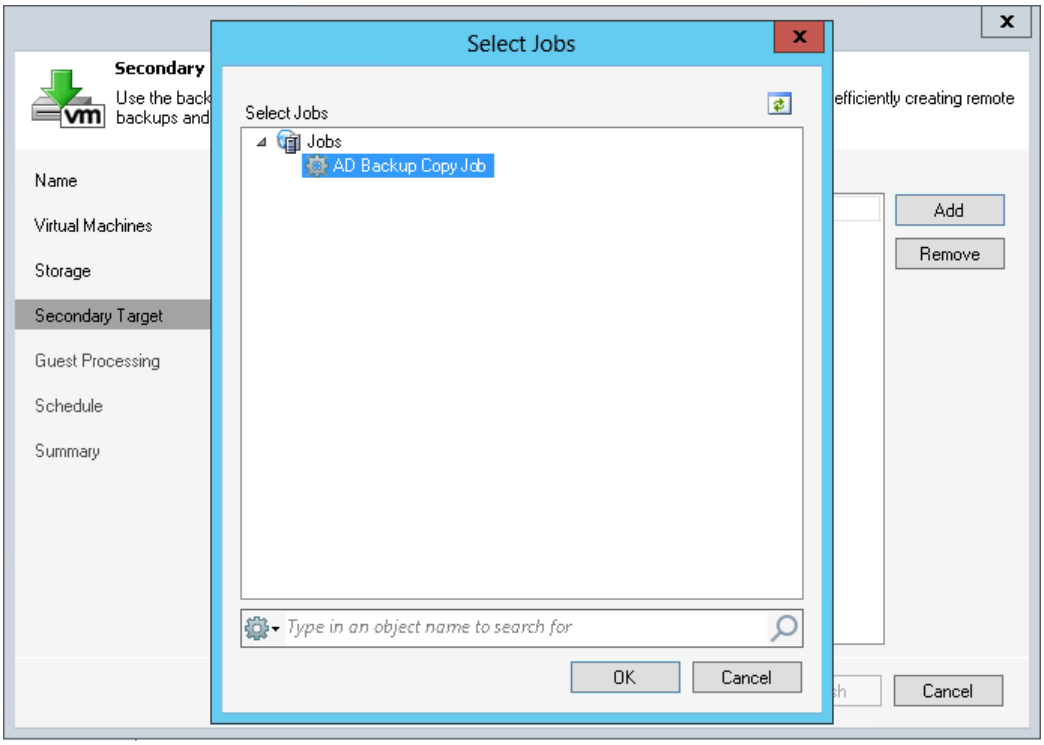

Step 9. Enable Application-Aware Image Processing and Indexing

At the **Guest Processing** step of the wizard, you can enable guest file indexing and select to create a transactionally consistent backup.

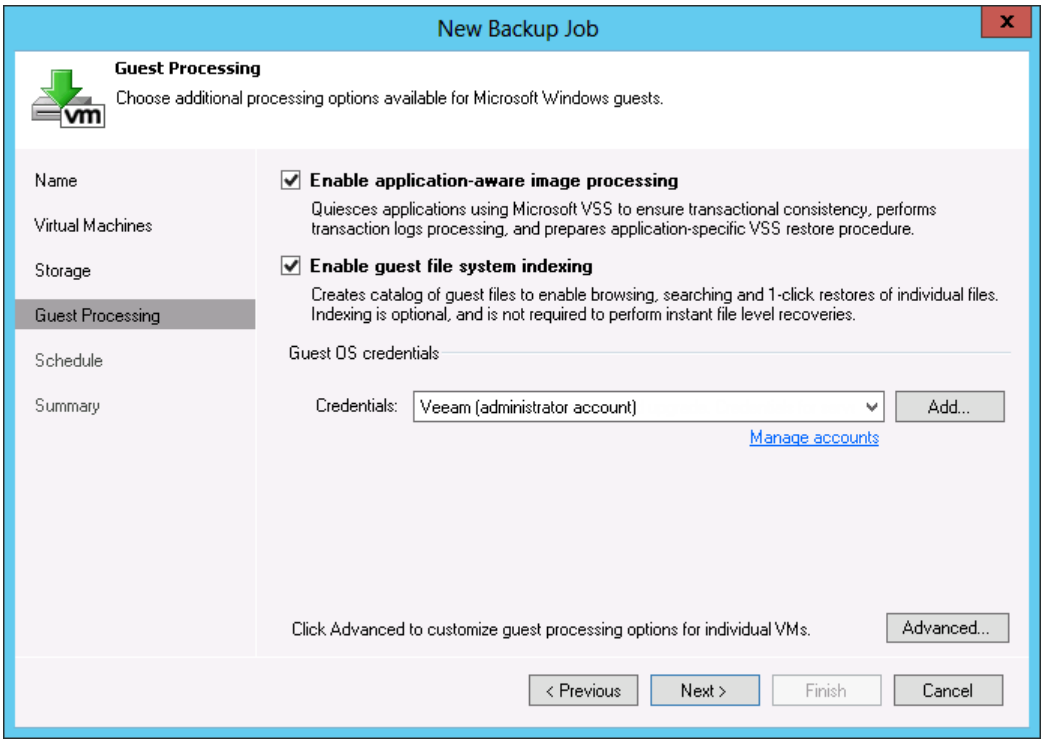

- If you want to create a transactionally consistent backup ensuring successful recovery of VM applications without any data loss, select the Enable application-aware image processing check box. For details, see Transaction Consistency.
- If you want to index guest files in a VM you back up, select the Enable guest file system indexing check box. Veeam Backup & Replication will perform file indexing and enable you to perform fast and accurate search for VM guest OS files via the Veeam Backup Enterprise Manager web UI, and restore these files with a single click.

To coordinate proper indexing and VSS activities, Veeam Backup & Replication injects a runtime process inside the VM. The process is run only during VSS quiescence and indexing procedures and is stopped immediately after processing is finished (depending on the selected option, during the backup job or after it is finished), thus producing low impact on VM performance and stability. In the **Guest OS credentials** section, specify an account with local administrative privileges for injecting the process. If you have not set up the necessary credentials beforehand, click the **Manage accounts** link at the bottom of the list or click **Add** on the right to add the necessary credentials. To learn more, see Managing Credentials.

The user name in credentials must be supplied in the *DOMAIN\USERNAME* format. The guest OS credentials you provide will be used for all VMs included in the backup job.

Click **Advanced** to specify advanced options for Veeam VSS and indexing processing. The **Advanced Options** section contains a list of VMs that will be processed with Veeam VSS and indexing tools.

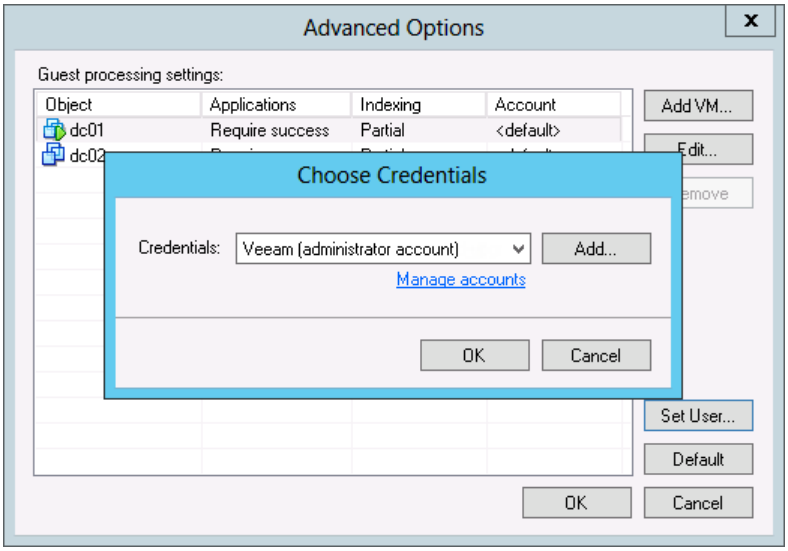

By default, for all VMs in the list Veeam Backup & Replication uses common credentials you provided in the **Guest OS credentials** section. If a different account should be used to inject the process into a specific VM, select the VM in the list, click **Set User** and enter custom guest OS credentials. To discard custom credentials for a VM, select it in the list and click **Default**.

If you want to define custom settings for a VM added as part of a container, include the VM in the list as a standalone instance. To do so, click **Add VM** and choose a VM whose settings you want to customize. Next, select the VM in the list and define the necessary custom settings. To discard custom settings of a VM, select the VM in the list and click **Remove**.

To provide granular quiescence and indexing options for a VM, select it in the list and click **Edit**.

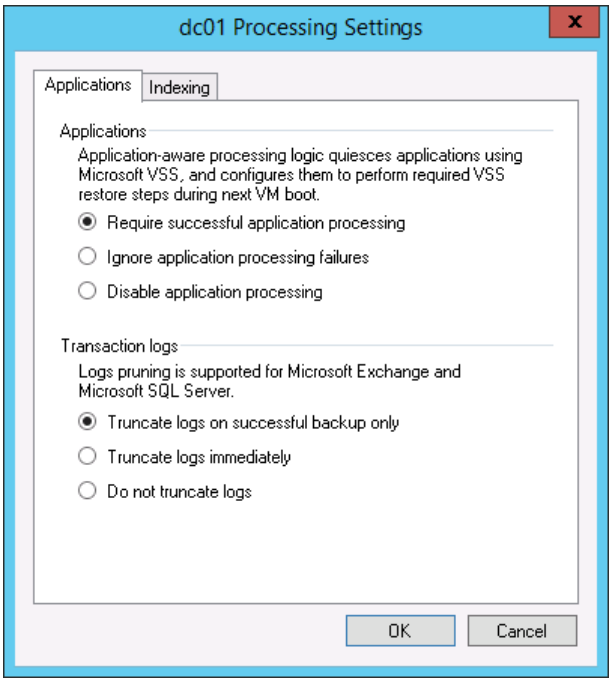

In the **Applications** section on the **Applications** tab, specify the VSS behavior scenario:

- Select **Require successful application processing** if you want Veeam Backup & Replication to stop the backup of the VM if any VSS errors occur.
- Select **Ignore application processing failures** if you want to continue backing up the VM even if VSS errors occur. This option is recommended to guarantee completion of the job. The created backup image will not be transactionally consistent, but rather crash consistent.
- Select **Disable application processing** if you do not want to enable quiescencing for the VM at all.

Use the **Truncation logs** section to define the scenario of transaction log handing:

- Select **Truncate logs on successful backup only** if you want Veeam Backup & Replication to trigger truncation of transaction logs only after the job is finished successfully. In this case, the Veeam runtime process will wait for the backup to complete, and then it will trigger truncation of transaction logs. If truncation of transaction logs is not possible for some reason, the logs will remain untouched in the VM guest OS until the next start of the Veeam runtime process.
- Select **Truncate logs immediately** if you want Veeam Backup & Replication to trigger truncation of transaction logs in any case, no matter whether the job finishes successfully or fails.
- Select **Do not truncate logs** if you do not want Veeam Backup & Replication to truncate logs at all. This option is recommended if you are using another backup tool to perform guest-level backup, and this tool maintains consistency of the database state. In such scenario, truncation of logs with Veeam Backup & Replication will break the guest-level backup chain and cause it to fall out of sync.

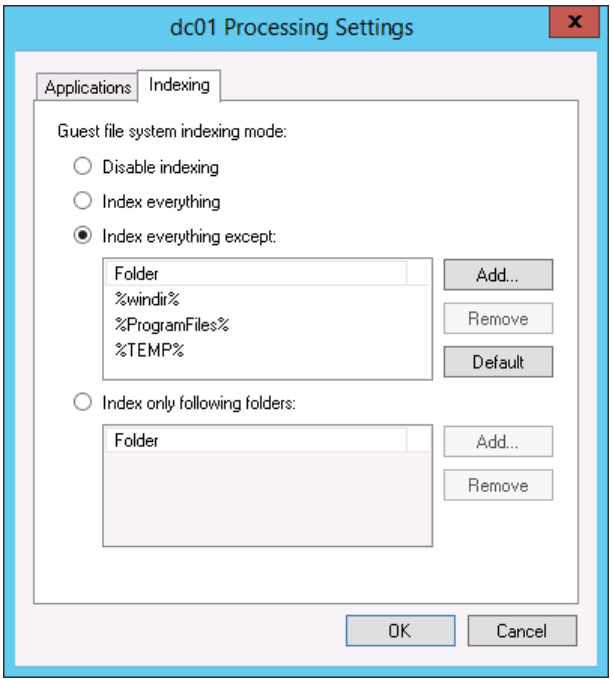

Click the **Indexing** tab to specify indexing options for the VM. Please keep in mind that file indexing is supported for Windows-based VMs only.

- Select **Disable indexing** if you do not want to index guest OS files of the VM.
- Select **Index everything** if you want to index all guest OS files inside the VM.
- Select **Index everything except** if you want to index all guest OS files except those defined in the list. By default, system folders are excluded from indexing. You can add or delete folders to exclude using the **Add** and **Remove** buttons on the right. You can use any system environment variables, for example: `%windir%`, `%Program Files%` and `%Temp%`.
- Use **Index only following folders** to select specific folders that you want to index. To form the list of folders, use the **Add** and **Remove** buttons.

Step 10. Define the Job Schedule

At the **Schedule** step of the wizard, you can select to only run the backup job manually, schedule the job to start at a specific time– (for example, the least busy hours to reduce the impact on the virtual environment) or define a schedule for the job to run on a regular basis.

To specify the job schedule, select the **Run the job automatically** check box. If this check box is not selected, the job is supposed to be run manually.

You can define the following scheduling settings for the job:

- You can choose to run the job at specific time on defined week days, monthly and with specific periodicity.
- You can choose to run the job continuously. In this case, the next run of the backup job will be started as soon as the previous one is completed, maintaining your backup always in the most recent state. To run the job continuously, select the **Periodically every** option and choose **Continuously** from the list on the right.
- You can choose to run the job repeatedly throughout a day with a set time interval. At the **Schedule** step of the wizard, select the **Periodically every** option, enter the necessary time interval and select the necessary time unit: *Hours* or *Minutes*. Click **Schedule** on the right and use the time table to define the permitted time window for the job. If you choose to run the job at an hourly interval, in the **Start time for hourly jobs** field, specify the exact time when the job should start.

For example, you want to start a job every 2 hours from 9 AM to 6 PM. At the **Schedule** step, select the **Periodically every** option, enter 2 in the field on the right and select *Hours* from the list. Click **Schedule** and use the **Permitted** and **Denied** options to mark the time window from 9 AM to 6 PM. In the **Start time for hourly jobs** field, specify the exact start time of the job, for example, 15 minutes. The job you have scheduled will start at 9:15 AM, 11:15 AM, 1:15 PM, 3:15 PM and 5:15 PM.

- You can chain jobs. In the common practice, jobs start one after another: when the job "A" finishes, the job "B" starts and so on. If you want to create a chain of jobs, you should define the time schedule for the first job in the chain. For the rest of the jobs in the chain, at the **Schedule** step of the wizard, select the **After this job** option and choose the preceding job from the list.

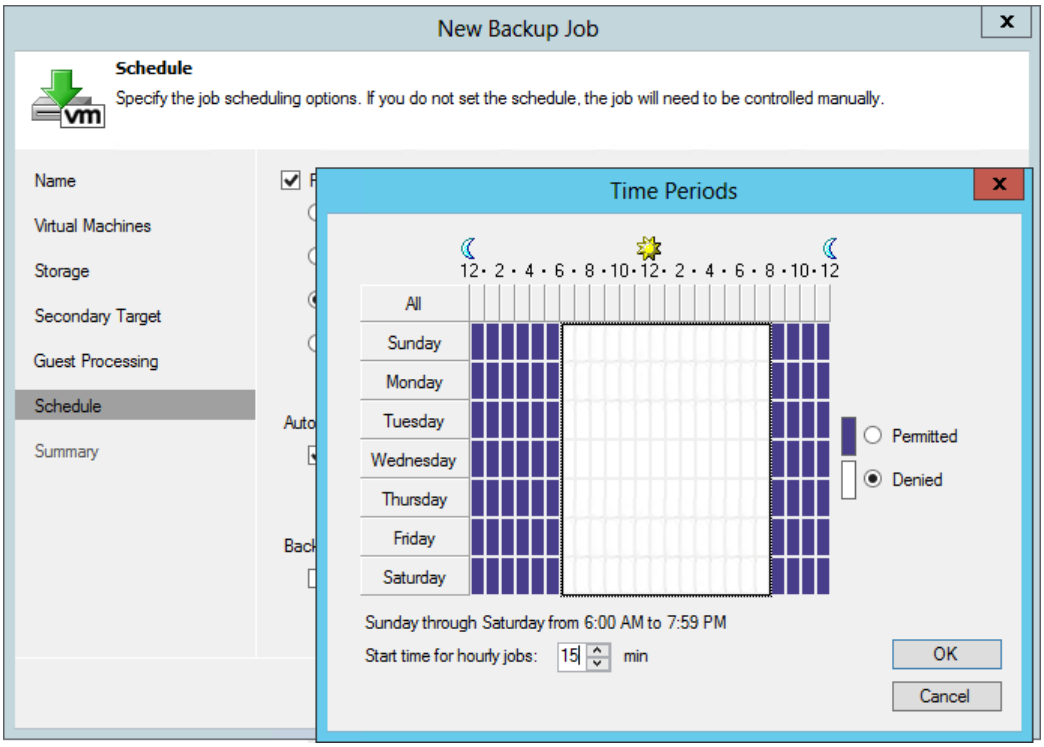

In the **Automatic retry** section, define whether Veeam Backup & Replication should attempt to run the backup job again in case it fails for some reason. A repeatedly run job will include failed VMs only. Enter the number of attempts to run the job and define time spans between them. If you select continuous backup, Veeam Backup & Replication will retry the job for the defined number of times without any time intervals between the job runs.

In the **Backup window** section, determine a time span within which the backup job must be completed. The backup window prevents the job from overlapping with production hours and ensures it does not provide unwanted overhead on your virtual environment. To set up a backup window for the job, select the **Terminate job if it exceeds allowed backup window** check box and click **Window**. In the **Time Periods** section, define the allowed window and prohibited hours for backup. If the job exceeds the allowed window, it will be terminated.

Note After you have created a scheduled job, you can temporarily disable it (that is, hold it for some time without changing the set time schedule). Right-click a job in the list and select **Disable Job** from the shortcut menu. To enable the job schedule, right-click the job and deselect **Disable Job** in the shortcut menu.

Step 11. Finish Working with the Wizard

After you have specified schedule settings, click **Create**. Select the **Run the job when I click Finish** check box if you want to start the created job right after you complete working with the wizard. Click **Finish** to close the wizard.

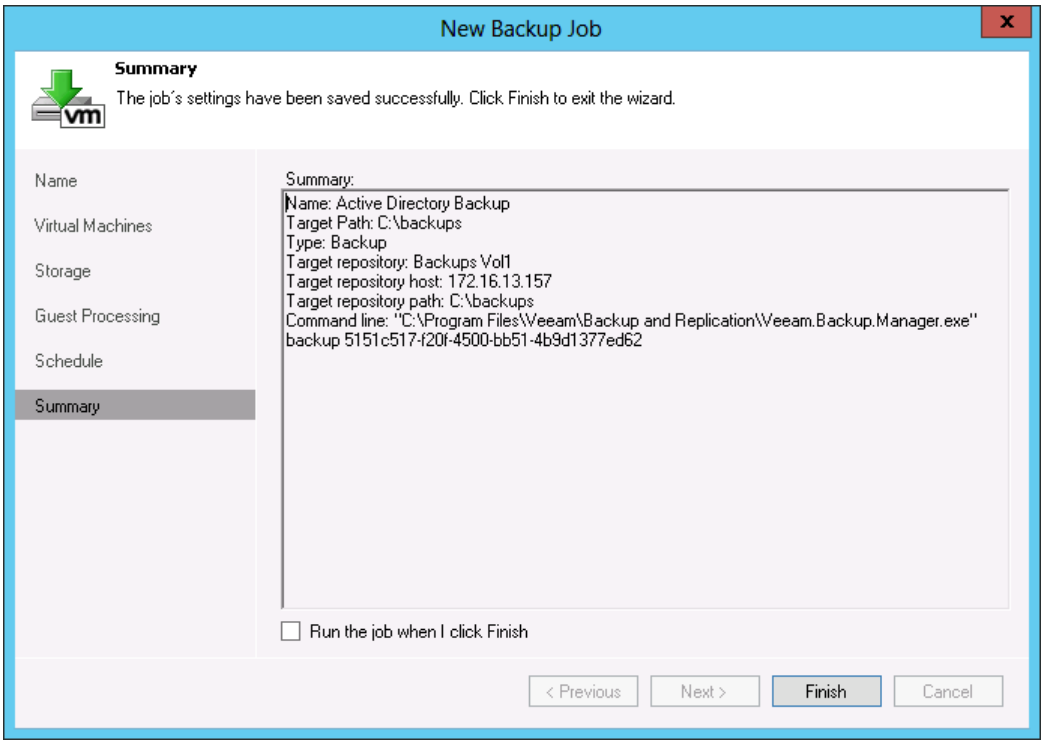

Creating Replication Jobs

To replicate VMs, you should create a replication job by means of the **New Replication Job** wizard. You can perform the created job immediately, schedule or save it. This section will guide you through all steps of the wizard and provide explanation on available options.

Step 1. Launch the New Replication Job Wizard

To run the **New Replication Job** wizard, do one of the following:

- On the **Home** tab, click **Replication Job** and select **VMware**.
- Open the **Backup & Replication** view, right-click the **Jobs** node and select **Replication** > **VMware**.
- Open the **Virtual Machines** view, select one or several VMs in the working area, click **Add to Replication** on the ribbon and select **New job**. Alternatively, you can right-click one or several VMs in the working area and select **Add to Replication Job** > **New job**. In this case, the selected VMs will be automatically included into the backup job. You can add other VMs to the job when passing through the wizard steps.

You can quickly include VMs to already existing jobs. To do that, open the **Virtual Machines** view, right-click necessary VMs in the working area and select **Add to Replication Job** > *name of a created job*.

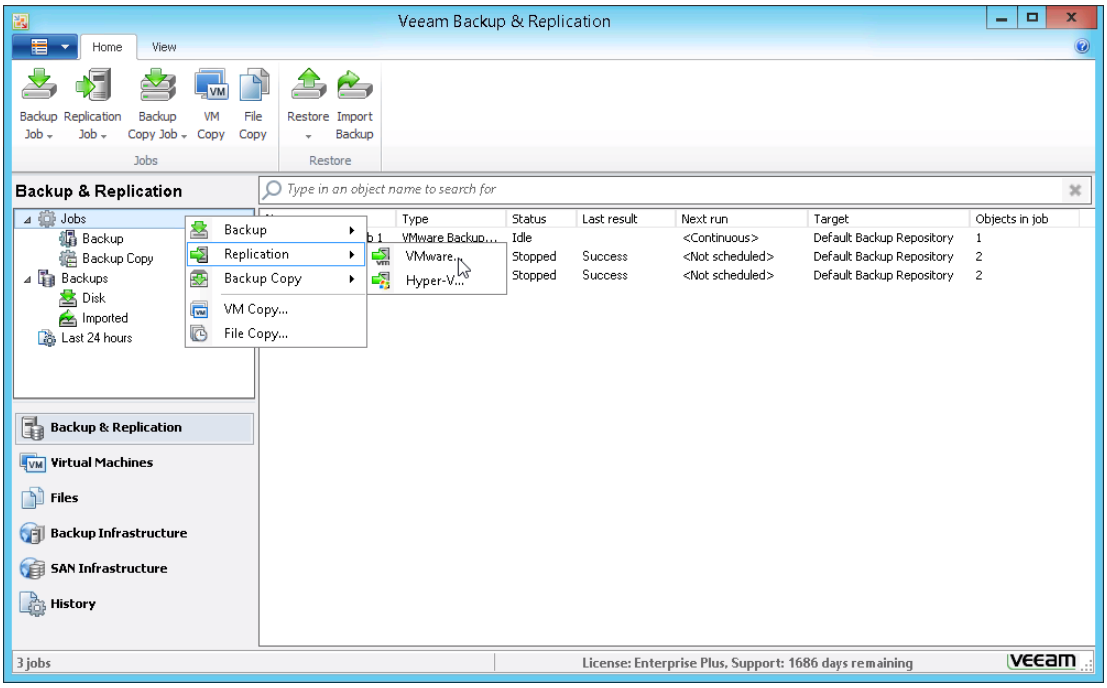

Step 2. Specify Job Name and Description

At the first step of the wizard, enter a name and description for the job. The default description contains information about the user who created the job, as well as the date and time when the job was created.

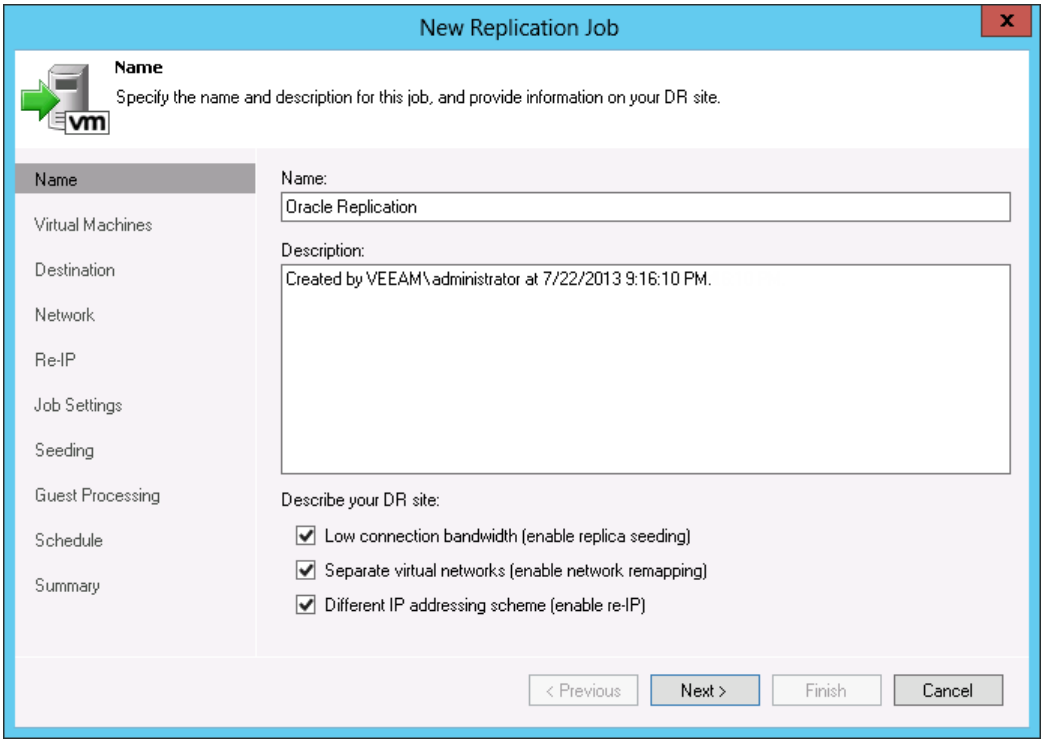

If you plan to replicate to a DR site, you can use a number of advanced configuration settings for the job:

- Select the **Low connection bandwidth** check box to enable the Seeding step in the wizard. Replica seeding can be used if you plan to replicate to a remote site and want to reduce the amount of traffic sent over network during the first run of the replication job.
- Select the **Separate virtual networks** check box to enable the Network step in the wizard. If the network in the DR site does not match the production network, you can resolve this mismatch by making up a network mapping table.
- Select the **Different IP addressing scheme** check box to enable the Re-IP step in the wizard. Re-IP possibilities can be used to automate reconfiguration of replica IP addresses for Windows-based VMs in case IP schemes in the DR and production sites do not match.

Step 3. Select Virtual Machines to Replicate

At this step, you should select an individual VM, multiple VMs or VM containers which you want to replicate. Jobs with VM containers are dynamic in their nature: if a new VM is added to the container after a replication job is created, the job will be automatically updated to include the added VM.

Note If you choose to replicate container objects, Veeam Backup & Replication will automatically exclude any existing VM templates from containers.

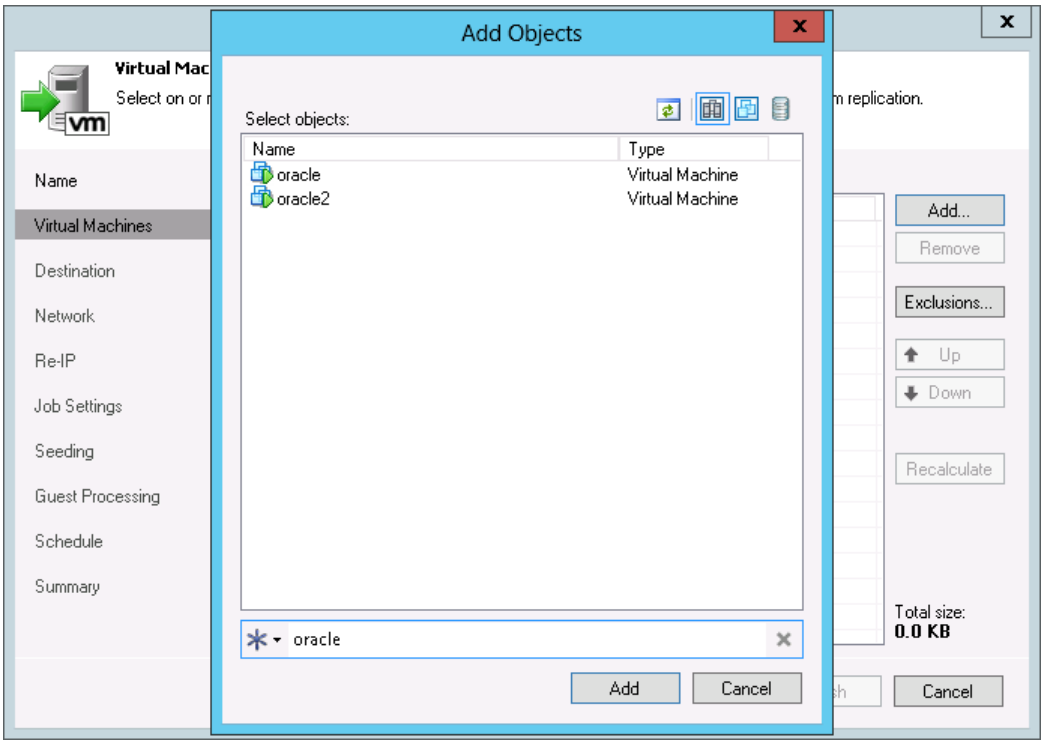

Click **Add** to browse to VMs and VM containers that should be replicated. In the displayed tree, select the necessary object and click **Add**.

To facilitate objects selection, you can switch between views by clicking **Hosts and Clusters**, **VMs and Templates** or **Datastores and VMs** at the top of the tree. In addition, you can use the search field at the bottom of the **Add Objects** window: click the button to the left of the field and select the necessary type of object to search for (**Everything**, **Folder**, **Cluster**, **Host** or **Virtual machine**), enter the object's name or a part of it and click the **Start search** button on the right or press **[ENTER]**.

Note Depending on the view you select, some objects may be not available: for example, if you select the **VMs and Templates** view, you will not be able to see and find resource pools.

To remove an object from the list, select it and click **Remove** on the right.

The initial size of VMs and VM containers added to a replication job is displayed in the **Size** column in the list. The total size of objects is displayed in the **Total size** field. Use the **Recalculate** button to refresh the total size value after you add a new object to the job.

Step 4. Exclude Objects from the Replication Job

After you have added VMs and VM containers to the list, you can specify which objects should be excluded from replication. Veeam Backup & Replication allows excluding the following types of objects: VMs from VM containers, as well as specific VM disks.

To select which objects should be excluded, click **Exclusions**.

- To exclude VMs from a VM container (for example, if you need to replicate the whole ESX(i) host excluding several VMs), click the **VMs** tab. Click **Add** on the right and select VMs that should be excluded. To facilitate objects selection, you can switch between the **Hosts and Clusters**, **VMs and Templates** and **Datastores and VMs** views, as well as use the search field to find necessary objects by their name.
- To exclude specific VM disks from replication, click the **Disks** tab, select the necessary VM in the list and click **Edit**. If you want to exclude disks of a VM added as part of a container, use the **Add** button to include the VM in the list as a standalone instance.
  You can choose to process all disks, 0:0 disks (typically, the system disks) or select custom disks. If you select the **Remove excluded disks from VM configuration** check box, Veeam Backup & Replication will modify the VMX file to remove excluded disks from VM configuration. If this option is used, you will be able to restore, replicate or copy the VM to a location where excluded disks are not accessible with the original paths. If you do not use this option, you will have to manually edit the VM configuration file to be able to power on the VM.

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

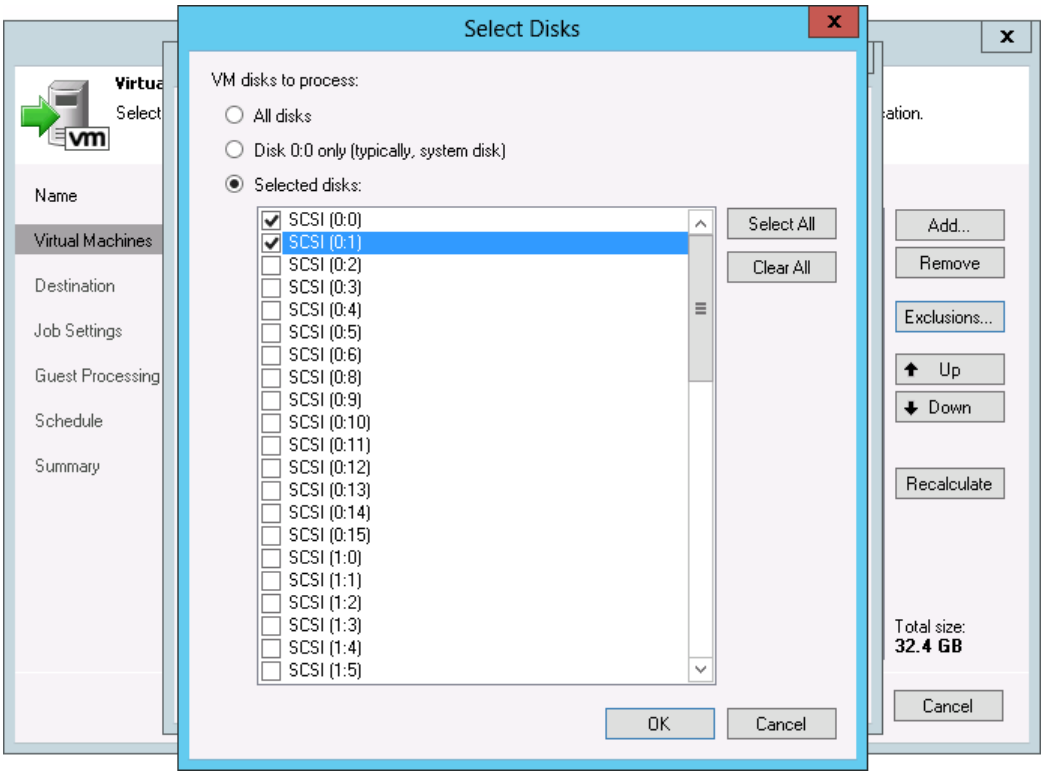

Step 5. Define VM Replication Order

If you want to replicate certain VMs before others, you can define the order in which the replication job must process VMs. VM replication order can be helpful if you want to ensure that replication of a VM does not overlap with other scheduled activities, or that replication is completed before a certain time.

To set VM replication order, select the necessary VMs and move them up or down the list using the **Up** and **Down** buttons on the right. In the same manner, you can set the replication order for containers in the backup list. Note, however, that if you choose to replicate a container, VMs inside a container will be processed at random. To ensure that VMs are processed in the defined order, you should add them as standalone VMs, not as part of a container.

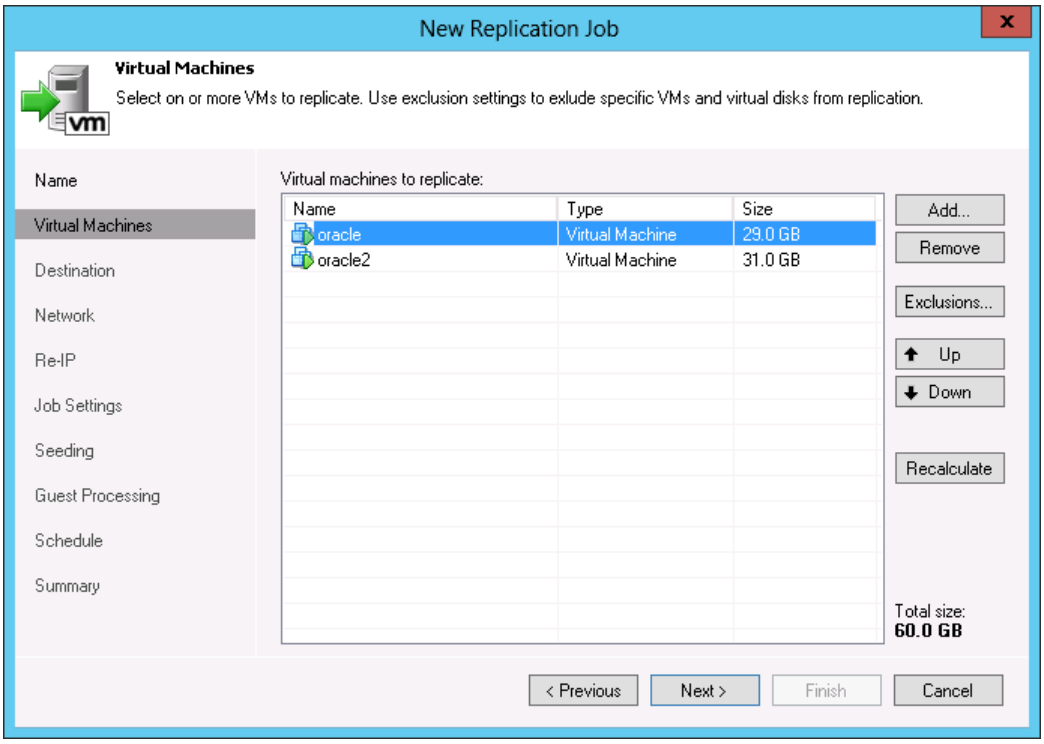

Step 6. Specify Replica Destination

At this step of the wizard, you should select destination for created replicas.

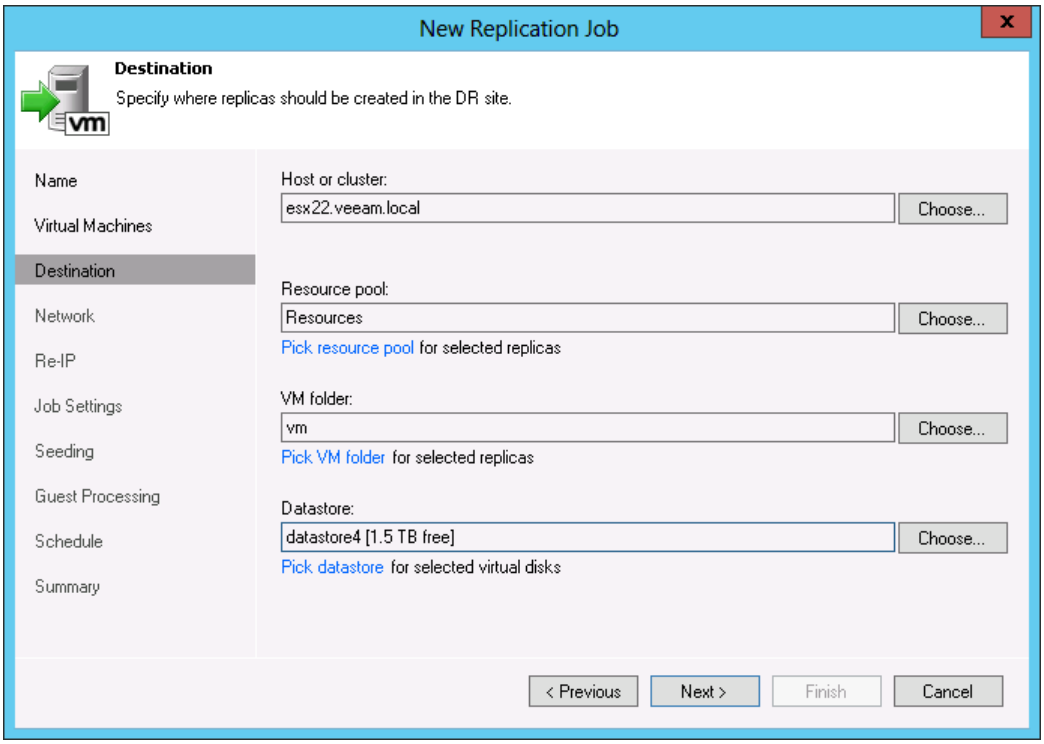

Host or cluster

Click **Choose** and select the host or cluster where replicas will be registered. Assigning a cluster as a target ensures more sustainable replication process — the replication job will be performed until there is at least one available host in the cluster.

To facilitate selection, use the search field at the bottom of the window: click the button on the left of the field to select the necessary type of object that should be searched for (**Cluster** or **Host**), enter an object's name or a part of it and click the **Start search** button on the right or press **[ENTER]**.

Resource pool

Specify a resource pool in which replicas will be included. If all or the majority of replicas should belong to the same resource pool, click **Choose** and select the destination resource pool. To facilitate selection, use the search field at the bottom of the window: enter the resource pool name or a part of it and press **ENTER**.

If you want to map replicas to different resource pools, click the **Pick resource pool for selected replicas** link. In the **Choose resource pool** section, click **Add VM** on the right and select the necessary VMs. To map the added VM to a resource pool, select it in the **Replica VM resource pool** list and click **Resource Pool** at the bottom of the window. From the list of available resource pools, choose the necessary one.

VM folder

Specify a folder to which replicas will be placed. If all or the majority of replicas should be stored in the same folder, click **Choose** and select the necessary folder. To facilitate selection, use the search field at the bottom of the window: enter the folder name or a part of it and click the **Start search** button on the right or press **[ENTER]**.

If you want to map replicas to different VM folders, click the **Pick VM folder for selected replicas** link. In the **Choose folder** section, click **Add VM** on the right and select the necessary VMs. To map an added VM to a VM folder, select it in the **Replica VM folder** list and click **VM Folder** at the bottom of the window. From the list of available folders, select the necessary one.

Datastore

Specify a datastore where VM replica files will be stored. If all or the majority of replicas will be stored on the same datastore, click **Choose** and select the necessary datastore. Note that Veeam Backup & Replication will display only those datastores that are accessible by the selected replication target. If you have chosen to replicate VMs to a cluster, Veeam Backup & Replication will display only shared datastores.

To facilitate selection, use the search field at the bottom of the window: enter the datastore name or a part of it and press **ENTER**.

If you want to place replicas to different datastores, click the **Pick datastore for selected virtual disks** link. In the **Choose VM Files Location** section, click **Add VM** on the right and select VMs that should be mapped to datastores. To map an added VM to a datastore, select it in the **Files location** list and click **Datastore** at the bottom of the window. From the list of available datastores, select the necessary one.

Additionally, you can choose to store replica configuration files and disk files in different locations. To do so, add a VM to the **Files location** list, expand it and select the required type of files. Click **Datastore** at the bottom of the window and choose the destination for the selected type of files.

By default, Veeam Backup & Replication preserves the format of replicated VM disks. If necessary, you can configure the job to change the disk format. For example, if the original VM uses thick disks, you can change the format of replica disks to thin provisioned, and thus, save on disk space required to store replica data. To change replica disk format, add the VM to the **Files location** list, select it and click **Disk type** at the bottom of the window. In the **Disk Type Settings** section, choose the format that will be used to restore replica disk files: same as the source disk, thin or thick. Please note that disk format change is available only for VMs using virtual hardware version 7 or later.

Step 7. Create a Network Map Table

This step is available if you have selected the **Separate virtual networks** option at the Name step of the wizard.

At this step, you should create a table that maps production (source) networks to DR site (target) networks. To add a network mapping entry, click **Add**. In the **Source network** field, define the production network where original VMs reside. In the **Target network** field, define the name of the network that will be substituted for the source network in the DR site.

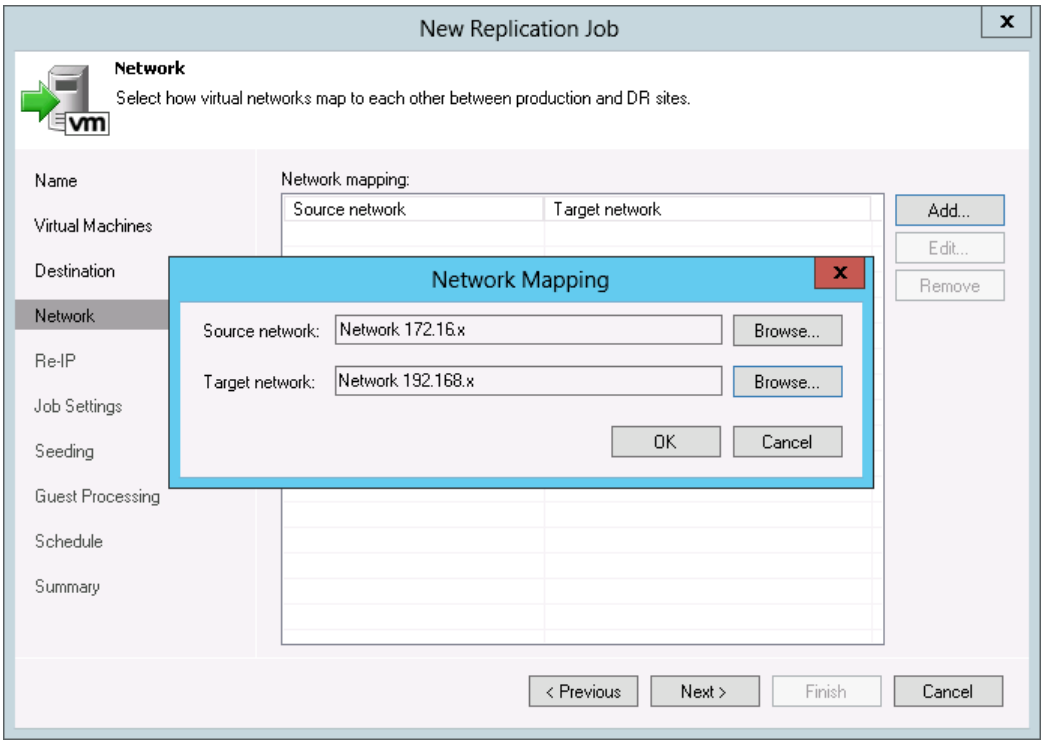

Veeam Backup & Replication checks the network mapping table during every job run and updates replicas' configuration files in accordance with the mapping table.

Step 8. Configure Re-IP Rules

This step is available if you have selected the **Different IP addressing scheme** option at the Name step of the wizard.

At this step, you should set up rules according to which Veeam Backup & Replication will reconfigure IP addresses of Windows-based replicas when you perform failover. To add a re-IP rule, click **Add**.

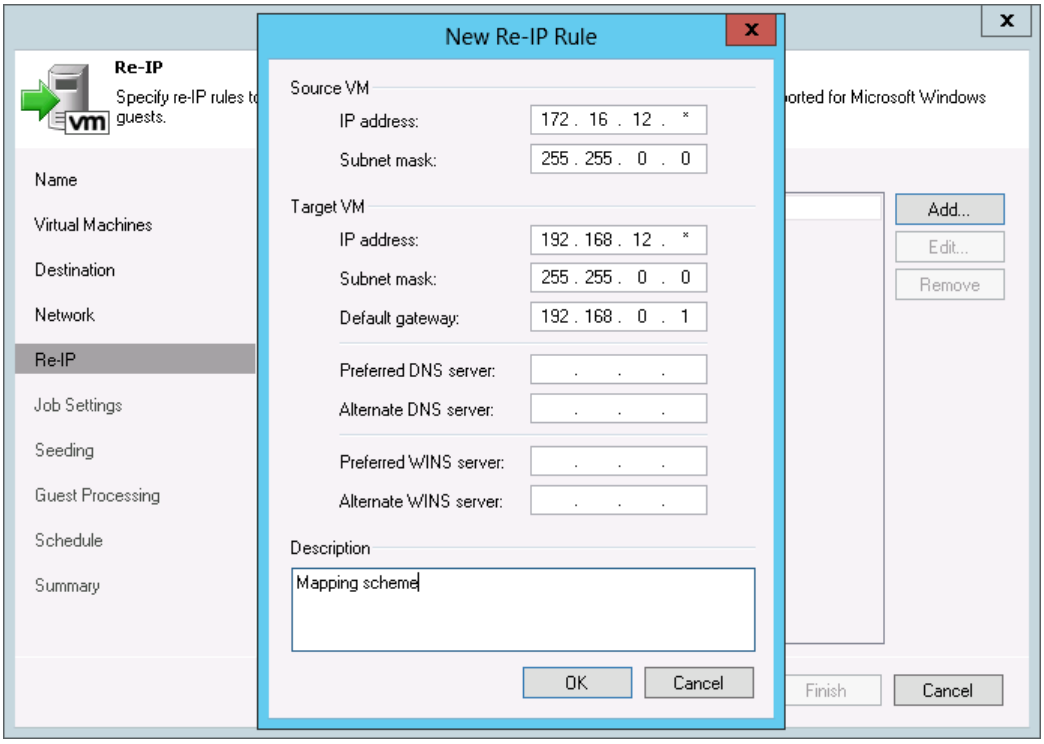

- In the **Source VM** section, describe an IP numbering scheme adopted in source site. To facilitate configuration, Veeam Backup & Replication detects an IP address and subnet mask for the machine where it is installed, and pre-populates these values.
- In the **Target VM** section, describe an IP numbering scheme adopted in the DR site – an IP address, subnet mask and default gateway that will be used for VM replicas. If necessary, define the DNS and WINS server addresses.
- In the **Description** field, enter a brief outline of the rule or any related comments.

Note To specify a range of IP addresses, use the asterisk character (*), for example: 172.16.17.* For a range of IP addresses 172.16.17.0-172.16.17.255.
Do not use 0 to specify a range of IP addresses. In Veeam Backup & Replication, value 172.16.17.0 will define a regular IP address 172.16.17.0, not an IP address range.

Step 9. Specify Replication Job Settings

At this step of the wizard, you should assign replication infrastructure components for the job and define replication job settings.

In the **Data transfer** section, select backup proxies that will be used to transfer VM data from source to target, and a repository to store replica metadata.

If you plan to replicate within one site, the same server can act as the source proxy and target proxy. For offsite replication, you must deploy at least one proxy server in each site to establish a stable connection for data transfer across sites. Click **Choose** next to the **Source proxy** and **Target proxy** fields to select backup proxies for the job. In the **Backup Proxy** section, you can choose automatic proxy selection or assign proxies explicitly.

- If you choose **Automatic selection**, Veeam Backup & Replication will detect backup proxies that are connected to the source and target datastores and will automatically assign optimal proxy resources for processing VM data.
  Veeam Backup & Replication assigns resources to VMs included in the replication job one by one. Before processing a new VM from the list, Veeam Backup & Replication checks available backup proxies. If more than one proxy is available, Veeam Backup & Replication analyzes transport modes that the proxies can use and the current workload on the proxies to select the most appropriate resource for VM processing.
- If you choose **Use the backup proxy servers specified below**, you can explicitly select proxies that the job can use. It is recommended that you select at least two proxies to ensure that the job will be performed should one of job proxies fail or lose its connectivity to the source datastore.

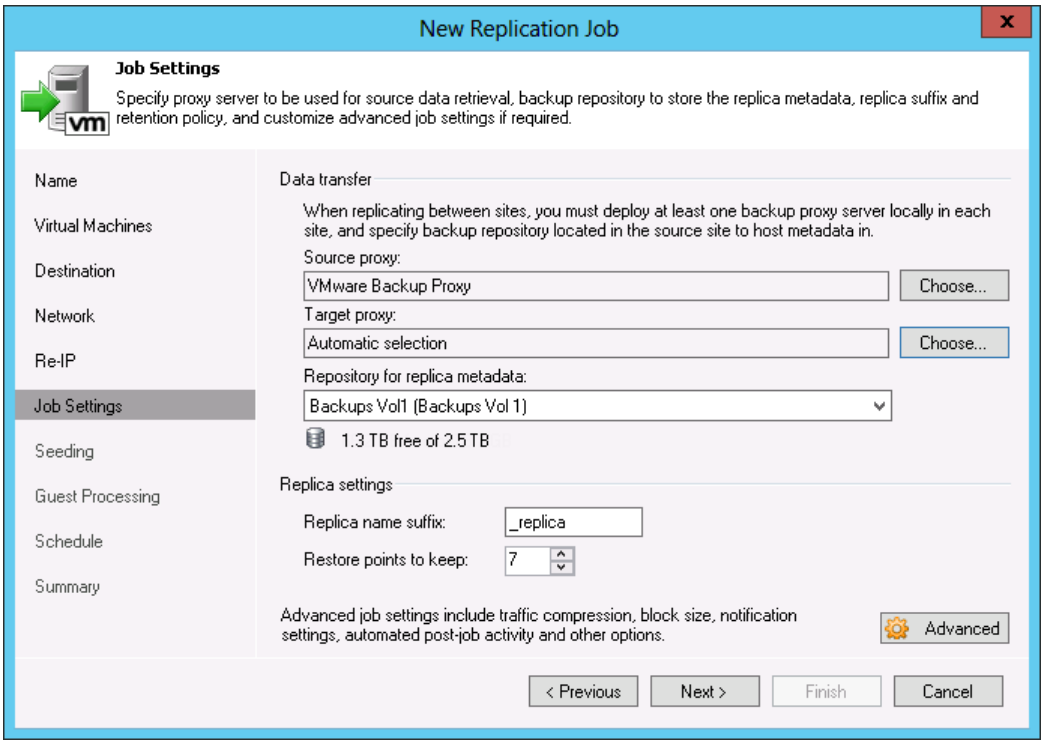

From the **Repository for replica metadata** list, select a repository that is located in the source site. The repository will store metadata (checksums of read data blocks) required to streamline incremental passes of the job.

In the **Replica name suffix** field, enter a suffix that will be appended to the name of the VM you are replicating. This name, with the suffix added, will be used to register the replicated virtual machine on the target server. Files of the replicated VM will be placed to the selected datastore to the `VMname(vm-ID)` folder.

In the **Restore points to keep** field, specify the number of restore points that should be maintained by the replication job. If this number is exceeded, the earliest restore point will be deleted.

Step 10. Specify Advanced Replica Settings

Click **Advanced** to specify advanced options for the created replication job.

Traffic settings

On the **Traffic** tab, specify compression and optimization settings for replication files that the job will provision.

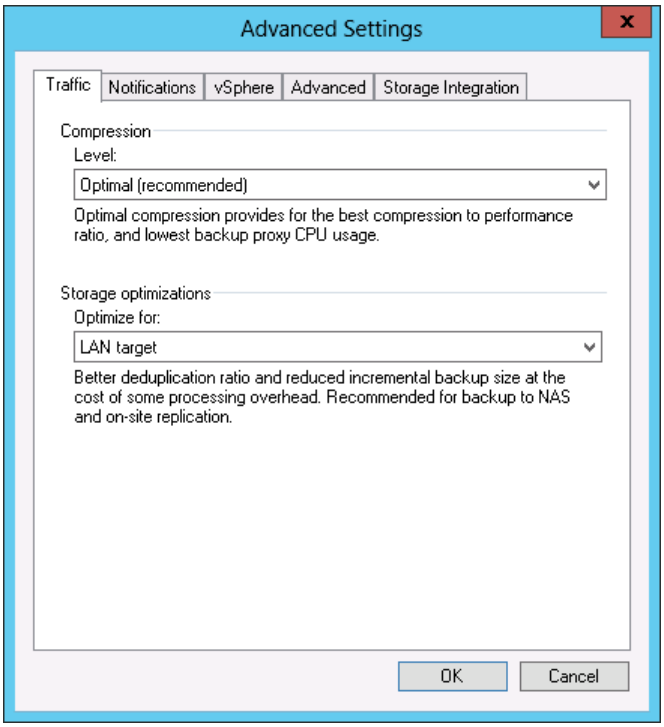

In the **Compression** section, specify a compression level for replica traffic sent over network: *None*, *Dedupe-friendly*, *Optimal*, *High* or *Extreme*.

Note Compression is applicable only if replicated data is transferred between two proxy servers. If one server acts as a source and target proxy, replicated data is not compressed at all.

In the **Storage optimizations** section, select the type of replication target you are planning to use. Depending on the chosen option, Veeam Backup & Replication will use data blocks of different sizes. When selecting the data block size, consider the following aspects:

- Veeam Backup & Replication writes information about every data block of the VM replica to the VM replica metadata stored on the backup repository. The smaller is the size of the data block, the more data blocks there are and the more netadata is written to the backup repository.
- When reading the VM image, Veeam Backup & Replication "splits" the VM image into blocks of the selected size. The more data blocks there are, the more time is required to process the VM image.
- During incremental job cycles, Veeam Backup & Replication uses CBT to define changed data blocks in the VM. The larger is the size of the found changed data block, the more amount of data needs to be transferred to the target side.

You can select the following data blocks for VM image processing:

- **Local target (16 TB + backup size)**. This option is recommended for replication jobs that can produce very large replicas — larger than 16 TB. With this option selected, Veeam Backup & Replication uses data blocks of a very large size, which reduces the job processing overhead.
- **Local target**. This option is recommended if you are planning to use SAN, DAS or local storage as a target. SAN identifies larger blocks of data and therefore can process larger quantities of data at a time. This option provides the fastest replication job performance but increases the amount of traffic at incremental replication cycles.
- **LAN target**. This option is recommended for NAS and onsite replication: it provides the best ratio between the time required for VM data processing and the amount of traffic going to the target side.

- **WAN target**. This option is recommended if you are planning to use WAN for offsite replication. Veeam Backup & Replication will use small data blocks, which will lead to significant processing overhead but will result in small amount of traffic going over the WAN at incremental replication cycles.

Notification settings

Use the **Notifications** tab if you want to be notified when the replication job is completed.

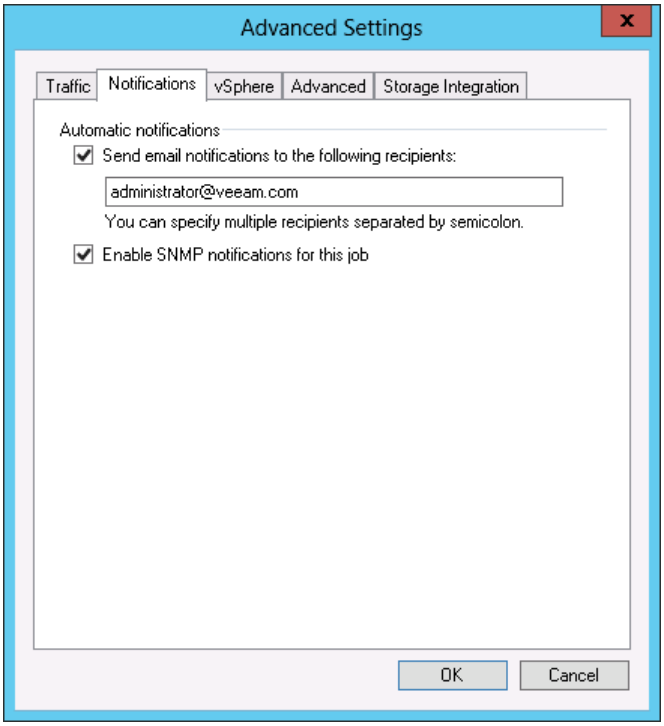

- Select the **Send email notifications to the following recipients** check box if you want to receive notifications by email in case of job failure or success. In the field below, specify a recipient's email address. You can enter several addresses separated by a semicolon. Email notifications will be sent only if you have selected the **Enable email notification** check box in the **Options** window and specified email notification settings (select **Tools** > **Options** from the main menu). For details, see Specifying Notification Settings.
- Select the **Enable SNMP notification for this job** check box if you want to receive SNMP traps when a job is completed successfully. SNMP traps will be sent if you configure SNMP settings in Veeam Backup & Replication and on the recipient's computer. For details, see Specifying SNMP Settings.

vSphere settings

On the **vSphere** tab, specify if file system freezing and changed block tracking should be used.

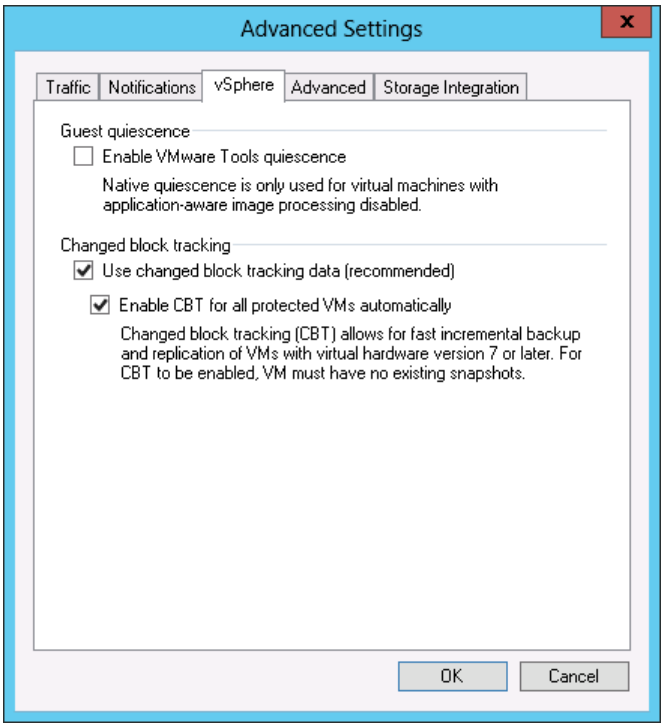

The **Enable VMware tools quiescence** option enables freezing of the file system for proper snapshot creation. With this option enabled, creation of the snapshot is performed with the help of the sync driver responsible for holding incoming I/O and flushing all dirty data to a disk, thus making the file systems consistent.

In the **Changed block tracking** section, specify if vSphere Changed Block Tracking (CBT) should be used. By default, this option is selected. If you want to force using changed block tracking even if CBT is disabled on the ESX(i) host, select the **Enable changed block tracking for all processed VMs** check box.

Important! You can use this option only for VMs using virtual hardware version 7 or later.

Advanced settings

On the **Advanced** tab, specify miscellaneous advanced settings for the job.

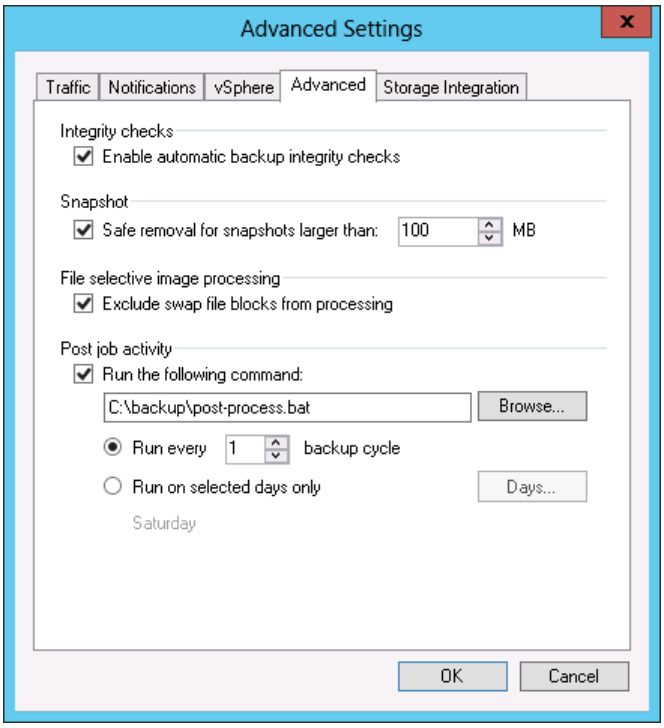

- Select the **Enable automatic backup integrity checks** check box if you want Veeam Backup & Replication to periodically check the VM replica. An automatic integrity check allows you to verify integrity of the VM replica files and avoid a situation when a VM replica is corrupted, making all further increments corrupted, too.
  An integrity check is performed every time the job is started. During the check, Veeam Backup & Replication verifies service information for the VM replica written to the replica metadata on the backup repository. If the check fails, Veeam Backup & Replication displays a notification message, prompting you to replicate the VM anew. During such replication cycle, no integrity check is performed.
- If you are running pre-ESX 3.5 Update 2 hosts, consider enabling the safe snapshot removal option. Because full image-level replication can take a long time depending on the VM size, the VM snapshot can grow very large. When a large snapshot is removed on a VM with heavy disk I/O, a consolidation helper snapshot may grow large too, and will then require a long time to be committed. While a helper snapshot is being committed into VM virtual disk files, the VM remains completely "frozen", and depending on the consolidation helper snapshot size, the freeze time may be so long that some applications running on a VM will time out. To prevent such situation, Veeam Backup & Replication offers a procedure of safe snapshot removal which includes creating an additional snapshot in cases when the "main" snapshot size is above the specified threshold. An additional snapshot is used to host writes while the "main" snapshot is being deleted. This ensures that a consolidation helper snapshot does not grow large. To use this option, select the **Safe removal for snapshots larger than ... MB** check box and specify a threshold for the size of a snapshot that should not be exceeded.
- In the **File selective image processing** section, define whether you want to exclude blocks of Windows page files from the replica. During replication, Veeam Backup & Replication checks the NTFS MTF file on Windows-based VMs to identify blocks of the Windows pagefile, and excludes these blocks from processing. Windows page files are dynamic in their nature and change intensively between replication job runs, even if VMs do not change much. Therefore, page file processing results in reduced replication performance and increased size of increments.

Clear the **Exclude swap file blocks from processing** check box if Windows page files should be processed during replication.

- Select the **Run the following command** check box if you want to execute post-replication actions. Use the **Browse** button to select an executable file.
  You can select to execute post-replication actions after a number of replication cycles or on specific week days. If you select the **Run every... backup cycle** option, specify the number of the replication job run after which the file should be executed. If you select the **Run on selected days only** option, click the **Days** button and specify week days when actions should be performed.

Storage integration settings

On the **Storage Integration** tab, define whether you want to use the Backup from Storage Snapshots technology or not. Backup from Storage Snapshots lets you leverage HP SAN snapshots for VM data processing. The technology dramatically improves RPOs and reduces impact of replication activities on the production environment.

By default, the **Use storage snapshots** option is enabled. If you do not want to use the Backup from Storage Snapshots technology, clear the check box. To learn more, see Performing Backup from Storage Snapshots.

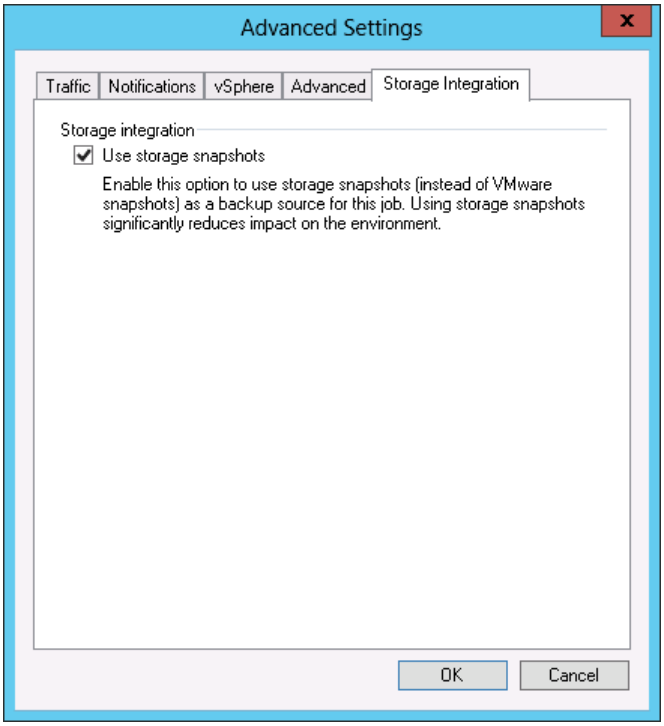

Step 11. Define Seeding and Mapping Settings

This step is available if you have selected the **Low connection bandwidth** option at the Name step of the wizard.

Configuring Replica Seeding

If you plan to replicate to a remote DR site over WAN or low-bandwidth network, consider initial replica seeding. To take advantage of replica seeding, you need to perform a number of preparatory tasks before configuring this step:

1. Create a backup (seed) of the VM you plan to replicate. To do that, configure a backup job that points to an onsite backup repository. Run the job to perform a full backup.
   If you have previously created a backup containing all necessary VMs, there is no need to configure and run a new backup job. For seeding, you can use any existing backup created with Veeam Backup & Replication v6. The backup must include .vbk and .vbm files. If you have a full backup and a chain of forward increments, you can use .vib files along with the .vbk and .vbm files. In this case, Veeam Backup & Replication will restore VMs from the seed to the latest available restore point.
2. Copy the backup from the backup repository in the production site to a backup repository in the DR site. If you do not have a backup repository in the DR site, you will need to create one. You can move the backup using a file copy job or any other appropriate method (for example, copy the backup to a removable storage device, ship the device to the DR site and copy backups to the backup repository in the DR site).
3. After the backup is copied to the backup repository in the DR site, perform rescan of this repository. Otherwise, Veeam Backup & Replication might not be able to detect the new backup.

When the preliminary steps are completed, you can configure replica seeding settings for the job. In the **Initial seeding** section, select the **Get seed from the following backup repository** check box. From the list of repositories, select a DR-site repository to which the seed (the full backup) was copied.

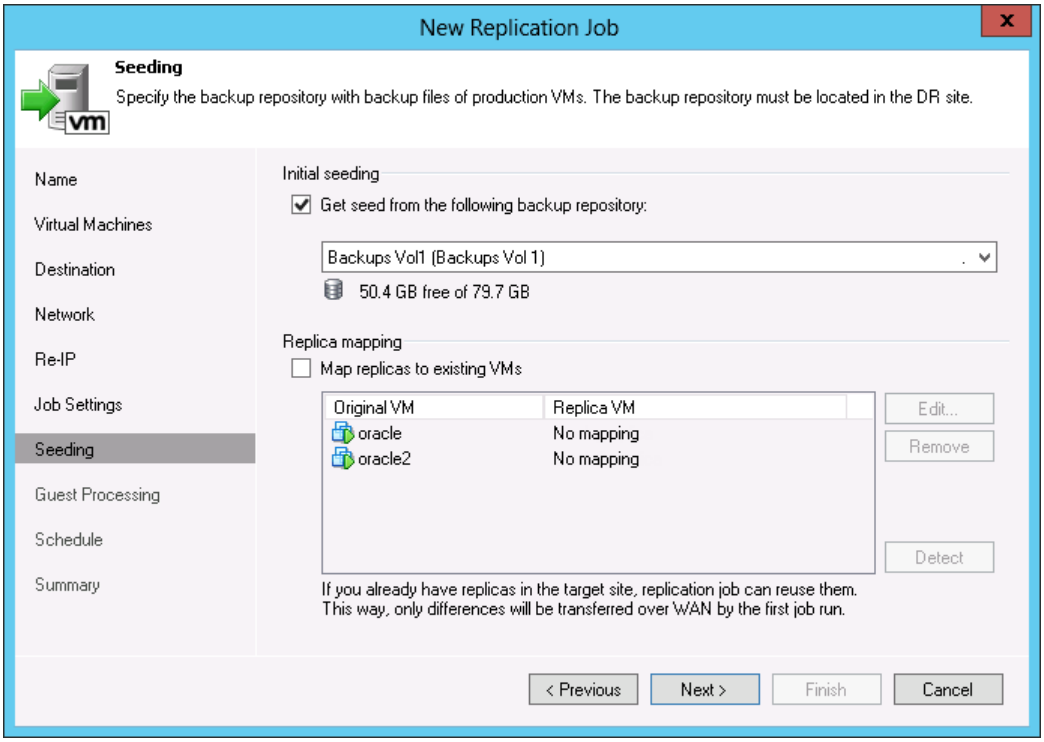

During the first run of the job, Veeam Backup & Replication will restore full VMs from the backup and then additionally synchronize them with the original VMs. All subsequent incremental replication runs will be performed in the regular course.

Note If you select the **Get seed from the following backup repository** check box, Veeam Backup & Replication will attempt to restore all VMs included in the job from the seed you specified. If a VM is not found in the seed, it will be skipped from replication.
Within the same job, you configure both replica seeding and replica mapping – for example, if a job includes 2 VMs, you can use seeding for one VM, and map the other VM to an existing replica. Note that if the **Get seed from the following backup repository** check box is selected, all VMs in the job must be covered with seeding or mapping – in case a VM is neither available in the seed, nor mapped, it will be skipped from replication. And on the contrary, if the same VM is both available in the seed and mapped to an existing replica, replication will be performed using replica mapping – as mapping has precedence over seeding.

Configuring Replica Mapping

To reduce the amount of traffic sent over the network during replication, you can map the VM you plan to replicate to its already existing replica. Veeam Backup & Replication will link the original VM to the existing replica, and the replication job will use this replica in a regular replication course.

If there is no existing replica, you can restore a VM from the backup and map it to the original VM. During the first pass of the job, Veeam Backup & Replication will compare the original VM to its mapped copy and will transfer only differences between the two VMs to synchronize the restored replica with the production VM.

To set up VM mapping, select the **Map replicas to existing VMs** check box and click the **Detect** button. Veeam Backup & Replication will scan the destination location for existing replicas and will populate the mapping table if any matches are found.

If Veeam Backup & Replication does not find a match, you can map a VM to a replica manually. To do so, select a production VM from the list, click **Edit** and choose an existing replica. To facilitate selection, use the search field at the bottom of the window.

To break a mapping association, select the VM in the list and click **Remove**.

Note If you use replica seeding or mapping, make sure that you correctly set up replication infrastructure components for the job (source-side backup repository for metadata and backup proxies). It is recommended that you explicitly assign backup proxies in the production site and in the DR site. For details, refer to step 9 of the procedure.

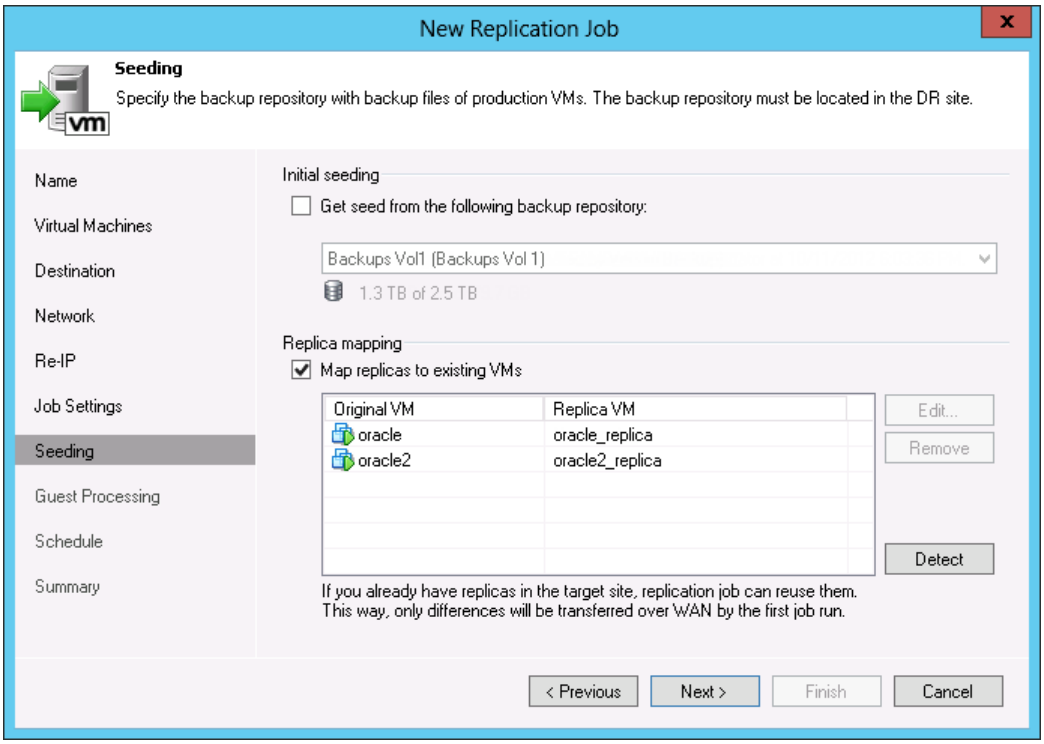

Step 12. Enable Application-Aware Image Processing

If you want to create a transactionally consistent replica ensuring successful recovery of VM applications without any data loss, select the **Enable application-aware image processing** check box.

To coordinate proper VSS activities, Veeam Backup & Replication injects a runtime process inside the VM. The process is run only during VSS quiescence procedure and stopped immediately after the processing is finished (depending on the selected option, during the replication job or after it is finished), thus producing low impact on VM performance and stability.

In the **Guest OS credentials** section, specify an account with administrative privileges for injecting the process. Please note that the user name must be supplied in the *DOMAIN\USERNAME* format. The guest OS credentials you provide will be used for all VMs included into the replication job.

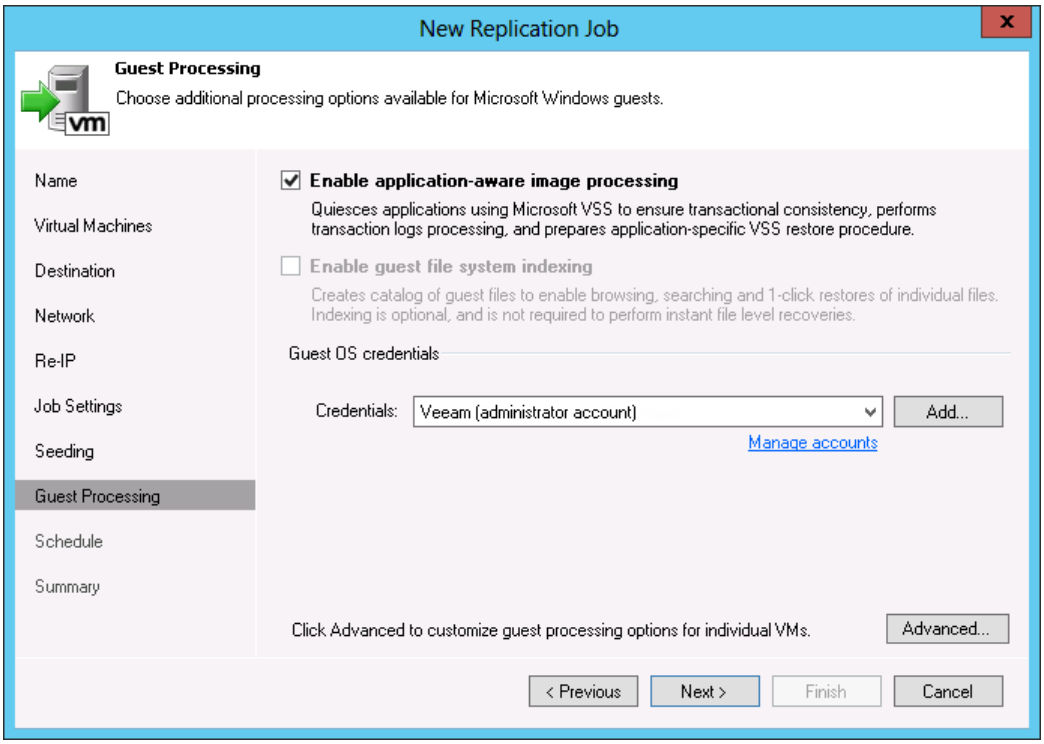

Click **Advanced** to specify advanced option for Veeam VSS processing. The **Advanced Options** section contains a list of VMs that will be processed with Veeam VSS.

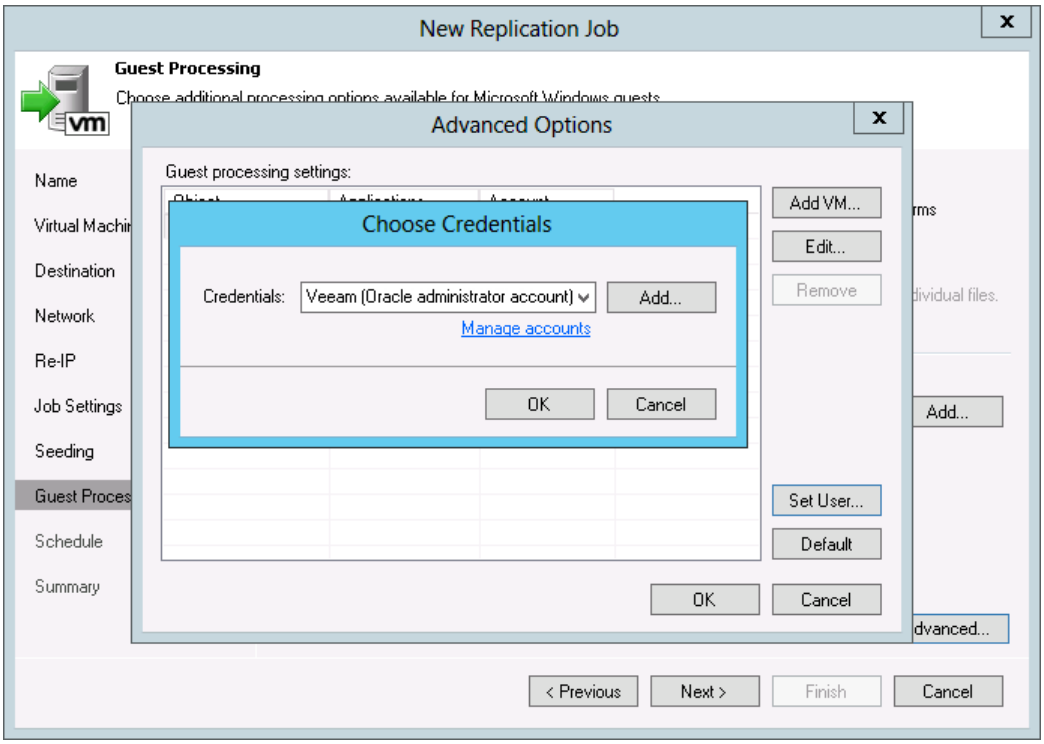

By default, for all VMs in the list Veeam Backup & Replication uses common credentials you provided in the **Guest OS credentials** section. If a different account should be used to inject the process into a specific VM, select the VM in the list, click **Set User** and enter custom guest OS credentials. To discard custom credentials for a VM, select it in the list and click **Default**.

If you want to define custom settings for a VM added as part of a container, include the VM in the list as a standalone instance. To do so, click **Add VM** and choose a VM whose settings you want to customize. Next, select the VM in the list and define the necessary custom settings. To discard custom settings of a VM, select the VM in the list and click **Remove**.

To provide granular quiescence and indexing options for a VM, select it in the list and click **Edit**.

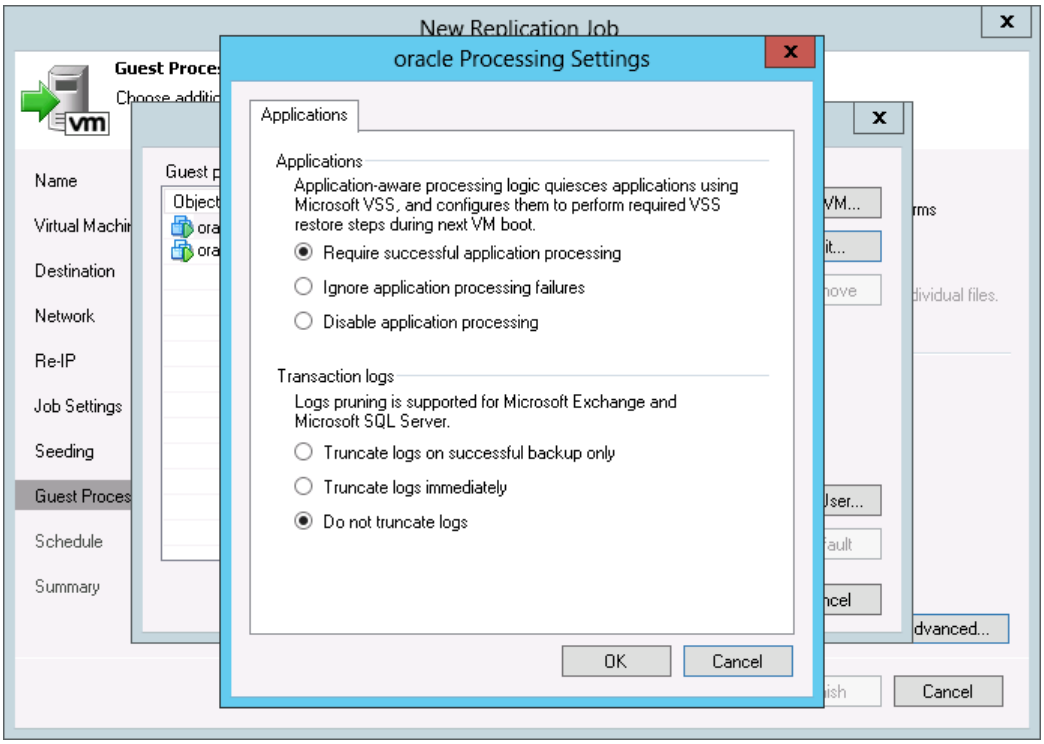

In the **Applications** section on the **Applications** tab, specify the VSS behavior scenario:

- Select **Require successful application processing** if you want Veeam Backup & Replication to stop replicating the VM if any VSS errors occur.
- Select **Ignore application processing failures** if you want to continue replicating the VM even if VSS errors occur. This option is recommended to guarantee completion of the job. The created replica will not be transactionally consistent, but crash consistent.
- Select **Disable application processing** if you do not want enable quiescencing for the VM.

Use the **Truncation logs** section to define the scenario of transaction log handing:

- Select **Truncate logs on successful backup only** if you want Veeam Backup & Replication to trigger truncation of logs only after the job is finished successfully. In this case, Veeam runtime process will wait for the replication to complete, and then will trigger truncation of transaction logs. If truncation of transaction logs is not possible for some reason, the logs will remain untouched in the VM guest OS till the next start of the Veeam runtime process.
- Select **Truncate logs immediately** if you want Veeam Backup & Replication to trigger truncation of logs in any case, no matter whether the job finishes successfully or fails.
- Select **Do not truncate logs** if you do not want Veeam Backup & Replication to truncate logs at all. This option is recommended if you are using another tool to perform guest-level replication and this tool maintains consistency of the database state. In such scenario, truncation of logs with Veeam Backup & Replication will break the guest-level replication chain and cause it to fall out of sync.

Step 13. Define the Job Schedule

At the **Schedule** step of the wizard, you can select to manually run the replication job, schedule the job to start at a specific time –(for example, the least busy hours to reduce impact on the virtual environment) or define a schedule for the job to run on a regular basis.

To specify the job schedule, select the **Run the job automatically** check box. If this check box is not selected, the job is supposed to be run manually.

You can define the following scheduling settings for the job:

- You can choose to run the job at specific time on defined week days, monthly and with specific periodicity.
- You can choose to run the job continuously. In this case, the next run of the backup job will be started as soon as the previous one is completed, maintaining your backup always in the most recent state. To run the job continuously, select the **Periodically every** option and choose **Continuously** from the list on the right.
- You can choose to run the job repeatedly throughout a day with a set time interval. At the **Schedule** step of the wizard, select the **Periodically every** option, enter the necessary time interval and select the necessary time unit: *Hours* or *Minutes*. Click **Schedule** on the right and use the time table to define the permitted time window for the job. If you choose to run the job at an hourly interval, in the **Start time for hourly jobs** field, specify the exact time when the job should start.
  For example, you want to start a job every 2 hours from 9AM to 6PM. At the **Schedule** step, select the **Periodically every** option, enter 2 in the field on the right and select *Hours* from the list. Click **Schedule** and use the **Permitted** and **Denied** options to mark the time window from 9AM to 6PM. In the **Start time for hourly jobs** field, specify the exact start time of the job, for example, 15 minutes. The job you have scheduled will start at 9:15 AM, 11:15 AM, 1:15 PM, 3:15 PM and 5:15 PM.
- You can chain jobs. In the common practice, jobs start one after another: when the job "A" finishes, the job "B" starts and so on. If you want to create a chain of jobs, you should define the time schedule for the first job in the chain. For the rest of the jobs in the chain, at the **Schedule** step of the wizard, select the **After this job** option and choose the preceding job from the list.

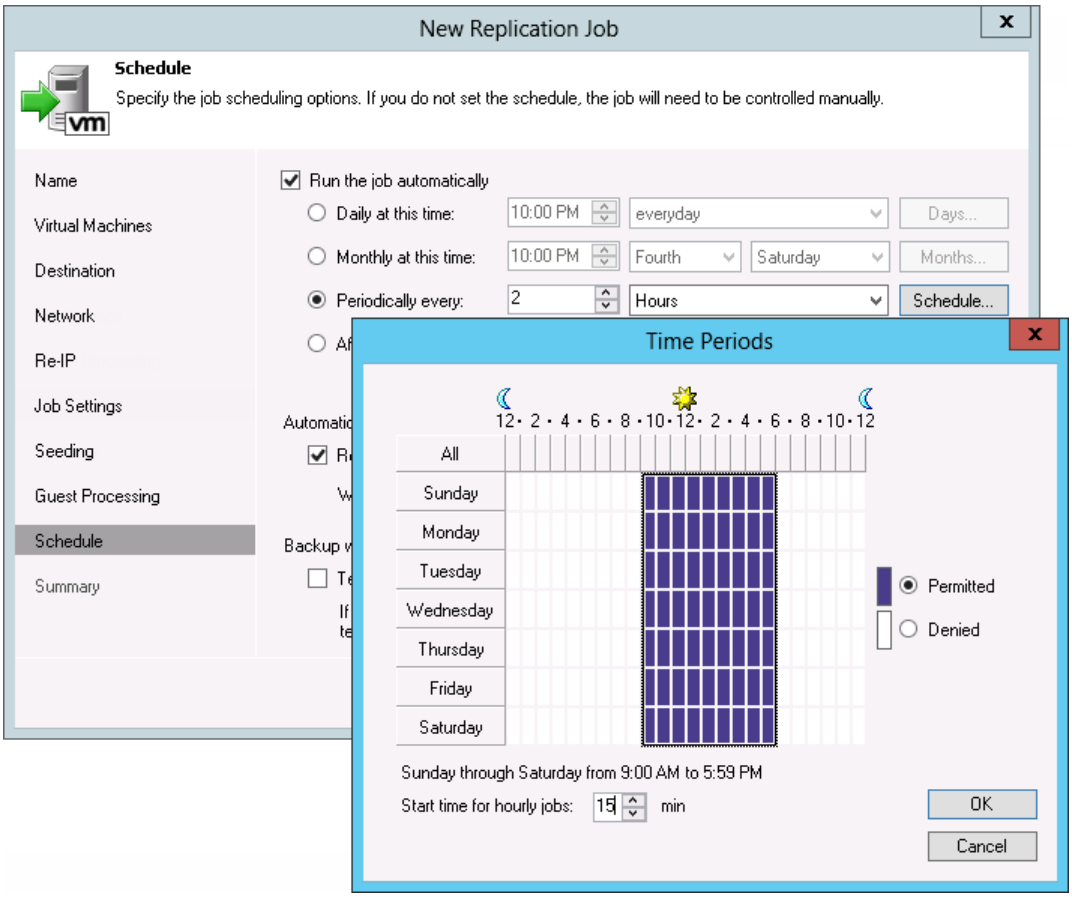

In the **Automatic retry** section, select to repeat an attempt to run the replication job in case it fails for some reason. A repeatedly run job will include failed VMs only. Enter the number of attempts to run the job and define time spans between them. If you select continuous replication, Veeam Backup & Replication will retry the job for the defined number of times without any time intervals between the job runs.

In the **Backup window** section, determine a time span within which the replication job must be completed. The backup window prevents the job from overlapping with production hours and ensures it does not provide unwanted overhead on your virtual environment. To set up a backup window for the job, select the **Terminate job if it exceeds allowed backup window** check box and click **Window**. In the **Time Periods** section, define the allowed window and prohibited hours for replication. If the job exceeds the allowed window, it will be terminated.

Note After you have created a scheduled job, you can temporarily disable it (that is, hold it for some time without changing the set time schedule). Right-click a job in the list and select **Disable Job** from the shortcut menu. To enable the job schedule, right-click the job and deselect **Disable Job** in the shortcut menu.

Step 14. Finish Working with the Wizard

After you have specified schedule settings, click **Create**. Select the **Run the job when I click Finish** check box if you want to start the created job right after you complete working with the wizard. Click **Finish** to close the wizard.

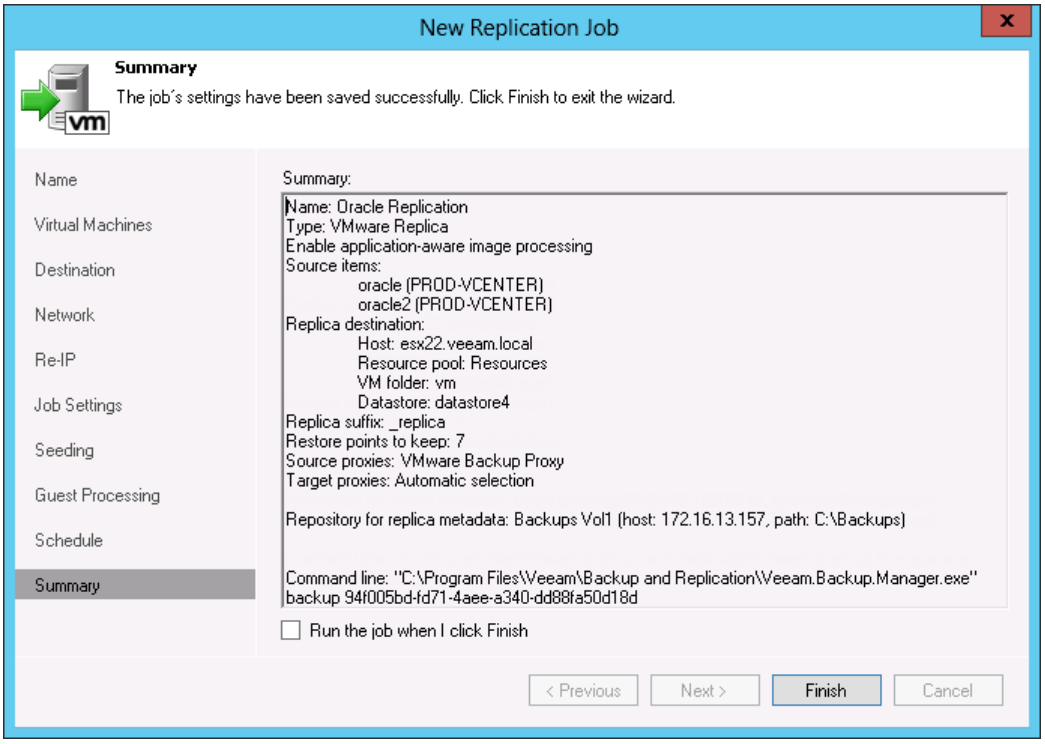

Creating VM Copy Jobs

With a VM copy job, you can create a fully-functioning copy of a VM that will require no manual editing and adjustments. VM copying can be helpful if you want to move your datacenter, mirror your production environment to test lab storage and so on. Just as backup and replication jobs, a VM copy job can be performed using the Direct SAN Access, Virtual Appliance and Network transport modes, supports VSS options, can be run on demand or scheduled.

This section will guide you through all steps of the VM Copy wizard and provide explanation on available options.

Before You Begin

- Prior to creating a VM copy job, make sure you have set up all necessary infrastructure components for the job. Open the Infrastructure view and check if source and destination hosts are available under the Managed servers node in the management tree, backup proxies and backup repositories are available under the Backup Proxies and Backup Repositories nodes and properly configured. You will not be able to add VM copy infrastructure components or change their configuration once the VM Copy Job wizard is launched. For detailed information on adding VM copy infrastructure components, see Setting Up Backup Infrastructure.
- During every job run, Veeam Backup & Replication checks disk space on the destination storage. If the disk space is below a specific threshold value, Veeam Backup & Replication will display a warning in the job session log. To specify the disk space threshold, select **Options** from the main menu. On the **Notifications** tab, specify the amount of free disk space required in percent.

To create a copy job for VMware virtual machines, follow the next steps:

Step 1. Launch the VM Copy Job Wizard

To run the **VM Copy Job** wizard, do either of the following:

- On the **Home** tab, click **VM Copy**.
- Open the **Virtual Machines** view, select one or several VMs in the working area, click **Add to VM Copy** on the ribbon and select **New job**. Alternatively, you can right-click one or several VMs in the working area and select **Add to VM Copy Job** > **New job**. In this case, the selected VMs will be automatically included into the VM copy job. You can add other VMs to the job when passing through the wizard steps.

You can quickly include VMs to already existing jobs. To do that, open the **Virtual Machines** view, right-click necessary VMs in the working area and select **Add to VM Copy Job** > *name of a created job*.

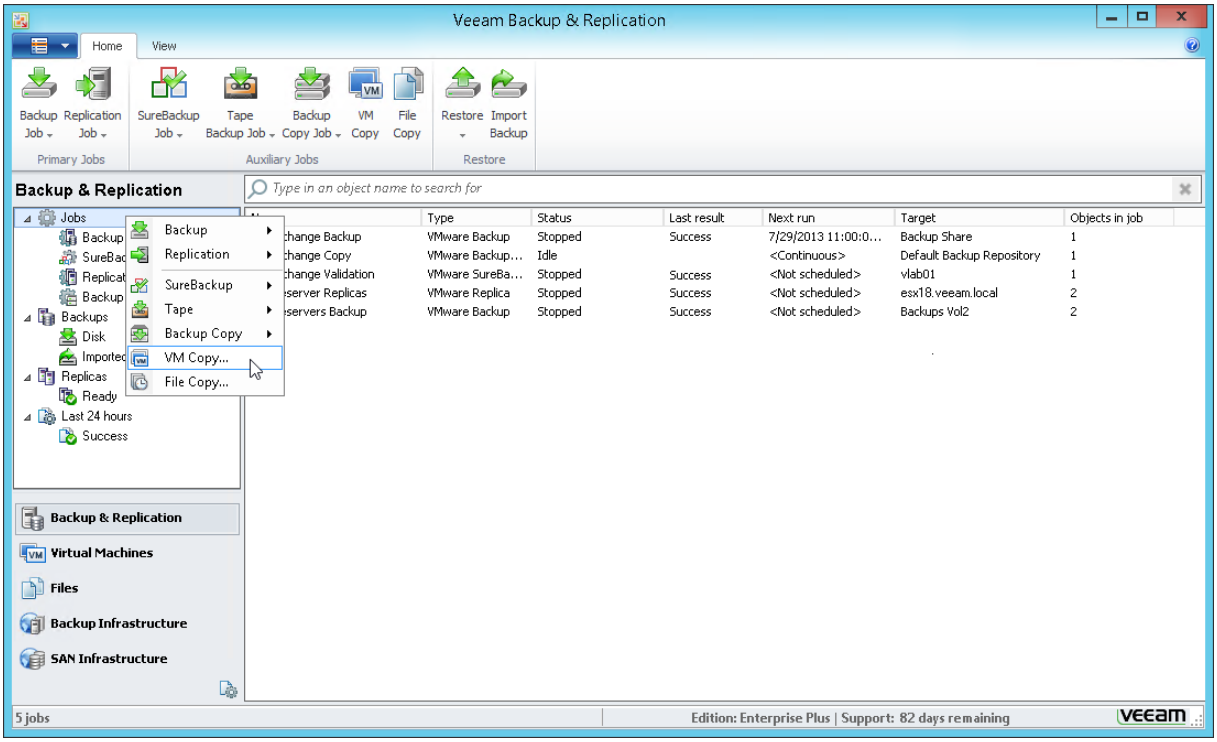

Step 2. Specify Job Name and Description

At the first step of the wizard, enter a name and description for the job. The default description contains information about the user who created the job, as well as the date and time when the job was created.

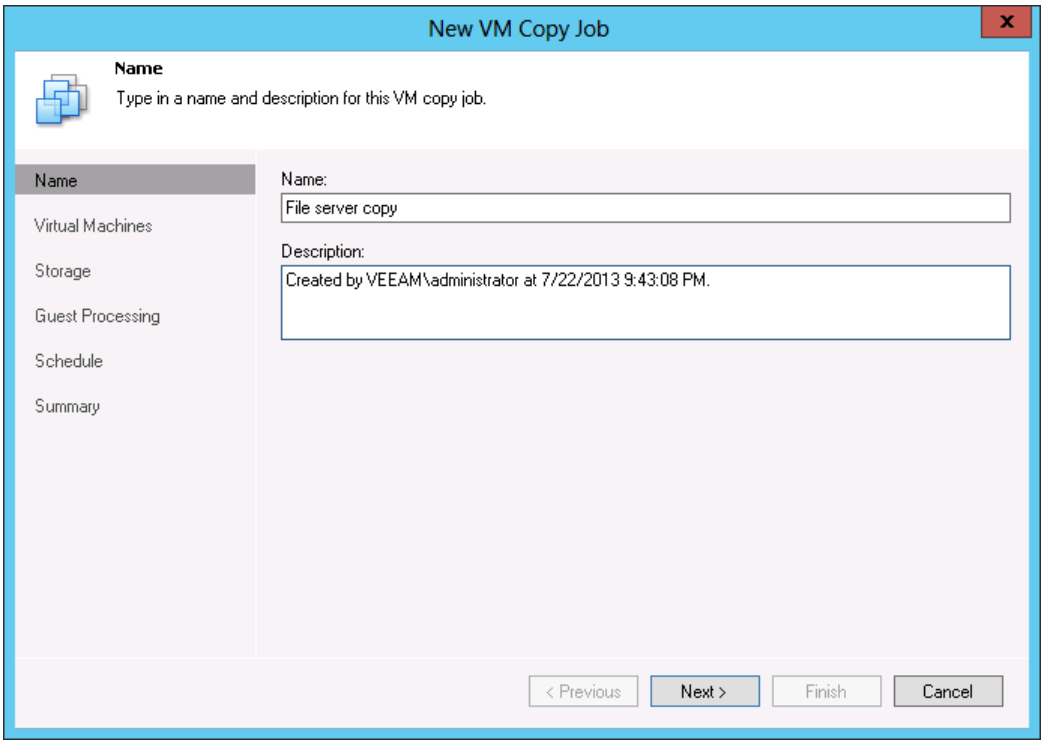

Step 3. Select Virtual Machines to Copy

At this step, you should select an individual VM, multiple VMs or VM containers you want to copy. Jobs with VM containers are dynamic in their nature: if a new VM is added to the container after a copy job is created, the job will be automatically updated to include the added VM.

Click **Add** to browse to VMs and VM containers that should be copied. In the displayed tree, select the necessary object and click **Add**.

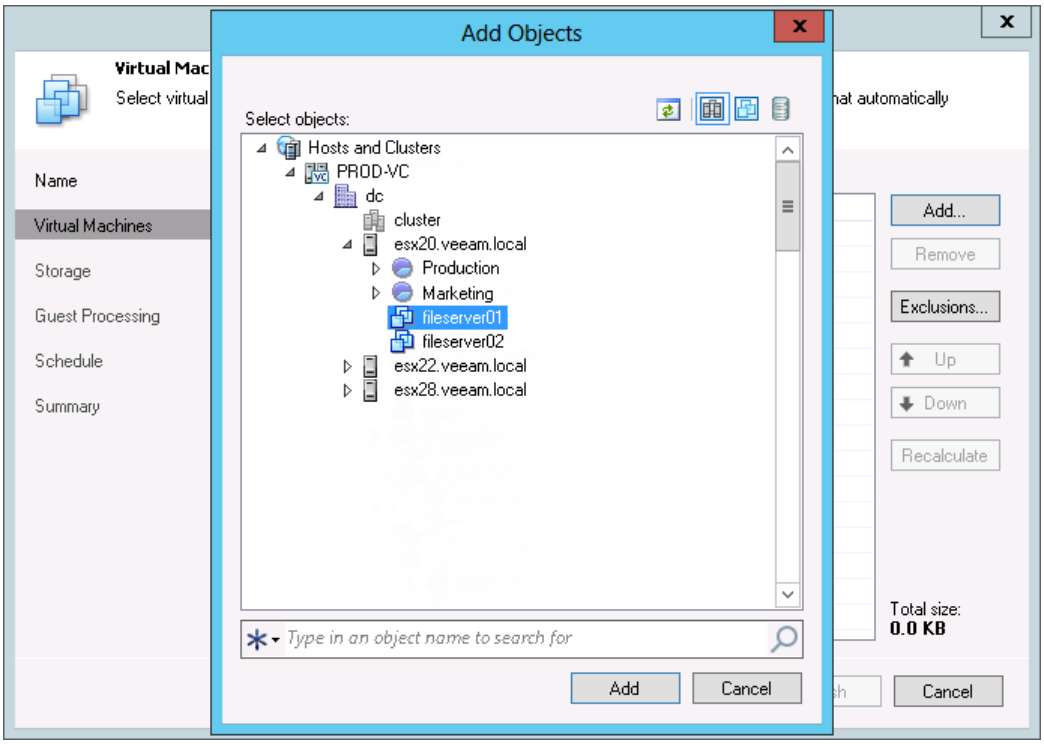

To facilitate objects selection, you can:

- Switch between views by clicking the **Hosts and Clusters**, **VMs and Templates** or **Datastores and VMs** buttons at the top of the tree.
- Use the search field at the bottom of the window: click the button on the left of the field to select the necessary type of object that should be searched for (**Everything**, **Folder**, **Cluster**, **Host**, **Resource Pool**, **Virtual Application** or **VM**), enter an object's name or a part of it and click the **Start search** button on the right or press **[ENTER]**.

Note Depending on the view you select, some objects may not be available (for example, if you select the **VMs and Templates** view, you will not be able to see and find resource pools).

To remove an object from the list, select it and click **Remove** on the right.

The initial size of VMs and VM containers added to the VM copy job is displayed in the **Size** column in the list. The total size of objects is displayed in the **Total size** field. Use the **Refresh** button to refresh the total size value after you add a new object to the job.

Step 4. Exclude Objects from the VM Copy Job

After you have added VMs and VM containers to the list, you can specify which objects should be excluded from the VM copy job. Veeam Backup & Replication allows excluding the following types of objects: VMs and VM templates from VM containers, as well as specific VM disks.

To select which objects should be excluded, click **Exclusions** on the right.

- To exclude VMs from a VM container (for example, if you need to copy the whole ESX host excluding several VMs running on this server), click the **VMs** tab. Click **Add** on the right and select VMs that should be excluded. To display all hosts added to Veeam Backup & Replication, select the **Show full hierarchy** check box. To facilitate objects selection, you can switch between the **Hosts and Clusters**, **VMs and Templates** and **Datastores and VMs** views, and use the search field just as in the main window of the wizard.
- To exclude specific VM disks from the copy job, open the **Disks** tab, select the necessary VM in the list and click **Edit**. If a VM is not in the list, you can add it by clicking **Add**. You can choose to process all disks, 0:0 disks (typically, the system disks) or select custom disks.

By default, the **Remove excluded disks from VM configuration** check box is selected, which means that Veeam Backup & Replication will modify VMX file to remove disks you want to skip from VM configuration. When this option is used, you will be able to restore, replicate or copy VM to a location where excluded disks are not accessible with the original paths. If you do not use this option, you will have to manually edit VM configuration file to be able to power on a VM.

- When processing VM containers, Veeam Backup & Replication copies VM templates along with VMs. If you do not want to copy VM templates, open the Templates and clear the Backup VM templates check box. The **Exclude templates from incremental backup** option allows you to process VM templates with a full copy job only.

Note Veeam Backup & Replication automatically excludes VM log files from a copy job to make copying process faster and reduce the size of the VM copy.

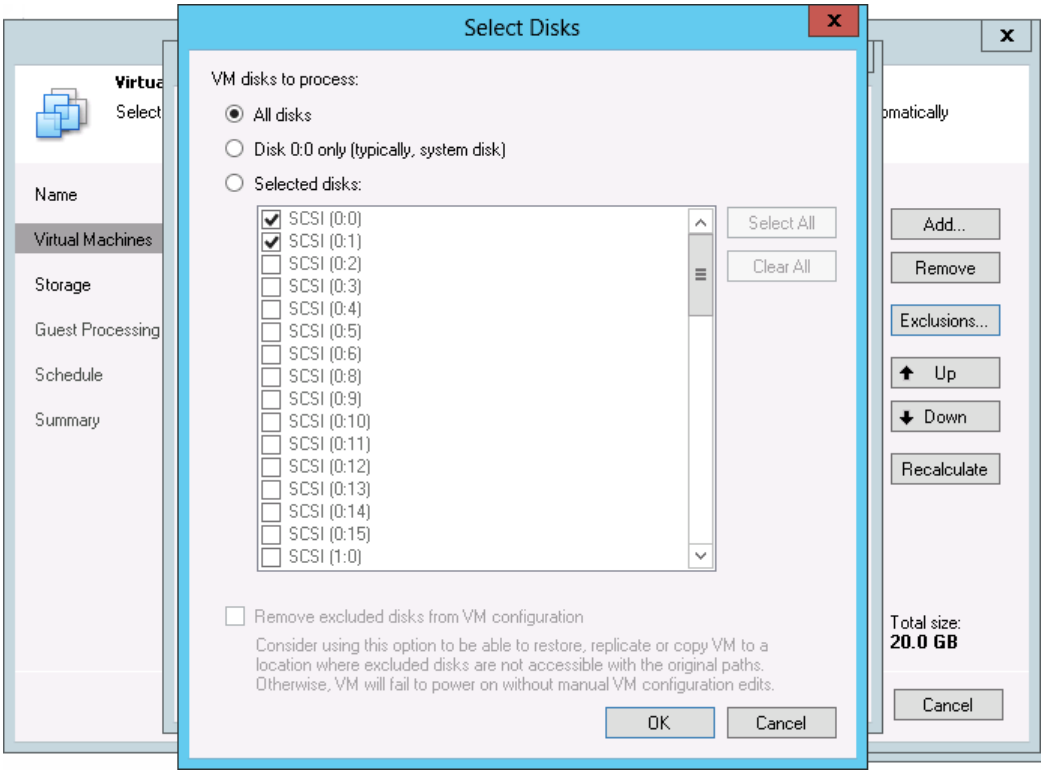

Step 5. Define VM Copy Order

If you want to copy certain VMs before others, you can define the order in which the VM copy job must process VMs. VM copy order can be helpful if you want to ensure that copying of a VM does not overlap with other scheduled activities, or that the copy process is completed before a certain time.

To set VM copy order, select the necessary VMs and move them up or down the list using the **Up** and **Down** buttons on the right. In the same manner, you can set the copy order for containers in the backup list. Note, however, that if you choose to copy a container, VMs inside the container will be processed randomly. To ensure that VMs are processed in the defined order, you should add them as standalone VMs, not as part of a container.

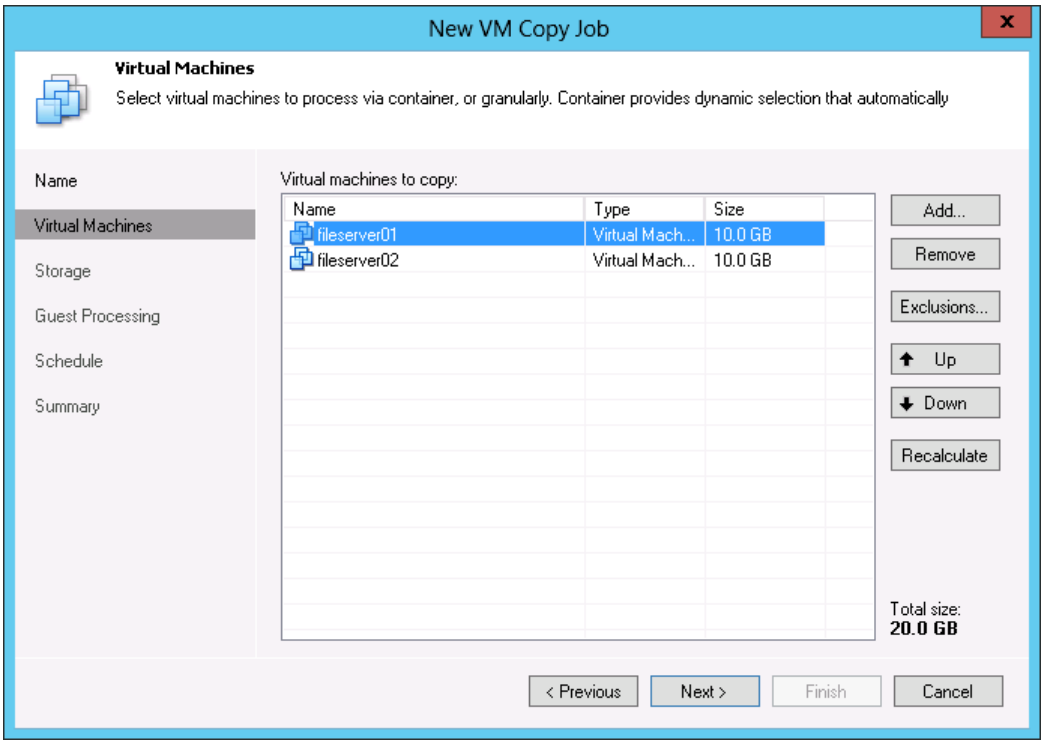

Step 6. Specify Copy Destination

At this step of the wizard, you should select VM copy infrastructure components and define storage destination for the created VM copy.

Click **Choose** next to the **Backup proxy** field to select a backup proxy for the job.

- If you choose **Automatic selection**, Veeam Backup & Replication will detect backup proxies that are connected to the source datastore and will automatically assign optimal proxy resources for processing VM data.
  VMs included in the copy job are processed one by one. Before processing a new VM from the list, Veeam Backup & Replication checks available backup proxies. If more than one proxy is available, Veeam Backup & Replication analyzes transport modes that the proxies can use for data retrieval and the current workload on the proxies to select the most appropriate resource for VM processing.
- If you choose **Use the backup proxy servers specified below**, you can explicitly select proxies that the job can use. It is recommended that you select at least two proxies to ensure that the VM copy job will be performed should one of job proxies fail or lose its connectivity to the source datastore.

From the **Backup repository** list, select the repository where the created VM copies should be stored. Make sure you have enough free space on the storage device.

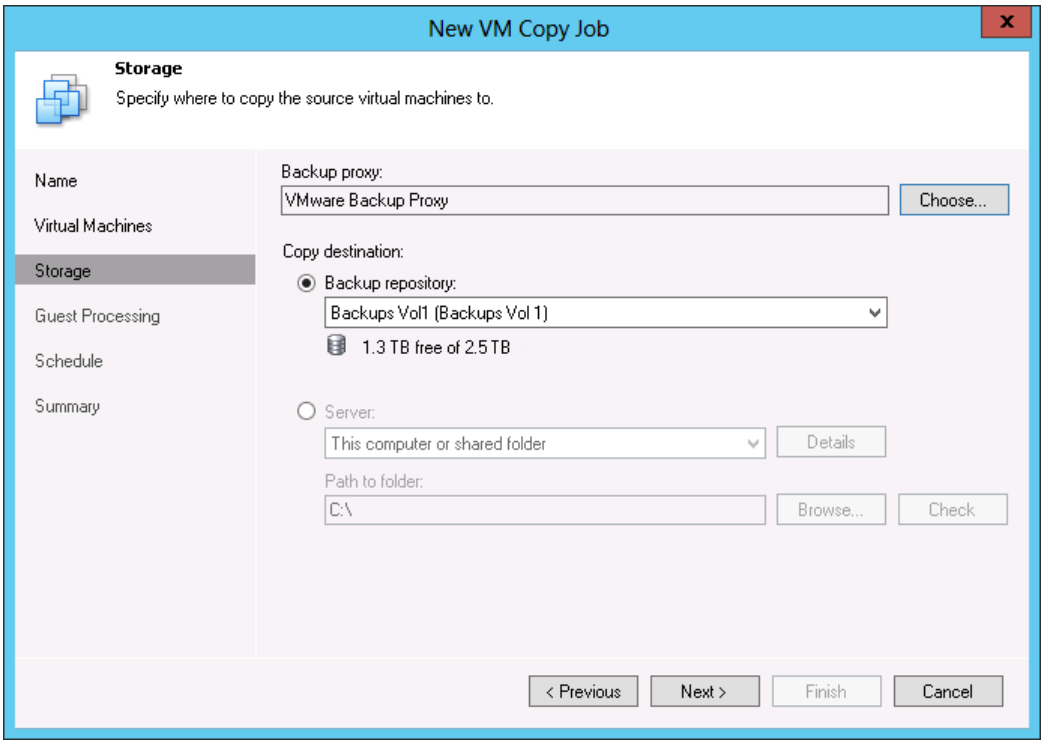

Alternatively, you can copy selected VMs and VM containers to any server connected to Veeam Backup & Replication. Select the **Server** option and choose the server where the created copies should be stored. You can store VM copies to the local host or to any servers added to the backup infrastructure. Use the **Details** button to view or change server connection and data transfer settings.

In the **Path to folder** field, specify the folder on the server where the created copy should be stored. Use the **Check** button to check how much free space is available on copy destination.

Step 7. Enable Application-Aware Image Processing

If you want to create a transactionally consistent VM copy ensuring successful recovery of VM applications without any data loss, select the **Enable application-aware image processing** check box.

To coordinate proper VSS activities, Veeam Backup & Replication injects a runtime process inside a VM. The process is run only during VSS quiescence procedure and stopped immediately after processing is finished (depending on the selected option, during the copy job or after it is finished), thus producing low impact on VM performance and stability.

In the **Guest OS credentials** section, specify an account with administrative privileges for injecting the process. Please note that the user name must be specified in the *DOMAIN\USERNAME* format. The guest OS credentials you provide will be used for all VMs included into the VM copy job.

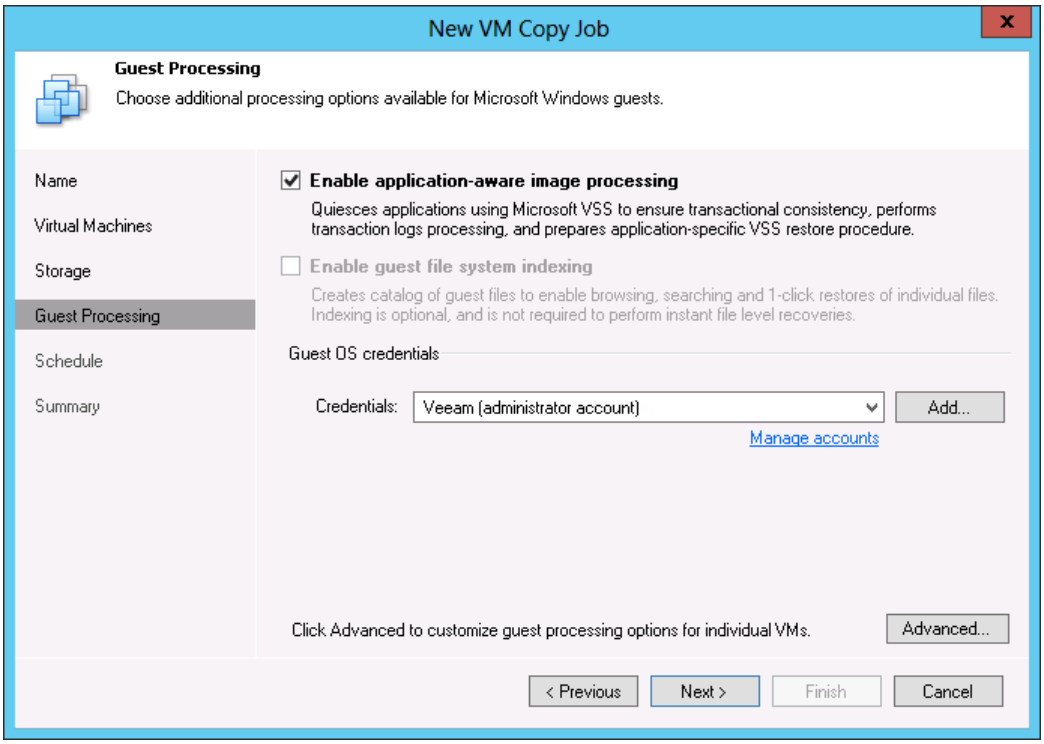

Click **Advanced** to specify advanced options for Veeam VSS processing. The **Advanced Options** window contains a list of VMs that will be processed with Veeam VSS.

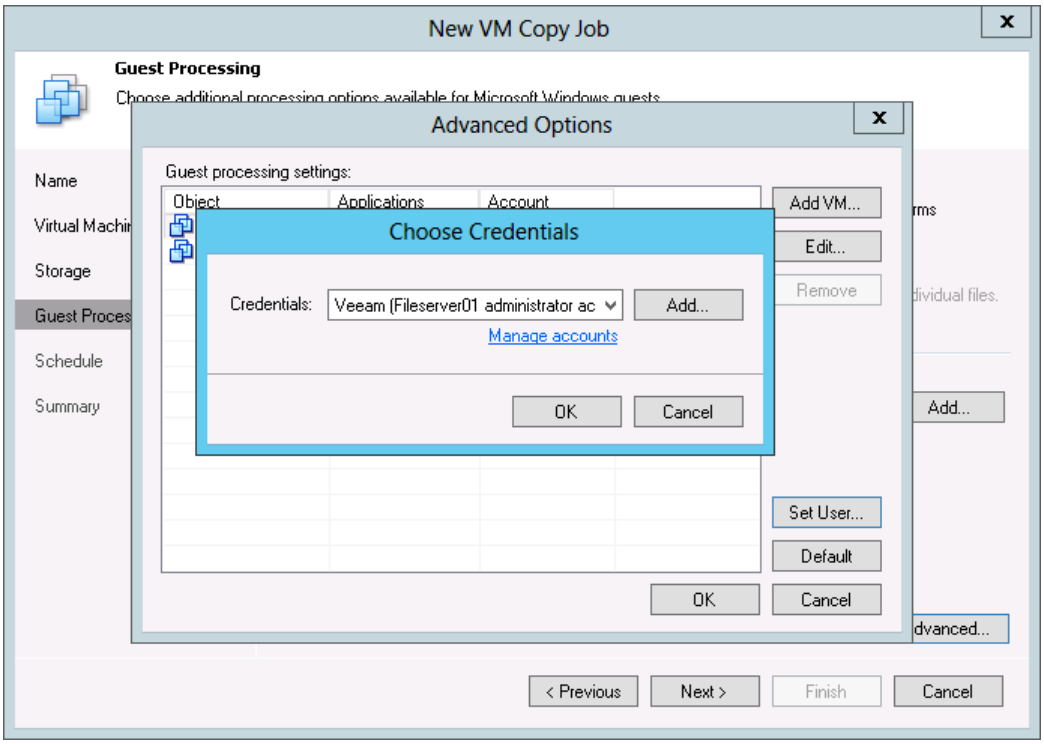

By default, for all VMs in the list Veeam Backup & Replication uses common credentials you provided in the **Guest OS credentials** section. If a different account should be used to inject the process into a specific VM, select the VM in the list, click **Set User** and enter custom guest OS credentials. To discard custom credentials for a VM, select it in the list and click **Default**.

If you want to define custom settings for a VM added as part of a container, include the VM in the list as a standalone instance. To do so, click **Add VM** and choose a VM whose settings you want to customize. Next, select the VM in the list and define the necessary custom settings. To discard custom settings of a VM, select the VM in the list and click **Remove**.

To provide granular quiescence and indexing options for a VM, select it in the list and click **Edit**.

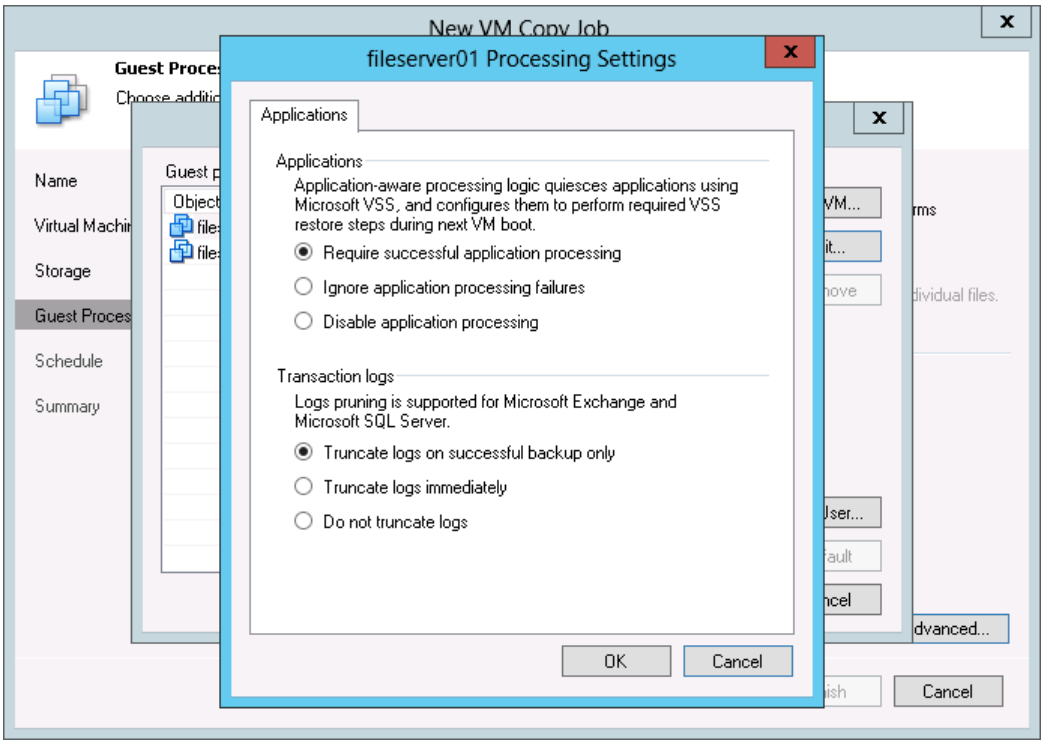

In the **Applications** section on the **Applications** tab, specify the VSS behavior scenario:

- Select **Require successful application processing** if you want Veeam Backup & Replication to stop copying the VM if any VSS errors occur. In this case,the Veeam runtime process will wait for the copy job to complete and then truncate transaction logs. If the runtime process does not manage to truncate transaction logs for some reason, it will remain in the VM guest OS till the next start of Veeam VSS.
- Select **Ignore application processing failures** if you want to continue copying the VM even if VSS errors occur. This option is recommended to guarantee completion of the job. The created copy will not be transactionally consistent, but crash consistent.
- Select **Disable application processing** if you do not want enable quiescencing for the VM.

If you are copying VMs running database systems that use transaction logs, you can select to truncate transaction logs after the job so that they do not overflow the storage space. Use the **Truncation logs** section to define the scenario of transaction log handing:

- Select **Truncate logs on successful backup only** if you want Veeam Backup & Replication to trigger truncation of logs only after the job is finished successfully. In this case, Veeam runtime process will wait for the job to complete, and then will trigger truncation of transaction logs. If truncation of transaction logs is not possible for some reason, the logs will remain untouched in the VM guest OS until the next start of the Veeam runtime process.
- Select **Truncate logs immediately** if you want Veeam Backup & Replication to trigger truncation of logs in any case, no matter whether the job finishes successfully or fails.
- Select **Do not truncate logs** if you do not want Veeam Backup & Replication to truncate logs at all. This option is recommended if you are using another tool to perform guest-level copy, and this tool maintains consistency of the database state. In such scenario, truncation of logs with Veeam Backup & Replication will break the guest-level copy chain and cause it to fall out of sync.

Step 8. Define the Job Schedule

At the **Schedule** step of the wizard, you can select to manually run the VM copy job, schedule the job to start at a specific time (for example, the least busy hours to reduce impact on the virtual environment) or define a schedule for the job to run on a regular basis.

To specify the job schedule, select the **Run the job automatically** check box. If this check box is not selected, the job is supposed to be run manually.

You can define the following scheduling settings for the job:

- You can choose to run the job at specific time on defined week days, monthly and with specific periodicity.
- You can choose to run the job continuously. In this case, the next run of the backup job will be started as soon as the previous one is completed, maintaining your backup always in the most recent state. To run the job continuously, select the **Periodically every** option and choose **Continuously** from the list on the right.
- You can choose to run the job repeatedly throughout a day with a set time interval. At the **Schedule** step of the wizard, select the **Periodically every** option, enter the necessary time interval and select the necessary time unit: *Hours* or *Minutes*. Click **Schedule** on the right and use the time table to define the permitted time window for the job. If you choose to run the job at an hourly interval, in the **Start time for hourly jobs** field, specify the exact time when the job should start.
  For example, you want to start a job every 2 hours from 9AM to 6PM. At the **Schedule** step, select the **Periodically every** option, enter 2 in the field on the right and select *Hours* from the list. Click **Schedule** and use the **Permitted** and **Denied** options to mark the time window from 9AM to 6PM. In the **Start time for hourly jobs** field, specify the exact start time of the job, for example, 15 minutes. The job you have scheduled will start at 9:15 AM, 11:15 AM, 1:15 PM, 3:15 PM and 5:15 PM.
- You can chain jobs. In the common practice, jobs start one after another: when the job "A" finishes, the job "B" starts and so on. If you want to create a chain of jobs, you should define the time schedule for the first job in the chain. For all the rest of the jobs in the chain, at the **Schedule** step of the wizard, select the **After this job** option and choose the preceding job from the list. Note that the list contains jobs of all types: *Backup*, *Replication, VM Copy* and *File Copy*.

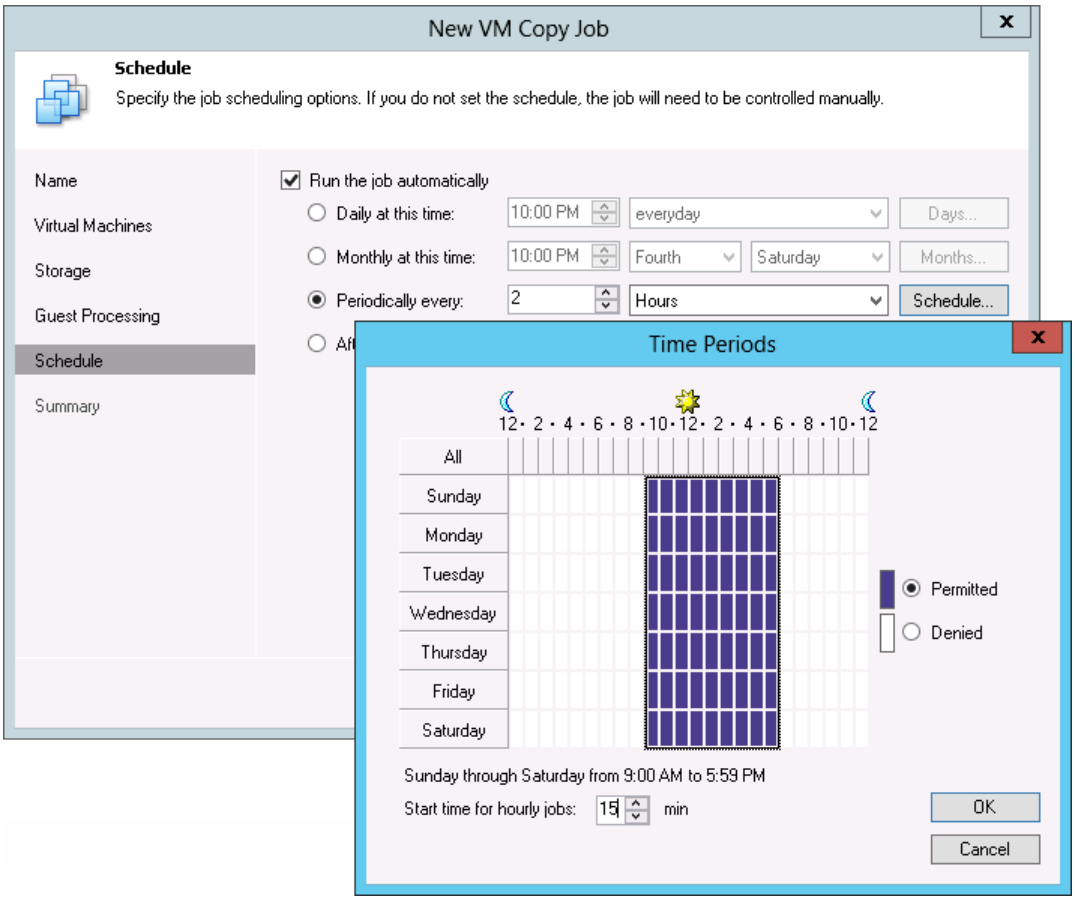

Note After you have created a scheduled job, you can temporarily disable it (that is, hold it for some time without changing the set time schedule). Right-click the job in the list and select **Disable Job** from the shortcut menu. To enable the job schedule, right-click the job and deselect **Disable Job** in the shortcut menu.

Step 9. Finish Working with the Wizard

After you have specified schedule settings, click **Create**. Select the **Run the job when I click Finish** check box if you want to start the created job right after you complete working with the wizard. Click **Finish** to close the wizard.

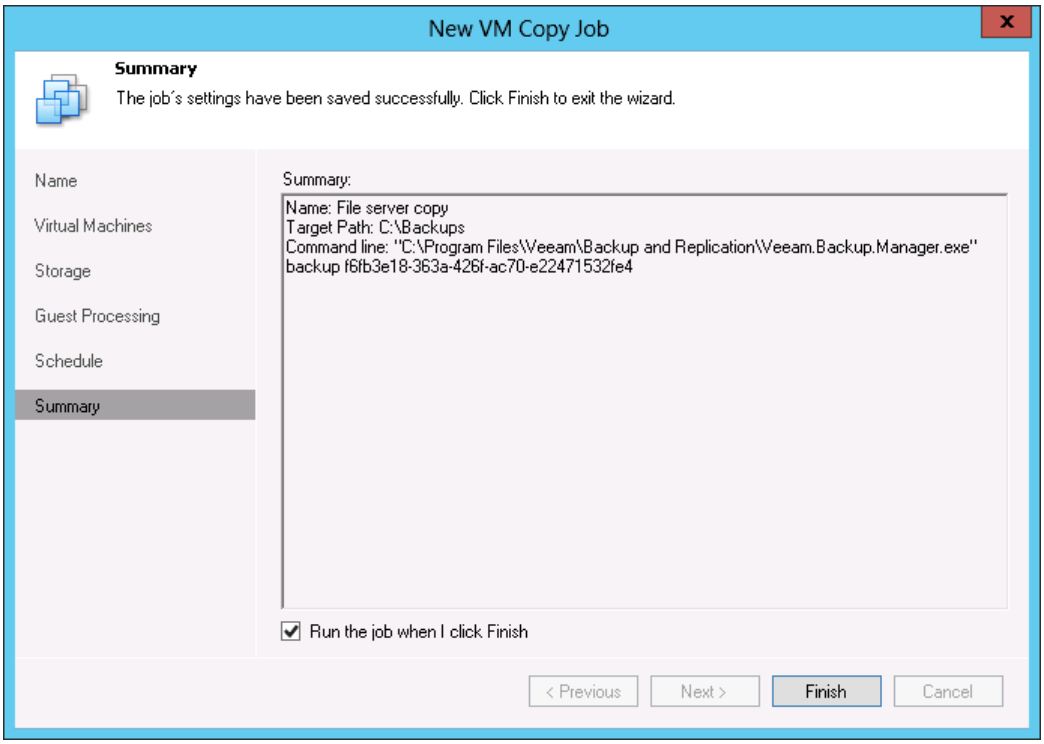

Creating VeeamZIP Files

With Veeam Backup & Replication, you can quickly perform backup of one or several VMs with VeeamZIP.

To create a backup of VMs with VeeamZIP, follow the next steps:

Step 1. Select VMs to Back Up

1. Open the **Virtual Machines** view.
2. In the infrastructure tree, click the ESX(i) host or a VM container in which the VM you want to back up resides.
3. In the working area, right-click one or several VMs you want to back up and select **VeeamZIP**. You can also select the necessary VM(s), click the **Virtual Machines** tab and click the **VeeamZIP** button on the ribbon.
   To quickly find the necessary VM, type the VM name or a part of it in the search field at the top of the working area and click the **Start search** button on the right or press **Enter** on the keyboard.

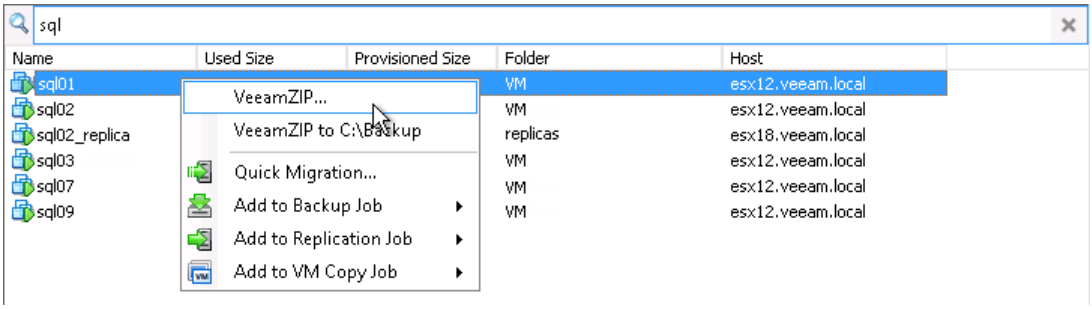

Step 2. Select Destination for the VeeamZIP File

In the **Destination** section, specify a location to which you want to store the VeeamZIP file.

- If you want to store the backup file to a specific backup repository, select the **Backup repository** option and choose the necessary repository from the list.

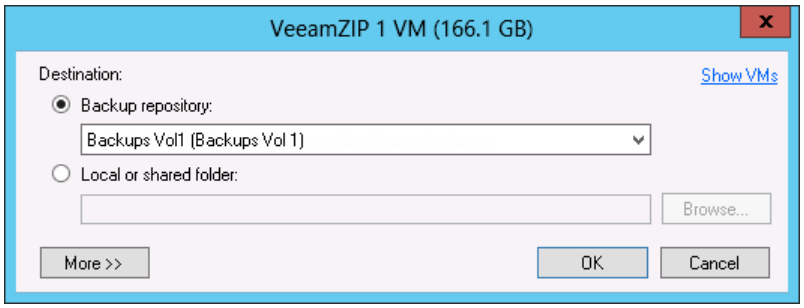

- If you want to store the backup file to a local folder on the Veeam backup server, select the **Local or shared folder** option, click **Browse** on the right and select the folder to which the backup file should be stored.

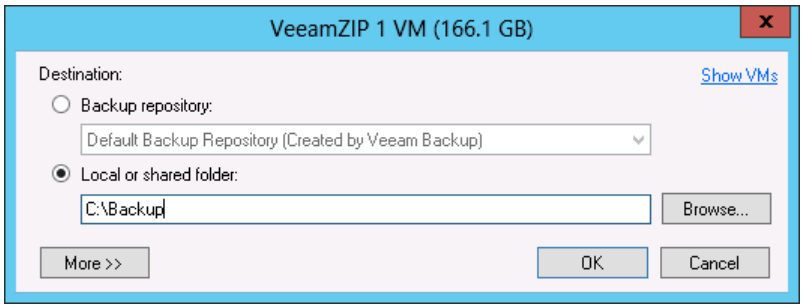

- If you want to store a backup file to a shared folder, select the **Local or shared folder** option and type in the UNC name of the shared folder to which the backup file should be stored. Keep in mind that the UNC name always starts with two back slashes (\\).
  If the shared folder requires authentication, select the necessary credentials from the **Credentials** list. If you have not specified credentials beforehand, click the **Manage accounts** link to set up credentials. Alternatively, you can click **Add** on the right to add credentials to the shared folder. To learn more, see Managing Credentials.

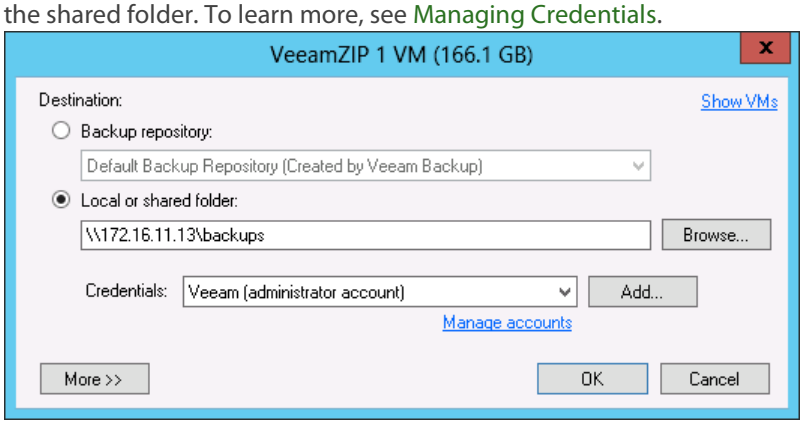

Step 3. Specify Backup Settings

To specify additional backup options, click **More**.

1. From the **Compression** level list, select a compression level for the created backup: *None*, *Dedupe-friendly*, *Optimal*, *High* or *Extreme*.
2. By default, Veeam Backup & Replication uses application-aware image processing to create a transactionally consistent backup of VMs running applications with VSS support. If you are backing up VMs that run other than Windows OS or applications without VSS support, you can disable this option by clearing **Disable guest quiescence** check box.

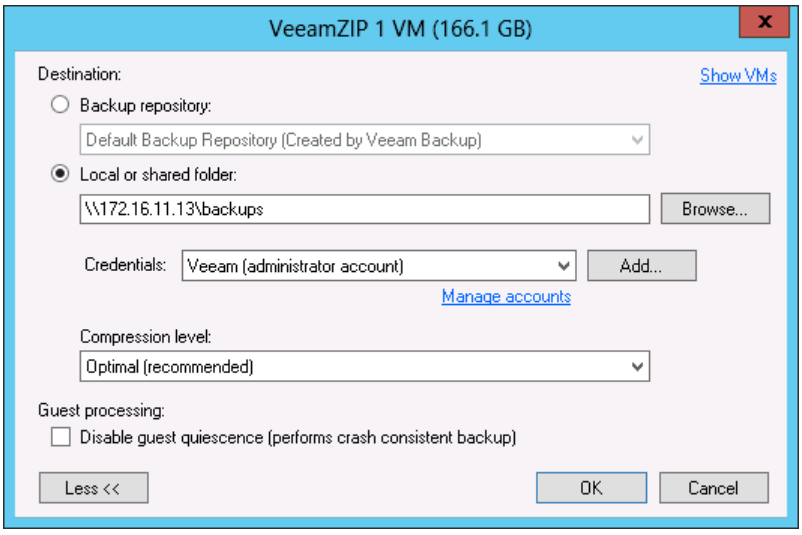

Step 4. Run the VeeamZIP Job

Click **OK**. The VeeamZIP job will start immediately.

As the job runs, you can track the job performance in the real-time mode. To see the job results once it completes, open the **History** view, expand the **Jobs** node, click **Backup;** then double-click the job session in the list.

As a result of the job processing, Veeam Backup & Replication will create a full backup file (.vbk) and store it to the specified destination. The VM name, date and time of the file creation are appended to the file name, so you can easily find the necessary VeeamZIP file afterwards.

Tip Veeam Backup & Replication keeps settings of the latest VeeamZIP job. To quickly create a VeeamZIP file with the same settings as the previous one and store it to the same location, right-click the necessary VM and select **VeeamZIP to** from the shortcut menu. The VeeamZIP job with the same settings as the previous job will start immediately.

Cloning Jobs

In Veeam Backup & Replication, you can add new jobs by means of job cloning. Job cloning allows you to create an exact copy of any job available in the job list. Configuration details of a created job copy are written to the same SQL database that stores configuration details of the original job. Thus, the job copy is available and can be managed via the Veeam Backup & Replication console.

If you want to create multiple jobs with similar settings, the recommended practice is to configure a set of jobs that will be used as 'job templates'. These job templates can then be cloned and edited as required.

To clone a job:

1. Open the **Backup & Replication** view.
2. Click the **Jobs** node in the inventory pane.
3. Select the necessary job in the working area and click **Clone** on the ribbon or right-click the job and select **Clone**.

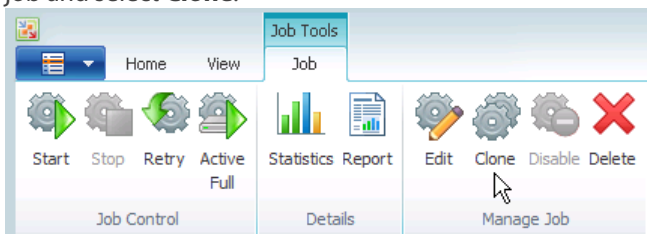

- The name of the cloned job is formed by the following pattern: *<job_name_clone1>*, where *job_name* is the name of the original job and *clone1* is a suffix added to the original job name. If you clone the same job again, the number in the name will be incremented, for example, *job_name_clone2*, *job_name_clone3* and so on. Once a job is cloned, you can edit all its settings, including its name.
- If the original job is scheduled to run automatically, the newly cloned job will be *Disabled*. To enable it, select it in the list and click **Disable** on the ribbon or right-click the job and select **Disable**.
- If the original job is configured to use a secondary target, the cloned job will be created without the secondary target settings.

Note The job cloning functionality in Veeam Backup & Replication is similar to the job cloning functionality in Veeam Backup Enterprise Manager. However, job cloning in Veeam Backup Enterprise Manager is limited to backup and replication jobs only; after you clone a job, you can edit only some of the job settings. In Veeam Backup & Replication, you can clone any type of job except SureBackup jobs and edit all settings of the cloned job.

Migrating Virtual Machines

To relocate one or more VMs to a new destination host or datastore, you can use Quick Migration. Note that quick migration is not job-driven, so it cannot be saved or scheduled to run later. Veeam Backup & Replication will start relocating VMs immediately after you finish working with the **Quick Migration** wizard.

This section will guide you through all steps of the **Quick Migration** wizard and provide explanation on available options.

Before You Begin

Prior to performing quick migration, make sure you have set up all necessary backup infrastructure components for it. Open the **Backup Infrastructure** view and check if the source and target hosts are available under the **Managed servers** node in the inventory pane and backup proxies are available under the **Backup Proxies** node and properly configured. You will not be able to add quick migration infrastructure components or change their configuration once the **Quick Migration** wizard is launched. For detailed information on adding Quick Migration infrastructure components, see Setting Up Backup Infrastructure.

To migrate VMware VMs, follow the next steps:

Step 1. Launch the Quick Migration Wizard

To run the **Quick Migration Wizard**, do either of the following:

- On the **Home** tab, click **Quick Migration**.
- Open the **Virtual Machines** view, right-click one or several VMs in the working area and select **Quick Migration**.

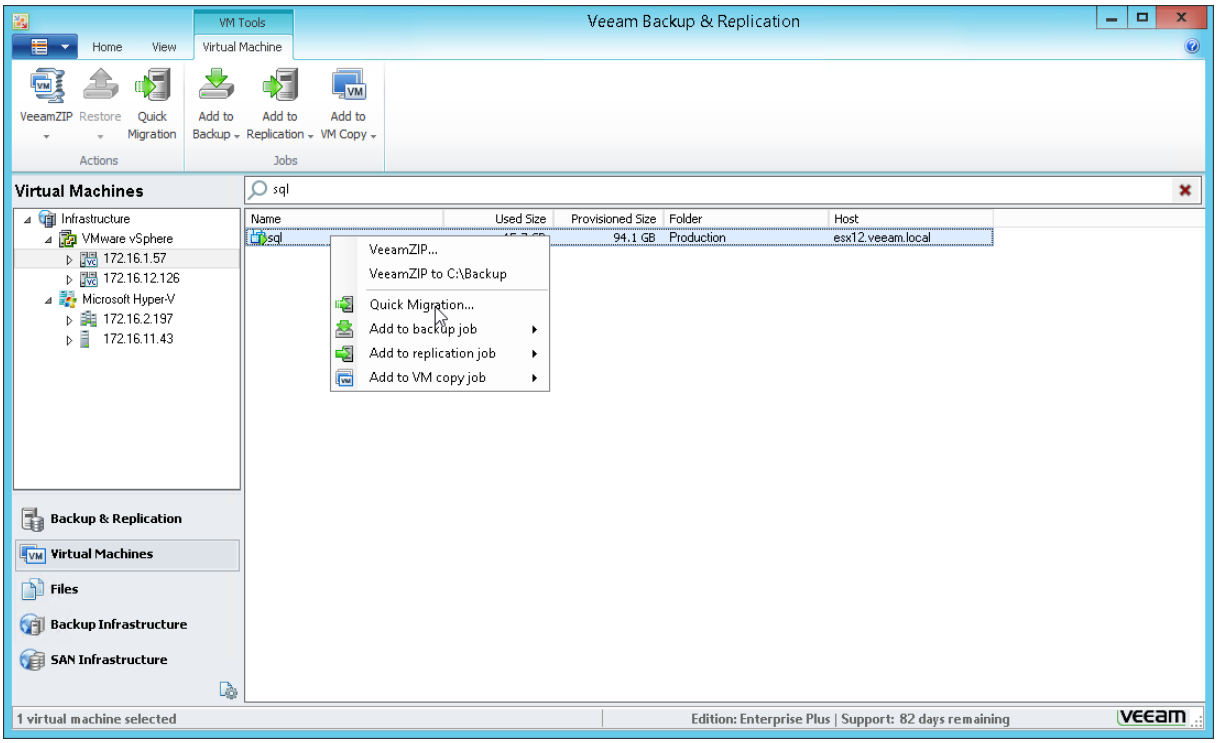

Step 2. Select Virtual Machines to Relocate

At this step, you should select individual VMs, multiple VMs or VM containers you want to relocate. Click **Add** to browse to VMs and VM containers that should be relocated. In the displayed tree, select the necessary object and click **Add**.

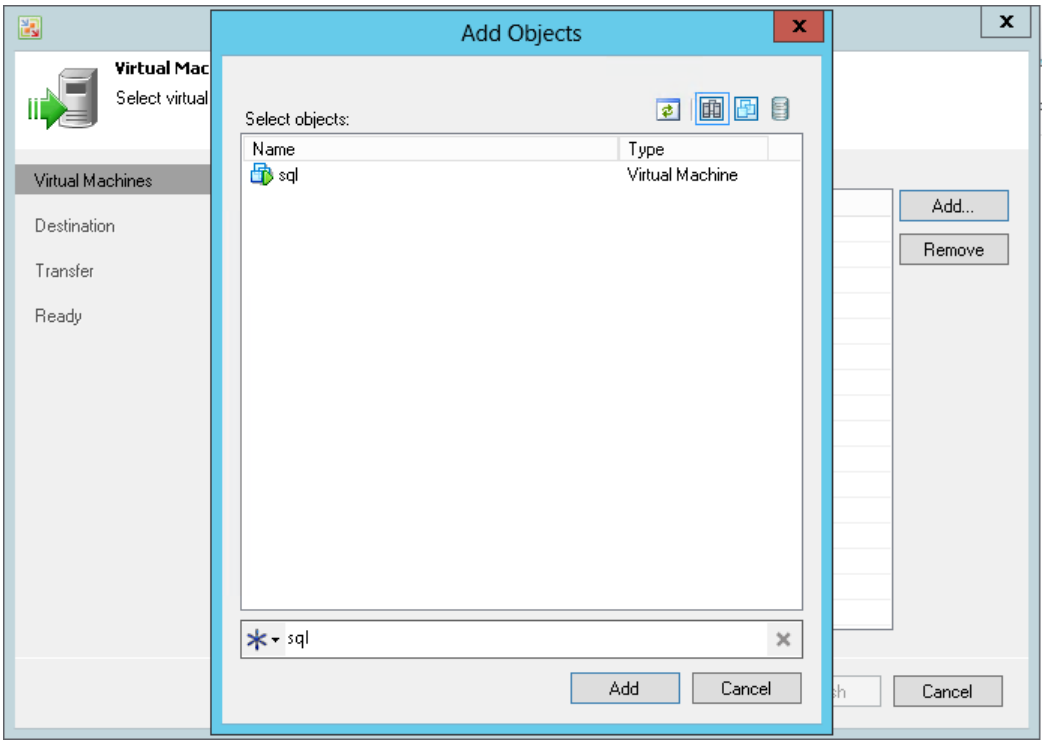

To facilitate objects selection, you can:

- Switch between views by clicking the **Hosts and Clusters**, **VMs and Templates** or **Datastores and VMs** buttons at the top of the tree.
- Use the search field at the bottom of the **Add Objects** window: click the button on the left of the field to select the necessary type of object that should be searched for (**Everything**, **Folder**, **Cluster**, **Host**, **Resource pool**, **VirtualApp** or **Virtual machine**), enter an object's name or a part of it and and click the **Start search** button on the right or press **[ENTER]**.

Note Depending on the view you select, some objects may not be available (for example, if you select the **VMs and Templates** view, you will not be able to see and find resource pools).

To remove an object from the list, select it and click **Remove** on the right.

Step 3. Specify VM Destination

At this step of the wizard, you should select the destination to which selected VMs should be relocated (host or cluster, resource pool, VM folder and datastore).

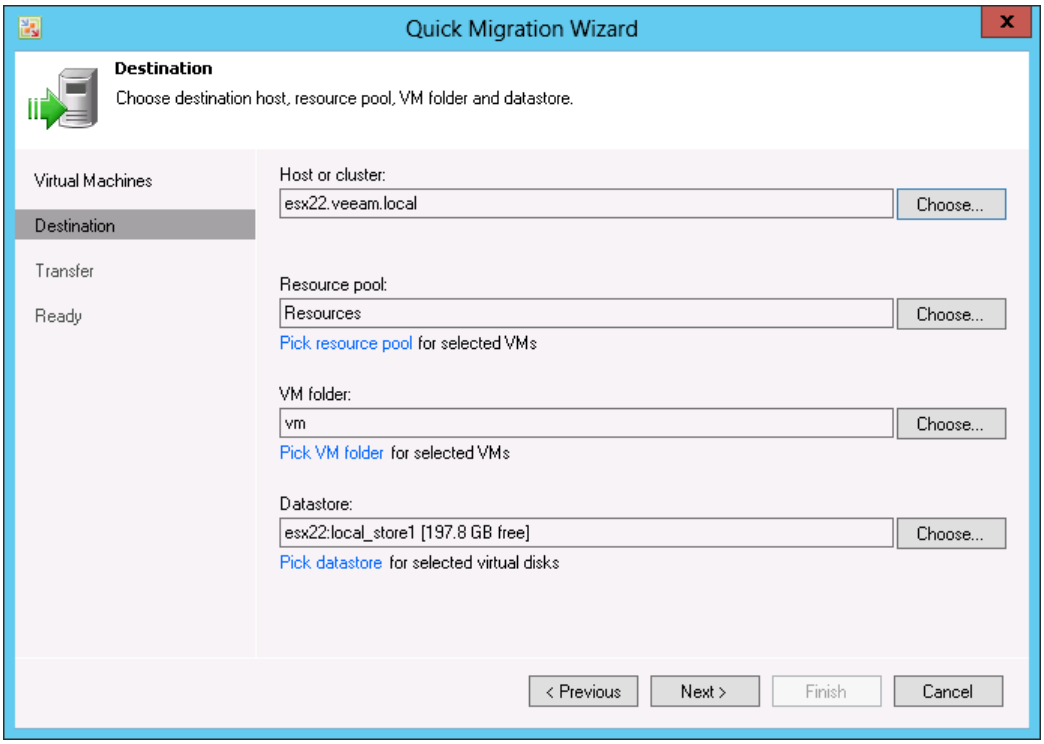

Host or cluster

Click **Choose** and select the host or cluster where relocated VMs should be registered. To facilitate selection, use the search field at the bottom of the window: click the button on the left of the field to select the necessary type of object that should be searched for (**Cluster** or **Host**), enter an object's name or a part of it and press **ENTER**.

Resource pool

Specify the resource pool in which relocated VMs should be included. If all or the majority of VMs should belong to the same pool, click **Choose** and select the destination resource pool. To facilitate selection, use the search field at the bottom of the window: enter the resource pool name or part of it and press **ENTER**.

If you want to relocate VMs to different resource pools, click the **Pick resource pool for selected VMs** link. In the **Choose resource pool** section, click **Add VM** on the right and select the necessary VMs. To map an added VM to a resource pool, select it in the **VM resource pool** list and click **Resource Pool** at the bottom of the window. From the list of available resource pools, choose the necessary one.

VM folder

Specify the folder to which relocated VMs should be placed. If all or the majority of VMs should be stored the same folder, click **Choose** and select the destination folder. To facilitate selection, use the search field at the bottom of the window: enter the folder name or a part of it and press **ENTER**.

If you want to relocate VMs to different VM folders, click the **Pick VM folder for selected VMs** link. In the **Choose Folder** section, click **Add VM** on the right and select the necessary VMs. To map an added VM to a VM folder, select it in the **VM folder** list and click **VM Folder** at the bottom of the window. From the list of available folders, select the necessary one.

Datastore

Specify the datastore where files of relocated VMs should be stored. If all or the majority of VMs will be stored on the same datastore, click **Choose** and select the destination datastore. Note that Veeam Backup & Replication will display only those datastores that are accessible by the selected ESX(i) host. If you have chosen to replicate VMs to a cluster, Veeam Backup & Replication will display only shared datastores.

To facilitate selection, use the search field at the bottom of the window: enter the datastore name or a part of it and press **ENTER**.

If you want to relocate VMs to different datastores, click the **Pick datastore for selected virtual disks** link. In the **Files location** section, click **Add VM** on the right and select VMs that should be mapped to datastores. To map an added VM to a datastore, select it in the **Files location** list and click **Datastore** at the bottom of the window. From the list of available datastores, select the necessary one.

Additionally, you can choose to store VM configuration files and disk files in different locations. To do so, add the VM to the **Files location** list, expand it and select the required type of files. Click **Datastore** at the bottom of the window and choose the destination for the selected type of files.

By default, Veeam Backup & Replication preserves the format of relocated VM disks. If necessary, you can configure the job to change the disk format. For example, if a VM uses thick disks, you can change the format to thin provisioned, and thus, save on disk space required to store the relocated VM. To change disk format, add a VM to the **Files location** list, select it and click **Disk type** at the bottom of the window. In the **Disk Type Settings** section, choose the format for disks of the relocated VM: same as the source disk, thin or thick. Please note that disk format change is available only for VMs using virtual hardware version 7 or later.

Step 4. Select Infrastructure Components for Data Transfer

At this step of the wizard, you should assign infrastructure components to migrate selected VMs. In the **Data transfer** section, select backup proxies that will be used to transfer VM data from source to target.

If you plan to perform migration of VMs within one site, the same server can act as the source proxy and target proxy. For offsite migration, you must deploy at least one proxy server in each site to establish a stable connection across the sites for data transfer. Click **Choose** next to the **Source proxy** and **Target proxy** fields to select backup proxies for migration. In the **Backup Proxy** section, you can choose automatic proxy selection or assign proxies explicitly.

- If you choose **Automatic selection**, Veeam Backup & Replication will detect backup proxies that are connected to the source datastore and will automatically assign optimal proxy resources for processing VM data.
  VMs selected for migration are processed one by one. Before processing a new VM in the VM list, Veeam Backup & Replication checks available backup proxies. If more than one proxy is available, Veeam Backup & Replication analyzes transport modes that the proxies can use for data retrieval and the current workload on the proxies to select the most appropriate resource for VM processing.
- If you choose **Use the backup proxy servers specified below**, you can explicitly select proxies that will be used to perform migration.

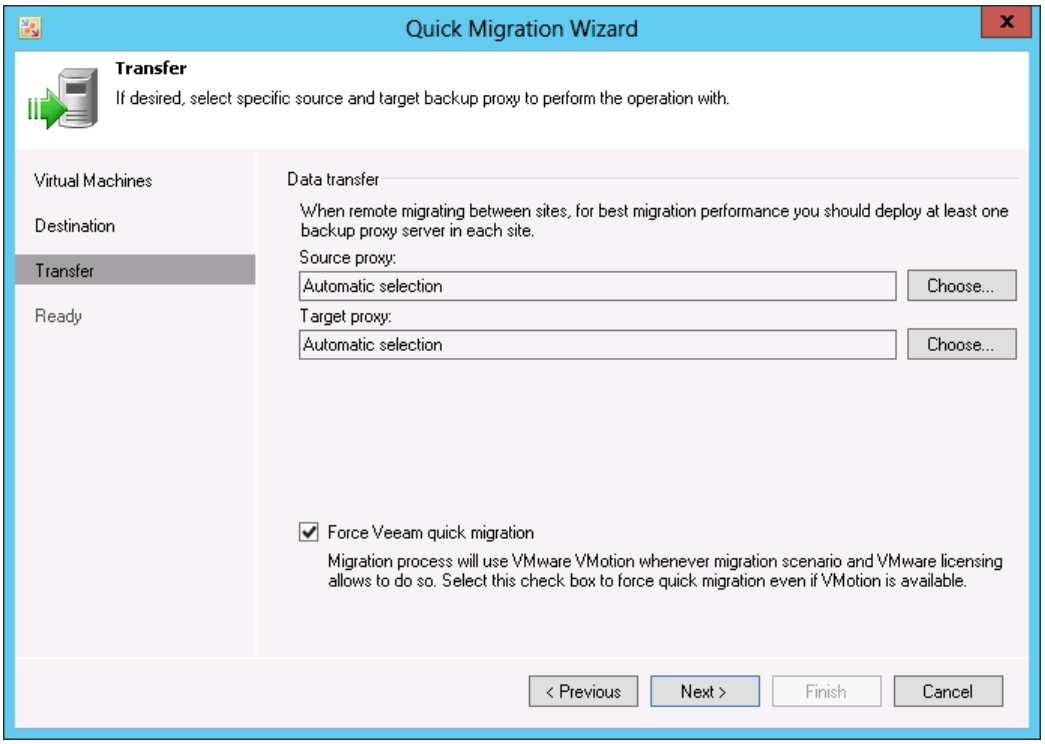

You can optionally select which migration mechanism to use: VMware vMotion or Veeam Quick Migration.

- If you want to use VMware vMotion to relocate the VM(s), leave the **Force Veeam quick migration** check box not selected. In this case, Veeam Backup & Replication will attempt to use the VMware vMotion mechanism to migrate the VM(s). If VMware vMotion cannot be used for some reason, Veeam Backup & Replication will use its own migration mechanism. Note that Veeam Backup & Replication can use VMware vMotion only if your VMware license covers this functionality.
- If you do not want to use VMware vMotion for some reason, select the **Force Veeam quick migration** check box. Veeam Backup & Replication will use its own migration mechanism.

Step 5. Finish Working with the Wizard

After you have specified migration settings, click **Next**. Veeam Backup & Replication will check whether selected VMs can be relocated.

By default, when VM migration completes successfully, Veeam Backup & Replication waits for a heartbeat signal from the VM on the target host. If the heartbeat is received, the original VM on the source host is deleted. If you do not want to delete the original VM on the source host, clear the **Delete source VM files upon successful migration** check box.

Click **Finish** to start migration.

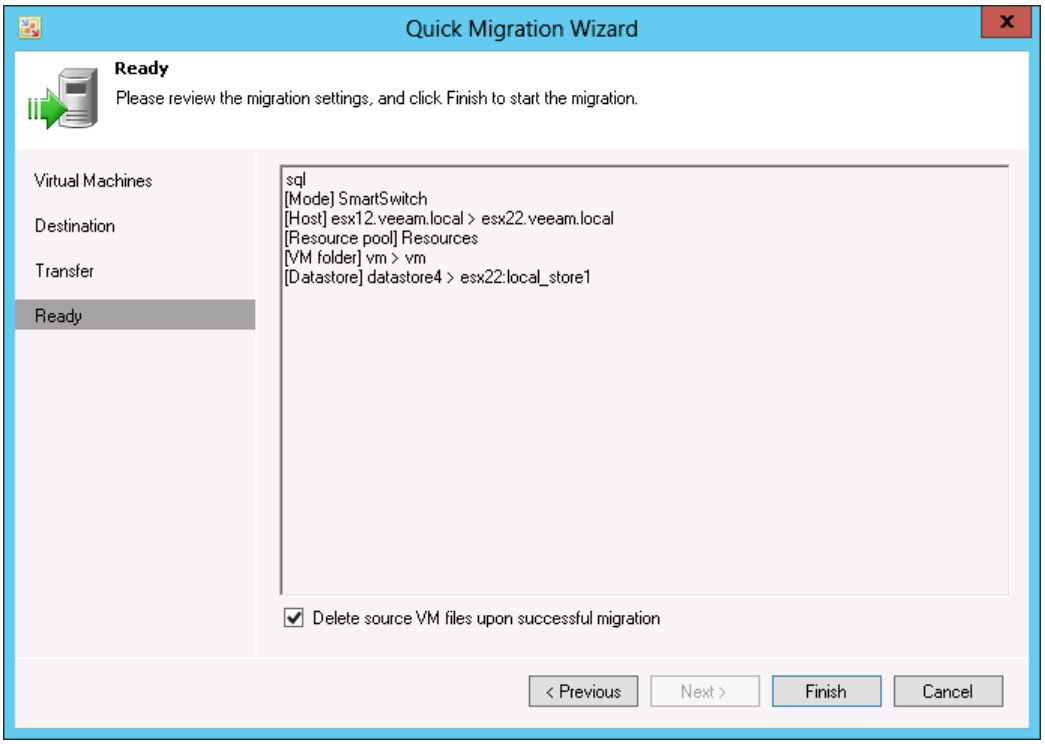

Performing Recovery Verification

SureBackup and SureReplica are Veeam's technologies developed to automate and simplify the recovery verification process — one of the most crucial parts of data management and protection. SureBackup and SureReplica let users verify the recoverability of every VM backup and replica, without additional hardware or administrative time and effort.

Verifying Backups and Replicas with SureBackup

To test VM backups and VM replicas with SureBackup recovery verification, complete the following tasks:

1. Create an application group for a verified VM
2. Create a virtual lab in which a verified VM from the backup or VM replica will be tested
3. Create and start SureBackup job

Important! You can verify only VM replicas in the *Normal* state. If a VM replica is in the *Failover* or *Failback* state, you will not be able to run a SureBackup job for it.

Creating an Application Group

To create a new application group, do either of the following:

- Open the **Backup Infrastructure** view, select the **Application Groups** node under **SureBackup** in the inventory pane and click **Add App Group** > **VMware** on the ribbon.
- Open the **Backup Infrastructure** view, right-click the **Application Groups** node under **SureBackup** in the inventory pane and select **Add Application Group** > **VMware**.

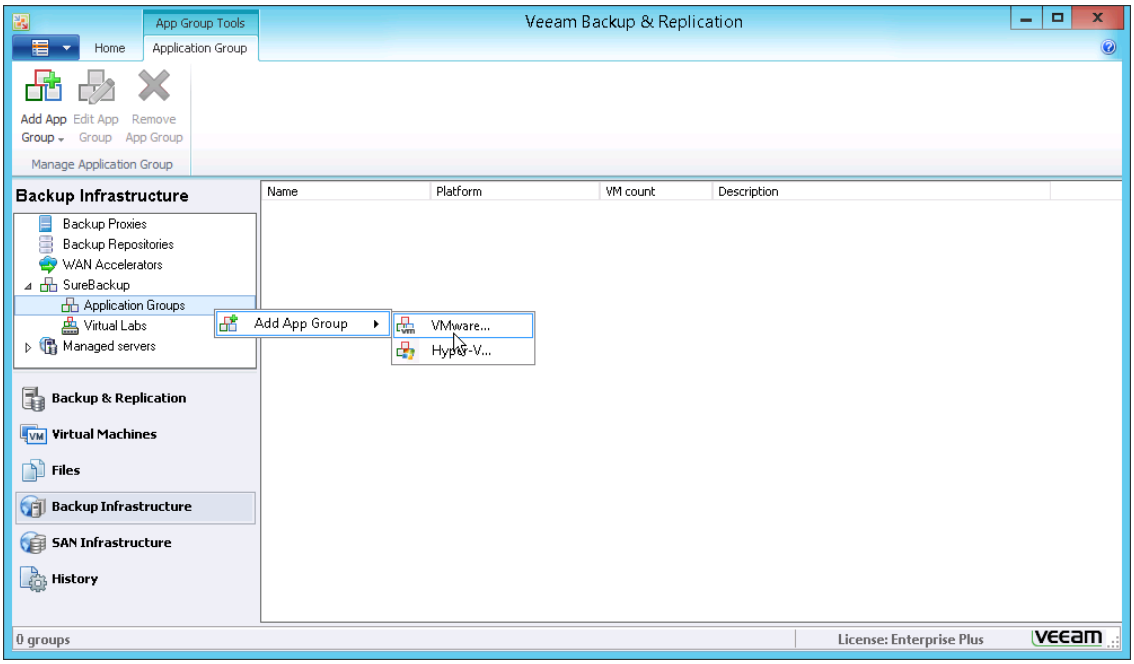

Then follow the steps of the **New Application Group** wizard:

Step 1. Specify an Application Group Name and Description

At the **Name** step of the wizard, enter a name and description for the new application group. The default description contains information about the user who created the group, as well as the date and time when the group was created.

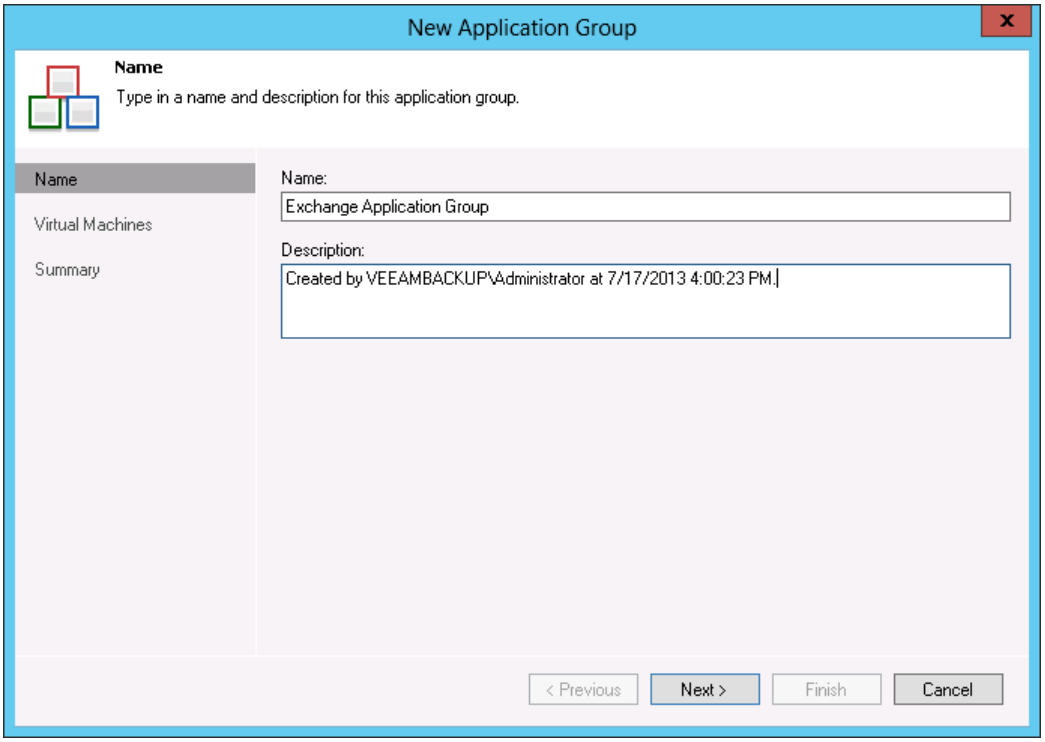

Step 2. Select VMs

At the **Virtual Machines** step of the wizard, you should add VMs to the created application group. You can add VMs from VM backups or VM replicas.

Important! Veeam Backup & Replication supports mixed application groups. You can add to the same application groups VMs from backups and VMs from replicas. Keep in mind that all VMs from the application group must belong to the same platform — VMware or Hyper-V, and must have at least one valid restore point created by the time the SureBackup job starts.

Adding VMs from VM Backups

To add a VM from the backup to the group, click **Add Backup** and select where to browse for VMs:

- **From Infrastructure** to browse the virtual environment. Because VMs from the application group are started from VM backups, you need to make sure that VMs you have selected for the group have been successfully backed up at least once by the time you plan to run a SureBackup job.
- **From Backups** to browse existing backups.

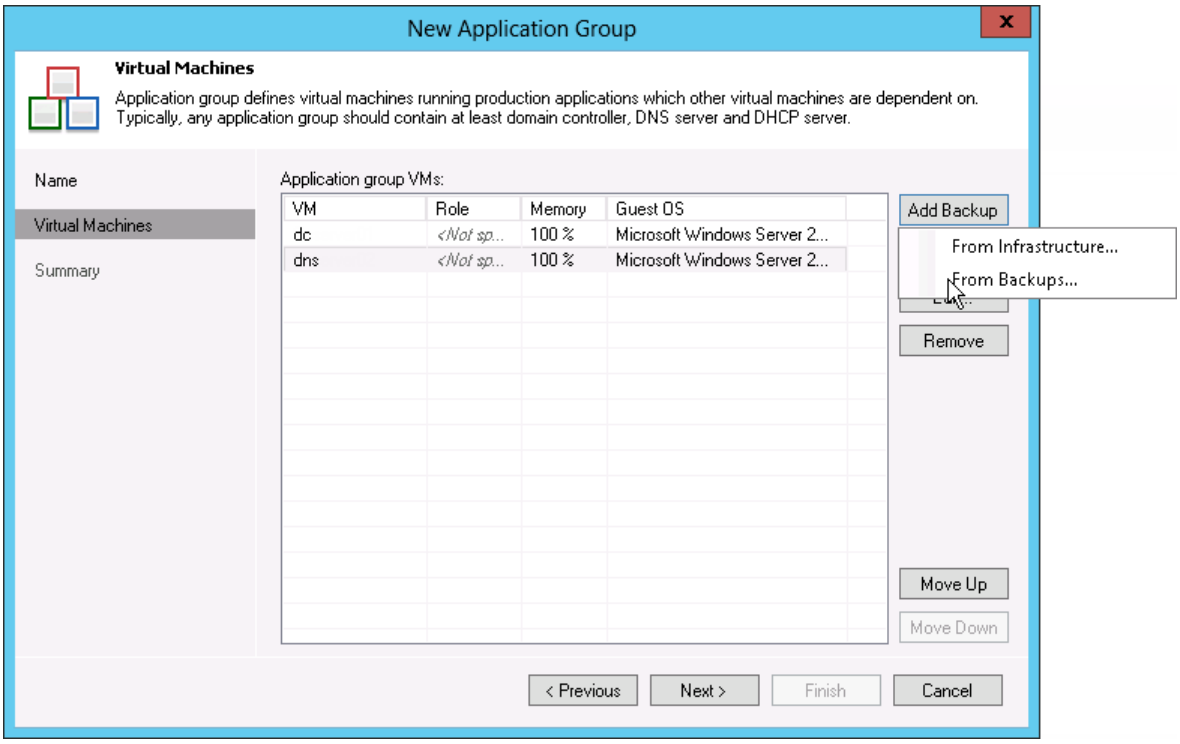

Adding VMs from VM Replicas

To add a VM replica to the group, click **Add Replica** and select where to browse for VMs:

- **From Infrastructure** to browse the virtual environment. Because VMs from the application group are started from VM replicas, you need to make sure that VMs you have selected for the group have been successfully replicated at least once by the time you plan to run a SureBackup job.
- **From Replicas** to browse existing VM replicas.

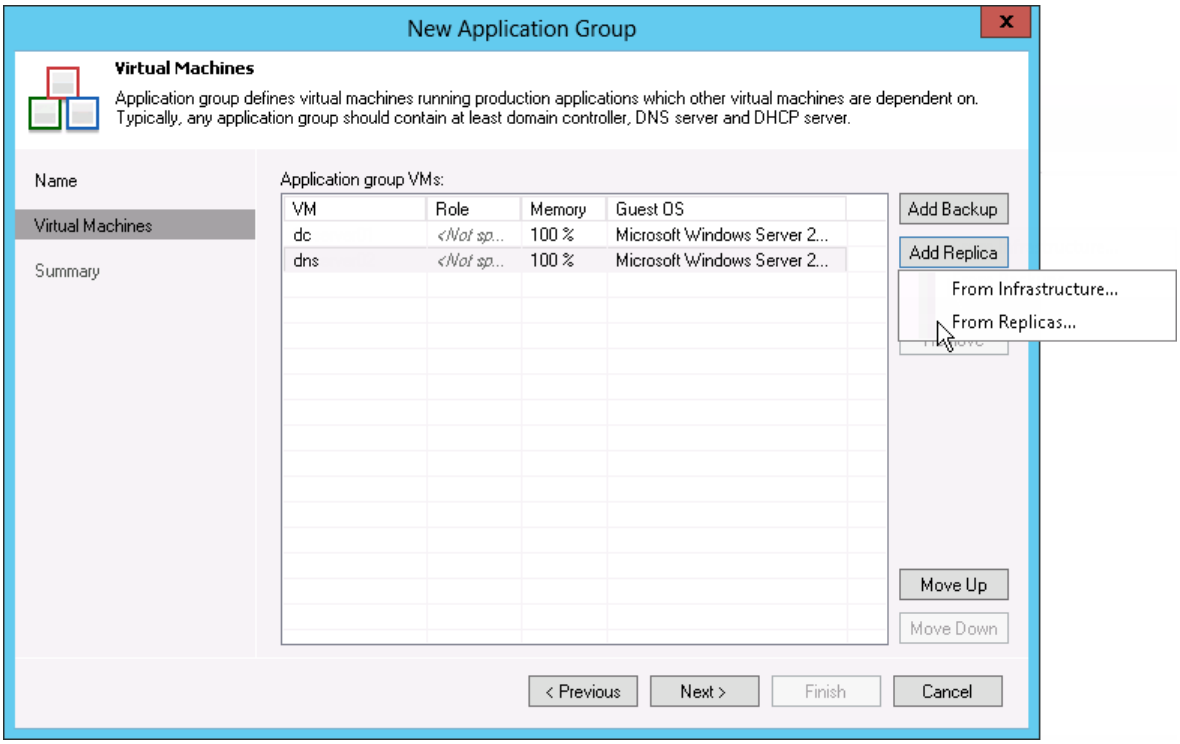

VMs in the list are specified in the order of their boot priority. To move a VM up or down in the list, select it and click the **Move Up** or **Move Down** button.

To remove a VM from the list, select it and click **Remove**.

Step 3. Specify Recovery Verification Options and Tests

After you have added necessary VMs to the application group, you should specify a role, VM startup options and select tests to be performed for each VM in the application group.

Important! To be able to perform tests, Veeam Backup & Replication requires VMware Tools to be installed on the verified VM. If VMware Tools are not installed, the VM will be started but tests will not be performed. VMs without VMware Tools can still be used as auxiliary VMs that should be started to enable proper work of verified VMs or VM replicas. In this case, you do not need to select any role for such auxiliary VM.

Select the necessary VM in the list and click **Edit** on the right.

Role settings

On the **Role** tab, select the role that a VM performs. Veeam Backup & Replication offers the following predefined roles for VMs:

- DNS Server
- Domain Controller
- Global Catalog
- Mail Server
- SQL Server
- Web Server

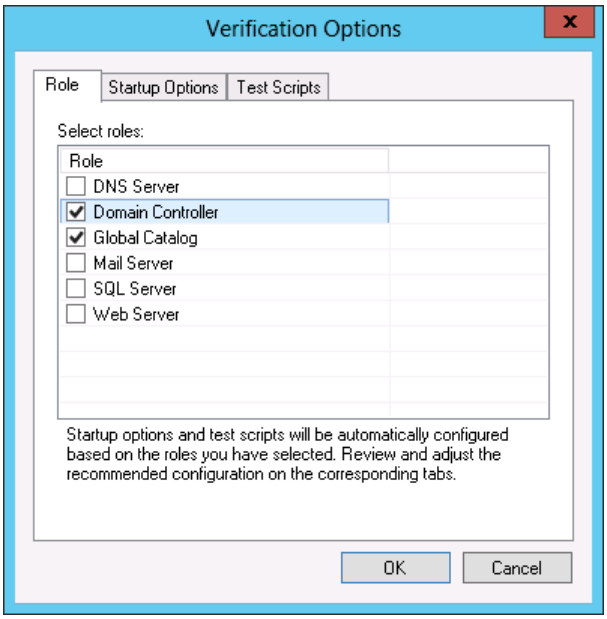

VM roles are described in XML files stored in the *%Program Files%\Veeam\Backup and Replication\Backup\SbRoles* folder. You can add your own roles. To do this, you need to create a new XML file and specify role and test scripts settings in it. For details, see Creating XML files with VM Roles Description.

Once you select the necessary role, Veeam Backup & Replication will automatically configure startup options and provide predefined test scripts applicable for the chosen role. You can use these settings or specify custom ones using the **Startup Options** and **Test Scripts** tabs.

To verify VMs that perform roles other than those specified in the list, you will have to manually configure startup options and specify test scripts to be run against these VMs.

Startup Options

On the **Startup Options** tab, specify VM startup settings:

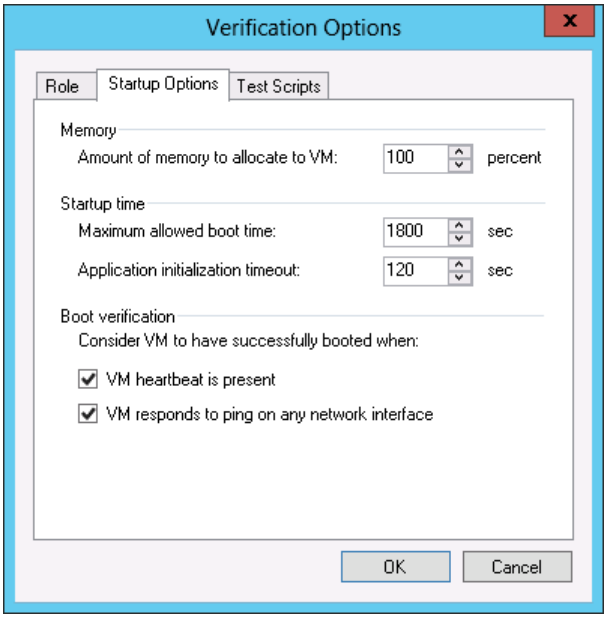

1. In the **Memory** section, specify the amount of memory that you want to pre-allocate to the VM on the system boot. The amount of pre-allocated memory is specified in percent. The percentage rate is calculated based on the system memory level available for the production VM. For example, if 1024 MB of RAM is allocated to the VM in the production and you specify 80% as a memory rate, 820 Mb of RAM will be allocated to the verified VM on startup.
2. In the **Startup time** section, specify the allowed boot time for the VM and timeout to initialize applications on it.
3. In the **Boot verification** section, specify when the VM should be considered to have been booted successfully:
    - VM heartbeat is present. If you select this option, Veeam Backup & Replication will perform a heartbeat test for the verified VM.
    - VM responds to ping on any network adapter. If you select this option, Veeam Backup & Replication will perform a ping test for the verified VM.
      Before you start a SureBackup job, make sure that the firewall on the tested VM allows ping requests.

Important! Be careful when specifying the **Maximum allowed boot time** value. Typically, a VM started by a SureBackup job requires more time to boot than a VM started regularly. If an application fails to be initialized within the specified interval of time, the recovery verification process will be finished with the timeout error. If such error situation occurs, you will need to increase the **Maximum allowed boot time** value and start the job again.

Test Scripts

On the **Test Scripts** tab, specify what test scripts should be run to verify a VM. When you select a VM role, Veeam Backup & Replication automatically assigns a predefined script that will be run to verify applications inside the VM.

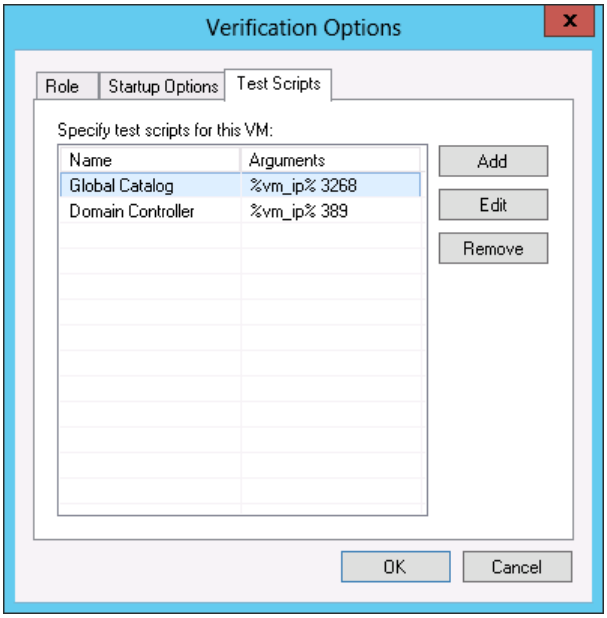

If you want to verify a VM that has some other role not listed on the **Role** tab, do the following:

1. Click **Add**.
2. In the **Test Script** window, select **Use the following test script**.
3. In the **Name** field, specify a name for the script.
4. In the **Path** field, define a path to an executable script file that should be run to verify the VM.
    - If you have your own custom script, define a path to it.
    - If you do not have a custom script, you can use Veeam's standard utility, *Veeam.Backup.ConnectionTester.exe*, that probes application communication ports. The utility is located in the installation folder of Veeam Backup & Replication: *%Program Files%\Veeam\Backup and Replication\Backup\Veeam.Backup.ConnectionTester.exe*. Specify this path in the **Path** field.
5. In the **Arguments** field, specify an IP address of the tested VM and the port that you want to probe. You can use the *%vm_ip%* variable to define the VM IP address or the *%vm_fqdn%* variable to define the VM fully qualified domain name.
6. Click **OK** to add the configured test.

To edit test settings, select the test in the list and click **Edit**. To delete a test, select it in the list and click **Remove**.

Note If a VM performs several roles running a number of applications at once, you can verify their work by adding several verification scripts. For such VMs, it is recommended to specify the maximum startup timeout value and allocate the greatest amount of memory.

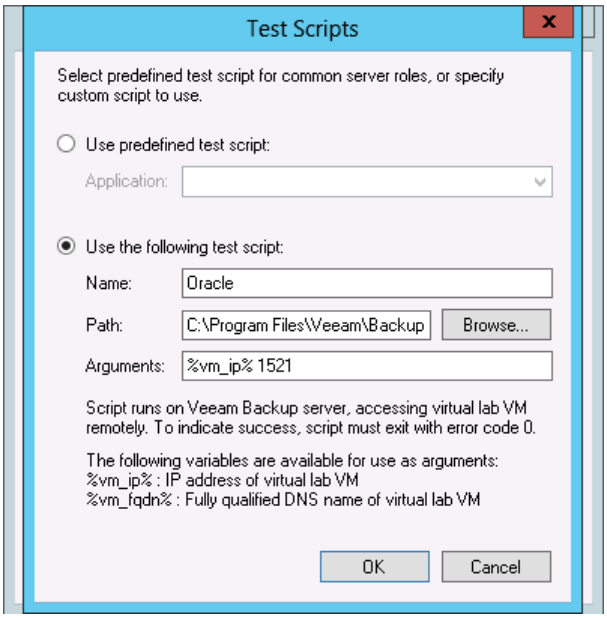

Step 4. Review the Application Group Summary and Finish Working with Wizard

After the group is created, review the application group summary and click **Finish** to exit the wizard.

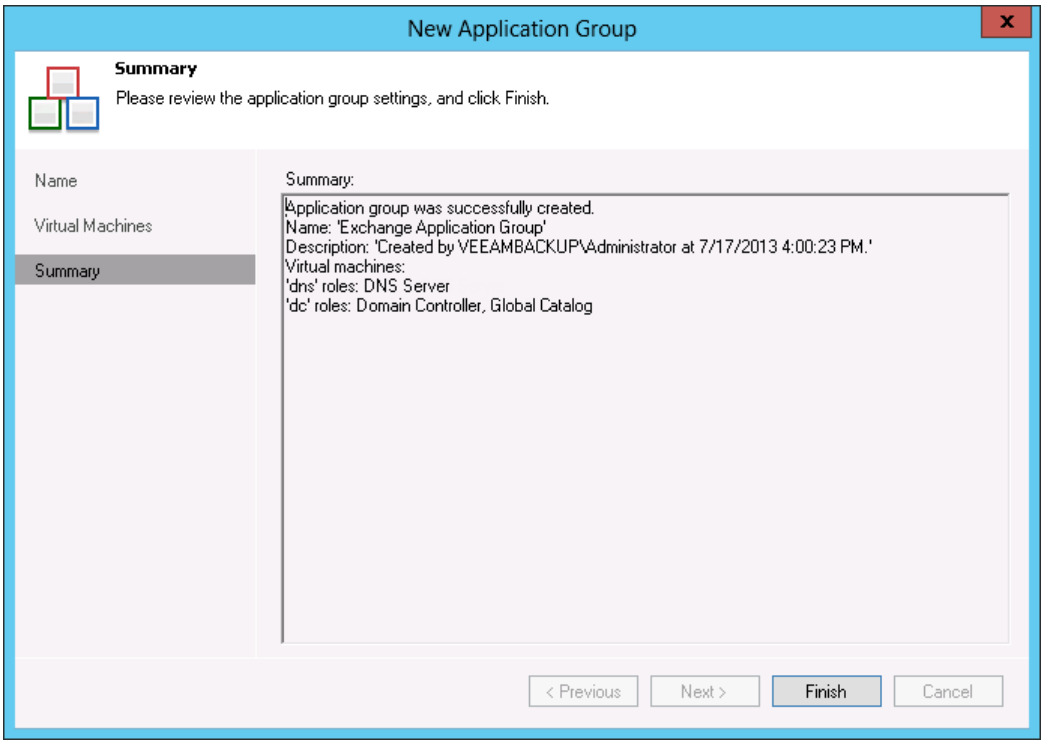

Creating a Virtual Lab

To create a new virtual lab, do either of the following:

- Open the **Backup Infrastructure** view, select the **Virtual Labs** node under **SureBackup** in the inventory pane and click **Add Virtual Lab** > **VMware** on the ribbon.
- Open the **Backup Infrastructure** view, right-click the **Virtual Labs** node under **SureBackup** in the inventory pane and select **Add Virtual Lab** > **VMware**.

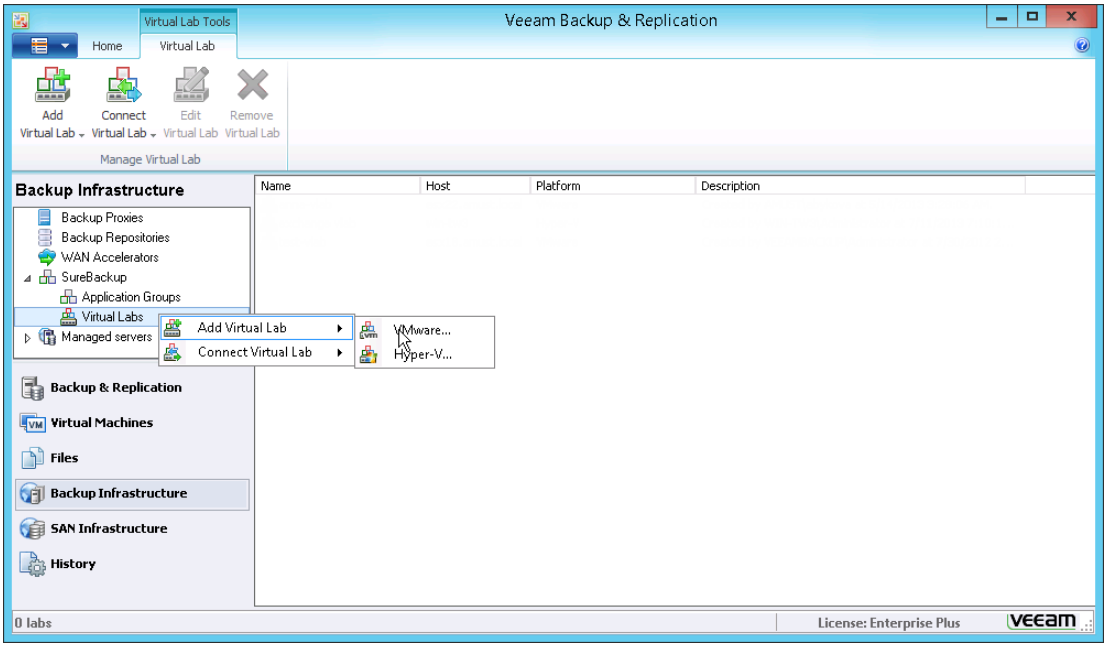

Follow the steps of the **New Virtual Lab** wizard:

Step 1. Specify a Name and Description

At the **Name** step of the wizard, enter a name and description for the new virtual lab. The default description contains information about the user who created the lab, as well as the date and time when the lab was created.

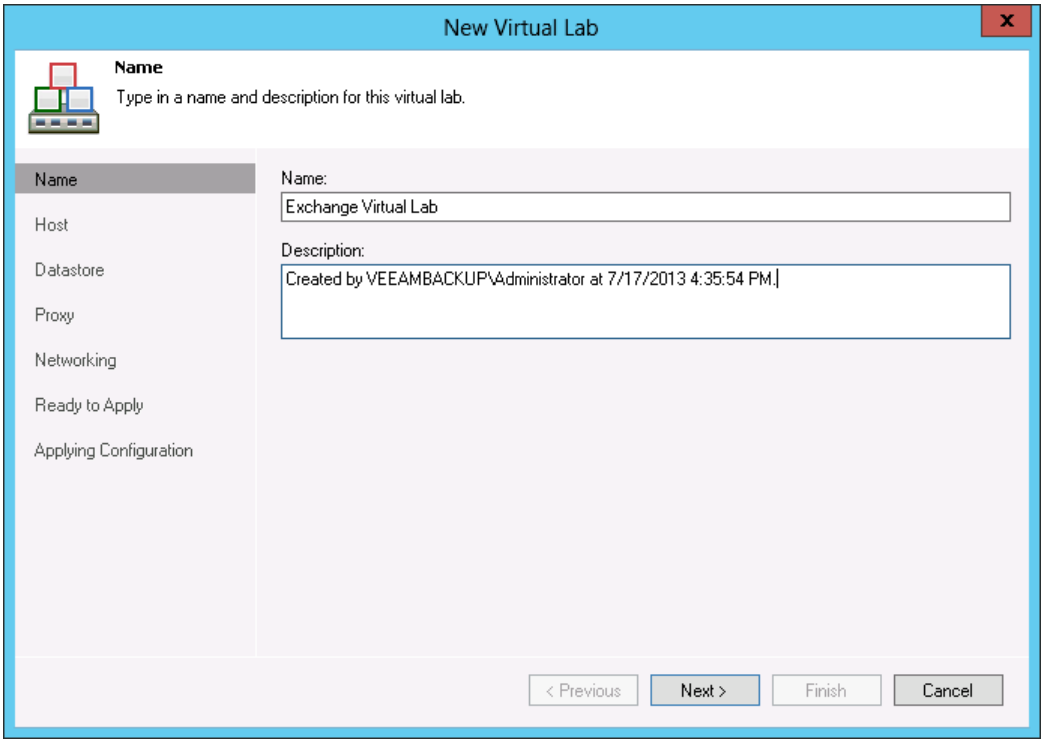

Step 2. Select a Host

At the **Host** step of the wizard, select an ESX(i) host on which the virtual lab should be created. This step differs for virtual labs used for VM backups and VM replicas verification.

Selecting an ESX(i) Host for VM Backups Verification

To select an ESX(i) host for VM backups recovery verification:

1. Click **Choose** to select an ESX(i) host on which the new virtual lab will be created. You can select a standalone ESX(i) host or an ESX(i) host being part of a cluster or vCenter Server hierarchy.
2. For every new virtual lab, Veeam Backup & Replication creates a dedicated folder and a resource pool to which all tested VMs, VMs from the application group and the proxy appliance are placed. By default, the folder and the pool have the same name as the virtual lab. To change the name of the destination folder and/or resource pool, click **Configure** and enter the necessary names in the **Destination Options** section.

Important! You cannot create resource pools in clusters with disabled DRS. If the destination host is a part of such a cluster, the **Create a designated resource pool** option will be disabled in the **Destination Options** window. For details, see http://kb.vmware.com/kb/1004098.
You cannot create folders on standalone ESX(i) hosts or ESX(i) hosts that are managed by the vCenter Severs but are added to Veeam Backup & Replication console as standalone hosts. To overcome this situation, add the corresponding vCenter Server to the Veeam Backup & Replication console.

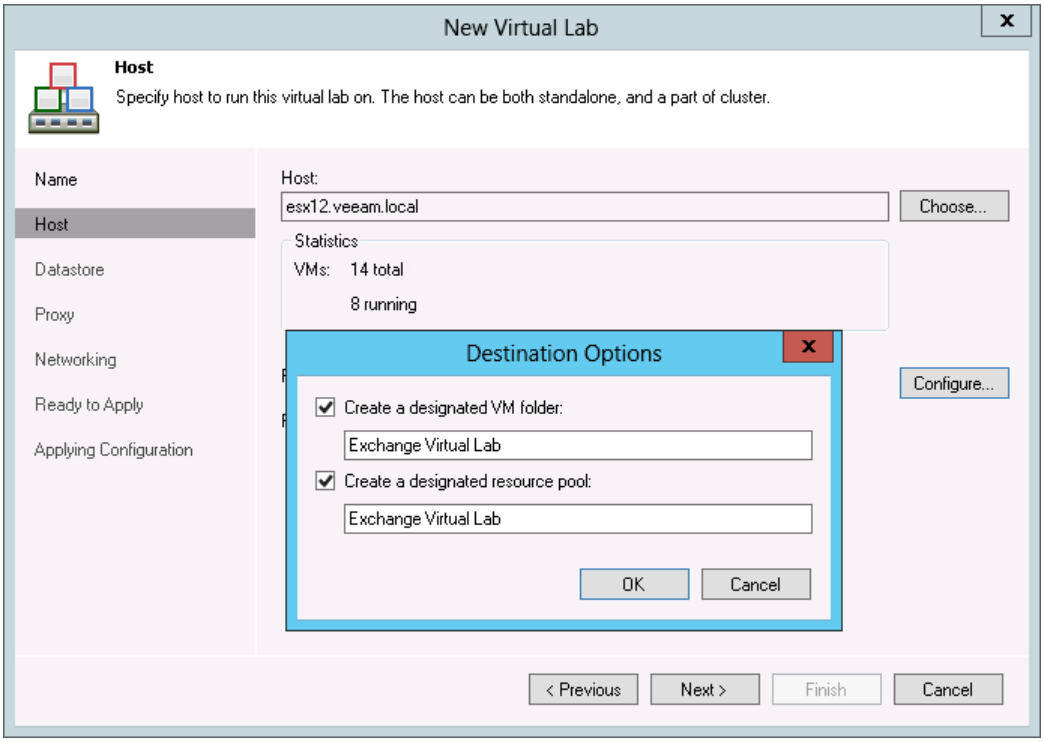

Selecting an ESX(i) Host for VM Replicas Verification

For VM replicas recovery verification, click **Choose** and select an ESX(i) host on which the new virtual lab should be created. You can select a standalone ESX(i) host or an ESX(i) host being part of a cluster or vCenter Server hierarchy.

When selecting an ESX(i) host for the virtual lab, mind the location of verified VM replicas and VM replicas from the application group:

- If all VM replicas you plan to verify and VM replicas from the application group are located on the same ESX(i) host, you should select the ESX(i) host on which these VM replicas are registered. In this case, the virtual lab, verified VM replicas and VMs from the application group will be started on the selected ESX(i) host. If the application group contains VMs added from VM backups, these VMs will also be started on the selected ESX(i) host.
  For this type of virtual lab configuration, you need to choose one of single-host networking modes: *Basic single-host* or *Advanced single-host*. To learn more, see Selecting a Networking Mode.
- If VM replicas you plan to verify and/or VM replicas from the application group are located on different ESX(i) hosts, you should select any ESX(i) host in your virtual environment. In this case, Veeam Backup & Replication will create the virtual lab on the selected ESX(i) host. Verified VM replicas and VM replicas from the application group will be started on ESX(i) hosts where they are registered and connected to the virtual lab utilizing VMware's DVS technology.

The ESX(i) host on which the virtual lab is created must meet the following requirements:

- The ESX(i) host must be located in the same datacenter where VM replicas are registered.
- The ESX(i) host must have enough CPU and RAM resources: in case the application group contains VMs that are started from backups, these VMs will be started on the selected ESX(i) host.

For this type of virtual lab configuration, you need to use the *Advanced multi-host* networking mode. To learn more, see Selecting a Networking Mode.

For every new virtual lab, Veeam Backup & Replication creates a dedicated folder and a resource pool on the selected ESX(i) host. By default, the folder and the pool have the same name as the virtual lab. To change the name of the destination folder and/or resource pool, click **Configure** and enter the necessary names in the **Destination Options** section.

Important! You cannot create resource pools in clusters with disabled DRS. If the destination host is a part of such a cluster, the **Create a designated resource pool** option will be disabled in the **Destination Options** window. For details, see http://kb.vmware.com/kb/1004098.
You cannot create folders on standalone ESX(i) hosts or ESX(i) hosts that are managed by the vCenter Severs but are added to Veeam Backup & Replication console as standalone hosts. To overcome this situation, add the corresponding vCenter Server to the Veeam Backup & Replication console.

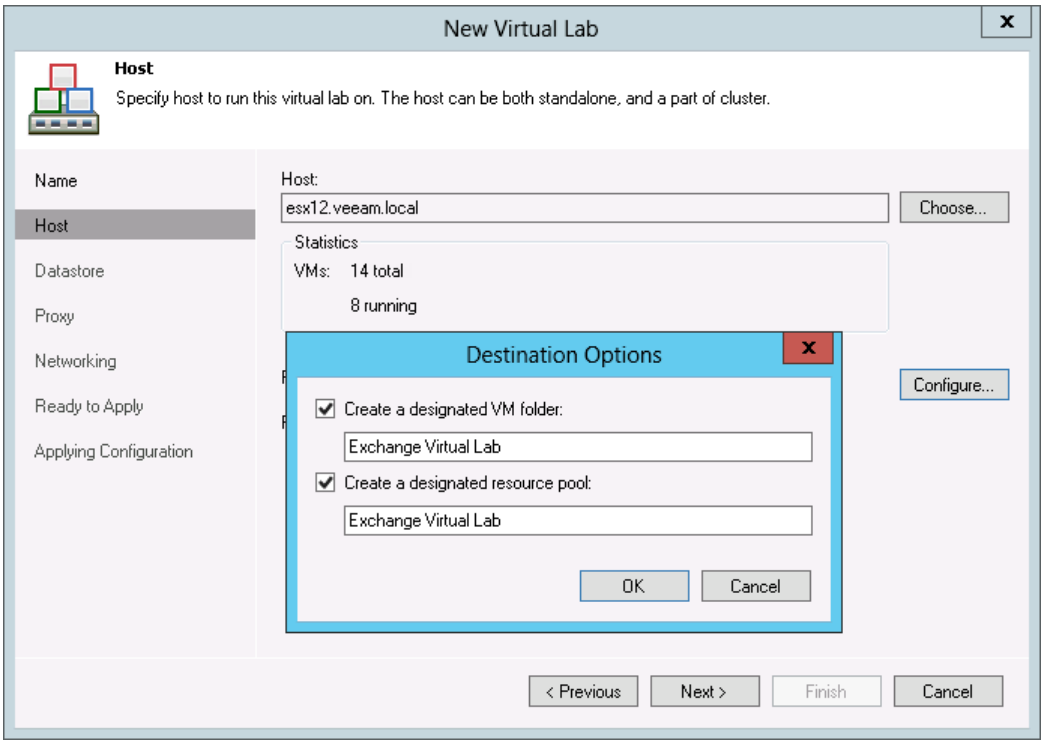

Step 3. Select Datastore

Click **Choose** to select a datastore on which redo logs for verified VMs should be placed. Redo logs are auxiliary files used to store all changes that take place when a VM runs from a read-only backup. As soon as the recovery verification job completes, redo logs are deleted.

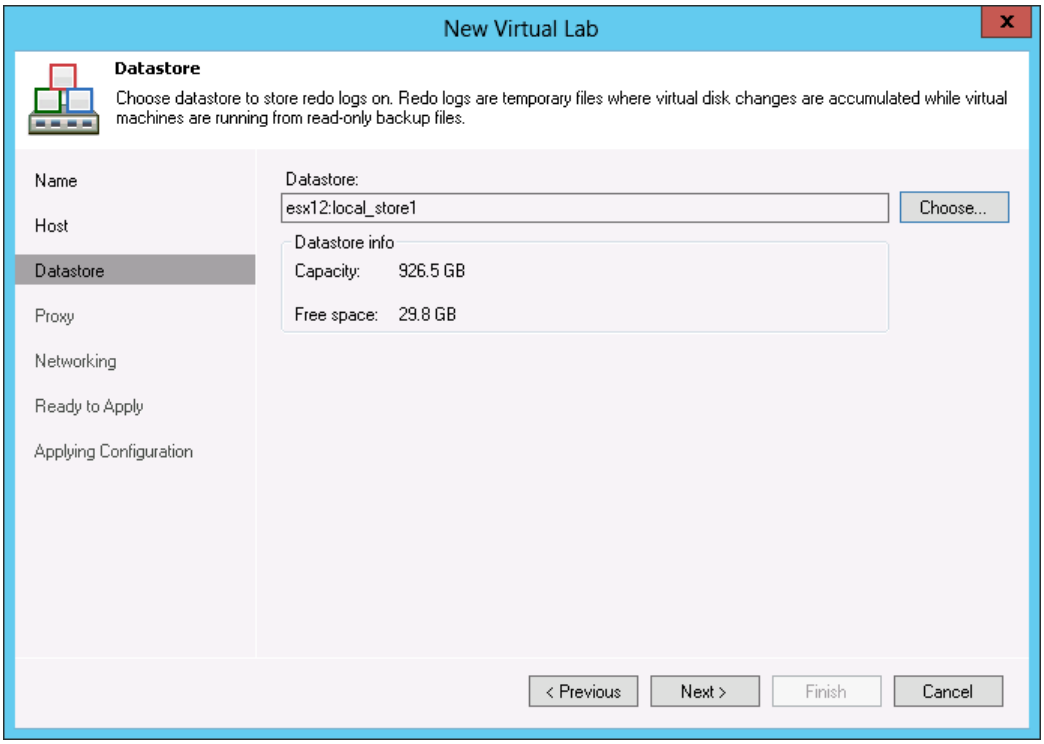

Step 4. Set Up a Proxy Appliance

At this step of the wizard, you should configure proxy appliance settings.

1. To enable automatic recovery verification of VMs, select the **Use proxy appliance in this virtual lab** check box. The proxy appliance acts as a gateway that provides access from the Veeam backup server to VMs started in the isolated virtual lab. If you do not select this check box, Veeam Backup & Replication will perform only heartbeat tests for VMs during verification. You will only be able to manually test VMs and perform manual item-level restore via the VM console.
2. To change a name of the proxy appliance, click **Configure** in the **Proxy appliance VM settings** section and specify the necessary name. By default, the proxy appliance uses the virtual lab name that you have specified at the **Name** step of the wizard.
3. To select a production network in which the proxy appliance should be created, click **Configure** in the **Production network connection** section. Specify an IP address for the proxy appliance in the production network and settings of the DNS server to be used. You can choose to automatically obtain an IP address and DNS server settings or set them manually.
4. If you want to allow access to the Internet for VMs in the virtual lab, select the **Allow proxy appliance to act as internet proxy for virtual machines in this lab** check box. In the **Port** field, specify a port for HTTP traffic. By default, port 8080 is used. In the **Production proxy** field, you can optionally specify an IP address or a fully qualified domain name of an Internet-facing proxy server that VMs should use to access the Internet.

Important! If you assign to the proxy appliance an IP address from the same network where the Veeam backup server is located, Veeam Backup & Replication will automatically add a new route to the routing table on the Veeam backup server. If you assign to the proxy appliance an IP address from a different network, you will have to manually add a new route to the routing table on the router in the production network. If you do not add a new route, tests and application scripts will fail and you will not be able to access VMs in isolated networks.

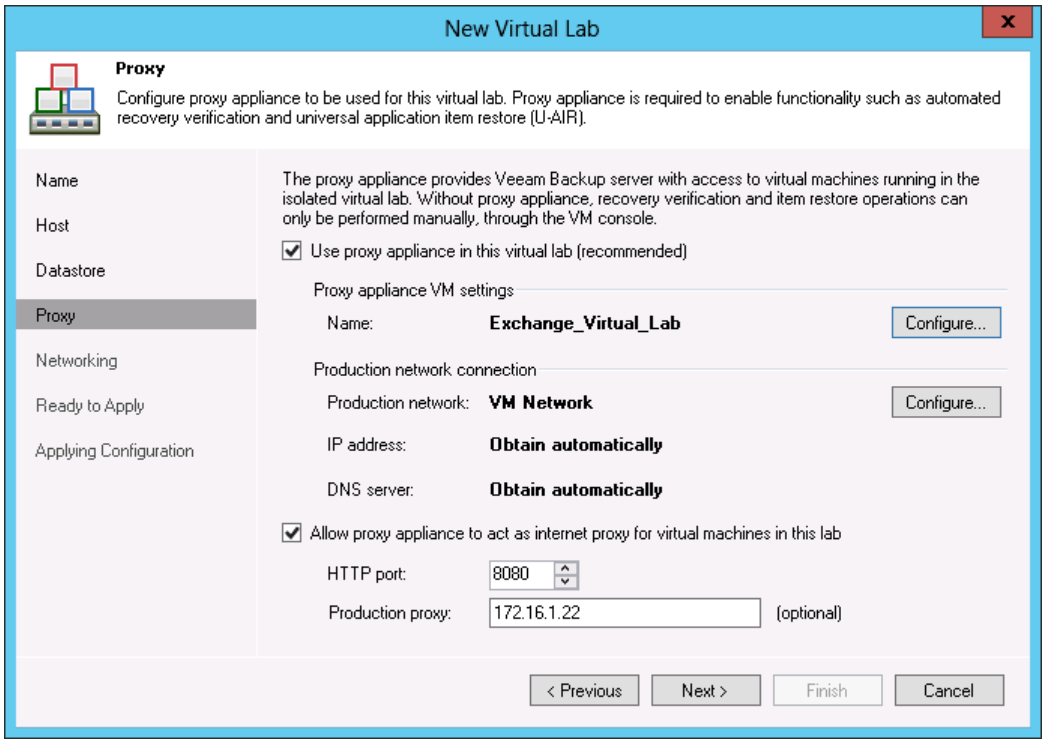

For more information on virtual lab architecture, see Virtual Lab.

Step 5. Select a Networking Mode

Select the type of network settings configuration. This step differs for virtual labs used for VM backups and VM replicas verification.

Selecting a Networking Mode for VM Backups Verification

Veeam Backup & Replication offers two networking modes for the virtual lab in which VMs from backups are verified:

- **Basic single-host (automatic configuration)**. This networking mode is recommended if all VMs you plan to verify, VMs from the application group and Veeam backup server are located in the same production network. In this case, Veeam Backup & Replication will automatically define all networking settings and create the virtual lab on an ESX(i) host.
- **Advanced single-host**. This networking mode is recommended if VMs you plan to verify and/or VMs from the application group are located in different networks. In this case, you will have to manually define settings for isolated networks in the virtual lab.
  Veeam Backup & Replication will use the specified parameters to create the virtual lab on an ESX(i) host.
  If this option is selected, the **New Virtual Lab** wizard will include additional steps for customizing network settings.

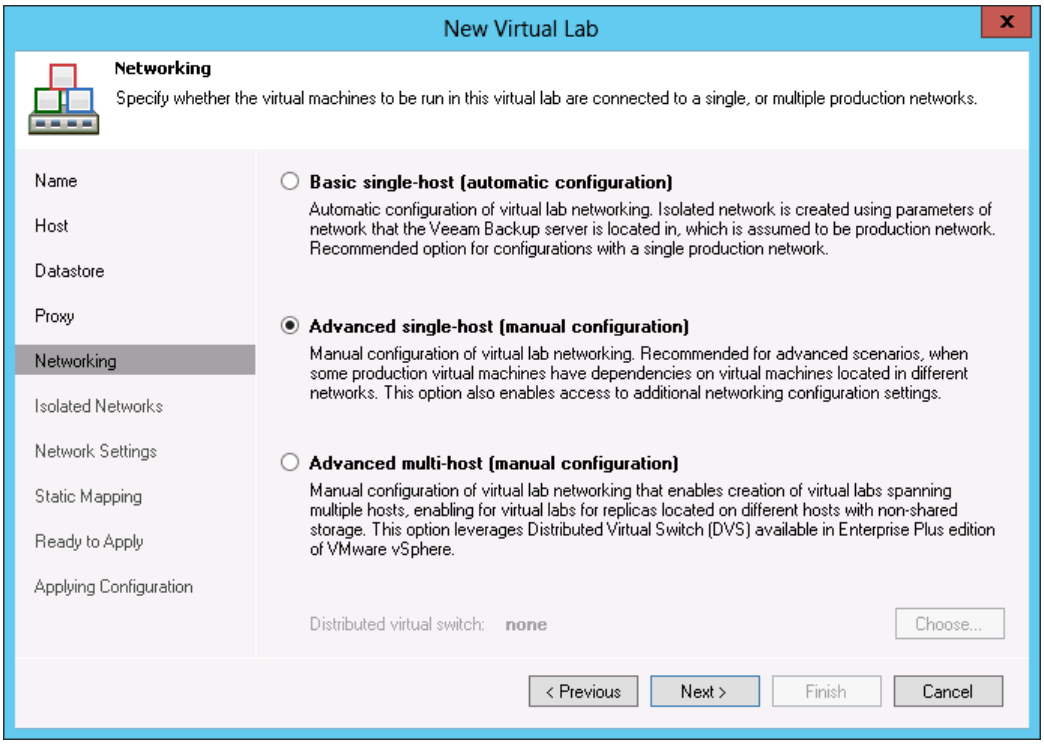

Selecting a Networking Mode for VM Replicas Verification

Veeam Backup & Replication offers three networking modes for the virtual lab in which VM replicas are verified:

- **Basic single-host (automatic configuration)**. This type of networking is recommended if VM replicas you plan to verify are located on the same ESX(i) host and are connected to the same production network. The Veeam backup server should also be located in this production network. In this case, Veeam Backup & Replication will automatically define all networking settings and create the virtual lab on the ESX(i) host that you have selected at the previous steps of the wizard.
- **Advanced single-host (manual configuration)**. This type of networking is recommended if VM replicas you plan to verify are located on the same ESX(i) host but connected to different networks. In this case, you will have to manually define settings for isolated networks in the virtual lab. Veeam Backup & Replication will use the specified parameters to create the virtual lab on the ESX(i) host that you have selected at the previous steps of the wizard.
- **Advanced multi-host (manual configuration)**. This type of networking is recommended if VM replicas you plan to verify are located on the different ESX(i) hosts. For the multi-host configuration of the virtual lab, Veeam Backup & Replication uses VMware's DVS technology. With this option selected, Veeam Backup & Replication will create isolated networks on a DVS configured in your virtual environment.
  If you have selected the **Advanced multi-host** option, click **Choose** and select the necessary DVS in your virtual environment. Note that Veeam Backup & Replication does not allow you to configure a DVS automatically: you need to configure it beforehand.

If the **Advanced single-host** or **Advanced multi-host** option is selected, the **New Virtual Lab** wizard will include additional steps for customizing network settings.

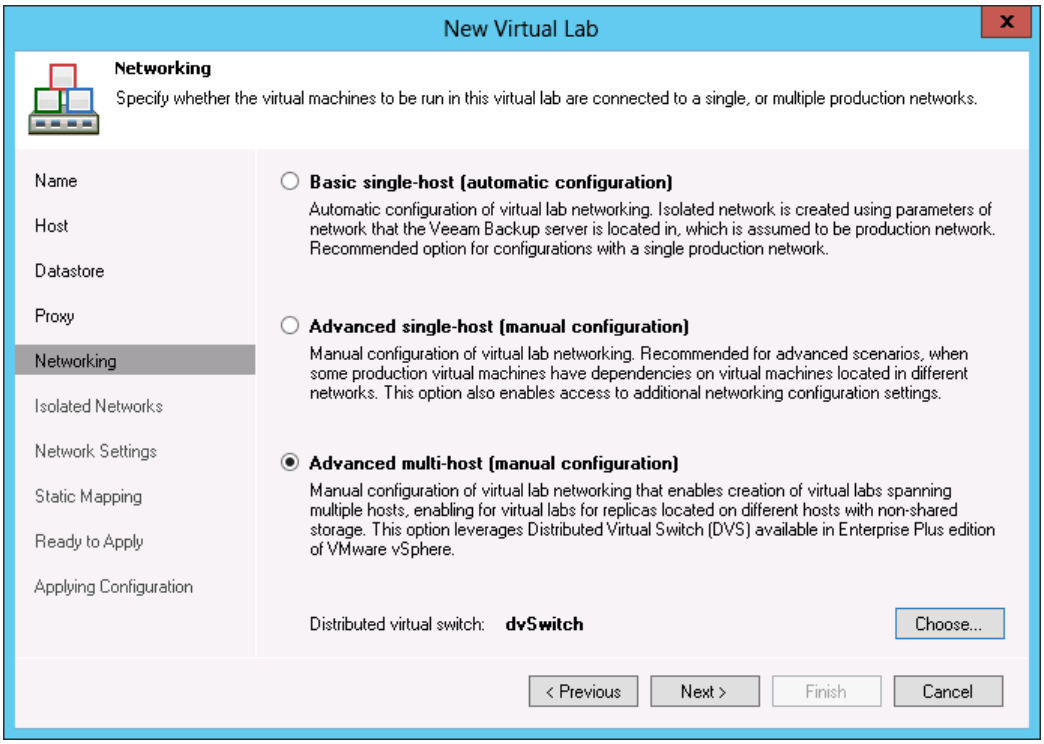

Step 6. Specify Isolated Networks

This step is available if you have selected the **Advanced networking** option at the **Networking** step of the wizard.

At this step of the wizard, you should create isolated networks to which verified VMs and VMs from the application group will be connected and map these networks to production networks where initial VMs are located.

To add a network:

1. Click **Add**.
2. From the **Production network** list, select a production network in which a VM from the application group or a verified VM resides.
3. In the **Isolated network** field, specify a name for an isolated network that should be mapped to this production network.
4. In the **VLAN ID** field, enter an identifier for the created network. In the advanced multi-host virtual lab, VLAN IDs help ensure that the created network is isolated from the production environment. Alternatively, you can manually connect the DVS you plan to use to the isolated network. To learn more, see Advanced Multi-Host Virtual Lab.

Note You can map several production networks to the same isolated network. The production networks you plan to map to the same isolated networks must have the same network masks and pools of IP addresses for mapping to be configured correctly.

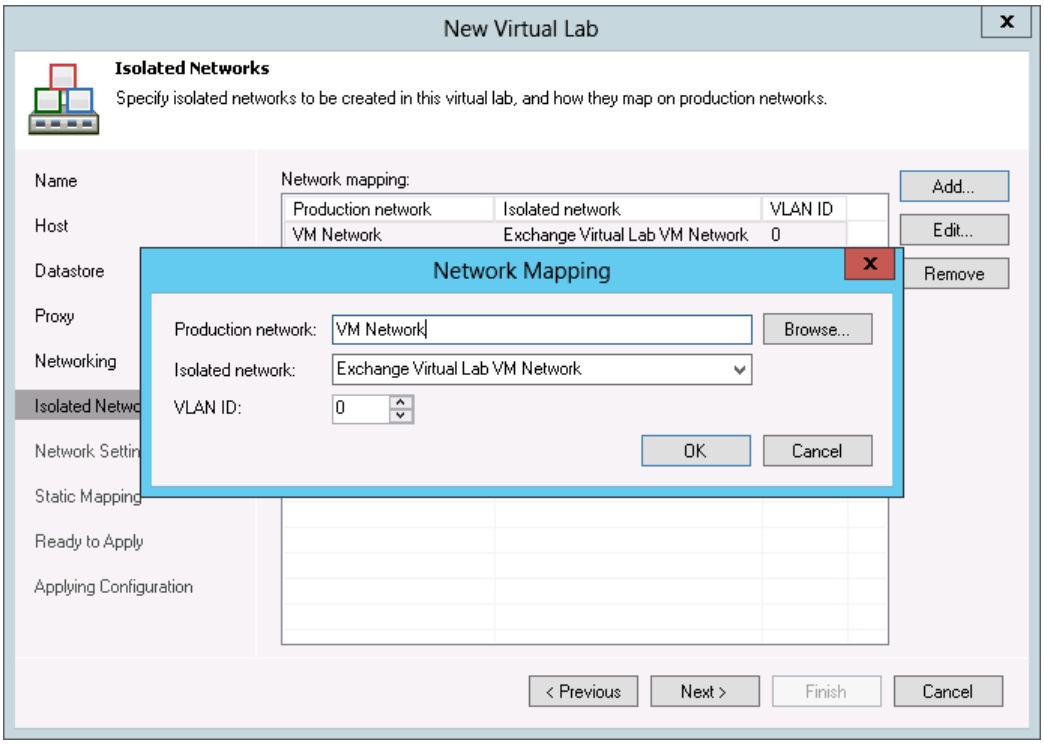

Step 7. Specify Network Settings

At this step of the wizard, you should specify settings for every created isolated network and define how production networks should map to networks in the isolated virtual lab.

Communication between production networks and isolated networks is carried out through vNIC adapters that are added to the proxy appliance. A new vNIC adapter is added for every isolated network.

To add a vNIC adapter for an isolated network:

1. At the **Network Settings** step of the wizard, click **Add**.
2. Select a network to which you want the vNIC adapter to be connected. Specify an IP address that the proxy appliance should have in the isolated network and a subnet mask of this isolated network. Typically, the IP address set for the proxy appliance coincides with the IP address of the gateway in the corresponding production network.
3. Once you specify the IP address, Veeam Backup & Replication will automatically configure a masquerade IP address for accessing VMs in the virtual lab from the production network. You can change the masquerade network IP address if necessary.
4. Select the **Enable DHCP service on this interface** check box and specify settings of a virtualized DNS server if necessary.
5. Click **OK**.
6. To enable communication between isolated networks, select the **Route network traffic between vNICs** check box. When you select this option, make sure that the IP address of the proxy appliance in the isolated network matches the IP address of the gateway in the production network.

Important! Network addresses for different network adapters should belong to different networks. For example, if the first network adapter has address 192.168.0.1 with mask 255.255.255.0, and the second one — 192.168.0.2 with mask 255.255.255.0, such configuration will not work. In this situation, you need to assign to the second adapter the IP address from a different network, for example, 172.16.0.1.

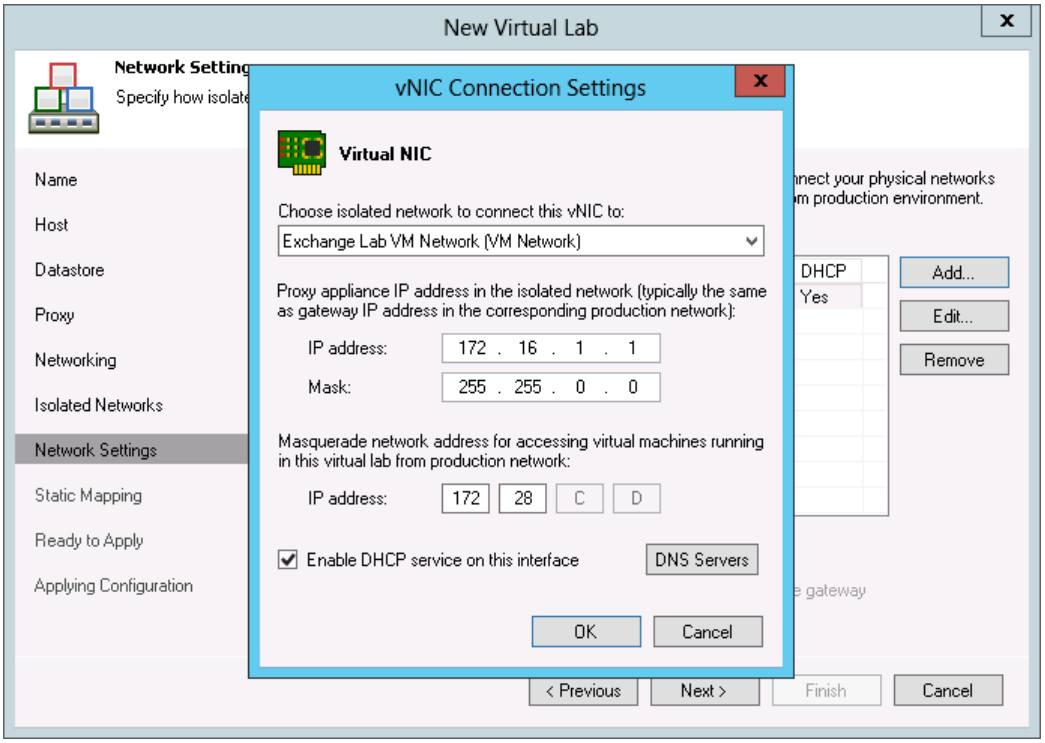

Step 8. Specify Static IP Mapping

At this step of the wizard, you can specify static IP address mapping rules to make VMs in the virtual lab accessible from any machine in the production network.

To add a new rule:

1. Click **Add**.
2. In the IP **Address Mapping** window, specify settings of a new rule:
    a. In the **Isolated IP** filed, specify a production IP address of a VM that will be started in the virtual lab and that you plan to access from the production environment.
    b. In the **Access IP** field, specify an IP address from the production network that you want to use to access this VM in the virtual lab. For a static IP address, you should use an IP address from the production network that is not yet allocated to any machine.

For example, the DNS server you plan to start in the virtual lab has IP address 192.168.1.2 in the production network. To set static mapping for the DNS server, in the Isolated IP field, you need to define its production IP address — 192.168.1.2. In the Access IP field, you need to define any unallocated IP address from the production network, for example, 192.168.1.48. After a virtual lab is created and VMs are started in the virtual lab, you will be able to access the DNS server in the virtual lab from the production environment using IP address 192.168.1.48.

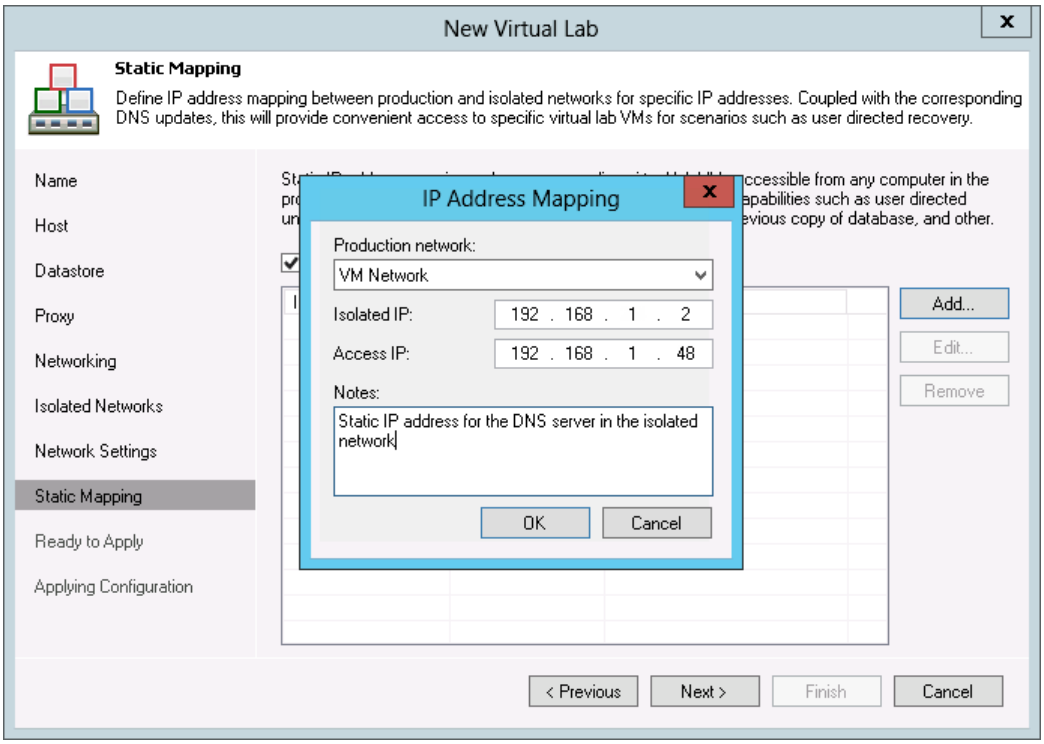

Step 9. Apply Parameters

Review the parameters of the virtual lab you create. You can go back to any previous step to adjust the parameters. If everything is fine, click **Next** to create the virtual lab.

Important! Always use Veeam Backup & Replication to modify or delete a virtual lab. If you change lab settings or delete any of its components from outside (for example, using vSphere Client), the lab will be corrupted and its component such as the created vSwitch, resource pool and so others will remain in the virtual infrastructure.

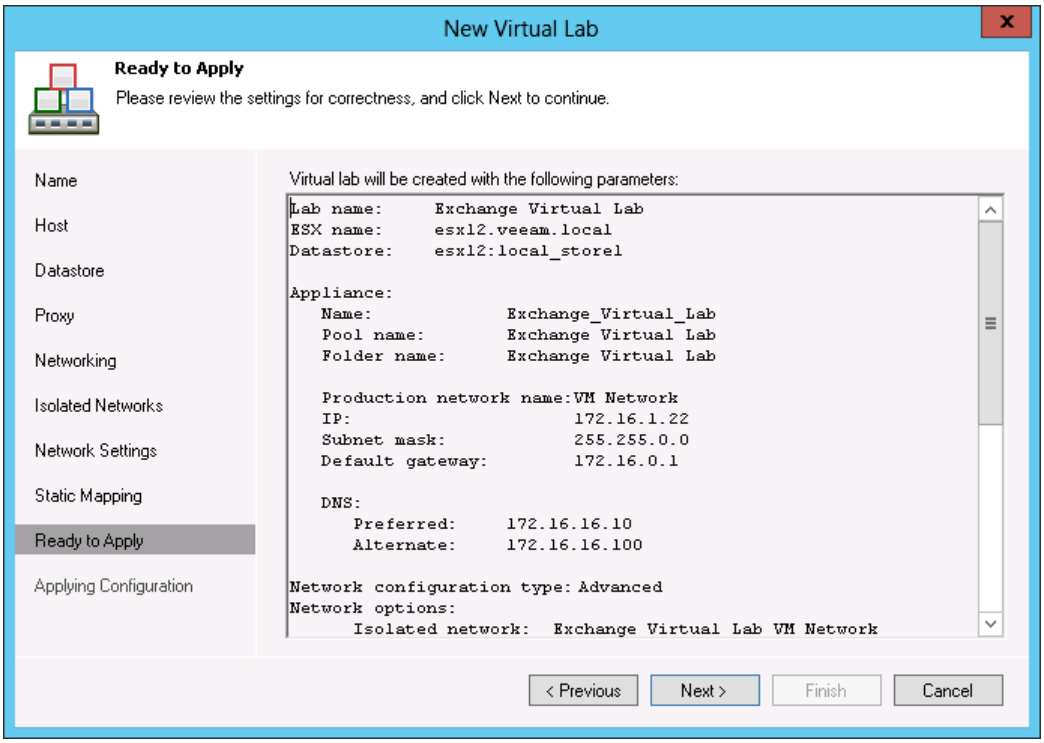

Connecting to an Existing Virtual Lab

To perform recovery verification, you can create a new virtual lab or connect to any existing virtual lab. For example, this can be a virtual lab created on another Veeam backup server.

To connect to a virtual lab, do either of the following:

- Open the **Backup Infrastructure** view, select the **Virtual Labs** node under **SureBackup** in the inventory pane and click **Connect Virtual Lab > VMware** on the ribbon.
- Open the **Backup Infrastructure** view, right-click the **Virtual Labs** node under **SureBackup** in the inventory pane and select **Connect Virtual Lab > VMware**.

Select the necessary virtual lab from the virtual environment and click **Connect**. To facilitate selection, use the search field at the bottom of the **Select Virtual Lab** window: enter a virtual lab name or a part of it in the field below and press **[ENTER]**.

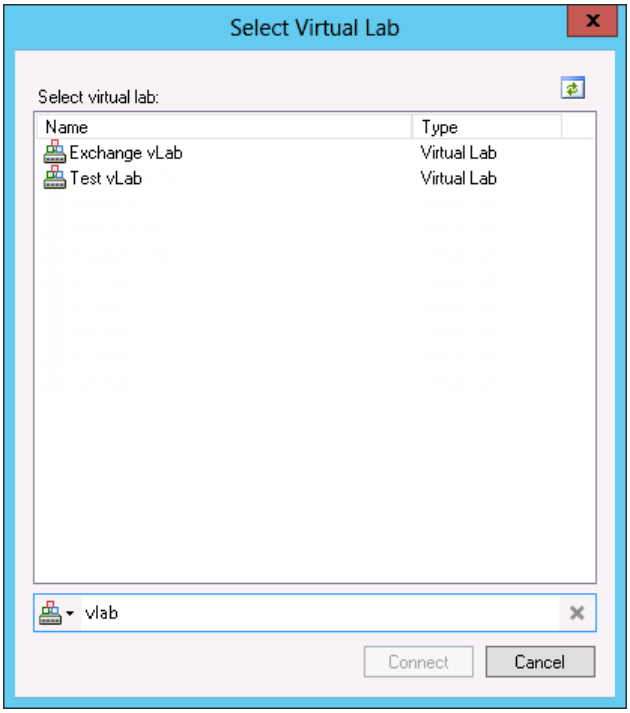

Creating a SureBackup Job

To create a new SureBackup job, do either of the following:

- Open the **Backup & Replication** view. On the **Home** tab, click **SureBackup Job > VMware** on the ribbon. Note that the **SureBackup Job** button becomes available only after you create or connect a virtual lab.
- Open the **Backup & Replication** view, right-click the **SureBackup** node under **Jobs** in the inventory pane and select **SureBackup > VMware**. You can use this method if you already have at least one SureBackup job. If there are no SureBackup jobs, the **SureBackup** node will not be available in the inventory pane. In this case, you can right-click the **Jobs** node in the inventory pane and select **SureBackup** > **VMware**.

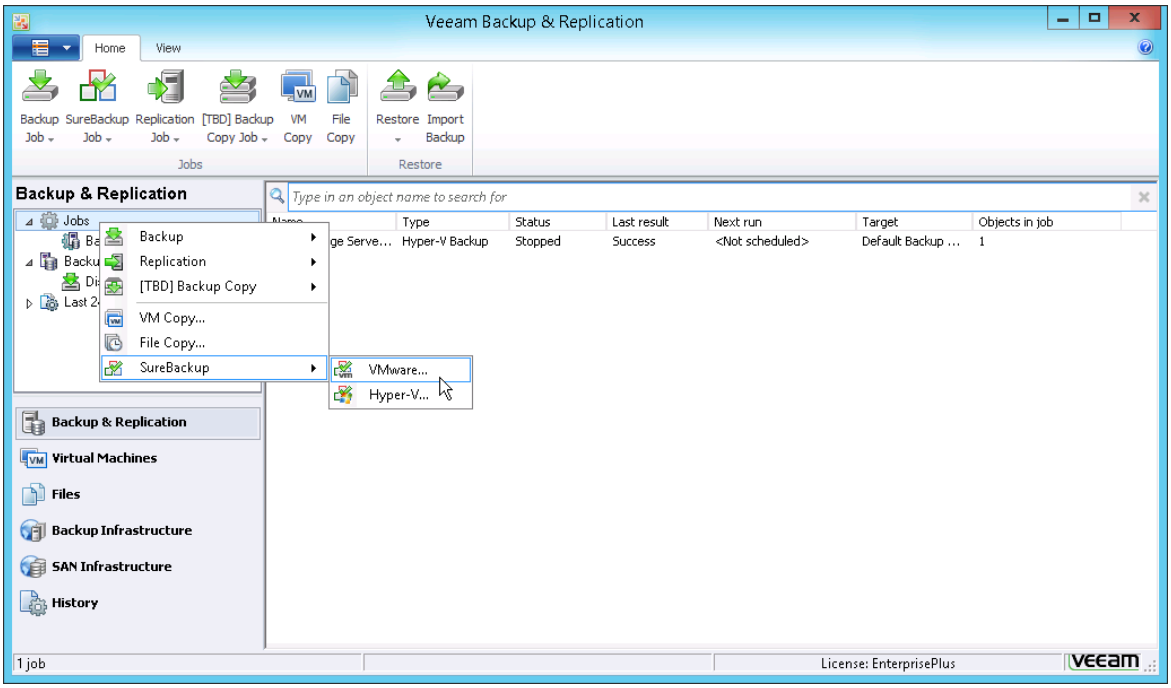

Follow the steps of the **New SureBackup Job** wizard:

Step 1. Specify Name and Description

Enter a name and description for the new SureBackup job. The default description contains the time at which the job was created and user who created it.

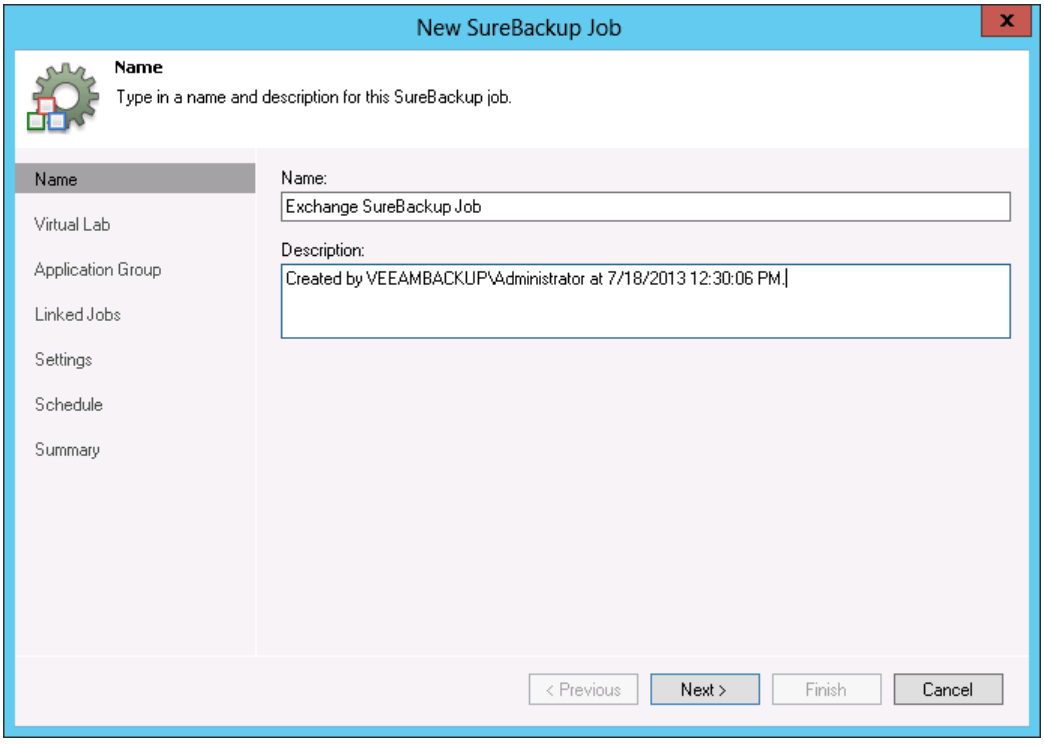

Step 2. Select a Virtual Lab

From the **Virtual lab** list, select one of existing virtual labs in which recovery verification should be performed. The list contains all virtual labs that were created or connected to the Veeam backup server. Information about the selected virtual lab is displayed in the **Virtual lab info** section.

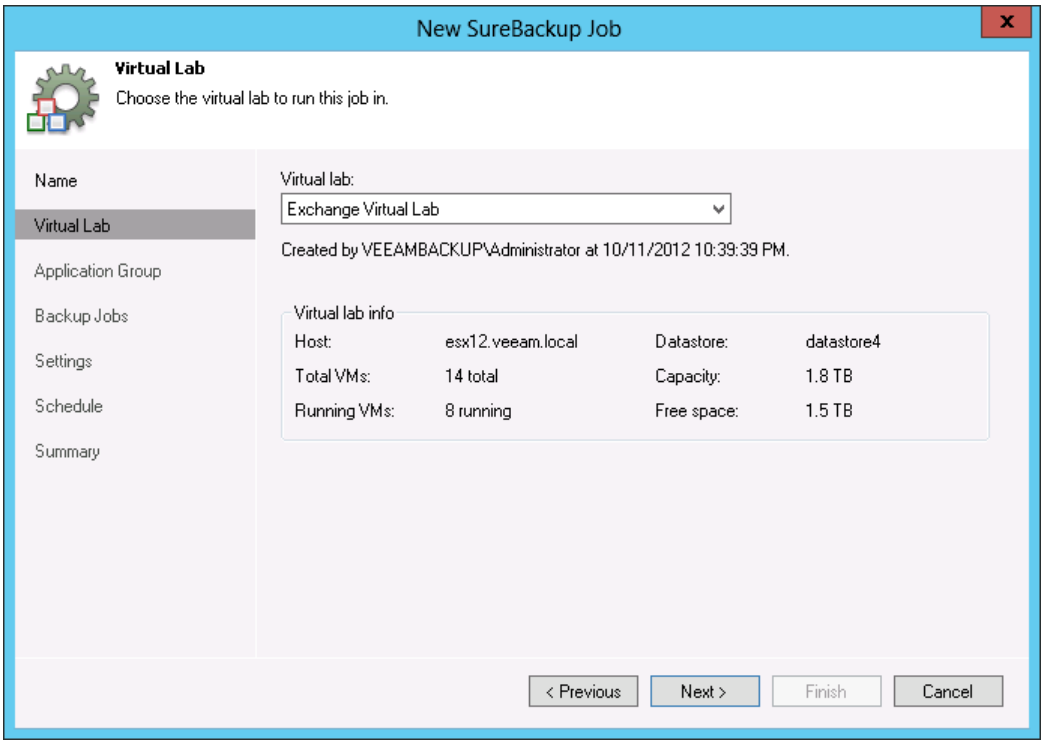

Step 3. Select an Application Group

At this step of the wizard, you should define an application group that should be used to enable full functionality of verified VMs.

You can select an application group for recovery verification or skip this step. If the application group is not selected, you must link a backup job to the created SureBackup job at the next step. In this case, when the SureBackup job is started, Veeam Backup & Replication will only run VMs from the linked backup job in the virtual lab and verify them.

To select an application group:

1. From the **Application group** list, select the application group containing all VMs required to properly run applications and services on VMs that you want to test. The list contains all application groups that were created on this Veeam backup server. Refer to the **Backup Status** column in the **Application group info** list to make sure that the backups for VMs in this group are available.
2. To leave VMs from the application group running after the SureBackup job is finished, select the **Keep the application group running once the job completes** check box. If you select this option, the lab will not be powered off when the SureBackup job completes and you will be able to perform application item-level restore (U-AIR) and manually test VMs started in the virtual lab.

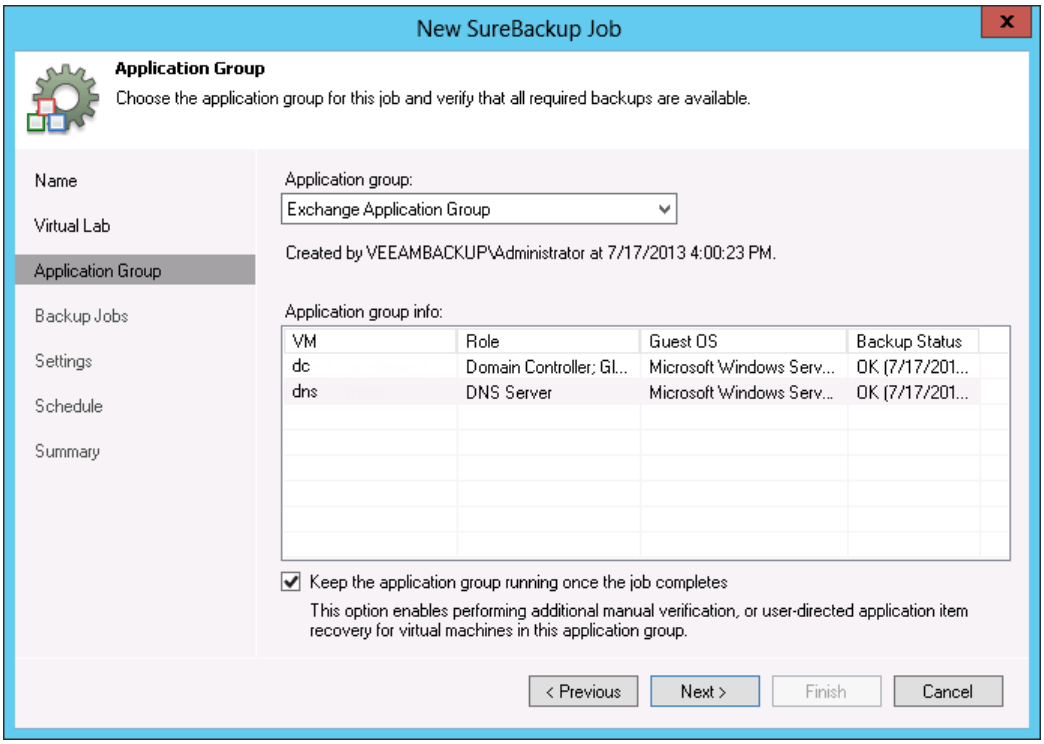

Step 4. Link a Backup or Replica Job to the SureBackup Job

At this step of the wizard, you should select VM backups or replicas that you want to verify with the created SureBackup job. Once you run a SureBackup job, Veeam Backup & Replication will start VMs from the application group in the required order and then boot and verify VMs from the linked backups or replicas.

You can link a backup or replica to the SureBackup job or skip this step. If you do not link a backup or replica to the SureBackup job, Veeam Backup & Replication will only start VMs from the application group in the virtual lab and verify them. You have an option not to link a backup or replica to the SureBackup job only in case you have selected an application group at the previous step of the wizard.

To link a backup or replica job to the SureBackup job:

1. Select the **Link jobs** check box.
2. Click **Add**.
3. In the **Select Job** window, select the necessary backup or replica job(s).
4. In the **Process simultaneously up to … VMs** field, specify the maximum number of VMs that can be started at the same time. For example, if you select to start thee VMs at the same time, Veeam Backup & Replication will create three streams in which one verified VM will be started. When the VM has been tested and powered off, the next VM will be started in the available stream. After all VMs are verified, VMs from the application group will be powered off or will be left running if the corresponding option has been selected at the previous step of the wizard.

To remove a backup or replica job from the list, select it and click **Remove**.

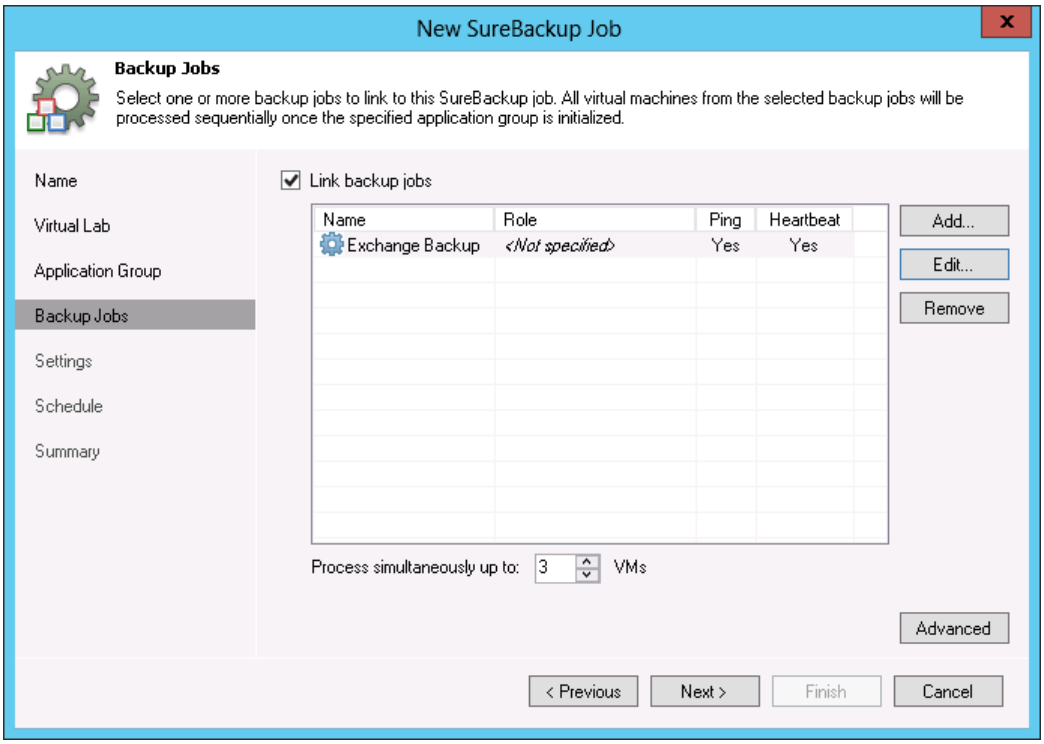

Step 5. Specify Recovery Verification Options and Tests

After you link a backup job with VMs you want to verify, you should define roles, specify startup options and select tests to be performed for these VMs.

If all VMs in the linked backup(s) perform the same role, you can specify startup options and test settings for the whole backup job in bulk. If VMs have different roles, you can granularly set startup options and select tests to be used for each VM in the backup job.

- To specify startup options and select tests for the whole backup job, select the job in the list and click **Edit** on the right.
- To specify startup options and select tests for each VM in the backup job separately, select the job in the list and click **Advanced** on the right. Click **Add** and select the necessary VM in the **Add Object** window. Select the added VM in the list, click **Edit** and specify recovery verification settings as described below.

Important! To be able to perform tests, Veeam Backup & Replication requires VMware Tools to be installed on the verified VM. If VMware Tools are not installed, the VM will be started but tests will not be performed. VMs without VMware Tools can still be used as auxiliary VMs that should be started to enable proper work of other VMs. In this case, you do not need to select any role for the VM.

Role settings

On the **Role** tab, select the role that a VM performs. Veeam Backup & Replication offers the following predefined roles for VMs:

- DNS Server
- Domain Controller
- Global Catalog
- Mail Server
- SQL Server
- Web Server

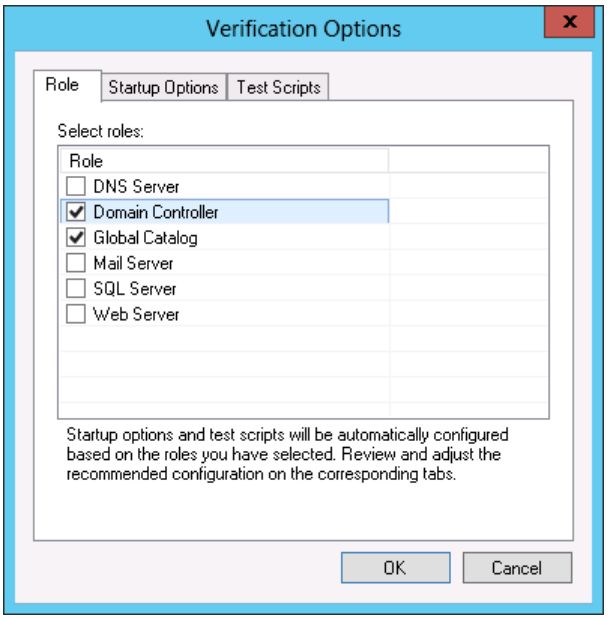

VM roles are described in XML files stored in the *%Program Files%\Veeam\Backup and Replication\Backup\SbRoles* folder. You can add your own roles. To do this, you need to create a new XML file and specify role and test scripts settings in it. For details, see Creating XML files with VM Roles Description.

Once you select the necessary role, Veeam Backup & Replication will automatically configure startup options and provide predefined test scripts applicable for the chosen role. You can use these settings or specify custom ones using the **Startup Options** and **Test Scripts** tabs.

To verify VMs that perform roles other than those specified in the list, you will have to manually configure startup options and specify test scripts to be run against these VMs.

Startup Options

On the **Startup Options** tab, specify VM startup settings:

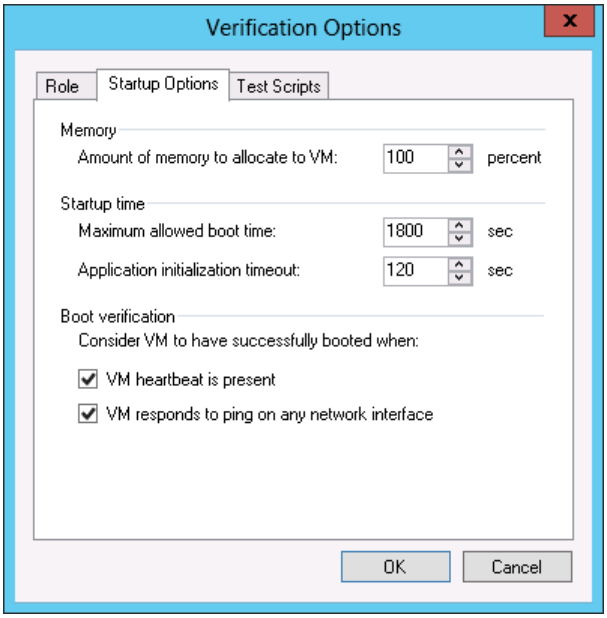

1. In the **Memory** section, specify the amount of memory that you want to pre-allocate to the VM on the system boot. The amount of pre-allocated memory is specified in percent. The percentage rate is calculated based on the system memory level available for the production VM. For example, if 1024 MB of RAM is allocated to the VM in the production and you specify 80% as a memory rate, 820 Mb of RAM will be allocated to the verified VM on startup.
2. In the **Startup time** section, specify the allowed boot time for the VM and timeout to initialize applications on it.
3. In the **Boot verification** section, specify when the VM should be considered to have been booted successfully:
    - VM heartbeat is present. If you select this option, Veeam Backup & Replication will perform a heartbeat test for the verified VM.
    - VM responds to ping on any network adapter. If you select this option, Veeam Backup & Replication will perform a ping test for the verified VM.
      Before you start a SureBackup job, make sure that the firewall on the tested VM allows ping requests.

Important! Be careful when specifying the **Maximum allowed boot time** value. Typically, a VM started by a SureBackup job requires more time to boot than a VM started regularly. If an application fails to be initialized within the specified interval of time, the recovery verification process will be finished with the timeout error. If such error situation occurs, you will need to increase the **Maximum allowed boot time** value and start the job again.

Test Scripts

On the **Test Scripts** tab, specify what test scripts should be run to verify a VM. When you select a VM role, Veeam Backup & Replication automatically assigns a predefined script that will be run to verify applications inside the VM.

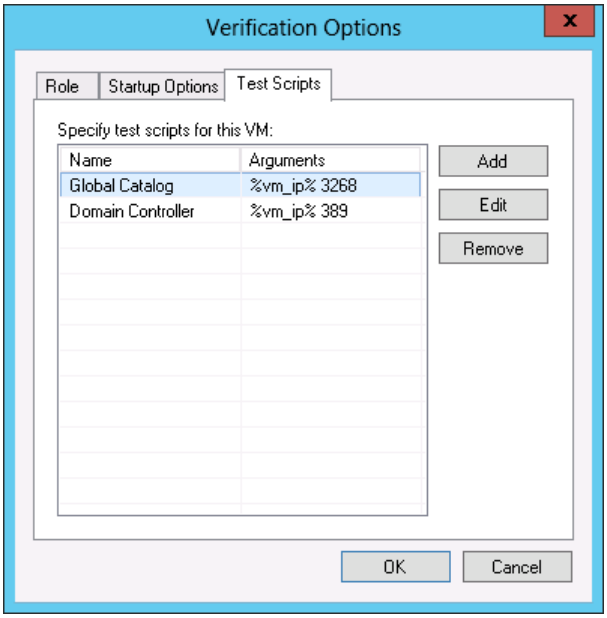

If you want to verify a VM that has some other role not listed on the **Role** tab, do the following:

1. Click **Add**.
2. In the **Test Script** window, select **Use the following test script**.
3. In the **Name** field, specify a name for the script.
4. In the **Path** field, define a path to an executable script file that should be run to verify the VM.

- If you have your own custom script, define a path to it.
- If you do not have a custom script, you can use Veeam's standard utility, *Veeam.Backup.ConnectionTester.exe*, that probes application communication ports. The utility is located in the installation folder of Veeam Backup & Replication: *%Program Files%\Veeam\Backup and Replication\Backup\Veeam.Backup.ConnectionTester.exe*. Specify this path in the **Path** field.

5. In the **Arguments** field, specify an IP address of the tested VM and the port that you want to probe. You can use the *%vm_ip%* variable to define the VM IP address or the *%vm_fqdn%* variable to define the VM fully qualified domain name.
6. Click **OK** to add the configured test.

To edit test settings, select the test in the list and click **Edit**. To delete a test, select it in the list and click **Remove**.

Note If a VM performs several roles running a number of applications at once, you can verify their work by adding several verification scripts. For such VMs, it is recommended to specify the maximum startup timeout value and allocate the greatest amount of memory.

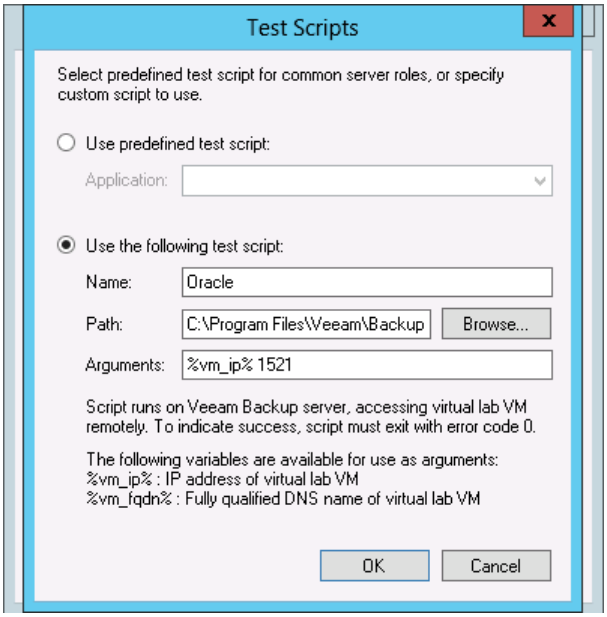

Step 6. Specify Additional Job Settings

On the **Settings** step, you can specify additional settings for the SureBackup job:

1. If you want to receive SNMP traps, select the **Send SNMP trap** check box. SNMP traps will be sent only if you configure SNMP settings in Veeam Backup & Replication and on the recipient's computer. To learn more, see Specifying SNMP Settings.
2. If you want to receive notifications by email, select the **Send email notifications to the following recipients** check box. In the field below, specify a recipient's email address. You can enter several addresses separated by a semicolon. Email notifications will be sent only if you configure general email notification settings in Veeam Backup & Replication. To learn more, see Specifying Email Notification Settings.

3. (For VM backups only) If you want to validate the backup file with a CRC check and make sure the file is not corrupted, select the **Validate consistency of virtual machines' backup files** check box. You can optionally exclude VMs being a part of the application group from this test. To do this, select the **Skip validation for application group** check box. To learn more, see Recovery Verification Tests.

Note If you select the **Keep the application group running once the job completes** option at the **Application Group** step of the wizard, the **Skip validation for application group** check box will be automatically selected and this option will be enabled.

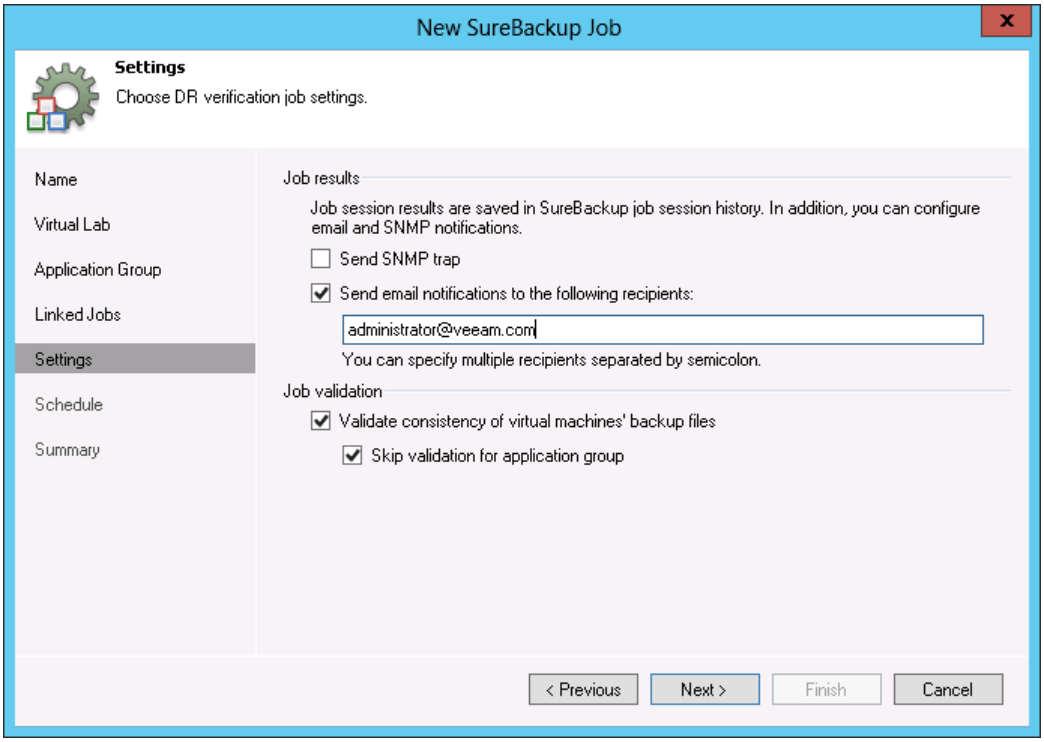

Step 7. Specify the Job Schedule

At the **Schedule** step of the wizard, you can select to manually run the SureBackup job or schedule the job at specific time, for example, after the corresponding backup job completes.

1. To specify the job schedule, select the **Run the job automatically** check box. If this check box is not selected, the job is supposed to be run manually.
2. Choose the necessary schedule option for the job:
   - **Daily at** to start the job at specific time every day, on week days or on specific days.
   - **Monthly at** to start the job once a month on the specified day.
   - **After this job** to chain the job. Typically, a SureBackup job should run after the linked backup job completes: in this case, the SureBackup job will verify the backup created by the corresponding backup job. To create a chain of jobs, you should define the time schedule for the first job in the chain. For the rest of the jobs in the chain, at the **Schedule** step of the wizard, select the **After this job** option and choose the preceding job from the list.
3. In some cases, the linked backup job may not complete until the SureBackup job starts. If Veeam Backup & Replication finds out that the backup job is still running, the SureBackup job will fail to start. To overcome this situation, select the **If some linked backup jobs are still running, wait up to** check box and specify the necessary time period in the field on the right.

In this case, if the linked backup job is still running, Veeam Backup & Replication will wait for the defined period of time and check the backup job after this period elapses.

- If the linked backup job is finished within the specified period, the SureBackup job will start.
- If the backup job is still running, the SureBackup job will not be started.

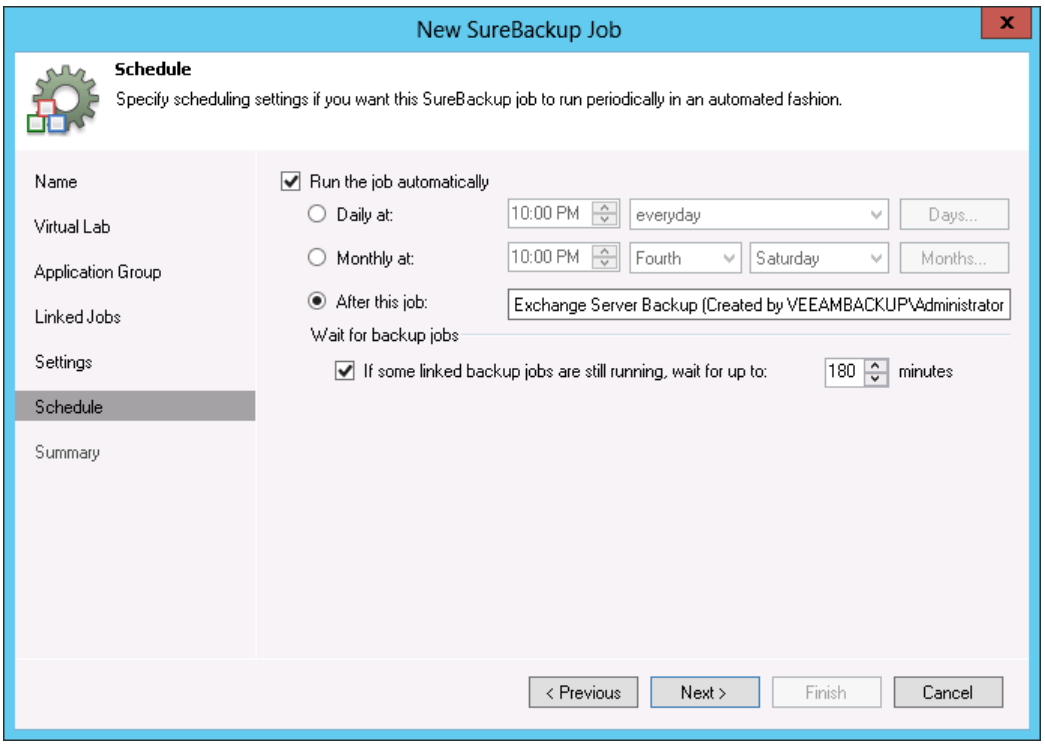

Step 8. Review the Job Summary and Finish Working with Wizard

Review the summary of the created recovery verification job. Select the **Run the job when I click Finish** check box to start the created job right after you finish working with the wizard; then click **Finish**.

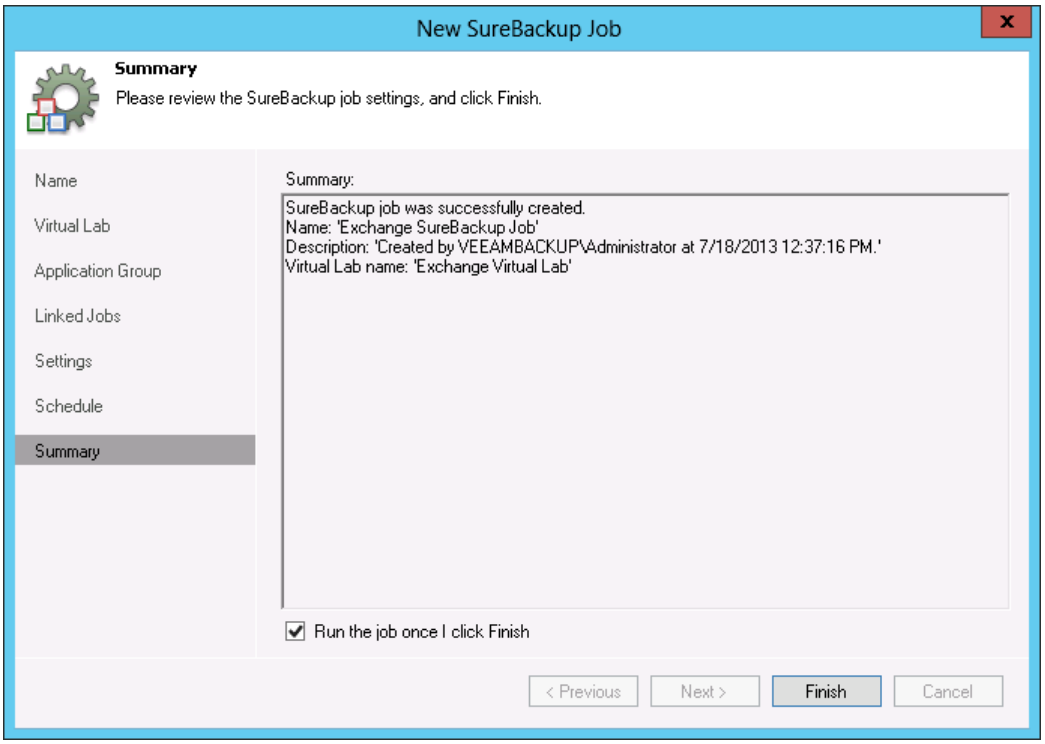

Viewing Recovery Verification Job Statistics

When a recovery verification job is running, you can monitor how tests for verified VMs are performed and see their results in the real-time mode. To see the status of VM tests:

1. Open the **Backup & Replication** view.
2. Select the **SureBackup** node under **Jobs** in the inventory pane.
3. Right-click the necessary recovery verification job in the working area and select **Statistics**. You can also simply double-click the job in the list.

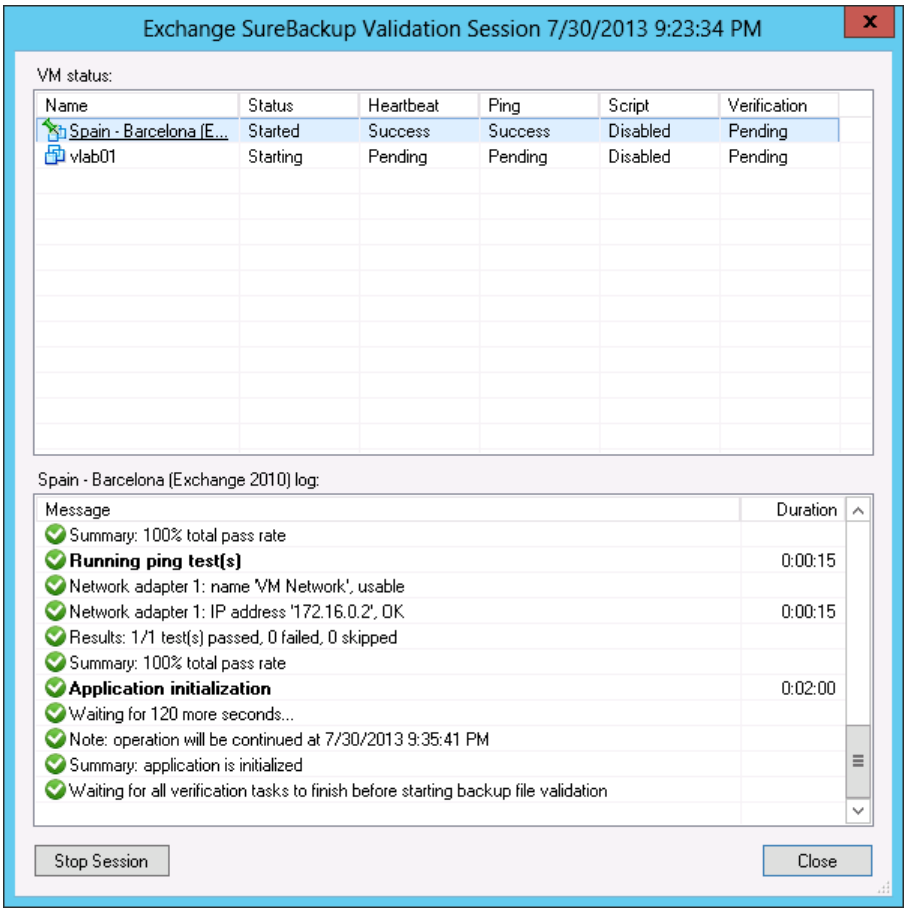

The job session window displays statistics for all VMs that are started during the SureBackup job: VMs from the application group in the specified order and VMs from the linked backup job(s). For your convenience, these VMs are marked with different icons.

Once the verified VM is powered on, its name is displayed as a hyperlink. You can click the link to open the VM console to see what is happening inside a VM or perform manual testing. To open the VM console, click the VM name link in the list of verified VMs.

After the verified VM is started and the application running inside is initialized, you can start U-AIR wizards right from the **Realtime statistics** window to perform granular application-item recovery. To do this, right-click the verified VM and select a corresponding command from the shortcut menu. Depending on the type of a running VM, you can start the **Active Directory item recovery** wizard, **Exchange item recovery** wizard or **SQL item recovery** wizard.

If some VM fails to be verified automatically, once it is powered off, you can start it by right-clicking it in the list and selecting **Start**. If the application group has already been powered off by that time, it will be started again. After that, you can open the VM console and perform verification and testing manually.

Creating SureBackup Session Reports

Veeam Backup & Replication allows you to generate HTML reports with statistics on the performed SureBackup job, a separate session or multiple sessions.

A report generated for the job contains detailed data on job sessions: job status, start and end time and details of the session performance, as well as the status of verified VMs and test results. You can generate a report for a SureBackup job or a specific job session.

The SureBackup job report contains data on all sessions initiated for a specific job. To make up a SureBackup job report:

1. Open the **Backup & Replication** view.
2. Select the **Jobs** node in the inventory pane.
3. Right-click a necessary SureBackup job in the working area and select **Report**.

The session report contains data on a single job session. To make up a session report:

1. Open the **History** view.
2. Select the **Jobs** node in the inventory pane.
3. Right-click a necessary session in the working area and select **Report**.

Creating XML Files with VM Roles Description

VM roles that you can assign to the verified VMs or VMs from the application group are described in XML files stored in the *%Program Files%\Veeam\Backup and Replication\Backup\SbRoles* folder. To add a new role, you should create a new XML file and save it to the *SbRoles* subfolder.

XML files describing VM roles have the following structure:

```
<SbRoleOptions>
 <Role>
  <SbRole>
   <Id>4CDC7CC4-A906-4de2-979B-E5F74C44832F</Id>
   <Name>Web Server</Name>
  </SbRole>
 </Role>
 <Options>
  <SbVerificationOptions>
   <ActualMemoryPercent>100</ActualMemoryPercent>
   <MaxBootTimeoutSec>300</MaxBootTimeoutSec>
   <AppInitDelaySec>120</AppInitDelaySec>
   <TestScripts>
    <TestScripts>
     <TestScript>
      <Name>Web Server</Name>
      <Type>Predefined</Type>

<TestScriptFilePath>Veeam.Backup.ConnectionTester.exe</TestScriptFileP
ath>
      <Arguments>%vm_ip% 80</Arguments>
     </TestScript>
    </TestScripts>
   </TestScripts>
   <HeartbeatEnabled>True</HeartbeatEnabled>
   <PingEnabled>True</PingEnabled>
  </SbVerificationOptions>
 </Options>
</SbRoleOptions>
```

Available XML tags are described in the table below:

Tag	Required /Optional	Description
<SbRoleOptions>	Required	Encapsulates the VM role file.
<Role>	Required	Parent tag for a role assigned to a VM. *<SbRole>*, *<Id>* and *<Name>* are children of this tag.
<SbRole>	Required	Encapsulates basic information for a VM role: ID and name.
<Id>	Required	Unique identifier of a VM role.
<Name>	Required	Name of a VM role that is displayed in the roles list on the **Role** tab.
<Options>	Required	Parent tag for startup and test script options to be used for the defined role. *<SbVerificationOptions>*, *<ActualMemoryPercent>*, *<MaxBootTimeoutSec>*, *<AppInitDelaySec>*, *<TestScripts>*, *<Name>*, *<Type>*, *<TestScriptFilePath>*, *<Arguments>*, *<HeartbeatEnabled>*, *<PingEnabled>* are children of this tag.
<SbVerificationOptions>	Required	Encapsulates options data for a VM role.
<ActualMemoryPercent>	Optional	Percent of the original memory level set for a production VM that should be pre-allocated to a verified VM on the system boot.
<MaxBootTimeoutSec>	Optional	Maximum allowed time to boot a VM.
<AppInitDelaySec>	Optional	Maximum allowed time to initialize an application inside the VM.
<TestScripts>	Optional	Encapsulates test script data for a VM role.
<Name>	Optional	Name of a VM role to be displayed on the **Test Scripts** tab.
<Type>	Optional	Type of the test script: *Predefined* or *Custom*
<TestScriptFilePath>	Optional	Path to an executable file with a test script to be performed. Can be absolute or relative.
<Arguments>	Optional	Arguments to be passed to the script. You can use two variables here: • *%vm_ip%* – IP address of a virtual lab VM • *%vm_fqdn%* – a fully qualified domain name of a virtual lab VM
<HeartbeatEnabled>	Required	Should a heartbeat test be enabled for this VM role: *True* or *False*.
<PingEnabled>	Required	Should a ping test be enabled for this VM role: *True* or *False*.

Performing Restore

Veeam Backup & Replication offers a variety of data recovery facilities to protect your virtual environment. If important data accidentally gets lost or corrupted, you can use Veeam Backup & Replication to restore entire VMs and specific VM files from backups or recover individual VM guest OS files and folders from backups and replicas.

Performing Instant VM Recovery

With Veeam Backup & Replication, you can immediately start a VM from a backup stored in a regular backup repository. Instant VM recovery accelerates VM restore, allowing you to improve recovery time objectives and decrease downtime of production VMs.

To perform instant recovery of a VMware VM, follow the next steps:

Step 1. Launch the Instant VM Recovery Wizard

To launch the **Instant VM Recovery** wizard, do one of the following:

- On the **Home** tab, click **Restore** and select **VMware**. In the **Restore from backup** section, select **Instant VM recovery**.
- Open the **Backup & Replication** view and select the **Backups** node. In the working area, expand the necessary backup job, select the VM(s) you want to restore and click **Instant VM Recovery** on the ribbon.
- Open the **Backup & Replication** view and select the **Backups** node. In the working area, expand the necessary backup job, right-click the VM(s) you want to restore and select **Instant Recovery**.

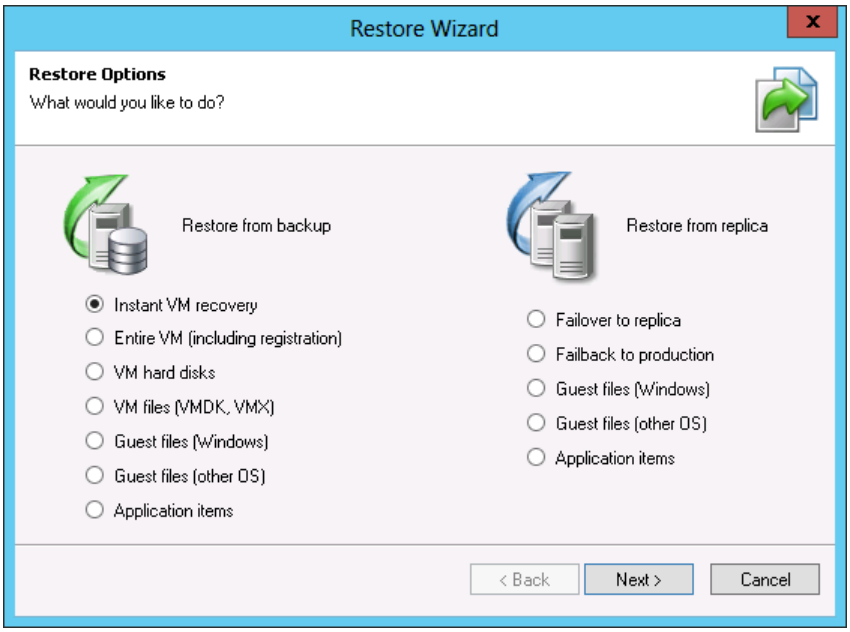

Step 2. Select a Virtual Machine

Select the necessary VM in the list of available backup jobs. You can instantly recover a VM from the backup that has been successfully created at least once. To quickly find the VM, use the search field at the bottom of the window: enter a VM name or a part of it and click the **Start search** button on the right or press **[ENTER]**.

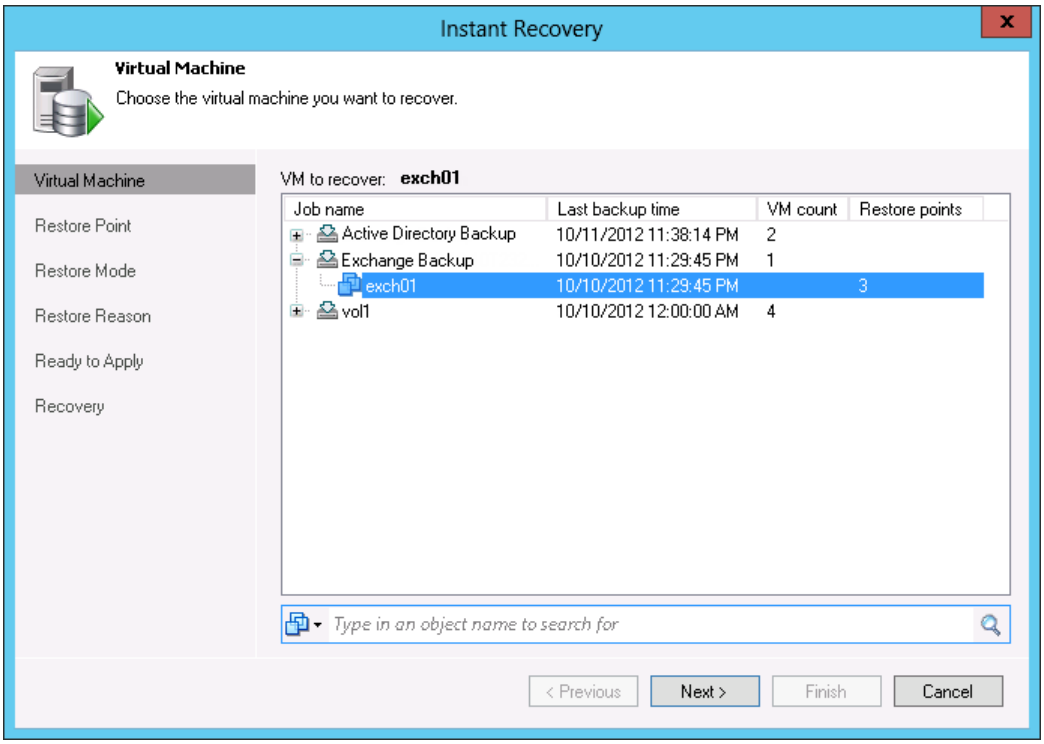

Step 3. Select a Restore Point

Select the necessary restore point for the virtual machine.

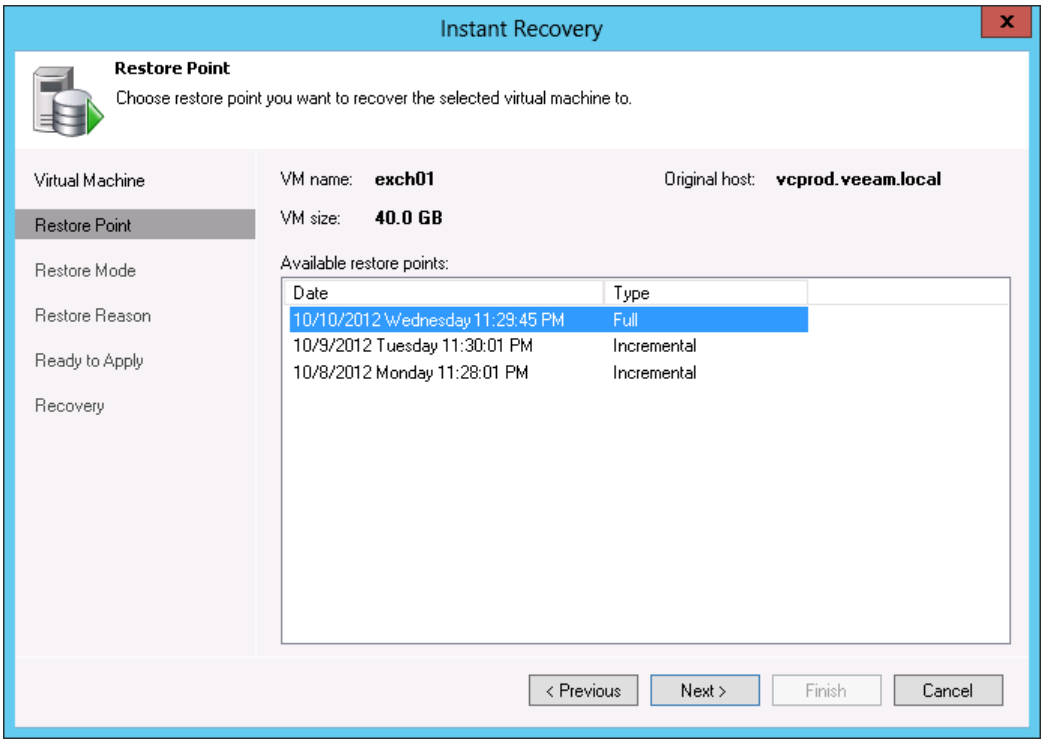

Step 4. Select a Restore Mode

Choose the necessary restore mode:

- Select **Restore to the original location** if you want to restore the VM with its initial settings and to its original location. If this option is selected, you will pass directly to the Reason step of the wizard.
- Select **Restore to a new location, or with different settings** if you want to restore the VM to a different location and/or with different settings (such as VM location, network settings, format of restored virtual disks and so on). If this option is selected, the **Instant Recovery** wizard will include an additional step for customizing VM settings.

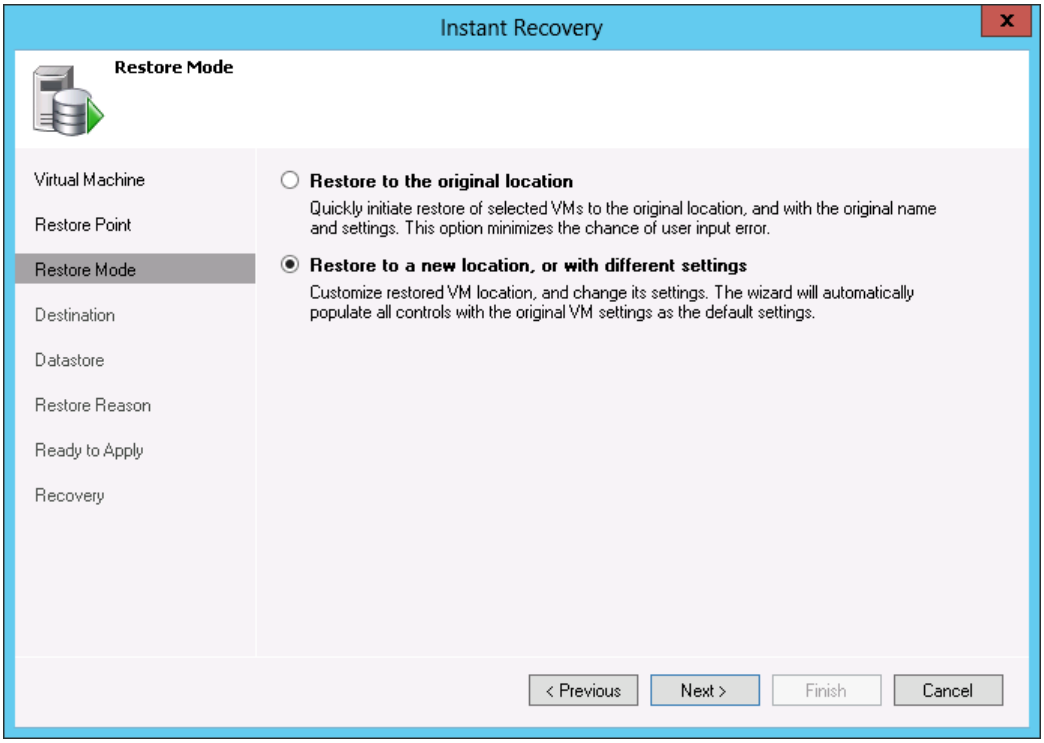

Step 5. Select Destination for the Restored VM

This step of the wizard is available if you have chosen to change the location and settings of the restored VM.

Select the destination for the recovered VM:

1. In the **Host** field, specify a host on which the VM should run.
2. In the **VM folder** field, specify the folder to which the restored VM should be placed.
3. In the **Restored VM name** field, enter a name under which the VM should be restored and registered. By default, the original name of the VM is used. If you are restoring the VM to the same ESX(i) host or the same datacenter where the original VM is registered and the original VM still resides there, it is recommended that you change the VM name to avoid conflicts.
4. In the **Resource pool** list, select a resource pool to which the VM should be recovered.

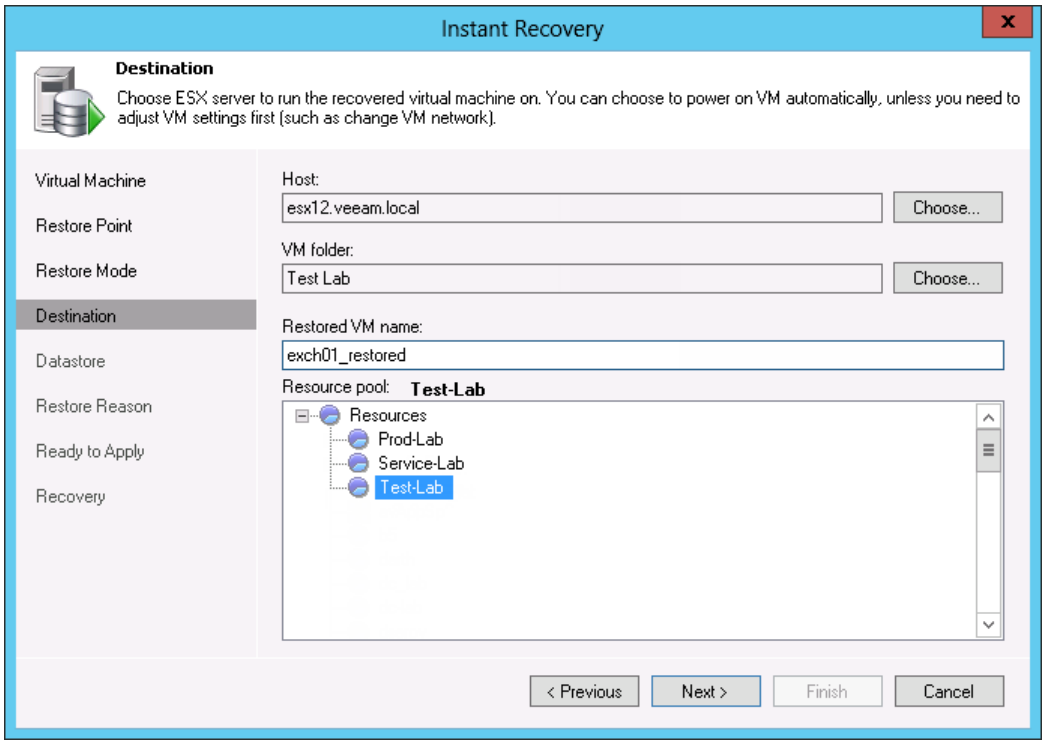

Step 6. Select Destination for Virtual Disk Updates

This step of the wizard is available if you have chosen to change the location and settings of the restored VM.

Select where disk changes should be written when a VM is restored. By default, disk changes are stored directly on the Veeam backup server. However, you can store disk changes to any datastore in your VMware environment. Select the **Redirect virtual disk updates** check box and choose the necessary datastore. Redirecting disk changes improves recovery performance but makes Storage vMotion not possible.

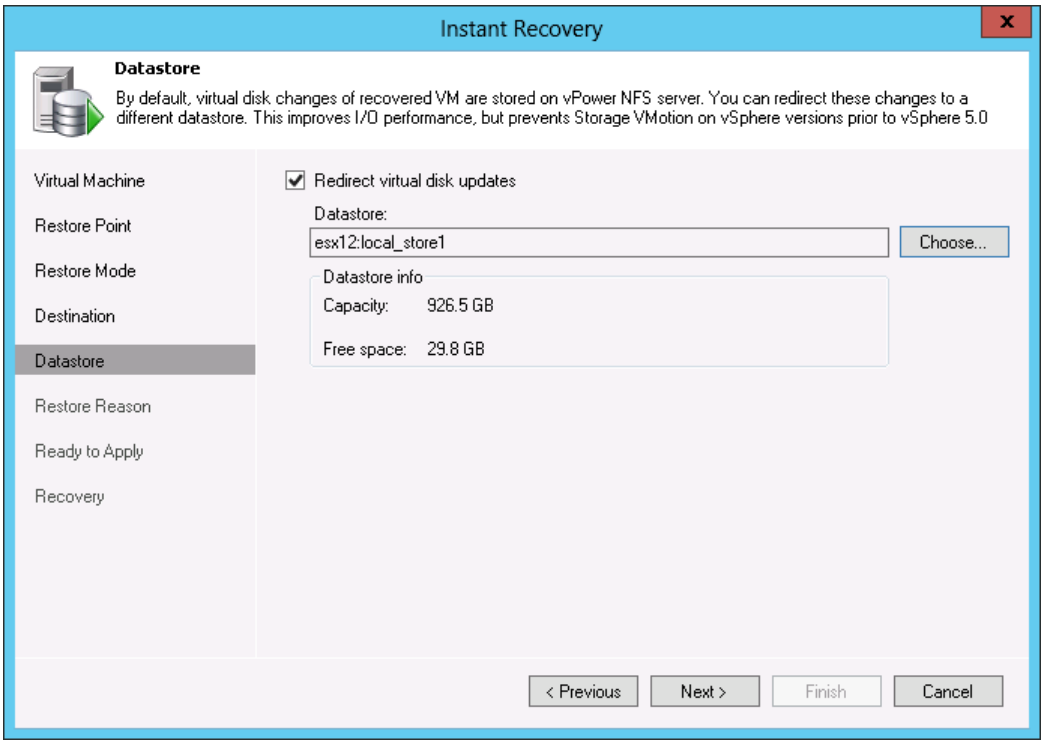

Step 7. Specify Restore Reason

If necessary, enter the reason for performing instant restore of selected VMs. The information you provide will be saved in the session history so that you can reference it later.

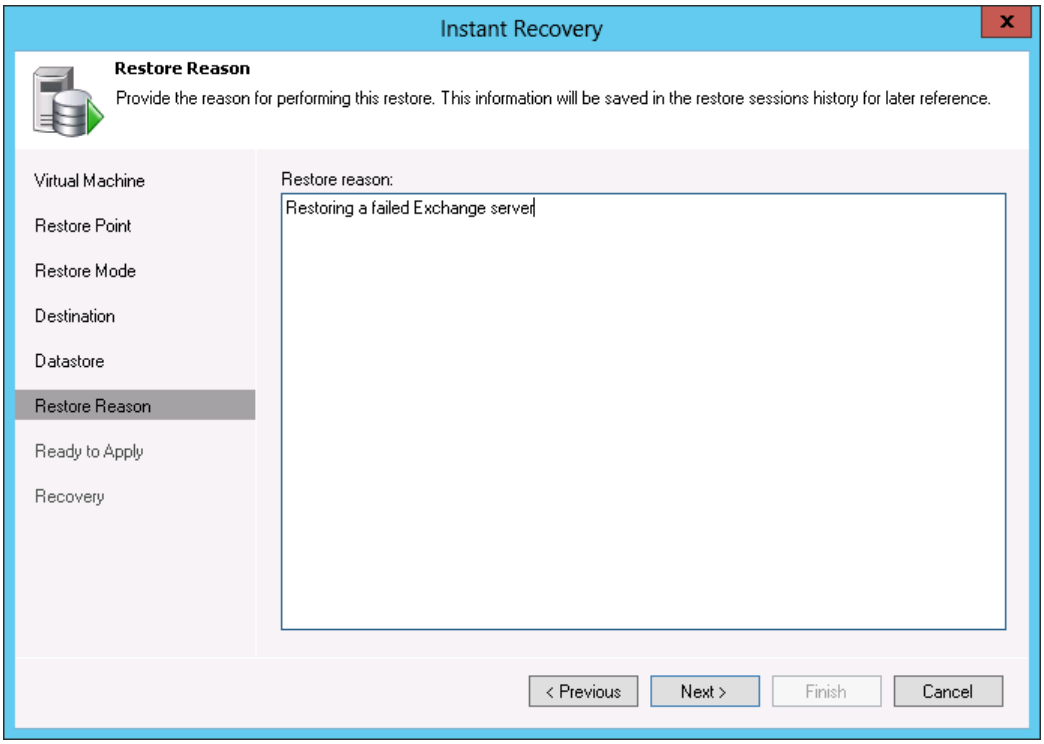

Step 8. Verify Instant VM Recovery Settings

Specify additional restore settings:

1. If you are recovering a production VM that has failed and want to restore it with initial network settings, select the **Connect VM to network** check box. If you are recovering a VM for testing disaster recovery while the initial VM is still running, leave this check box not selected. Before you power on a VM, you will have to manually change VM network configuration: disconnect it from the production network and connect it to an isolated non-production network to avoid conflicts.
2. To start a VM immediately after recovery, select the **Power on VM automatically** check box. If you are recovering the VM to the production network, make sure that the initial VM is powered off to avoid conflicts.
3. Check the settings for instant recovery and click **Next**. Veeam Backup & Replication will restore the VM on the selected ESX(i) host.

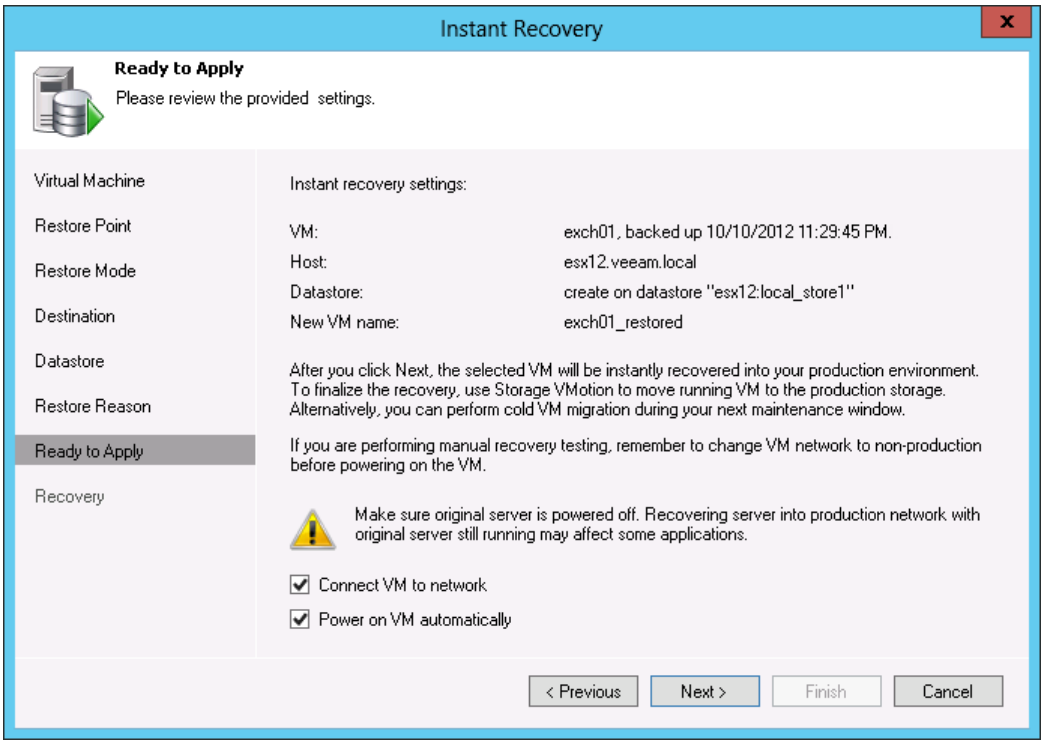

Step 9. Finalize Instant VM Recovery

After the VM has been successfully restored, you can finalize Instant VM Recovery: migrate the restored VM to production or remove the restored VM.

To migrate the restored VM to production:

1. Open the **Backup & Replication** view.
2. In the inventory pane, select the **Instant Recovery** node.
3. Right-click the VM in the working area and select **Migrate to production**. As a result, the Quick Migration wizard will be launched. During migration, Veeam Backup & Replication will restore a VM instance from the backup file and then additionally move the changes that were made while the VM was running in the Instant Recovery mode.

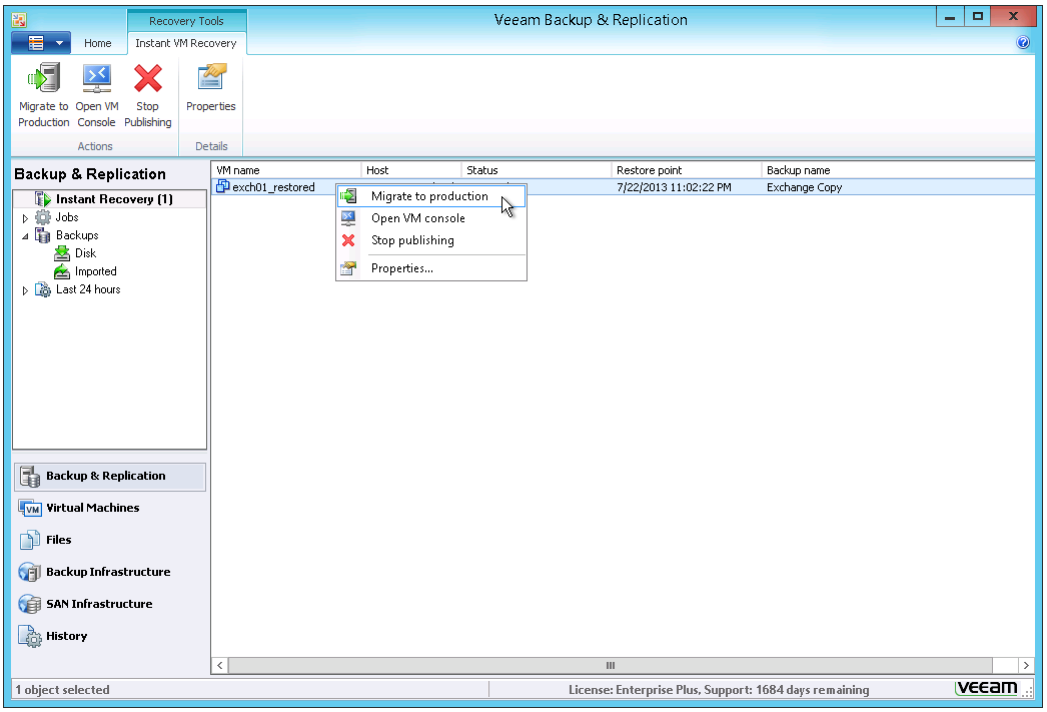

To remove the recovered VM:

1. Open the **Backup & Replication** view.
2. In the inventory pane, select the **Instant Recovery** node.
3. Right-click the necessary VM in the working area and select **Stop publishing**.

Performing Full VM Restore

With the **Full VM Restore** wizard, you can restore the entire VM and start it on the target host if necessary. This section will guide you through all steps of the **Full VM Restore** wizard and provide explanation on available options.

To perform full recovery of a VMware VM, follow the next steps:

Step 1. Launch the Restore Wizard

To launch the **Restore** wizard, do one of the following:

- On the **Home** tab, click **Restore** and select **VMware**. In the **Restore from backup** section, select **Entire VM (including registration)**.
- Open the **Backup & Replication** view and select the **Backups** node. In the working area, expand the necessary backup job, select the VM(s) you want to restore and click **Entire VM** on the ribbon.
- Open the **Backup & Replication** view and select the **Backups** node. In the working area, expand the necessary backup job, right-click the VM(s) you want to restore and select **Restore entire VM**.

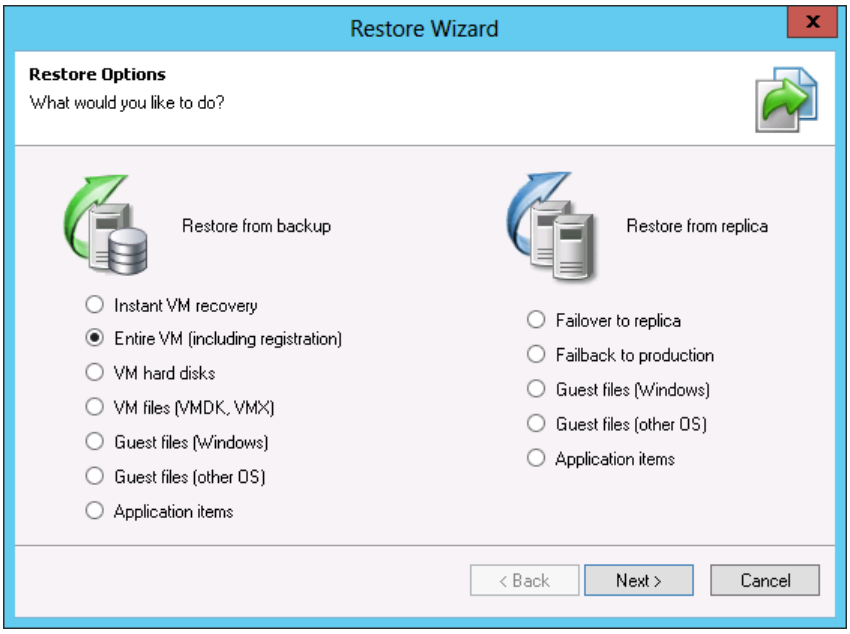

Step 2. Select a Virtual Machine

At this step, you should select one or more VMs to restore. To add a VM or a VM container, click **Add VM** and select where to browse for the machines:

- **From Infrastructure** — browse the virtual environment and select VMs or VM containers to restore. If you choose a VM container, Veeam Backup & Replication will expand it to a plain VM list.
  To facilitate selection, use the search field at the bottom of the **Add Objects** window: click the button to the left of the field and select the necessary type of object to search for (**Everything**, **Folder**, **Cluster**, **Host**, **Resource pool**, **VirtualApp** or **VM**), enter an object's name or a part of it and click the **Start search** button on the right or press **[ENTER]**.
  Make sure that VMs you select from the virtual environment have been successfully backed up at least once.
- **From Backup** — browse existing backups and select VMs under backup jobs. To quickly find VMs, use the search field at the bottom of the **Select Objects** window: enter a VM name or a part of it and click the **Start search** button on the right or press **[ENTER]**.

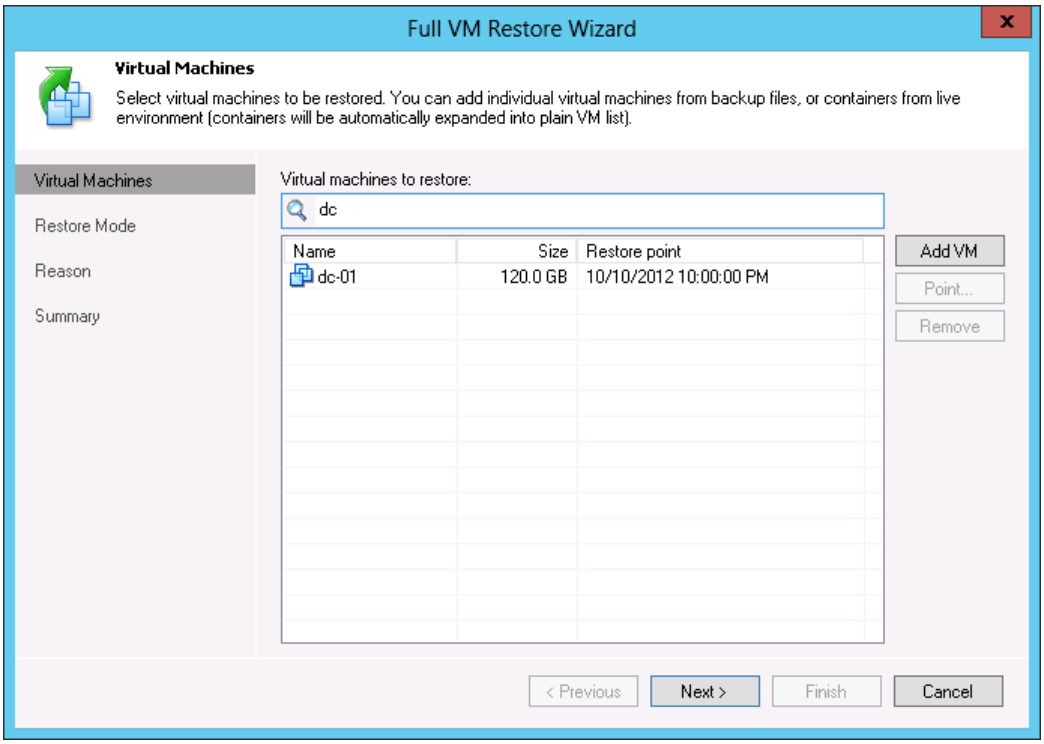

Alternatively, you can use the search field at the top of the window: enter a VM name or a part of it in the search field. Veeam Backup & Replication will search existing backups for the specified VM and display matching results. To add a VM, double-click it in the list of search results. If a VM is not found, click the **Show more** link to browse existing backups and choose the necessary VM.

To remove a VM from the list, select it and click **Remove** on the right.

Step 3. Select a Restore Point

At this step, you should select the necessary restore point for the virtual machine.

By default, Veeam Backup & Replication uses the latest good restore point to recover a VM. However, if you want to restore a VM to an earlier state, select a VM in the **Virtual machines to restore** list and click **Point** on the right. In the **Restore Points** section, select a restore point that should be used to recover the VM.

If you have chosen to restore multiple VMs, you can select a different restore point for every VM specifically.

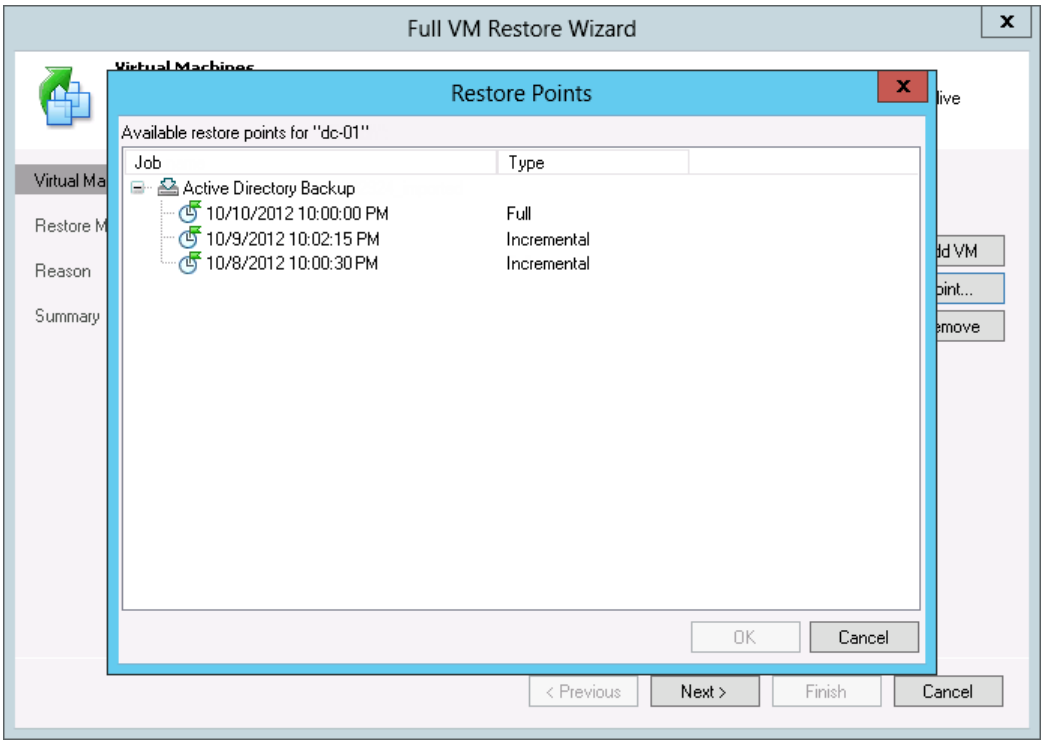

Step 4. Select a Restore Mode

At this step of the wizard, you should select where you want to restore selected VMs.

- Select **Restore to original location** if you want to restore VMs with their initial settings and to their original location. If this option is selected, you will immediately pass to the Reason step of the wizard.
- Select **Restore to a new location, or with different settings** if you want to restore VMs to a different location and/or with different settings (such as, VM location, network settings, the format of restored virtual disks and so on). If this option is selected, the **Full VM Restore** wizard will include additional steps for customizing VMs settings.

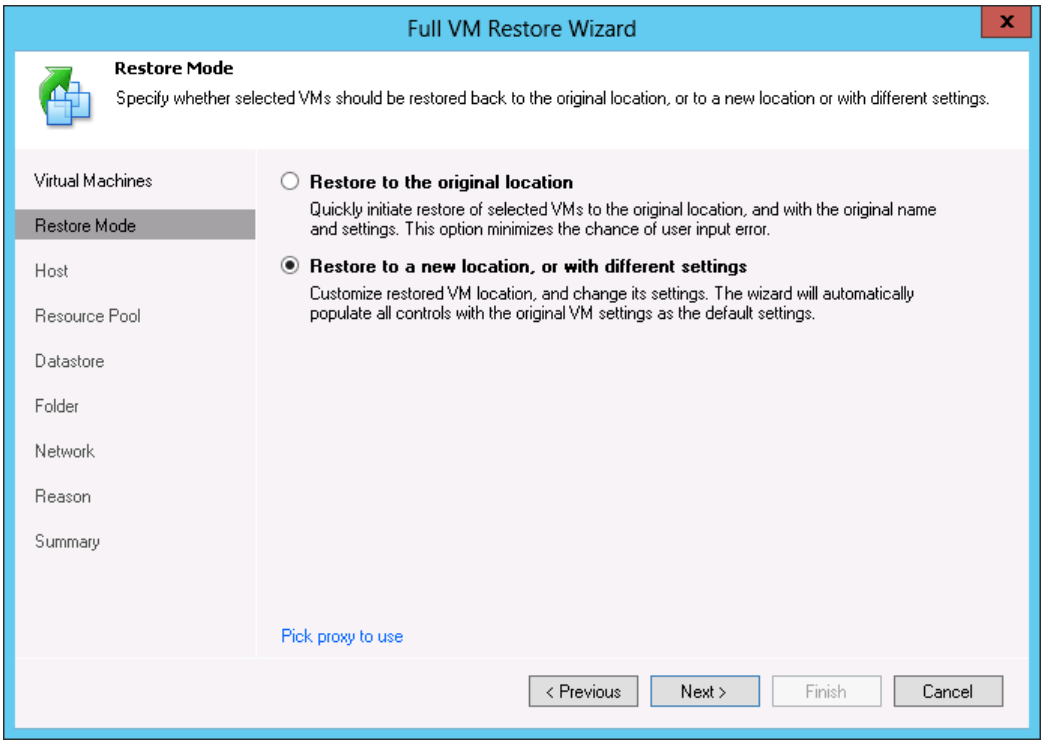

Click the **Pick proxy to use** link to select backup proxies for restore. In the **Backup Proxy** section, you can choose automatic proxy selection or assign proxies explicitly.

- If you choose **Automatic selection**, Veeam Backup & Replication will detect backup proxies that are connected to the source datastore and will automatically assign optimal proxy resources for processing VM data.
  VMs selected for recovery are processed simultaneously. Before restoring, Veeam Backup & Replication checks available backup proxies. If more than one proxy is available, Veeam Backup & Replication analyzes transport modes that the proxies can use for writing data to target, the current workload on the proxies to select the most appropriate resources for VMs processing.
- If you choose **Use the backup proxy servers specified below**, you can explicitly select proxies that will be used to perform restore. It is recommended to select at least two proxies to ensure that recovery will be performed should one of job proxies fail or lose its connectivity to the source datastore.

Step 5. Select a Destination Host for Restored VM

This step of the wizard is available if you have chosen to change the location and settings for the restored VM. To specify a destination host, select a VM in the list and click **Host**. From the virtual environment, choose a host or cluster where the selected VM should be registered.

To facilitate selection, use the search field at the bottom of the window: click the button on the left of the field to select the necessary type of object that should be searched for (**Cluster** or **Host**), enter an object's name or a part of it and click the **Start search** button on the right or press **[ENTER]**.

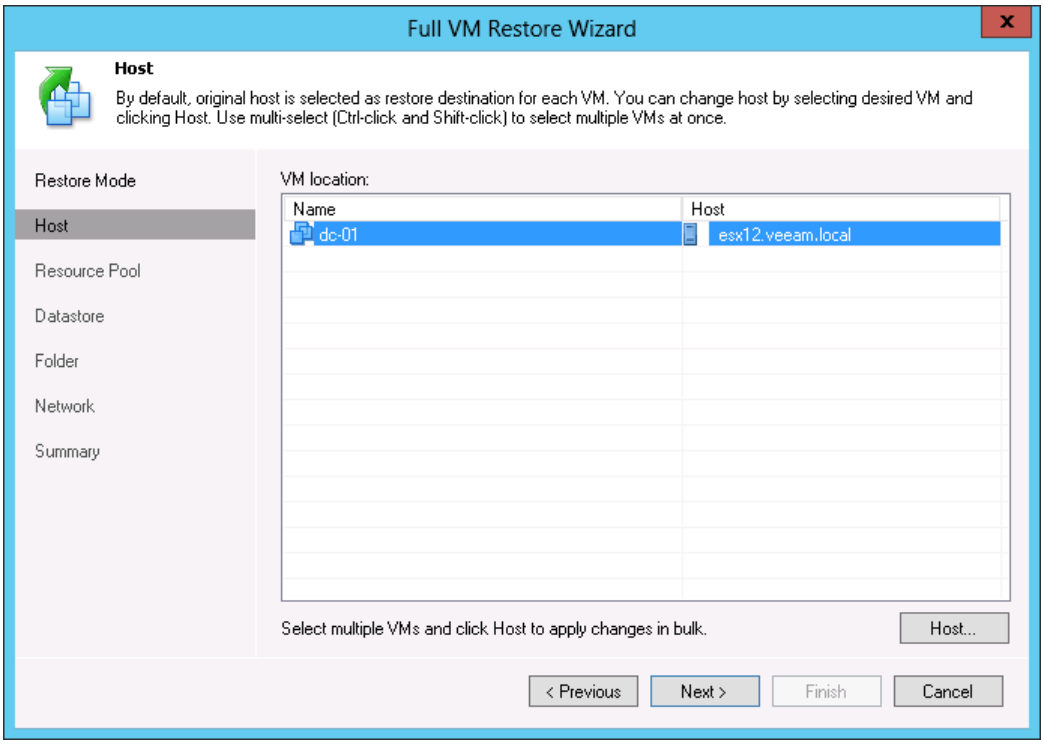

Step 6. Select a Destination Resource Pool

This step of the wizard is available if you have chosen to change the location and settings for the restored VM. To specify a destination resource pool, select a VM in the list and click **Pool**. Select a resource pool to which the selected VM will belong.

To facilitate selection, use the search field at the bottom of the window: enter a resource pool name or a part of it and click the **Start search** button on the right or press **[ENTER]**. If required, you can also select a vApp to which the restored VM will be included.

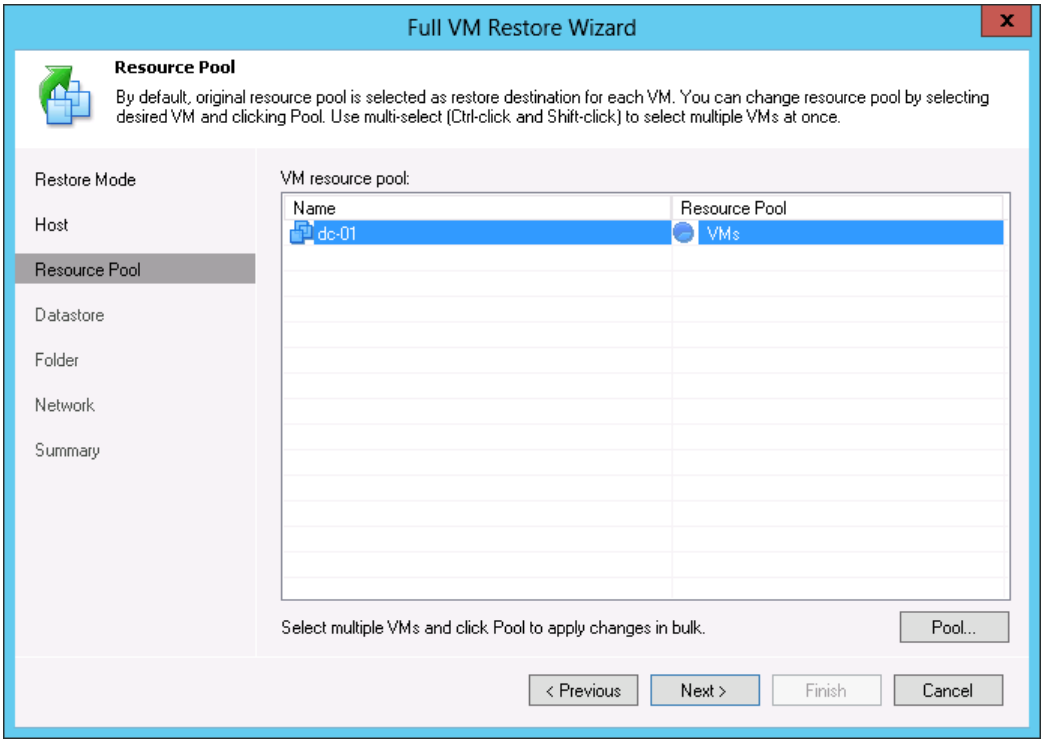

Step 7. Select a Destination Datastore

This step of the wizard is available if you have chosen to change the location and settings for the restored VM. You can place an entire VM to a particular datastore or choose to store configuration files and disk files of the restored VM in different locations.

To specify a destination datastore, select a VM in the list and click **Datastore**. If configuration and disk files of the VM should be placed to different datastores, expand the VM in the list, select the necessary file type and click **Datastore**. Select a datastore to which the selected objects will be stored. To facilitate selection, use the search field at the bottom of the window: enter a datastore name or a part of it and click the **Start search** button on the right or press **[ENTER]**.

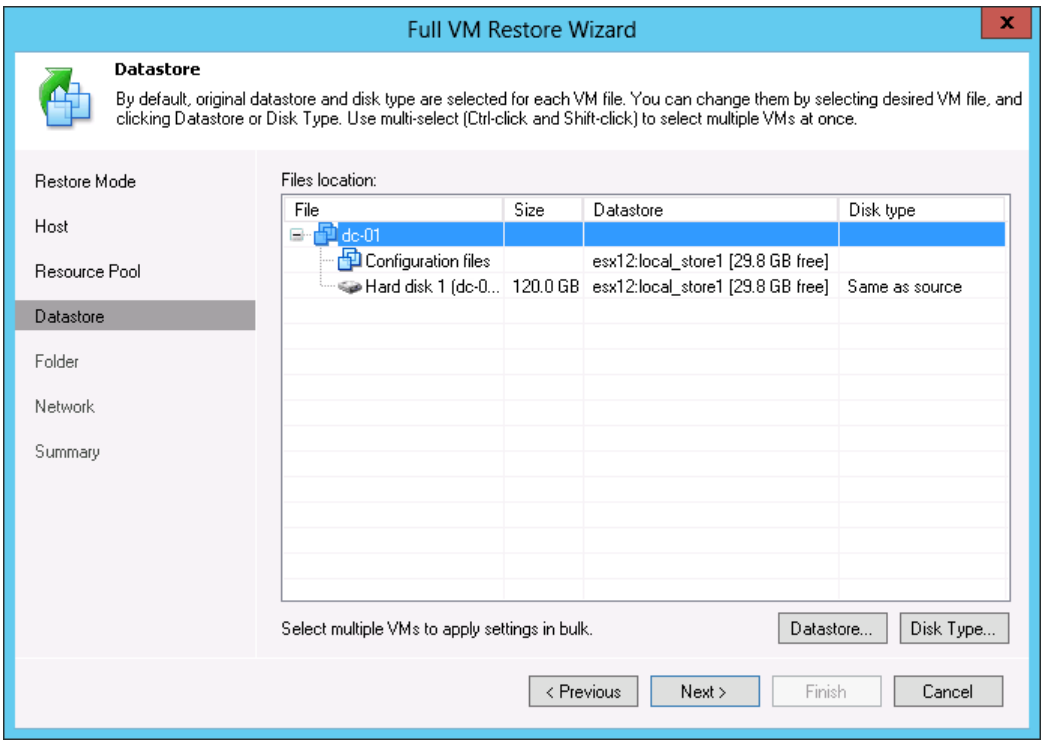

By default, Veeam Backup & Replication preserves the format of restored VM disks, so that if the disks of the original VM were provisioned as thick, Veeam Backup & Replication will restore the VM with thick disks. However, if necessary, you can change the disk format of a restored VM. To do so, expand a VM in the list, select the necessary disk and click **Disk Type**. In the **Disk Type Settings** section, choose the format that will be used to restore virtual disks of the VM — same as the source disk, thin or thick.

Note Disk format changing is available only for VMs using virtual hardware version 7 or later.

Step 8. Select a Destination Folder and Change VM Names

This step of the wizard is available if you have chosen to change the location and settings for the restored VM.

To specify a destination VM folder, select a VM in the list and click **Folder**. From the virtual environment, choose a folder to which the selected VM will belong. To facilitate selection, use the search field at the bottom of the window: enter a folder name or a part of it and click the **Start search** button on the right or press **[ENTER]**.

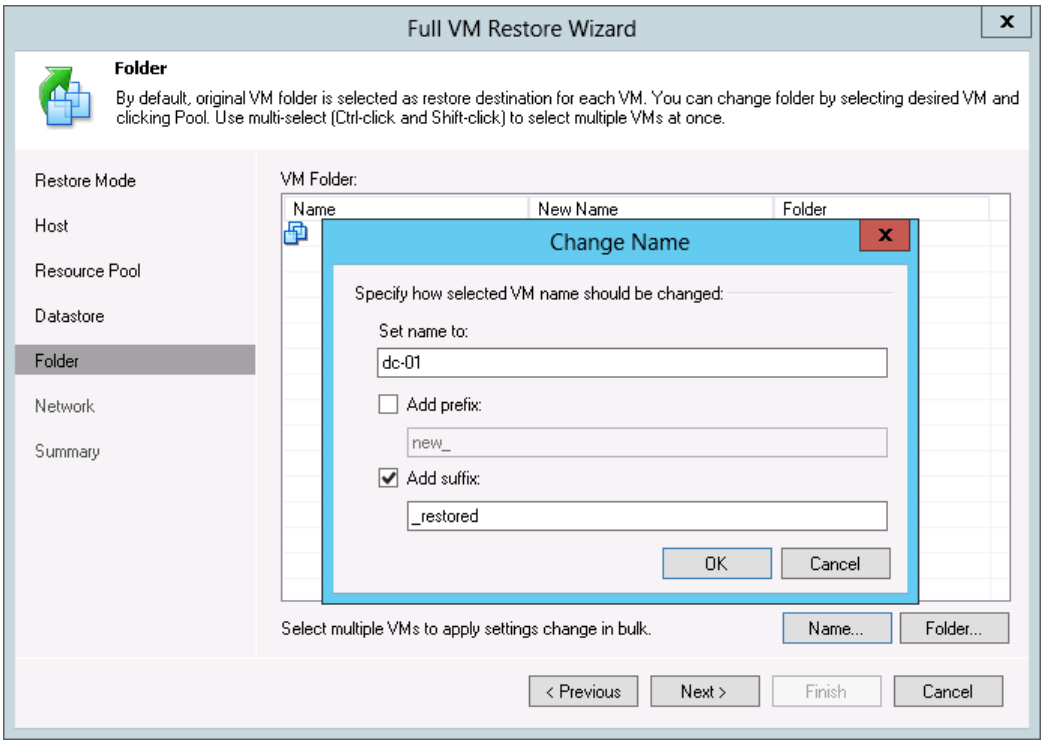

By default, Veeam Backup & Replication restores a VM with its original name. However, you can change the name of the restored VM. For example, if you restore a VM to its original location, you may need to change its name to avoid potential problems.

To change the VM name, select a VM in the list and click **Name**. In the **Change Name** section, you can enter a new name explicitly or specify a change name rule by adding a prefix and/or suffix to the regular VM name. Alternatively, you can change the VM name directly in the list: select a VM, click the **New Name** field and enter the name to be assigned to the restored VM.

Step 9. Select a Destination Network

This step of the wizard is available if you have chosen to change the location and settings for the restored VM. If you plan to restore a VM to a new location (for example, another site with a different set of networks), you can map source site networks to target site networks.
Veeam Backup & Replication will use the network mapping table to update configuration files of the VM on the fly, during the restore process.

To change networks to which the restored VM will be connected, select a VM in the list and click **Networks**. If a VM is connected to multiple networks, expand the VM, select the network to map and click **Network**. The **Select Network** section displays all networks to which the destination host or cluster is connected. From the list of available networks, choose a network to which the selected VM should have access upon restore. To facilitate selection, use the search field at the bottom of the window: enter a network name or a part of it and click the **Start search** button on the right or press **[ENTER]**.

To prevent the restored VM from accessing any network, select the VM or its network connections in the list and click **Disconnected**.

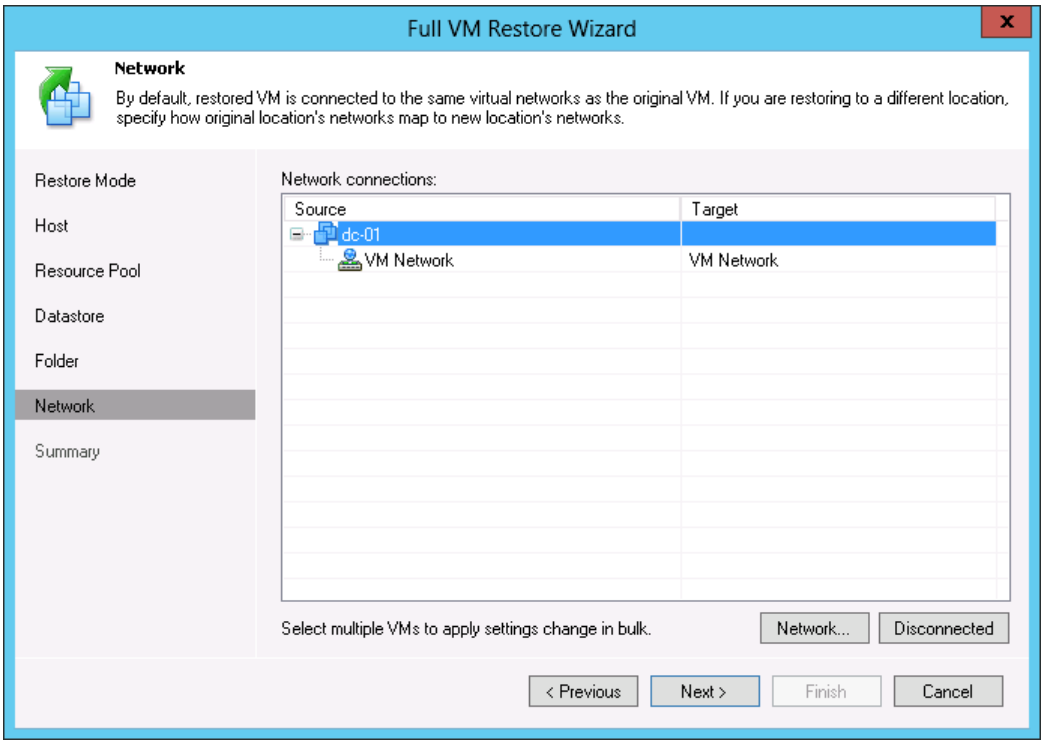

Step 10. Specify a Restore Reason

If necessary, enter the reason for performing restore of selected VMs. The information you provide will be saved in the session history so that you can reference it later.

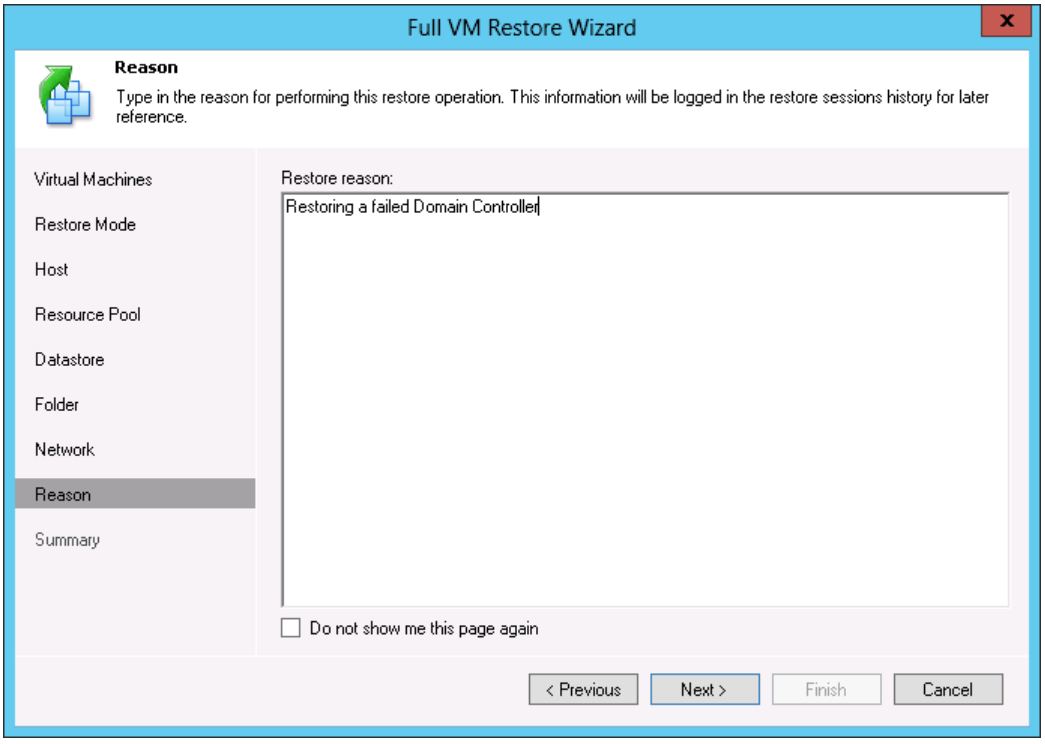

Step 11. Verify Recovery Settings

If you want to start the virtual machine after the work with the wizard is complete, select the **Power on VM after restoring** check box under the list of restore points.

Check specified settings for full VM recovery of a VM and click **Finish**. Veeam Backup & Replication will restore selected VMs in the specified destination.

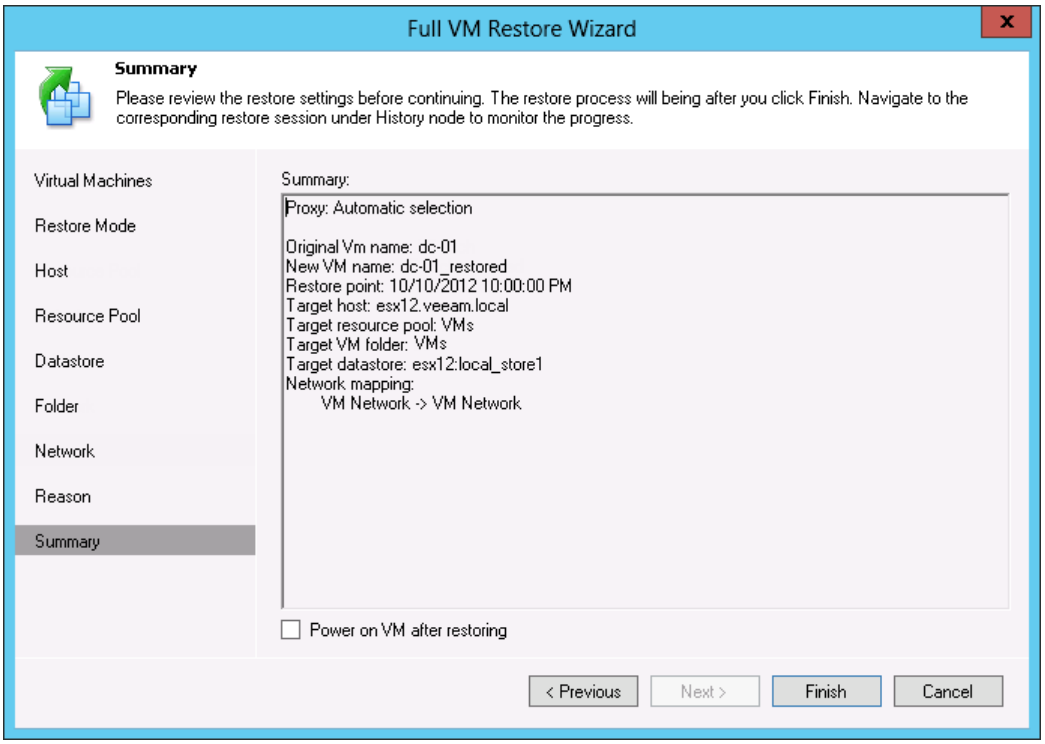

Restoring VM Files

The **Restore** wizard allows you to restore specific VM files (.vmdk, .vmx and others). You can use Veeam VM files recovery to replace deleted or corrupted VM files. This section will guide you through all steps of the **VMware Restore** wizard and provide explanation on available options.

To restore files of a VMware VM, follow the next steps:

Step 1. Launch the Restore Wizard

To launch the **Restore** wizard, do one of the following:

- On the **Home** tab, click **Restore** and select **VMware**. In the **Restore from backup** section, select **VM files (VMDK, VMX)**.
- Open the **Backup & Replication** view and select the **Backups** node. In the working area, expand the necessary backup job, click the VM(s) whose files you want to restore and click **VM Files** > **VM Files** on the ribbon.
- Open the **Backup & Replication** view and select the **Backups** node. In the working area, expand the necessary backup job, right-click the VM(s) whose files you want to restore and select **Restore VM files**.

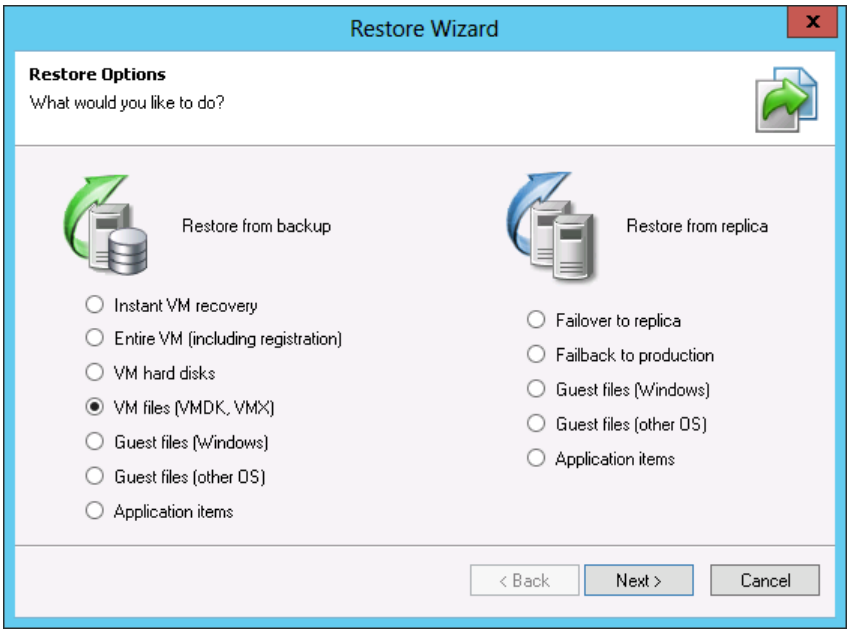

Step 2. Select a Virtual Machine

Select the necessary virtual machine in the list of available jobs. To quickly find VMs in jobs, use the search field at the bottom of the window.

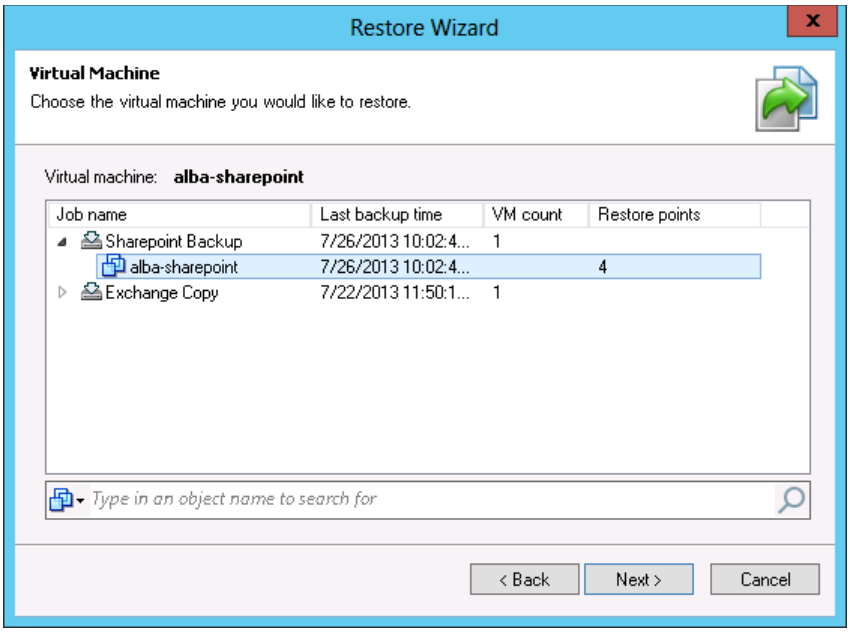

Step 3. Select a Restore Point

Select the necessary restore point for the virtual machine.

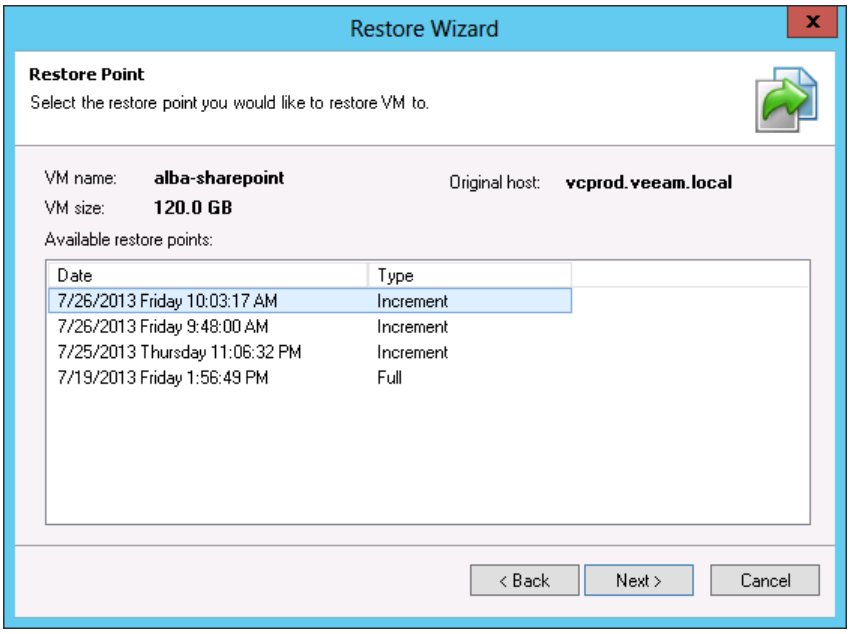

Step 4. Select VM Files and Destination

At this step of the wizard, you should select the VM files you want to restore and the destination where the restored files should be stored. From the **Destination** list, select where to store VM files: to an ESX(i) host, the local machine or any Windows server connected to Veeam Backup & Replication.

Use the **Host Summary** button to view information on available storage resources. In the **Server Properties** section, click **Populate** to load the list of storage locations, their capacity and available space.

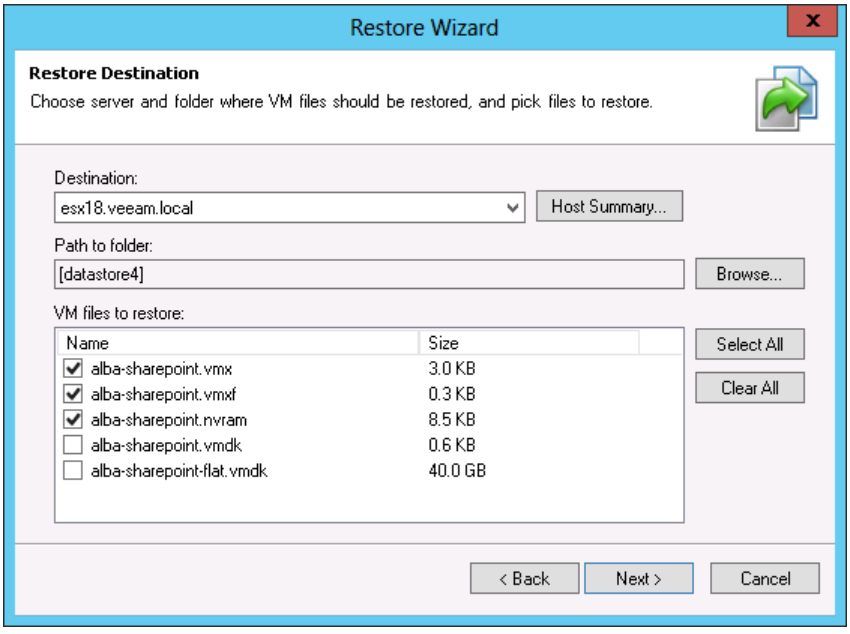

In the **Path to folder** section, specify the path to the folder on the selected host where files should be restored. To create a dedicated folder for restored files, use the **Make New Folder** button at the bottom of the window.

In the **VM files to restore** section, select check boxes next to files that should be restored. By default, all VM files are selected.

Step 5. Specify a Restore Reason

If necessary, enter the reason for performing VM file recovery. The information you provide will be saved in the session history so that you can reference it later.

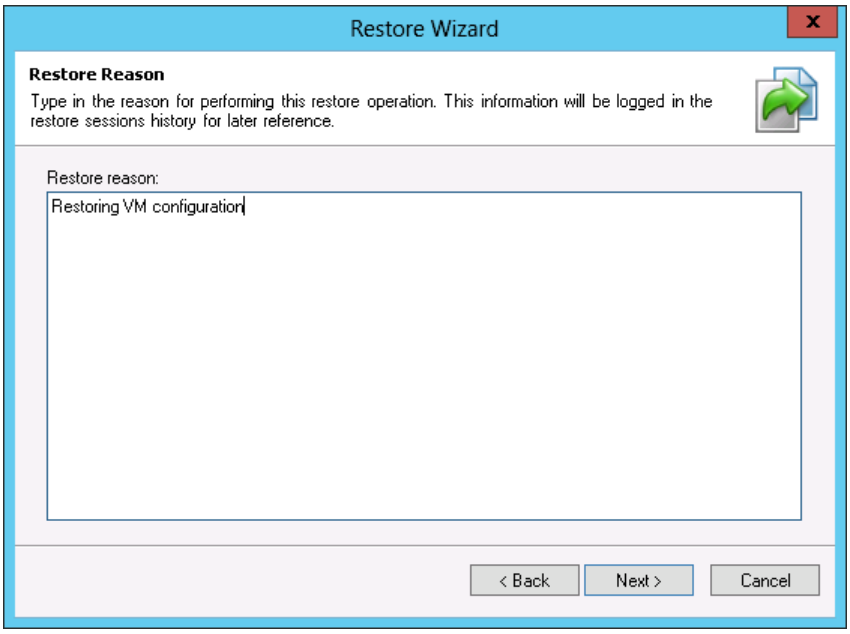

Step 6. Finish Working with the Wizard

Click **Finish** to start restoring the VM files.

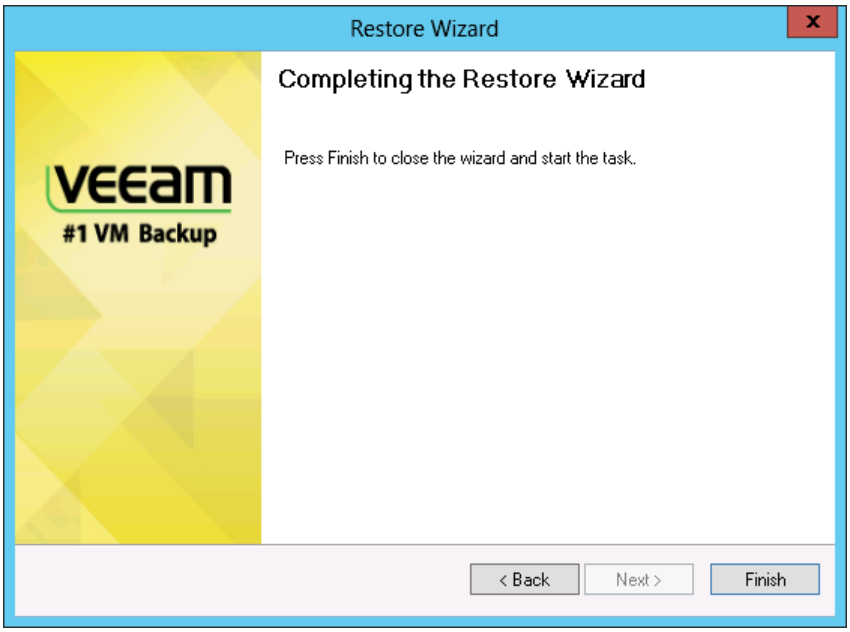

Restoring VM Hard Disks

The **Hard Disk Restore** wizard allows you to restore virtual hard drives of a VM. You can attach restored disks either to the original VM (for example, if you need to replace a corrupted disk) or map them to any other VM. Note that during the virtual hard disk restore, Veeam Backup & Replication turns off the target VM to reconfigure its settings and attach restored disks. For this reason, it is recommended to stop all active processes on the VM for the restore period.

This section will guide you through all steps of the **Hard Disk Restore** wizard and provide explanation on available options.

To restore hard disks of a VMware VM, follow the next steps:

Step 1. Launch the Restore Wizard

To launch the **Hard Disk Restore** wizard, do one of the following:

- On the **Home** tab, click **Restore** and select **VMware**. In the **Restore from backup** section, select **VM hard disks**.
- Open the **Backup & Replication** view and select the **Backups** node. In the working area, expand the necessary backup job, select the VM whose disks you want to restore and click **VM Files** > **Virtual Disk** on the ribbon.
- Open the **Backup & Replication** view and select the **Backups** node. In the working area, expand the necessary backup job, right-click the VM whose disks you want to restore and select **Restore VM hard disk**.

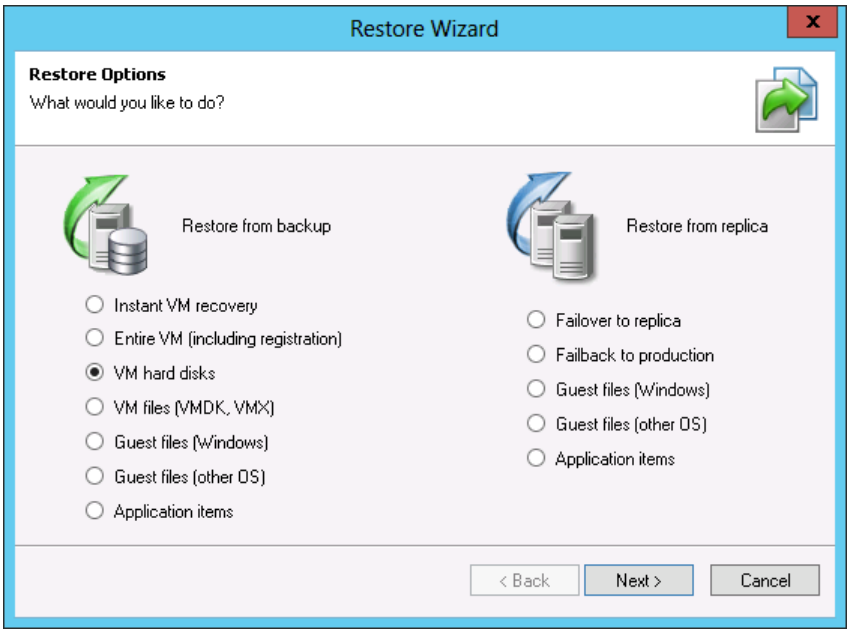

Step 2. Select a Virtual Machine

Select the necessary virtual machine in the list of available jobs. To quickly find VMs in jobs, use the search field at the bottom of the window.

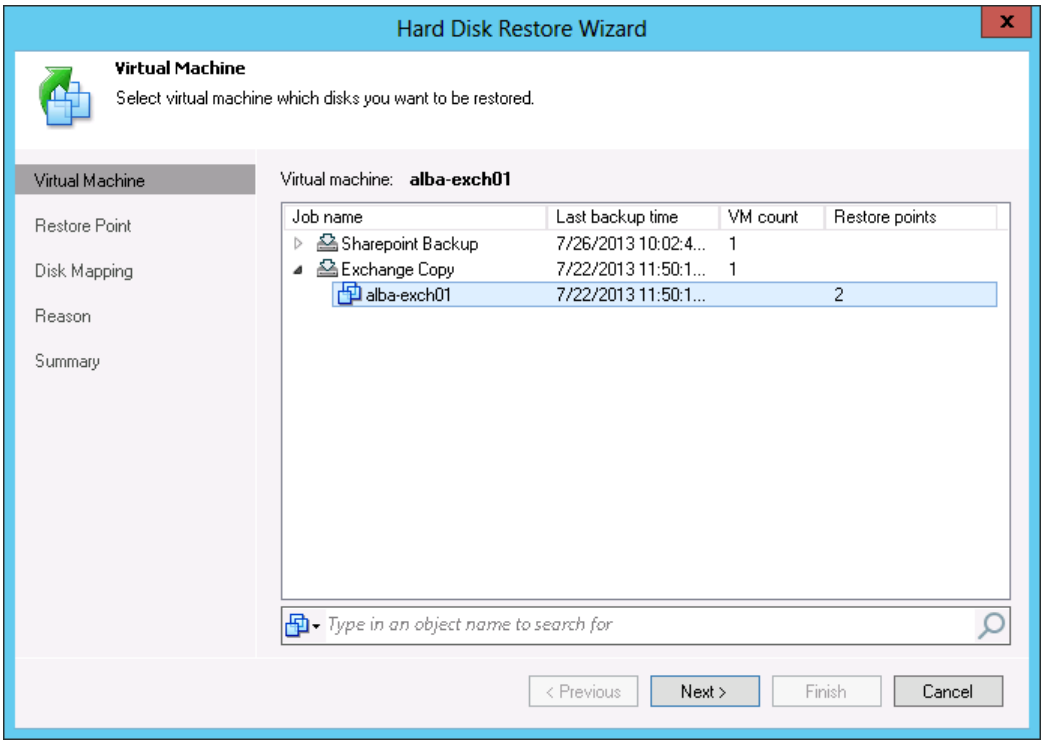

Step 3. Select a Restore Point

Select the necessary restore point for the virtual machine.

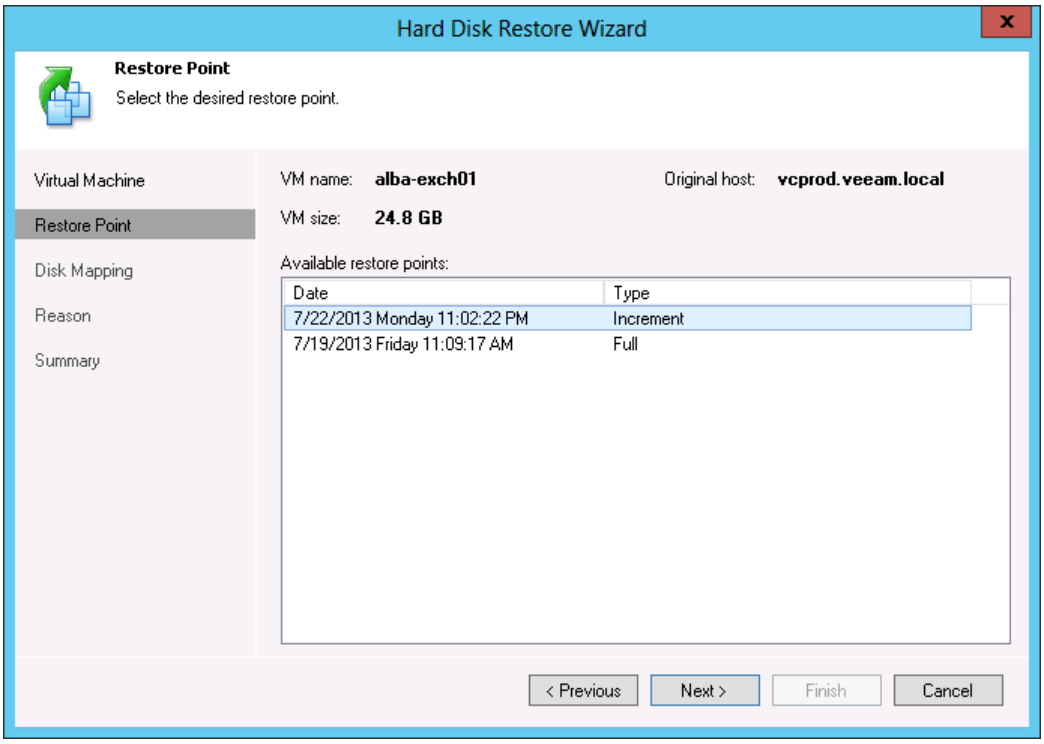

Step 4. Select Virtual Hard Disks to Restore

At this step, you should select virtual hard disks to restore, choose a VM to which the disks will be attached and define additional restore properties.

By default, Veeam Backup & Replication maps restored disks to the original VM. If the original VM was relocated or if you want to attach disks to another VM, you need to select the target VM manually.

Click **Browse** and select the necessary VM from the virtual environment. To facilitate selection, use the search field at the bottom of the window: click the button on the left of the field to select the necessary type of object that should be searched for, enter a VM name or a part of it and click the **Start search** button on the right or press **[ENTER]**.

Select check boxes next to virtual hard disks that you want to restore. To define virtual disk properties, select a disk in the list and click **Change**. In the **Virtual Disk Properties** section, pick a datastore where the restored hard disk will be located and select a virtual device node.

- If you want to replace an existing virtual disk, select an occupied virtual node.
- If you want to attach the restored disk to the VM as a new drive, select a node that is not yet occupied.

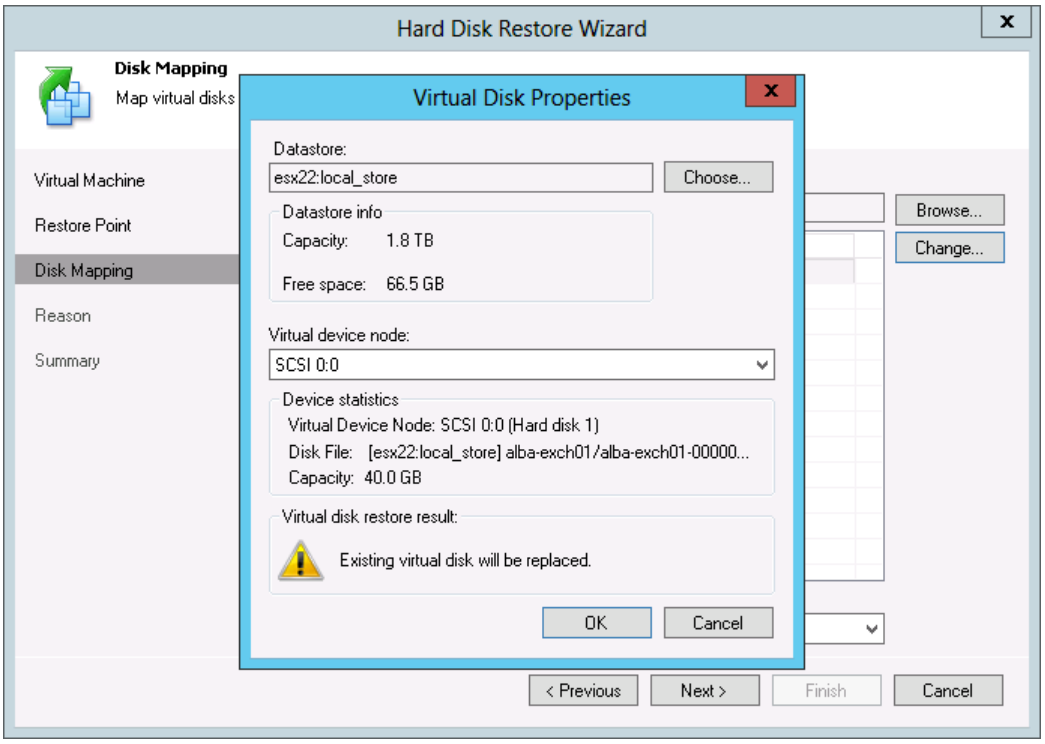

Veeam Backup & Replication preserves the format of the restored virtual hard disks. To change disk format, select the required option from the **Restore disks** list – same as on the original VM, force thin or force thick. Please note that disk format changing is available only for VMs using virtual hardware version 7 or later.

Step 5. Specify a Restore Reason

If necessary, enter the reason for performing VM file recovery. The information you provide will be saved in the session history so that you can reference it later.

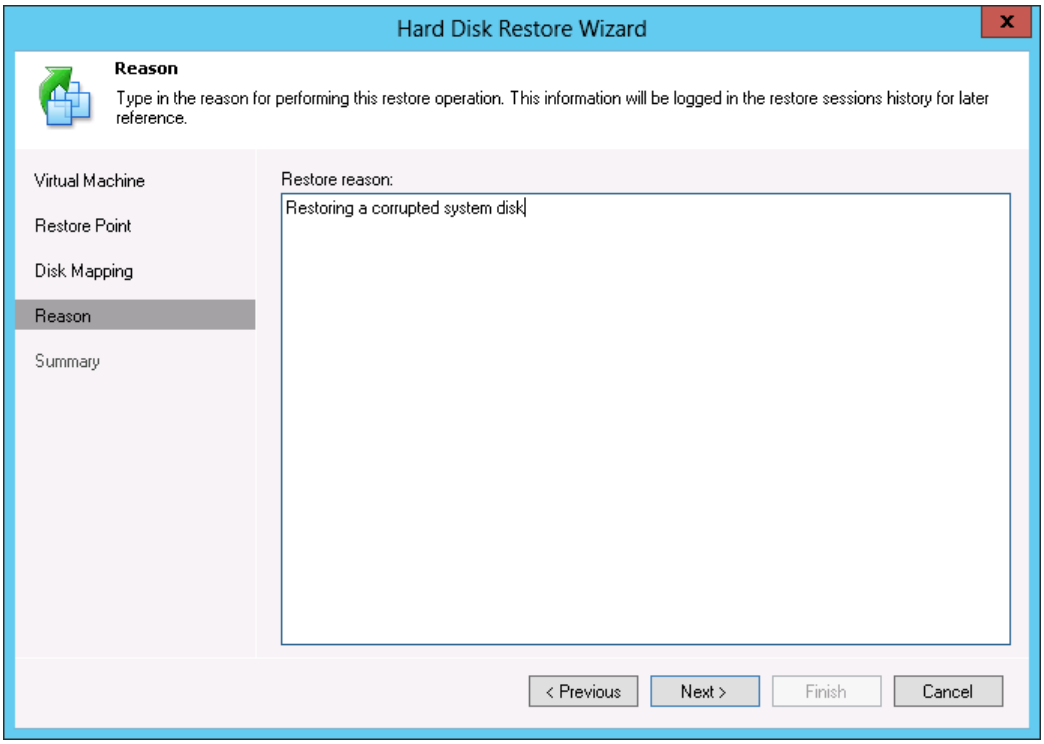

Step 6. Finish Working with the Wizard

To start a VM immediately after hard disk recovery, select **Power on VM after restoring**. Then click **Finish** to start restoring the VM files.

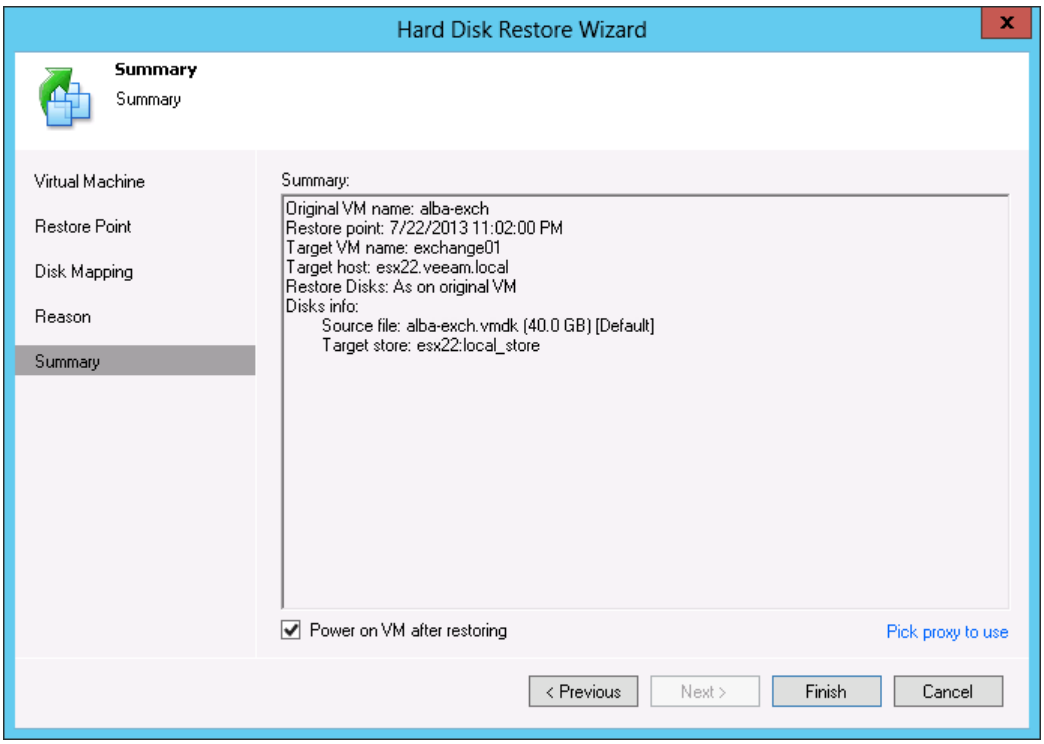

Restoring VM Guest Files

With the **Restore** wizard, you can restore individual Windows guest OS files from any successfully created backup or replica of a Windows-based VM.

When you perform file-level recovery, the VM image is not extracted from the backup. The content of a backup file is mounted directly to the Veeam backup server (to the `C:\veeamflr\<vmname>` folder) and displayed in the inbuilt Veeam Backup browser. For mounting file systems of VM guest OS'es, Veeam Backup & Replication uses its proprietary driver. After the file system is mounted, you can copy necessary files and folders to their initial location, to your local machine drive, save them anywhere within the network or simply point any applications to the files and use them normally.

Important! File-level restore has the following limitations:

- You cannot restore files from a running replica, or if the replication job with the necessary VM is being performed.
- You cannot restore files from a backup created in the reversed incremental mode if the backup job is being performed. However, if the backup is created in the incremental backup mode and the backup job is being performed, you can restore files from any available restore point.
- Guest OS file-level restore for ReFS is supported only if Veeam Backup & Replication is installed on Windows 2012 Server.

To restore guest OS files from a VMware VM, follow the next steps:

Step 1. Launch the Restore Wizard

To launch the **Restore** wizard, do one of the following:

- On the **Home** tab, click **Restore** and select **VMware**. In the **Restore from backup** section, select **Guest files (Windows)**.
- Open the **Backup & Replication** view and select the **Backups** node. In the working area, expand the necessary backup job, select the VM whose guest OS files you want to restore and click **Guest Files** > **Guest Files (Windows)** on the ribbon.
- Open the **Backup & Replication** view and select the **Backups** node. In the working area, expand the necessary backup job, right-click the VM whose guest OS files you want to restore and select **Restore guest files (Windows)**.

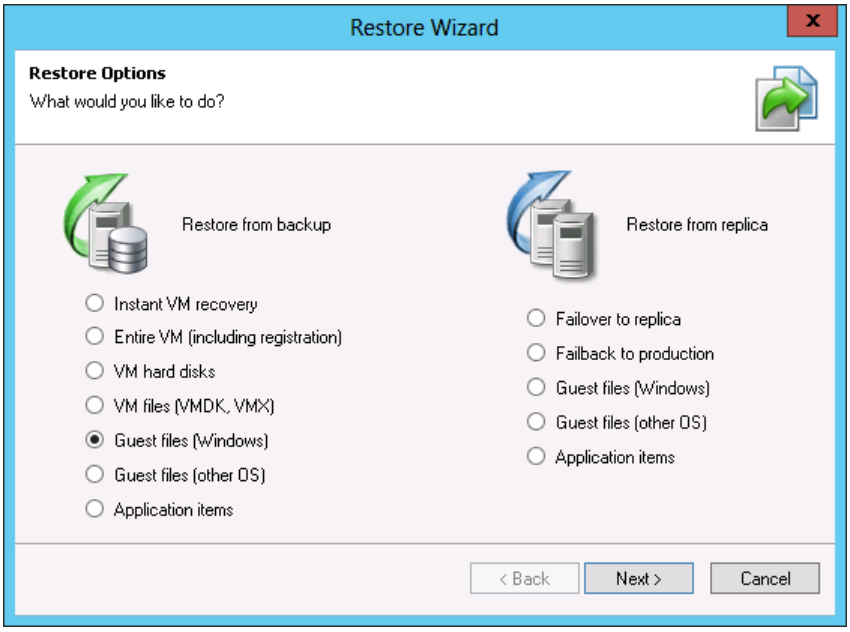

Step 2. Select a Virtual Machine

In the list of available jobs, select the necessary virtual machine. To quickly find VMs in jobs, use the search field at the bottom of the window.

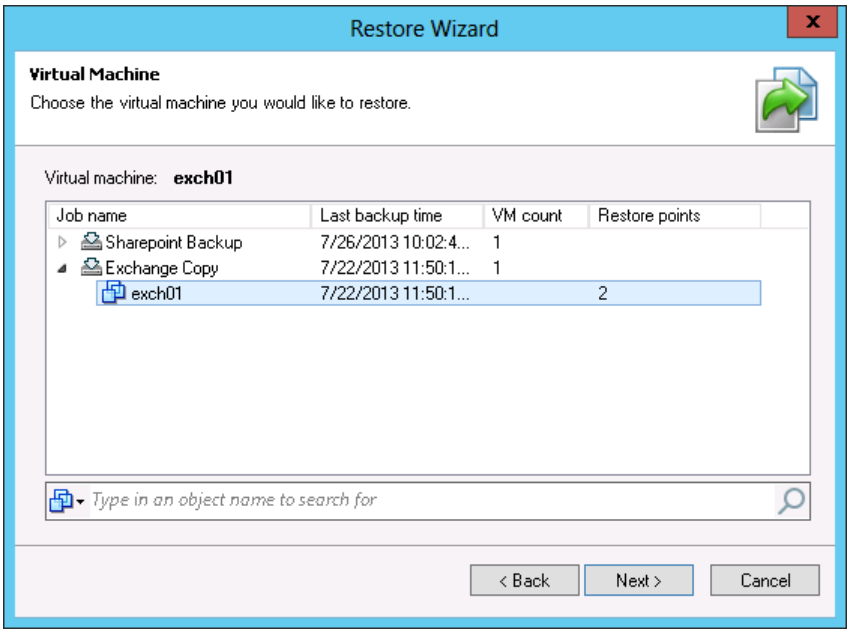

Step 3. Select a Restore Point

Select the necessary restore point for the virtual machine.

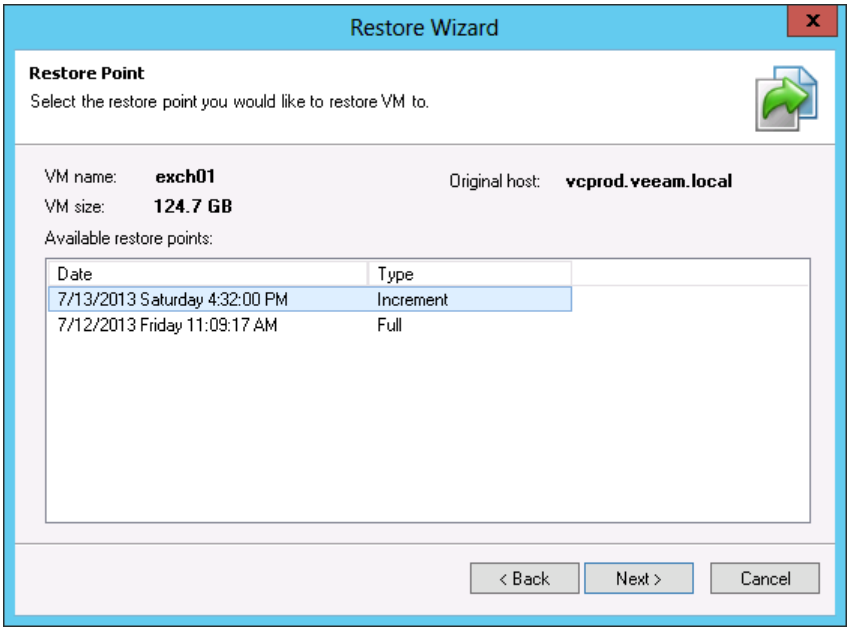

Step 4. Specify Restore Reason

If necessary, enter the reason for performing VM guest file restore. The information you provide will be saved in the session history so that you can reference it later.

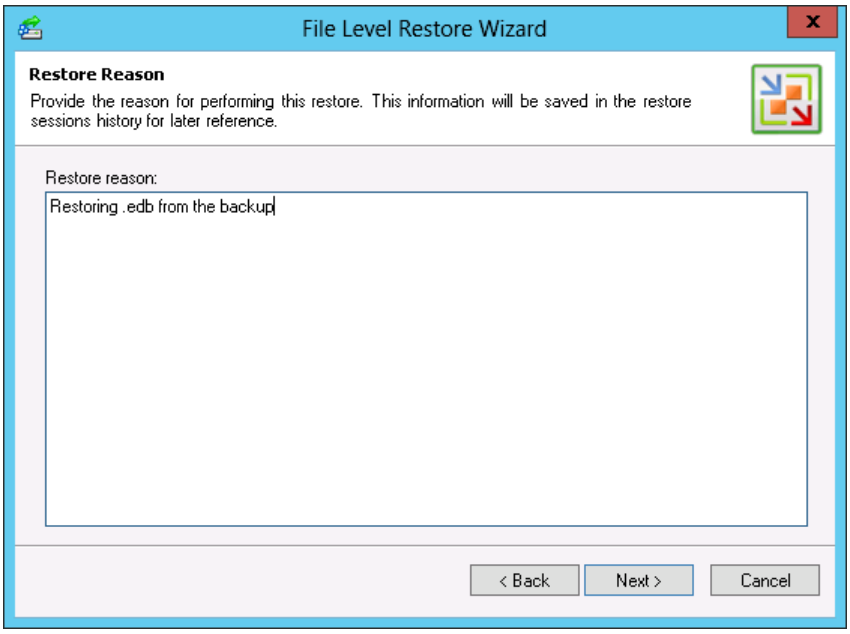

Step 5. Finish Working with the Wizard

Click **Finish** to start restoring files from a backup or replica. Once restoring is completed, Veeam Backup & Replication will open a file browser displaying the file system tree of the restored VM. Please note that the names of the restored machine drives may differ from the original ones.

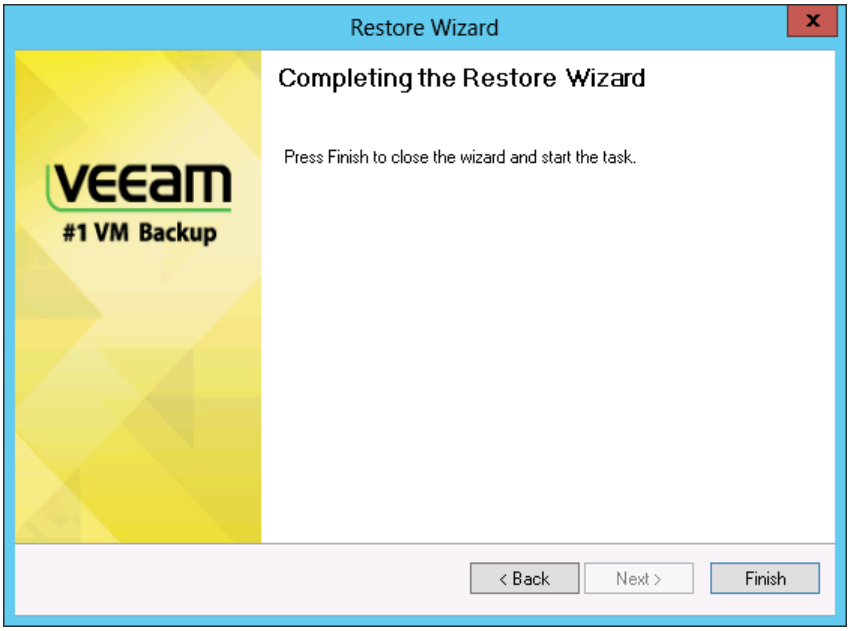

Step 6. Save Restored Files

You can save guest OS files to their initial location, to any folder on the local machine or within the network or open Windows Explorer for work with files.

Note	You can browse to the VM guest OS files mounted to the Veeam backup server only while the Veeam Backup browser with the restored files is open. After the Veeam Backup browser is closed, the VM disks will be unmounted from the Veeam backup server.

Saving Files to the Initial Location

To save files or folders to their initial location:

1. Right–click the necessary file or folder in the file system tree or in the details pane on the right and select **Restore**.
2. In the **Credentials** window, specify credentials of the account that will be used to connect to the initial VM. When you restore files to their initial location, Veeam Backup & Replication deploys a small runtime process in the initial VM. The process is used to control restore operations.
   To deploy the process, you need to connect to the initial VM under an account having administrator permissions on this VM. You can use the account under which you are currently logged on or choose another account.
3. Click **OK** to start the restore process.

Important! Restore to the initial location may fail for the following reasons:

- VMware Tools are not installed on the target VM
- You have excluded the system disk from the VM backup

To restore guest OS files in such situation, you can use 1-click file-level restore or copy files to the selected folder and then move them to the original location.

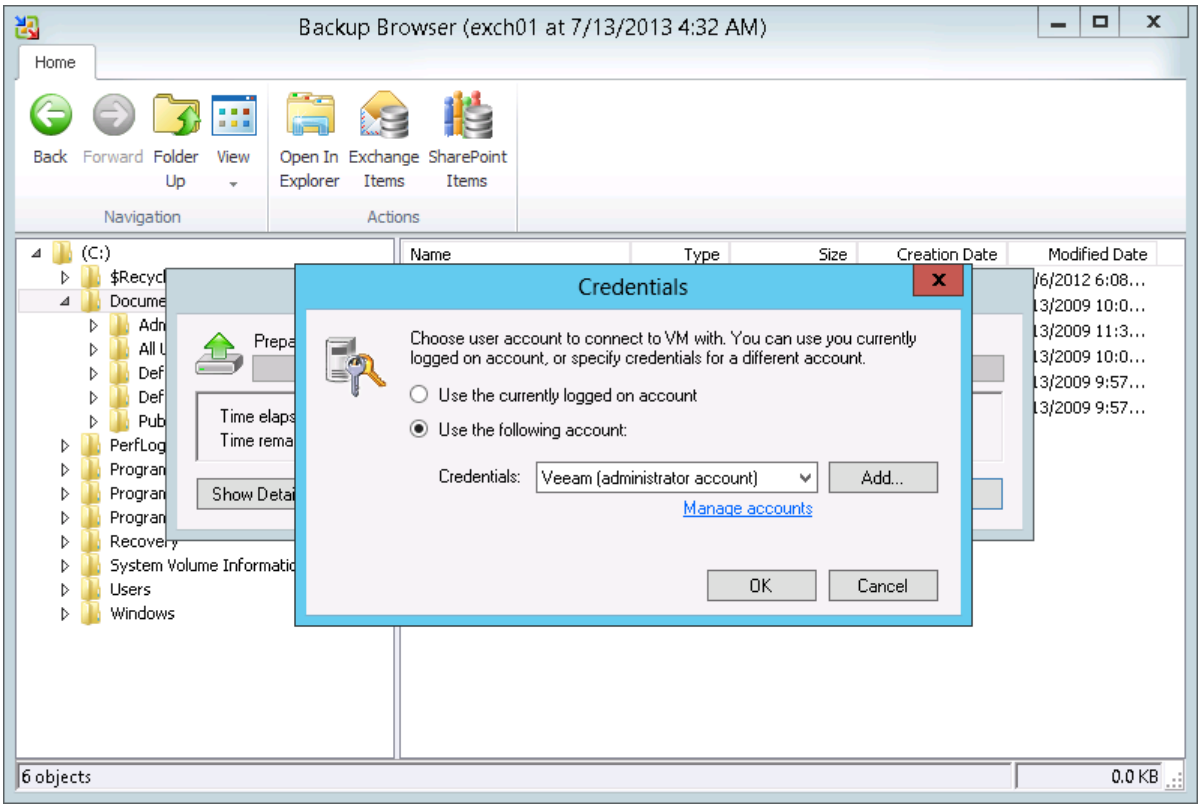

Saving Files to a New Location

To save restored files or folders on the local machine or within the network, right–click the necessary file or folder in the file system tree or in the details pane on the right and select **Copy To**.

When restoring file objects, you can choose to preserve their original NTFS permissions:

- Select the **Preserve permissions and ownership** check box to keep the original ownership and security permissions for restored objects. Veeam Backup & Replication will copy selected files and folders along with associated Access Control Lists, preserving granular access settings.
- Leave the **Preserve permissions and ownership** check box cleared if you do not want to preserve the original ownership and access settings for restored objects. In this case, Veeam Backup & Replication will change security settings: the user who launched the Veeam Backup & Replication console will be set as the owner of the restored object, while access permissions will be inherited from the folder to which the restored object is copied.

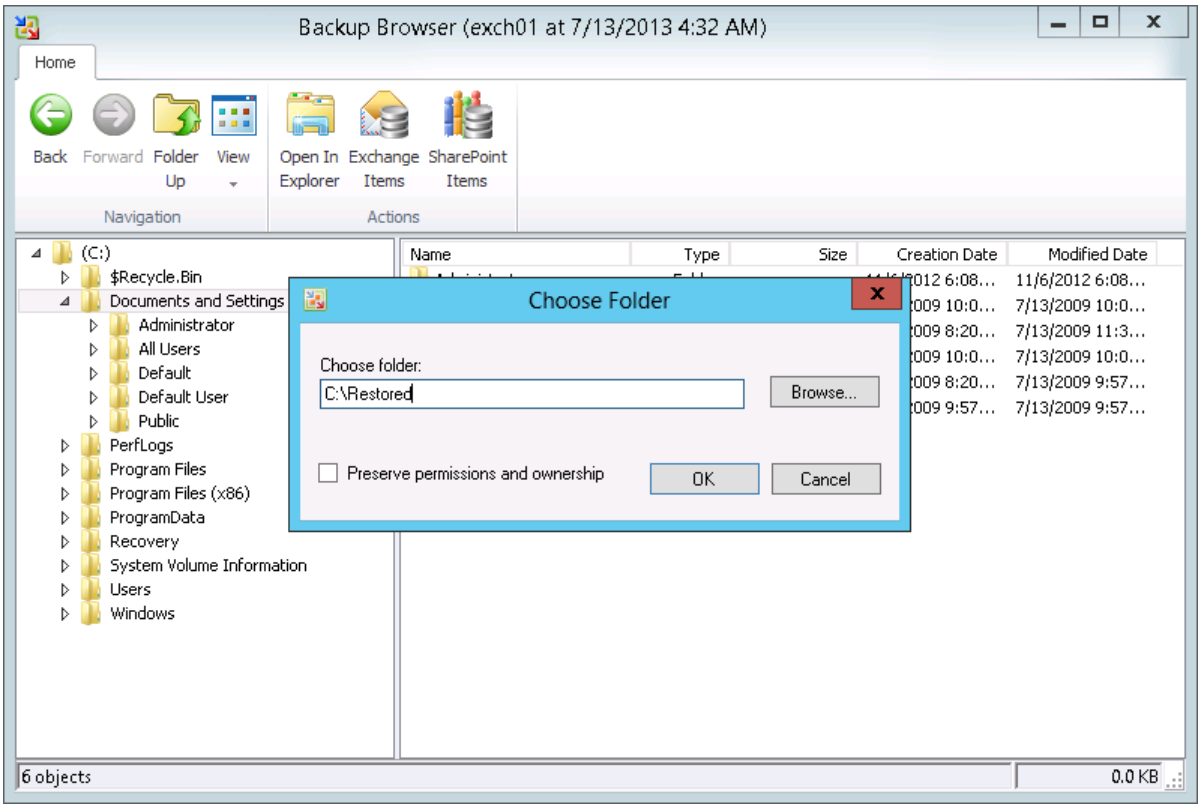

Tip If you are restoring guest OS files of the virtualized Microsoft Exchange server or Microsoft Sharepoint server, you can launch Veeam Explorer for Exchange and Veeam Explorer for SharePoint directly from the Veeam Backup browser:

- To start Veeam Explorer for Exchange, browse to the Exchange database file (EDB) in the Veeam Backup browser, select it and click **Exchange Items** on the **Home** tab or simply double-click the EDB file.
- To start Veeam Explorer for SharePoint, browse to the Microsoft SharePoint content database (MDF) in the Veeam Backup browser, select it and click **SharePoint Items** on the **Home** tab or simply double-click the MDF file.

Working with Windows Explorer

Beside copying files via the Veeam Backup browser, you can use Windows Explorer to work with restored files. Click **Explore** on the ribbon in the Veeam Backup browser or right-click the necessary folder and select **Explore**. Veeam Backup & Replication will launch Windows Explorer so that you can browse to VM guest OS files.

You can also start Windows Explorer as usually and browse to the necessary files. VM disks are mounted under the `C:\veeamflr\<vmname>\<volume n>` folder of the Veeam backup server.

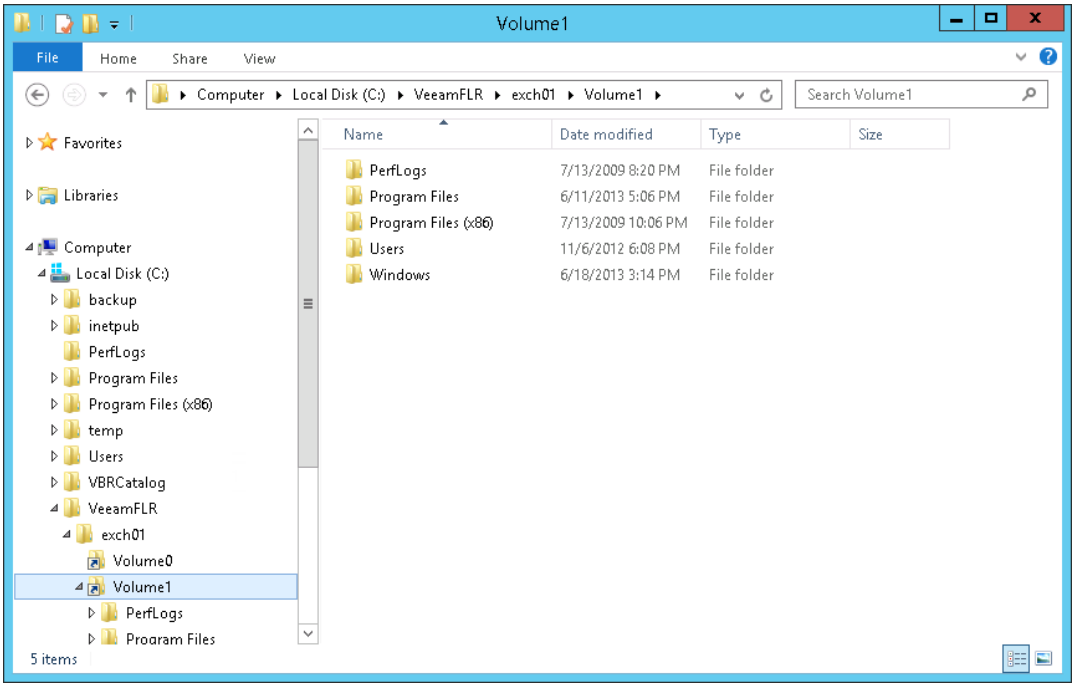

Multi-OS File Level Recovery

To let you recover guest OS files, Veeam Backup & Replication uses a specific proxy appliance — a helper VM. The proxy appliance is very small — around 20 MB and takes only 10-20 seconds to boot. Veeam Backup & Replication automatically starts the proxy appliance on the host in the virtual environment and mounts disks of the restored VM to the proxy appliance as virtual hard drives. VM files are mounted directly from backup files, without prior extraction of the backup file content. After that, you can copy necessary files and folders to your local machine drive or save them anywhere within the network.

To perform multi-OS file-level restore, follow the next steps:

Step 1. Launch the Veeam File Level Restore Wizard

To launch the **Restore** wizard, do one of the following:

- On the **Home** tab, click **Restore** and select **VMware**. In the **Restore from backup** section, select **Guest files (other OS)**.
- Open the **Backup & Replication** view and select the **Backups** node. In the working area, expand the necessary backup job, select the VM whose guest OS files you want to restore and click **Guest Files** > **Guest Files (other OS)** on the ribbon.
- Open the **Backup & Replication** view and select the **Backups** node. In the working area, expand the necessary backup job, right-click the VM whose guest OS files you want to restore and select **Restore guest files (other OS)**.

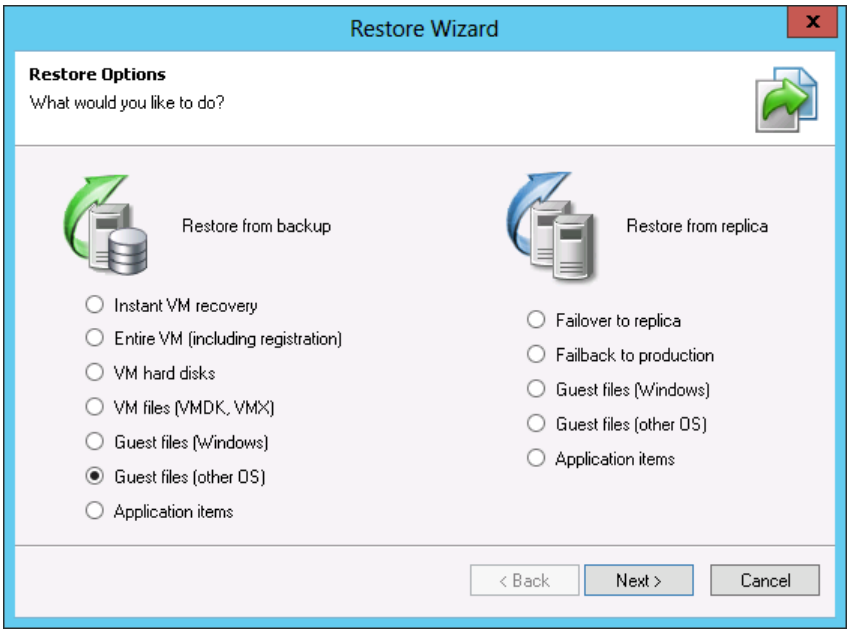

The welcome screen of the wizard will be displayed. If you do not want to see the welcome screen at subsequent launches of the wizard, select the **Don't show this step again** check box at the bottom of the screen.

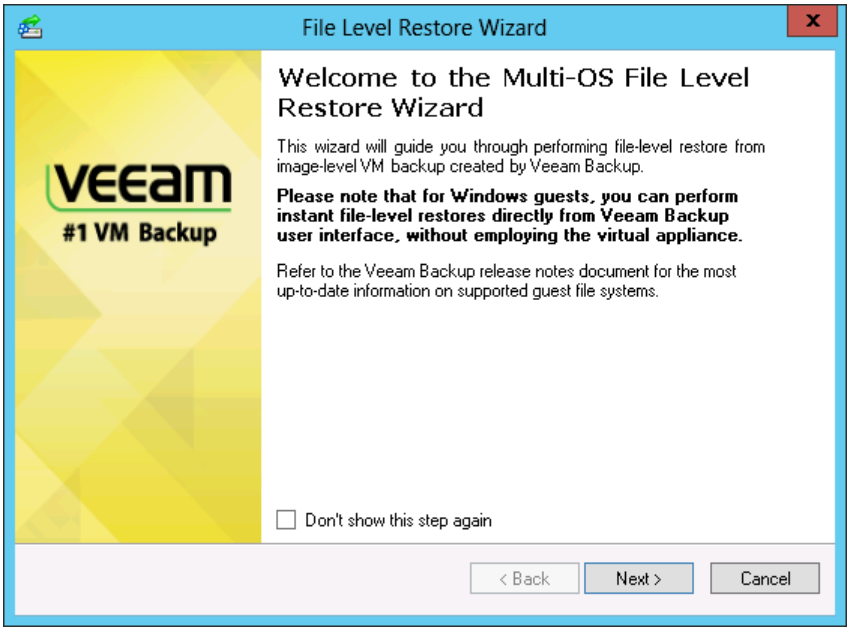

Step 2. Select a Virtual Machine

At the **Virtual Machine** step of the wizard, select the necessary virtual machine.

Tip	To quickly find the necessary VM, use the search field at the bottom of the window: enter the VM name or a part of it and press **[ENTER]**.

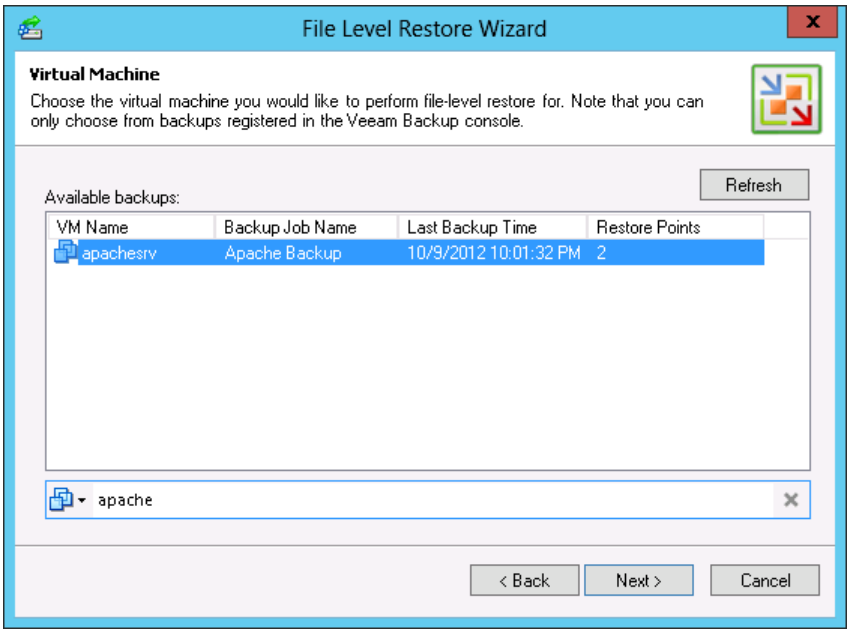

Step 3. Select a Restore Point

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

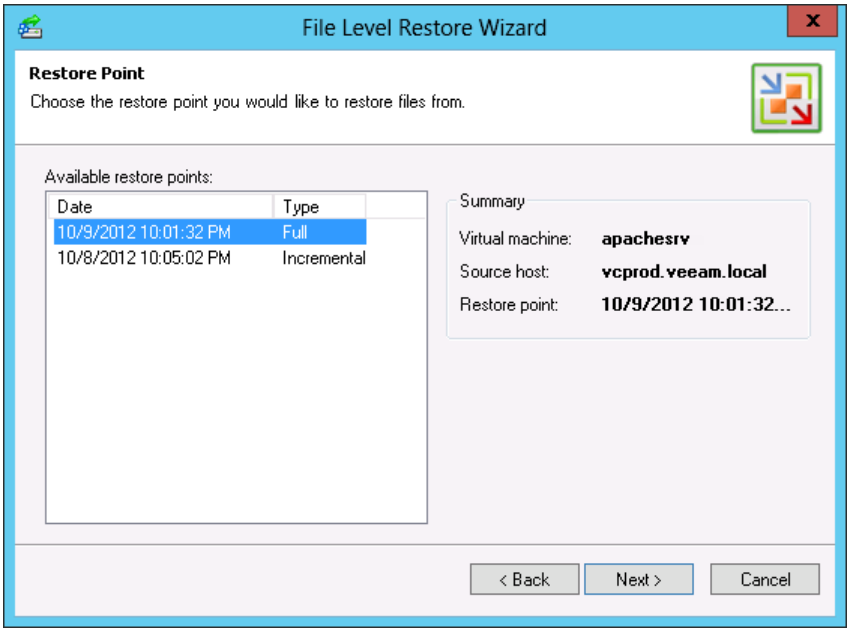

Step 4. Specify a Restore Reason

At the **Restore Reason** step of the wizard, enter the reason for restoring files if necessary. The information you provide will be saved in the session history so that you can reference it later.

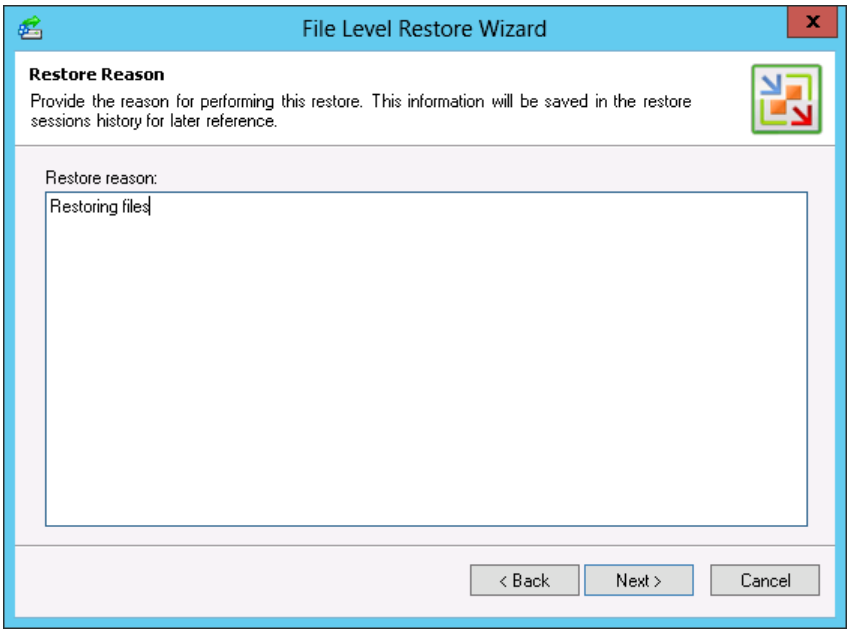

Step 5. Select Location for the Proxy Appliance

At the **Ready** step of the wizard, you should select an ESX(i) host for placing the proxy appliance. When the restore process starts, Veeam Backup & Replication will register the proxy appliance on the selected ESX(i) host and mount disks of the restored VM to this proxy appliance. The file system tree of the restored VM will be displayed in the Veeam Backup browser. After you restore necessary files and finish working with the Veeam Backup browser, the proxy appliance will be deleted from the ESX(i) host.

To locate the appliance, do the following:

1. Click **Customize** at the bottom of the window.
2. In the **FLR Appliance Configuration** window, select the ESX(i) host on which the proxy appliance will be registered.
3. Specify the resource pool and network in which the proxy appliance will be run.
4. Select between a static or dynamic IP address for the proxy appliance and specify the necessary network settings for the proxy appliance.
5. To enable FTP access to the restored file system, select the **Enable FTP server on appliance** check box. As a result, your users will be able to access the proxy appliance via FTP, browse the file system of the restored VM and download necessary files on their own.
6. If you are performing restore of a VM with the Novell Storage services file system, select the **Restore Novell Storage services filesystem** check box. In this case, Veeam Backup & Replication will deploy a specific proxy appliance supporting the Novell file system.
7. Click **OK**.

Important! When choosing an ESX(i) host for the Novell file system proxy appliance, make sure that it allows running VMs with 64-bit guest OSs.

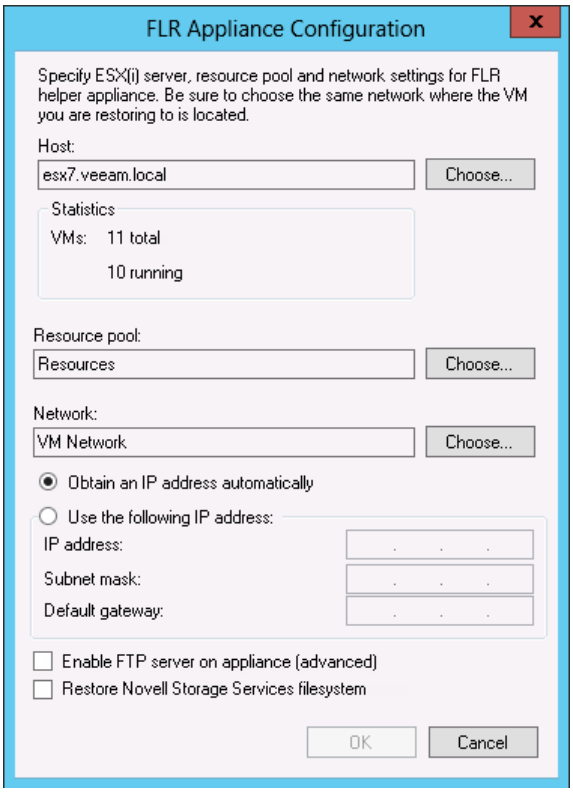

Step 6. Finish Working with the Wizard

Click **Finish** to start restoring files from a backup or replica. Please note that the file-level restore appliance may take about 30-40 seconds to boot.

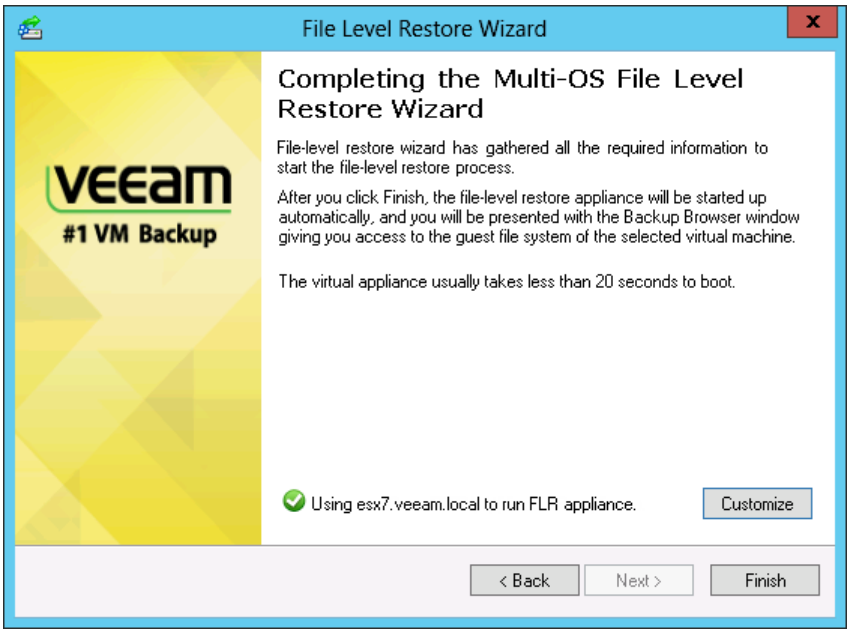

Step 7. Save Restored Files

Once the restore process is completed, a file browser displaying the file system tree of the restored virtual machine will be opened. To save restored files or folders on the local machine or within the network, right–click the necessary file or folder and select **Copy to** from the shortcut menu and select the necessary destination and folder on the local or remote host. The file or folder will be saved at the specified folder on the host.

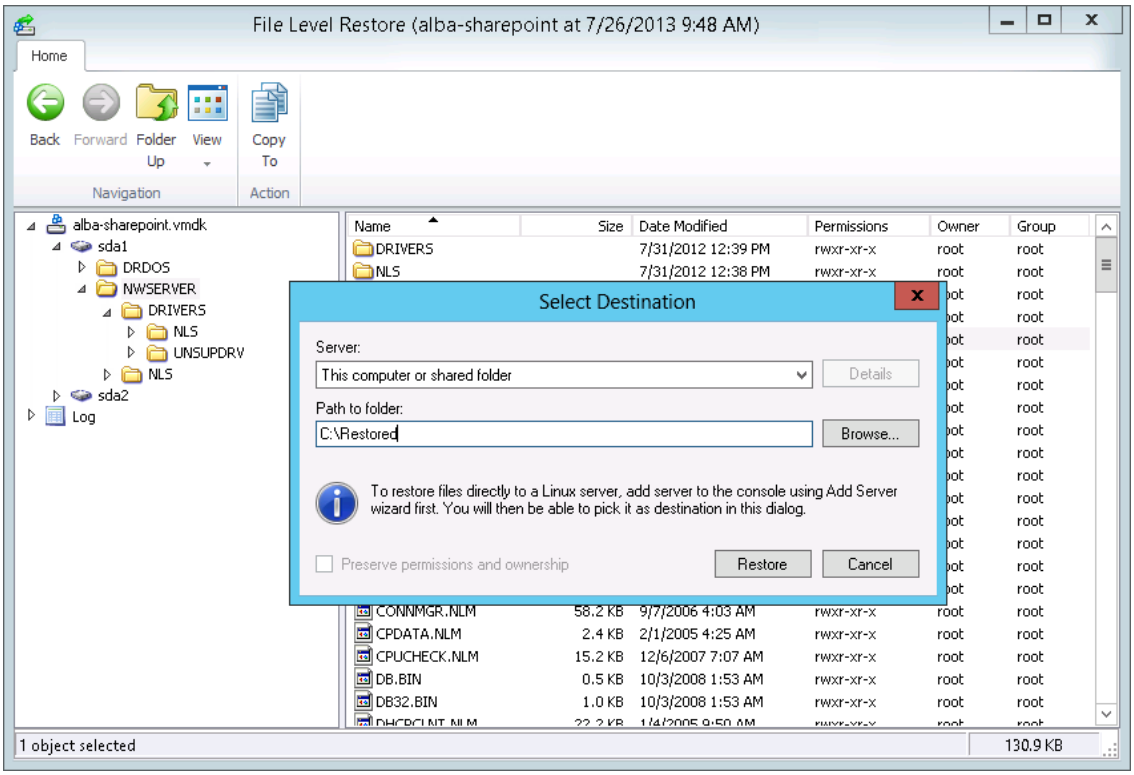

If you are recovering files to the original Linux host, you can preserve file permissions. Note in this case, the Linux host must be added to the list of servers managed by Veeam Backup & Replication in advance. For details, see the Adding a Linux Server section. Select the **Preserve permissions and ownership** check box to keep original permission settings for recovered files. Ownership settings are restored only if you have privileges to change the owner at the remote Linux host where files are restored.

If you have chosen to enable FTP server on the FLR appliance, the restored file system will also be available over FTP at *ftp://<FLR_appliance_IP_address>*. Other users in the same network can access the FLR appliance to restore the files they need.

Performing Replica Failover and Failback

With the virtual machine replica failover and failback possibilities, you can recover a corrupted virtual machine in case of software or hardware malfunction. The failover option can be used for any virtual machine replicas that were successfully created at least once.

To learn more about the purpose of each operation and associated background processes, see Replica Failover and Failback.

The following operations can be performed as part of the failover and failback workflow.

Performing Failover

During failover, Veeam Backup & Replication rolls back the replica to the required restore point and recovers a fully functional VM on the target host. Failing over to replicas is performed by means of the **Failover** wizard. This section will guide you through all steps of the wizard and provide explanation on offered options.

To fail over to a replica, follow the next steps:

Step 1. Launch the Failover Wizard

To launch the **Failover** wizard, do one of the following:

- On the **Home** tab, click **Restore** and select **VMware**. In the **Restore from replica** section, select **Failover to replica**.
- Open the **Backup & Replication** view and select the **Replicas** node. In the working area, expand the necessary replication job, select the VM and click **Failover Now** on the ribbon.
- Open the **Backup & Replication** view and select the **Replicas** node. In the working area, expand the necessary replication job, right-click the VM and select **Failover Now**.
- Open the **Backup & Replication** view and select **Ready** under the **Replicas** node. In the working area, select the necessary replica and click **Failover Now** on the ribbon or right-click the replica and select **Failover Now**.

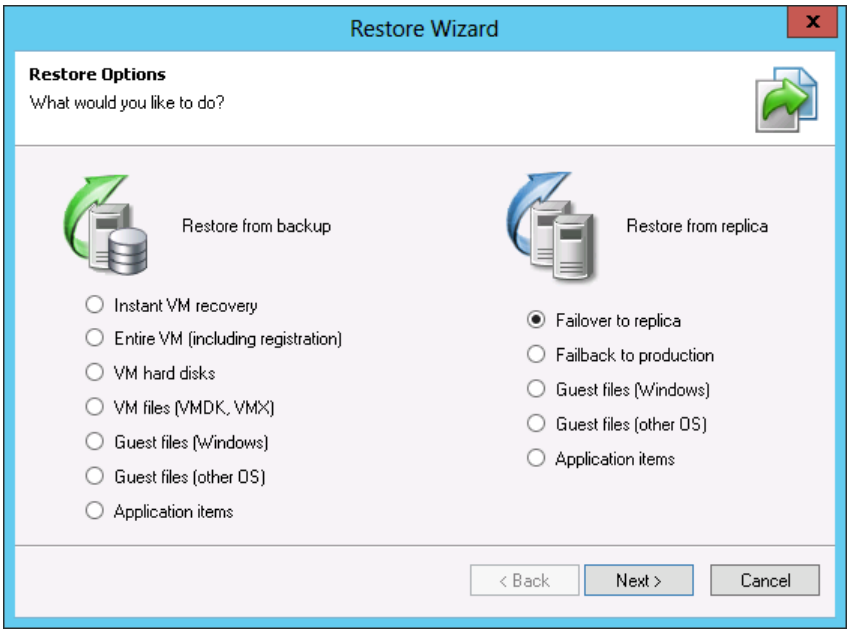

Step 2. Select Virtual Machines

At this step, you should select one or more VMs that you want to fail over. To add a VM or a VM container, click **Add VM** and select where to browse for the machines:

- **From Infrastructure** — browse the virtual environment and select VMs or VM containers to fail over. If you choose a VM container, Veeam Backup & Replication will expand it to a plain VM list.
  To facilitate selection, use the search field at the bottom of the **Add Objects** window: click the button to the left of the field and select the necessary type of object to search for (**Everything**, **Folder**, **Cluster**, **Host**, **Resource pool**, **VirtualApp** or **Virtual machine**), enter an object's name or a part of it and press **[ENTER]**.
  Make sure that VMs you select from the virtual environment have been successfully replicated at least once.

- **From Replica** — browse existing replicas and select VMs under replication jobs. To quickly find VMs, use the search field at the bottom of the **Select Objects** window: enter a VM name or a part of it and press **[ENTER]**.

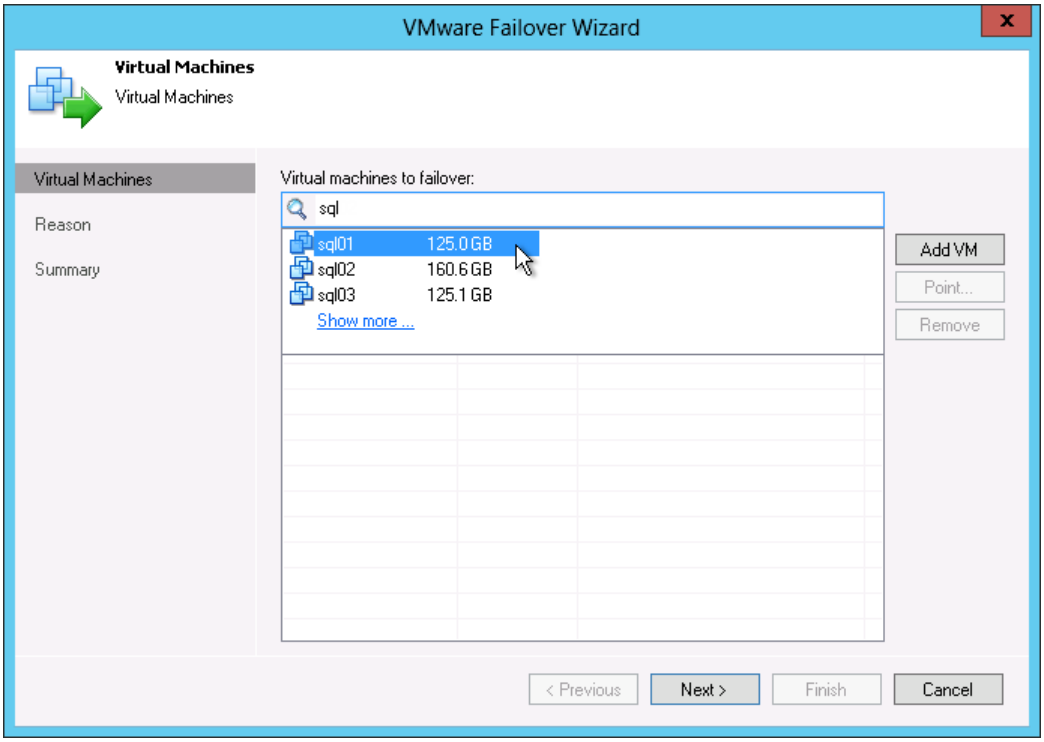

Alternatively, you can use the search field at the top of the window: enter a VM name or a part of it in the search field. Veeam Backup & Replication will search existing replicas for the specified VM and display matching results. To add a VM, double-click it in the list of search results. If a VM is not found, click the **Show more** link to browse existing replicas and choose the necessary VM.

To remove a VM from the list, select it and click **Remove** on the right.

Step 3. Select a Restore Point

At this step, you should select the necessary restore point to which you want to fail over.

By default, Veeam Backup & Replication uses the latest good restore point to recover a VM. However, if you want to fail over to an earlier replica state, select a VM in the **Virtual machines to failover** list and click **Point** on the right. In the **Restore Points** section, select a restore point that should be used to fail over the VM.

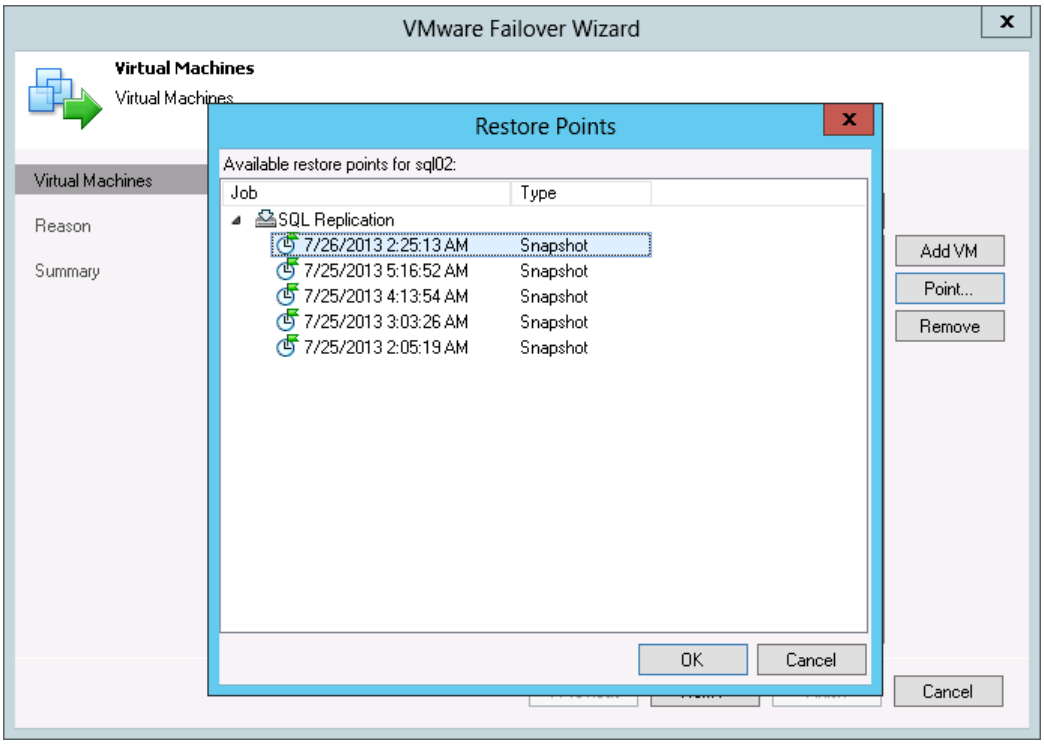

Step 4. Specify a Failover Reason

If necessary, enter the reason for performing failover of selected VMs. The information you provide will be saved in the session history so that you can reference it later.

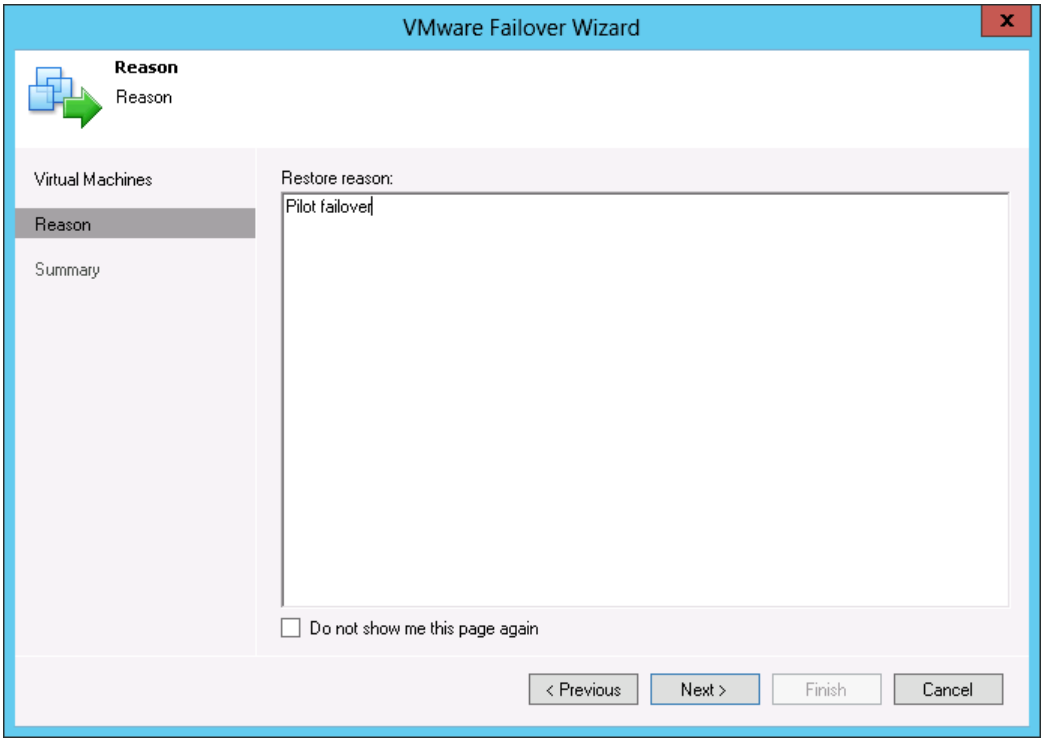

Step 5. Review Summary and Finish Working with the Wizard

Review the list of VMs to fail over and click **Finish** to start the failover procedure. Once the failover is complete, the VM replicas will be started on the target hosts.

Performing Permanent Failover

The Permanent failover option finalizes failover to a VM replica. As a result of the permanent failover, the VM replica on the target host ceases to exist as a replica and takes on the role of the original VM.

To perform permanent failover, do either of the following:

- Open the **Backup & Replication** view and select the **Replicas** node. In the working area, expand the necessary replication job, select the VM and click **Permanent Failover** on the ribbon.
- Open the **Backup & Replication** view and select the **Replicas** node. In the working area, expand the necessary replication job, right-click the VM and select **Permanent Failover**.
- Open the **Backup & Replication** view and select **Active** under the **Replicas** node. In the working area, select the necessary replica and click **Permanent Failover** on the ribbon or right-click the replica and select **Permanent Failover**.

In the displayed dialog box, click **Yes** to confirm the operation.

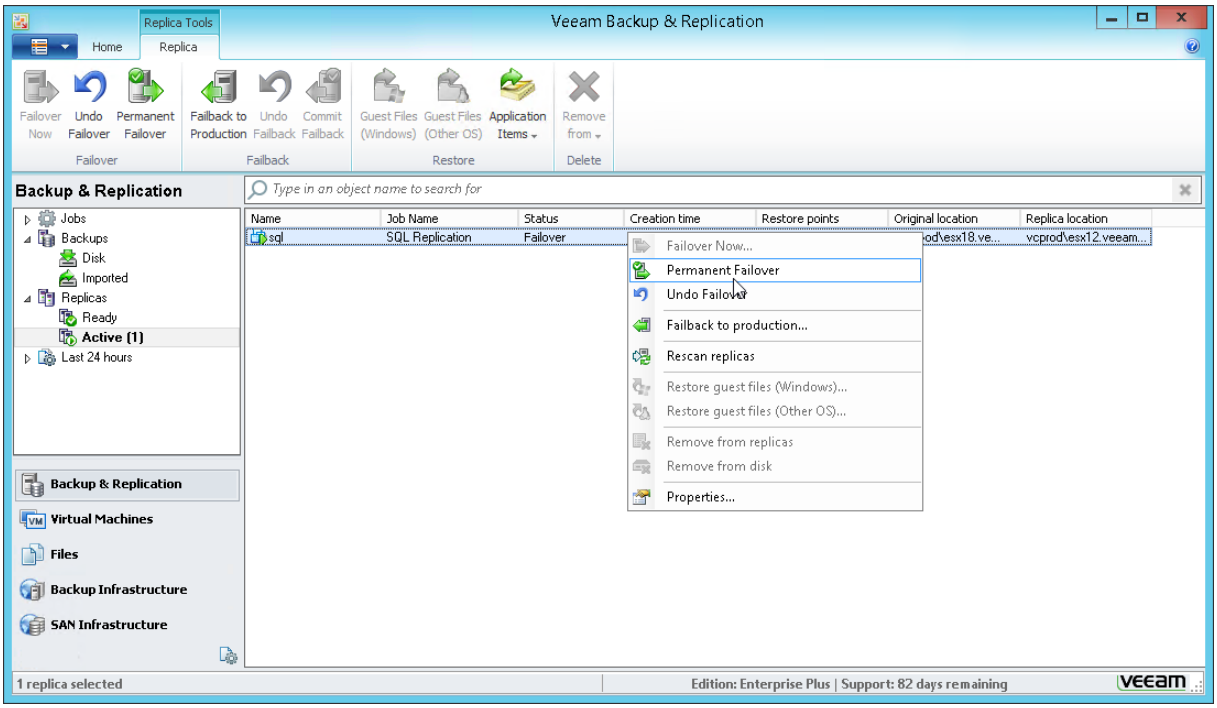

To protect the VM replica from corruption after performing a permanent failover, Veeam Backup & Replication removes the VM replica from the **Replicas** list. Additionally, Veeam Backup & Replication reconfigures the replication job and adds the original VM to the list of exclusions. When the replication job that processes the original VM starts, the VM will be skipped from processing, and no data will be written to the working VM replica.

Undoing Failover

The **Undo failover** option allows powering off running VM replicas on target hosts and rolling back to their initial state.

To undo failover, do either of the following:

- On the **Home** tab, click **Restore**. In the **Restore** from replica section, select **Undo previously performed failover**.
- Open the **Backup & Replication** view and select the **Replicas** node. In the working area, expand the necessary replication job, select the VM and click **Undo Failover** on the ribbon.
- Open the **Backup & Replication** view and select the **Replicas** node. In the working area, expand the necessary replication job, right-click the VM and select **Undo Failover**.
- Open the **Backup & Replication** view and select **Active** under the **Replicas** node. In the working area, select the necessary replica and click **Undo Failover** on the ribbon or right-click the replica and select **Undo Failover**.

In the displayed dialog box, click **Yes** to confirm the operation.

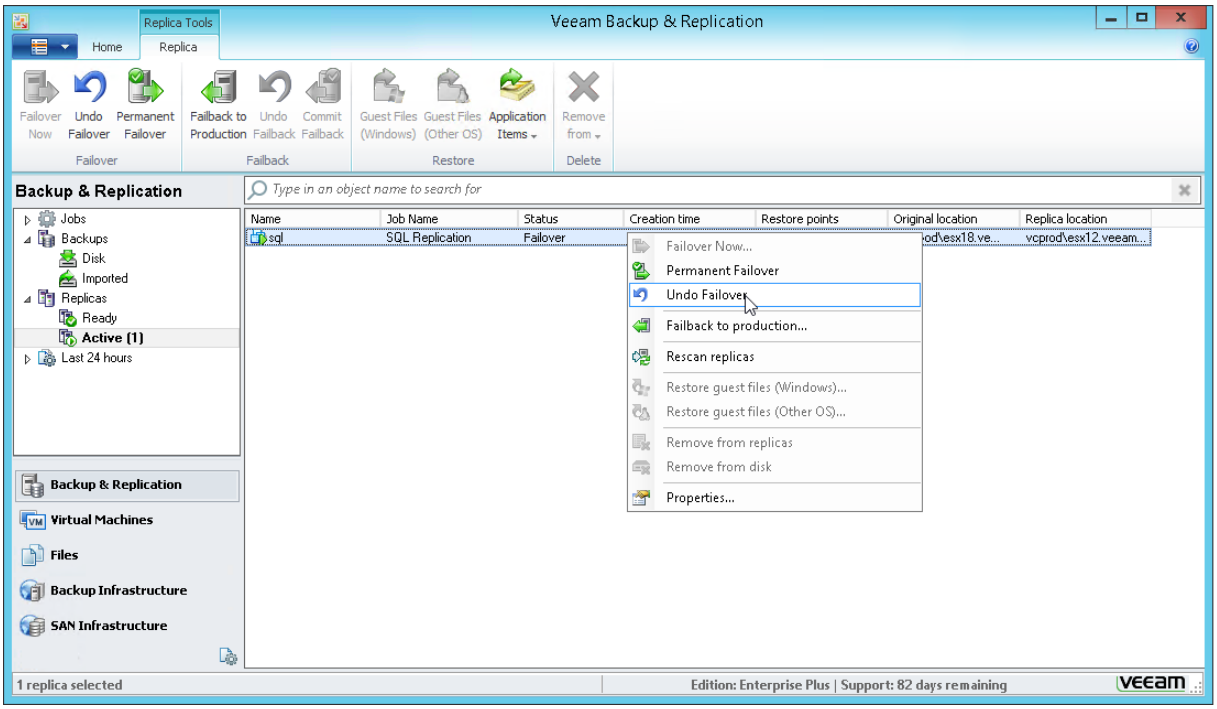

Performing Failback

The **Failback** option allows you to switch from a VM replica back to the original VM or restore a VM from a replica in a new location. Failback is performed by means of the **Failback** wizard. This section will guide you through all steps of the wizard and provide explanation on offered options.

Important! You can perform failback for a VM replica in the Failover state. The VM replica is put to the Failover state when you fail over to it from the original VM. To see all VMs in the Failover state, open the Backup & Replication view and select the **Active** node under **Replicas** in the inventory pane.

To perform failback, follow the next steps:

Step 1. Launch the Failback Wizard

To launch the **Failback** wizard, do one of the following:

- On the **Home** tab, click **Restore** and select **VMware**. In the **Restore from replica** section, select **Failback to production**.
- Open the **Backup & Replication** view and select the **Replicas** node. In the working area, expand the necessary replication job, select the VM and click **Failback to production** on the ribbon.
- Open the **Backup & Replication** view and click the **Replicas** node. In the working area, expand the necessary replication job, right-click the VM and select **Failback to production**.
- Open the **Backup & Replication** view and select **Active** under the **Replicas** node. In the working area, select the necessary replica and click **Failback to production** on the ribbon or right-click the replica and select **Failback to production**.

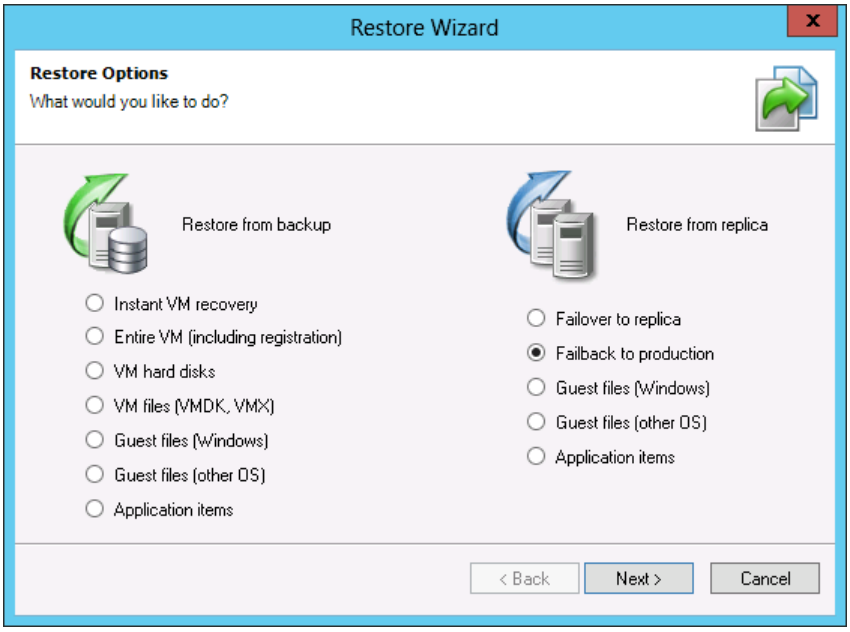

Step 2. Select VM Replicas to Fail Back

At this step, you should select one or more VM replicas from which you want to fail back. Click **Populate** to display all existing replicas in the *Failover* state. Leave check boxes selected only for those replicas from which you want to fail back.

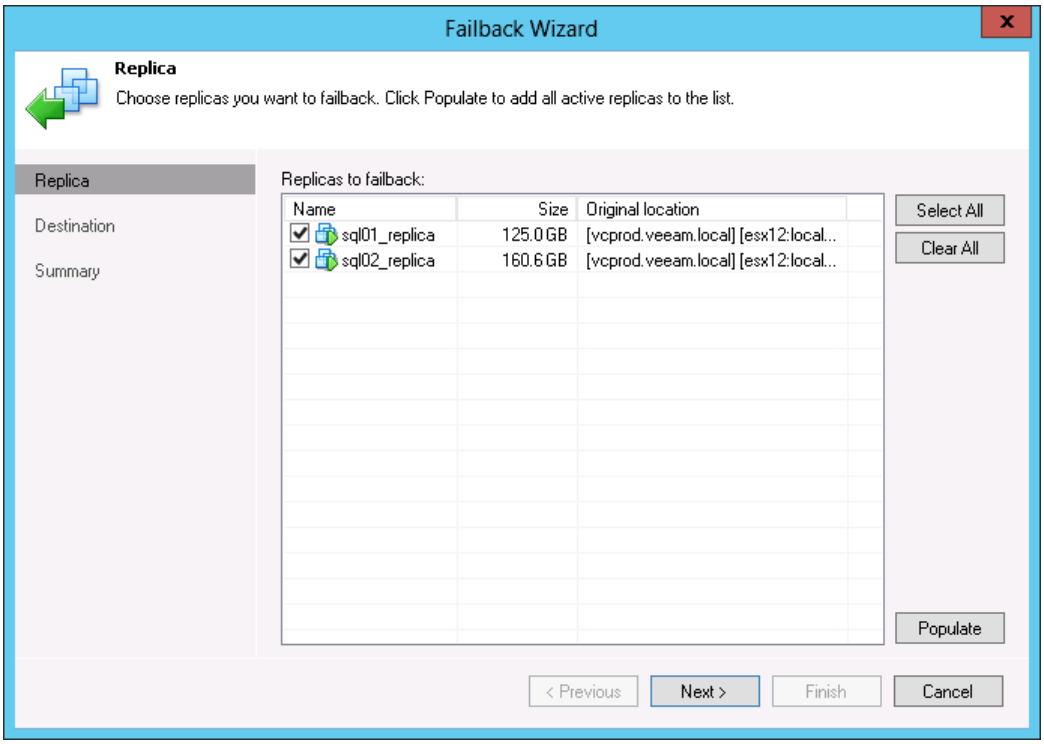

Step 3. Select Failback Destination

At this step of the wizard, you should select failback destination and backup proxies that will be used to perform failback.

Veeam Backup & Replication supports three possible failback destination variants. Note that the Failback wizard displays a different set of steps for every failback variant.

- Select **Failback to the original VM** if you want to fail back to the original VM residing on the source host. In this case, Veeam Backup & Replication will restore the original VM to the current state of its replica.
  If this option is selected, you will pass to the Summary step of the wizard.
- Select **Failback to the original VM restored in a different location** if you have recovered the original VM from a backup in a new location, and you want to switch to it from the replica. In this case, Veeam Backup & Replication will restore the recovered VM to the current state of the replica.
  If this option is selected, you will pass to the Target VM step of the wizard.
- Select **Failback to the specified location** if you want to restore the original VM from a replica – in a new location and/or with different settings (such as VM location, network settings, virtual disk and configuration files path and so on).
  If this option is selected, you will need to complete all further steps of the wizard.

Note that if you fail back to the original VM or to the original VM restored in a new location, only differences between the existing virtual disks and their state will be transferred over to the original VM. Veeam Backup & Replication will not transfer replica configuration changes, such as a different IP address or network settings (if replica re-IP and network mapping were applied), new hardware or virtual disks added while the replica was in the Failover state.

If you choose to perform advanced failback, the entire VM replica, including its configuration and virtual disks content, will be restored in the selected location.

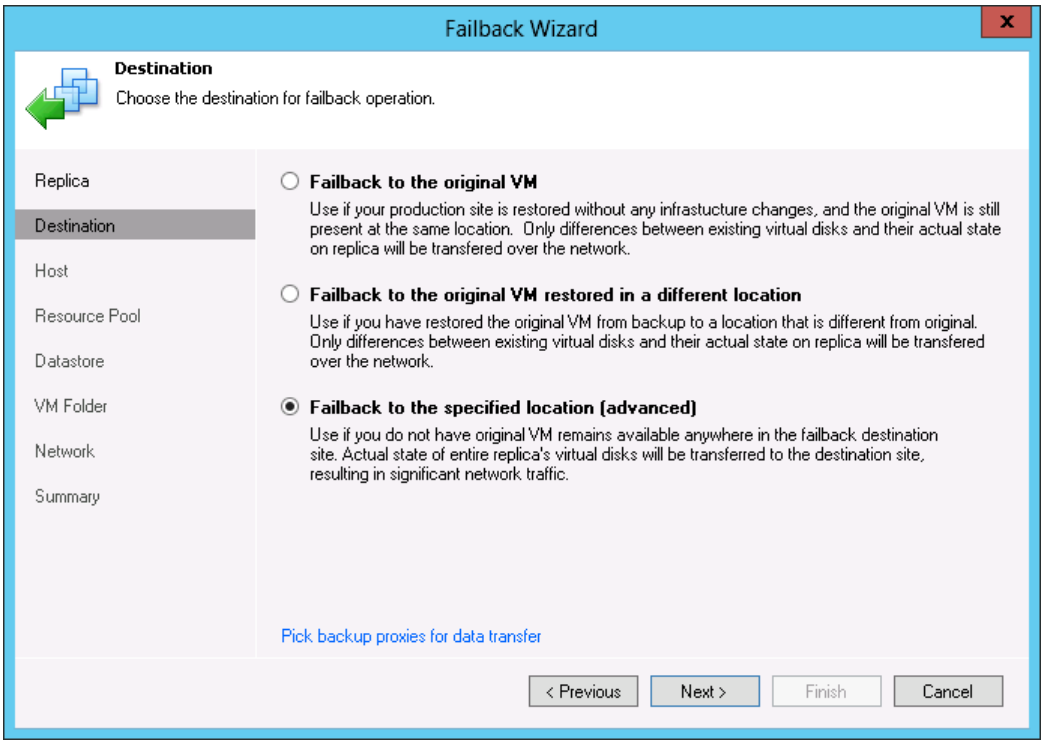

Click the **Pick backup proxies for data transfer** link to select backup proxies to perform data transfer during failback. In the offsite replication scenario, you should select one backup proxy in the production site and one proxy in the DR site. In the onsite replication scenario, you can use the same server as the source and target proxy.

In the **Choose backup Proxy** section, click **Choose** to assign a backup proxy. You can use automatic proxy selection or assign proxies explicitly.

- If you choose **Automatic selection**, Veeam Backup & Replication will detect backup proxies that are connected to the source datastore and will automatically assign optimal proxy resources for processing VM data.
  VMs selected for failback are processed one by one. Before processing a new VM in the VM list, Veeam Backup & Replication checks available backup proxies. If more than one proxy is available, Veeam Backup & Replication analyzes transport modes that the proxies can use, the current workload on the proxies to select the most appropriate resource for VM processing.
- If you choose **Use the backup proxy servers specified below**, you can explicitly select proxies that will be used to perform data transfer. It is recommended to select at least two proxies to ensure that failover will be performed should one of job proxies fail or lose its connectivity to the source or target datastore.

Step 4. Select a Failback Destination Host

This step of the wizard is only available if you have chosen to perform advanced failback. To specify a destination host, select one or more VMs in the list and click **Host**. From the virtual environment, choose a host or cluster where the selected VMs should be registered.

To facilitate selection, use the search field at the bottom of the window: click the button on the left of the field to select the necessary type of object that should be searched for (**Cluster** or **Host**), enter an object's name or a part of it and click the **Start search** button on the right or press **[ENTER]**.

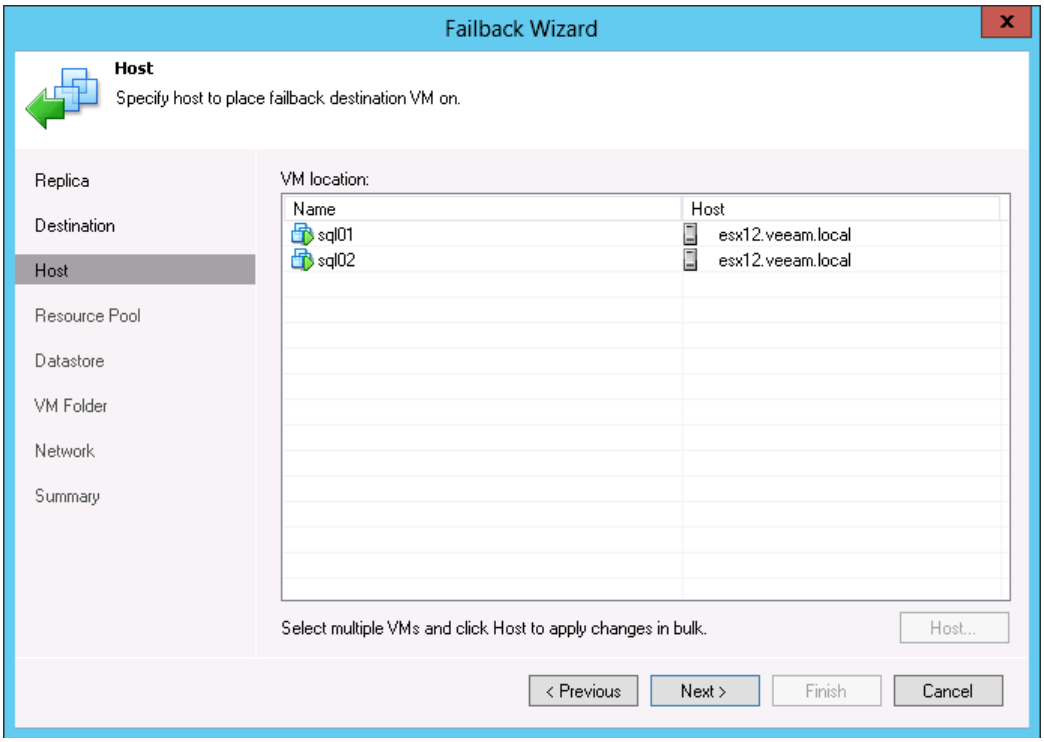

Step 5. Select a Failback Destination Resource Pool

This step of the wizard is only available if you have chosen to perform advanced failback. To specify a destination resource pool, select one or more VMs in the list and click **Pool**. From the virtual environment, choose a resource pool to which the selected VMs will belong.

To facilitate selection, use the search field at the bottom of the window: enter a resource pool name or a part of it and click the **Start search** button on the right or press **[ENTER]**. If required, you can also select a vApp to which the restored VM will be included.

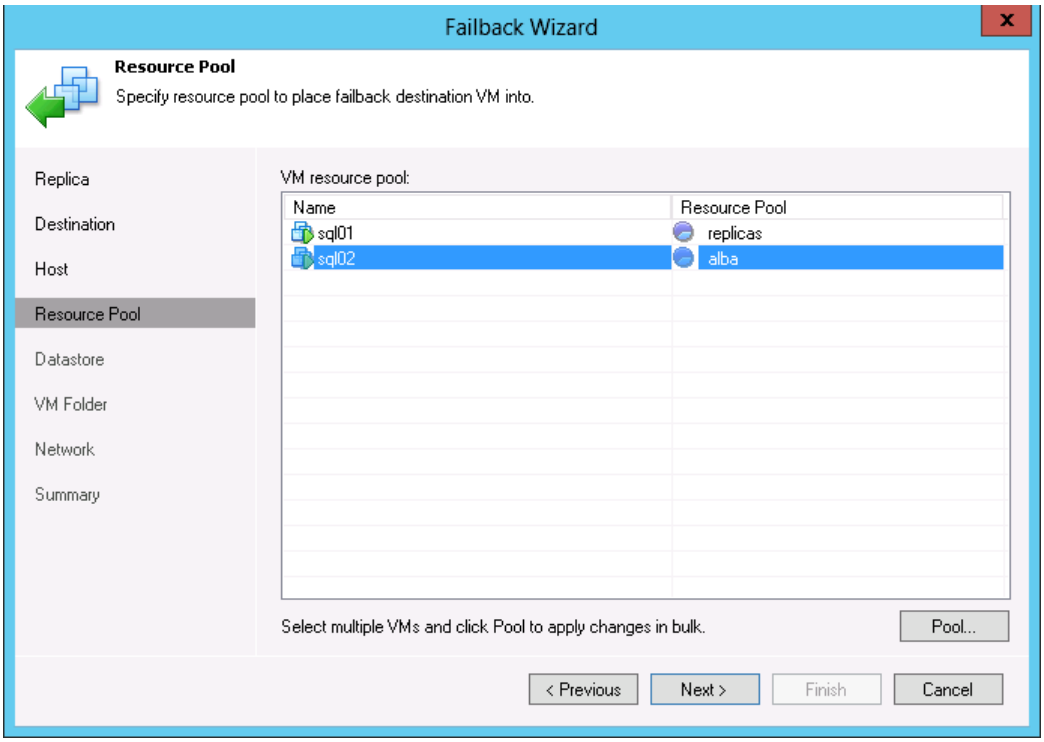

Step 6. Select a Failback Destination Datastore

This step of the wizard is only available if you have chosen to perform advanced failback. When restoring a VM from a replica, you can place an entire VM to a particular datastore or choose to store configuration files and disk files of a restored VM in different locations.

To specify a destination datastore, select one or more VMs in the list and click **Datastore**. If configuration and disk files of a VM should be placed to different datastores, expand the VM in the list, select the necessary file type and click **Datastore**. From the virtual environment, choose a datastore to which the selected objects will be stored. To facilitate selection, use the search field at the bottom of the window: enter a datastore name or a part of it and click the **Start search** button on the right or press **[ENTER]**.

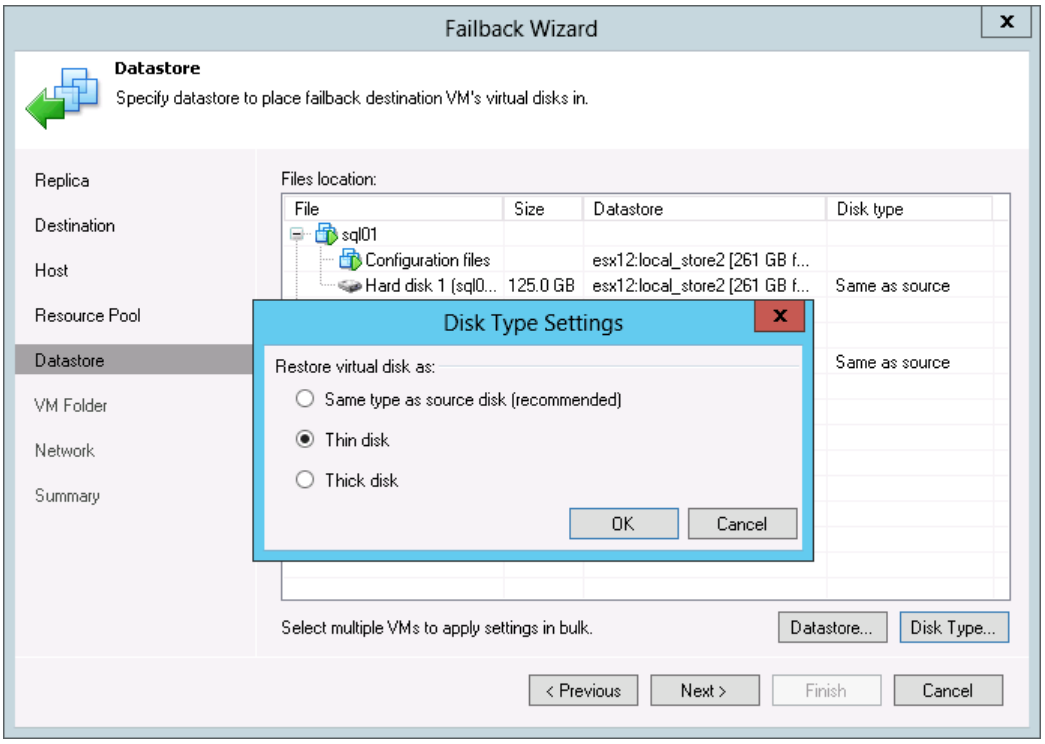

By default, Veeam Backup & Replication preserves the format of restored VM disks, so that if disks of the VM replica were provisioned as thick, Veeam Backup & Replication will restore the VM with thick disks. However, if necessary, you can change the disk format of a restored VM. To do so, expand a VM in the list, select the necessary disk and click **Disk Type**. In the **Disk Type Settings** section, choose the format that will be used to restore virtual disks of the VM — same as the source disk, thin or thick. Please note that changing disk format is available only for VMs using virtual hardware version 7 or later.

Step 7. Select a Failback Destination Folder

This step of the wizard is only available if you have chosen to perform advanced failback. To specify a destination VM folder, select one or more VMs in the list and click **Folder**. From the virtual environment, choose a folder to which the selected VMs will belong. To facilitate selection, use the search field at the bottom of the window: enter a folder name or a part of it and click the **Start search** button on the right or press **[ENTER]**.

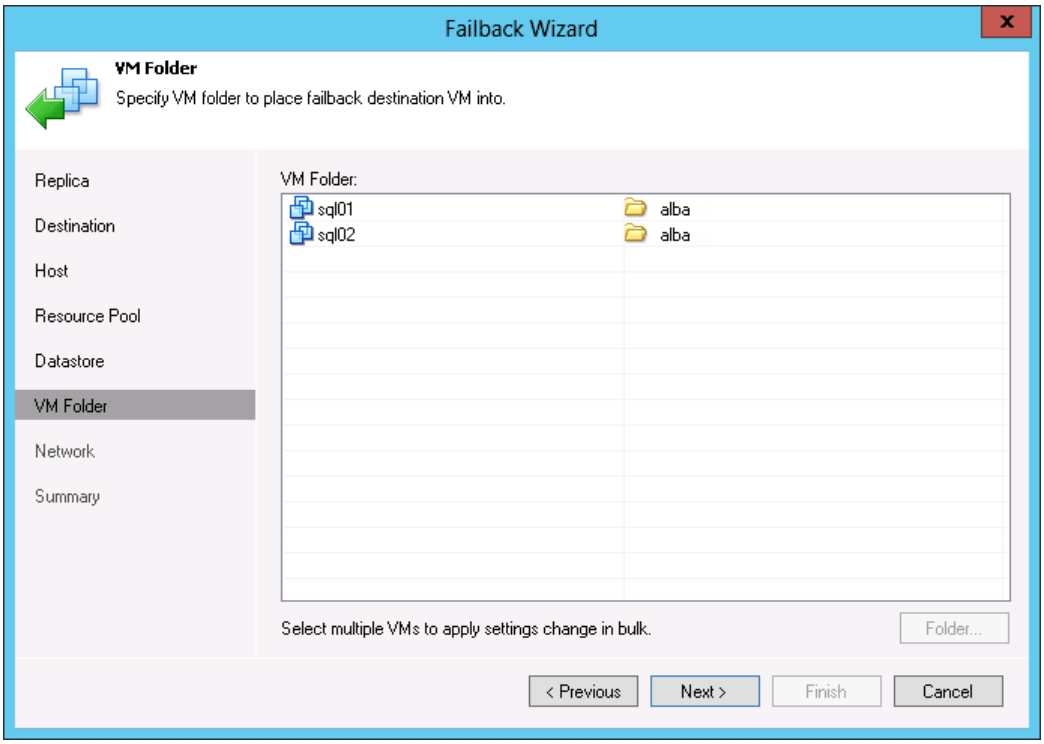

Step 8. Select a Destination Network

This step of the wizard is only available if you have chosen to perform advanced failback. If you plan to fail back to VMs to a new location (for example, another site with a different set of networks), you can map DR site networks to production site networks. Veeam Backup & Replication will use the network mapping table to update configuration files of VMs on the fly, during the restore process.

To change networks to which restored VMs will be connected, select one or more VMs in the list and click **Networks**. If a VM is connected to multiple networks, expand the VM, select the network to map and click **Network**. The **Select Network** section displays all networks to which the destination host or cluster is connected. From the list of available networks, choose a network to which the original VMs should have access upon failback. To facilitate selection, use the search field at the bottom of the window: enter a network name or a part of it and click the **Start search** button on the right or press **[ENTER]**.

To prevent the original VM from accessing any network upon failback, select the VM or its network connections in the list and click **Disconnected**.

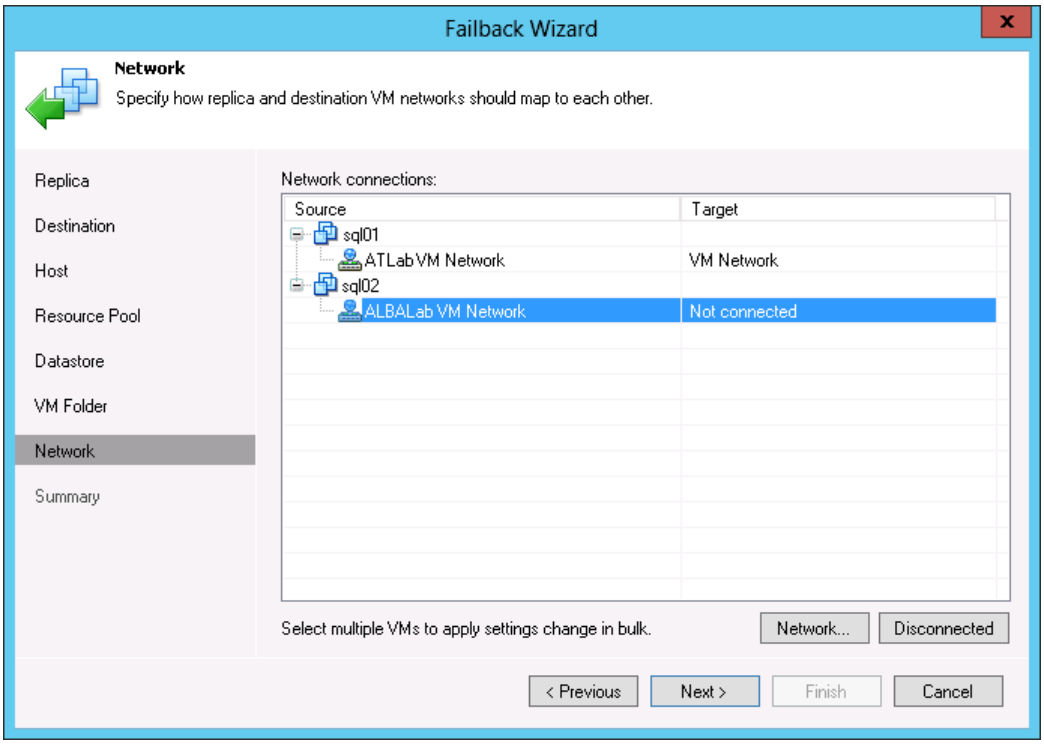

Step 9. Map the Replica to the Restored VM

This step of the wizard is only available if you have chosen to fail back to the original VM restored in a different location. At this step, you should define how VM replicas map to VMs restored from backup.

To create a mapping association, select a replica in the list and click **Edit**. Select the restored VM from the VI. To facilitate selection, use the search field at the bottom of the **Add Objects** window: click the button to the left of the field and select the necessary type of object to search for (*Everything*, *Folder*, *Cluster*, *Host*, *Resource pool*, *VirtualApp* or *Virtual machine*), enter an object's name or a part of it and click the **Start search** button on the right or press **[ENTER]**.

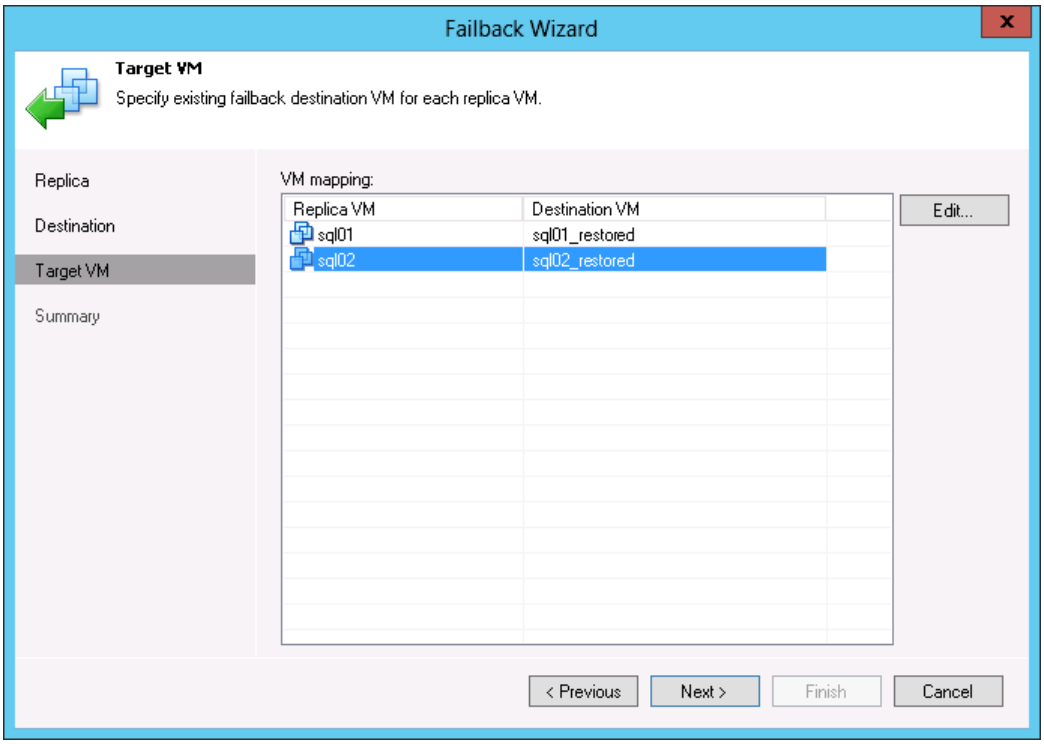

Step 10. Review Summary and Finish Working with the Wizard

If you want to start the original VM after the work with the Failback wizard is complete, select the **Power on VM after restoring** check box.

Check specified settings for failback and click **Finish**. Veeam Backup & Replication will restore the original VMs to the state of corresponding VM replicas.

Committing Failback

The **Commit failback** option finalizes failback from the VM replica to the original VM.

To commit failback, do either of the following:

- On the **Home** tab, click **Restore**. In the **Restore from replica** section, select **Commit failback**.
- Open the **Backup & Replication** view and select the **Replicas** node. In the working area, expand the necessary replication job, select the VM and click **Commit Failback** on the ribbon.
- Open the **Backup & Replication** view and select the **Replicas** node. In the working area, expand the necessary replication job, right-click the VM and select **Commit Failback**.
- Open the **Backup & Replication** view and select **Active** under the **Replicas** node. In the working area, select the necessary replica and click **Commit Failback** on the ribbon or right-click the replica and select **Commit Failback**.

In the displayed dialog box, click **Yes** to confirm the operation.

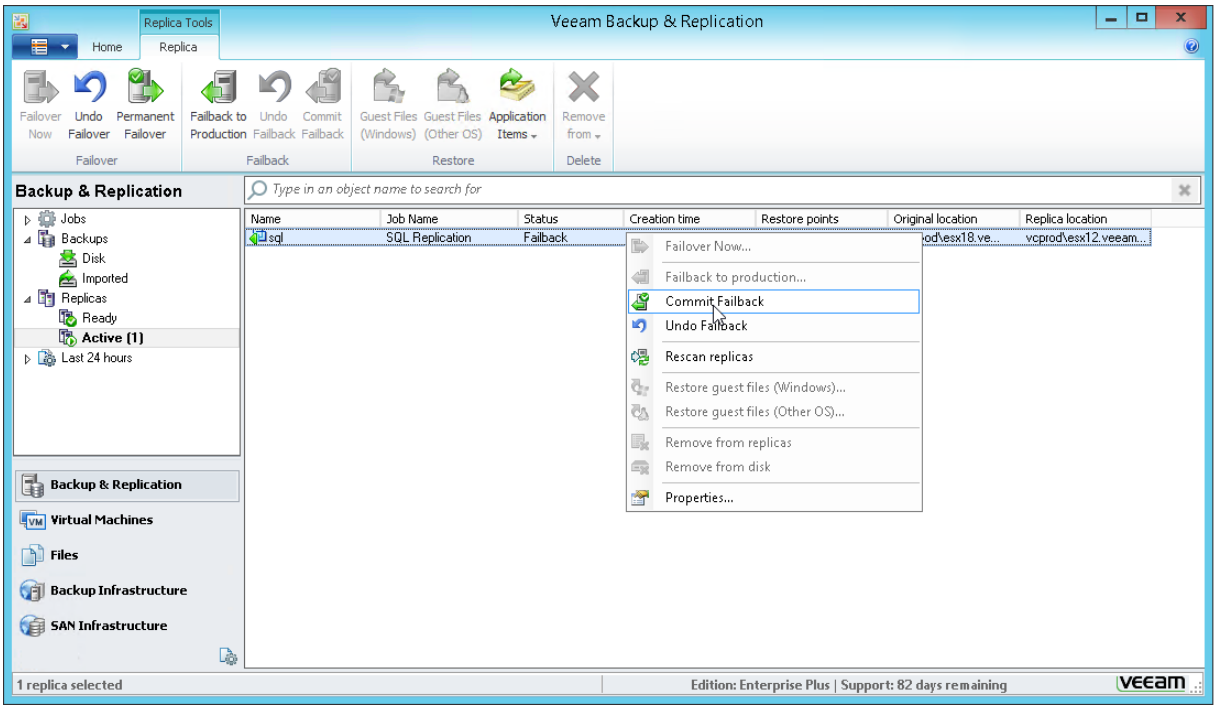

Depending on the location to which the VM is failed back, Veeam Backup & Replication performs the following finalizing operations after failback is committed:

- If the VM replica is failed back to a new location, Veeam Backup & Replication additionally reconfigures the replication job and adds the former original VM to the list of exclusions. The VM restored in the new location takes the role of the original VM, and is included into the replication job instead of the excluded VM. When the replication job starts, Veeam Backup & Replication will skip the former original VM from processing, and will replicate the newly restored VM instead.
- If the VM replica is failed back the original location, the replication job is not reconfigured. When the replication job starts, Veeam Backup & Replication will process the original VM in the normal mode.

Undoing Failback

The **Undo failback** option allows you to switch from the original VM back to the VM replica and roll back the replica to the failover state.

To undo failback, do either of the following:

- Open the **Backup & Replication** view and select the **Replicas** node. In the working area, expand the necessary replication job, select the VM and click **Undo Failback** on the ribbon.
- Open the **Backup & Replication** view and select the **Replicas** node. In the working area, expand the necessary replication job, right-click the VM and select **Undo Failback**.
- Open the **Backup & Replication** view and select **Active** under the **Replicas** node. In the working area, select the necessary replica and click **Undo Failback** on the ribbon or right-click the replica and select **Undo Failback**.

In the displayed dialog box, click **Yes** to confirm the operation.

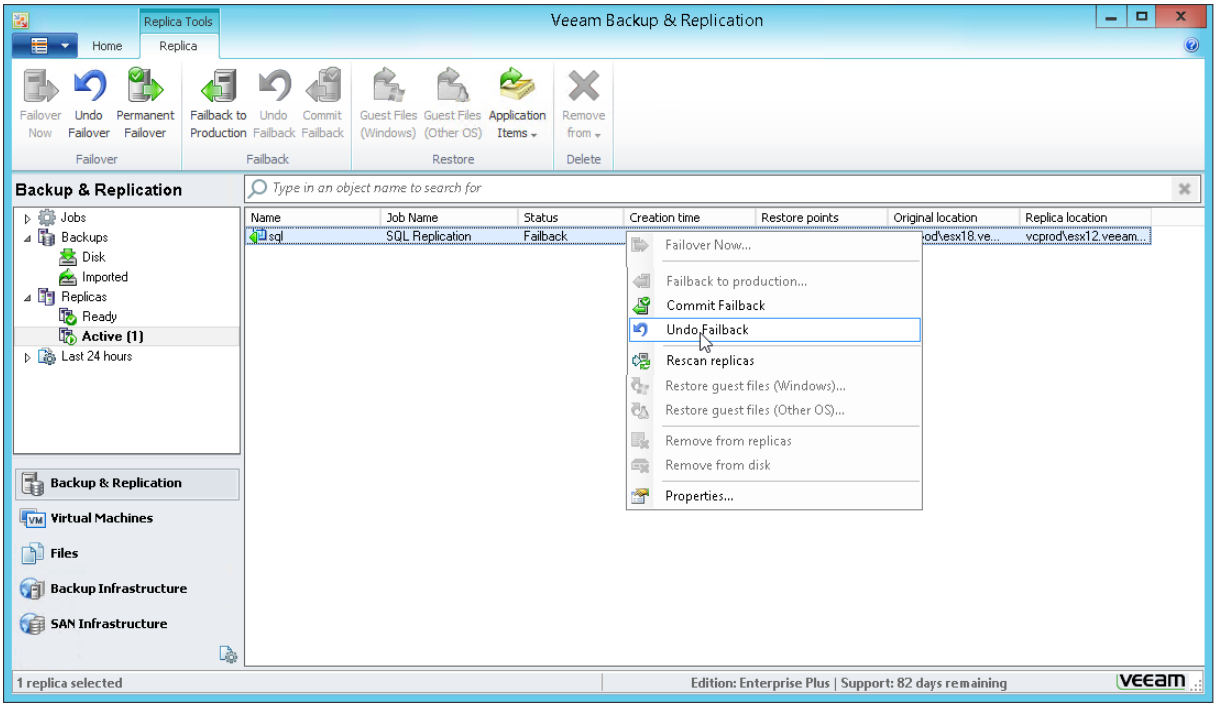

Managing Backups and Replicas

Veeam Backup & Replication offers the following management options for your backups and replicas: removing from backups/replicas, deleting from disks and viewing properties. All options are available from the shortcut menu.

- The **Remove from Backups or Replicas** option is used when you want to remove records about backup and replica files from the Veeam Backup configuration database. Please note that all backup files (.vbk, .vib, .vrb, .vbm) will stay safe on the destination backup storage, so you can easily import these files later to the Veeam Backup & Replication console for restore operations if needed. As for replicas, all references will be removed from the Veeam Backup & Replication console; however, all your replicated VMs will still reside on your target hosts, so you can start them manually after the **Remove from replicas** option is performed.
- In addition to removing records about backup and replica files from the Veeam Backup configuration database, the **Delete from disk** option also removes actual backups and replicas from the destination storage. Note that you should avoid deleting backup files manually from your destination storage, otherwise all subsequent job sessions will be failing. You can use this option for all VMs in the backup or replication job or for each VM separately.
- The **Properties** option for backups is used to view summary information on backups you made. It contains information on compression and de-duplication ratios, available restore points for a particular backup, as well as date, data size and backup size.

Importing Backups

Importing backups can be useful if you need to restore backups from tape or from .vbk files of other Veeam Backup & Replication versions or instances, if you happened to delete the server with which the backup was associated from the management tree, or in case the application has been uninstalled. You can also use the import option to work with VeeamZIP files: if you have created VeeamZIP files, you can import them to the Veeam Backup & Replication console and use them for data restore as usual backup. For details, see Creating VeeamZIP Files.

To import backups to Veeam Backup & Replication:

1. On the **Home** tab, click **Import Backup**.
2. From the **Computer** list, select the host on which the backup you want to import is stored.
3. Click **Browse** and select the necessary .vbm or .vbk file. Note that the import process is notably faster if you select the .vbm file. Therefore, it is recommended to use the .vbk files for import only if no corresponding .vbm file is available.
4. By default, index data of the guest OS file system is not imported with the backup file to speed up the import process. However, if it is necessary, select the **Import guest file system index** check box.
5. Click **OK** to import the selected backup. The imported backup will become available in the **Backup & Replication** view, under the **Backups** > **Imported** node in the inventory pane. Backups are imported using the original name of the backup job with the `_imported` suffix appended.

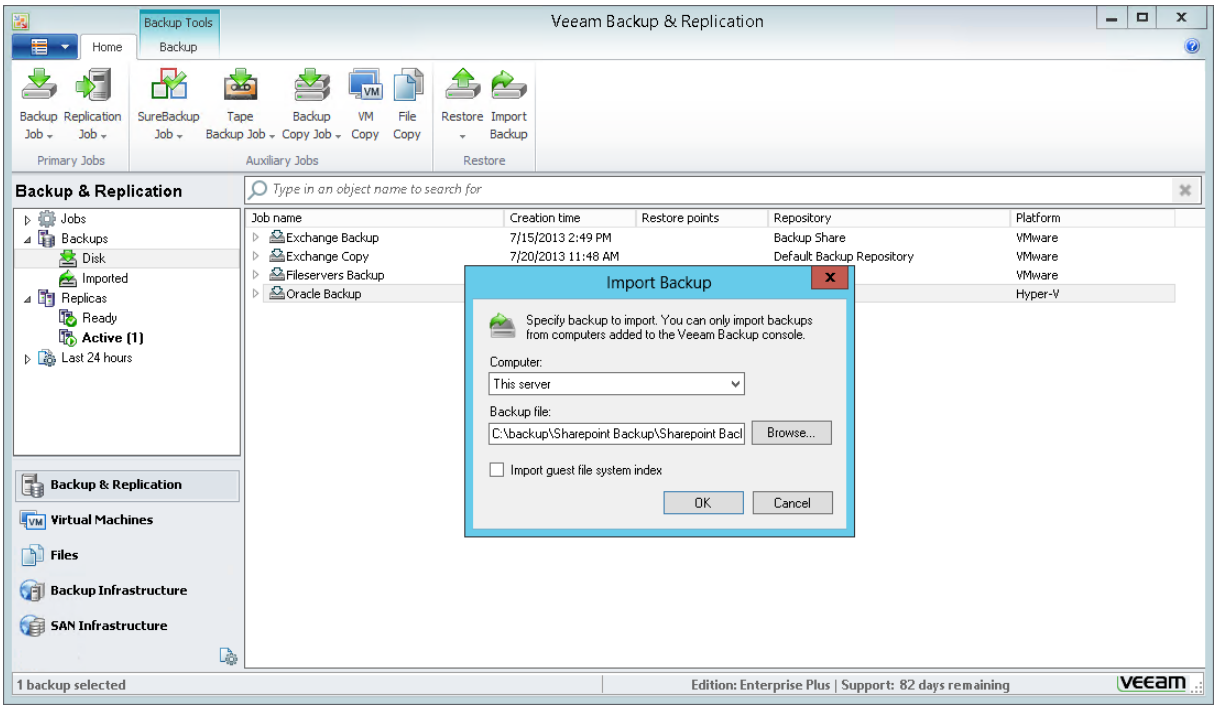

Important! To be able to perform any restore operation from previous points in time (rollbacks) for your backed up VM, before importing a full backup file to the Veeam Backup & Replication console, make sure that you have all required increments (either forward or reverse) in the same folder.

Working with vCloud Director VMs

This section describes administrative tasks you can perform for VMs managed by vCloud Director.

Viewing vCloud Director VMs

After you have added the vCloud Director server, you can view the vCloud Director hierarchy directly in the Veeam Backup & Replication console and work with VMs managed with vCloud Director. To open the vCloud Director hierarchy:

1. In Veeam Backup & Replication, open the **Virtual Machines** view.
2. Click the **View** tab on the ribbon.
3. On the **View** tab, click **vCloud View**.

The hierarchy of the added vCloud Director server will become available in the inventory pane, VMs managed by vCloud Director will be displayed in the working area. You can work with these VMs just as if you worked with VMs managed by vCenter Servers or registered on ESX(i) hosts added to Veeam Backup & Replication.

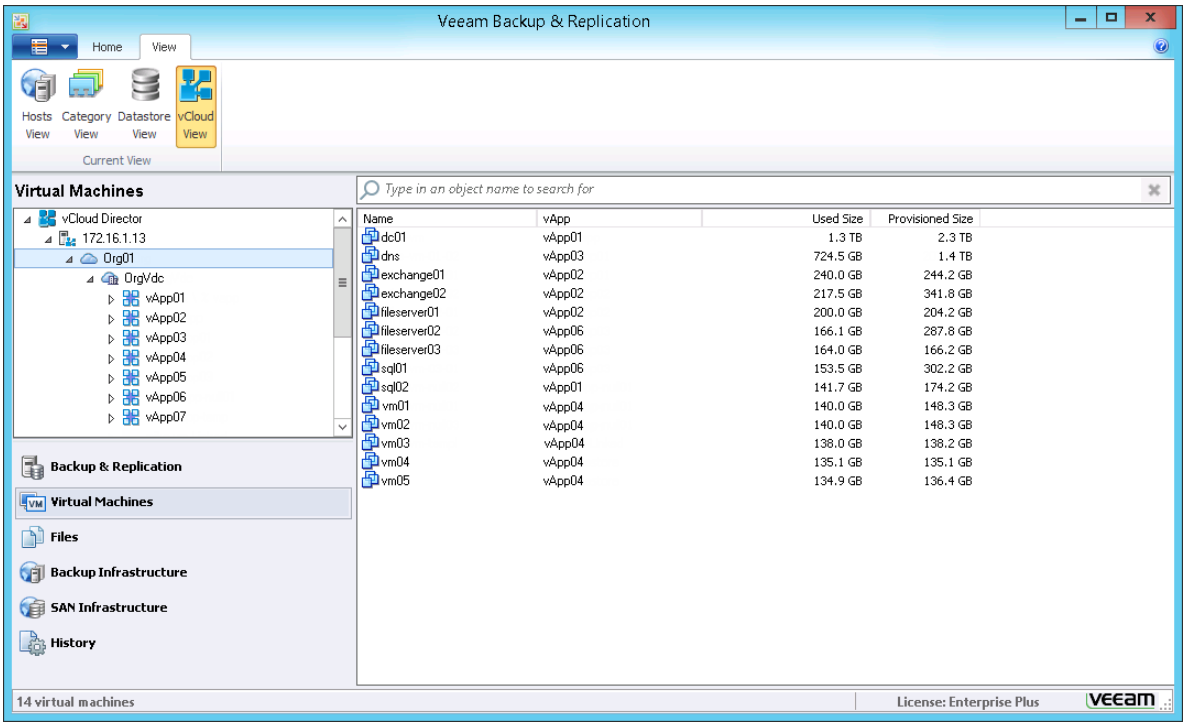

Performing Backup of vCloud Director VMs

From the user's side, the vCD backup is practically the same as a regular VM backup. The vCD backup job aggregates main settings for the backup task and defines when, what, how and where to back up.

You can perform the vCD backup job for single VMs and for VM containers, which, in terms of vCloud Director, are the following:

- vApp
- Organization vDC
- Organization
- vCloud Director instance

As well as a regular backup job, the vCD backup job can be scheduled or run manually. To create a vCD backup job, do one of the following:

- On the **Home** tab, click **Backup Job** and select **vCloud**.
- Open the **Backup & Replication** view, right-click the **Jobs** node and select **Backup** > **vCloud**.
- Open the **Virtual Machines** view, click the **View** tab and click **vCloud View** on the ribbon. Expand the **vCloud Director** hierarchy in the inventory pane, select one or several VMs in the working area, click **Add to Backup** on the ribbon and select **New job**. Alternatively, you can right-click one or several VMs in the working area and select **Add to backup job** > **New job**. In this case, the selected VMs will be automatically included into the new vCD backup job. You can add other VMs to the job when passing through the wizard steps.

You can quickly include VMs to already existing vCD backup jobs. To do that, in the **Virtual Machines** view, right-click necessary VMs in the working area and select **Add to backup job** > *name of a created job*.

The **New vCD Backup Job** wizard offers the same options as a **New Backup Job** wizard. To learn what settings you can specify for the backup job, see the Creating Backup Jobs section.

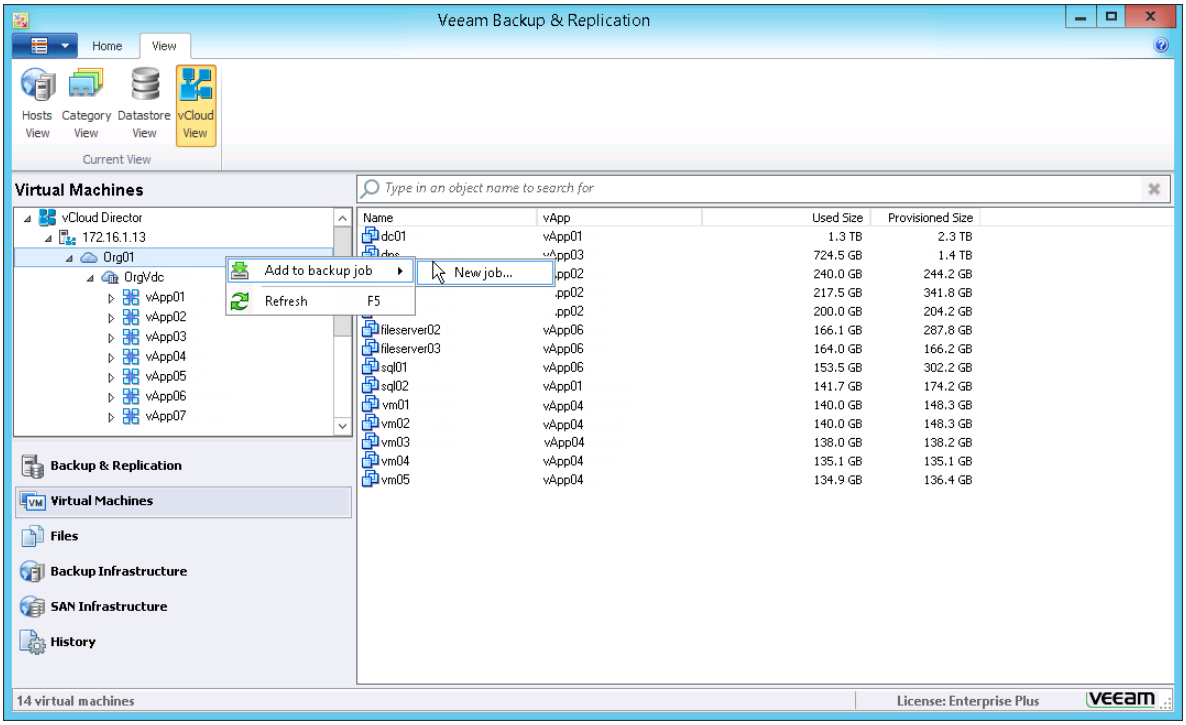

Performing Restore from vCD Backups

Veeam Backup & Replication offers a number of options for restoring VM data from vCD backups. You can:

- Restore a VM with Instant VM Recovery
- Restore an entire vApp back to the vCloud Director hierarchy
- Restore a full VM back to the vCloud Director hierarchy
- Restore an entire VM to the VMware vSphere infrastructure
- Restore VM hard disks
- Restore VM files (VMX, VMDK and so on)
- Restore VM guest OS files

Performing Instant VM Recovery for VMs

Veeam Backup & Replication provides two options for Instant VM Recovery of vCD VMs:

- You can instantly recover a VM to a vApp in the vCloud Director hierarchy.
- You can instantly recover a VM to the virtual infrastructure. In this case, the VM will be restored at the level of the underlying vCenter Server and the Instant VM Recovery process will be the same as for regular VMware VMs.

When you instantly restore a VM to vCloud Director, Veeam Backup & Replication uses the vPower NFS datastore, just as with other VMware VMs. However, to import the VM to the vApp, Veeam Backup & Replication needs to associate the vPower NFS datastore with some storage profile. To do that, Veeam Backup & Replication creates in the underlying vCenter Server an auxiliary storage profile — *Veeam-InstantVMRecovery* and displays it in vCloud Director.

The created storage profile is added to the Provider vDC and Organization vCD hosting the vApp to which the VM is restored. When the vPower NFS datastore is mounted to the ESX(i) host, the vPower NFS datastore is associated with the *Veeam-InstantVMRecovery* storage profile. After that, the VM is instantly restored in the regular course and imported to the selected vApp.

When an Instant VM Recovery session is finished, the storage profile is not deleted from the Provider vDC; it remains in vCenter Server. This helps speed up all subsequent Instant VM Recovery operations. However, the storage profile is deleted from the Organization vDC as Organization vDC settings can be accessed by Organization administrators.

Restoring vCD VMs with Instant VM Recovery to vCloud Director

To restore VMs to vCloud Director hierarchy using Instant VM Recovery, follow the next steps:

Step 1. Launch the vCloud Instant VM Recovery Wizard

To launch the **vCloud Instant VM Recovery** wizard, do one of the following:

- Open the **Backup & Replication** view and select the **Backups** node. In the working area, expand the necessary backup job, select the VM you want to restore and click **Instant VM Recovery** > **Into vCloud vApp** on the ribbon.
- Open the **Backup & Replication** view and select the **Backups** node. In the working area, expand the necessary backup job, right-click the VM you want to restore and select **Instant VM recovery** > **Into vCloud vApp**.
- Open the **Virtual Machines** view. On the **View** tab, click **vCloud View**. In the inventory pane, expand the **vCloud Director** hierarchy. In the working area, right-click the VM you want to restore and select **Restore > Instant VM recovery** > **Into vCloud vApp**.

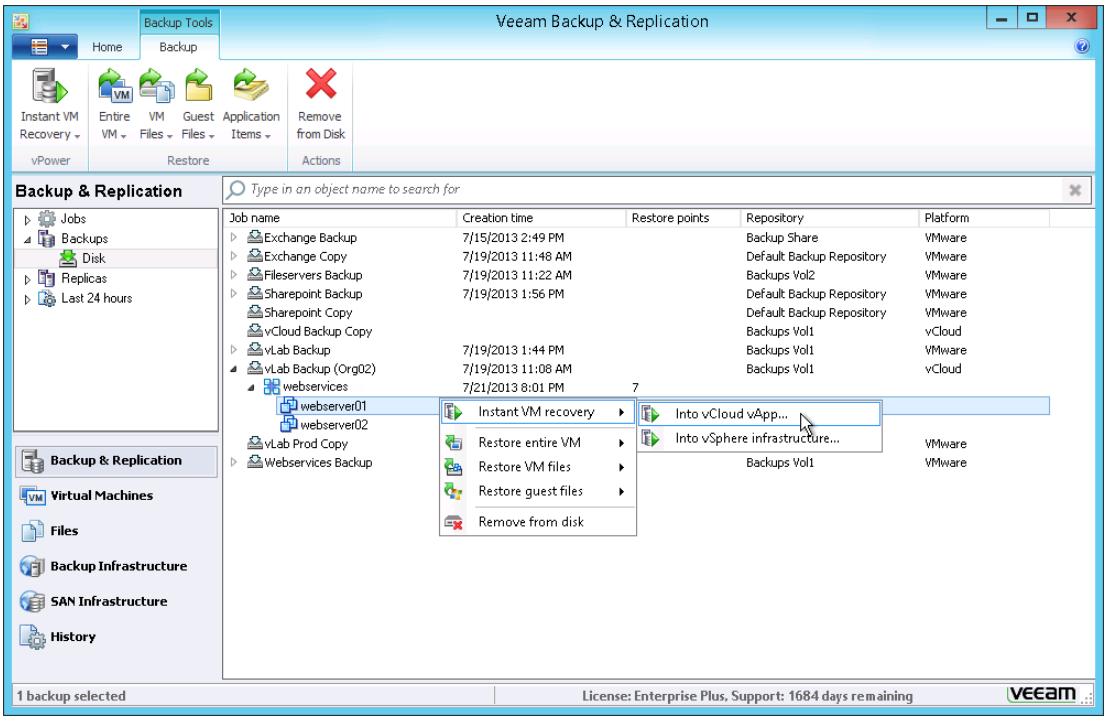

Step 2. Select a Restore Point

At this step of the wizard, you should select the necessary restore point for the vCD VM.

By default, Veeam Backup & Replication uses the latest good restore point to recover a VM. However, if you want to restore a VM to an earlier state, select from the list a necessary restore point that should be used to recover the VM.

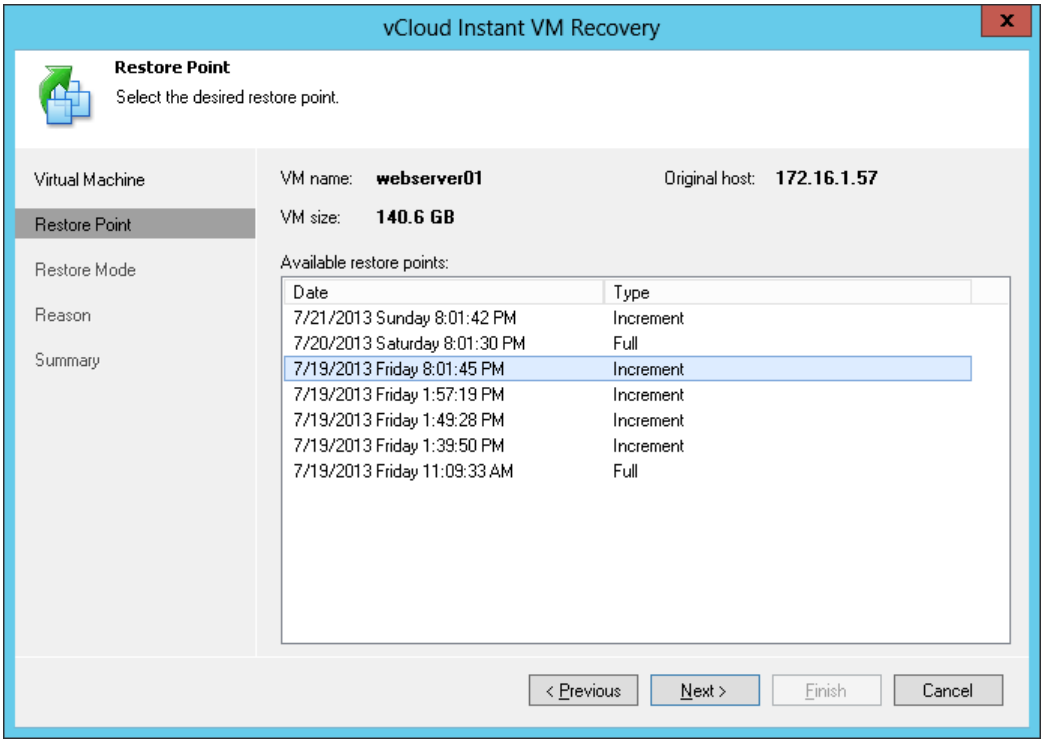

Step 3. Select a Restore Mode

Choose the necessary restore mode:

- Select **Restore to the original location** if you want to restore the VM with its initial settings and to its original location. If this option is selected, you will pass directly to the Reason step of the wizard.
- Select **Restore to a new location, or with different settings** if you want to restore the VM to a different location and/or with different settings (such as VM location, network settings, format of restored virtual disks and so on). If this option is selected, the **vCloud Instant VM Recovery** wizard will include an additional step for customizing VM settings.

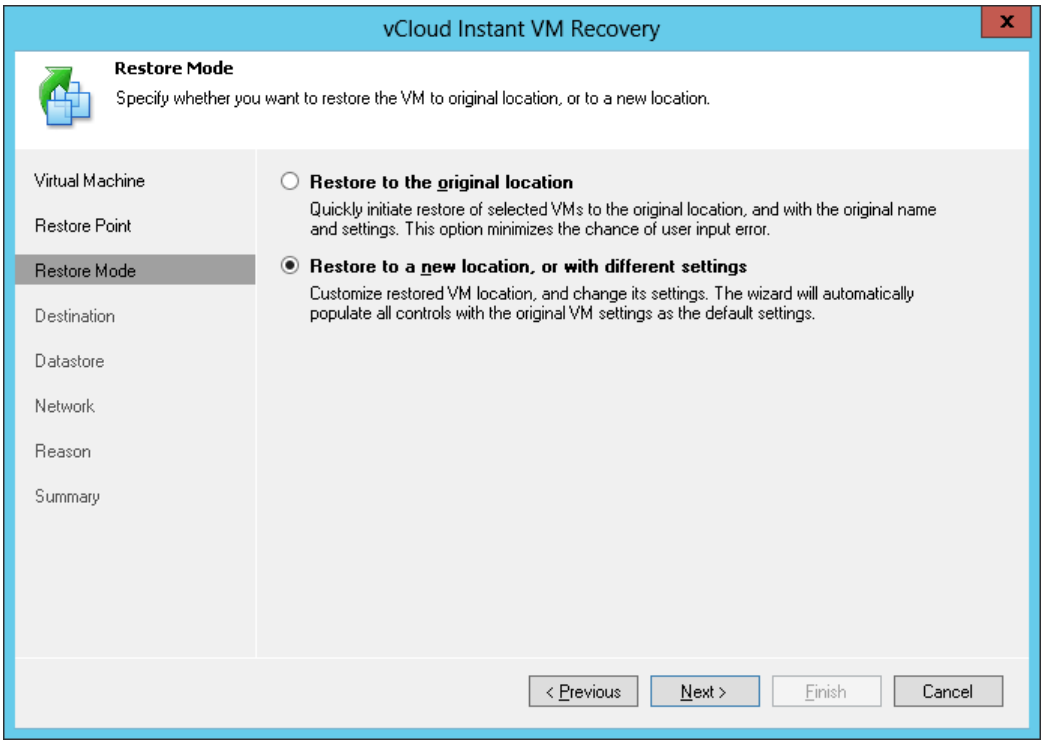

Step 4. Select Destination for the Restored VM

This step of the wizard is available if you have chosen to change the location and settings of the restored VM.

Select the destination for the recovered VM and specify the name for the restored VM:

1. In the **vApp** field, specify the vApp to which the VM should be restored. By default, Veeam Backup & Replication restores the VM to its initial vApp.
2. In the **Restored VM name** field, enter a name under which the VM should be restored and registered. By default, the original name of the VM is used. If you are restoring the VM to the same vApp where the original VM is registered and the original VM still resides there, it is recommended that you change the VM name to avoid conflicts.

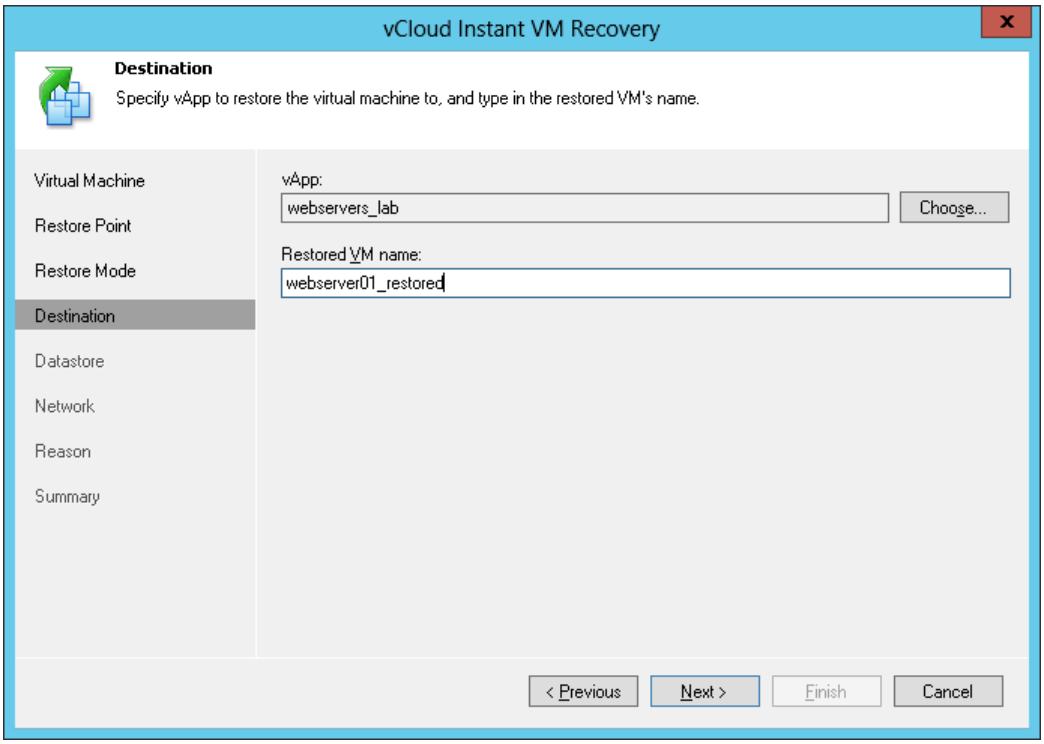

Step 5. Select Destination for Virtual Disk Updates

This step of the wizard is available if you have chosen to change the location and settings of the restored VM.

Select the location for holding the VM disk changes when the VM is restored. By default, disk changes are stored directly on the vPower NFS server. However, you can store disk changes on any datastore in your VMware environment. To do that, select the **Redirect virtual disk updates** check box and choose the necessary datastore. Redirecting disk changes improves recovery performance but makes Storage vMotion not possible for ESX 4.x and earlier.

Note	You can select only a datastore that is available in the Organization vCD hosting the vApp to which the VM is restored.

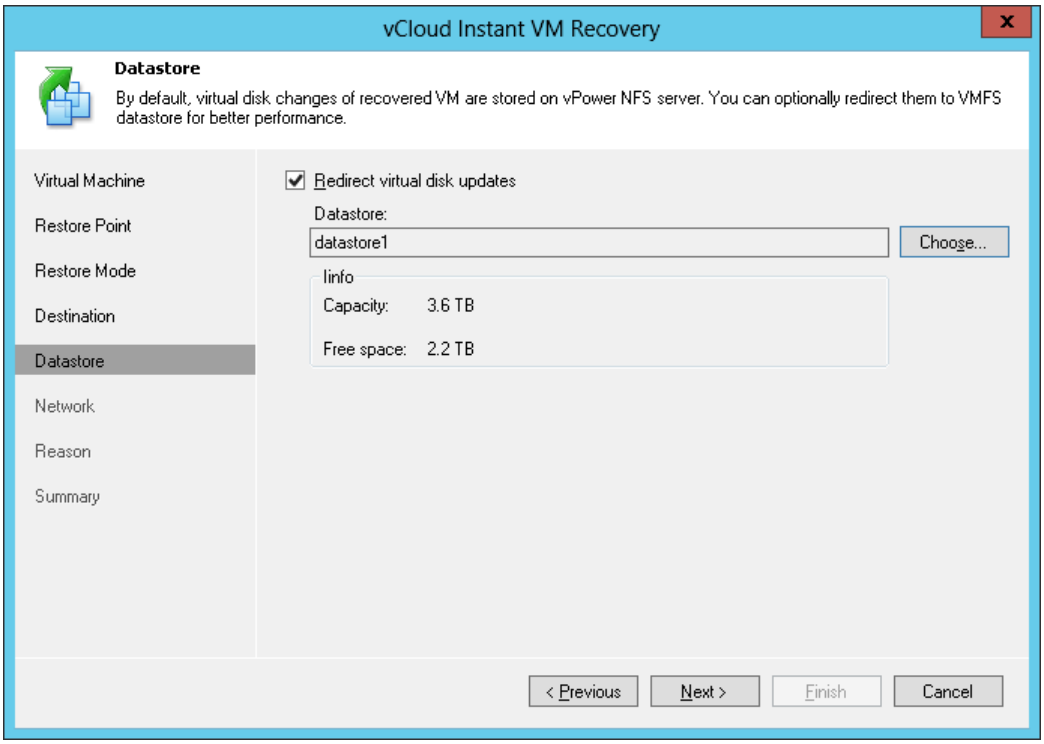

Step 6. Select a Destination Network

This step of the wizard is available if you have chosen to change the location and settings of the restored VM.

To select networks to which restored VMs will be connected:

1. Select the VM in the list and click **Network**.
2. The **Select Network** window displays all networks that are configured for the destination vApp. From the list of available networks, choose a network to which selected VM should have access upon restore.

To prevent the restored VM from accessing any network, select it in the list and click **Disconnect**.

Tip To facilitate selection, use the search field at the bottom of the window: enter a network name or a part of it and click the **Start search** button on the right or press **[ENTER]**.

Veeam Backup & Replication maps the network settings you define and network settings of the initial VM. If necessary, Veeam Backup & Replication makes changes to the network settings of the restored VM. For example, if the initial VM was connected to the network using the static IP mode and you have selected to connect a restored VM to a network using the dynamic IP mode,
Veeam Backup & Replication will change the network settings to the dynamic mode.

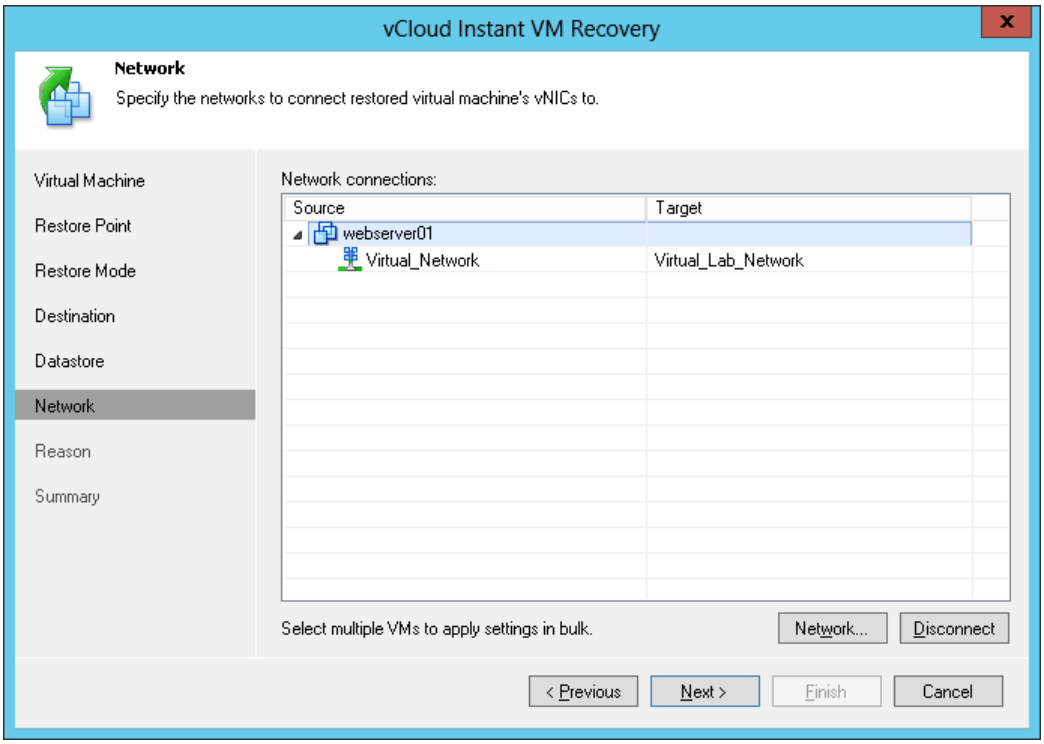

Step 7. Specify a Restore Reason

If necessary, enter the reason for performing instant restore of the VM. The information you provide will be saved in the session history so that you can reference it later.

Note Veeam Backup & Replication checks the lease term for the vApp to which the VM is restored. In case the lease period has expired, the lease will be automatically updated.

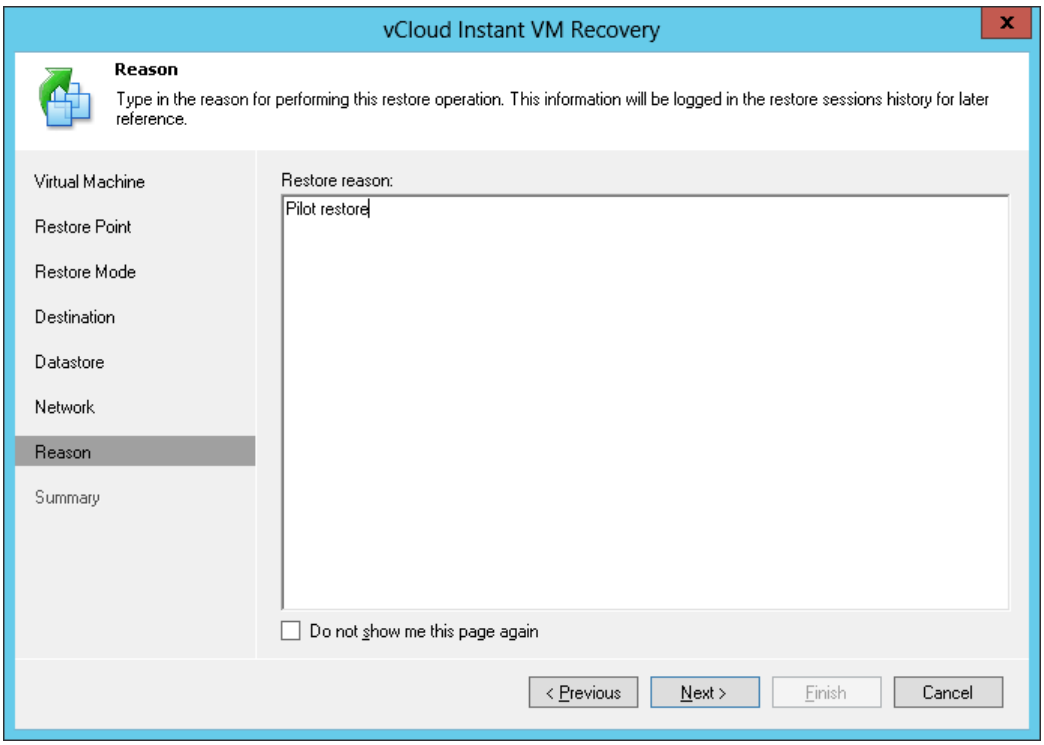

Step 8. Verify Instant VM Recovery Settings

At the last step of the wizard, check the settings you have specified for Instant VM Recovery.

To start a VM immediately after the restore process is successfully complete, select the **Power on VM automatically** check box. If you are recovering the VM to the production network, make sure that the initial VM is powered off to avoid conflicts.

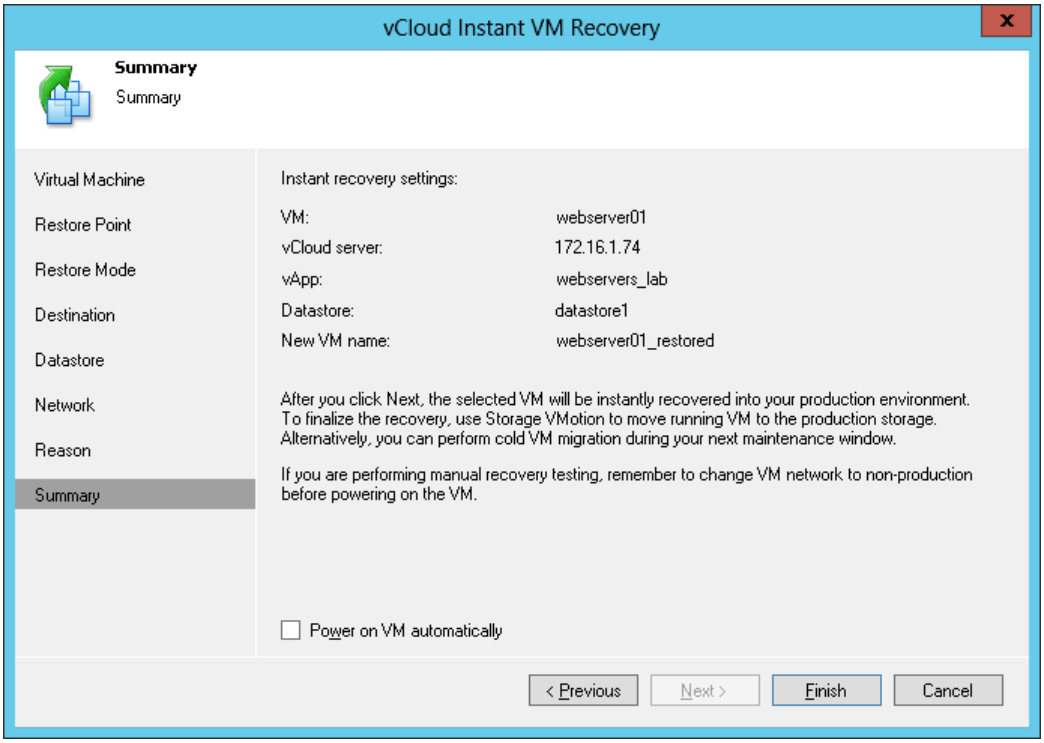

Step 9. Finalize Instant VM Recovery

After the VM has been successfully restored, you can finalize Instant VM Recovery: migrate the restored VM to production or remove the restored VM.

To migrate the restored VM to production:

1. Open the **Backup & Replication** view.
2. In the inventory pane, select the **Instant Recovery** node.
3. Right-click the VM in the working area and select **Migrate to production**. As a result, the **Quick Migration** wizard will be launched. During migration, Veeam Backup & Replication will restore a VM instance from the backup file and then additionally move the changes that were made while the VM was running in the *Instant Recovery* mode. To learn more, see Migrating Virtual Machines.

Note Quick Migration will be performed using VMware vMotion technology.

To remove the recovered VM:

1. Open the **Backup & Replication** view.
2. In the inventory pane, select the **Instant Recovery** node.
3. Right-click the necessary VM in the working area and select **Stop publishing**.

Restoring vCD VMs with Instant VM Recovery to Virtual Infrastructure

To launch the **Instant Recovery** wizard, do one of the following:

- Open the **Backup & Replication** view and select the **Backups** node. In the working area, expand the necessary backup job, select the VM you want to restore and click **Instant VM Recovery** > **Into virtual infrastructure** on the ribbon.
- Open the **Backup & Replication** view and select the **Backups** node. In the working area, expand the necessary backup job, right-click the VM you want to restore and select **Instant recovery** > **Into vSphere Infrastructure**.
- Open the **Virtual Machines** view. On the **View** tab, click **vCloud Director View**. In the inventory pane, expand the vCloud Director hierarchy. In the working area, right-click the VM you want to restore and select **Restore** > **Instant VM Recovery** > **Into vSphere infrastructure**.

The process of Instant VM Recovery for vCD VMs does not differ from the regular Instant VM Recovery process. To learn what steps you should perform and what settings you can specify, see the Performing Instant VM Recovery section.

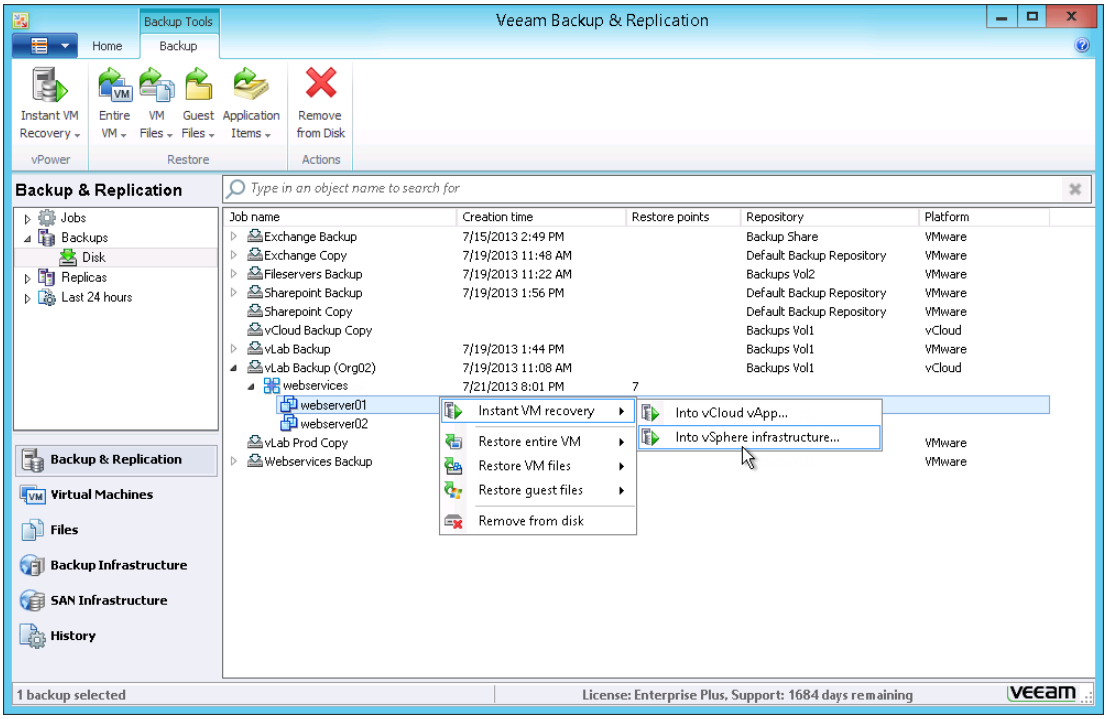

Restoring vApps to vCloud Director

In addition to vCD VMs restore, Veeam Backup & Replication lets you restore vApps from backups back to vCloud Director.

vApps can be restored to their Organization vDC or to any other Organization vDC of your choice. You can restore the vApp that already exists — for example, in case the initial vApp is corrupted or you want to revert to an earlier state of the vApp, or the vApp that no longer exists — for example, if it was deleted by mistake. If you restore a vApp that already exists, the vApp is overwritten with that from the vCD backup.

To restore a vApp to vCloud Director, follow the next steps:

Step 1. Launch the Full vApp Restore Wizard

To launch the **Full vApp Restore** wizard, do one of the following:

- On the **Home** tab, click **Restore** and select **vCloud**. At the **Object Type** step, select the object you would like to restore: *vApp*.
- Open the **Backup & Replication** view and select the **Backups** node. In the working area, expand the necessary backup job, select the vApp in the working area and click **Restore vCloud vApp** on the ribbon.
- Open the **Backup & Replication** view and select the **Backups** node. In the working area, expand the necessary backup job, right-click the vApp in the working area and select **Restore vCloud vApp**.

Step 2. Select a vApp to Restore

At this step of the wizard, you should select the vApp you want to restore. To add a vApp, click **Add vApp** and select where to browse for the vApps:

- **From Infrastructure** — browse the vCloud Director hierarchy and select a vApp to restore. Note that the vApp you select from the vCloud Director hierarchy must be successfully backed up at least once.
- **From Backup** — browse existing backups and select the vApp under backup jobs.

Tip To facilitate selection, use the search field at the bottom of the **Select VMs** window: enter an object's name or a part of it and click the **Start search** button on the right or press **[ENTER]**.

To add a vApp to the list, you can also use the search field at the top of the window:

1. Enter a vApp name or a part of it in the search field and Veeam Backup & Replication will search existing backups for the specified vApps and display matching results.
2. To add the vApp to the list, double-click it in the list of search results.
3. If the necessary vApp is not found, click the **Show more** link to browse existing backups and choose the necessary vApp.

To remove a vApp from the list, select it and click **Remove** on the right.

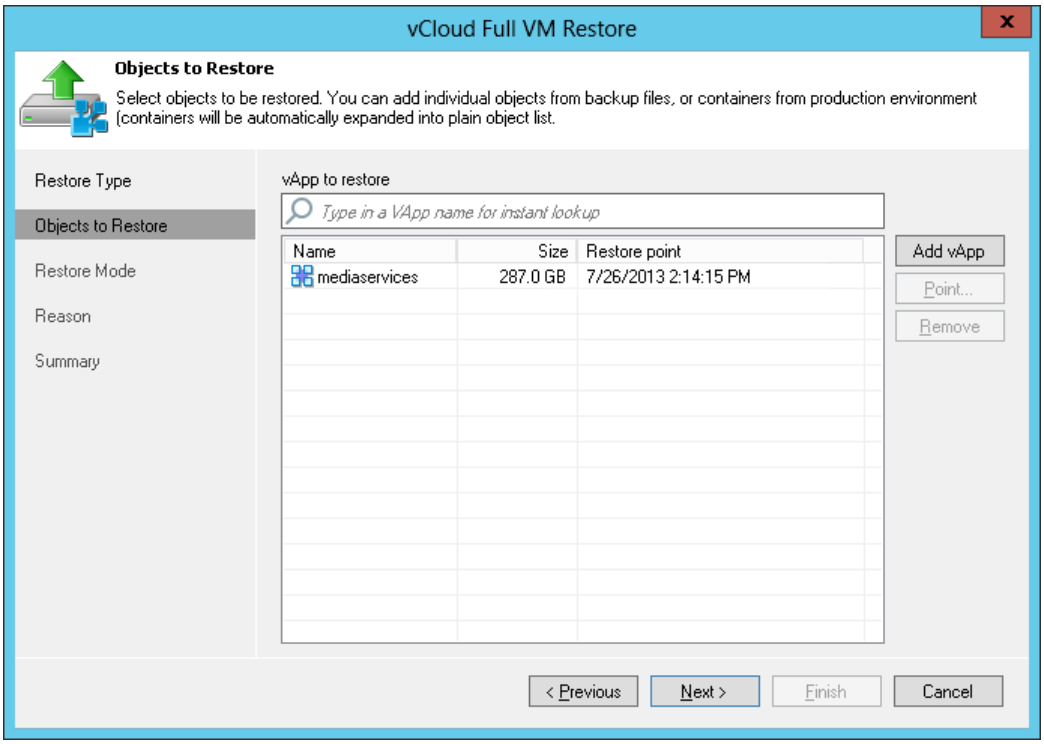

Step 3. Select a Restore Point

At this step of the wizard, you should select the necessary restore point for the vApp.

By default, Veeam Backup & Replication uses the latest good restore point to recover a vApp. However, if you want to restore a vApp to an earlier state, select it in the list and click **Point** on the right. In the **Restore Points** window, select a restore point that should be used to recover the vApp.

Note If you run a vCD backup job for the vApp, the job is considered to finish with the *Success* status and the restore point is created only if all VMs in the vApp are backed up successfully. If any VM in the job fails, the restore point for the vApp is not created and you will not be able to restore the vApp from such restore point.

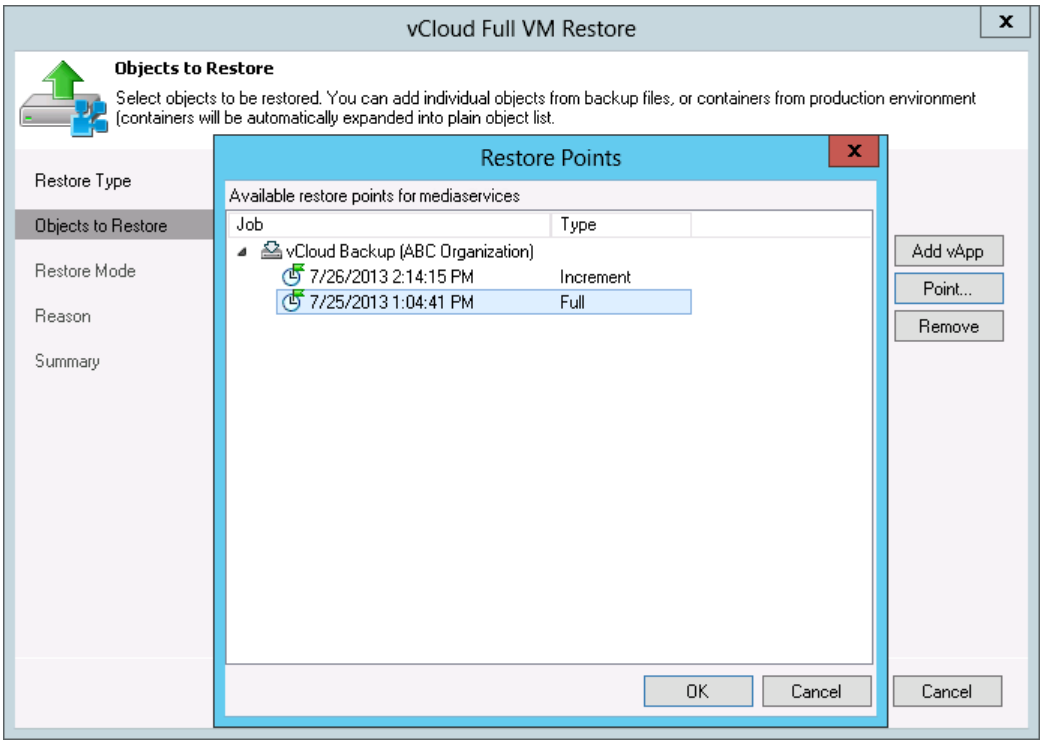

Step 4. Select a Restore Mode

At this step of the wizard, you should select the destination location for the restored vApp.

- Select **Restore to original location** if you want to restore the vApp with its initial settings and to its original location. If this option is selected, you will immediately pass to the Summary step of the wizard.
- Select **Restore to a new location, or with different settings** if you want to restore the vApp to a different location and/or with different settings (such as Organization vDC, network settings, fast provisioning settings and so on). If this option is selected, the **Full vApp Restore** wizard will include additional steps for customizing vApp settings.

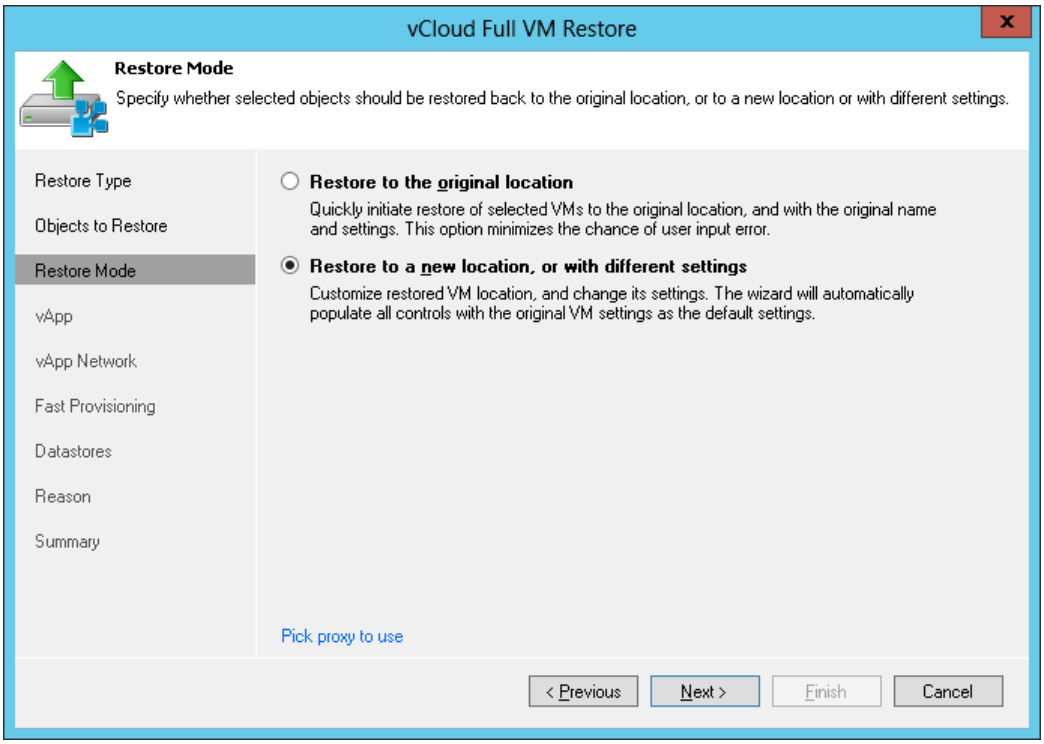

Step 5. Select the vApp Location

This step of the wizard is available if you have chosen to change the location and settings of the restored vApp.

By default, Veeam Backup & Replication restores the vApp to its initial location. To restore the vApp to a different location:

1. Select the App in the list and click **vDC**.
2. From the vCloud Director hierarchy, choose an Organization vDC where the selected vApp should be registered.

Tip To facilitate selection, use the search field at the bottom of the window: enter an object's name or a part of it and click the **Start search** button on the right or press **[ENTER]**.

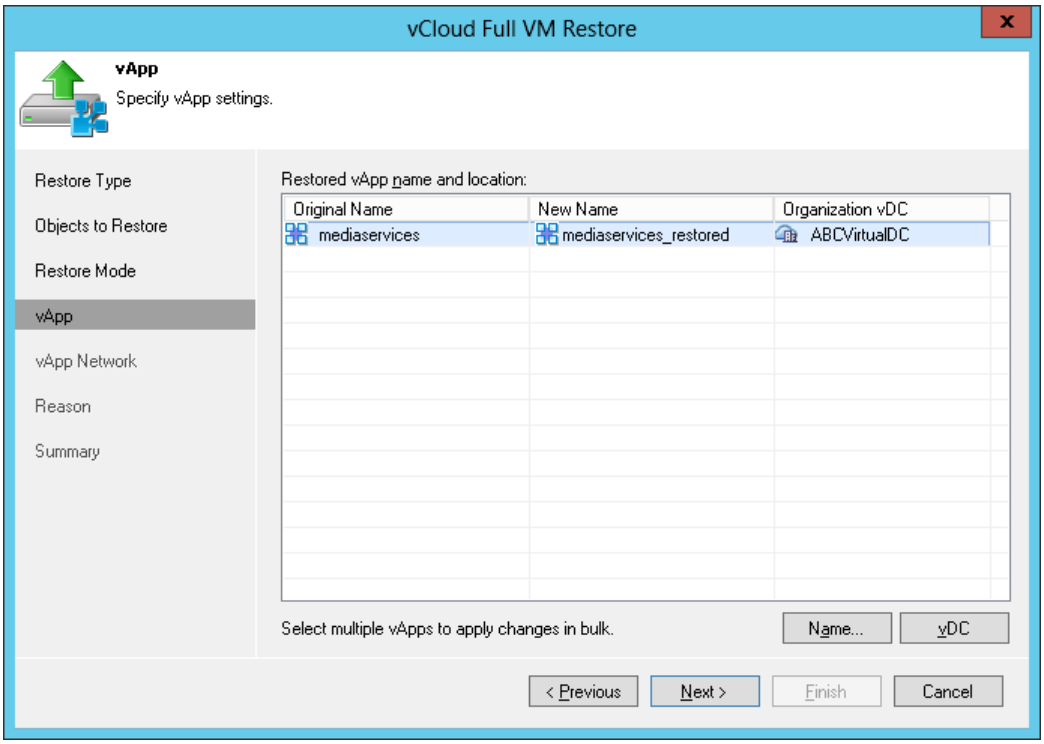

Step 6. Select the Destination Network

This step of the wizard is available if you have chosen to change the location and settings of the restored vApp.

To select networks to which the restored vApp should connected:

1. Select the vApp in the list and click **Network**.
2. The **Select Network** window displays all networks that are configured for the destination Organization vDC. From the list of available networks, choose a network to which selected vApp should have access upon restore.
   To facilitate selection, use the search field at the bottom of the window: enter a network name or a part of it and click the **Start search** button on the right or press **[ENTER]**.

To prevent the restored vApp from accessing any network, select it in the list and click **Disconnect**.

Note When Veeam Backup & Replication backs up a vApp, it saves information about organization networks to which the vApp is connected along with vApp networks data. If you restore the vApp to the initial vDC organization and do not change the organization network settings, Veeam Backup & Replication attempts to connect the vApp to all source organization networks to which the vApp was connected at the moment of backup. If one or several source organization networks are not detected, for example, if they have been changed or removed by the time of restore, Veeam Backup & Replication will not be able to restore the vApp to its initial organization. In this case, you will need to change the network settings for the restored vApp: map the vApp to other organization network(s) or disconnect it from organization network(s) at all.

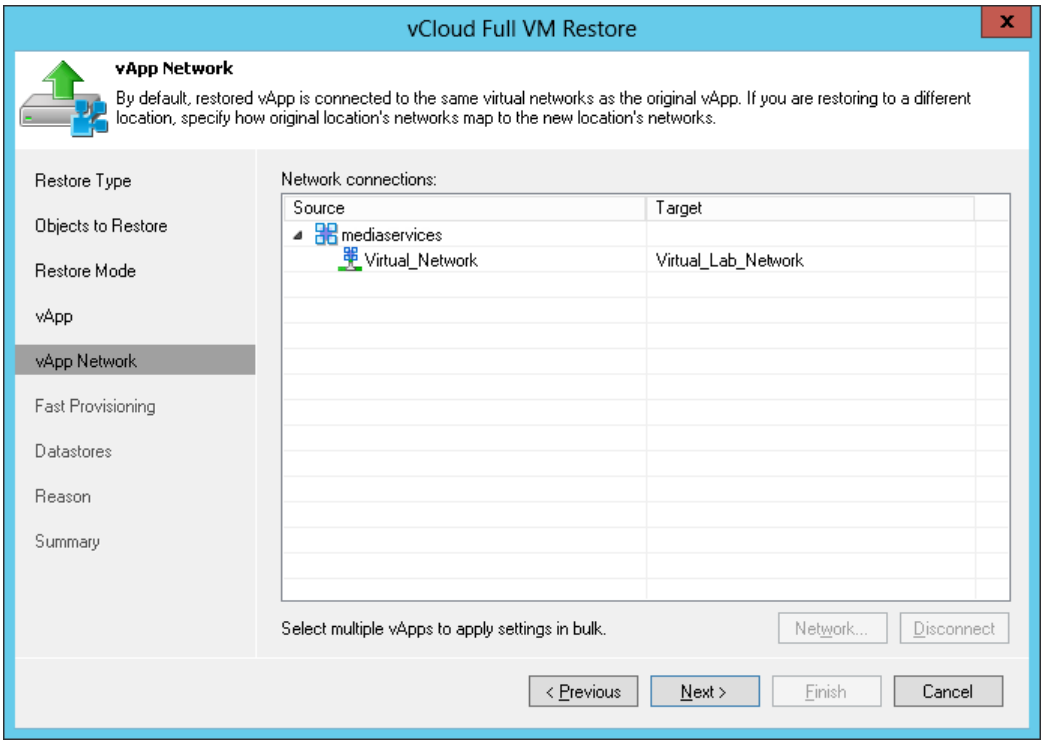

Step 7. Select a Template to Link

This step of the wizard is available if you have chosen to change the settings of the restored vApp, for example, its name or location.

To select a VM template:

1. Select the VM in the list and click **Set Template**.
2. From the vCloud Director hierarchy, choose a template to which the VMs from the restored vApp should be linked.
   To facilitate selection, use the search field at the bottom of the window: enter a VM template name or a part of it and click the **Start search** button on the right or press **[ENTER]**.

If you want to disable fast provisioning for the VM and restore it as a regular VM, select the VM in the list and click **Disable**.

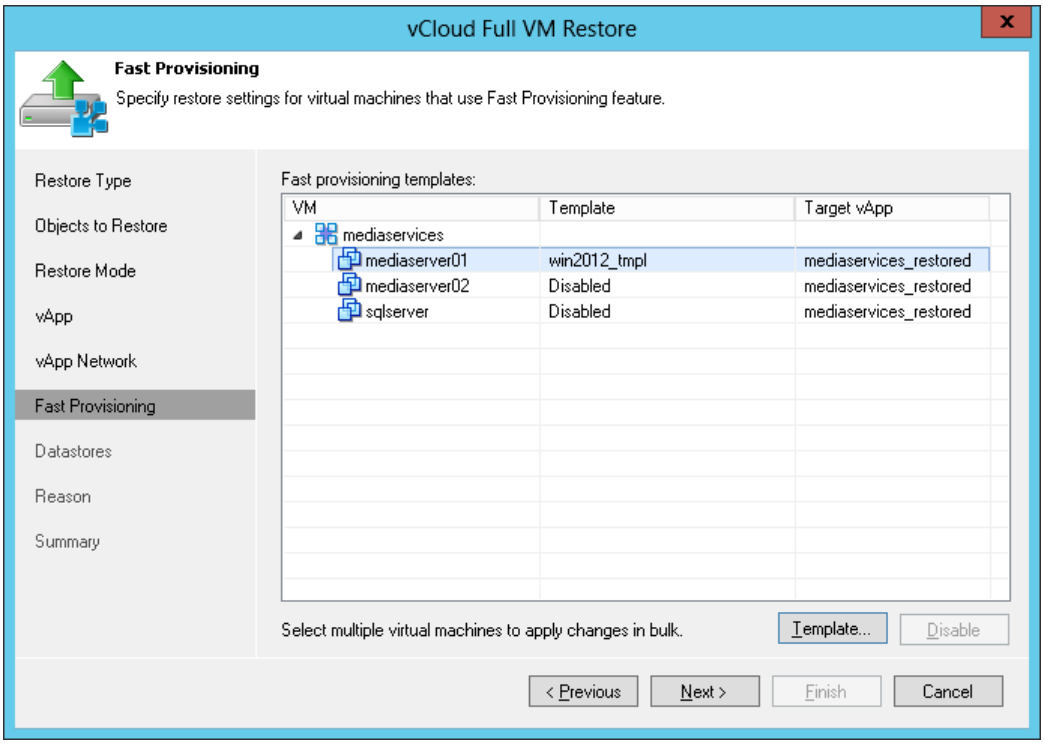

Step 8 . Select a Storage Profile and Datastores

This step of the wizard is available if you have chosen to change the settings of the restored vApp, for example, its name or location.

To select a storage profile for the vApp:

1. Select the vApp in the list and click **Profile**.
2. In the displayed window, select the necessary profile for the vApp.

If you have selected to disable fast provisioning at the previous step, you should select a datastore on which the disks of restored VMs should be placed. To do that:

1. Select VM or vApp in the list and click **Datastore**.
2. In the displayed window, select the datastore on which the disks of the VM should be placed.

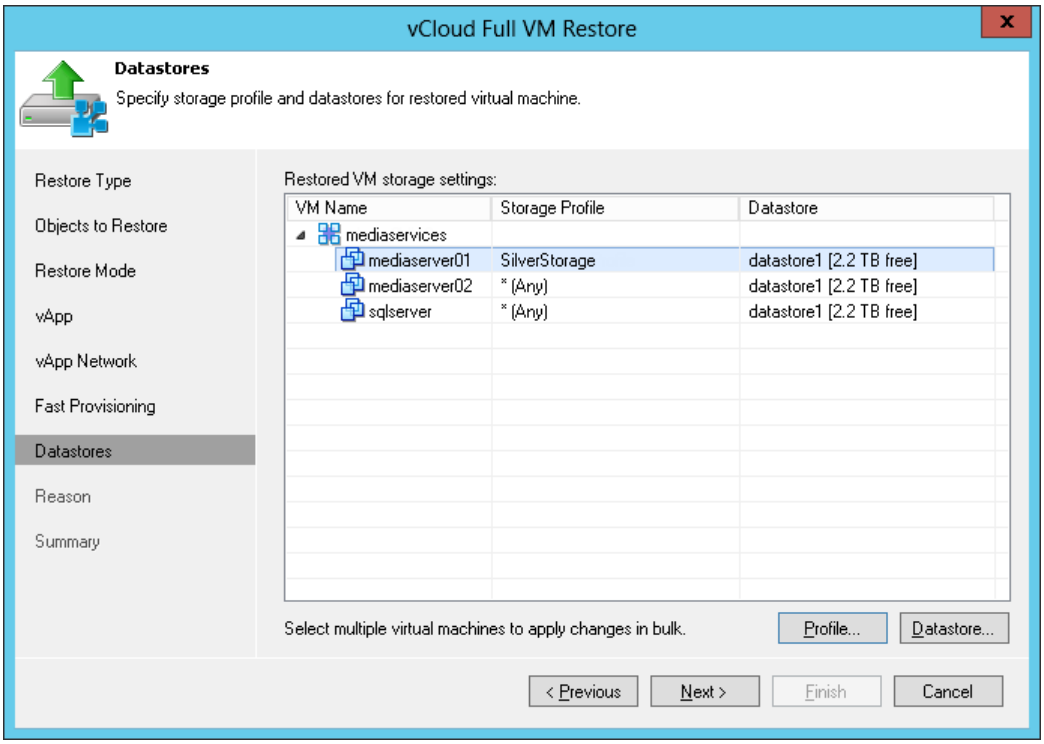

Step 9. Specify a Restore Reason

If necessary, enter the reason for performing restore of the selected vApp. The information you provide will be saved in the session history so that you can reference it later.

Note Veeam Backup & Replication checks the lease term for the restored vApp. In case the lease period has expired, the lease will be automatically updated.

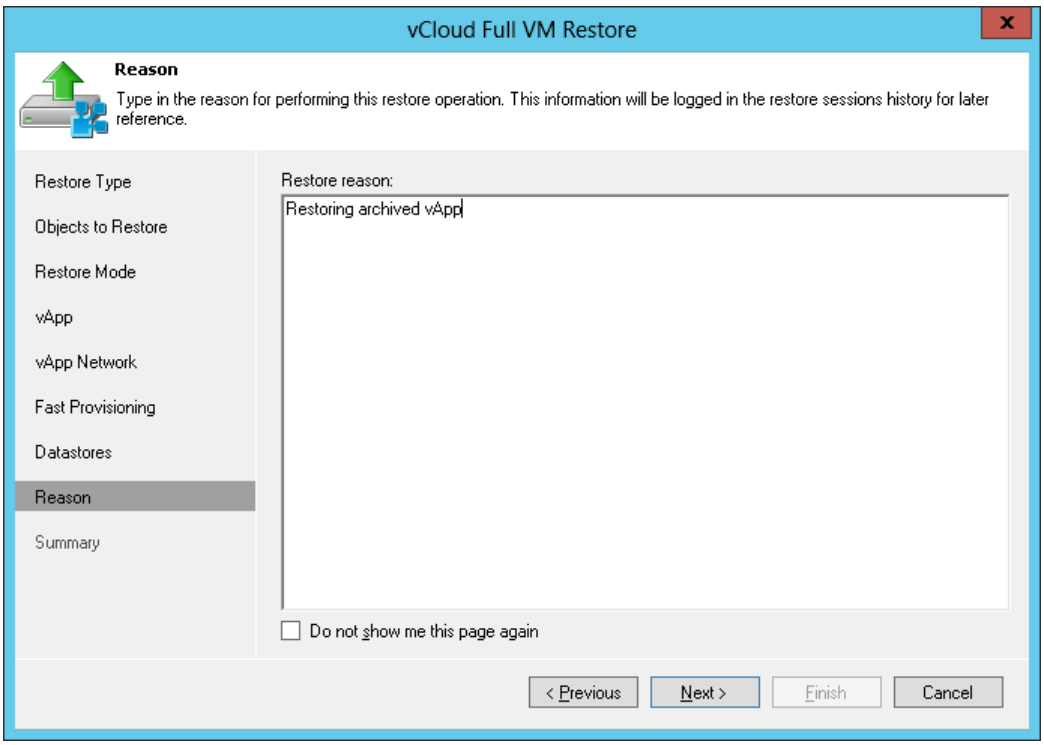

Step 10. Verify Recovery Settings

If you want to start VMs in the vApp after the restore process is successfully complete, select the **Power on VM after restoring** check box under the list of restore points.

Check the specified settings and click **Finish**. Veeam Backup & Replication will restore the selected vApp to the specified destination.

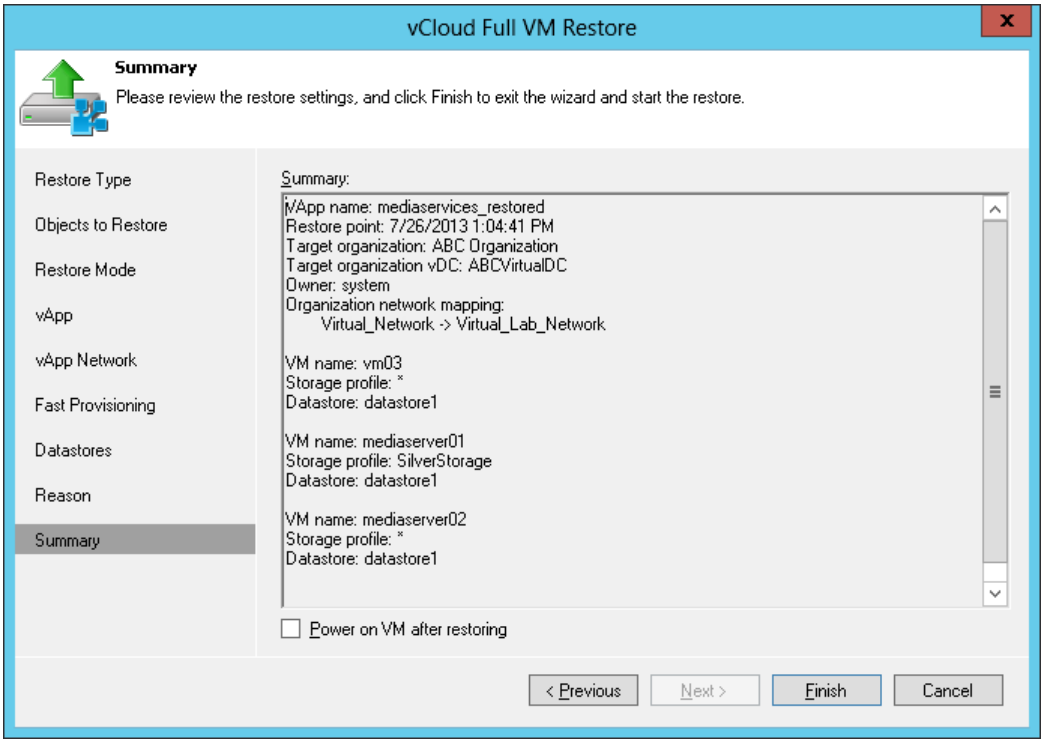

Restoring VMs to vCloud Director

Veeam Backup & Replication lets you restore VMs from vCD backups back to the vCloud Director hierarchy. You can restore a single VM or a number of VMs at once.

The vCD VM can be restored to its initial location or to any other location of your choice. You can restore a VM that already exists — for example, in case the initial VM is corrupted or you want to revert to an earlier state of the VM, or a VM that no longer exists — for example, if it was deleted by mistake. If you restore a VM that already exists, the initial VM is overwritten with that from the vCD backup.

Note When restoring VMs to the vCloud Director hierarchy, make sure that you select the **Restore into vCloud vApp** option. If you select the **Restore into VI** option, the VM will be restored at the level of the underlying vCenter Server. To get a fully functional VM managed by vCloud Director, you will need to manually import the restored VM to the vCloud Director hierarchy.

To restore a VM to the vCloud Director hierarchy, follow the next steps:

Step 1. Launch the vCloud Full VM Restore Wizard

To launch the **vCloud Full VM Restore** wizard, do one of the following:

- On the **Home** tab, click **Restore** and select **vCloud**. At the **Object Type** step, select the object you would like to restore: *VM*.
- Open the **Backup & Replication** view and select the **Backups** node. In the working area, expand the necessary backup job, select the VM you want to restore and click **Entire VM** > **Into vCloud vApp** on the ribbon.

- Open the **Backup & Replication** view and select the **Backups** node. In the working area, expand the necessary backup job, right-click the VM you want to restore and select **Restore entire VM** > **Into vCloud vApp**.
- Open the **Virtual Machines** view. On the **View** tab, click **vCloud Director View**. In the inventory pane, expand the vCloud Director hierarchy. In the working area, right-click the VM you want to restore and select **Restore** > **Restore VM into vCloud vApp**.

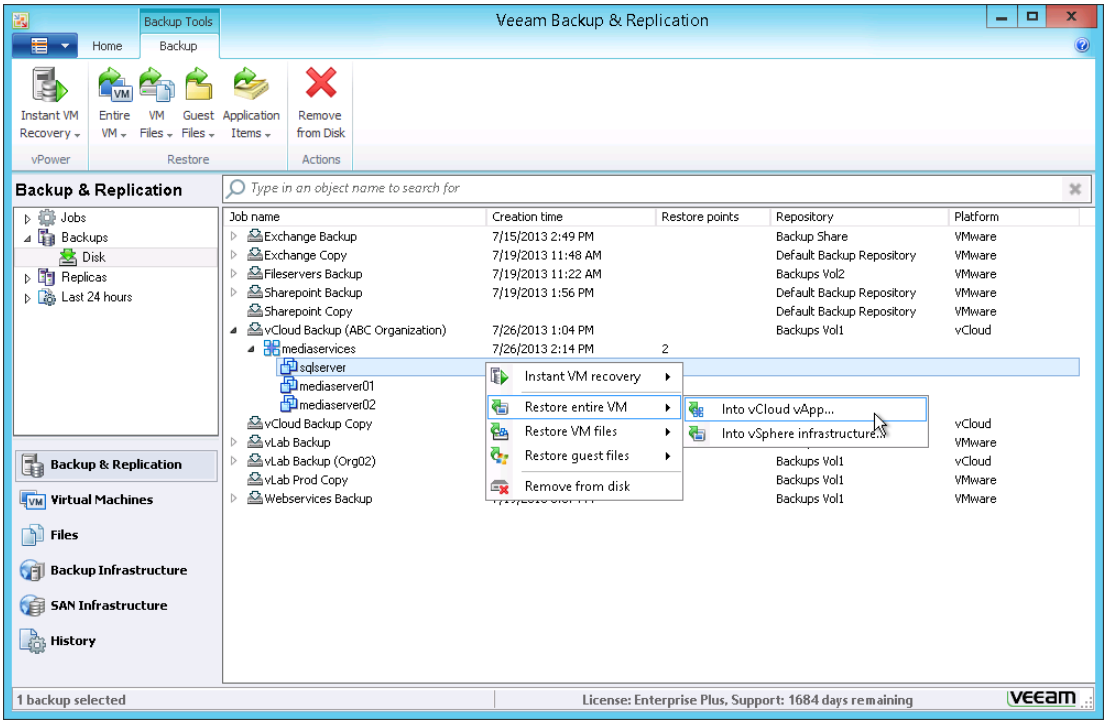

Step 2. Select a VM to Restore

At this step of the wizard, you should select one or more VMs to restore. To add a VM, click **Add VM** and select where to browse for VMs:

- **From Infrastructure** — browse the vCloud Director hierarchy and select VMs to restore. Note that the VM you select from the vCloud Director hierarchy must be successfully backed up at least once.
- **From Backup** — browse existing backups and select VMs under backup jobs.

Tip To facilitate selection, use the search field at the bottom of the **Select VMs** window: enter an object's name or a part of it and click the **Start search** button on the right or press **[ENTER]**.

To add VMs to the list, you can also use the search field at the top of the window:

1. Enter a VM name or a part of it in the search field and Veeam Backup & Replication will search existing backups for the specified VM and display matching results.
2. To add the VM to the list, double-click it in the list of search results.
3. If the necessary VM is not found, click the **Show more** link to browse existing backups and choose the necessary VM.

To remove a VM from the list, select it and click **Remove** on the right.

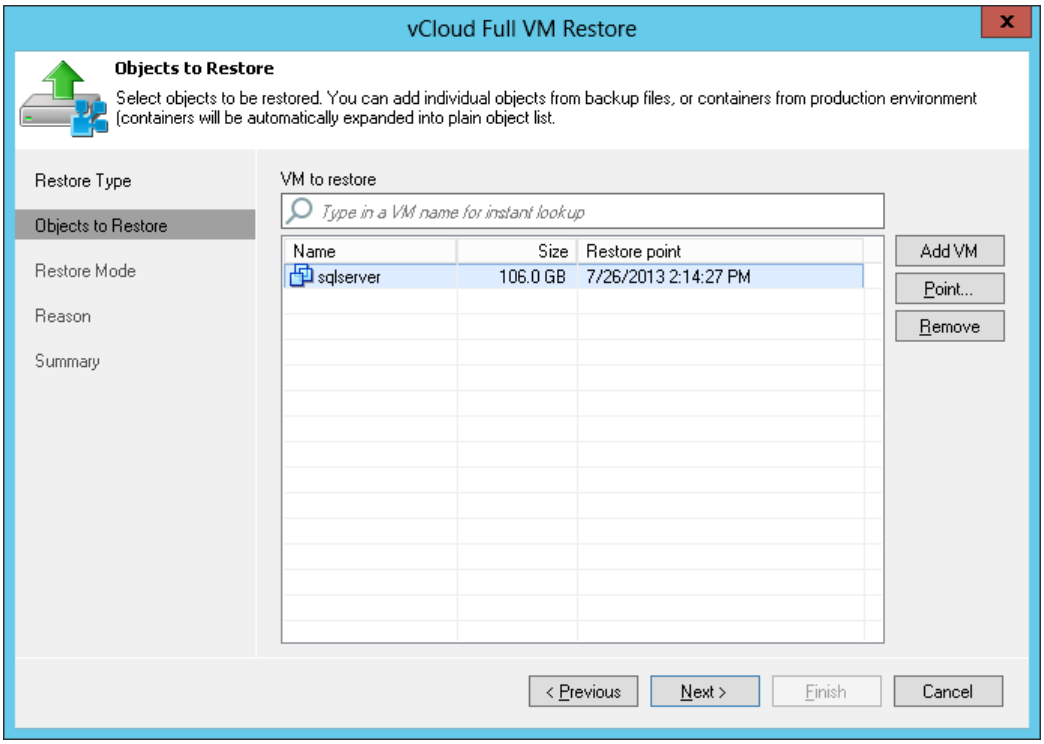

Step 3. Select a Restore Point

At this step of the wizard, you should select the necessary restore point for the VM.

By default, Veeam Backup & Replication uses the latest good restore point to recover a VM. However, if you want to restore the VM to an earlier state, select it in the **VMs to restore** list and click **Point** on the right. In the **Restore Points** window, select a restore point that should be used to recover the VM.

If you have chosen to restore multiple VMs, you can select a different restore point for every VM specifically.

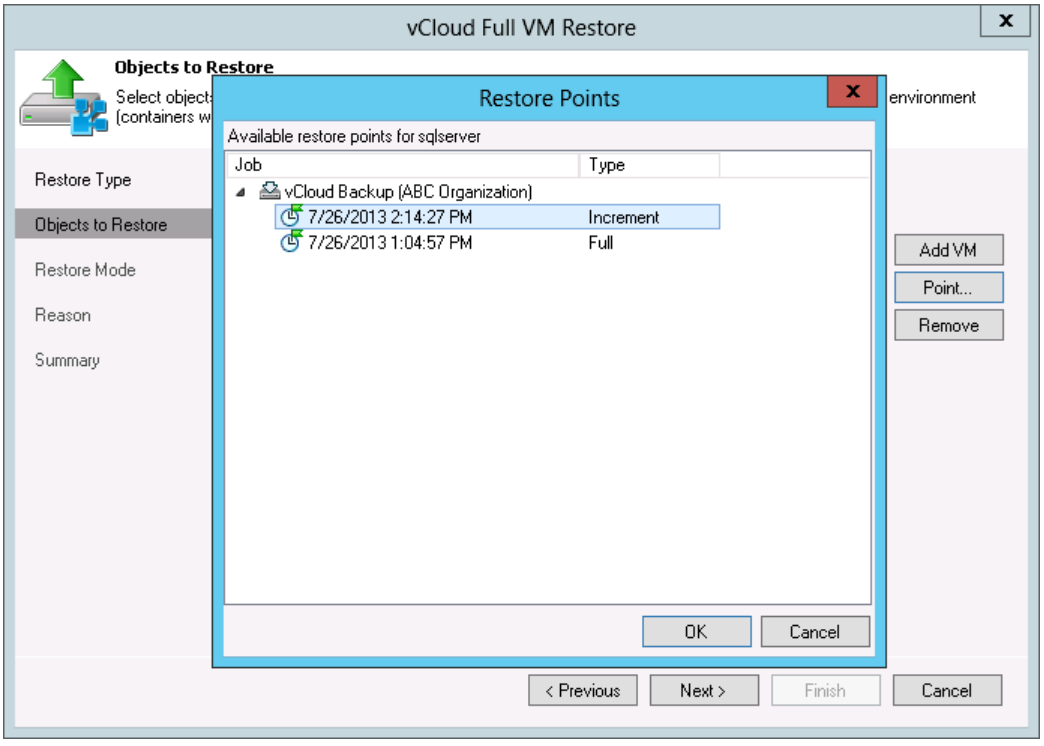

Step 4. Select a Restore Mode

At this step of the wizard, you should select the destination location for the restored VM.

- Select **Restore to original location** if you want to restore the VM with its initial settings and to its original location. If this option is selected, you will immediately pass to the Summary step of the restore process.
- Select **Restore to a new location, or with different settings** if you want to restore the VM to a different location and/or with different settings (such as VM location, network settings, fast provisioning settings and so on). If this option is selected, the wizard will include additional steps for customizing VM settings.

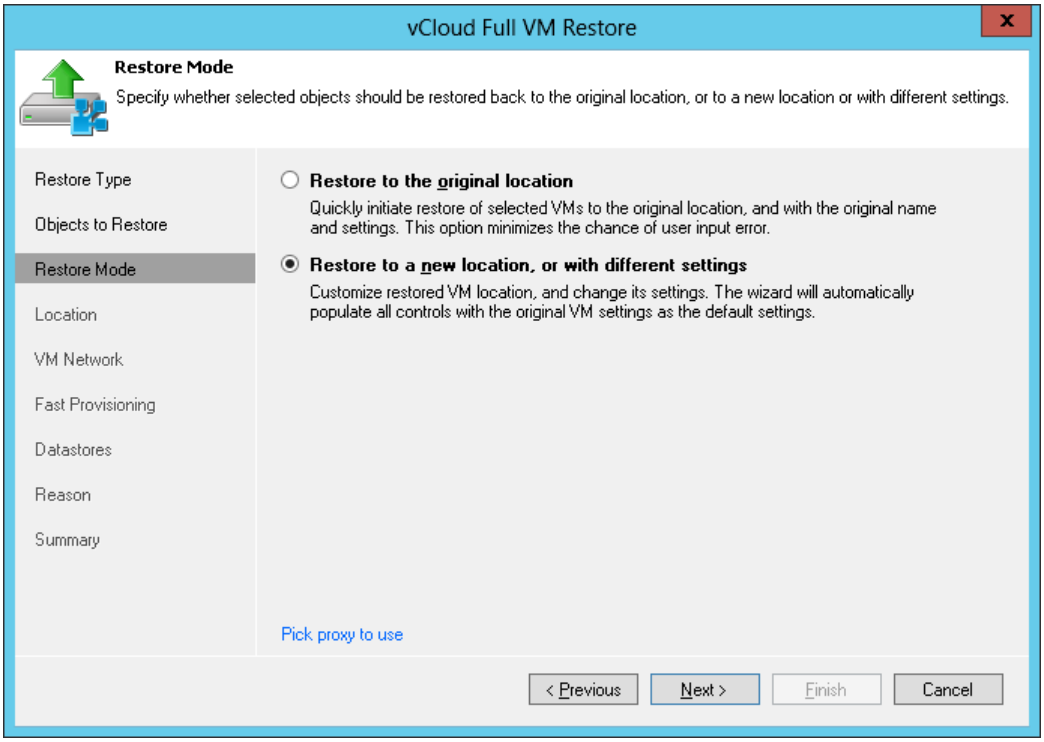

Step 5. Select the VM Location

This step of the wizard is available if you have chosen to change the location and settings for the restored VM.

By default, Veeam Backup & Replication restores the VM to its initial location. To restore the VM a different location:

1. Select the VM in the list and click **vApp**.
2. From the vCloud Director hierarchy, choose the vApp in which the restored VM should be registered.
   To facilitate selection, use the search field at the bottom of the window: enter the vApp name or a part of it and click the **Start search** button on the right or press **[ENTER]**.

Important! If you are restoring a linked clone VM to a different location, make sure that fast provisioning is enabled at the level of the target Organization vDC. Otherwise Veeam Backup & Replication will restore the VM as a regular VM.

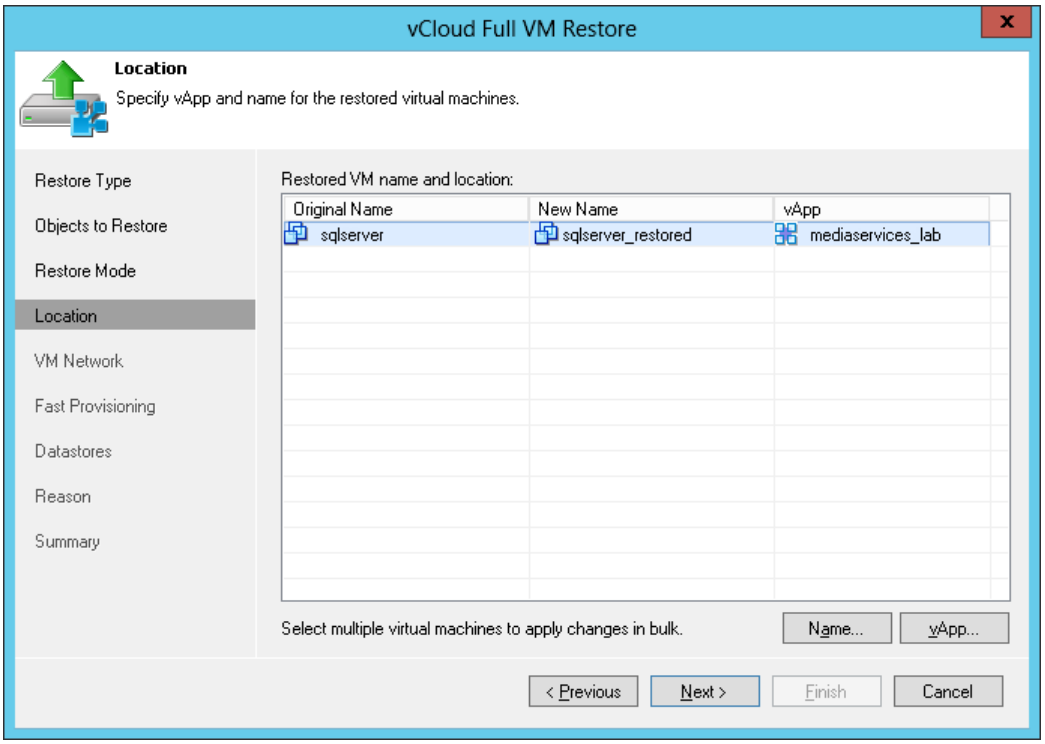

Step 6. Select the Destination Network

This step of the wizard is available if you have chosen to change the location and settings of the restored VM.

To select networks to which the restored VM should be connected:

1. Select the VM in the list and click **Networks**.
2. The **Select Network** window displays all networks that are configured for the destination vApp. From the list of available networks, choose a network to which the restored VM should have access upon restore.
   To facilitate selection, use the search field at the bottom of the window: enter a network name or a part of it and click the **Start search** button on the right or press **[ENTER]**.

To prevent the restored VM from accessing any network, select it in the list and click **Disconnected**.

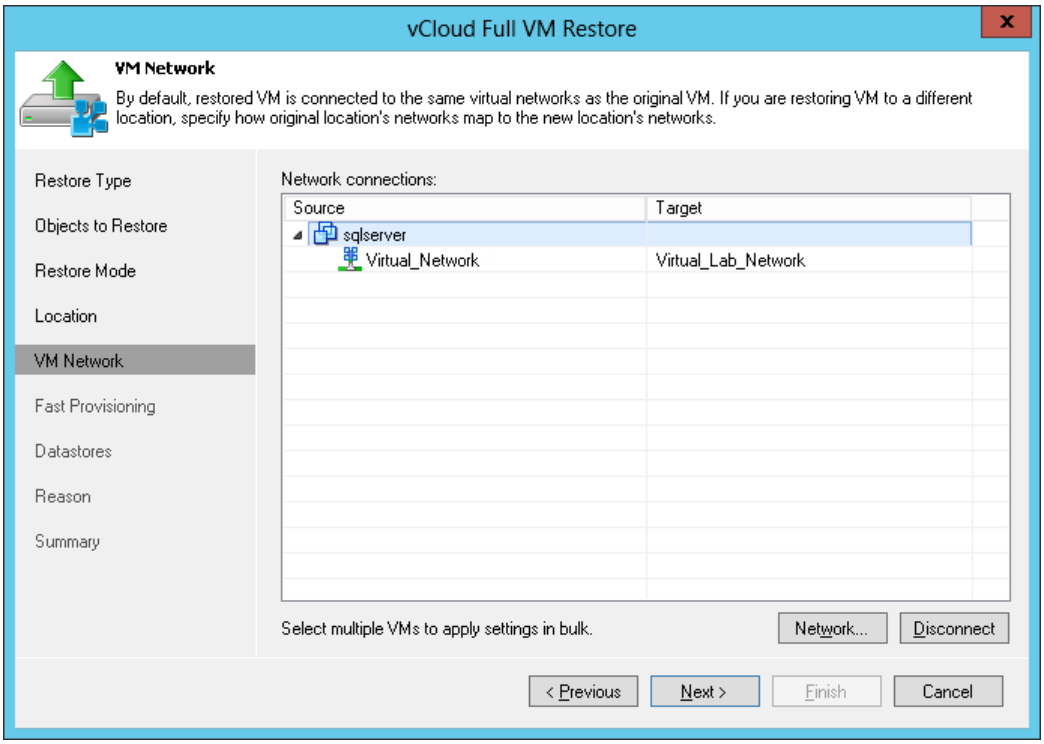

Step 7. Select a Template to Link

This step of the wizard is available if you have chosen to change the settings of the restored VM, for example, its name or location.

To select a VM template:

1. Select the VM in the list and click **Set Template**.
2. From the vCloud Director hierarchy, choose a template to which the restored VM should be linked.
   To facilitate selection, use the search field at the bottom of the window: enter a VM template name or a part of it and click the **Start search** button on the right or press **[ENTER]**.

If you want to disable fast provisioning for the VM and restore it as a regular VM, select the VM in the list and click **Disable**.

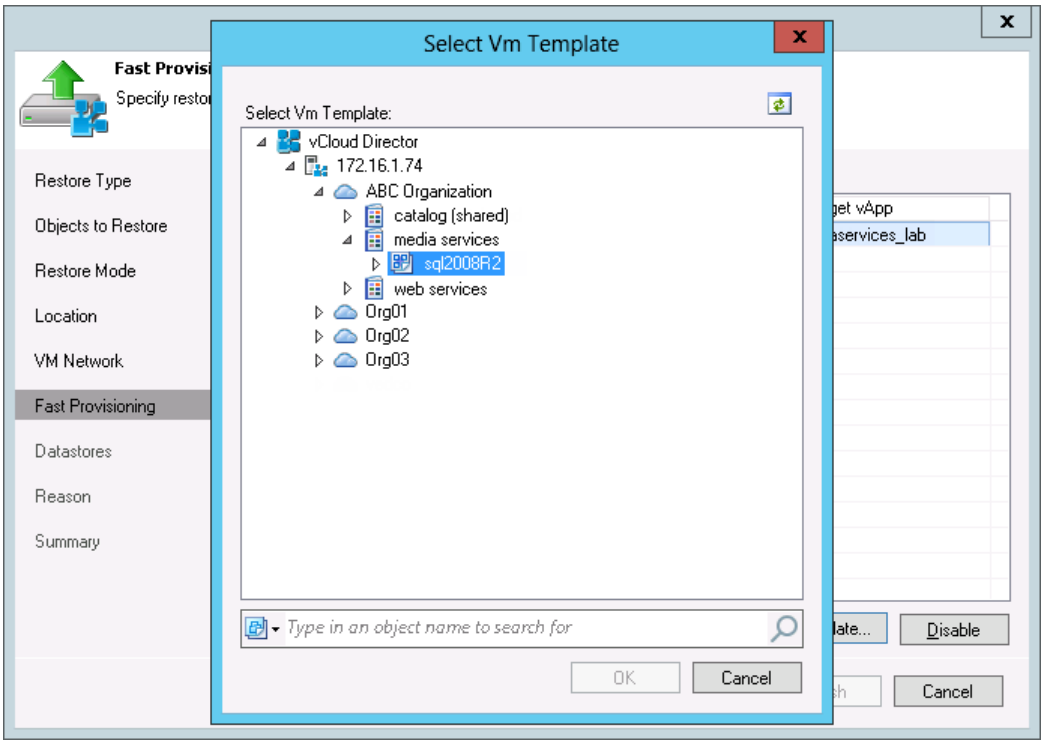

Step 8. Select a Storage Profile and Datastores

This step of the wizard is available if you have chosen to change the settings of the restored VM, for example, its name or location.

To select a storage profile for the restored VM:

1. Select the VM in the list and click **Storage Profile**.
2. In the displayed window, select the necessary profile for the VM.

If you have selected to disable fast provisioning at the previous step, you should select a datastore on which disks of the restored VM should be placed.

1. Select the VM in the list and click **Datastore**.
2. In the displayed window, select the datastore on which the VM disks will be located.

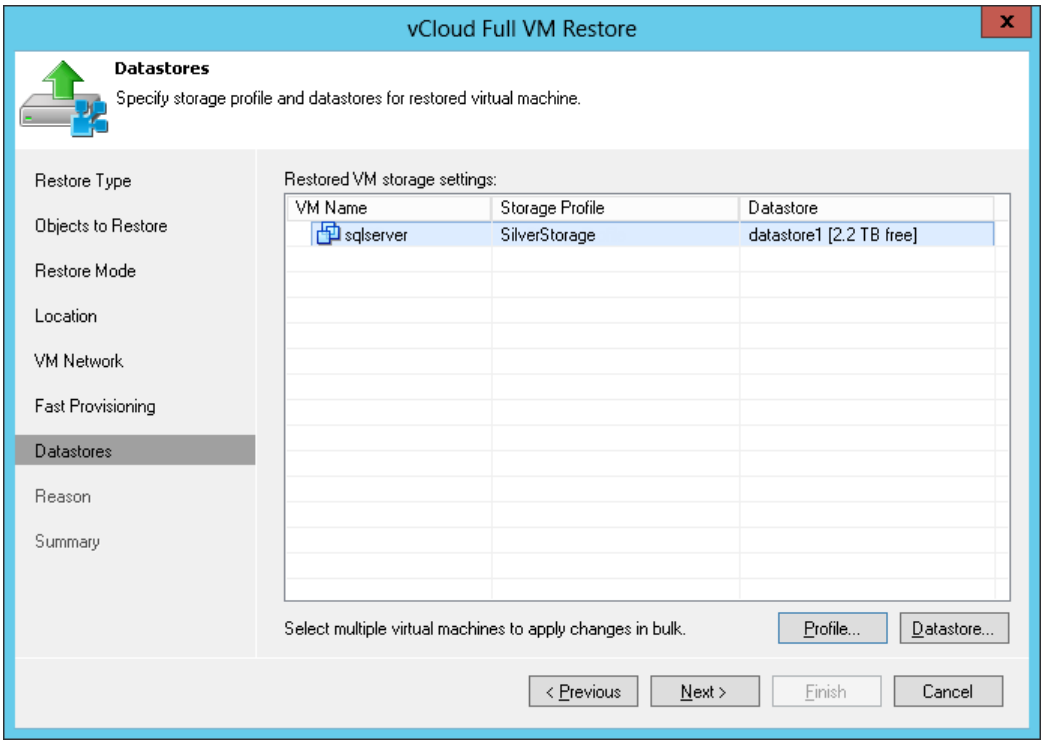

Step 9. Specify a Restore Reason

If necessary, enter the reason for performing restore of the selected VM. The information you provide will be saved in the session history so that you can reference it later.

Note Veeam Backup & Replication checks the lease term for the restored VM. In case the lease period has expired, the lease will be automatically updated.

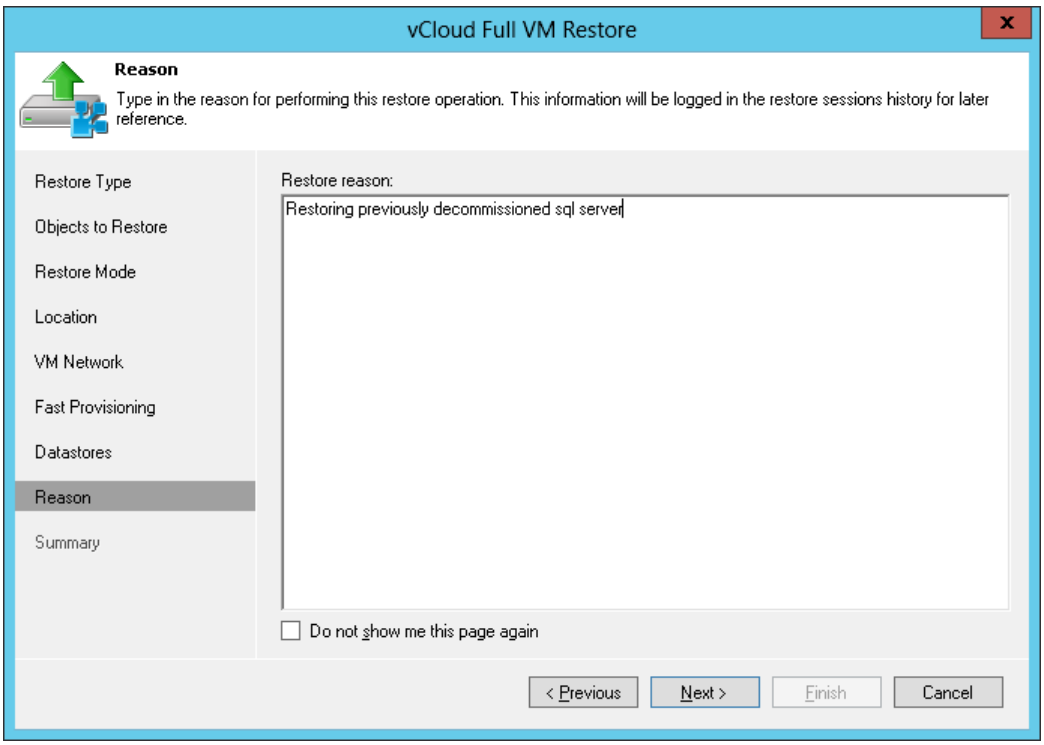

Step 10. Verify Recovery Settings

If you want to start the VM after the restore process is successfully complete, select the **Power on VM after restoring** check box under the list of restore points.

Check the specified settings for the full VM recovery and click **Finish**. Veeam Backup & Replication will restore selected VM to the specified vApp.

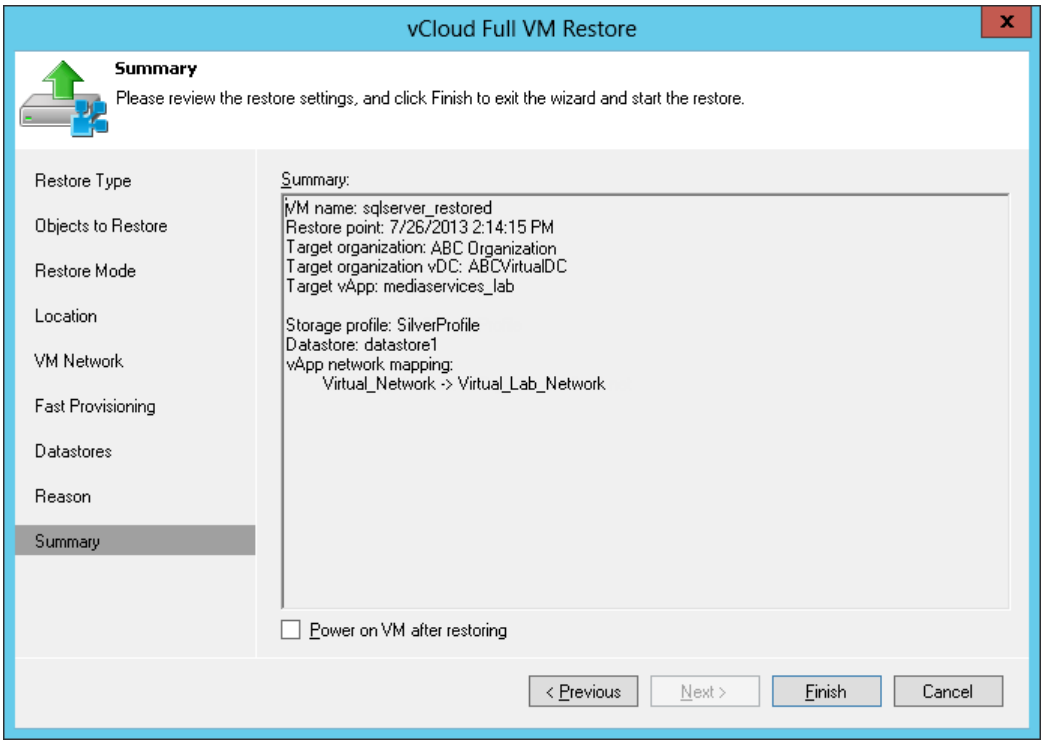

Restoring VMs to VMware vSphere Infrastructure

In addition to restoring a vCD VM to the vCloud Director hierarchy, you can also restore it to the VMware vSphere infrastructure. In this situation, Veeam Backup & Replication neglects the vApp metadata saved to the backup file and performs a regular full VM restore process: the VM is restored to the vCenter Server or ESX(i) host and is not registered in vCloud Director.

vCloud-specific features such as fast provisioning are not supported for such type of restore.

To launch the **Full VM Restore** wizard, do one of the following:

- On the **Home** tab, click **Restore** and select **VMware**. In the **Restore from backup** section, select **Entire VM (including registration)**.
- Open the **Backup & Replication** view and select the **Backups** node. In the working area, expand the necessary backup job, select the VM(s) you want to restore and click **Entire VM** > **Into virtual infrastructure** on the ribbon.
- Open the **Virtual Machines** view. In the inventory pane, expand the vCloud Director hierarchy and select the necessary vCenter Server. In the working area, right-click the VM you want to restore and select **Restore entire VM** > **Into vSphere infrastructure**.

Full VM restore of VMs managed by vCloud Director does not differ from full VM restore of regular VMs. To learn what settings you can specify for restore, see Performing Full VM Restore.

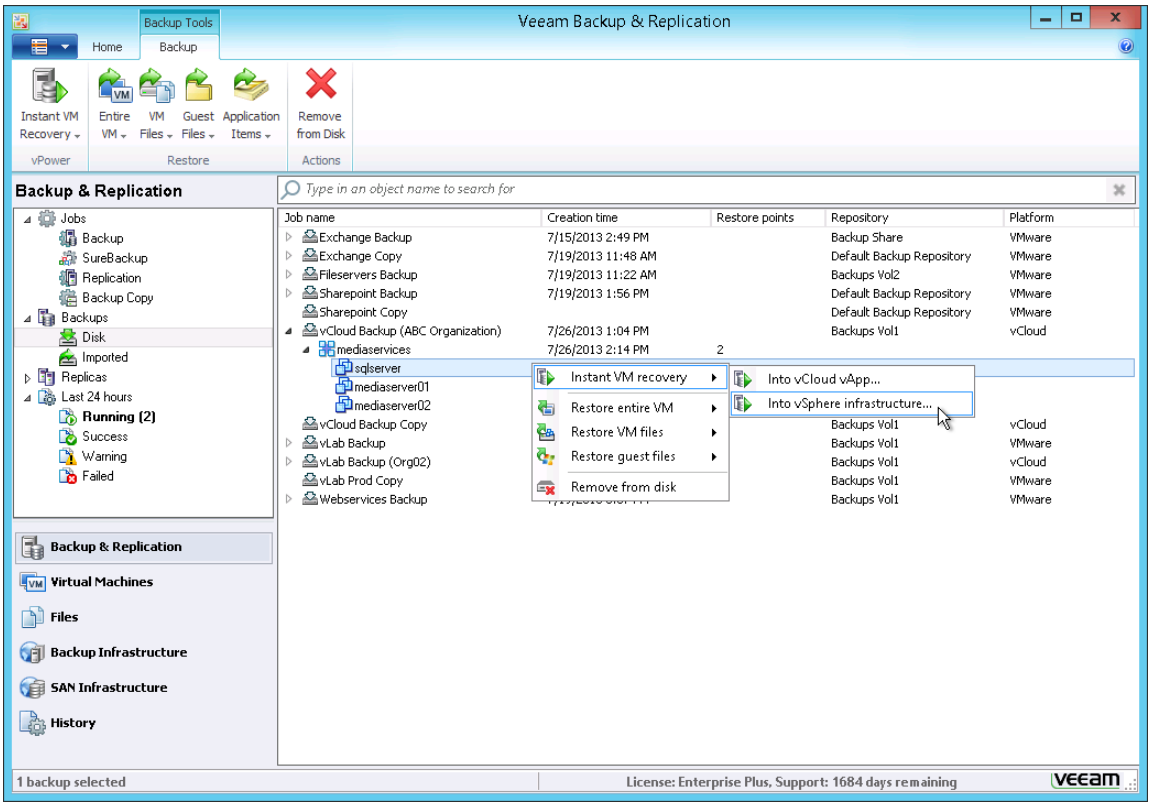

Restoring VM Files

The process of VM files restore for VMs managed by vCloud Director does not differ from that for regular VMs. To learn what settings you can specify for restore, see Restoring VM Files.

Restoring VM Hard Disks

The process of VM hard disks restore for VMs managed by vCloud Director does not differ from that for regular VMs. To learn what settings you can specify for restore, see Restoring VM Hard Disks.

Restoring VM Guest OS Files

The process of VM hard disks restore for VMs managed by vCloud Director does not differ from that for regular VMs. To learn what settings you can specify for restore, see Restoring VM Guest Files and Multi-OS File Level Recovery.

Working with HP SAN Storages

With Veeam Explorer for SAN Snapshots, you can create backups of VMware VMs leveraging HP SAN snapshots and restore VMware VM data directly from HP SAN snapshots. Veeam Backup & Replication supports the following HP SAN storage systems:

- HP StoreServ (3PAR)
- HP StoreVirtual (LeftHand/P4000 series)
- HP StoreVirtual VSA (Virtual Storage Appliance)

Setting Up Backup Infrastructure for HP SAN Snapshots

Before you start working with HP SAN snapshots in Veeam Backup & Replication, make sure that you have properly configured your backup infrastructure:

1. The Veeam Explorer for SAN Snapshots plug-in for the corresponding HP SAN storage is enabled.
2. You have properly configured a VMware backup proxy. The VMware backup proxy is required for volume rescan of HP StoreServ storage systems and backup from storage snapshots for HP StoreVirtual and HP StoreServ storage systems.
3. You have added to Veeam Backup & Replication a vCenter Server managing ESX(i) hosts with VMs whose disks are located on the HP SAN storage.
4. (For HP StoreServ only) You have checked that the HP 3PAR Web Services API server is enabled.
5. You have added the HP SAN storage system on which VM disks are located to the **SAN Infrastructure** view in Veeam Backup & Replication.

Note Work with HP SAN storage in Veeam Backup & Replication is performed in the **SAN Infrastructure** view. Right after the installation, the **SAN Infrastructure** view is not displayed. To display it, you should add at least one vCenter Server or ESX(i) host to Veeam Backup & Replication.

Enabling and Disabling Veeam Explorer for SAN Snapshots Plug-Ins

Starting from version 7.0, Veeam Explorer for SAN Snapshots is available as a plug-in.
Veeam Backup & Replication offers separate plug-ins for different types of HP SAN storage:

- Veeam Explorer for SAN Snapshots plug-in for HP StoreVirtual and HP StoreVirtual VSA
- Veeam Explorer for SAN Snapshots plug-in for HP StoreServ storage systems

Veeam Explorer for SAN Snapshots plug-ins are enabled after the installation of Veeam Backup & Replication.

In some cases, you may want to disable Veeam Explorer for SAN Snapshots plug-ins: for example, if you have moved to another SAN storage system.

To disable a plug-in:

1. Select **Manage Plugins** from the main menu of Veeam Backup & Replication.
2. In the displayed window, select the necessary plug-in and click **Disable**.

Note Before you disable a Veeam Explorer for SAN Snapshots plug-in, you must remove the corresponding HP SAN storage from the **SAN Infrastructure** view in Veeam Backup & Replication.

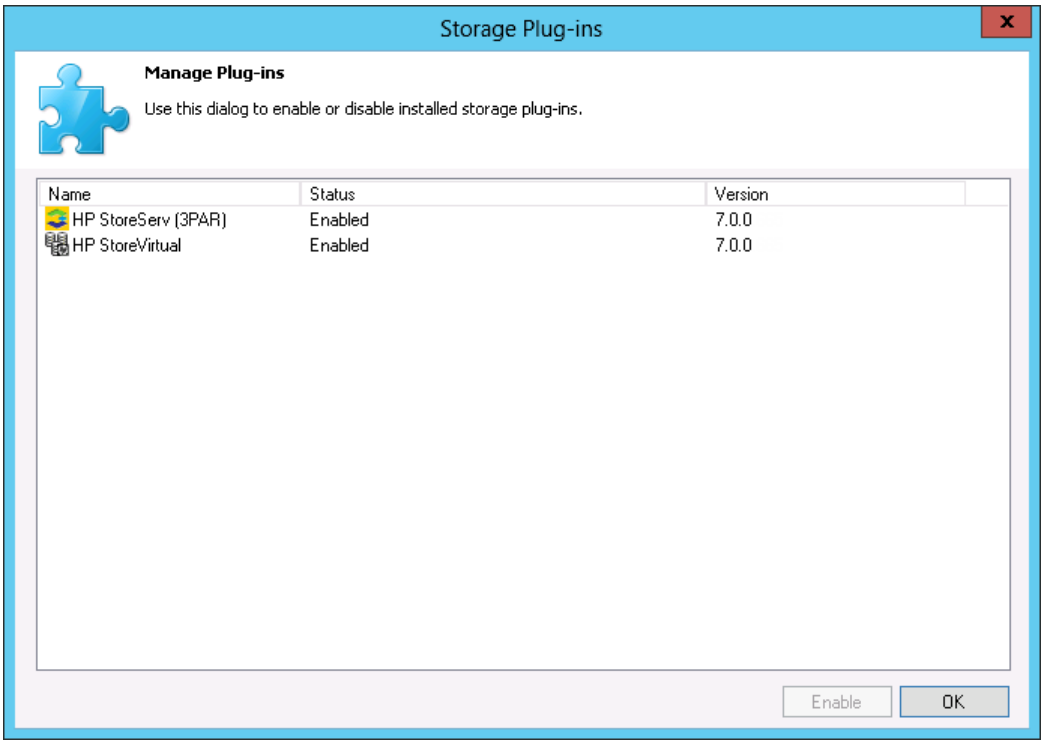

Configuring a VMware Backup Proxy for HP SAN Snapshots

For work with HP SAN snapshots, Veeam Backup & Replication uses a VMware backup proxy. The VMware backup proxy is used for two aims:

- Performing Backup from Storage Snapshots (for both HP StoreVirtual and HP StoreServ storage systems)
- Re-scanning VMFS on HP SAN volumes. Note that the backup proxy is required for rescan of HP StoreServ storage system only. Rescan of HP StoreVirtual SAN storage is performed without a backup proxy.

To perform backup from HP SAN storage snapshots and rescan VMFS on HP SAN volumes, Veeam Backup & Replication needs to read the content of the SAN volumes and volume snapshots. To do that, it uses a Windows-based machine — VMware backup proxy — as a helper. HP SAN volumes and volume snapshots are mounted to the VMware backup proxy as new volumes. As a result, Veeam Backup & Replication can access mounted HP SAN volumes and volume snapshots via the VMware backup proxy and can read the VM data from them.

The VMware backup proxy must meet the following requirements:

1. The role of the VMware backup proxy must be assigned to a Windows-based machine. This can be a dedicated machine or the Veeam backup server performing a role of the default VMware backup proxy.
2. The VMware backup proxy must have a direct access to the HP SAN storage via Fibre Channel or iSCSI. To connect to the HP SAN storage, the VMware Backup proxy must have an iSCSI initiator enabled or a Fibre Channel adapter installed.
   Note that the HP StoreServ SAN system supports only VMware backup proxies with a Fibre Channel connection. iSCSI SAN connections are not supported. Make sure that the Fibre Channel adapter is installed on the VMware backup proxy you plan to use for backup and storage rescan.

3. If you plan to use a Fibre Channel connection, Fibre Channel devices must be properly installed and zoned.
4. (For HP StoreServ SAN storage only) The backup proxy server must be added to the list of servers having access to the HP StoreServ SAN storage.

Important! If Backup from Storage Snapshots cannot be used due to incorrect infrastructure setup, Veeam Backup & Replication will not fail over to the regular processing mode and the backup or replication job will be finished with the *Failed* status.

Enabling the HP 3PAR Web Services API Server

To work with the HP StoreServ storage system, Veeam Backup & Replication uses HP 3PAR Web Services API server. When you add the HP StoreServ storage to the Veeam Backup & Replication console, Veeam Backup & Replication attempts to enable the server automatically.

Important! The HP 3PAR Web Services API is a part of the HP 3PAR Operating System starting from v3.1.2. Veeam Backup & Replication will not work with earlier versions of HP 3PAR Operating System.

In case Veeam Backup & Replication fails to enable the server automatically due to lack of privileges, you need to enable the server manually.

To verify that the HP 3PAR Web Services API server is enabled and running on the HP StoreServ storage system:

1. Log on to the Processor with administrator privileges:

```
#ssh <administrator account>@<SP IP Address>
```

2. View the current state of the Web Services API Server:

```
#showwsapi
-- -State- -HTTP_State-
HTTP_Port -HTTPS_State- HTTPS_Port -Version-
Enabled  Active Enabled    8008
Enabled    8080     1.1
```

3. If the Web Services API Server is disabled, start it:

```
#startwsapi
```

If the HTTP or HTTPS state is disabled, enable one of them:

```
#setwsapi -http enable
```

or

```
#setwsapi -https enable
```

Adding HP SAN Storage

To start working with Veeam Explorer for SAN Snapshots, you should add the HP SAN storage to Veeam Backup & Replication. Veeam Backup & Replication works with the following types of HP SAN storage:

- HP StoreVirtual and HP StoreVirtual VSA
- HP StoreServ Storage

When you add the HP SAN storage, Veeam Backup & Replication automatically performs indexing of the HP SAN storage content. As part of indexing, Veeam Backup & Replication performs the following operations:

1. Veeam Backup & Replication retrieves the information about the HP SAN storage topology from the HP SAN storage.
2. Veeam Backup & Replication retrieves information about datastores from the vCenter Server and matches SAN volumes to these datastores. Matching helps Veeam Backup & Replication understand what VMs are located on SAN volumes. Additionally, it helps
Veeam Backup & Replication get the information about VMs on SAN snapshots:
Veeam Backup & Replication assumes that VMs located on SAN volumes are also located on SAN snapshots.
3. Veeam Backup & Replication indexes SAN snapshots. As a result of snapshot indexing, Veeam Backup & Replication knows exactly what VMs are located on SAN snapshots. If for some reason Veeam Backup & Replication fails to index all SAN snapshots, it indexes the last SAN snapshot in the hierarchy and propagates indexing data to all other SAN snapshots.

After you add the HP storage to Veeam Backup & Replication, the HP storage topology appears in the **SAN Infrastructure** view of Veeam Backup & Replication. SAN snapshots also appear under the **SAN** node in the inventory pane of the **Backup & Replication** view.

Note Before you add the HP SAN storage to Veeam Backup & Replication, make sure that you have added a vCenter Server managing ESX(i) hosts with VMs whose disks are located on this HP SAN storage. In the opposite case, Veeam Backup & Replication will not be able to map SAN volumes to VMware datastores and will not display these VMs in the **SAN Infrastructure** view. If you have first added the HP SAN storage and then added the vCenter Server, perform storage rescan to map SAN volumes to VMware datastores.
For HP 3PAR StoreServ Storage, you also need to configure a VMware backup proxy.

Adding HP StoreVirtual SAN Storage

To add the HP StoreVirual SAN storage to Veeam Backup & Replication, follow the next steps:

Step 1. Launch the New HP StoreVirtual Storage Wizard

To launch the wizard, do one of the following:

- Open the **SAN Infrastructure** view, right-click the **SAN Infrastructure** node in the inventory pane and select **Add Storage**. In the displayed window, click **HP StoreVirtual**.
- Open the **SAN Infrastructure** view, select the **HP StoreVirtual** node in the inventory pane and click **Add Storage** on the ribbon.
- Open the **SAN Infrastructure** view, right-click the **HP StoreVirtual** node in the inventory pane and select **Add Storage**.
- Open the **SAN Infrastructure** view, click the **HP StoreVirtual** node in the inventory pane, right-click anywhere in the working area and select **Add Storage**.

Step 2. Specify the HP Storage Name or Address

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

1. Click **Browse** next to the **Management server DNS name or IP address** field and select the necessary HP storage management group. Alternatively, you can explicitly specify a DNS name or IP address of the HP storage management server or storage cluster by typing it into the field.
2. In the **Description** field, provide a description for future reference. The default description contains information about the user who added the HP management group as well as the date and time when the HP management group was added.

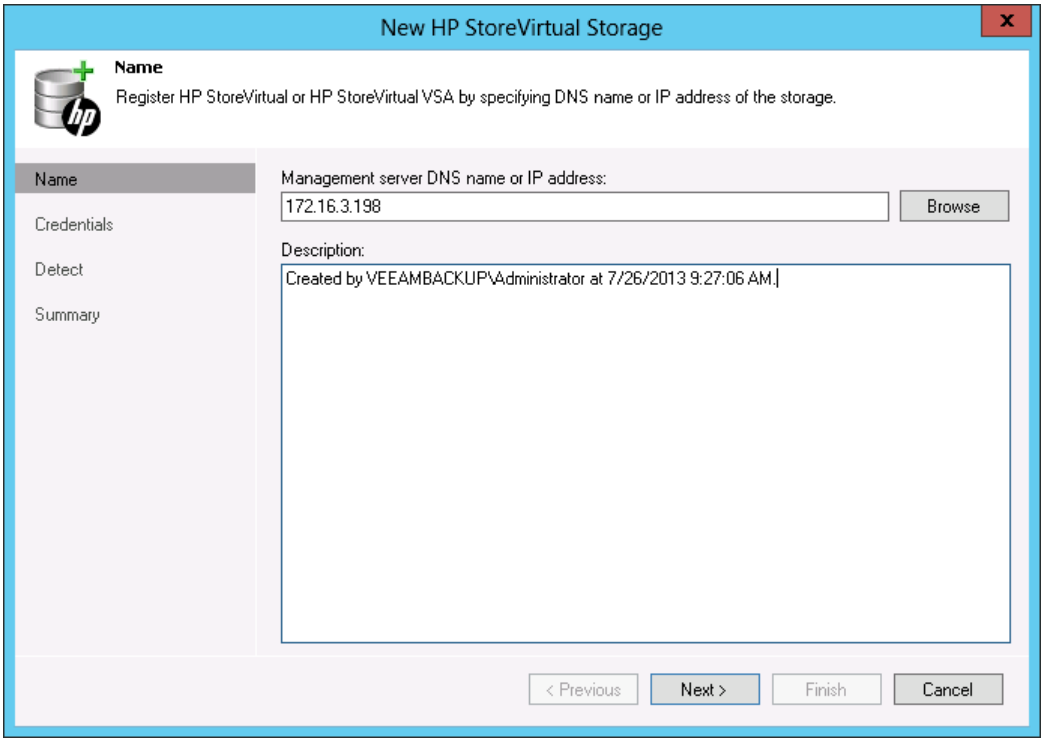

Step 3. Specify Credentials

At the **Credentials** step of the wizard, specify credentials for the user account with administrator privileges on the HP management group you are adding.

From the **Credentials** list, select credentials to connect to the specified HP management group. If you have not set up the necessary credentials beforehand, click the **Manage accounts** link at the bottom of the list to add the necessary credentials. To learn more, see Managing Credentials.

Important! The user name and password values for credentials are case-sensitive.

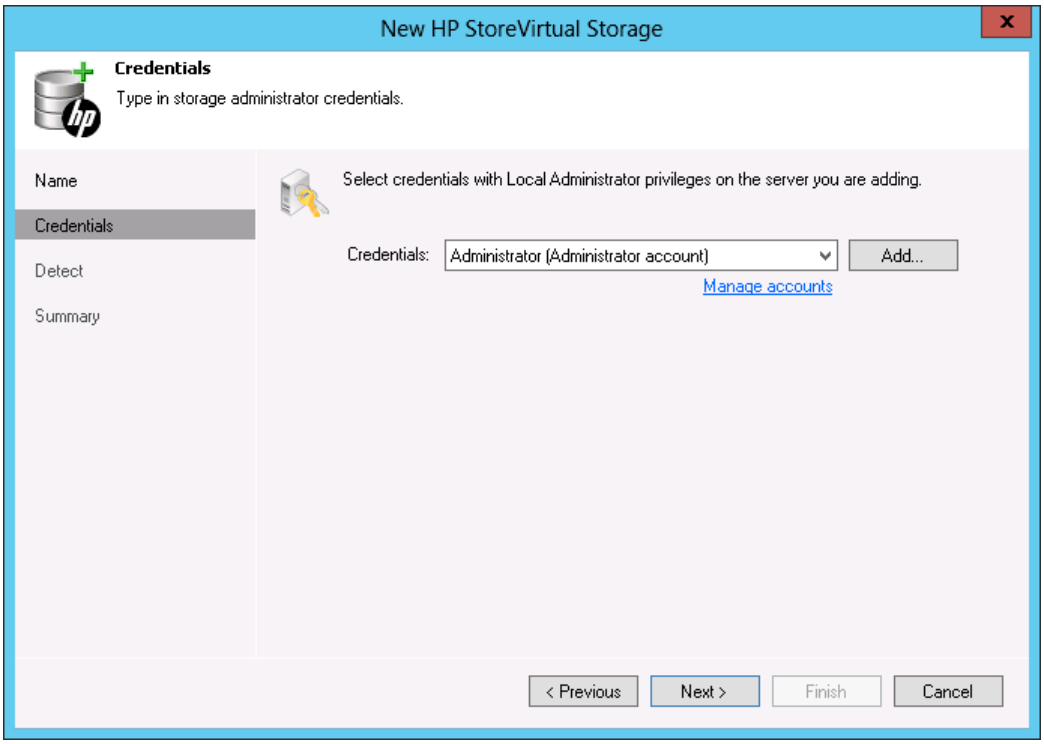

Step 4. Scan the HP Storage

At the **Detect** step of the wizard, Veeam Backup & Replication will re-scan the HP storage in the real-time mode. During the re-scan process, Veeam Backup & Replication will retrieve information about datastores from VMware servers and match HP SAN volumes to the datastores.

The rescan process may take a while. Wait for the operation to complete and click **Next**.

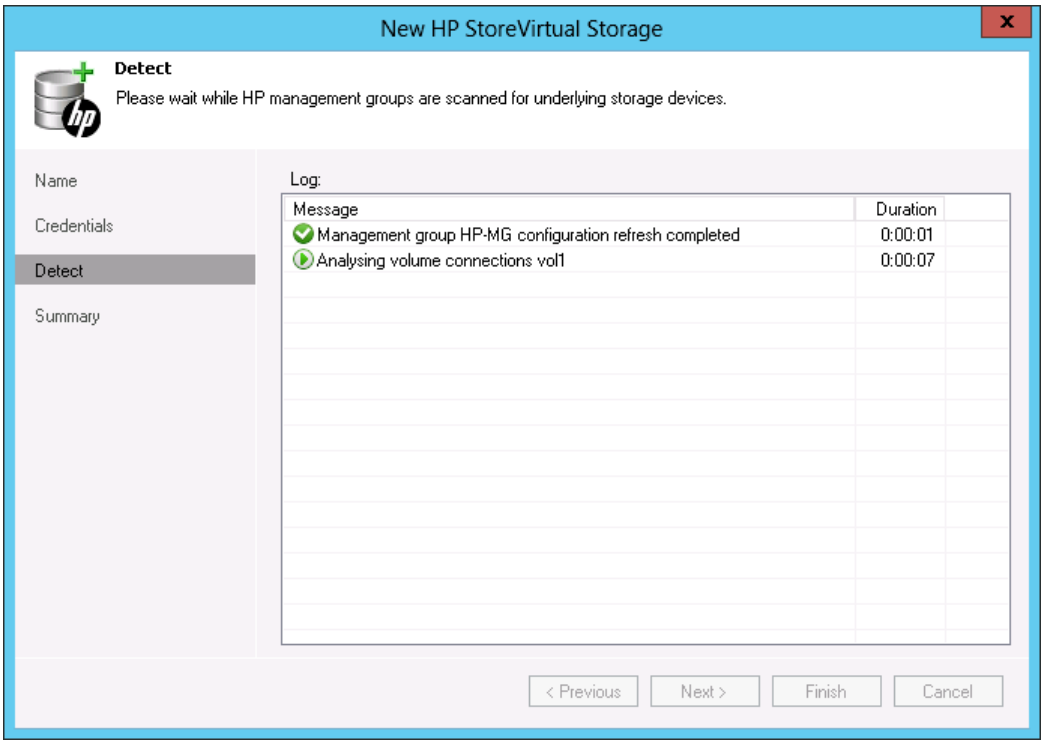

Step 5. Finish Working with the Wizard

At the **Summary** step of the wizard, review the summary information and click **Finish**. After you click **Finish**, Veeam Backup & Replication will perform indexing of HP SAN snapshots.

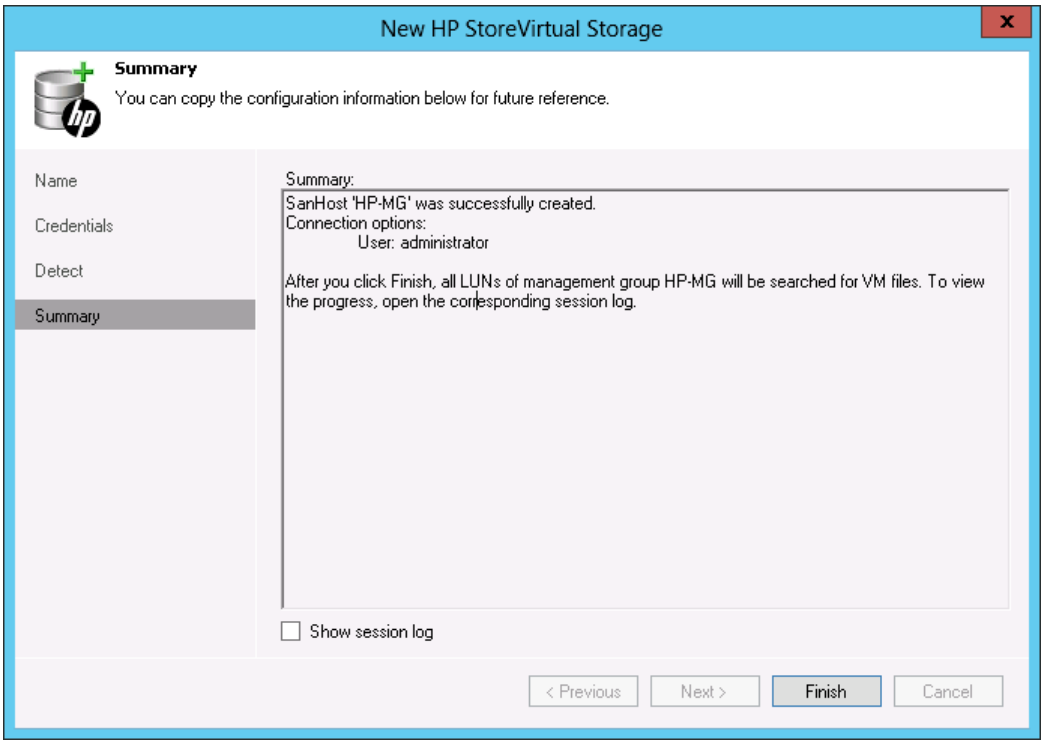

Adding HP StoreServ Storage

To add the HP StoreServ storage system to Veeam Backup & Replication, follow the next steps:

Important! The HP StoreServ storage system you plan to add must meet the following requirements:

- The license for the HP StoreServ storage system must support 3PAR Virtual Copy.
- The HP StoreServ storage system must have the HP 3PAR Web Services API Server enabled. The HP 3PAR Web Services API Server is a part of the HP 3PAR Operating System starting from v3.1.2. For this reason, earlier versions of the HP 3PAR Operating System are not supported by Veeam Backup & Replication.
  See also: Enabling the HP 3PAR Web Services API Server.

Step 1. Launch the New HP StoreServ Storage Wizard

To launch the wizard, do one of the following:

- Open the **SAN Infrastructure** view, right-click the **SAN Infrastructure** node in the inventory pane and select **Add Storage**. In the displayed window, click **HP StoreServ**.
- Open the **SAN Infrastructure** view, select the **HP StoreServ** node in the inventory pane and click **Add Storage** on the ribbon.
- Open the **SAN Infrastructure** view, right-click the node in the inventory pane and select **Add Storage**.
- Open the **SAN Infrastructure** view, click the **HP StoreServ** node in the inventory pane, right-click anywhere in the working area and select **Add Storage**.

Step 2. Specify the HP 3PAR Web Services API Address

Veeam Backup & Replication works with the HP StoreServ storage system via the HP 3PAR Web Services API. The HP 3PAR Web Services API delivers a programming interface for performing storage management tasks with HP StoreServ storage systems.

At the **Name** step of the wizard, provide the following information for the HP 3PAR Web Services API:

1. In the **DNS name or IP address** field, enter a full DNS name or IP address of the HP 3PAR Web Services API Server.
2. In the **URL** field, enter the URL of the HP 3PAR Web Services API. By default, Veeam Backup & Replication uses the following URL: https://< *websapiserver*>:*8080*, where <*websapiserver*> is the name or IP address of the HP 3PAR Web Services API Server that you have specified in the field above.
3. In the **Description** field, provide a description for future reference. The default description contains information about the user who added the server as well as the date and time when the server was added.

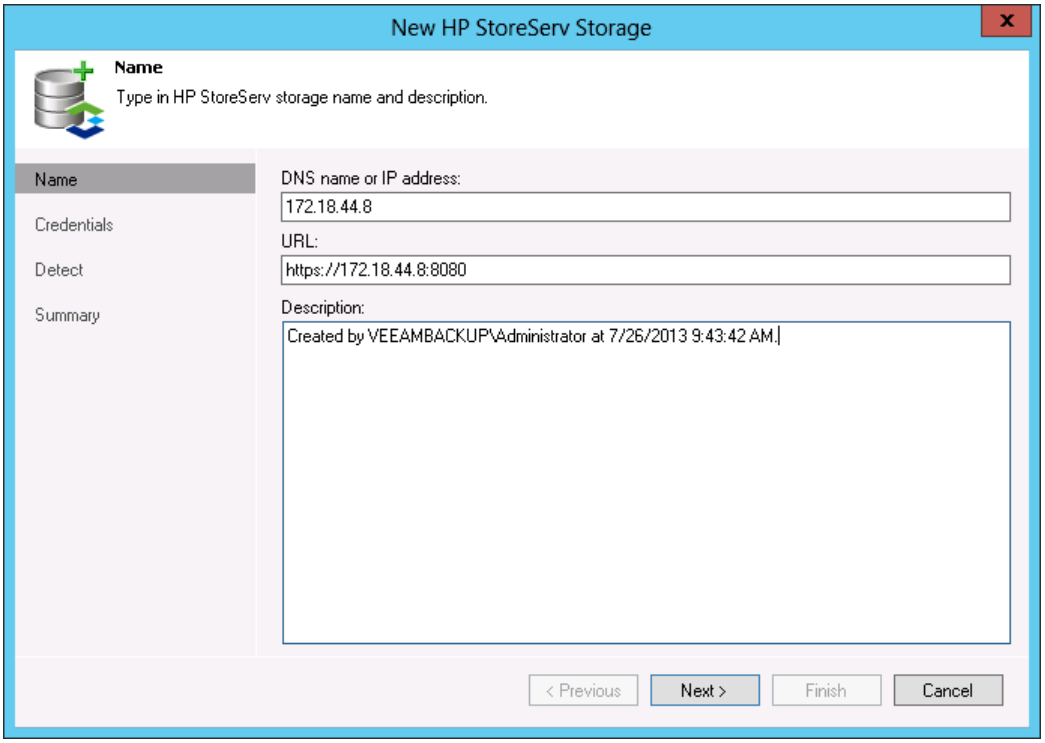

Step 3. Specify Credentials

At the **Credentials** step of the wizard, specify credentials for the user account with administrator privileges on the HP 3PAR Web Services API Server you are adding.

From the **Credentials** list, select credentials to connect to the HP 3PAR Web Services API Server.If you have not set up the necessary credentials beforehand, click the **Manage account** link at the bottom of the list to add the necessary credentials. To learn more, see Managing Credentials.

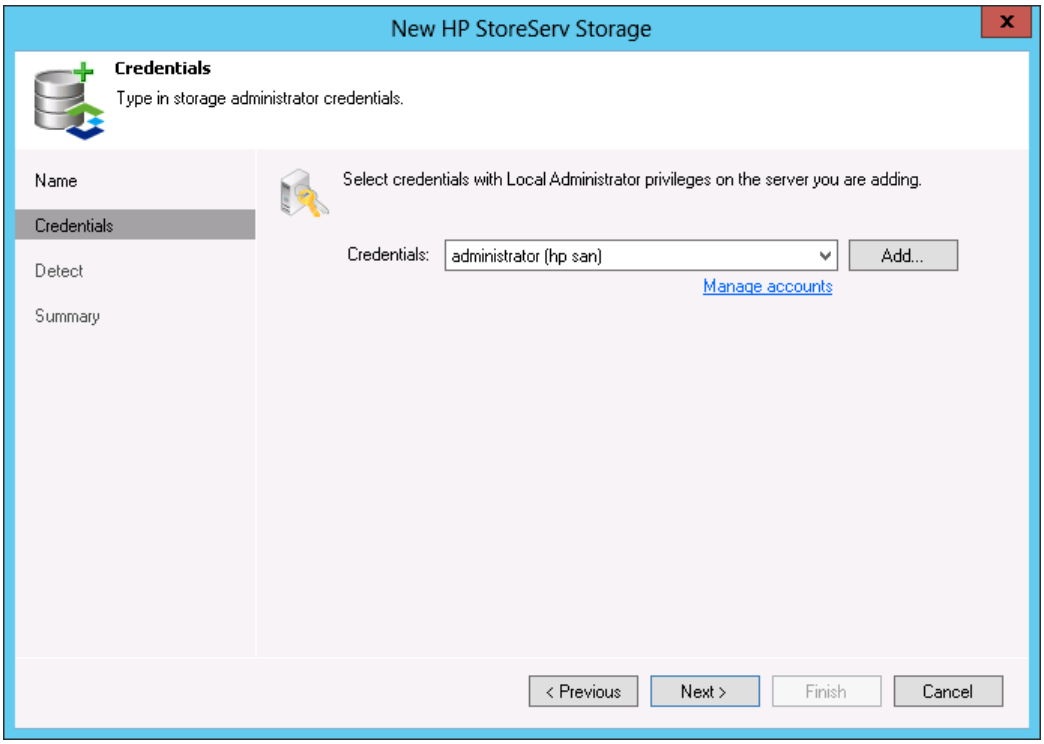

Step 4. Scan the HP StoreServ SAN Storage

At the **Detect** step of the wizard, Veeam Backup & Replication will re-scan the HP StoreServ storage system in the real-time mode. During the re-scan process, Veeam Backup & Replication will retrieve information about datastores from VMware servers and match SAN volumes to the datastores.

The rescan process may take a while. Wait for the operation to complete and click **Next**.

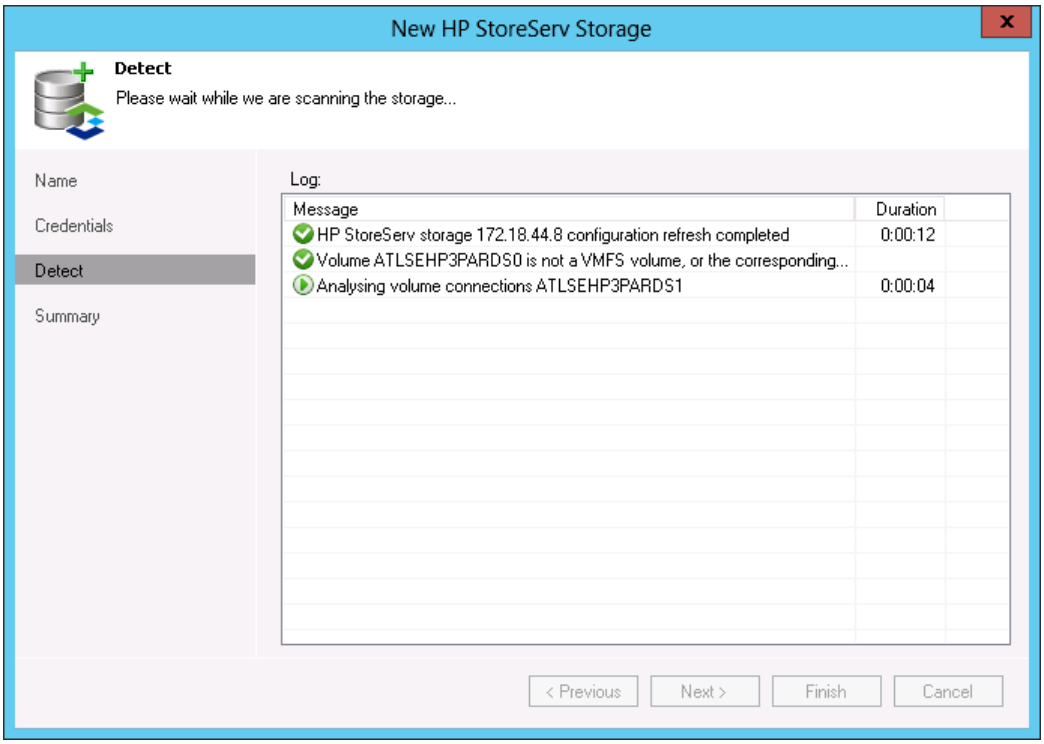

Step 5. Finish Working with the Wizard

At the **Summary** step of the wizard, review the summary information and click **Finish**. After you click **Finish**, Veeam Backup & Replication will perform indexing of HP SAN snapshots.

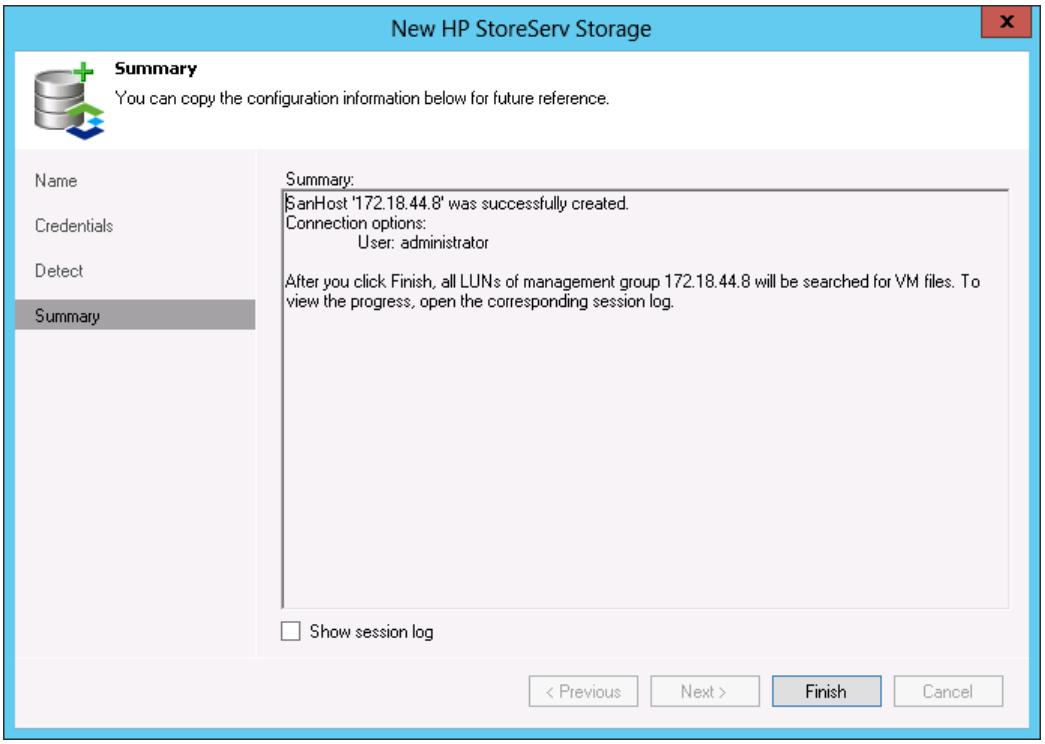

Working with HP SAN Snapshots

After you have properly configured your backup infrastructure and added the HP SAN storage to Veeam Backup & Replication, you can perform the following operations:

Re-Scanning HP SAN Storages

Veeam Backup & Replication periodically runs the *Storage discovery* command against the added HP storage systems to update the information about the storage topology: discover new volumes and snapshots, check for deleted volumes and snapshots and so on. The command is started automatically every 3 minutes.

If necessary, you can start the *Storage discovery* command manually. Manual storage discovery may be necessary, for example, if you have created a new volume snapshot and want to display it in the **SAN Infrastructure** view. You can perform re-scan of the whole HP storage system, HP management group, 3PAR Web Services API server or a specific HP SAN volume.

1. Open the **SAN Infrastructure** view.
2. In the inventory pane, select the **HP StoreServ** or **HP StoreVirtual** node.
3. Expand the HP SAN storage system tree.
4. In the working area, select the necessary node: HP management group, SAN volume and so on.
5. Click **Rescan** on the ribbon. Alternatively, you can right-click the necessary node in the working area and select **Rescan**.

Tip To quickly find the necessary object, use the search field at the top of the window: enter the object name or a part of it and press **[ENTER]**.

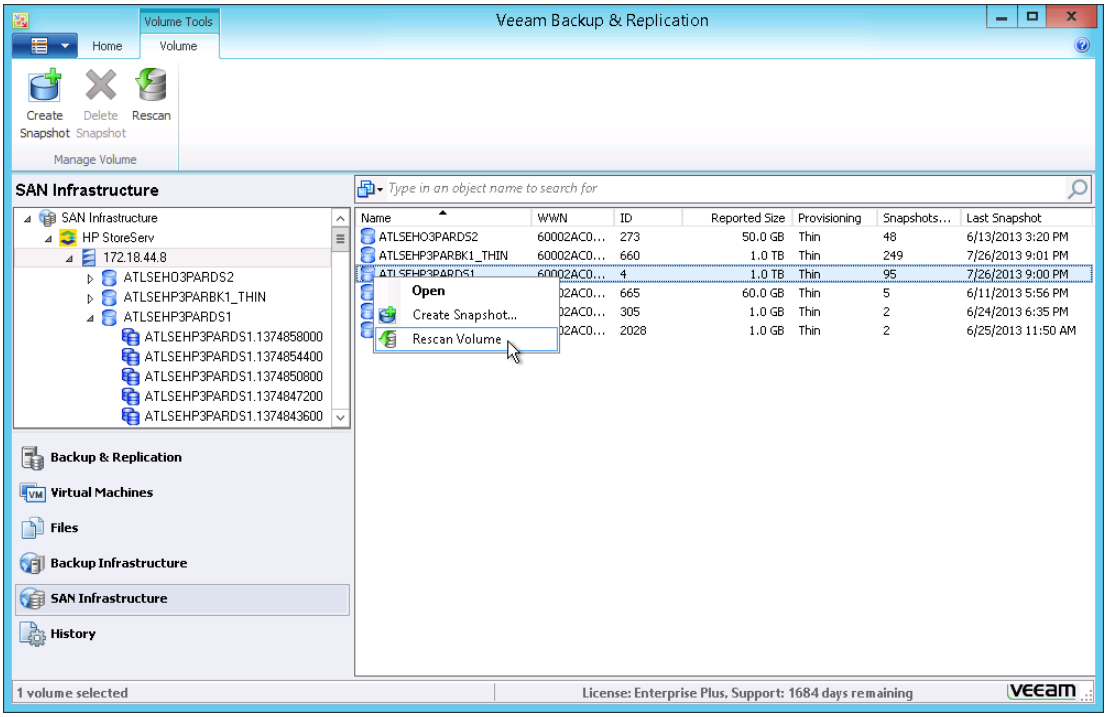

Creating and Deleting Snapshots

Veeam Explorer for SAN Snapshots allows you to create and delete snapshots of SAN volumes directly from the Veeam Backup & Replication interface.

To create a volume snapshot:

1. Open the **SAN Infrastructure** view.
2. In the inventory pane, expand the necessary HP storage system tree.
3. In the working area, right-click the necessary volume and select **Create Snapshot**.
4. In the **New SAN Snapshot** window, specify a name for the created snapshot and provide a snapshot description.
5. (*For HP P4000 storage only*) If you want to quiesce VMs on the volume, select the **Create application-managed snapshot** check box. An application-managed snapshot can be compared to a transactionally consistent backup. When the snapshot is taken, the SAN triggers a command to the vCenter Server. The vCenter Server quiesces VMs on the volume using VMware Tools to bring the VM data to a consistent state. If this option is not selected, Veeam Backup & Replication will create a point-in-time snapshot.

Important! You can create application-managed snapshots only if you have installed and properly configured the Application AwareSnapshot Manager with the HP P4000 SAN solution. If the Application AwareSnapshot Manager is not installed, Veeam Backup & Replication will report an error and the snapshot will not be created at all. To learn more, see the HP P4000 Application-Aware Snapshot Manager Deployment Guide at http://h10032.www1.hp.com/ctg/Manual/c03037557.pdf.

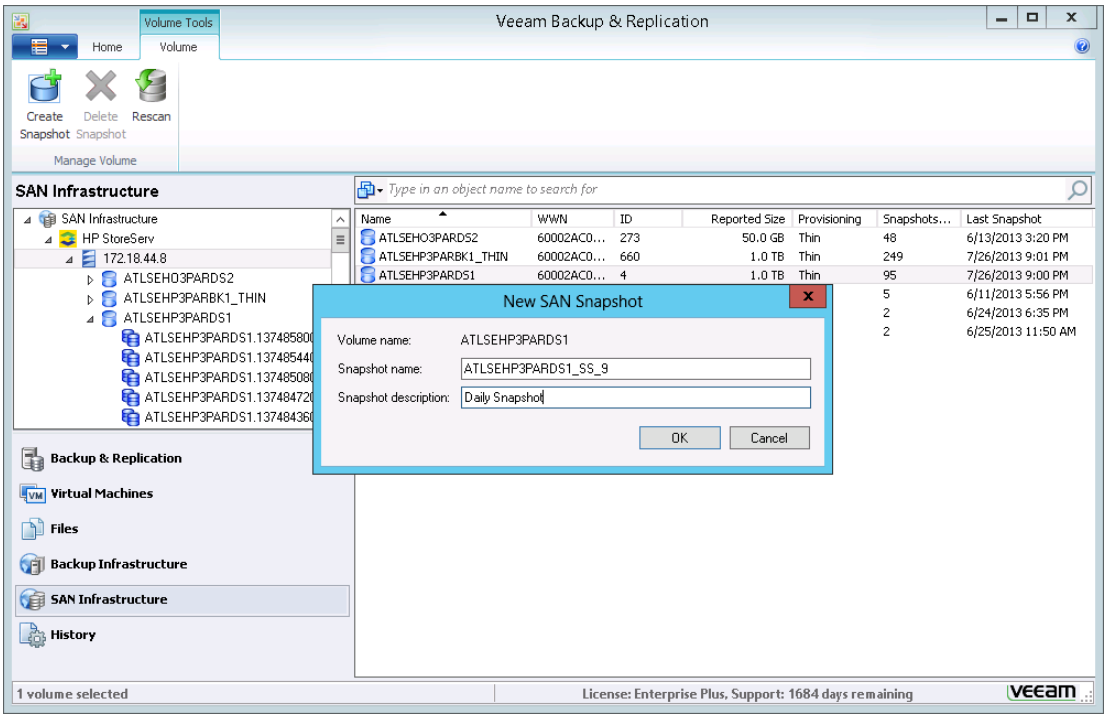

To delete a volume snapshot:

1. Open the **SAN Infrastructure** view.
2. In the inventory pane, expand the necessary HP storage system tree.
3. Right-click the necessary snapshot and select **Delete Snapshot**.

Performing Backup from Storage Snapshots

To perform backup from storage snapshots, you need to perform the following steps:

1. Configure a VMware backup proxy to which SAN snapshots will be mounted.
2. In the backup or replication job settings, select the configured VMware backup proxy to be used for data transport. You can assign the VMware backup proxy explicitly or choose the automatic mode of proxy selection.
3. Make sure that the **Use storage snapshots** option is enabled in the backup or replication job settings. By default, the **Use storage snapshots** option is selected in backup and replication job settings and Veeam Backup & Replication uses it when applicable.

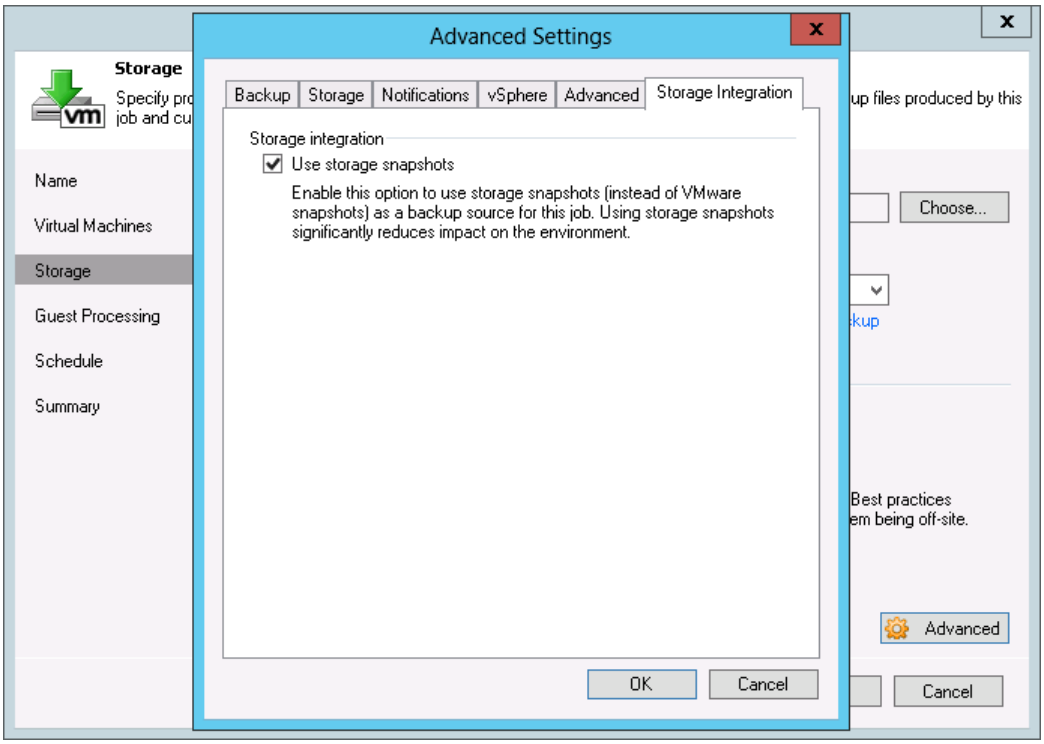

The **Use storage snapshots** option is applied only to those VMs whose disks are located on the HP SAN storage.

- If the job includes a number of VMs whose disks are located on different types of storage, Veeam Backup & Replication will apply the **Use storage snapshots** option only to VMs whose disks are located on the HP SAN storage.
- If a VM has several disks, some on the HP SAN and some on the other type of storage, Veeam Backup & Replication will not apply the **Use storage snapshots** option to such VM.

If the Backup from Storage Snapshots technology cannot be used due to incorrect infrastructure setup, Veeam Backup & Replication will not fail over to the regular processing mode and the backup or replication job will be finished with the *Failed* status. To successfully perform backup or replication in such situation, you will need to properly configure the backup infrastructure or disable the **Use storage snapshots** option in the job properties.

Note The Backup from Storage Snapshots functionality is available only in the Enterprise Plus Edition of Veeam Backup & Replication.

Restoring VM Data from SAN Snapshots

If important data accidentally gets lost or corrupted, you can use Veeam Backup & Replication to restore VM data from HP SAN snapshots. The restore process is similar to restore from image-level backups of VMs and is performed in the Veeam Backup & Replication interface.

Important! If you perform restore from HP StoreServ storage snapshots, for mounting SAN snapshots you must select an ESX(i) host that meets the following requirements:

- The ESX(i) host must have access to HP StoreServ storage over a Fibre Channel connection.
- The ESX(i) host must be added to the list of servers having access to the HP StoreServ storage snapshots.

Veeam Backup & Replication offers the following restore options for HP SAN snapshots:

- Performing Instant VM Recovery
- Restoring VM guest OS files (Windows)
- Restoring VM guest OS files (Linux, Unix and other)
- Restoring individual objects from Microsoft Exchange
- Restoring individual objects from Microsoft SharePoint

Performing Instant VM Recovery

With Veeam Backup & Replication, you can immediately start a VM that resides on a volume snapshot without prior de-staging and intermediate restores. Instant VM Recovery accelerates the VM restore, allowing you to improve RTOs and decrease downtime of production VMs.

To perform instant recovery of a VM on the SAN snapshot, follow the next steps:

Step 1. Launch the Instant Recovery Wizard

To launch the **Instant Recovery** wizard, do one of the following:

- On the **Home** tab, click **Restore** and select **VMware**. In the **Restore from backup** section, select **Instant VM recovery**.
- Open the **SAN Infrastructure** view. In the inventory pane, expand the HP storage system tree and select the necessary volume snapshot. In the working area, right-click the VM you want to restore and select **Instant VM recovery**. In this case, you will pass immediately to step 3.
- Open the **Backup & Replication** view and select the **SAN** node in the inventory pane. In the working area, expand the necessary volume, select the VMs you want to restore and click **Instant VM Recovery** on the ribbon. In this case, you will pass immediately to step 3.
- Open the **Backup & Replication** view and select the **SAN** node in the inventory pane. In the working area, expand the necessary volume, right-click the VMs you want to restore and select **Instant VM recovery**. In this case, you will pass immediately to step 3.

Tip To quickly find the necessary VM, use the search field at the top of the window: enter the VM name or a part of it and press **[ENTER]**.

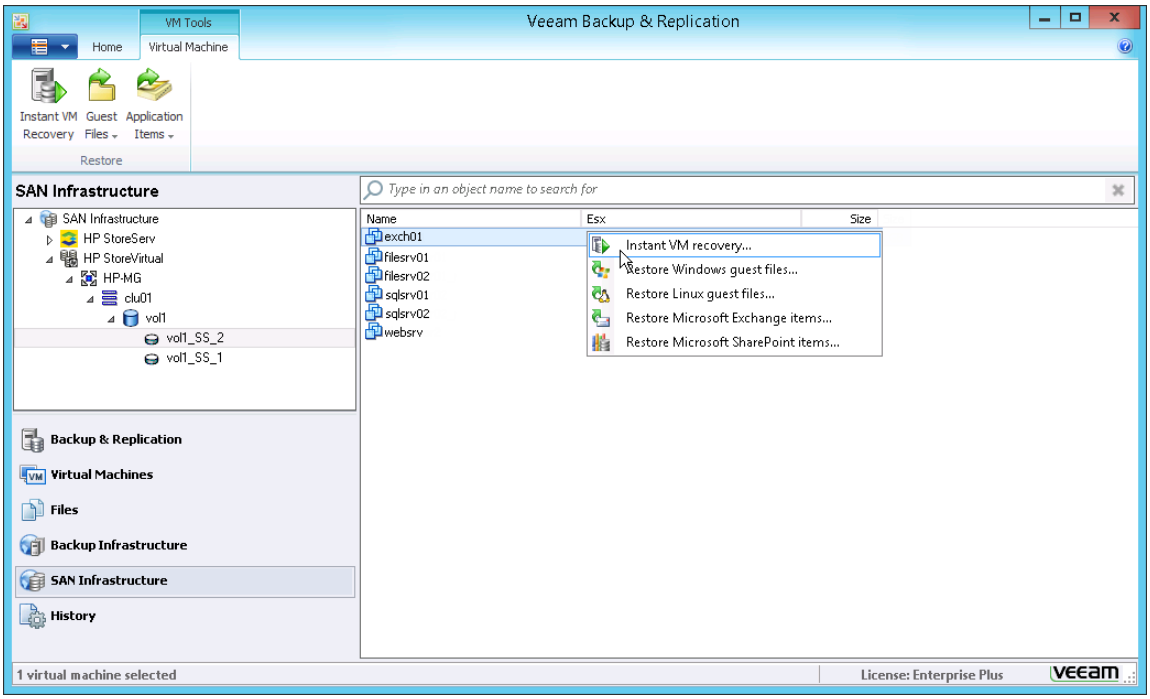

Step 2. Select a Virtual Machine

At the **Virtual Machine** step of the wizard, expand the necessary volume snapshot and select the VM you want to restore.

Tip To quickly find the necessary VM, use the search field at the bottom of the window: enter the VM name or a part of it and press **[ENTER]**.

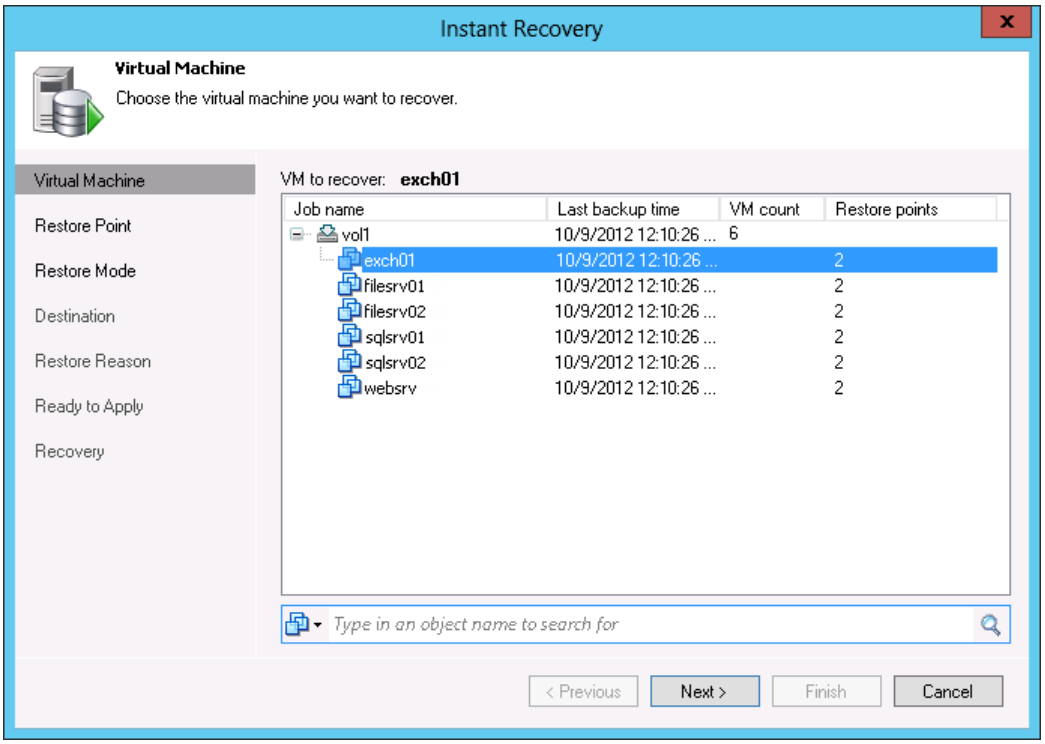

Step 3. Select a Restore Point

At the **Restore Point** step of the wizard, select the necessary restore point for the VM. Every snapshot of a volume acts as an independent restore point.

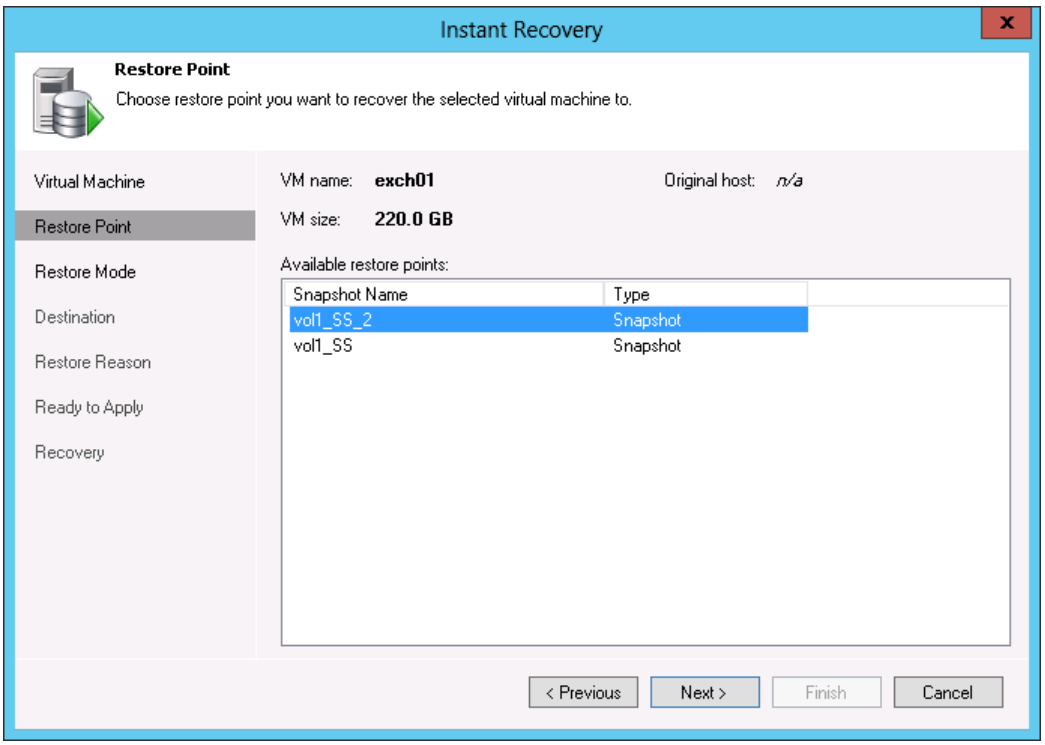

Step 4. Select a Restore Mode

At the **Restore Mode** step of the wizard, choose the necessary restore mode:

- Select **Restore to the original location** if you want to restore the VM with its initial settings and to its original location. If this option is selected, you will pass directly to step 6.
- Select **Restore to a new location, or with different settings** if you want to restore the VM to a different location and/or with different settings (such as VM location and VM name). If this option is selected, the **Instant Recovery** wizard will include an additional step for customizing VM settings.

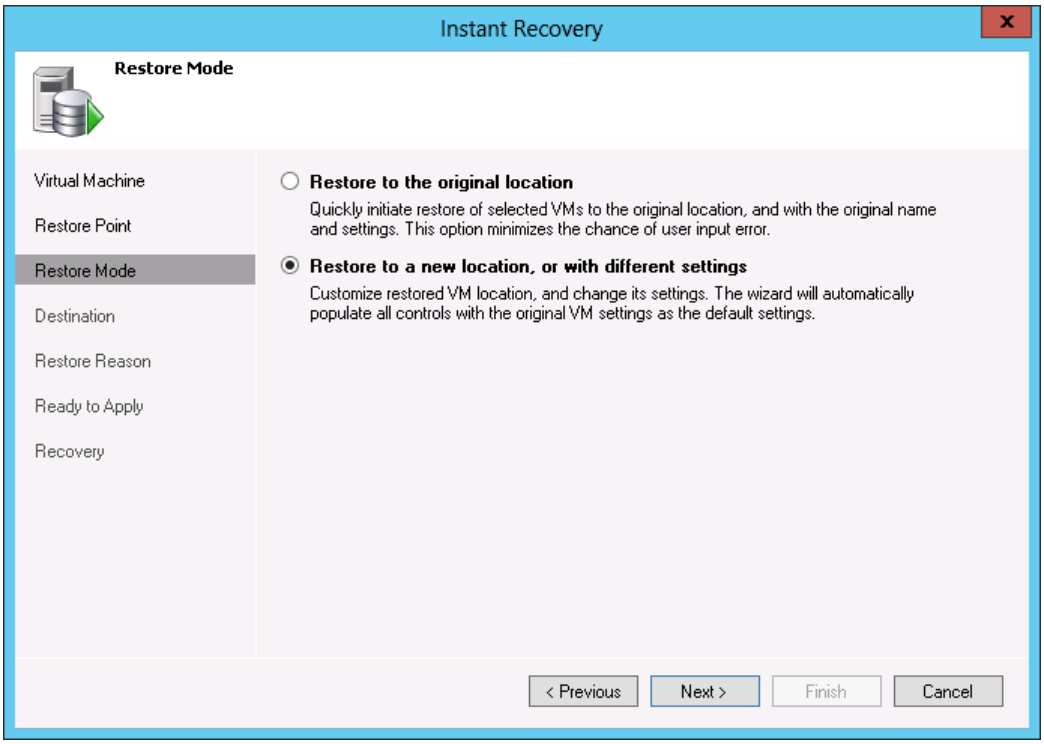

Step 5. Select Destination for the Recovered VM

The **Destination** step of the wizard is available if you have chosen to change the location and settings of the restored VM.

Select the destination for the recovered VM:

1. In the **Host** field, specify a host on which the VM should run. Veeam Backup & Replication will create a clone/virtual copy of the SAN snapshot, mount it to the selected ESX(i) host and start the VM on this ESX(i) host.
2. In the **VM folder** field, specify the folder to which the restored VM should be placed.
3. In the **Restored VM name** field, enter a name under which the VM should be restored and registered. By default, Veeam Backup & Replication uses the original name of the VM. If you are restoring the VM to the same ESX(i) host or the same datacenter where the original VM is registered and the original VM still resides there, it is recommended that you change the VM name to avoid conflicts.
4. In the **Resource pool** list, select a resource pool to which the VM should be recovered.

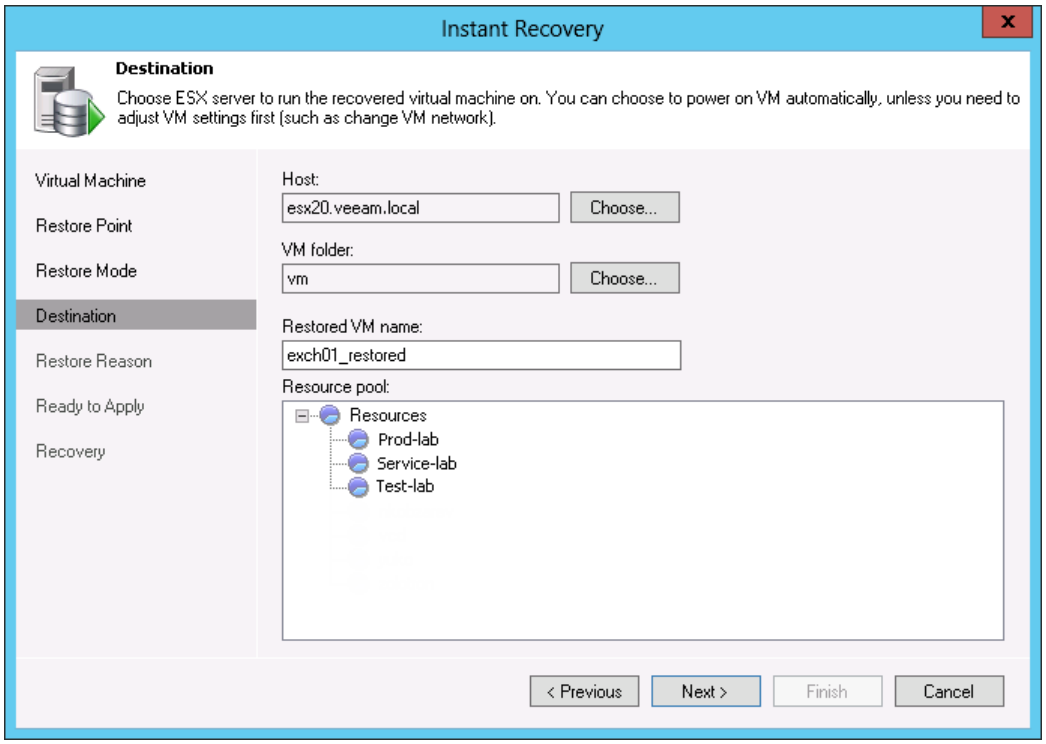

Step 6. Specify a Restore Reason

At the **Restore Reason** step of the wizard, enter the reason for performing instant restore of the selected VM if necessary. The information you provide will be saved in the session history so that you can reference it later.

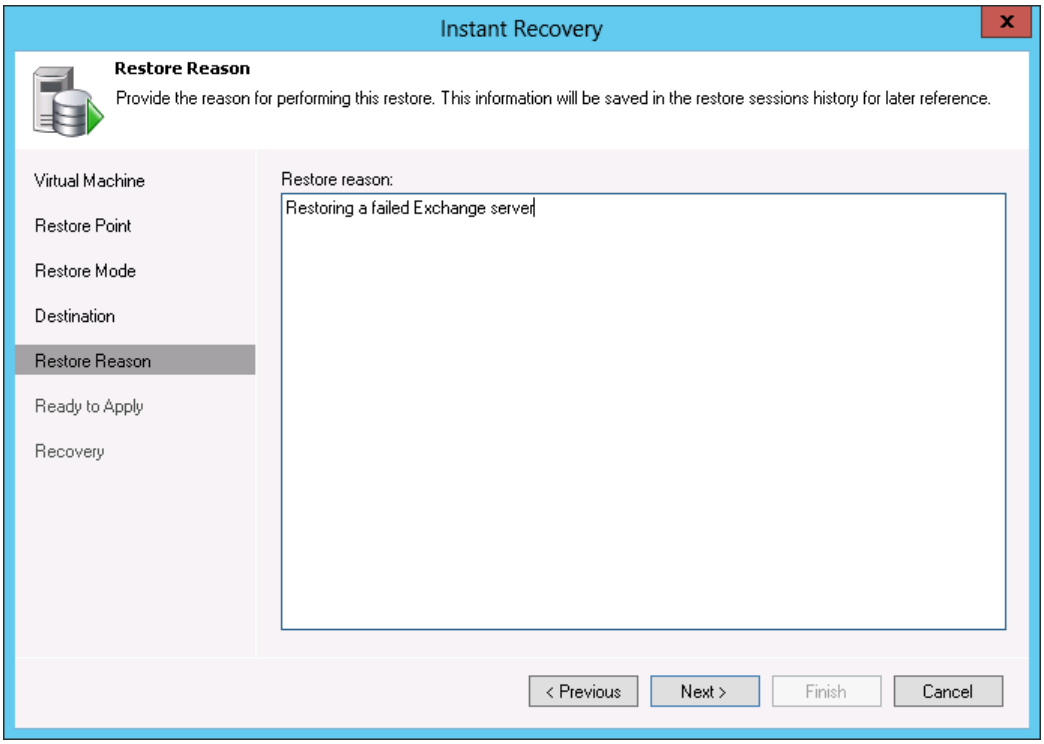

Step 7. Verify Instant VM Recovery Settings

At the **Ready to Apply** step of the wizard, specify additional restore settings:

1. If you are recovering a production VM that has failed and want to restore it with initial network settings, select the **Connect VM to network** check box. If you are recovering a VM for testing disaster recovery while the initial VM is still running, leave this check box not selected. Before you power on a VM, you will have to manually change VM network configuration: disconnect it from the production network and connect it to an isolated non-production network to avoid conflicts.
2. To start a VM immediately after recovery, select the **Power on VM automatically** check box. If you are recovering the VM to the production network, make sure that the initial VM is powered off to avoid conflicts.
3. Check the settings for instant recovery and click **Next**. Veeam Backup & Replication will restore the VM on the selected ESX(i) host.

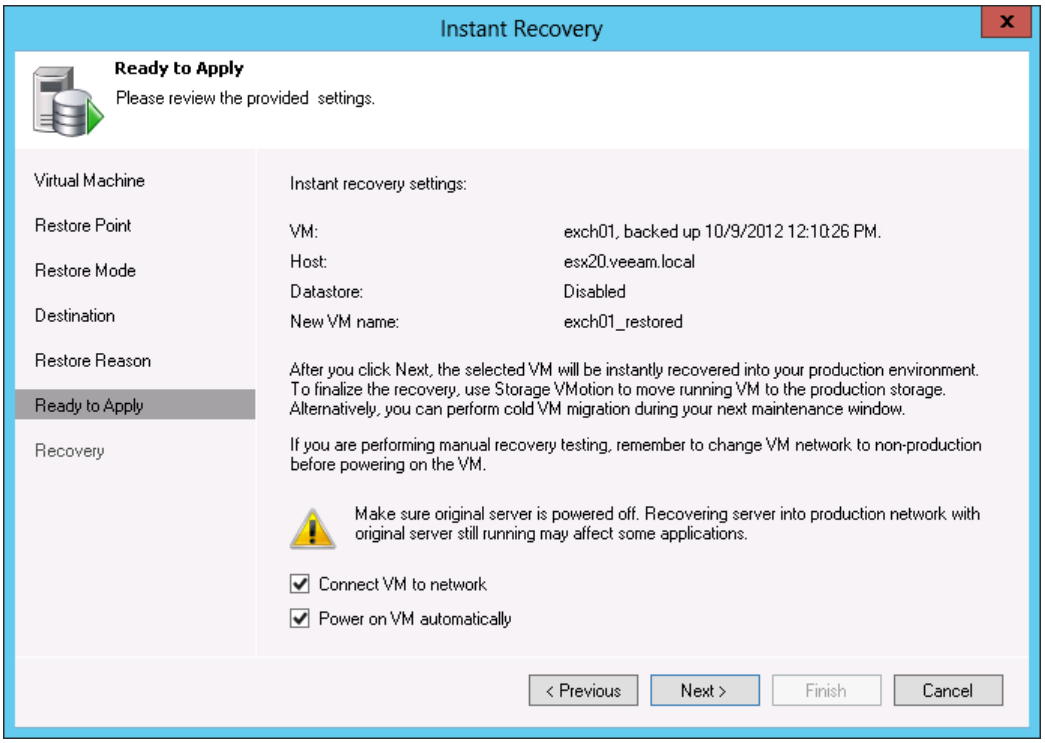

Step 8. Finalize Instant VM Recovery

All currently running Instant VM Recovery sessions are displayed in the **Backup & Replication** view, under the **Restore** > **Instant VM** recovery node.

To check the progress of the Instant VM Recovery session and view session details:

1. Open the **Backup & Replication** view.
2. Select the **Instant Recovery** node.
3. Right-click the necessary session in the working area and select **Properties**. You can also double-click the necessary session in the list.

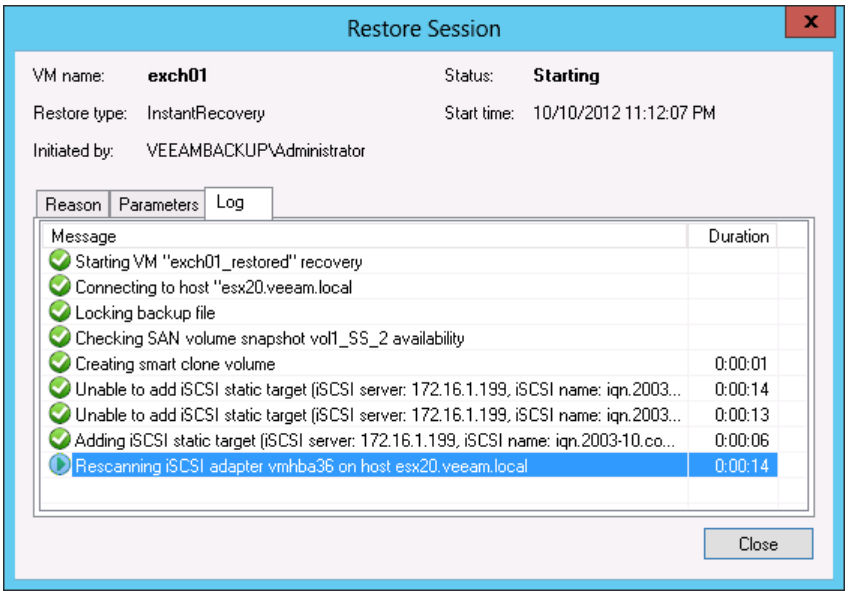

After the VM has been successfully restored, you can finalize Instant VM Recovery in two ways: migrate the restored VM to production or unpublish the restored VM.

To migrate the restored VM to production:

1. Open the **Backup & Replication** view.
2. Select the **Instant Recovery** node.
3. Right-click the necessary Instant VM Recovery session in the working area and select **Migrate to production**. As a result, the Quick Migration wizard will be launched. During migration, Veeam Backup & Replication will restore a VM instance from the backup file and then additionally move the changes that were made while the VM was running in the Instant Recovery mode.

To unpublish the restored VM:

1. Open the **Backup & Replication** view.
2. Select the **Instant Recovery** node.
3. Right-click the necessary session in the working area and select **Stop publishing**.

You can also unpublish the VM using the **History** view:

1. Open the **History** view.
2. Select the **Restore > Instant VM recovery** node.
3. Right-click the necessary session in the working area and select **Stop session**.

Restoring VM guest OS files (Windows)

With the **Restore** wizard, you can restore individual Windows guest OS files from any Windows-based VM located on the SAN snapshot.

To let you perform file−level recovery for Windows-based VMs, Veeam Backup & Replication creates a clone of the SAN snapshot and mounts the snapshot clone to the selected ESX(i) host. After that, it accesses the configuration file of the VM (VMX) on the snapshot clone and uses this configuration file to register a temporary VM on the ESX(i) host. Veeam Backup & Replication then mounts disks of the restored VM to the temporary VM. After that, you can copy necessary guest OS files and folders to their initial location, to your local machine drive or save them anywhere within the network.

Note	Guest OS file-level restore for ReFS is supported only if Veeam Backup & Replication is installed on Windows 2012 Server.

To recover VM guest OS files from the SAN snapshot, follow the next steps:

Step 1. Launch the Restore Wizard

To launch the **Restore** wizard, do one of the following:

- On the **Home** tab, click **Restore** and select **VMware**. In the **Restore from backup** section, select **Guest files (Windows)**.
- Open the **SAN Infrastructure** view. In the inventory pane, select the necessary volume snapshot. In the working area, right-click the VM whose guest OS files you want to restore and select **Restore Windows guest files**. In this case, you will pass immediately to step 5.
- Open the **Backup & Replication** view and select the **SAN** node in the inventory pane. In the working area, expand the necessary volume, select the VM whose guest OS files you want to restore and click **Guest Files > Guest Files (Windows)** on the ribbon. In this case, you will pass immediately to step 5.
- Open the **Backup & Replication** view and select the **SAN** node in the inventory pane. In the working area, expand the necessary volume, right-click the VMs whose guest OS files you want to restore and select **Restore guest files (Windows)**. In this case, you will pass immediately to step 5.

Tip	To quickly find the necessary VM, use the search field at the top of the window: enter the VM name or a part of it and press **[ENTER]**.

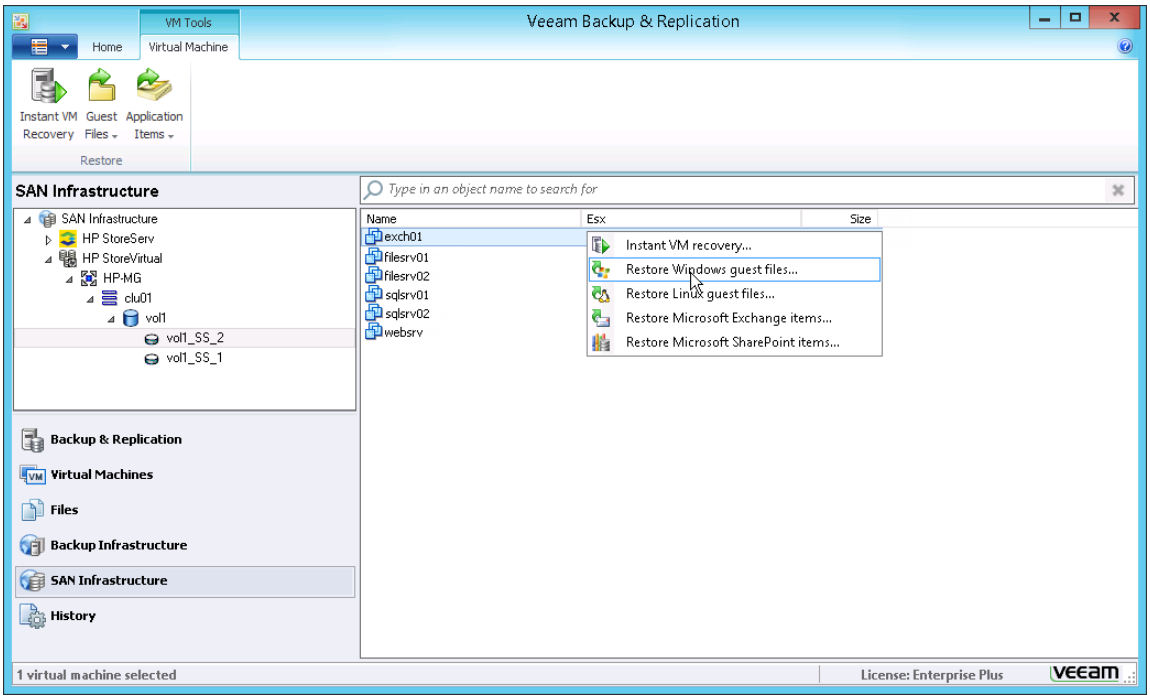

Step 2. Select a Virtual Machine

At the **Virtual Machine** step of the wizard, expand the necessary volume snapshot and select the VM whose guest OS files you want to restore.

Tip	To quickly find the necessary VM, use the search field at the top of the window: enter the VM name or a part of it and press **[ENTER]**.

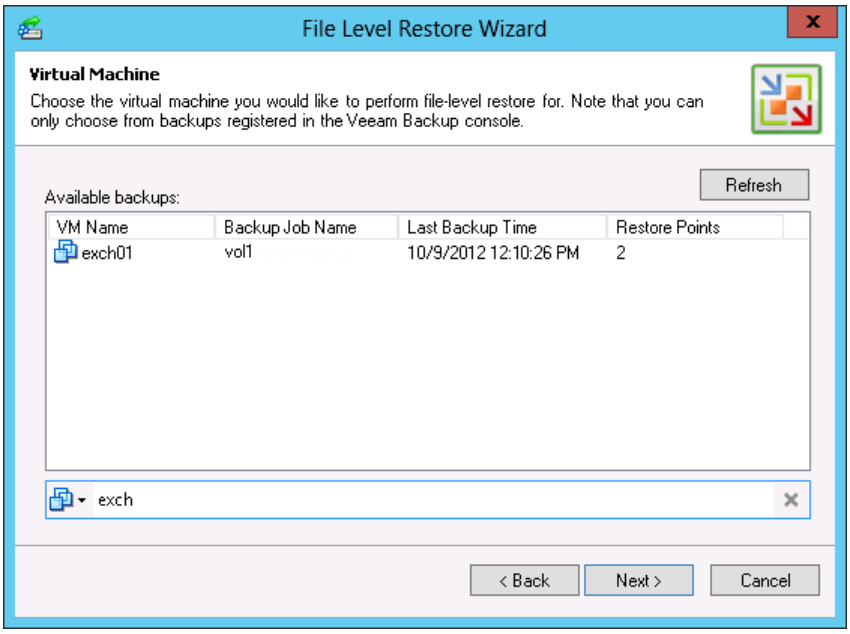

Step 3. Select a Restore Point

At the **Restore Point** step of the wizard, select the necessary restore point for the VM. Every snapshot of a volume acts as an independent restore point.

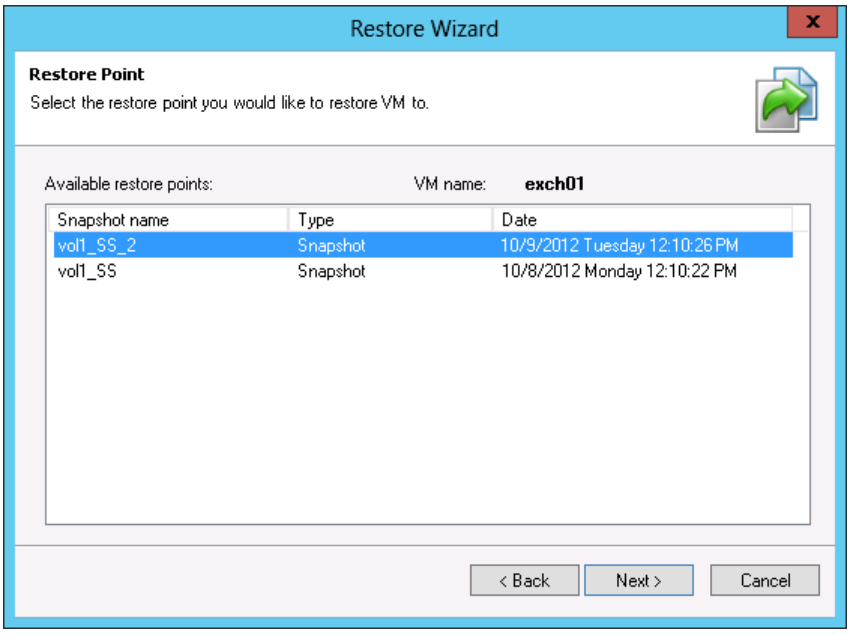

Step 4. Specify a Restore Reason

At the **Restore Reason** step of the wizard, enter the reason for performing VM guest file restore if necessary. The information you provide will be saved in the session history so that you can reference it later.

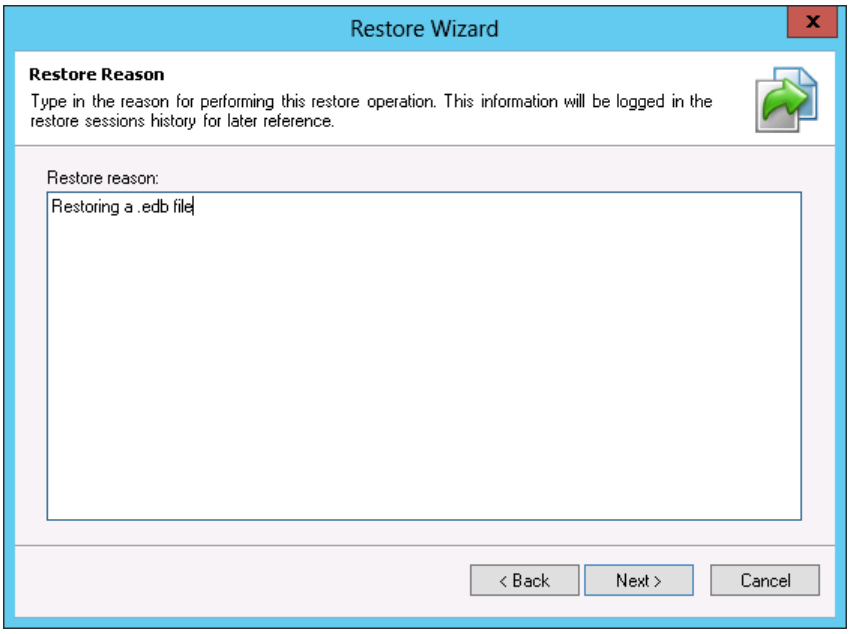

Step 5. Select an ESX(i) Host for Snapshot Mounting

At the **Ready** step of the wizard, you should select an ESX(i) host to which the clone of the SAN snapshot will be mounted. On the selected ESX(i) host, Veeam Backup & Replication will create a temporary VM and mount the disks of the restored VM to this temporary VM. After you restore the necessary files and finish working with the wizard, the temporary VM will be deleted and the snapshot clone will be unmounted from the ESX(i) host.

To specify the destination for the snapshot clone and a temporary VM:

1. Click **Choose** next to the **Host** field and select the ESX(i) host to which the snapshot clone should be mounted and on which the temporary VM should be created.
2. Click **Choose** next to the **Resource pool** field and select the resource pool to which the temporary VM should be placed.
3. Click **Choose** next to the **Folder** field and select the folder to which the temporary VM should be placed.
4. Click **OK**.
5. Click **Start** to begin the restore process.

Once restore is completed, Veeam Backup & Replication will open a file browser displaying the file system tree of the restored VM. Please note that the names of the restored machine drives may differ from the original ones.

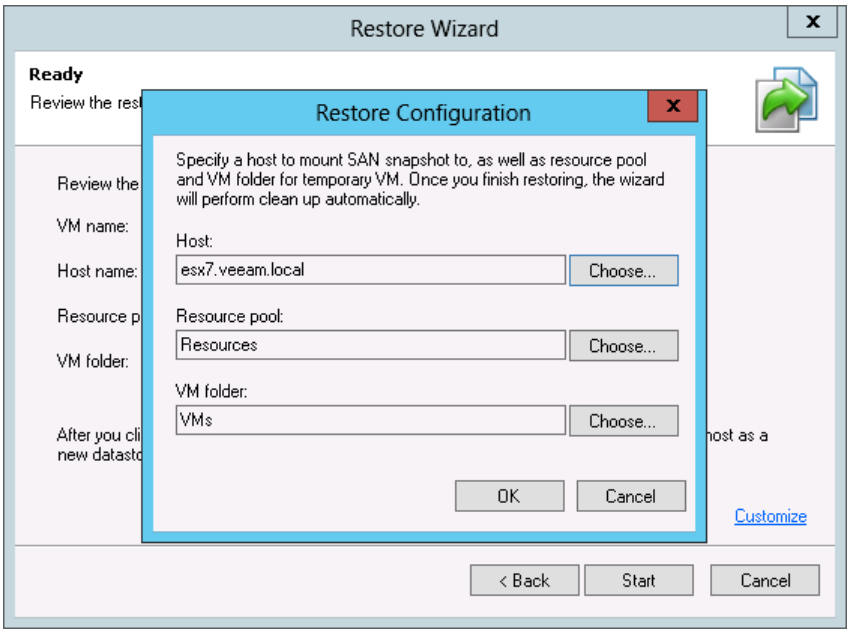

Step 6. Save Restored Files

You can save guest OS files to their initial location, to any folder on the local machine or within the network or open Windows Explorer for work with files.

Note You can browse to the VM guest OS files only while the Veeam Backup browser with the restored files is open. After the Veeam Backup browser is closed, the VM disks will be unmounted from the temporary VM. The temporary VM will be deleted and the SAN snapshot will be unmounted from the ESX(i) host (unless this SAN snapshot is used by other restore operations).

Saving Files to the Initial Location

To save files or folders to their initial location:

1. Right–click the necessary file or folder in the file system tree or in the details pane on the right and select **Restore**.
2. In the **Credentials** window, specify credentials of the account that will be used to connect to the initial VM. When you restore files to their initial location, Veeam Backup & Replication deploys a small runtime process in the initial VM. The process is used to control restore operations.
   To deploy the process, you need to connect to the initial VM under an account having administrator permissions on this VM. You can use the account under which you are currently logged on or choose another account.
3. Click **OK** to start the restore process.

Important! Before you start restoring files to the initial location, make sure that VMware Tools are installed on the target VM. Otherwise the restore process will fail.

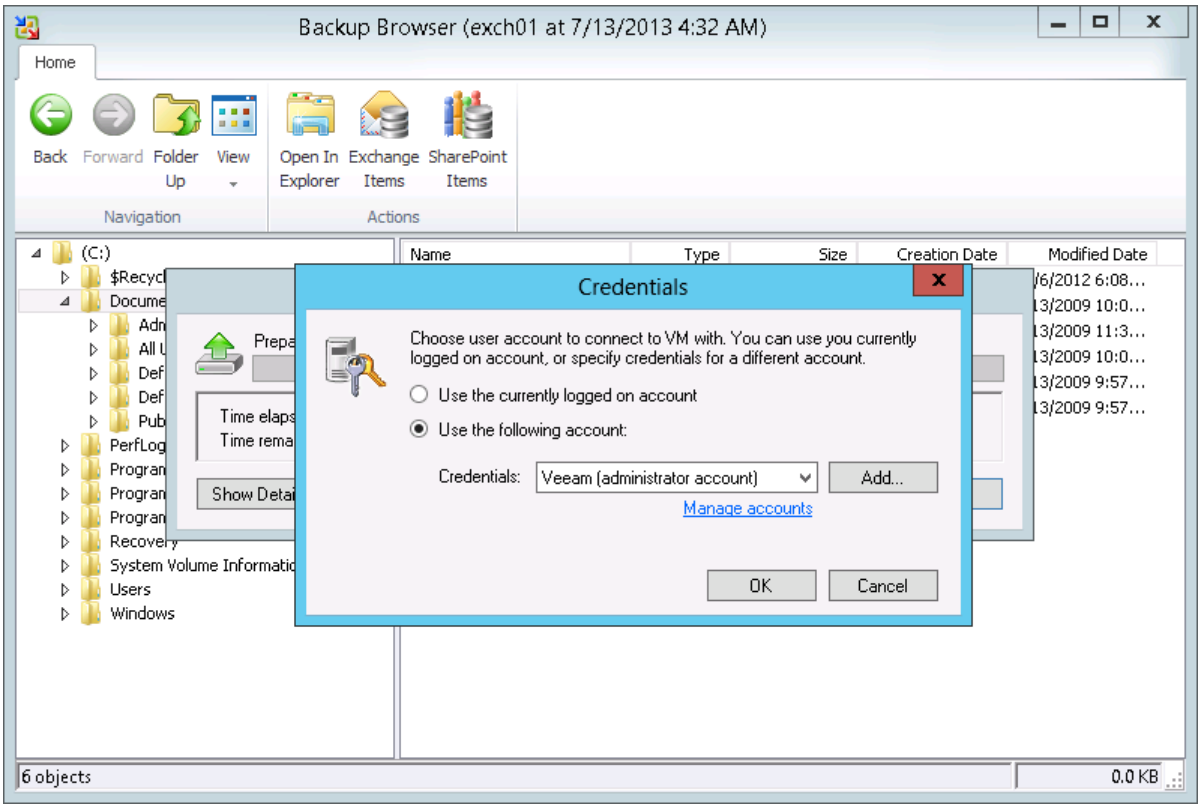

Saving Files to a New Location

To save restored files or folders on the local machine or a network share, right–click the necessary file or folder in the file system tree or in the details pane and select **Copy To**.

When restoring file objects, you can choose to preserve their original NTFS permissions:

- Select the **Preserve permissions and ownership** check box to keep the original ownership and security permissions for restored objects. Veeam Backup & Replication will copy selected files and folders along with associated Access Control Lists, preserving granular access settings.
- Leave the **Preserve permissions and ownership** check box not selected if you do not want to preserve the original ownership and access settings for restored objects. In this case, Veeam Backup & Replication will change security settings: the user who launched the Veeam Backup & Replication console will be set as the owner of the restored object, while access permissions will be inherited from the folder to which the restored object is copied.

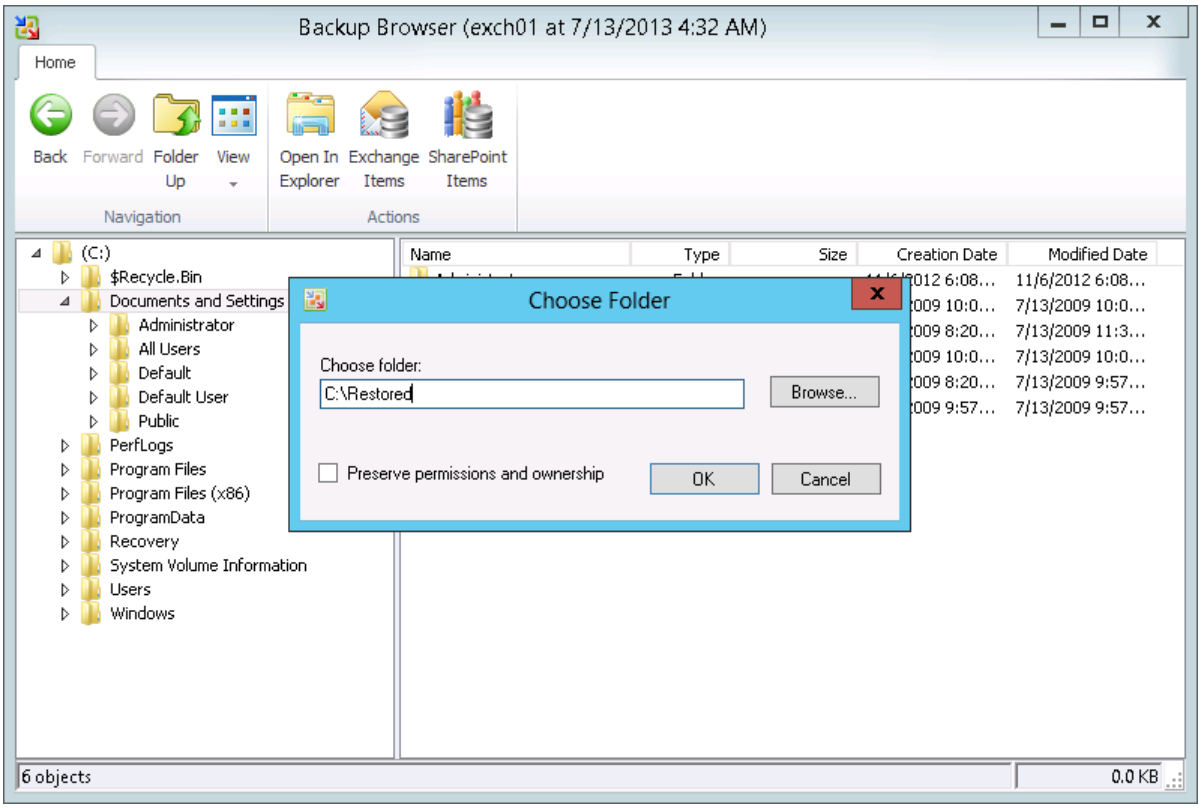

Working with Windows Explorer

Beside copying files via the Veeam Backup browser, you can use Windows Explorer to work with restored files. Click **Explore** on the ribbon in the Veeam Backup browser or right-click the necessary folder and select **Explore**. Veeam Backup & Replication will launch Windows Explorer so that you can browse to VM guest OS files.

You can also start Windows Explorer manually and browse to the necessary files. VM disks are mounted under the `C:\veeamflr\<vmname>\<volume n>` folder of the Veeam backup server.

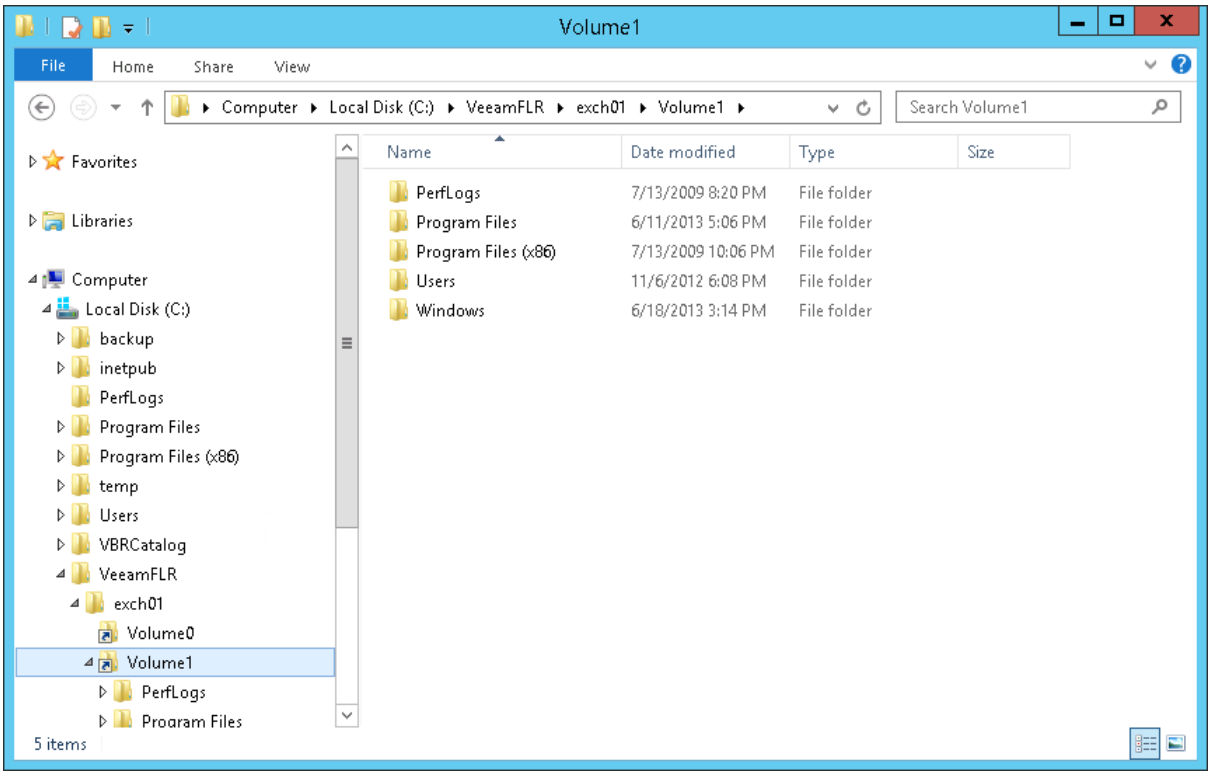

Tip If you are restoring guest OS files of the virtualized Microsoft Exchange server or Microsoft Sharepoint server, you can launch Veeam Explorer for Exchange and Veeam Explorer for SharePoint directly from the Veeam Backup browser:

- To start Veeam Explorer for Exchange, browse to the Exchange database file (EDB) in the Veeam Backup browser, select it and click **Exchange Items** on the **Home** tab or simply double-click the EDB file.
- To start Veeam Explorer for SharePoint, browse to the Microsoft SharePoint content database (MDF) in the Veeam Backup browser, select it and click **SharePoint Items** on the **Home** tab or simply double-click the MDF file.

Restoring VM guest OS files (Linux, Unix and other)

The **Veeam File Level Restore** wizard allows you to restore VM guest files and folders for 15 most commonly used file systems on Windows, Linux, Solaris, BSD, Unix and Novell Storage Services.

To let you perform file−level recovery, Veeam Backup & Replication creates a clone of the SAN snapshot and mounts the snapshot clone to the selected ESX(i) host. To the snapshot clone, it copies an ISO of the proxy appliance, a helper VM. The proxy appliance is very small — around 20 MB and takes only 10-20 seconds to boot.

Veeam Backup & Replication automatically starts the proxy appliance on the ESX(i) host and mounts disks of the restored VM to the proxy appliance as virtual hard drives. VMDK files are mounted directly from backup files, without prior extraction of the backup file content. After that, you can copy necessary files and folders to your local machine drive or save them anywhere within the network.

To perform instant multi-OS restore, follow the next steps:

Step 1. Launch the Veeam File Level Restore Wizard

To launch the **Restore** wizard, do one of the following:

- On the **Home** tab, click **Restore** and select **VMware**. In the **Restore from backup** section, select **Guest files (other OS)**.
- Open the **SAN Infrastructure** view. In the inventory pane, select the necessary volume snapshot. In the working area, right-click the VM whose guest OS files you want to restore and select **Restore Linux guest files**. In this case, you will pass immediately to step 4.
- Open the **Backup & Replication** view and select the **SAN** node in the inventory pane. In the working area, expand the necessary volume, select the VM whose guest OS files you want to restore and click **Guest Files > Guest Files (other OS)** on the ribbon. In this case, you will pass immediately to step 4.
- Open the **Backup & Replication** view and select the **SAN** node in the inventory pane. In the working area, expand the necessary volume, right-click the VMs whose guest OS files you want to restore and select **Restore guest files (other OS)**. In this case, you will pass immediately to step 4.

Tip To quickly find the necessary VM, use the search field at the top of the window: enter the VM name or a part of it and press **[ENTER]**.

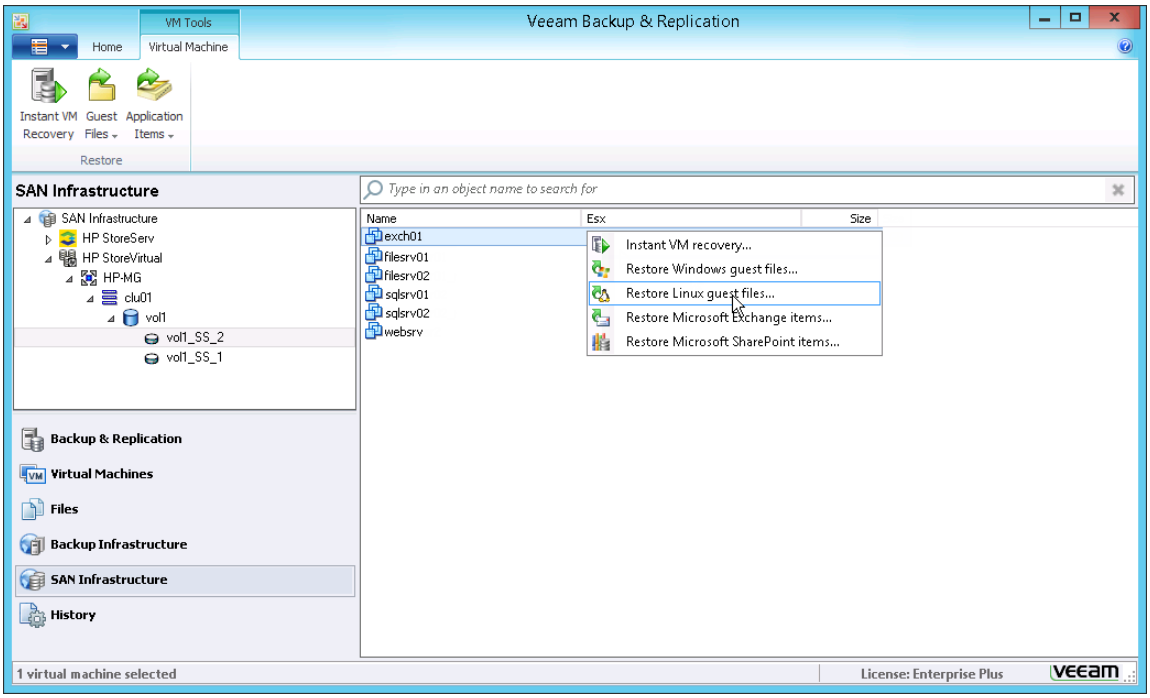

Step 2. Select a Virtual Machine

At the **Virtual Machine** step of the wizard, expand the necessary volume snapshot and select the VM whose guest OS files you want to restore.

Tip To quickly find the necessary VM, use the search field at the top of the window: enter the VM name or a part of it and press **[ENTER]**.

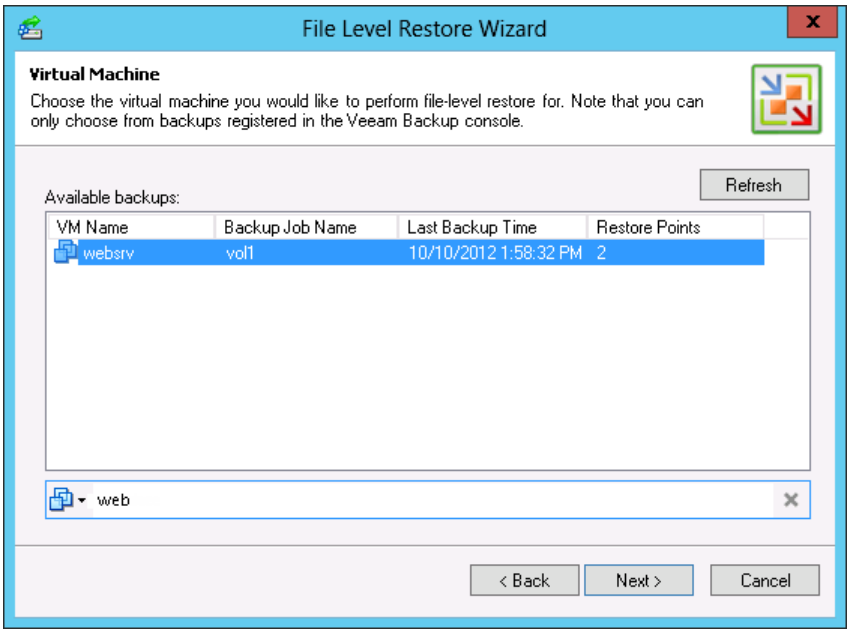

Step 3. Select a Restore Point

At the **Restore Point** step of the wizard, select the necessary restore point for the VM. Every snapshot of a volume acts as an independent restore point.

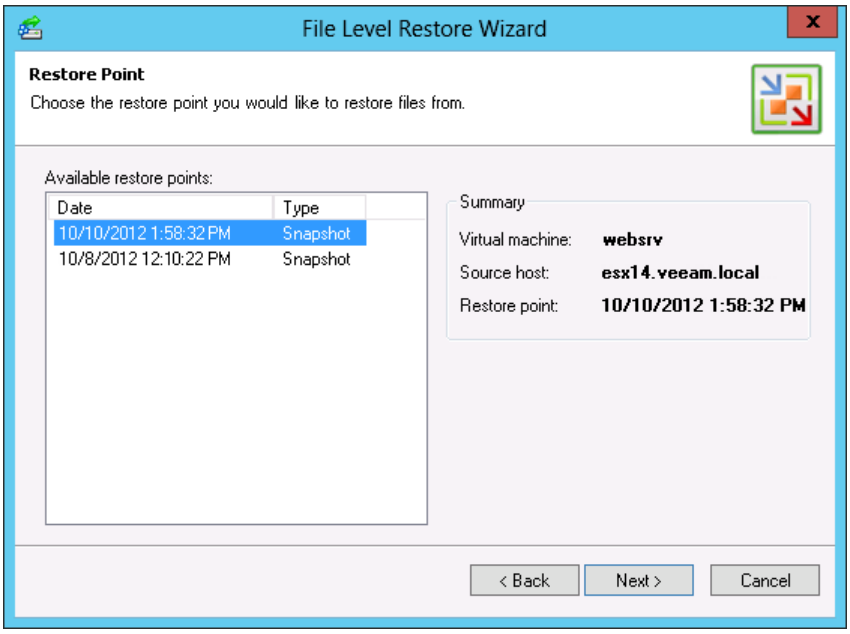

Step 4. Select the Location for the Proxy Appliance

At the **Ready** step of the wizard, you should select an ESX(i) host to which the clone of the SAN snapshot will be mounted. On the selected ESX(i) host, Veeam Backup & Replication will also place a proxy appliance and mount disks of the restored VM to this proxy appliance. After you restore the necessary files and finish working with the Veeam Backup browser, the proxy appliance will be deleted and the snapshot clone will be unmounted from the ESX(i) host.

To locate the appliance and select an ESX(i) host for snapshot mounting:

1. Click the **Customize** link at the bottom of the window.
2. In the **FLR Appliance Configuration** window, select the ESX(i) host to which the snapshot clone will be mounted and on which the proxy appliance will be registered.
3. Specify the resource pool and network in which the proxy appliance will be run.
4. Select between a static or dynamic IP address for the proxy appliance and specify the necessary network settings for the proxy appliance.
5. To enable FTP access to the restored file system, select the **Enable FTP server on appliance** check box. As a result, your users will be able to access the proxy appliance via FTP, browse the file system of the restored VM and download necessary files on their own.
6. If you are performing restore of a VM with the Novell Storage services file system, select the **Restore Novell Storage services filesystem** check box. In this case, Veeam Backup & Replication will deploy a specific proxy appliance supporting the Novell file system.
7. Click **OK**.

Important! When choosing an ESX(i) host for the Novell file system proxy appliance, make sure that it allows running VMs with 64-bit guest OSs.

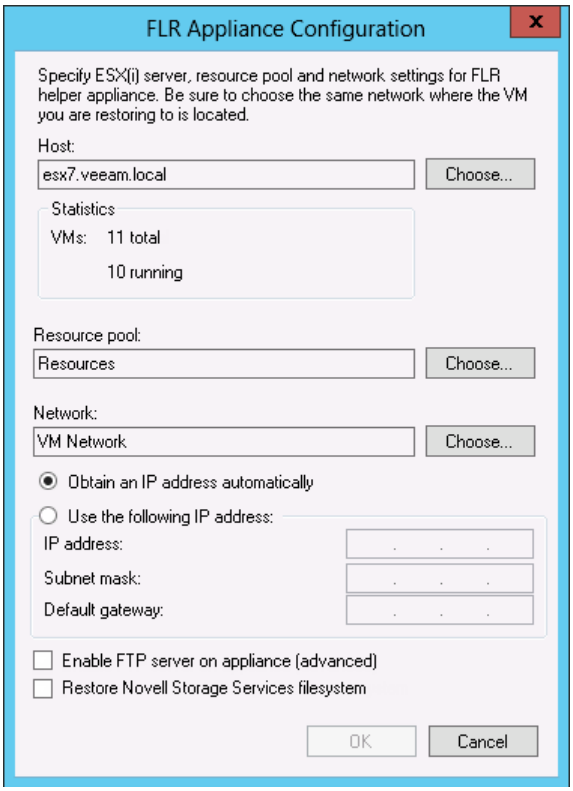

Step 5. Start the Restore Process

Click **Next** to start the restore process. Please note that the file-level restore appliance may take about 10-20 seconds to boot.

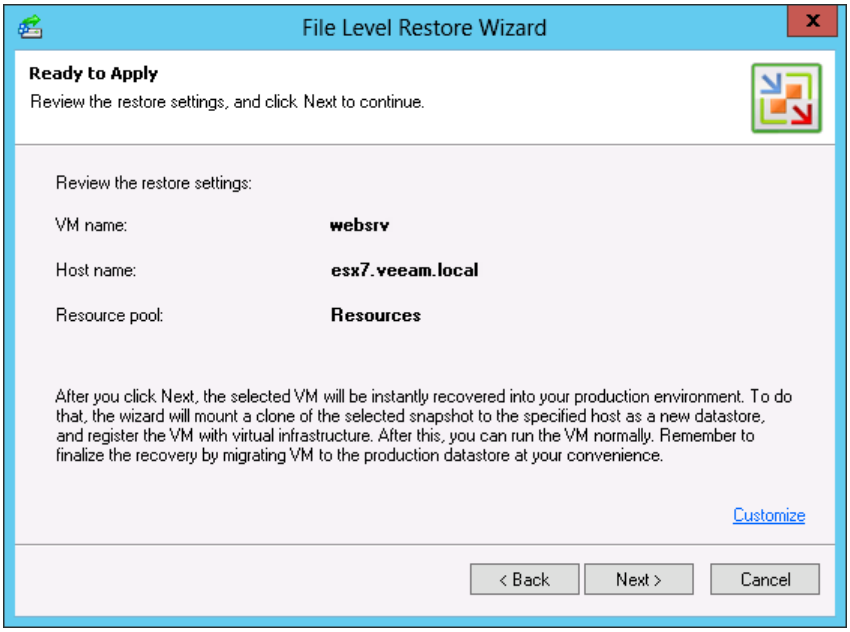

Use the **File Level Restore** wizard to trace the progress of the file-level restore process:

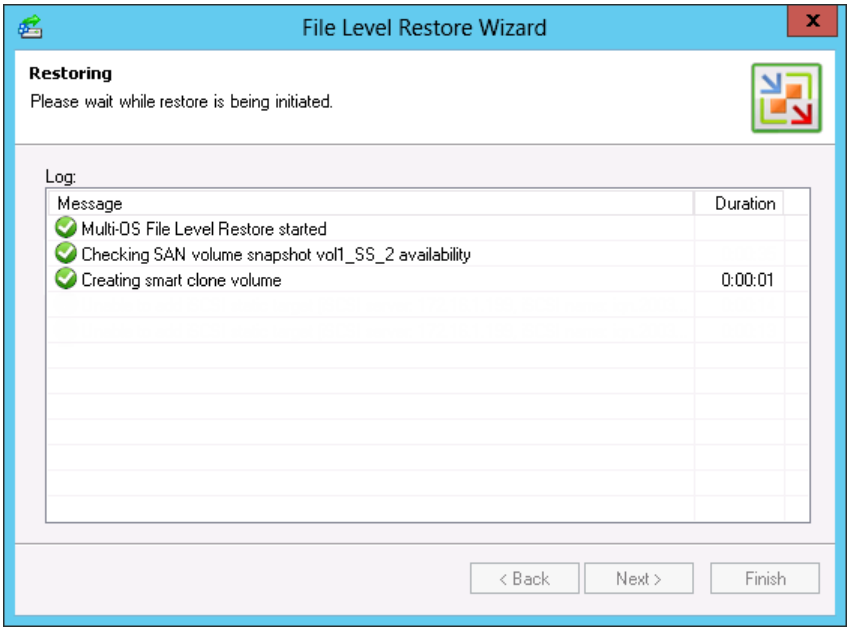

Step 6. Save Restored Files

Once the restore process is completed, Veeam Backup & Replication will display a file browser with the file system tree of the restored VM.

To save restored files or folders on the local machine or within the network:

1. Right-click the necessary file or folder and select **Copy to**.
2. Select the destination folder on the local or remote host.
3. If you are recovering files to the original Linux host, you can preserve file permissions. Note in this case, the Linux host must be added to the list of servers managed by Veeam Backup & Replication in advance. For details, see Adding a Linux Server. Select the **Preserve permissions and ownership** check box to keep original permission settings for recovered files. Ownership settings are restored only if you have privileges to change the owner at the remote Linux host where files are restored.
4. Click **Restore**. The file or folder will be saved at the specified folder on the host.

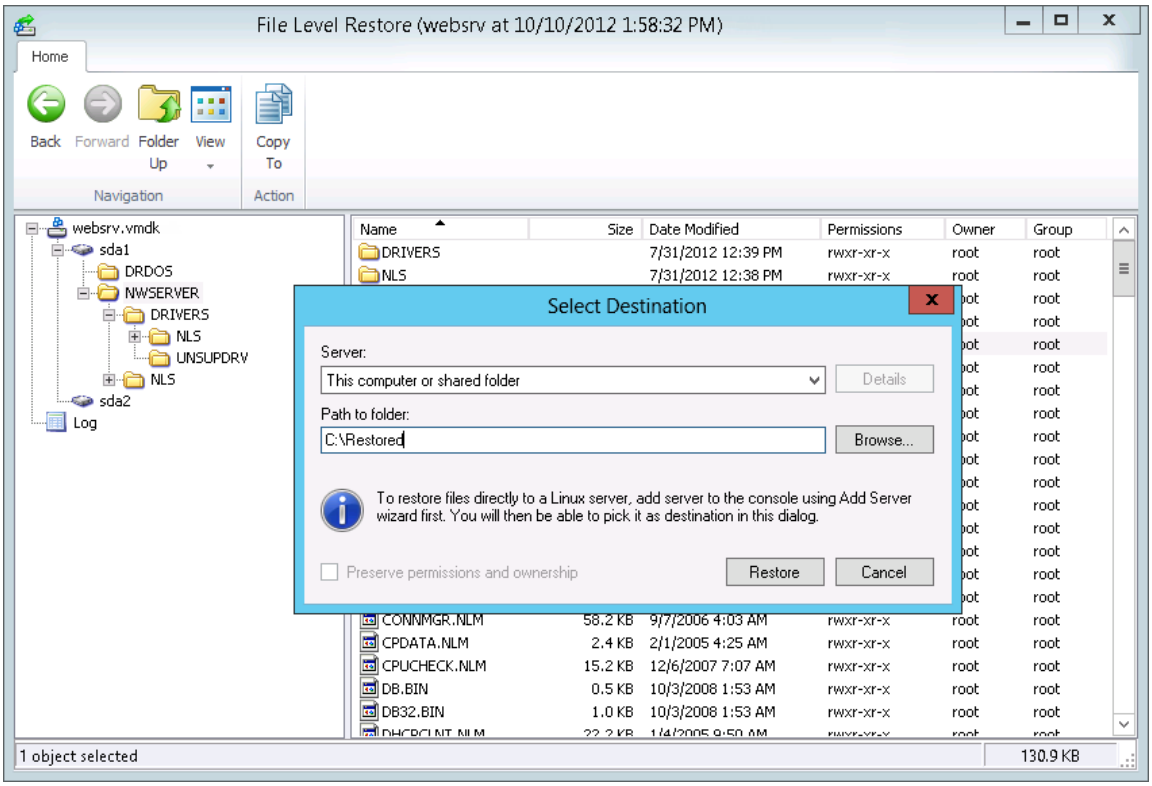

If you have chosen to enable FTP server on the FLR appliance, the restored file system will also be available over FTP at *ftp://<FLR_appliance_IP_address>*. Other users in the same network can access the FLR appliance to restore the files they need.

Note You can browse to the VM guest OS files only while the Veeam Backup browser with the restored files is open. After the Veeam Backup browser is closed, the VM disks will be unmounted from the proxy appliance. The proxy appliance will be deleted and the SAN snapshot will be unmounted from the ESX(i) host, unless this SAN snapshot is used by other restore operations.

Restoring Individual Items from Microsoft Exchange

Veeam Backup & Replication integrates with Veeam Explorer for Exchange — Veeam's tool that lets you browse Exchange mailbox stores inside Veeam backups. It features an easy-to-use interface and allows you to quickly locate the mailboxes or items you need.

Before you can start working with Veeam Explorer for Exchange, you need to extract the Exchange database (EDB file) from the HP SAN snapshot. You can do it in one of the following ways:

- You can use the Microsoft Exchange Item Level Restore wizard. In this case, Veeam Backup & Replication will automatically extract the Exchange database from the HP SAN snapshot and open it in Veeam Explorer for Exchange.
- You can perform guest OS files recovery for the virtualized Microsoft Exchange server, manually locate the restored Exchange database and open it in Veeam Explorer for Exchange.

Using the Microsoft Exchange Item Level Restore Wizard

When you run the **Microsoft Exchange Item Level Restore** wizard, Veeam Backup & Replication automatically extracts the Microsoft Exchange database (EDB file) from the HP SAN snapshot and opens it in Veeam Explorer for Exchange.

As part of this procedure, Veeam Backup & Replication performs the following steps:

1. Veeam Backup & Replication creates a clone of the SAN snapshot and mounts the snapshot clone to an ESX(i) host.
2. Veeam Backup & Replication accesses the configuration file of the virtualized Microsoft Exchange server (VMX) on the snapshot clone and uses this configuration file to register a temporary VM on the ESX(i) host.
3. Veeam Backup & Replication mounts disks of the restored Microsoft Exchange server to the temporary VM.
4. Veeam Backup & Replication scans the Microsoft Exchange Event log to locate the EDB file and the `ese.dll` file on the mounted disks. The `ese.dll` file is required for configuring Veeam Explorer for Exchange.
5. After the EDB file is successfully located, Veeam Backup & Replication opens it in Veeam Explorer for Exchange so that you can browse it to find necessary items. If for some reason Veeam Backup & Replication fails to locate and open the EDB file automatically, you can manually extract the EDB file and open it in Veeam Explorer for Exchange. To learn more, see Opening the EDB File Manually.

To restore Microsoft Exchange objects from the HP SAN snapshot, follow the next steps:

Step 1. Launch the Microsoft Exchange Item Level Restore Wizard

To launch the **Microsoft Exchange Item Level Restore** wizard, do one of the following:

- On the **Home** tab, click **Restore** and select **VMware**. In the **Restore from backup** section, select **Application items** and click **Next**. At the **Select Application** step of the wizard, select **Microsoft Exchange**.
- Open the **SAN Infrastructure** view. In the inventory pane, select the necessary volume snapshot. In the working area, right-click the VM running Microsoft Exchange and select **Restore Microsoft Exchange items**. In this case, you will pass immediately to step 5.
- Open the **Backup & Replication** view and select the **SAN** node in the inventory pane. In the working area, expand the necessary volume, select the VM running Microsoft Exchange and click **Application Items > Exchange Items** on the ribbon. In this case, you will pass immediately to step 3.
- Open the **Backup & Replication** view and select the **SAN** node in the inventory pane. In the working area, expand the necessary volume, right-click the VM running Microsoft Exchange and select **Restore Microsoft Exchange items**. In this case, you will pass immediately to step 5.

Tip To quickly find the necessary VM, use the search field at the top of the window: enter the VM name or a part of it and press **[ENTER]**.

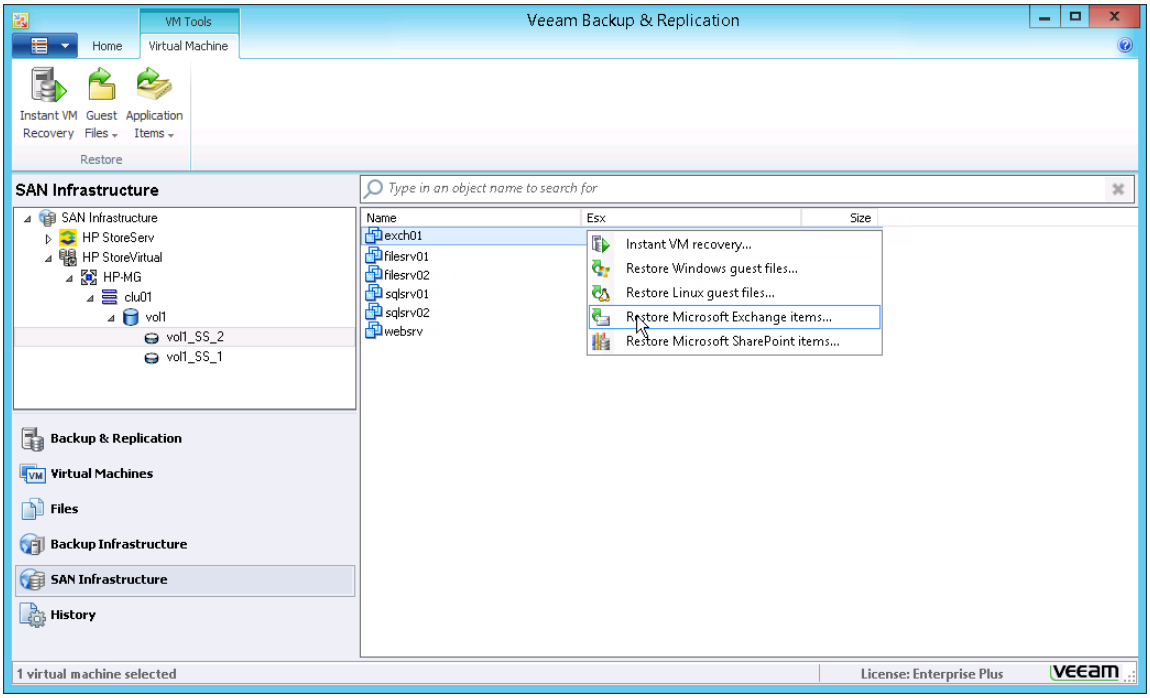

Step 2. Select a Virtual Machine

At the **Virtual Machine** step of the wizard, expand the necessary volume snapshot and select the VM running Microsoft Exchange.

Tip To quickly find the necessary VM, use the search field at the top of the window: enter the VM name or a part of it and press **[ENTER]**.

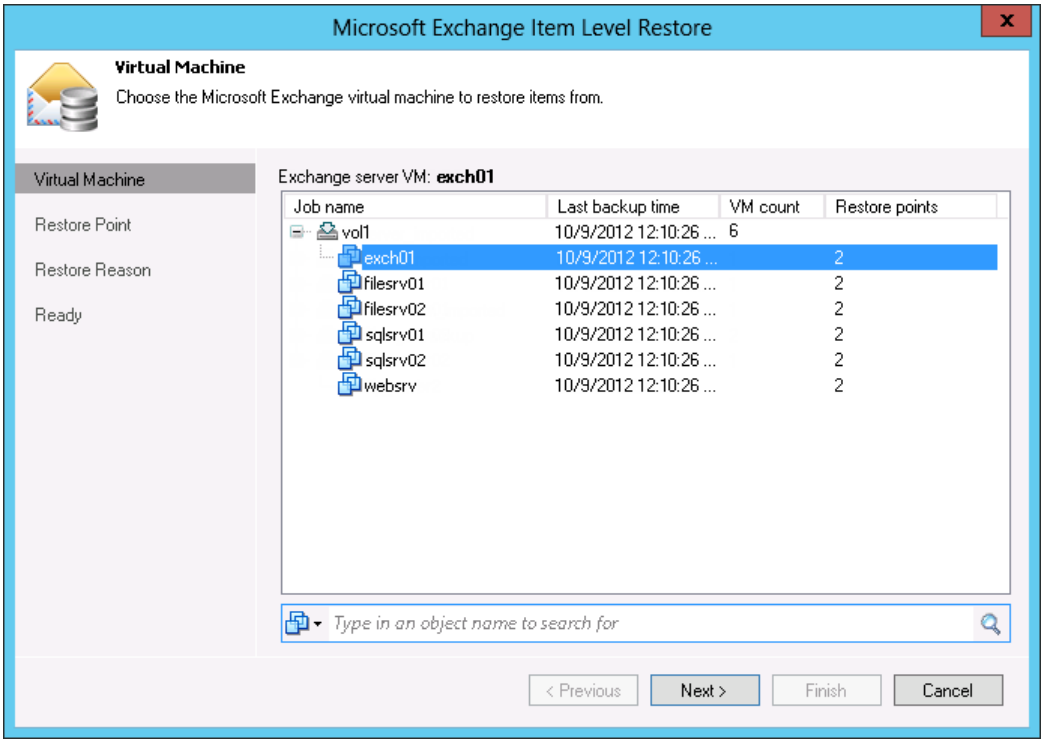

Step 3. Select a Restore Point

At the **Restore Point** step of the wizard, select the necessary restore point for the VM. Every snapshot of a volume acts as an independent restore point.

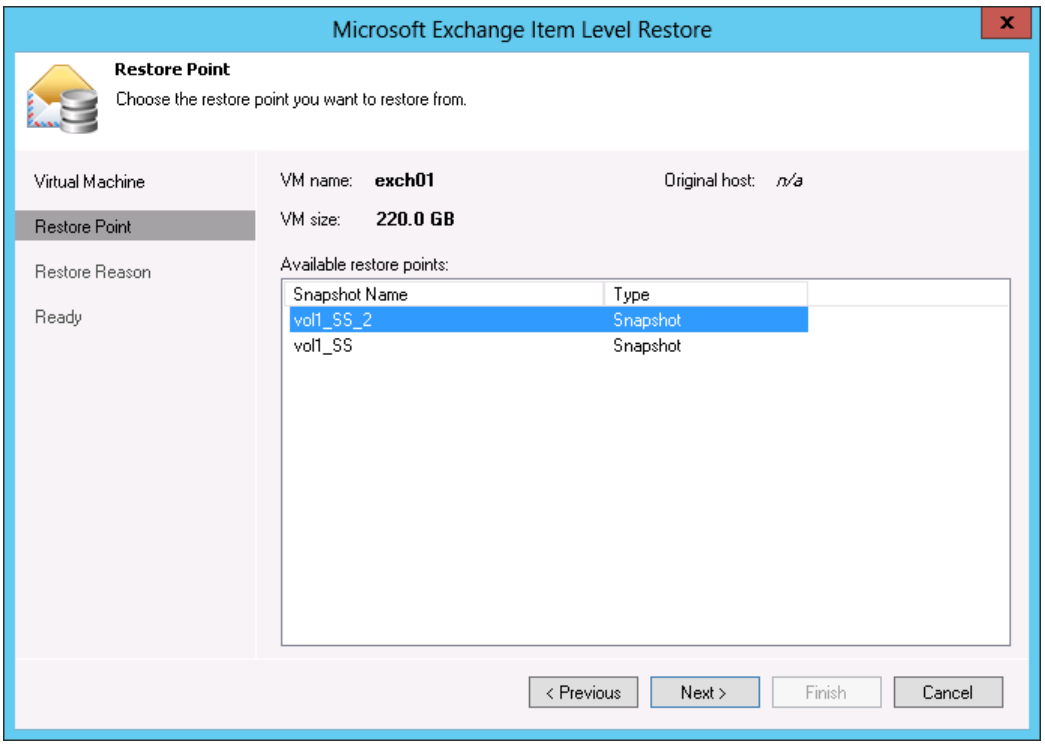

Step 4. Specify a Restore Reason

At the **Restore Reason** step of the wizard, enter the reason for restoring data from Microsoft Exchange if necessary. The information you provide will be saved in the session history so that you can reference it later.

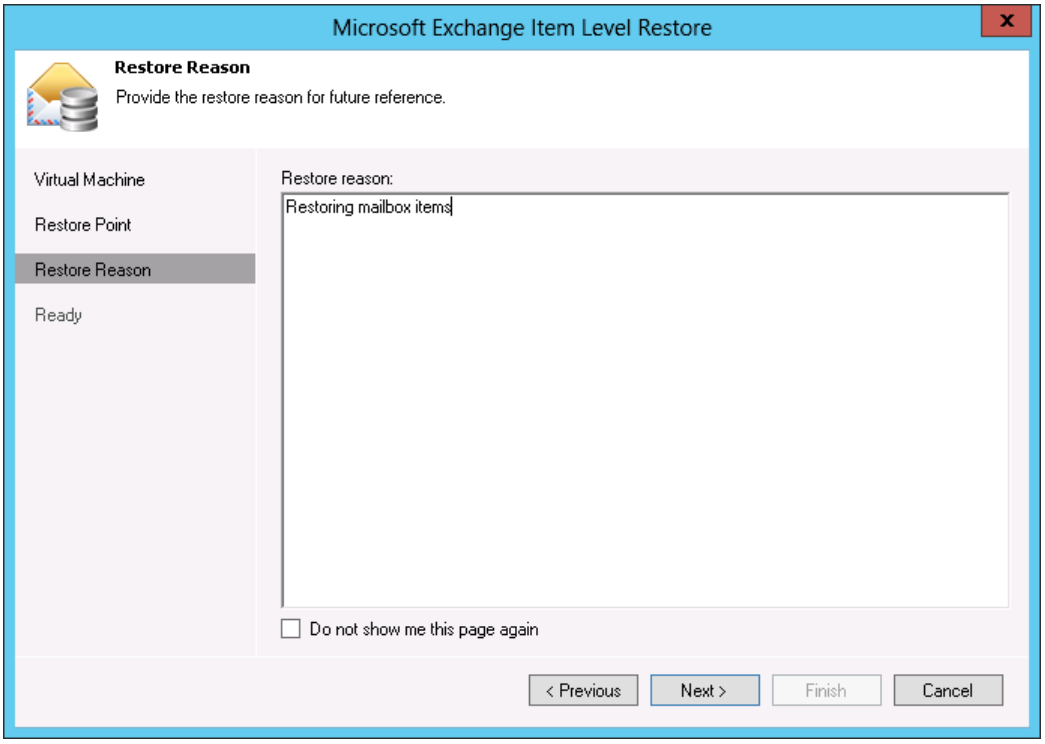

Step 5. Select an ESX(i) Host for Snapshot Mounting

At the last step of the wizard, you should select an ESX(i) host to which the clone of the SAN snapshot will be mounted. On the selected ESX(i) host, Veeam Backup & Replication will create a temporary VM and mount the disks of the virtualized Microsoft Exchange Server to this temporary VM.

To specify destination for the snapshot clone and a temporary VM:

1. Click **Choose** next to the **Host** field and select the ESX(i) host to which the snapshot clone should be mounted and on which the temporary VM should be created.
2. Click **Choose** next to the **Resource pool** field and select a resource pool to which the temporary VM should be placed.
3. Click **Choose** next to the **Folder** field and select a folder to which the temporary VM should be placed.
4. Click **OK**.
5. Click **Start** to begin the restore process.

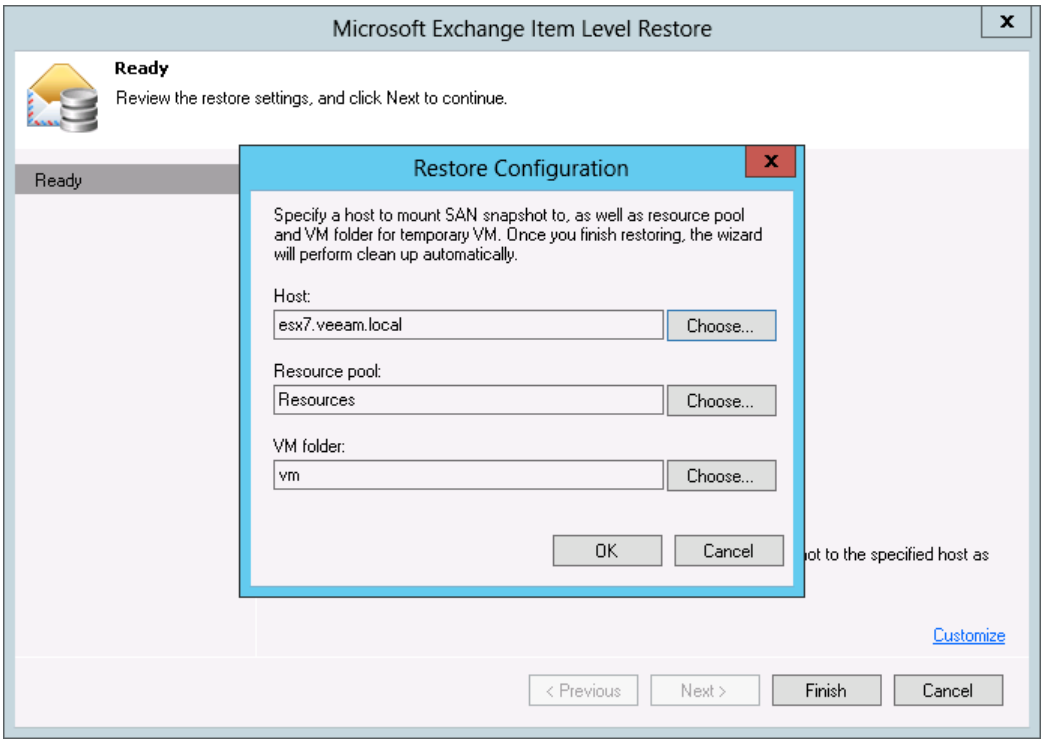

Step 6. Open the EDB File in Veeam Explorer for Exchange

Veeam Backup & Replication will automatically locate the Microsoft Exchange database and open it in Veeam Explorer for Exchange. After that, you can browse the database and restore the items you need.

To learn more, see Working with Veeam Explorer for Exchange.

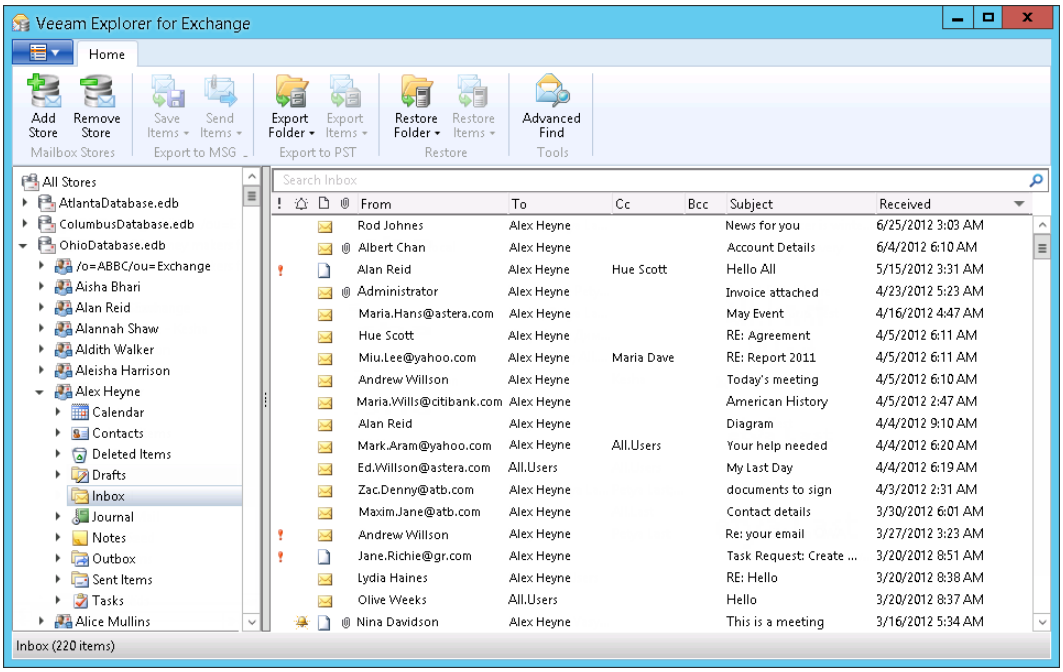

Opening the EDB File Manually

If Veeam Backup & Replication fails to locate the Microsoft Exchange database on the restored VM for some reason, you can restore it with guest OS files restore and open it manually.

1. Perform guest OS files restore for the virtualized Microsoft Exchange server.
2. In the Veeam Backup browser, double-click the EDB file or click **Exchange Items** on the ribbon.
3. Veeam Backup & Replication will open the selected database in Veeam Explorer for Exchange. After that, you can browse the database and restore the items you need.

To learn more, see Working with Veeam Explorer for Exchange.

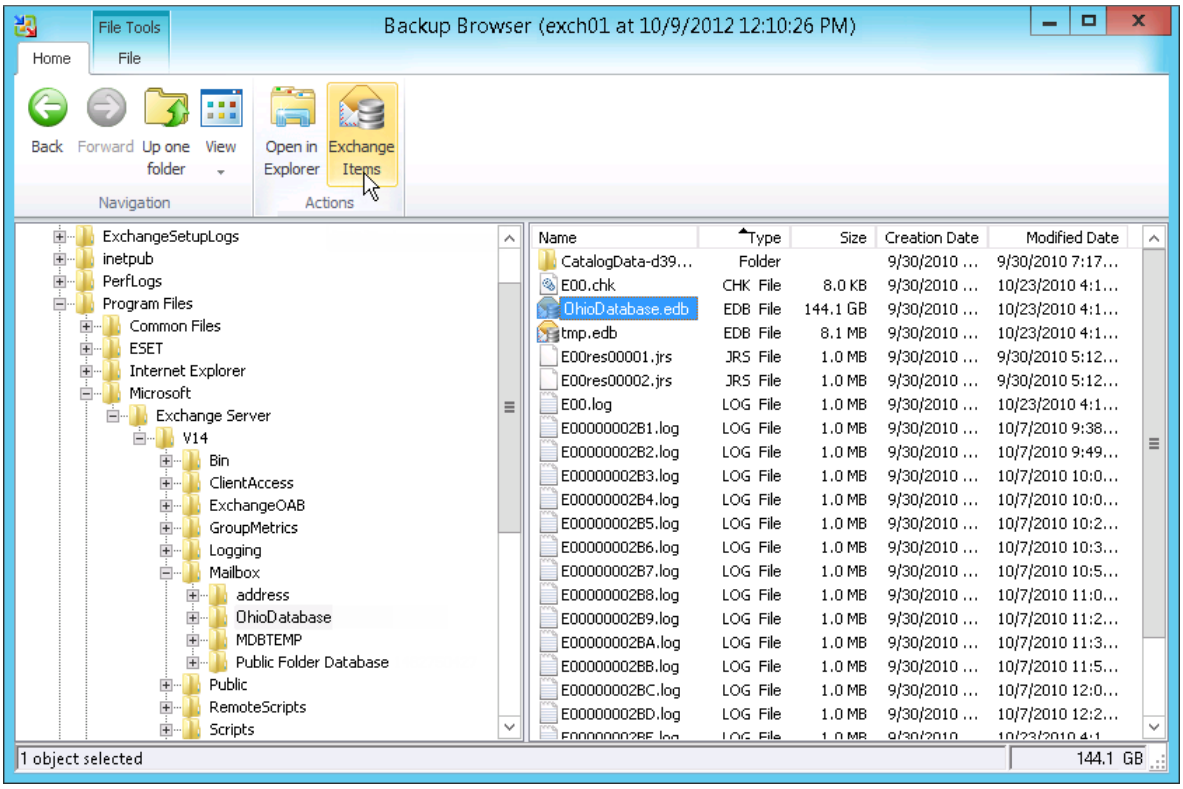

Restoring Individual Items from Microsoft SharePoint

Veeam Backup & Replication integrates with Veeam Explorer for SharePoint — Veeam's tool that lets you browse SharePoint content databases in VM disks located on HP SAN snapshots. Veeam Explorer for SharePoint features a familiar, easy-to-use interface and allows you to quickly locate the documents, items and document libraries you need; several recovery options (saving, e-mailing and others) are provided for Microsoft SharePoint 2010.

Before you can start working with Veeam Explorer for SharePoint, you should perform the following steps:

1. Recover the Microsoft SharePoint content database (MDF) from the HP SAN snapshot.
2. Add the recovered Microsoft SharePoint content database files to Veeam Explorer for SharePoint scope. To learn more, see Adding Content Database to Veeam Explorer's Scope.

Using the Microsoft SharePoint Farm Restore Wizard

When you run the **Microsoft Sharepoint Item Level Restore** wizard, Veeam Backup & Replication automatically extracts the Microsoft SharePoint content database from the HP SAN snapshot and opens it in Veeam Explorer for SharePoint.

As part of this procedure, Veeam Backup & Replication performs the following steps:

1. Veeam Backup & Replication creates a clone of the SAN snapshot and mounts the snapshot clone to an ESX(i) host.
2. Veeam Backup & Replication accesses the configuration file of the virtualized Microsoft SharePoint server (VMX) on the snapshot clone and uses this configuration file to register a temporary VM on the ESX(i) host.
3. Veeam Backup & Replication mounts disks of the restored Microsoft SharePoint server to the temporary VM.
4. Veeam Backup & Replication opens the Veeam Explorer for SharePoint so that you can add database and browse it to find necessary items.

To restore Microsoft SharePoint items from the HP SAN snapshot, follow the next steps:

Step 1. Launch the Microsoft SharePoint Farm Restore Wizard

To launch the **Microsoft SharePoint Farm Restore** wizard, do one of the following:

- On the **Home** tab, click **Restore** and select **VMware**. In the **Restore from backup** section, select **Application items** and click **Next**. At the **Select Application** step of the wizard, select **Microsoft SharePoint**.
- Open the **SAN Infrastructure** view. In the inventory pane, select the necessary volume snapshot. In the working area, right-click the VM running Microsoft Exchange and select **Restore Microsoft SharePoint items**.
- Open the **Backup & Replication** view and select the **SAN** node in the inventory pane. In the working area, expand the necessary volume, select the VM running Microsoft Exchange and click **Application Items > SharePoint Items** on the ribbon.
- Open the **Backup & Replication** view and select the **SAN** node in the inventory pane. In the working area, expand the necessary volume, right-click the VM running Microsoft Exchange and select **Restore Microsoft SharePoint items**.

Tip To quickly find the necessary VM, use the search field at the top of the window: enter the VM name or a part of it and press **[ENTER]**.

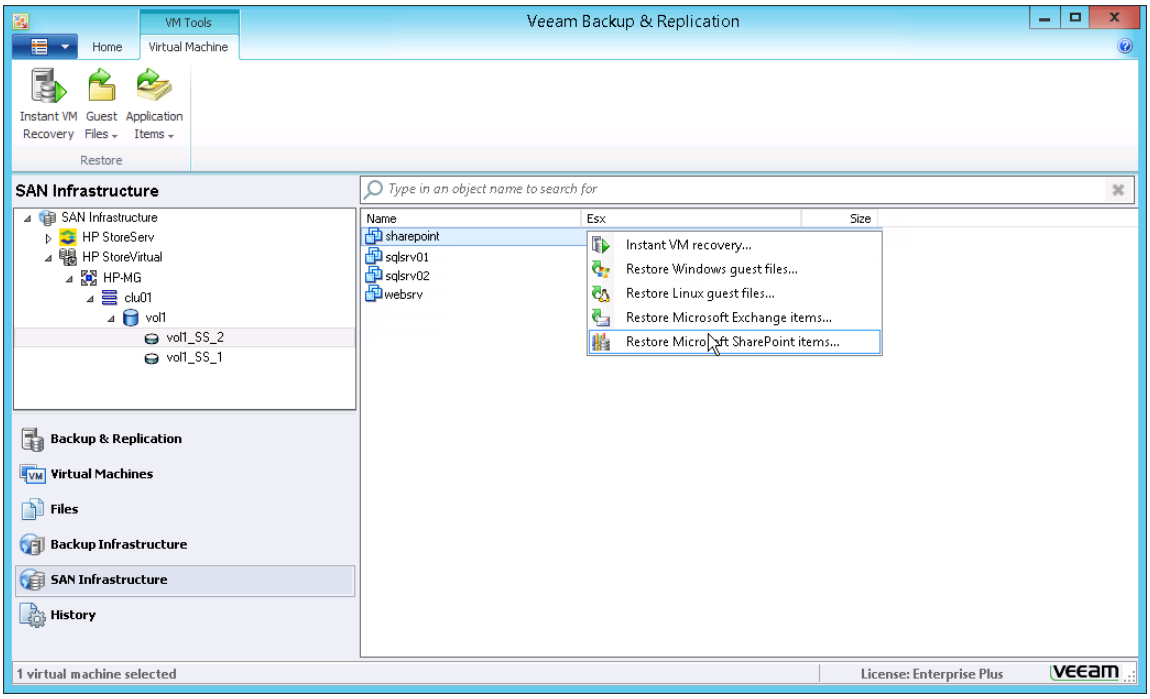

Step 2. Select a Restore Point

At the **Restore Point** step of the wizard, select the necessary restore point for the VM. Every snapshot of a volume acts as an independent restore point.

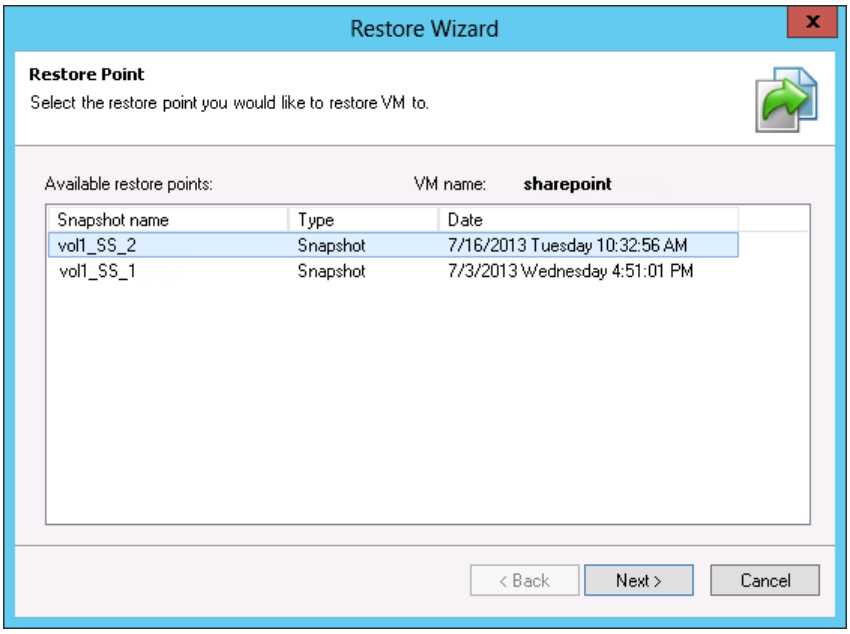

Step 3. Specify a Restore Reason

At the **Restore Reason** step of the wizard, enter the reason for restoring data from Microsoft Exchange if necessary. The information you provide will be saved in the session history so that you can reference it later.

Step 4. Select an ESX(i) Host for Snapshot Mounting

At the last step of the wizard, you should select an ESX(i) host to which the clone of the SAN snapshot will be mounted. On the selected ESX(i) host, Veeam Backup & Replication will create a temporary VM and mount the disks of the virtualized Microsoft SharePoint Server to this temporary VM.

To specify destination for the snapshot clone and a temporary VM:

1. Click **Choose** next to the **Host** field and select the ESX(i) host to which the snapshot clone should be mounted and on which the temporary VM should be created.
2. Click **Choose** next to the **Resource pool** field and select a resource pool to which the temporary VM should be placed.
3. Click **Choose** next to the **Folder** field and select a folder to which the temporary VM should be placed.
4. Click **OK**.
5. Click **Start** to begin the restore process.

Step 5. Open the Database in Veeam Explorer for SharePoint

Veeam Backup & Replication will perform guest OS files restore for the virtualized Microsoft SharePoint server and will mount the disks of the Microsoft SharePoint server to the Veeam backup server. After that, Veeam Backup & Replication will open the Veeam Explorer for SharePoint.

Add the content database file to Veeam Explorer for SharePoint directly from the mounted VM file system:

1. In Veeam Explorer for SharePoint, click **Add Database**.
2. Choose the folder where content database files reside. By default, these files reside by the following path:
   `%ProgramFiles%\Microsoft Office Servers\14.0\Data\MSSQL10.SHAREPOINT\MSSQL\DATA`

To learn more, see Working with Veeam Explorer for SharePoint.

Working with Backup Copy Jobs

To let you adopt the 3-2-1 backup strategy, Veeam Backup & Replication offers backup copying capabilities. Backup copy jobs allow you to create several instances of the same backup file in different locations, whether onsite or offsite. Copied backup files have the same format as those created by backup jobs and you can use any data recovery option for them.

Creating Backup Copy Jobs

To copy backup files to another location, you should create a backup copy job using the **New Backup Copy** Job wizard. This section will guide you through all steps of the wizard and provide explanation on available options.

To create a backup copy job, follow the next steps:

Step 1. Launch the New Backup Copy Job Wizard

To run the **New Backup Copy Job** wizard, do either of the following:

1. On the **Home** tab, click **Backup Copy Job** and select the necessary platform: *VMware or vCloud*.
2. Open the **Backup & Replication** view, right-click the **Jobs** node and select **Backup Copy** > *VMware or vCloud*.

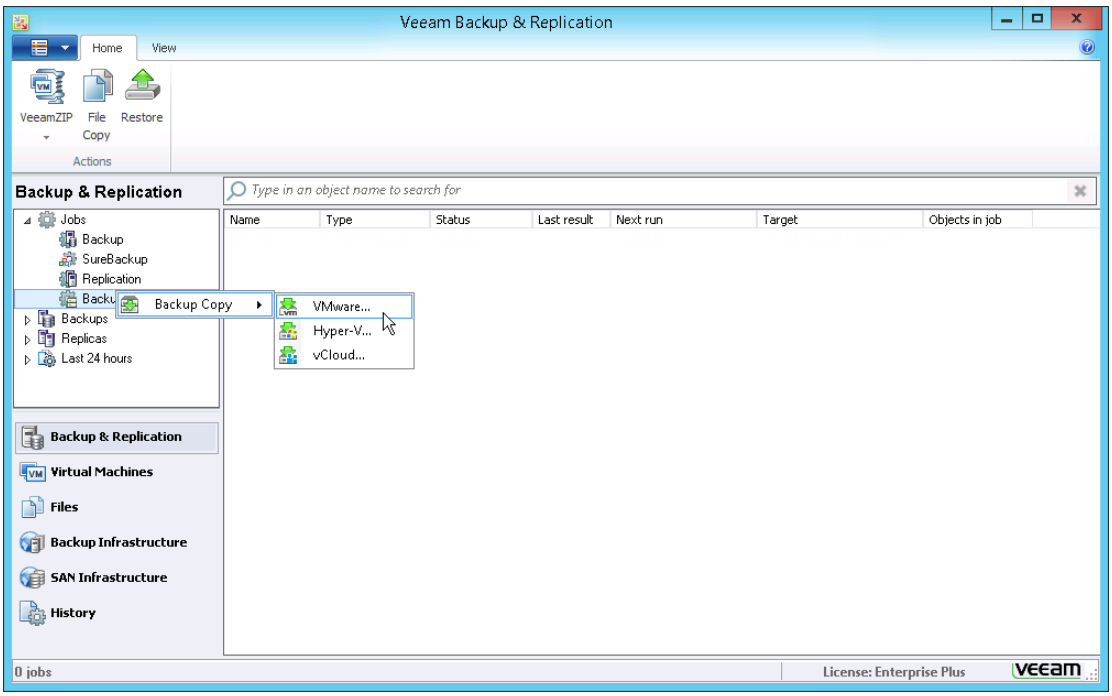

Step 2. Specify the Job Name and Description

At the **Job** step of the wizard, you should define basic settings for the created backup copy job.

1. In the **Name** field, enter a name for the created job.
2. In the **Description** field, enter a description of the created job. The default description contains information about the user who created the job, date and time when the job was created.

3. The backup copy job starts the synchronization process at specific time intervals. During this synchronization interval, Veeam Backup & Replication copies new restore points from the source backup repository to the target backup repository.
   In the **Copy every** field, specify the time interval according to which the synchronization process must be started.
   By default, the synchronization interval is set to 1 day. This means that the backup copy job will create a new synchronization interval once a day. Veeam Backup & Replication will check if a new restore point is available in the source backup repository. If a new restore point is found, it will be copied to the target backup repository within the synchronization interval. To learn more, see Synchronization Interval.
4. If you have selected a daily synchronization interval, specify the start time for it. By default, the daily synchronization interval starts at 12:00 AM

Important! In some cases, the defined synchronization interval may not be enough to copy a VM restore point. If such situation occurs, Veeam Backup & Replication will display a warning in the job session results. In this case, it is recommended that you increase the synchronization interval time.

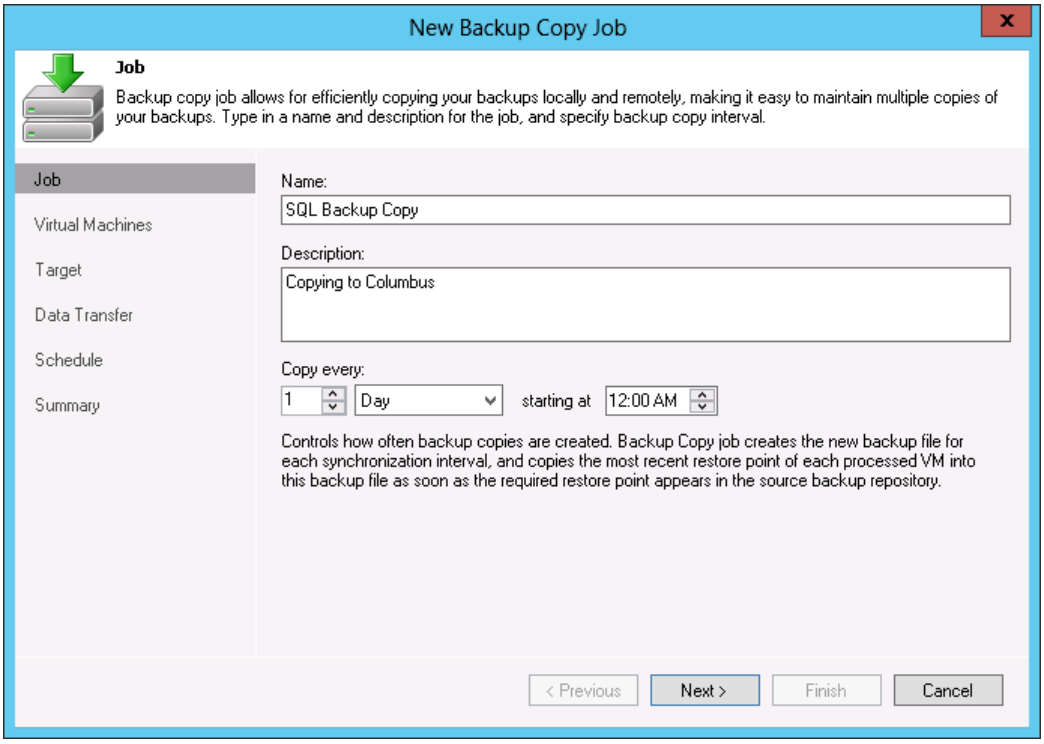

Step 3. Select VMs to Process

At the **Virtual Machines** step of the wizard, you should select VMs whose restore points you want to copy to the target backup repository.

Click **Add** and select the VM(s) that you want to process with the created backup copy job. You can browse VMs in the following sources:

- **From Infrastructure**. Using this option, you can browse the virtual infrastructure to add single VMs or VM containers to the job. When a backup copy job is run, Veeam Backup & Replication will search for restore points of selected VM(s) in all backup repositories connected to Veeam Backup & Replication. You can limit the search scope by selecting only specific repositories for the backup copy job.
- **From Backup**. Using this option, you can select VMs from available backups. When a backup copy job is run, Veeam Backup & Replication will search for restore points of selected VM(s) in all backups created on the Veeam backup server. You can limit the search scope by selecting only specific repositories for the backup copy job.

- **From Jobs**. Using this option, you can select VMs from available backup jobs. When a backup copy job is run, Veeam Backup & Replication will search for restore points of selected VM(s) in backups created for the selected jobs.

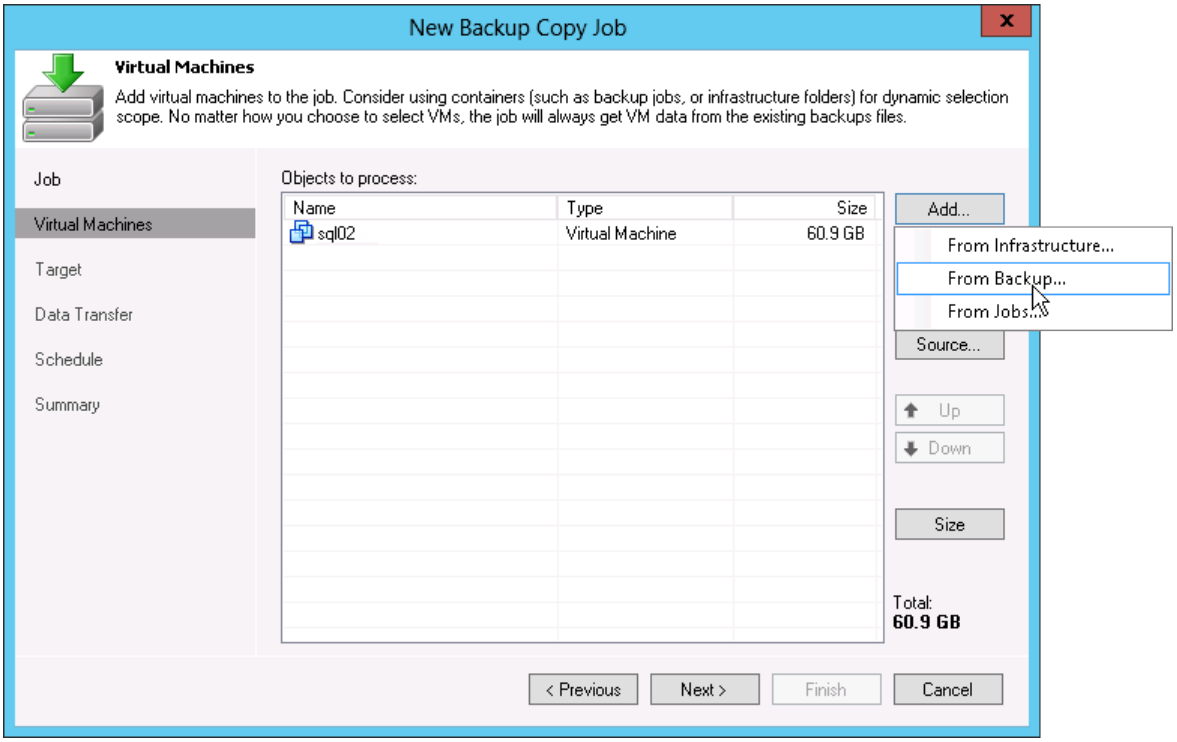

Step 4. Exclude Objects from the Backup Copy Job

If you have added VM containers to the list of processed VMs, you can specify which objects should be excluded from the backup copy job.

1. At the **Virtual Machines** step of the wizard, click **Exclusions**.
2. Click **Add** on the right and select the object that should be excluded.

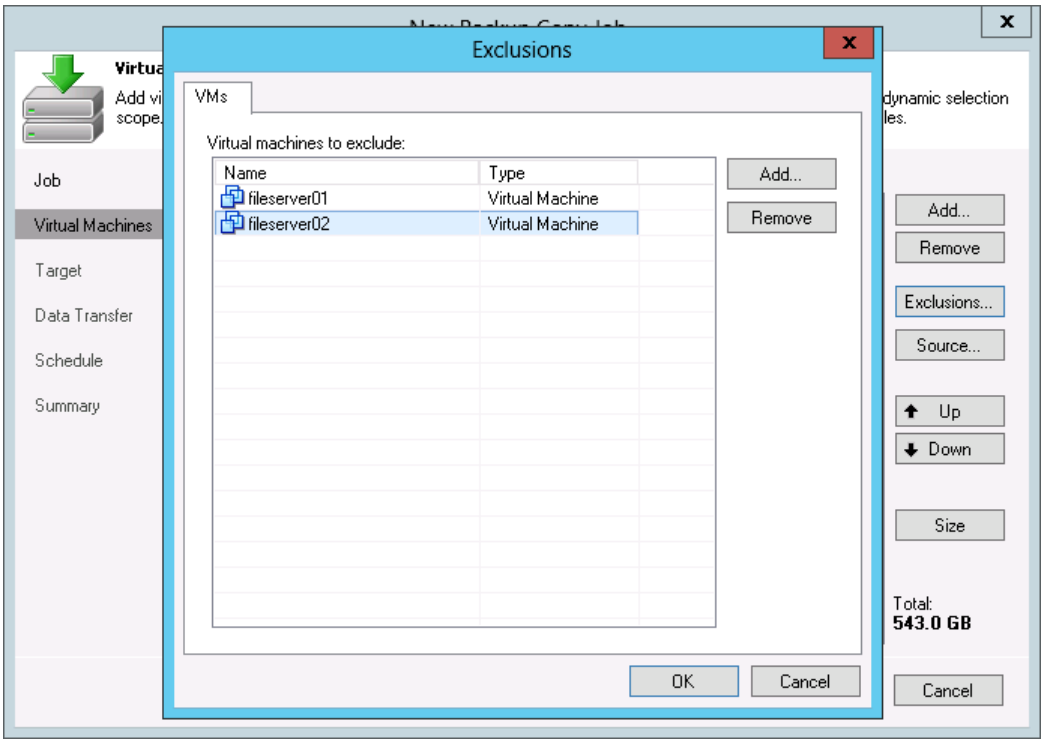

Step 5. Select Backup Repositories

Bu default, Veeam Backup & Replication searches for restore points in all backup repositories connected to the Veeam backup server. However, you can select backup repositories in which Veeam Backup & Replication should search for restore points of selected VM(s).

1. At the **Virtual Machines** step of the wizard, click **Source**.
2. Choose backup repositories in which restore points should be searched for. You can select all backup repositories connected to the Veeam backup server or define specific backup repositories.

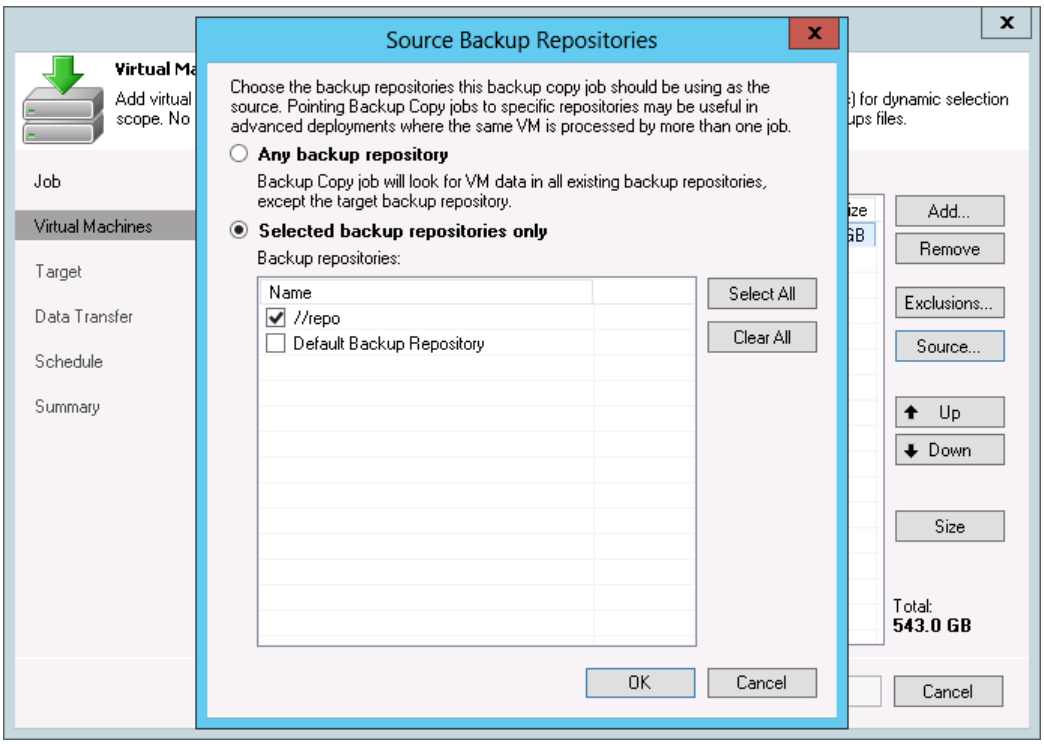

Step 6. Define VM Processing Order

If you want to copy restore points of some VMs before others, you can define the order in which the backup copy job must process VMs. VM copy order can be helpful, for example, if you want to ensure that the backup copying process does not overlap other scheduled activities or is completed before certain time.

To define the VM backup copy order:

1. Select the necessary VM in the list.
2. Move the VM up or down in the list using the **Up** and **Down** buttons on the right. In the same manner, you can set the processing order for VM containers added to the list.

Note VMs inside the container are processed at random. To ensure that VMs are processed in the defined order, you should add them as standalone VMs, not as a part of the container.

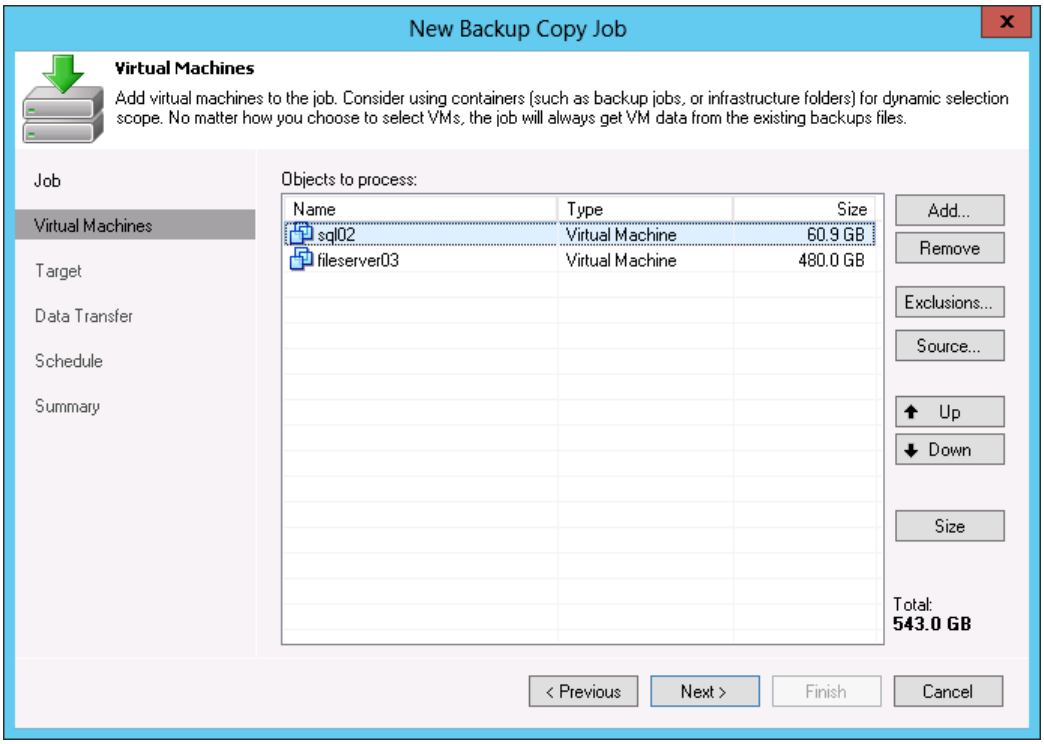

Step 7. Define the Backup Copy Target

At the **Target** step of the wizard, you should define the target backup repository for the backup copy job and define retention policy settings.

1. From the **Backup repository** list, select a backup repository in which copied restore points should be stored. When you select a target backup repository, Veeam Backup & Replication automatically checks how much free space is available on it. Make sure that you have enough free space to store copied backups.
2. In the **Restore points to keep** field, specify the number of restore points that should be retained on the target backup repository. When this number is exceeded, the earliest restore point will be removed from the backup chain. To learn more, see Simple Retention Policy.
3. To use the GFS (Grandfather-Father-Son) retention scheme, select the **Keep full backups for archival purposes** check box. In the fields below, define the number of daily, weekly, monthly, quarterly and yearly full intervals for which backups should be retained. Use the **Schedule** button to define the time schedule by which GFS full backups should be created. To learn more, see GFS Retention Policy.

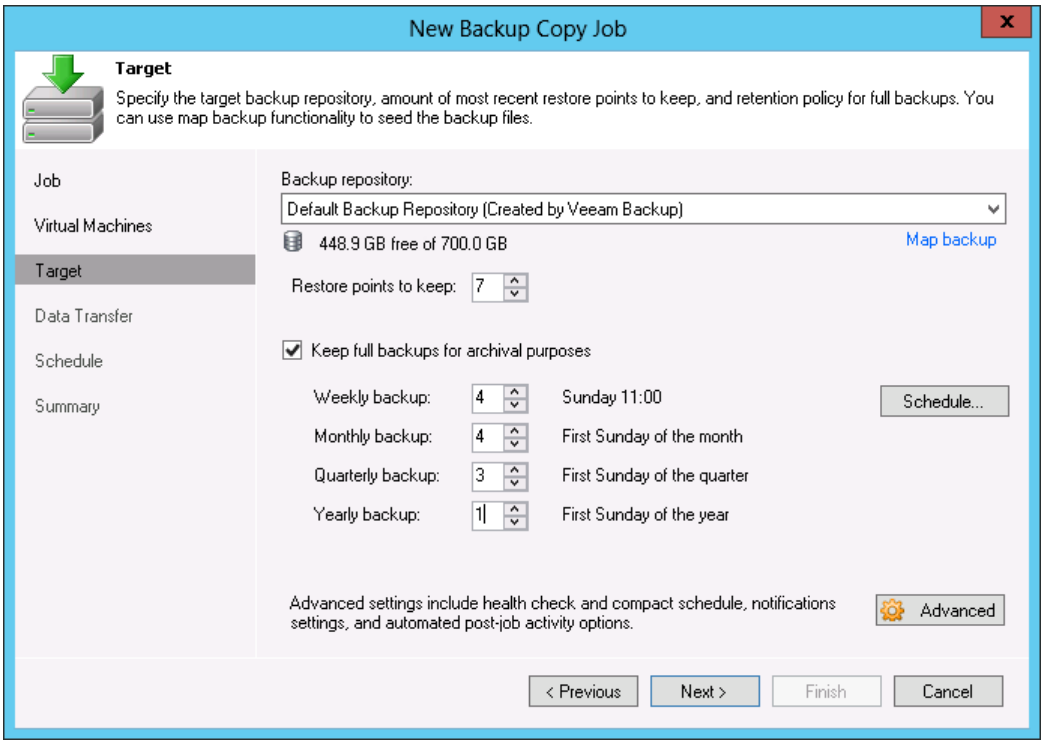

Step 8. Map a Backup File

If you plan to copy VM restore points over the WAN and slow connections, you can use backup mapping.

Backup mapping can only be used if you already have a full backup file for the processed VM on the target backup repository. In this case, you can point the backup copy job to this backup file. This full backup will be used as a "seed" for the backup copy job and you will need to transfer small incremental changes over the network. To learn more, see Mapping Backup Copy Jobs.

To map a backup copy job to the backup file:

1. Click the **Map backup** link.
2. Point the backup copy job to the necessary backup on the target backup repository. Backups stored on the target backup repository can be easily identified by backup job names. To facilitate search, you can also use the search field at the bottom of the window.

Important! The backup copy job can be mapped to the backup only if the backup chain you plan to use as a "seed" contains one restore point — a full backup file. If the chain contains a number of restore points, Veeam Backup & Replication will fail to map the backup copy job to the selected backup. To overcome this situation, you can create a backup "seed" by means of an auxiliary backup copy job on the target repository. To learn more, see Creating a Seed for the Backup Copy Job.

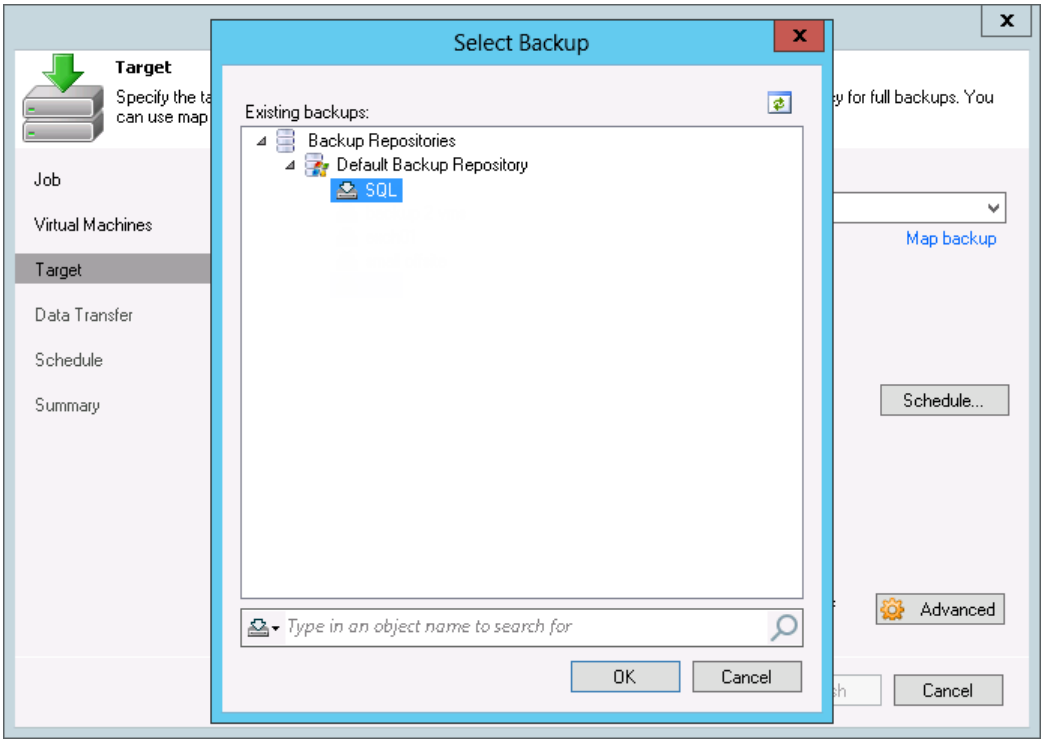

Step 9. Specify Advanced Settings

At the **Target** step of the wizard, click **Advanced** to specify advanced options for the backup copy job.

Backup Settings

On the **Backup** tab, specify advanced settings for the restore points that will be stored on the target backup repository.

1. In the **VM retention** section, define the time period for which restore points for deleted VMs should be stored. After a VM restore point has been copied and Veeam Backup & Replication has performed transform operations on the target side, Veeam Backup & Replication checks the list of VMs included in the job. If a VM is no longer available, for example, it was deleted from the virtual infrastructure or moved to another location, Veeam Backup & Replication will keep its data on the target backup repository for the specified period. When this retention period is over, data of the deleted VM will be removed from the target backup repository.
2. If you want to periodically perform a health check of the most recent restore point, select the **Health check** check box and specify the time schedule for the health check. An automatic health check allows you to avoid a situation when a restore point gets corrupted, making all further increments corrupted, too. If Veeam Backup & Replication detects corrupted data blocks in the restore point during the health check, it will transfer these data blocks to the target backup repository during the next synchronization interval and store them in the newly copied restore point. By default, the health check is performed on the last Sunday of every month. To learn more, see Health Check for Copied Backups.
3. To periodically compact a full backup, select the **Compact full backup periodically** check box and specify the schedule for the compacting operation. Note that this option can be enabled only if you have not specified the GFS settings.

During the compacting operation, Veeam Backup & Replication creates a new empty VBK file and copies to it all data blocks from the full backup file. As a result, the full backup file gets defragmented, its size reduces and the speed of writing and reading to/from the file increases. To learn more, see Compacting a Full Backup File.

Note The **Deleted VM data retention period** option is applied only for regular backup chains. Veeam Backup & Replication does not remove deleted VMs from weekly, monthly, quarterly and yearly backups.

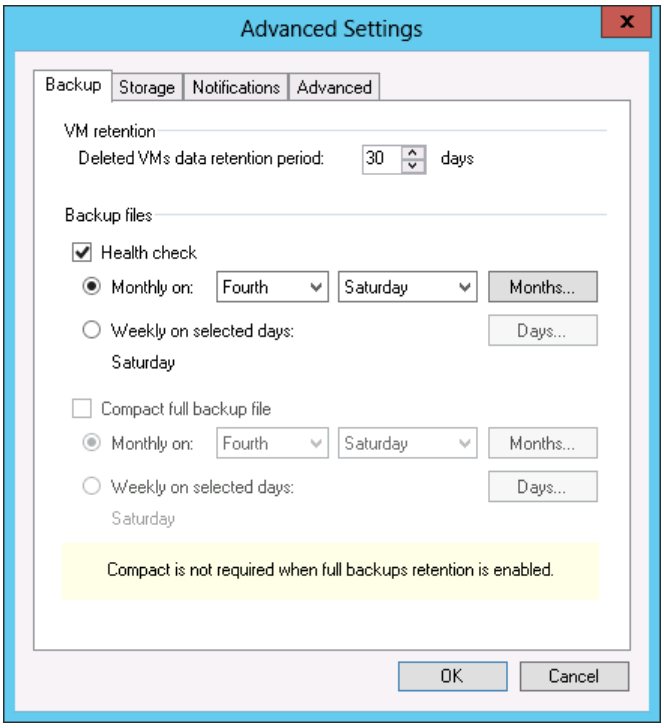

Storage Settings

On the **Storage** tab, specify compression and deduplication settings for the backup copy job.

1. By default, Veeam Backup & Replication performs deduplication before storing VM data on the target backup repository. Deduplication provides a smaller size of the resulting backup file but may reduce the job performance.
   You can disable deduplication at all by clearing the **Enable inline data deduplication** check box.
2. In the **Compression** section, specify a compression level to be used: *None*, *Dedupe-friendly*, *Optimal*, *High* or *Extreme*.

To learn more, see Compression and Deduplication.

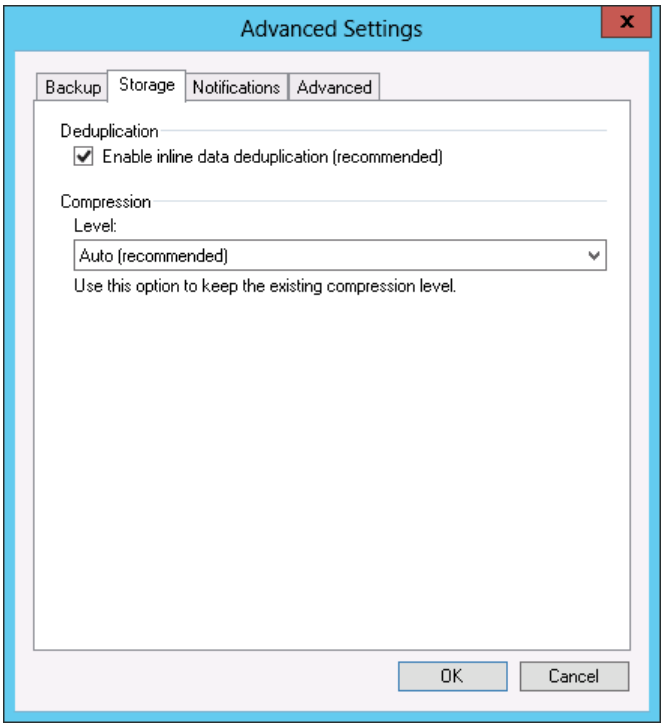

Notification Settings

On the **Notifications** tab, specify notification settings for the backup copy job.

1. Select the **Send email notifications to the following recipients** check box if you want to receive notifications informing about the backup copy job results by email.
   In the field below, specify a recipient's email address. You can enter several addresses separated by a semicolon.
   Veeam Backup & Replication sends a consolidated email notification once for the specified synchronization interval. Even if the synchronization process is started several times within the interval, for example, due to retries, only one email notification will be sent.
   Email notifications can be sent only if you configure general email notification settings in Veeam Backup & Replication. To learn more, see Specifying E-Mail Notification Settings.
2. Select the **Enable SNMP notification for this job** check box if you want to receive SNMP traps when the backup copy job completes. SNMP traps can be sent only if you configure SNMP settings in Veeam Backup & Replication and on the recipient's computer. For details, see Specifying SNMP Settings.

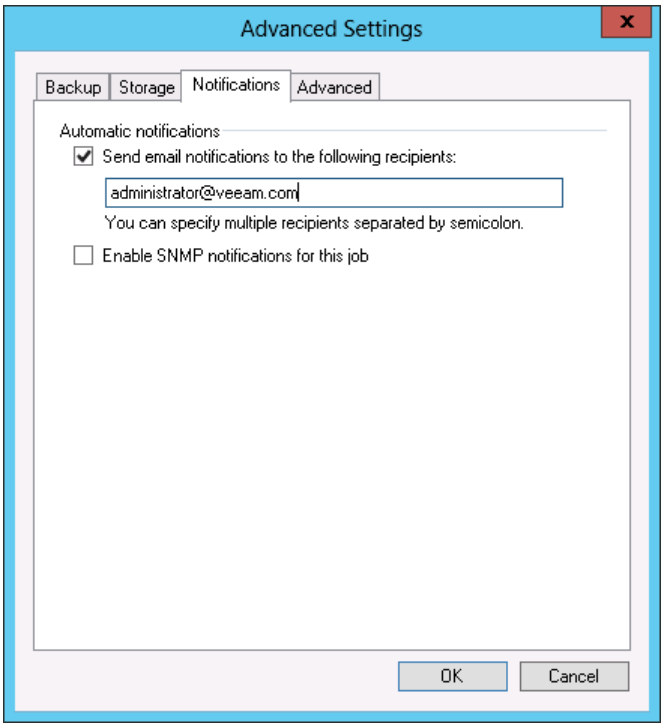

Advanced settings

On the **Advanced** tab, specify miscellaneous advanced settings for the backup copy job.

1. Select the **Run the following command** check box if you want to execute post-job actions.
2. In the field below, specify a path to an executable script file.
3. You can select to execute post-job actions after a number of synchronization intervals or on specific week days:
   - If you select the **Run every... backup cycle** option, specify the number of synchronization intervals after which post-job actions should be performed.
   - If you select the **Run on selected days only** option, click **Days** and specify week days on which post-job actions should be performed.

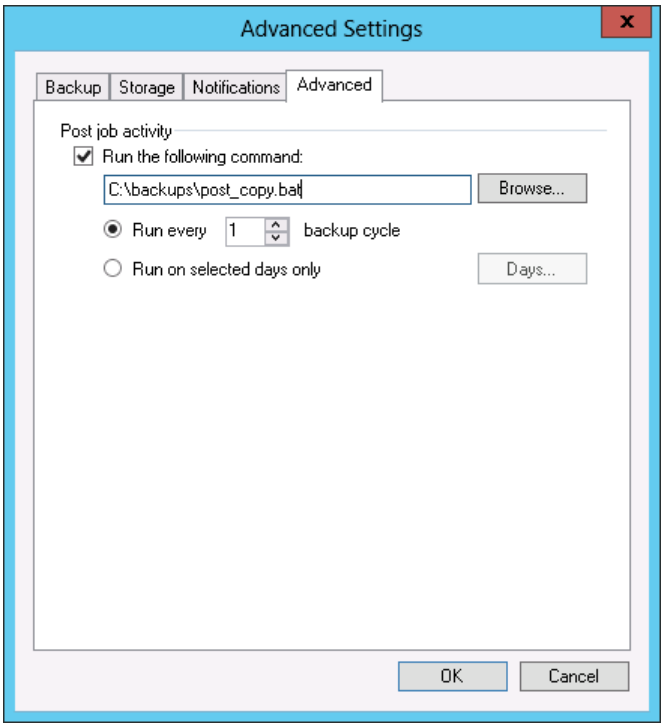

Step 10. Specify WAN Optimization Settings

By default, during the backup copy job Veeam Backup & Replication transports VM data directly from the source backup repository to the target backup repository. This type of transport is recommended if you plan to copy backup files over fast connections.

However, if you plan to transport backup files over the WAN or slow connections, it is recommended to configure a pair of WAN accelerators in your backup infrastructure and copy VM backups via these WAN accelerators. WAN accelerators perform global data deduplication, eliminating the need to transport redundant blocks of data and reducing the load on the WAN. To learn more, see WAN Acceleration.

Important! The WAN optimization option is available in the Enterprise Plus Edition of Veeam Backup & Replication.

To use WAN acceleration for the backup copy job:

1. Select the **Through built-in WAN accelerators** option.
2. From the **Source WAN accelerator** list, select the WAN accelerator configured on the source site.
3. From the **Target WAN accelerator** list, select the WAN accelerator configured on the target site.

Be extremely careful when assigning WAN accelerators to the backup copy job. If you make a mistake and assign the WAN accelerator in the target site to be used as the source one, VM data will go in the backward direction and the load on the WAN will increase.

To learn more, see Adding WAN Accelerators.

Important! You cannot assign one source WAN accelerator to several backup copy jobs that you plan to run simultaneously. The source WAN accelerator requires a lot of CPU and RAM resources and cannot be shared by a number of backup copy jobs. As an alternative, you can create one backup copy job for all VMs you plan to process via one source WAN accelerator.
The target WAN accelerator, however, can be assigned to several backup copy jobs.

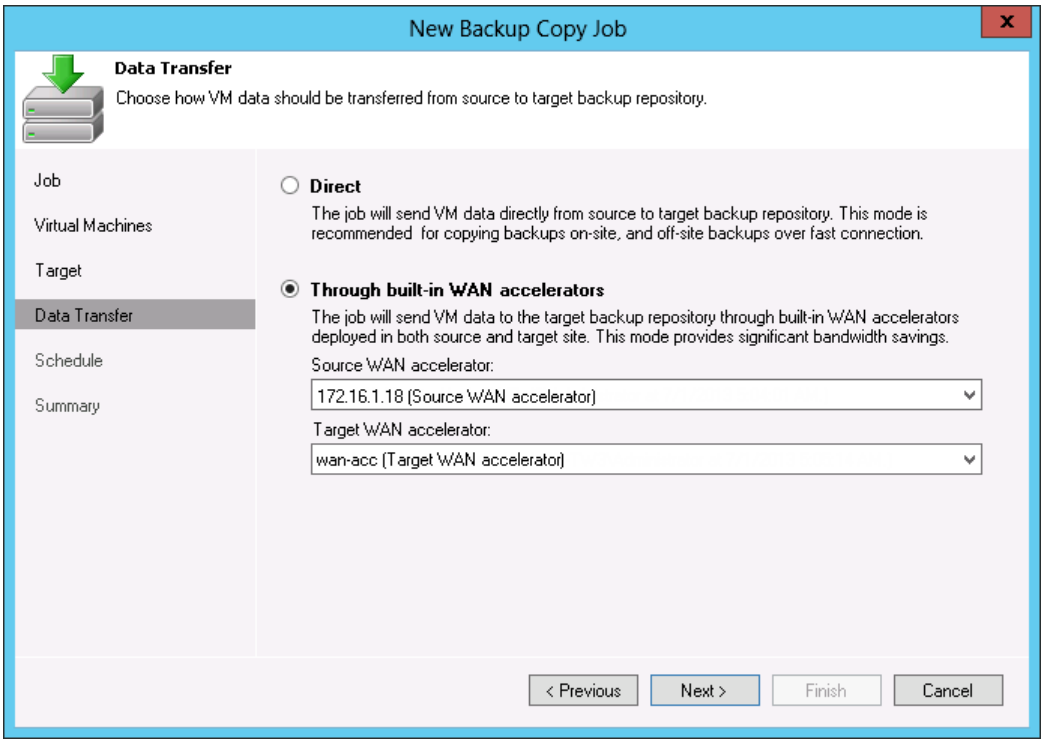

Step 11. Define the Backup Copy Window

At the **Schedule** step of the wizard, you can define the time span in which the backup copy job must not transport data over the network. To learn more, see Backup Copy Window.

To define a backup window for the backup copy job, do the following:

1. Select the **During the following time periods only** option.
2. In the schedule box, select the desired time area.
3. Use the **Enable** and **Disable** controls to mark the selected area as allowed or prohibited for the backup copy job.

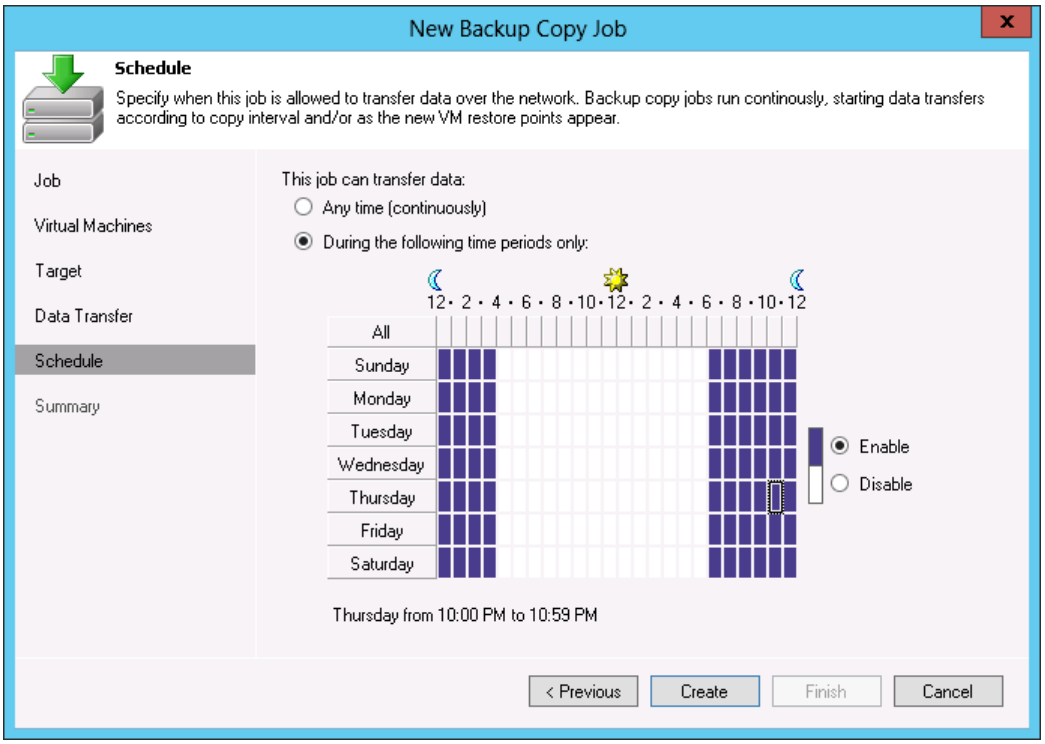

Step 12. Finish Working with the Wizard

After you have specified schedule settings, click **Create**. Select the **Enable the job when I click Finish** check box if you want to start the created backup copy job right after you complete working with the wizard. Click **Finish** to close the wizard.

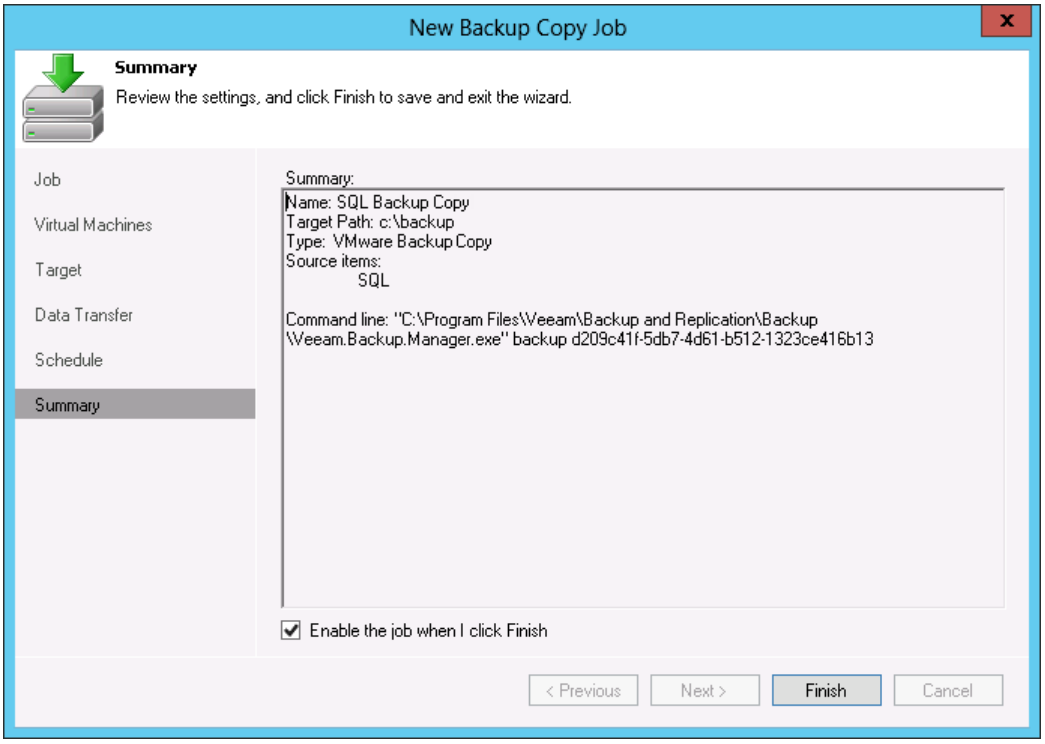

Linking Backup Jobs to Backup Copy Jobs

Veeam Backup & Replication provides an option for linking backup jobs to backup copy jobs. This option lets you automatically create a second instance of the backup file in some other location.

When you link a backup job to the backup copy job, Veeam Backup & Replication automatically updates properties of the corresponding backup copy job and includes to it the backup job as a source of data. As a result, the backup copy job starts monitoring the backup job linked to it. At every synchronization interval, the backup copy job checks the source backup repository for new restore points. As soon as a backup job session is finished and a new restore point is created, the backup copy job automatically copies this restore point to the target backup repository.

You can point a backup job to an existing backup copy job using the **Backup Job** wizard. Perform the following steps:

1. Open the backup job settings and navigate to the **Storage** step. Select the **Configure secondary destination for this job** check box.

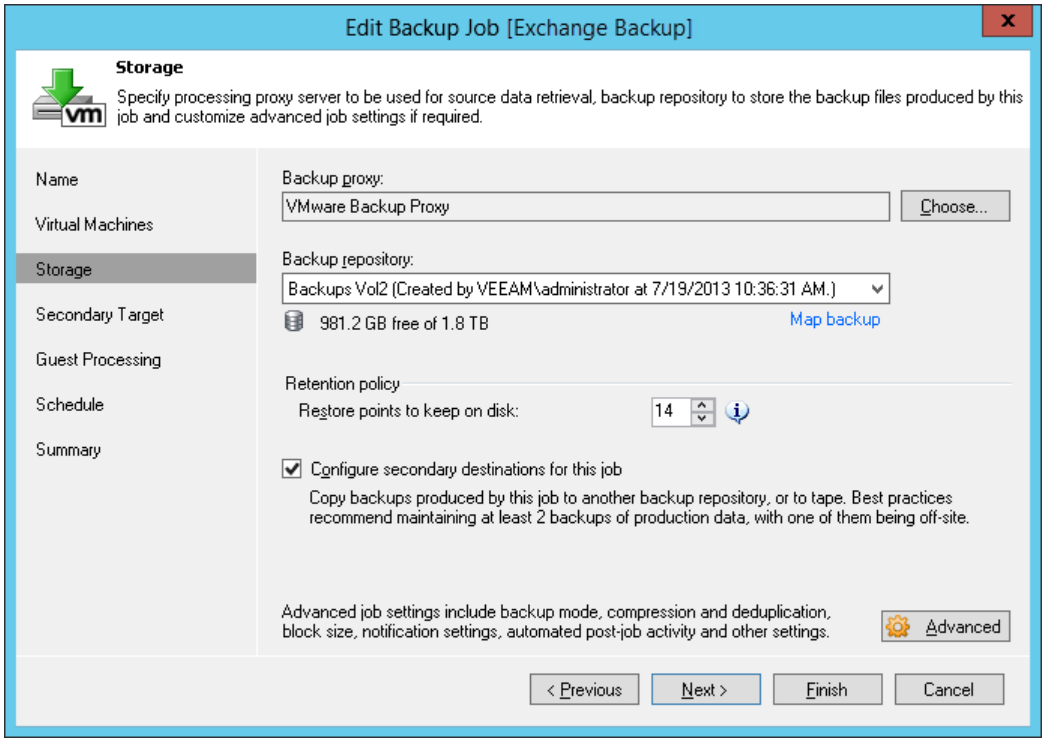

2. At the **Secondary Target** step, click **Add** and choose a backup copy job to which the backup job should be linked. Note that the backup copy job must be already configured by the moment you link it to the backup job.

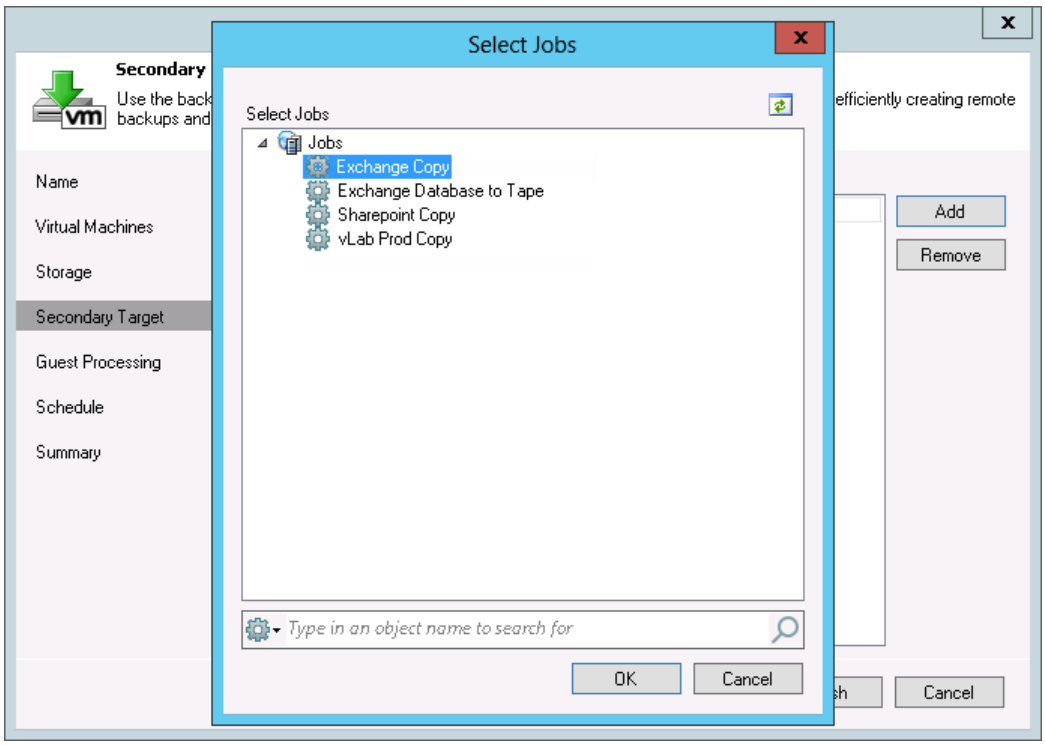

Starting the Synchronization Cycle Manually

As soon as you create a backup copy job and start it, Veeam Backup & Replication will automatically launch it. Data synchronization will be performed automatically according to the specified synchronization interval. To learn more, see Synchronization Interval.

However, you can start the synchronization process manually. This can be helpful, for example, if a new restore point has already been created on the source backup repository but the previous synchronization interval has not yet elapsed.

To start a new data synchronization cycle manually:

1. Open the **Backup & Replication** view.
2. In the inventory pane, select the **Backup Copy** node under **Backup**.
3. In the working area, right-click the backup copy job and select **Sync Now**. Alternatively, you can click the **Sync Now** button on the ribbon.

Note When you manually start the synchronization process, Veeam Backup & Replication creates a new synchronization interval.

- In case of backup copy jobs with minutely and hourly intervals, this synchronization interval is equal to those that are created automatically by the schedule. As a result, the start time of backup copy processing shifts forward.
- In case of backup copy jobs with the daily synchronization intervals, the manual synchronization process must always be finished before a new scheduled synchronization interval begins. As a result, when a new synchronization interval starts by the schedule, Veeam Backup & Replication forces to finish the manual synchronization process, even if it has not completed yet.

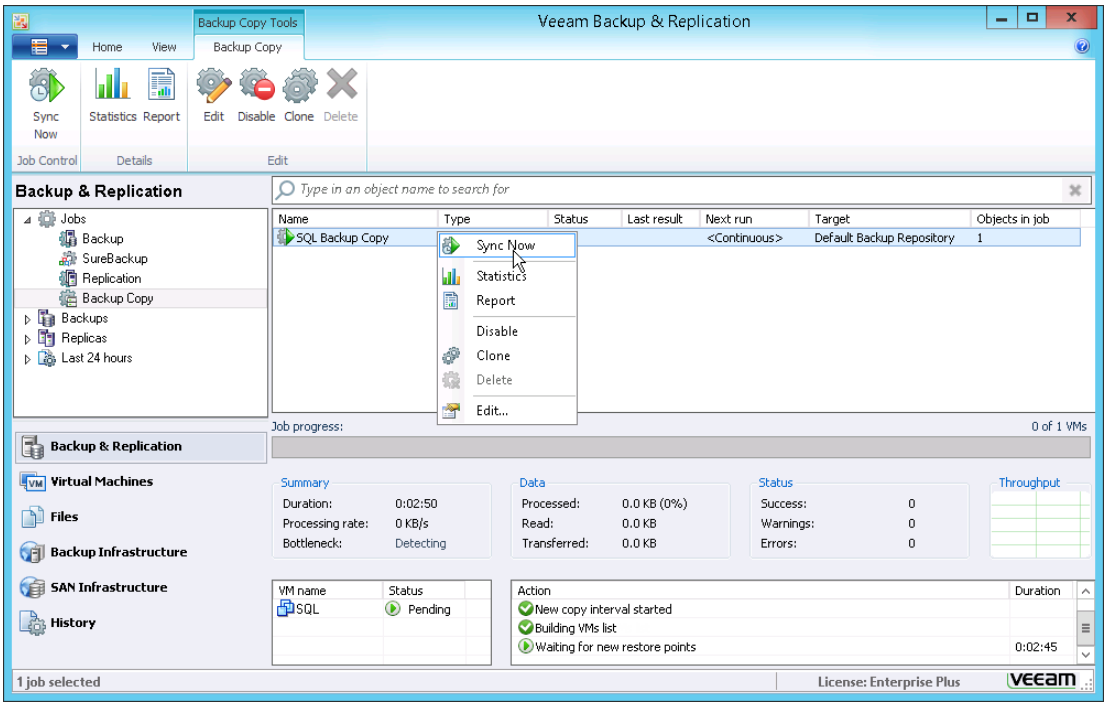

Disabling and Removing Backup Copy Jobs

If you want to put data synchronization on hold, you can disable a backup copy job. The disabled job is not deleted from the console, it is simply stopped for some period of time. You can enable a disabled job at any time later.

To disable a job:

1. Open the **Backup & Replication** view.
2. In the inventory pane, select the **Backup Copy** node under **Backup**.
3. In the working area, right-click the job and select **Disable**.

To enable a disabled job, right-click it in the list and select **Disable** once again.

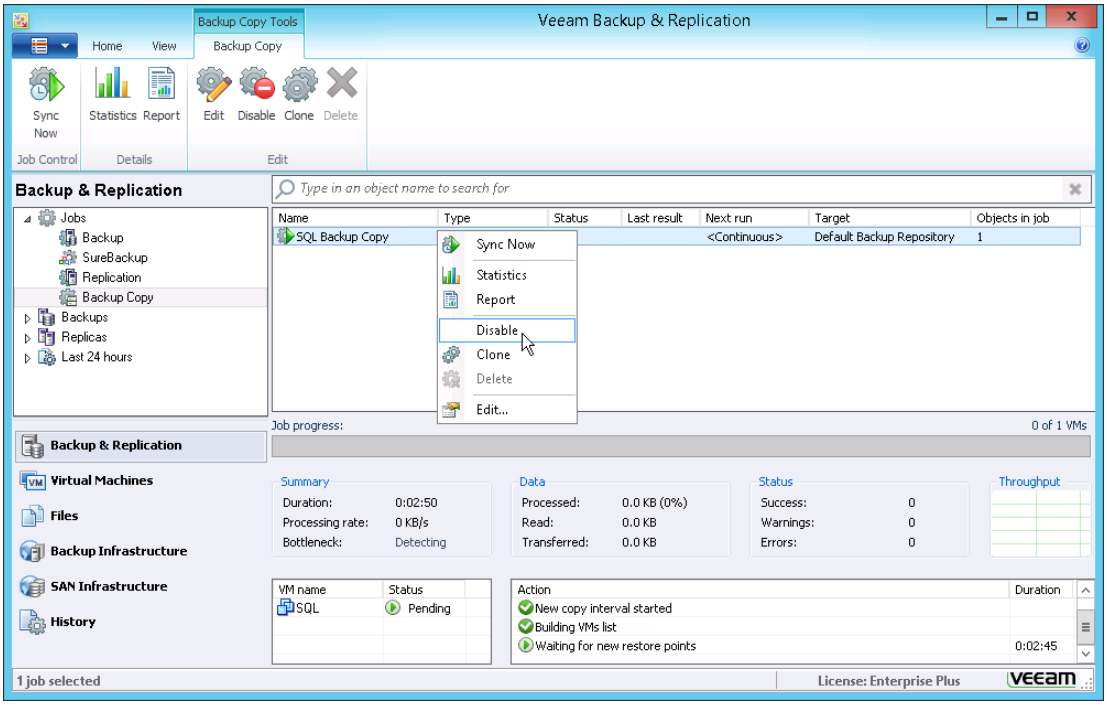

If you want to permanently remove a backup copy job, you first need to stop the synchronization process by disabling the job.

To remove a job:

1. Disable the backup copy job as described above.
2. In the working area, right-click the job and select **Remove**.

As a result, the backup copy job will be removed from the Veeam Backup & Replication console and from the Veeam Backup & Replication database.

Removing Backups from the Target Repository

You can remove backups created by the backup copy job from the Veeam Backup & Replication console and permanently delete backup chains from the target backup repository.

To remove a backup from the console:

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

Veeam Backup & Replication will remove the backups from the console. The backup files will still remain on the target backup repository and the backup copy job will be available in the list of jobs.

To permanently remove backup chains from the target backup repository:

1. Open the **Backup & Replication** view.
2. In the inventory pane, select **Backups**.
3. In the working area, right-click the necessary backup copy job and select **Remove from disk**.
4. If you want to remove all weekly, monthly, quarterly and yearly backups, select the **Delete archived full backups** check box.

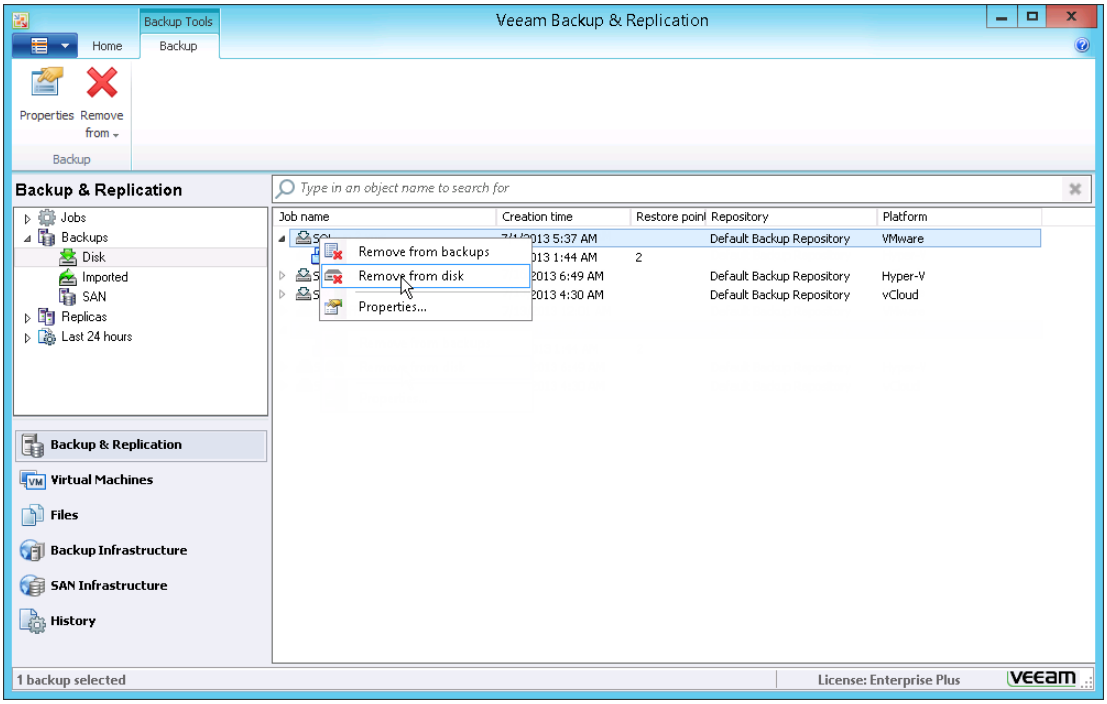

Working with Tape Media

Veeam Backup & Replication allows working with to tape devices that are directly attached to the Veeam backup server. Tape devices can be connected over Fibre Channel (FC), Serial Attached SCSI (SAS), SCSI. You can also use the Microsoft iSCSI initiator on the Veeam backup server to connect to the tape device on a remote server via iSCSI.

Both physical and virtual tape libraries and standalone drives are supported.

Prerequisites

To ensure tape device visibility to the Veeam Backup server and to manage tape media from the Veeam Backup & Replication console, it is recommended that you take these steps prior to installing Veeam Backup & Replication:

1. Enable the connection between a tape device and the Veeam Backup server.
   Veeam Backup & Replication allows working with tape devices that are directly attached to the Veeam Backup server. Tape devices can be connected over Fibre Channel (FC), Serial Attached SCSI (SAS), SCSI. You can also use the Microsoft iSCSI initiator on the Veeam Backup server to connect to the tape device on a remote server via iSCSI.
2. Install an appropriate device driver on the machine where Veeam Backup & Replication is installed.
   Please consider that only OEM (original equipment manufacturer) drivers are supported; make sure you have the latest driver version available.
   Also, if multiple driver installation modes are supported for your storage device, make sure the driver is installed in the mode that allows for multiple open handles from a host to a drive to exist at the same time.
   For example, if installing a driver for IBM System Storage TS3200, you should use the *install_nonexclusive.exe* installer as described in the product Readme.

Please refer to your storage system manufacturer recommendations on choosing the appropriate setup option.

Getting Started with Tapes

To start working with tapes in Veeam Backup & Replication, you need to complete the following steps:

1. Make sure that Veeam Backup & Replication is installed on a machine with directly attached tape devices connected over Fibre Channel (FC), Serial Attached SCSI (SAS), SCSI, or remotely via iSCSI. Check that all required tape device drivers are installed.

   When you start Veeam Backup & Replication, it will perform auto-discovery — it will scan attached tape devices and will display all discovered tape libraries and tape drives under the **Tape** node in the **Backup Infrastructure** view. Afterward, the auto-discovery process will be performed periodically every 30 seconds.

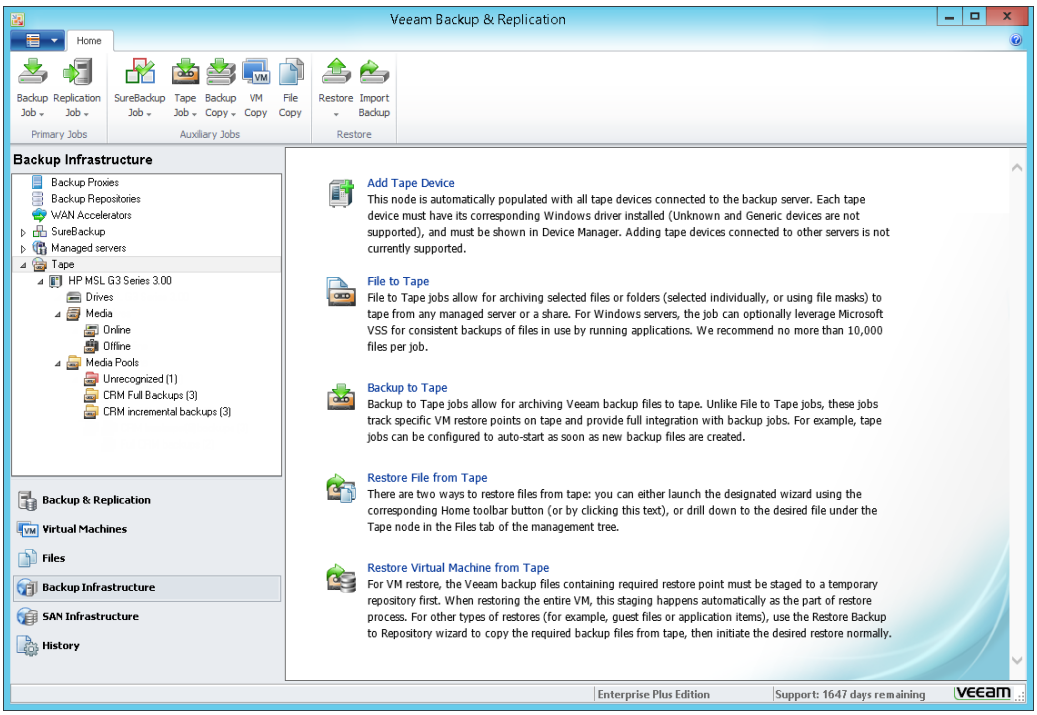

2. Load tapes to the tape device (if not yet loaded). All newly loaded tapes will be available in the **Unrecognized** media pool. To prepare tapes for data archiving and restore, do the following:
    - If you are using empty tapes, mark these tapes as free to move them to the **Free** media pool.
    - If you are using non-empty tapes with outdated contents that can be erased, you can either erase the tapes or mark them as free to move the tapes to the **Free** media pool.
    - If you are using non-empty tapes with backup content that should be preserved (for example, the tapes store backups of files that you want to restore), select these tapes and run the tape catalog job. Tape catalog job will scan the contents on tapes, move the tapes to the **Imported** pool and update the file catalog in the Veeam Backup & Replication database. After the tape catalog job finished, you will be able to restore files from tape.
3. Create one or more custom media pools that will be used as targets for backup to tape and files to tape jobs.
4. Configure and run backup to tape or files to tape jobs.

Managing Tape Media

Veeam Backup & Replication automatically discovers tape devices connected to the Veeam backup server and displays all discovered tape libraries and tape drives under the **Tape** node in the **Backup Infrastructure** view.

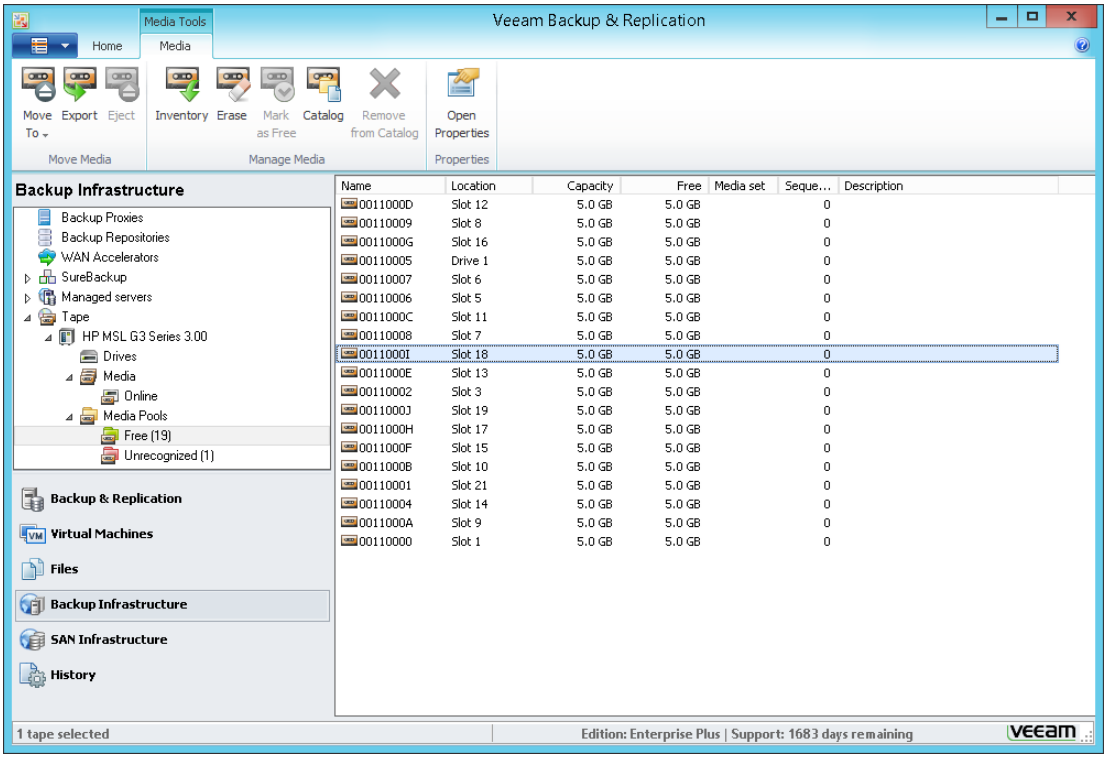

The following nodes in the **Tapes** hierarchy refer to physical entities operating within tape media:

- Tape Library
- Tape Drives
- Media (magnetic tapes)
- Media Pools (logical groups of tapes)

You can see detailed properties for each entity available in the **Tapes** hierarchy. To view properties, right-click the entity and choose **Properties**.

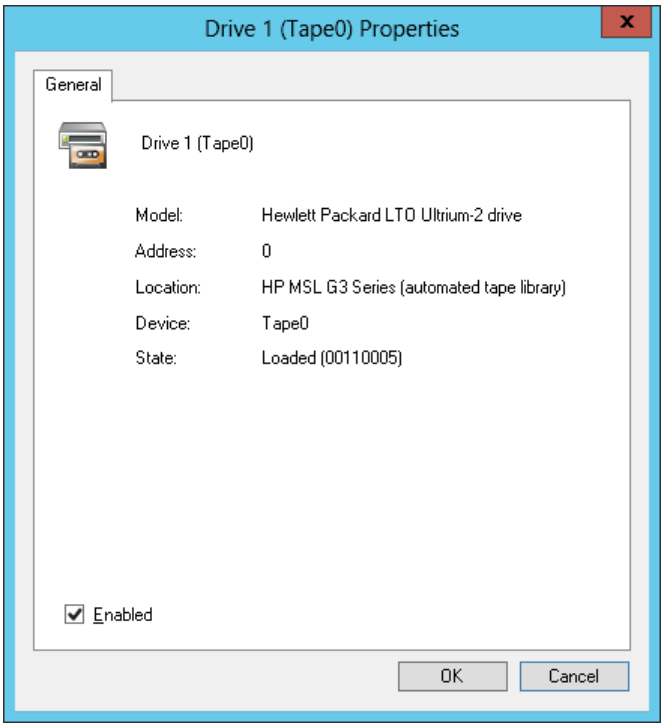

Working with Media Pools

All tape media are divided into media pools — logical groups of tapes. There are two types of media pools in Veeam Backup & Replication: predefined media pools and custom media pools.

Predefined media pools are service pools created by Veeam Backup & Replication. The following predefined media pools are available:

- **Free** — a media pool containing empty tapes
- **Unrecognized** — a media pool containing tapes that are not yet identified by the inventory or catalog job
- **Imported** — a media pool containing non-empty tapes; these are tapes identified by the tape inventory or catalog job
- **Retired** — a media pool containing retired tapes that reached the maximal number of re-writes. This media pool may also contain tapes with some mechanical breakdown.

Custom media pools serve as targets for backup to tape and files to tape jobs. Custom media pools describe media set and retention settings that are applied to all tapes in the pool. You can allocate to custom media pools a limited set of tapes, or create replenishable media pools.

To be able to configure backup to tape and files to tape jobs, you need to first create custom media pools.

Creating Custom Media Pools

To create a custom media pool, use the **New Media Pool** wizard. This section will guide you through all steps of the wizard and provide explanation on available options.

Step 1. Launch the New Media Pool Wizard

To run the **New Media Pool** wizard, do either of the following:

- Open the **Backup Infrastructure** view, expand the **Tape** > *LibraryName* node and select the **Media Pools** node. Click **Add Media Pool** on the ribbon.
- Open the **Backup Infrastructure** view, expand the **Tape** > *LibraryName* node. Right-click the **Media Pools** node and choose **Add Media Pool**.

Tip You can also launch the **New Media Pool** wizard when configuring archiving jobs (that is, directly from the **New Backup To Tape Job** wizard and **New File To Tape Job** wizard. For more details, see Creating Backup to Tape Jobs and Creating File to Tape Jobs.

Step 2. Specify Media Pool Name

At the **Name** step of the wizard, you should define basic description for the new media pool.

1. In the **Name** field, enter a name for the created media pool.
2. In the **Description** field, enter a description of the new media pool. The default description contains information about the user who created the media pool, date and time when the media pool was created.

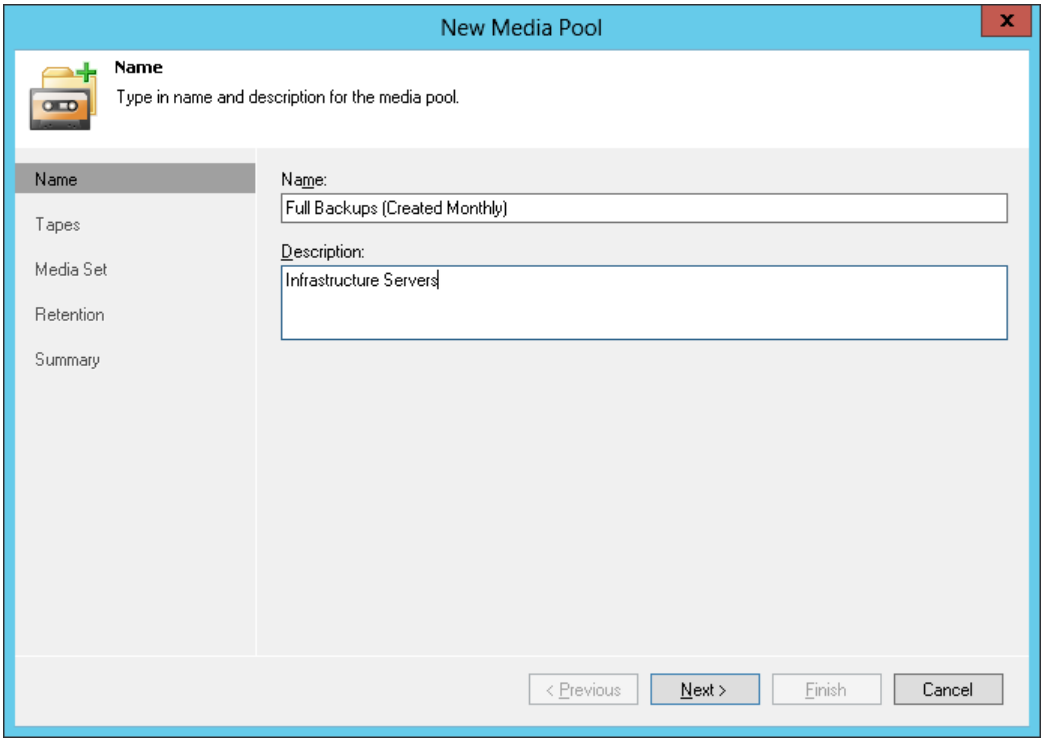

Step 3. Add Tapes to Media Pool

At the **Tapes** step of the wizard, you should allocate tapes for the pool.

1. From the **Tape library** list, select the library from which the tapes will be allocated for the media pool. Note that this option is available only if you have launched the **New Media Pool** wizard directly from the **New Backup To Tape Job** wizard or **New File To Tape Job** wizard. Otherwise, the option is disabled.

2. To allocate specific tapes from the library, click the **Add** button on the right and select tapes that should be added to the media pool. Allocated tapes will be reserved for the created media pool; other custom media pools will not be able to use these tapes. The capacity and free space on the allocated tapes will be displayed in the bottom right corner.
3. To make the media pool replenishable, select the **Add tapes from Free media pool automatically when more tapes are required** check box.
With this option enabled, additional tapes will be allocated from the **Free** media pool when needed. That is, when a backup to tape or files to tape job uses all available tapes from this media pool, Veeam Backup & Replication will automatically add the required number of tapes from the **Free** media pool to let the job complete. If the option is disabled, the job will pause and prompt the backup administrator to add new tapes to the media pool.

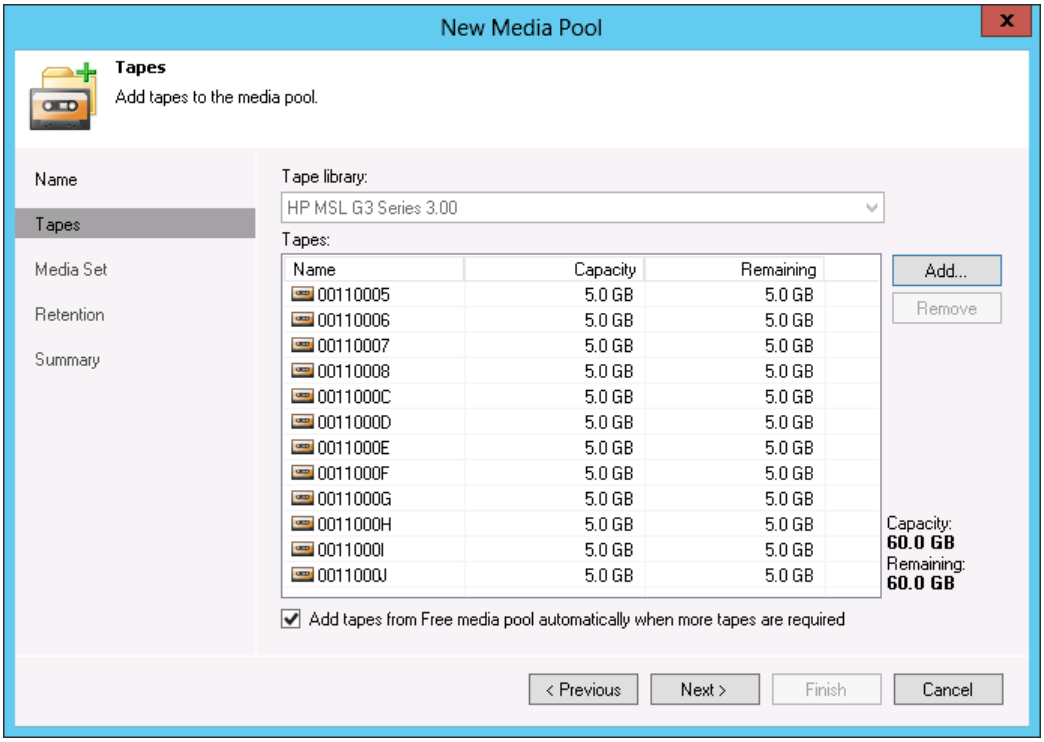

Step 4. Specify Media Set Options

For each media pool, you should specify how new media sets are created. A *media set* is consequent data stream that can span several tapes (for example, a weekly backup stored on tapes). At the **Media Set** step of the wizard, specify how this data stream will be organized.

1. In the **Media set name** field, define the pattern according to which created media set(s) will be named.
2. In the **Automatically create new media set** section, specify conditions for creating new media sets on tapes allocated to the media pool. The following options are available:
   - **Do not create, continue using the current media set**. If this option is selected, each subsequent backup session will write its backup set to the existing media set: it will append backup content to the content that was written to tape with a previous backup session. If, however, a backup set is started with a new tape, Veeam Backup & Replication will create a new media set for it.
   - **Create a new media set for every backup session**. If this option is selected, a new media set will be created for each new backup session. Each backup session will write its backup set starting with a new tape.

- **Daily at**. If this option is selected, you can specify day and time when new media sets should be created. For example, if at the end of the week you send weekly media sets to offsite storage, you can schedule creation of new media sets at the beginning of each week.

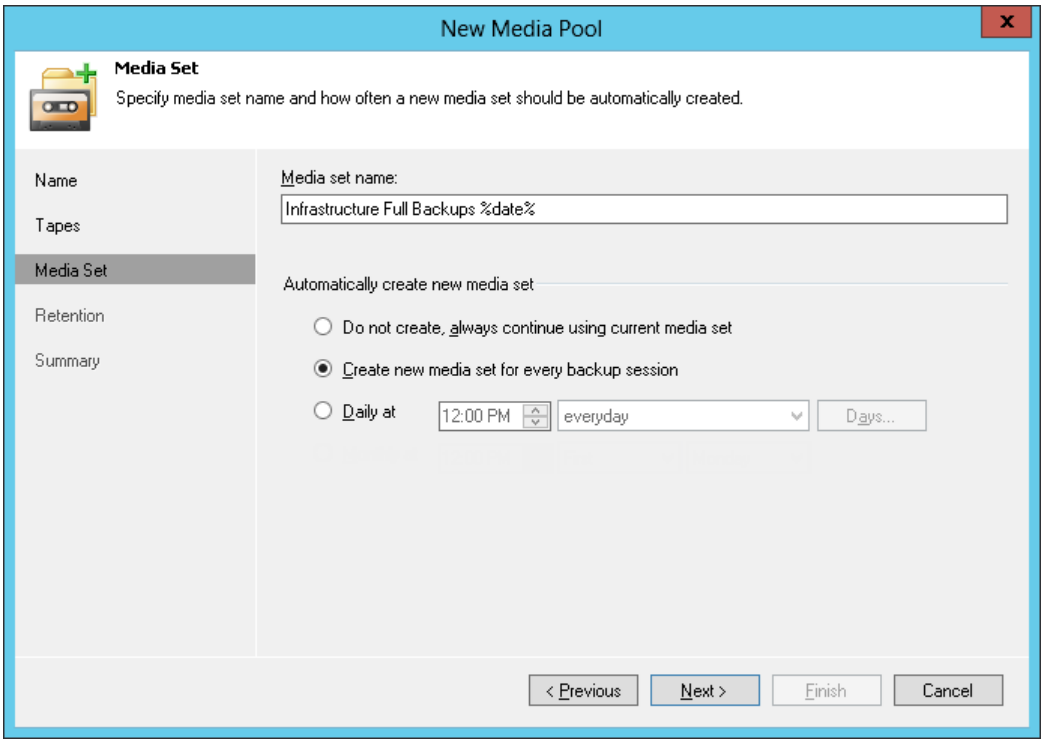

Step 5. Specify Retention Settings

At the **Retention** step of the wizard, specify overwrite rules for cases when all tapes allocated to the media pool are full (and there are no more free tapes available). You can select one of the following options:

- **Do not protect data (cyclic ally overwrite tapes as required)**. If this option is selected, tapes allocated to the pool will be overwritten, starting with the tape that stores the oldest archive.
- **Protect data for (time interval)**. If this option is selected, archives on tapes will be preserved for the specified period. When this period is over, data will be overwritten, starting with the tape that stores the oldest archive.
- **Never overwrite data.** If this option is selected, data on tapes will not be overwritten. If there is not enough tape capacity for the archiving job to complete, Veeam Backup & Replication will pause the job and prompt the backup administrator to add new tapes to the media pool.

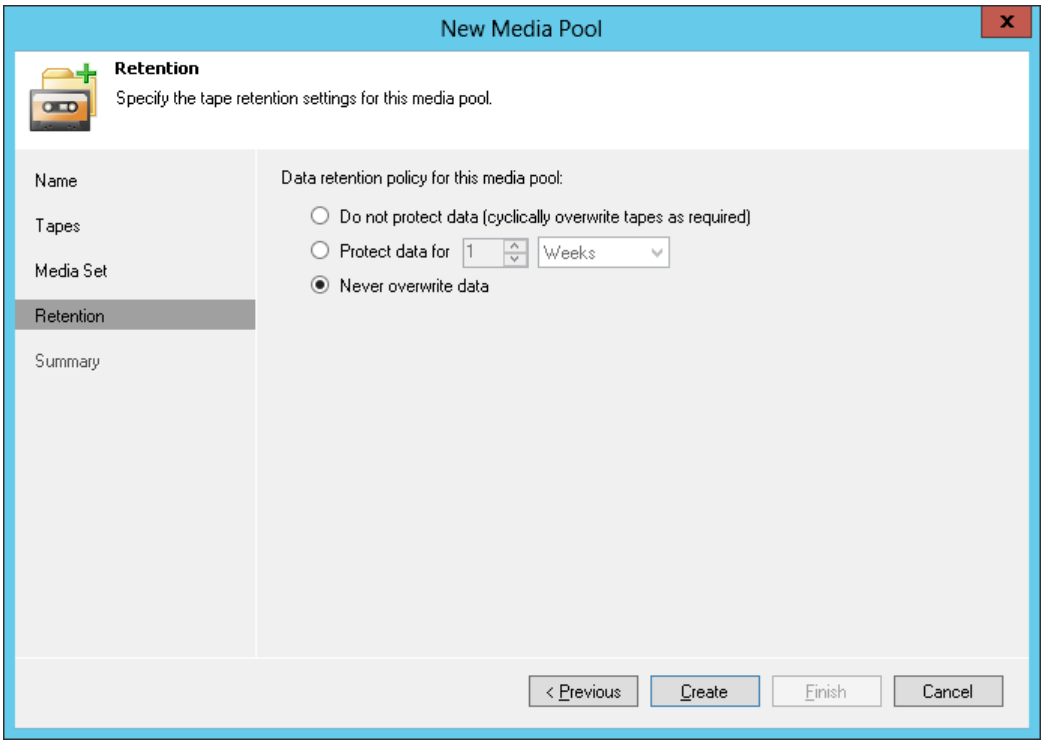

Step 6. Finish Working with the Wizard

Review the media pool settings and click **Finish** to complete the wizard.

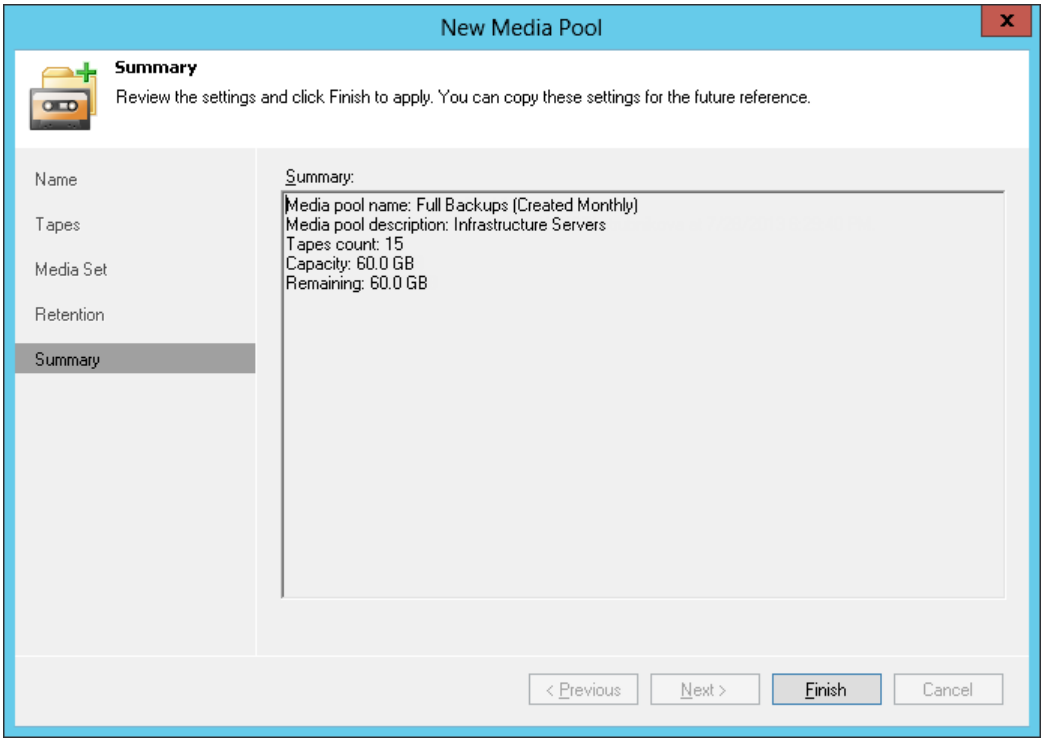

A new media pool will be available under the **Tape** > *LibraryName* > **Media Pools** node in **Backup Infrastructure** view.

Modifying Media Pools

If necessary, you can modify settings of a media pool. Note that you can only change custom media pools; predefined media pools cannot be modified.

To modify media pool settings:

1. Open the **Backup Infrastructure** view.
2. Right-click the necessary media pool and choose **Properties**. Alternatively, select a media pool and click **Edit Media Pool** on the ribbon.
3. Go through the **Media Pool Wizard** to change the necessary settings.
4. Apply changes.

Deleting Media Pools

If you no longer need a media pool, you can delete it. Mind the following limitations:

- You can only delete custom media pools; predefined media pools cannot be deleted.
- You cannot delete a media pool that is used in a backup to tape or files to tape job. To be able to delete such a pool, first point corresponding jobs to other custom media pools.
- You cannot delete a media pool that contains tapes. To be able to delete such a pool, first move tapes from this pool to other media pools.

To delete a media pool:

1. Open the **Backup Infrastructure** view.
2. Right-click the necessary media pool and choose **Remove media pool** from the shortcut menu. Alternatively, select a media pool and click **Remove Media Pool** on the ribbon.
3. In the displayed dialog box, click **OK** to confirm deletion.

Working with Tapes

Tape media in Veeam Backup & Replication are displayed in the **Backup Infrastructure** view, under the **Tapes** > **Media** node. You can work with both online and offline tapes:

- Tapes that are currently loaded to the tape device are available under the **Online** node.
- Tapes that have been unloaded from the tape device are shown under the **Offline** node

Note Veeam Backup & Replication can use only online tapes for backup to tape and files to tape jobs. If you work with a standalone tape drive and all its tapes are offline, Veeam Backup & Replication will display a message informing that you need to insert a tape into the drive. At the same time, Veeam Backup & Replication will hint what tape has been recently used for archiving. You can insert any tape into the tape drive:

- If the tape you have inserted has been used last for archiving and the current media set can still be used, Veeam Backup & Replication will continue writing to this media set and append the new content to the content recently written on the tape.
- If the tape you have inserted has not been used last, Veeam Backup & Replication will mark this tape as free, create a new media set and start writing data to this new media set.

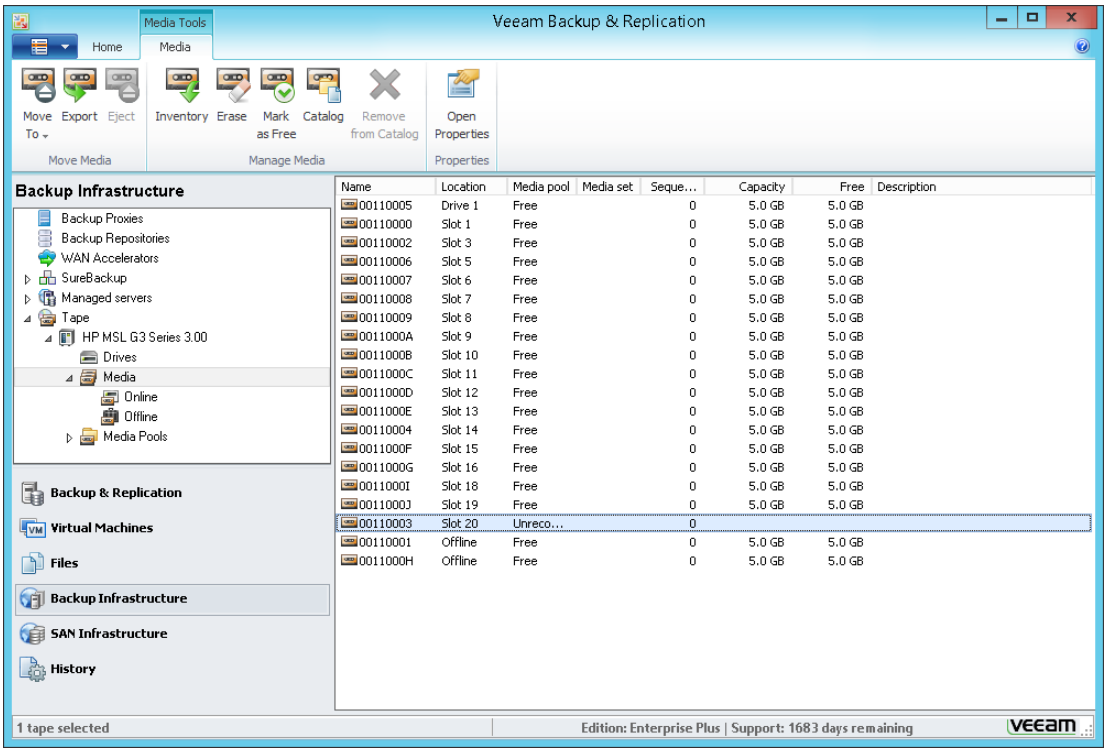

All tapes are grouped to predefined and custom media pools available under the **Tapes** > **Media** > **Media Pools** node.

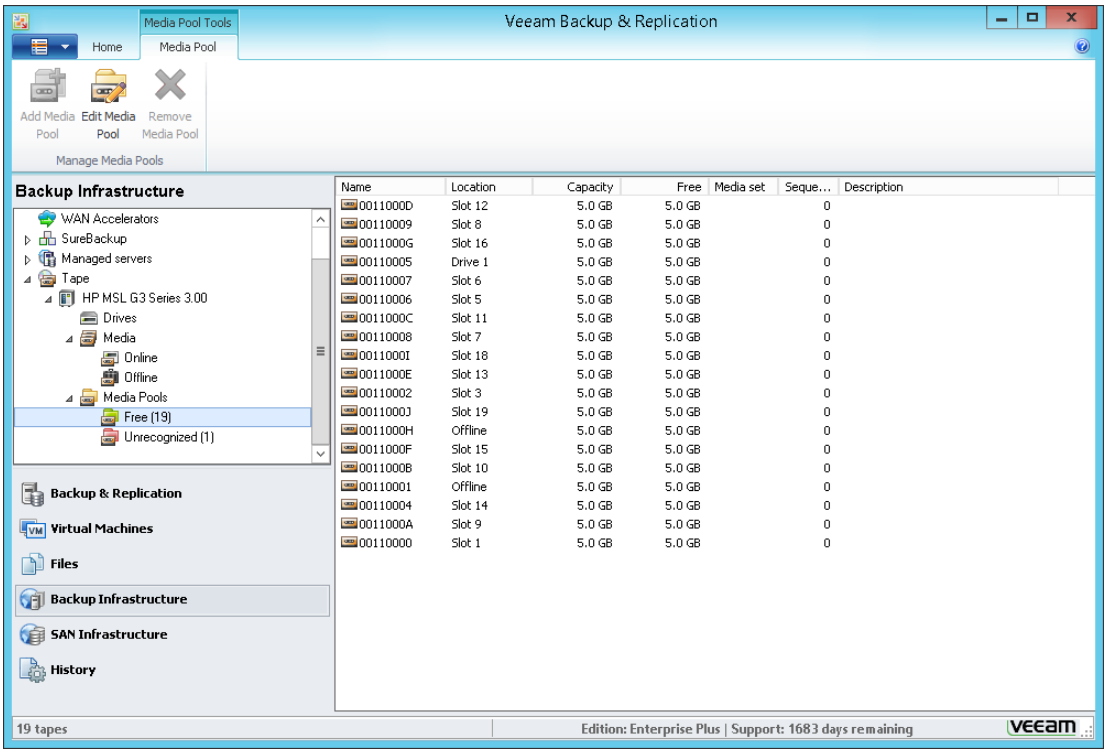

Veeam Backup & Replication supports the baseline set of tape management operations that can be performed via the Veeam Backup console.

Inventorying Tapes

When you load new tapes in your tape device for the first time, these tapes are presented to Veeam Backup & Replication as **Unrecognized**. To identify unrecognized tapes, you need to run tape inventory job against them.

Tape inventory job is a relatively fast process of reading metadata written on tape with the aim of detecting name of the media set and the sequence number for the tape. Tape inventory jobs helps Veeam Backup & Replication identify empty tapes and detect non-empty tapes belonging to specific media set.

You can perform tape inventory for a whole tape library or run the job against selected tapes only.

To start tape inventory, for a whole tape library:

1. Open the **Backup Infrastructure** view.
2. Select the tape library node under **Tapes** and click **Inventory Library** on the ribbon. Alternatively, you can right-click the tape library node and select **Inventory Library** from the shortcut menu. Veeam Backup & Replication will perform inventory for all online tapes in the library.

To start tape inventory, for selected tapes:

1. Open the **Backup Infrastructure** view.
2. Navigate to the list of tapes either under the **Media** > **Online** or under the **Media Pools** > *MediaPoolName* node.
3. Select the necessary tapes in the list and click **Inventory** on the ribbon. Alternatively, you can right-click the selected tapes and choose **Inventory Tape**.

The inventory log will display job session results. To access the inventory session details, you can open the **History** view and locate the necessary session under the **Jobs** > **Tape** node.

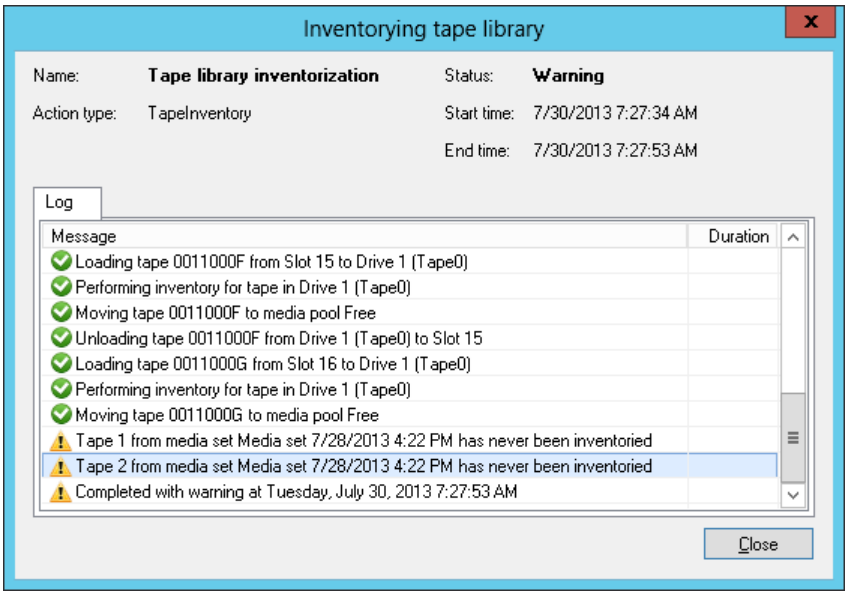

As a result of inventory, Veeam Backup & Replication places tapes in predefined pools:

- Empty tapes are moved to the **Free** pool. You can use these tapes for archiving backups and files
- Tapes that contain data written by 3rd party applications or tapes with data written on another Veeam backup server remain in the **Unrecognized** pool.
  Veeam Backup & Replication displays the detected media set name and sequence number for these tapes. If you want to restore data from a specific media set, you need to run the tape catalog job for all tapes in this media set.

After the tapes are inventoried, you can decide on further steps: for example, you can mark unrecognized tapes as free if you do not need archived content, or allocate tapes from the **Free** pool to a custom pool.

Cataloging Tapes

To streamline search for archived content and speed up data recovery process, Veeam Backup & Replication maintains tape and backup catalogs in the Veeam Backup & Replication database.

- **Tape Catalog** stores information about files and folders archived to tape media, as well as information about archived VBK and VIB backup files
- **Backup Catalog** stores information about VMs in backups that were archived to tape media

Tape and backup catalogs can be updated as follows:

- Files to tape and backup to tape jobs automatically update the catalog with details on created backup sets after each job session.
- You can manually run the tape catalog job to rescan tapes and update the Veeam Backup & Replication database. Running the job manually is required when there is no information about tape contents in the Veeam Backup & Replication database — for example, if a tape stores data in the MTF format written with a 3rd party backup solution or if the archive on tape was created on another Veeam backup server. To be able to restore from such tapes, you need to run the catalog job against all tapes in the necessary media set.

When a catalog job is performed, Veeam Backup & Replication first performs tape inventory, reads the information about backup contents on tape from the on tape catalog information, scans tape contents and updates the database with details of new detected backup sets.

You can perform tape catalog job for a whole tape library or for selected tapes only. To perform tape catalogization, for a whole tape library:

1. Open the **Backup Infrastructure** view.
2. Select the tape library node under **Tapes** and click **Catalog Library** on the ribbon. Alternatively, you can right-click the tape library node and select **Catalog Library**.

To perform tape catalogization for selected tapes:

1. Open the **Backup Infrastructure** view.
2. Navigate to the list of tapes either under the **Media** > **Online** or under the **Media Pools** > *MediaPoolName* node.
3. Select the necessary tapes in the list and click **Catalog** on the ribbon. Alternatively, you can right-click the selected tapes and choose **Catalog Tape**.

The catalog log will display job session results. To access the catalog session details, you can open the **History** view and locate the necessary session under the **Jobs** > **Tape** node.

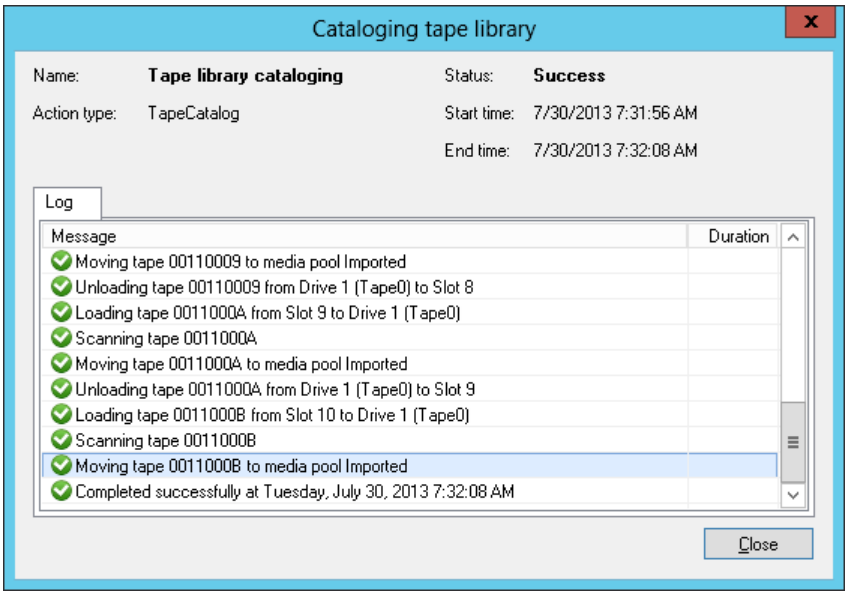

Rescanning of tapes during the catalog job may take a lot of time. To speed up the catalogization process, you might do one of the following:

- First, run inventory job to identify tape media sets and decide which tapes should be included in catalogization process. Run the catalog job only against tapes in the necessary media set.
- If you work with a tape library, you can run the catalog job against the whole media set at once.
- If you work with a standalone drive, start catalogization from the last tape in the media set (as this tape usually stores on tape catalog information).

Identifying Tapes

When performing periodic auto-rescan, Veeam Backup & Replication reads barcodes on tapes. If a tape is identified with a barcode, the barcode is displayed as the **Name** property of the tape. If a tape does not have a barcode, Veeam Backup & Replication will automatically create the tape name using to the pattern '*Tape N*'.

If necessary, you can change the name of a tape:

1. Open the **Backup Infrastructure** view.
2. Navigate to the list of tapes either under the **Media** > **Online/Offline** or under the **Media Pools** > *MediaPoolName* node.
3. Select the tape you want to rename and click **Open Properties** on the ribbon. Alternatively, you can right-click the tape and choose **Properties**.
4. Change tape name and description as required.
5. Click **OK** to save changes.

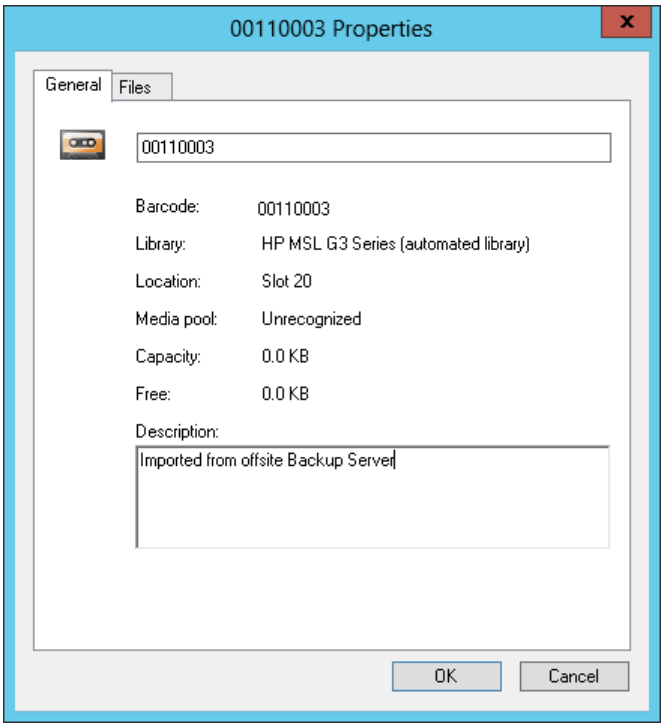

Removing Tapes from Catalog

If you do not want to store information about a tape and contents on this tape in Veeam Backup & Replication database, you can remove the tape from the catalog. Removing from catalog can be performed for offline tapes. For example, you can perform removal from catalog if a physical tape no longer exists, and the tape should no longer be displayed among media in the Veeam Backup & Replication console.

To remove one or more tapes from the catalog:

1. Open the **Backup Infrastructure** view.
2. Navigate to the list of tapes either under the **Media** > **Offline** or under the **Media Pools** > *MediaPoolName* node.
3. Select offline tapes you want to remove from the catalog and click **Remove from Catalog** on the ribbon. Alternatively, you can right-click selected tapes and choose **Remove from Catalog** from the shortcut menu.
4. In the opened dialog box, click **Yes** to confirm removal.

Moving Tapes to a Custom Media Pool

Veeam Backup & Replication allows you to move tapes between custom pools. To move tapes from one custom media pool to another:

1. Open the **Backup Infrastructure** view.
2. Navigate to the list of tapes either under the **Media** > **Online** or under the **Media Pools** > *MediaPoolName* node.
3. Select tapes you want to move and click **Move** to on the ribbon. Choose the target media pool from the list.
   Alternatively, you can right-click selected tapes and choose **Move to**. Next, choose the target media pool from the list.

Note Veeam Backup & Replication also allows you to move tapes from any pool to the **Free** media pool.

After you move a tape to another custom pool or move a tape to the **Free** pool, Veeam Backup & Replication marks this tape as free.

Erasing Tapes

If you do not need contents stored on tape, you can erase tapes. Veeam Backup & Replication supports two options for erasing data:

- **Short erase (fast)** — use this option to clear tape metadata only (refer to hardware properties to see if this option is supported)
- **Long erase (slow)** – use this option to clear all data written to tape

To erase tapes:

1. Open the **Backup Infrastructure** view.
2. Navigate to the list of tapes either under the **Media** > **Online** or under the **Media Pools** > *MediaPoolName* node.
3. Select tapes you want to erase and click **Erase** on the ribbon. Choose the type of erase and click **OK**.
   Alternatively, you can right-click selected tapes and choose **Erase**. Next, choose how the tape should be erased and click **OK**.

Marking Tapes as Free

Instead of erasing tapes, you can mark tapes as free. During this operation, Veeam Backup & Replication deletes from backup and tape catalogs information about backup contents stored on tape. Data written to tape remains intact.

To mark tapes as free:

1. Open the **Backup Infrastructure** view.
2. Navigate to the list of tapes either under the **Media** > **Online** or under the **Media Pools** > *MediaPoolName* node.
3. Select tapes you want to mark as free and click **Mark as Free** on the ribbon.
   Alternatively, you can right-click selected tapes and choose **Mark as Free** from the shortcut menu.
4. In the displayed dialog box, click **Yes**.

After a tape is marked as free, Veeam Backup & Replication removes from the catalog information about contents on this tape.

Ejecting Tapes

Veeam Backup & Replication allows you to eject a tape from the media drive and place the tape into a slot. For example, if you want to pull out from the library a tape that is currently in the tape drive, you first need to eject it and then export it to the I/E port.

To eject a tape from a drive:

1. Open the **Backup Infrastructure** view.
2. Navigate to the list of tapes either under the **Media** > **Online** or under the **Media Pools** > *MediaPoolName* node.
3. Select a tape you want to eject and click **Eject** on the ribbon. Alternatively, you can right-click selected tape and choose **Eject Tape**.

Tip	You can also eject tapes at the drives level. To do so, select the **Drives** node under **Tapes** in the inventory pane, right-click the necessary drive in the working area and select **Eject**.

Importing/Exporting Tape Media

Veeam Backup & Replication allows you to import and export tapes to import/export (I/E) ports (or Mail slots):

- The **Import** command places a tape from I/E slot into a standard slot of your library.
- The **Export** command will move the tape from the standard slot to the I/E slot.

To export a tape to an I/E slot:

1. Open the **Backup Infrastructure** view.
2. Navigate to the list of tapes either under the **Media** > **Online** or under the **Media Pools** > *MediaPoolName* node.
3. Select a tape you want to export and click **Export** on the ribbon.
   Alternatively, you can right-click selected tape and choose **Export Tape** from the shortcut menu.

Note Import/export commands are available only for the devices that support corresponding operations and include I/E slot.

To import a tape:

1. Open the **Backup Infrastructure** view.
2. Select the tape library node under **Tapes** and click **Import Library** on the ribbon.
   Alternatively, you can right-click the tape library node and select Import from the shortcut menu.

Viewing Backup Sets on Tape

To view backup contents that is currently stored on tape:

1. Open the **Backup Infrastructure** view.
2. Navigate to the list of tapes either under the **Media** > **Online/Offline** or under the **Media Pools** > *MediaPoolName* node.
3. Select the necessary tape and click **Open Properties** on the ribbon.
   Alternatively, you can right-click the tape and choose **Properties**.
4. Open the **Files** tab.

Alongside with backup archives and archived files, Veeam Backup & Replication stores on tape VTM files that contain metadata required for catalogization and restore.

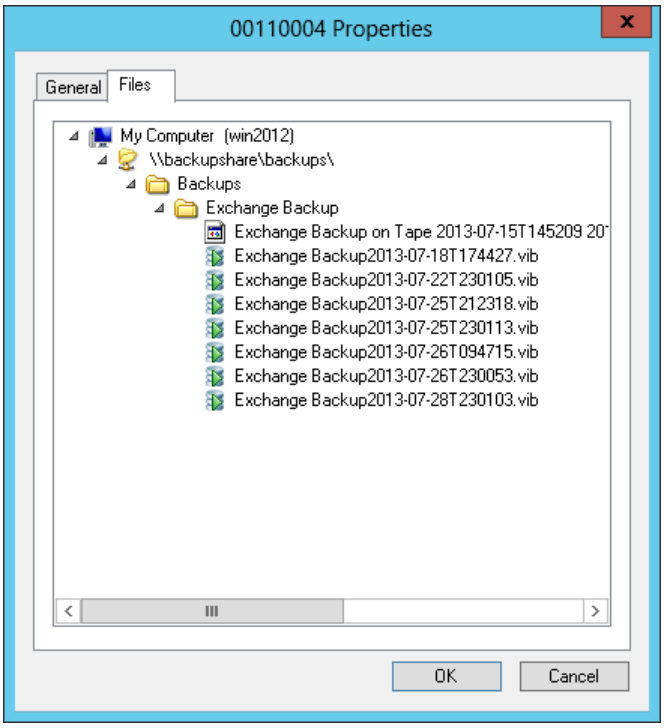

Archiving Data to Tape

Archiving data to tape in Veeam Backup & Replication is a job-driven process:

- Backup to tape jobs allow you to archive full backups or forward incremental backup chains stored on disk.
- Files to tape jobs allow you to archive files from any Windows or Linux connected to the Veeam backup server.

Creating Backup to Tape Jobs

To archive backups to tape, you should create a backup to tape job using the **New Backup To Tape Job** wizard. This section will guide you through all steps of the wizard and provide explanation on available options.

Before You Begin

Before you configure a backup to tape job, it is recommended that you:

- Create backup job(s) that produce the backup for archiving
- Create one or more custom media pool with the necessary media set and retention settings

Note	Note that backup to tape job processes only VBK (full backups) and VIB files (forward increments). Reversed increments (VRB) are skipped from processing.

To create a backup to tape job, follow the next steps:

Step 1. Launch the New Backup To Tape Job Wizard

To run the **New Backup to Tape Job** wizard, do either of the following:

- On the **Home** tab, click **Tape Job** and select **Backups**.
- Open the **Backup & Replication** view, right-click the **Jobs** node and select **Tape Job** > **Backups**.

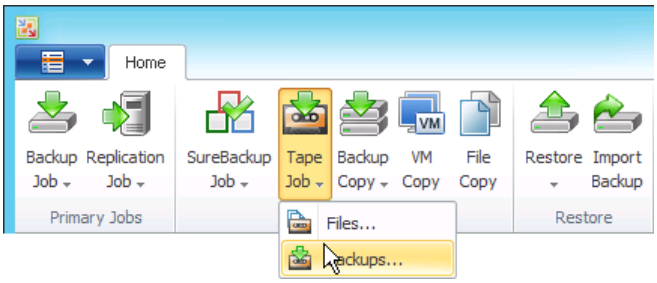

Step 2. Specify Job Name and Description

At the Name step of the wizard, you should define basic settings for the created backup to tape job.

1. In the **Name** field, enter a name for the created job.
2. In the **Description** field, enter a description of the created job. The default description contains information about the user who created the job, date and time when the job was created.

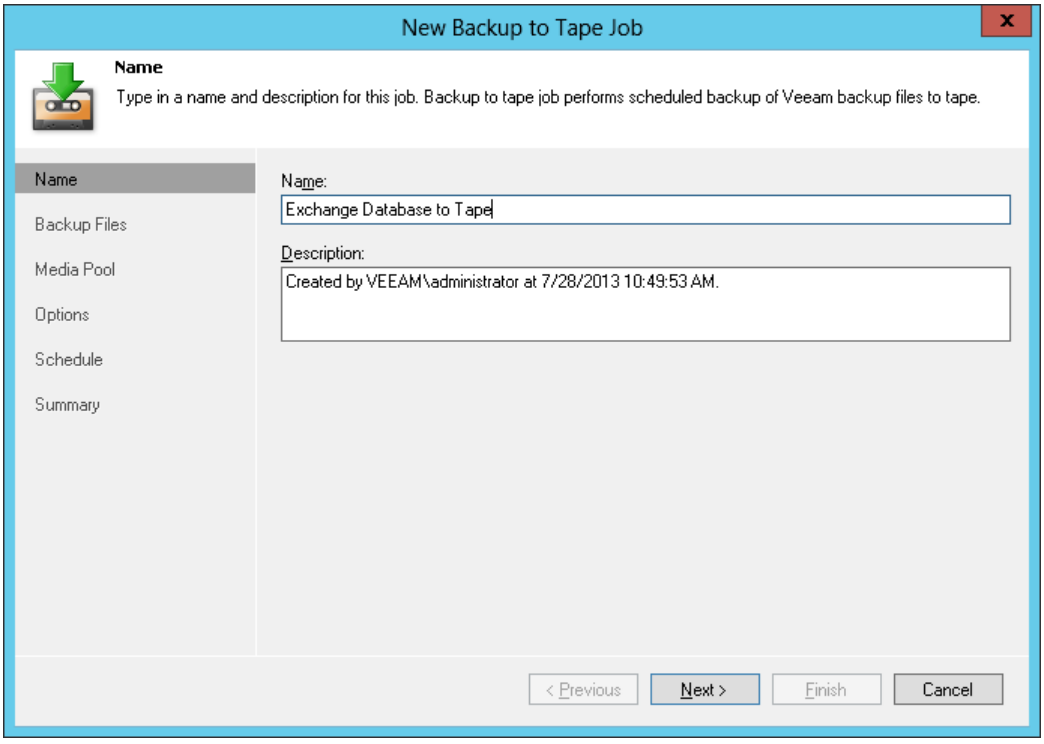

Step 3. Choose Backups to Archive

At the **Backup Files** step of the wizard, select backups that you want to archive to tape with the created job.

Click **Add** and select the necessary backups. You can choose backups from the following sources:

- **Backup jobs**. Using this option, you can select backups from available backup jobs. When a backup to tape job is run, Veeam Backup & Replication will search and archive restore points from the backups created by the selected backup jobs. If you choose to link to a backup job, Veeam Backup & Replication will update the secondary target settings for the backup job. In the same backup to tape job, you can combine backup jobs for different platforms (VMware vSphere, vCloud Director and Microsoft Hyper-V).
- **Backup repositories**. Using this option, you can select whole backup repositories. When a backup to tape job is run, Veeam Backup & Replication will search and archive restore points from all backups stored in the chosen backup repositories.

Note If you choose to archive data from backup repositories, the backup to tape job will process only the backups that were created with backup jobs configured on this Veeam backup server. Imported backups and configuration backups will be skipped from processing.

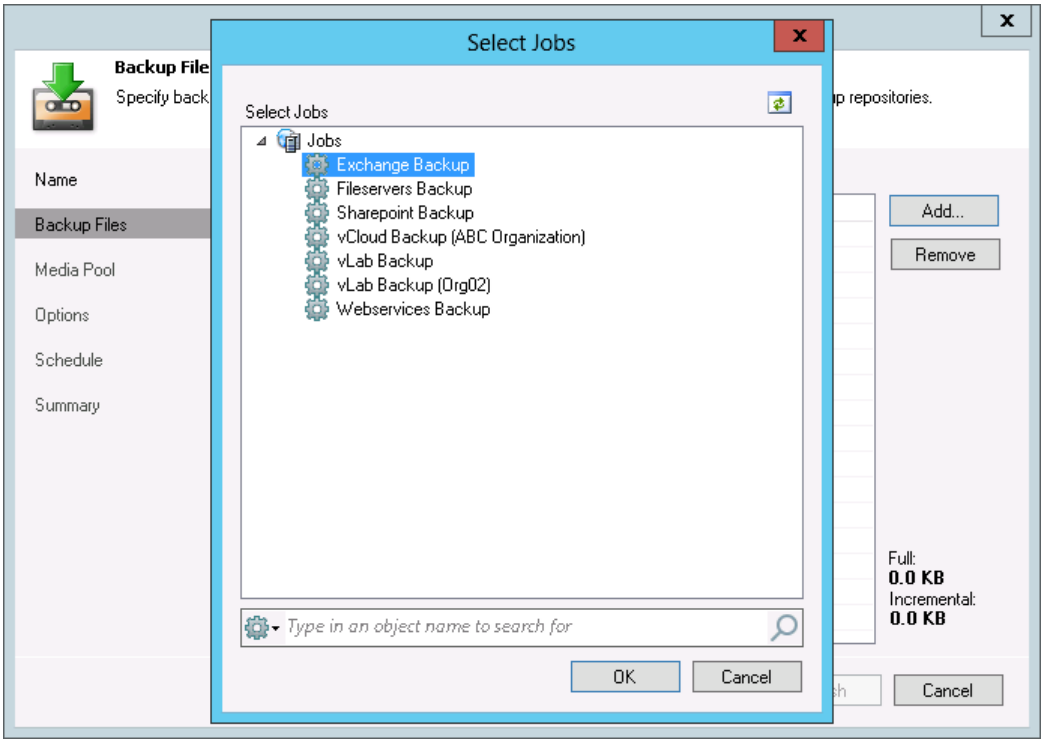

The total size of full backups added to the backup to tape job is displayed in the **Full** field. The total size of incremental backups added to the backup to tape job is displayed in the **Incremental** field.

To remove a backup from the list, select it and click **Remove** on the right.

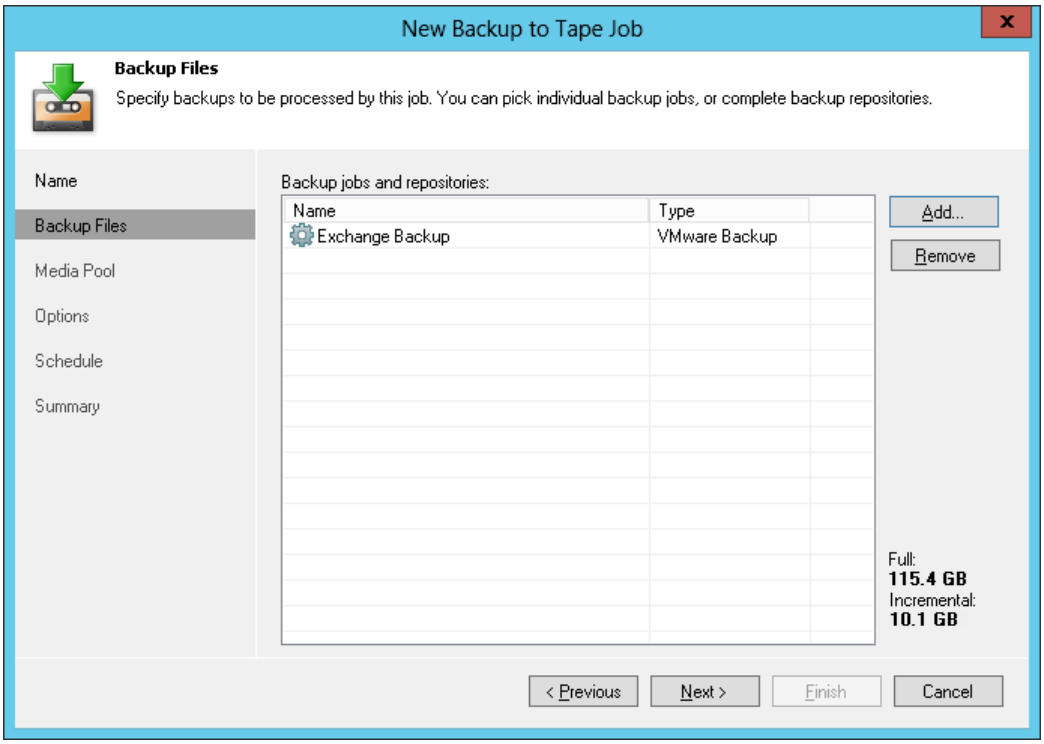

Step 4. Choose Media Pools

At the **Media Pool** step of the wizard, choose media pools that will be used for archiving backups.

1. From the **Media pool for full backups** list, choose a media pool that will be used for archiving full backup files.
2. From the **Media pool for incremental backups** list, choose a media pool that will be used for archiving incremental backup files.

Veeam Backup & Replication allows you to choose different media pools for full backups and incremental backups. This can be required if you use different media set or retention settings for archiving full backups and increments.

If you want the job to archive full backups only, clear the **Process incremental backup files** check box. If this option is disabled, the backup to tape job will archive only VBK files and will skip VIB files from processing.

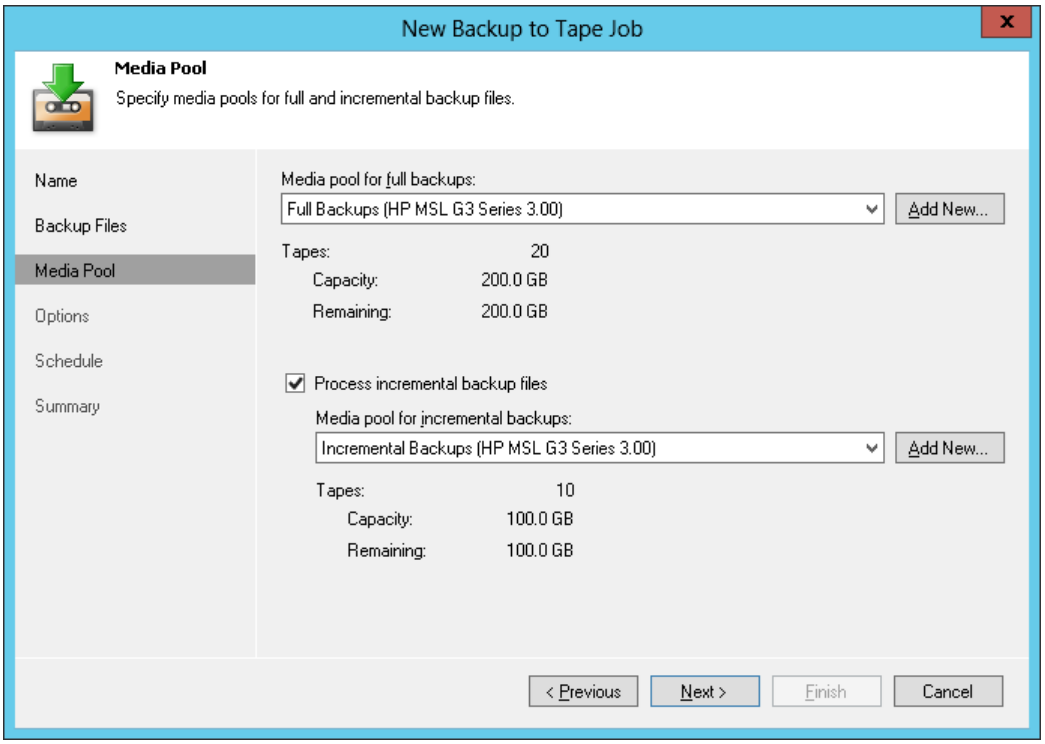

Tip If you have not previously created a media pool with the required settings, you can click the **Add New** button and create a new media pool without closing the job wizard. For more details, see Creating Custom Media Pools.

Step 5. Specify Archiving Options

At the **Options** step of the wizard, specify archiving and media automation options:

1. Disable the **Use hardware compression if available** check box.
2. Select the **Eject media once the job finishes** check box if upon job completion the tape should be automatically ejected from the tape drive and placed into a slot.
3. Select the **Export current media set once the job finishes** check box if upon job completion the tapes belonging to the media set should be placed to Import/Export (Mail) slot for further export (for example, for export to a remote office).

If you want to export tapes on specific days only (for example, every Saturday, when a backup to tape job completes), click **Days** and schedule export on the necessary days.

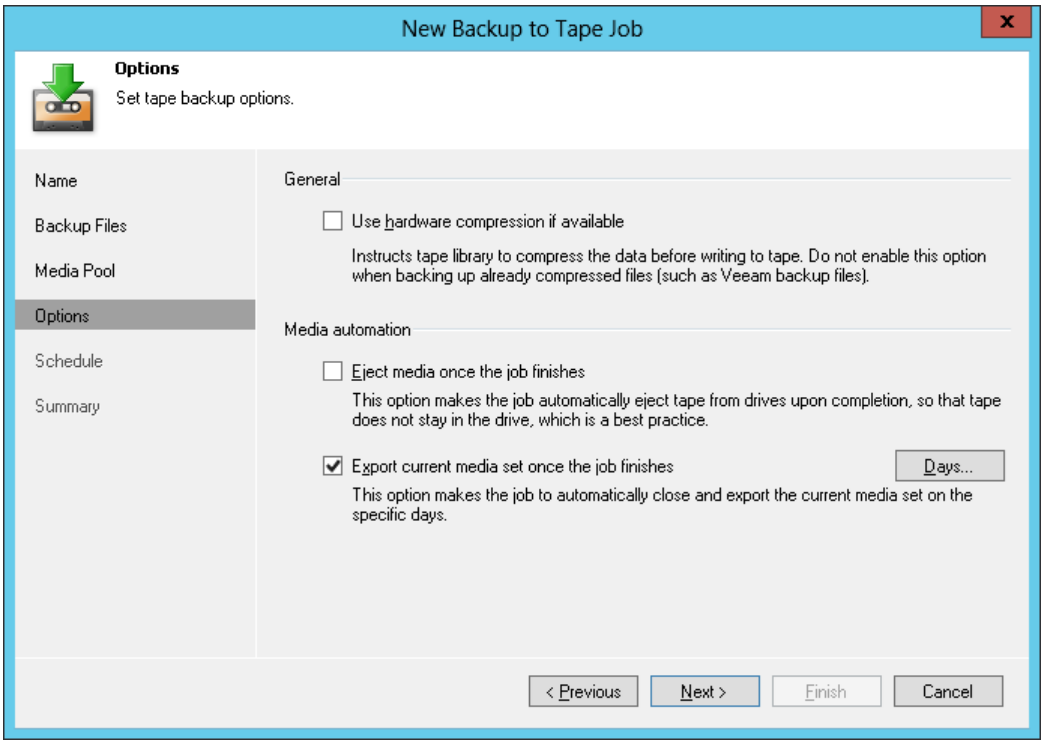

Step 6. Define the Job Schedule

At the **Schedule** step of the wizard, you can select to run the job manually or define a schedule for the job to run on a regular basis.

To specify the job schedule, select the **Run the job automatically** check box. If this check box is not selected, the job is supposed to be started manually.

You can define the following scheduling settings for the job:

- You can choose to run the job at specific time on defined week days, monthly and with specific periodicity.
- You can schedule the backup to tape job when a corresponding backup job completes. To do so, select the **After this job** option and choose the preceding backup job from the list.
- You can schedule the job to periodically check the source for new backups and archive new backups to target media. To do so, select the **As new backup files appear** option. If this option is selected, the backup to tape job will constantly remain in the *Idle* state, monitoring for new backups to appear. As soon as new backups are created, the job will start archiving these backup to tape.
  If necessary, you can define the time interval during which the backup to tape job must not archive data. These can be hours when backup repositories are busy with other tasks (backup jobs writing to repositories or backup copy jobs reading from repositories). To define prohibited time for the backup to tape job, click the **Schedule** button and define the time when the job is allowed and prohibited to run.

If you have scheduled the job to run at the specific time daily or monthly, consider configuring wait timeout value. Select the **If some linked backup jobs are still running, wait for up to … minutes** check box and specify the new timeout. When a backup to tape job starts,
Veeam Backup & Replication checks the status of the linked backup jobs. If a linked backup job is still writing data to the source repository, the backup to tape job will wait for the specified time interval.

If the timeout is disabled, the backup job will terminate without waiting for the backup job to complete.

The timeout option is unavailable if you schedule the backup to tape job to run after a backup job or if you schedule the backup to tape job to start when new backups appear.

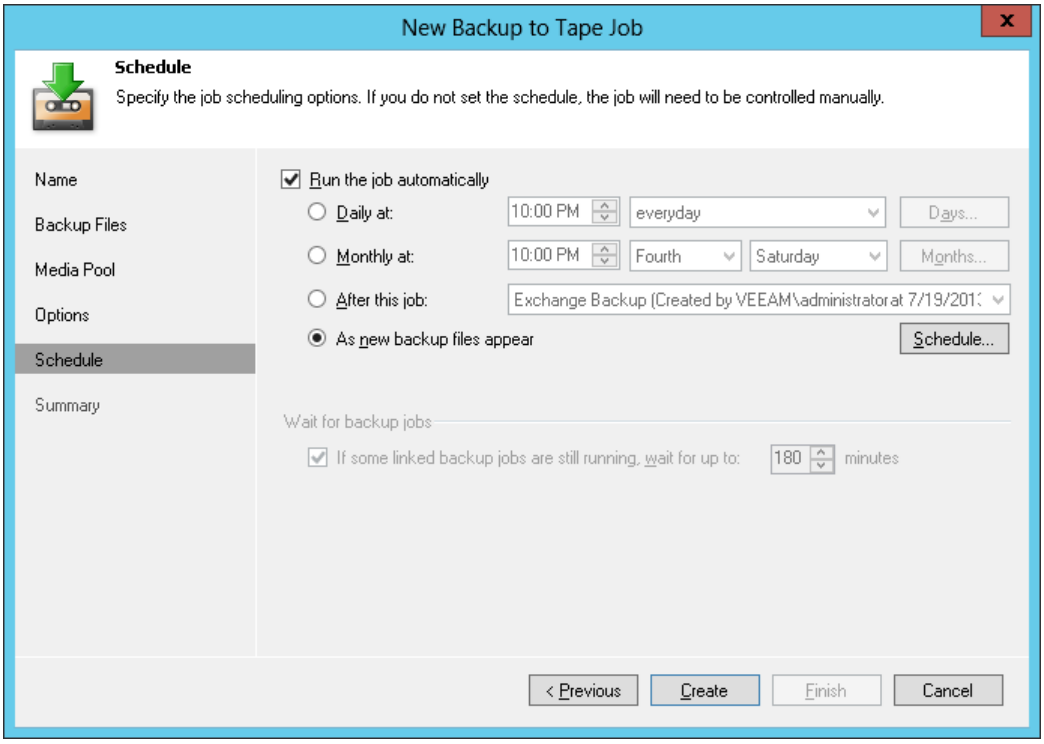

Step 7. Finish Working with the Wizard

After you have specified schedule settings, click **Create**. Select the **Run the job when I click Finish** check box if you want to start archiving backups to tape job right after you complete working with the wizard. Click **Finish** to close the wizard.

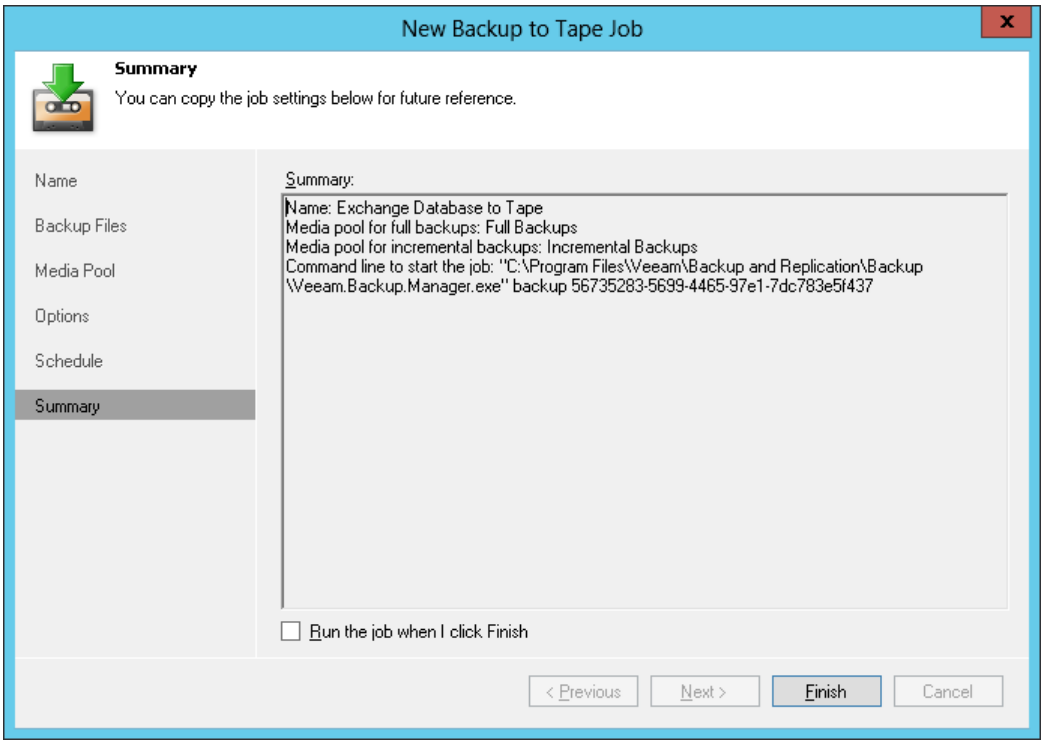

View Backups on Tape

After the backup to tape job completes, you can view the created archive on tape:

1. Open the **Backup & Replication** view.
2. Expand the **Backups** > **Tape** node and locate the backup archive under the name of a corresponding job. You can use these archives for full VM recovery.

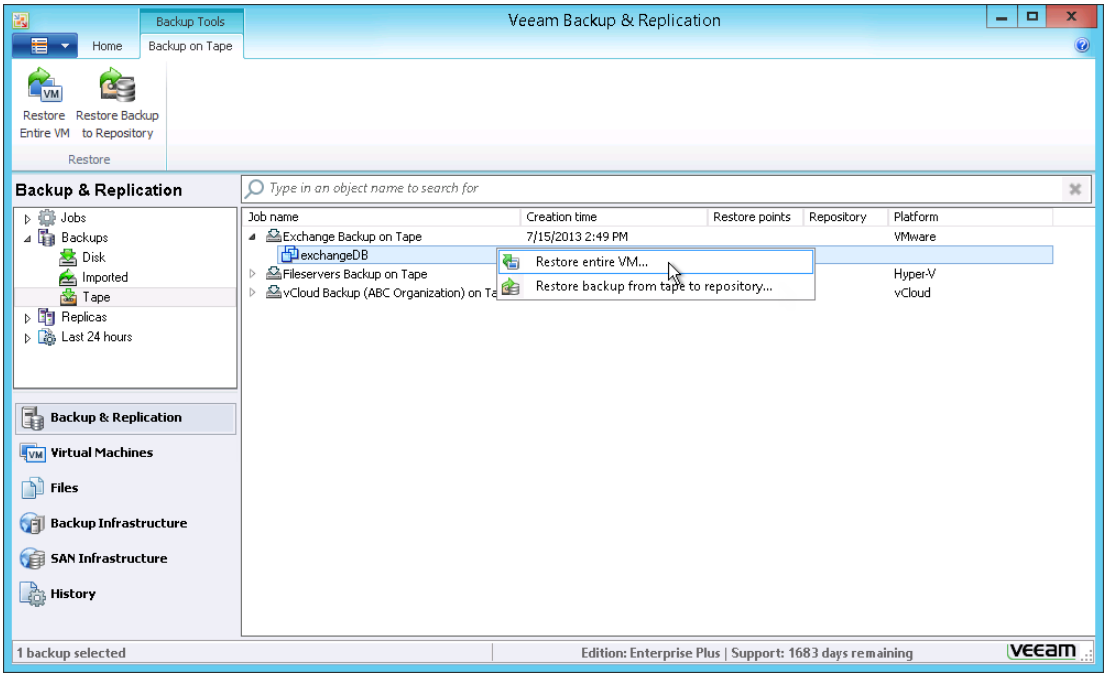

Linking Backup Jobs to Backup to Tape Jobs

Veeam Backup & Replication provides two options for linking backup jobs to backup to tape jobs:

- If you already have backup jobs configured, you can choose the necessary jobs in the **Backup to Tape Job Wizard**. For details, see Creating Backup to Tape Jobs
- Alternatively, you can point a backup job to an existing backup to tape job using the **Backup Job** wizard.

To point a backup job to an existing backup to tape job, perform the following steps:

1. Open the backup job settings and navigate to the **Storage** step. Select the **Configure secondary destination for this job** check box.

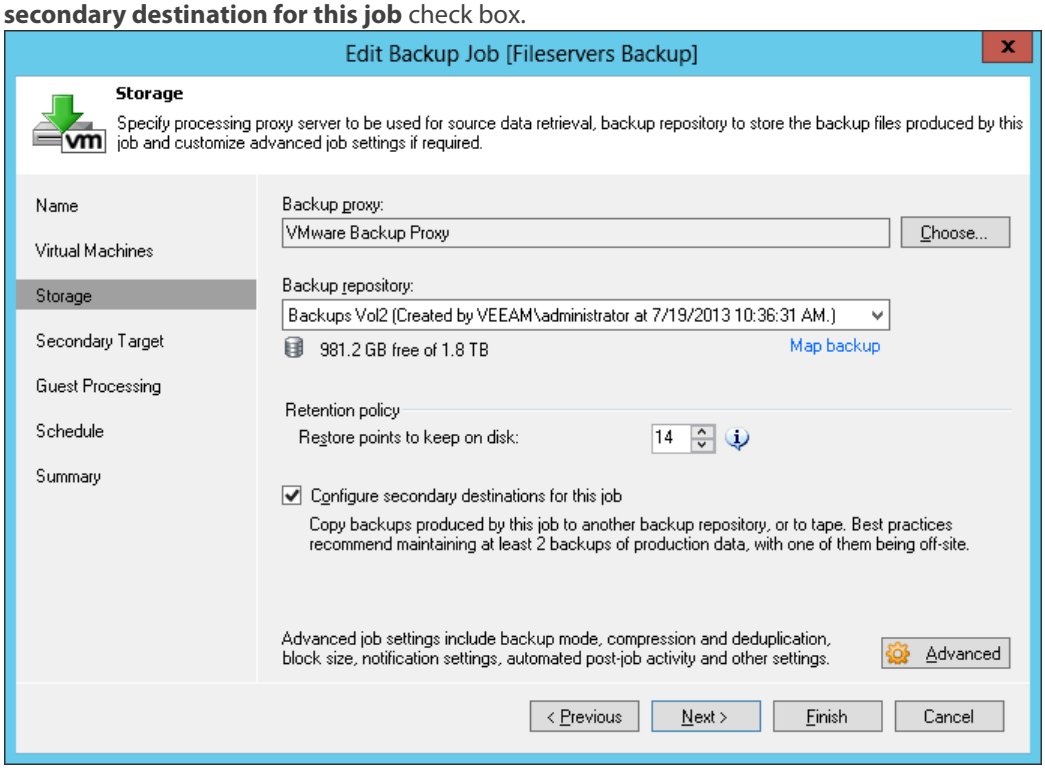

2. At the **Secondary Target** step, click **Add** and choose a backup to tape job to which the backup job should be linked.

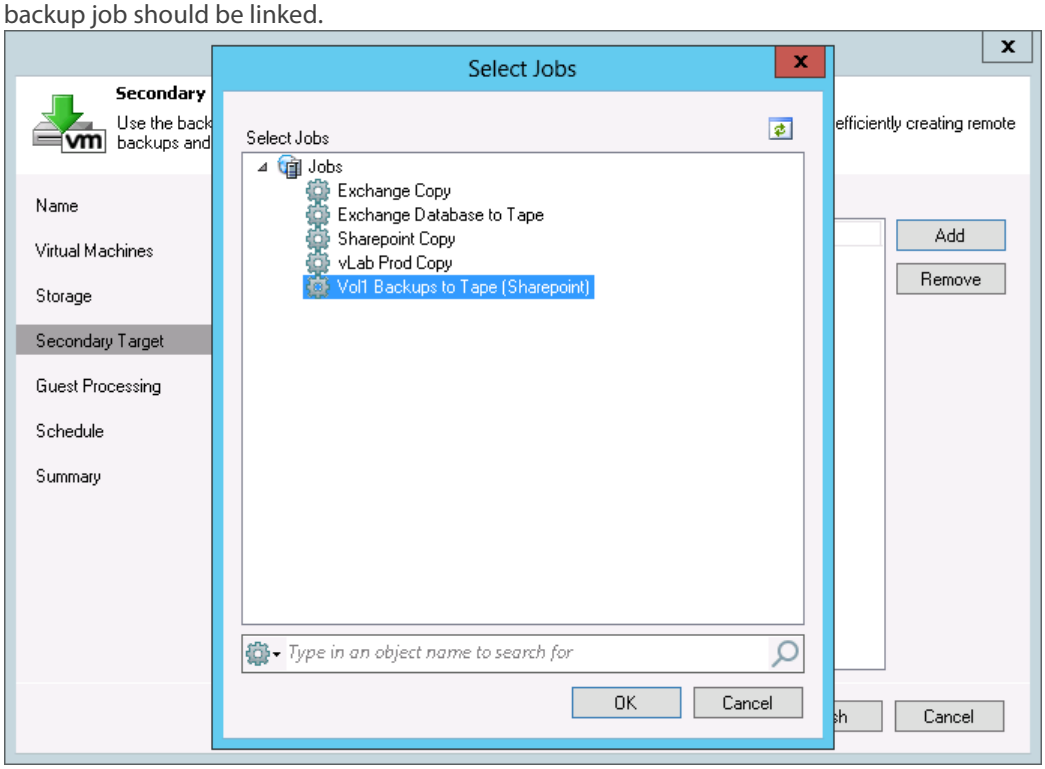

3. Save settings. Veeam Backup & Replication will automatically update backup files settings of the chosen backup to tape job.

Creating File to Tape Jobs

To archive files to tape, you should create a file to tape job using the **New File To Tape Job** wizard. This section will guide you through all steps of the wizard and provide explanation on available options.

To create a file to tape job, follow the next steps:

Step 1. Launch the New File To Tape Job Wizard

To run the **New File to Tape Job** wizard, do either of the following:

- On the **Home** tab, click **Tape Job** and select **Files**.
- Open the **Backup & Replication** view, right-click the **Jobs** node and select **Tape Job** > **Files**

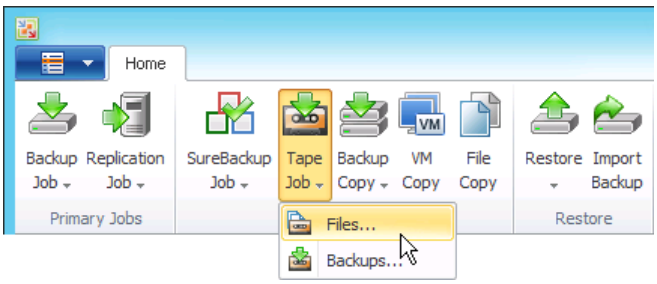

Alternatively, you can:

- Open the **Files** view, browse to the necessary files, select the files and choose **Add to Tape Job** > **New job** from the ribbon menu.
- Open the **Files** view, browse to the required files, right-click the necessary files and choose **Add to Tape Job** > **New job**.

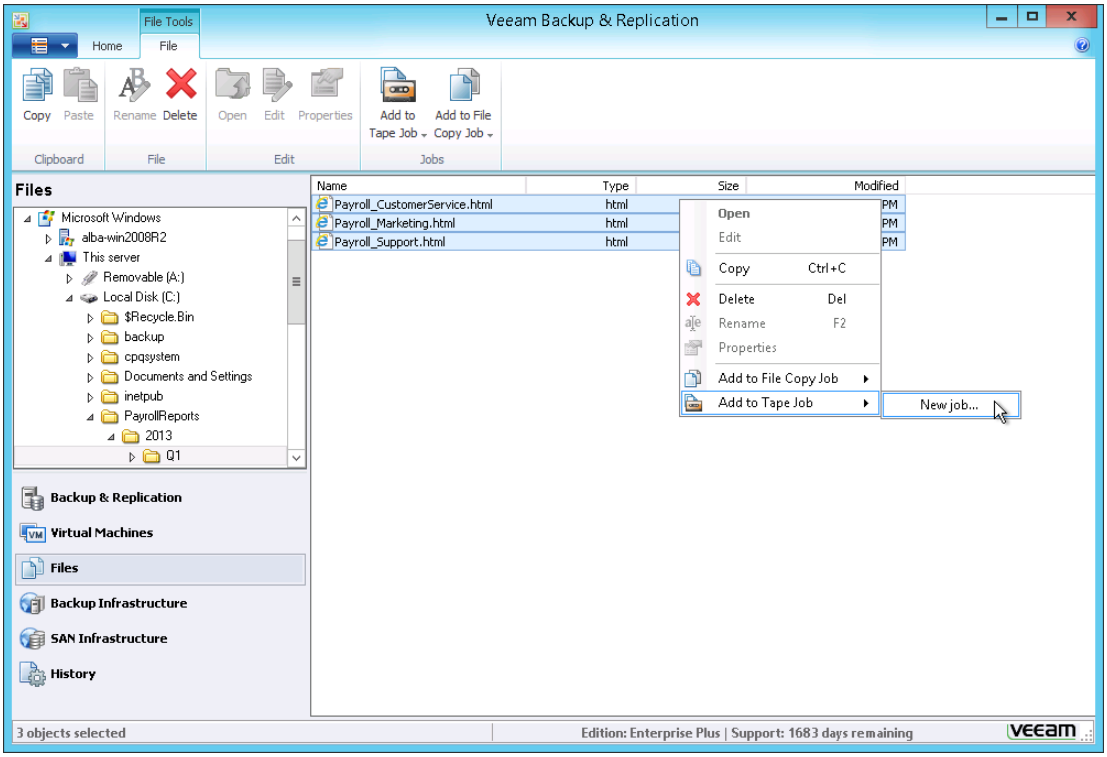

Step 2. Specify Job Name and Description

At the **Name** step of the wizard, you should define basic settings for the created files to tape job.

1. In the **Name** field, enter a name for the created job.
2. In the **Description** field, enter a description of the created job. The default description contains information about the user who created the job, date and time when the job was created.

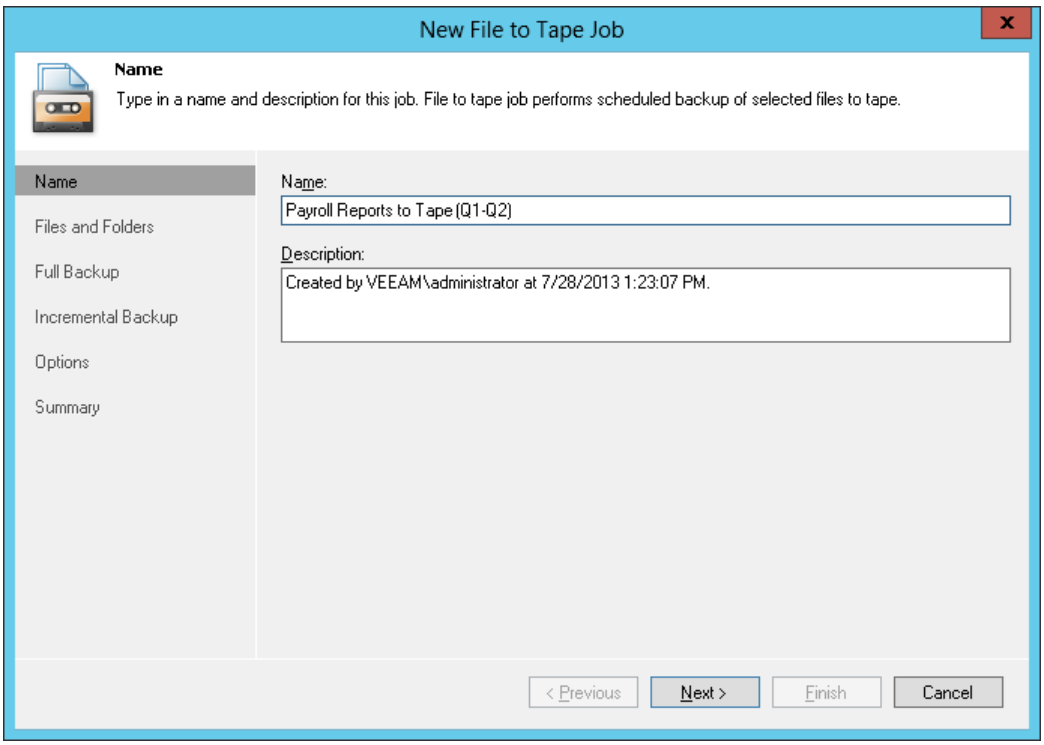

Step 3. Choose Files and Folders to Archive

At the **Files and Folders** step of the wizard, select files and folders that you want to archive.

From the **Server** list, choose a machine on which the necessary files or folders reside. Click **Add** and browse to the file or folder that should be archived. The selected item will be added to the list. To remove a file or folder from the list, select it and click **Remove**.

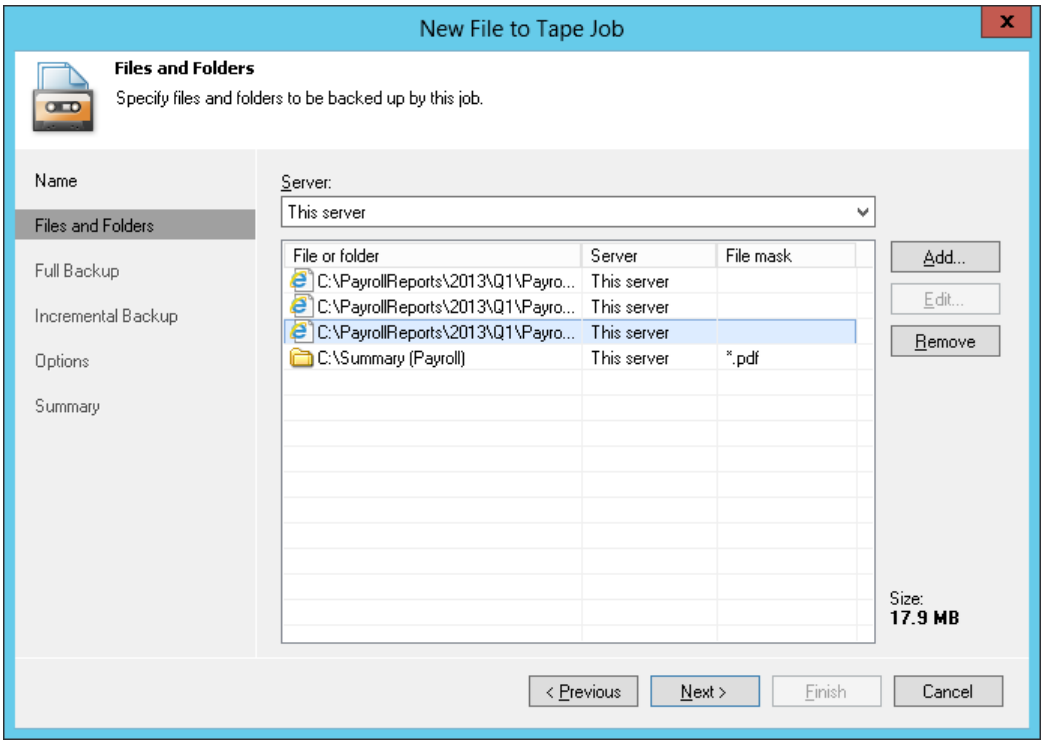

If you include a folder to the job, all of the folder contents will be archived. If necessary, you can choose only specific files from the included folder. To do so, select a folder in the list and click **Edit**. In the **File Masks** window, specify names of files in the folder that should be archived. You can use exact file names of create name masks (for example, *.evt or *.pdf). Separate file names and masks with semicolons.

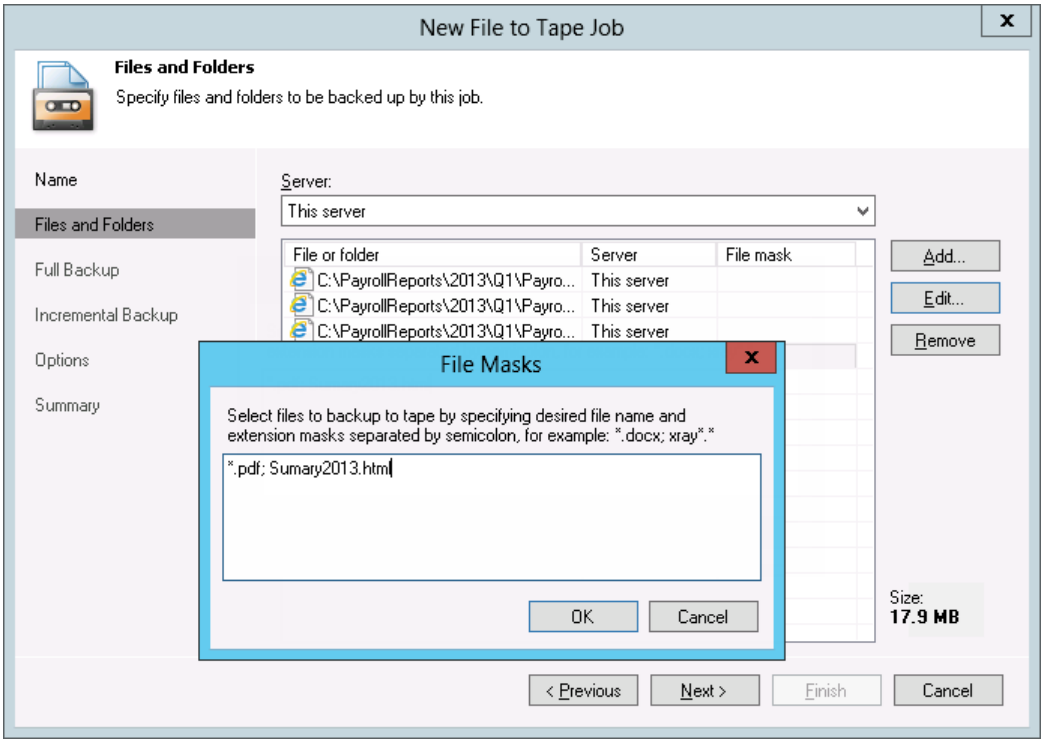

Step 4. Choose Media Pool for Full Backup

At the **Full Backup** step of the wizard, choose a media pool that will be used for archiving full backups of the selected files and create a schedule for full file backups.

1. From the **Media pool for full backups** list, choose a media pool that will be used for full file backups.
2. To schedule periodic creation of full file backups, select the **Run the full backup automatically** check box and specify the schedule according to which the job will create full file backups. If this option is disabled, you will need to start the job manually to create full backups of files.

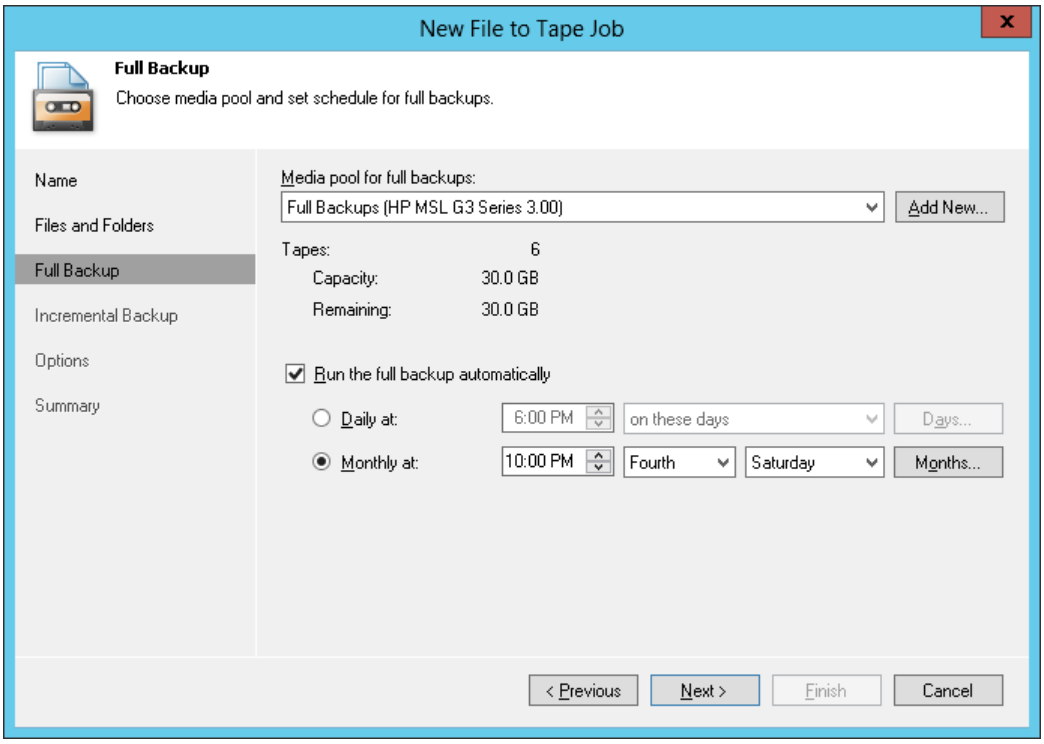

Tip If you have not previously created a media pool with the required settings, you can click the **Add New** button and create a new media pool without closing the job wizard. For more details, see Creating Custom Media Pools.

Step 5. Specify Media Pool for Increments

At the **Incremental Backup** step of the wizard, choose a media pool that will be used for archiving incremental backups of the selected files and create a schedule for incremental file backups.

1. From the **Media pool for incremental backups** list, choose a media pool that will be used for incremental file backups.
2. To schedule periodic creation of incremental file backups, select the **Run incremental backup automatically** check box and specify the schedule according to which the job will create incremental file backups. If this option is disabled, you will need to start the job manually to create incremental backups of files.

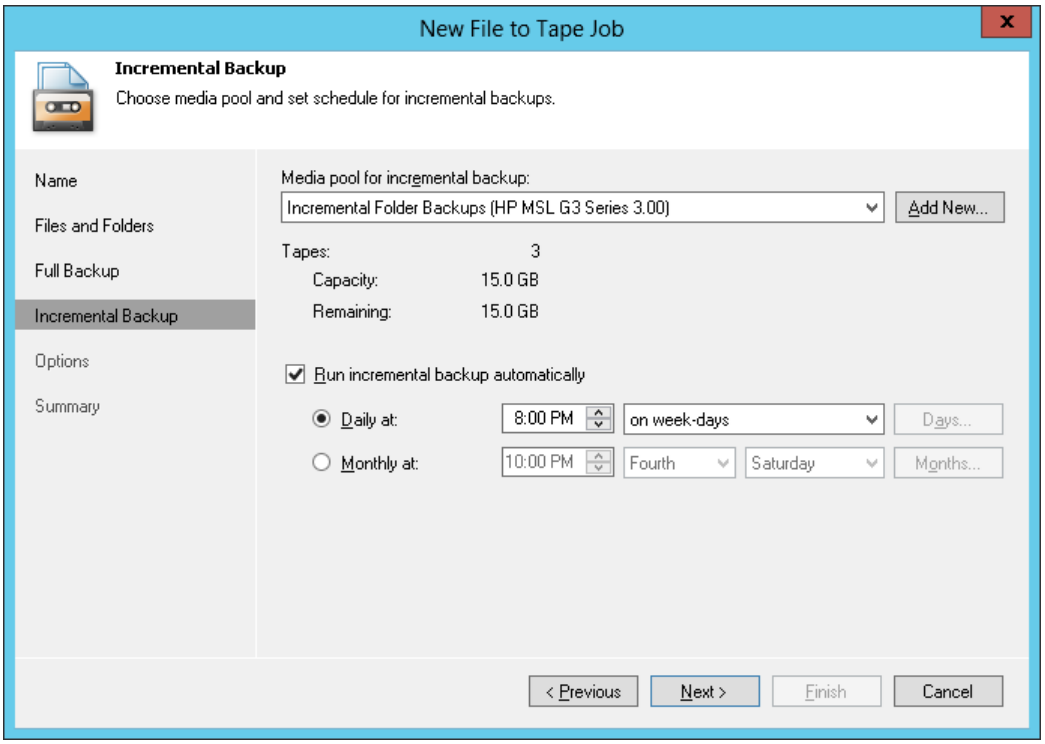

Tip If you have not previously created a media pool with the required settings, you can click the **Add New** button and create a new media pool without closing the job wizard. For more details, see Creating Custom Media Pools.

Step 6. Specify Archiving Options

At the **Options** step of the wizard, specify archiving and media automation options:

1. Select the **Use Microsoft volume shadow copy** check box to enable backup of files with the help if Microsoft shadow volume copies. This option enables backup of files locked by application and provides file-level quiescence. This possibility can only be applied for files from servers running under Windows XP or later Windows-family OSs.
2. Select the **Use hardware compression if available** check box if the tape drive should compress file data before archiving it to tape.

Important! Enable the hardware compression only if your tape library provides support for hardware compression.

3. Select the **Eject media once the job finishes** check box if upon job completion, the tape should be automatically ejected from the tape drive and placed into a slot.
4. Select the **Export current media set once the job finishes** check box if upon job completion the tapes belonging to the media set should be placed to Import/Export (Mail) slot for further export (for example, for export to a remote office). If you want to export tapes on specific days only (for example, every Saturday, when a backup to tape job completes), click **Days** and schedule export on the necessary days).

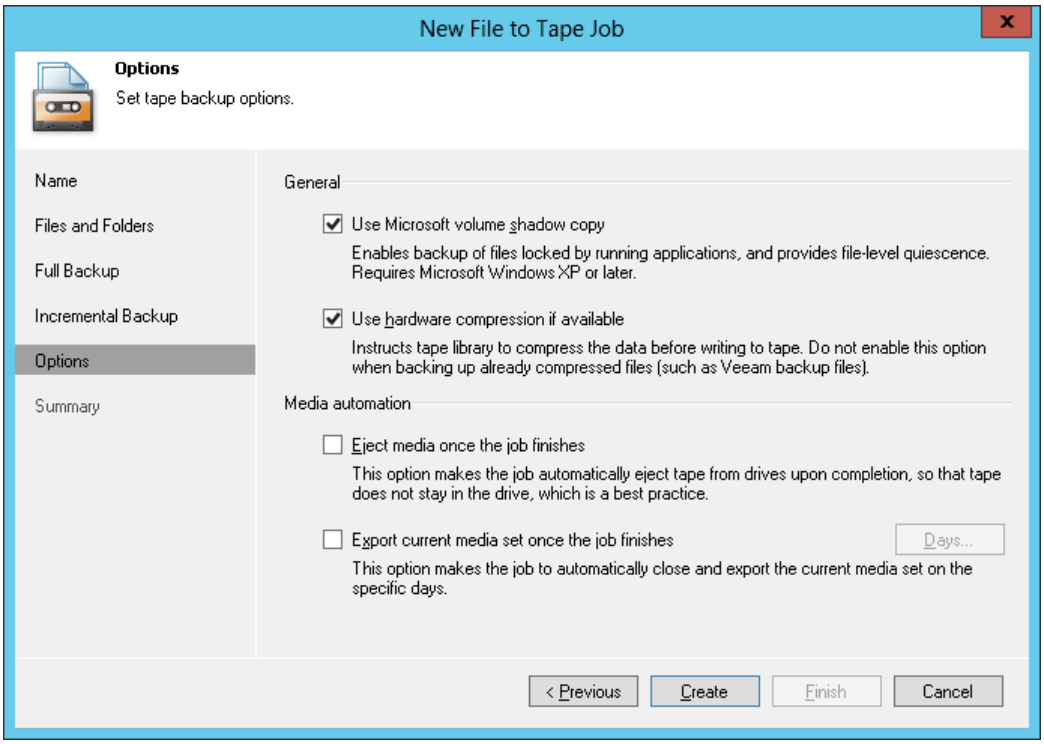

Step 7. Finish Working with the Wizard

After you have specified schedule settings, click **Create**. Select the **Run the job when I click Finish** check box if you want to start archiving files to tape job right after you complete working with the wizard. Click **Finish** to close the wizard.

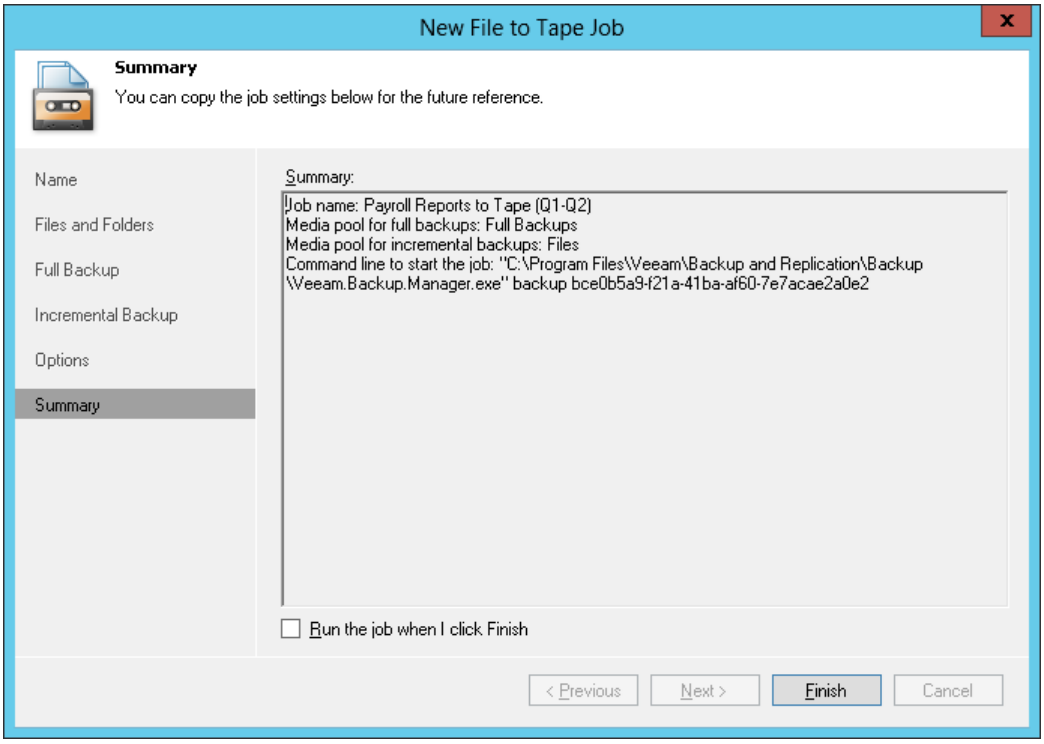

View Files on Tape

After the files to tape job completes, you can view the created archive on tape:

1. Open the **Files** view and press F5 to refresh it.
2. Expand the **Tape** node and locate the files archive. Veeam Backup & Replication preserves the source hierarchy of folders for archived files. You can use the created archive for file recovery.

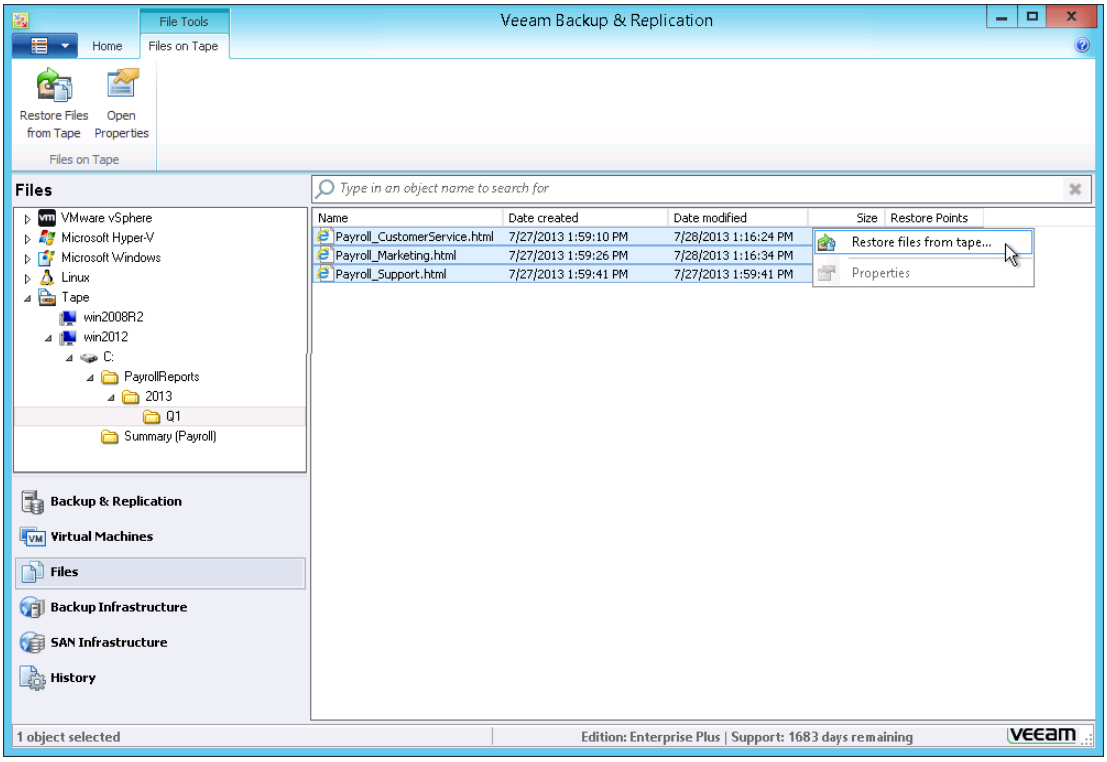

Restoring Data from Tape

Veeam Backup & Replication offers the following options for restoring data from tape:

- Restoring full backups or backup chains to disk
- Performing full VM recovery from backup on tape into the virtual infrastructure
- Restoring files and folders to their original location or any folder of your choice
- Restoring files backed up to tape with 3rd party solutions

With Veeam Backup & Replication, you can quickly search the catalog for VMs, files and folders stored on tape. If the necessary tape is offline at the moment of restore, Veeam Backup & Replication will prompt you to load the necessary tape to the device.

Restoring Backups from Tape

Veeam Backup & Replication allows you to restore backups from tape to a repository or a folder on disk. To restore backups from tape, use the **Restore Backup from Tape to Repository** wizard. This section will guide you through all steps of the wizard and provide explanation on available options.

To restore files from tape, follow the next steps:

Step 1. Launch the Files Restore Backup from Tape to Repository Wizard

To run the **Restore Backup from Tape to Repository** wizard, on the **Home** tab, click **Restore** and choose **Tape** > **Restore Backups**.

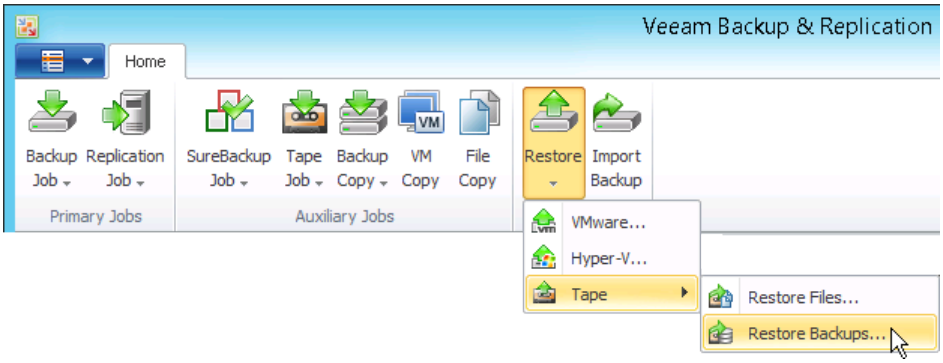

Alternatively, you can:

- Open the **Backup & Replication** view, expand the **Backups** > **Tapes** node. Select the necessary VMs in backup and click **Restore backup to Repository** on the ribbon.
- Open the **Backup & Replication** view, expand the **Backups** > **Tapes** node. Right-click the necessary VMs in backup and choose **Restore backup from tape to repository**.

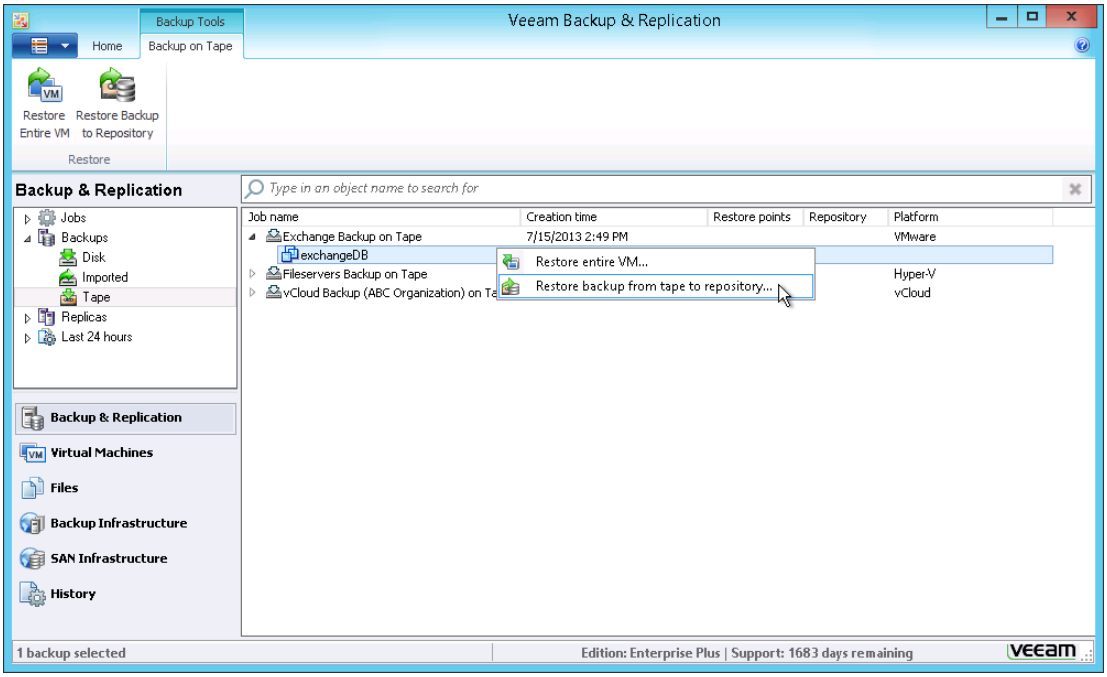

Step 2. Choose VMs to Restore

At the **Source** step of the wizard, select one or more VMs for which backup files should be restored. If you have chosen VMs to restore from archives on tape, the list of objects to restore will be populated with selected VMs.

To add one or more VMs to the list, click **Add VM** and select where to browse for the machines:

- **From Infrastructure** — browse the virtual environment and select VMs to restore. If you choose a VM container, Veeam Backup & Replication will expand it to a plain VM list. To quickly find a VM, use the search field at the top of the list: enter the VM name or a part of it and click the search button on the right or press **[ENTER]**.
  Make sure that VMs you select from the virtual environment have been successfully archived to tape at least once.
- **From Backup** — browse existing backups on tape and select VMs under backup to tape jobs. To quickly find VMs, use the search field at the bottom of the **Select Objects** window: enter a VM name or a part of it and click the **Start search** button on the right or press **[ENTER]**.

To remove a VM, select it in the list and click **Remove** on the right.

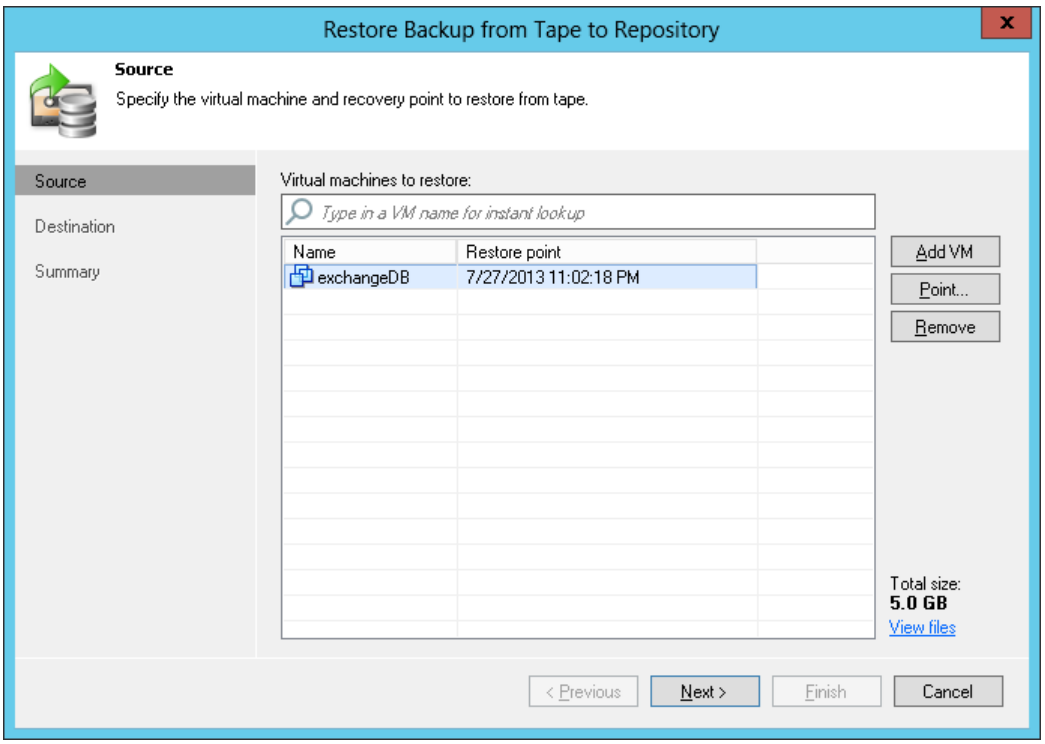

Step 3. Select a Restore Point

By default, Veeam Backup & Replication will restore backup with the latest state of the archived VM. However, if you want to restore a backup for the VM to an earlier state, select a VM in the list and click **Point** on the right. In the **Restore Points** section, select a restore point that should be used to restore VM backup.

Note If you choose a full backup point in the list, Veeam Backup & Replication will restore only this full backup. If you choose an increment, Veeam Backup & Replication will restore a chain consisting of a full backup and forward increments, necessary to restore VMs to the required point-in-time.

If you have chosen to restore multiple VMs, you can select a different restore point for every VM specifically.

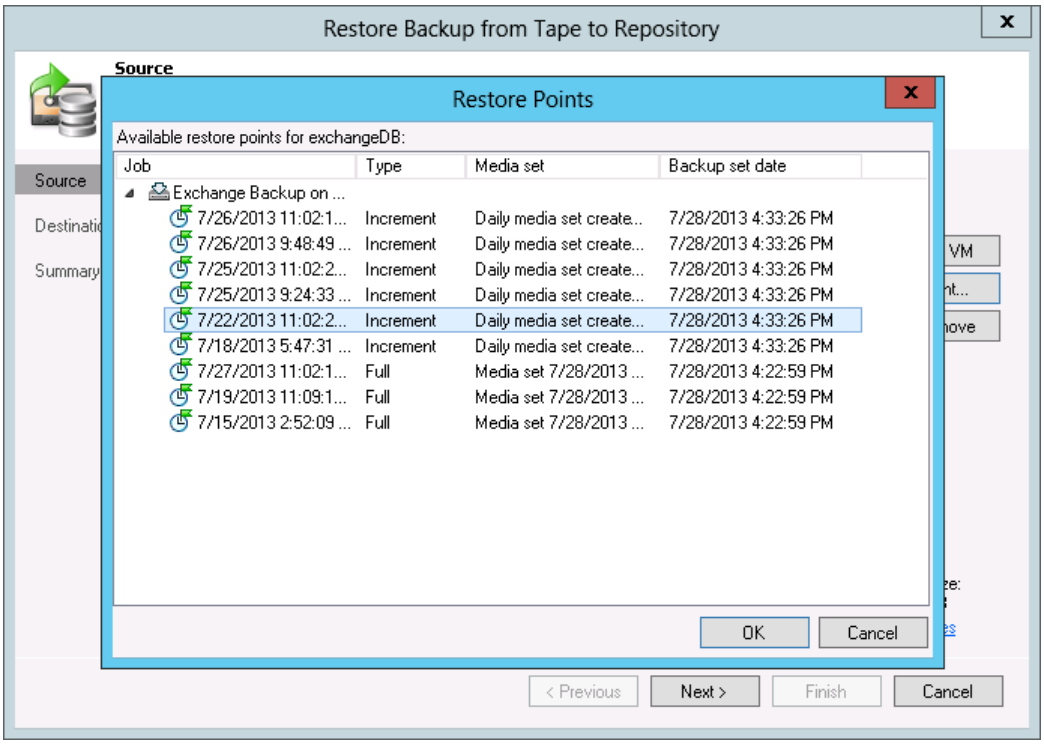

Step 4. Choose Backup Destination

At the **Destination** step of the wizard, select where the backup files for the selected VMs should be restored:

- To restore VM backup files to a repository, select the **Backup Repository** option and choose the necessary repository from the list.
- To restore VM backup files to a server connected to Veeam Backup & Replication, select the **Server** option, choose the necessary server from the list and specify path to the target folder in the **Path to folder** field.
- To restore VM backup files to network share, select the **Server** option, choose **This computer or shared folder** from the list and specify path to the shared folder in the **Path to folder** field. If you choose to restore backup files to a shared folder, make sure that the account under which Veeam Backup Service runs has write permissions to the target folder. If the account does not have sufficient permissions, Veeam Backup & Replication will prompt you to enter credentials for the account that can be used for writing to the target folder.

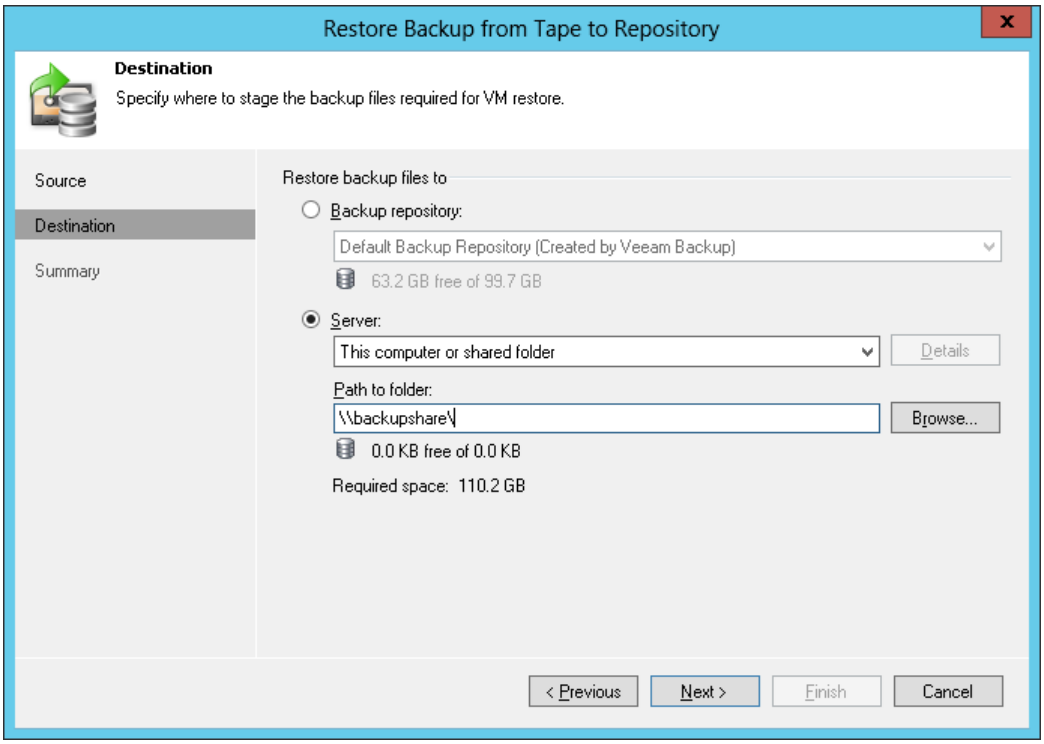

Step 5. Finish Working with the Wizard

Review the settings and click **Finish** to restore selected VM backups.

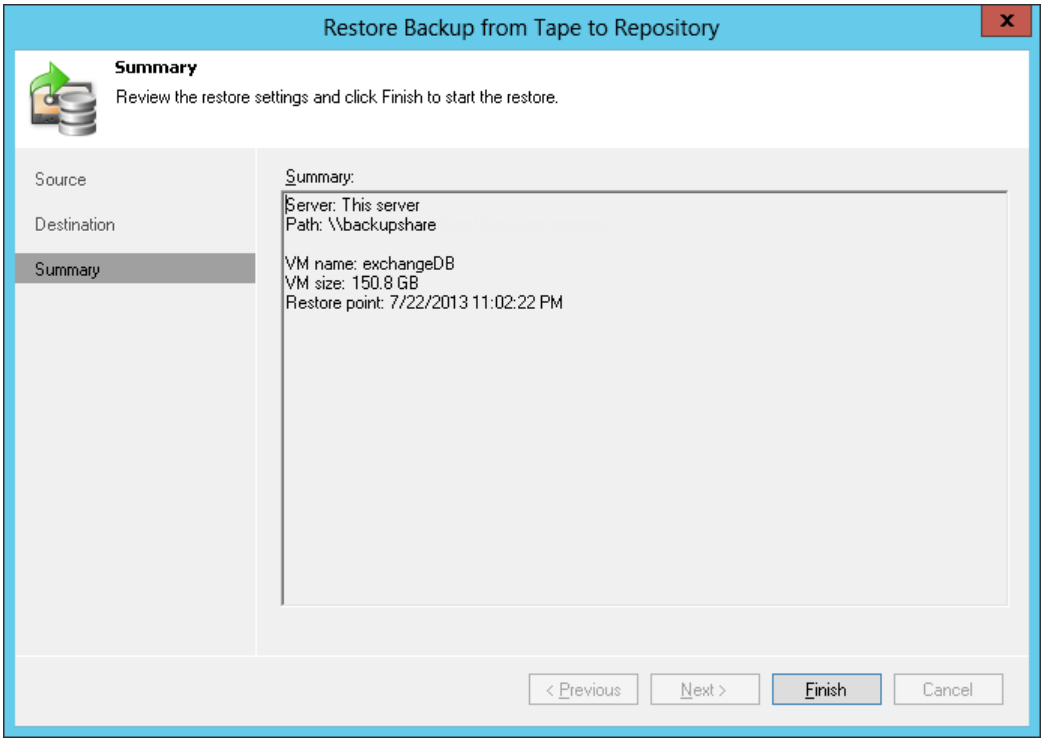

After backups are restored from tape, they are displayed as imported backups in the **Backup & Replication** view > **Backups** > **Imported**. You can use the restored to disk backup for regular data recovery, including full VM recovery, recovery of VM files, guest OS files recovery and other.

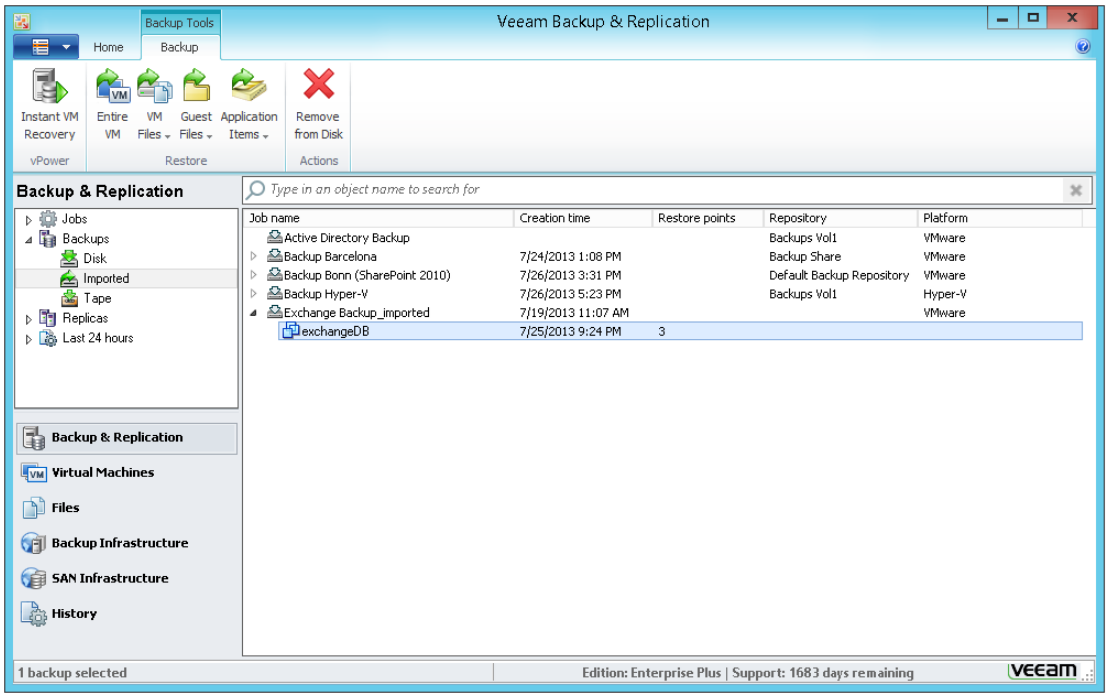

Restoring VMs from Tape to Virtual Infrastructure

Veeam Backup & Replication allows you to recover full VMs from archives on tape. The process of full VM restore from tape includes two stages:

1. First, Veeam Backup & Replication restores a backup file to a staging location (a backup repository or a folder).
2. Next, Veeam Backup & Replication restores the VM into the virtual infrastructure. After the VM is restored, the backup is deleted from the staging location.

To restore full VMs from tape to virtual infrastructure, use the **Full VM Restore** wizard. This section will guide you through all steps of the wizard and provide explanation on available options.

To restore files from tape, follow the next steps:

Step 1. Launch the Full VM Restore wizard

To run the **Full VM Restore** wizard, do one of the following:

- Open the **Backup & Replication** view, expand the **Backups** > **Tapes** node. Select the necessary VMs in backup and click **Restore entire VM on the Backup on Tape** on the ribbon.
- Open the **Backup & Replication** view, expand the **Backups** > **Tapes** node. Right-click the necessary VMs in backup and choose **Restore entire VM**.

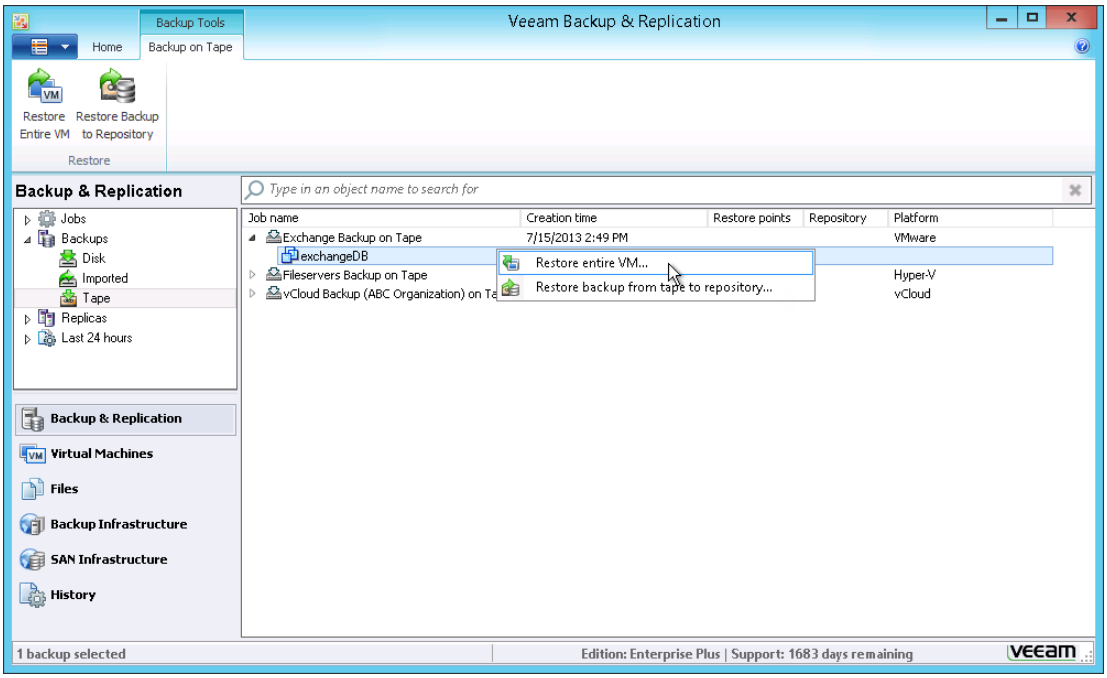

Step 2. Choose Virtual Machines to Restore

At the **Virtual Machines** step of the wizard, review VMs that should be restored. To add one or more VMs to the list, click **Add VM** and select where to browse for the machines:

- **From Infrastructure** — browse the virtual environment and select VMs to restore. If you choose a VM container, Veeam Backup & Replication will expand it to a plain VM list. To quickly find a VM, use the search field at the top of the list: enter the VM name or a part of it and click the search button on the right or press **[ENTER]**.
  Make sure that VMs you select from the virtual environment have been successfully archived to tape at least once.
- **From Backup** — browse existing backups and select VMs under backup to tape jobs. To quickly find VMs, use the search field at the bottom of the **Select Objects** window: enter a VM name or a part of it and click the **Start search** button on the right or press **[ENTER]**.

To remove a VM, select it in the list and click **Remove** on the right.

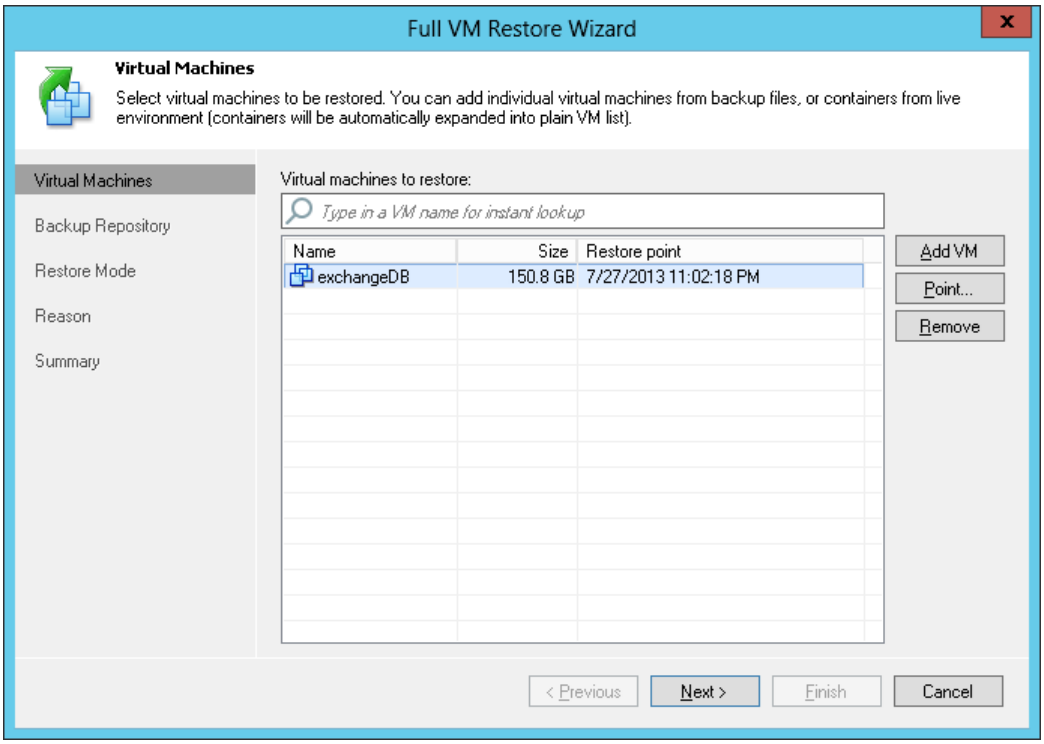

Step 3. Select a Restore Point

By default, Veeam Backup & Replication will restore VMs to their latest state archived to tape. However, if you want to restore a VM to an earlier state, select a VM in the list and click **Point** on the right. In the **Restore Points** section, select a restore point that should be used for full VM recovery.

If you have chosen to restore multiple VMs, you can select a different restore point for every VM specifically.

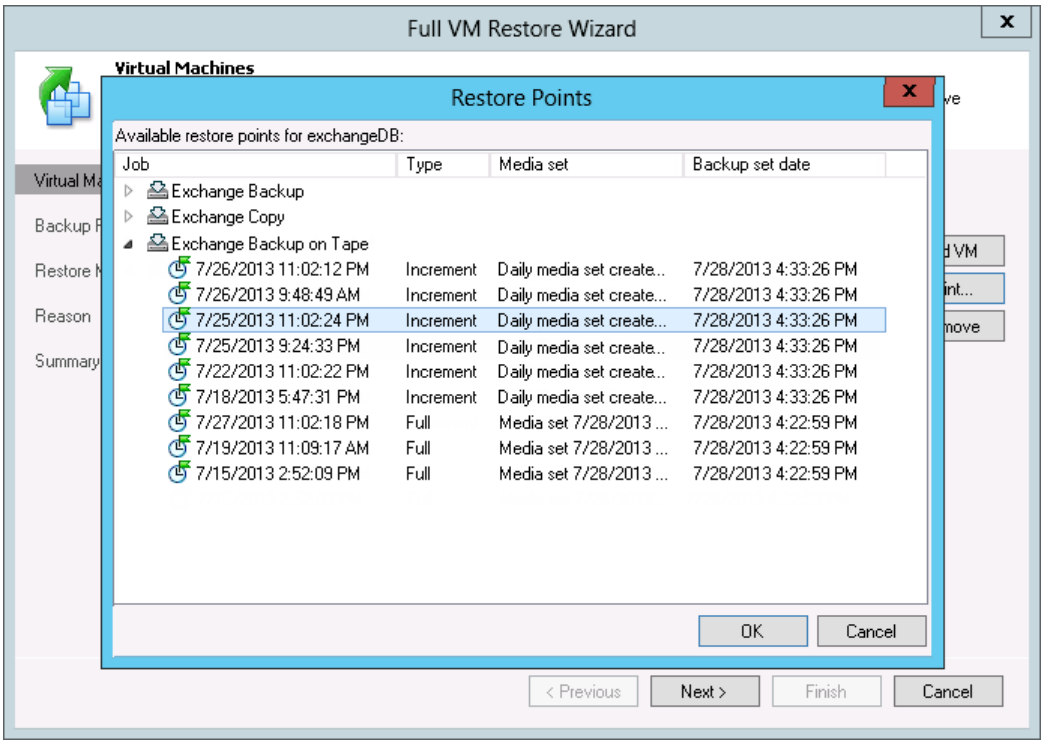

Step 4. Choose Backup Repository

Because tape does not support random access, Veeam Backup & Replication restores VM backups from tape to a staging area (backup repository or another location) and after that performs full VM recovery from disk.

From the **Backup repository** list, select the repository that should be used as a temporary storage (staging area) for VM backup before the VM is restored to the virtual infrastructure.

You can also select a target folder on any server connected to Veeam Backup & Replication. To do so, choose the **Select folder** option from the **Staging area** list and choose the location to which backups should be restored before full VM recovery.

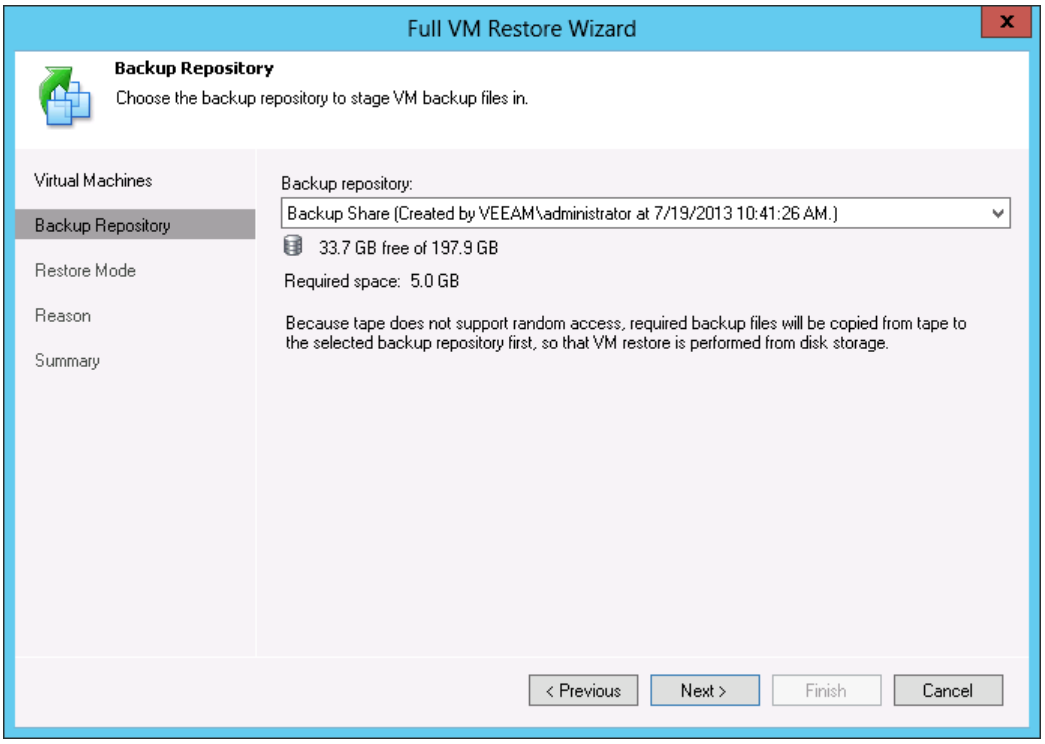

Step 5. Specify Restore Mode and Other Recovery Options

Go through the remaining steps of the **Full VM Restore** wizard. The procedure is identical to full VM recovery. For more details, refer to Performing Full VM Restore (VMware vSphere).

Restoring Files from Tape

Veeam Backup & Replication allows you to restore from tape files and folders. You can restore items that were previously archived to tape either with Veeam Backup & Replication, or with any other backup solution, provided these items were written to tape in the native MTF format.

To restore files backed up to tape, use the **Files from Tape Restore** wizard. This section will guide you through all steps of the wizard and provide explanation on available options.

To restore files from tape, follow the next steps:

Step 1. Launch the Files from Tape Wizard

To run the **Files from Tape** wizard, on the **Home** tab, click **Restore** and choose **Tape** > **Restore Files**.

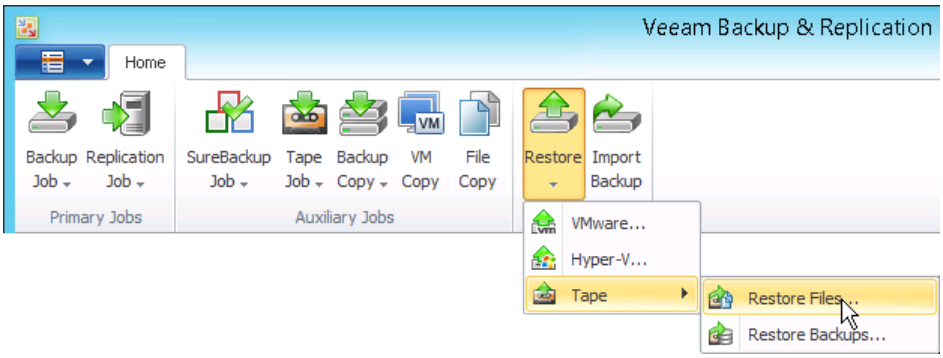

Alternatively, you can:

- Open the **Files** view, expand the **Tapes** node and browse to the necessary files on tape. Select the files and click **Restore files from tape** on the ribbon
- Open the **Files** view, expand the **Tapes** node and browse to the necessary files on tape. Select the files and choose **Restore files from tape**

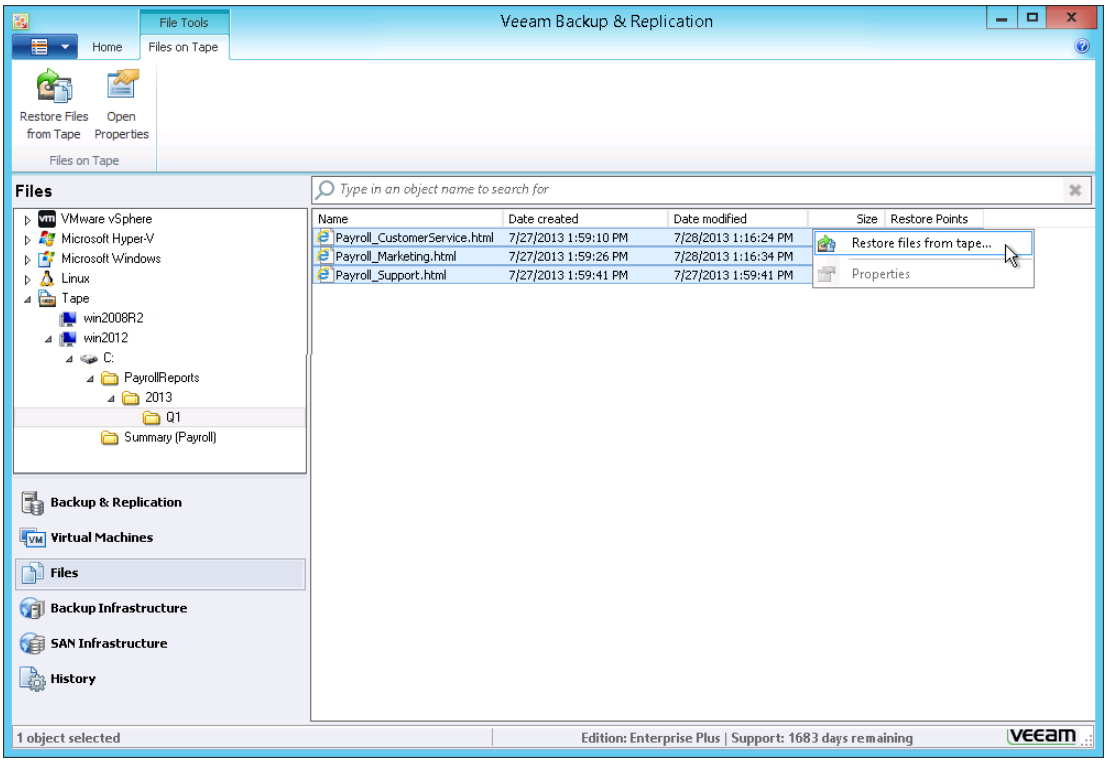

Step 2. Choose Files to Restore

At the **Objects to Restore** step, choose files and folders that you want to restore. Click **Add** and browse to the file or folder that should be restored. The selected item will be added to the list. To quickly find file or folder, use the search field at the top of the list: enter an object name or a part of it and click the search button on the right or press **[ENTER]**.

If you have chosen files to restore in the **Files** view, the list of objects to restore will be populated with selected files.

To remove a file or folder from the list, select it and click **Remove**.

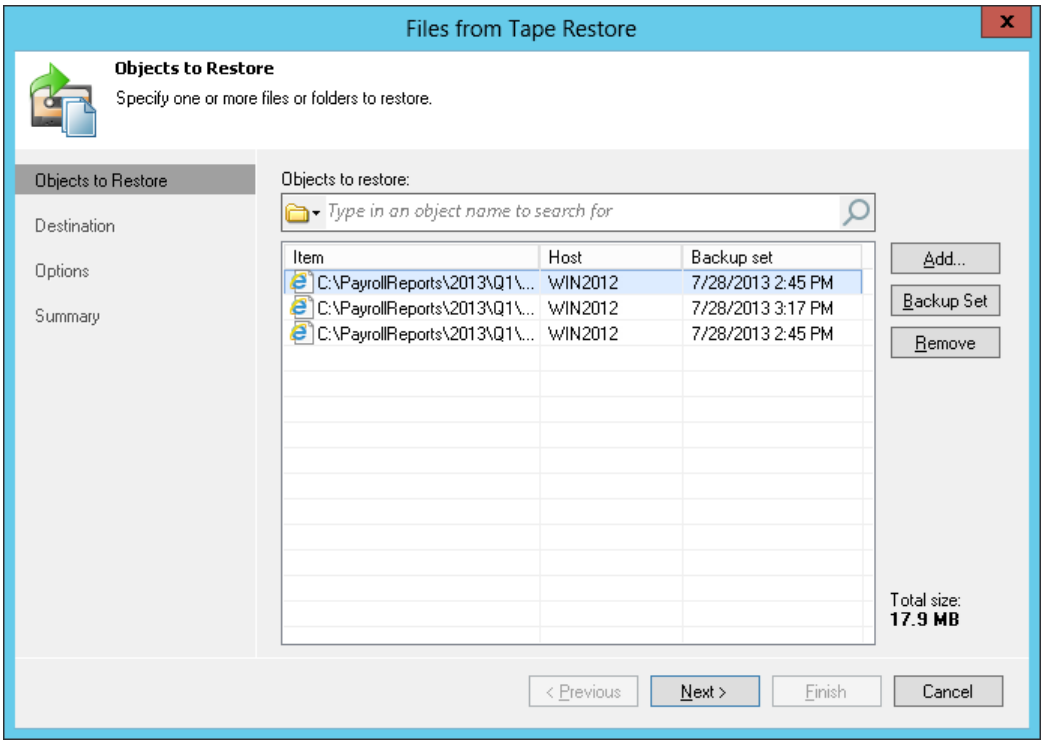

By default, Veeam Backup & Replication will restore the latest version of files available on tape. If you want to restore files from another restore point, select the necessary file and click **Backup Set**. In the list of available backup sets, select the necessary archiving session and click **OK**.

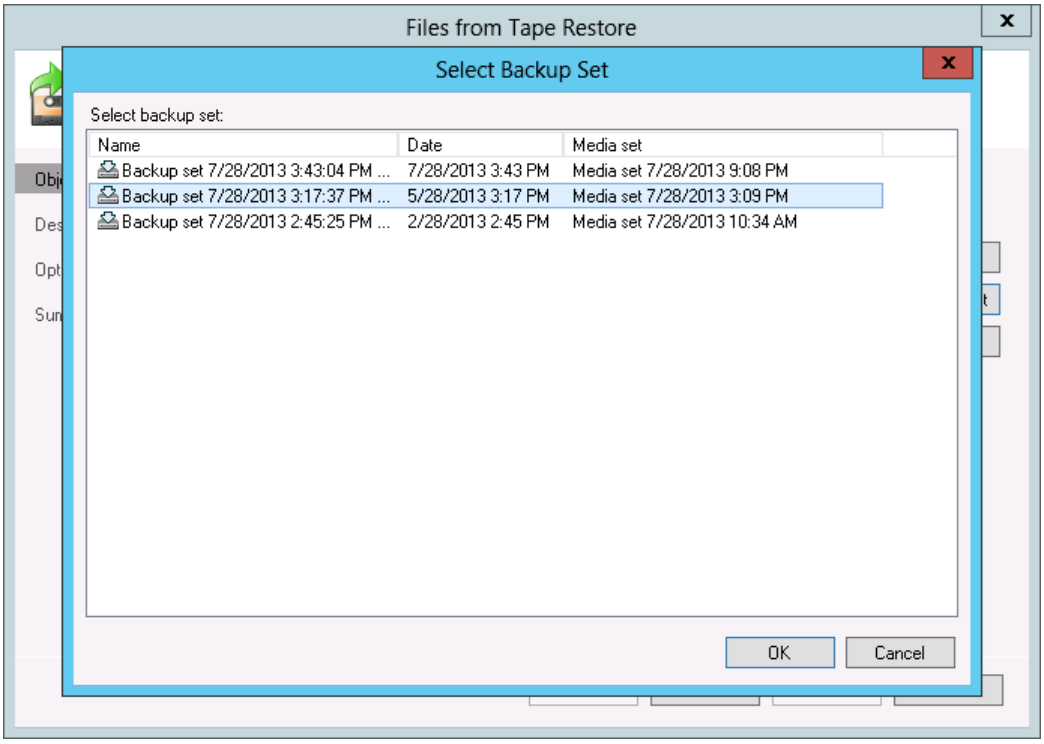

Note When you restore an entire folder from tape, Veeam Backup & Replication restores all files that have ever existed in the folder and been archived to tape. To restore files in the folder, Veeam Backup & Replication scans the selected backup set and backup sets that were created previously. If a file is not found in the selected backup set, Veeam Backup & Replication will restore the most recent version of the file from a backup set preceding the selected backup set.

Step 3. Specify Restore Destination

At the **Destination** step of the wizard, specify destination where the archived files will be restored. You can use one of the following options:

- **Original location**. Use this option to restore a file to the location where the original file resides (or resided). This type of restore is only possible if the original machine is connected to Veeam Backup & Replication and powered on.
- **This server**. Use this option if you want to restore the file to the Veeam backup server, shared folder or to any other machine connected to Veeam backup server. From the server list, choose a machine to which files should be restored and specify path to the target folder.

  If you choose to restore files to a shared folder, make sure that the account under which Veeam Backup Service runs has write permissions to the target folder. If the account does not have sufficient permissions, Veeam Backup & Replication will prompt you to enter credentials for the account that can be used for writing to the target folder.

When restoring files to a location other than original, Veeam Backup & Replication preserves the folder hierarchy. To restore files to the specified target folder without keeping the folder structure, clear the **Preserve folder hierarchy** check box.

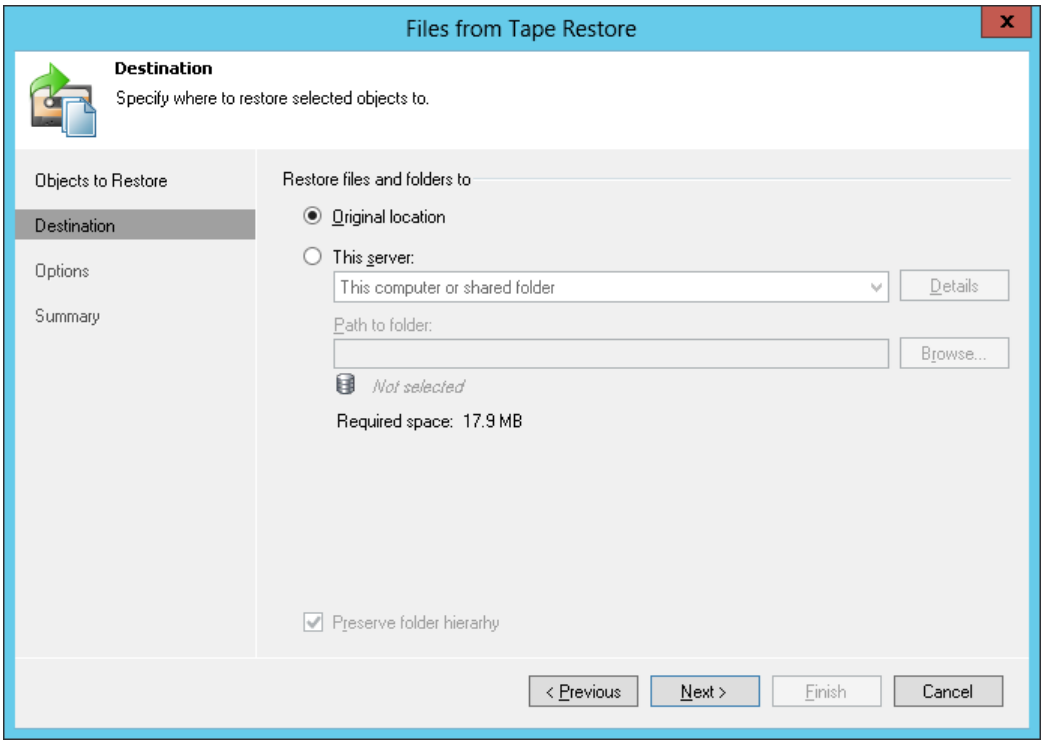

Step 4. Specify Restore Options

At the **Options** step of the wizard, specify overwrite options in case the file already exists in the target folder:

- **Leave the existing file**. Select this option if you do not want to overwrite the existing file with the restored one
- **Overwrite the existing file if older than the backed up file**. Select this option if you want to overwrite the existing file only if it is older than the restored file
- **Always overwrite the existing file**. Select this option if you want to overwrite the existing file with the restored file in all cases

Select the **Restore file and folder security** check box if you want the restored files to keep their original ownership and security permissions. If this option is disabled, Veeam Backup & Replication will change security settings: the user who launched the Veeam Backup & Replication console will be set as the owner of the restored objects, while access permissions will be inherited from the target folder to which the objects are restored.

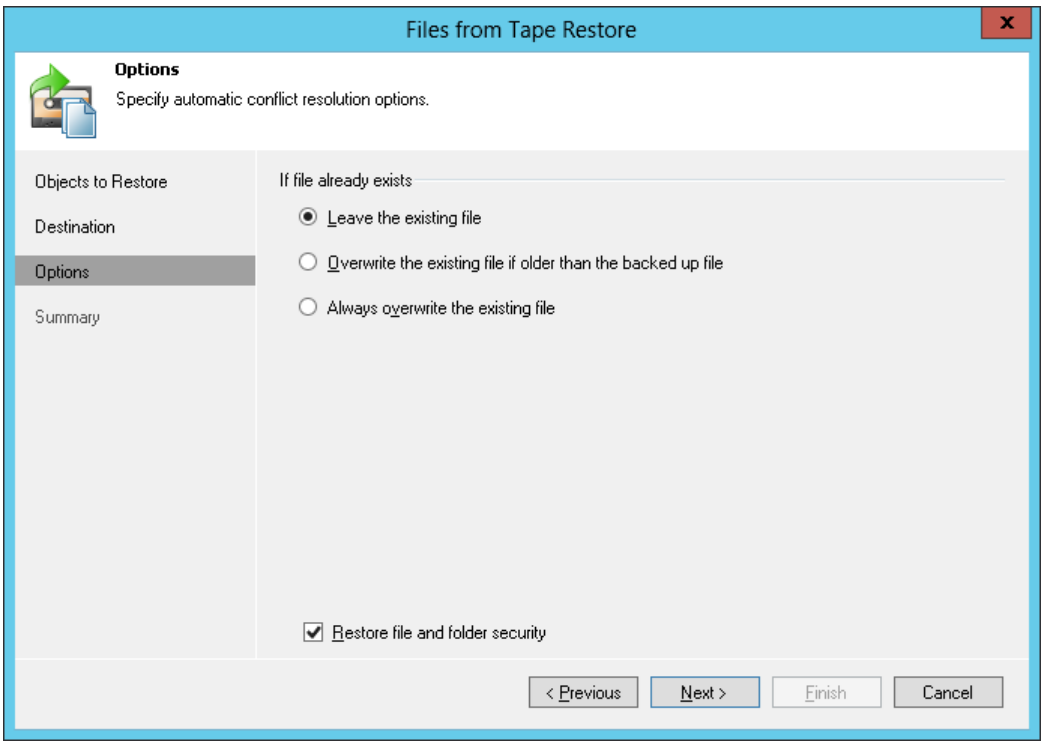

Step 5. Finish Working with the Wizard

Review the settings and click **Finish** to restore selected files and folders.

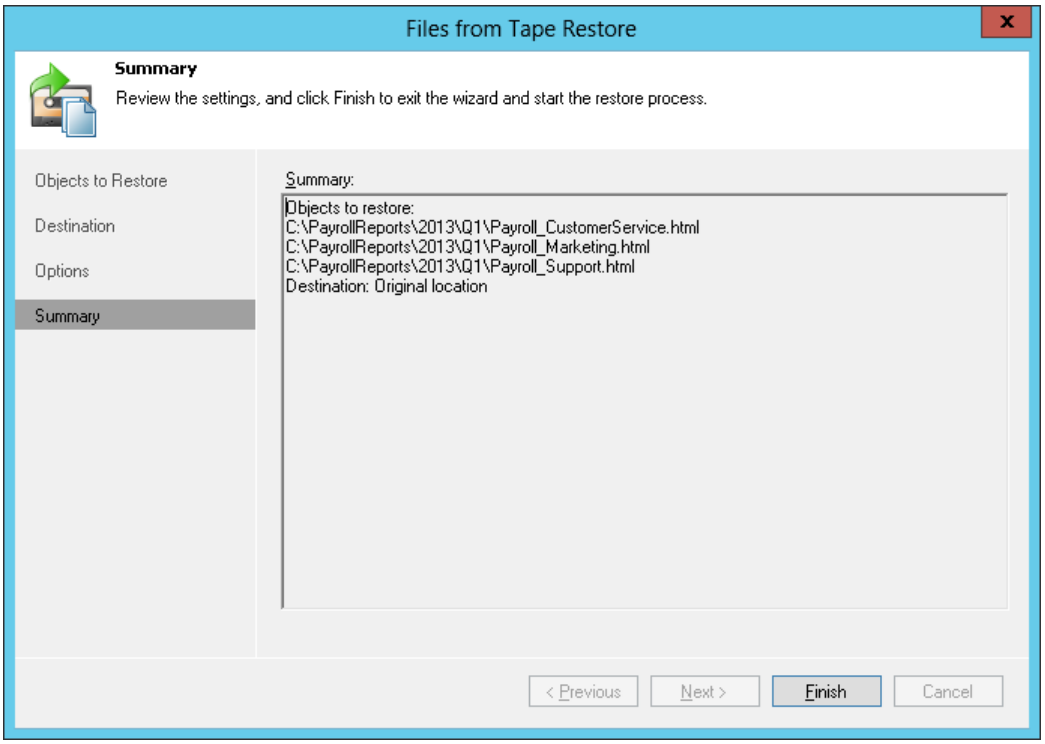

Performing File Copy Operations

As an added benefit, Veeam Backup & Replication provides you with a possibility to perform file copying operations.

File copying is the most natural way to deliver an ISO file to a host or exchange VMs and templates between hosts, Windows and Linux servers. Veeam Backup & Replication ensures security by using a one-time password feature, and works over 6 times faster than SCP.

Adding Servers and Folders

To take advantage of managing files with Veeam Backup & Replication, you should first add VirtualCenter, ESX(i), Linux or Windows servers you are going to work with. For details, see Adding Servers.

You can also add folders to the management tree, which will help you to better organize your work.

To add a server or a folder:

1. Open the **Files** view.
2. Right-click on the blank area in the inventory pane and select **Add New Folder** or **Add Server**.

Copying Files and Folders

Using Veeam Backup & Replication, you can move files and folders between and within ESX(i), Windows and Linux hosts:

1. Open the **Files** view.
2. Expand the file tree in the inventory pane.
3. Right-click the item you want to copy and select **Copy**. In the inventory pane, right-click a destination folder and select **Paste**. Alternatively, you can copy files and folders with drag-and-drop operations.

You can also automate or postpone the copy job. For details, see Creating File Copy Jobs.

Creating File Copy Jobs

To schedule a copying process of files and folders, you should create a copy job by means of the **New File Copy Job** wizard. You can perform the created job immediately, schedule or save it. This section will guide you through all steps of the **New File Copy Job** wizard and provide explanation on the offered options.

To copy files and folders, follow the next steps:

Step 1. Launch the New File Copy Job Wizard

To launch the **New File Copy Job** wizard, open the **Home** tab and click **File Copy**.

Step 2. Specify Job Name and Description

At the first step of the wizard, enter the name and description of the created job. By default, the following description is initially provided for the created job: time at which the job was created and user who created the job.

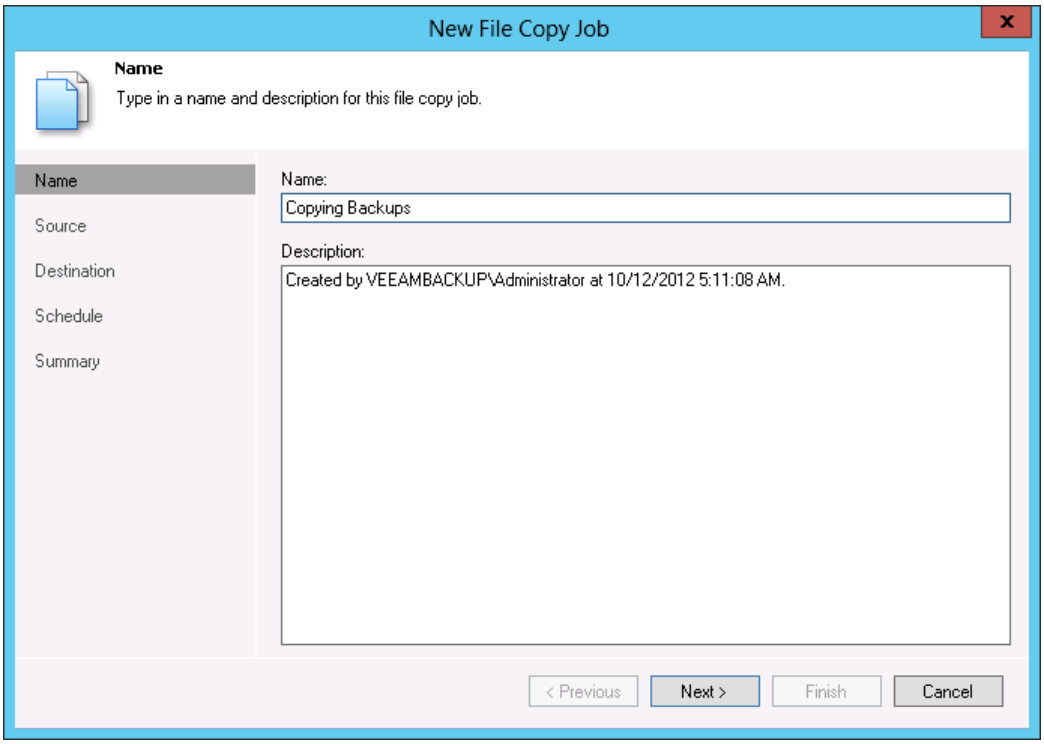

Step 3. Select Files and Folders to Be Copied

At this step, you should select files and folders that you want to copy. From the **Host** list, choose a host on which the necessary file or folder resides. Click **Add** to browse to the file or folder that should be copied. The selected item will be added to the list. To remove a file or folder from the list, select it and click **Remove**.

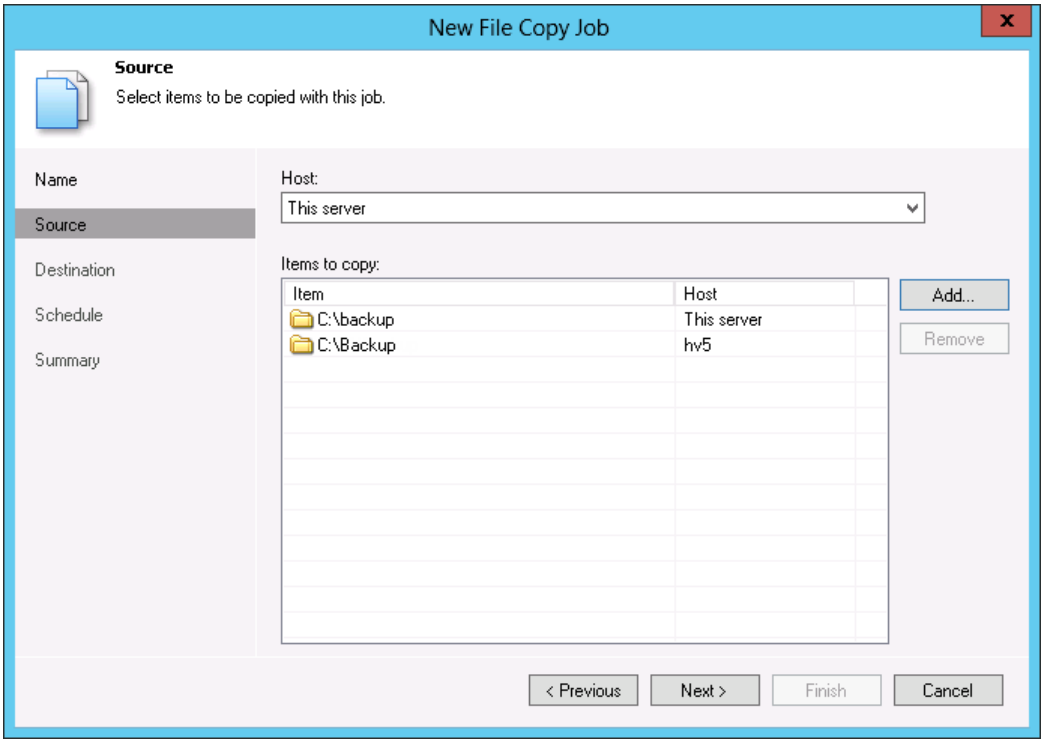

Step 4. Select Destination for Copying

Select a destination host and location to which source files or folders should be copied. Click **Details** on the right of the **Server** field to view or edit the properties of the server.

Click **Browse** next to the **Path to folder** field to browse to a folder where copied items should be stored. To create a dedicated folder for copied files or folders, use the **Make New Folder** button at the bottom of the **Select Folder** window.

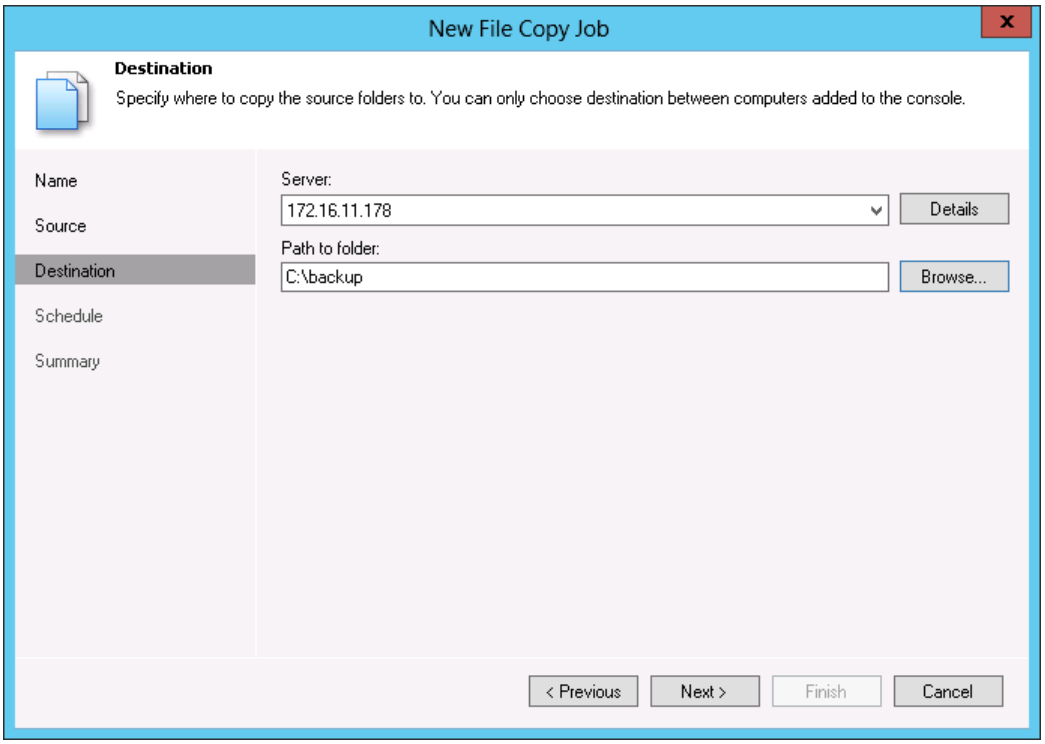

Step 5. Define the Job Schedule

At the **Schedule** step of the wizard, you can select to manually run the file copy job, schedule the job to start at a specific time — for example, the least busy hours to reduce impact on the virtual environment, or define a schedule for the job to run on a regular basis.

To specify the job schedule, select the **Run the job automatically** check box. If this check box is not selected, the job is supposed to be run manually.

You can define the following scheduling settings for the job:

- You can choose to run the job at specific time on defined week days, monthly and with specific periodicity.
- You can choose to run the job continuously. In this case, the next run of the file copy job will be started as soon as the previous one is completed, maintaining your files always in the most recent state. To run the job continuously, select the **Periodically every** option and choose **Continuously** from the list on the right.
- You can choose to run the job repeatedly throughout a day with a set time interval. At the **Schedule** step of the wizard, select the **Periodically every** option, enter the necessary time interval and select the necessary time unit: *Hours* or *Minutes*. Click **Schedule** on the right and use the time table to define the permitted time window for the job. If you choose to run the job at an hourly interval, in the **Start time for hourly jobs** field, specify the exact time when the job should start.
  For example, you want to start a job every 2 hours from 9AM to 6PM. At the **Schedule** step, select the **Periodically every** option, enter 2 in the field on the right and select *Hours* from the list. Click **Schedule** and use the **Permitted** and **Denied** options to mark the time window from 9AM to 6PM. In the **Start time for hourly jobs** field, specify the exact start time of the job, for example, 15 minutes. The job you have scheduled will start at 9:15 AM, 11:15 AM, 1:15 PM, 3:15 PM and 5:15 PM.
- You can chain jobs. In the common practice, jobs start one after another: when the job "A" finishes, the job "B" starts and so on. If you want to create a chain of jobs, you should define the time schedule for the first job in the chain. For the rest of the jobs in the chain, at the **Schedule** step of the wizard, select the **After this job** option and choose the preceding job from the list.

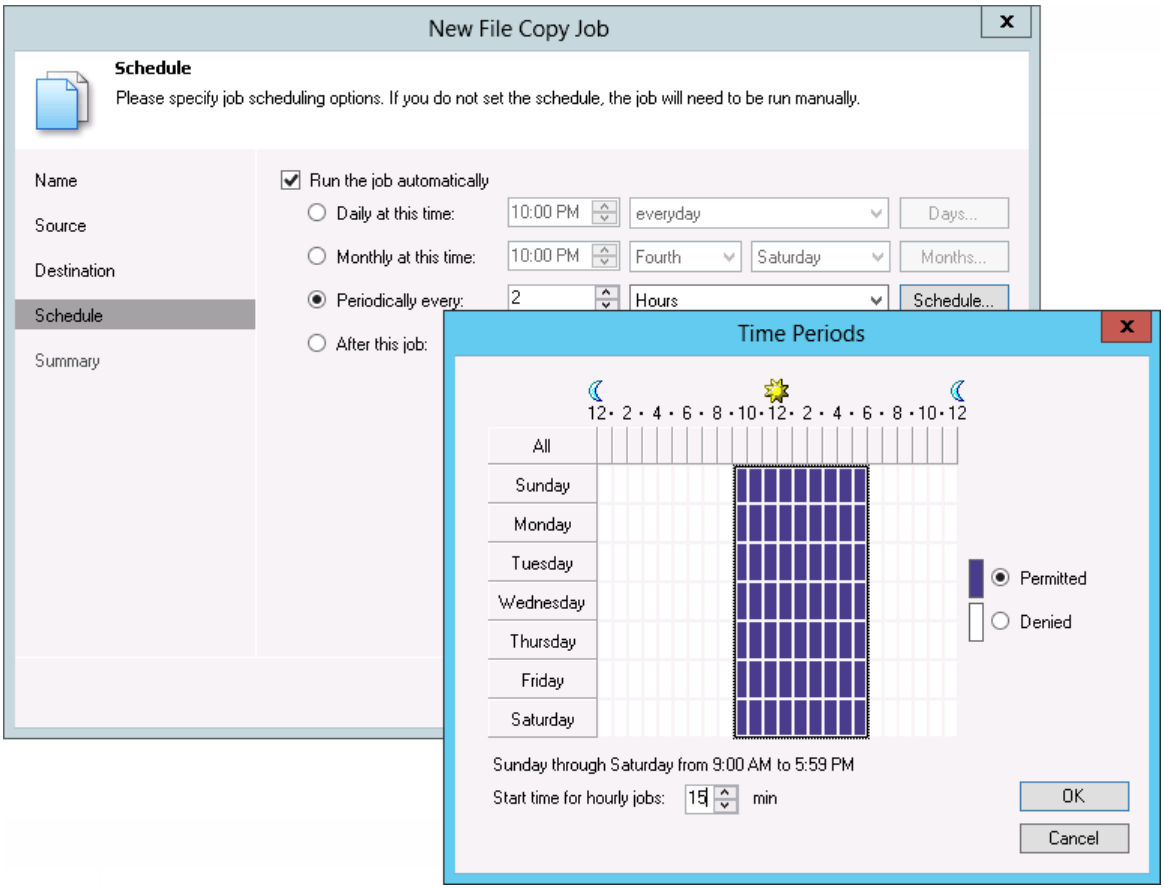

Step 6. Finish Working with the Wizard

After you have specified the schedule settings, click **Create**. Select the **Run the job when I click Finish** check box if you want to start the created job right after you complete working with the wizard. Click **Finish** to close the wizard.

Changing Server Connection Settings

The server connection properties are defined when you add a server. To view or change the server connection properties, use the server properties option:

1. Open the **Infrastructure** or **Files** view.
2. Right-click the server and select **Properties**.
3. Navigate the tabs and configure the connection settings as desired.

The available settings vary depending on the server type (ESX(i) host, vCenter Server, Windows server, Linux server). The settings and their default values are listed below.

Server Type	Connection Property/ Option	Default Value	Comments
ESX, Linux	SSH connection port	22	—
	SSH channel connection timeout (default length of time before terminating the incomplete task)	20000 ms	Customize the option when you receive a timeout error.
	Elevate account to root	Unmarked	Check the option to run the sudo command. This will enable you to get root privileges in case the remote root access to the desired server is denied.
	Add account to the sudoers file automatically	Enabled along with *Elevate account to root* checked	If providing root credentials is a security issue, disable the option. Ensure the user data have been previously added to the sudoers configuration file.
ESX	Web port	443	When entering other than root user credentials, ensure the user is granted shell access in Virtual Infrastructure Client.
ESX Linux Windows Hyper–V	Data channel port range	2500 – 5000	Only 1 port is required for data transfer.
	Packet size (Kb)	64	Adjust this option if you have any stability issues when copying.
	Run server on this side for copying between servers	Unmarked	Check the option to copy files to the server behind the NAT or router.
Virtual Center	Web port	443	—

Important! SSH Connection port and timeout, Data Transfer port range and transfer packet size affect the execution of copying procedures dramatically. If you encounter any problems with copying or connection, first check these parameters.

Editing File Attributes

If necessary, you can change user (owner), group, and other read, write and execute permissions for ESX(i) and Linux files and folders directly from Veeam Backup & Replication:

1. Open the **Files** view.
2. Click the necessary ESX(i) or Linux file or folder and click the **Properties** button on the ribbon. You can also right-click the necessary ESX(i) or Linux folder or file in the inventory pane and select **Properties**.
3. Change the permission settings as necessary.

Specifying Veeam Backup & Replication Options

This section provides a detailed description about general Veeam Backup & Replication options.

Specifying E-Mail Notification Settings

With Veeam Backup & Replication you can select to receive email messages in case of success or failure of a created backup or replication job. To be able to receive email notifications, you should configure general email notification settings and select to receive a notification when creating a corresponding job.

Tip To receive email notification about all performed jobs at once, use Veeam Backup Enterprise Manager.

Configuring General E-Mail Notification Settings

To configure general email notification settings, select **Options** from the main menu. Select the **Enable email notification** check box and specify email notification settings:

1. In the **SMTP Server** field, enter the DNS name or IP address of the SMTP server that will be used for sending email messages.
2. Use the **Advanced** button to specify user credentials and connection options – port number and connection timeout. Additionally, you can enable the **Connect using SSL** option to use the secure SSL connection for email operations.
3. In the **From** field, specify the email from which email notifications should be sent.
4. In the **To** field, specify the recipient address(es). Use a semicolon to separate multiple addresses. Recipient(s) specified in this field will receive notification about every job managed by the Veeam backup server. You can leave the field empty if required.
   For every particular job, you can also specify additional recipients (for details, see the Configuring Job Notification Settings section).
5. In the **Subject** field, specify the subject for a sent message. You can use the following variables in the subject: %Time% (completion time), *%JobName%*, *%JobResult%*, *%VmCount%* (the number of VMs in the job) and *%Issues*% (the number of VMs in the job that have been processed with the *Warning* or *Failed* status).
6. Select the **Notify on success**, **Notify on warning** and/or **Notify on failure** check boxes to receive email notification in case a job is run successfully, not successfully or with a warning.

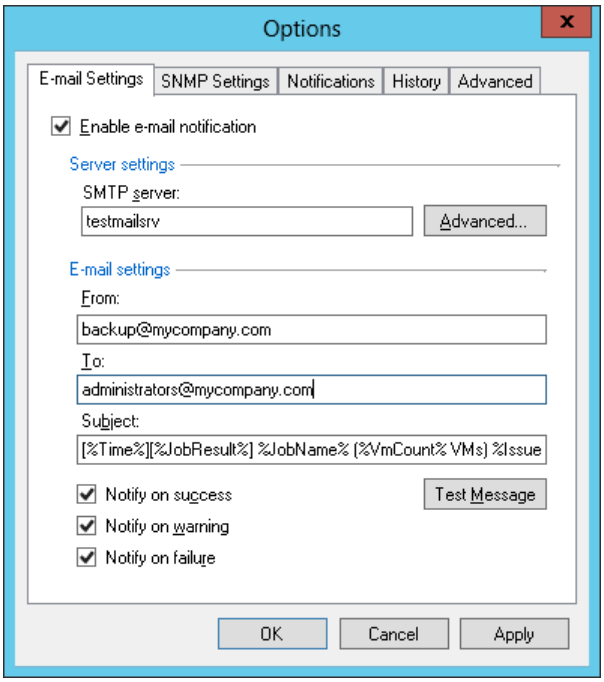

Veeam Backup & Replication allows sending a test email to check if all settings have been configured correctly: click **Test Message** to receive a test email.

Configuring Job Notification Settings

To configure job notification settings:

1. At the step of specifying storage settings for the created job, click **Advanced**.
2. On the **Notifications** tab, select the **Send email notifications to the following recipients** check box.
3. In the field below enter an email to which a notification should be sent. Use a semicolon to separate multiple addresses.

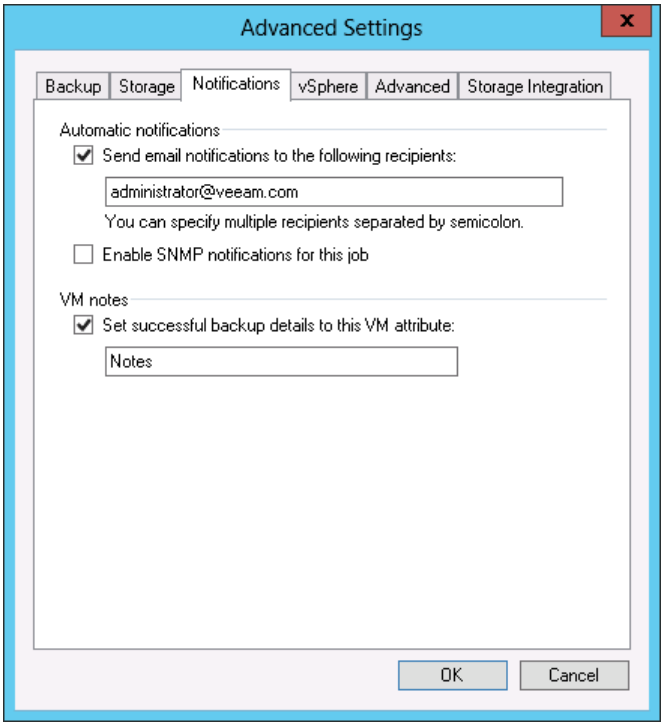

Specifying SNMP Settings

Veeam Backup & Replication provides a possibility to monitor execution of backup and replication jobs using SNMP traps. You can select to receive SNMP notifications once each job is completed and backup or replica is created. SNMP traps can be used to feed data into other popular system monitors, such as CA Unicenter, BMC Patrol, IBM Tivoli or HP OpenView.

To be able to receive SNMP traps, you should:

- Configure General SNMP Settings
- Configure SNMP Service Properties
- Specify SNMP Settings for Jobs

Configuring General SNMP Settings

To configure general SNMP settings:

1. Select **Options** from the main menu.
2. Click the **SNMP Settings** tab.
3. In the **Receiver** field, specify an IP address of the SNMP recipient.
4. In the field on the right, enter the port number to be used.
5. In the **Community String** field, enter the community identifier.

Trap notifications can be sent to 5 different destinations.

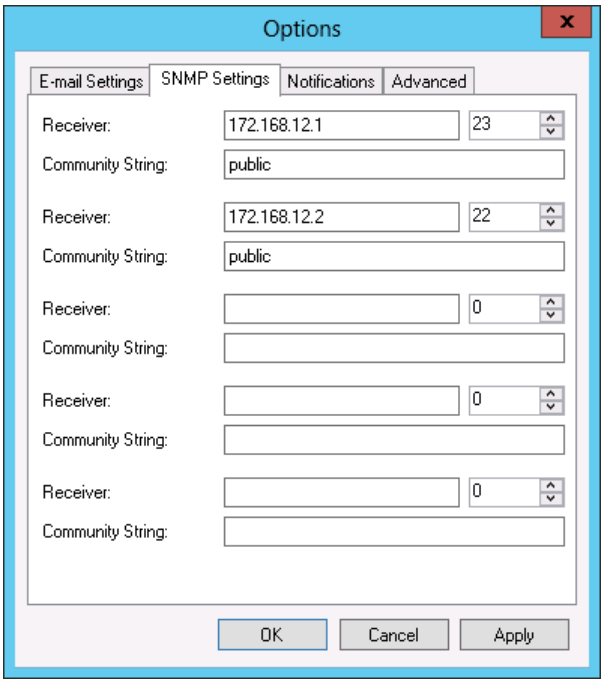

Configuring SNMP Service Properties

To configure SNMP service properties on the trap recipients' computers:

1. Install standard Microsoft SNMP agent from the Windows distribution.
2. From the **Start** menu, select **Control Panel** > **Administrative Tools** > **Services**.
3. Double-click **SNMP Service** to open the **SNMP Service Properties** window.
4. Click the **Traps** tab.

5. Add the public string to the **Community name** list and the necessary host name to the **Trap destinations** list.
6. Click the **Security** tab.
7. Make sure the **Send authentication trap** option is selected.
8. Add the public string to the **Accepted community names** list.
9. Select the **Accept SNMP packets from any hosts** option.
10. Click **OK** to accept changes.

Specifying SNMP Settings for Jobs

To be able to receive SNMP traps with results for a specific job:

1. At the step of specifying storage settings for the created job, click **Advanced**.
2. On the **Notifications** tab, select the **Enable SNMP notifications for this job** check box.

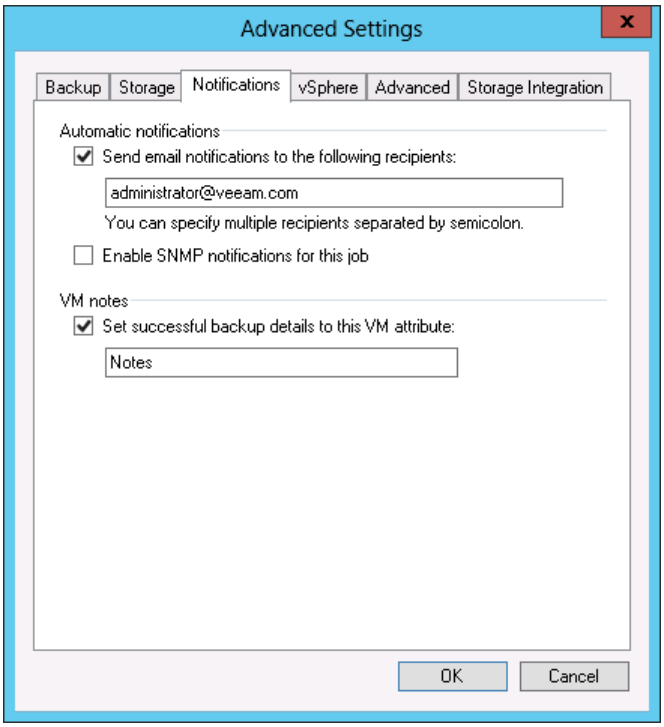

Specifying Other Notification Settings

When a job is run, Veeam Backup & Replication checks disk space on the backup storage and on production datastores. If the disk space is below a specific value, Veeam Backup & Replication will display a warning message in the job session details.

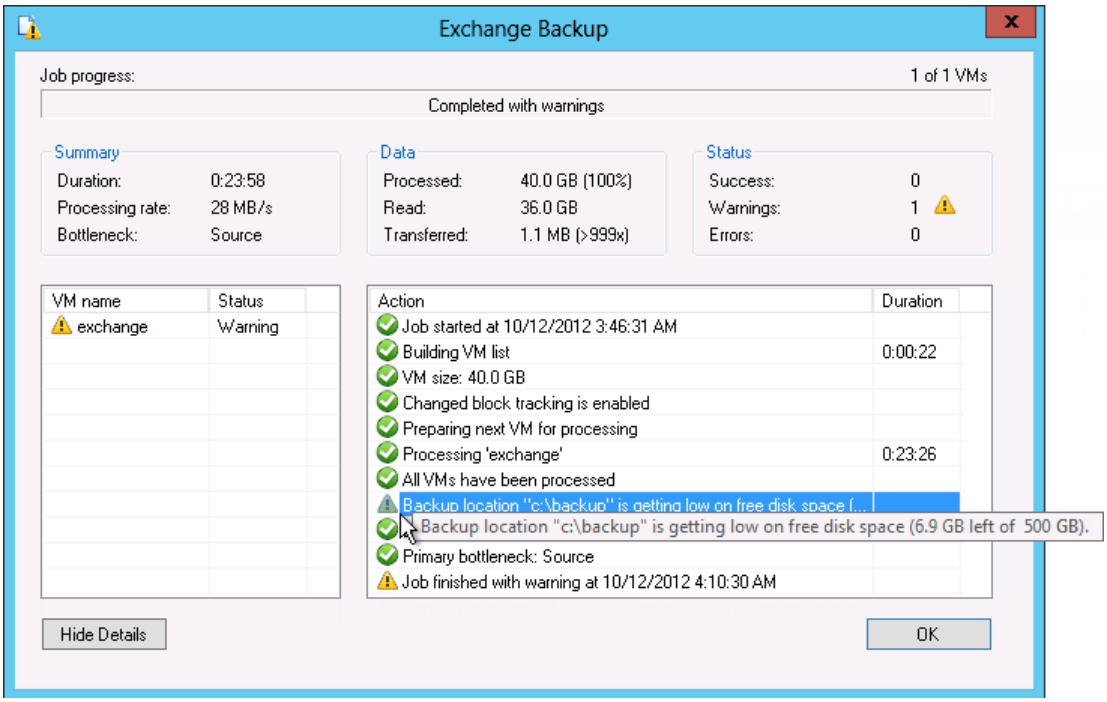

To specify the disk space threshold:

1. Select **Options** from the main menu.
2. Click the **Notifications** tab.
3. In **Backup storage** and **Production datastores** sections, select the **Warn me if free disk space is below ... percent/GB** options and specify a desired disk space threshold.

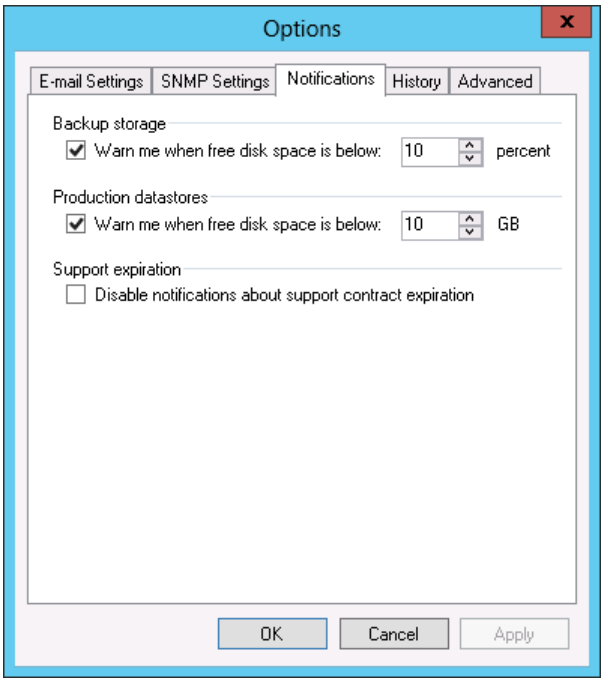

By default, e-mail recipients specified on the E-mail Settings tab will be informed about the support expiration date in every received email notification. Veeam Backup & Replication will start sending such notifications 14 days before the expiration date. Expiration information will be also shown on the splash screen and on the **License Information** dialog shown after you select **Help** > **License** from the main menu.

To stop receiving the notification, select **Options** from the main menu, and on the **Notifications** tab select the **Disable notifications about support contract expiration** check box.

Specifying Session History Settings

Using advanced settings of Veeam Backup & Replication, you can specify session history settings:

1. Select **Options** from the main menu.
2. Click the **History** tab.
3. In the **Sessions** and **Session history retention** sections, specify the number of sessions to display in the **Sessions** list and the number of sessions to keep in the database.

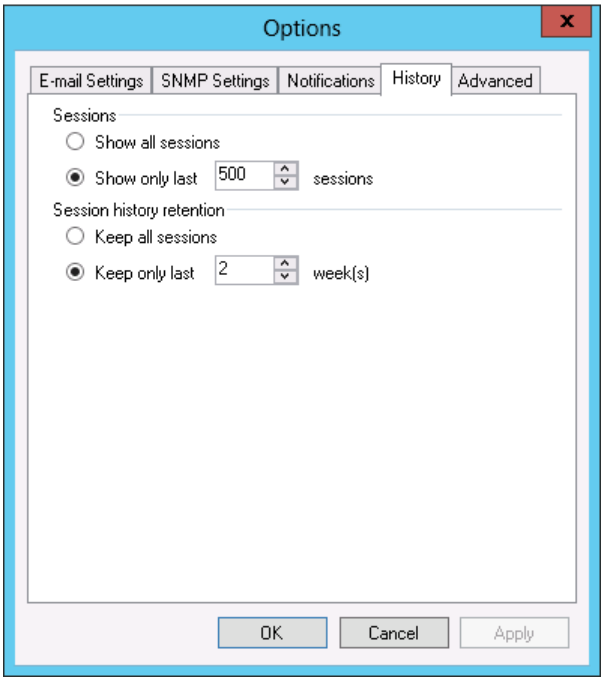

Enabling Parallel Processing

In previous versions, VMs and VM disks used to be processed one by one within the same job, that is, the tasks of the job were accomplished sequentially; each job was handled by two CPU cores (recommended configuration).

Note Each task means one VM disk being processed.

Starting with Veeam Backup & Replication 7.0, multiple VMs and VM disks can be processed in parallel, optimizing your backup infrastructure performance and increasing the efficiency of resource usage; each data processing task within a job requires one CPU core - consider this value when configuring job settings.

Important! This is a global setting, so if configured, it will take effect for all backup & replication jobs.

To use this capability, on the **Advanced** tab of Veeam Backup & Replication **Options**, select the **Enable parallel VM and virtual disk processing** check box:

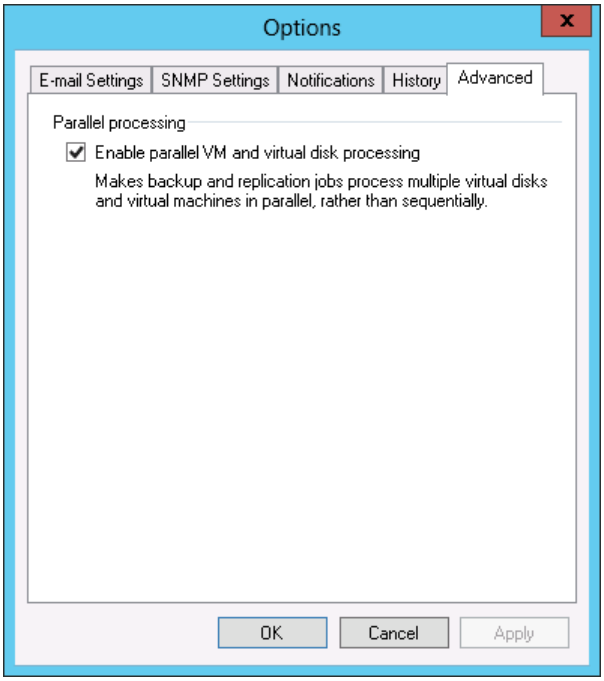

Note that if you upgrade to Veeam Backup & Replication 7.0 from an earlier version of the product, this option is turned off by default.

Reporting

When a job is being run, jobs statistics and operation data is written to the Veeam Backup & Replication database. Veeam Backup & Replication allows viewing realtime statistics for any performed job and generating reports with statistics for any job or separate job session.

Realtime Statistics

To view real-time statistics for a job that is being run, do one of the following:

- Open the **Backup & Replication** view and select the **Jobs** node. Double-click the necessary job in the working area.
- Open the **Backup & Replication** view and select the **Jobs** node. Right-click the job in the working area and select **Statistics**.

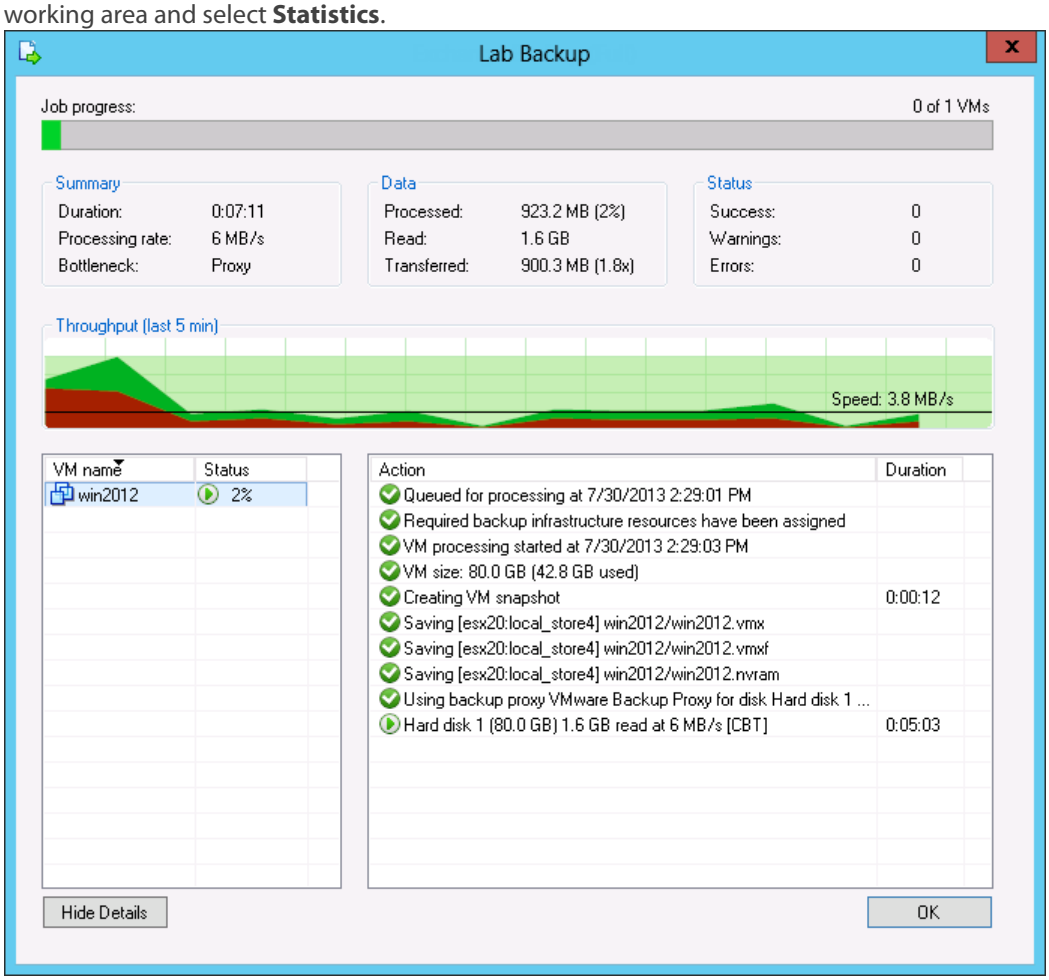

A report generated for a job contains detailed data on job sessions: job progress, duration, processing rate, performance bottlenecks, the amount of data processed, read and transferred, and details of the session performance (for example, errors that have occurred in the process of operation).

You can double-click on the graph to view data rate for the last 5 minutes or the whole processing period.

- Green area – amount of data read from source
- Brown area – amount of data written to target
- Horizontal line – current data processing speed

In addition to overall job statistics, it contains detailed data on each object processed within the frames of a job (that is, a virtual machine). To view backup progress for a specific VM, select it in the list on the left.

Veeam Backup & Replication also allows you to view detailed statistics on every job session. To view statistics for a selected job session, do either of the following:

- Open the **History** view, select the **Jobs** node and double-click the necessary session in the working area.
- Open the **History** view, select the **Jobs** node, right-click the session and select **Details**.

Tip To switch between past job sessions in the **Statistics** window, use left and right arrow keys on the keyboard.

Session Report

The session report contains data on a single job session:

- Cumulative session statistics – session duration details, details of the session performance, amount of read, processed and transferred data, backup size, compression and deduplication ratios.
- Detailed statistics for every VM processed within the session – processing duration details, backup data size, the amount of read and transferred data, the list of warnings and errors (if any).

To generate a report:

1. Open the **History** view.
2. Select the **Jobs** node.
3. Right-click the necessary session in the working area and select **Report**.

Job Report

The job report contains data on all sessions initiated for a specific job. To generate a job report:

1. Open the **Backup & Replication** view.
2. Select the **Jobs** node.
3. Right-click the necessary session in the working area and select **Report**.

Users and Roles

There are four levels of security that can be granted to users who work with Veeam Backup & Replication:

- Veeam Restore Operators
- Veeam Backup Viewers
- Veeam Backup Operators
- Veeam Backup Administrators

The security scheme in Veeam Backup & Replication is mainly used for work with Veeam Backup Enterprise Manager. To learn more about security settings in Veeam Backup Enterprise Manager, see the Veeam Backup Enterprise Manager documentation.

In Veeam Backup & Replication, security settings are checked for managing (starting and stopping) jobs and performing restore operations.

Role	Operations
Veeam Restore Operator	Can perform restore operations using existing backups and replicas.
Veeam Backup Viewer	Has the "read-only" access to Veeam Backup & Replication — can view existing and performed jobs and review the job session details.
Veeam Backup Operator	Can start and stop existing jobs.
Veeam Backup Administrator	Can perform all administrative activities in Veeam Backup & Replication.

To specify user security settings:

1. Select **Users and Roles** from the main menu.
2. Click **Add**.
3. In the **User name** field, enter the name of a user or group in the *DOMAIN\USERNAME* format.
4. From the **Role** list, select the necessary role to be assigned.

Note By default, during installation the *Veeam Backup Administrator* role is assigned to users listed in the local *Administrators* group.

Logging

Veeam Backup & Replication provides detailed logging of performed activities, initiated jobs, Veeam transport service work and so on. On the Veeam Backup & Replication server, log files are stored in the following folder:

- For Windows XP and Windows Server 2003: `%allusersprofile%\Application Data\Veeam\Backup`
- For Windows Vista and later: `%allusersprofile%\Veeam\Backup`

Veeam Backup & Replication keeps a separate log file for each of its components: *Veeam Shell*, *Veeam Backup Service*, *Veeam Backup Catalog Service*, *Veeam vPower NFS Service*, *Veeam Installer*, *Veeam Transport* and performed jobs.

In addition to logs stored on the Veeam Backup & Replication server, log files are also stored on all servers managed by Veeam Backup & Replication:

- On Linux servers and ESX(i) hosts, logs are stored in the following directory: `/var/log/VeeamBackup/`
- On Windows servers and Hyper−V hosts, logs are stored as follows:
  - For Windows XP and Windows Server 2003: `%allusersprofile%\Application Data\Veeam\Backup`
  - For Windows Vista and later: `%allusersprofile%\Veeam\Backup`

To browse to the log files, select **Help** > **Support Information** from the main menu. As a result, the **Export Logs** wizard will be launched.

Exporting Logs

Use log files to submit a support ticket. It is recommended that you send the whole content of the logs folders to ensure that overall and comprehensive information is provided to the support team.

To aggregate all log files in the same location, you can use the **Export Logs** wizard. To launch the wizard, select **Help** > **Support Information** from the main menu.

Then follow the next steps:

Step 1. Select Virtual Infrastructure Scope

At this step of the wizard, you should define the scope for logs export. You can export logs for the following entities:

- Specific jobs on the Veeam backup server
- Specific VMs in the virtual environment
- Specific components in the backup infrastructure

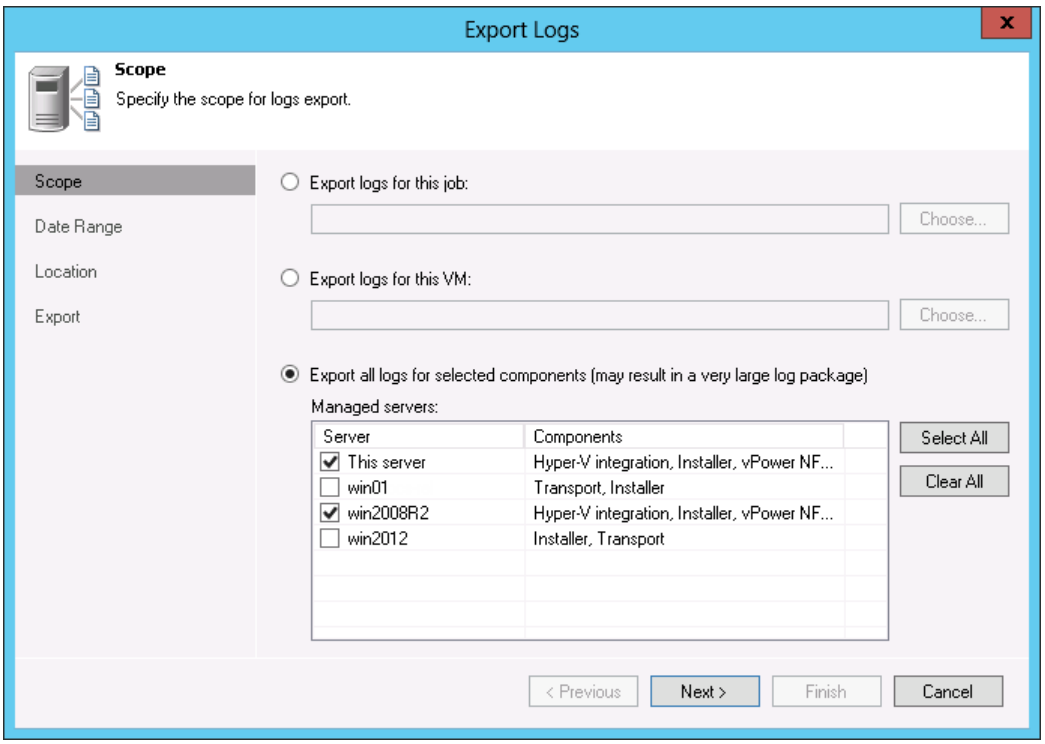

Step 2. Specify Time Interval

At this step of the wizard, you should define the time interval for which logs should be collected. You can select one of the following options:

- Collect logs for the last N days
- Collect logs for a specific period of time
- Collect all available logs

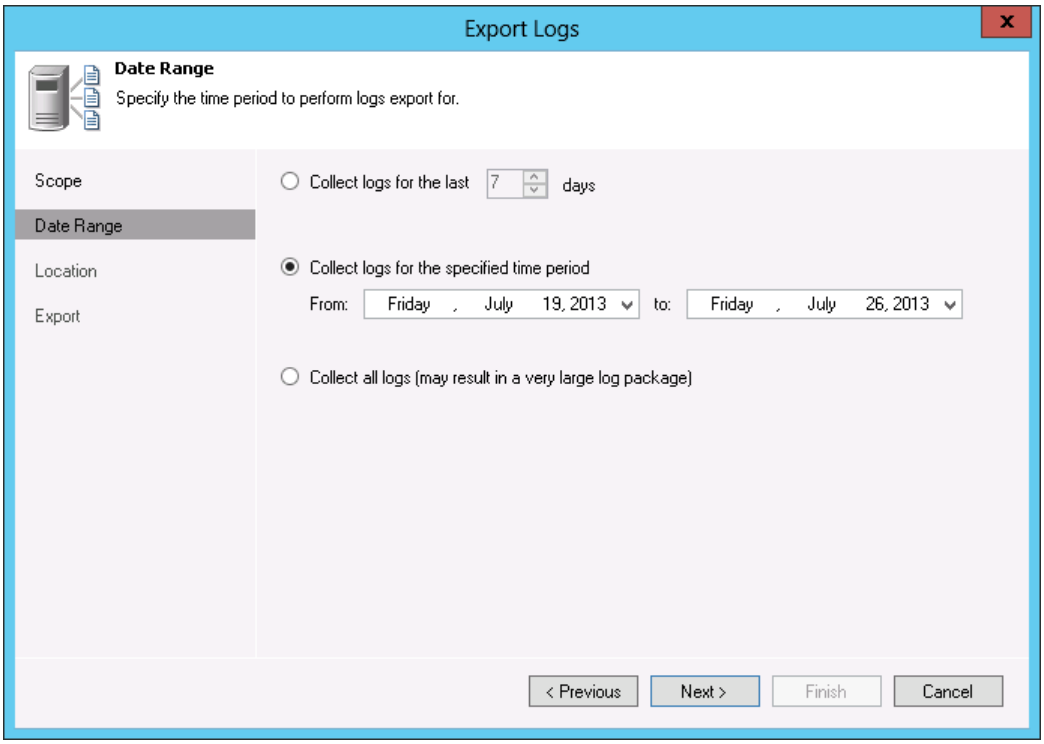

Step 3. Specify a Destination Folder

At this step of the wizard, you should specify the destination folder to which the logs will be exported. To create an archive with exported logs, which is generally required by Veeam support, select the **Prepare logs package for technical support** check box.

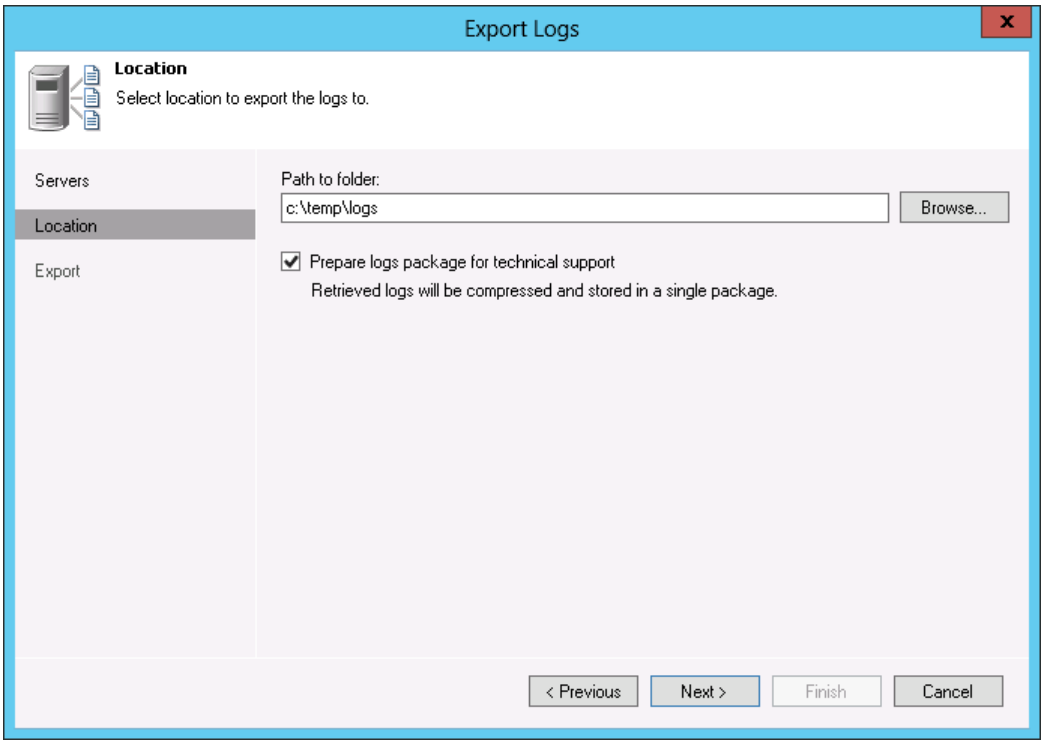

Step 4. Review the Results

When the export completes, review the results and click the **Open folder** link to browse to exported log files and log package.

Performing Configuration Backup and Restore

With Veeam Backup & Replication, you can create a configuration backup of the Veeam backup server.

When you create a configuration backup, you export the configuration data from the Veeam Backup SQL database and save it into a backup file on the repository. If the Veeam backup server fails for some reason, you can re-install the Veeam backup server and then quickly restore its configuration from the backup file. Alternatively, you can apply the configuration of one Veeam backup server to any other Veeam backup server in your backup infrastructure.

It is recommended that you regularly create a configuration backup for every Veeam backup server in your backup infrastructure. Periodic configuration backups reduce the possibility of data loss and minimize the administrative overheard if any problem with Veeam backup server(s) occurs.

Creating Configuration Backups

When you perform configuration backup, Veeam Backup & Replication retrieves configuration data for the Veeam backup server from the SQL database, writes this data into a set of .xml files and archives these .xml files to a .bco file.

Veeam Backup & Replication exports configuration data for all Veeam Backup & Replication objects:

- Backup infrastructure configuration data: vSphere hosts, HP SAN storages, backup proxies, backup repositories, virtual labs and other
- Jobs configuration: backup, replication and other jobs, registered backups, replicas, restore points and other
- Veeam Backup & Replication settings: user roles, SMTP settings and so on
- Additionally, Veeam Backup & Replication retrieves data for dynamic objects: for example, folders in the server hierarchy or snapshots in the SAN storage topology

Note When storing credentials for objects, Veeam Backup & Replication does not keep passwords. During the restore process, you will have to specify passwords manually.

The configuration backup is job-driven. Just like any other job, you can schedule it to run regularly or start it manually. You can choose the backup repository to which the configuration backup should be stored and specify the necessary retention policy settings.

By default, Veeam Backup & Replication is configured to create a configuration backup daily and store it to the default backup repository: the `C:\backup\VeeamConfigBackup\%BackupServer%` folder on the Veeam backup server.

However, for security's sake, it is recommended that you store configuration backups on the backup repository other than the default one. In this case, configuration data of the Veeam backup server(s) will be available for recovery even if the Veeam backup server fails.

Scheduling Configuration Backups

To schedule a configuration backup:

1. From the main menu of Veeam Backup & Replication, choose **Configuration Backup**.
2. Make sure that the **Enable configuration backup** check box is selected in the **Export Configuration** window.
3. From the **Backup repository** list, choose the repository to which the configuration backup should be written.
4. In the **Schedule** section, click **Configure** and specify the time schedule according to which the configuration backup should be created.
5. In the **Retention policy** section, specify the number of backups to keep on disk.

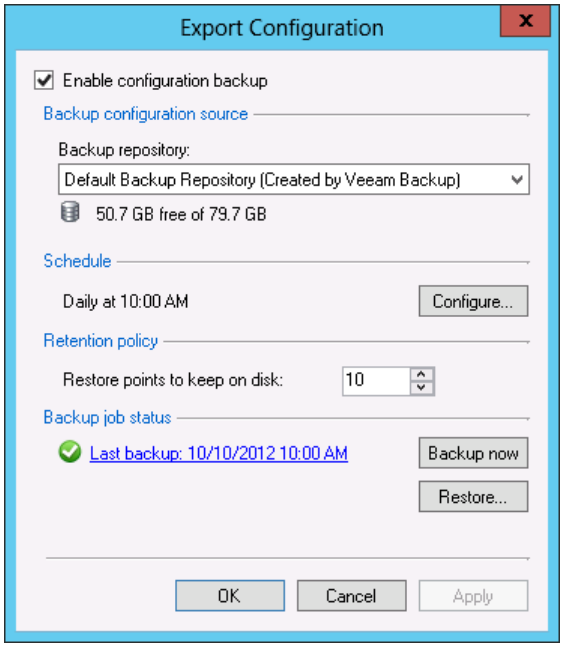

Running Configuration Backups Manually

To create a configuration backup manually:

1. From the main menu of Veeam Backup & Replication, choose **Configuration Backup**.
2. Make sure that the **Enable configuration backup** check box is selected in the **Export Configuration** window.
3. From the **Backup repository** list, choose the repository to which the configuration backup should be written.
4. In the **Backup job status** section, click **Backup now**. Veeam Backup & Replication will create a new configuration backup and store it to the chosen repository.

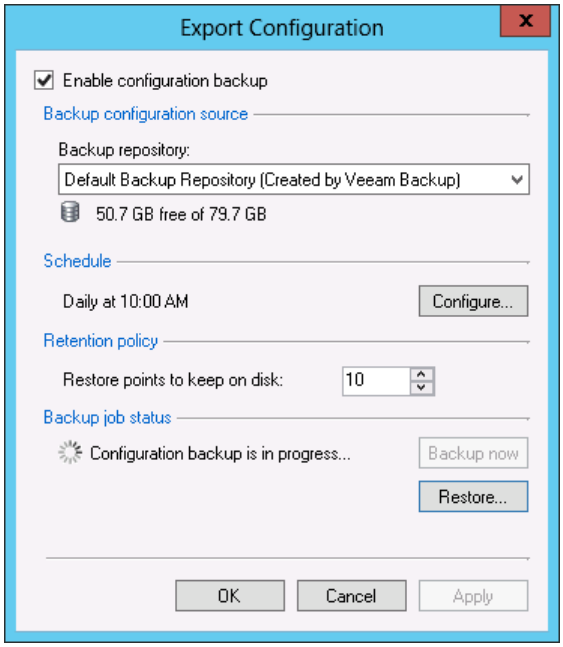

Restoring Configuration Data

If a Veeam backup server fails, you can re-deploy the Veeam backup server, restore configuration data for the Veeam backup server from the backup and apply it to the re-built server. Alternatively, you can apply configuration data to any other Veeam backup server in your backup infrastructure. After the restore process is finished, the Veeam backup server is ready for work. In terms of configuration, you get a replica of the Veeam backup server you had, without additional adjustments and fine-tuning.

During the restore process, Veeam Backup & Replication retrieves configuration data from the `.bco` file and writes it to the Veeam Backup & Replication SQL database used by the target Veeam backup server. You can write configuration data to a new Veeam Backup & Replication SQL database or restore configuration data to the current Veeam Backup & Replication SQL database.

- If you select to write configuration data to a new database, Veeam Backup & Replication will create a new database on the SQL server and populate it with data from the `.bco` file. Note that you should have sufficient permissions to create a new database on the SQL server.
- If you select to write configuration data to an existing database, Veeam Backup & Replication will first delete the current database schema and all data from the existing database. After that, it will populate the clean database with data from the .bco file. To protect the existing database from any kind of errors that can occur during restore, you can additionally select to create a SQL backup of the existing database before starting the restore process.

After the configuration data is imported into the Veeam Backup & Replication database, Veeam Backup & Replication starts the rescan process for the following objects:

- Hosts and servers
- Repositories
- Backups that had been imported to the Veeam Backup & Replication console but not stored on registered backup repositories
- Replicas

Performing Restore

Performing Restore

Before you start the restore process, make sure that you have performed the following tasks:

1. Make sure that the repository with a configuration backup (.bco) you plan to use for restore is added to the Veeam Backup & Replication console. To learn more, see Adding Backup Repositories.
2. Stop all jobs that are currently running. During restore of configuration, Veeam Backup & Replication temporary stops the Veeam Backup services and jobs.
3. Check the version of the Veeam backup server. You can restore the backup configuration on the Veeam backup server of the same version.

To restore configuration data of the Veeam backup server, follow the next steps:

Step 1. Launch the Configuration Database Restore Wizard

To launch the **Configuration Database Restore** wizard:

1. From the main menu of Veeam Backup & Replication, choose **Configuration Backup**.
2. In the **Backup job status** section, click **Restore**. The **Configuration Database Restore** wizard will be launched.

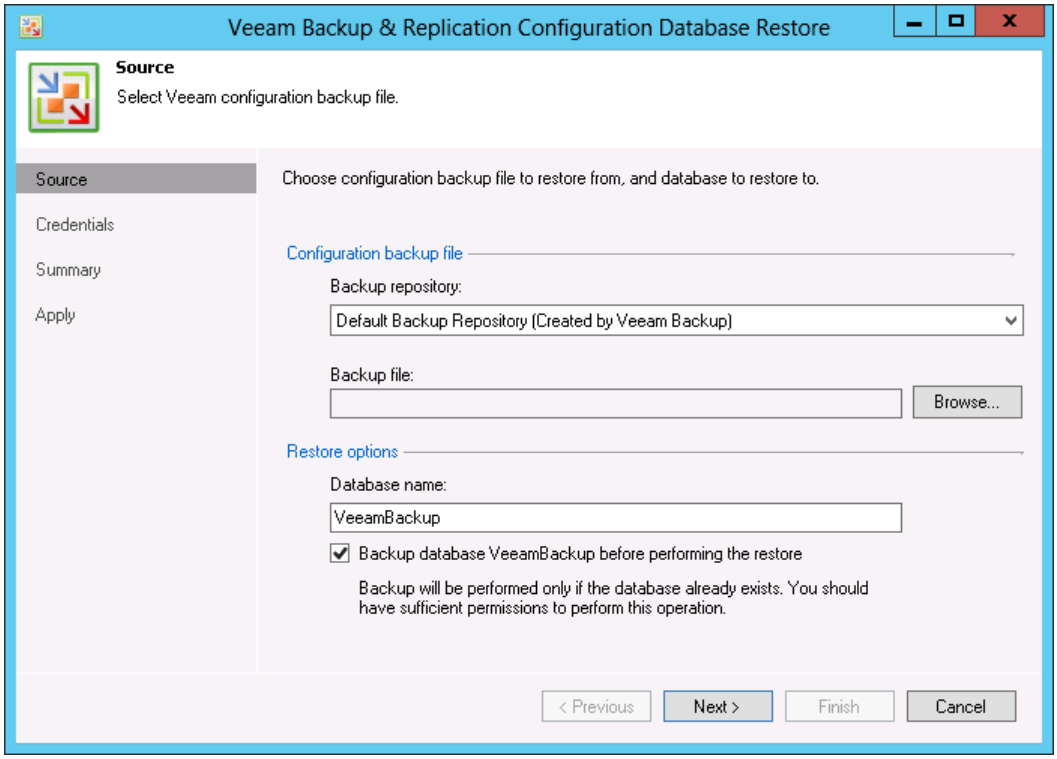

Step 2. Select the Backup File and SQL Database

At the first step of the wizard, select the configuration backup file and the target SQL database:

1. From the **Backup repository** list, choose the repository with the configuration backup file (.bco).
2. Click **Browse** next to the **Backup file** field and choose the necessary .bco file. By default, configuration backups are stored to the `C:\backup\VeeamConfigBackup\%BackupServer%` folder on the Veeam backup server.
3. In the **Database name** field, specify the name of the database into which data from the configuration backup should be imported. You can import data to the existing database or to a new database.
4. If you decide to write configuration data to an existing Veeam Backup & Replication SQL database, select the **Backup database <DatabaseName> before restore** check box. This option will help you protect the current database from accidental errors during the restore process. With this option selected, Veeam Backup & Replication will first back up the current database using the SQL native tools. After that, it will purge the current database and import data from the configuration backup to it. In such scenario, if an error occurs during the restore process, you will be able to restore the current database from the SQL backup using SQL Management Studio or SQL scripts.
   The SQL database backup is named by the following pattern: `VeeamBackup<DatabaseName><date>.bak` and stored to the default SQL backups location, for example: `%Program Files%\Microsoft SQL Server\MSSQL11.MSSQLSERVER\MSSQL\Backup\.`

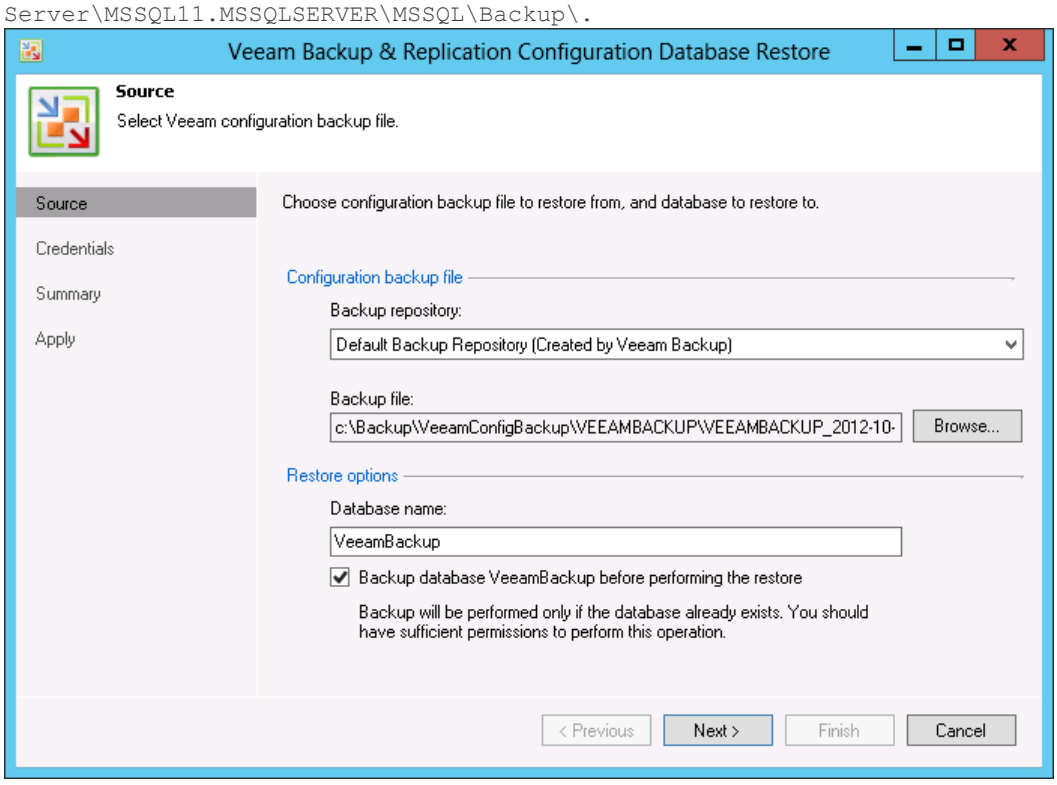

Note The backup of the current database will be created only if the user account under which you perform restore has sufficient permissions to create a database on the SQL server.

Step 3. Set Credentials for Restored Objects

At the **Credentials** step of the wizard, specify a password for every object in the list. When Veeam Backup & Replication exports configuration data, it does not save passwords for objects. Therefore, you need to specify them anew.

If you do not provide passwords at this step, you will need to manually define them in the properties of every object after the restore process is complete.

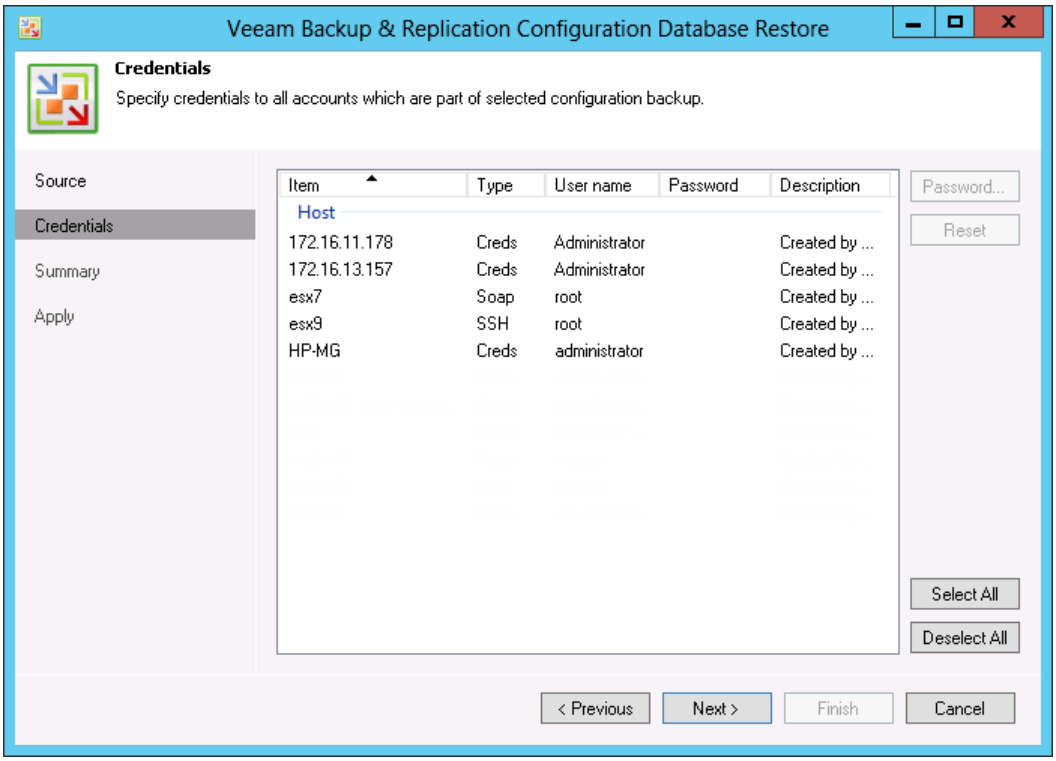

Important! It is strongly recommended that you specify passwords for hosts and repositories referenced by backup and replication jobs. In the opposite case, Veeam Backup & Replication will fail to rescan these hosts and repositories. As a result, Veeam Backup & Replication will not display restore points for such backups and replicas. To overcome this situation, you can do either of the following:

- Perform the configuration restore once again and specify passwords for corresponding hosts and repositories.
- After restore, map backups to existing jobs (note that such backups will be displayed under the **Backups** > **Imported** node) and perform rescan for replicas. In the **Backup & Replication** view, select the **Replicas** node in the inventory pane, right-click the necessary replica in the working area and select **Rescan replicas**.

Step 4. Start Configuration Data Restore

At the **Summary** step of the wizard, review the configuration data and select additional options:

1. If you are restoring a configuration backup on the Veeam backup server with the *Resticted* PowerShell execution policy, select the **Enable required PowerShell execution policy on SCVMM** check box. With this option selected, you will be able to run backup and replication jobs on the target Veeam backup server immediately after the restore process and will not need to change the PowerShell execution policy manually. Essentially, enabling this option is identical to running the '*Set-ExecutionPolicy RemoteSigned*' command on the Veeam backup server.

2. If necessary, select the **Launch the application automatically** check box. If the configuration restore process completes successfully, Veeam Backup & Replication using the new configuration data will be automatically started.
3. Click **Restore** to start the configuration restore process.

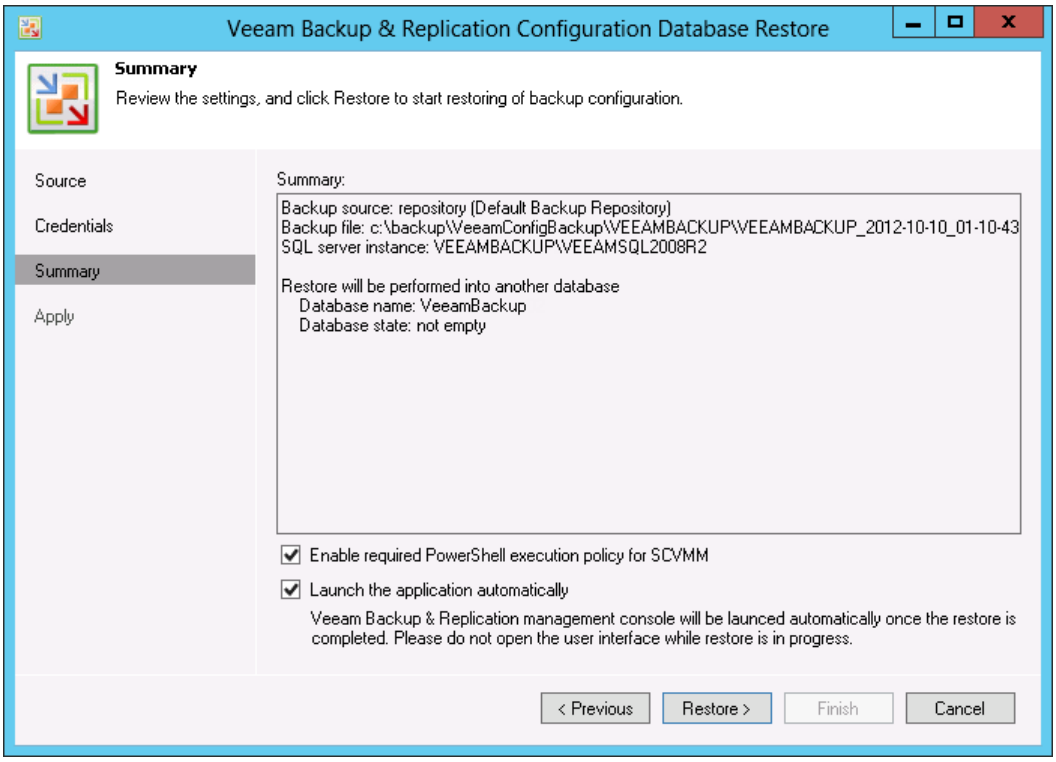

Working with Veeam Explorer for Exchange

This section describes administrative tasks you can perform with Veeam Explorer for Exchange. Veeam Explorer for Exchange is a free tool that you can use to browse Exchange mailbox stores inside Veeam backups. It features a familiar, easy-to-use interface and allows you to quickly locate the mailboxes or items you need.

Overview

Veeam Explorer for Exchange is a free tool available to users of Veeam Backup & Replication. It allows you to browse Microsoft Exchange database files and restore necessary items, such as mailboxes, folders, messages, tasks, contacts and so on. Instead of fully restoring and starting the virtual machine (VM) with the Microsoft Exchange Server, you can use Veeam Backup & Replication capabilities to extract the necessary Microsoft Exchange database from the backup file and then use Veeam Explorer for Exchange to browse and restore items.

You can use granular browsing and searching capabilities to find any item or a bunch of items stored in any number of Microsoft Exchange database files. Restore options include:

- Exporting mailbox folders and items as Personal Folder Files (.pst)
- Saving mailbox items as Microsoft Exchange Mail Documents (.msg)
- Sending mailbox items as attachments via email
- Restoring mailbox folders and items (available only with Veeam Backup & Replication Enterprise and Enterprise Plus Editions)

Note Veeam Explorer for Exchange requires full access to Microsoft Exchange database files for item recovery. This level of access is usually granted to a very limited number of employees within the organization. If you would like to allow less privileged users to perform recovery of Microsoft Exchange items from backups, you can use the Application-Item Recovery (AIR) wizard for Microsoft Exchange. For more information, refer to the Veeam Application-Item Recovery Wizards User Guide on the product Resources web page.

System Requirements

Veeam Explorer for Exchange is a free tool and does not require a separate license. It is installed with Veeam Backup & Replication version 7.0 (all editions including Veeam Backup Free Edition are supported).

The table below contains the list of system requirements necessary to use Veeam Explorer for Exchange.

Specification	Requirement
OS	Only 64-bit versions of the following operating systems are supported: • Microsoft Windows Server 2008 SP2 • Microsoft Windows Server 2008 R2 SP1 • Microsoft Windows 7 SP1 • Microsoft Windows 8 • Microsoft Windows Server 2012
Microsoft Exchange	Veeam Explorer for Exchange only supports database files (.edb) created with the 64-bit version of Microsoft Exchange 2010 and Microsoft Exchange 2013. To open database files, Veeam Explorer for Exchange requires a service dynamic link library (**ese.dll**) which is installed together with Microsoft Exchange.
Software	The feature for restoring folders and items into their original location is available only to users of Veeam Backup & Replication Enterprise and Enterprise Plus Editions. If you are planning to export folders and items as Personal Folder Files (.pst), it is necessary to have a 64-bit version of Microsoft Outlook 2010 or Microsoft Outlook 2013 installed on the system. The following software is required (included in the setup): • Microsoft .NET Framework 4.0 or later • Microsoft Visual C++ Runtime Library

Required Permissions

Veeam Explorer for Exchange requires full access to Microsoft Exchange database and its log files for item recovery. That is, you need both read and write permissions to all files in the folder with the database.

To restore folder(s)/item(s) to Exchange server, the account used for connection to that server will need sufficient access rights. They can be granted using Exchange Management PowerShell cmdlets.

Typically, access rights are provided through impersonation, as described in Configuring Exchange Impersonation. If you need to provide granular permissions, consider the examples below. To be able to restore folder(s)/item(s) to Public Folder or Mailbox, the account can be assigned an appropriate role on target Exchange server by running the following cmdlet:

```
Add-RoleGroupMember "Organization Management" –Member
"<user_account>"
```

To be able to restore items to a Mailbox, the account can be granted Full access rights for that Mailbox by running the following cmdlet:

```
Add-MailboxPermission –Identity "<mailbox>" -User
"<user_account>" -AccessRights FullAccess –InheritanceType All
```

Administration

To start Veeam Explorer for Exchange, you can:

- Pass through the Microsoft Exchange Item Level Restore wizard
- Restore the .edb file manually and open it from the Veeam Backup browser
- Select **All Programs** > **Veeam** > **Veeam Explorer for Exchange** from the Windows **Start** menu

The basic procedure of work with Veeam Explorer for Exchange involves the following steps:

1. Perform initial configuration of Veeam Explorer for Exchange.
2. Restoring the database file (.edb) from the backup or replica
3. Add one or several database files to Veeam Explorer for Exchange
4. Find necessary items
5. Restore items

Configuring Veeam Explorer for Exchange

When you launch Veeam Explorer for Exchange for the first time, you need to perform initial configuration.

- Specify the location of the `ese.dll` file used by Microsoft Exchange to create the database file
- Configure email settings

To perform these configuration tasks, open the main menu (top left corner of the **Veeam Explorer for Exchange** main window) and select **Options**.

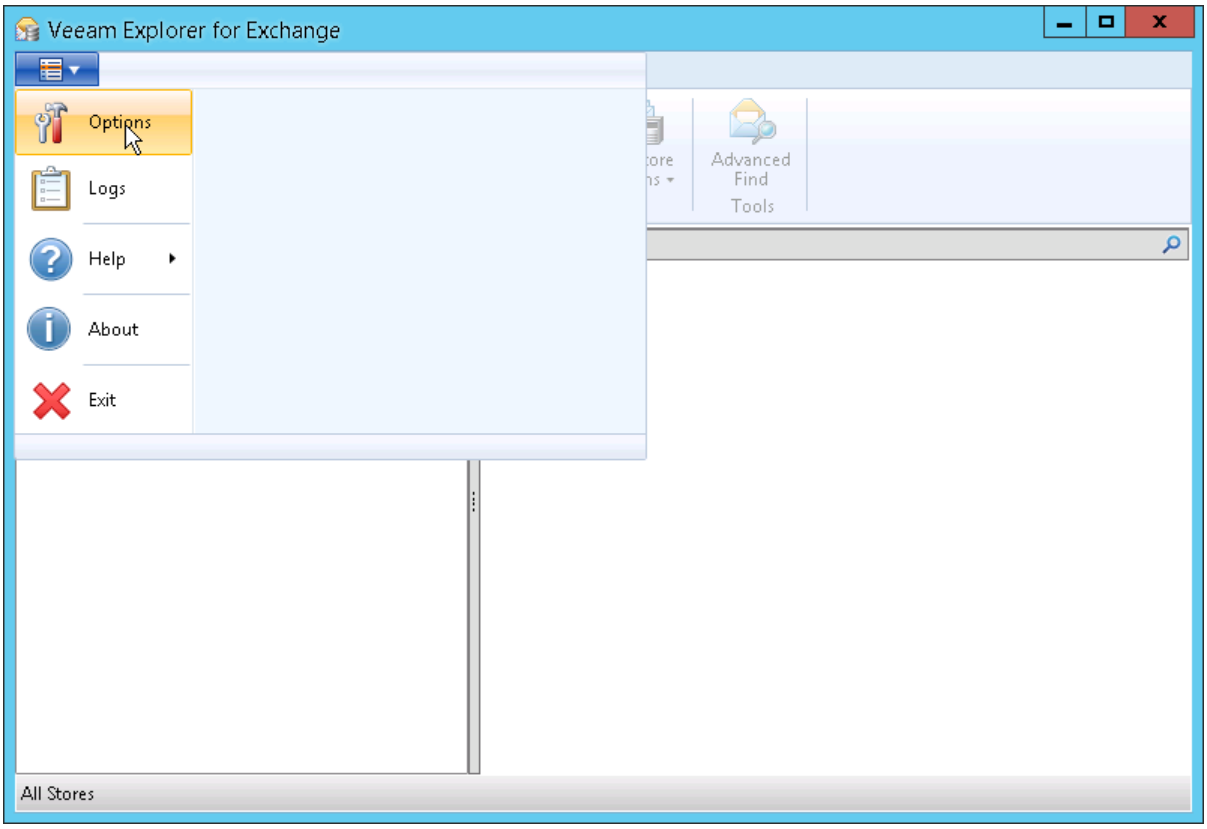

SMTP Settings

To send items recovered from an Exchange database as email attachments, it is necessary to provide mail server information.

To configure the email settings:

1. Open the **SMTP Settings** tab in the **Options** window.
2. Select the **Use SMTP settings** check box.
3. In the **SMTP Server** field, specify the DNS name or IP address of the mail server to be used for sending emails. If necessary, you can change the port number. By default, port number 25 is used.
4. In the **From** field, specify the email address, usually that of the administrator responsible for Microsoft Exchange item recovery, from which emails with attached items should be sent. This email will be used by default when sending restored items, though you can specify a different email address every time.
5. If necessary, select the **Use authentication** check box and provide credentials for the account that will be used to send emails with attached items.
6. If security is an issue, select the **Enable SSL security** check box to use SSL encryption for transferred data.
7. Click **Test Mail** to send a test email message to the account under which Veeam Explorer for Exchange was launched.

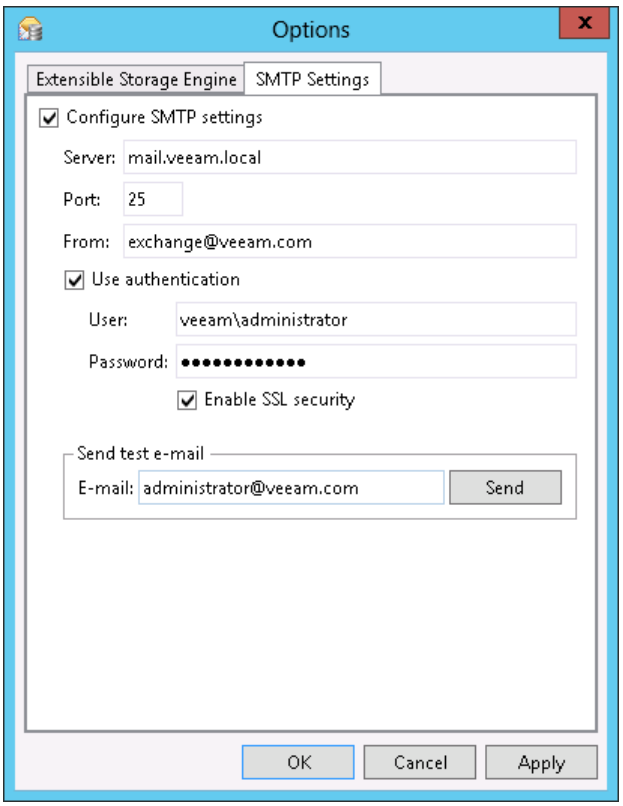

Extensible Storage Engine

To work with database files, Veeam Explorer for Exchange requires a special dynamic link library — `ese.dll`, supplied with Microsoft Exchange. The `ese.dll` file should be of the same version as Microsoft Exchange that was used to create database files. Currently, Veeam Explorer for Exchange supports Microsoft Exchange 2010 SP1, SP2 and SP3 and Microsoft Exchange 2013.

Note When you run the Exchange Items Restore wizard on a VM backed up using VSS, the `ese.dll` file is added automatically. For more information, see Using the Restore Exchange Items Wizard.

To specify the path to the `ese.dll` file:

1. Open the **Extensible Storage Engine** tab in the **Options** window.
2. Click **Browse** and specify the path to the `ese.dll` file. The file can be found on the Microsoft Exchange Server distribution CD at `X:\Setup\ServerRoles\Common\ese.dll`, or in the installation directory of Microsoft Exchange Server:
    - for Exchange 2010 default path is `%ProgramFiles%\Microsoft\Exchange Server\V14\Bin`
    - for Exchange 2013 default path is `%ProgramFiles%\Microsoft\Exchange Server\V15\Bin`

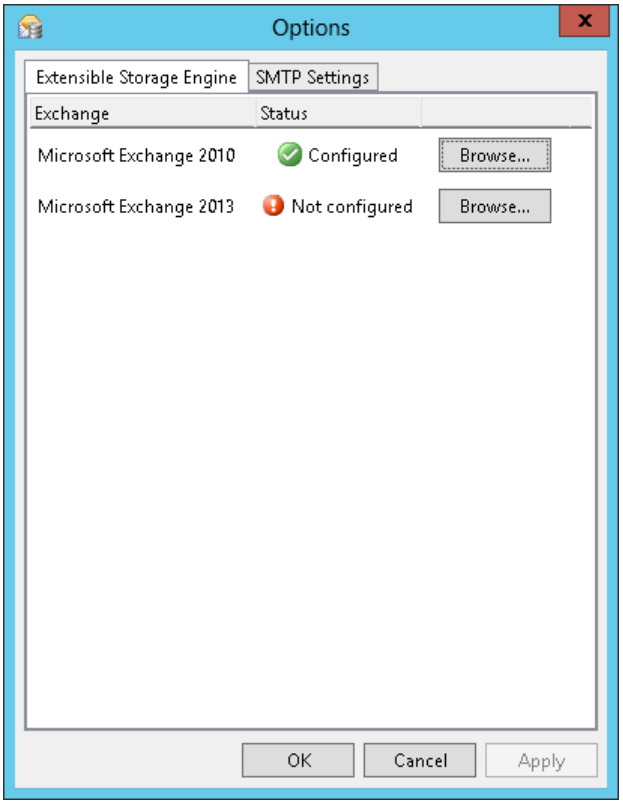

Restoring the Database File (.Edb) from the Backup

Before you can start working with Veeam Explorer for Exchange, you need to extract from the backup the Exchange database — `.EDB`. You can do it in two ways:

- You can use the Microsoft Exchange Item Level Restore wizard. In this case, Veeam Backup & Replication will automatically extract the Exchange database from the backup and open it in Veeam Explorer for Exchange.
- You can manually recover the database from the backup, locate the restored Exchange database and open it in Veeam Explorer for Exchange.

Using the Microsoft Exchange Item Level Restore Wizard

To start the **Microsoft Exchange Item Level Restore** wizard, do one of the following:

- In Veeam Backup & Replication, open the **Backup & Replication** view. Click **Restore > VMware** on the toolbar. In the **Restore Options** window, select **Application items** and click **Next**. Select **Microsoft Exchange** and click **Next**. Note that you can use this option for VMware VMs only.
- In Veeam Backup & Replication, open the **Backup & Replication** view. Select the **Backups** node, expand the backup job with the Exchange server in the working area. Select the Exchange VM and click **Application items > Microsoft Exchange** on the toolbar. In this case, you will pass to the Restore Point step.
- In Veeam Backup & Replication, open the **Backup & Replication** view. Select the **Backups** node, expand the backup job with the Exchange server in the working area. Right-click the Microsoft Exchange VM and select **Restore Exchange items**. In this case, you will pass to the Restore Point step. Note that you can use this method only for Exchange backups that were created with VSS-aware image processing enabled.

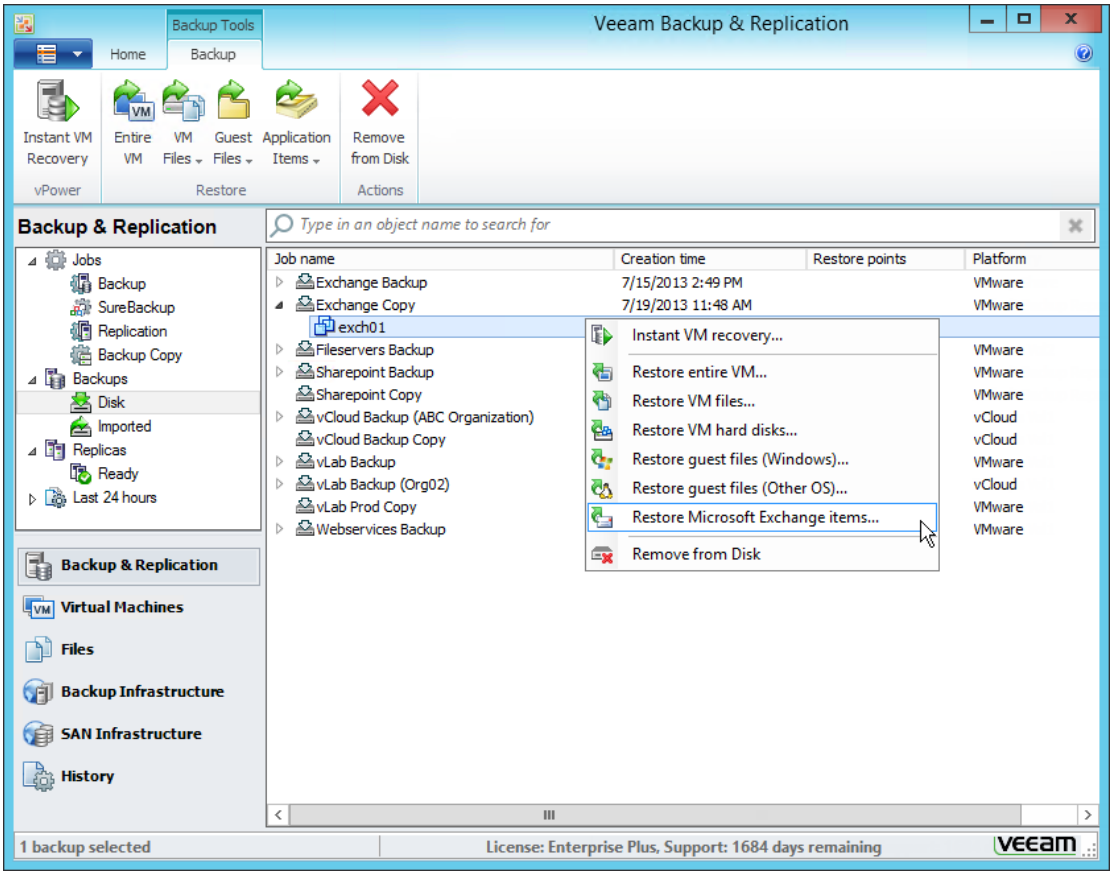

Step 1. Select the Backup with Microsoft Exchange Server

In the list of available jobs, select the necessary virtual machine. To quickly find VMs in jobs, use the search field at the bottom of the window.

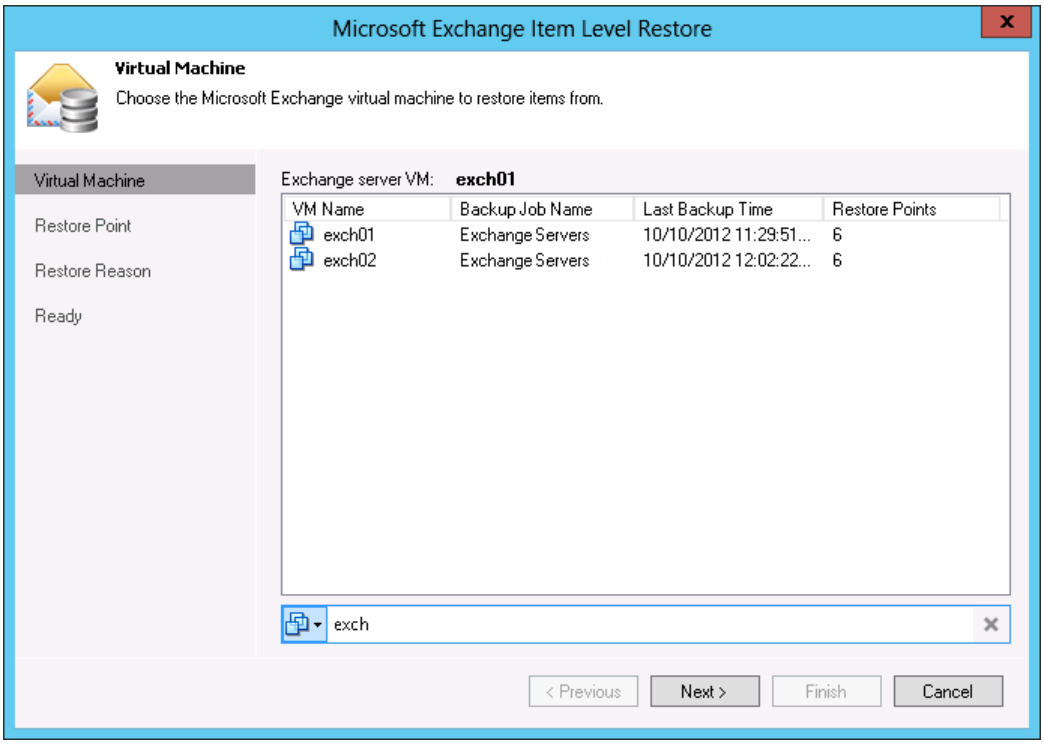

Step 2. Select the Restore Point

Select the necessary restore point for the virtual machine.

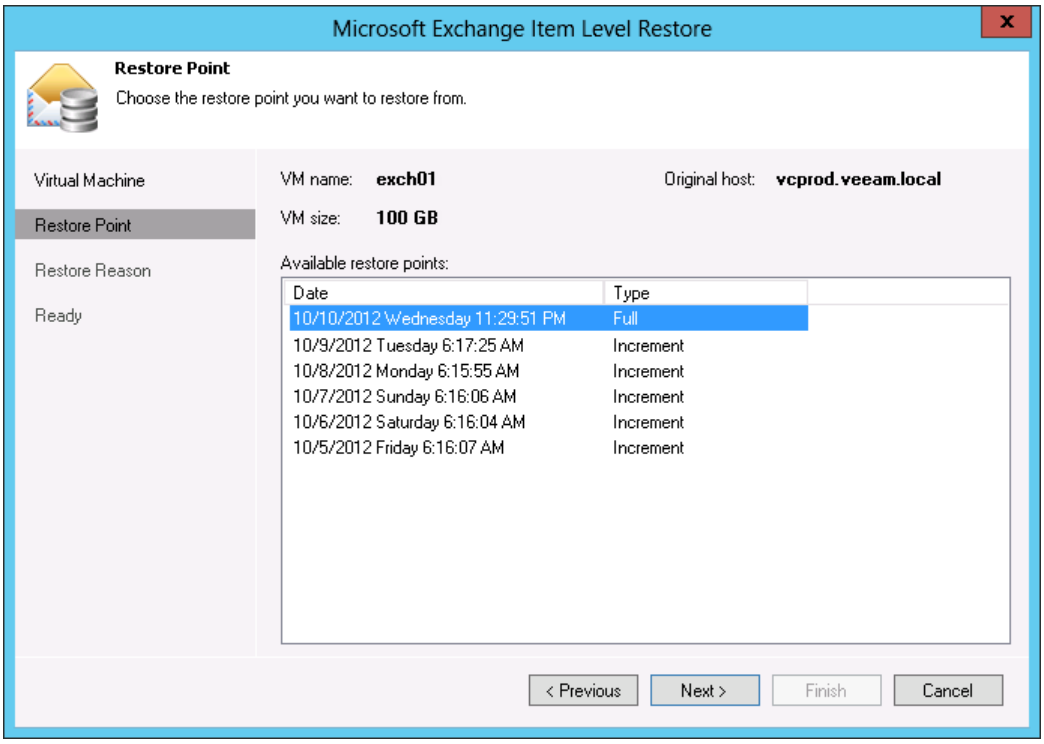

Step 3. Enter a Restore Reason

If necessary, enter the reason for performing VM guest file restore and click **Next**. The information you provide will be saved in the session history so that you can reference it later.

If you do not want Veeam Backup & Replication to display the **Reason** step next time, select the **Do not show me this page again** check box at the bottom of the wizard.

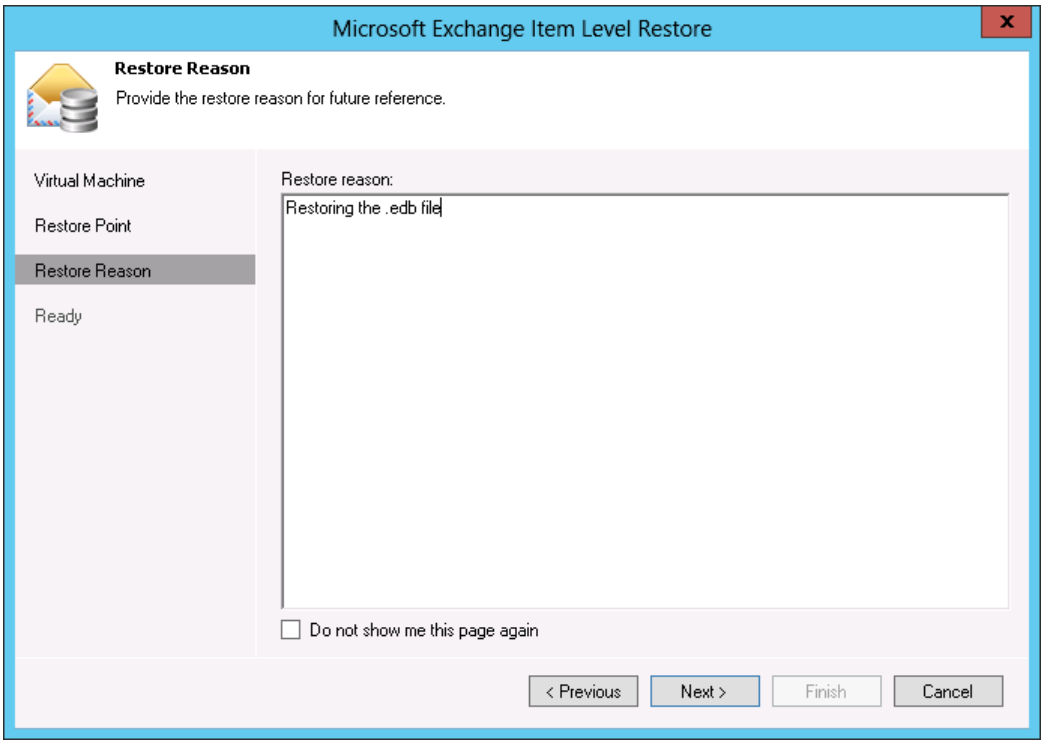

Step 4. Restore the EDB File

On the last step, read the summary and click **Finish** to start the restore process.

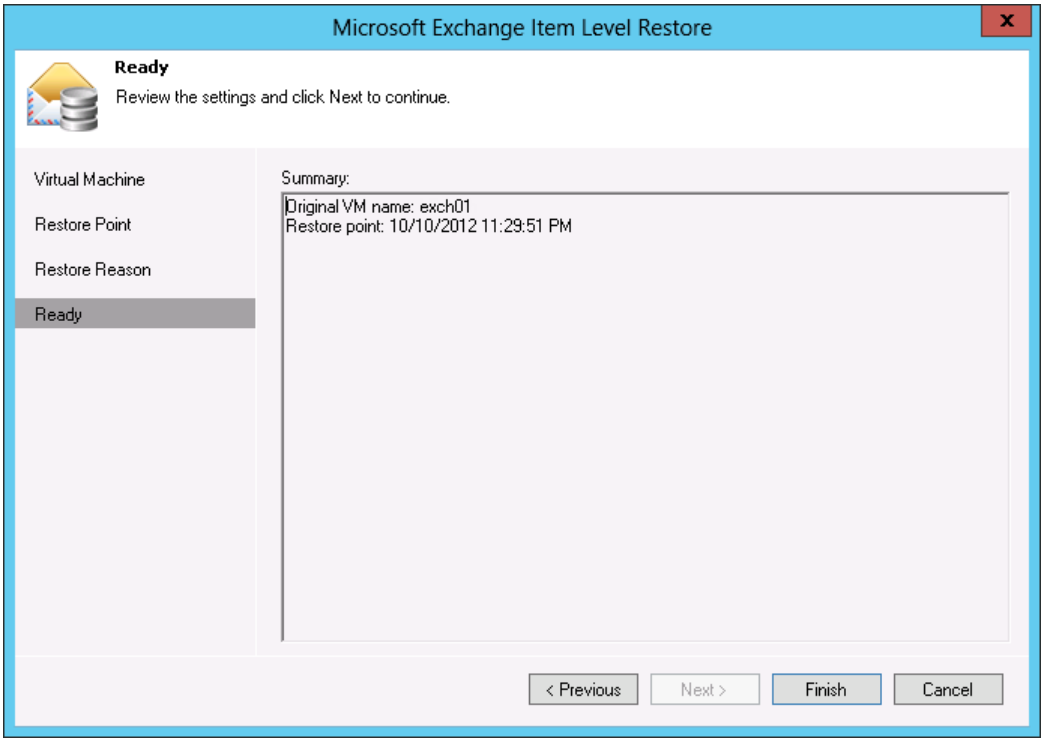

The restore process will depend on whether the backup job had application-aware image processing (using VSS) enabled or not:

- If the Exchange Server VM was backed up using VSS, the wizard will automatically restore the database file from the Exchange Server via File-Level Restore (FLR), launch Veeam Explorer for Exchange and add the restored database to it. You will then only have to manually find and restore necessary Exchange items.
- If the Exchange Server VM was backed up without using VSS, the wizard will run FLR, mount the file system of the VM and open the backup browser window and Veeam Explorer for Exchange. It may be necessary to manually configure Veeam Explorer for Exchange (that is, specify the path to the `ese.dll` file), then find and add the database file. After that you can find and restore necessary Exchange items.

Restoring Database Files Manually

To restore the database from the VM backup, you can use any data recovery feature available in Veeam Backup & Replication, such as Instant VM Recovery, full VM restore, restore of guest OS files from a replica and so on. However, the most convenient method is to use the guest OS file-level recovery option.

To restore a Microsoft Exchange database manually:

1. Perform guest OS files restore for the virtualized Microsoft Exchange server.
2. In the Veeam Backup browser, double-click the .edb file or click **Exchange Items** on the toolbar.
3. Veeam Backup & Replication will open the selected database in Veeam Explorer for Exchange. After that, you can browse the database and restore the items you need. To learn more, see Browsing, Searching and Viewing Items.

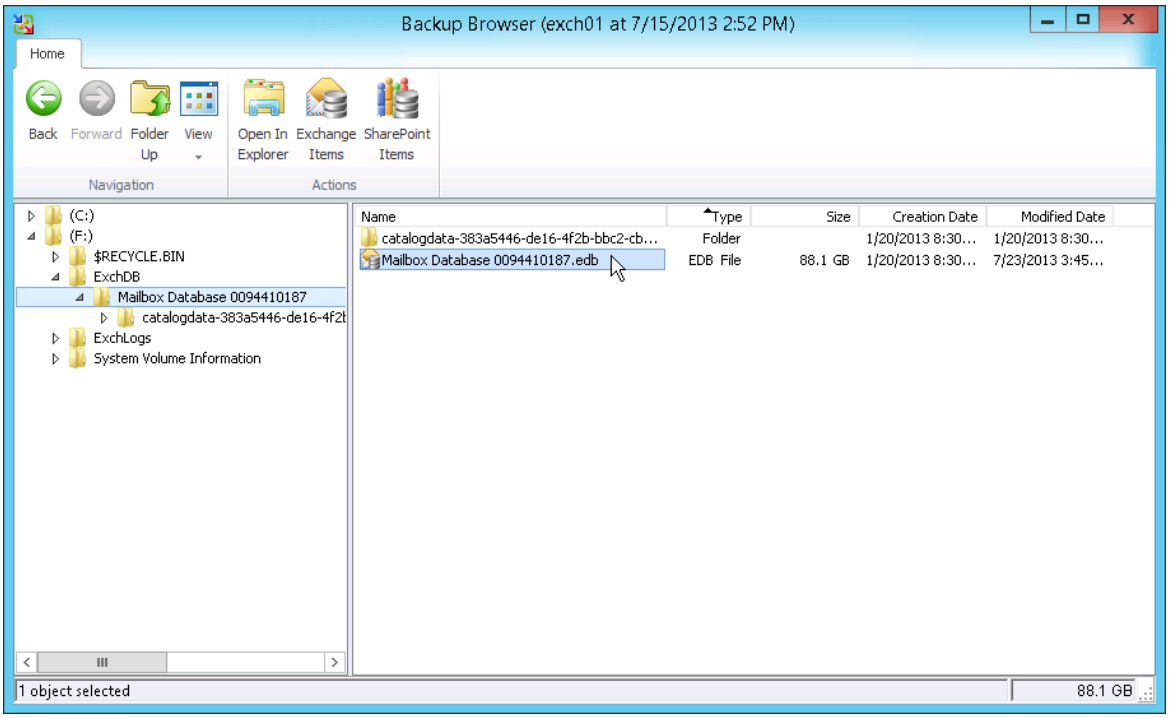

Adding and Removing Database Files

After you restore the database file(s) from backup, there are two common usage scenarios:

- If you know in which database the necessary item or items are located, you will only need to add one database.
- If you are not sure in which database the item or items are located, or they are scattered across a number of databases, Veeam Explorer for Exchange allows you to add and work with several databases at the same time.

Adding Database

You can only add database files created with Microsoft Exchange Server 2010 or Microsoft Exchange Server 2013. Also, it is necessary to make sure that Veeam Explorer for Exchange has access to `Ese.dll` supplied with the same version of Microsoft Exchange. For details, see Configuring Veeam Explorer for Exchange.

To add a database file to the Veeam Explorer for Exchange console:

1. Click **Add Store** on the toolbar, or right-click **All Stores** in the navigation pane and select **Add Store**.
2. Click **Browse** and specify a path to the Microsoft Exchange database file (.edb).
3. Click **Browse** below and specify a path to the Exchange logs folder.

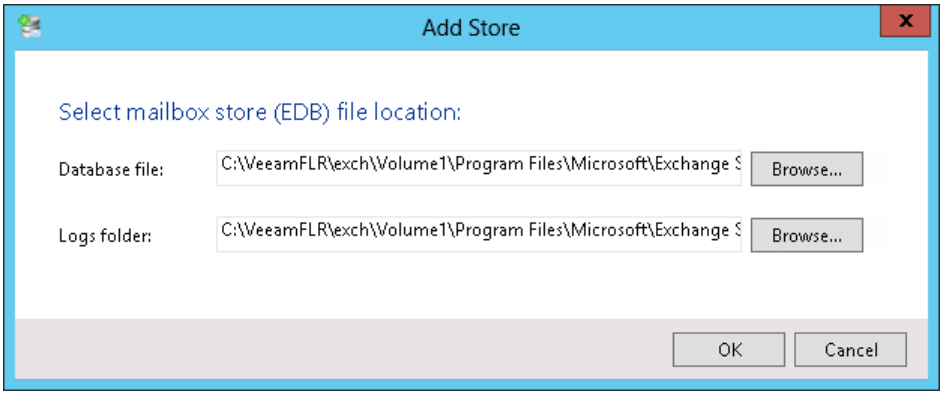

Note If you are using file-level restore to mount the contents of the backup file with the database file to the Veeam Backup server, VM disks are mounted under the `C:\veeamflr\<vmname>\<volume n>` folder.

Alternatively, you can double-click the necessary Exchange database file to automatically start Veeam Explorer for Exchange and add the database to the console. If you mounted the file system of the VM with the database file via FLR, simply double-click the file within the Backup Browser or use Windows Explorer to browse to the necessary database file and then double-click it.

If the database is in "dirty state", Veeam Explorer for Exchange will display a warning. In this case, you will have to recover the database before adding it.

To recover the database:

1. Check the **Logs folder** field and make sure you have specified the correct folder storing the Exchange database logs.
2. Click **Recover** to recover the damaged database.

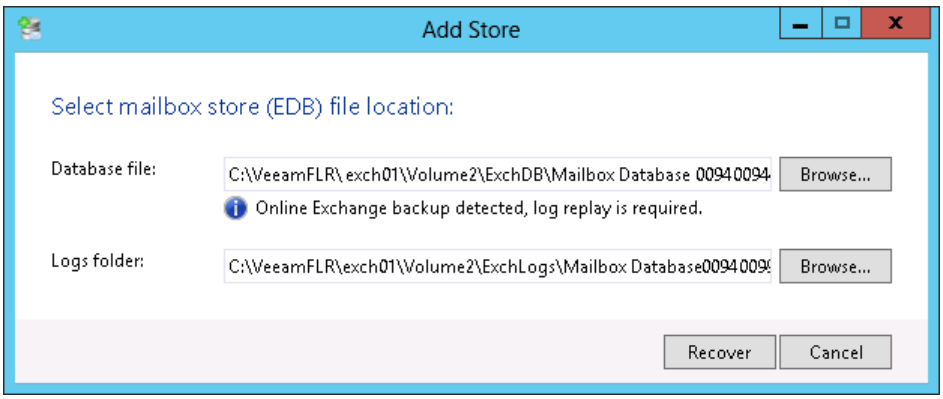

Important! You need to have write permissions for the database to be able to apply replay logs to the database.

Removing Database

To remove a database from the Veeam Explorer for Exchange console (that is, to unlink it), select the database in the navigation pane and click **Remove Store** or right-click the database name and select **Remove Store**.

Browsing, Searching and Viewing Items

After you add one or several database files to the Veeam Explorer for Exchange console, you can browse and search through the contents of the database(s) to find necessary items.

Browsing

In the navigation pane, you can see all the containers (added databases with mailboxes and folders). When you select a folder, you can browse its contents in the main working pane (messages in the **Inbox** folder, tasks in the **Tasks** folder and so on).

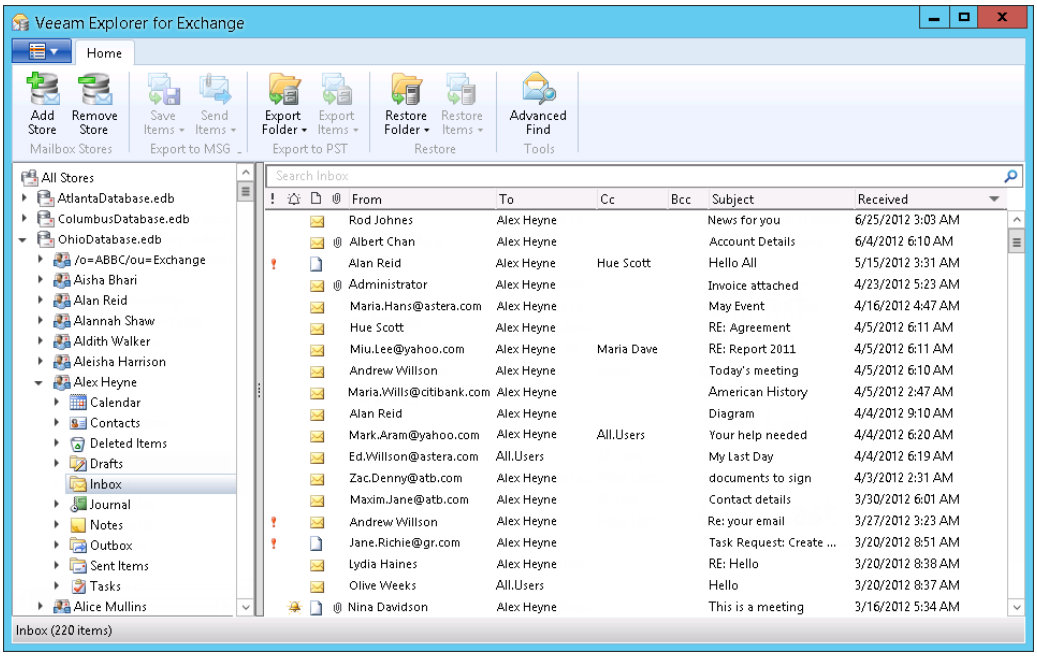

Searching

At the top of the main working pane, there is a search field that allows you to find items in the selected container that match a specified search term. For example, you can select a user's mailbox and search for *veeam software* to find all the items that contain the words starting with "veeam" and "software" in the body text, address field, subject field or name of attachments. Search results for databases and mailboxes will include all items (that is, email messages, tasks, contacts and other items).

Note When you search a phrase, the search will by default locate items that contain words in the phrase, not the phrase itself. To find the exact phrase, use quotes, for example, “veeam software”.

You can select the whole database or several databases to get more results, or select a specific folder (for example, **Inbox**) inside a specific database and mailbox to narrow a search.

You can further narrow your search results by specifying various search criteria inside the box in the format "*criteria:value*" (for example, *from:John*, *hasattachments:yes*, *messagesize:<10 KB*, *received:yesterday* and so on). You can also use logical operators such as *AND*, *OR* and *NOT* (must be typed in uppercase letters) and wildcard characters such as * and *?*. These search criteria are similar to those used for searching in Microsoft Outlook 2010. For more information, see this Outlook help topic.

Alternatively, you can click **Advanced Find** on the toolbar to make up a list of necessary search criteria using predefined dropdown menus.

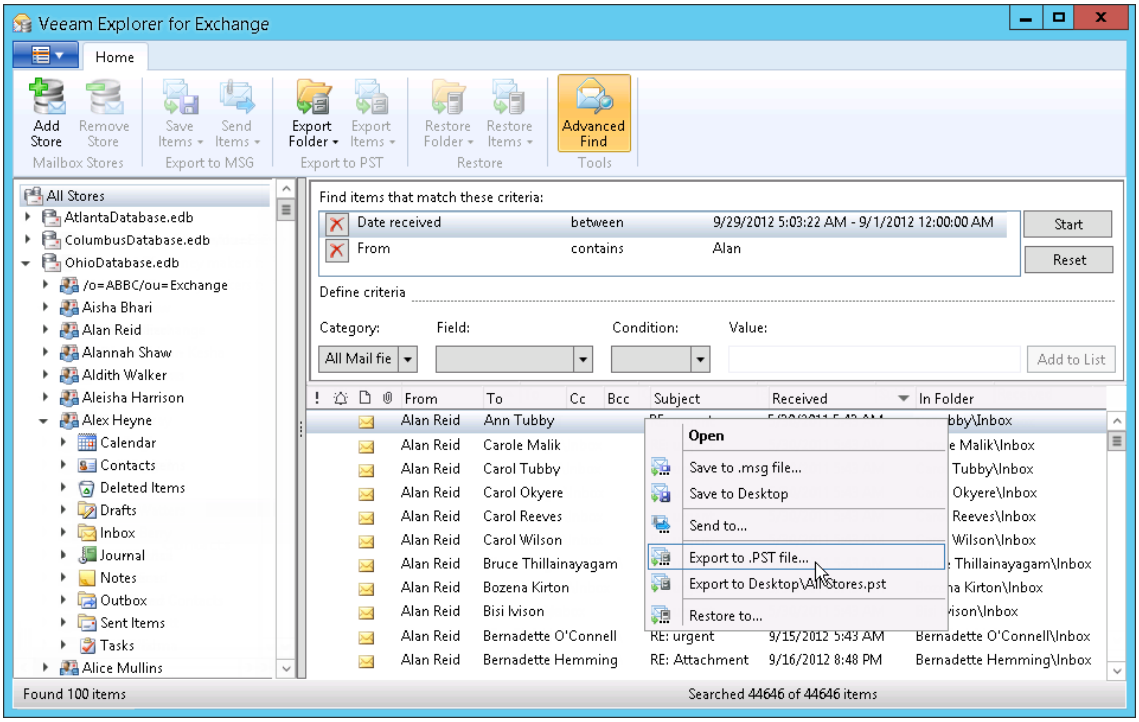

Viewing Items

Veeam Explorer for Exchange allows you to quickly view any item by double-clicking it. This opens the item in a separate window, showing all the necessary details (the address fields, the subject field, the body text and so on).

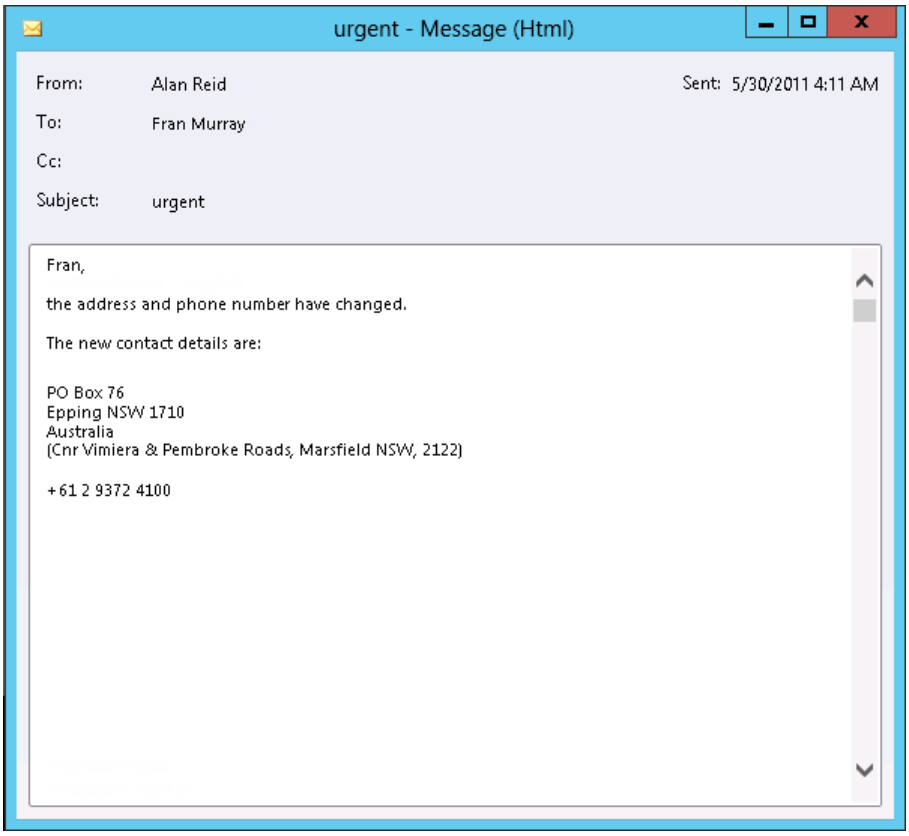

Restore Scenarios

Veeam Explorer for Exchange provides the following restore options:

- Exporting folders and items as Personal Folder Files (.pst)
- Saving items as Microsoft Exchange Mail Documents (.msg)
- Sending items as email attachments
- Restoring folders and items (available only with Veeam Backup & Replication Enterprise and Enterprise Plus Editions)

Exporting Folders and Items

If you have a 64-bit version of Microsoft Outlook 2010 or Microsoft Outlook 2013 installed on the computer running Veeam Explorer for Exchange, you can export folders (or mailboxes) and items as Personal Folder Files (.pst). Otherwise, you can use other restore options (see Saving Items and Sending Items).

Important! Windows Search can conflict with exporting of items and folders. If the .pst file to which you are exporting is indexed at the same time, the exporting process may hang up. To avoid this, it is necessary to exclude .pst files from the indexing scope. For this, use the search box in the Windows **Start** menu to open the **Indexing Options** window, click **Modify** to open the **Indexed Locations** window and clear the **Microsoft Outlook** check box.
Alternatively, you can disable Windows Search completely in the **Windows Features** dialog box.
Also note that if you are exporting .pst files to a shared folder, it is necessary to exclude Outlook files or disable Windows Search on the computer where the shared folder is located.

To export a folder as a single .pst file:

1. Select the folder in the navigation pane, click **Export Folder** on the toolbar and select **Export to .PST file**. You can also right-click the folder in the navigation pane and select **Export to .PST file**.
2. Specify the name and location for the file and click **Save**.

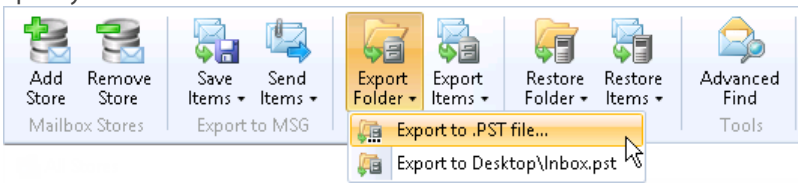

To export several items as one .pst file:

1. Select the items in the main working area (use **SHIFT** and **CTRL** to select multiple items), click **Export Items** on the toolbar and select **Export to .PST file**. You can also right-click the items in the working area and select **Export to .PST file**.
2. Specify the name and location for the file and click **Save**.

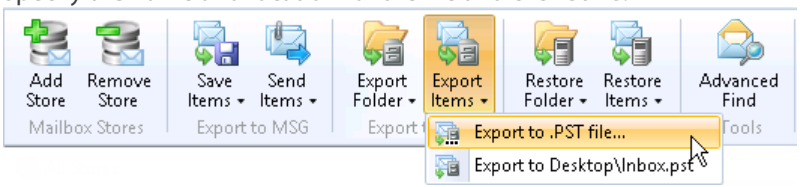

Note For convenience, Veeam Explorer for Exchange allows you to quickly export .pst files directly to the Desktop with a default name. For example, you can select the **Tasks** folder and then click **Export Folder** > **Export to 'Desktop\Tasks.pst'** on the toolbar.

Saving Items

Veeam Explorer for Exchange allows you to save any item as a Microsoft Exchange Mail Document (.msg file) to a specific location or directly to the user's desktop.

To save an item:

1. Browse to the necessary folder or use the search field to see a list of items in the main working area.
2. Select the required item (or use **SHIFT** and **CTRL** keys to select multiple items) in the main working area, click **Save Items** on the toolbar and select **Save to .msg file**. You can also right-click the item(s) in the working area and select **Save to .msg file**.
3. Specify the location for the file or files and click **OK**.

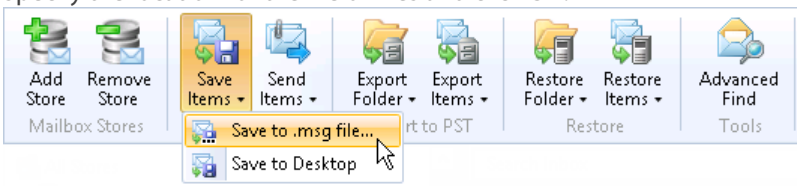

Note For convenience, Veeam Explorer for Exchange allows you to quickly save .msg files directly to the Desktop with a default name. For example, you can select an item and then on the toolbar click **Save Items** > **Save to Desktop**. By default, the name of the file will be the same as the subject field of the item.

Sending Items

The most common scenario involves finding an email message and delivering it to the owner. Veeam Explorer for Exchange allows you to send items as attachments to emails. To be able to send items as attachments, it is necessary to specify email settings in the **Configure Options** window. For details, see Configuring Veeam Explorer for Exchange.

To send an item as an email attachment:

1. Select the item (or use **SHIFT** and **CTRL** keys to select multiple items) in the main working area, click **Send Items** on the toolbar and select **Send to**. You can also right-click the item(s) in the working area and select **Send to**.
2. Specify the email addresses of the sender and recipient as well as the subject and text of the email message. To edit the body text, click **More details**. Review the attached items and click **Send**.

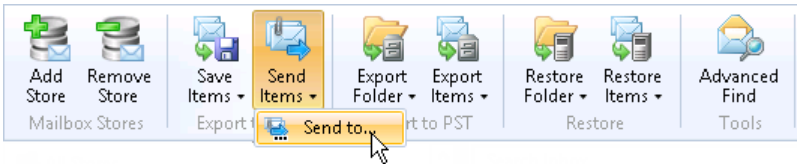

Restoring Folders and Items

Veeam Explorer for Exchange allows you to restore mailbox folders and items into their original location (that is, directly to the production Exchange server). This is a convenient way of restoring lost data without any additional actions.

Important! The feature for restoring folders and items into their original location is available only to users of Veeam Backup & Replication Enterprise and Enterprise Plus Editions.

To restore a folder or mailbox to the original location:

1. Select the folder or item (or use **SHIFT** and **CTRL** keys to select multiple folders or items) in the main working area, click **Restore Folder** or **Restore Items** on the toolbar and select **Restore to**. You can also right-click the item(s) in the working area and select **Restore to**.

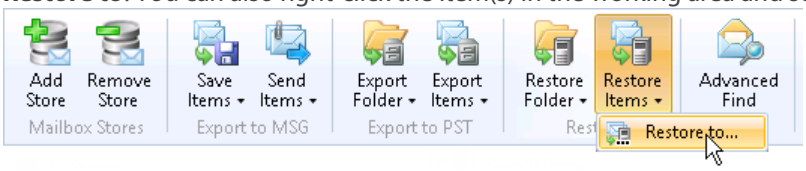

2. Specify the target mailbox and domain account to be used. You can either use the current account or specify a different one.

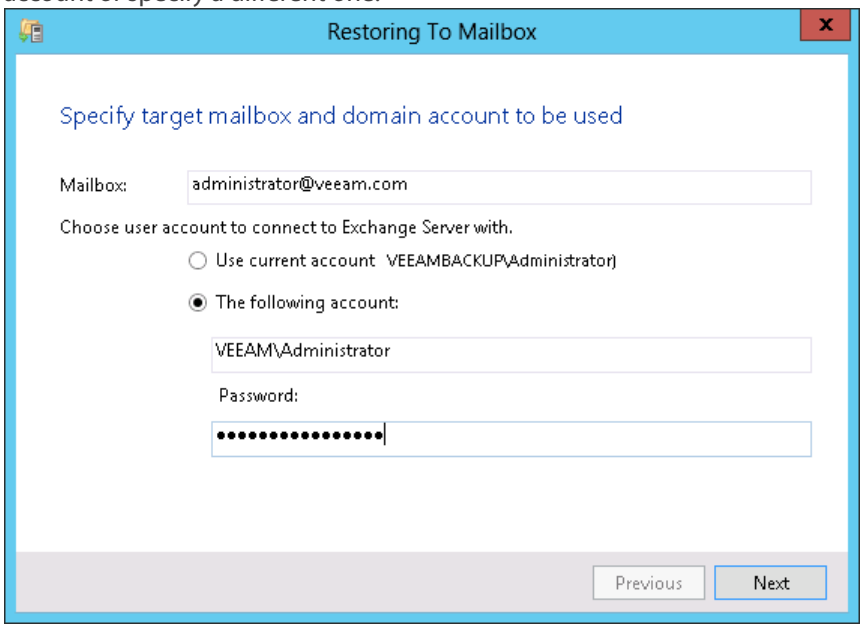

3. Specify the target mailbox server and folder. You can restore to original folder or specify a different one.

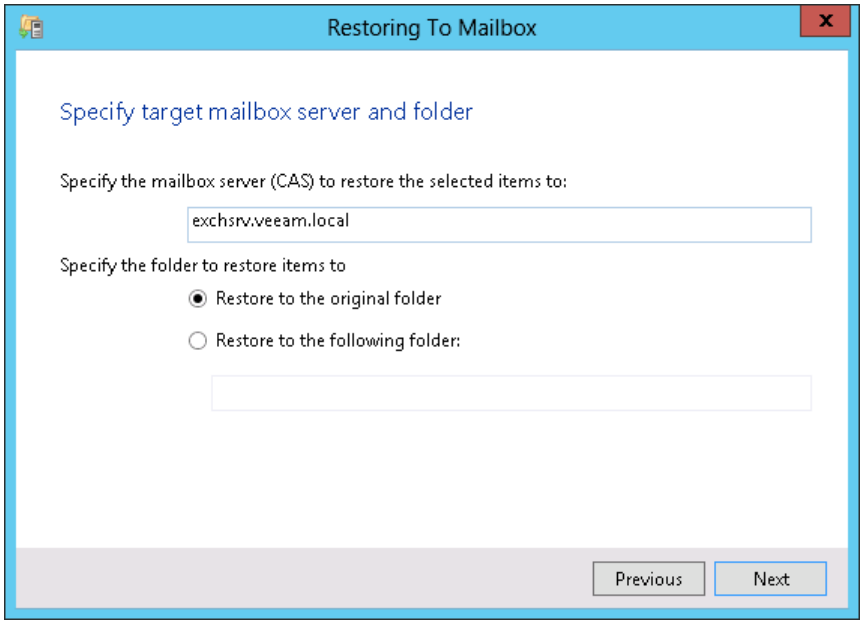

4. Specify additional restore options and click **Restore** to begin the restore process.

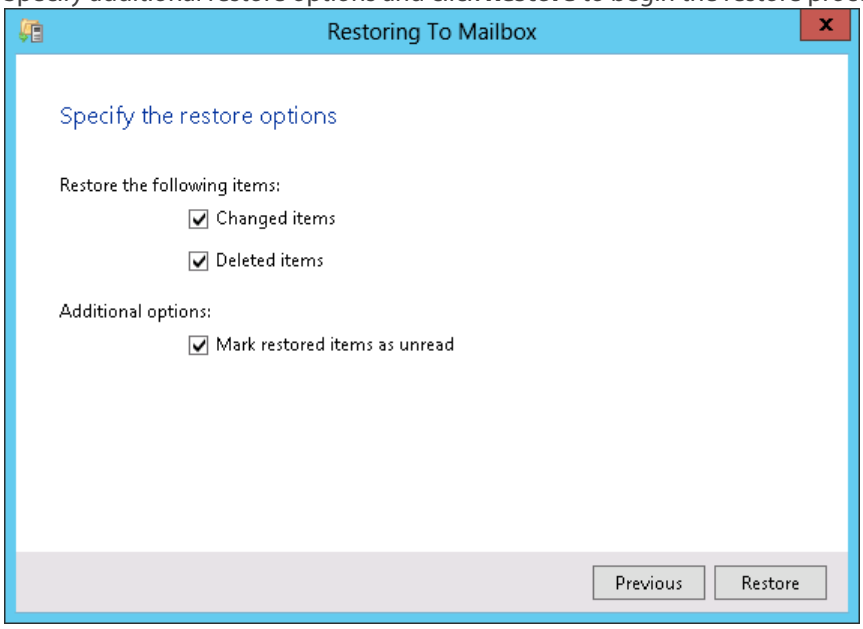

Working with Veeam Explorer for SharePoint

This section describes administrative tasks you can perform with Veeam Explorer for SharePoint - a tool that you can use to browse SharePoint content databases inside Veeam backups. It features a familiar, easy-to-use interface and allows you to quickly locate the documents, items and document libraries you need; several recovery options (saving, e-mailing and others) are provided for Microsoft SharePoint 2010.

Overview

Veeam Explorer for SharePoint is a new tool that extends the functionality of Veeam Backup & Replication, automating the operation of restoring documents from virtualized Microsoft SharePoint Servers.

Read more about Veeam Explorer for SharePoint features and benefits and learn how it works.

About Veeam Explorer for SharePoint

Veeam Explorer for SharePoint allows you to browse Microsoft SharePoint content and recover the necessary items (such as library documents, images, webpages and so on) without a need to fully restore and start the virtual machine hosting SharePoint content database. Instead, you can use Veeam Backup & Replication data recovery options to quickly extract the necessary Microsoft SharePoint content database file (.MDF) from the virtualized server image-level backup, and then use Veeam Explorer for SharePoint to find and restore Microsoft SharePoint documents you need.

Available with all editions of Veeam Backup & Replication, Veeam Explorer for SharePoint provides granular browsing and searching capabilities to find any item or a bunch of items stored in any number of Microsoft SharePoint content databases. You can save items/lists to a local folder or network drive, or send restored items as e-mail attachments. With Veeam Backup & Replication Enterprise and Enterprise Plus edition, you can also use Veeam Explorer for SharePoint to restore Microsoft SharePoint items in the way you need; authorized users (for example, Microsoft SharePoint administrators) can also import content database files (exported by Veeam Explorer for SharePoint) to SharePoint database using PowerShell cmdlets. For more details, refer to Licensing and Editions.

How It Works

Veeam Backup & Replication allows backup administrators to have image-level backup of Microsoft SQL Server VM (hosting Microsoft SharePoint content database) always at hand, and Veeam Explorer for SharePoint makes it possible to find and restore a particular SharePoint item that has been deleted or modified.

The basic procedure of searching and restoring Microsoft SharePoint items with Veeam Explorer for SharePoint and Veeam Backup & Replication then includes the following steps:

1. The backup administrator uses Veeam Backup & Replication restore options to extract SharePoint content database files (.mdf) from the server's backup - through mounting it to Veeam Backup & Replication server.
2. Then Veeam Explorer for SharePoint is used to automatically attach content database to a staging Microsoft SQL Server, so that SharePoint content (libraries, webpages, etc.) become available for browsing, search and other operations.
3. Now Veeam Explorer for SharePoint users can easily locate and restore SharePoint documents they need: save them to selected location, e-mail to the specified recipients, export to XML and restore to production system.
4. After a user finishes working with SharePoint content and closes the Veeam Explorer for SharePoint console, content database will be detached from the staging Microsoft SQL Server.

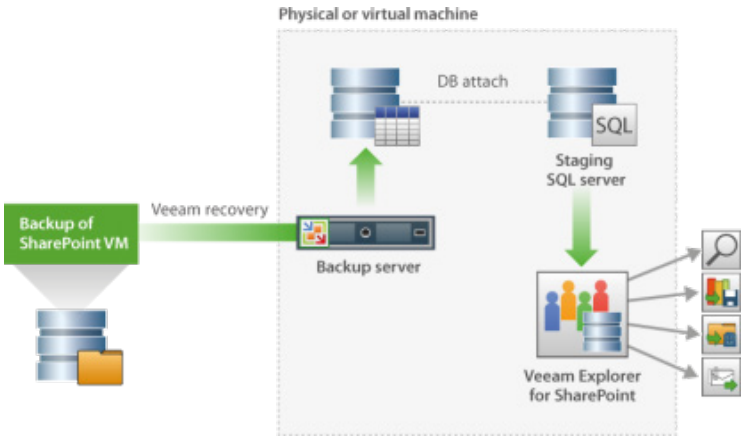

Prerequisites

This section describes system requirements, supported platforms and configurations, required permissions and licensing.

Supported SharePoint Platforms

Veeam Explorer for SharePoint currently provides item recovery possibilities for the following versions and editions of Microsoft SharePoint Server (virtualized either on VMware or Hyper-V platform):

- Microsoft SharePoint Server 2010 Foundation
- Microsoft SharePoint Server 2010 Standard
- Microsoft SharePoint Server 2010 Enterprise

System Requirements

To successfully install and use Veeam Explorer for SharePoint as a component of Veeam Backup & Replication, make sure that all necessary system requirements for Veeam Backup & Replication are met.

Specification	Requirement
OS	Only 64-bit versions of the following operating systems are supported: ▪ Microsoft Windows 7 SP1 ▪ Microsoft Windows Server 2008 SP2 ▪ Microsoft Windows Server 2008 R2 SP1 ▪ Microsoft Windows 8 ▪ Microsoft Windows Server 2012
Software	Veeam Explorer for SharePoint is installed on the machine running Veeam Backup & Replication. The following versions and editions of Veeam Backup & Replication are supported: ▪ Veeam Backup & Replication 7.0 (Standard, Enterprise and Enterprise Plus Edition) ▪ Veeam Backup Free Edition 7.0 ▪ Veeam Backup & Replication 7.0 Cloud Edition For requirements to Veeam Backup & Replication software, see Veeam Backup & Replication documentation.
Microsoft SQL Server	See Staging Microsoft SQL Server.

Staging Microsoft SQL Server

To perform Microsoft SharePoint item recovery, Veeam Explorer for SharePoint requires a Microsoft SQL Server that will be used as a staging system. On this staging system, Veeam Explorer for SharePoint creates temporary Microsoft SharePoint content databases by attaching restored content database files. Below are the requirements for staging server:

1. The staging Microsoft SQL Server must run on the machine where Veeam Explorer for SharePoint is installed (that is, on the machine running Veeam Backup server).
2. The staging system must run the same or a later version of Microsoft SQL Server as the server that hosts restored Microsoft SharePoint content databases. For example, if the Microsoft SharePoint server uses Microsoft SQL Server 2008, then the staging system can run Microsoft SQL Server 2008 or later.

To be able to work with remote BLOB stores (RBS) of the SharePoint content database, also make sure the staging SQL Server configuration meets the following requirements:

1. **FILESTREAM** should be enabled on the database server, and **filestream** settings should be enabled at database level, as described in Microsoft documentation:
    - For SQL Server 2012: http://msdn.microsoft.com/en-us/library/cc645923.aspx
    - For SQL Server 2008R2: http://msdn.microsoft.com/en-us/library/cc645923%28v=sql.105%29.aspx

2. **RBS Client Library** should be installed on the database server; use corresponding Microsoft SQL Server Remote Blob Store setup package (**RBS.msi**) available at Microsoft website:
    - For SQL Server 2012: http://www.microsoft.com/en-us/download/details.aspx?id=35580
    - For SQL Server 2008R2: http://www.microsoft.com/en-us/download/details.aspx?id=16978

By default, the Microsoft SQL Server that hosts the Veeam Backup & Replication configuration database will act as the staging system. If required, you can configure Veeam Explorer for SharePoint to use another Microsoft SQL Server (co-installed with Veeam Backup server).

Note As the staging system, you can use the Microsoft SQL Server Express 2008 R2 SP1 that is shipped with the Veeam Backup & Replication setup. However, consider that content databases that exceed 10 GB cannot be attached to this SQL Server.

Permissions

The account used for working with Veeam Explorer for SharePoint requires membership in the **sysadmin** fixed server role on the staging Microsoft SQL Server.

The account used for connection with target SharePoint server where document item(s)/list will be restored needs the following:

- If permissions of the item being restored are inherited from the parent item (list) - **Full Control** for that list is required.
- If permissions are not inherited, and restored item will replace an existing item - then **Contribute** for the item and **Full Control** for its parent list are required.

Licensing and Editions

Veeam Explorer for SharePoint requires no special license to purchase; Veeam Explorer for SharePoint is shipped with all editions of Veeam Backup & Replication, including Veeam Backup Free Edition. However, the full range of restore capabilities (including recovery to Microsoft SharePoint server in production environment) will be available with Veeam Backup & Replication Enterprise and Enterprise Plus Edition only.

Compare editions:

Feature	Free	Standard	Enterprise	Enterprise Plus
Browse	Available	Available	Available	Available
Search	Available	Available	Available	Available
Save	Available	Available	Available	Available
Send	Available	Available	Available	Available
Export	Available	Available	Available	Available
Restore to original location	Not available	Not available	Available	Available

Administration

This section describes the tasks you should perform to successfully use Veeam Explorer for SharePoint:

Initial Configuration Settings

To start working with Veeam Explorer for SharePoint, select it from the Start menu (for example, for Windows Server 2012) or choose **All Programs** > **Veeam** > Veeam Explorer for SharePoint. When you launch Veeam Explorer for SharePoint for the first time, you need to perform its initial configuration, as described below:

1. Open the main menu and select **Options**.

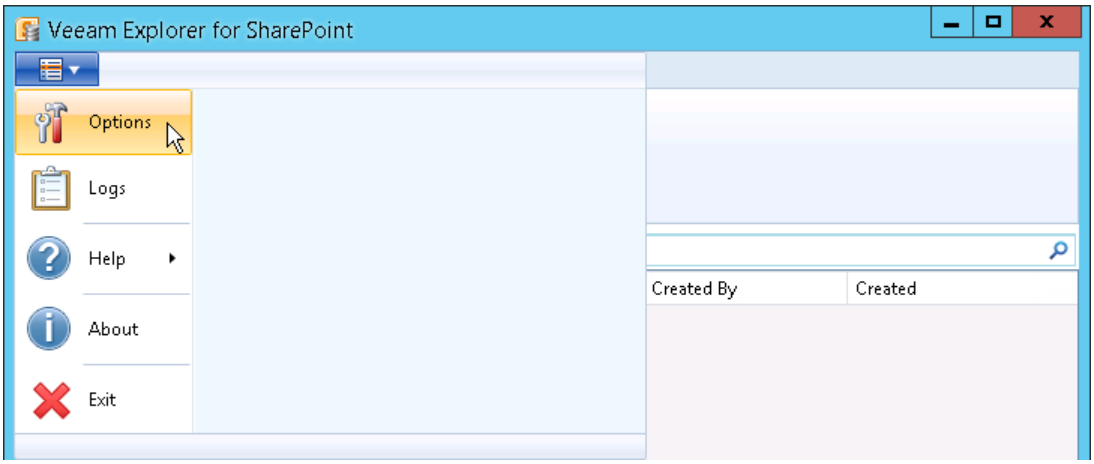

2. Specify the location of the staging Microsoft SQL Server. For details, see Staging SQL Server Settings.
3. Configure SMTP settings. For details, see SMTP Settings.

Staging SQL Server Settings

To perform Microsoft SharePoint item recovery, Veeam Explorer for SharePoint requires a Microsoft SQL Server that will be used as a staging system. To specify the location of the staging Microsoft SQL Server, do the following:

1. In the **Options** window, open the **SQL Server settings** tab.
2. In the **Use the following SQL server** field, specify the location of the staging Microsoft SQL Server.

Important! The staging Microsoft SQL Server must run on the machine where Veeam Explorer for SharePoint is installed (that is, the machine running Veeam Backup server).

By default, this field is populated with the location of the Microsoft SQL Server that hosts the Veeam Backup & Replication configuration database.

Note The specified system must run the same or a later version of Microsoft SQL Server as the server that hosts Microsoft SharePoint content databases. For details, see Staging Microsoft SQL Server.

To be able to work with remote BLOB stores (RBS) of the SharePoint content database, also make sure the staging SQL Server configuration meets the following requirements:

1. FILESTREAM should be enabled on the database server, and filestream settings should be enabled at database level, as described in Microsoft documentation:
   - For SQL Server 2012: http://msdn.microsoft.com/en-us/library/cc645923.aspx
   - For SQL Server 2008R2: http://msdn.microsoft.com/en-us/library/cc645923%28v=sql.105%29.aspx
2. RBS Client Library should be installed on the database server; use corresponding Microsoft SQL Server Remote Blob Store setup package (RBS.msi) available at Microsoft website:
   - For SQL Server 2012: http://msdn.microsoft.com/en-us/library/cc645923%28v=sql.105%29.aspx
   - For SQL Server 2008R2: http://www.microsoft.com/en-us/download/details.aspx?id=16978

SMTP Settings

If you want to send items recovered from a SharePoint content database as e-mail attachments, you need to configure SMTP server settings, as follows:

1. In the **Options** window, open the **SMTP settings** tab.
2. Select the **Use SMTP settings** check box.
3. In the **Server** field, specify the DNS name or IP address of the mail server that will be used for sending e-mails. Change the SMTP communication port if needed (by default, port number 25 is used).
4. In the **From** field, specify the e-mail address of the sender (for example, e-mail address of the administrator responsible for Microsoft SharePoint item recovery). When you choose to send restored items as e-mail attachments, this e-mail address will be used by default in the **From** field of the message form.
5. If your SMTP server requires SMTP authentication for outgoing mail, select the **Use authentication** check box and provide credentials for the account that will be used to send e-mails with attached items.
6. For SMTP server with SSL support, select the **Enable SSL security** check box to enable SSL data encryption.
7. To test if e-mail settings have been configured correctly, enter an e-mail address in the **Send test e-mail** section and click the **Send** button. Veeam Explorer for SharePoint will send a test e-mail message to the specified address.

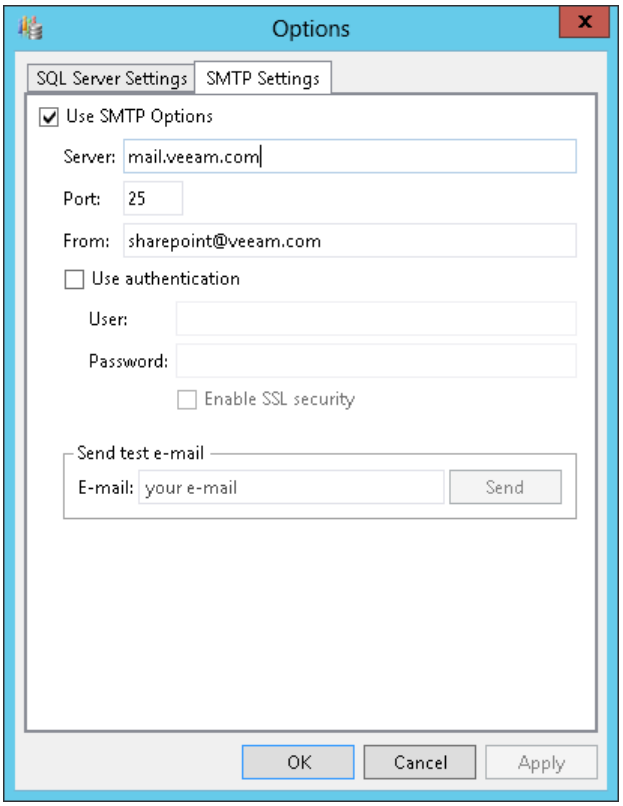

Before You Start Exploring

Veeam Explorer for SharePoint allows you to restore sites, lists and separate list objects from Microsoft SharePoint content database files. To be able to browse and search for the necessary item within the content database files, you should perform the following steps:

1. Recover Microsoft SharePoint content database (.MDF) from the image-level VM backup. For details, see Recovering Content Database Files from Backup.
2. Add the recovered Microsoft SharePoint content database files to Veeam Explorer for SharePoint scope. For details, see Adding Content Database to Veeam Explorer's Scope.

Recovering Content Database Files from Backup

Before adding content database files to Veeam Explorer for SharePoint, it is necessary to extract them from the backup or replica file. To do this, you can use any data recovery feature available in Veeam Backup & Replication: Instant VM recovery, full VM restore, replica failover and so on. Data recovery options are described in detail in Veeam Backup & Replication User Guide. When you perform file-level restore, the VM image is not extracted from the backup. The content of a backup file is mounted directly to the Veeam Backup server and displayed in the inbuilt Veeam Backup Browser. After the backup file is mounted to the Veeam Backup server, you can either add the necessary SharePoint content database files to Veeam Explorer for SharePoint directly, or copy the files and then add them to Veeam Explorer for SharePoint.

In particular, SharePoint Farm Restore Wizard can be launched if you select **Application Recovery** option in the File-Level Restore Wizard). Another option is to use Guest OS Restore Wizard.

Using SharePoint Farm Restore Wizard

If you need to locate and/or recover SharePoint items or lists from the VM backup, you can use the SharePoint Farm Restore Wizard, as described in this section.

During its work, the wizard will build SharePoint topology for the available backups done with application-aware image processing enabled (to read more about this option, refer to Veeam Backup & Replication documentation). Then Veeam Backup & Replication will be able to perform auto-discovery of the SharePoint web sites and corresponding SQL server VMs hosting content databases.

So, if you plan to use this automated workflow, make sure your SharePoint backup jobs have application-aware image processing option enabled.

Step 1. Launch the Wizard

To automatically select and restore Microsoft SharePoint content database files using SharePoint Farm restore wizard, do the following:

1. In the Veeam Backup & Replication console, go to the **Home** tab, click **Restore** and choose **VMware** or **Hyper-V**, depending on the platform you need.
2. In the **Restore Wizard** window, select **Application Items**.

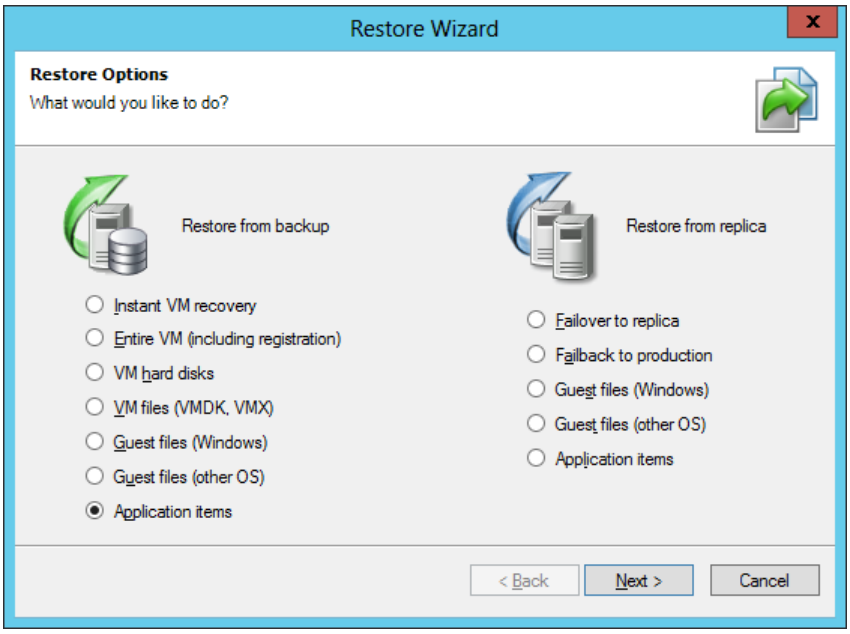

Step 2. Select Application

In the **Select Application** step of the wizard, select **Microsoft SharePoint**.

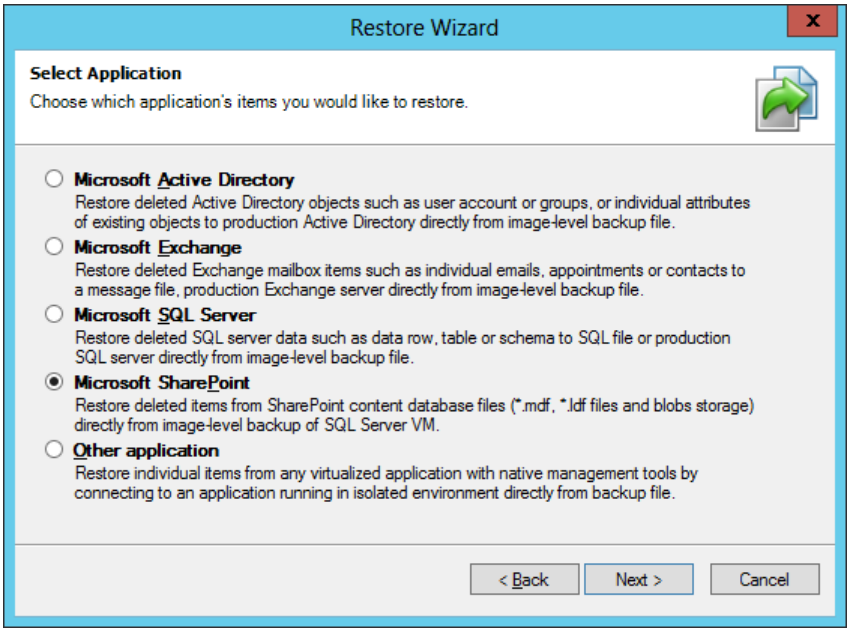

Wait while Veeam Backup & Replication builds SharePoint topology for available SharePoint farm backups.

Step 3. Select SharePoint Site

In this step, you can select the necessary Microsoft SharePoint site from the automatically populated list of available sites; use the search field to look for the site you need.

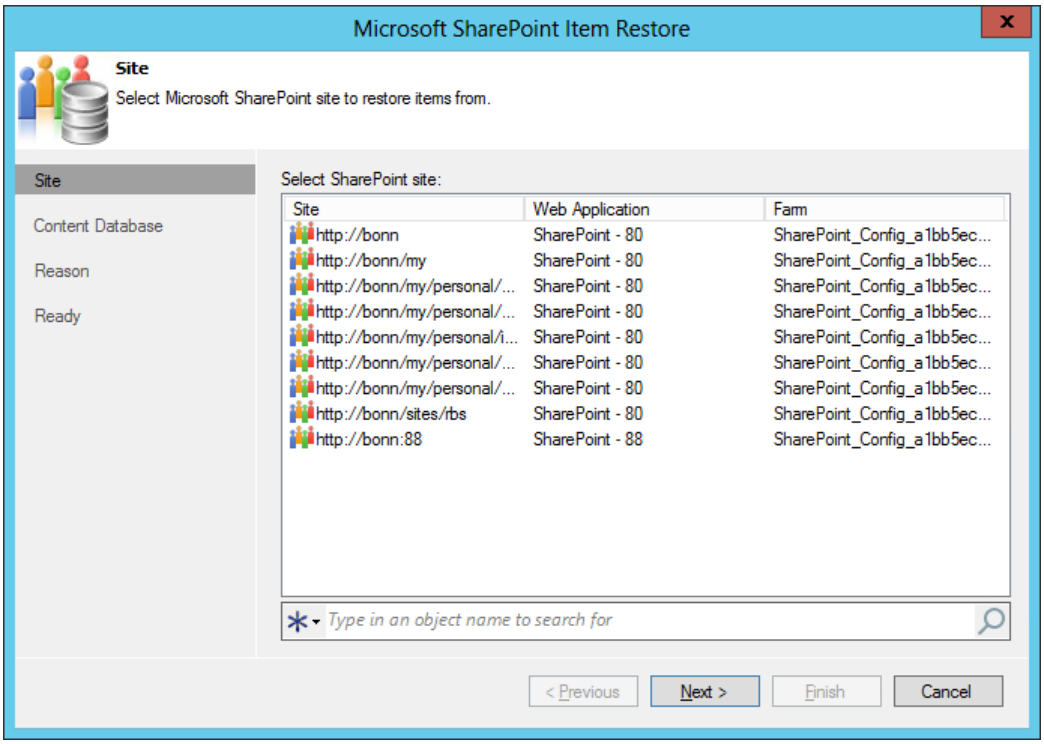

Note Veeam Backup & Replication performs auto-discovery for the SharePoint farms that were backed up with application-aware image processing enabled.

During auto-discovery, Veeam Backup & Replication retrieves information about SharePoint sites, corresponding database server VMs and their restore points. If, for some reason, database VM cannot be discovered, Veeam Backup & Replication will display a warning, notifying you that database should be recovered manually using Veeam Backup & Replication file-level restore capabilities. Then you can manually locate the content database and open it in Veeam Explorer for SharePoint.

Step 4. Select Database Restore Point

With auto-discovery, the VM hosting SQL Server instance with the SharePoint content database will be picked automatically. From the list of available restore points, select the restore point containing content you want to restore.

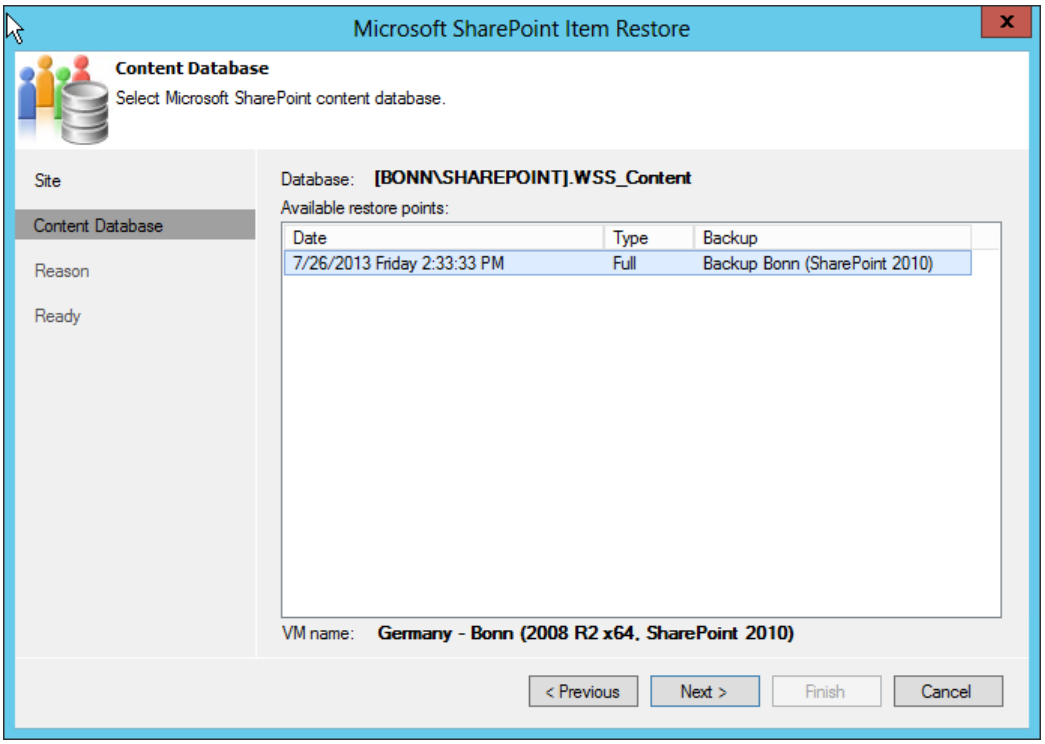

Step 5. Specify Restore Reason

If necessary, enter the reason for performing restore. The information you provide will be saved in the session history so that you can reference it later.

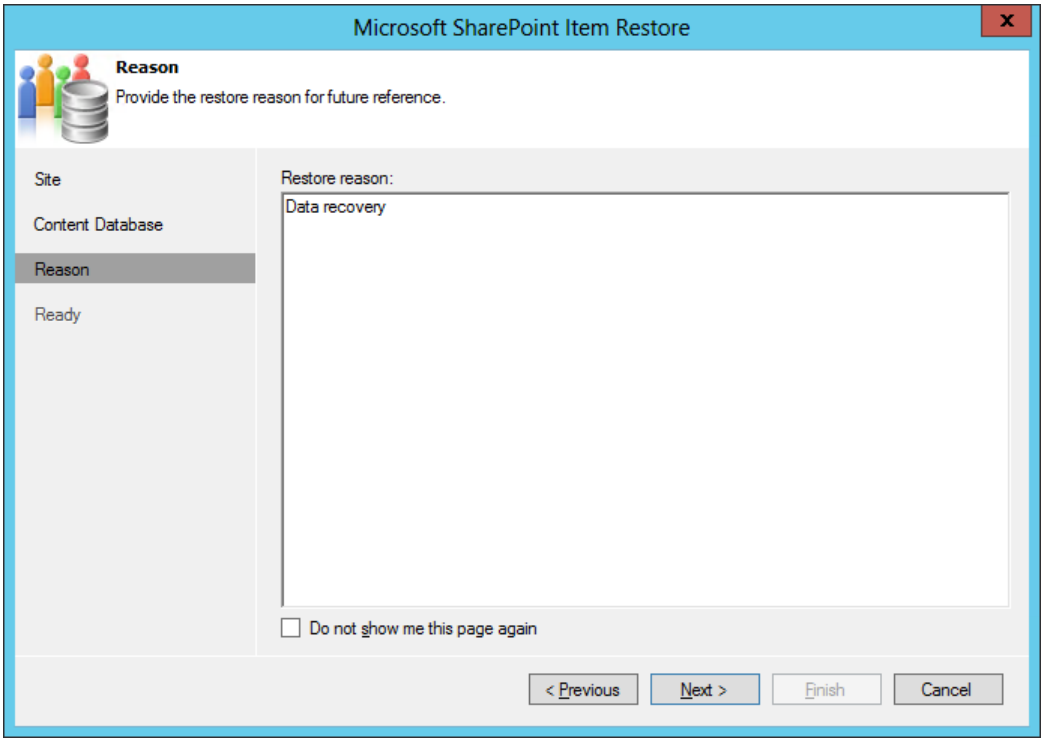

Step 6. Finish Working with the Wizard

Finally, review the restore settings. Due to auto-discovery of SharePoint topology, detailed information about SharePoint farm, including primary content database (.mdf), associated log file (.ldf) and BLOB storage (if any) will be presented:

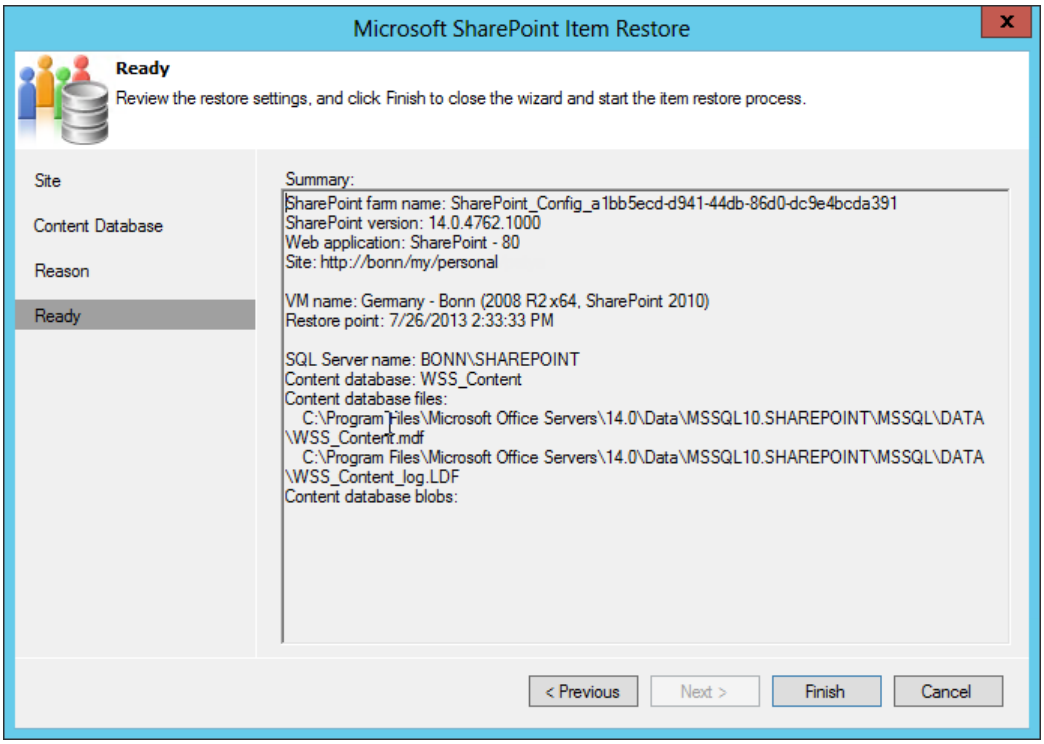

Click **Finish** to start the restore process. Veeam Backup & Replication will perform the restore, and then required database can be added to Veeam Explorer for SharePoint scope. Alternatively, you can manually locate the content database and open it in Veeam Explorer for SharePoint.

After that, you can browse the database and restore the items you need.

Using Guest OS Restore Wizard

Another way to extract a content database file from the VM backup is to perform guest OS file recovery using Veeam Backup & Replication File-Level Restore (FLR) wizard. When you perform file-level restore, the VM image is not extracted from the backup. The content of a backup file is mounted directly to the Veeam Backup server and displayed in the inbuilt Veeam Backup Browser. After the backup file is mounted to the Veeam Backup server, you can either add the necessary SharePoint content database files to Veeam Explorer for SharePoint directly, or copy the files and then add them to Veeam Explorer for SharePoint.

This section will guide you through the steps of the File-Level Restore (FLR) wizard necessary to mount the backup file to the Veeam Backup server and extract the content database files.

Step 1. Launch the Restore Wizard

To restore Microsoft SharePoint content database files using guest OS restore wizard, open the Veeam Backup & Replication console and do any of the following:

- On the **Home** tab, click **Restore**. In the **Restore Wizard** window, select **Guest Files > Guest Files (Windows)** .
- Open the **Backup & Replication** view and select the **Backups** node. In the working area, expand the necessary backup job, select the VM whose guest OS files you want to restore and click **Application Items > Microsoft SharePoint** on the toolbar. You will immediately pass to Restore Point of the wizard.
- Open the **Backup & Replication** view and select the **Backups** node. In the working area, expand the necessary backup job, right-click the VM whose guest OS files you want to restore and select **Restore guest files (Windows)**. You will immediately pass to Restore Point of the wizard.

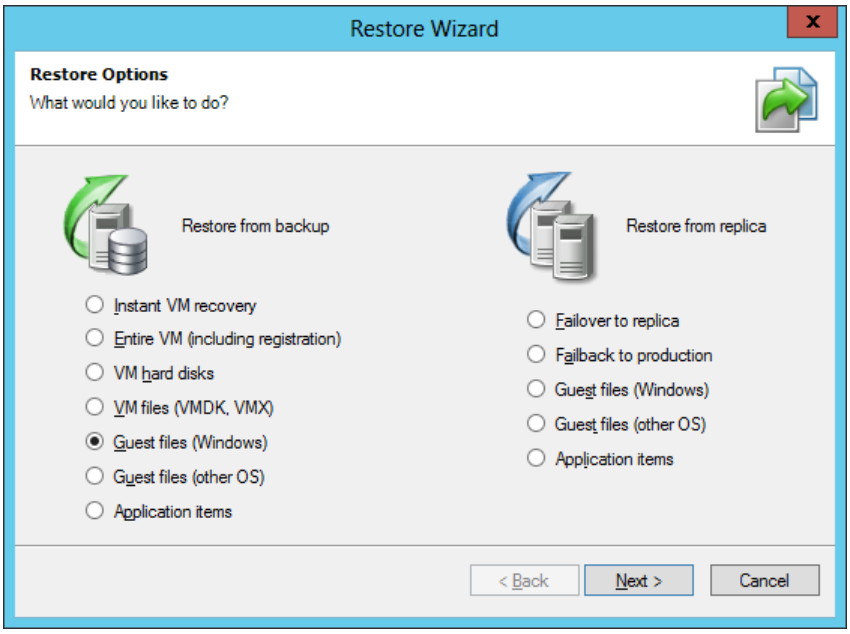

Step 2. Select a VM

In the list of available jobs, select the backup of the VM from which the content database should be restored. To quickly find VMs in jobs, use the search field at the bottom of the window.

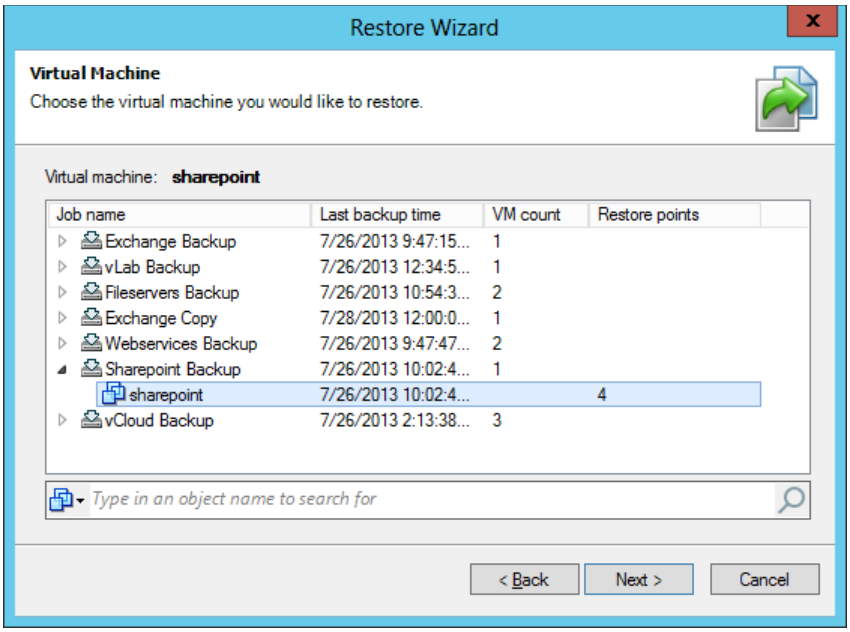

Step 3. Select a Restore Point

Select the necessary restore point for the virtual machine.

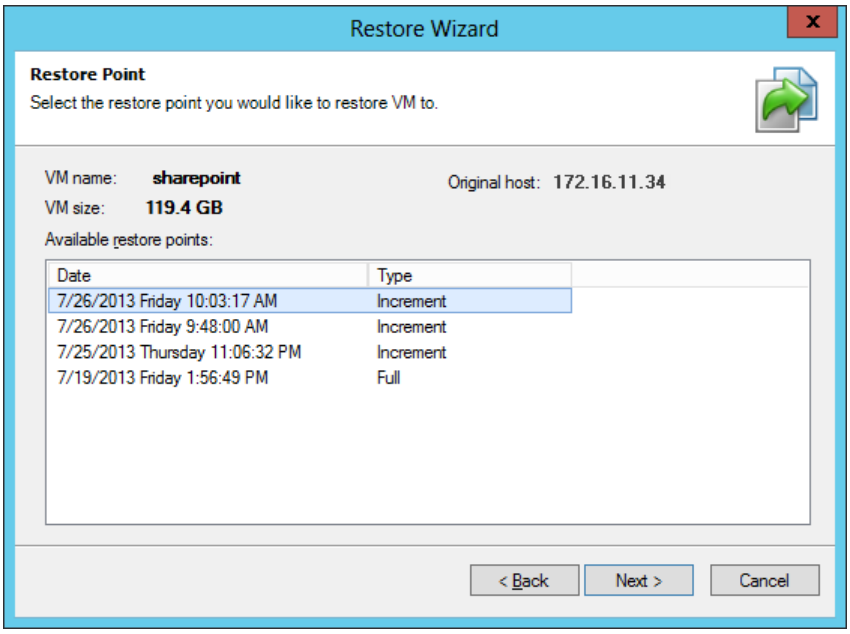

Step 4. Specify Restore Reason

If necessary, enter the reason for performing restore. The information you provide will be saved in the session history so that you can reference it later.

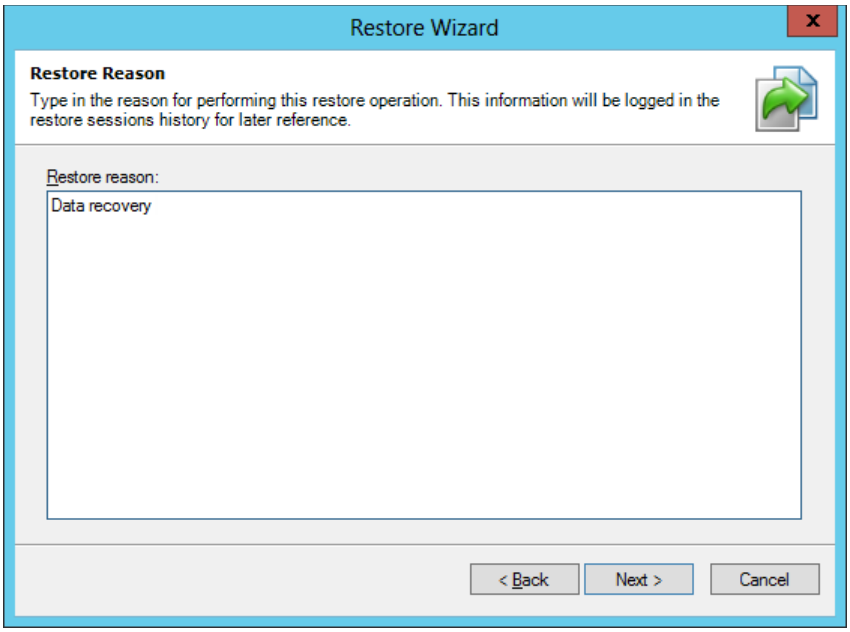

Step 5. Finish Working with the Wizard

Click **Finish** to start restoring files. Once restore is completed, Veeam Backup & Replication will open a Backup Browser displaying the file system tree of the restored VM. Please note that the names of the restored machine drives may differ from the original ones.

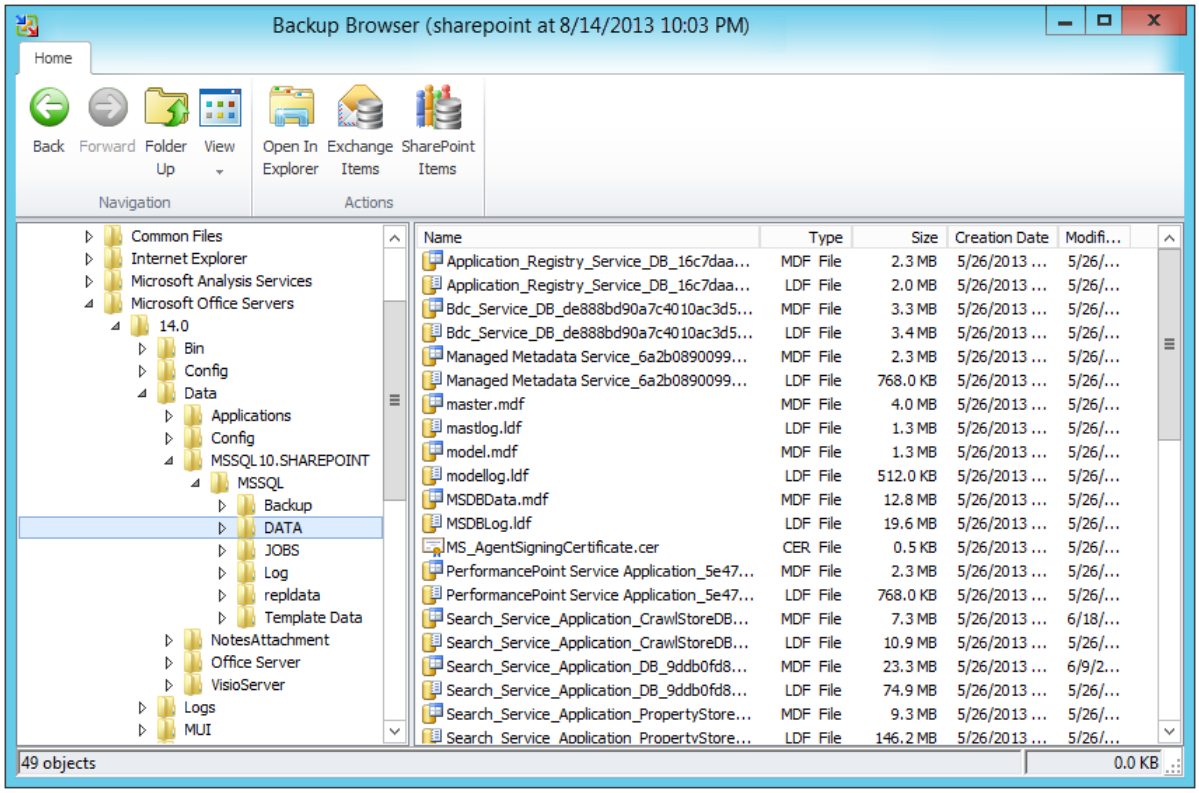

Then you can click **SharePoint Items** on the Backup Browser toolbar to launch Veeam Explorer for SharePoint and mount the discovered database files automatically, or manually locate the required database(s) and then add them to Veeam Explorer for SharePoint scope.

See also: Locating Content Database File, Adding Content Database to Veeam Explorer's Scope

Locating Content Database File

Once you have mounted the VM hosting content database using Guest OS Restore Wizard, you can locate the database files with the content which you want to restore. By default, these files reside by the following path:

```
%ProgramFiles%\Microsoft Office
Servers\14.0\Data\MSSQL10.SHAREPOINT\MSSQL\DATA
```

Next, to make the content database file available to Veeam Explorer for SharePoint, you can perform the following steps:

1. In the navigation tree of the Backup Browser, locate the folder where content database files reside.
2. Select the .mdf file you need and click **SharePoint Items** on the toolbar. Veeam Explorer for SharePoint will be launched and content database added to its scope.

Note If you choose to restore SharePoint items directly from the mounted VM file system, do not close the Backup Browser until you finish restore operations. Closing the Backup Browser automatically unmounts the file system of the backed up VM from the Veeam Backup server.

If you need to keep SharePoint content databases hosted on the staging Microsoft SQL Server for a period of time that is longer than a File-Level Restore session, it is recommended that you copy the database files to a different location rather than restore them directly from the mounted VM file system. To copy files, do the following:

1. In the Backup Browser, open the **File** view.
2. In the navigation tree, locate the folder where content database files reside; select Microsoft SharePoint content database files (.mdf) and associated log files (.ldf).
3. On the Backup Browser toolbar, click **Copy To**.
4. In the **Choose Folder** window, specify the location to which content database files will be copied and click **OK**.

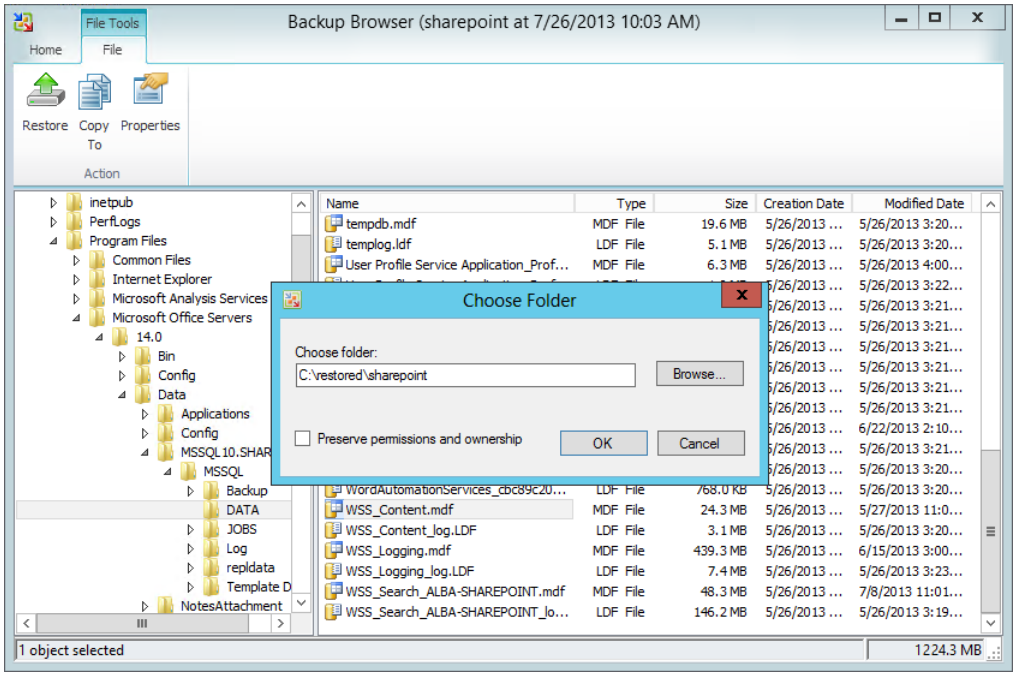

6. Open the folder to which files were copied and copy the full path to the target directory to clipboard. You can now launch Veeam Explorer for SharePoint and add the database file to its scope, as described in Adding Content Database to Veeam Explorer's Scope.

Adding Content Database to Scope

For users to be able to browse and restore the SharePoint items, Microsoft SharePoint content database should be added to the Veeam Explorer for SharePoint scope. When this operation is performed, Veeam Explorer for SharePoint automatically attaches the database to the staging Microsoft SQL server, creating a temporary Microsoft SharePoint content database from which you can recover the necessary items.

In case you have used SharePoint Farm Restore Wizard to perform the restore, the content database will be added to Veeam Explorer for SharePoint scope automatically and become available for browsing. Alternatively, you can add the required database to Veeam Explorer for SharePoint scope manually, as described below.

- If you know in which database the necessary Microsoft SharePoint items are located, you only need to add one .MDF file.
- If you are not sure in which database the necessary Microsoft SharePoint items are located, or these items are scattered across a number of databases, you can add multiple .MDF files to Veeam Explorer for SharePoint scope (no limitations are applied).

To add a database file to the Veeam Explorer for SharePoint scope manually and make it visible in the console:

1. Open Veeam Explorer for SharePoint and click **Add Database** on the toolbar.
2. In the displayed window, specify the location of the Microsoft SharePoint primary content database file (MDF); corresponding secondary database and transaction log file (LDF) will be also added. If necessary, you can also add the remote BLOB (binary large objects) stores (RBS).

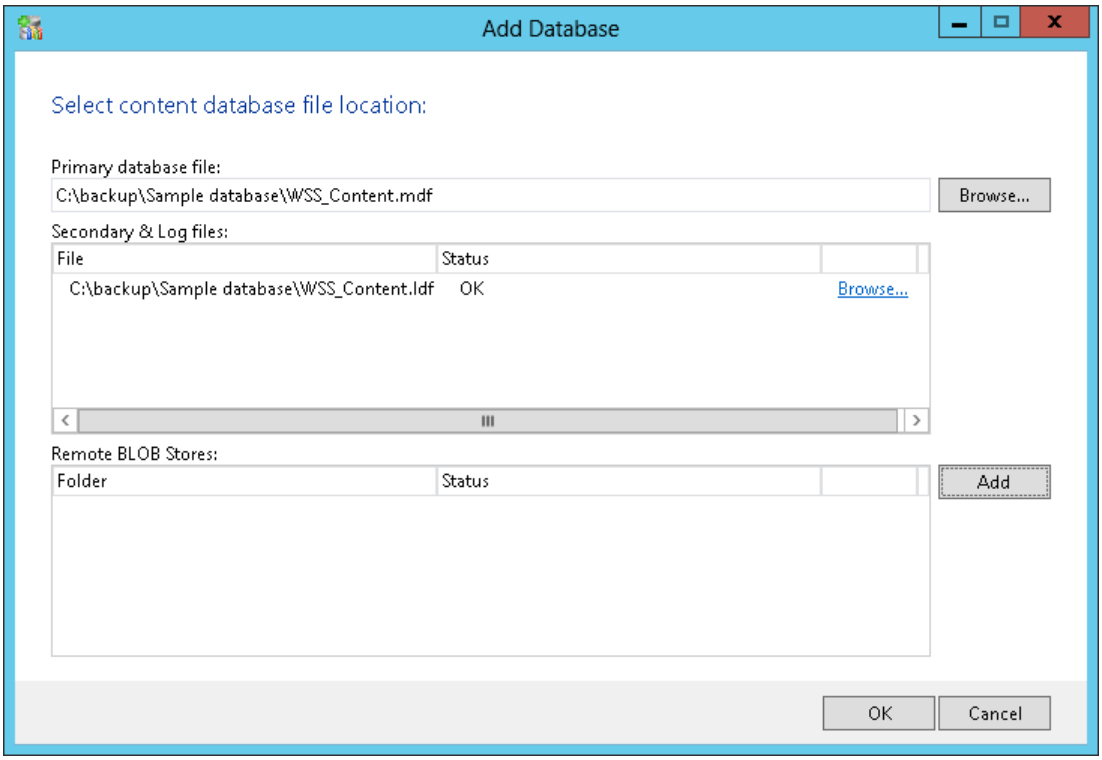

3. Click **OK** and wait for the operation to complete.

Note To successfully attach the database to staging SQL server, you should ensure that SQL server service account has sufficient rights to access the database files. Otherwise, you will get the following message displayed: "*SQL server cannot access sharepoint database file. SQL Server account: <service_account>*". If so, provide the access rights to the specified account.

Now selected database is attached to the staging server. It is displayed in the Veeam Explorer for SharePoint console, and you can view its content and carry out the actions you need (search, export, restore and so on).

You can add more than one content database to Veeam Explorer for SharePoint scope using the same steps.

To remove a content database from the Veeam Explorer for SharePoint scope:

1. Select the database in the navigation pane.
2. Click **Remove Database** on the toolbar. The database will be removed from the Veeam Explorer for SharePoint scope, no longer shown in the console and automatically detached from the staging SQL server.

Searching and Restoring SharePoint Items

After you add one or several Microsoft SharePoint content databases to the Veeam Explorer for SharePoint scope, you can browse through the contents of these database(s), find and restore necessary Microsoft SharePoint items.

- Browsing, Searching and Viewing Microsoft SharePoint Items (supported for all editions of Veeam Backup & Replication)
- Restoring Microsoft SharePoint Items

Browsing, Searching and Viewing Microsoft SharePoint Items

In the Veeam Explorer for SharePoint console, you can browse and search through the contents of the database(s) to find the necessary items.

Browsing

In the Veeam Explorer for SharePoint navigation pane, you can see all the containers (added content databases, sites, subsites, lists and list objects). When you select a container node, you can browse its contents in the working area (library documents, discussion boards, agenda notes and so on).

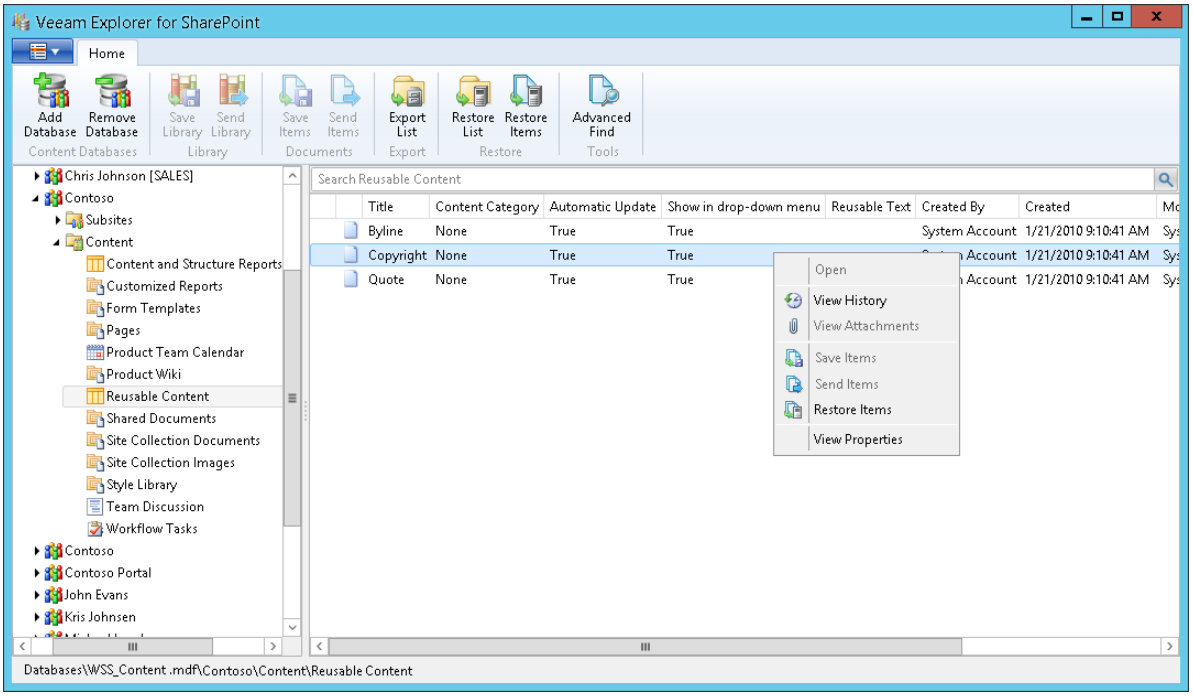

Veeam Explorer for SharePoint allows you to browse to associated Microsoft SharePoint items — attached files and document versions. For example, if you maintain version history for documents, you can right-click the item you need and select **View History** to examine the list of all stored document versions. If you want to open the list of attachments for an item, from its shortcut menu select **View Attachments**.

Searching

At the top of the working area, there is a search field that allows you to find items that match a specified search term. For example, you can select a document library and search for *veeam software* to find all the documents that contain the words starting with *'veeam'* or *'software'* in the document name or document text. Search results for content databases and sites will include all child items (that is, subsites, lists and list objects).

Note When you search a phrase, the search will by default locate items that contain any search terms in the phrase, not the phrase itself. To find the exact phrase, use double quotes, for example, *"veeam software"*.

You can select the whole content database or several databases at once to get more results. You can select a specific node (for example, a document library) inside a specific database and a site to narrow the search.

You can further narrow your search results by specifying various search criteria inside the search box in the *criteria:value* format. For example, to find in the list of decisions all items that require approval, you can use the following search query: *status:pending approval*. You can also use logical operators such as *AND*, *OR* and *NOT* (must be typed in uppercase letters) and wildcard characters such as * and *?*. These search criteria are similar to those used for searching in Microsoft SharePoint 2010. For more information, see this Microsoft SharePoint help topic.

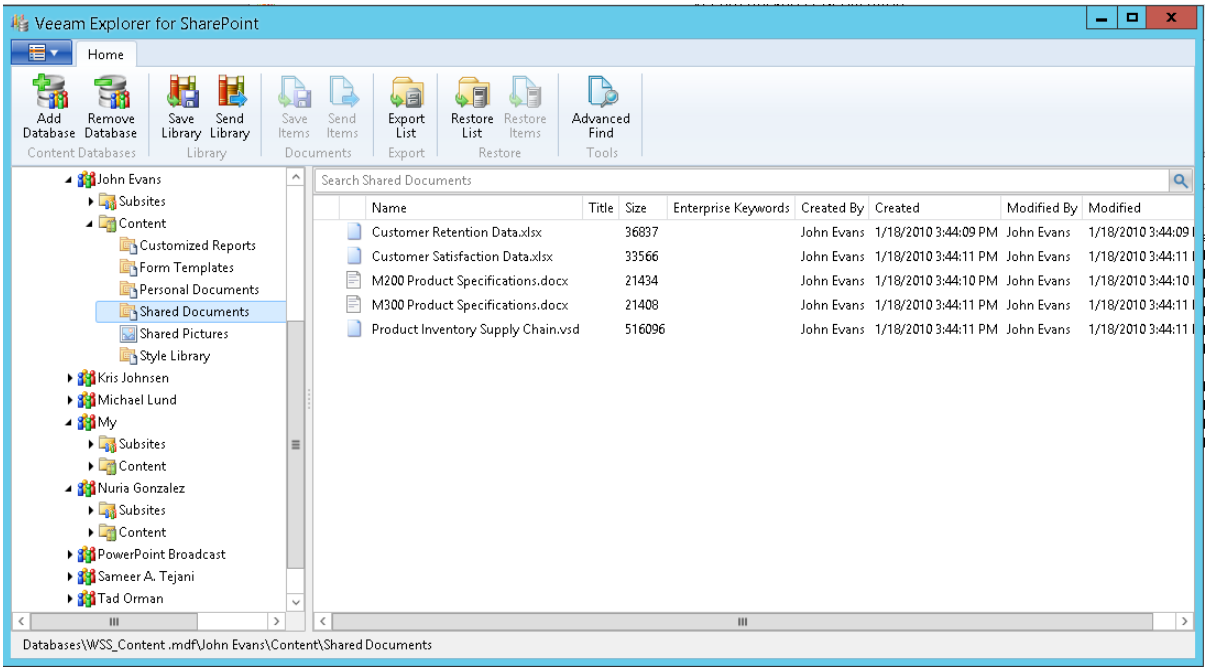

You can also use the **Advanced Find** command, which is available from the ribbon menu. With this option, you can apply flexible search criteria to selected document library/item list, like document author, creation or modification date, file extension, and so on, easily building any search query you need. For example, to find all files modified prior to the certain date, you can do the following:

1. Select the required node from the content tree on the left, then click **Advanced Find**.
   In the **Define search criteria** section of the search window on the right, select the **Category** for the new search filter. This will instruct search to look through the corresponding content property fields. In this example, to filter by date, choose **Date/Time fields**.
2. From the **Field** list, select the date to filter by – this can be **Date Created** or **Date Modified** (chosen for this example).
3. From the **Condition** list, select the one you need – this can be less than (chosen for this example), equal to, between, and so on.
4. Specify the **Value** for the date to look for, and click **Add to List**. Configured filter will be shown in the above pane.

Click **Start** to search using the new filter.

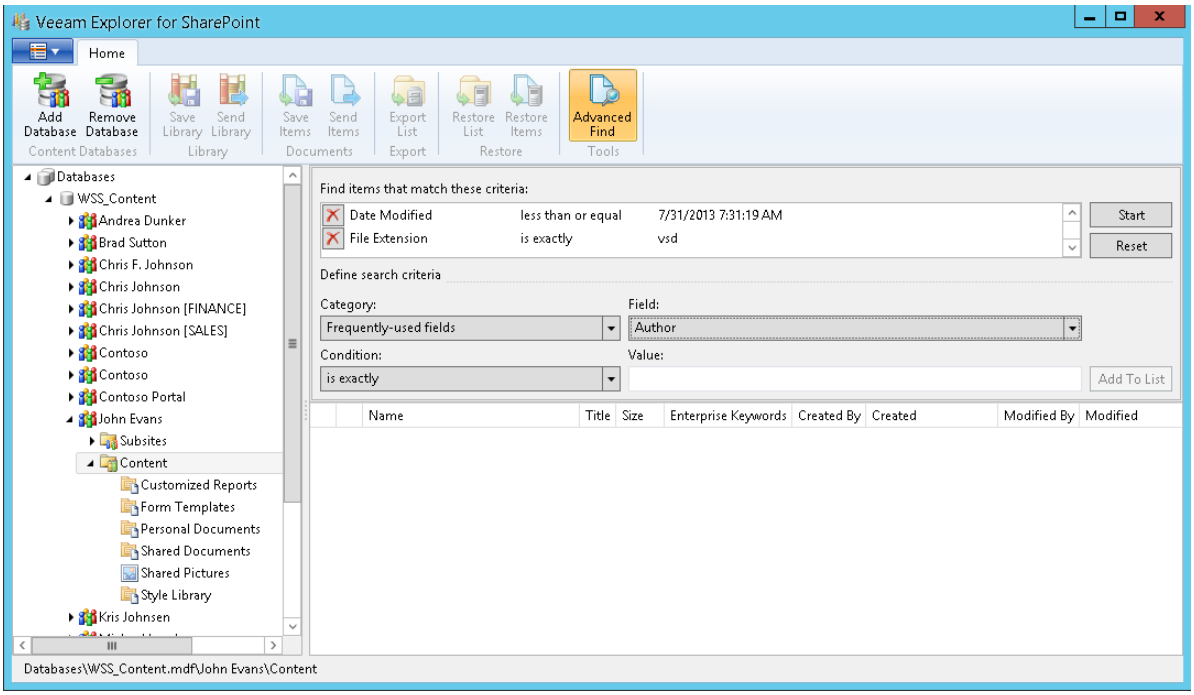

To remove a filter, click on the cross mark next to it; to remove all configured filters, click **Reset**.

Viewing Item Details

Veeam Explorer for SharePoint allows you to quickly view properties of any Microsoft SharePoint item. To view item properties, right-click an item in the list and choose **View Properties**. The properties of the chosen item will be displayed in a separate window.

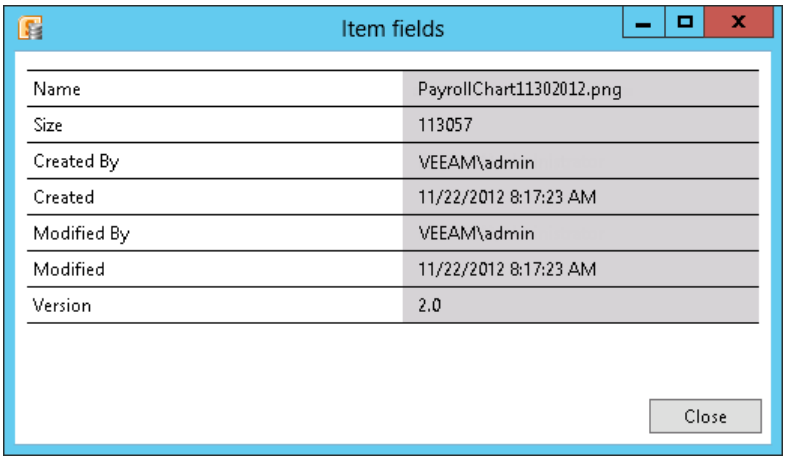

Veeam Explorer for SharePoint also provides possibilities for opening library documents. To open a document, right-click it in the list and choose **Open**. The document will be opened in the associated application.

Restoring Microsoft SharePoint Items

Veeam Explorer for SharePoint provides the following restore options for Microsoft SharePoint items:

- Saving documents and lists to the selected folder
- Sending documents and lists as e-mail attachments
- Exporting libraries and lists for further import to Microsoft SharePoint
- Restoring documents and lists to Microsoft SharePoint (if used with Enterprise or Enterprise Plus Edition of Veeam Backup & Replication)

Saving Documents and Lists

Veeam Explorer for SharePoint allows you to save any Microsoft SharePoint document library/list or any document/item to a specific location or directly to your desktop. To save a document library or a list of items, do the following:

1. Browse to the necessary document library or list, or use the search field to locate the one you need.
2. Select the required library and click **Save Library** on the toolbar; you can also use the shortcut menu command.

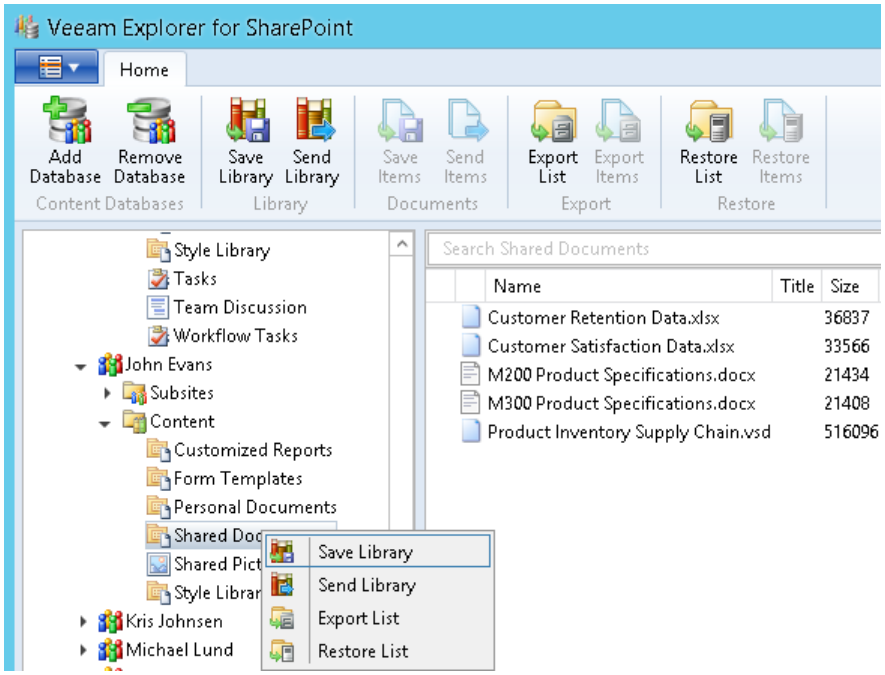

To save a document or a list item, do the following:

1. Browse to the necessary item or use the search field to see a list of items in the working area.
2. Select the required item (or use **[SHIFT]** and **[CTRL]** keys to select multiple items) in the working area and click **Save Items** on the toolbar. You can also right-click the selected item(s) in the working area and select **Save Item**.

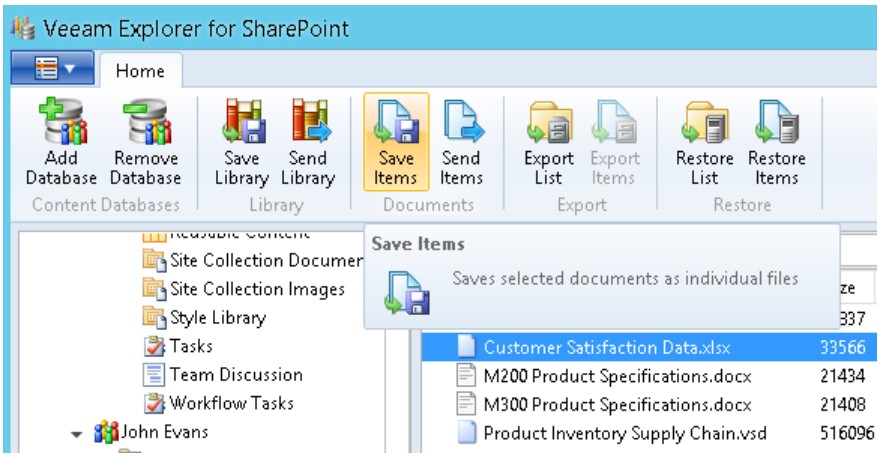

3. Specify the location for the file or files and click **OK**.

Note Veeam Explorer for SharePoint does not keep original ownership and access settings for restored documents. Access permissions for the saved document will be inherited from the folder to which the restored document is copied.

E-mailing Documents and Lists

Veeam Explorer for SharePoint allows you to send restored document libraries/lists and documents/list items to their owners as e-mail attachments. For that, it is necessary to configure e-mail settings in the **Options** window, as described in SMTP Settings.

To send a document library/list as an e-mail attachment, do the following:

1. Browse to the necessary library/list in the navigation tree.
2. Select the library you need and click **Send Library** on the toolbar; you can also use the shortcut menu command for the selection.

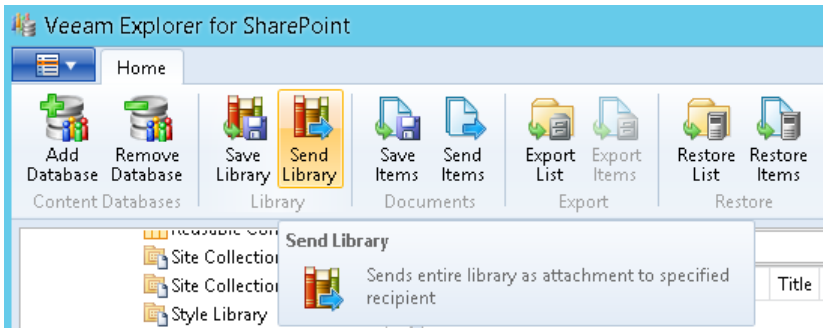

3. Specify the e-mail addresses of the sender and recipient as well as the subject of the e-mail message. To edit the body text, click **More details**.
4. Click **Send**.

To send a document/item as an e-mail attachment, do the following:

1. To display available documents/items in the working area, browse to the necessary document library/item list in the tree on the left, or use the search field on the right.
2. Select the required item (or use **[SHIFT]** and **[CTRL]** keys to select multiple items) in the working area and click **Send Items** on the toolbar. You can also right-click the selected item(s) in the working area and select **Send Items**.

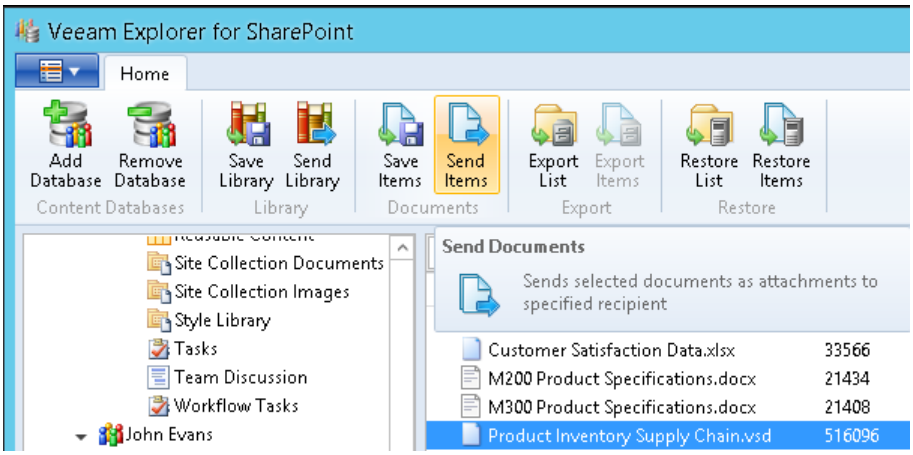

3. Specify the e-mail addresses of the sender and recipient as well as the subject of the e-mail message. To edit the body text, click **More details**.
4. Click **Send**.

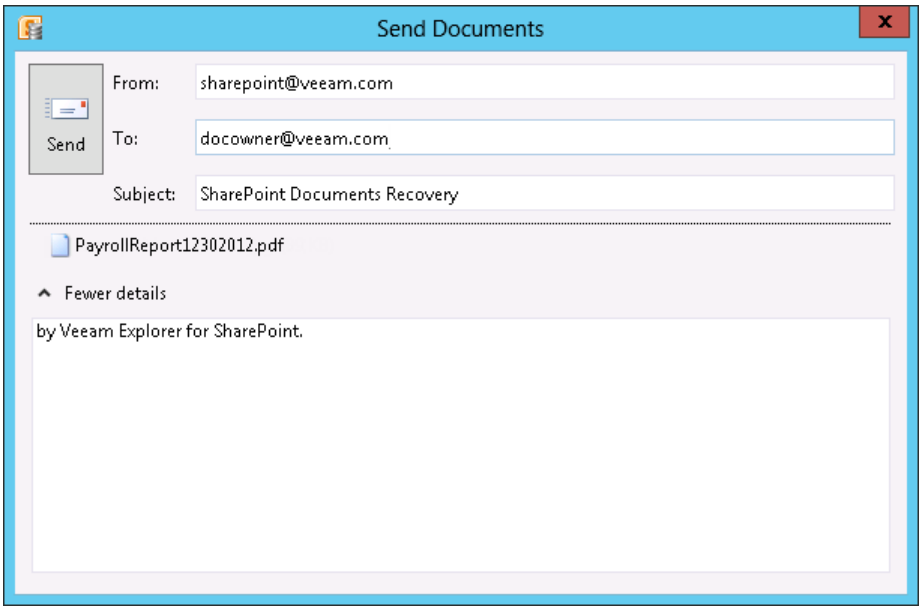

Exporting Document Libraries and Lists

With Veeam Backup & Replication, you can export the whole document library or item list to a folder in the specified location (export of items is not supported in this version). Exported content will be saved in XML files and can be then imported to the SharePoint database of your choice using PowerShell cmdlets.

To export a library/list, do the following:

1. Select the required library/list in the navigation pane of Veeam Explorer for SharePoint.
2. Click **Export Library** on the toolbar.
3. Specify destination location.
4. Click **OK** and wait for the export to complete.

Exported content appears as .DAT and .XML files in the specified destination folder:

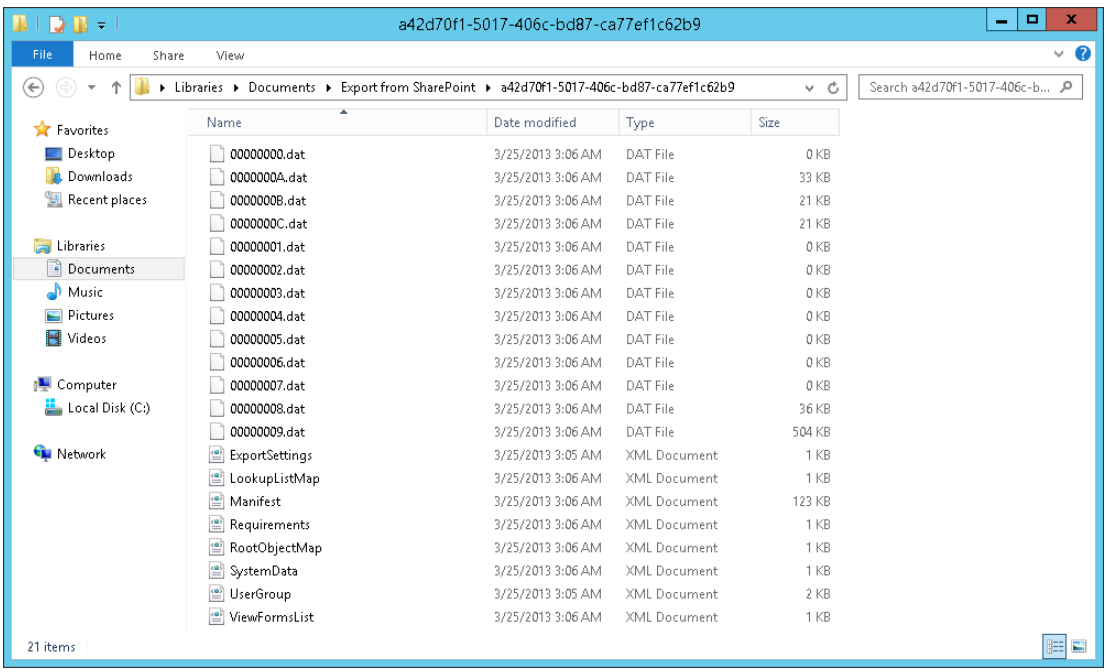

You can make this folder content available to SharePoint server you need, importing it by means of PowerShell command.

See also: Importing Exported Content

Importing Exported Content

To import document library/list you have exported from the SharePoint content database, run the appropriate PowerShell cmdlet locally on the SharePoint server, as described below:

If using PowerShell snap-in, run the following:

```
Add-PsSnapin Microsoft.SharePoint.PowerShell

Import-SPWeb -Identity
"http://<web_server_name>/sites/<destination_site>" -Path
"C:\<export_folder>" -NoFileCompression –IncludeUserSecurity
```

If using SharePoint 2010 Management Shell, run the following:

```
Import-SPWeb -Identity
"http://<web_server_name>/sites/<destination_site>" -Path
"C:\<export_folder>" -NoFileCompression –IncludeUserSecurity
```

where:

<*web_server_name*> - destination web server;

<*destination_site*> - destination web site;

<*export_folder*> - source folder containing exported library/list content.

To get extended Help on the **Import-SPWeb** command, run the following PowerShell cmdlet:

```
Get-Help Import-SPWeb -full
```

See also: Exporting Document Libraries and Lists

Restoring Documents and Lists to Microsoft SharePoint

You can use Veeam Explorer for SharePoint to restore a SharePoint document/item or the whole document library/item list to the specified location on the SharePoint server in your production environment. Note that this capability is supported for Veeam Backup & Replication Enterprise and Enterprise Plus Editions only.

Restoring Document Libraries and Lists

To restore a document library/list, do the following:

1. Select the required library/list in the navigation pane of Veeam Explorer for SharePoint.
2. Click **Restore List** on the toolbar; you can also use the shortcut menu command.

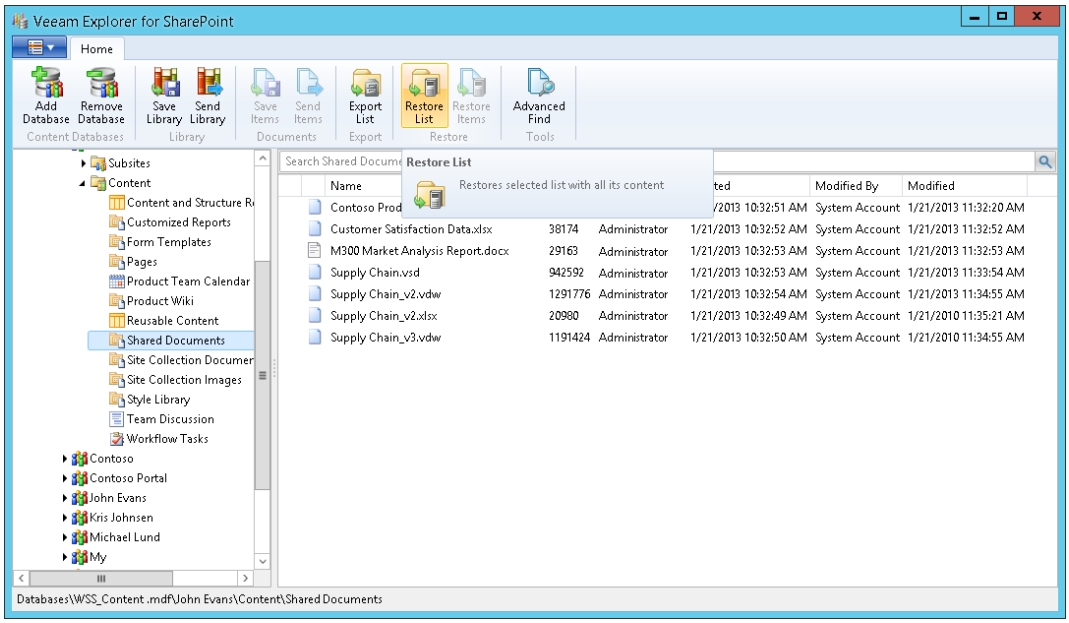

3. The restore wizard is launched; follow its steps to specify restore options for the document library/item list.

Step 1. Specify Target SharePoint

Specify target settings to be used for restoring SharePoint content:

- Target SharePoint server's URL (as `http://server_name`) and site path.
- Domain account to be used for connection. You can use the account under which you are running Veeam Explorer for SharePoint, or specify another account in the domain\username format. Make sure this account has sufficient rights to access the specified server (see Permissions).

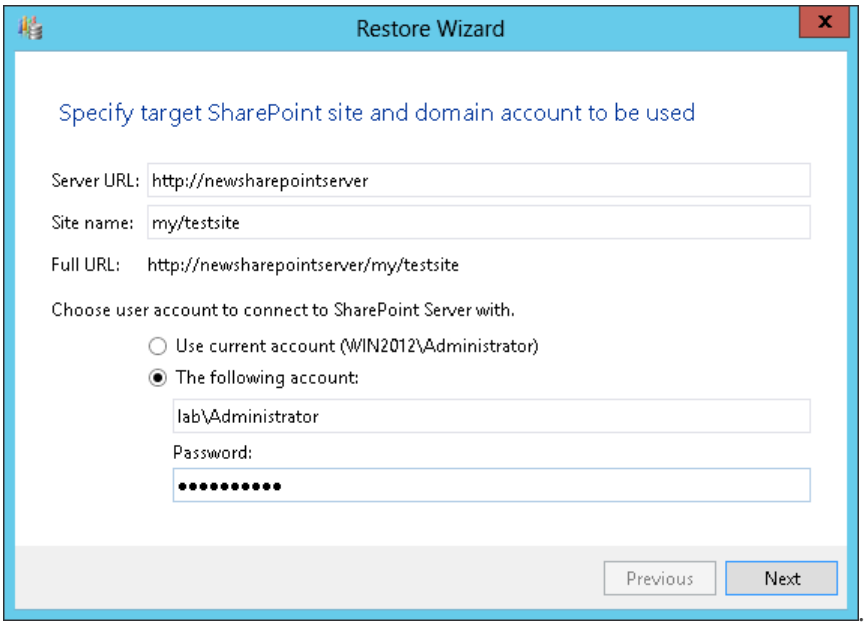

Step 2. Specify Target List

Next, specify whether the list should be restored to the same list/document library as the original one, or to a different list.

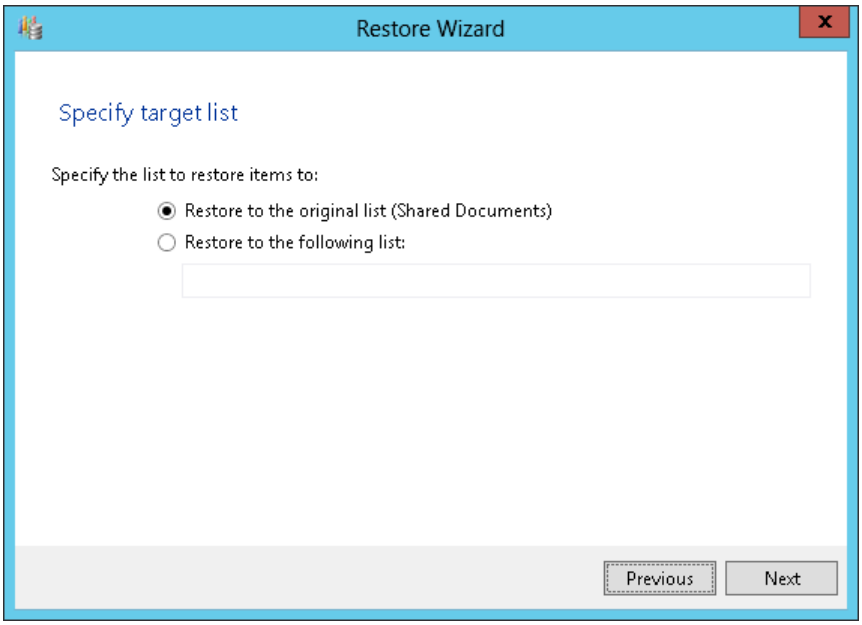

Step 3. Specify Restore Options

Then specify the restore options you want to be applied — you can select to restore **Changed items** and/or **Deleted items**.

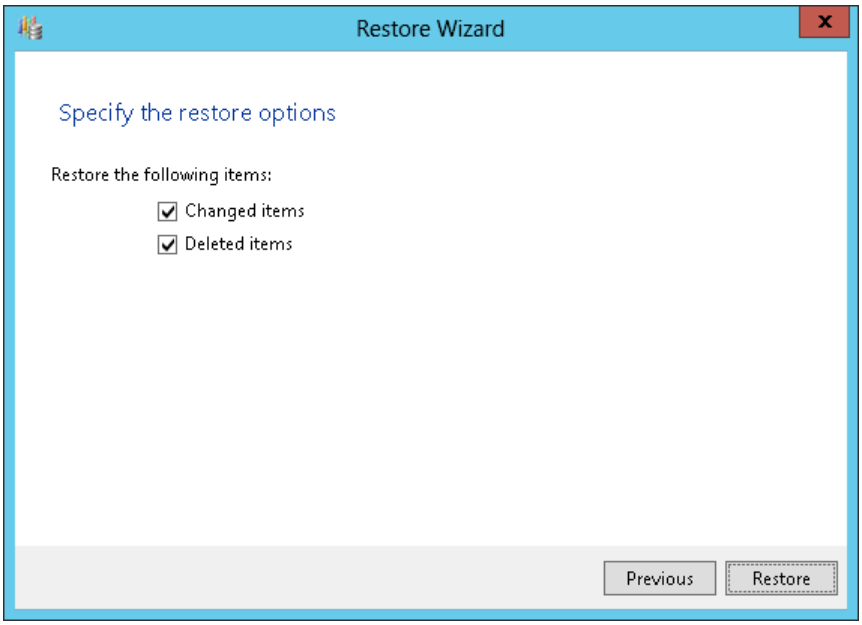

Note When configuring restore options, consider some peculiarities described in the Recovery Specials section.

Click **Restore** and wait for the operation to complete.

Restoring Documents and List Items

To restore a document/item, take the following steps:

1. To display available documents/items in the working area, browse to the necessary document library/item list in the tree on the left, or use the search field on the right.
2. Select the required item (or use **[SHIFT]** and **[CTRL]** keys to select multiple items) in the working area and click **Restore Items** on the toolbar. You can also right-click the selected item(s) in the working area and select **Restore Items**.

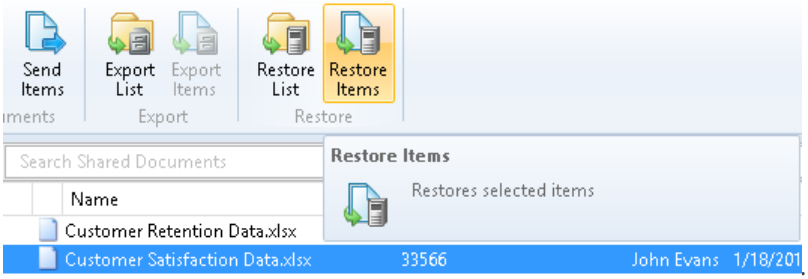

3. The restore wizard is launched; follow its steps to specify restore options for the document library/item list.

Step 1. Specify Target SharePoint

Specify target settings to be used for restoring SharePoint content:

a. Target SharePoint server's URL (as http://server_name) and site path.

b. Domain account to be used for connection. You can use the account under which you are running Veeam Explorer for SharePoint, or specify another account in the domain\username format. Make sure this account has sufficient rights to access the specified server (see Permissions).

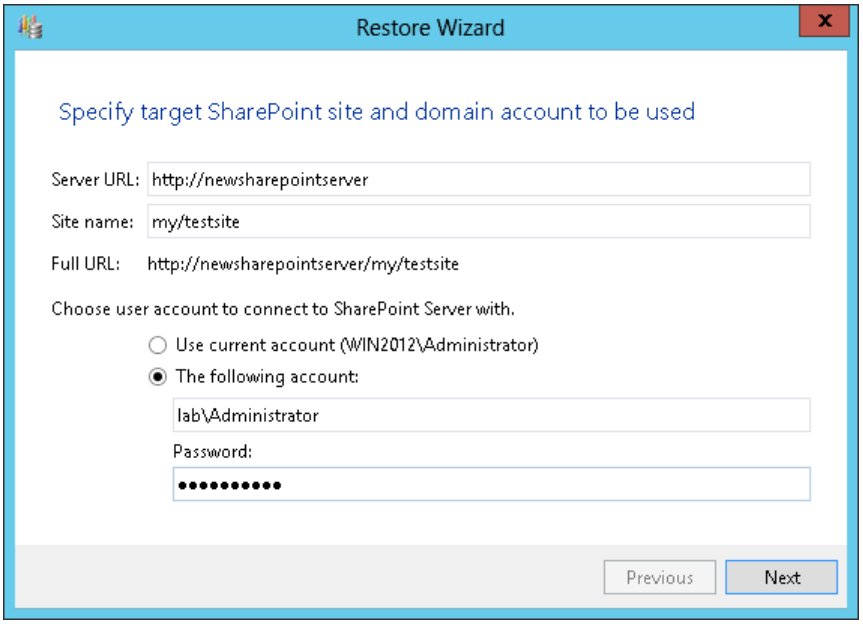

Step 2. Specify Target Location

Next, specify whether the document/item should be restored to the original list or to a different list.

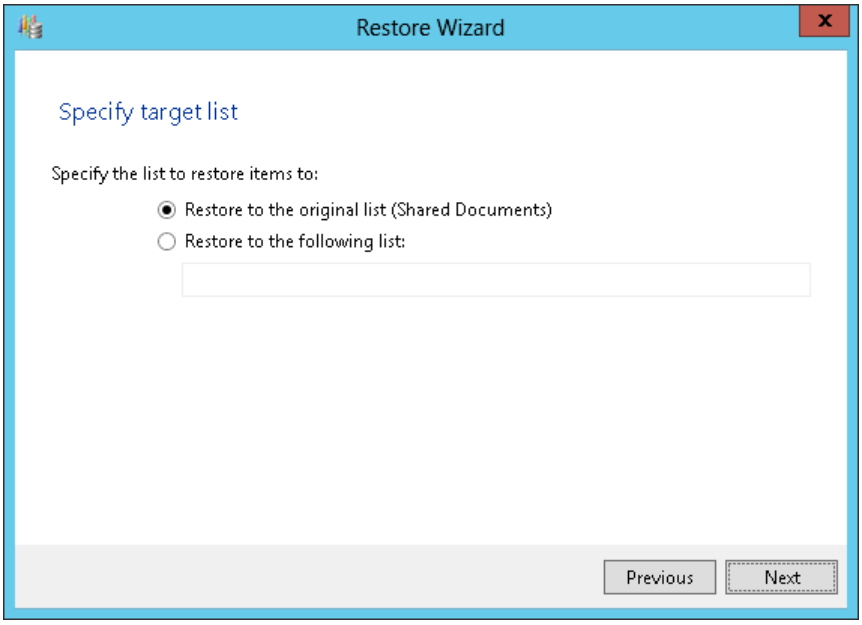

Click **Restore** and wait for the operation to complete.

Step 3. Specify Restore Options

Then specify the restore options you want to be applied — you can select to restore **Changed items** and/or **Deleted items**.

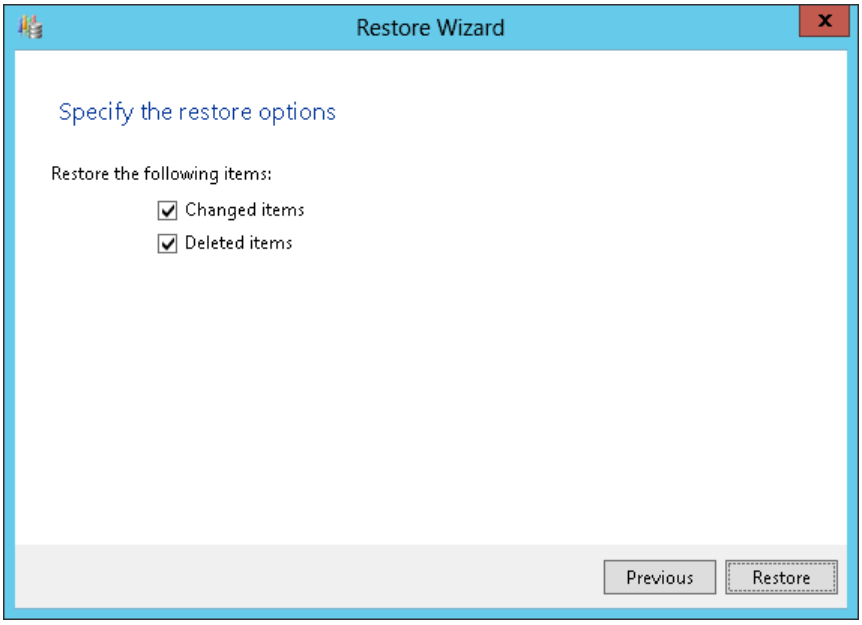

Note	There are some peculiarities of SharePoint item recovery support with Veeam Explorer for SharePoint. For example, you may need to restore the item(s) originally belonging to a list with no content approval required, to another list. If you try to restore such item(s) to a list that requires content approval, item version and status will be modified. For more details, refer to the Recovery Specials section.

Click **Restore** and wait for the operation to complete.

Status Recovery Limitations

Consider the following when planning for the recovery of the documents/list items:

- If a document/item was in *Check Out* state when the backup was created, item's last version will not be restored to the target SharePoint but will be available for viewing only; previous versions (if any) will be restored.
- If *Declare this item as a record* action was originally applied to list item, the corresponding status will not be preserved; instead, restored item status will be set in accordance with the target list/library content approval workflow.
- Original status *On Hold* will not be restored

Restoring List Items with Links (Attachments)

Consider the following when planning for the recovery of list items with links (attachments):

- If the retention policy for target list/document library was configured to **Declare record** automatically, only the last version of the item will be restored to target list/library. Target retention policy settings will be applied to restored item (**Declare record**). However, links (attachments) will not be restored.
- Alternatively (with different retention policy settings), all versions of the original item will be restored to target list/library; item links (attachments) will be restored only if such item does not exist on target SharePoint.
- If *Declare this item as a record* action was originally applied to list item, such item will not be restored.

Restoring Surveys

Consider the following limitations when planning for the recovery of surveys, survey questions and responses:

- Survey item(s) can be restored to a new survey, created automatically by Veeam Explorer for SharePoint in the specified destination instead of the previously deleted survey. However, if a new survey is created by user from scratch (not replacing a deleted one) – items cannot be recovered to such a survey.
- A survey can be restored to an existing target survey only if that target survey includes at least one item (question) same as survey questions stored in the content database.
- If a survey question was not answered completely in the source survey, after restore the response status in the target survey will be set to *Completed*, anyway.
- When restoring a single response to a survey, target response item with the same number will be deleted, and restored item will be placed in the target survey after the last numbered response.
  For example, if the target survey has responses #1-15, and you try to restore a response that used to be #6 on source – then target response #6 will be deleted, and restored response will be assigned #16.

Data Type Limitations

Consider the column (field) data type when planning for the recovery of your SharePoint libraries/lists:

- If source column (field) data type was set to *Lookup*, but the referenced list/library was deleted, such columns (fields) will not be restored even if you recover that referenced list. The reason is that if referenced list is deleted, the reference (link) to that list is no longer valid.
- If source column (field) data type was set to *Managed Metadata*, such columns (fields) will not be restored.

Workflow-related Considerations

You may need to restore the item(s) originally belonging to a list with no content approval required, to another list. If you try to restore such item(s) to a list that requires content approval, item version and status will be modified in the following way:

a. If a target list is configured to include major versions only – then all versions of restored item will become *major* (despite the original versioning); item status will be set to *Pending*;

b. If a target list is configured to include both major and minor versions – then all versions of restored item will become *minor* (despite the original versioning); item status will be set as follows:

   - if the last version (original) was major – status will be set to *Pending*;
   - if the last version was minor – status will be set to *Draft*.

Also, consider the following when planning for the recovery of list items (with or without content approval originally required):

a. If the retention policy for target list/document library is configured to **Declare record** automatically, only the last version of the item will be restored to target list/library. Target retention policy settings will be applied to restored item (**Declare record**).
Besides, if **Require content approval for submitted items** was enabled for the original list, then after recovery item status will be set to *pending*.

b. Alternatively (with different retention policy settings), all versions of the original item will be restored to target list/library.
Besides, if **Require content approval for submitted items** was enabled for the original list, then after recovery item status in the content approval workflow will be also restored, except for the states listed (see "Status Recovery Limitations" above).

Working with Veeam Backup & Replication Utilities

Veeam Backup & Replication comes with the extract utility that has been designed for a specific purpose. In this section, you will find a detailed description of this utility and its functionality.

Extract Utility

Veeam Backup & Replication comes with an extract utility that can be used to recover VMs from a full backup file — VBK. The utility can be used as an independent tool on Linux and Windows computers as it does not require any interaction with Veeam Backup & Replication.

The extract utility can be helpful, for example, if it is written to the tape next to backup files. In this case, you get a possibility to recover VMs from backups at any moment of time even if backups are removed from Veeam Backup & Replication or the application is uninstalled at all.

The extract utility can be utilized via two interfaces:

- Graphic user interface
- Command-line interface working in the interactive or regular mode

The installation folder of Veeam Backup & Replication (by default: *%PROGRAMFILES%\Veeam\Backup and Replication\Backup*) contains three files for the extract utility:

- *Veeam.Backup.Extractor.exe* — the utility working via the graphic user interface, can be used on Microsoft Windows machines only
- *extract.exe* — the utility working via the command-line interface, version for Microsoft Windows
- *extract* — the utility working via the command-line interface, version for Linux

Important! The extract utility always restores a VM from the full backup file.

Using the Extract Utility via the GUI

To start the extract utility in this mode, perform the following steps:

1. Run the `Veeam.Backup.Extractor.exe` file from the installation folder of Veeam Backup & Replication. The extract utility will be started.
2. In the **VBK file** field, specify a path to the full backup file from which you want to restore VM(s).
3. In the **Extract folder** field, specify a path to the destination folder where the VM files should be extracted.
4. From the **Virtual machines** list, select VMs that should be extracted.
5. Click **Extract**. The VM files will be extracted to the specified folder.

Important! The extract utility working via the graphic user interface has the following limitations:

- The extract utility can be started on Microsoft Windows machines only.
- If you plan to start the extract utility on the machine other than the Veeam backup server, make sure that you copy the `Veeam.Backup.Extractor.exe` file together with the `extract.exe` file from the *%PROGRAMFILES%\Veeam\Backup and Replication\Backup* folder and store these files to the same folder on the destination machine. In the opposite case, the extract utility will fail to start.

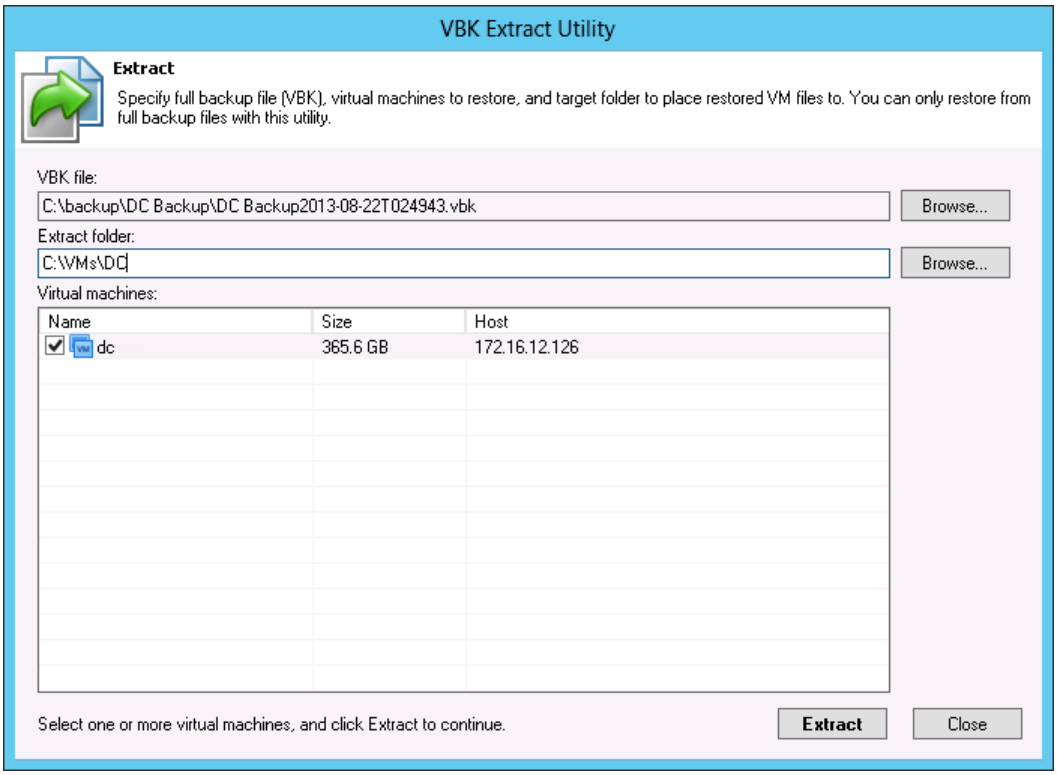

Using the Extract Utility in the Interactive Mode

To start the extract utility in the interactive mode, run the `extract.exe` file from the installation folder of Veeam Backup & Replication (in case of a Linux machine, run the `extract` file).

You will have to sequentially enter the following arguments:

1. Path to the full backup file from which the VM should be restored. After you enter the path, the restore utility will display a list of all VMs included in the backup and their description.
2. Name of the VM you want to restore. If there is more than one VM with the specified name in the backup, you will be asked to specify the host on which the backed up VM resides. If you want to restore all VMs from the backup, press **[ENTER]**.
3. Output directory to which VMs should be restored. If you want to restore VM(s) to the current directory, press **[ENTER]**.
4. Operation confirmation. Press **[Y]** on the keyboard to restore a VM to the directory you selected. If you want to abort the operation, press **[ENTER]**.

Using the Extract Utility from the Command Line

If you run the extract utility from the command line, you can perform the following actions:

- Run the extract utility in the interactive mode
- Display help information for the utility usage
- Display the list of all VMs in the backup file
- Restore all or selected VMs from the backup

Running the Extract Utility in the Interactive Mode

This command runs the extract utility in the interactive mode.

Syntax

```
extract.exe [pathtovbk]
```

Parameters

Parameter	Description	Required/Optional
pathtovbk	Path to the backup file from which VM(s) should be restored.	Optional

Displaying Help Information for the Utility Usage

This command prints all variants of the extract utility usage along with required and optional parameters.

Syntax

```
extract.exe -help
```

Displaying the List of VMs in the Backup

This command displays the list of all VMs in the backup file from which you want to perform restore.

Syntax

```
extract.exe -dir [-vm vmname] [-host hostname] pathtovbk
```

Parameters

Parameter	Description	Required/Optional
vm	Name of the VM that you want to restore. Use this parameter to filter VMs in the backup.	Optional
host	Name of the host on which the initial VM resides. This parameter should be specified if the `vm` parameter is used. Use this parameter to filter VMs that have the same name but reside on different hosts.	Optional
pathtovbk	Path to the backup file from which the VM should be restored.	Required

Restoring VMs from Backup

This command restores files for all VMs or for the selected VM(s) from the backup file.

Syntax

```
extract.exe -restore [-vm vmname] [-host hostname] pathtovbk
[outputdir]
```

Parameters

Parameter	Description	Required/Optional
vm	Name of a VM that you want to restore. Use this parameter to filter VMs in the backup. If you want to restore all VMs from the backup file, do not specify this parameter.	Optional
host	Name of the host on which the initial VM resides. This parameter should be specified if the `vm` parameter is used. Use this parameter to filter VMs that have the same name but reside on different hosts.	Optional
pathtovbk	Path to the backup file from which the VM should be restored.	Required
outputdir	Path to the directory to which the VM files should be restored. If this parameter is not specified, the VM will be restored to the current directory.	Optional