



VBR Core Concepts & Best Practices



Speaker
Title



Speaker
Title

Agenda

1. Introduction
2. Architecture & Core Components
3. Advanced Features & Optional Components
4. General Product Configuration
5. Jobs Concept & Configuration
6. Restore Capabilities
7. Security
8. Tips & Tricks
9. Additional Resources

Section 01

Introduction

Core Capabilities, Supported Platforms, Licensing

Veeam Data Platform

Recovery Orchestration

Monitoring & Analytics

Backup & Recovery

Native APIs

Platform
Extensions

aws AWS

Azure

Google Cloud

Kubernetes



Cloud



Virtual



Physical



Apps



SaaS


Microsoft 365

Salesforce

On-Premises • In the Cloud • XaaS

Supported Platforms

by Veeam Backup & Replication v12.1



Agentless (virtual)

- VMware vSphere 6.x – 8.0U2
- VMware Cloud Director 10.1 – 10.5.x
- Microsoft Windows Server Hyper-V 2012 – 2022
- Azure Stack HCI
 - *Microsoft Hyper-V Server (free hypervisor) is supported
 - *Server Core installations are fully supported
- Red Hat Virtualization 4.4 SP1
- Nutanix AOS 6.5 and later

Supported Platforms

by Veeam Backup & Replication v12.1



Agents (physical or virtual workloads)

- Microsoft Windows Server 2008 R2 SP1 – 2022
- Microsoft Windows 7 SP1 – 11
 - *Server Core installations are supported
- Majority of popular Linux Distros ([full list](#))
- Mac OS 10.13.6 High Sierra – 14 Sonoma
- IBM AIX 7.1 – 7.3 TL1
- Oracle Solaris 10 1/13, 11.0 – 11.4
 - *SPARK and x86 are supported

Supported Platforms

by Veeam Backup & Replication v12.1



Unstructured data

- File Shares (SMB and NFS)
*Including Enterprise NAS systems: NetApp Data ONTAP, Lenovo ThinkSystem DM Series, Dell PowerScale (formerly Isilon), Nutanix Files Storage
- Object storage repositories:
 - S3 compatible object storage repositories
 - Amazon S3 object storage
 - Microsoft Azure Blob storage

Licensing

Two **models** of licensing

Per-Socket

Legacy model. Still available for existing customers.

Veeam Backup & Replication is licensed by the number of CPU sockets on the protected hosts.

License is required only for source hosts — hosts on which VMs that you back up or replicate reside. Target hosts (for replication and migration jobs) do not need to be licensed.

Instance-based (VUL)

New model. It is more flexible and can be utilized across different products.

Instances are units (or tokens) that you can use to protect your virtual, physical or cloud-based workloads.

You must obtain a license with the total number of instances for workloads that you plan to protect in Veeam Backup & Replication.

Licensing

Two **types** of licensing

Perpetual

Permanent license. Support and maintenance period is included with the license and is specified in months or years.

Perpetual License includes Production or Basic Support and Maintenance agreement for the first year.

Maximum prepaid term for Support and Maintenance for a Perpetual License is three years.

Subscription

License that expires at the end of the subscription term.

Subscription License includes Production Support and Maintenance agreement for the full term of the license.

The Subscription license term is normally 1–3 years from the date of license issue.

Maximum prepaid term for a Subscription License is five years.

Section 02

Architecture & Core Components

Components Overview, Installation Process Review, Backup Workflow, In-depth Components Breakdown, Design

Components Overview

Architecture & Core Components

Overview



VBR
Console



Backup Proxy



Backup Server



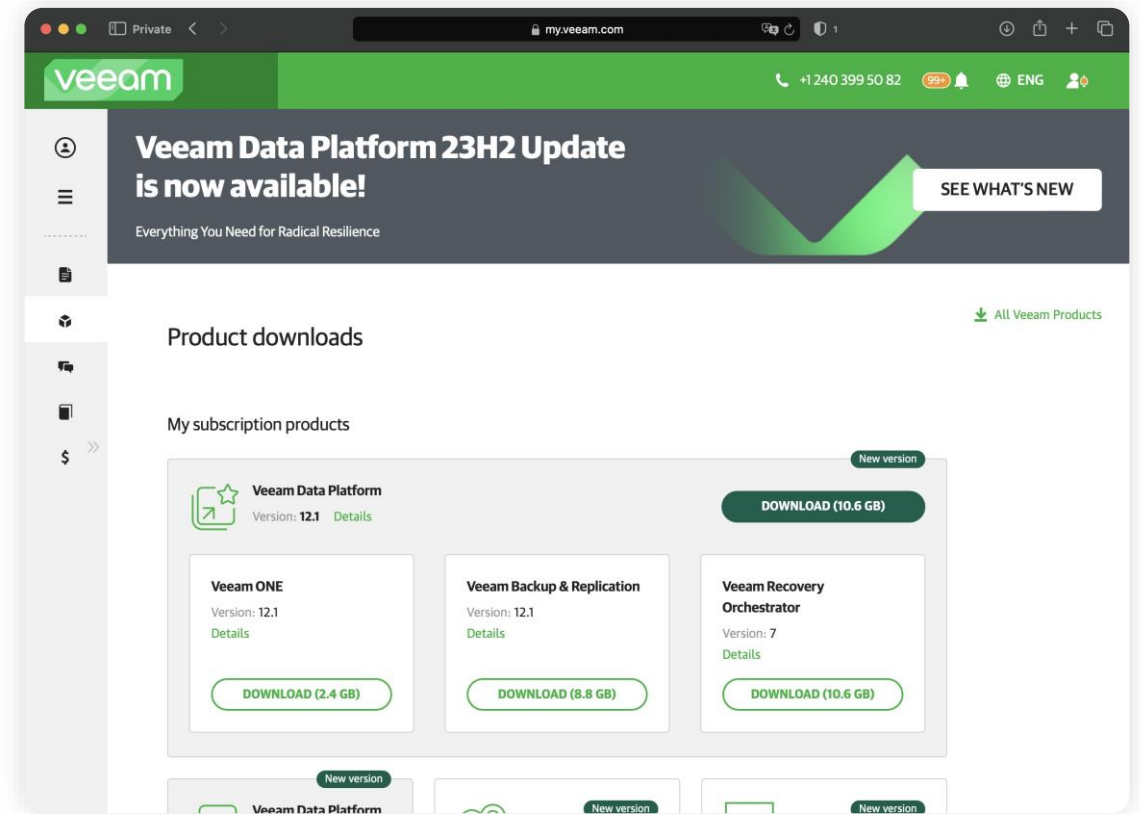

Backup Repository

Installation

Architecture & Core Components

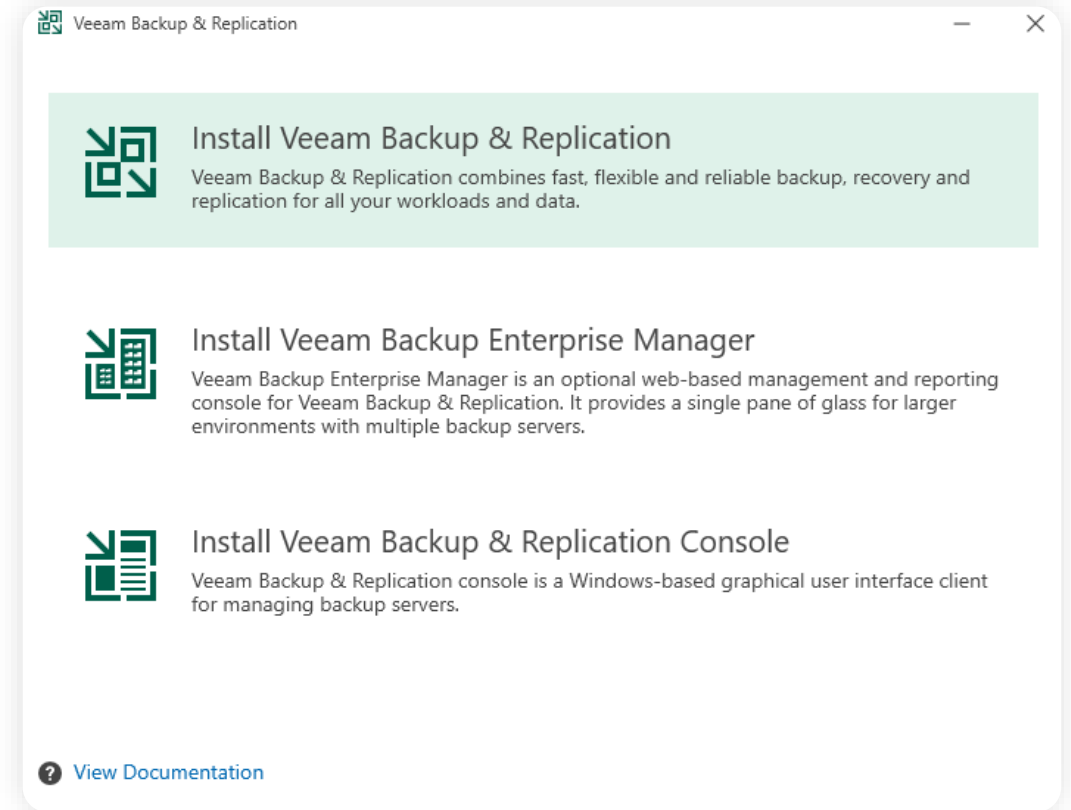
Installation

Download the most recent version of .ISO file from <https://my.veeam.com/my-products>. Sign in if necessary.



Architecture & Core Components

Installation



Architecture & Core Components



Installation

Veeam Backup & Replication

License

Provide license file for Veeam Backup & Replication.

Select license provisioning method:

 [Sign in with Veeam](#) |  [Browse license file](#)

License details:
Subscription, Suite, 60 Instances, License expires on 03/08/2024


Update license automatically (enables usage reporting)









Download and install new license automatically when you renew or expand your contract. This requires sending the license ID, the installation ID, and workload usage counters to Veeam servers periodically. Successful usage reporting doubles the number of workloads you can exceed your installed license by.

Veeam Backup & Replication

System Configuration Check

System is being verified for potential installation problems.

 Setup could not automatically install required system prerequisites. Please install missing components and click "Retry" to continue.

Requirement	Status
Microsoft .NET Framework 4.7.2	 Passed
Microsoft Visual C++ 2015-2019 Redistributable	 Passed
Microsoft System CLR Types for SQL Server 2014	 Failed
Microsoft Report Viewer Redistributable 2015	 Passed
Microsoft PowerShell v5.1	 Passed
Microsoft Universal C Runtime	 Passed
Microsoft .NET Runtime 6.0.10	 Passed
Microsoft ASP.NET Core Shared Framework 6.0.10	 Passed

Architecture & Core Components

Installation

Veeam Backup & Replication

Ready to Install

Installation will begin with the following settings.

Installation folder:	C:\Program Files\Veeam\Backup and Replication
vPower cache folder:	C:\ProgramData\Veeam\Backup\IRCach
Guest catalog folder:	C:\VBRCatalog
Service account:	LOCAL SYSTEM
Database engine:	PostgreSQL
SQL Server:	repo32:5432
Database name:	VeeamBackup
Catalog service port:	9393
Service port:	9392
Secure connections port:	9401
REST API Service Port:	9419
Check for updates:	Automatically

Customize Settings

Back Install Cancel

Veeam Backup & Replication

Database

Choose database engine and instance for Veeam Backup & Replication.

Use following database engine: PostgreSQL

Install new instance

Use existing instance (HOSTNAME:PORT)

repo32:5432

Database name: VeeamBackup

Connect to PostgreSQL server using:

Windows authentication credentials of service account

Native authentication using the following credentials:

Username: postgres

Password:

Back Next Cancel

Configuration database

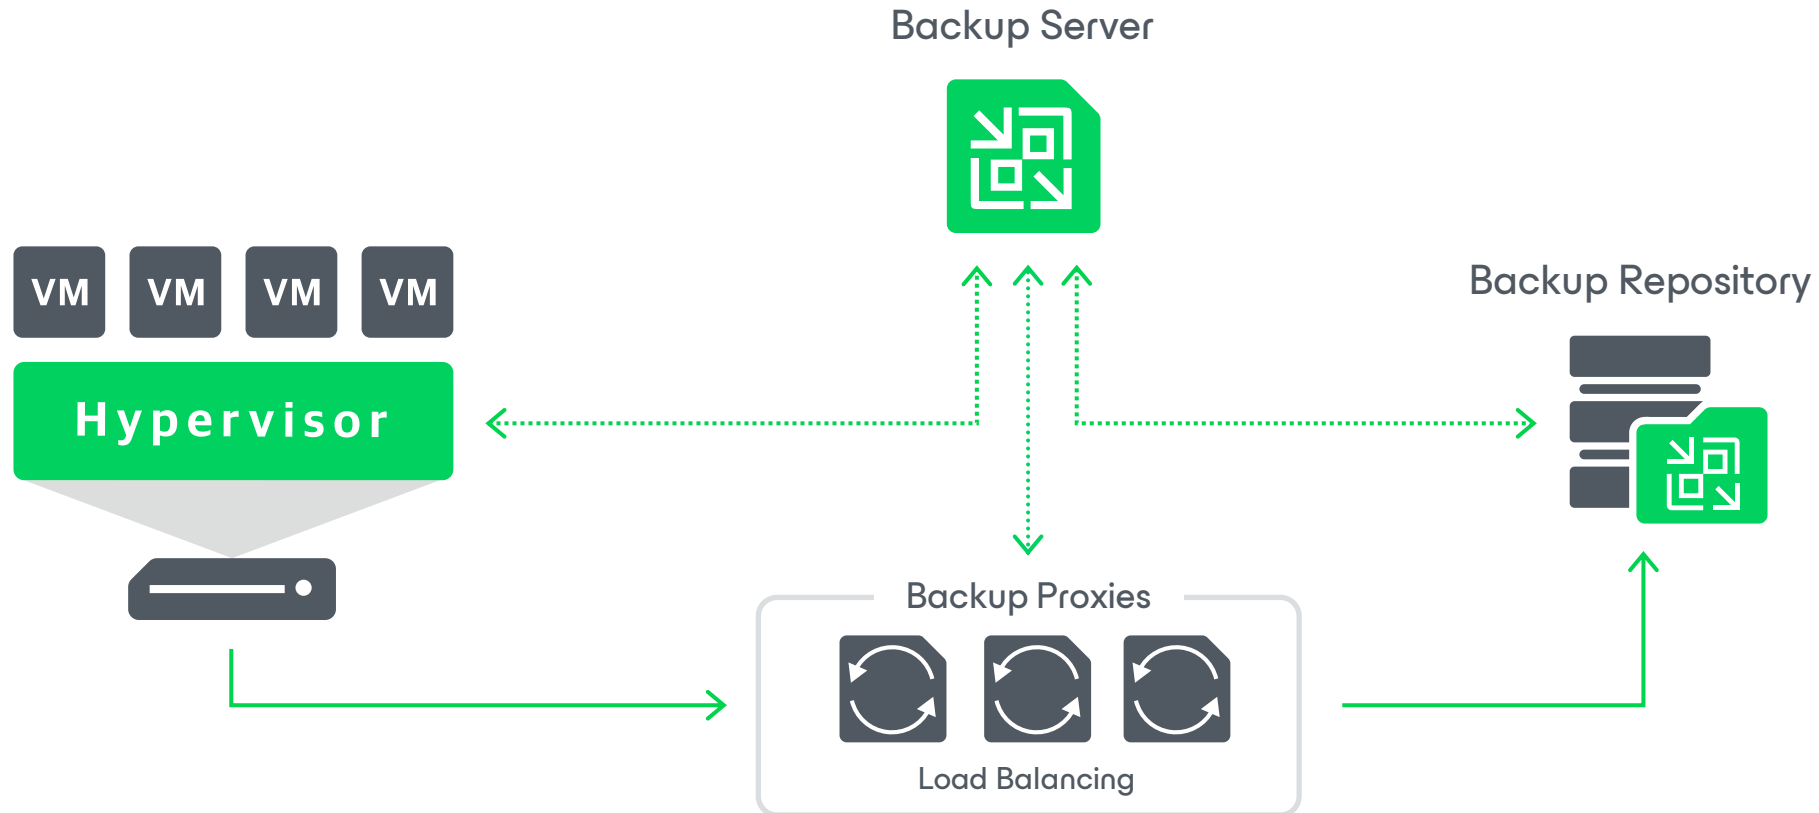
Configuration database

- Stores data about the backup infrastructure, jobs, sessions.
- **Since v.12**, VBR installer by default installs **PostgreSQL**. One can change it to Microsoft SQL.
- DB can be migrated from MSSQL to PostgreSQL and vice versa.
- DB can be installed either **locally** (on the same machine with the backup server) **or remotely**.
- Configuration DB backup is the way how VBR 'backs itself up.'
- DB backup can be restored to another VBR server.

In-depth Components Breakdown

Architecture & Core Components

Backup Workflow Example



→ Backup Traffic
↔ Management Traffic

Backup Proxy

Architecture & Core Components

Backup Proxy

Types

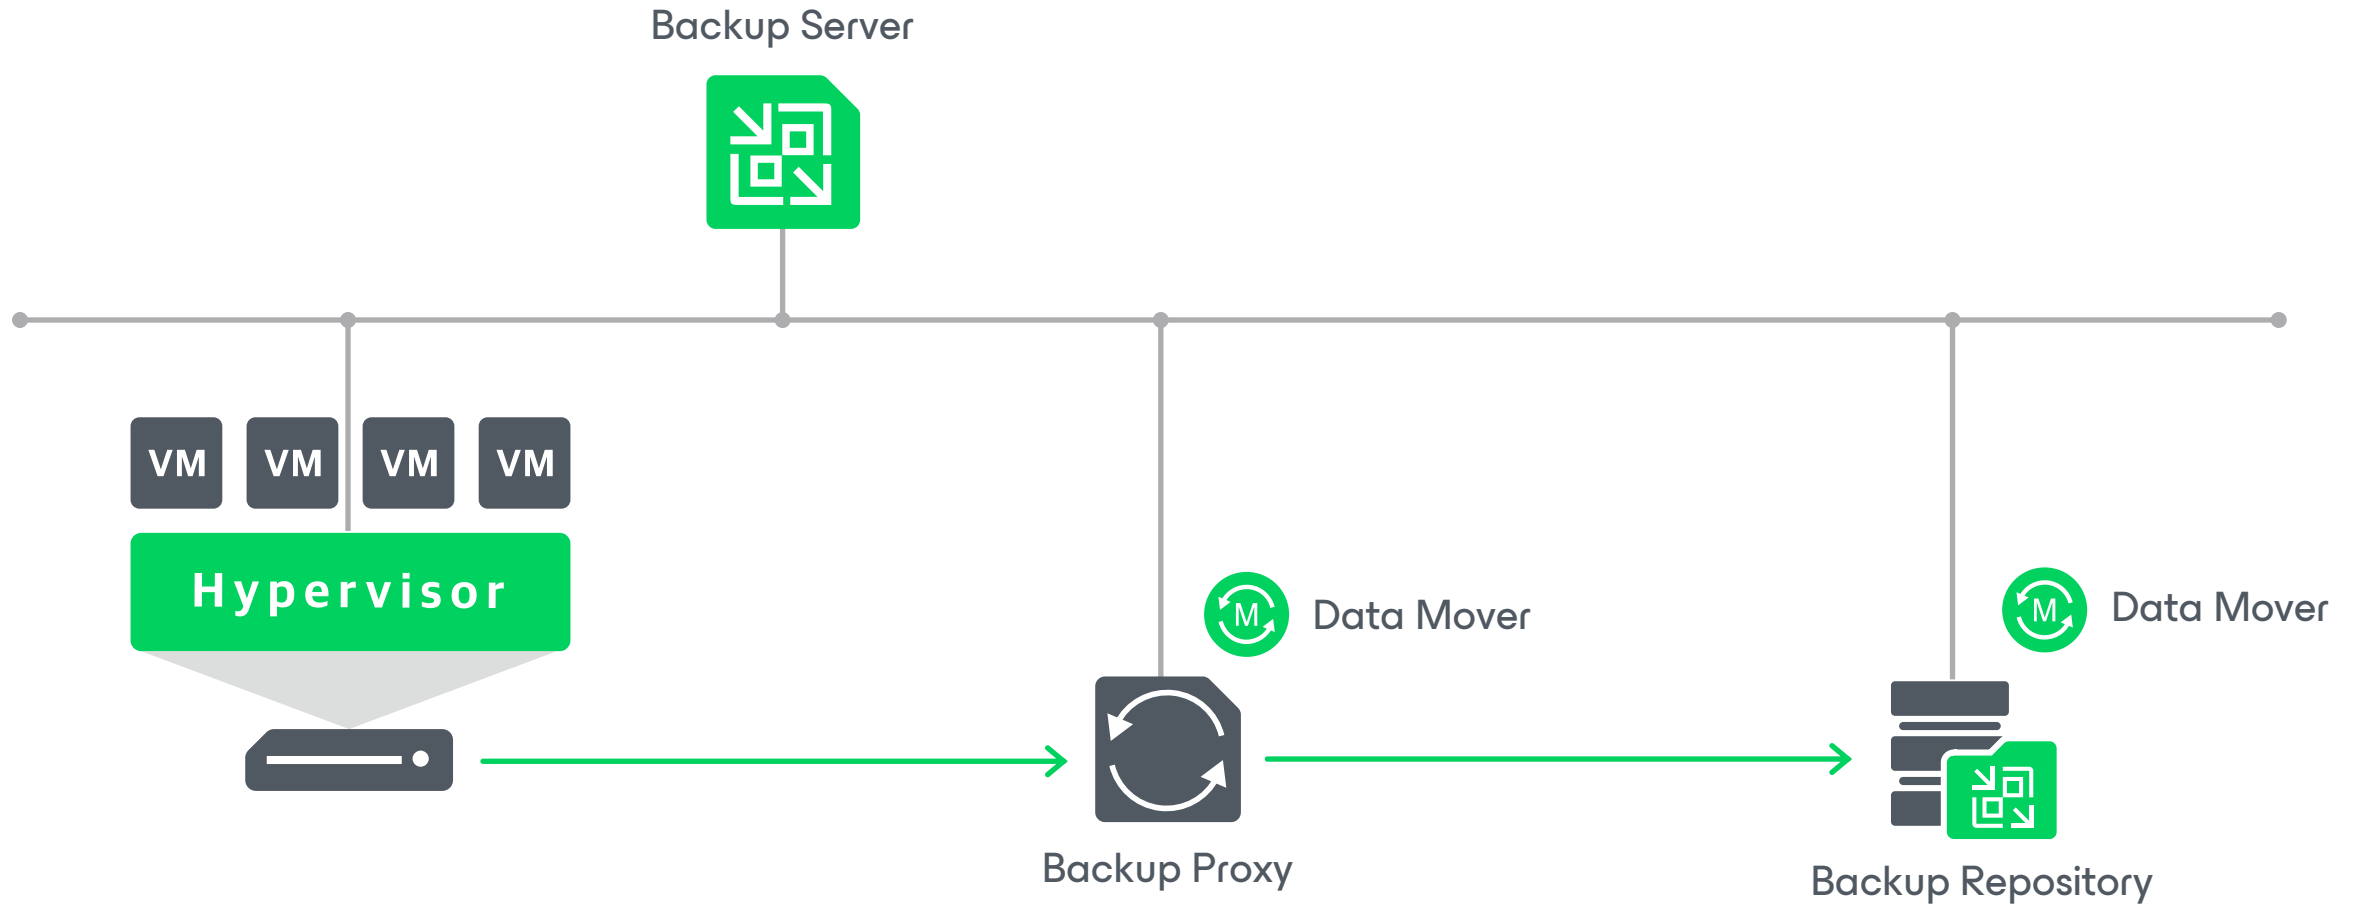
- VMware
- VMware CDP
- Hyper-V On-host
- Hyper-V Off-host
- Nutanix AHV (Appliance)
- Red Hat Virtualization (Appliance)
- General Purpose (NAS, File shares, Physical Servers Off-host backup)

Responsibilities

- Retrieving VM data from the production storage
- Compressing
- Deduplicating
- Encrypting
- Sending it to the backup repository (backup) or another backup proxy (replication)

Architecture & Core Components

Backup Proxy



Backup Proxy (VMware)

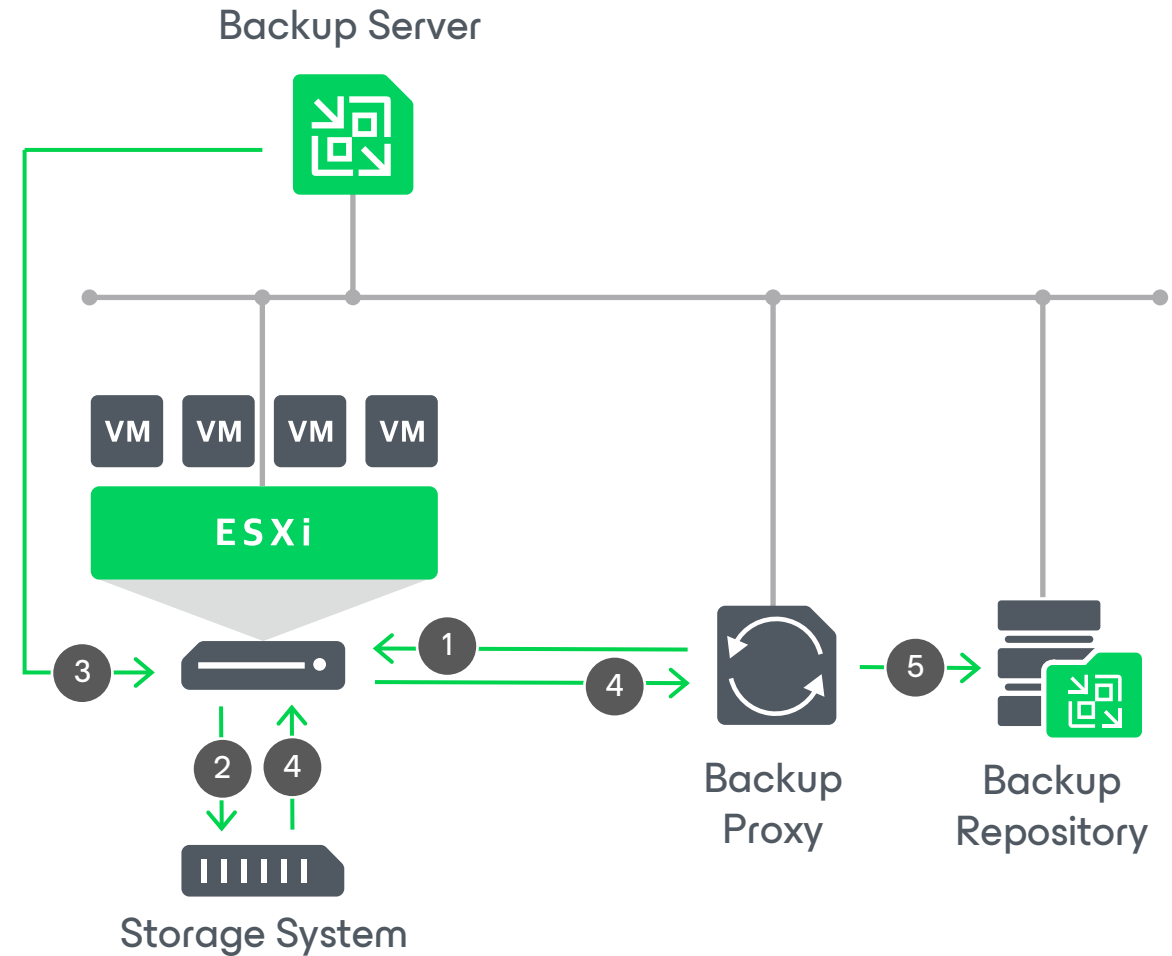
Architecture & Core Components

VMware Backup Proxy Transport Modes

Network Mode

In this mode, data is retrieved through the ESXi host over LAN using the Network Block Device protocol (NBD).

1. The Proxy sends a request to ESXi host to locate the VM on the datastore.
2. The ESXi host locates the VM.
3. VBR instructs VMware vSphere to create a snapshot.
4. The ESXi host copies VM data blocks from the source storage and sends them to the proxy over LAN.
5. The Proxy sends the data to the Repository.



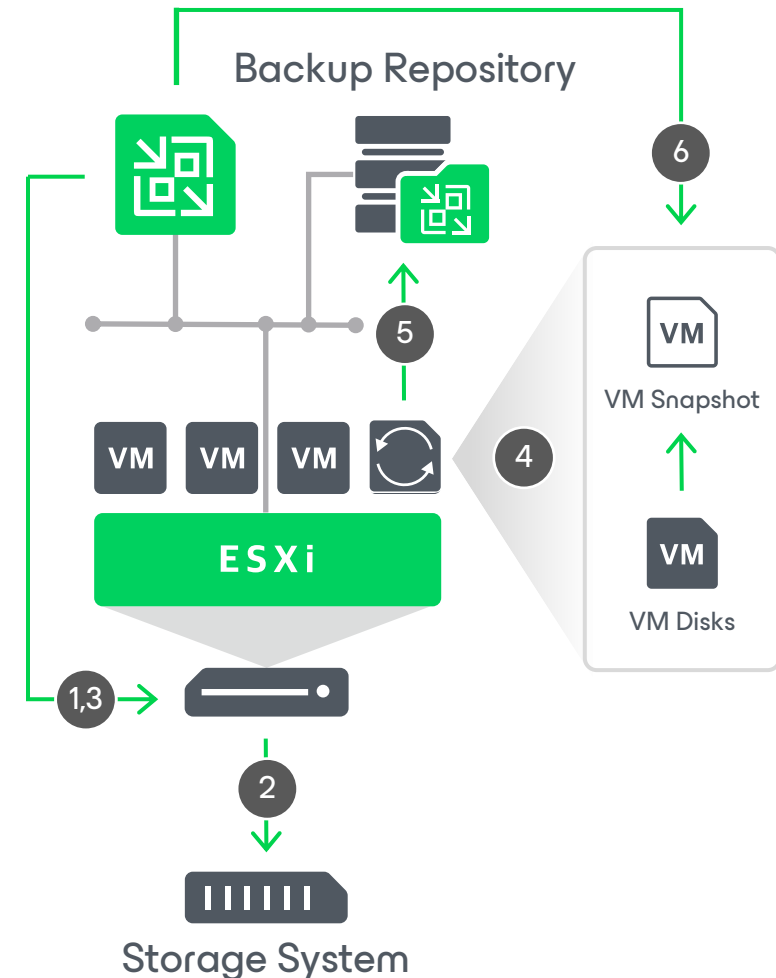
Architecture & Core Components

VMware Backup Proxy Transport Modes

Virtual Appliance

This mode provides better performance than the Network mode.

1. The backup server sends a request to the ESXi host to locate the VM on the datastore.
2. The ESXi host locates the VM.
3. VBR triggers VMware vSphere to create a VM snapshot.
4. VM disks are attached (hot-added) to the Proxy.
5. VBR reads data directly from disks attached to the Proxy.
6. When the VM processing is complete, VM disks are detached from the Proxy and the VM snapshot is deleted.



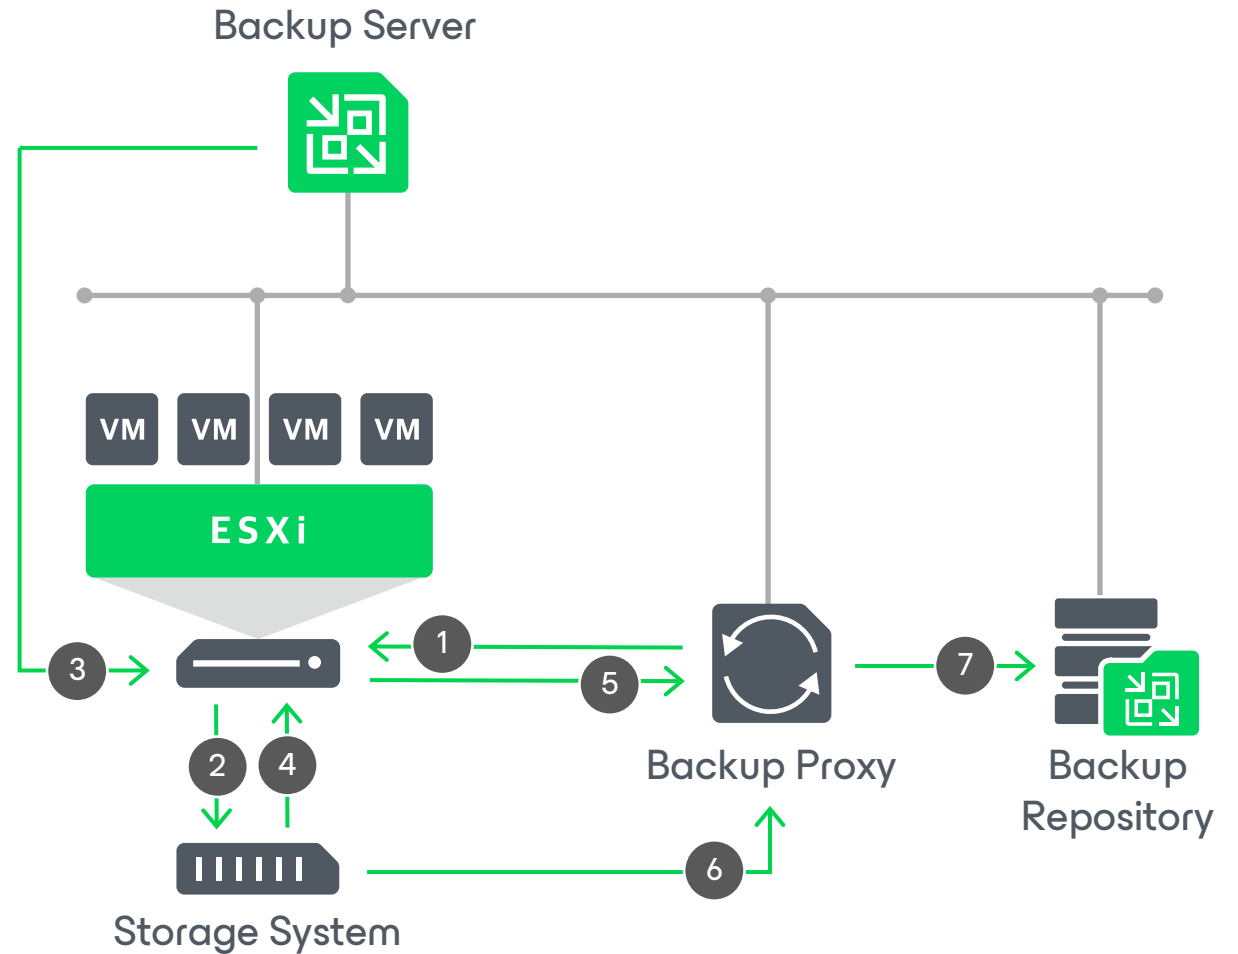
Architecture & Core Components

VMware Backup Proxy Transport Modes

Direct Storage Access

In this mode, VBR uses VMware VADP for direct VM data transport over SAN (FC, FCoE, iSCSI).

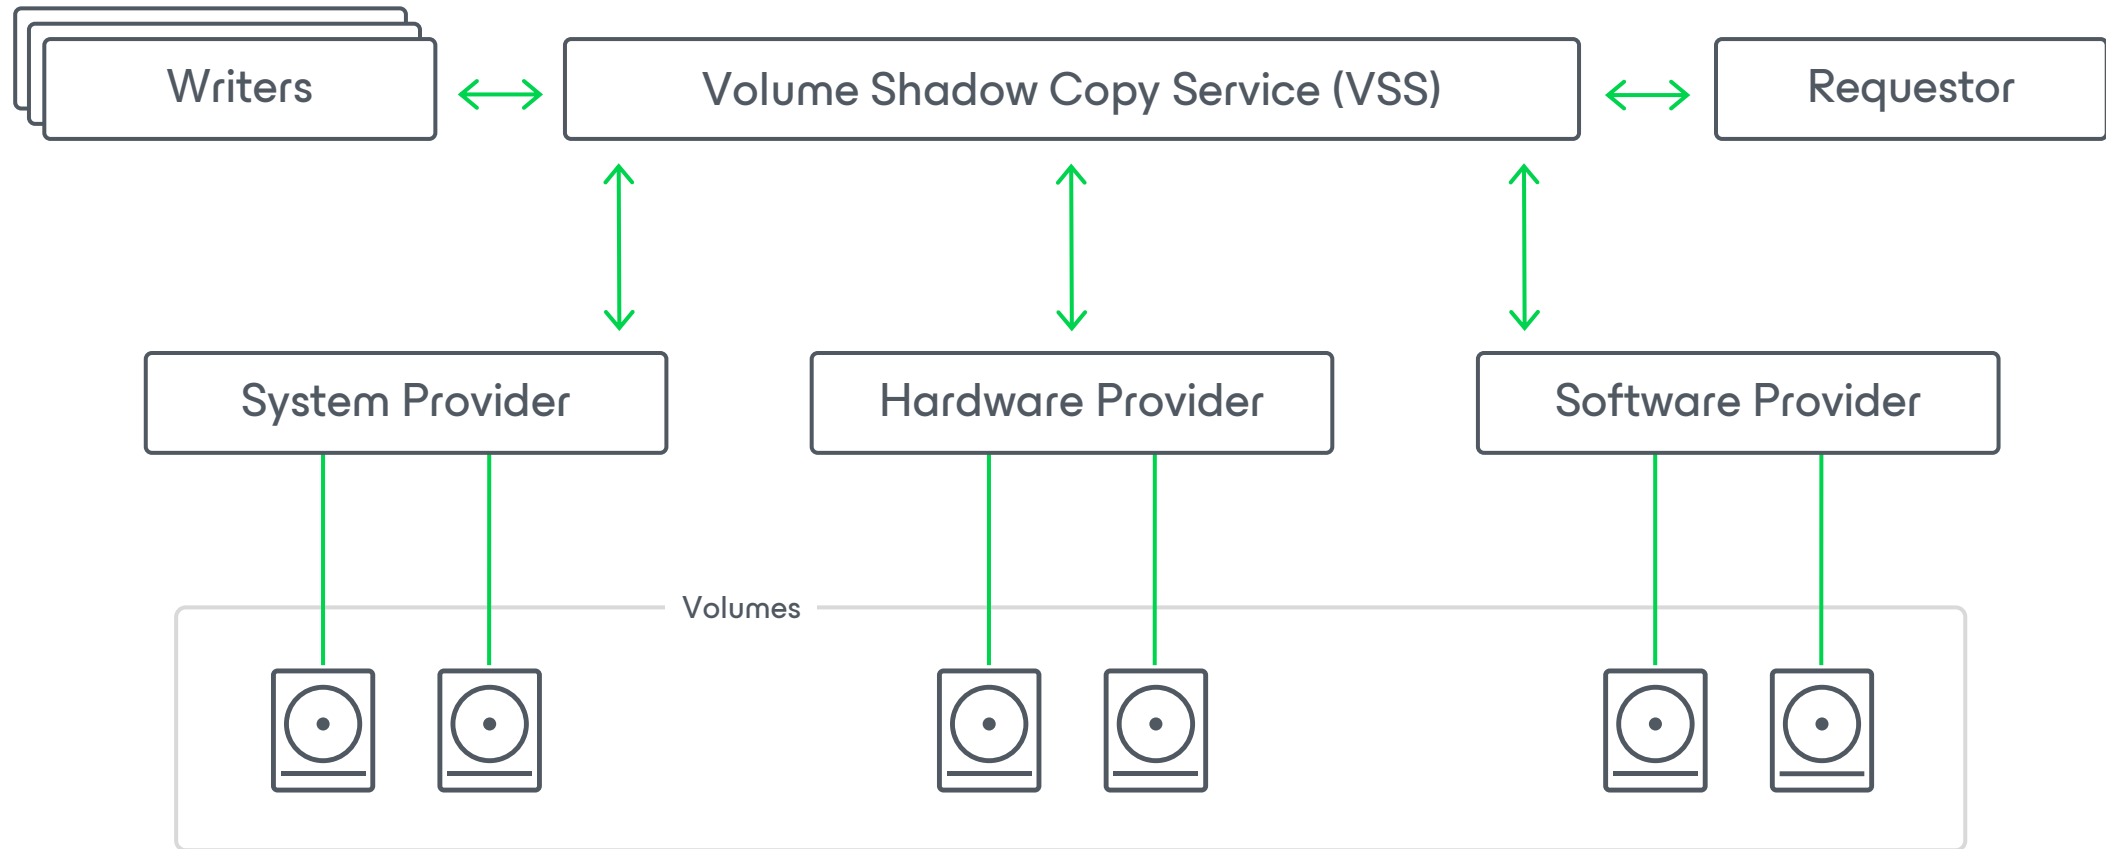
1. The Proxy requests VM location from ESXi.
2. The ESXi host locates the VM.
3. The VBR triggers a VM snapshot.
4. The ESXi host retrieves metadata about the layout of VM disks on the storage.
5. The ESXi host sends metadata to the Proxy.
6. The Proxy uses metadata to copy VM data blocks directly from the source storage.
7. The Proxy processes copied data blocks and sends them to the target.



Backup Proxy (Hyper-V)

Architecture & Core Components

Volume Shadow Copy Service (VSS)



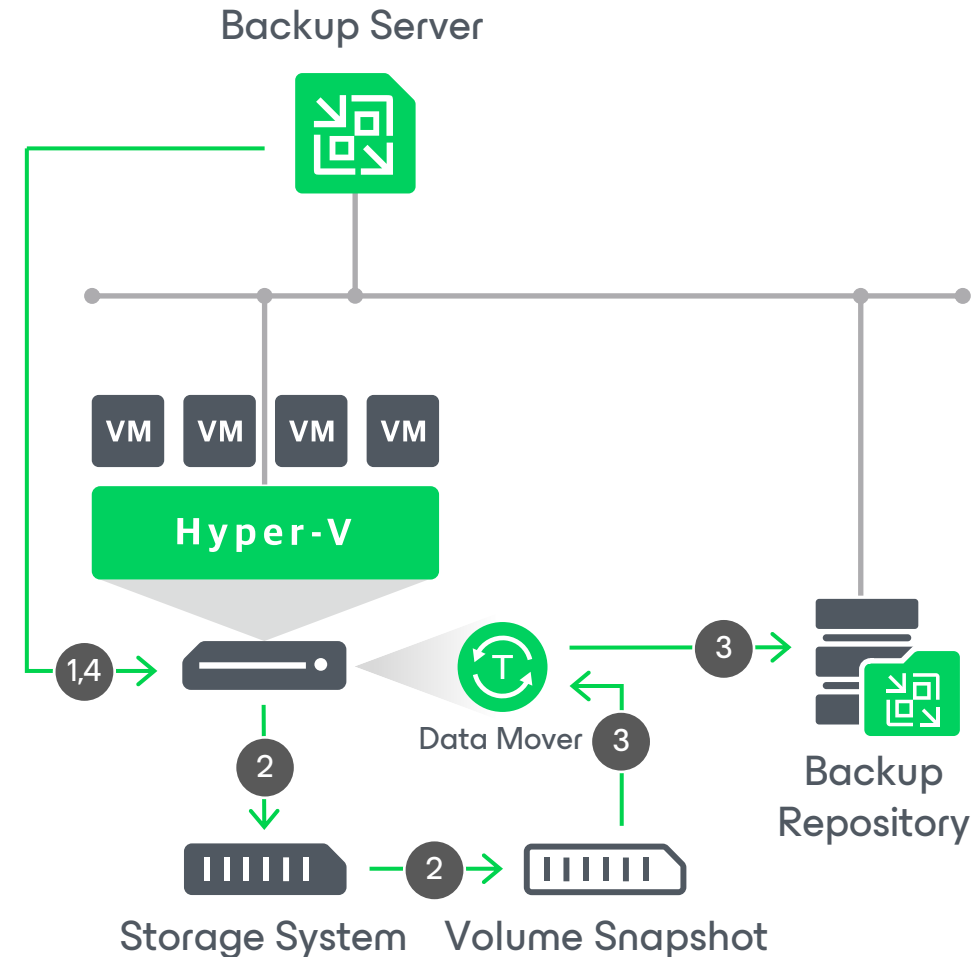
Architecture & Core Components

Hyper-V Backup Modes

On-Host

The processing of backup data occurs on the Hyper-V node hosting the Virtual Machines, utilizing non-transportable shadow copies by using software VSS provider.

1. VBR queries VM and virtualization host details from Microsoft Hyper-V.
2. VBR instructs Microsoft Hyper-V VSS to create a point-in-time VM copy, using a volume snapshot or VM checkpoint.
3. Veeam Transport Service reads VM data from the snapshot or checkpoint and transfers it to the backup repository.
4. After data transfer, VBR prompts Microsoft Hyper-V VSS for clean-up operations.



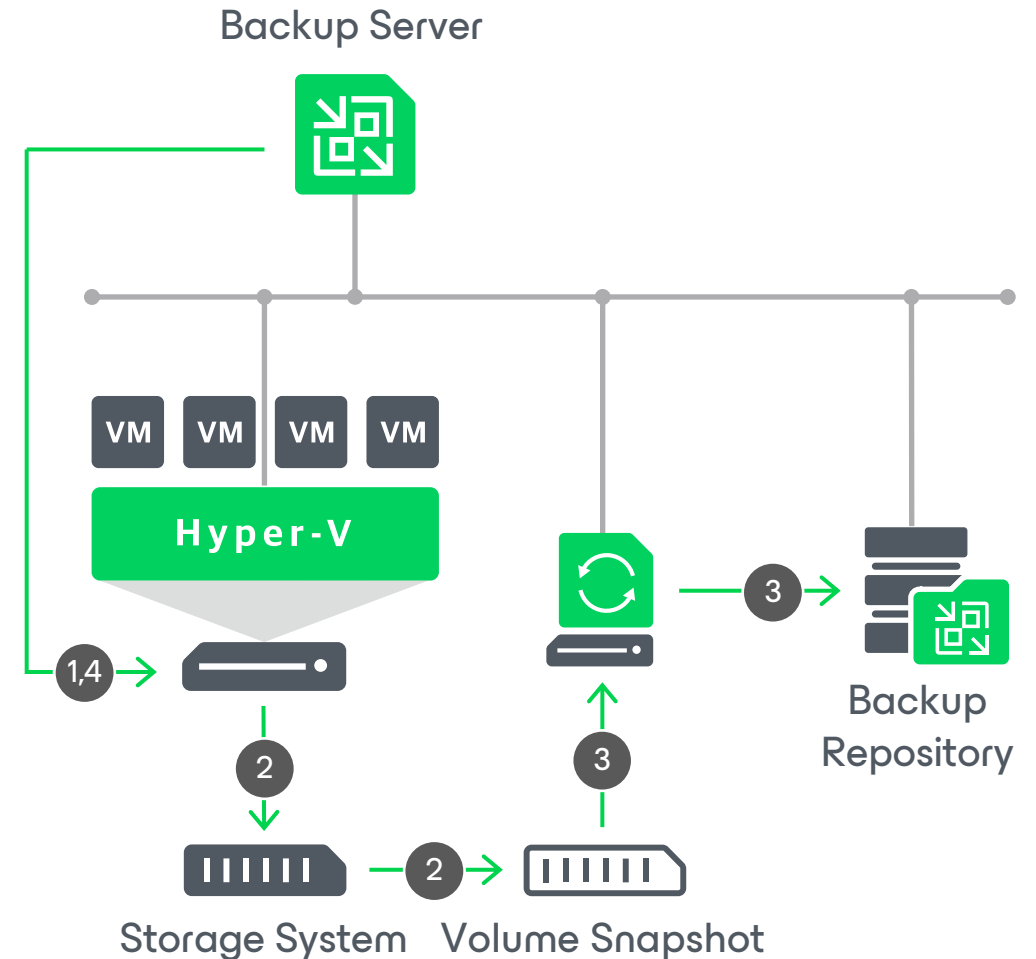
Architecture & Core Components

Hyper-V Backup Modes

Off-Host

Data processing for backups is delegated to a separate, non-clustered Hyper-V node, utilizing transportable shadow copies through the Hardware VSS provider offered by the SAN storage vendor.

1. VBR queries VM and virtualization host details from Microsoft Hyper-V.
2. VBR instructs Microsoft Hyper-V VSS to create a point-in-time VM copy, using a volume snapshot or VM checkpoint.
3. Veeam Off-host Proxy reads VM data from the snapshot or checkpoint and transfers it to the backup repository.
4. After data transfer, VBR prompts Microsoft Hyper-V VSS for clean-up operations.



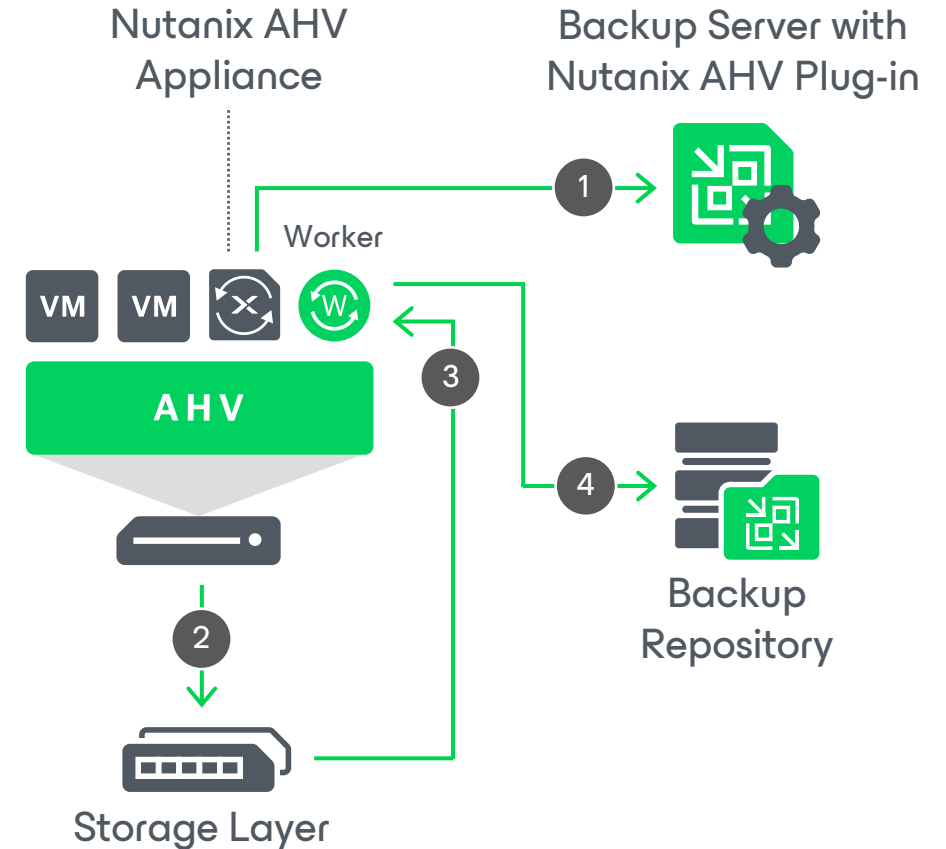
Backup ~~Proxy~~ Appliance (AHV & RHV)

Architecture & Core Components

Nutanix AHV

Veeam Backup for Nutanix AHV uses Nutanix AHV's native capabilities to create image-level backups and doesn't install agent software inside VMs to retrieve data.

1. The Nutanix AHV backup appliance initiates a backup job, transmitting session data to VBR.
2. Utilizing Nutanix REST API, it links with the Nutanix AHV cluster, generating snapshots for all designated VMs or protection domains. After, it establishes a volume group on the Nutanix AHV cluster.
3. The Nutanix AHV backup appliance deploys a worker, mounting VM disks to the worker via iSCSI.
4. The worker retrieves VM data at the block level, applying compression and deduplication, and transmits the data to the backup repository.

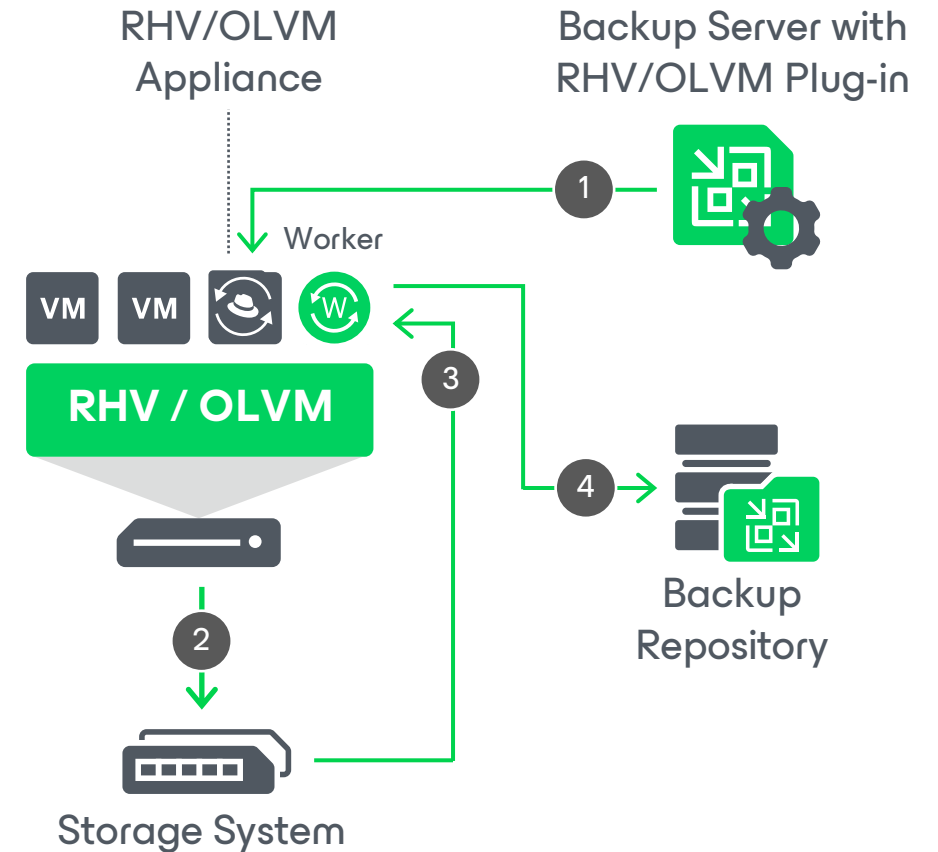


Architecture & Core Components

Red Hat Virtualization & Oracle Linux Virtualization Manager (oVirt)

Veeam Backup for RHV/OLVM uses native oVirt capabilities to create image-level backups and doesn't install agent software inside VMs to retrieve data.

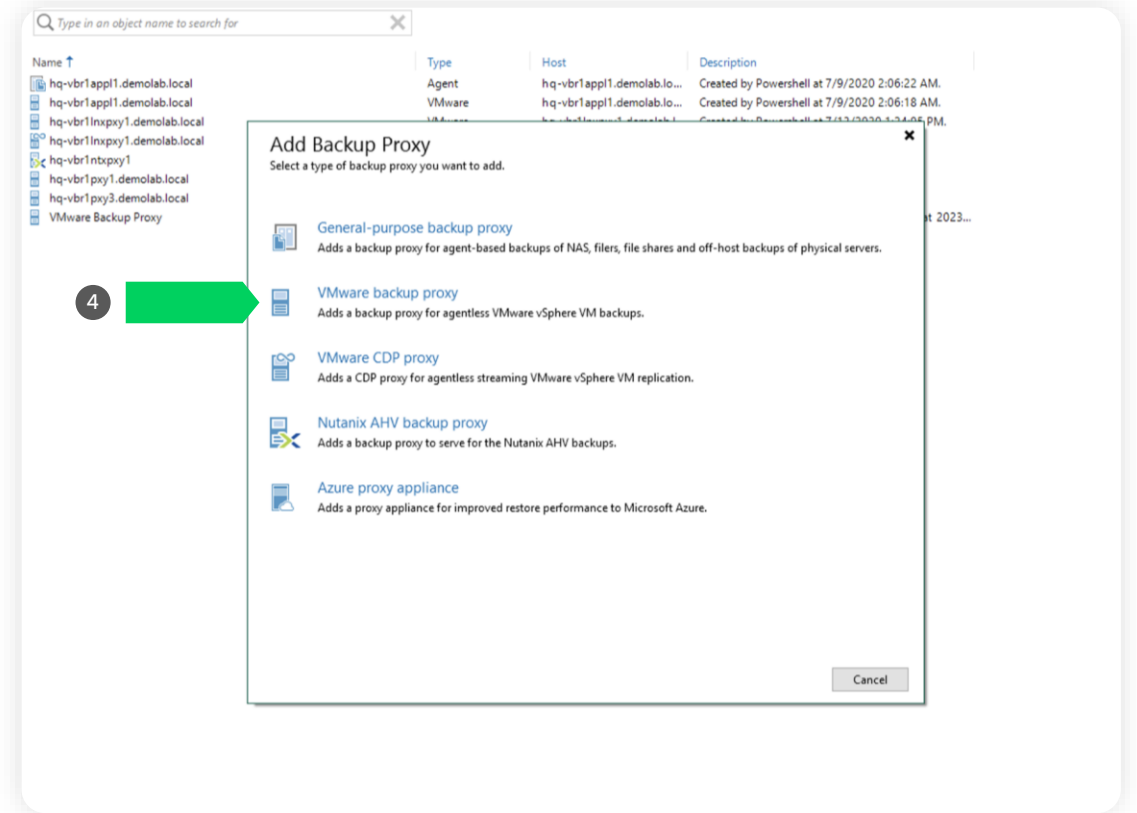
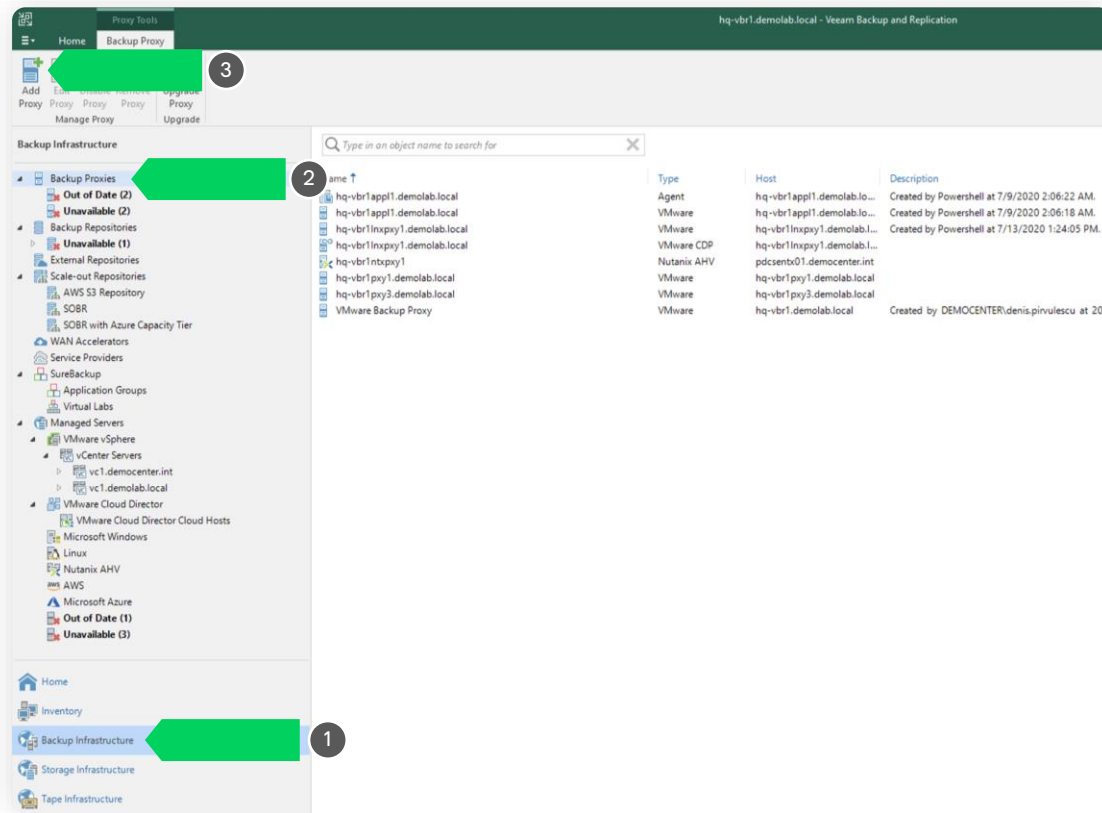
1. VBR initiates a backup job, transmitting session data to the backup appliance.
2. The backup appliance interfaces with the virtualization manager via REST API, generating snapshots for designated VMs. It subsequently establishes an image transfer session, supplying the corresponding URL.
3. The backup appliance deploys a worker, utilizing the provided URL to retrieve VM data.
4. The worker applies compression and deduplication to the VM data before transmitting it to the designated backup repository.



How to add a Backup Proxy?

Architecture & Core Components

How to add a Proxy? VMware example.



Architecture & Core Components

How to add a Proxy? VMware example.

New VMware Proxy

Server
Choose a server for VMware backup proxy. You can choose between any Microsoft Windows or Linux servers added to the Managed Servers which are not assigned a VMware backup proxy role already.

Server
Choose server: hq-vbr1.demolab.local (Backup server) Add New...

Proxy description:

Transport mode: Automatic selection Choose...

Connected datastores: Automatic detection (recommended) Choose...

Max concurrent tasks: 16 ✓

< Previous Next > Finish Cancel

Transport Mode

Veeam Backup & Replication creates one task per every VM disk.

Automatic selection
Data retrieval mode is selected automatically by analyzing backup proxy configuration and reachable VMFS and NFS datastores. Transport modes allowing for direct storage access will be used whenever possible.

Direct storage access
Data is retrieved directly from shared storage, without impacting production hosts. For block storage, backup proxy server must be connected into SAN fabric via hardware or software HBA, and have VMFS volumes mounted.

Virtual appliance
Data is retrieved directly from storage through hypervisor I/O stack by hot adding backed up virtual disks to a backup proxy VM. Datastores containing protected VMs must be connected to a host running backup proxy VM.

Network
Data is retrieved from storage through hypervisor network stack using NBD protocol over host management interface. This mode has no special setup requirements. Recommended for 10 Gb Ethernet or faster.

Options:
 Failover to network mode if primary mode fails, or is unavailable

OK Cancel

For example, for a 4-core CPU, it is recommended that you specify maximum 8 concurrent tasks, for an 8-core CPU - 16 concurrent tasks.

Backup Repository

Architecture & Core Components

Backup Repository

A repository is a storage location where backup files, copies of VM data, and metadata necessary for restoration are stored.

The repository can be a local or remote location, and it serves as a centralized storage space for backup and replication data.

Direct Attached Storage

- Microsoft Windows server
- Linux server
- Hardened Repository

Network Attached Storage

- SMB (CIFS) share
- NFS share

Deduplicating Storage Appliances

- Dell Data Domain
- ExaGrid
- Fujitsu ETERNUS CS800
- HPE StoreOnce
- Infinidat InfiniGuard
- Quantum Dxi

Object Storage

- S3 Compatible
- Amazon S3
- Google Cloud Storage
- IBM Cloud Object Storage
- Microsoft Azure Storage
- Wasabi Cloud Storage

Architecture & Core Components

Backup Repository

There are many file types you can find in the Backup Repository. Here is the list of the most used:

.vbm – backup chain metadata file

.vbk – full backup file

.vib – incremental backup file

.vrb – reverse incremental backup file

.bco – configuration backup file

.vsb – virtual full backup for tapes

.vlb – archived log backup file

.vsm – Microsoft SQL transaction log backup

.vom – Oracle database log file

.vpm – PostgreSQL transaction log backup

.vab – enterprise plug-ins backup file

.vasm – enterprise plug-ins backup metadata file

.vacm – enterprise plug-ins backup job metadata file


.vblob – file that stores data from the file share (NAS) backup job

.vindex – binary metadata file that describes backup files (names and versions) of file share backup job

.vslice – binary metadata that describes allocation of data in VBLOB backup files

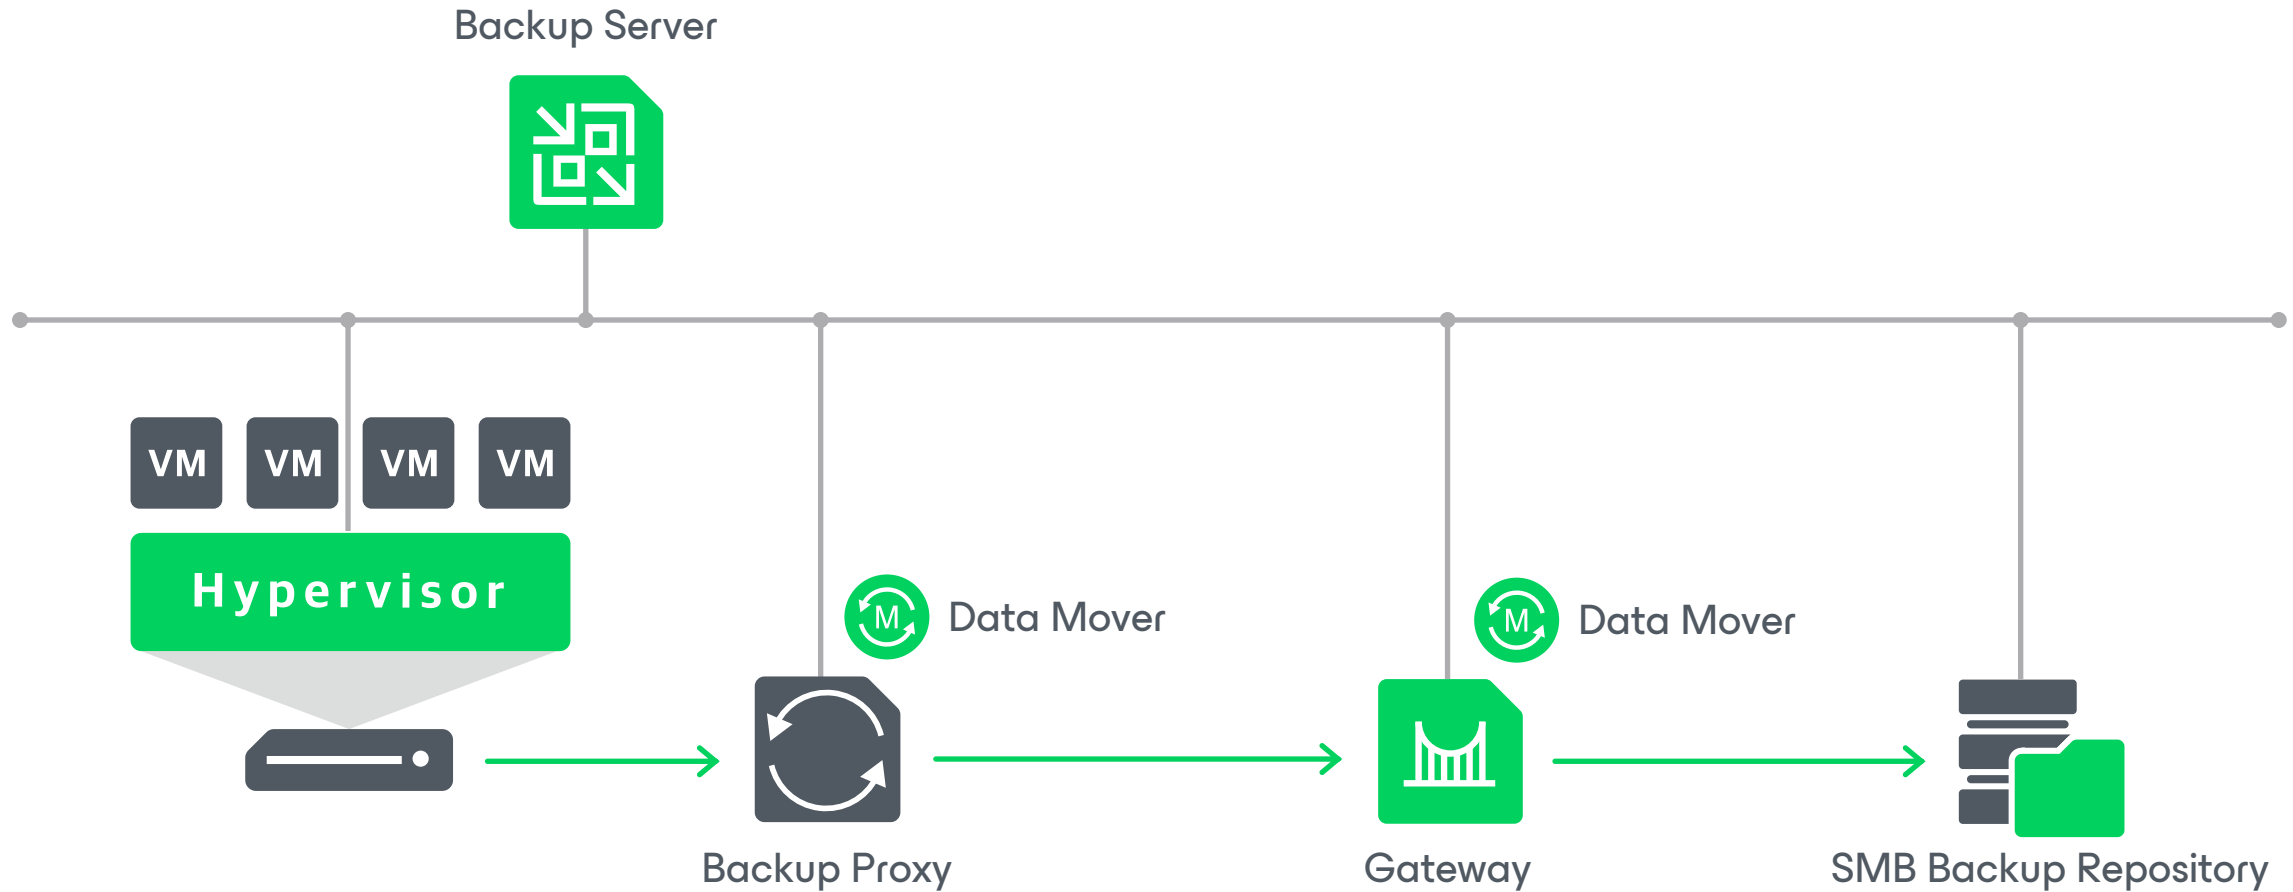
Architecture & Core Components

Windows Backup Repository Example



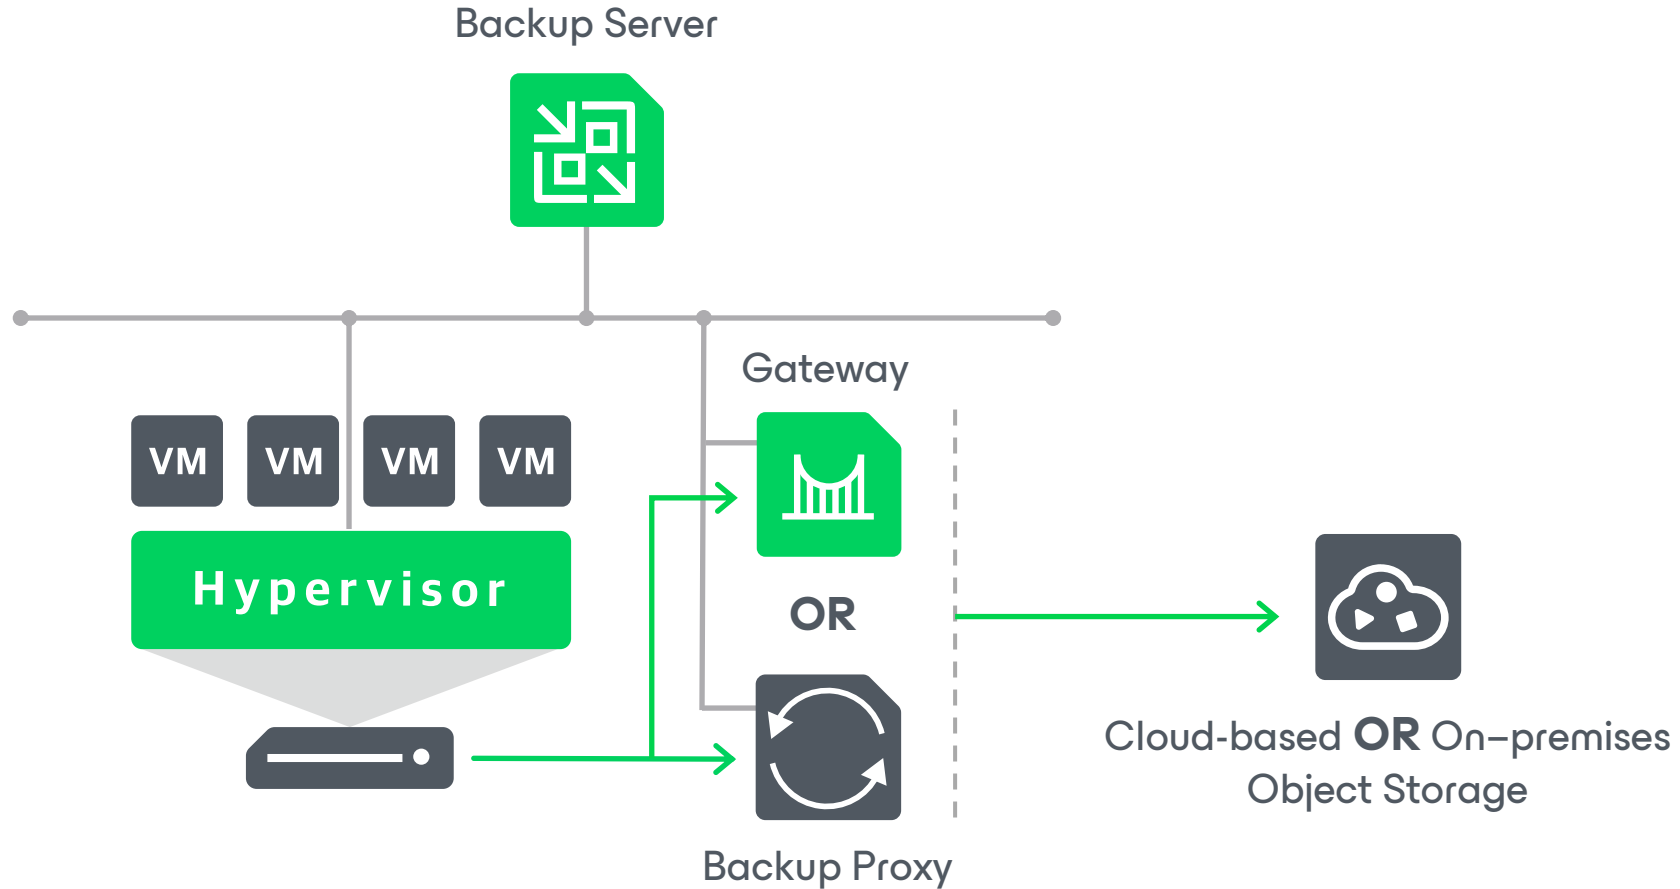
Architecture & Core Components

SMB Backup Repository Example



Architecture & Core Components

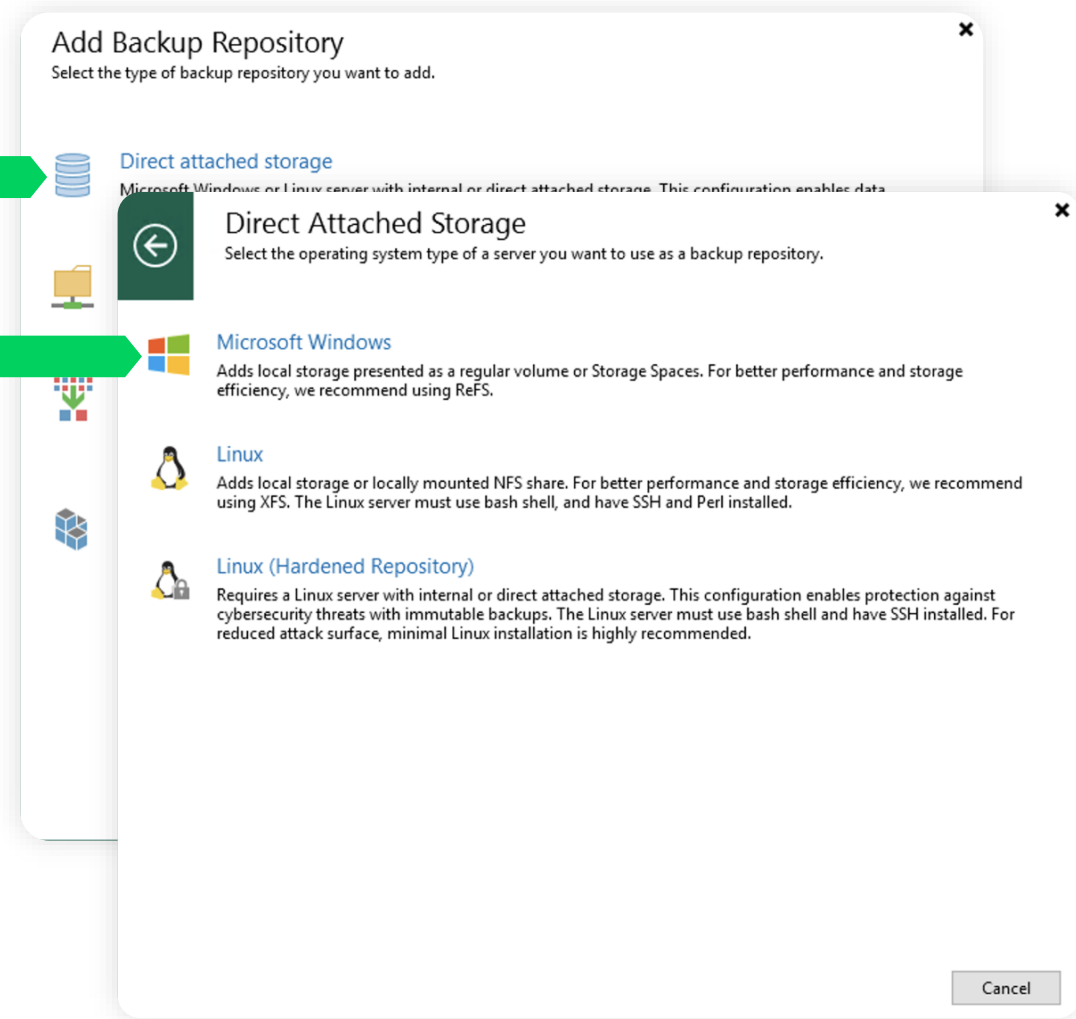
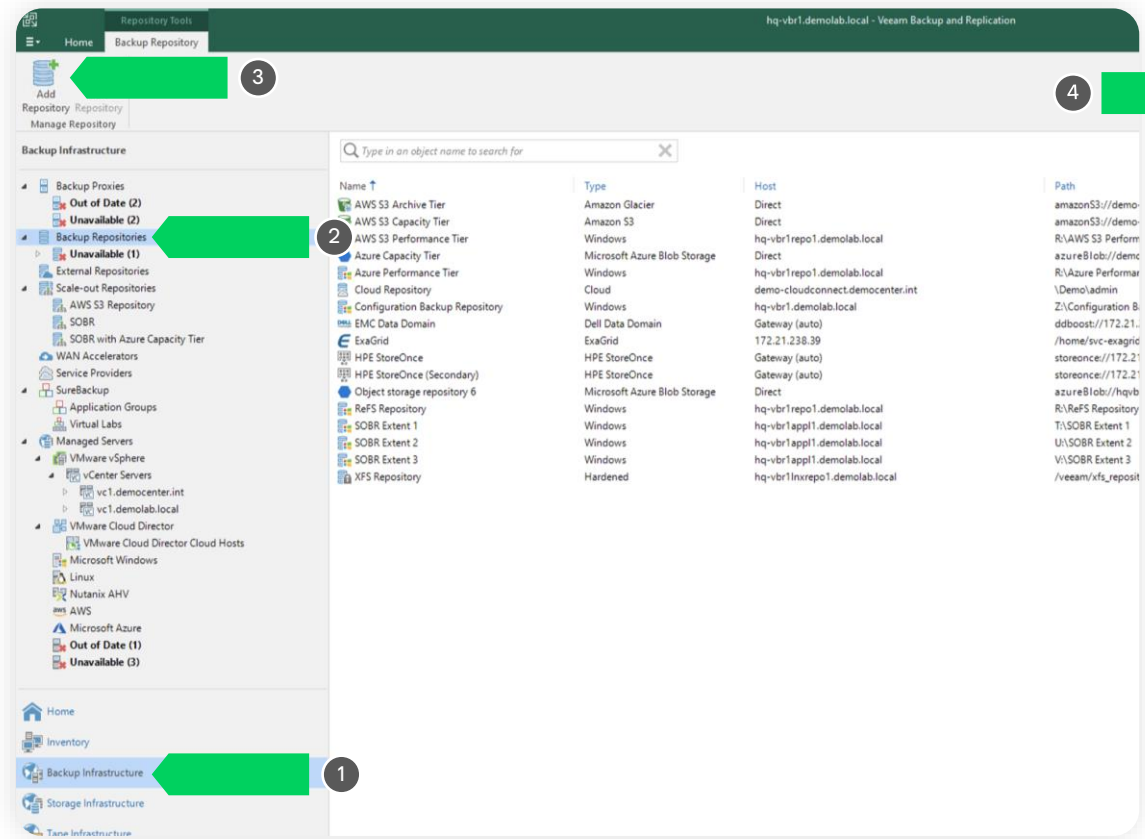
Object Storage Repository Example



How to add a Backup Repository

Architecture & Core Components

How to add a Repository? Direct Attached Storage example.



Architecture & Core Components

How to add a Repository? Direct Attached Storage example.

New Backup Repository

Repository
Type in path to the folder where backup files should be stored, and set repository load control options.

Name

Location
Path to folder: C:\ Browse...

Capacity: <Unknown> Populate
Free space: <Unknown>

Load control
Running too many concurrent tasks against the repository may reduce overall performance, and cause I/O timeouts. Control storage device saturation with the following settings:

Limit maximum concurrent tasks to: 4

Limit read and write data rate to: 1 MB/s

Click Advanced to customize repository settings. Advanced...

< Previous Next > Finish Cancel

Storage Compatibility Settings

Align backup file data blocks (recommended)
Significantly improves backup and restore performance while reducing storage CPU usage by avoiding unaligned I/O. Increases backup size by less than 2%.

Decompress backup file data blocks before storing
Source data mover compresses data according to the backup job compression settings to minimize LAN traffic. Uncompressing the data before storing allows for better deduplication ratio on most deduplicating storage appliances.

This repository is backed by rotated drives
Backup jobs pointing to this repository will tolerate the disappearance of previous backups by creating a new full, and track the repository volume location across unintentional drive letter changes.
When a drive is changed:
Continue existing backup chains (if present)

Use per-machine backup files (recommended)
Improves backup performance for storage devices benefiting from multiple I/O streams, such as enterprise grade block storage and deduplicating storage appliances. Enables additional backup management functionality.

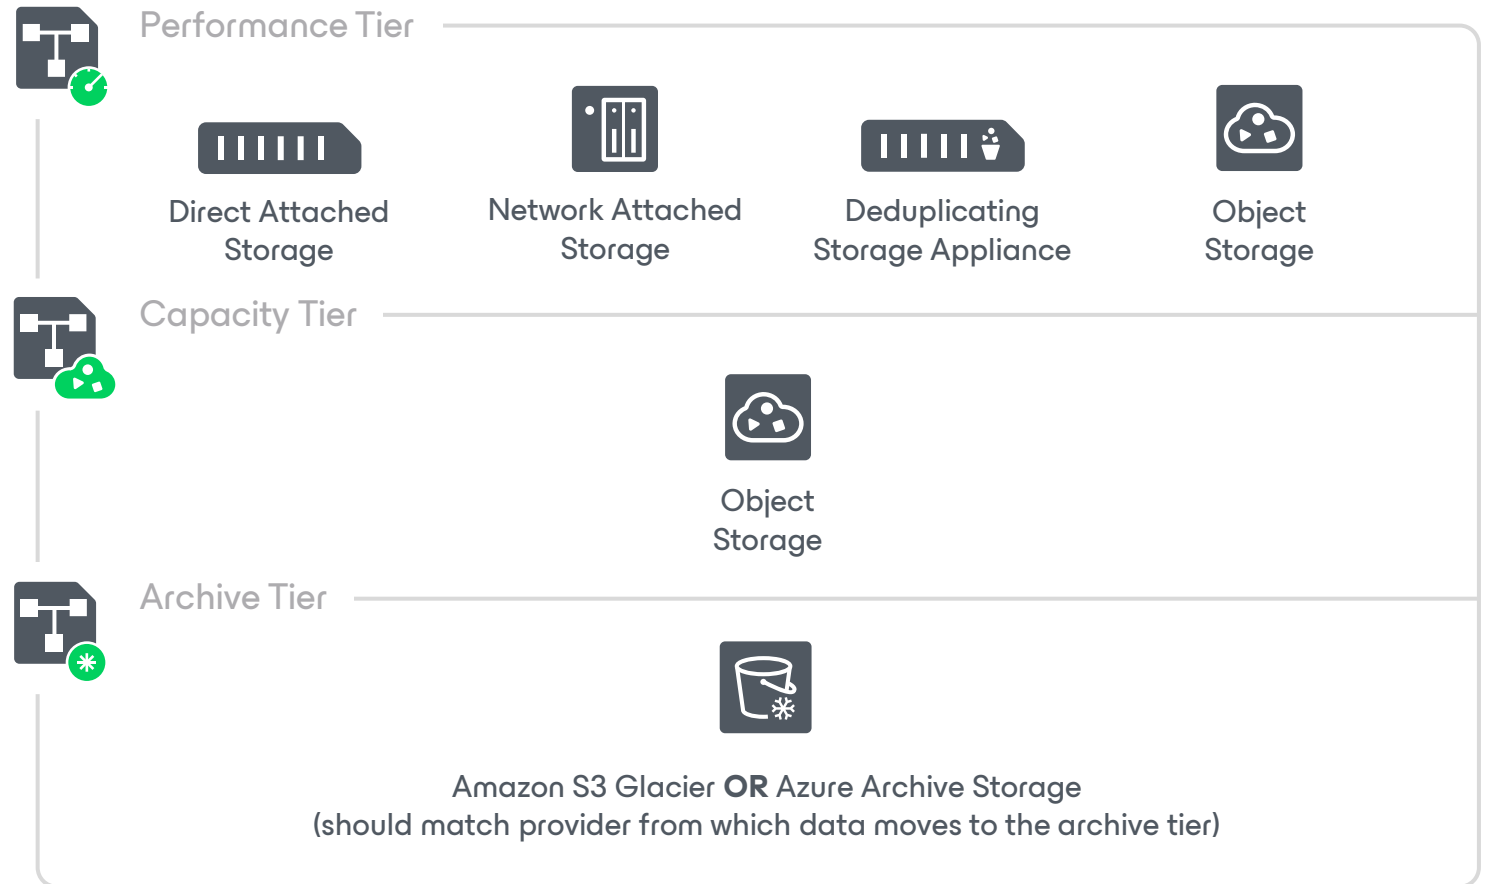
OK Cancel

Scale-Out Backup Repository

Architecture & Core Components

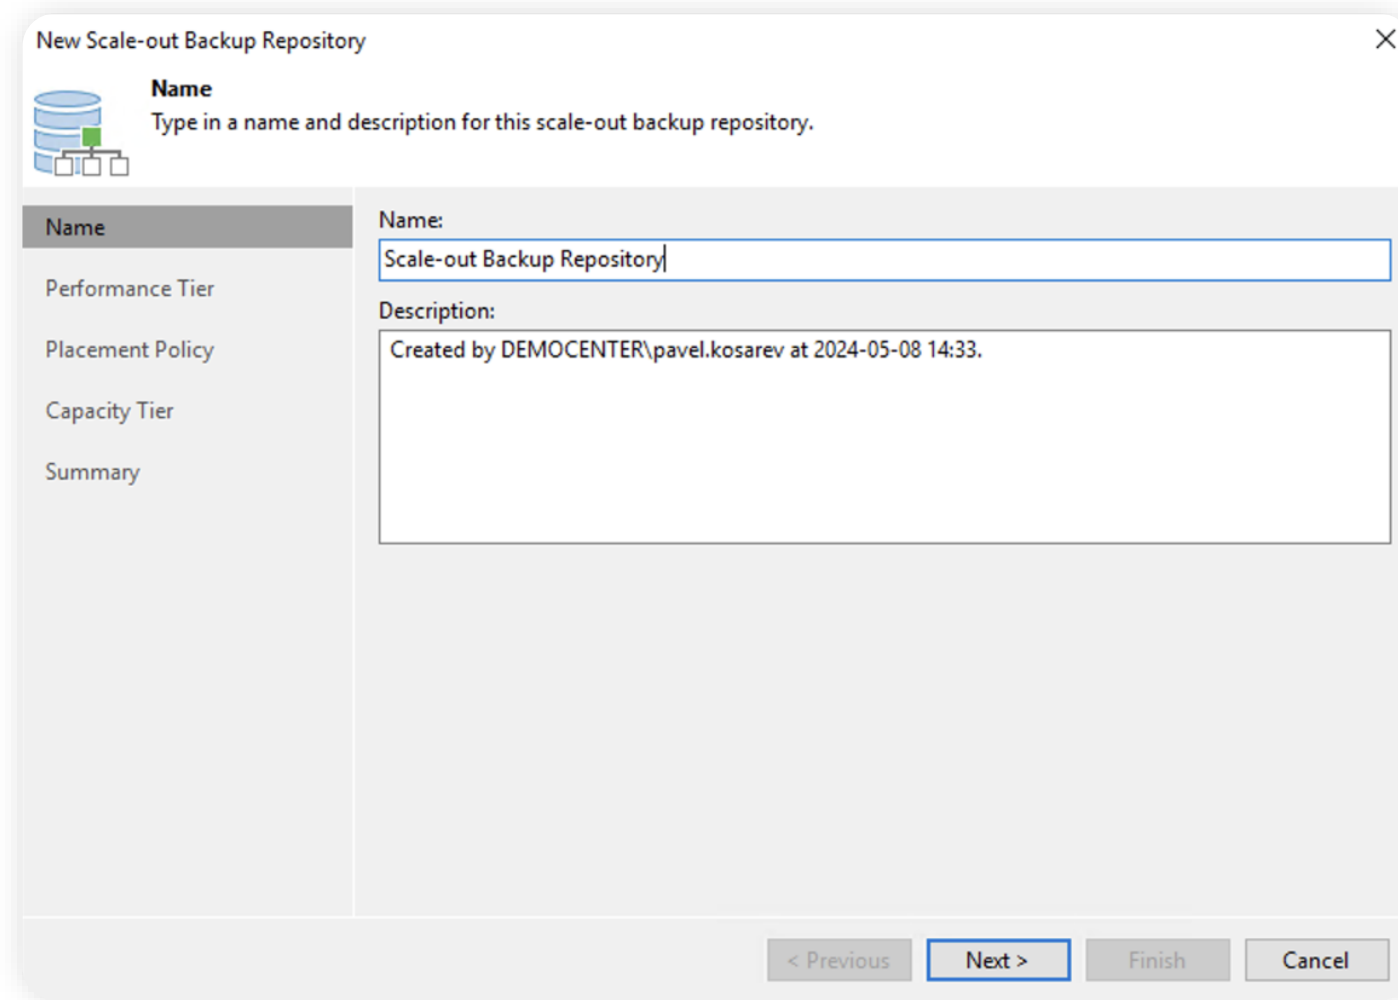
Scale-Out Backup Repository

A scalable repository system with **multi-tier** storage support. Includes **performance** tier (local or shared storage) and can be extended with **capacity** and **archive** tiers, providing horizontal scaling for diverse storage needs.




Architecture & Core Components

How to add/configure Scale-Out Backup Repository?



New Scale-out Backup Repository

 **Name**
Type in a name and description for this scale-out backup repository.

Name

Performance Tier

Placement Policy

Capacity Tier

Summary

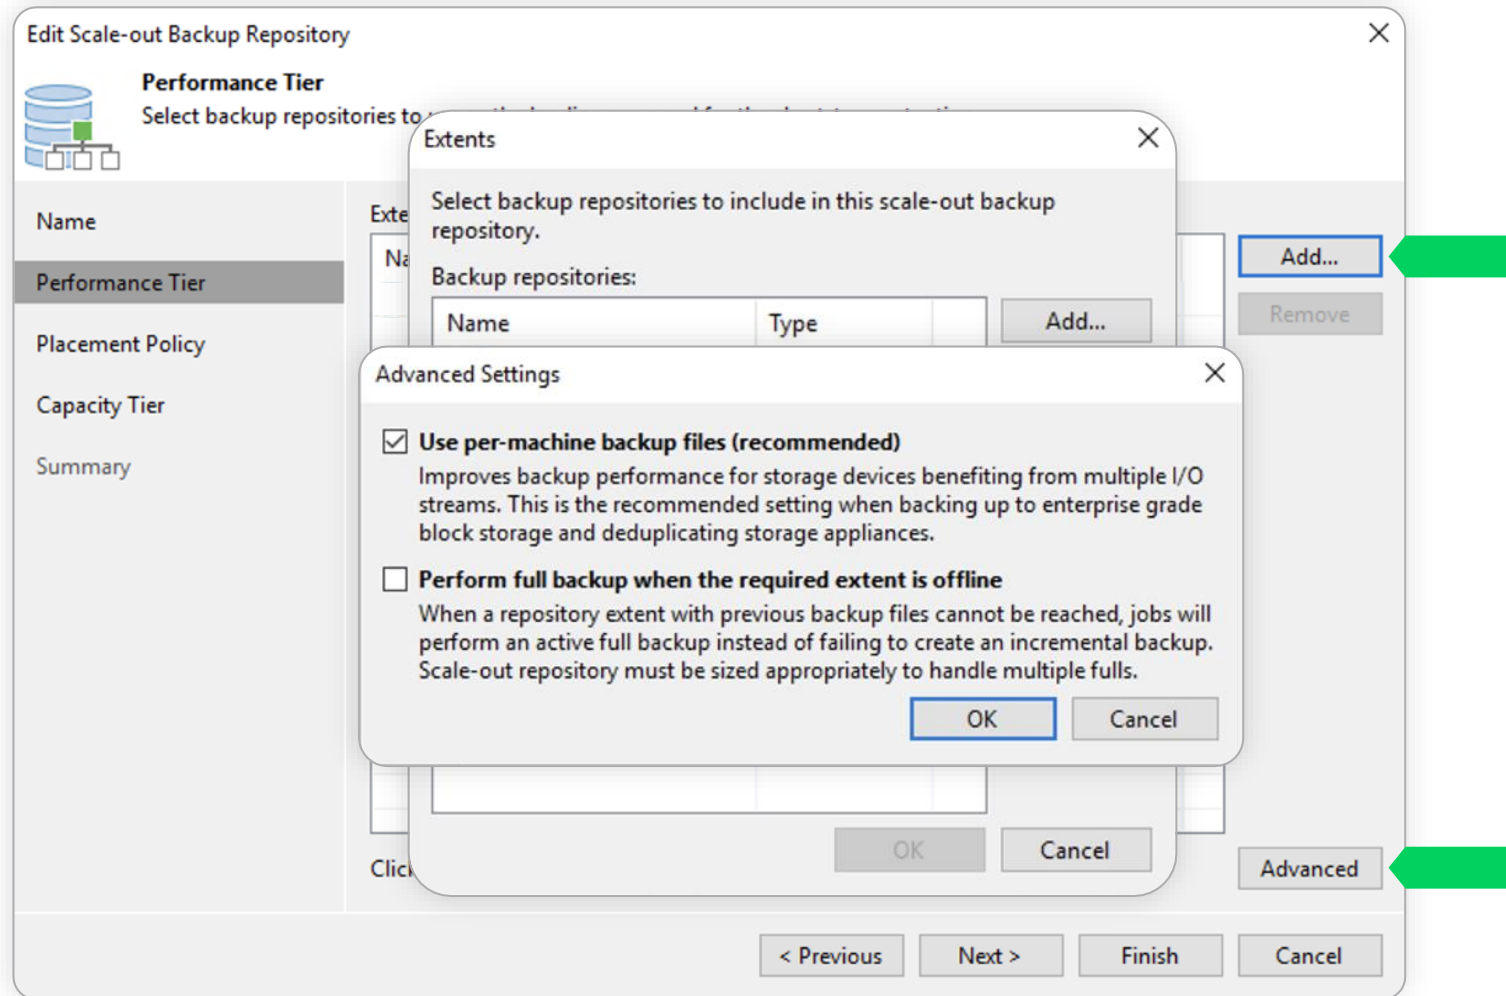
Name:
Scale-out Backup Repository

Description:
Created by DEMOCENTER\pavel.kosarev at 2024-05-08 14:33.

< Previous **Next >** Finish Cancel

Architecture & Core Components

How to add/configure Scale-Out Backup Repository?



Architecture & Core Components

How to add/configure Scale-Out Backup Repository?

Edit Scale-out Backup Repository

Placement Policy
Choose a backup files placement policy for this performance tier. When more than one extent matches the placement policy, backup job will choose the extent with the most free disk space available.

Name [Redacted] **Data locality**
All dependent backup files are placed on the same extent. For example, incremental backup files will be stored together with the corresponding full backup file. However, the next full backup file can be created on another extent (except extents backed by a deduplicating storage).

Performance Tier

Placement Policy [Redacted] **Performance**
Incremental backup files are placed on a different extent from the corresponding full backup file, providing better performance. Note that this policy may not be supported on all backup devices.

Capacity Tier

Summary

Strip...
By to...
suc...

Backup Placement Settings

Name	Allowed backups
SOBR Extent 1	All backups

SOBR Extent 1 Extent Settings

Allowed backup files:

- Full backup files
- Incremental backup files

Buttons: OK, Cancel

Architecture & Core Components

How to add/configure Scale-Out Backup Repository?

Edit Scale-out Backup Repository

Capacity Tier
Specify object storage to copy backups to for redundancy and DR purposes. Older backups can be moved to object storage completely to reduce long-term retention costs while preserving the ability to restore directly from offloaded backups.

Name: [Redacted] Extend scale-out backup repository capacity with object storage:
Performance Tier: <Click Choose to pick object storage> Choose...

Placement Policy: [Redacted] Copy backups to object storage as soon as they are created

Capacity Tier
 Offload backup files sooner if scale-out backup repository is reaching capacity
Offload until used space is below: 90 %

Summary
Offload window: Any time
Health check: Disabled

Buttons: < Previous, Apply, Finish, Cancel, Override...

SureBackup, SureReplica & Virtual Lab

Advanced Features


SureBackup, SureReplica & On-Demand Sandbox

SureBackup

Allows you to automatically perform recoverability testing, verification and content scans for the successfully performed backups.

How?

1. Starts VMs in an isolated Virtual Lab environment.
2. Performs a set of tests.
3. Sends a status report to your mailbox.



Advanced Features

SureBackup Verification Modes

Full recoverability testing

Veeam Backup & Replication published the machines in the isolated environment and performs the verification according to the scripts selected.

Backup verification and content scan only

Without publishing the machines to a Virtual Lab, Veeam Backup & Replication performs backup integrity check and its content analysis to detect traces of malware or any other unwanted or sensitive data.

New SureBackup Job

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

Name

Virtual Lab

Application Group

Linked Jobs

Settings

Schedule

Summary

Name: Exchange SureBackup Job Full

Description: Daily Verification Job

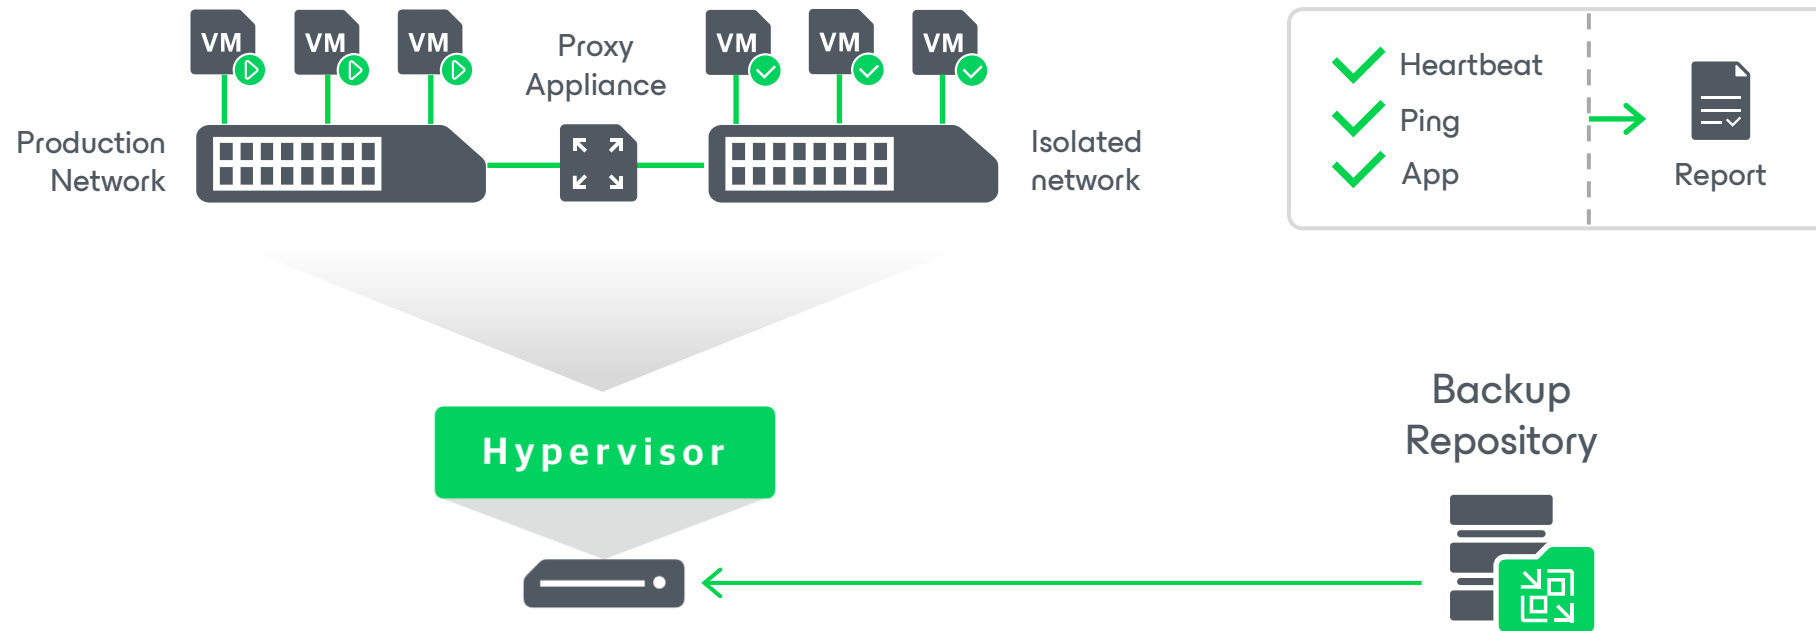
Backup verification mode:

- Full recoverability testing (recommended)
Runs machines in an isolated environment directly from backup and performs tests against live applications. This ensures recoverability of your production workloads in a DR event.
- Backup verification and content scan only
Performs backup integrity check and its content analysis to detect traces of malware or any other unwanted or sensitive data. These tests do not require setting up a virtual lab.

< Previous Next > Finish Cancel

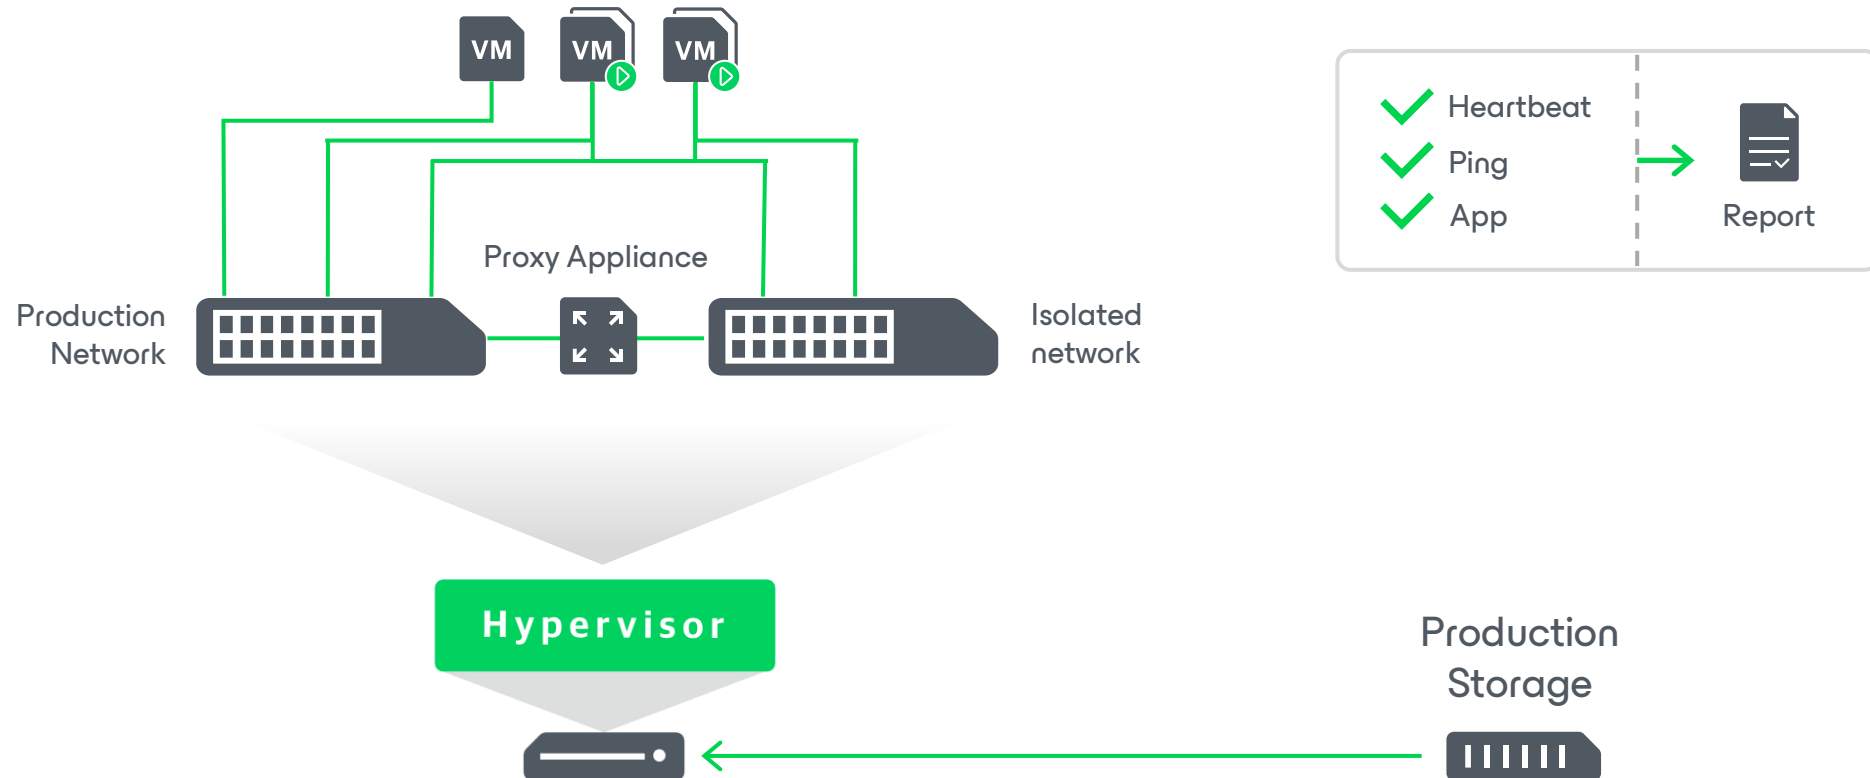
Advanced Features

SureBackup Workflow



Advanced Features

SureReplica Workflow (VMware only)




Advanced Features

On-Demand Sandbox

Start a copy of your production environment at any time for a variety of **testing, security, training** or **troubleshooting** purposes.

All changes made to VMs are written to redo logs (for VM backups and storage snapshots) or saved to delta files (for VM replicas). Redo logs and delta files are deleted after you finish working with the On-Demand Sandbox and power it off.



Immutability

Immutability

Definition of Immutability:

- Immutability refers to the state of data that prevents it from being modified or deleted.

Benefits of Immutability:

- Ensures data integrity and security.
- Provides protection against ransomware and accidental deletions.

Supported types of immutable repositories

- Amazon, S3-compatible and Azure object storage repositories.
- Hardened Repository.
- HPE StoreOnce.
- Dell Data Domain

Immutability

Adding Linux Hardened Repository

At the SSH Connection step of the wizard, specify **single-use credentials** to connect to the Linux server and deploy Veeam Data Mover.

Veeam Backup & Replication **does not store these credentials** in the configuration database.

NOTE: The user account you use must be a non-root account. Also, it must have the **home** directory created on the Linux server.

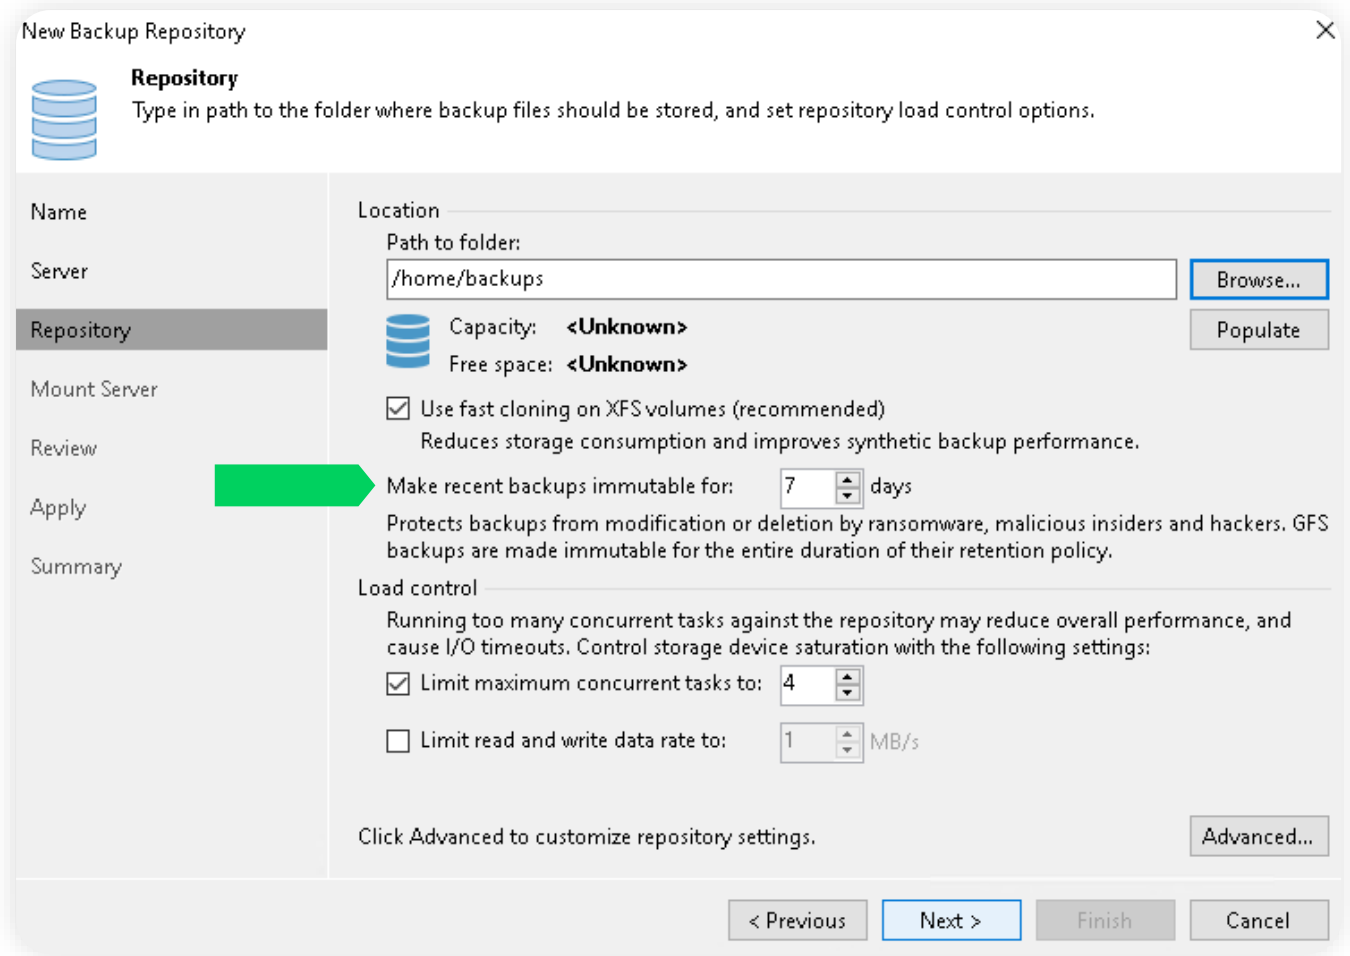
The screenshot shows the 'Edit Linux Server' wizard in the 'SSH Connection' step. A 'Credentials' dialog box is open, allowing the user to configure connection details. The dialog includes fields for Username (set to 'user1'), Password, SSH port (set to 22), and Root password. Under the 'Non-root account' section, there are three checkboxes: 'Elevate account privileges automatically' (checked), 'Add account to the sudoers file' (unchecked), and 'Use "su" if "sudo" fails' (checked). A Description field contains the text 'user1'. The dialog has 'OK' and 'Cancel' buttons at the bottom. The main wizard window has a sidebar with 'Name', 'SSH Connection', 'Review', 'Apply', and 'Summary' options. At the bottom of the wizard, there are '< Previous', 'Next >', 'Finish', and 'Cancel' buttons.

Immutability

Making backups immutable @ LHR

When you add a hardened repository, you specify the time period while backup files must be immutable.

During this period, backup files stored in this repository **cannot be modified or deleted**.



New Backup Repository

Repository
Type in path to the folder where backup files should be stored, and set repository load control options.

Name

Server

Repository

Mount Server

Review

Apply

Summary

Location

Path to folder: /home/backups [Browse...](#)

Capacity: <Unknown> [Populate](#)

Free space: <Unknown>

Use fast cloning on XFS volumes (recommended)
Reduces storage consumption and improves synthetic backup performance.

Make recent backups immutable for: 7 days
Protects backups from modification or deletion by ransomware, malicious insiders and hackers. GFS backups are made immutable for the entire duration of their retention policy.

Load control

Running too many concurrent tasks against the repository may reduce overall performance, and cause I/O timeouts. Control storage device saturation with the following settings:

Limit maximum concurrent tasks to: 4

Limit read and write data rate to: 1 MB/s

Click Advanced to customize repository settings. [Advanced...](#)

< Previous Next > Finish Cancel

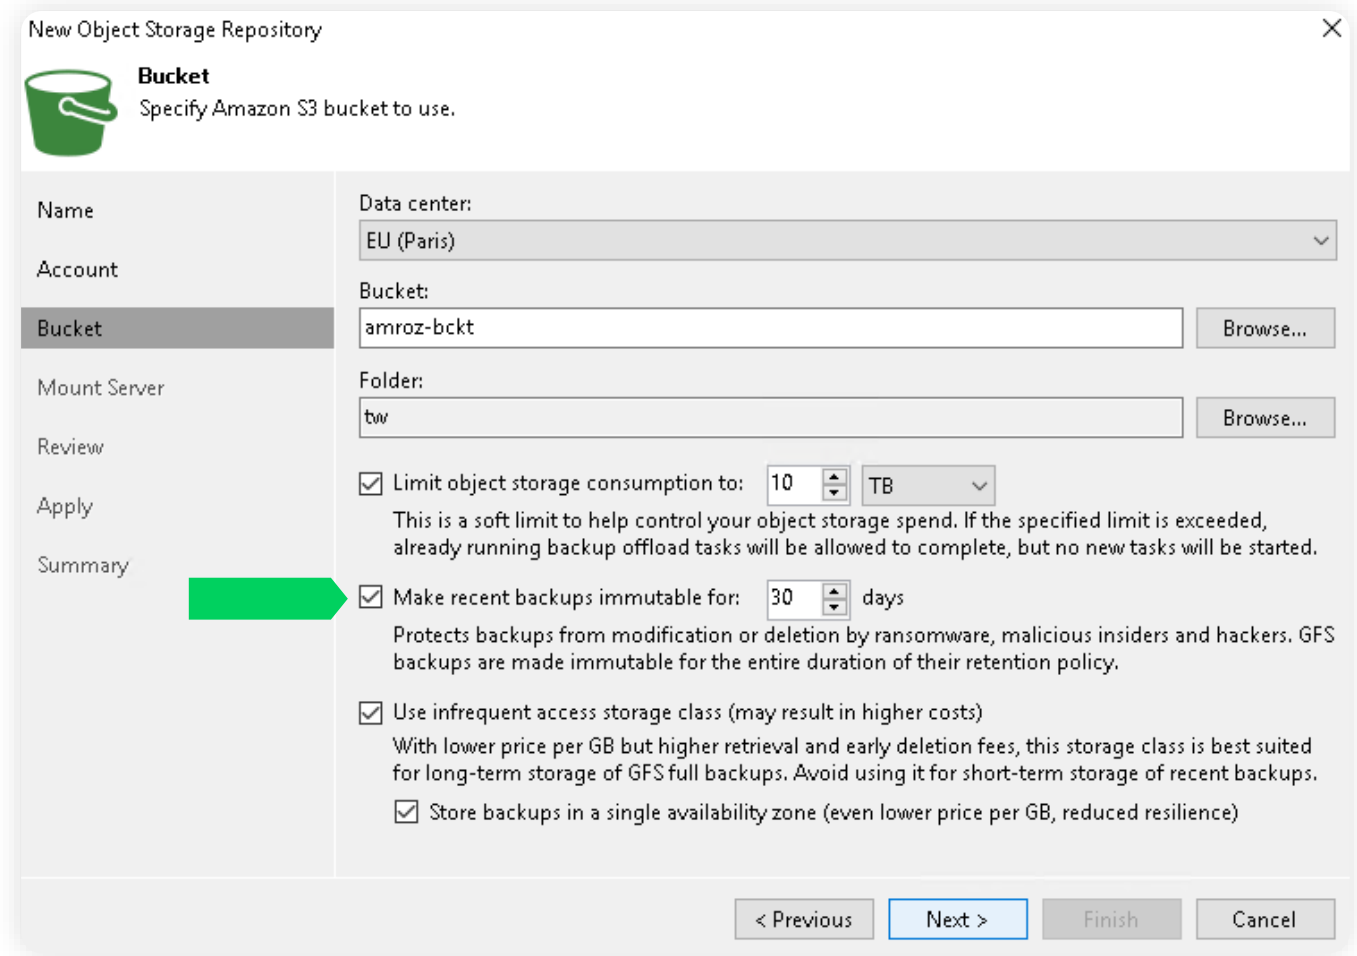
Immutability

Making backups immutable @ Object Storage

To make data immutable, Veeam Backup & Replication utilizes the technology that prevents data from deletion and allows you to keep several versions of objects.

The selected **technology depends on the type of object storage:**

- Object lock and Versioning – for Amazon S3 Storage, S3 Compatible, IBM Cloud, Wasabi Cloud.
- Version-level WORM and blob versioning – for Azure Storage.



New Object Storage Repository

Bucket
Specify Amazon S3 bucket to use.

Name

Account

Bucket

Mount Server

Review

Apply

Summary

Data center:
EU (Paris)

Bucket:
amroz-bckt

Folder:
tw

Limit object storage consumption to: 10 TB
This is a soft limit to help control your object storage spend. If the specified limit is exceeded, already running backup offload tasks will be allowed to complete, but no new tasks will be started.

Make recent backups immutable for: 30 days
Protects backups from modification or deletion by ransomware, malicious insiders and hackers. GFS backups are made immutable for the entire duration of their retention policy.

Use infrequent access storage class (may result in higher costs)
With lower price per GB but higher retrieval and early deletion fees, this storage class is best suited for long-term storage of GFS full backups. Avoid using it for short-term storage of recent backups.

Store backups in a single availability zone (even lower price per GB, reduced resilience)

< Previous Next > Finish Cancel

Storage Integration

Storage Integration

Integration types

VMware Integration

Backup VMware vSphere VMs.

Orchestrate snapshots across storage arrays.

Restore data from storage snapshots with Veeam Backup & Replication.

NAS Integration

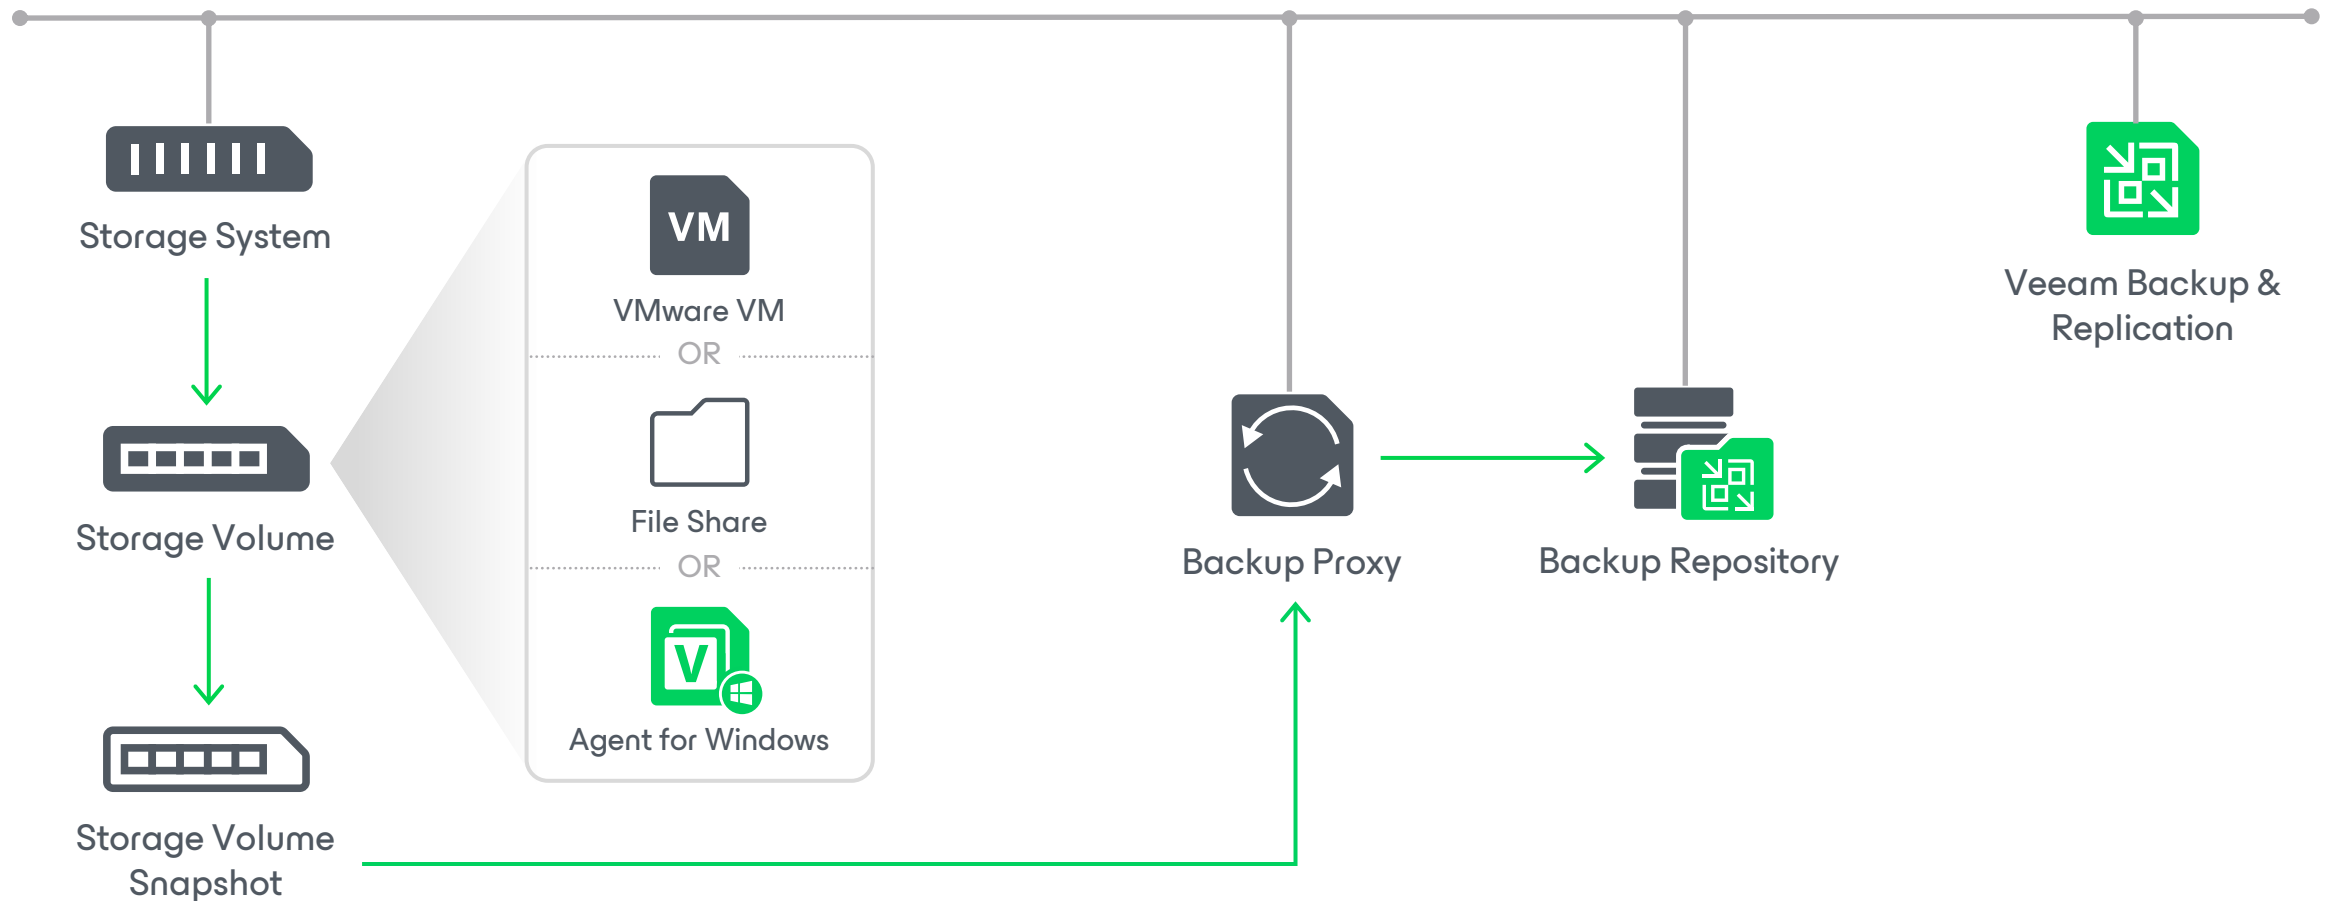
Backup NAS file shares effortlessly.

Veeam Agent for Windows Integration

Create Veeam Agent backups for Windows computers.

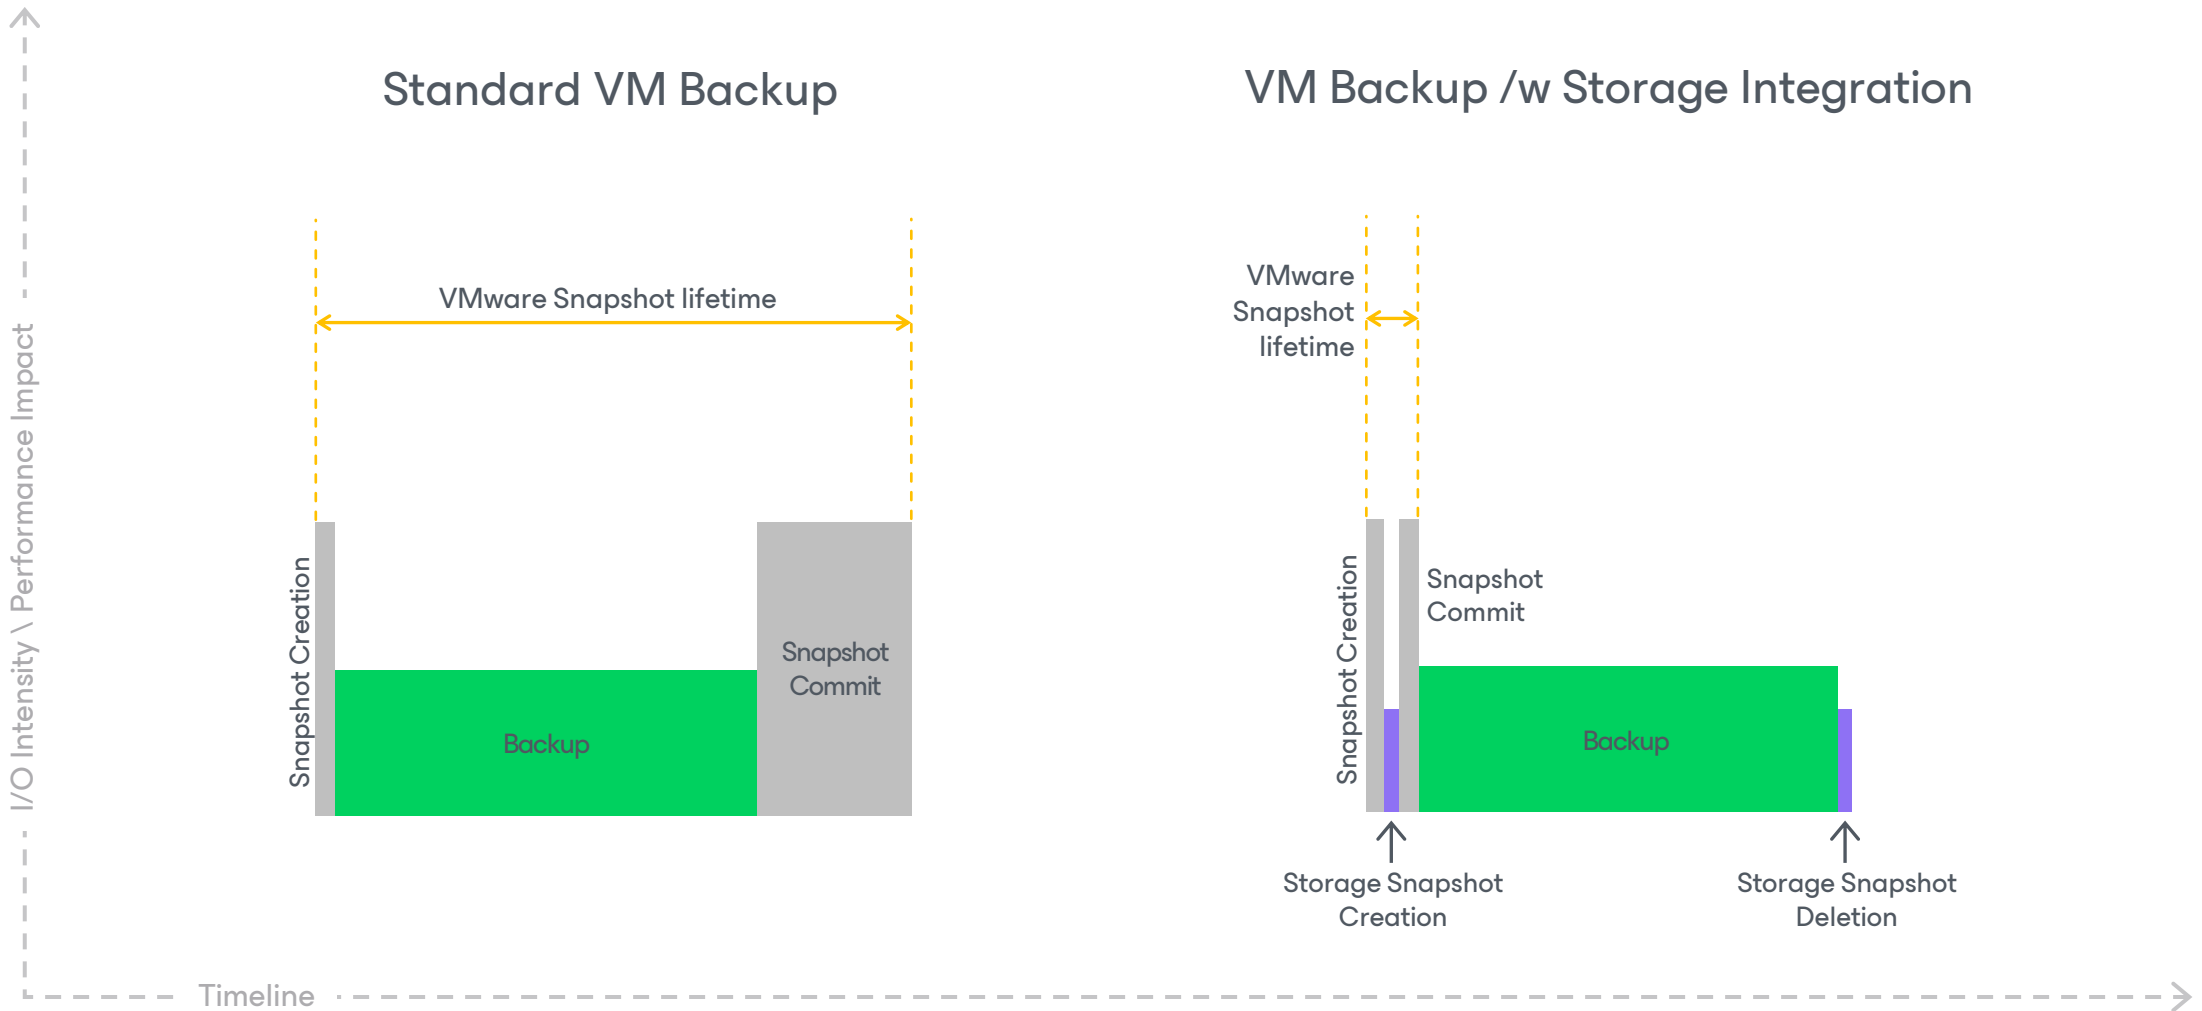
Storage Integration

Infrastrurcture Overview



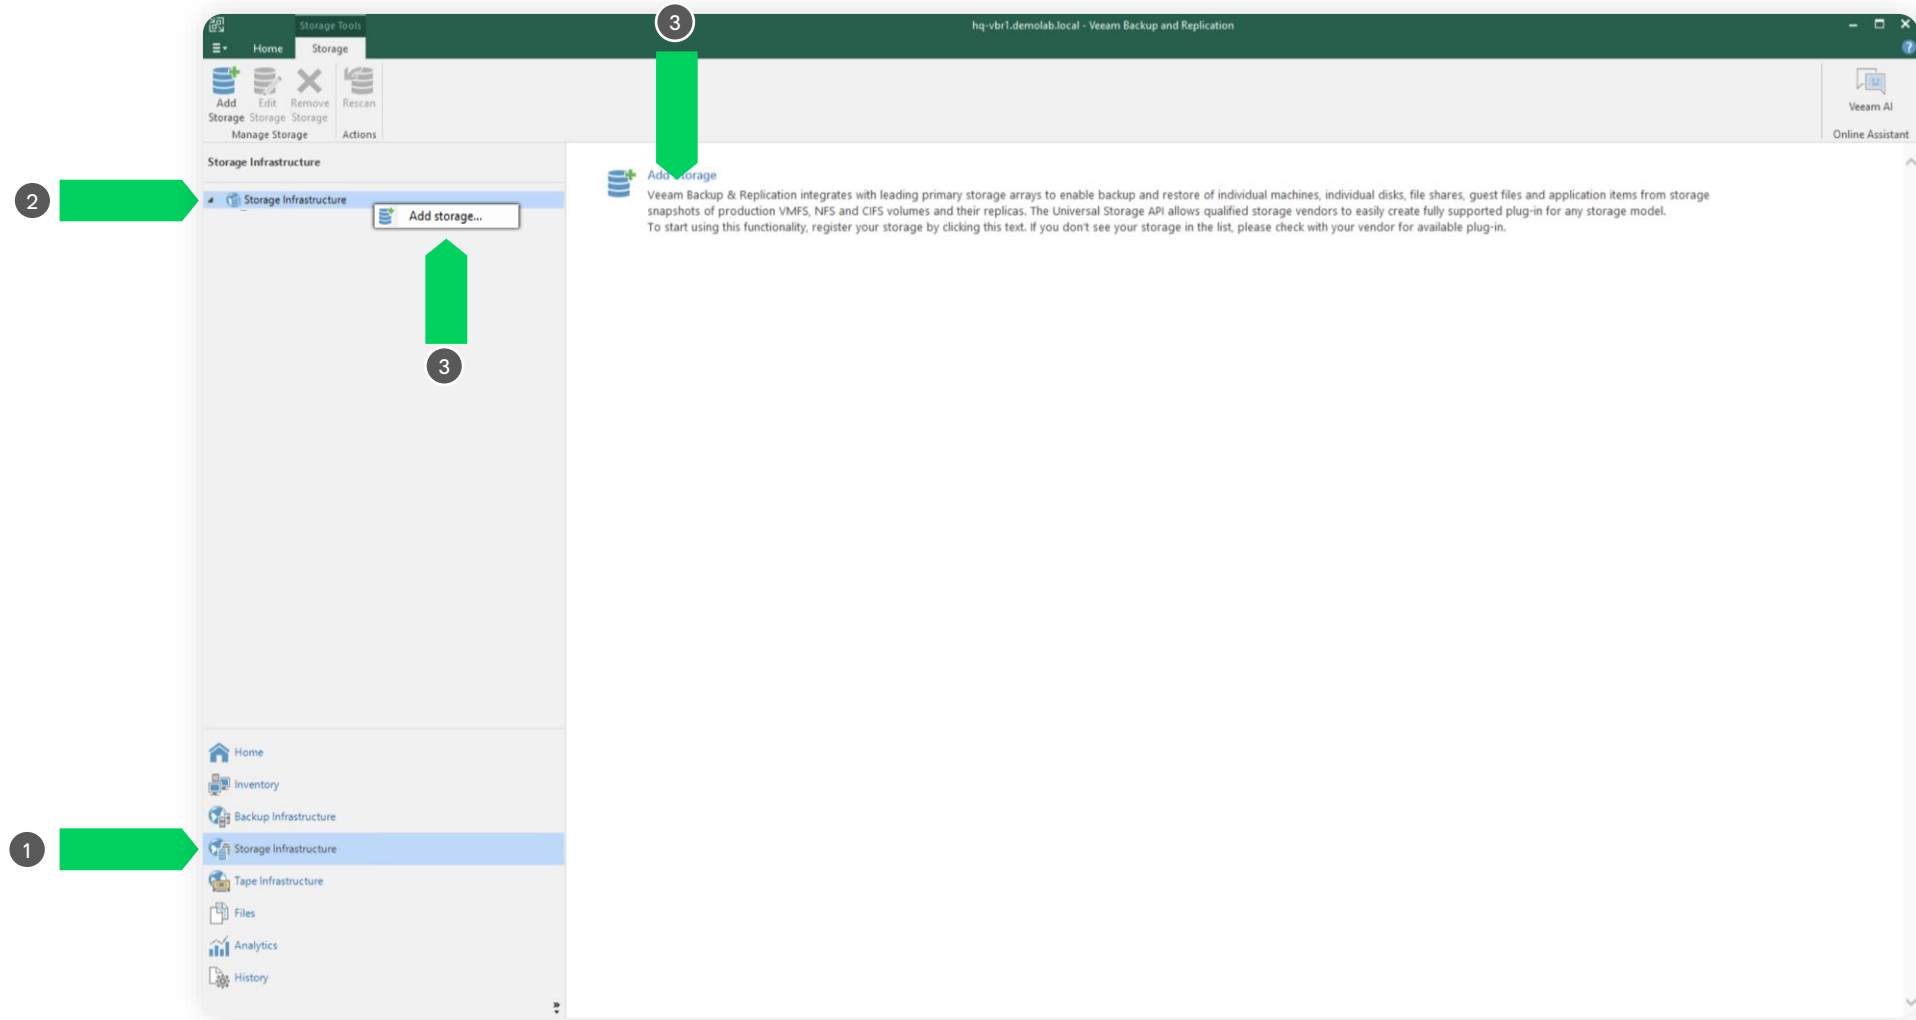
VMware Integration

Timeline Representation



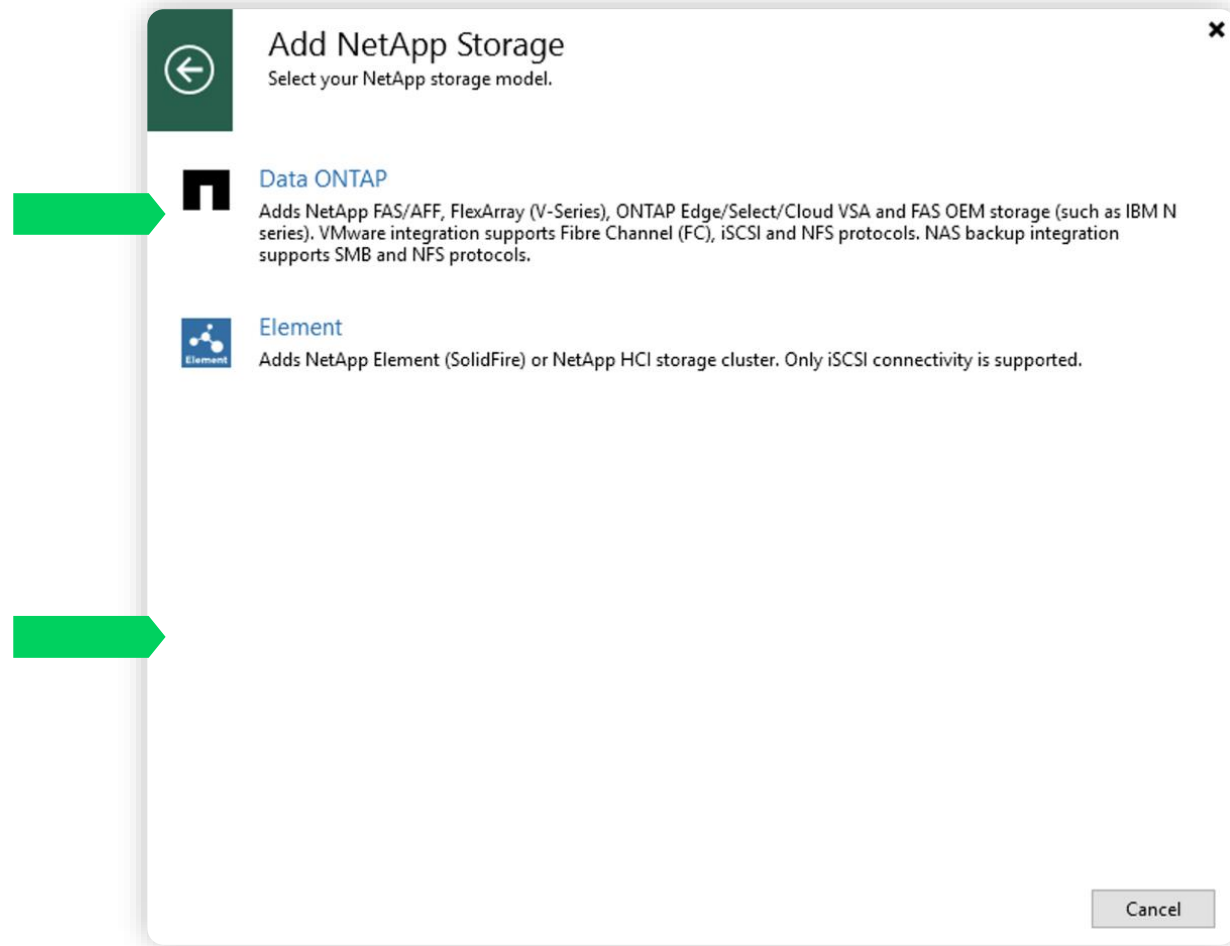
Storage Integration

How to configure?



Storage Integration


How to configure?



Storage Integration

How to configure?

Edit NetApp Data ONTAP Storage

 **Name**
Register NetApp Data ONTAP storage by specifying DNS name or IP address.

Name	Management server DNS name or IP address: <input type="text" value="172.21.238.72"/>
Credentials	Description: <input type="text"/>
VMware vSphere	Role: <input checked="" type="checkbox"/> Block or file storage for VMware vSphere <input type="checkbox"/> Block storage for Microsoft Windows servers <input checked="" type="checkbox"/> NAS filer
NAS Filer	
Apply	
Summary	

< Previous **Next >** Finish Cancel

Storage Integration

How to configure?

Edit NetApp Data ONTAP Storage

Credentials
Specify account with storage administrator privileges.

Name

Credentials

VMware vSphere

NAS Filer

Apply

Summary

Credentials:

svc-vbr (NetApp access account, last edited: 1349 days ago)

Add...

Manage accounts

Protocol: HTTPS

Port: 443

< Previous Next > Finish Cancel

Storage Integration

How to configure?

The screenshot shows a configuration window titled "Edit NetApp Data ONTAP Storage" with a close button (X) in the top right corner. The window is divided into two main sections. On the left is a sidebar with a tree view containing the following items: "Name", "Credentials", "VMware vSphere" (which is selected and highlighted), "NAS Filer", "Apply", and "Summary". The main area on the right is titled "VMware vSphere" and contains the instruction "Specify how this storage can be accessed by VMware vSphere backup jobs." Below this, there are several configuration options: "Protocol to use:" with three radio buttons: "Fibre Channel (FC)", "iSCSI", and "NFS" (which is selected); a checked checkbox "Create required export rules automatically"; "Volumes to scan:" with a text box containing "All volumes" and a "Choose..." button; "Backup proxies to use:" with a text box containing "hq-vbr1appl1.demolab.local" and a "Choose..." button; and "Mount server:" with a dropdown menu showing "hq-vbr1appl1.demolab.local (Created by Powershell at 7/9/2020 2:05:00 AM.)" and an "Add New..." button. At the bottom of the window are four buttons: "< Previous", "Next >" (highlighted with a blue border), "Finish", and "Cancel". Several green arrows point to the "Choose..." and "Add New..." buttons, and a green arrow points to the "VMware vSphere" section header.

Storage Integration

How to configure?

Edit NetApp Data ONTAP Storage

NAS Filer
Specify how this storage can be accessed by file backup jobs.

Name

Credentials

VMware vSphere

NAS Filer

Apply

Summary

Protocol to use:

SMB

NFS

Create required export rules automatically

Volumes to scan:

All volumes Choose...

Backup proxies to use:

Automatic selection Choose...

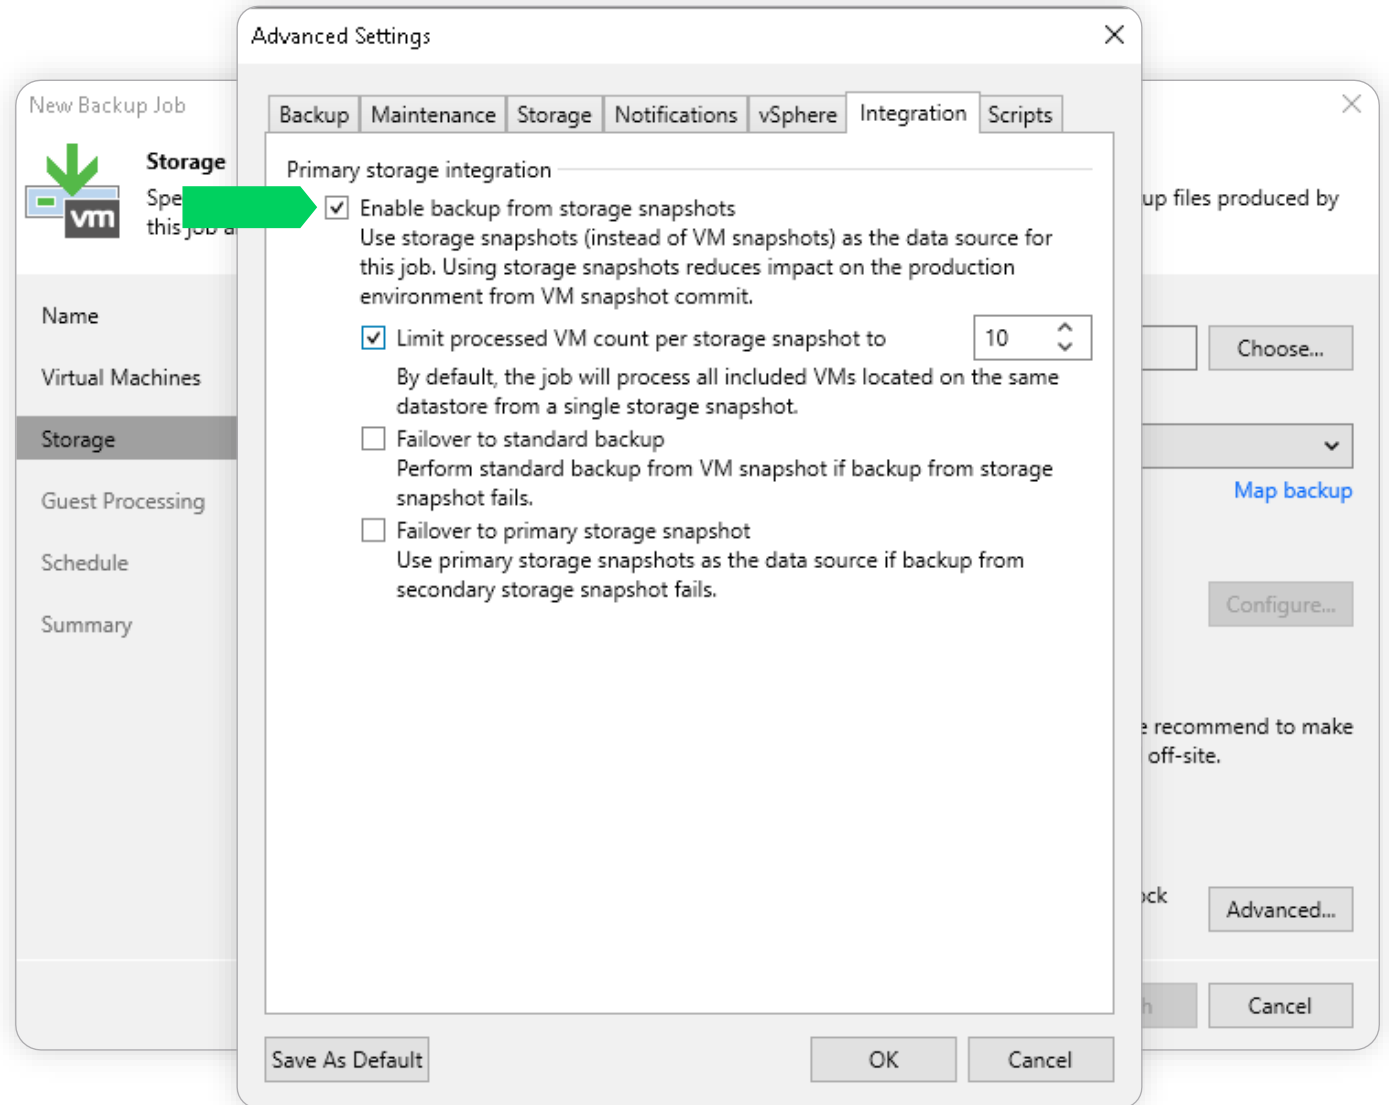
< Previous Apply Finish Cancel

Storage Integration

The **Integration** tab in the backup job settings allows you to determine if you want to use the Backup from Storage Snapshot technology to create the backup. It's **enabled by default**.

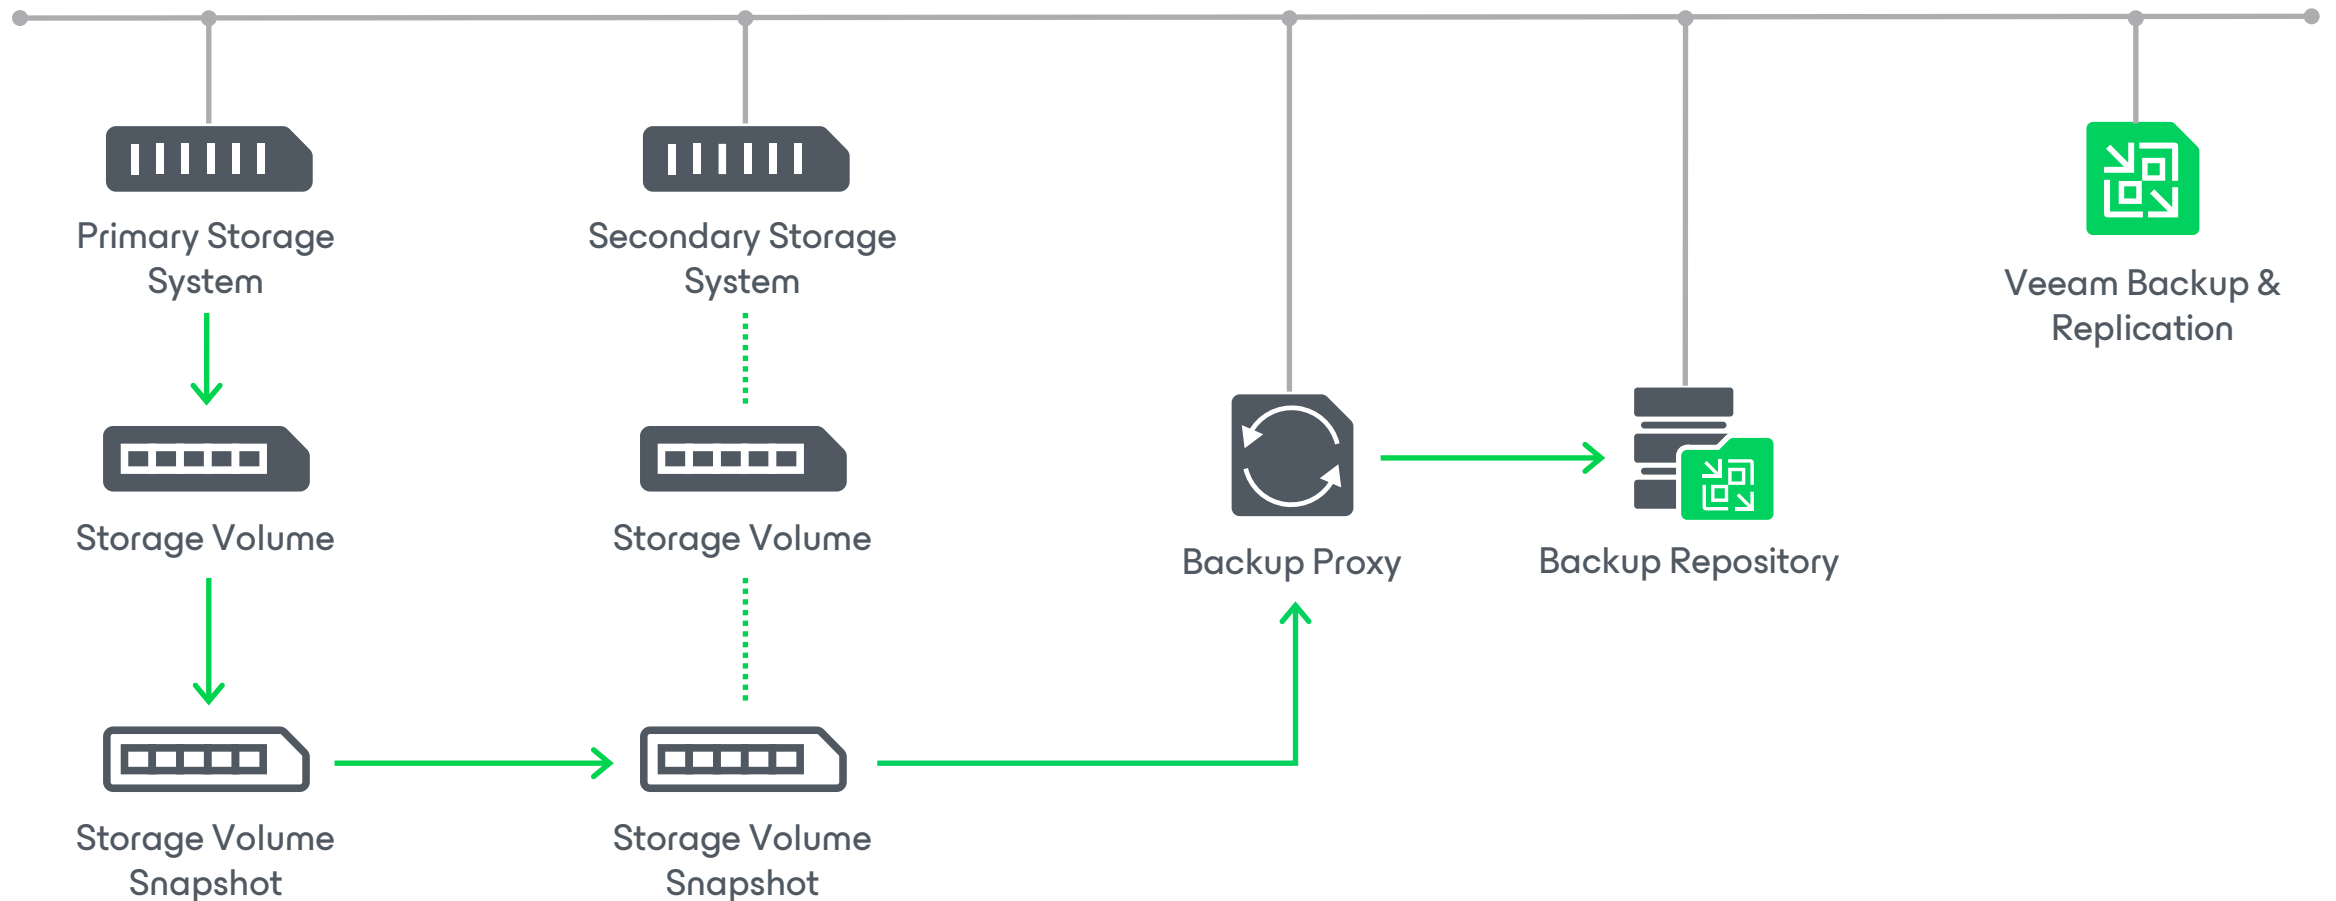
Backup from Storage Snapshots lets you leverage storage snapshots for VM data processing.

The technology **improves RPOs and reduces the impact of backup activities** on the production environment.

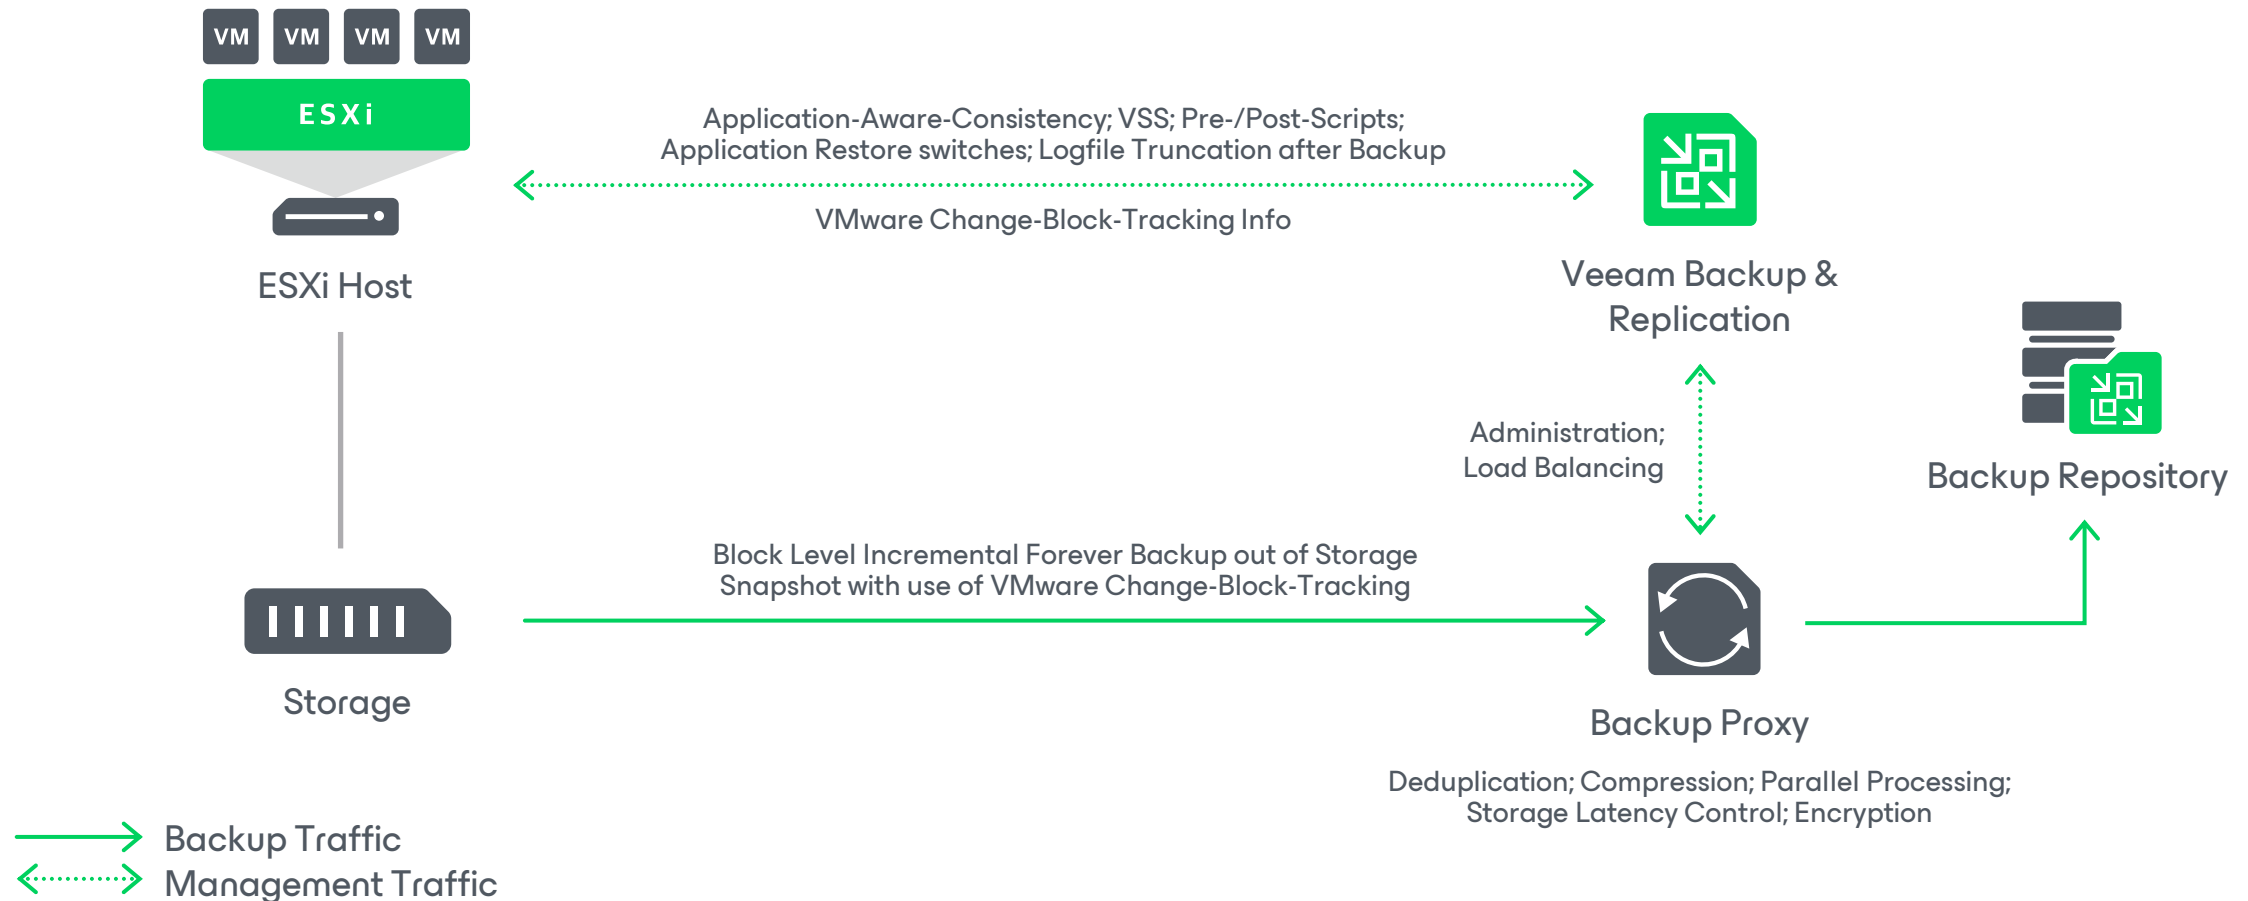


Storage Integration

Backup from Secondary Storage Snapshot



Storage Snapshot Based Backup



Tape (briefly)

Tape

Veeam Backup & Replication supports native tape operations.

Supported tape devices: Linear Tape-Open (LTO3 – LTO9) and IBM 3592 (TS1160 & TS1170) libraries.

Purpose: long-term data archiving and/or compliance requirements.

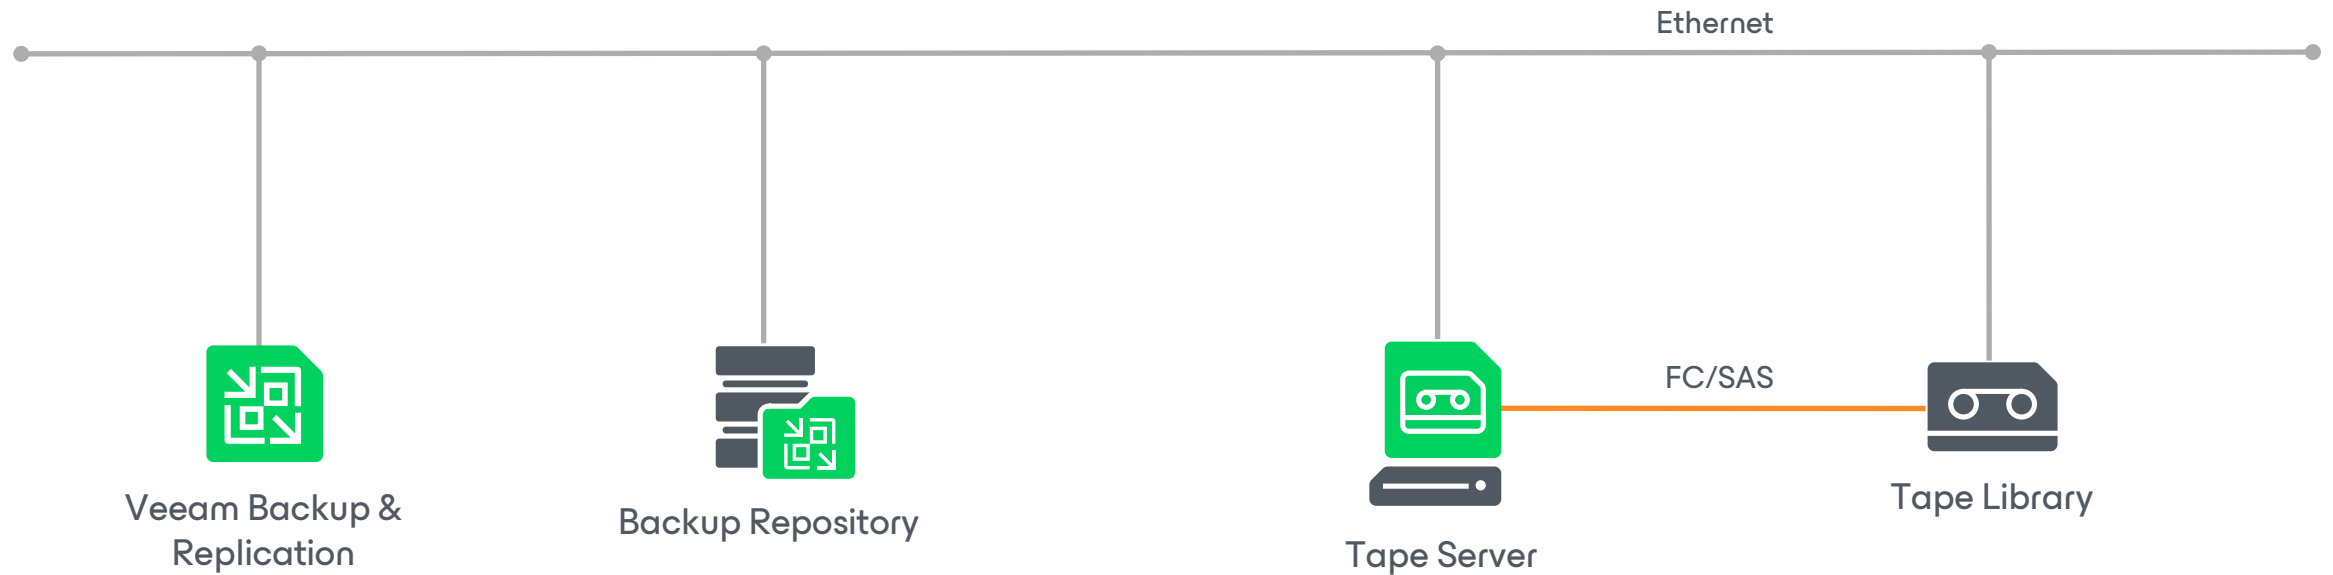
Backup Strategy: adheres to '3-2-1' backup principle: 3 copies, 2 media types, 1 offsite location.

Supported Data Types for Archiving:

- VM backups
- Veeam Agent backups
- Entire Backup repositories, including Hardened Linux Repositories
- Unstructured data:
 - Storage device volumes via NDMP
 - SMB & NFS file shares
 - Object Storage

Tape

Infrastructure overview



WAN Accelerator

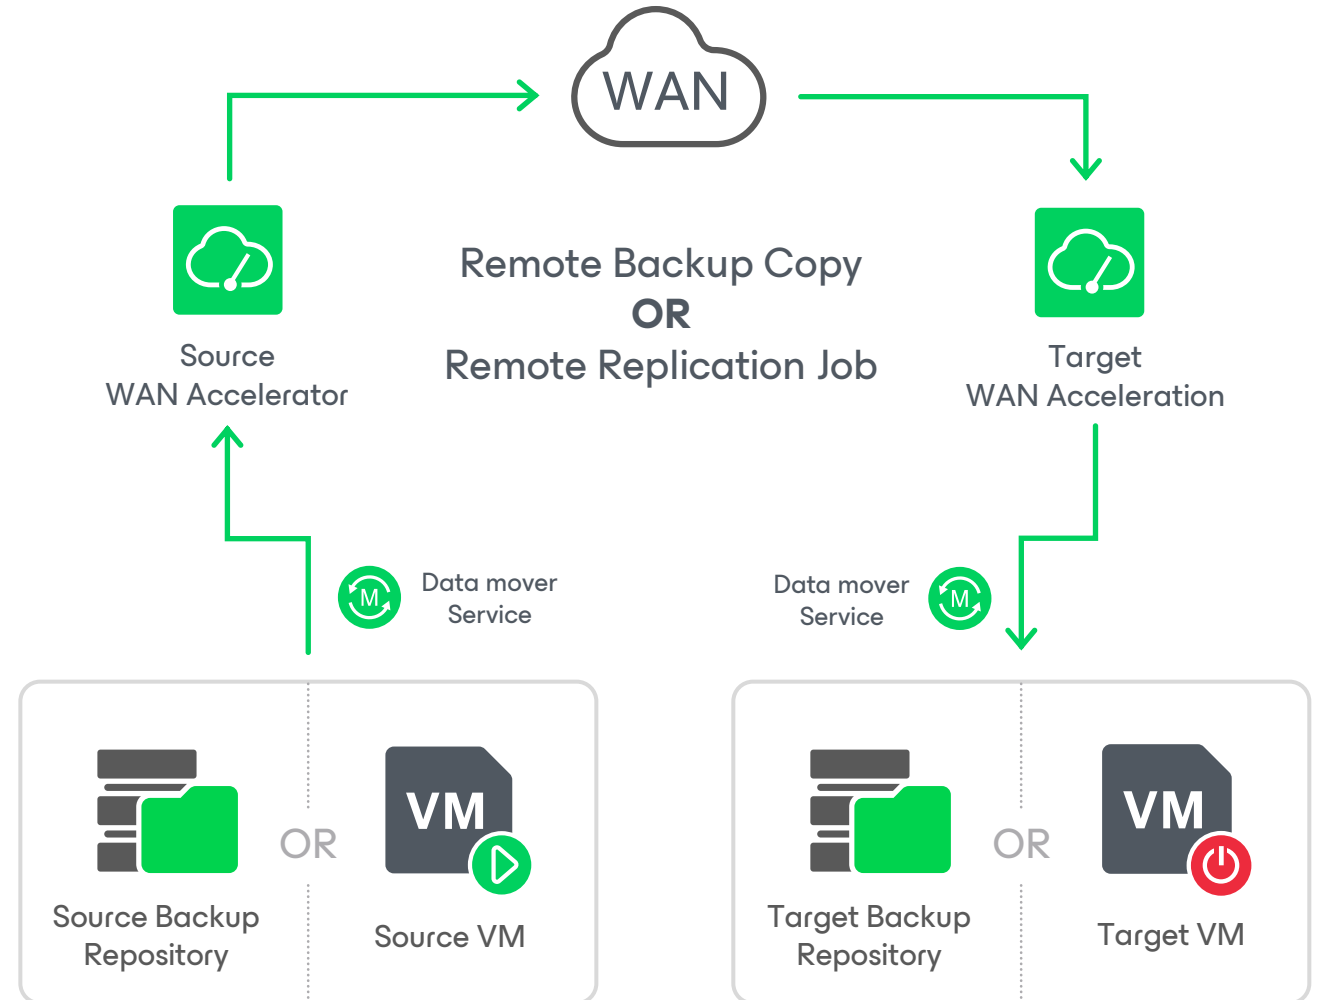
WAN Accelerator

Veeam Backup & Replication offers the WAN acceleration technology that combines:

- Network traffic compression
- Multi-streaming upload
- Global data deduplication
- Variable block size deduplication

This combination of technologies reduces the volume of transferred data, which helps transport data faster.

- **High Bandwidth Mode** is designed for WAN connections faster than 100Mbps.
- **Low Bandwidth Mode** is designed for WAN connections slower than 100Mbps.



WAN Accelerator

How to deploy?

The screenshot shows the Veeam Backup and Replication console interface. The top navigation bar includes 'Home' and 'WAN Accelerator'. The left sidebar shows the 'Backup Infrastructure' tree, with 'WAN Accelerators' highlighted. A green arrow labeled '1' points to 'Backup Infrastructure' in the sidebar. The main pane shows a table of WAN Accelerators with columns for Name, Host, and Description. A green arrow labeled '2' points to the 'WAN Accelerators' folder in the sidebar. At the top, there are buttons for 'Add WAN Accelerator', 'Upgrade WAN Accelerator', and 'Manage WAN Accelerators'. A green arrow labeled '3' points to the 'Add WAN Accelerator' button.

The screenshot shows the 'New WAN Accelerator' dialog box in the 'Review' stage. The dialog has a 'Review' tab selected. The settings are as follows:

- Server name: **hq-vbr1.demolab.local**
- Server type: **Virtual (VMware)**
- Cache size: **10 GB**
- Cache path: **C:\VeeamWAN**

The dialog also shows a table of components to be processed on the server:

Component name	Status
Transport	already exists
WAN Accelerator	will be installed

At the bottom, there are buttons for '< Previous', 'Apply', 'Finish', and 'Cancel'. A green arrow points down to the 'Apply' button.

WAN Accelerator

How to start using it?

The image shows two overlapping configuration windows from Veeam Backup & Replication. The background window is titled "Edit Backup Copy Job" and the foreground window is titled "Edit Replication Job". Both windows have a "Data Transfer" section. In the "Edit Backup Copy Job" window, the "Data Transfer" section is active, and a green arrow points to the "Through built-in WAN accelerators" radio button. The "Direct" option is currently selected. In the "Edit Replication Job" window, the "Data Transfer" section is also active, and a green arrow points to the "Through built-in WAN accelerators" radio button. The "Direct" option is currently selected. Both windows have a sidebar on the left with navigation options: Job, Objects, Target, Data Transfer, Schedule, and Summary. The "Data Transfer" option is highlighted in both. At the bottom of the "Edit Replication Job" window, there are four buttons: "< Previous", "Next >", "Finish", and "Cancel".

Edit Backup Copy Job

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

Direct
Object data will be sent directly from source to target backup repository, copying backups on-site, and off-site over a fast link.

Through built-in WAN accelerators
Object data will be sent to target repository through built-in WAN accelerators at both source and target sites. This mode provides significant bandwidth savings.

Source WAN accelerator:

Target WAN accelerator:

Edit Replication Job

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

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

Source proxy:

Target proxy:

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

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

Source WAN accelerator:

▼

Target WAN accelerator:

▼



< Previous **Next >** Finish Cancel

Enterprise Manager

Veeam Enterprise Manager

Installation

Veeam Enterprise Manager is included in the same .ISO as Veeam Backup & Replication.




Veeam Enterprise Manager

Installation

Veeam Backup Enterprise Manager

System Configuration Check

System is being verified for potential installation problems.

 To finalize required components installation, system reboot is required. Installation process will continue automatically.

Requirement	Status
Microsoft .NET Framework 4.7.2	✔ Passed
Microsoft Visual C++ 2015-2019 Redistributable	⚠ Reboot required
Microsoft System CLR Types for SQL Server 2014	✔ Passed
IIS URL Rewrite Module 2	✔ Passed
Microsoft IIS	✔ Passed
Default Document Component	✔ Passed
Directory Browsing Component	✔ Passed
HTTP Errors Component	✔ Passed
Static Content Component	✔ Passed
Windows Authentication Component	✔ Passed
WebSocket Protocol Component	✔ Passed
ASP.NET 4.5 Component	✔ Passed



[Back](#) [Reboot](#) [Cancel](#)

Veeam Backup Enterprise Manager

License

Provide license file for Veeam Backup Enterprise Manager.

Select license provisioning method:

 [Sign in with Veeam](#) |  [Browse license file](#)

License details:
Subscription, Suite, 100 Instances, License expires on 31/12/2024

Update license automatically (enables usage reporting)

Download and install new license automatically when you renew or expand your contract. This requires sending the license ID, the installation ID, and workload usage counters to Veeam servers periodically. Successful usage reporting doubles the number of workloads you can exceed your installed license by.

[Back](#) [Next](#) [Cancel](#)

Veeam Enterprise Manager

Installation

Veeam Backup Enterprise Manager

Ready to Install

Installation will begin with the following settings.

Installation folder:	C:\Program Files\Veeam\Backup and Replication
Guest catalog folder:	C:\VBRCatalog
Service account:	LOCAL SYSTEM
Database engine:	PostgreSQL
Database name:	VeeamBackupReporting
SQL server:	winsrv88:5432
Catalog service port:	9393
Service port:	9394
Web UI ports:	9080 (HTTP), 9443 (HTTPS)
REST API service ports:	9399 (HTTP), 9398 (HTTPS)
Certificate:	Self-signed certificate will be generated automatically
Check for product updates:	Automatically

Customize Settings

Back **Install** Cancel

Veeam Backup Enterprise Manager

Database

Choose database engine and instance for Veeam Backup Enterprise Manager.

Use following database engine: PostgreSQL Server

Install new instance

Use existing instance (HOSTNAME:PORT)

winsrv88:5432

Database name:
VeeamBackupReporting

Connect to PostgreSQL Server using:

Windows authentication credentials of service account

Native authentication using the following credentials:

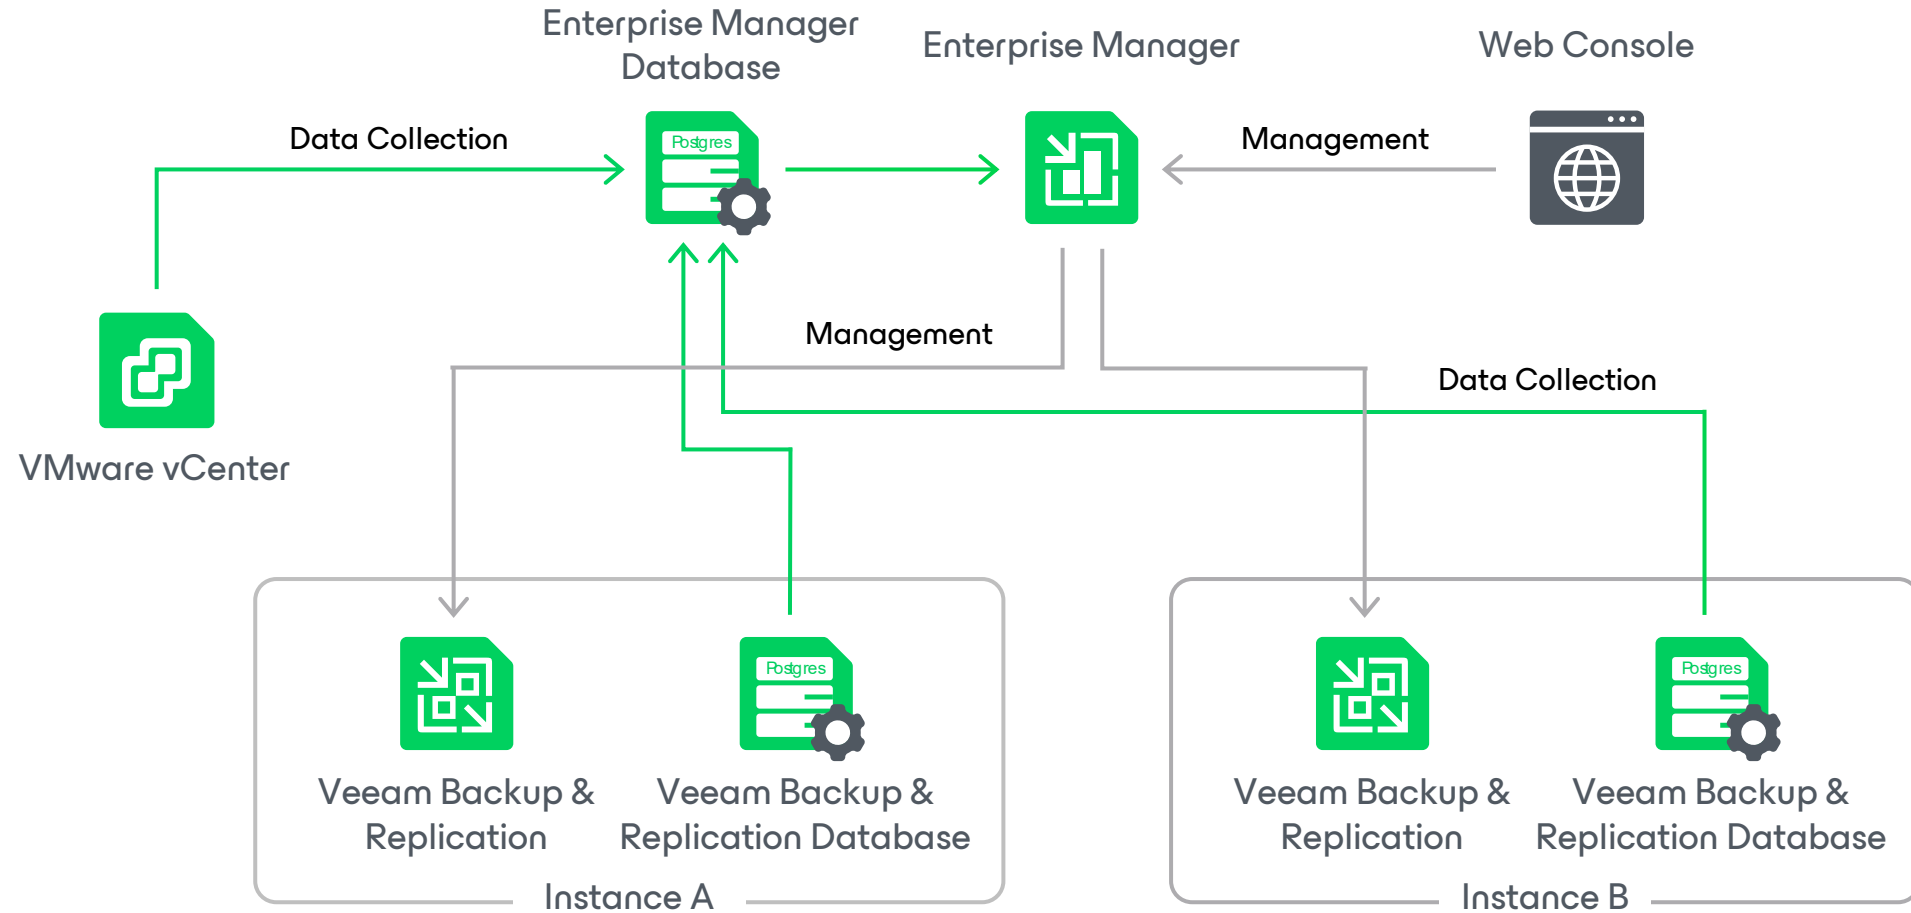
Login: postgres

Password:

Back **Next** Cancel

Veeam Enterprise Manager

Deployment example (logical relationships)



Veeam Enterprise Manager

Capabilities

The screenshot displays the Veeam Enterprise Manager interface. The top navigation bar includes 'Dashboard', 'Reports', 'Jobs', 'Policies', 'Unstructured Data', 'Machines', 'Files', 'Items', and 'Requests'. The user is logged in as 'DEMOCENTER\achraf.akrouf'. The 'Items' tab is active, showing sub-tabs for 'Mailbox Items', 'SQL Database', 'Oracle Database', and 'PostgreSQL Instance'. The 'Mailbox Items' sub-tab is selected. The interface features a search box for 'Username' with the placeholder 'Enter username'. Below it, the 'Accounts' section is empty, with a message: 'Available accounts will be shown after specifying a username'. To the right, there is a 'History' link, a 'Restore point' section with 'No restore point selected', and an 'Items' section with checkboxes for 'Mail', 'Calendar', and 'Contacts'. A checkbox for 'Only restore missing items created or received' is also present, with a 'Yesterday' dropdown menu. A 'Restore' button is located at the bottom.

Veeam Enterprise Manager

Capabilities. Configuration.

The screenshot displays the Veeam Enterprise Manager configuration interface. The top navigation bar includes Dashboard, Reports, Jobs, Policies, Unstructured Data, Machines, Files, Items, and Requests. The user is logged in as DEMOCENTER\achraf.akrou. The left sidebar shows navigation options: Exit Configuration, Backup Servers, vCenter Servers, Self-Service, Sessions, Roles, Settings, Licensing, Notifications, and About. The main area shows a search bar for account names and a list of account types. A modal window titled 'Add' is open, showing fields for Type (User), Account, Repository, Quota, Job scheduling, Job priority, vCenter scope, and vSphere tags. A 'Delegation Mode' dialog is also open, showing three options: vSphere tags (selected), vSphere role, and VM privilege. The background shows a table with columns for Quota and Per-user.

Quota	Per-user
100 GB	No
100 GB	No
200 GB	No
20 GB	No
10 GB	No
100 GB	No
100 GB	No
100 GB	No
100 GB	No
100 GB	No

Veeam Enterprise Manager

Capabilities. Configuration.

Dashboard Reports Jobs **Licenses** Unstructured Data Machines Files Items Requests

Exit Configuration Summary **Instances**

Backup Servers vCenter Ser Self-Servic Sessions Roles Settings **Licensing** Notification About

hq-vbr1.demolab.local (46 instances)

Virtual Machines (22 instances)

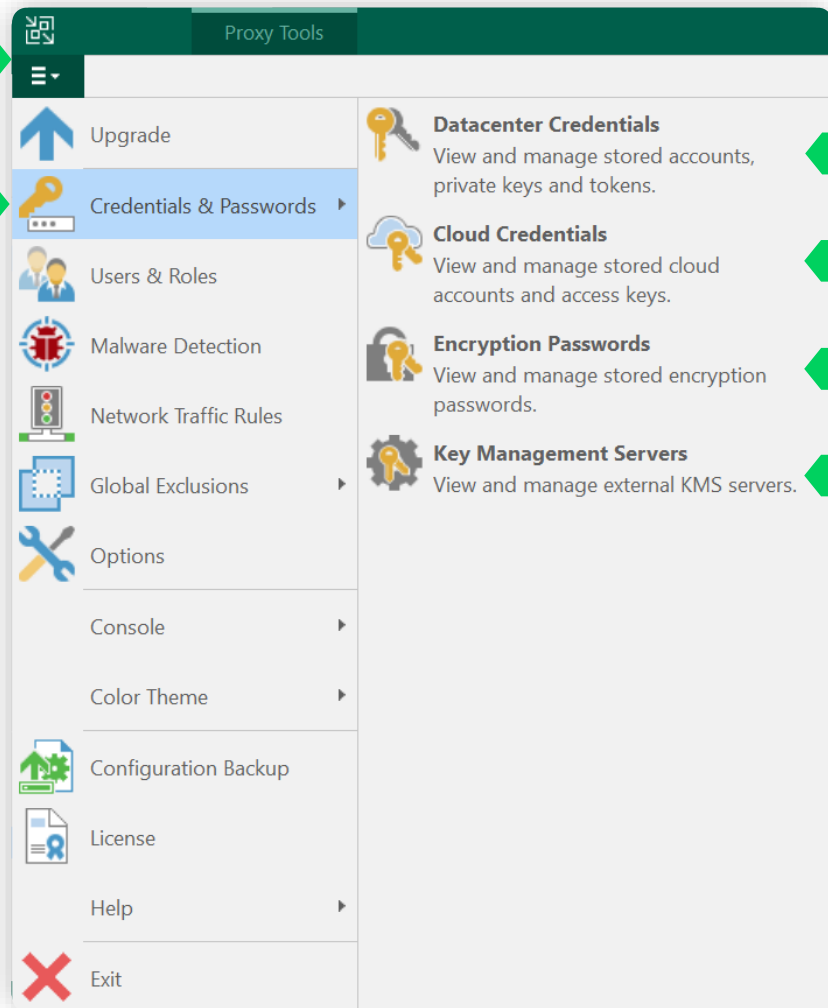
Name	Instances	Type	Job name	Last processed
demo-linuxvm101	1	vSphere	VMware - Backup from 3PAR Snapshot	01/08/2024
demo-linuxvm102	1	vSphere	VMware - Backup from IBM Snapshot	01/08/2024
demo-linuxvm103	1	vSphere	VMware - Backup from NetApp Snapshot	01/08/2024
demo-linuxvm107	1	vSphere	VMware - Backup from VNX Snapshot	01/08/2024
demo-linuxvm109	1	vSphere	VMware - Backup from ETERNUS Snapshot	01/08/2024
demo-sql106	1	vSphere	VMware - Backup from Pure Snapshot	01/08/2024
demo-winvm101	1	vSphere	VMware - Backup from 3PAR Snapshot	01/08/2024
demo-winvm102	1	vSphere	VMware - Backup from IBM Snapshot	01/08/2024
demo-winvm103	1	vSphere	VMware - Backup from NetApp Snapshot	01/08/2024
demo-winvm105	1	vSphere	VMware - Backup from Nimble Snapshot	01/08/2024
demo-winvm106	1	vSphere	VMware - Backup from Pure Snapshot	01/08/2024
demo-winvm107	1	vSphere	VMware - Backup from VNX Snapshot	01/08/2024
demo-winvm108	1	vSphere	VMware - Backup from VSAN	01/08/2024
demo-winvm109	1	vSphere	VMware - Backup from ETERNUS Snapshot	01/08/2024
exch1	1	vSphere	VMware - Backup to XFS Repository	01/08/2024
fc1	1	vSphere	VMware - Backup to XFS Repository	01/08/2024

Install license

Main Menu

Main Menu

Credentials & Passwords

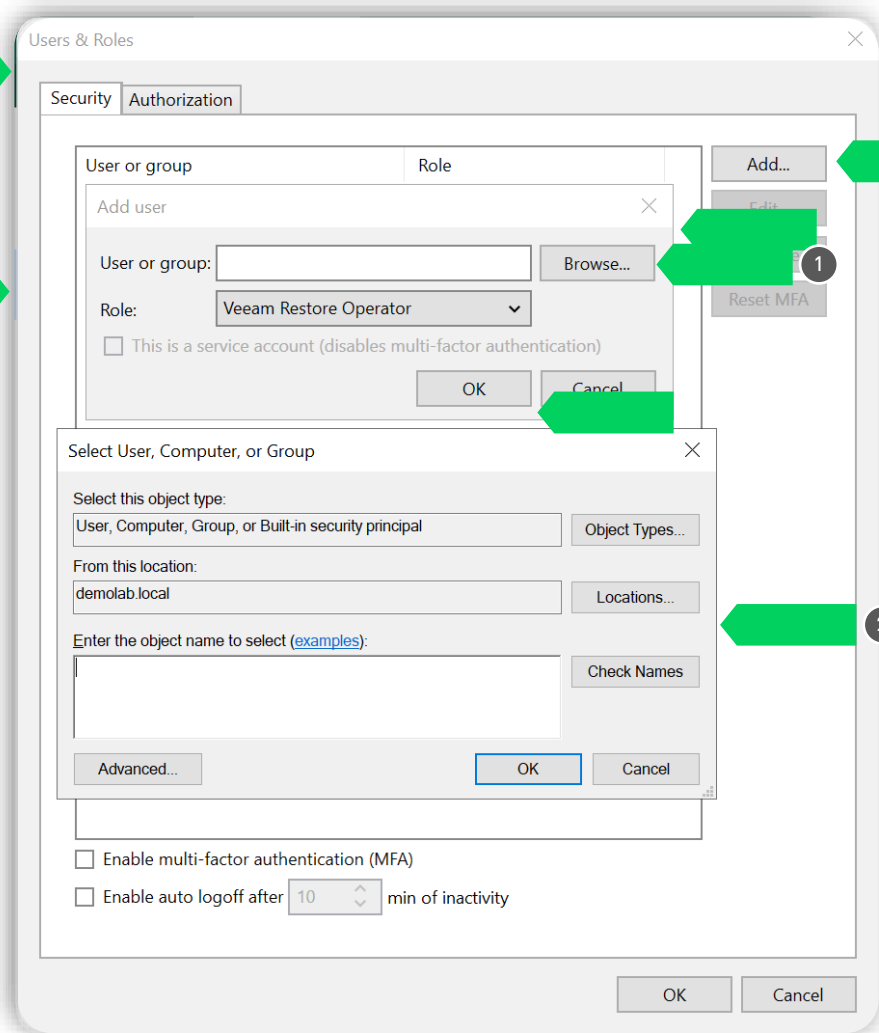


You can use this option to create and maintain a list of credentials records that you plan to use to connect to cloud services.

You can use this option to create and maintain a list of passwords that you plan to use for data encryption. You can use this option to add KMS servers and use KMS keys.

Main Menu

Users & Roles



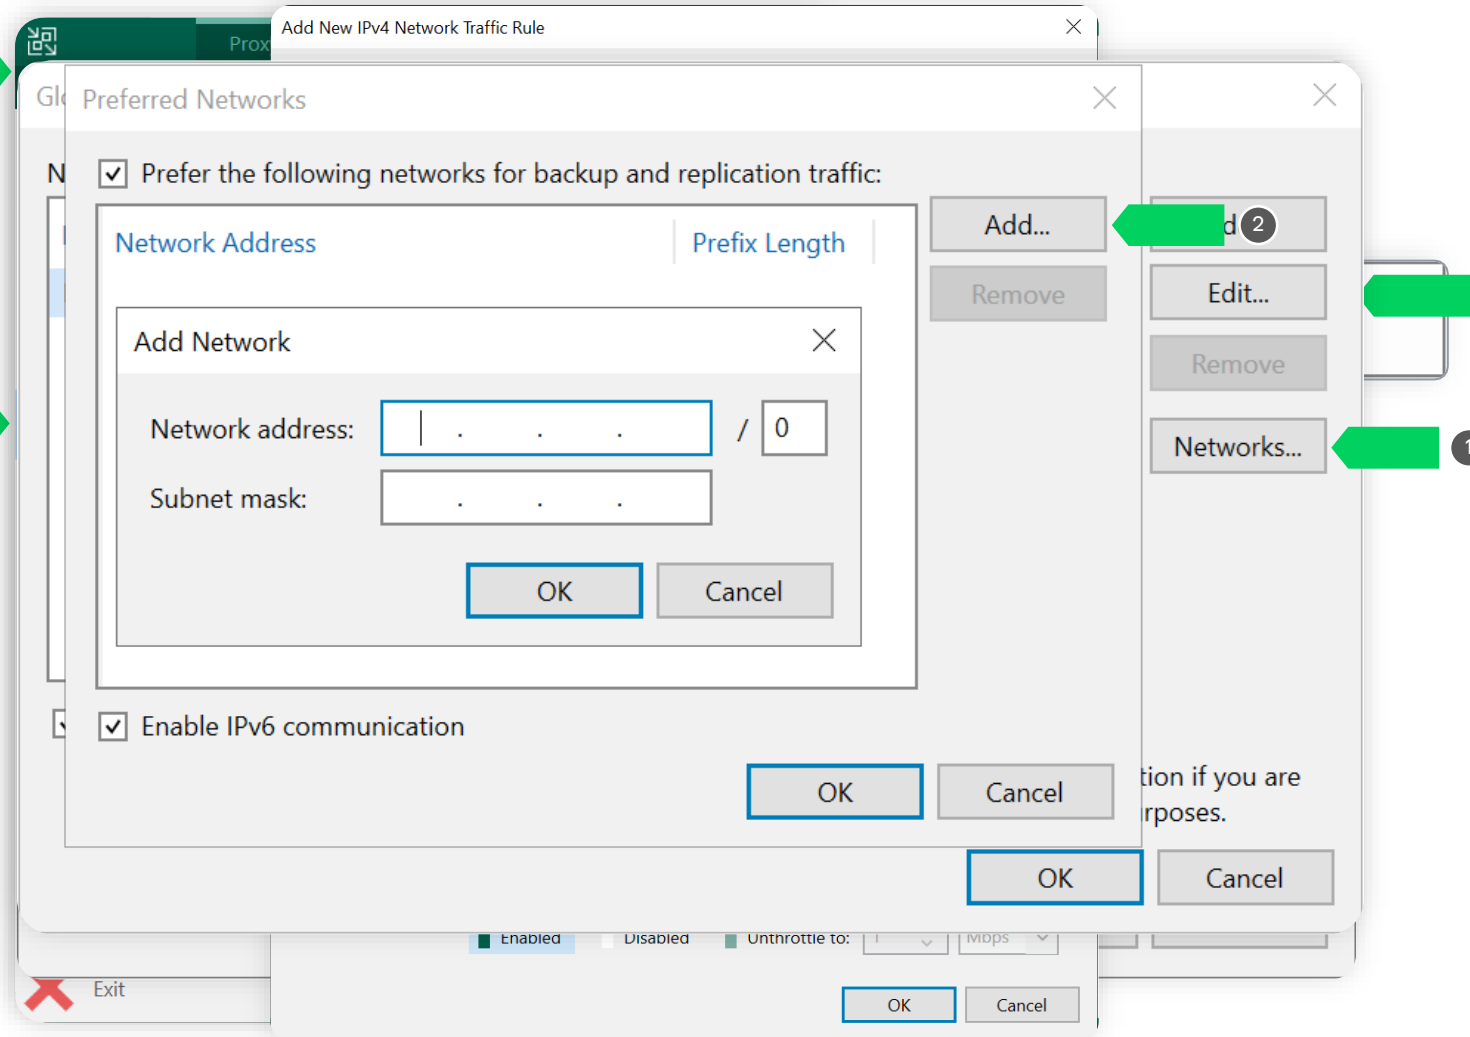
Editing of all existing users or user groups with a new specifying of the role. You can add new users or user groups. Roles can be observed; for notes will define the Veeam

Backup & Replication operations that they are entitled to perform. When it comes to roles, there are five options we can go for:

- **Restore operator** – which allows restoration processes to be performed by the specific user/user group
- **Backup operator** – responsible for starting and stopping backup jobs, as well as exporting and copying backups and creating Veeam ZIP backups
- **Backup administrator** – able to perform all administrative activities and with full access to everything added to the backup infrastructure
- **Tape operator** – responsible for any tape operation and starting or stopping tape jobs
- **Backup viewer** – which has the “read-only” access to the console

Main Menu

Global Network Rules

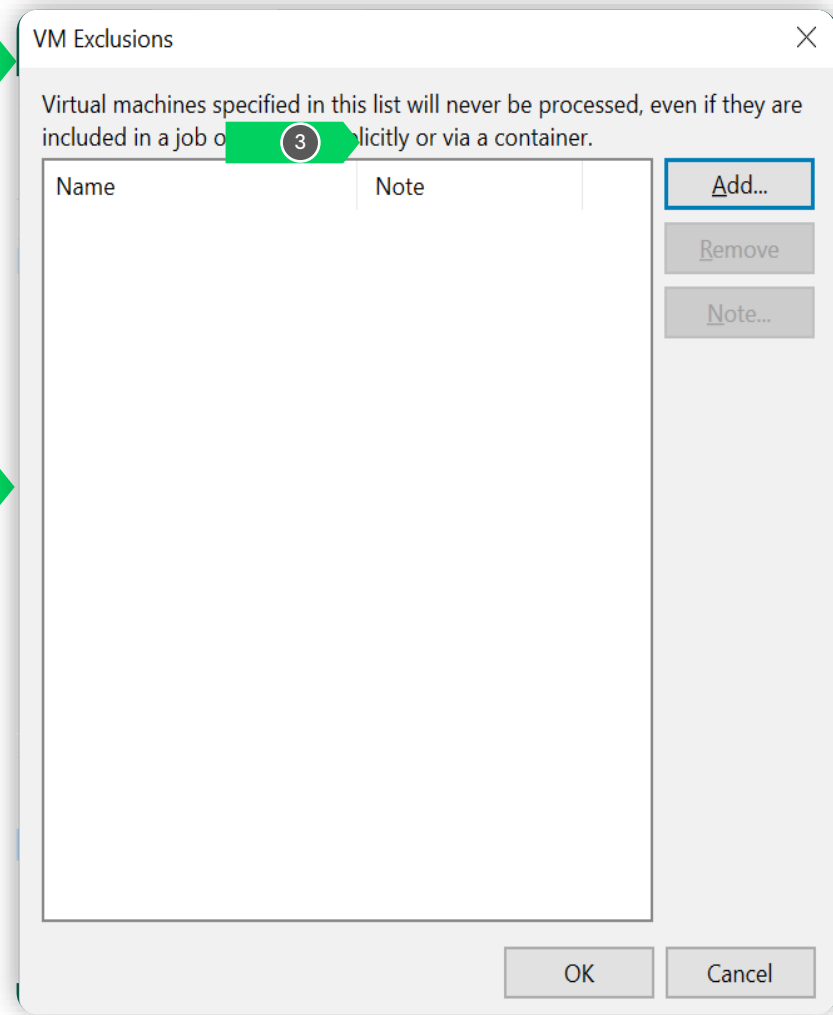


The results of the impossibility to replicate refer to the possibility to add different IPv4/6 rules by specifying the range for the data transfer replication profession, source and target and enforcing either encryption of the network TCP/IP connections or throttling it during specific periods.

1 However, if you schedule several jobs to run at the same time, load on the network may be heavy. If the network capacity is not sufficient to support multiple data transfer connections, you can disable multithreaded data transfer or change the number of TCP/IP connections.

Main Menu

VM Exclusions



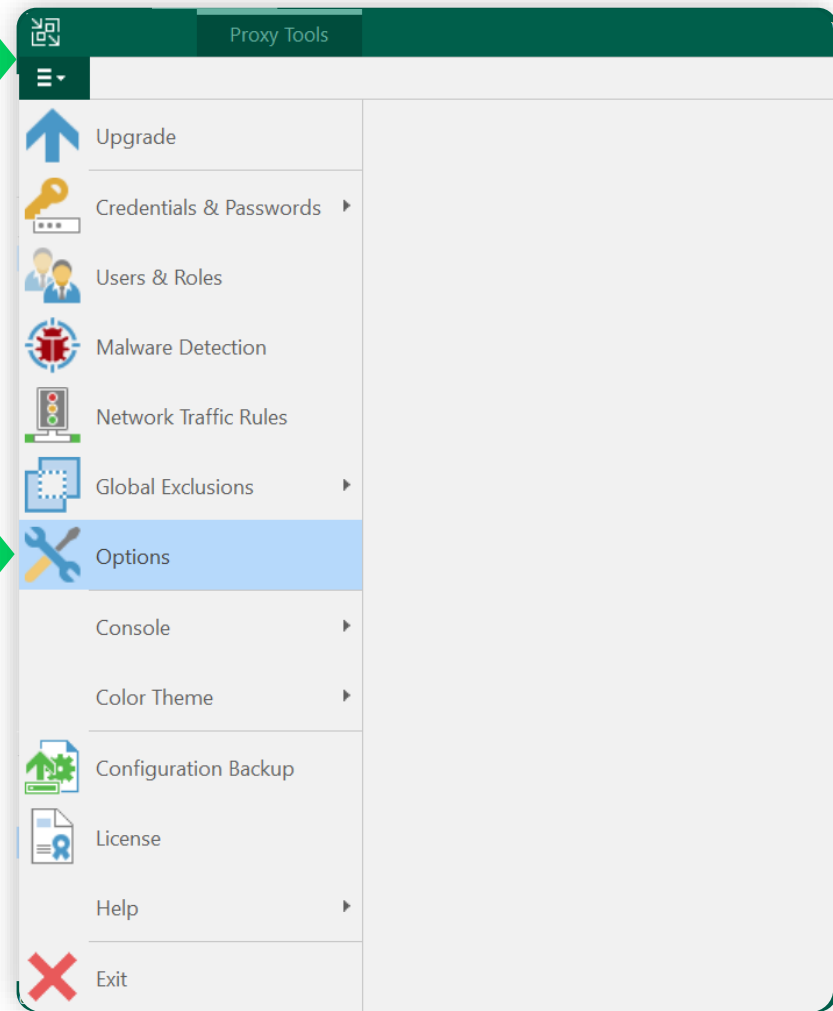
Veeam Backup & Replication offers the possibility to add virtual machines to an exclusion list.

VMs included in this list will not be processed in any type of jobs, except backup copy jobs and SureBackup jobs.

In order to resume the processing of the VMs, they need to be removed from this list first.

Main Menu

Options



Main Menu

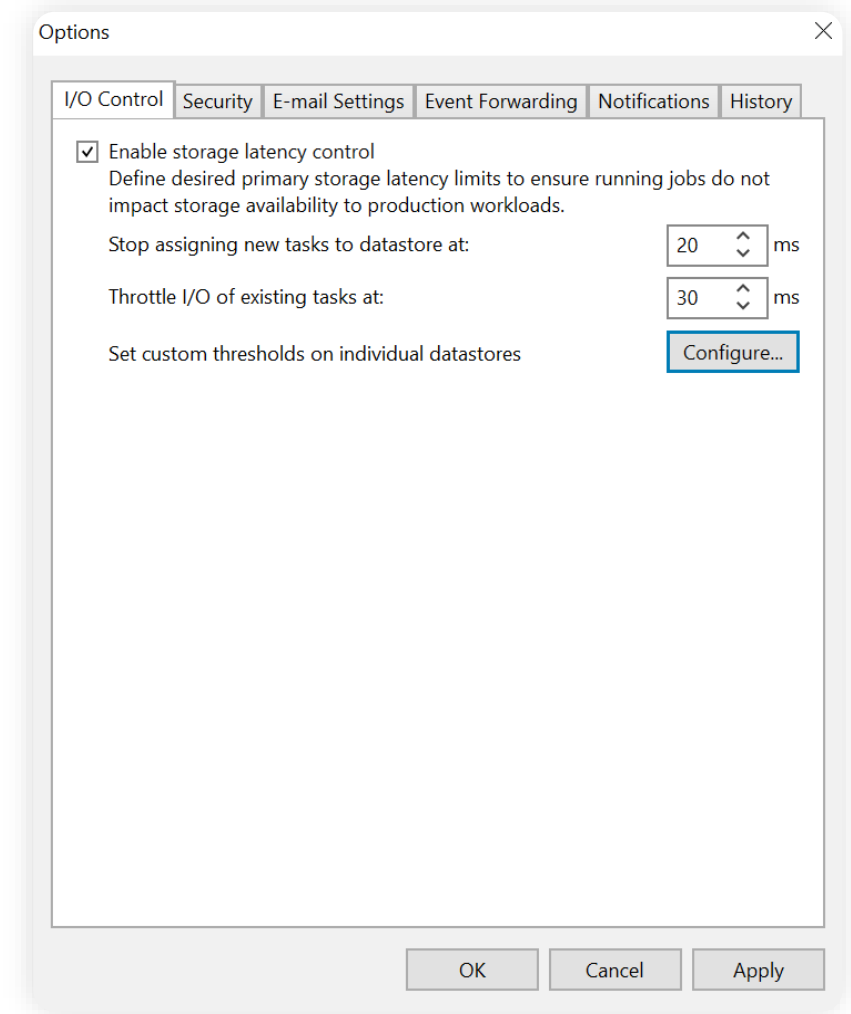
Options. I/O Control.

When storage latency control is enabled, Veeam Backup & Replication monitors storage read latency on production datastores during data protection and disaster recovery activities. To monitor the storage latency, Veeam Backup & Replication uses real-time metrics from the hypervisor where VMs reside.

Stop assigning new tasks to datastore at – the I/O latency limit at which Veeam Backup & Replication must not assign new tasks targeted at the datastore.

Throttle I/O of existing tasks at – the I/O latency limit at which Veeam Backup & Replication must decrease the speed of data retrieval or writing to/from the datastore.

Same options can be configured for individual datastores.



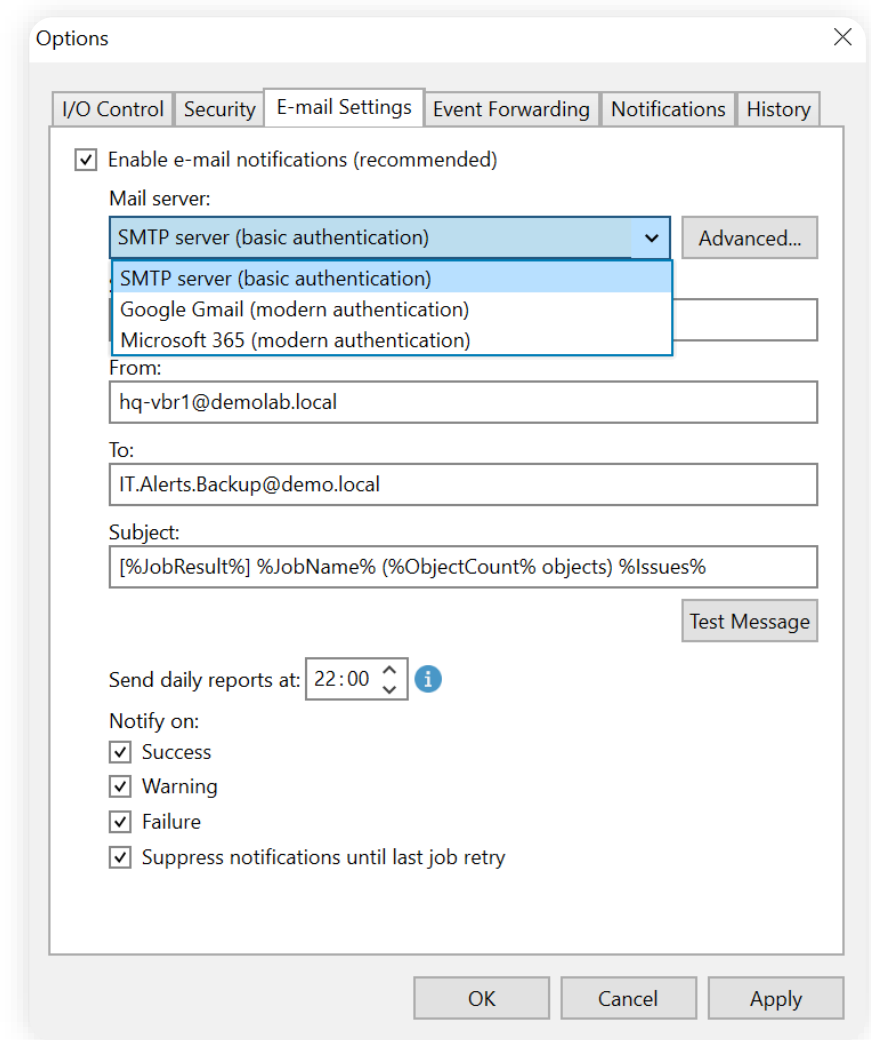
Main Menu

Options. E-mail Settings.

Veeam Backup & Replication also offers the possibility to send e-mails that contain daily reports on the last status of the backup jobs.

At the moment, we support SMTP, Google Gmail and Microsoft 365 as the mail servers that can be used for this purpose.

The subject of the e-mail, the recipients, the e-mail address of the sender and what the report should contain in terms of the last status that should be included can all be configured from this tab.

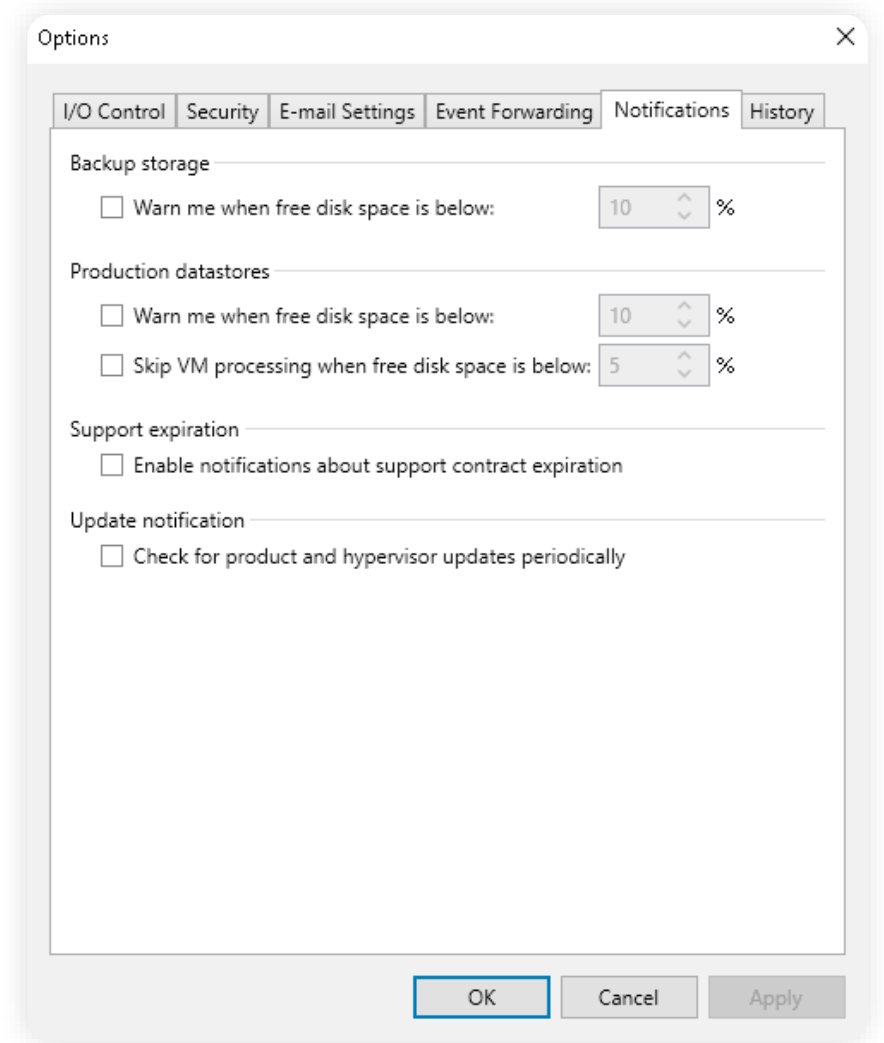


The screenshot shows the 'Options' dialog box with the 'E-mail Settings' tab selected. The 'Enable e-mail notifications (recommended)' checkbox is checked. The 'Mail server' dropdown menu is open, showing options: 'SMTP server (basic authentication)', 'SMTP server (basic authentication)', 'Google Gmail (modern authentication)', and 'Microsoft 365 (modern authentication)'. The 'From:' field contains 'hq-vbr1@demolab.local', the 'To:' field contains 'IT.Alerts.Backup@demo.local', and the 'Subject:' field contains '[%JobResult%] %JobName% (%ObjectCount% objects) %Issues%'. There is a 'Test Message' button. The 'Send daily reports at:' field is set to '22:00' with an information icon. The 'Notify on:' section has four checked options: 'Success', 'Warning', 'Failure', and 'Suppress notifications until last job retry'. At the bottom are 'OK', 'Cancel', and 'Apply' buttons.

Main Menu

Options. Notifications.

Veeam Backup & Replication can also display notifications directly inside the console whenever thresholds are hit on backup storage on the repositories or productions datastores, there are 14 days left before the end of the support for the installed license and update for the hypervisors are available.




Main Menu

Options. History.

The **Sessions** area controls how many entries should be visible in the History tab of the Veeam Backup & Replication console.

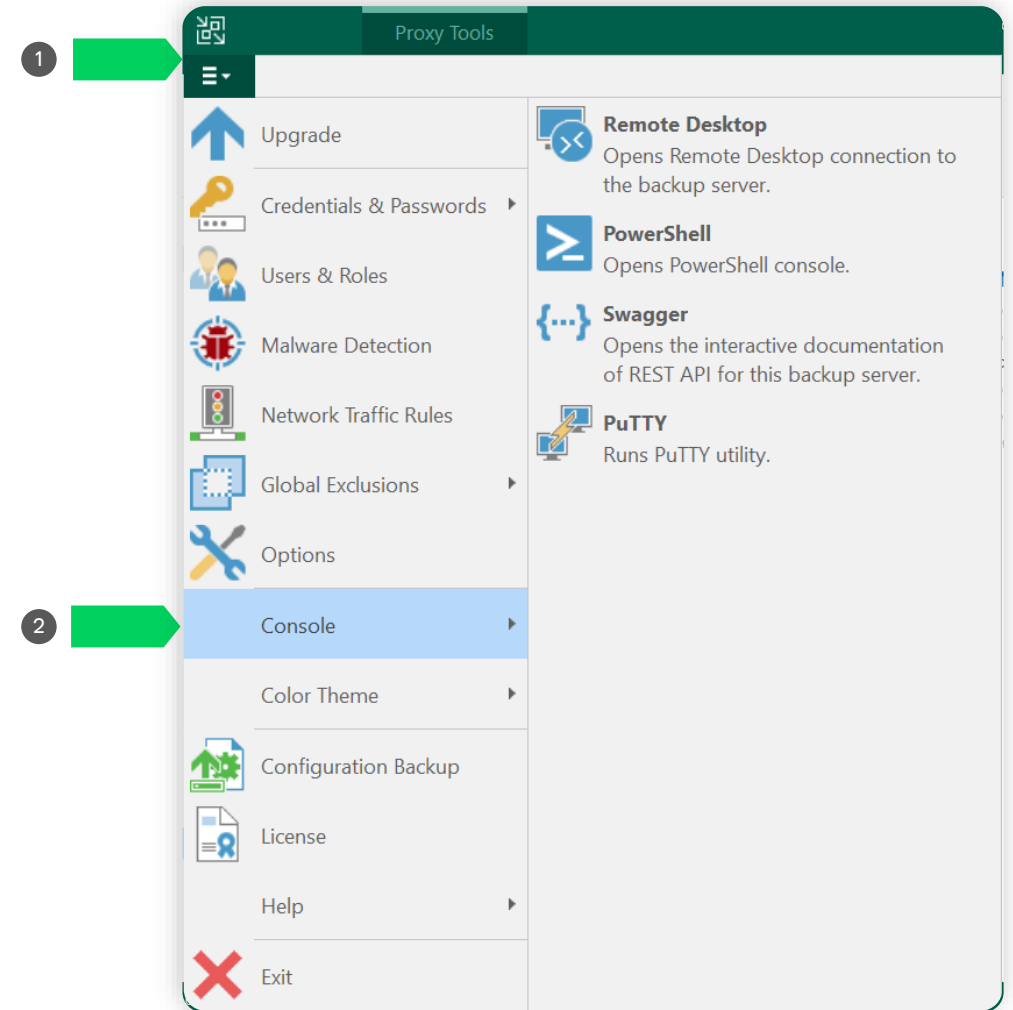
The **Session history retention** area dictates for how long session information should be kept in the database.



Main Menu

Console.

Veeam Backup & Replication offers the possibility to open a Remote Desktop connection to the backup server, open the PowerShell console, run the PuTTY utility and open the interactive documentation or REST APIs for the backup server directly from the console.



Main Menu

Configuration Backup.

1

Configuration Backup Settings

Backup

Enable configuration backup to the following repository:

Configuration Backup Repository (Created by Ansible)

7.9 GB free of 10 GB

Restore points to keep: 10

Perform backup on: Daily at 10:00

Last successful backup: No backup

Enable configuration backup file encryption

Password: Configuration backup password (Last edited: 284 days ago)

Loss protection enabled [Manage passwords](#)

2

Restore

Restore the configuration backup to this server:

OK Cancel Apply

Veeam Backup & Replication backs up the configuration to a repository that is stored inside a console on a specific server as a configuration backup. configuration backup can be sent

automatically.

If the main Veeam Backup & Replication server fails for some reason, the backup to be performed automatically. Another good practice is to protect the configuration through encryption. If the encryption password is lost, only Veeam Backup Enterprise Manager can assist in retrieving it. This method is also advisable if the Veeam Backup & Replication server should be replaced or changed in the future.


configuration backup can be performed only from this window in order to upload the configuration to a new server.

Jobs

The concept of the Jobs

In Veeam Backup and Replication, jobs encompass a range of data management tasks, including **backup**, **replication**, and **restore** operations. Once configured, jobs will automate the process, ensuring that critical data is regularly and reliably protected against loss or corruption.

When it comes to backup infrastructure resource allocation, **the Resource Scheduler Service** within Veeam Backup & Replication identifies the priority of the jobs awaiting free resources based on their type, Priority, and scheduled start time.



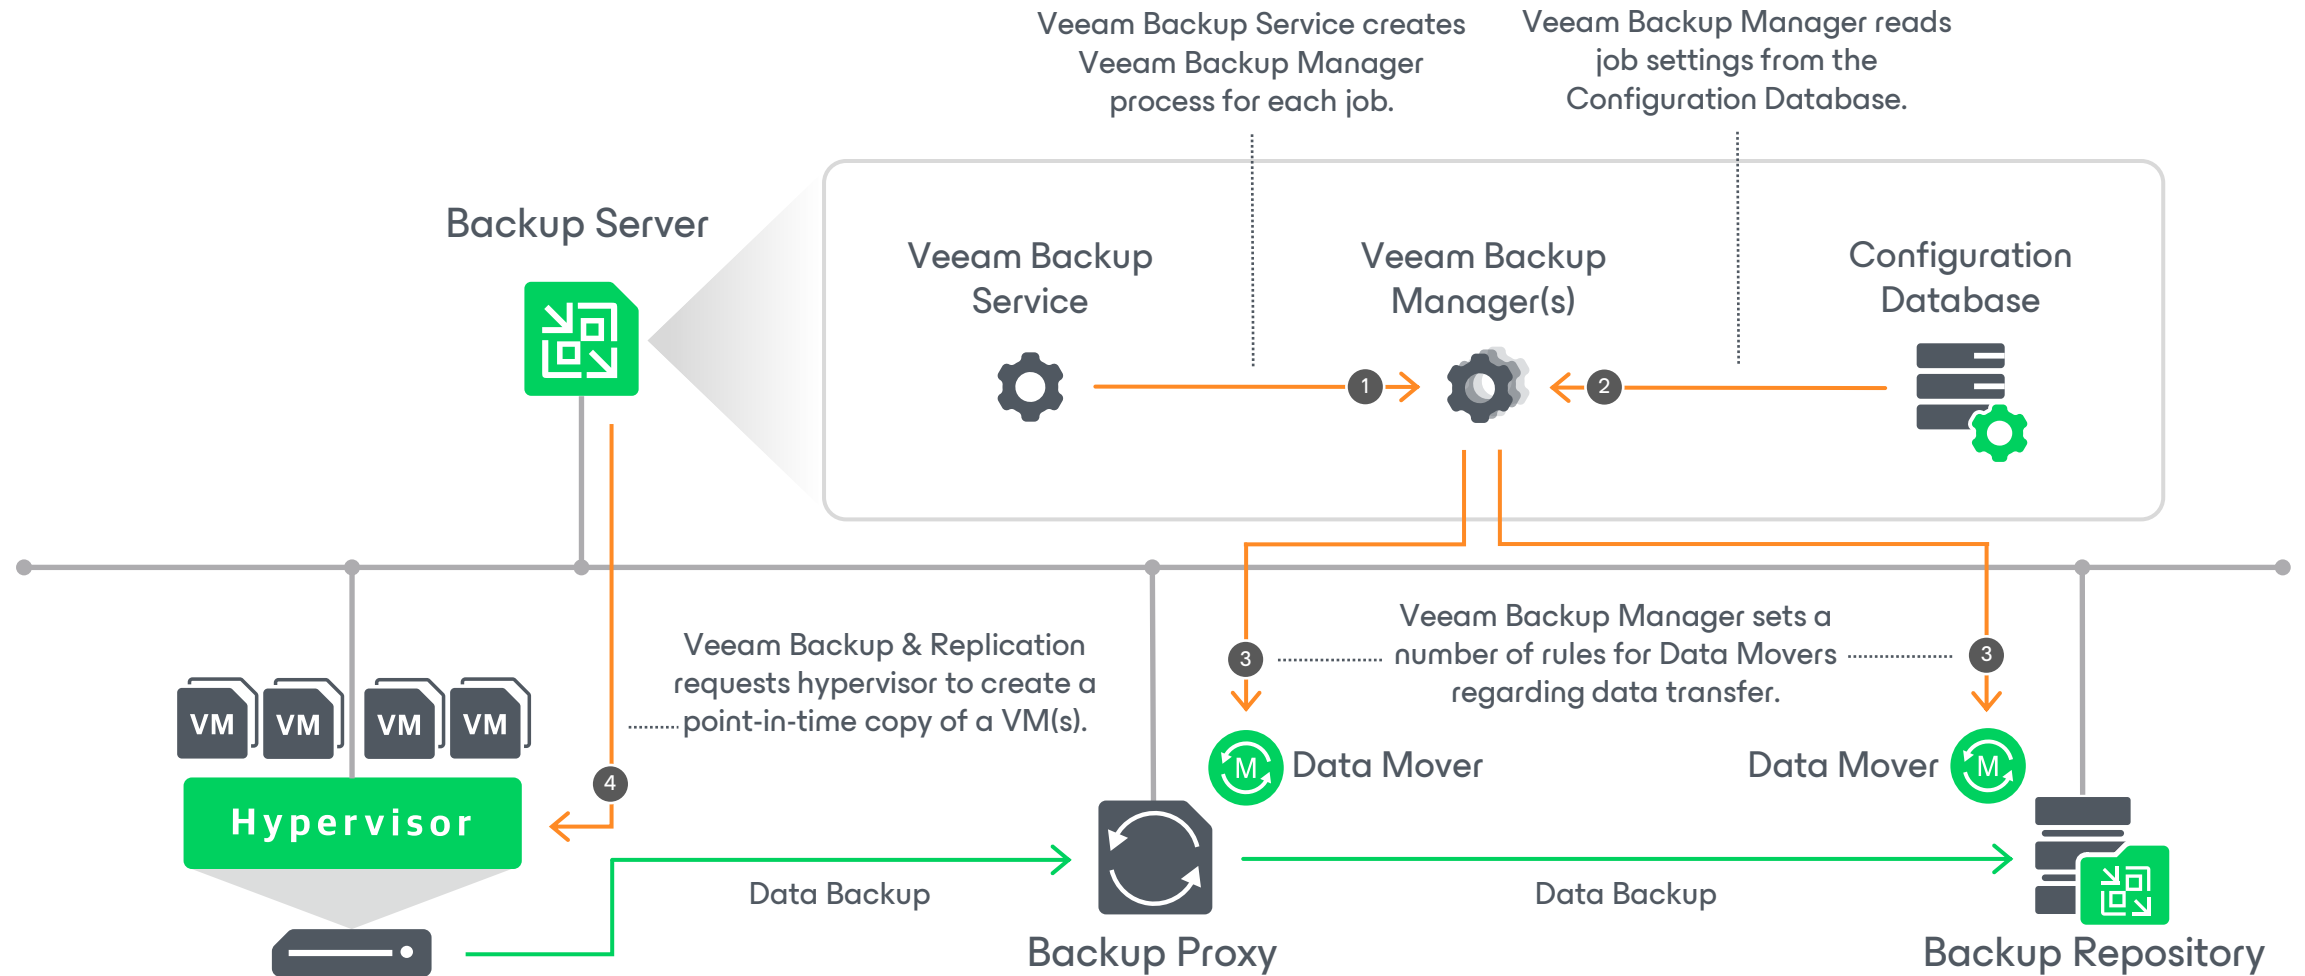
Jobs prioritization

Here is a **simplified list** in order from high-priority to low-priority:

1. **Data Restore jobs** — these jobs have the highest priority and are the first to get free system resources.
2. **Continuous Data Protection policies.**
3. **Background system jobs** (SnapshotDeleter for example).
4. **Quick Backup jobs** — these jobs have the second highest priority.
5. **High Priority jobs** — jobs with the enabled High Priority option have the third highest priority. It's possible to enable the High priority option for the following jobs: backup jobs, replication jobs, agent jobs managed by backup server, file backup jobs.
6. **Regular Backup jobs and Replication jobs.**
7. **Backup Copy jobs, and Archive jobs** — these jobs have the lowest priority and are the last to get free system resources.

Jobs


Example of a backup process



Jobs

Example of a backup job creation process

New Backup Job

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

Name	Name: Backup Job Example
Virtual Machines	Description: Created by DEMOCENTER\User at 2024-04-09 15:54.
Storage	
Guest Processing	
Schedule	
Summary	

High priority
Backup infrastructure resources are offered to high priority jobs first. Use this option for jobs sensitive to the start time, or jobs with strict RPO requirements.

< Previous Next > Finish Cancel

Jobs

Example of a backup job creation process

New Backup Job

Virtual Machines
Select virtual machines to process via container, or granularly. Container provides dynamic selection that automatically changes as you add new VMs into the container.

Virtual machines to backup:

Name	Type	Size	
			Add...
			Remove
			Exclusions...
			Up
			Down
			Recalculate

Total size: 0 B

< Previous Next > Finish Cancel

Add Objects

Select objects:


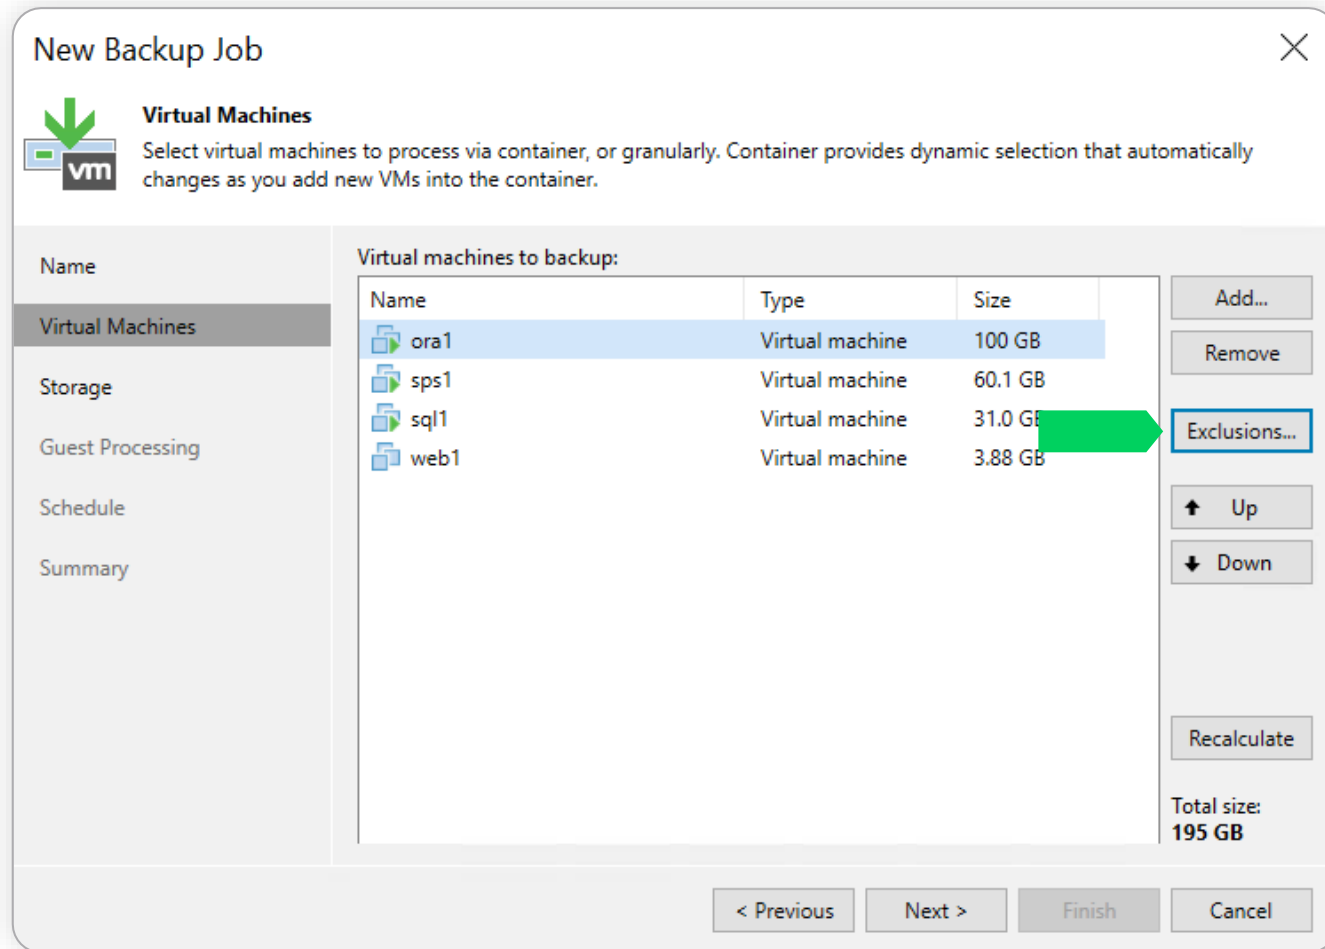
- VMs and Tags
 - vc1.democenter.int
 - vc1.demolab.local
 - AA VMs with No Backup -Business-View
 - Application
 - ASH-TEST01
 - VAO Recovery Location
 - DR Site
 - HQ Site

Type in an object name to search for

Add Cancel

Jobs

Example of a backup job creation process



Jobs

Example of a backup job creation process

New Backup Job

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

Name Backup proxy: Automatic selection

Virtual Machines

Storage AWS S3 Repository (Created by Powershell at 7/17/2020 1:05:23 PM.)
1.63 TB free of 1.99 TB

Guest Processing

Schedule Retention policy: 14 days

Summary Keep certain full backups longer for archival purposes 12 monthly
 Configure secondary destinations for this job
Copy backups produced by this job to another backup repository, or tape. We recommend to make at least one copy of your backups to a different storage device that is located off-site.

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

Backup Proxy

Choose backup proxy servers for this job. For redundancy, we recommend to select at least two proxies. When multiple proxies are available, selection will be performed on per-VM basis, taking into account proxy connectivity and current load.

Automatic selection
The job will automatically select the most suitable backup proxy server from all available backup proxy servers.


Use the selected backup proxy servers only
The job will automatically select the most suitable backup proxy server from the following list of proxy servers.

Name	<input type="checkbox"/>	<input type="checkbox"/>
hq-vbr1appl1.demolab.local	<input type="checkbox"/>	<input checked="" type="checkbox"/>
hq-vbr1lnxpxy1.demolab.local	<input type="checkbox"/>	<input checked="" type="checkbox"/>
hq-vbr1pxy1.demolab.local	<input type="checkbox"/>	<input checked="" type="checkbox"/>
hq-vbr1pxy3.demolab.local	<input type="checkbox"/>	<input checked="" type="checkbox"/>

Jobs

Example of a backup job creation process

New Backup Job

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

Name

Virtual Machines

Storage

Guest Processing

Schedule

Summary

Backup proxy:
Automatic selection Choose...

Backup repository:
AWS S3 Repository (Created by Powershell at 7/17/2020 1:05:23 PM.) ▼

1.63 TB free of 1.99 TB Map backup

Retention policy: 14 ▼ days !

Keep certain full backups longer for archival purposes Configure...
12 monthly

Configure secondary destinations for this job
Copy backups produced by this job to another backup repository, or tape. We recommend to make at least one copy of your backups to a different storage device that is located off-site.

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

< Previous Next > Finish Cancel

Advanced Settings

Backup Maintenance Storage Notifications vSphere Integration Scripts

Backup mode

Reverse incremental (slower)
Increments are injected into the full backup file, so that the latest backup file is always a full backup of the most recent VM state.

Incremental (recommended)
Increments are saved into new files dependent on previous files in the chain. Best for backup targets with poor random I/O performance.

Create synthetic full backups periodically on:
Saturday Configure...

Active full backup

Create active full backups periodically on:
Saturday Configure...

Reverse incremental mode will become deprecated in vNext. It will continue to work but will not be available for newly created jobs.

Save As Default OK Cancel

Jobs

Example of a backup job creation process

New Backup Job

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

Name

Virtual Machines

Storage

Guest Processing

Schedule

Summary

Backup proxy:
Automatic selection Choose...

Backup repository:
AWS S3 Repository (Created by Powershell at 7/17/2020 1:05:23 PM.)
1.63 TB free of 1.99 TB Map backup

Retention policy: 14 days !

Keep certain full backups longer for archival purposes
12 monthly Configure...

Configure secondary destinations for this job
Copy backups produced by this job to another backup repository, or tape. We recommend to make at least one copy of your backups to a different storage device that is located off-site.

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

< Previous Next > Finish Cancel

Advanced Settings

Backup Maintenance Storage Notifications vSphere Integration Scripts

Storage-level corruption guard

Perform backup files health check (detects and auto-heals corruption) on:
At 22:00 on Last Friday of every month Configure...

Full backup file maintenance

Remove deleted items data after 3 days


Defragment and compact full backup file on:
Last Saturday of every month Configure...

Save As Default OK Cancel

Jobs

Example of a backup job creation process

New Backup Job

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

Name

Virtual Machines

Storage

Guest Processing

Schedule

Summary

Backup proxy:
Automatic selection

Backup repository:
AWS S3 Repository (Created by Powershell at 7/17/2020 1:05:23 PM.)

1.63 TB free of 1.99 TB

Retention policy: 14 days

Keep certain full backups longer for archival purposes
12 monthly

Configure secondary destinations for this job
Copy backups produced by this job to another backup repository, or tape. We recommend to make at least one copy of your backups to a different storage device that is located off-site.

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

Advanced Settings

Backup Maintenance **Storage** Notifications vSphere Integration Scripts

Data reduction

Enable inline data deduplication (recommended)

Exclude swap file blocks (recommended)

Exclude deleted file blocks (recommended)

Compression level:
Optimal (recommended)
Provides for the best compression to performance ratio, lowest backup proxy CPU usage and fastest restore.

Storage optimization:
1MB (recommended)
Delivers the optimal combination of backup speed, granular restore performance and repository space consumption.

Encryption

Enable backup file encryption
Password:

Jobs

Example of a backup job creation process

New Backup Job

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

Name
Virtual Machines

Storage
Backup proxy: Automatic selection Choose...
Backup repository: AWS S3 Repository (Created by Powershell at 7/17/2020 1:05:23 PM.)
1.63 TB free of 1.99 TB Map backup
Retention policy: 14 days !
 Keep certain full backups longer for archival purposes 12 monthly Configure...
 Configure secondary destinations for this job
Copy backups produced by this job to another backup repository, or tape. We recommend to make at least one copy of your backups to a different storage device that is located off-site.

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

< Previous Next > Finish Cancel

Advanced Settings

Backup Maintenance Storage Notifications vSphere Integration Scripts

Send SNMP notifications for this job

Send e-mail notifications to the following recipients:
Type in one or more e-mail addresses separated by semicolon

Use global notification settings

Use custom notification settings specified below:
Subject: [%JobResult%] %JobName% (%ObjectCount% machines) %Issues%

Notify on success

Notify on warning

Notify on error

Suppress notifications until the last retry

Set successful backup details to this VM attribute:
Notes

Append to the existing attribute's value

Save As Default OK Cancel

Jobs

Example of a backup job creation process

New Backup Job

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

Name

Virtual Machines

Storage

Guest Processing

Schedule

Summary

Backup proxy:
Automatic selection Choose...

Backup repository:
AWS S3 Repository (Created by Powershell at 7/17/2020 1:05:23 PM.) Map backup

1.63 TB free of 1.99 TB

Retention policy: 14 days !

Keep certain full backups longer for archival purposes
12 monthly Configure...

Configure secondary destinations for this job
Copy backups produced by this job to another backup repository, or tape. We recommend to make at least one copy of your backups to a different storage device that is located off-site.

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

< Previous Next > Finish Cancel

Advanced Settings

Backup Maintenance Storage Notifications vSphere Integration Scripts

Guest quiescence

Enable VMware Tools quiescence
Native quiescence will only be used for virtual machines with application-aware image processing disabled.

Changed block tracking

Use changed block tracking data (recommended)

Enable CBT for all protected VMs automatically
Changed Block Tracking (CBT) is a VMware feature that allows for faster incremental backup and replication. For CBT to be enabled, VM must have no existing snapshots.

Reset CBT on each Active Full backup automatically
Provides safety net against potential CBT issues at the cost of slightly increased backup window.

Save As Default OK Cancel

Jobs

Example of a backup job creation process

New Backup Job

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

Name

Virtual Machines

Storage

Guest Processing

Schedule

Summary

Backup proxy:
Automatic selection Choose...

Backup repository:
AWS S3 Repository (Created by Powershell at 7/17/2020 1:05:23 PM.) Map backup

1.63 TB free of 1.99 TB

Retention policy: 14 days !

Keep certain full backups longer for archival purposes
12 monthly Configure...

Configure secondary destinations for this job
Copy backups produced by this job to another backup repository, or tape. We recommend to make at least one copy of your backups to a different storage device that is located off-site.

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

< Previous Next > Finish Cancel

Advanced Settings

Backup Maintenance Storage Notifications vSphere Integration Scripts

Primary storage integration


- Enable backup from storage snapshots
Use storage snapshots (instead of VM snapshots) as the data source for this job. Using storage snapshots reduces impact on the production environment from VM snapshot commit.
- Limit processed VM count per storage snapshot to 10
- Failover to standard backup
Perform standard backup from VM snapshot if backup from storage snapshot fails.
- Failover to primary storage snapshot
Use primary storage snapshots as the data source if backup from secondary storage snapshot fails.

Save As Default OK Cancel

Jobs

Example of a backup job creation process

New Backup Job

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

Name

Virtual Machines

Storage

Guest Processing

Schedule

Summary

Backup proxy:
Automatic selection Choose...

Backup repository:
AWS S3 Repository (Created by Powershell at 7/17/2020 1:05:23 PM.)
1.63 TB free of 1.99 TB Map backup

Retention policy: 14 days !

Keep certain full backups longer for archival purposes
12 monthly Configure...

Configure secondary destinations for this job
Copy backups produced by this job to another backup repository, or tape. We recommend to make at least one copy of your backups to a different storage device that is located off-site.

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

< Previous Next > Finish Cancel

Advanced Settings

Backup Maintenance Storage Notifications vSphere Integration Scripts

Job scripts

Run the following script before the job:
 Browse...

Run the following script after the job:
 Browse...

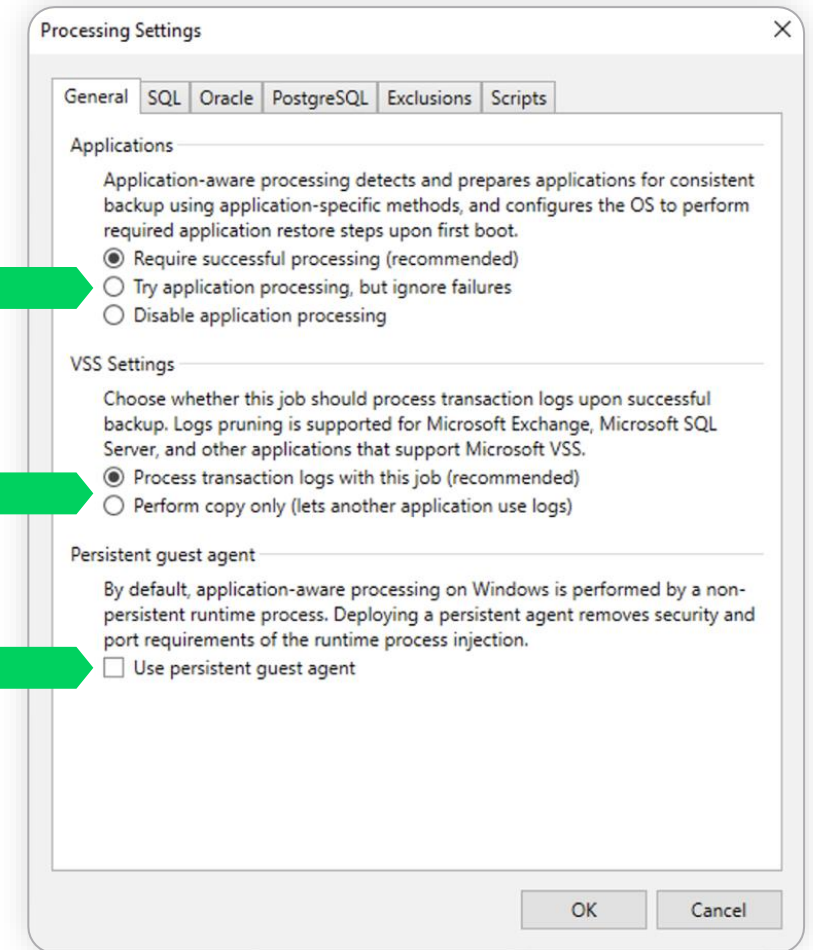
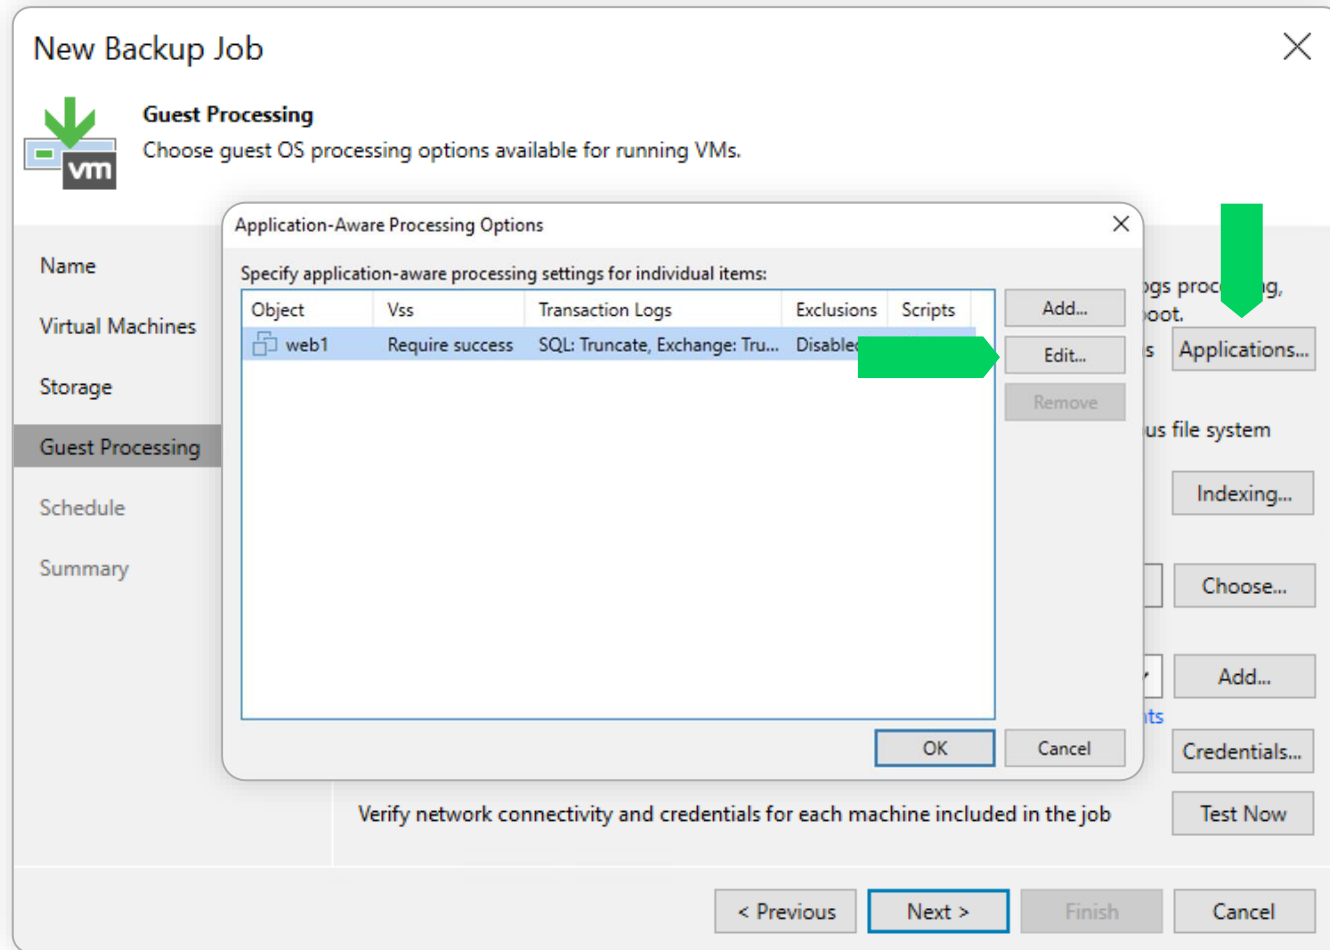
Run scripts every 1 backup session

Run scripts on the selected days only
Saturday Days...

Save As Default OK Cancel

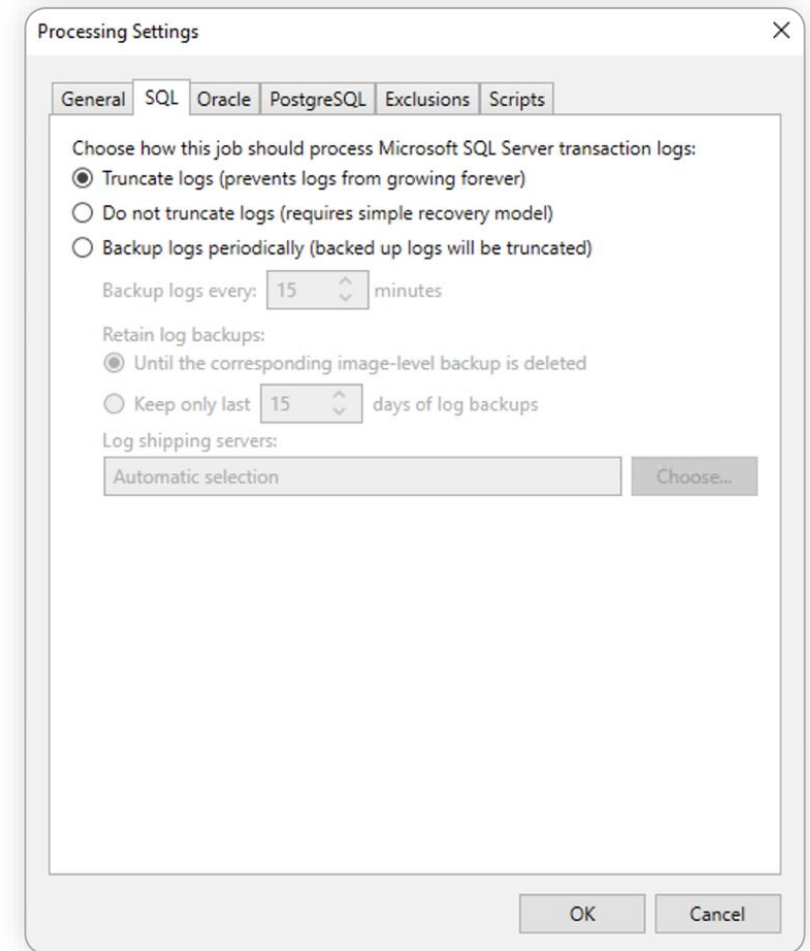

Jobs

Example of a backup job creation process



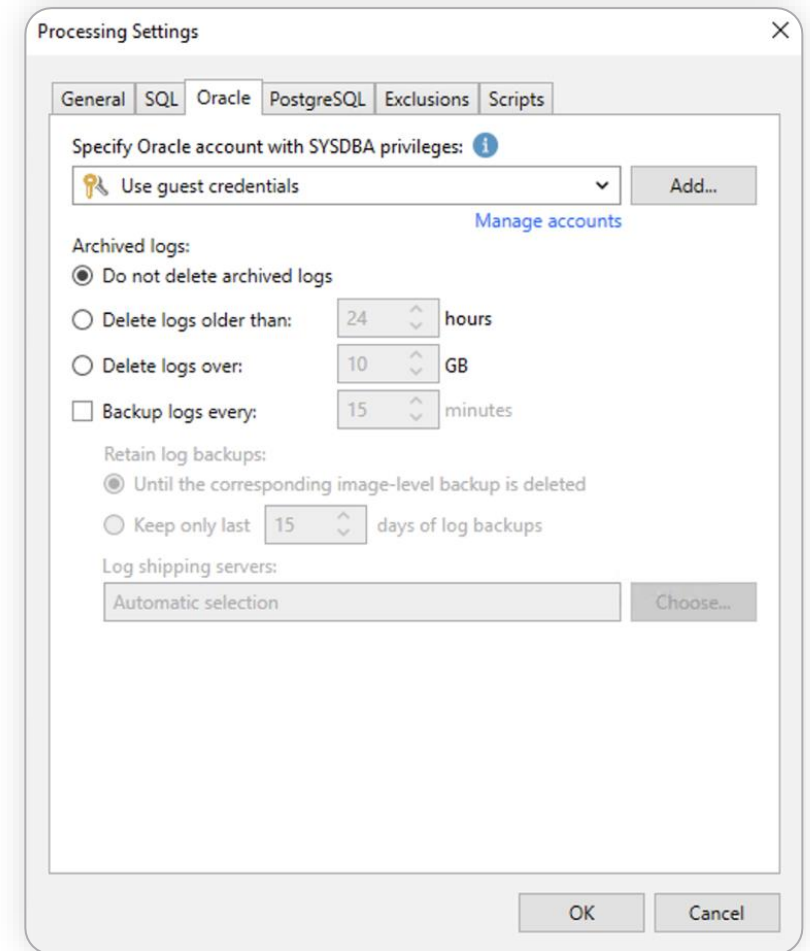
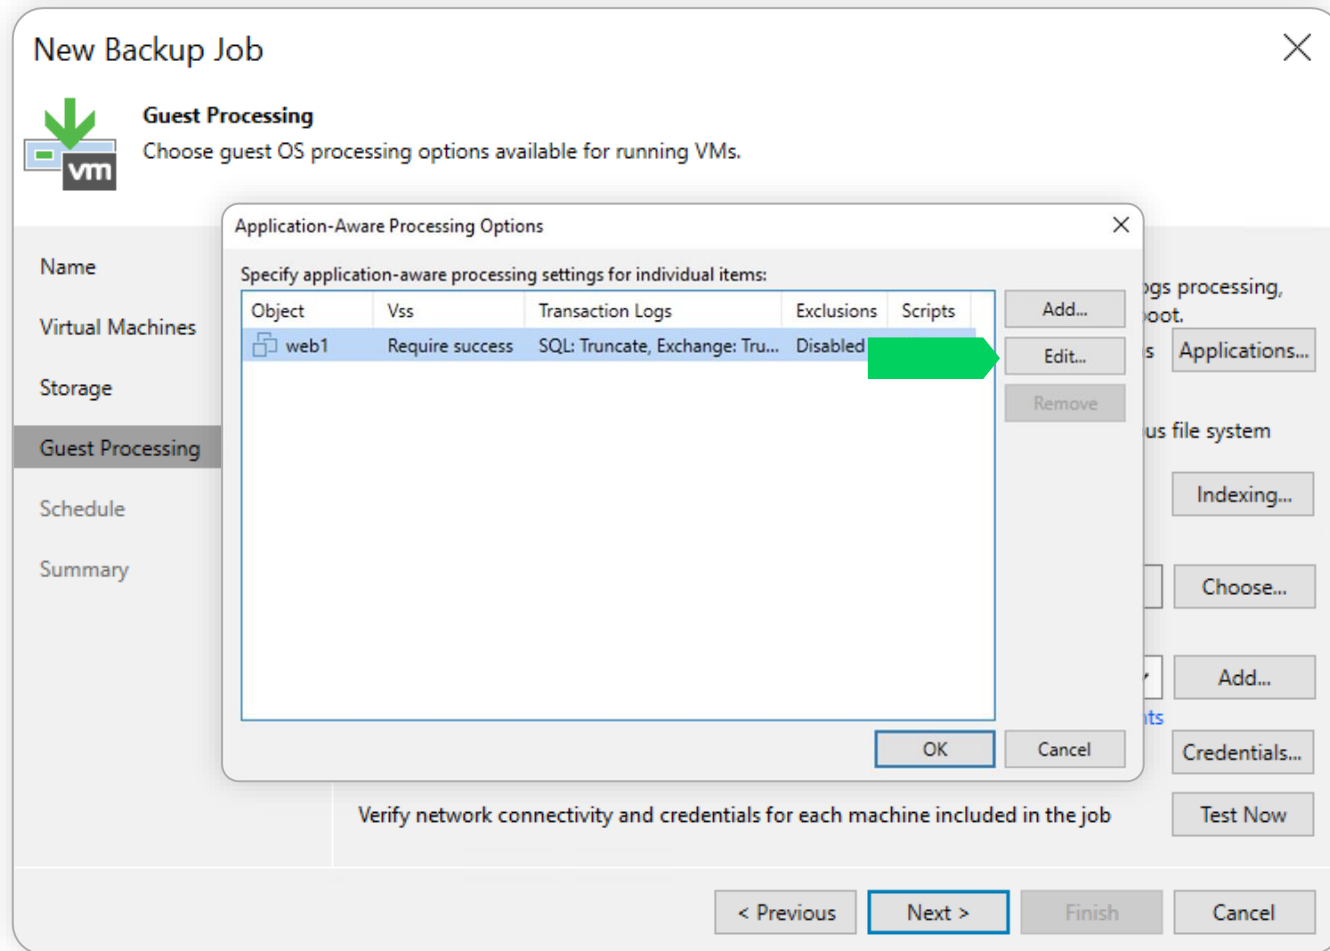
Jobs

Example of a backup job creation process



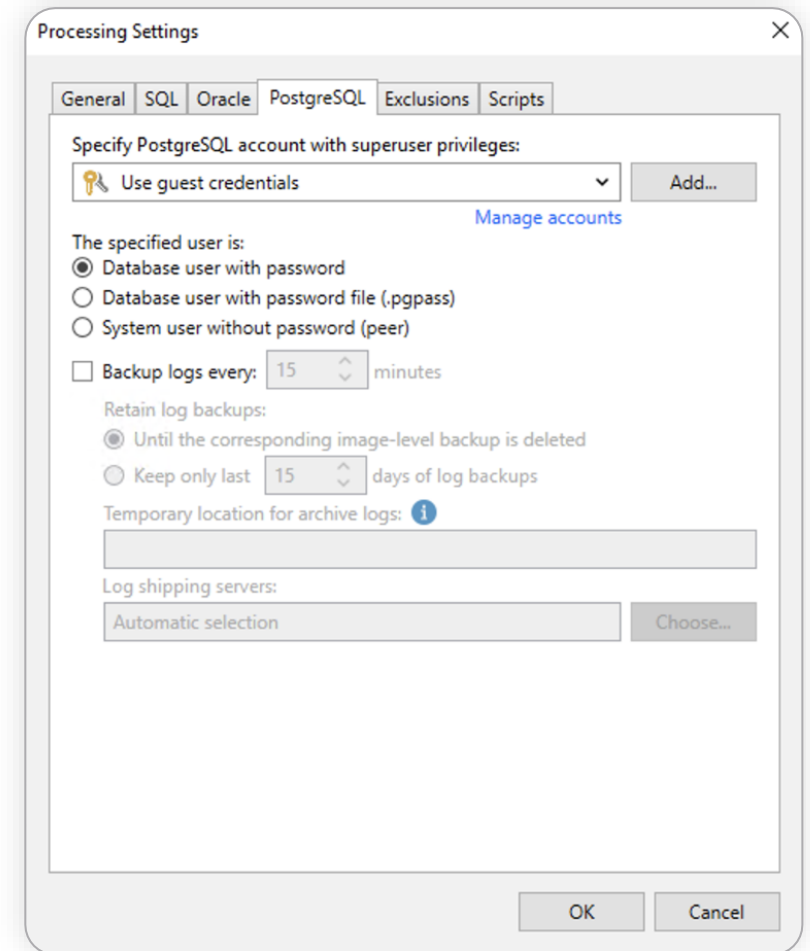
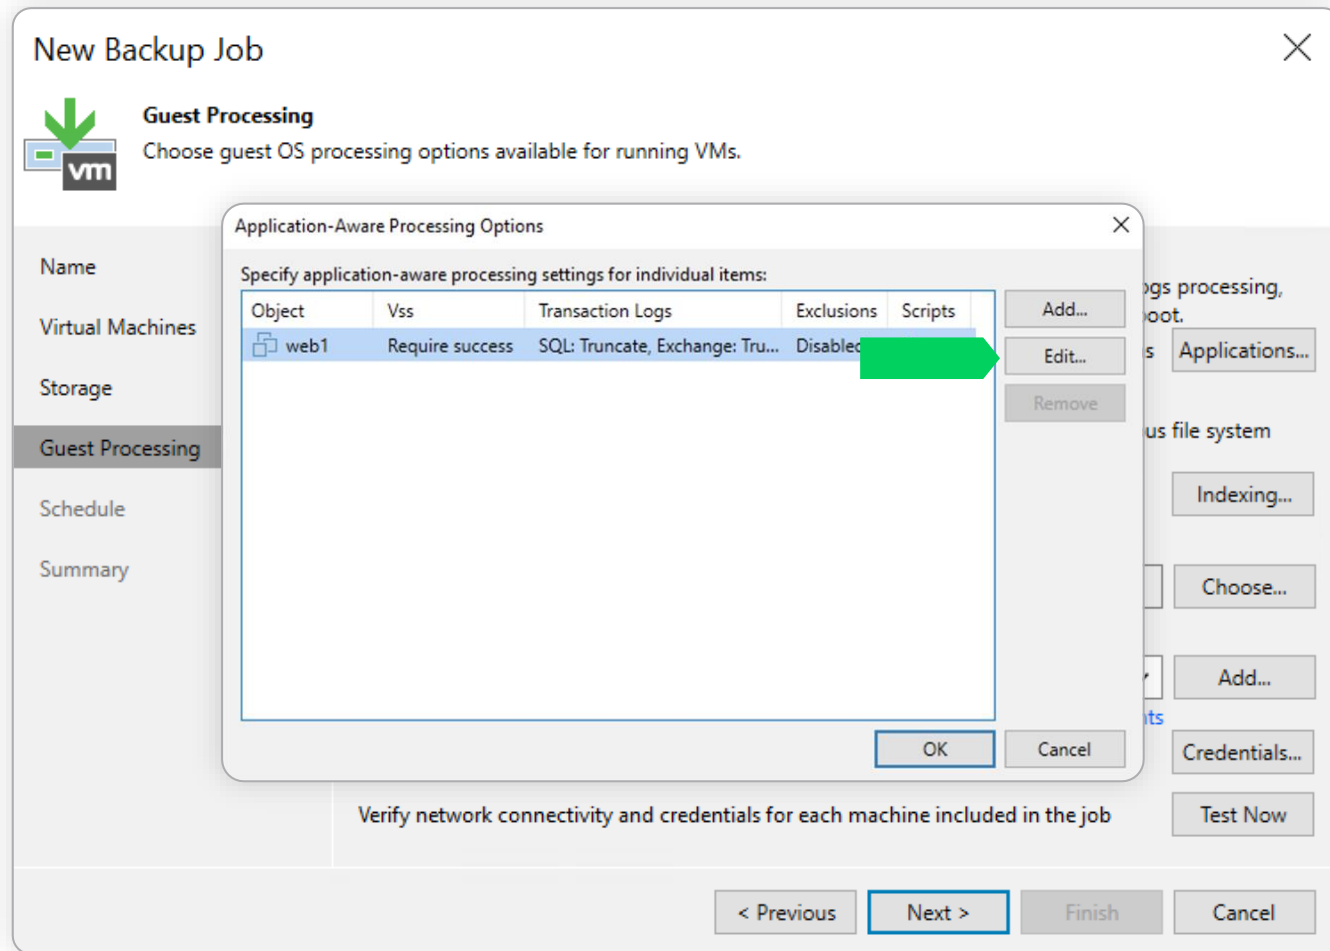
Jobs

Example of a backup job creation process



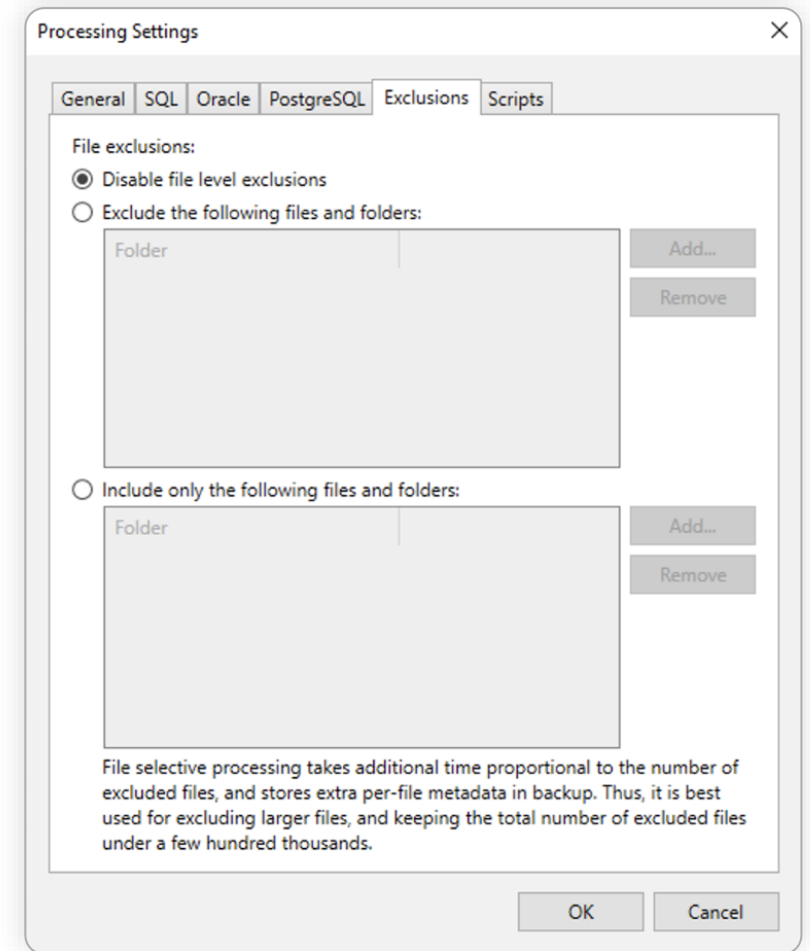

Jobs

Example of a backup job creation process



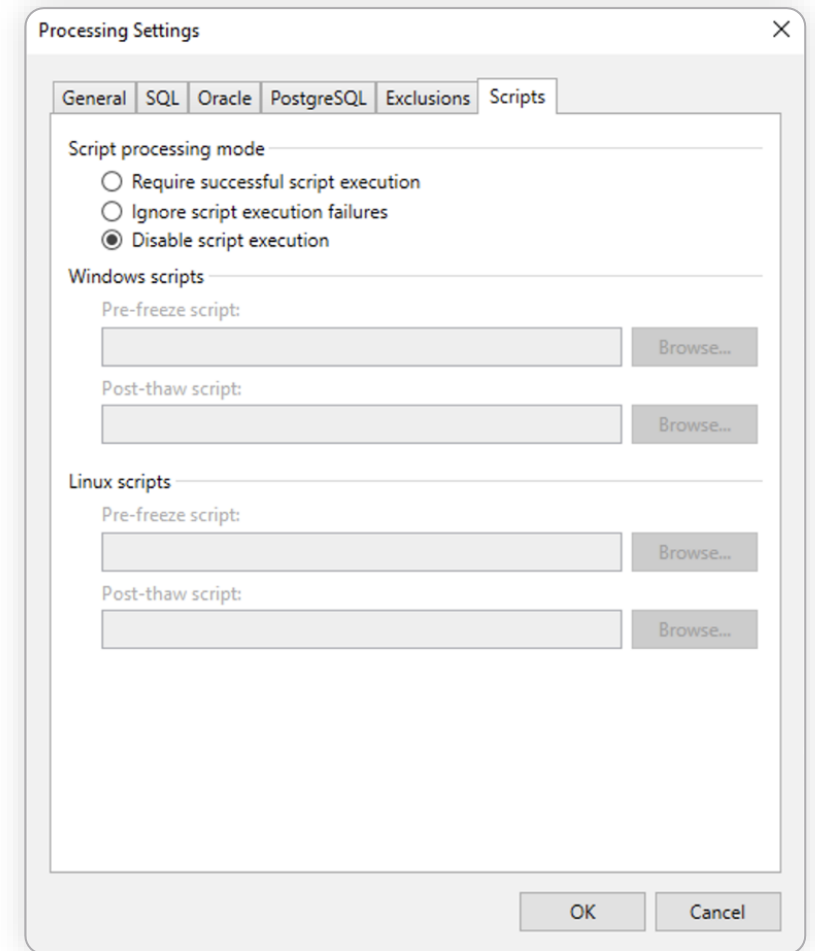
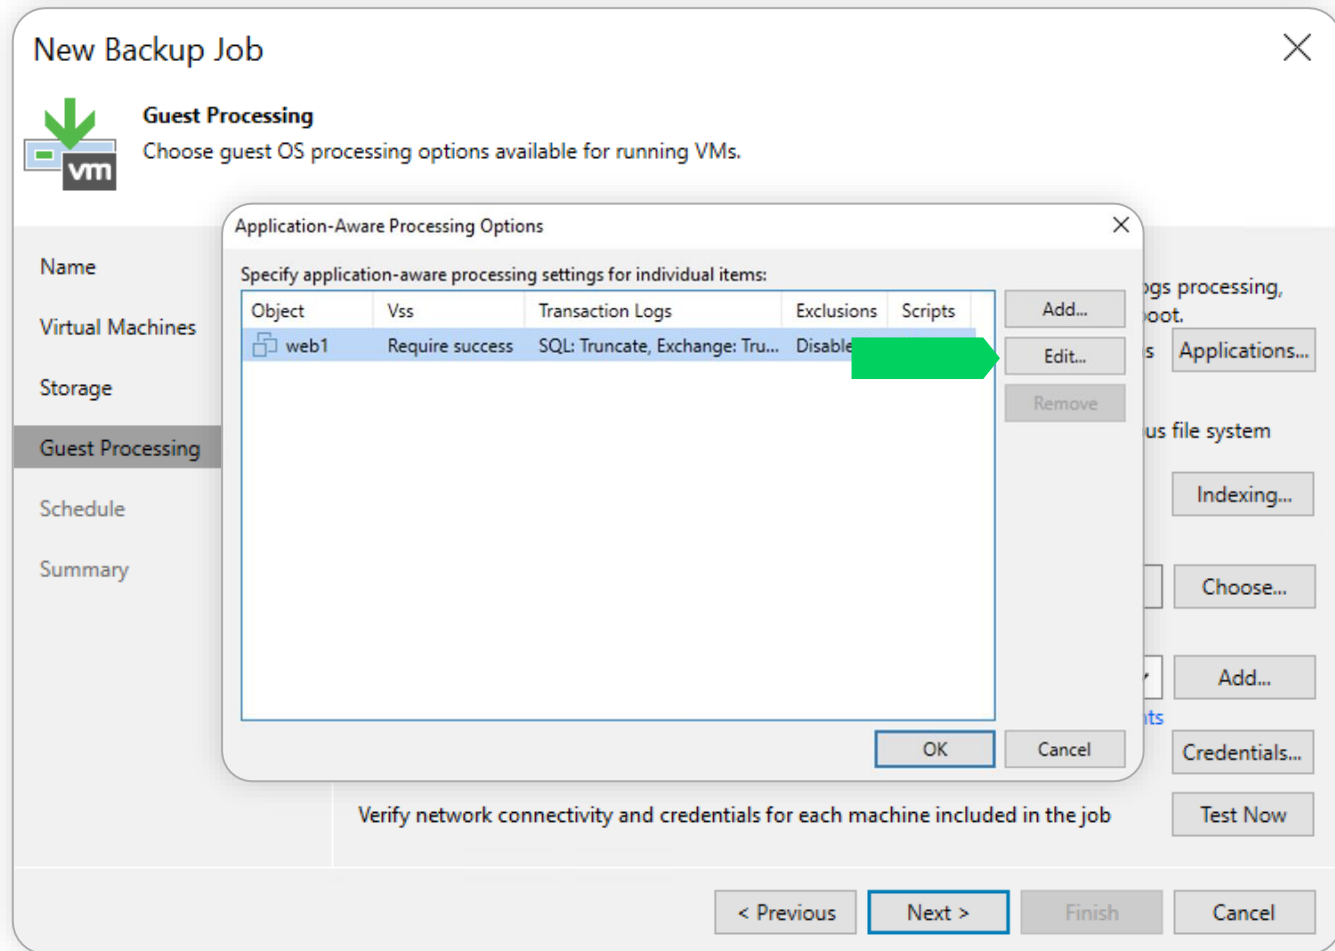
Jobs

Example of a backup job creation process



Jobs

Example of a backup job creation process



Jobs

Example of a backup job creation process

New Backup Job

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

Name

Virtual Machines

Storage

Guest Processing

Schedule

Summary

Enable application-aware processing
Detects and prepares applications for consistent backup, performs transaction logs processing, and configures the OS to perform required application restore steps upon first boot.
Customize application handling options for individual machines and applications [Applications...](#)

Enable guest file system indexing and malware detection
Indexing enables global file search functionality, automatic detection of suspicious file system activity and known malware files.
Customize advanced guest file system indexing options for individual machines [Indexing...](#)

Guest interaction proxy:
 [Choose...](#)

Guest OS credentials:
 [Add...](#)
[Manage accounts](#)

Customize guest OS credentials for individual machines and operating systems [Credentials...](#)

Verify network connectivity and credentials for each machine included in the job [Test Now](#)

[< Previous](#) [Next >](#) [Finish](#) [Cancel](#)

Jobs

Example of a backup job creation process

New Backup Job

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

Name: [Redacted]

Virtual Machines

Storage

Guest Processing

Schedule

Summary

Time Periods

All None

Day	0	2	4	6	8	10	12	14	16	18	20	22	24
Monday	Permitted	Permitted	Permitted	Permitted	Permitted	Permitted	Permitted	Permitted	Permitted	Permitted	Permitted	Permitted	Permitted
Tuesday	Permitted	Permitted	Permitted	Permitted	Permitted	Permitted	Permitted	Permitted	Permitted	Permitted	Permitted	Permitted	Permitted
Wednesday	Permitted	Permitted	Permitted	Permitted	Permitted	Permitted	Permitted	Permitted	Permitted	Permitted	Permitted	Permitted	Permitted
Thursday	Permitted	Permitted	Permitted	Permitted	Permitted	Permitted	Permitted	Permitted	Permitted	Permitted	Permitted	Permitted	Permitted
Friday	Permitted	Permitted	Permitted	Permitted	Permitted	Permitted	Permitted	Permitted	Permitted	Permitted	Permitted	Permitted	Permitted
Saturday	Permitted	Permitted	Permitted	Permitted	Permitted	Permitted	Permitted	Permitted	Permitted	Permitted	Permitted	Permitted	Permitted
Sunday	Permitted	Permitted	Permitted	Permitted	Permitted	Permitted	Permitted	Permitted	Permitted	Permitted	Permitted	Permitted	Permitted

Permitted

Denied

OK Cancel

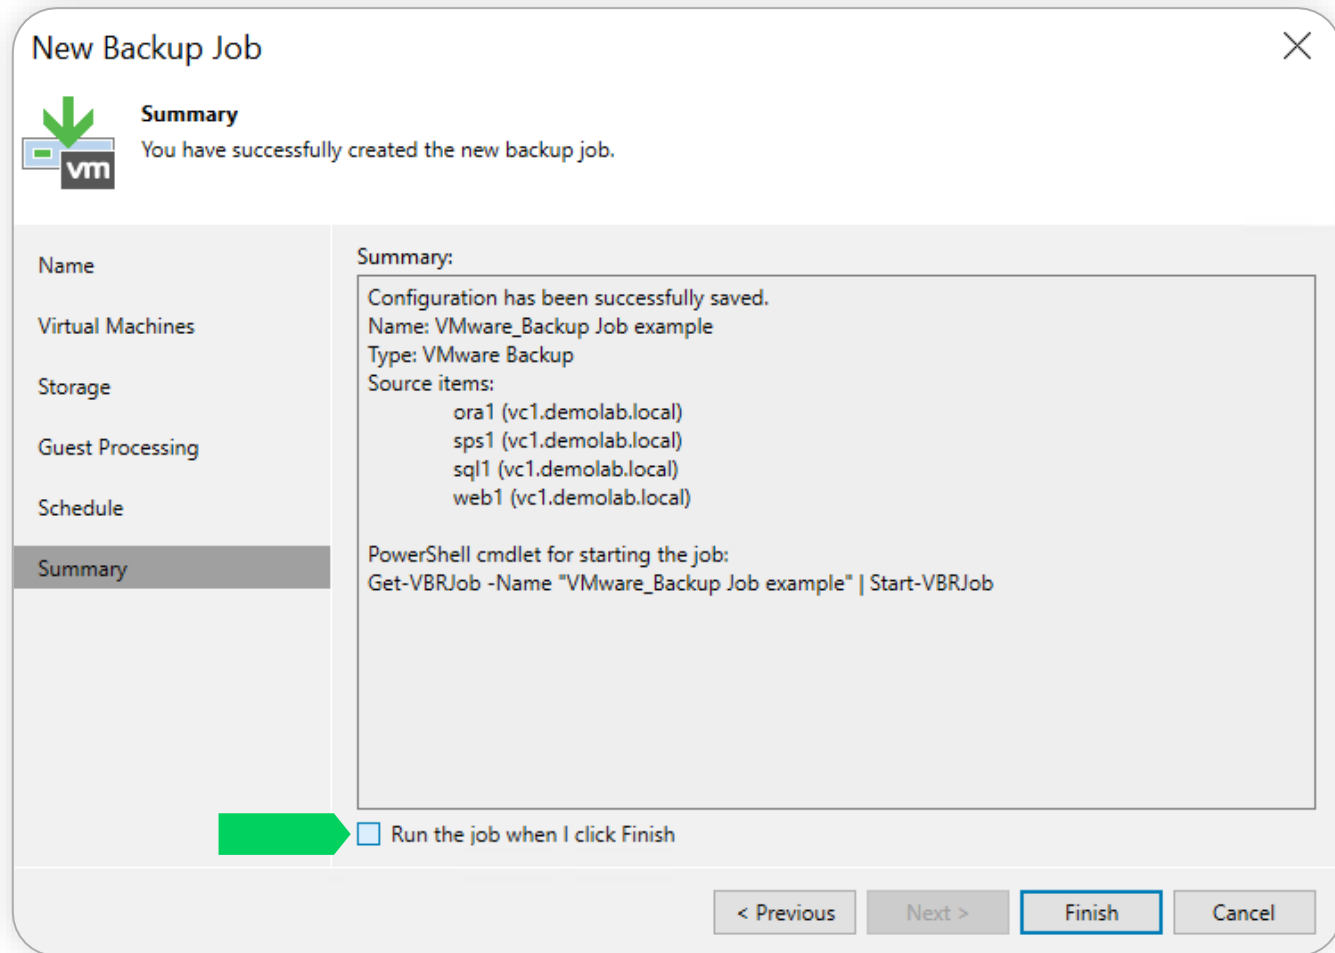
Window...

< Previous Apply Finish Cancel

Long running or accidentally started jobs will be terminated to prevent impact on your production infrastructure during busy hours.

Jobs

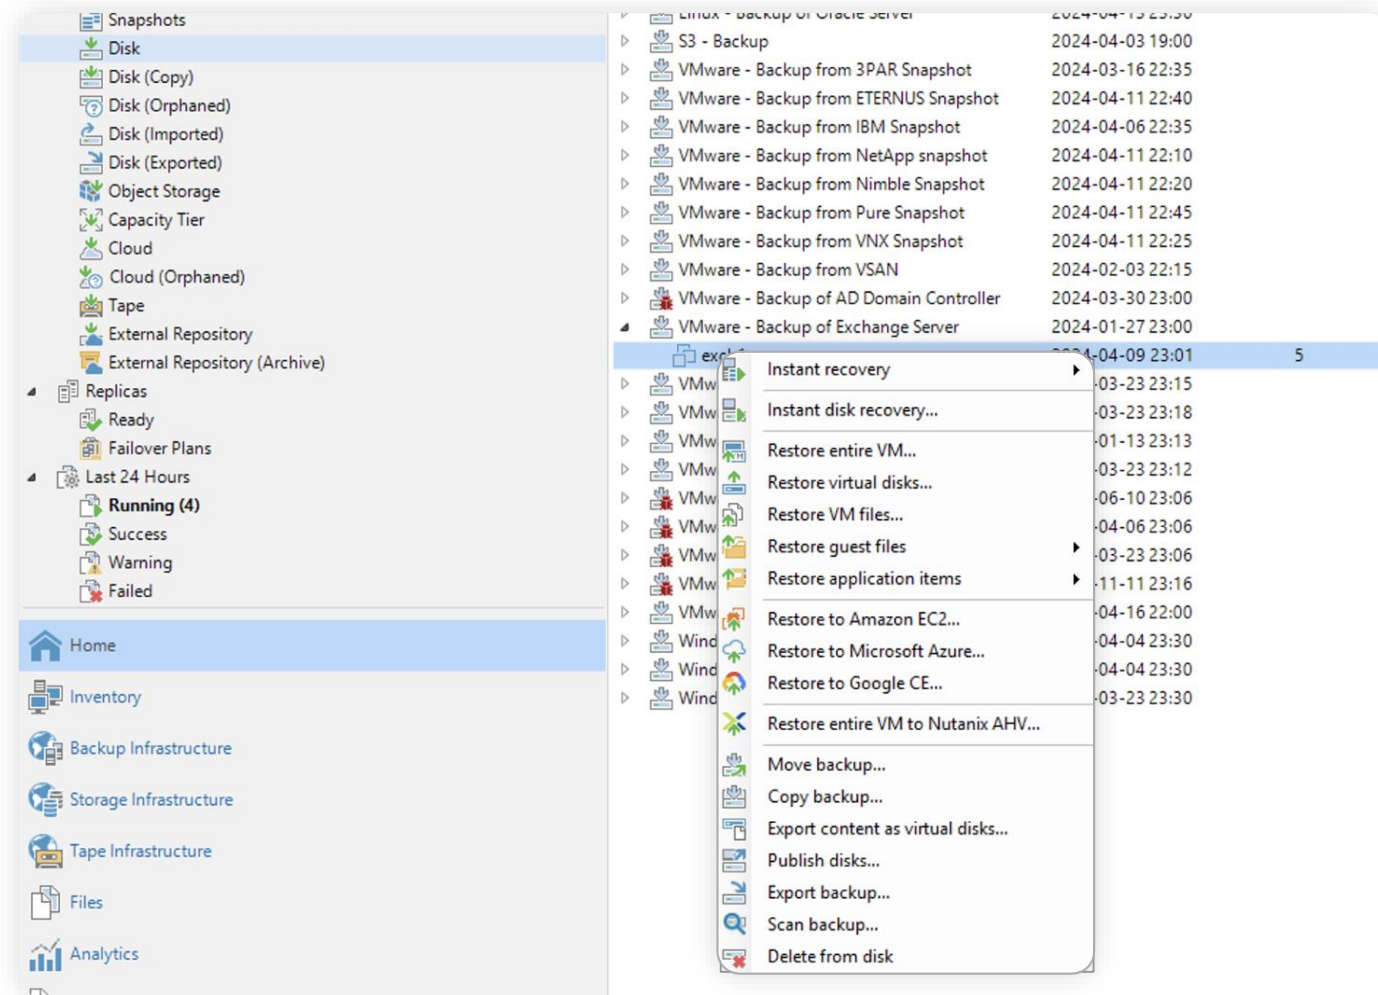
Example of a backup job creation process



Restore

Restore

Numerous restore options to reduce the time to restore.



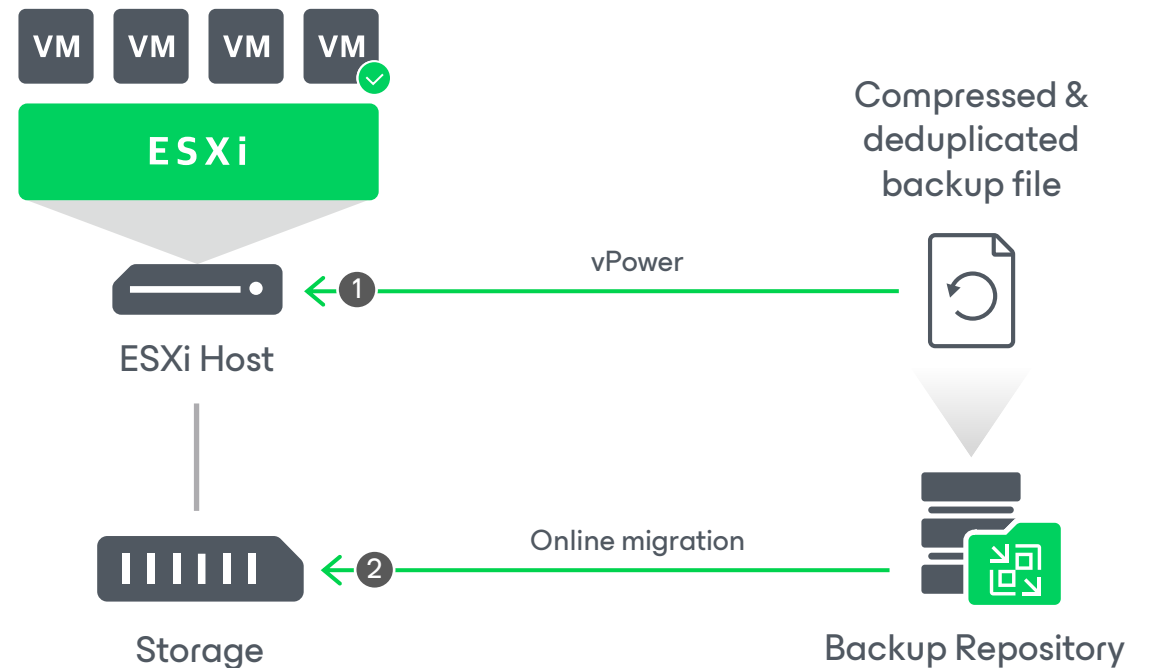
- Entire VM restore
- Instant recovery
- Virtual disks restore
- Instant disk recovery
- VM files restore
- Guest files restore
- Application items restore
- Restore to Amazon EC2; Azure; Google CE
- Restore to Nutanix AHV
- Export content as virtual disks
- Publish disks

Instant Recovery

With Instant Recovery, you can almost **immediately** restore a VM into your production environment by running it directly from the compressed and deduplicated backup file.

Instant Recovery helps improve recovery time objectives, minimize disruption and downtime of production VMs.

Instant Recovery can be used to restore data in a **cross-hypervisor** manner, as well as to recover **from a physical server to a VM**.

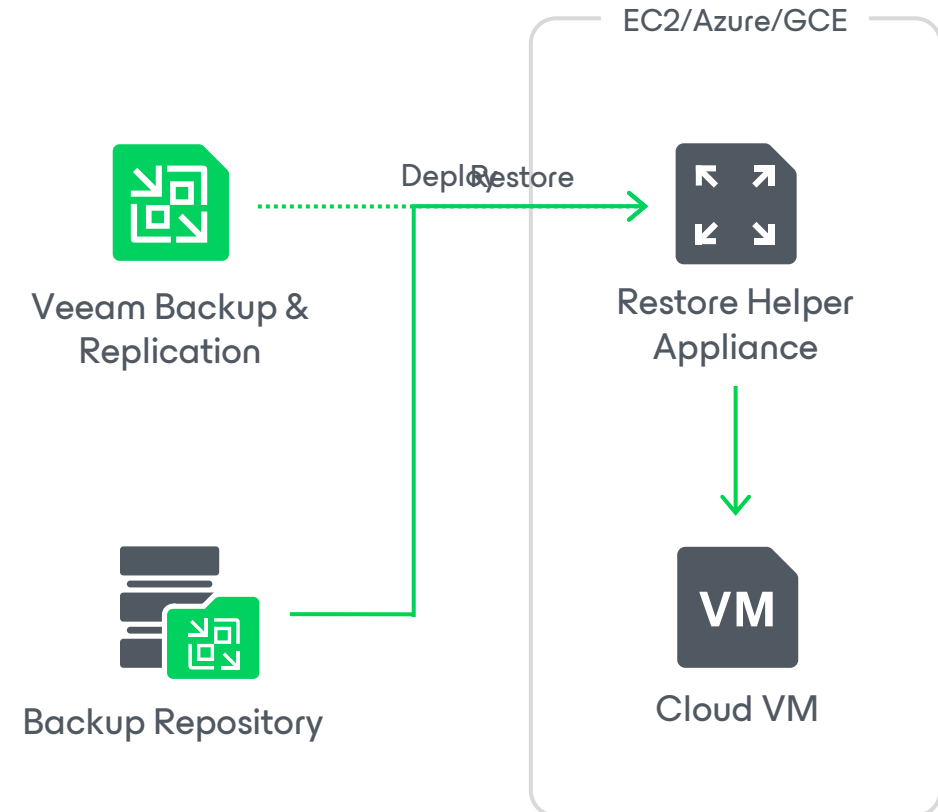


Restore to Amazon EC2/Azure/Google CE

Allows you to restore **VMs and physical workloads** to the cloud.

In case of **Amazon EC2** and **Google CE**, VBR automatically deploys the helper appliance only for the duration of the restore process and removes it immediately after that. The helper appliance **is optional**, however it may significantly **improve restore performance**.

Azure helper appliances **are mandatory and persistent**. After the restore process finishes, helper appliances get powered off and remain in Azure. The appliances remain in the powered off state until you start a new restore process.



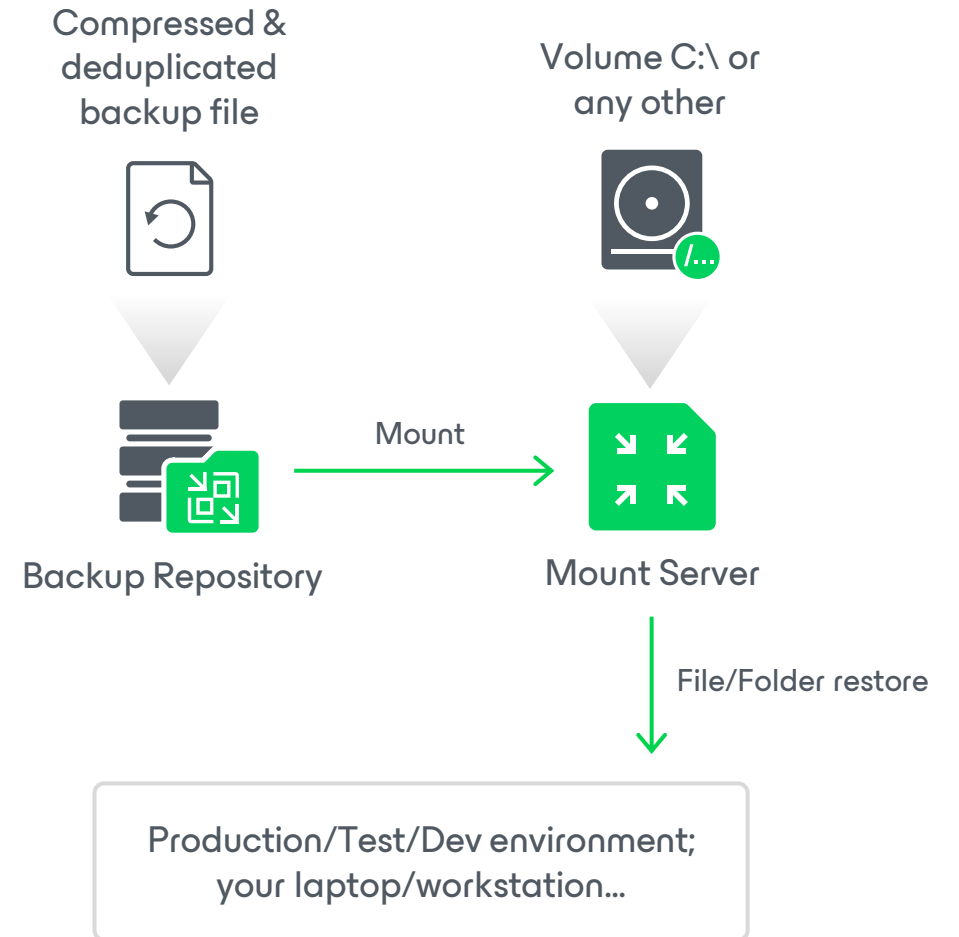
Guest files restore (FLR)

It's possible to restore **specific files** directly from the guest OS, without restoring the entire VM or physical server.

Supported recovery sources: Backups, Replicas, Storage Snapshots, Nutanix AHV Snapshots

Supported file systems:

- **Windows:** FAT, FAT32, NTFS, ReFS
- **Linux:** ext2, ext3, ext4, ReiserFS, JFS, XFS, Btrfs, NTFS
- **BSD, Mac, Solaris** are also supported



Application Items restore

Explorers are available for the following applications:

Microsoft Active Directory

Microsoft SQL Server

Oracle PostgreSQL


Microsoft Exchange

Microsoft SharePoint

Microsoft OneDrive for Business

Microsoft Teams

To use them, make sure your backups and replicas are created with application-aware processing turned on.



How to restore a VM?

The screenshot shows the Veeam Backup & Replication console interface. The top navigation bar includes 'Home' and 'Backup'. The left sidebar contains a tree view with categories like 'Jobs', 'Backups', 'Replicas', and 'Running (7)'. A green arrow labeled '1' points to the 'Entire VM' option in the 'Restore' menu. A second green arrow labeled '2' points to the 'Restore entire VM...' option in the context menu that appears over a selected backup job.

Job Name	Creation Time	Restore Points
Agent Backup Policy AIX	2024-04-17 22:00	
Oracle RMAN		
Agents		
AHV - Backup	2023-08-22 22:00	
Files - Backup	2024-03-27 19:00	
Linux - Backup of MySQL Server	2024-04-13 23:30	
Linux - Backup of Oracle Server	2024-04-20 23:30	
Linux - Backup of PostgreSQL Server	2024-04-13 23:30	
S3 - Backup	2024-04-09 19:00	
VMware - Backup from 3PAR Snapshot	2024-03-23 22:35	
VMware - Backup from ETERNUS Snapshot	2024-04-17 22:40	
VMware - Backup from IBM Snapshot	2024-04-13 22:35	
VMware - Backup from NetApp snapshot	2024-04-17 22:10	
VMware - Backup from Nimble Snapshot	2024-04-17 22:20	
VMware - Backup from Pure Snapshot	2024-04-17 22:45	
VMware - Backup from VNX Snapshot	2024-04-18 22:25	
VMware - Backup from VSAN	2024-02-03 22:15	
VMware - Backup of AD Domain Controller	2024-04-06 23:00	
VMware - Backup of Exchange Server	2024-01-27 23:00	
VMware - Backup of File Server	2024-03-30 23:15	
VMware - Backup of MS SQL Server	2024-03-30 23:18	
sql1 - SQL Server transaction log backup	2024-04-24 12:59	10120
sql1	2024-04-23 23:19	19

How to restore a VM?

Entire VM Restore

Virtual Machines
Select virtual machines to be restored. You can add individual virtual machines from backup files, or containers from live environment (containers will be automatically expanded into plain VM list).

Restore Points

Available restore points for sql1:

Job	Type	Location
VMware - Backup of MS SQL Server - sql1 (SOBR)		
less than a day ago (23:19 Tuesday 2024-04-23)	Increment	Performance Tier
1 day ago (23:19 Monday 2024-04-22)	Increment	Performance Tier
2 days ago (23:19 Sunday 2024-04-21)	Increment	Performance Tier
3 days ago (23:19 Saturday 2024-04-20)	Full (W)	Performance Tier
4 days ago (23:19 Friday 2024-04-19)	Increment	Performance Tier
5 days ago (23:19 Thursday 2024-04-18)	Increment	Performance Tier
6 days ago (23:19 Wednesday 2024-04-17)	Increment	Performance Tier
7 days ago (23:23 Tuesday 2024-04-16)	Increment	Performance Tier
8 days ago (23:19 Monday 2024-04-15)	Increment	Performance Tier
9 days ago (23:19 Sunday 2024-04-14)	Increment	Performance Tier
10 days ago (23:19 Saturday 2024-04-13)	Full (W)	Performance Tier
11 days ago (23:28 Friday 2024-04-12)	Increment	Performance Tier
12 days ago (23:19 Thursday 2024-04-11)	Increment	Performance Tier
13 days ago (23:19 Wednesday 2024-04-10)	Increment	Performance Tier
14 days ago (23:21 Tuesday 2024-04-09)	Increment	Performance Tier
15 days ago (23:20 Monday 2024-04-08)	Increment	Performance Tier
16 days ago (23:21 Sunday 2024-04-07)	Increment	Performance Tier

Buttons: Add..., Point..., Remove, OK, Cancel

Navigation: < Previous, Next >, Finish, Cancel

How to restore a VM?

Entire VM Restore

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

Virtual Machines

Restore Mode

Host

Resource Pool

Datastore

Folder

Network

Secure Restore

Reason

Summary

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

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

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

[Pick proxy to use](#)

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

< Previous Next > Finish Cancel

How to restore a VM?

Entire VM Restore

Datastore
By default, original datastore and disk type are selected for each VM file. You can change them by selecting desired VM file, and clicking Datastore or Disk Type. Use multi-select (Ctrl-click and Shift-click) to select multiple VMs at once.

Virtual Machines

Restore Mode

Host

Resource Pool

Datastore

Folder

Network

Secure Restore

Reason

Summary

Files location:

File	Size	Datastore	Disk type
sql1			
Configurati		[688.1 GB f...	
Hard disk 1		[688.1 GB f...	Same as source
Hard disk 2		[688.1 GB f...	Same as source
Hard disk 3		[688.1 GB f...	Same as source

Disk Type Settings

Restored VM disk type:

- Same as source
- Thin
- Thick (lazy zeroed)
- Thick (eager zeroed)

OK Cancel

Select multiple VMs to apply changes in bulk.

Datastore... Disk Type...

< Previous Next > Finish Cancel

How to restore a VM?

Entire VM Restore

Folder
By default, original VM folder is selected as restore destination for each VM. You can change folder by selecting desired VM and clicking Folder. Use multi-select (Ctrl-click and Shift-click) to select multiple VMs at once.

Virtual Machines

Restore Mode

Host

Resource Pool

Datastore

Folder

Network

Secure Restore

Reason

Summary

VM Folder:

Name	New Name	Folder
sql1	sql1	vm

Select multiple VMs to apply settings change in bulk.


Restore VM tags
Select this option to restore VM tags that were assigned to the VM when backup was taken.

< Previous Next > Finish Cancel


Name... Folder...

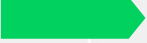
How to restore a VM?

Entire VM Restore ✕

 **Secure Restore**
Scan the selected backup for malware, such as computer viruses or ransomware, prior to performing the restore. This requires a compatible antivirus installed on the mount server specified for the corresponding backup repository.

Virtual Machines


Restore Mode  Scan the restore point with an antivirus engine

Host  Scan the restore point with the following YARA rule:

Resource Pool

Host

[Copy YARA rules location to clipboard](#)

Datastore  Scan options:

Folder

Network

Secure Restore


Reason

Summary

< Previous Next > Finish Cancel

How to restore a VM?

Entire VM Restore ✕

 **Reason**
Type in the reason for performing this restore operation. This information will be logged in the restore sessions history for later reference.

Virtual Machines

Restore Mode

Host

Resource Pool

Datastore

Folder

Network

Secure Restore

Reason

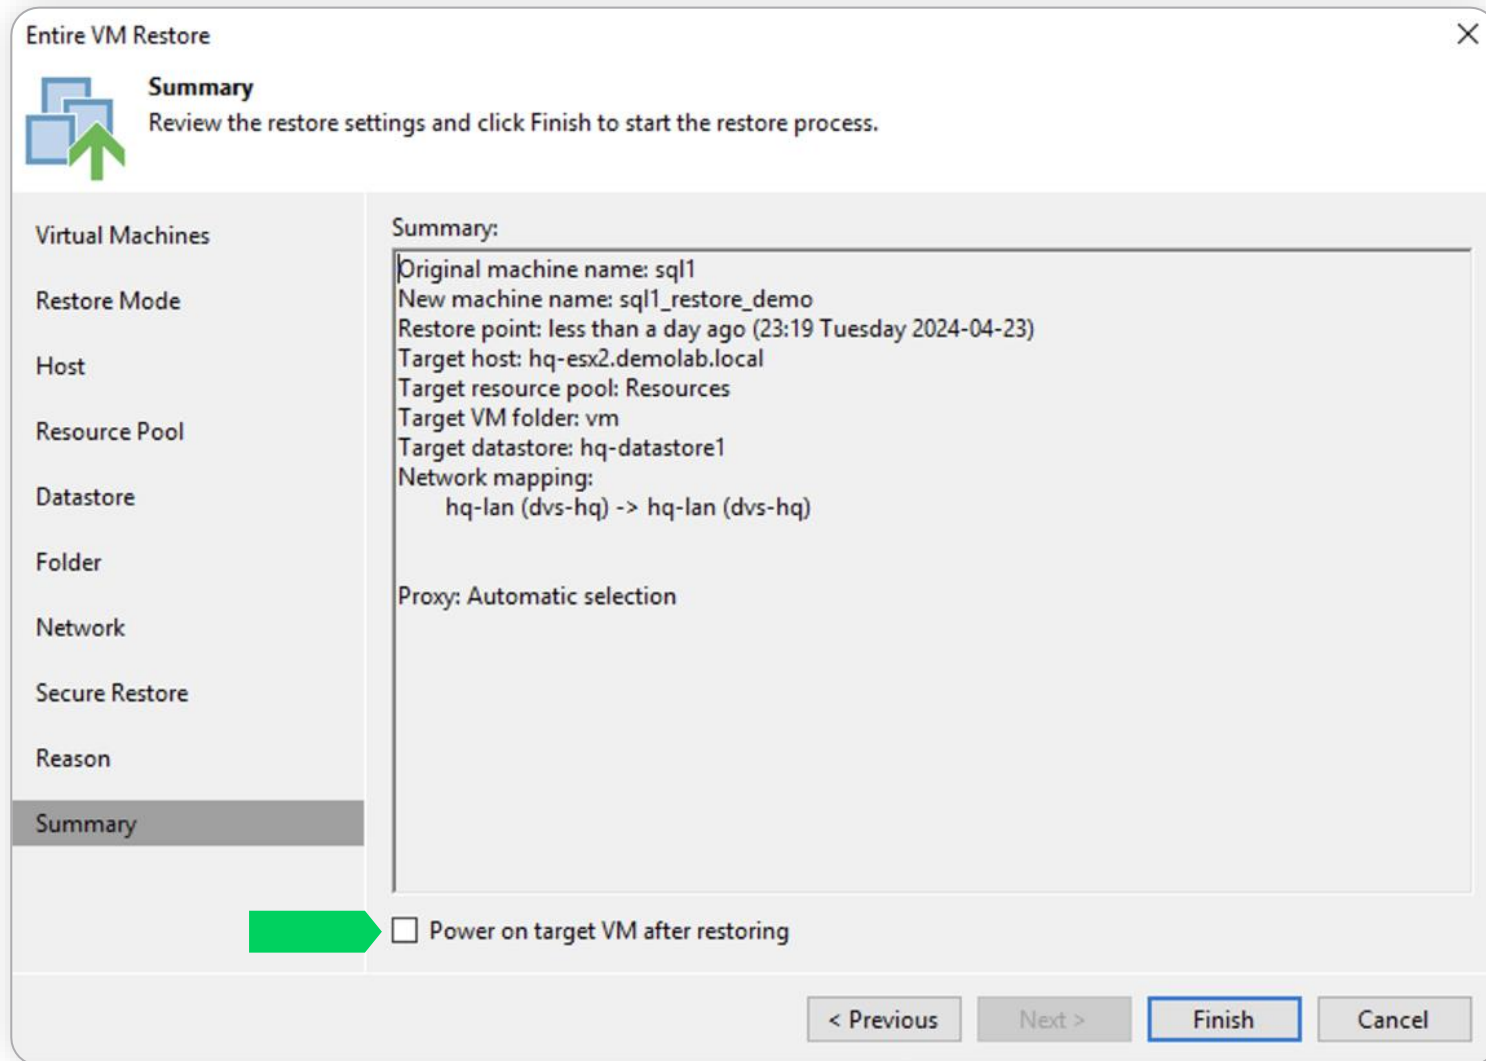
Summary

Restore reason:

Do not show me this page again

< Previous Next > Finish Cancel

How to restore a VM?



The screenshot shows the 'Entire VM Restore' dialog box with the 'Summary' tab selected. The dialog box has a title bar with 'Entire VM Restore' and a close button. Below the title bar is a 'Summary' section with a blue icon of three overlapping squares and a green arrow pointing up, and the text 'Review the restore settings and click Finish to start the restore process.' The main area is divided into a left sidebar and a right content area. The sidebar lists various settings: Virtual Machines, Restore Mode, Host, Resource Pool, Datastore, Folder, Network, Secure Restore, Reason, and Summary (which is highlighted). The right content area displays the following summary information: Original machine name: sql1, New machine name: sql1_restore_demo, Restore point: less than a day ago (23:19 Tuesday 2024-04-23), Target host: hq-esx2.demolab.local, Target resource pool: Resources, Target VM folder: vm, Target datastore: hq-datastore1, Network mapping: hq-lan (dvs-hq) -> hq-lan (dvs-hq), and Proxy: Automatic selection. At the bottom of the dialog box, there is a checkbox labeled 'Power on target VM after restoring' which is currently unchecked. A green arrow points to this checkbox. To the right of the checkbox are four buttons: '< Previous', 'Next >', 'Finish', and 'Cancel'. The 'Finish' button is highlighted with a blue border.

Entire VM Restore

Summary
Review the restore settings and click Finish to start the restore process.

Virtual Machines

Restore Mode

Host

Resource Pool

Datastore

Folder

Network

Secure Restore

Reason

Summary

Summary:

Original machine name: sql1
New machine name: sql1_restore_demo
Restore point: less than a day ago (23:19 Tuesday 2024-04-23)
Target host: hq-esx2.demolab.local
Target resource pool: Resources
Target VM folder: vm
Target datastore: hq-datastore1
Network mapping:
 hq-lan (dvs-hq) -> hq-lan (dvs-hq)

Proxy: Automatic selection

Power on target VM after restoring

< Previous Next > **Finish** Cancel

Security

Best Practices: Security & Resiliency

Security & Compliance Analyzer

Veeam Backup & Replication provides a built-in tool to ensure that your backup server configuration follows **security best practices** for Veeam backup infrastructure components based on Microsoft Windows Server and Linux operating systems.

It includes **11 Security Checks** & **19 Product Configuration Checks**

The Security & Compliance Analyzer window displays the following best practices and their status:

Best Practice	Status
Backup infrastructure security	
Remote Desktop Service (TermService) should be disabled	Not implemented
Remote Registry service (RemoteRegistry) should be disabled	Not implemented
Windows Remote Management (WinRM) service should be disabled	Not implemented
Windows Firewall should be enabled	Passed
WDigest credentials caching should be disabled	Passed
Web Proxy Auto-Discovery service (WinHttpAutoProxySvc) should be disabled	Not implemented
Deprecated versions of SSL and TLS should be disabled	Not implemented
Windows Script Host should be disabled	Not implemented
SMBv1 protocol should be disabled	Passed
Link-Local Multicast Name Resolution (LLMNR) should be disabled	Not implemented
SMBv3 signing and encryption should be enabled	Not implemented
Product configuration	
MFA for the backup console should be enabled	Not implemented
Immutable or offline (air gapped) media should be used	Not implemented
Password loss protection should be enabled	Passed
Backup server should not be a part of the production domain	Unable to detect
Email notifications should be enabled	Not implemented
All backups should have at least one copy (the 3-2-1 backup rule)	Not implemented
Reverse incremental backup mode is deprecated and should be avoided	Passed
Unknown Linux servers should not be trusted automatically	Not implemented
The configuration backup must not be stored on the backup server	Not implemented
Host to proxy traffic encryption should be enabled for the Network transport mode	Passed
Hardened repositories should not be hosted in virtual machines	Passed
Network traffic encryption should be enabled in the backup network	Passed
Linux servers should have password-based authentication disabled	Passed
Backup services should be running under the LocalSystem account	Passed
Configuration backup should be enabled and use encryption	Not implemented
Credentials and encryption passwords should be rotated at least annually	Passed
Hardened repositories should have the SSH Server disabled	Passed
S3 Object Lock in the Governance mode doesn't provide true immutability	Passed
Backup jobs to cloud repositories should use encryption	Passed

Summary table from the interface:


Status	Last Run
Stopped	16 hours ago
Stopped	16 hours ago
Enabled	
Running	
Stopped	21 minutes ago
Stopped	1 day ago
Stopped	1 day ago
Stopped	15 hours ago
Stopped	16 hours ago
Stopped	16 hours ago
Stopped	
Stopped	14 hours ago
Stopped	9 hours ago
Stopped	19 hours ago
Stopped	18 hours ago
Stopped	2 hours ago
Stopped	2 days ago
Stopped	14 hours ago
Stopped	14 hours ago
Stopped	14 hours ago
Stopped	14 hours ago
Idle	16 hours ago

Best Practices: Security & Resiliency

Multi-Factor Authentication

Veeam Backup & Replication supports **multi-factor authentication (MFA)** for additional user verification.

A **one-time password (OTP)** generated in the mobile authenticator application is used as a second verification method. Combined with **login and password credentials**, it creates a more secure environment and protects user accounts from being compromised.

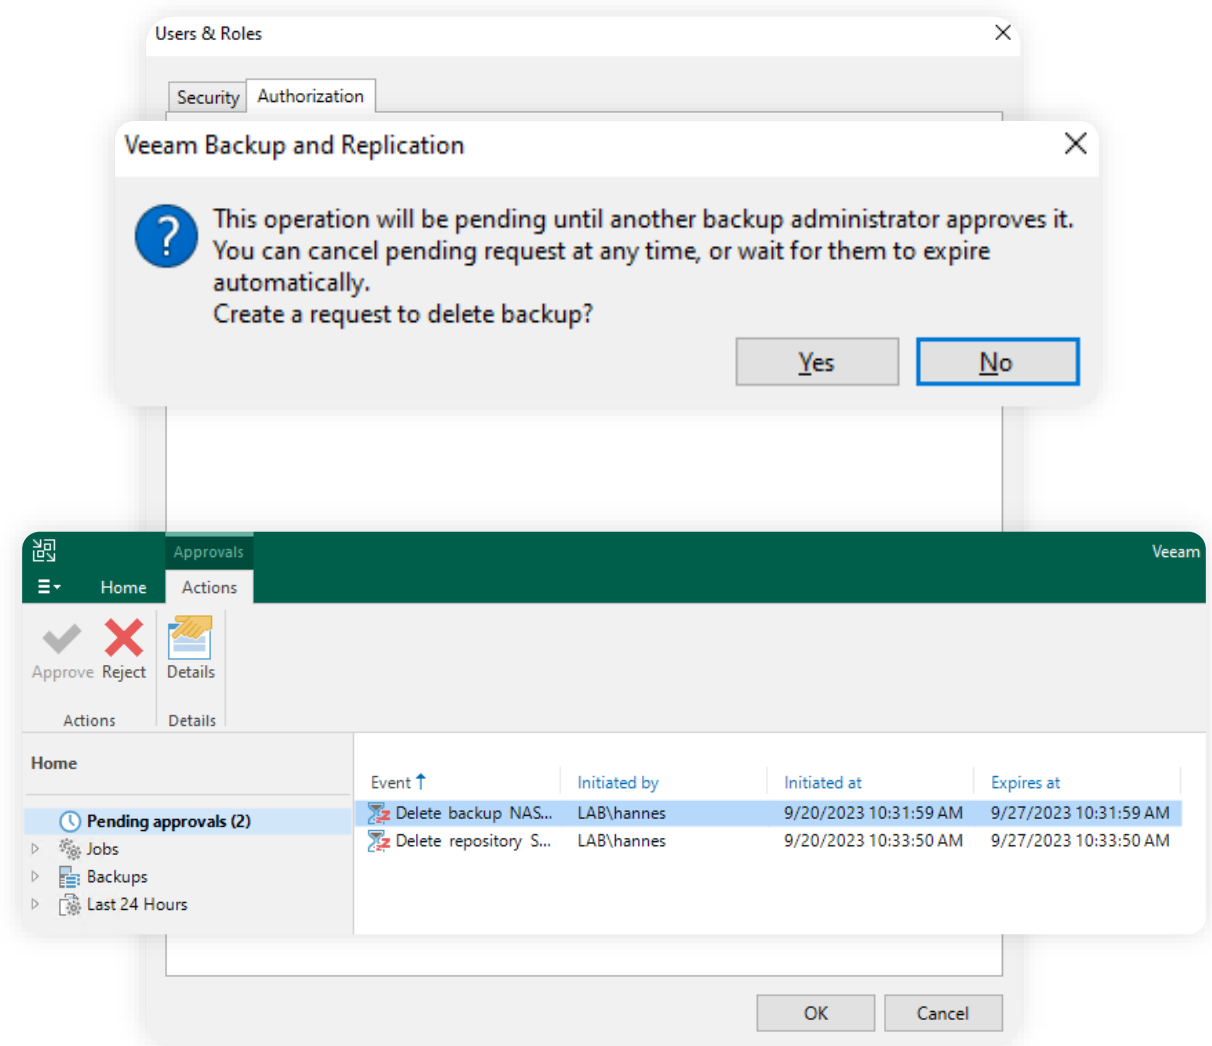


Best Practices: Security & Resiliency

Four-Eyes Authorization

This feature prevents accidental or erroneous **deletion of backups** or **entire backup repositories, changes to users, roles and other access settings** using a backup console by requiring an **approval from a second Veeam Backup Administrator** before requested changes to these sensitive backup server settings can be applied.

All events related to four-eyes authorization are displayed in the **History view**, under the **Authorization Events** node.

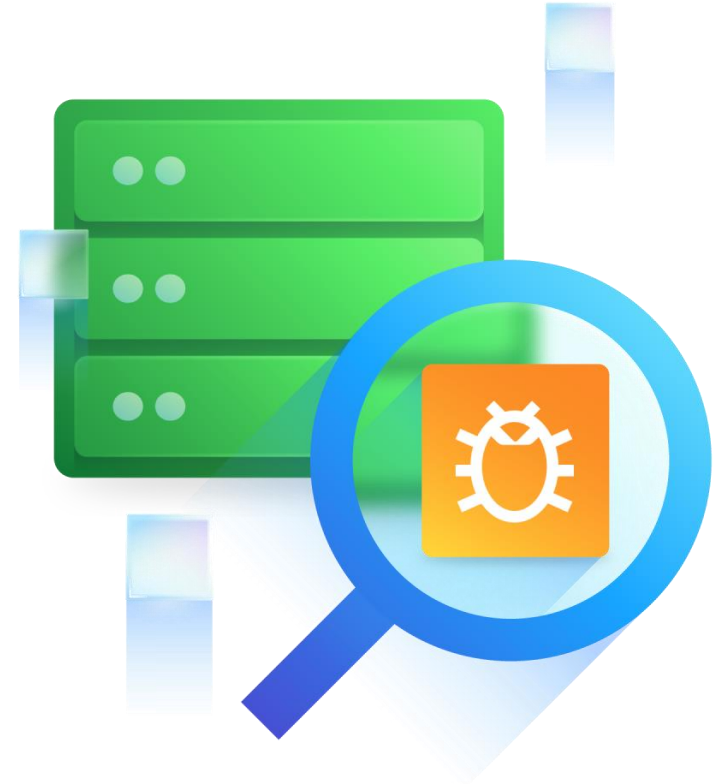


Best Practices: Security & Resiliency

Malware Detection

Malware Detection with Veeam is focused on detecting malware and suspicious activity in both production and backup infrastructure & sending out notifications about it.

Starting Veeam version 12.1, we can identify even the sleeping malware in production. Scanning the already performed and stored backups on the repository is also possible.



Best Practices: Security & Resiliency

Malware Detection. On-Demand Scan For Malware & Content.

Leverages **Antivirus** and/or **YARA** scan.

Provides **three scan modes** with multiple ranges:

- Find (first) clean backup(s)
- Find clean backup in range
- Find content (e.g., credit card numbers)

The scan happens on **mount server**.

Scan Backup

Performs an ad-hoc scan of you backups with an antivirus or the YARA engine to find the latest malware-free restore point or to detect the presence of specific entries, such as personal information.

Scan mode:

- Find the last clean restore point
Restore points will be scanned sequentially starting from the most recent one until the first malware-free one is found. Use this options when a cyber-attack is known to have started recently.
- Find the last clean restore point in range
Restore points will be scanned in an optimal order to identify the last clean backup in range with least number of scans possible. Use this option if you are not sure when the attack started, or when dealing with a known sleeping malware.
- Scan all restore points in range for content analysis
All restore points in range will be scanned sequentially. Use this option for backup content analysis with an applicable YARA rule, for example to look for personally identifiable information (PII), personal health information (PHI) or payment card industry (PCI) data.

Scan engine:

- Scan restore points with an antivirus
- Scan restore points with the following YARA rule:

YARA rules location: C:\Program Files\Veeam\Backup and Replication\Backup\YaraRules\

Scan range:

From:

- Most recent restore point
- Start date:

To:

- Oldest available restore point
- End date:

Continue scanning all remaining files after the first occurrence

[Hide scan range](#)

Best Practices: Security & Resiliency

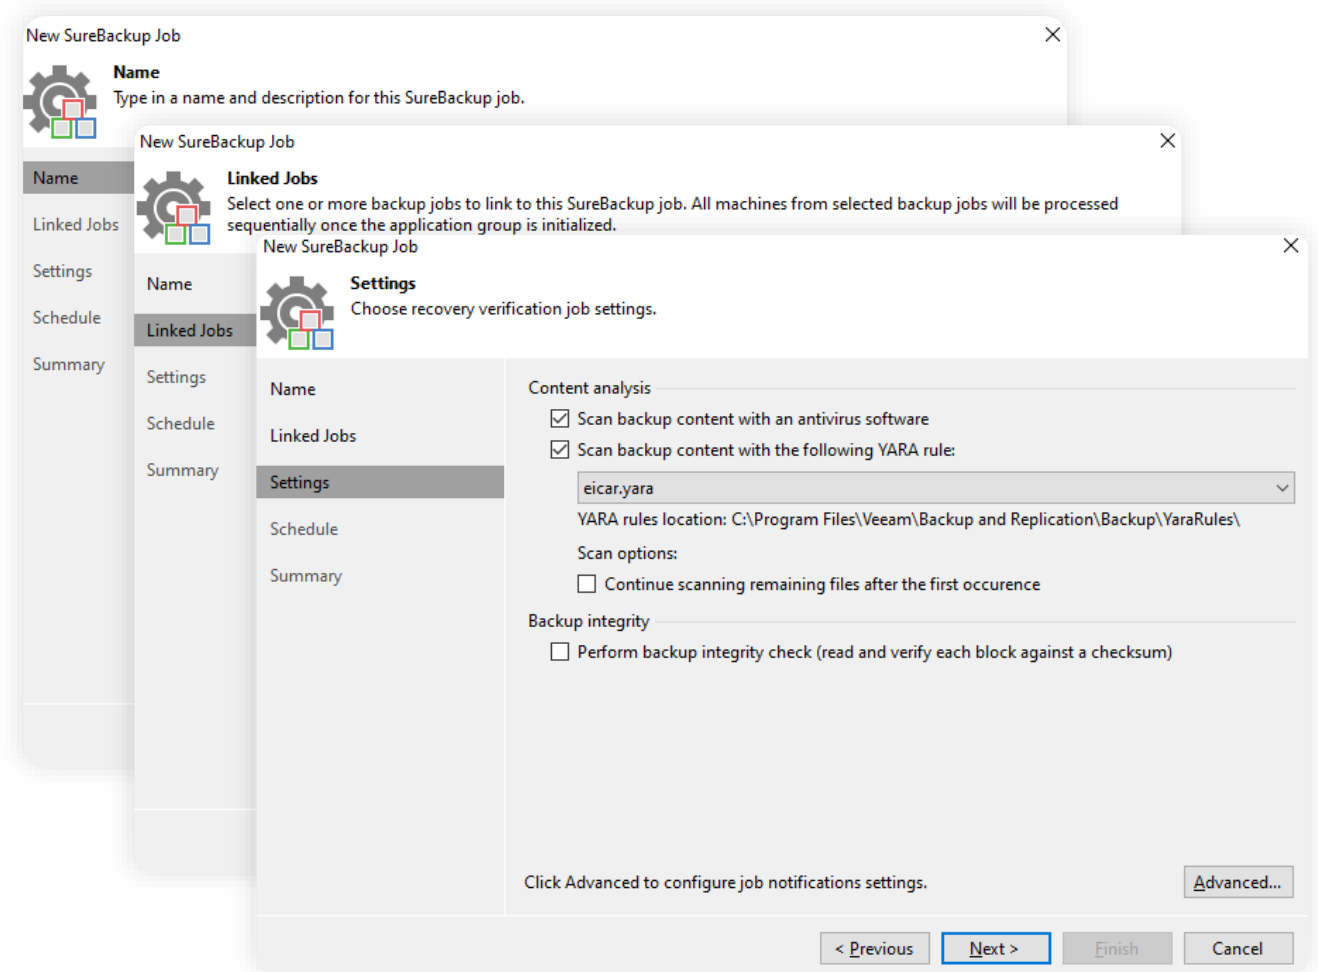
Malware Detection. Automated Malware & Content Scans.

To be activated in a **SureBackup Job**.

Leverages **Antivirus** and/or **YARA** scan.

Scans the **entire backup job** (exclusions are possible).

The scan happens on **mount server**.



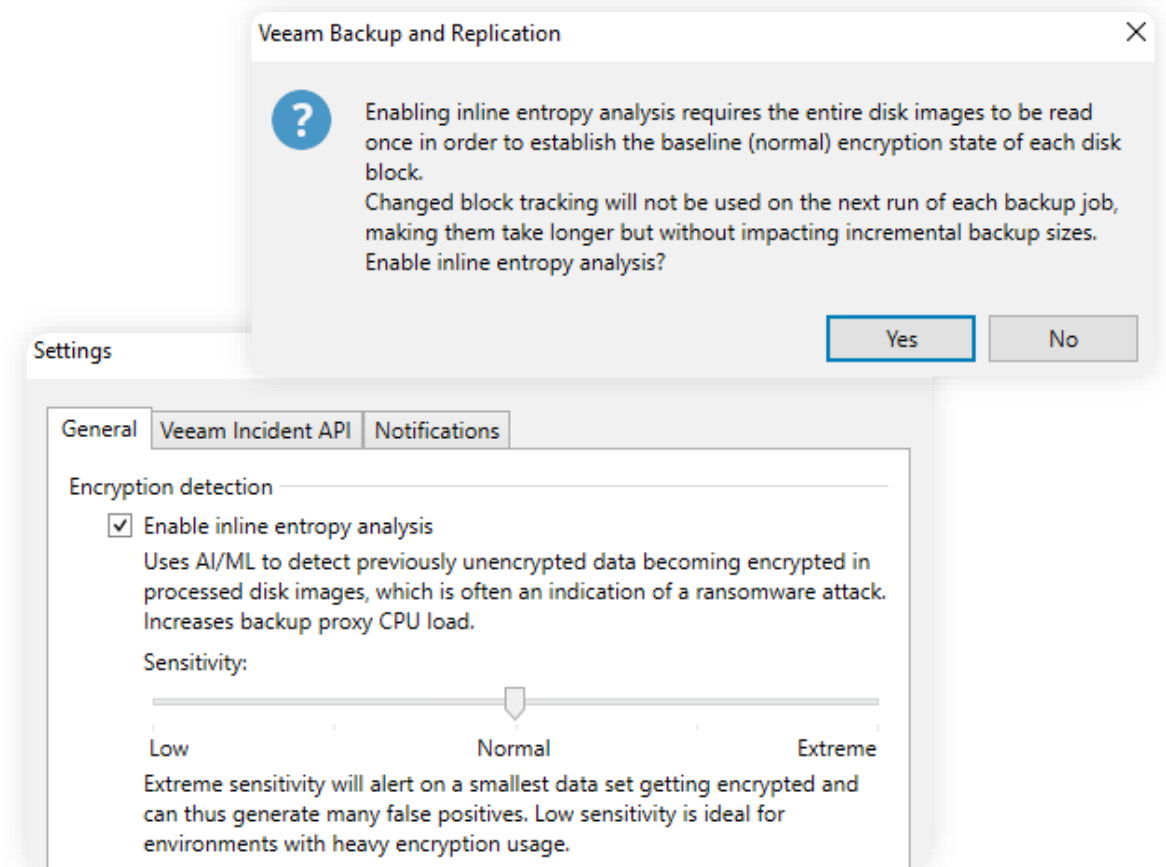
Best Practices: Security & Resiliency

Malware Detection. Inline Scan: Encryption & Text Analysis.

Analyzes block-level data during backup.

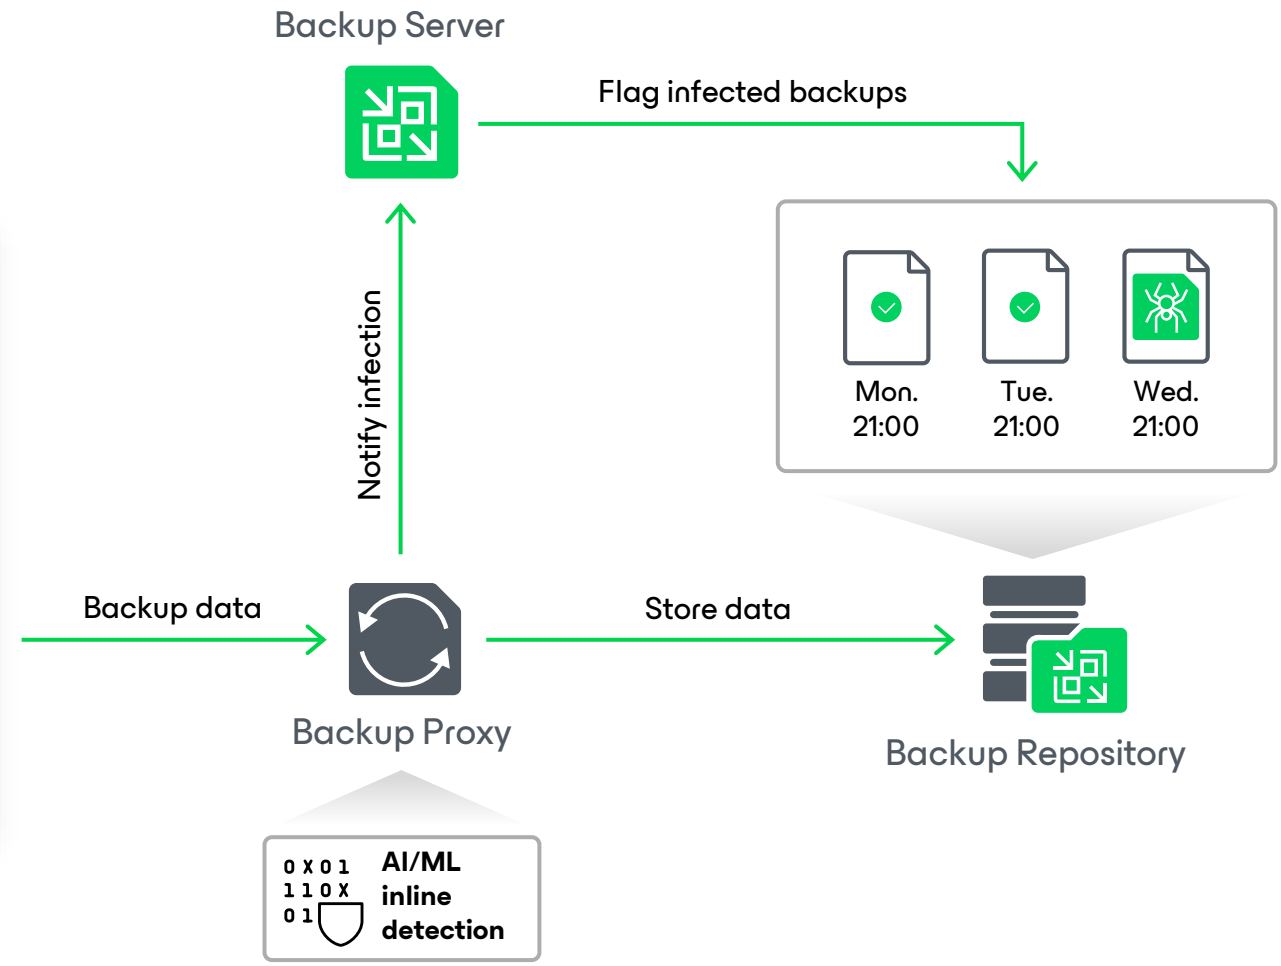
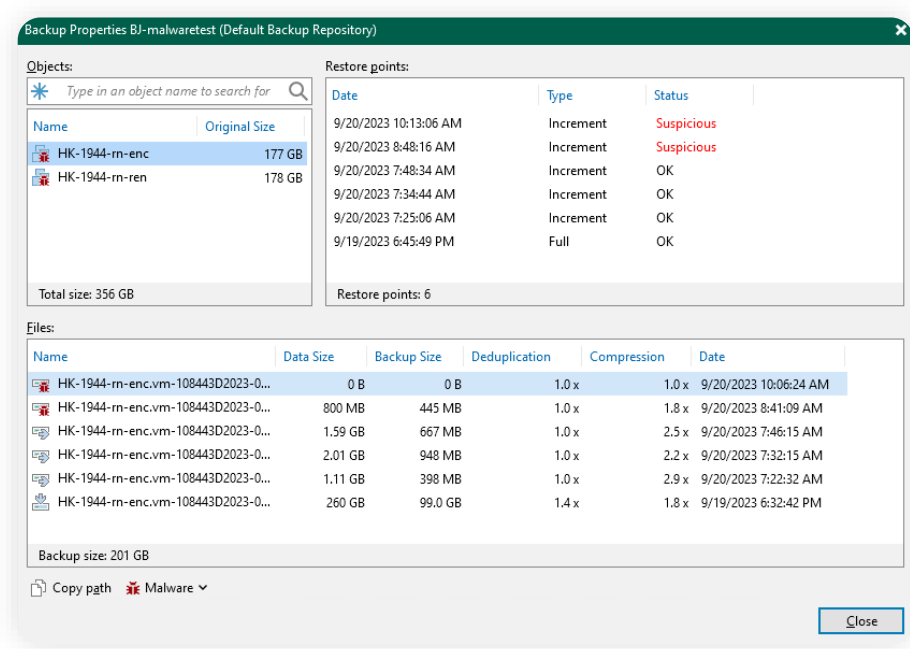
Using a **trained Machine Learning** model and a low-impact **entropy analysis**, Veeam can **detect data encrypted by ransomware**.

The same engine **detects other signs of malware or cyberattack** such as Onion links, directly in the backup stream.



Best Practices: Security & Resiliency

Malware Detection. How Inline Scan works.

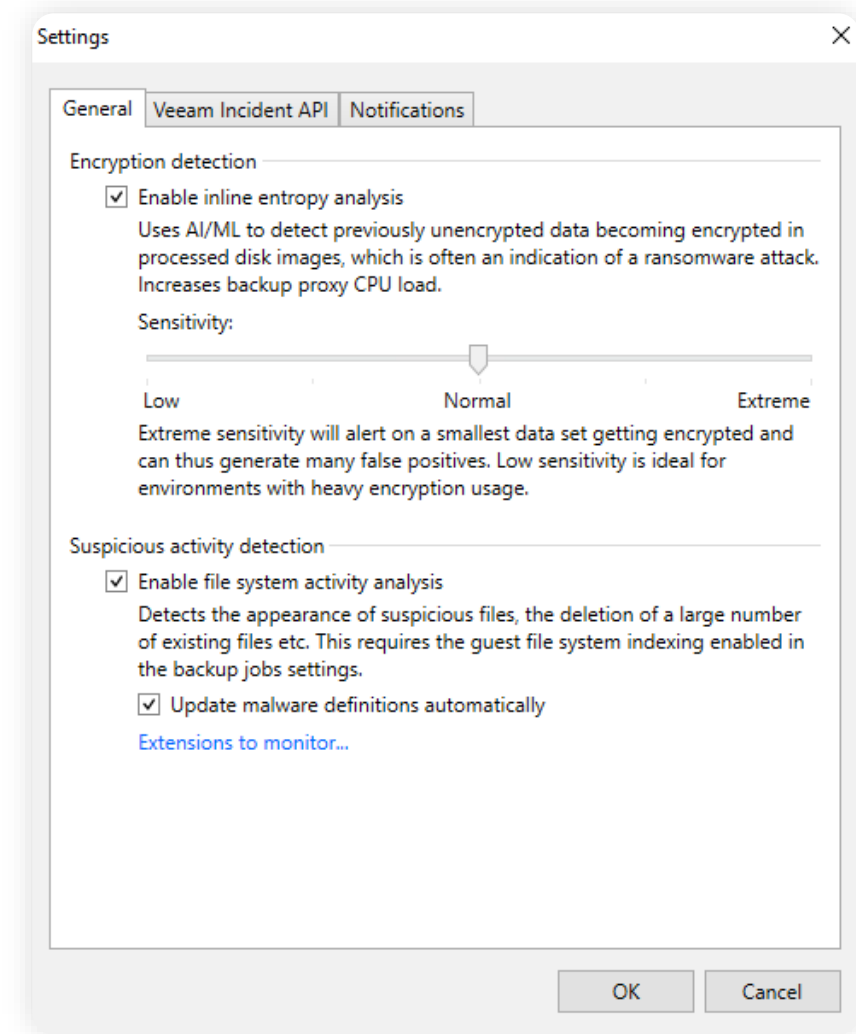


Best Practices: Security & Resiliency

Malware Detection. Detection via Guest Index.

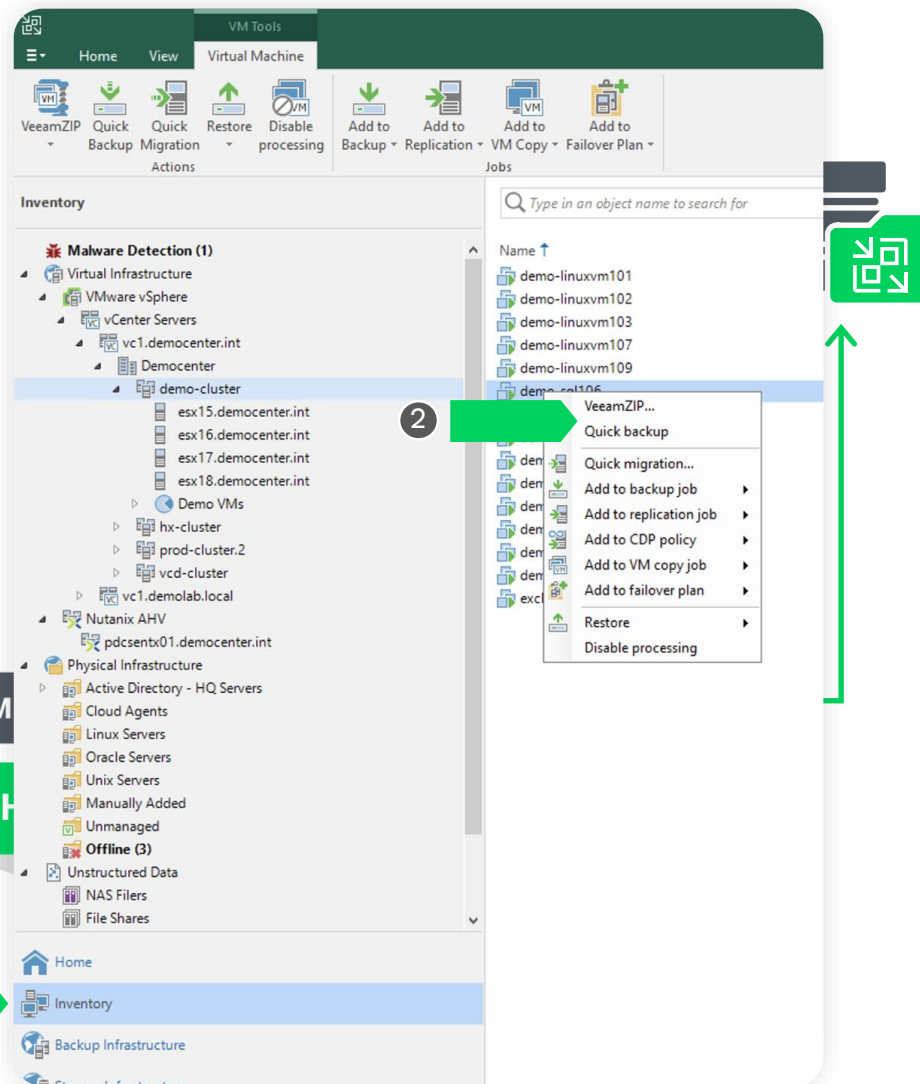
Also named **Suspicious File System Activity Detection** for its capabilities of searching for files with known malware file extensions, ransom notes, and similar flags of malware presence.

Analyzes file system activity by comparing guest indexes to detect suspicious changes like bulk deletions or renamings of known files.



Tips and tricks

Quick Backup & VeeamZIP



Quick Backup offers on-demand incremental backup for VMs without setting up a new job. It adds an extra restore point to existing backup chains, suitable for both incremental and reverse incremental backups. However, it requires a previously successful full backup for the VMs.

VeeamZIP functions like a full VM backup, creating an independent full backup file (.vbk). You can store this file in various locations, but note: VeeamZIP backup files appear under "Backups → Disk (Exported)" in the Home view.

Quick Backup and VeeamZIP can be created from the **Inventory view**.

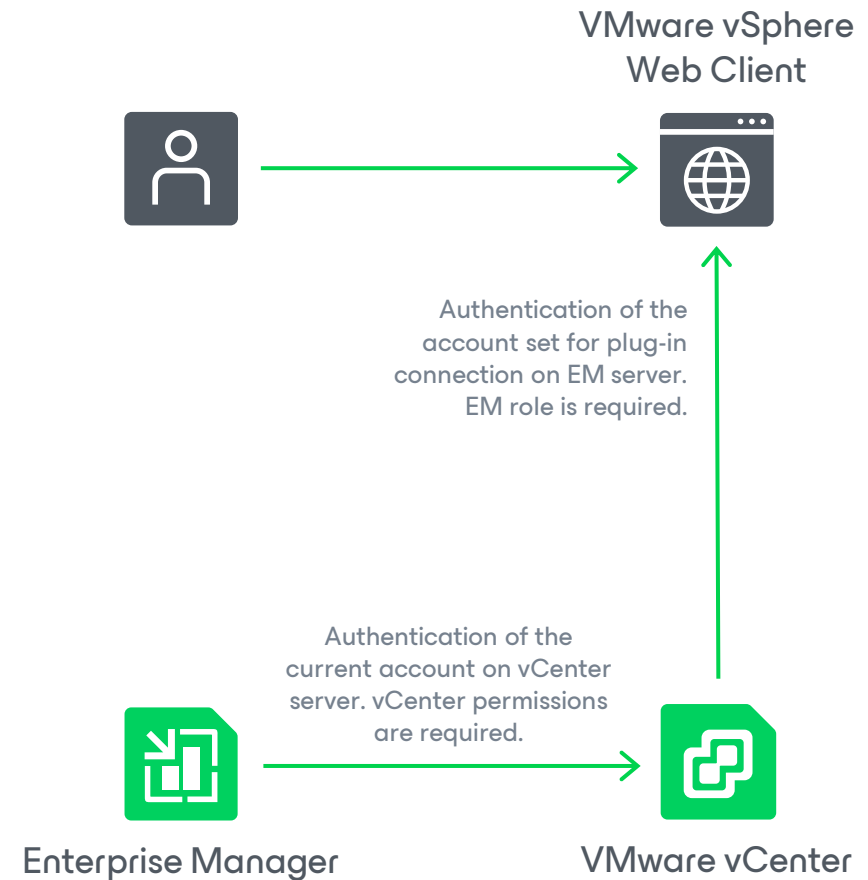
Enterprise Manager + vCenter

Veeam Plug-in for VMware vSphere Client **streamlines backup management for vSphere admins**. It lets them check Veeam Backup & Replication status, monitor job outcomes, and view storage stats directly within vSphere Client.

Admins can **identify unprotected VMs, plan capacity, and create restore points** using VeeamZIP and Quick Backup.

Configuration details:

- vSphere Client 7.0.0 or earlier – local install on vCenter Server.
- vSphere Client 7.0.1 or later – remote install on Veeam Backup Enterprise Manager server.



Enterprise Manager + vCenter

The screenshot shows the Veeam Backup & Replication Summary page within the vSphere Client. The interface includes a top navigation bar with the vSphere Client logo, a menu, a search bar, and refresh/help icons. The main content area is titled 'Veeam Backup & Replication' and has two tabs: 'Summary' (selected) and 'Settings'. On the left, there is a green icon representing the Veeam logo and a list of system components: Backup servers (2), Proxy servers (6), Repository servers (5), Running jobs (0), and Scheduled jobs (13). On the right, a bar chart shows the status of VM backups: Successful VM Backups (15, 83%), VMs with warnings (3, 17%), and Failed VMs (0, 0%). Below this, there are two main sections: 'VMs Overview' and 'Job Statistics'. The 'VMs Overview' section contains a table with the following data:

Protected VMs:	11
Backed Up	10
Replicated	1
Restore points:	14
Full backup size	68.35 MB
Incremental backup size	5.41 GB
Replica restore points size	32.00 bytes
Source VMs size	160.30 GB
Successful backup sessions ratio	100%

At the bottom of this section, there are buttons for 'LAST 24 HOURS' and 'VIEW PROTECTED VMS REPORT...'. The 'Job Statistics' section contains a table with the following data:

Running jobs:	0
Scheduled jobs:	13
Backup	10
Replica	3
Total jobs runs:	13
Successful jobs	9
Jobs with warnings	3
Jobs with errors	1
Max job duration:	3 hours 52 mins

At the bottom of this section, there are buttons for 'LAST 24 HOURS' and 'VIEW LATEST BACKUP JOB STATUS REPORT...'. The bottom of the interface shows 'Recent Tasks' and 'Alarms' tabs, and a home icon.

Copy and Move Backup

Veeam Backup & Replication allows you to move all backups of a backup job to another repository or to move specific workloads and their backups to another job.

Use case examples:

Move backups to different repository

Move backups to another job

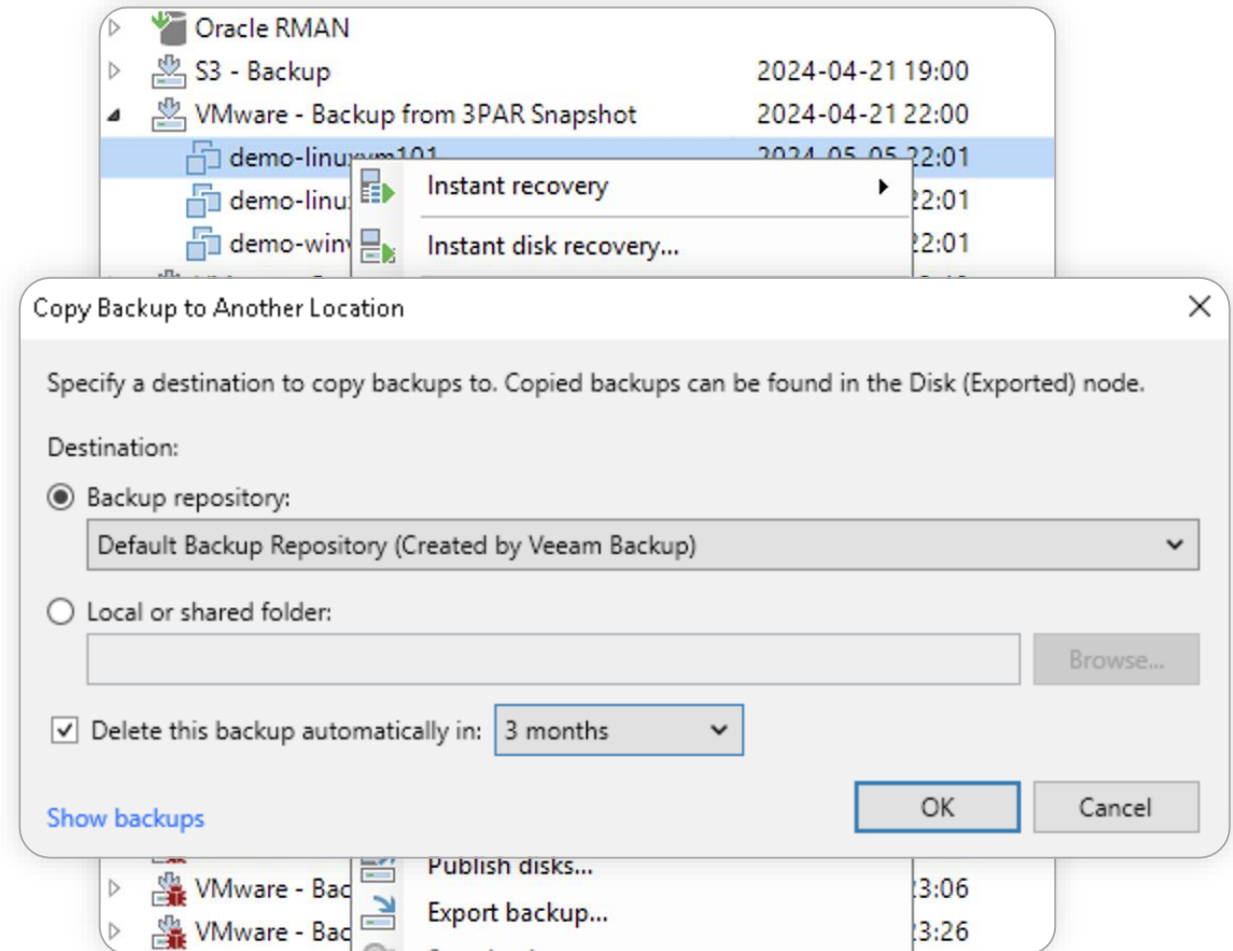
Copy backups to different repository

Migrate ReFS to XFS for Hardened Repository

Migrate NTFS to ReFS

Re-balance Scale-Out Repository

Scale-Out Repository extent evacuation

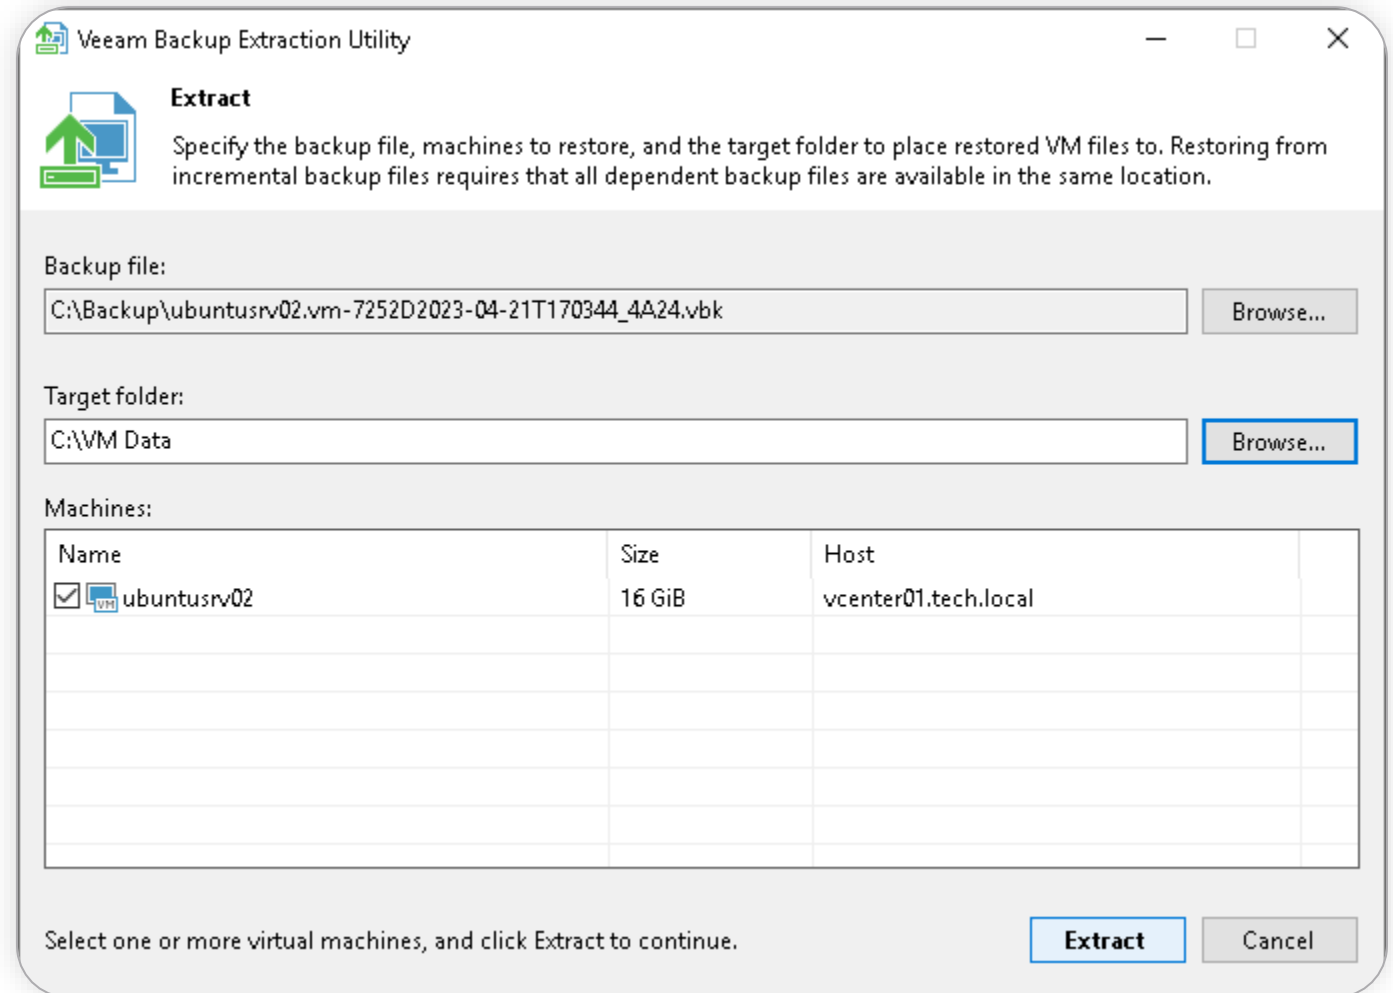


Veeam Extract Utility

- Independent tool for recovery
- GUI or CLI

Use case:

Recover machines even if backups are removed from Veeam Backup & Replication or Veeam Backup & Replication is not installed.



We are almost
done...

Want some hands-on experience?

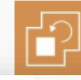
 **Intelligent offsite backups with Veeam Data Platform**


 **Embrace DevOps with Extensibility and**

 **Ransomware Defense with Immutable Veeam**

 **Test drive Veeam Backup and Replication V12**

 **Getting Started With Veeam-Powered Managed**

 **Test Drive Veeam Recovery Scenarios**

 **Machine Learning, Terraform, Microsoft Teams BOT & Ransomware clean room with Veeam API**

Learn how to achieve advanced capabilities with the Veeam REST API. Topics covered include how to address the malware challenge with Veeam clean room, leveraging Machine Learning for backup data analysis, and how to meet DevOps challenges with integrating Terraform with Veeam, amongst other topics.

Product: Veeam Backup and Replication
Use Cases: Data Recovery, Data Security, Data Mobility
Technical Level: Advanced

[REQUEST LAB](#)

 **Avoid platform lock-in with Veeam Flexible data recovery**

Learn how to achieve data recovery and data mobility flexibilities with Veeam Data Platform. Topics include Microsoft Active Directory backups and Recovery with Veeam, recover data from on-premises to AWS, instant recovery from VMware vSphere to Microsoft Hyper-V, Secure Restore for Anti Virus and Malware scans and NAS data recovery, amongst other topics.

Product: Veeam Backup and Replication
Use Cases: Data Mobility, Data Flexibility
Technical Level: Intermediate

[REQUEST LAB](#)

 **AI Powered Ransomware Threat Management with Veeam Data Platform**

Learn about Artificial Intelligence powered Ransomware threat management capabilities now available with Veeam Data Platform (12.1 update). Key topics include AI powered inline malware detection, YARA rules support for malware detection, automated malware detection with YARA rules & Anti-virus software and immutable data backups to combat Ransomware challenge.

Product: Veeam Backup and Replication
Use Cases: Data Security
Technical Level: Intermediate

[REQUEST LAB](#)



<https://go.veeam.com/hands-on-lab-experience#renewal>

Veeam Certified Engineer (VMCE)

- Instructor or On – Demand
- Available in 100 countries
- Theory and Labs
- <https://www.veeam.com/support/training/vmce-training.html>

Additional Resources

- Calculators (<https://www.veeam.com/calculators>)
- Help Center (<https://www.veeam.com/support/help-center-technical-documentation.html>)
- VBR Best Practices (<https://bp.veeam.com/vbr/>)
- KB: Security (<https://www.veeam.com/knowledge-base.html?type=security>)
- R&D Forums (<https://forums.veeam.com/>)

The Veeam logo is centered on a green background. It consists of the word "veeam" in a white, lowercase, sans-serif font. The text is contained within a white-outlined rectangular box with rounded corners and a notch on the right side. Behind the box, there are two large, semi-transparent, light-green shapes that resemble stylized arrows or chevrons pointing towards each other.

veeam

Follow us!



Join the community hub:

