

HPE REFERENCE CONFIGURATION FOR VEEAM BACKUP & REPLICATION VERSION 10 ON HPE STOREONCE SYSTEMS WITH HPE STOREONCE CATALYST



CONTENTS

Executive summary.....	3
Solution overview.....	3
HPE StoreOnce and advanced Veeam restore features	4
Solution components.....	6
Hardware.....	6
Software	6
Best practices and configuration guidance for the solution.....	7
Software and security requirements.....	7
Supported topologies and use cases for Catalyst-based repositories.....	7
Creating a Catalyst store for Veeam backups.....	10
HPE StoreOnce Catalyst advanced settings	12
Configuring access permissions to an HPE StoreOnce Catalyst store.....	12
How to connect to an HPE StoreOnce Catalyst store via IP or Fibre Channel	15
Initial Fibre Channel configuration.....	15
Deploying a Veeam Gateway server.....	16
Configuring the Veeam Gateway Server for CoFC.....	18
Creating an HPE StoreOnce Catalyst-based Veeam backup repository.....	18
Creating multiple Veeam backup repositories inside the same HPE StoreOnce Catalyst store.....	25
Creating a Veeam backup job for VMs.....	27
Creating a Veeam backup job for Windows and Linux computers	36
Copying Veeam backups to a remote site – Backup replication.....	40
HPE Cloud Volume Backups with Veeam Backup & Replication	50
Recovering data from an HPE StoreOnce Catalyst repository	53
Migrating Veeam Backup repositories to/from an HPE StoreOnce Catalyst store.....	53
Summary.....	60
Resources and additional links.....	61



EXECUTIVE SUMMARY

Veeam® Backup & Replication™ offers a data protection solution for backup and replication together in a single software solution. [HPE StoreOnce Systems with HPE StoreOnce Catalyst](#), coupled with Veeam Backup & Replication software, provide a deduplication-based backup architecture to significantly reduce storage space and increase performance compared to traditional NAS protocol connectivity. With available source-side or client-side deduplication capabilities, HPE StoreOnce with Veeam Backup & Replication provides deduplication technology to effectively backup, replicate, and restore virtual and physical environments.

Veeam Backup & Replication integrates with HPE StoreOnce via NAS and HPE StoreOnce Catalyst protocols. This Reference Configuration focuses on HPE StoreOnce Catalyst, hereinafter referred to as “Catalyst,” which provides these main advantages in comparison to traditional NAS protocols:

- Source-side deduplication across LAN, WAN, and Fibre Channel (FC) connections. This leads to more efficient backup data transfer over the network, reduced storage space, and faster performance in comparison to backup targets with NAS protocol (CIFS and NFS) connectivity
- Virtual Synthetic Full backup jobs for lighter and faster backup
- Virus and Ransomware protection. Data written to a Catalyst store is invisible for the Windows® OS hosting Veeam. For this reason, even in the case of Ransomware infection, your backup files are invisible and protected from corruption
- HPE StoreOnce Catalyst Copy as storage-based replication that is managed by Veeam

This Reference Configuration provides information on the setup and configuration of Veeam Backup & Replication with HPE StoreOnce Catalyst for both on-site and off-site backup copies. For off-site copies of backup data, three use cases are compared.

1. Veeam-managed storage-based replication – it does not require Veeam Gateway servers and is a fast and efficient methodology based on HPE StoreOnce Catalyst Copy.
2. Direct write from a server on the source site to a remote HPE StoreOnce appliance – it is a Backup Copy Job (BCJ) based on a highly deduplicated Catalyst communication over WAN.
3. Traditional replication is based on two Veeam Gateway servers – one at the local site and one at the remote site.

Target audience: Presales consultants, solution architects, storage operators, data center managers, enterprise architects, and deployment and implementation engineers. Working knowledge of Veeam Backup & Replication and HPE StoreOnce is recommended.

Document purpose: The purpose of this document is to describe a Reference Configuration that includes best practices and highlights recognizable benefits to technical audiences. It provides information for setting up and managing data availability infrastructures that include Veeam Backup & Replication and HPE StoreOnce backup appliances configured with HPE StoreOnce Catalyst backup targets. This document is current to Veeam Backup & Replication version 10 and HPE StoreOnce VSA version 4.2.4. For more details about Veeam Backup & Replication, see the [Veeam Backup & Replication User Guide](#). For HPE StoreOnce information, please visit hpe.com/storage/StoreOnce.

This Reference Configuration describes solution testing performed by Hewlett Packard Enterprise in July 2020.

SOLUTION OVERVIEW

This Reference Configuration focuses on the benefits of deploying an HPE StoreOnce Catalyst backup repository with Veeam Backup & Replication version 10. The solution environment layout illustrated in [Figure 1](#) consists of several Veeam, VMware®, Microsoft Hyper-V, and Hewlett Packard Enterprise components. A Veeam Backup server interfaces with a VMware vCenter® Server or System Center Virtual Machine Manager (SCVMM) to conduct numerous tasks, including backups, storage snapshot orchestration, and restores. Veeam Proxy services read data from the VMware vSphere® Hypervisor or Microsoft Hyper-V infrastructure before sending the data to a Veeam Gateway server, and then on to the HPE StoreOnce Catalyst repository.

The Veeam Backup & Replication architecture consists of two data mover services: one service on the source proxy and one on the backup repository. In the case of HPE StoreOnce where Veeam does not install a data mover on the HPE StoreOnce appliance, Veeam uses a gateway server to host this service. The gateway server sits in front of the HPE StoreOnce appliance in the data flow and enables source-side deduplication with advanced management of Veeam backup files. The automatic selection option allows Veeam to choose a gateway server—a role that can be assigned to any Windows server that has access to the HPE StoreOnce appliance. If no additional Windows servers are added, Veeam will use the default backup proxy for this role. To ensure that, only servers capable of serving this role are used and a minimum of



network device hops occur, it is recommended to explicitly select this server. This is a requirement if using FC connectivity, in which case the chosen server must be zoned to the HPE StoreOnce appliance.

A Windows agent or a Veeam Gateway server can be used for Catalyst over Ethernet (CoE) or Catalyst over Fibre Channel (CoFC) backups. As shown in Figure 1, Veeam can run a backup to a Catalyst store in several ways:

- Using only a Veeam Agent
- Using only a Veeam Gateway server
- Using a Veeam proxy server and a Veeam Gateway server

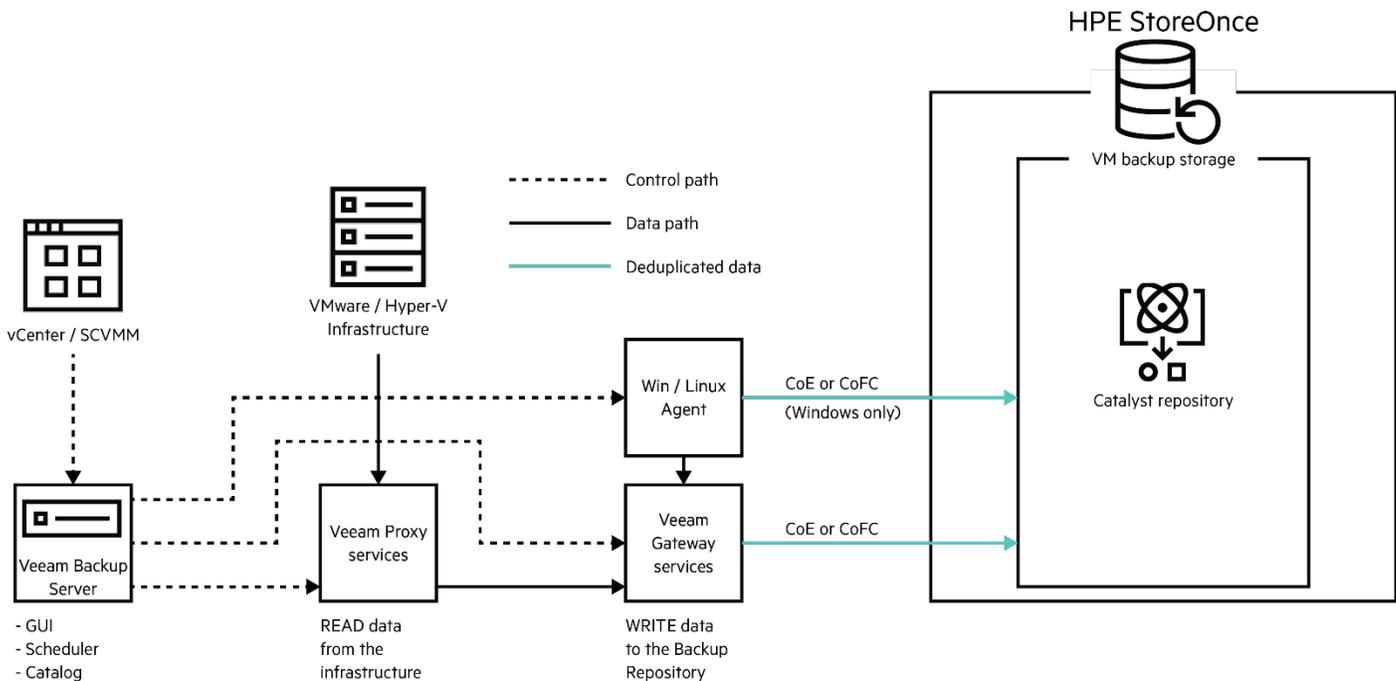


FIGURE 1. Overview of the main components in a Veeam Backup & Replication and HPE StoreOnce deployment

NOTE

Windows Agent can also run the gateway component that executes HPE StoreOnce Catalyst, which can be desirable for large, highly transactional servers.

HPE StoreOnce and advanced Veeam restore features

HPE StoreOnce supports the use of advanced Veeam restore features that require a higher random I/O ratio, such as Instant VM Recovery (IVMR) and DataLabs. If these features are going to be used intensively, additional considerations can be found in the [How to use HPE StoreOnce Systems with Veeam Backup and Replication](#) guide.

If the Veeam advanced restore features will be used frequently, there are two options to ensure restore performance requirements: Veeam snapshot orchestration and two-tier data protection.

Veeam snapshot orchestration

Veeam's snapshot orchestration on HPE primary storage takes advantage of the majority of Veeam's restore features, and then backups can be sent directly to an HPE StoreOnce appliance. This accomplishes the same goals of the two-tier approach while still utilizing HPE StoreOnce as the primary backup target.



Figure 2 demonstrates IVMR from a Veeam storage snapshot.

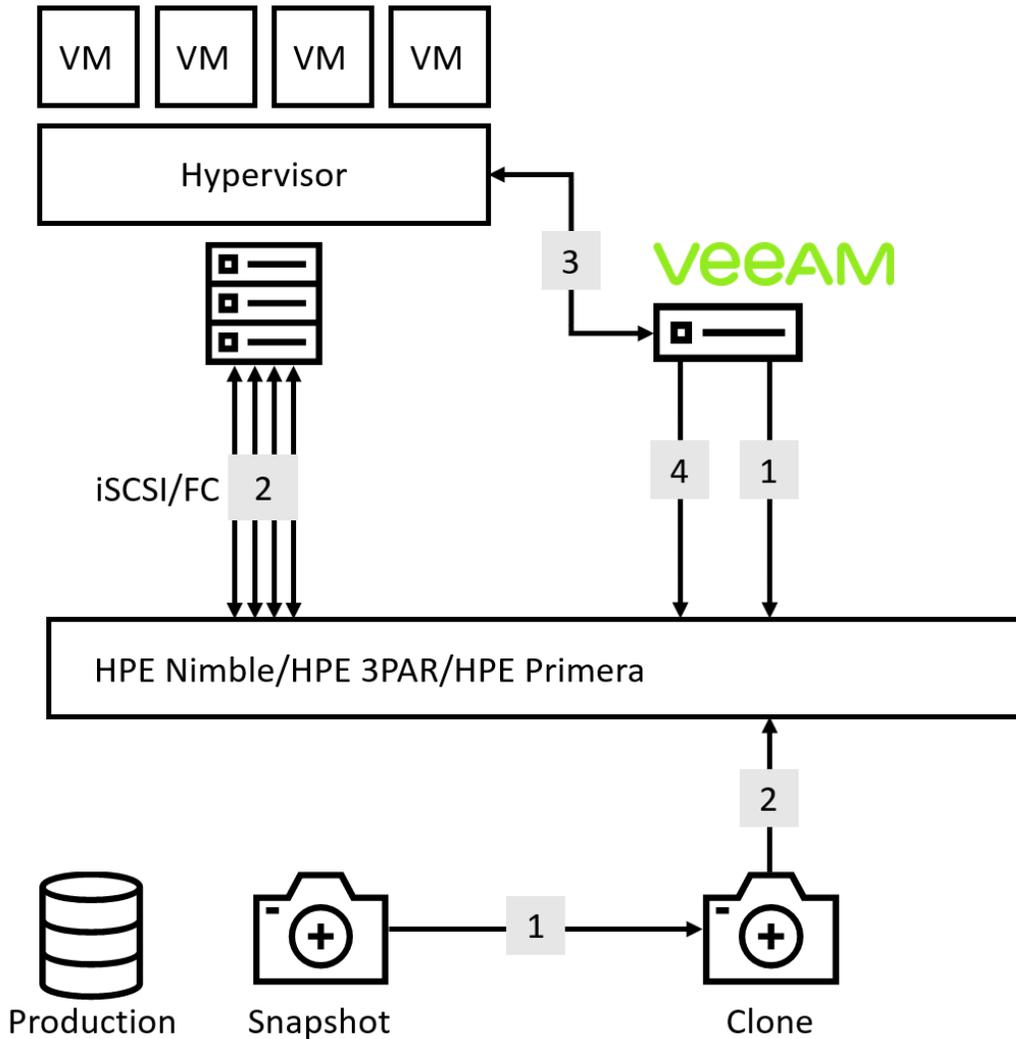


FIGURE 2. IVMR from a storage snapshot overview

The storage snapshot integration illustration shown in Figure 2 can be outlined as follows:

1. A storage snapshot writable clone is created. This preserves the original snapshot integrity.
2. The cloned snapshot is attached to the hypervisor as a new SAN volume. It is not a vPower NFS, but a storage volume.
3. Veeam initiates recovery operations against the chosen VM in the cloned snapshot.
4. After recovery, a clean-up process is performed on the hypervisor and the storage array.

For more details regarding official Veeam architectural guidelines for deduplicating storage systems, including HPE StoreOnce, see the Veeam Knowledge Base Article at veeam.com/kb2660.

For additional information on data protection with snapshot orchestration with HPE 3PAR, HPE Primera, and HPE Nimble, see:

- HPE 3PAR/HPE Primera, hpe.com/V2/GetDocument.aspx?docname=a00061747enw
- HPE Nimble, hpe.com/v2/getdocument.aspx?docname=a00079582enw



Two-tier data protection

A two-tier data protection solution should be considered if advanced Veeam restore features will be used frequently. For example, in a two-tier data protection solution, Veeam could target HPE Nimble or HPE Apollo storage as the primary backup target before using a Backup Copy job to copy the backup to the HPE StoreOnce appliance. This would enable best-performance to restore options from the primary backup target and long-term, highly deduplicated retention on HPE StoreOnce. To see a representation of this data flow, see Figure 3.

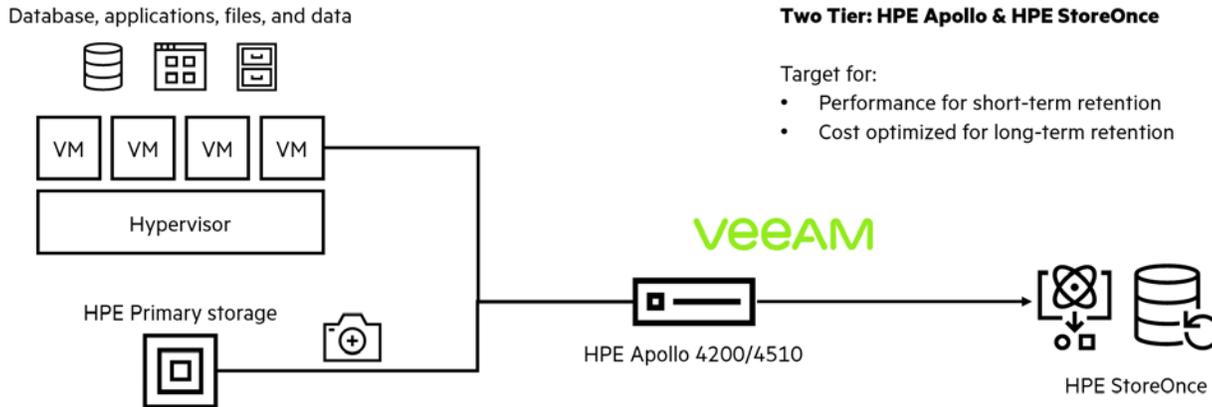


FIGURE 3. Veeam best practice for high-frequency use of full-featured restore options with random I/O characteristics

For additional information on two-tier data protection with HPE Nimble and HPE Apollo, see:

- HPE Nimble, hpe.com/v2/getdocument.aspx?docname=a00079582enw
- HPE Apollo, hpe.com/V2/GetDocument.aspx?docname=a00064646enw

SOLUTION COMPONENTS

This section describes the hardware and software components used to create this Reference Configuration.

Hardware

HPE ProLiant DL580 server

This solution is based on two identical [HPE ProLiant DL580 servers](#), each configured with four Intel® Xeon® E7-4830 v3 CPU @ 2.10 GHz and 3 TB of memory. LAN connectivity is through an HPE Ethernet 1 Gb 4-port 331FLR Adapter and an HPE Ethernet 10 Gb 2-port 530SFP+ Adapter. Storage connectivity is through two HPE SN1100E 16 Gb 2P FC HBAs. The onboard HPE ProLiant integrated Lights Out (iLO) server management requires a separate network connection and IP address.

HPE 3PAR StoreServ 9450 storage

Storage capacity is provided by an [HPE 3PAR StoreServ 9450 array](#) that was interconnected via a 16 Gb FC fabric. VMware datastores used storage from the HPE 3PAR volumes, including the datastore used by the vSphere Hypervisor. [HPE Nimble Storage](#) and [HPE Primera](#) arrays can also be used.

Software

HPE StoreOnce VSA

VMs often have a high deduplication potential because operating system core files and other application data do not change significantly day-to-day. HPE StoreOnce is an imperative part of this solution. HPE StoreOnce Catalyst allows clients to deduplicate Virtual Machine Disk (VMDK) files—in addition to physical server backup data—before the data is transferred to an HPE StoreOnce appliance.

For this Reference Configuration, an HPE StoreOnce VSA was deployed and configured with software version 4.2.4. HPE StoreOnce appliances can also be used.



vSphere Hypervisor

VMware vCenter Server 6.7 and VMware ESXi™ 6.7 were installed on the host servers. Using vCenter Server, a VMware cluster is created and the two ESXi hosts are added to it. VMware vSphere® High Availability and VMware vSphere® Distributed Resource Scheduler™ are enabled on the cluster. As well as being a prerequisite for vSphere HA, deploying the HPE StoreOnce VSA into a VMware cluster offers these advantages:

- Increases availability by enabling VMware vSphere® Storage vMotion® to re-host VMs to reduce unavailability caused by host maintenance
- Proactively migrates a VM off a host where alerts show it is at risk of failure or has degraded performance

Veeam Backup & Replication

Veeam Backup & Replication version 10.0.0.4461 was used in this solution. Veeam Backup & Replication runs as an application on a Windows server.

For information on platform support, see helpcenter.veeam.com/docs/backup/vsphere/platform_support.html?ver=100.

BEST PRACTICES AND CONFIGURATION GUIDANCE FOR THE SOLUTION

This section provides best practices and guidance to build this Reference Configuration solution environment. Tuning parameters are also provided in some instances.

Software and security requirements

Hewlett Packard Enterprise recommends the following links for proper setup of the software and security components to adequately run Veeam Backup & Replication with HPE StoreOnce and HPE StoreOnce Catalyst:

- hpe.com/Storage/StoreOnceSupportMatrix
- helpcenter.veeam.com/docs/backup/vsphere/platform_support.html?ver=100
- helpcenter.veeam.com/docs/backup/vsphere/deduplicating_appliance_storeonce.html?ver=100
- helpcenter.veeam.com/docs/backup/agents/system_requirements.html?ver=100

NOTE

Starting with HPE StoreOnce generation 4 systems (October 2018), Catalyst is pre-enabled on all systems and does not require any additional HPE StoreOnce licensing.

Hewlett Packard Enterprise has hidden the Catalyst store from Ransomware in plain sight—and behind an application programming interface (API) that both enhances and simplifies the process of backing up and deduplicating data while making it practically impossible for Ransomware to attack it directly. HPE StoreOnce Catalyst stores do not prevent the rest of the enterprise from being compromised by malware, but they will protect the mission-critical data stored from being either targeted or affected. Ransomware cannot encrypt what it cannot see, and because the Catalyst store does not use standard operating system command instructions for its operations, malware cannot become active. For more information about protecting data against malware with HPE StoreOnce Catalyst, see hpe.com/V2/getpdf.aspx/A0004-2003ENW.pdf.

Supported topologies and use cases for Catalyst-based repositories

HPE StoreOnce can be used as a Veeam Backup & Replication target in many different use cases. Common topologies are shown below.

NOTE

These solutions support HPE 3PAR, HPE Primera, and HPE Nimble storage snapshot integration. For more information on backup from storage snapshots, see the [Creating a Veeam backup job for VMs](#) section.



Backup over LAN

Figure 4 demonstrates a backup-over-LAN setup where VMs are backed up to a Catalyst store, potentially using the same IP address as the HPE StoreOnce Management IP address.

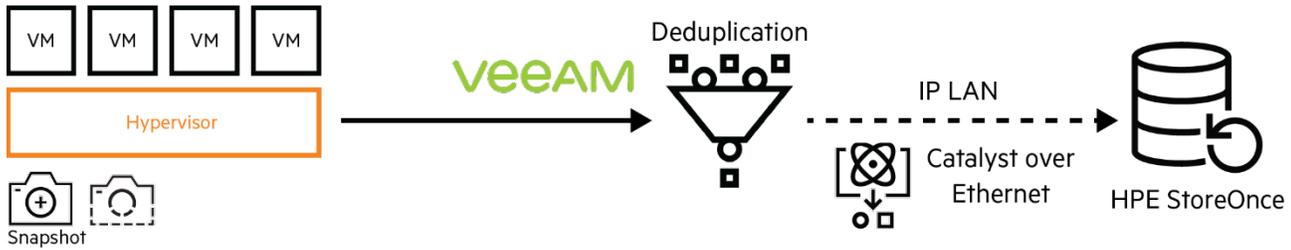


FIGURE 4. Veeam using Catalyst for source-side deduplication over LAN

Backup over FC

HPE StoreOnce Catalyst can also be configured over FC. Figure 5 illustrates backups using Catalyst over Fibre Channel (CoFC) in a SAN.

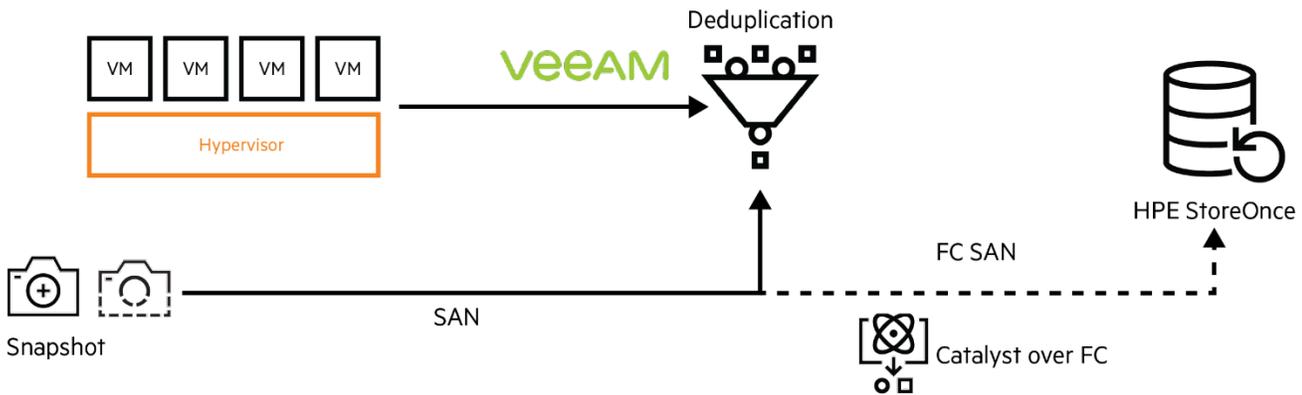


FIGURE 5. Veeam using Catalyst for source-side deduplication over Fibre Channel

Backup over WAN

Additionally, HPE StoreOnce Catalyst backups can be configured over a WAN, as shown in Figure 6. In this case, backup data is deduplicated and compressed before being sent over the WAN, effectively reducing the amount of bandwidth required on the WAN.

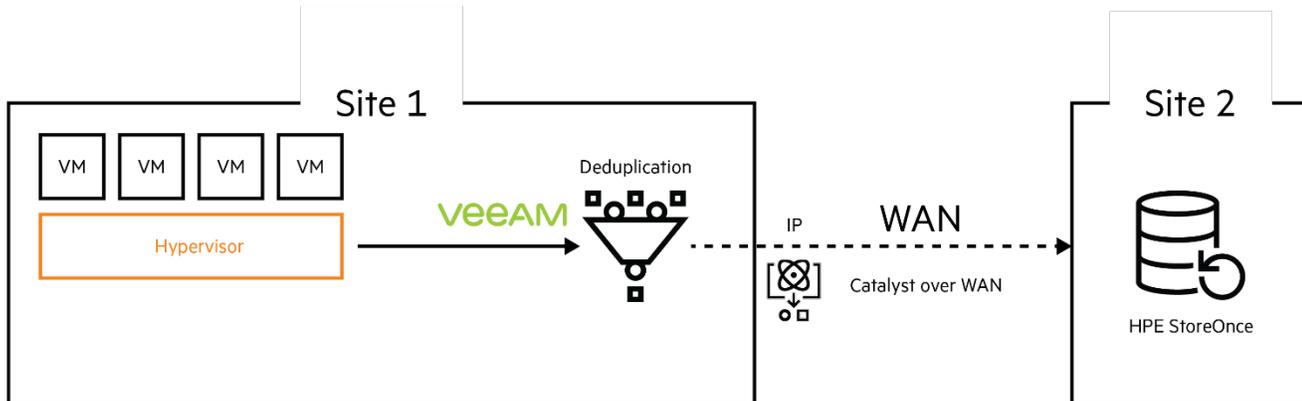


FIGURE 6. Veeam using Catalyst for source-side deduplication over a WAN for direct backup or backup copy to create a secondary copy of local backup data



Backup Copy jobs versus Catalyst Copy jobs

There are two options for copying data from a source repository to a target repository. The first option is a Backup Copy job, demonstrated in Figure 7, which utilizes backup data from the source repository to construct a backup on the target repository device. For more information on Backup Copy jobs, see helpcenter.veeam.com/docs/backup/vsphere/backup_copy.html?ver=100.

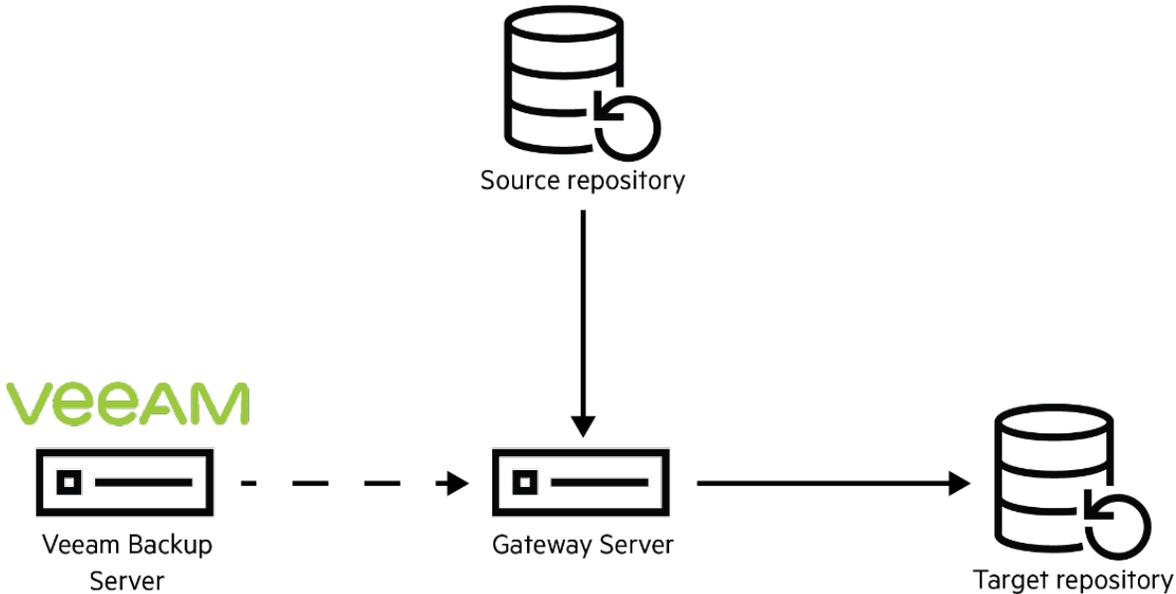


FIGURE 7. Veeam Backup Copy job

The second option, introduced in Veeam Backup & Replication version 10, is to conduct a Catalyst Copy job from the Veeam server to copy deduplicated HPE StoreOnce backup files directly from a source HPE StoreOnce appliance to a target HPE StoreOnce appliance. In this scenario, illustrated in Figure 8, Veeam utilizes HPE StoreOnce APIs to instruct HPE StoreOnce to conduct a Catalyst Copy job to another HPE StoreOnce appliance. This operation can be done in parallel or cascade. For more information, see the [Veeam-managed HPE StoreOnce Catalyst Copy job](#) section.

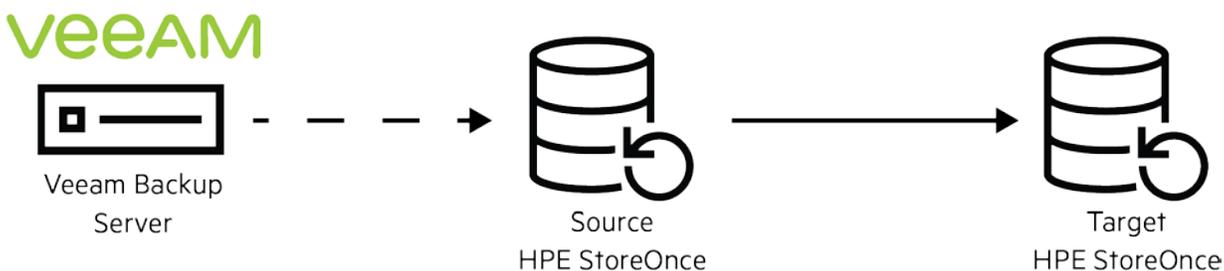


FIGURE 8. Veeam-managed HPE StoreOnce Catalyst Copy job

When copying data between HPE StoreOnce appliances, Catalyst Copy Jobs are recommended over Backup Copy jobs. Backup Copy jobs with multiple HPE StoreOnce appliances involve a gateway server to move data between the appliances. Catalyst Copy jobs more efficiently copy data directly between the two appliances. Additionally, Backup Copy jobs between HPE StoreOnce appliances involve rehydrating deduplicated data on the source HPE StoreOnce appliance and then dehydrating the data on the target HPE StoreOnce appliance. In contrast, Catalyst Copy jobs move data directly without having to rehydrate or dehydrate the data.



Creating a Catalyst store for Veeam backups

Before configuring a Veeam backup repository, an HPE StoreOnce Catalyst store must be created on an HPE StoreOnce appliance that will store the backups. To do this:

1. Log in to the HPE StoreOnce Management Console (HPE StoreOnce GUI) with administrator privileges.
2. As demonstrated in Figure 9, navigate to **Data Services** → **Catalyst Stores** → **Create Store**.

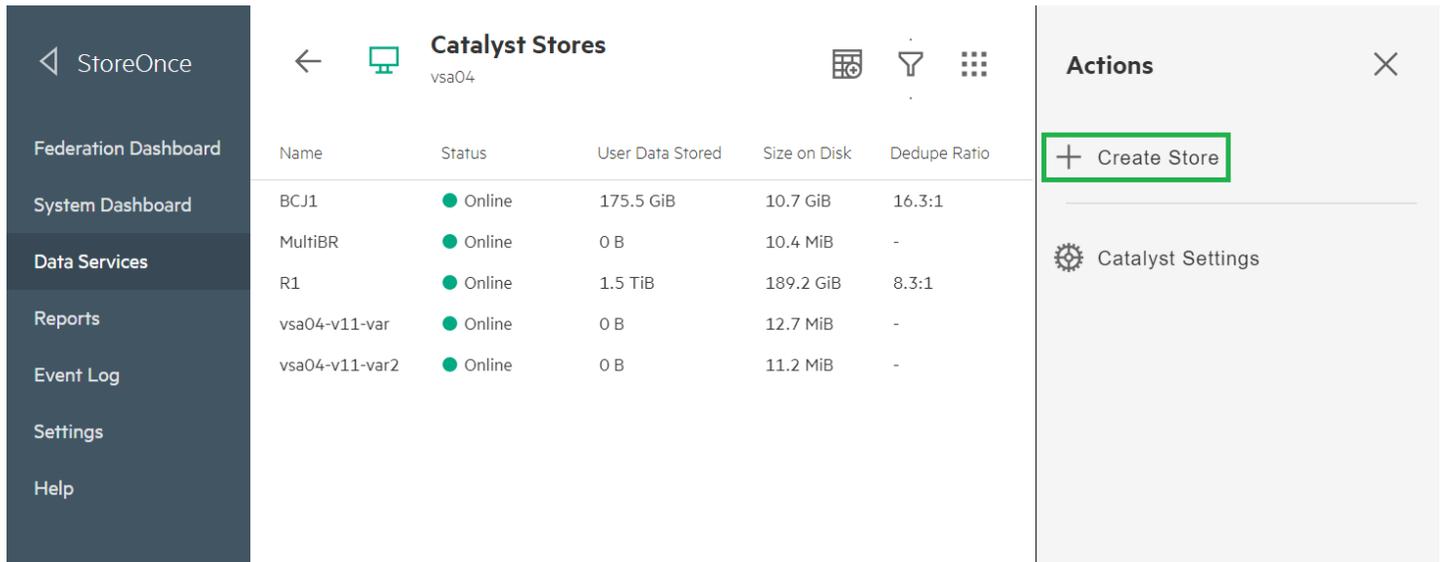


FIGURE 9. Creating a Catalyst Store in the HPE StoreOnce GUI

3. Enter the **Catalyst Store Name** as shown in Figure 10.

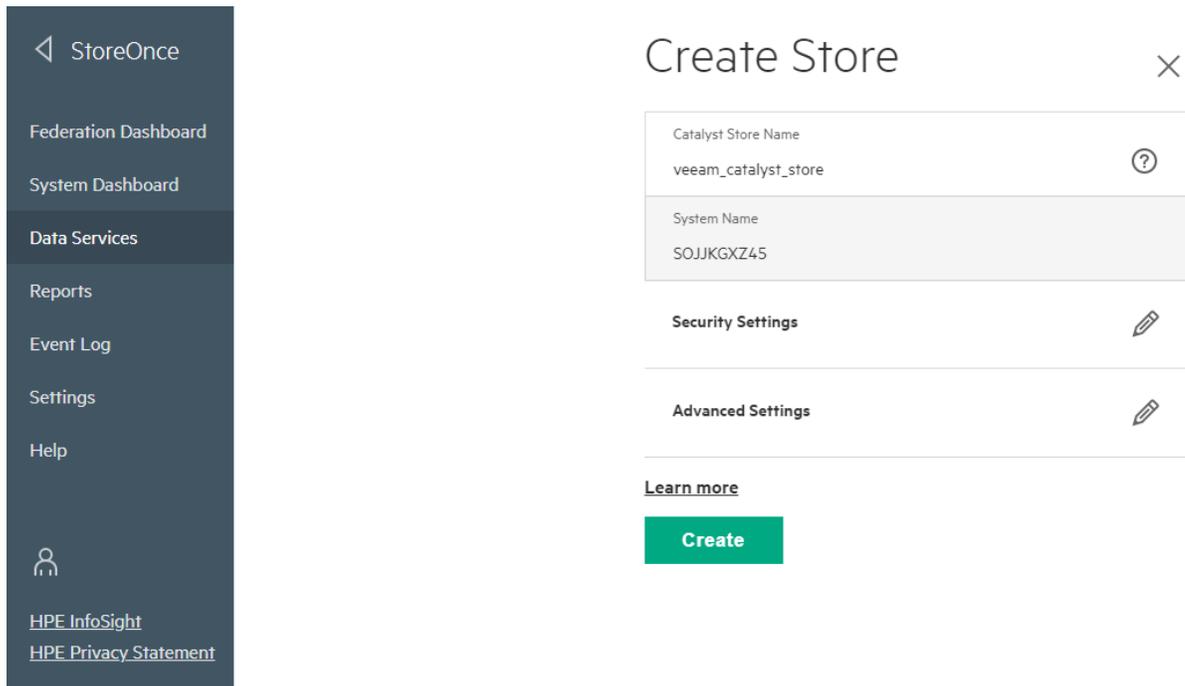


FIGURE 10. New catalyst store configuration settings in the HPE StoreOnce GUI



- If necessary, change the store **Security Setting** details; otherwise, keep the default settings. For more information, see the [HPE StoreOnce Systems Installation Guide](#).

NOTE

To allow HPE StoreOnce encryption, enable **Store Encryption** by editing the **Security Settings** of the Catalyst store, as shown in Figure 11. HPE StoreOnce encryption is a free license included on all HPE StoreOnce Gen 4 models and available in most countries, except for specific countries where the license is not available because of export restrictions.

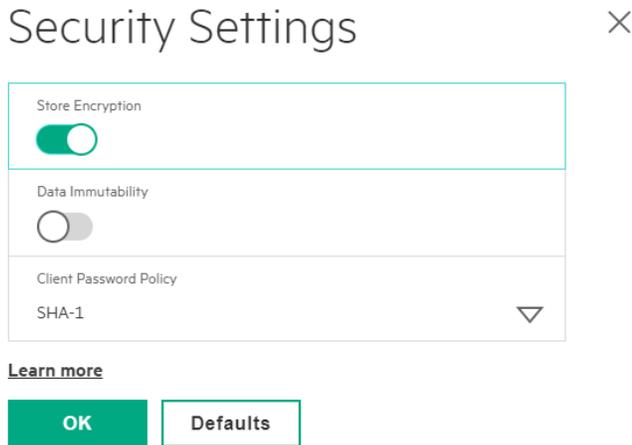


FIGURE 11. Enabling encryption in the HPE StoreOnce Catalyst store GUI

- If necessary, change the store **Advanced Settings** as shown in Figure 12, otherwise, keep the default setting. See the [HPE StoreOnce Catalyst Advanced Settings](#) section for more information on HPE StoreOnce Catalyst advanced settings.

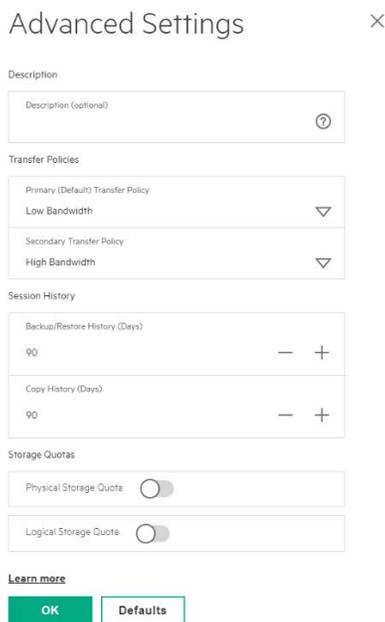


FIGURE 12. HPE StoreOnce GUI showing a Catalyst store configured with primary and secondary transfer policies



HPE StoreOnce Catalyst advanced settings

There are two possible values for the **Transfer Policy**: Low Bandwidth (LBW) or High Bandwidth (HBW). Specifying the value of the transfer policy determines how the backup data is transferred between the Veeam Gateway server and the HPE StoreOnce appliance.

- **LBW**: LBW transfer processes the backup data on the Veeam Gateway server to identify and only transfer unique backup data to the HPE StoreOnce appliance. This is called “source-side deduplication.” Source-side deduplication significantly reduces the physical backup data transferred over the network. It can also reduce backup time by increasing throughput, particularly when there are multiple concurrent backup jobs. It should also be used if the backup or backup copy job is over a WAN because of the significant reduction in bandwidth requirements.
- **HBW**: In contrast, HBW transfer sends all the backup data from the Veeam Gateway server to the HPE StoreOnce appliance. Duplicate data is then identified and deduplicated on the HPE StoreOnce appliance. This is called “target-side deduplication.” Target-side deduplication does not reduce the backup data transfer over the network.

In most situations, the recommended configuration is to keep the default source-side deduplication (LBW) because it increases the overall backup throughput. If a backup job with a single VM is used for testing, an HBW transfer policy might be slightly faster than an LBW transfer policy. As more VM backups are running in parallel, the speed advantage of an LBW transfer policy will increase. LBW is recommended for most use cases and should be used unless there is a reason not to use it.

The network load reduction on the first backup of a VM is lower than for subsequent backups. Starting with the second backup, the network load reduction will be measurable. For restore/read activities, Catalyst does not produce any bandwidth reduction and data is sent rehydrated to the Veeam server. Table 1 demonstrates the necessary transfer policies for source-side and target-side deduplication.

TABLE 1. Source-side and Target-side transfer policies

Desired behavior	Primary (default) transfer policy	Secondary transfer policy
Source-side deduplication	Low Bandwidth	Any value
Target-side deduplication	High Bandwidth	Any value

Storage Quota

Storage quota is used to set a limit on the current Catalyst store size. The **Logical Storage Quota** is a set limit on the amount of data written by Veeam independently of the deduplication. The **Physical Storage Quota** sets a limit on the HPE StoreOnce capacity utilization and is affected by the deduplication.

Session History

The **Session History** defines the duration that session logs are kept for an HPE StoreOnce appliance.

NOTE

Do not create too many Catalyst stores. The maximum number of Catalyst stores is reported in the [HPE StoreOnce Systems User Guide](#). There is no size limit for each Catalyst store, but it is generally better to keep a store at less than 150 TB of used capacity (that is about 2.4 PB as the logical capacity for a 16:1 deduplication).

Configuring access permissions to an HPE StoreOnce Catalyst store

After a Catalyst store has been created, it is possible to set its access permission and restrict its access to specific client and password couplets. It is a good practice to restrict Catalyst store access only to clients that should write and read backup data from the store. HPE StoreOnce Catalyst Client access permission checking is used to control access and is implemented by creating Catalyst client identifiers. Two modes can be set on HPE StoreOnce Catalyst stores for access control:

- **Public Access** allows all clients to have access to the HPE StoreOnce Catalyst stores
- **Client Access** allows a selected group of client identifiers to have access to the HPE StoreOnce Catalyst stores



To create an HPE StoreOnce Catalyst user, in the HPE StoreOnce GUI, navigate to **Settings** → **Catalyst Settings** → **Add Client**, demonstrated in Figure 13.

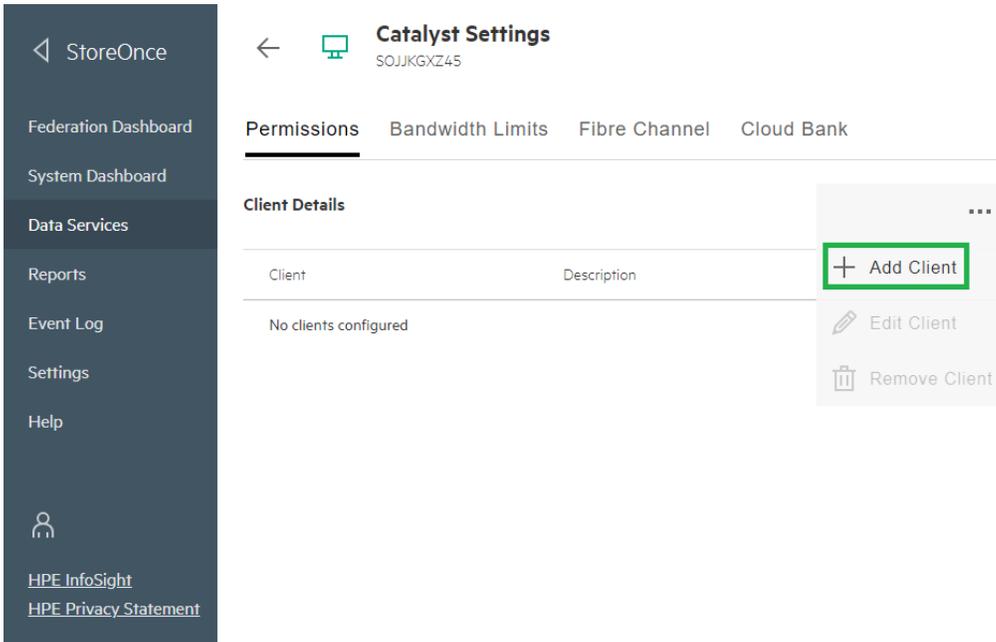


FIGURE 13. Creating a new Client (user) in the HPE StoreOnce GUI

Figure 14 is an example of the **Add Client** screen. Specify the **Catalyst Client Name**, an optional **Description**, and a **Password**.

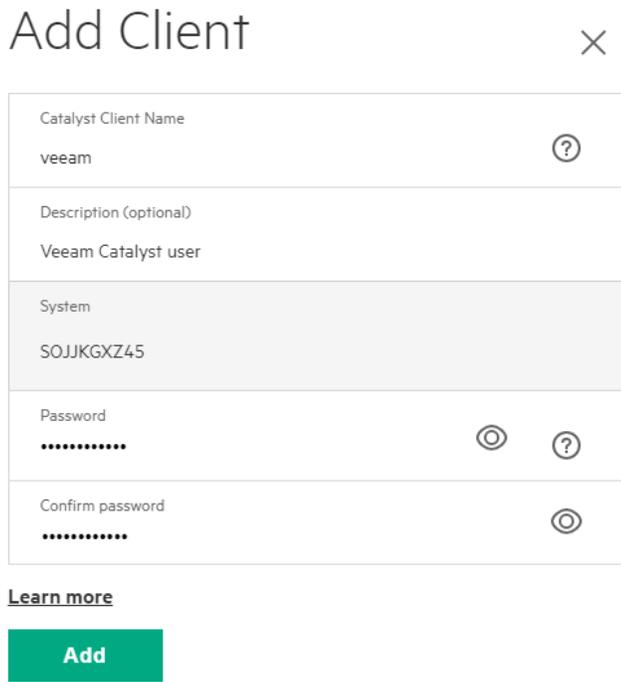


FIGURE 14. Viewing and assigning a password to Catalyst Client (user) in the HPE StoreOnce GUI



NOTE

Take note of the **Catalyst Client Name** and **Password** that has access to the Catalyst store. These credentials are used when you create a Backup Repository based on this Catalyst store. This step is demonstrated in [Creating an HPE StoreOnce Catalyst-based Veeam backup repository](#) section.

To grant the “veeam” client access to a particular store, navigate to **Data Services**, select the desired Catalyst store, click the **Permissions** tab, select the “veeam” client, and then click **Edit Access** as shown in Figure 15.

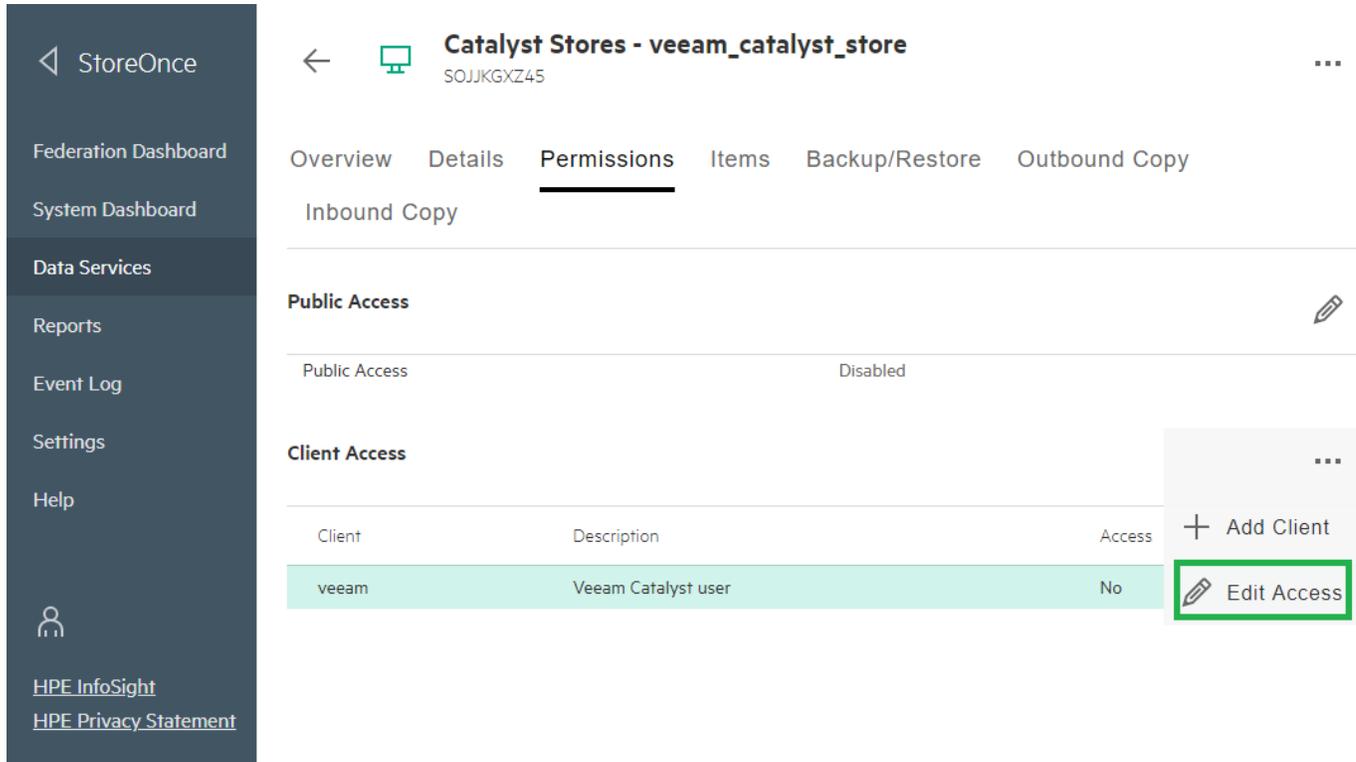


FIGURE 15. Granting a Catalyst client permission to a Catalyst store

After clicking **Edit Access**, the contents of Figure 16 will be displayed. Make sure that **Access** is enabled, then click **Update**.

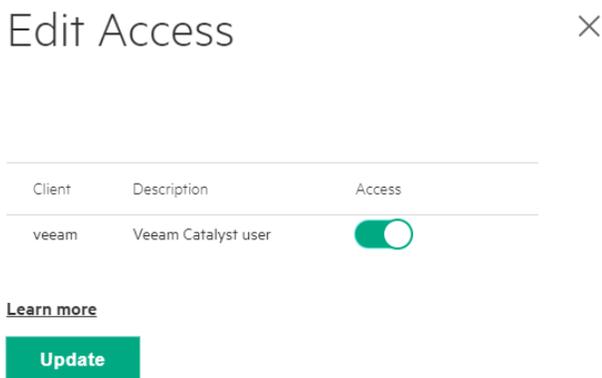


FIGURE 16. Enabling client (user) access to a Catalyst store in the HPE StoreOnce GUI



How to connect to an HPE StoreOnce Catalyst store via IP or Fibre Channel

There are two options for connecting Veeam servers to an HPE StoreOnce Catalyst store.

1. **IP/Ethernet:** CoE is the most commonly used option and it is available with minimal configuration. Make sure that the network infrastructure has enough bandwidth to support the backup load, and that the Veeam Gateway server has access to this network.
2. **Fibre Channel (FC):** CoFC functions the same way as HPE StoreOnce CoE. However, some additional configuration is required to set up the backup and restore connections between the ports on the HPE StoreOnce appliance and the ports on the client servers. This is done using the **Fibre Channel** settings tab of the HPE StoreOnce GUI, shown in Figure 17. This GUI page is only available if CoFC is enabled on the HPE StoreOnce appliance.



FIGURE 17. Configuring FC access to an HPE StoreOnce Catalyst store

Initial Fibre Channel configuration

The following are prerequisites for running CoFC backups:

- The Veeam Gateway server and the HPE StoreOnce appliance are connected to the Fibre Channel network.
- The Windows server that will run the Veeam Gateway service can communicate with the HPE StoreOnce appliance over a Fibre Channel network. To do so, the storage administrator must make sure that any network segregation, such as zoning, is set up to handle the required connectivity between server and appliance.

The FC settings of the HPE StoreOnce appliance are shown in Figure 18, which is located at **Settings → Catalyst Settings → Fibre Channel**.

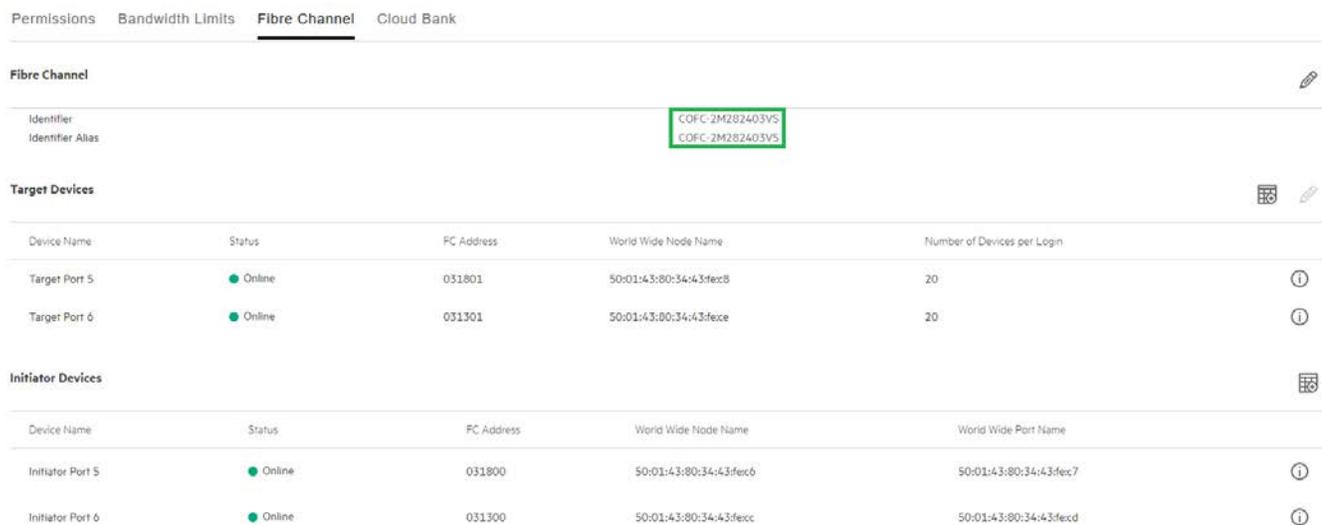


FIGURE 18. The Catalyst “Fibre Channel” settings tab shows Fibre Channel settings and the number of target ports and FC devices



The following information should be obtained from this **Fibre Channel** settings tab:

- Locate the **FC Identifier**. This is the Fibre Channel address of the HPE StoreOnce appliance that is configured in Veeam Backup & Replication. It is in the format “COFC-*<device-id>*,” which will be input into Veeam instead of an Ethernet address.
- Hewlett Packard Enterprise recommends using the default values in the target ports section.
- Locate the **World Wide Port Name** for each port on the HPE StoreOnce appliance. This is the information needed to zone the Veeam Gateway server Fibre Channel ports with the HPE StoreOnce appliance Fibre Channel ports.
- HPE StoreOnce Catalyst Copy over Fibre Channel is a two-way protocol. Zone the source initiator WWN with the destination target WWN, and zone the destination initiator WWN with the source target WWN. The source and destination must be able to communicate with each other over Fibre Channel.
- Locate the **Number of Devices per Login** for each port. This value determines the number of backup and restore connections that are allowed from each Veeam Gateway server Fibre Channel port on each available path to the ports on the HPE StoreOnce appliance.
 - The number of possible FC connections between a Veeam server and an HPE StoreOnce appliance is calculated as:
*number of client ports zoned * number of HPE StoreOnce node ports zoned * number of devices per login.*
Table 2 illustrates how to count the maximum number of FC connections between the Veeam Gateway server FC ports and the HPE StoreOnce appliance FC ports.
 - In most situations, this value should be increased to the maximum number of possible connections available in the Veeam Backup & Replication GUI. The only reason to limit the FC connection is when the HPE StoreOnce unit is shared among multiple protocols and multiple backup applications, and the administrators want to avoid a single protocol or backup application using too many resources.
 - If this value is changed, a device-file rescan on the Veeam Gateway server is needed before the change is recognized.

TABLE 2. Example relationships between device per initiator port and number of data sessions

Veeam server	HPE StoreOnce appliance	Number of devices per login (user modifiable)	Resulting in the maximum number of FC connections between Veeam Gateway server FC ports and HPE StoreOnce FC ports
1 port zoned to	4 ports	1	4
2 ports zoned to	4 ports	1	8
1 port zoned to	1 port	64	64
2 ports zoned to	2 ports	4	16
2 ports zoned to	4 ports	8	64

Deploying a Veeam Gateway server

A Veeam Gateway server is an auxiliary backup infrastructure component that bridges a backup server and backup repository. The role of a gateway server can be assigned to a Windows machine and added as a component in the Veeam Backup server. For more information, see helpcenter.veeam.com/docs/backup/vsphere/gateway_server.html?ver=100. The following are prerequisites for a Veeam Gateway server:

- Administrator privileges are required to access OS-specific device files associated with HPE StoreOnce CoFC devices.
- Fibre Channel users are advised to consult the HPE StoreOnce support matrix (available at hpe.com/Storage/StoreOnceSupportMatrix) to be sure that Veeam Gateway server host bus adapters (HBAs), switches, Fibre Channel driver, and firmware versions are supported.



To configure a gateway server, navigate to **Backup Infrastructure**, right-click **Microsoft Windows**, and then click **Add server...**, as shown in Figure 19.

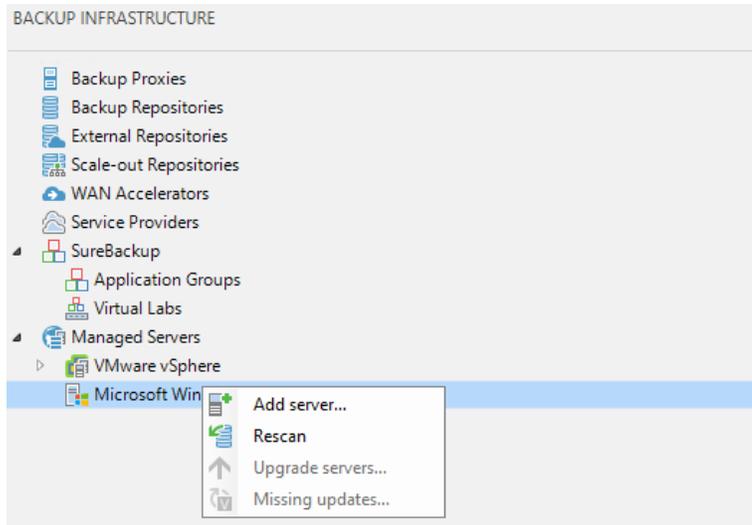


FIGURE 19. Adding a Veeam Gateway server role to a Microsoft Windows server

Supply the DNS name or IP address of a Microsoft Windows server, as shown in Figure 20.

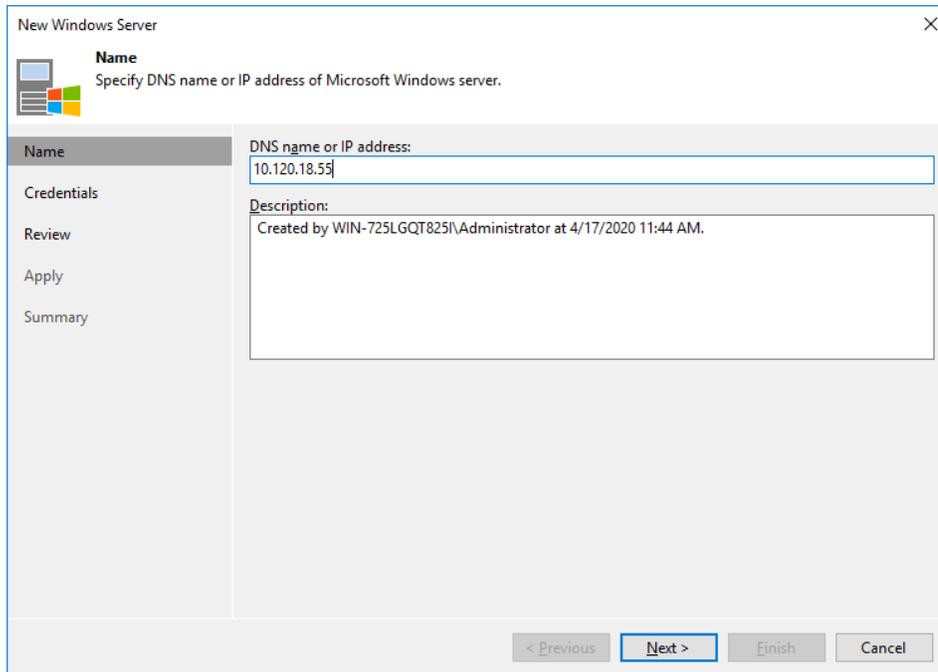


FIGURE 20. Adding a gateway server role to a Microsoft Windows server

Review and apply the settings to the Microsoft Windows server, and then click **Finish** to apply the settings.



Configuring the Veeam Gateway Server for CoFC

HP StoreOnce CoFC presents a device type of “Processor.” These devices are shown as **Other devices** in Windows Device Manager. After zoning the devices or changing the number of devices per initiator port, right-click **Other devices** and click **Scan for hardware changes** to detect the new devices. This is demonstrated in Figure 21.

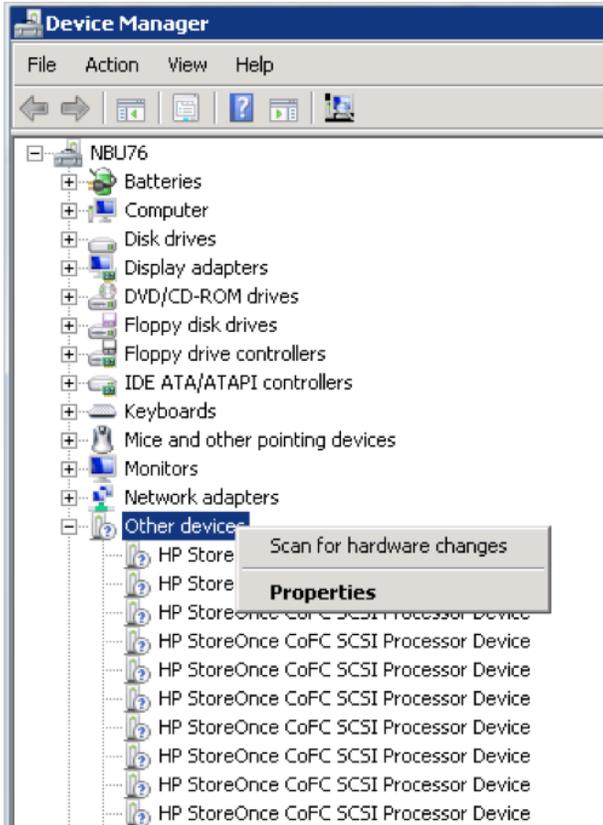


FIGURE 21. HPE StoreOnce CoFC presents a device type of Processor and is shown under “Other devices” in Windows Device Manager

Creating an HPE StoreOnce Catalyst-based Veeam backup repository

A Veeam backup repository is simply defined as storage for Veeam backup files created by Veeam Backup & Replication. A Veeam backup repository receives its data from the Veeam Backup Proxy, where data is read from the infrastructure before being sent to a Veeam Gateway server and then an HPE StoreOnce system. To create a Catalyst-enabled Veeam backup repository, perform the following steps:

1. Click **Backup Infrastructure** in Veeam Backup & Replication.



2. Right-click **Backup Repositories** and select **Add backup repository...**, as shown in Figure 22.

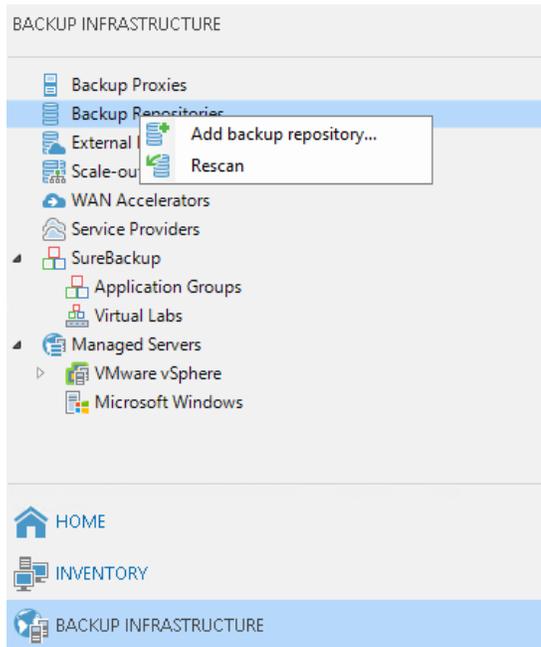


FIGURE 22. Veeam GUI showing the creation of a new backup repository

3. Select **Deduplicating storage appliance** as the repository type, shown in Figure 23.

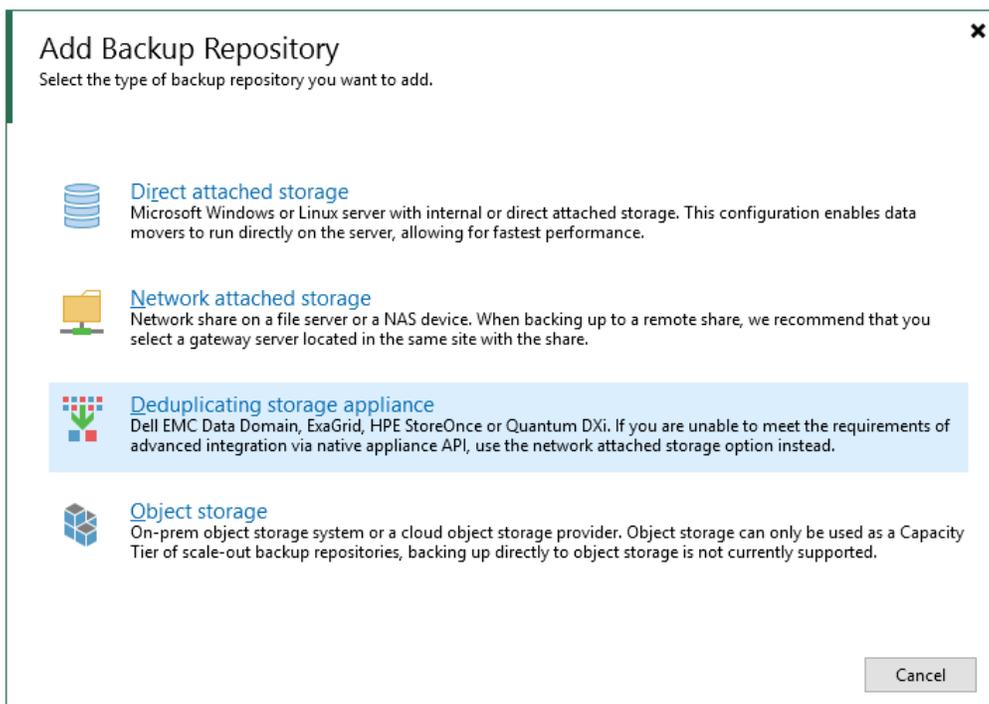


FIGURE 23. Specifying "Deduplicating storage appliance" as the type of backup repository in Veeam Backup & Replication



4. Select **HPE StoreOnce** as the **Deduplicating Storage Appliance**, illustrated in Figure 24.

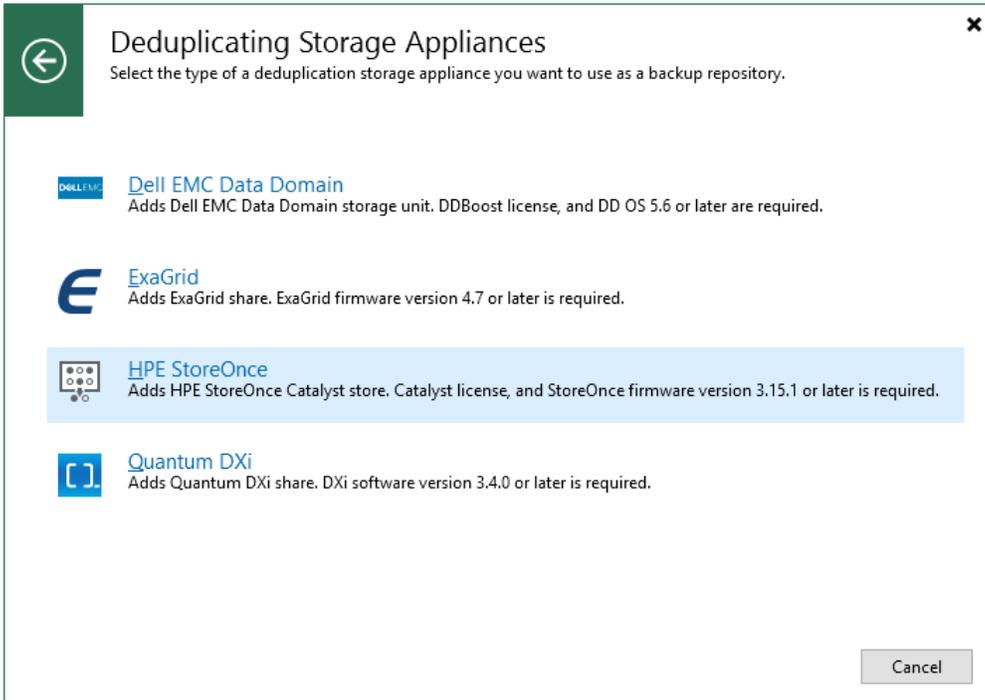


FIGURE 24. Specifying "HPE StoreOnce" as the deduplication storage appliance

5. The HPE StoreOnce backup repository can now be configured. As shown in Figure 25, a name and an optional description can be given to the repository.

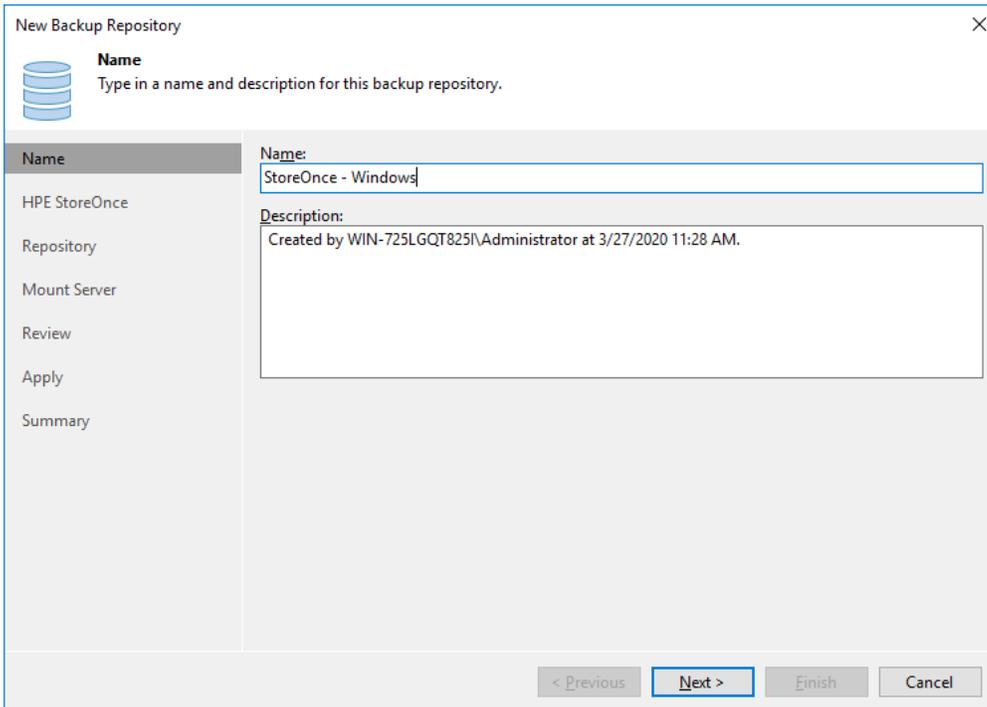


FIGURE 25. Defining the name and description of the backup repository



6. **HPE StoreOnce server name** –The HPE StoreOnce server name can be specified in two forms:

- IP or DNS name
 - Use this format when you want to access an HPE StoreOnce appliance via CoE.
 - See Figure 26 as an example of CoE.

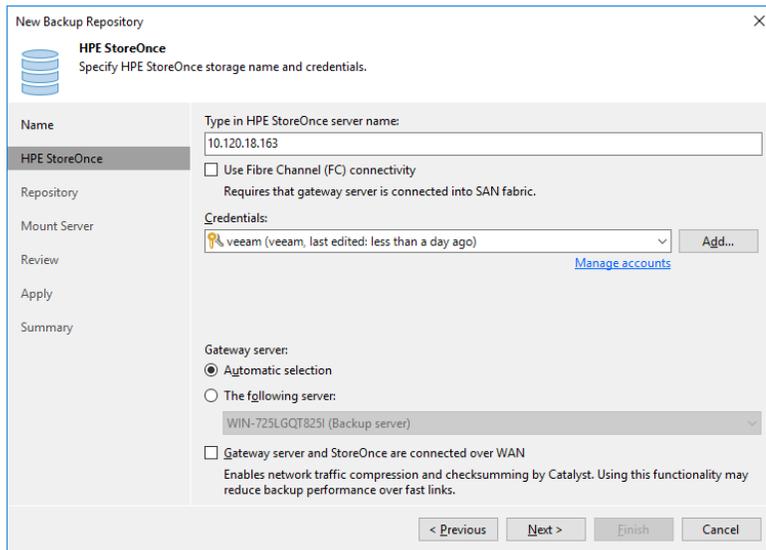


FIGURE 26. Specifying HPE StoreOnce server name for Catalyst over Ethernet (CoE)

- FC identifier
 - Use this format when you want to access an HPE StoreOnce appliance via CoFC
 - See Figure 27 as an example of CoFC
 - In this case, you have to select the Veeam Gateway server that you want to use as the gateway to the HPE StoreOnce appliance. For more configuration details, see the [Deploying a Veeam Gateway server](#) section

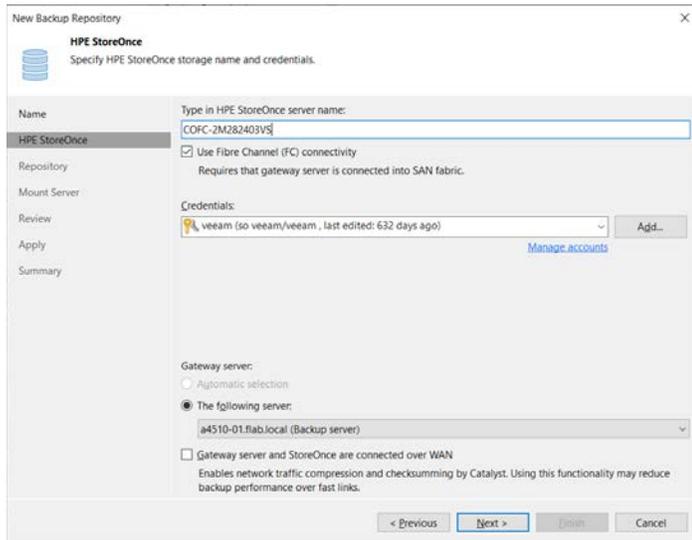


FIGURE 27. Specifying HPE StoreOnce server name for Catalyst over FC (CoFC)



As shown in Figure 26 and Figure 27, there are other important configuration fields:

Credentials – Credentials to access HPE StoreOnce are specified here. If **Client Access** has been enabled (see the [Configuring access permissions to an HPE StoreOnce Catalyst store](#) section), those credentials will be used here. Otherwise, any user or admin credentials can be used. New credentials can be added by clicking **Add...** and entering the appropriate credentials that will be stored in the Veeam credentials database for future use. In this instance, you should use the same credentials that were created in Figure 14.

- **Gateway server** – This is a Windows server running the Veeam Gateway server service. The gateway server reads and writes data to an HPE StoreOnce appliance. For more information on assigning a gateway server, see [Deploying a Veeam Gateway server](#). For CoFC, the gateway server must be specified. In contrast, for CoE, you can choose **Automatic selection** or specify **The following server**, displayed in Figure 27.

NOTE

Hewlett Packard Enterprise recommends configuring both proxy and gateway services to run on the same server whenever possible. This eliminates an extra hop in LAN for backup traffic and reduces backup proxy CPU utilization. To ensure that the gateway and proxy services are on the same server, you can specify the same server name on **The following server**: as shown in Figure 27, and on the **Backup proxy**: in the **Backup Job** configuration wizard. See the [Creating a Veeam backup job](#) section for more information on the **Backup Job** configuration wizard.

- **Gateway server and StoreOnce are connected over WAN** – Catalyst has been designed to work over LAN as well as WAN connections. When the communication between the Veeam Gateway server and HPE StoreOnce appliance is across a WAN link, a best practice is to enable the **Gateway server and StoreOnce connected over WAN** checkbox. This optimization option enables Catalyst native compression and checksum algorithms. Compression further reduces the WAN bandwidth utilization and checksum makes the communication more resilient to communication errors that are more frequent on WAN links. To avoid this additional workload, do not activate this functionality. When the **Gateway server and StoreOnce connected over WAN** checkbox is disabled, the compression will be executed by HPE StoreOnce before writing data to storage. This checkbox does not control the deduplication, which can happen at the source (a best practice) or the destination, as described in [HPE StoreOnce Catalyst Advanced Settings](#).

7. The **Repository** settings in the **New Backup Repository** wizard are illustrated in Figure 28.

The screenshot shows the 'New Backup Repository' wizard window. The 'Repository' step is selected in the left sidebar. The main area displays the following settings:

- Name:** HPE StoreOnce
- Location:** Catalyst store: veeam_catalyst_store (dropdown menu)
- Capacity:** (with a 'Populate' button)
- Free space:** (with a 'Populate' button)
- 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

At the bottom, there are navigation buttons: '< Previous', 'Next >', 'Finish', and 'Cancel'. An 'Advanced...' button is also present at the bottom right.

FIGURE 28. “New Backup Repository” wizard specifying Catalyst store and load control settings



Aspects of the settings are:

- **Catalyst Store** – Select the Catalyst store to associate with this Veeam repository. To enable deduplication across different backup repositories, it is possible to create multiple backup repositories on the same Catalyst store. This configuration is fully supported starting from Veeam V. 9.5-U4. For details on how to create multiple Veeam backup repositories, see the [Creating multiple Veeam backup repositories inside the same HPE StoreOnce Catalyst](#) section.
- **Load Control** – It is possible to limit the system resources that Veeam uses for reading and writing to a backup repository. Veeam Backup & Replication has a built-in scheduler to ensure that the active data sessions do not exceed the HPE StoreOnce appliance's maximum number of data sessions, so this setting is not mandatory and it is generally left disabled to maximize performance. You can enable the load control when you have multiple jobs running concurrently and you want to fine-tune the job resources among the different jobs.

Figure 29 illustrates the **Storage Compatibility Settings** in the **New Backup Repository** wizard.

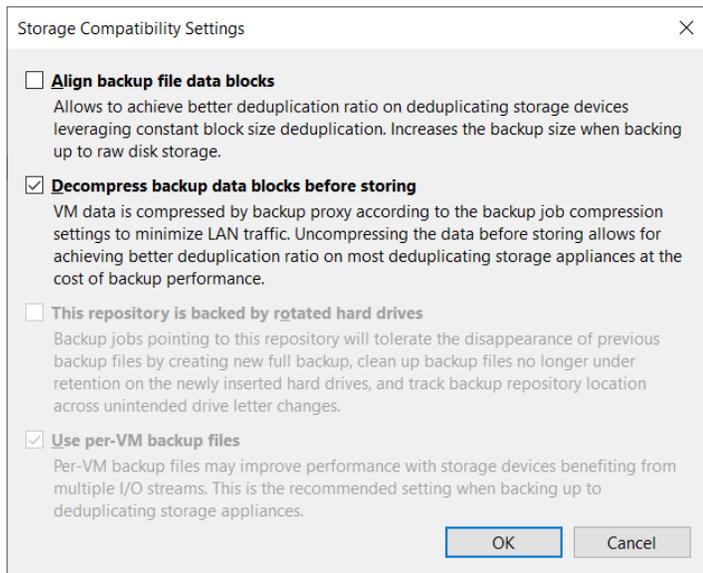


FIGURE 29. "Storage Compatibility Settings in the "New Backup Repository" wizard

NOTE

Do not change the default setting as shown in Figure 29. Advanced settings should not be adjusted for HPE StoreOnce-based backup repositories.

Storage compatibility settings can be configured as follows:

- **Align backup file data blocks** – This should be disabled for better HPE StoreOnce deduplication.
- **Decompress backup data blocks before storing** – This should be enabled for better Catalyst deduplication. Hewlett Packard Enterprise recommends keeping compression set to "Optimal" on the backup job and **Decompress backup data blocks before storing** on the backup repository for better performance.



8. The **Mount Server** settings in the **New Backup Repository** wizard are illustrated in Figure 30.

FIGURE 30. Configuring the mount server for the Veeam backup repository

Mount Server settings can be configured as follows:

- **Mount server** – In the case of HPE StoreOnce, a mount server is used to mount backup files for Instant VM Recovery, SureBackup®, and On-Demand Sandbox™. For granular restores, the backup server will automatically play this role. In most situations, it is a good practice to select the same server for the mount server and the gateway server to avoid an extra hop in LAN.
- **vPower NFS** – When a Veeam Instant VM Recovery is performed, backups are mounted to the hypervisor and powered on in a read-only state. All new data will be written to a temporary cache stored at the path specified in Figure 30.

NOTE

See the HPE StoreOnce Gen3 and Gen4 support matrix for release compatibilities with IVMR: [HPE StoreOnce support matrix \(SPOCK\)](#).



9. The **Summary** page is displayed, as shown in Figure 31.

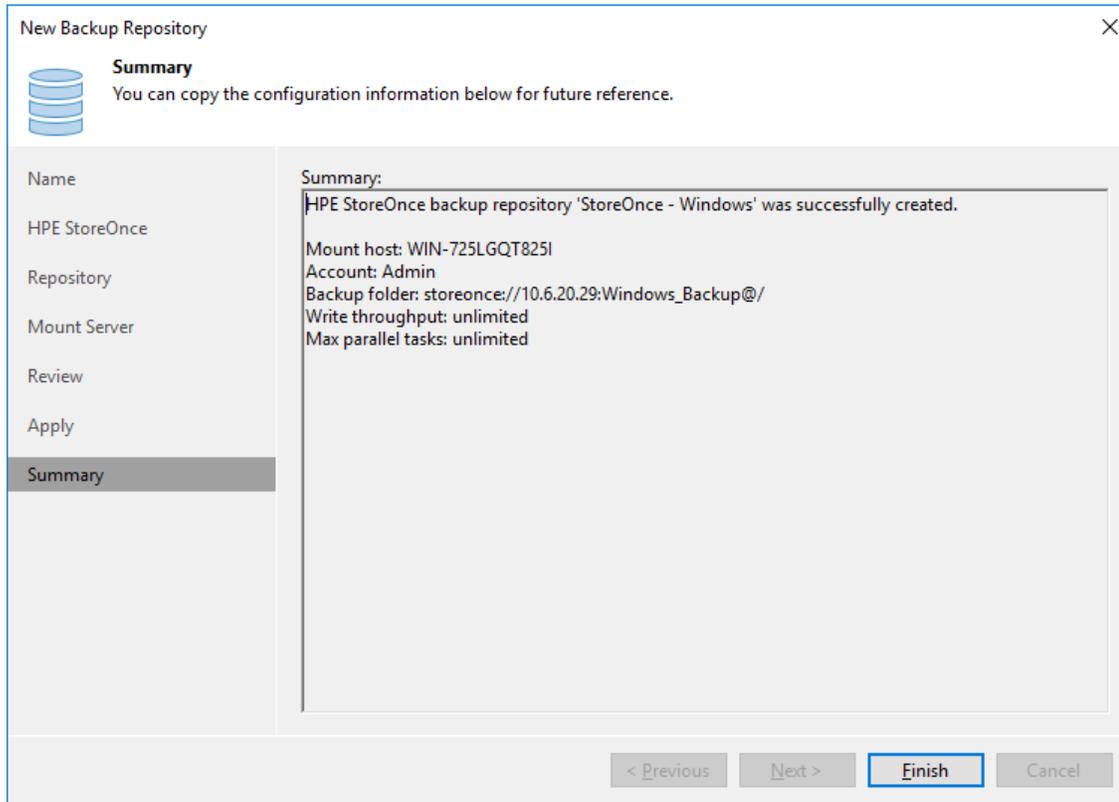


FIGURE 31. The “Summary” page in the “New Backup Repository” wizard

10. Review all settings and click **Finish**.

Creating multiple Veeam backup repositories inside the same HPE StoreOnce Catalyst store

There are situations where it is useful to create multiple Veeam backup repositories inside the same Catalyst store. The common reasons for creating multiple Veeam backup repositories in the same Catalyst store are:

- Better deduplication – Each Catalyst store is an independent deduplication domain. To enable cross-deduplication among multiple backup repositories, it is possible to create them inside the same Catalyst store.
- Workload distribution across multiple gateway servers – Each backup repository can be connected to a specific gateway server and cannot be accessed in write mode by more than one gateway server at a time. For better workload distribution, it is useful to have multiple backup repositories to associate them with different gateway servers.
- Avoiding extra hops in LAN between Proxy and gateway server – It is possible to specify an independent backup repository for each job and associate the same server for the proxy and gateway services. This configuration requires multiple backup repositories and it can be useful to have them in the same Catalyst store for better deduplication.
- Catalyst Copy granularity – As described in the [Veeam-managed HPE StoreOnce Catalyst Copy job](#) section, the Veeam Catalyst Copy job copies the contents of an entire backup repository to other HPE StoreOnce appliances. When multiple jobs write to the same backup repository, the Catalyst Copy job will replicate the backup data of all backup jobs. There are situations where it is necessary to tailor the replication parameters to the systems protected by specific jobs. This configuration requires multiple backup repositories, potentially one per job, and it can be useful to have them in the same Catalyst store to get better deduplication.
- Migration – Veeam provides an easy methodology for migrating entire backup repositories to new storage platforms (see the [Migrating Veeam Backup repositories to/from an HPE StoreOnce Catalyst store](#) section for details). If a storage platform becomes full, and you want to



migrate a subset to new storage, then a solution design based on multiple backup repositories offers more flexibility than a solution based on a single large backup repository.

IMPORTANT

Veeam recommends not to use the Veeam GUI for creating multiple backup repositories on the same Catalyst store. The configuration made using the GUI can create instability and should be avoided. Veeam recommends using the following supported PowerShell commands for creating multiple BRs instead of using the GUI. Once created, the backup repository can be managed by the GUI with the exception of the **Repository** tab. For more details, see veeam.com/kb2987.

After it is created by the CLI process, a backup repository can be conveniently managed by the GUI as with any other backup repository. The only parameters that cannot be changed via the GUI are the “Load Control.” The following examples demonstrate how to properly set up Veeam Backup Repositories via the CLI.

To open the Veeam PowerShell screen, click the icon on the top left Veeam GUI corner, select **Console**, and then click **PowerShell**, as shown in Figure 32.

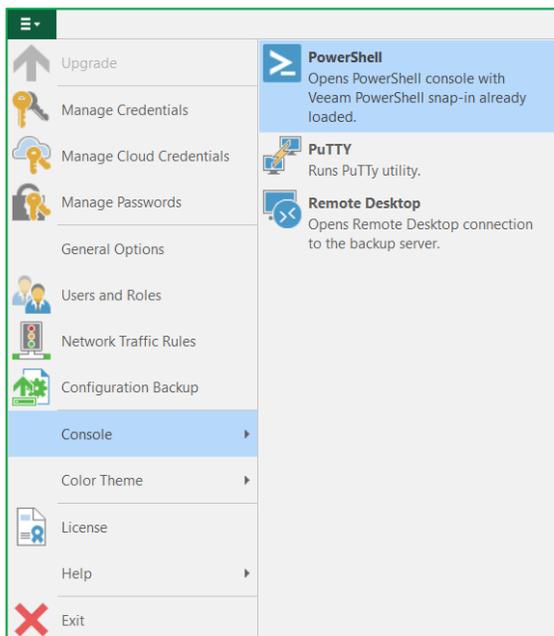


FIGURE 32. How to open Veeam PowerShell command windows

Example 1 – Create the additional repository “BR-1” in the Catalyst store “MultiBR:”

```
Add-VBRBackupRepository -Name BR-1 -Folder storeonce://vsa04:MultiBR@/Sub-1-Type HPStoreOnceIntegration
-StoreOnceServerName vsa04.flab.local -UserName veeam -Password veeam
```

Example 2 – Remove the limits for “maximum concurrent tasks” and for “read and write data rate.” This change is only possible via the CLI:

```
$repository = Get-VBRBackupRepository -Name "BR-1"
set-VBRBackupRepository -Repository $repository -DataRateLimit 12
-LimitDataRate:$false -MaxConcurrentJobs 12 -LimitConcurrentJobs:$false
```



Figure 33 illustrates the PowerShell examples:

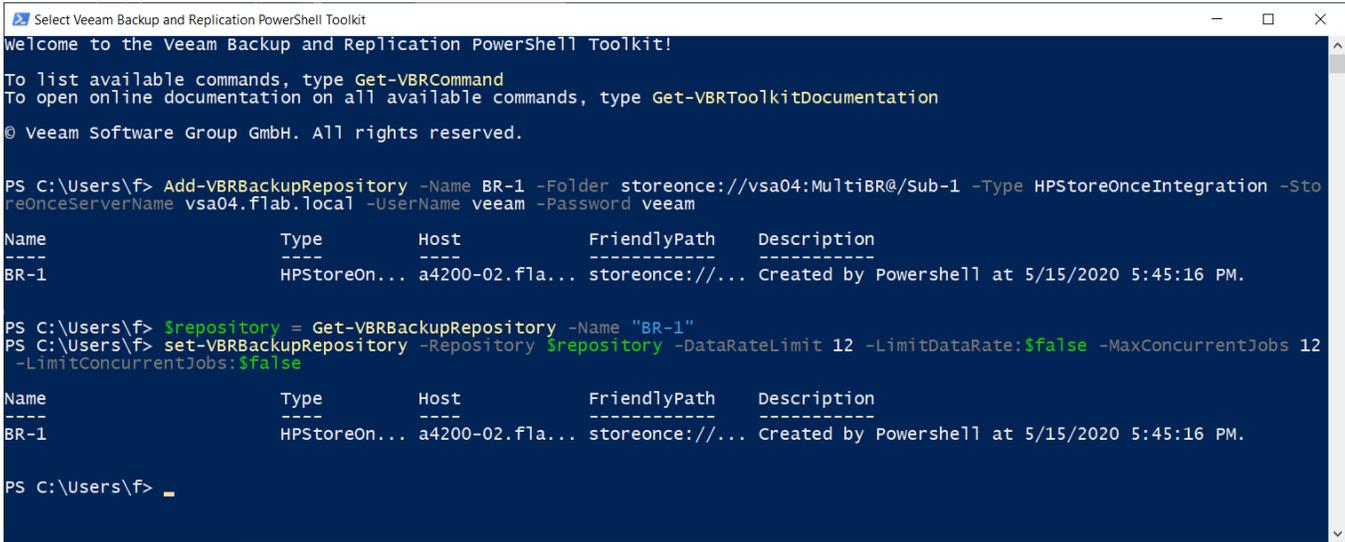


FIGURE 33. PowerShell CLI command examples

Creating a Veeam backup job for VMs

A Veeam backup job acts as a policy for protecting one or more VMs. A backup job holds specific settings for compression, deduplication, block transmission size, application quiescence, and execution scheduling. A backup job also defines the backup mode, such as forward incremental backups with periodic full backups. There are many types of Veeam backup jobs that support HPE StoreOnce Catalyst. See helpcenter.veeam.com/docs/backup/vsphere/deduplicating_appliance_storeonce.html for backup deployment options.

This section provides the recommended settings when creating a backup job writing to an HPE StoreOnce Catalyst-based backup repository. To add VMs to a backup job, perform these steps:

1. In the Veeam Backup & Replication console, click **HOME** and then **Jobs** to see the **Backup Job** selection in the ribbon, as demonstrated in Figure 34.

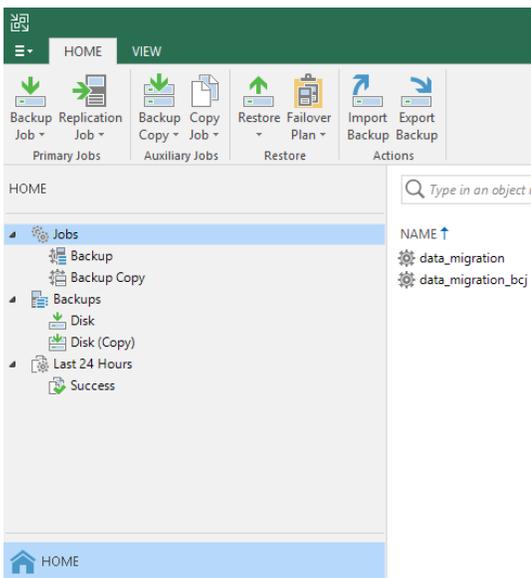


FIGURE 34. Navigating to a Veeam Backup Job



- Using the ribbon menu at the top of the user interface, click **Backup job**, then select the Backup job type. Figure 35 is an example of the **Virtual machine...** backup type.

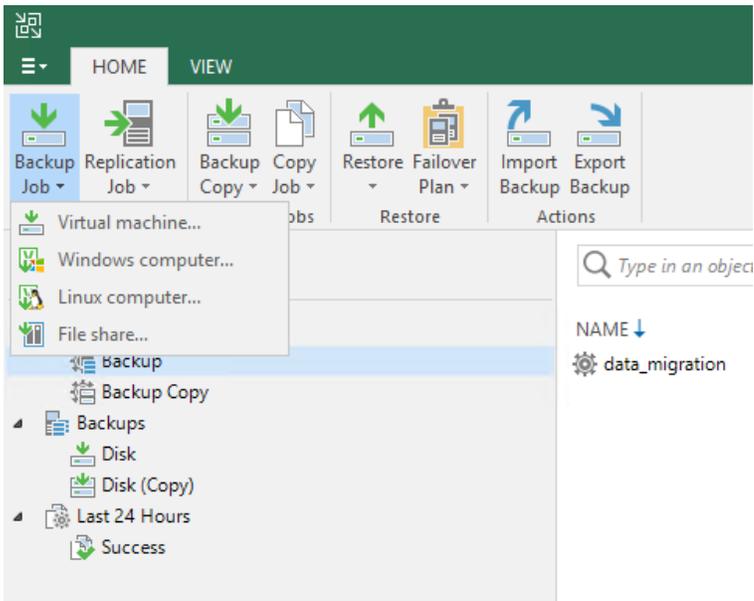


FIGURE 35. Selecting the “Virtual machine” backup type from the Veeam GUI

- As shown in Figure 36, give the job a meaningful **Name** and optional **Description**, and then click **Next**.

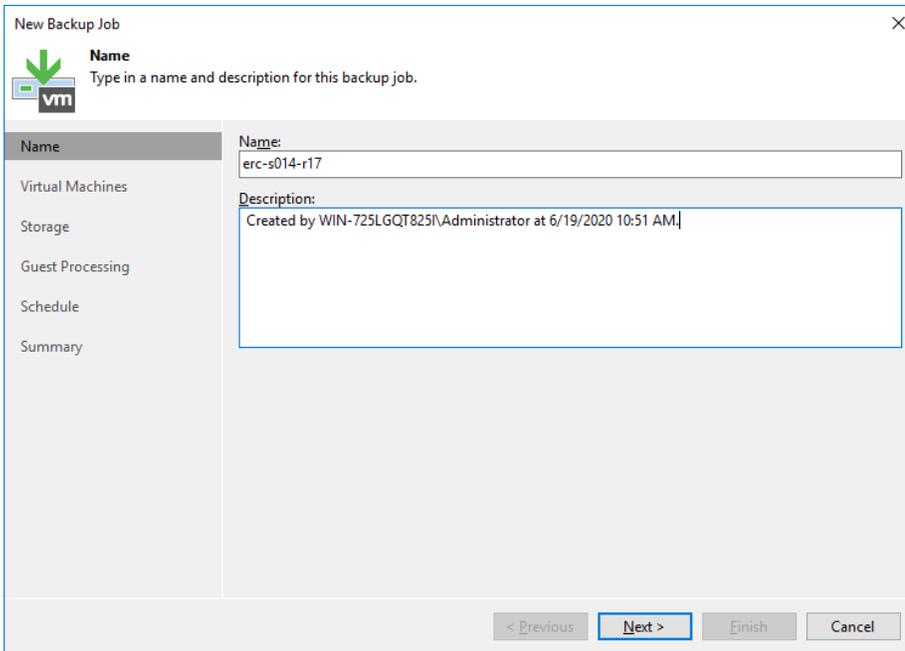


FIGURE 36. Provide a name and description to the Veeam backup job



4. On the **Virtual Machines** wizard section, click **Add** and select the VMs to add to this backup job, as demonstrated in Figure 37.

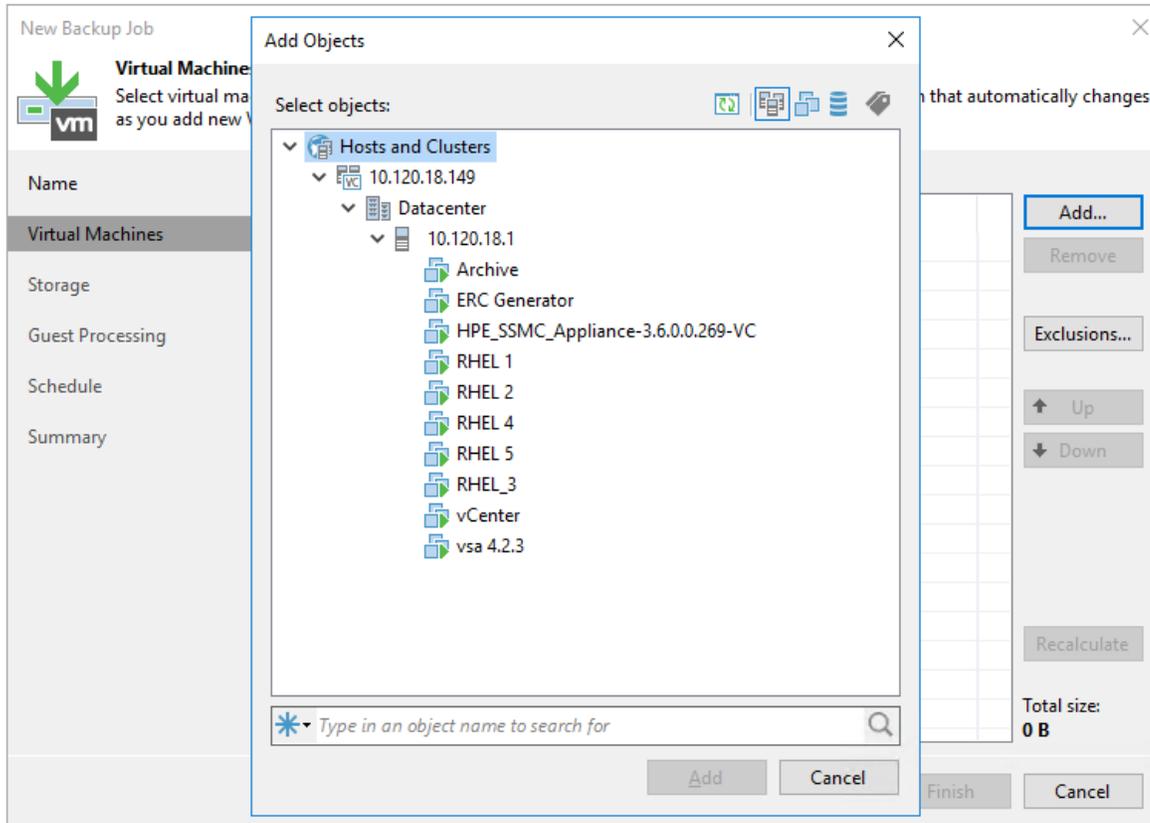


FIGURE 37. Adding VMs to a backup job in the Veeam GUI

NOTE

If a hierarchy object such as “Resource Pool” or “Cluster” is chosen, Veeam will back up all VMs under this object. Also, new VMs added under the object will be backed up on subsequent backup job runs, or removed from the job if they are deleted from the object.

TIP

If the infrastructure is designed for storage snapshot integration, it is generally useful to use the “Volume” view and select entire volumes.



5. The **Storage** wizard section displays the configuration elements shown in Figure 38.

FIGURE 38. Specifying proxy, repository, and retention parameters in the “New Backup Job” wizard

After adding VMs to the backup job, the storage settings shown in Figure 38 can now be specified. Storage settings can be summarized as follows:

- **Backup proxy** – Veeam’s intelligent load balancing between multiple backup proxies can be leveraged by leaving this as **Automatic**. Whenever possible, it is recommended that the backup proxy server be the same as the gateway server selected in the **Backup Repository** configuration wizard. This will reduce the number of hops in the data path and increases performance.
- **Backup repository** – Choose the HPE StoreOnce Catalyst-based backup repository created in previous steps.
- **Retention policy** – The retention policy can be specified in **restore points** or **days**. The frequency of the Veeam backup job is an important element to consider in the design.
- **Keep certain full backups for longer for archival purposes** – The long-term or Grandfather-Father-Son (GFS) retention policy allows you to store selected backup files for extended periods of time, including weeks, months, and years. Veeam Backup & Replication does not need to create any new backup files—it uses backup files already created and marked them with specific GFS flags. For Catalyst-based backup repositories, GFS retention policies are generally more efficient than using a Backup Copy job. For more information, see helpcenter.veeam.com/docs/backup/vsphere/gfs_retention_policy.html.

NOTE

HPE StoreOnce deduplication highly reduces the storage requirement for additional restore points and longer daily retention. For example, if the retention is doubled from 2 to 4 weeks, the storage utilization on HPE StoreOnce does not double but grows only by ~19%. If retention is extended from 2 to 8 weeks, the storage utilization on HPE StoreOnce grows only by ~57%. These are average values and in most environments the differences are minimal, but the deduplication effect can vary, and it is highly influenced by the protected data type.



For better deduplication, whenever possible, Hewlett Packard Enterprise recommends to write backup data from similar VMs to the same Catalyst store. This is explained in more detail in the [Creating multiple Veeam backup repositories inside the same HPE StoreOnce Catalyst store](#) section. This can be achieved in several ways:

- Group VMs with similar data in the same job
- Use the same backup repository for multiple jobs containing VMs with similar data
- Create multiple backup repositories writing to the same Catalyst store, and use those repositories for multiple jobs containing VMs with similar data

As shown in Figure 38, click **Advanced** to open the backup job **Advanced Settings** (Figure 39).

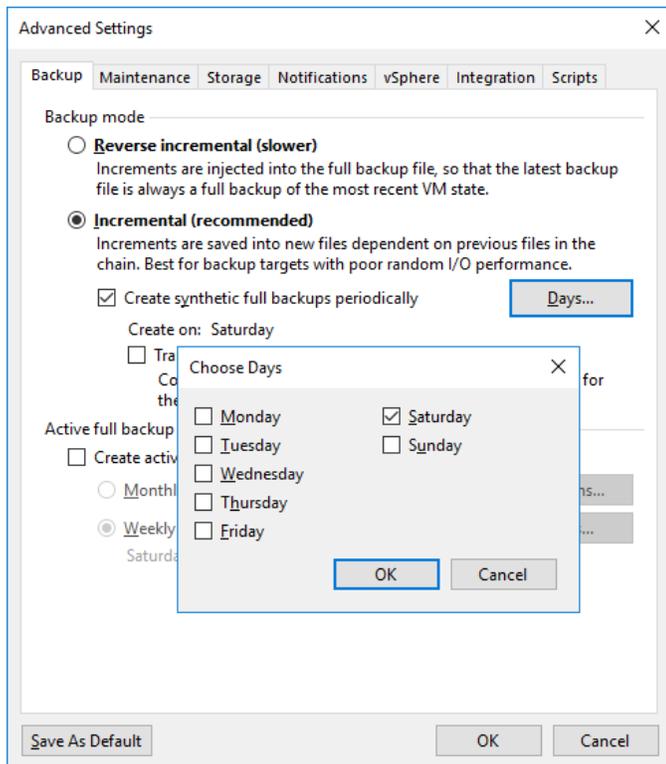


FIGURE 39. Specifying backup mode and synthetic full backups in the “Advanced Settings” of the “New Backup Job” wizard

The settings in the **Backup** tab shown in Figure 39 should be configured as follows:

- **Backup mode** – Incremental
- **Synthetic Full backups** – It is necessary to schedule periodic full backups to minimize the backup chain length. Synthetic full backups provide new “checkpoints” for backup chains by rolling up the latest existing full backup with the subsequent incremental backups to create a new full backup. Veeam leverages the Catalyst API to offload (virtualize) this process from the Veeam servers to the HPE StoreOnce appliance, reducing the load on the infrastructure and the time required to create the synthetic full backup.

This process takes the following steps:

- Veeam executes an incremental backup. A full backup is no longer required, and this reduces about 90% of the workload on production storage and servers.
- Veeam sends the “instructions” to the HPE StoreOnce appliance for merging the previous full with all the subsequent incremental backups to finally generate the new synthetic full backup.



- HPE StoreOnce executes this task at high speed because it does not need to move the actual data, but only the pointers to the stored deduplicated data segments. To enable synthetic full backups, select the checkbox **Create synthetic full backup periodically**, and choose when this should occur. A synthetic full backup generally completes in 50% or less of the time required for a traditional active full backup.

NOTE

It is required that either a periodic synthetic full backup or active full backup is selected when backing up to an HPE StoreOnce Catalyst store.

The settings in the **Maintenance** tab shown in Figure 40 should be configured as follows:

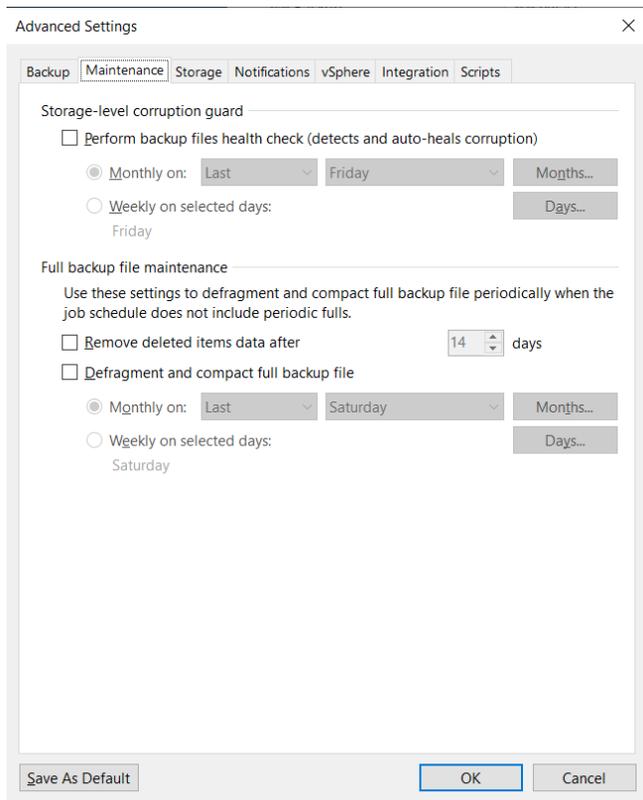


FIGURE 40. Maintenance settings are not required for Catalyst-based repositories

- **Storage-level corruption guard** settings can be left unchecked. HPE StoreOnce has internal data integrity checking and self-healing mechanisms.
- **Full backup file maintenance** settings can be left unchecked because the backup policy already includes periodic full backup.



For better performance and deduplication, Hewlett Packard Enterprise recommends the following **Storage** advanced settings, as illustrated in Figure 41.

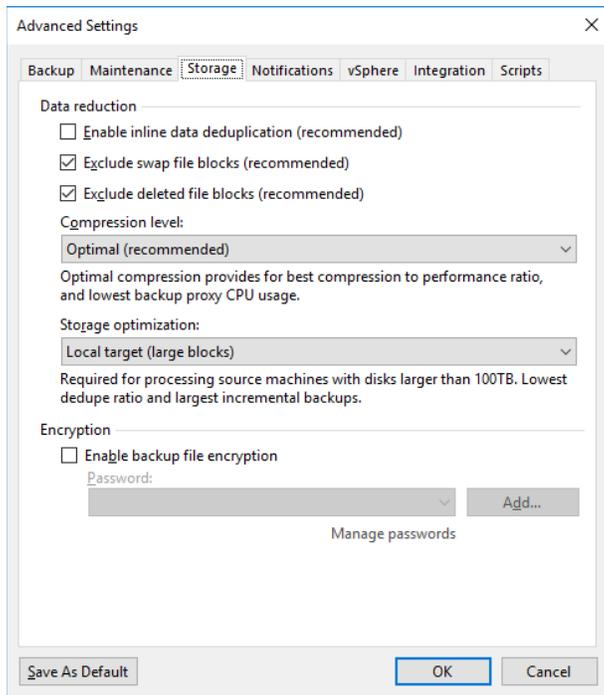


FIGURE 41. Specifying deduplication, compression, optimization, and encryption options

Data deduplication – Disabled

Compression level – Optimal

Storage Optimization – Local target (large blocks)

Encryption – Disabled

NOTE

If data-at-rest encryption is required, use HPE StoreOnce embedded encryption rather than Veeam encryption because Veeam encryption will reduce HPE StoreOnce data deduplication. For more information, see the [Creating a Catalyst store for Veeam backups](#) section.



The **Integration** tab backup job advanced settings (Figure 42) can be used to configure backups from HPE Primary Storage snapshots rather than from the VMware host itself. To enable backups from HPE Primary Storage snapshots, additional configuration is required. More information can be found in the HPE Reference Configuration at hpe.com/v2/getdocument.aspx?docname=a00061747enw.

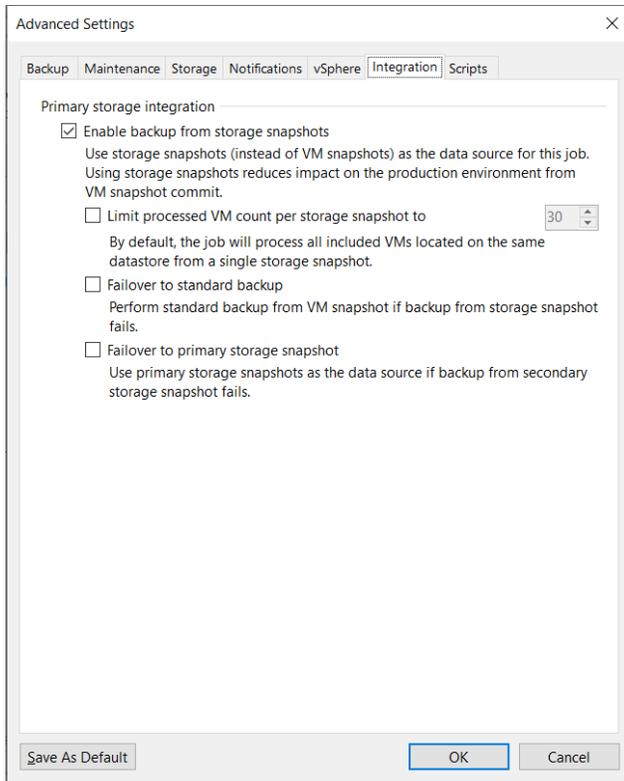


FIGURE 42. Configuring backups from storage snapshots

NOTE

Veeam attempts to concurrently process all of the VMs in a storage snapshot. If necessary, the number can be limited. See Figure 42 for **Limit processed VM count per storage snapshot to**. In general, it is better to process all of the VMs concurrently to avoid taking multiple HW snapshots. With HW snapshot integration, the overhead of processing multiple VMs concurrently is limited because each VM snapshot has a short life and, consequently, a lower impact on production infrastructure.



On the **Guest Processing** tab, shown in Figure 43, it is possible to enable transaction consistency for VM backups and indexing of VM guest file systems for quick searching through the optional Veeam Backup Enterprise Manager. To enable these capabilities, simply check the appropriate box and supply administrator credentials. Additional details on the configuration of these features can be found at helpcenter.veeam.com.

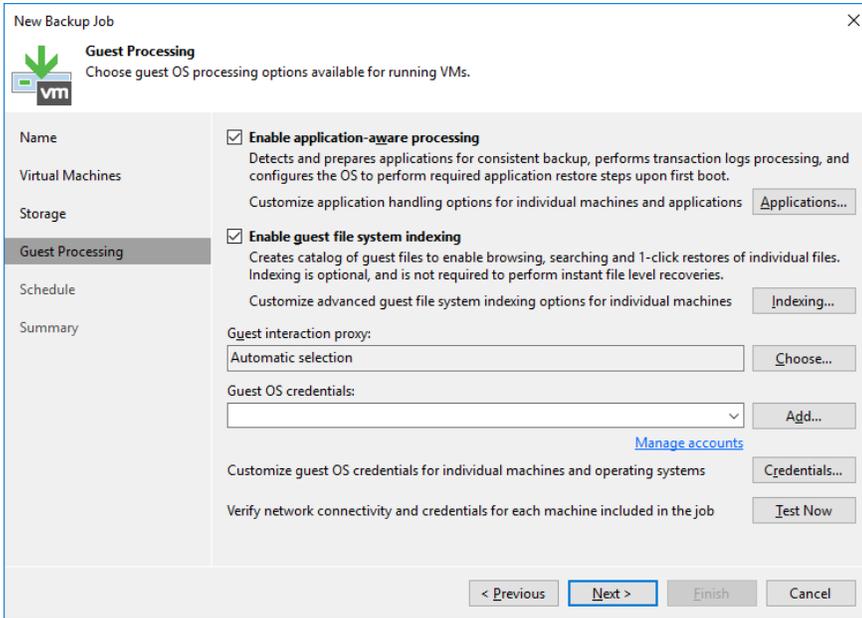


FIGURE 43. Processing and file indexing options in the “New Backup Job” wizard

On the **Schedule** tab, shown in Figure 44, it is possible to control the backup job scheduling. Usually, backup jobs are run daily during the night. If more frequent recovery points are required, a more frequent synthetic full backup schedule will need to be used. For instance, if three restore points per day are required, and a backup chain of under 14 restore points needs to be maintained, then there will need to be a synthetic full backup twice per week.

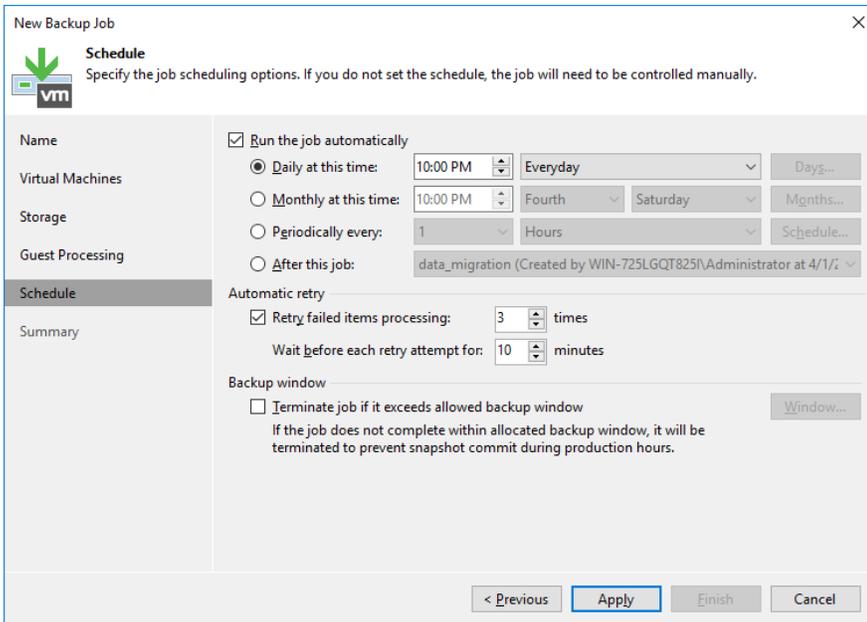


FIGURE 44. Set the job schedule to determine the RPO for the backup job



Creating a Veeam backup job for Windows and Linux computers

Veeam Agents can also be utilized to backup data from Windows and Linux® computers to an HPE StoreOnce Catalyst store target, which can yield the same level of protection and deduplication of virtual machine data.

In the **Agent Backup Job** configuration wizard, navigate to the **Storage** section, and do the following:

1. Select a Catalyst-based backup repository, as shown in Figure 45. For more details on how to create a Catalyst-based backup repository, see the [Creating a Catalyst store for Veeam backups](#) section.

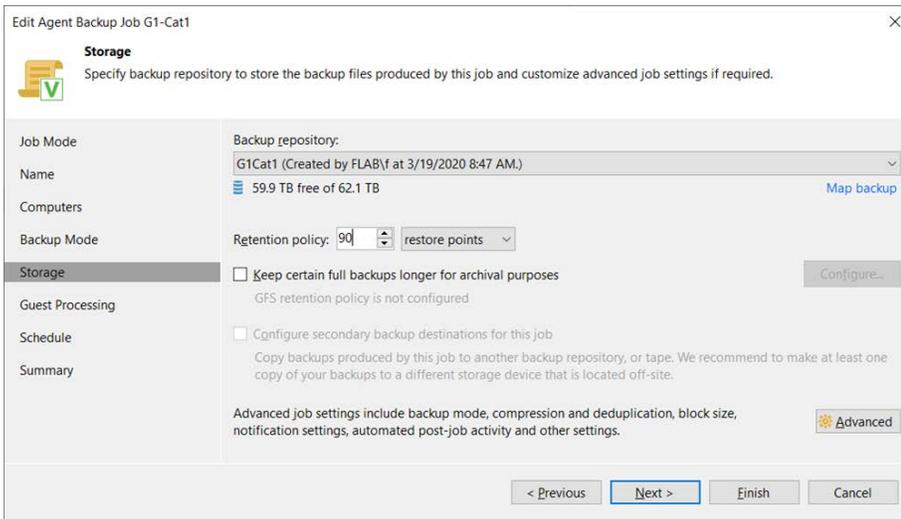


FIGURE 45. A Veeam Agent Backup Job can utilize an HPE StoreOnce Catalyst store

2. Open the **Advanced** section for more settings.
3. For the **Advanced** settings, the following **Backup** configuration is recommended, illustrated in Figure 46.

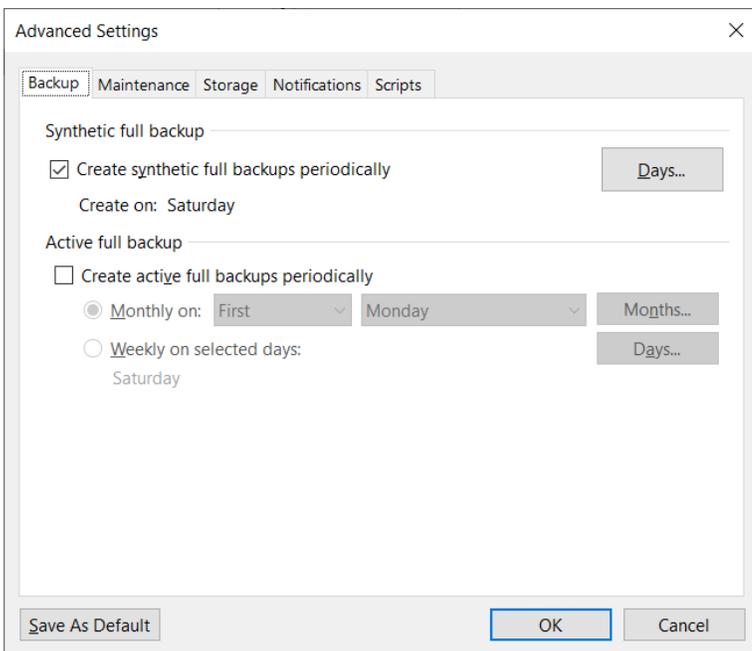


FIGURE 46. Advanced backup settings for a Veeam Agent Backup Job



- Enable **Create synthetic full backups periodically** and define how often they should occur. If you have multiple jobs, distribute the synthetic full on different weekdays to avoid workload peaks on the HPE StoreOnce appliance.
- Do not enable **Create active full backups periodically** because this is not needed for HPE StoreOnce and creates an unnecessary additional workload.

4. The settings in the **Maintenance** tab shown in Figure 47 should be configured as follows:

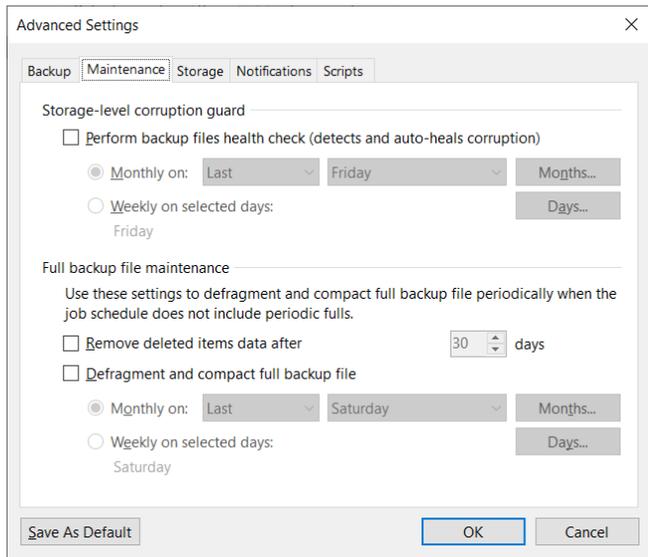


FIGURE 47. Advanced maintenance settings for a Veeam Agent Backup Job

- Do not enable **Storage-level corruption guard**. HPE StoreOnce has internal data integrity checking and a self-healing mechanism.
- Do not enable **Full backup file maintenance** because the backup policy already includes periodic full backups.

5. The settings in the **Storage** tab shown in Figure 48 should be configured as follows:

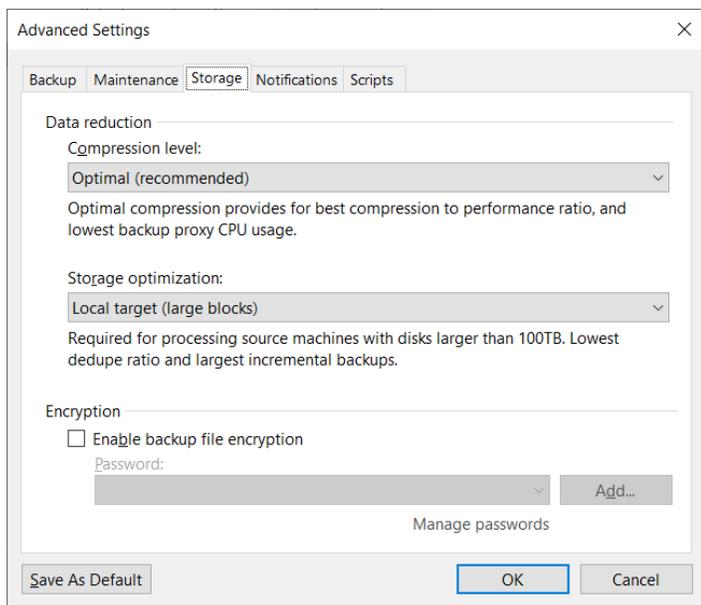


FIGURE 48. Advanced storage settings for a Veeam Agent Backup Job



- For **Compression level**, select **Optimal**. Remember to set un-compress before writing on the backup repository as explained in the [Creating an HPE StoreOnce Catalyst-based Veeam backup repository](#) section.
- For **Storage Optimization**, select **Local target (large blocks)** because it is faster and does not use more HPE StoreOnce capacity.
- For **Encryption**, do not select **Enable backup file encryption**. If you want encryption for data-at-rest, use HPE StoreOnce encryption as described in the [Creating a Catalyst store for Veeam backups](#) section.

Backup data from a Veeam Agent can be transferred to HPE StoreOnce in two ways. The first is the server-based backup, demonstrated in Figure 49, which supports both Windows and Linux servers, and requires a Veeam Gateway service to write to HPE StoreOnce.

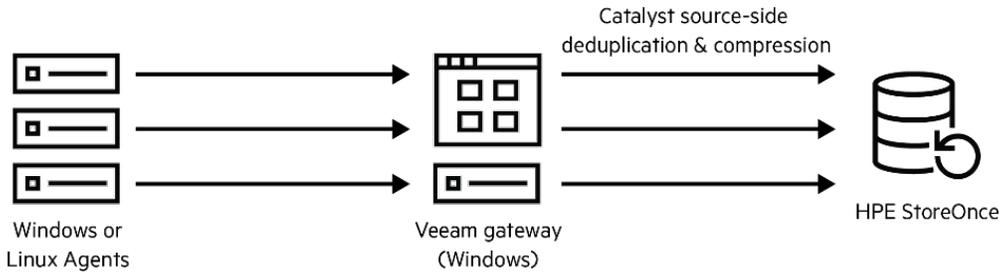


FIGURE 49. Veeam Agent server-based backup

The second is the direct-client backup, illustrated in Figure 50, which supports only Windows servers. This backup is generally faster because it removes the extra hop in LAN. This is the suggested configuration for large clients where performance is important.

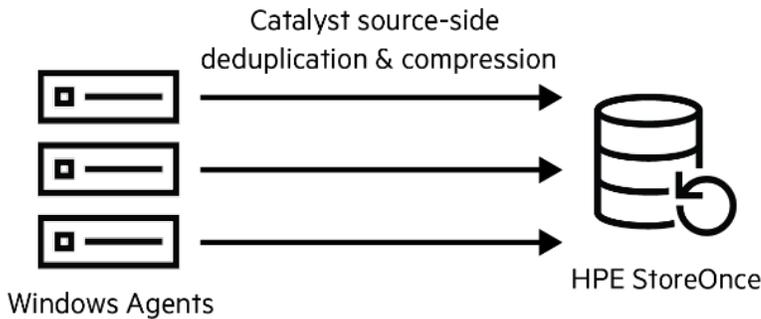


FIGURE 50. Veeam Agent direct-client backup

For configuring the Direct backup there are few simple additional configuration steps:



1. Add each Windows Client to the Veeam-managed server list, demonstrated in Figure 51. This allows Veeam to install the gateway service for the next step. There is no need to make it a proxy.

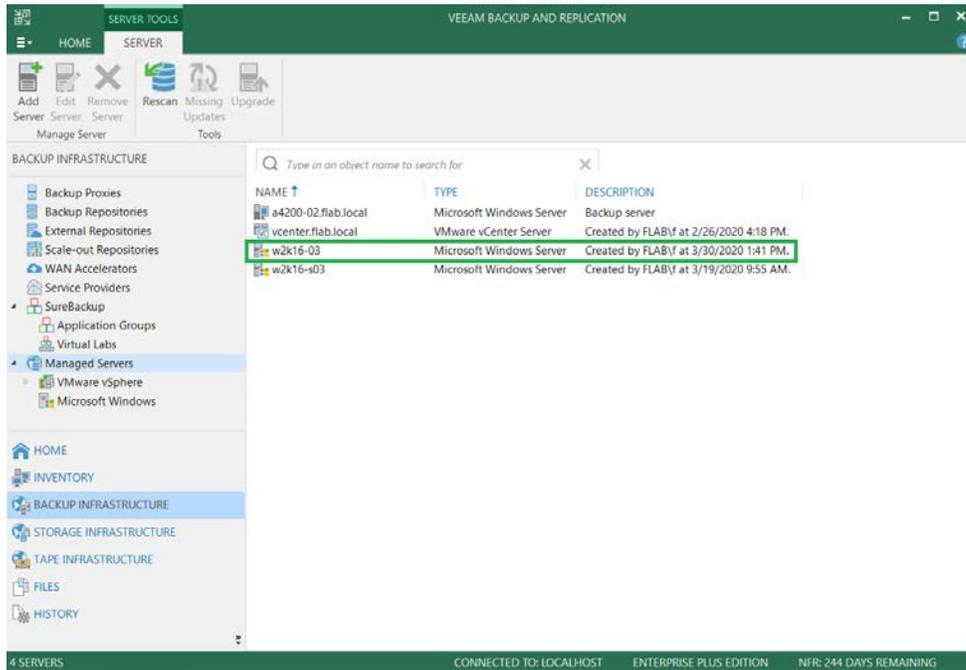


FIGURE 51. Adding a Windows client as a Veeam-managed server

2. In reference to Figure 52, create a new backup repository for each Windows client and set the client as the gateway server. If you need to create many BRs, it might be useful to have all of them writing to the same Catalyst store to increase deduplication. For a full description, see [Creating an HPE StoreOnce Catalyst-based Veeam backup repository.](#)

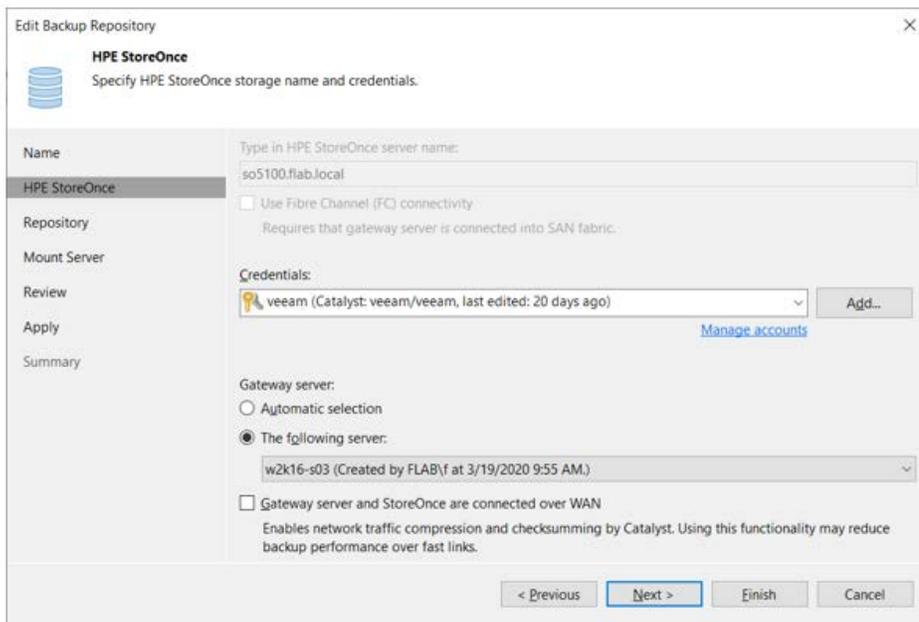


FIGURE 52. Adding a new backup repository using the Windows client as the Veeam Gateway server



3. Finally, create a Veeam Agent Backup Job for each client computer and assign the associated BR as its backup target.

With the above configuration, each Veeam Agent running on a computer client will write to a backup repository where the same computer acts as a gateway. Therefore, the same computer will write to a local or remote HPE StoreOnce appliance with all the benefits of direct communication with a high level of data reduction. After the first backup, the average data reduction for each new data transfer is about 17:1 for incremental and 72:1 for a full backup. These are average values and in most environments the differences are minimal, but the deduplication effect can vary, and it is highly influenced by the protected data type.

Copying Veeam backups to a remote site – Backup replication

Maintaining multiple copies of backups is critical to an organization's data availability strategy. The use of Backup Copy jobs plays a key part in complying with 3-2-1 backup best practices—three copies of data on two different media types, with at least one off-site copy. Veeam with HPE StoreOnce enables this by distributing copies of Veeam backups on geographically separated HPE StoreOnce appliances. The following three architectures available to execute backup copies are described:

1. Veeam-managed HPE StoreOnce Catalyst Copy job – Storage-based data replication from an HPE StoreOnce appliance to another HPE StoreOnce appliance. Granularity is controlled by Veeam Backup & Replication, but fully executed by a direct HPE StoreOnce-to-HPE StoreOnce communication. It does not require any server as a data-mover. This process offers maximum efficiency, speed, and bandwidth reduction.
2. Veeam Backup Copy Job with Catalyst write over WAN – Server-based data replication to a remote HPE Catalyst-based repository. This requires one single server as a data-mover and supports any source repository, such as a local file system or HPE StoreOnce appliance. It requires an HPE StoreOnce appliance at the remote destination. This design is often used for small branch offices that do not have a local HPE StoreOnce appliance and want to replicate data to a central HPE StoreOnce destination. The process offers excellent bandwidth reduction.
3. Traditional Veeam configuration – Server-based data replication from/to any type of remote repository. It requires two servers as data-movers. This design does not leverage any HPE StoreOnce bandwidth reduction and it is not often used when the destination is an HPE StoreOnce appliance because the designs described above offer more advantages.

Veeam-managed HPE StoreOnce Catalyst Copy job

With the release of Veeam Backup & Replication version 10, HPE StoreOnce Catalyst Copy operations can be conducted directly through Veeam. Veeam sends Catalyst Copy commands to an HPE StoreOnce appliance, and then HPE StoreOnce executes the commands. This configuration has a number of advantages:

- Data is never rehydrated because the design does not require a Veeam server to read the data that has to be transferred. This removes the need to deploy an additional server and greatly reduces the workload on the source HPE StoreOnce appliance because it does not need to rehydrate any data.
- Parallel processing – HPE StoreOnce native replication has a wide queue length optimized for the WAN latency. This optimization increases the replication throughput even when there is a single large file to be replicated.
- Best bandwidth reduction – data is replicated with the best possible deduplication and compression level. After the first backup, the average data reduction for each new data transfer is about 17:1 for incremental and 72:1 for a full backup. These are average values and in most environments the differences are minimal, but the deduplication effect can vary, and it is highly influenced by the protected data type.
- It does not require any additional data manipulation, such as rolling-full merge operations or periodic defragmentation. Data integrity checks are possible, but they are not required because HPE StoreOnce replication already implements internal consistency and corruption controls.
- Every replication operation is controlled by Veeam and executed by HPE StoreOnce. Veeam defines what data to replicate and when to replicate that data and the user can configure and monitor the process from the Veeam GUI.
- The remote repository content is shown on the Veeam GUI and it is directly accessible for any restore operation.



- It supports multi-site replication and multiple copies. It supports multiple parallel replications as well as cascade replication. Figure 53 shows the multiple replication models that can be used for Veeam-managed HPE StoreOnce Catalyst Copy jobs. This means that any source repository, including source repositories that were originally target repositories, can be propagated to a different target repository.

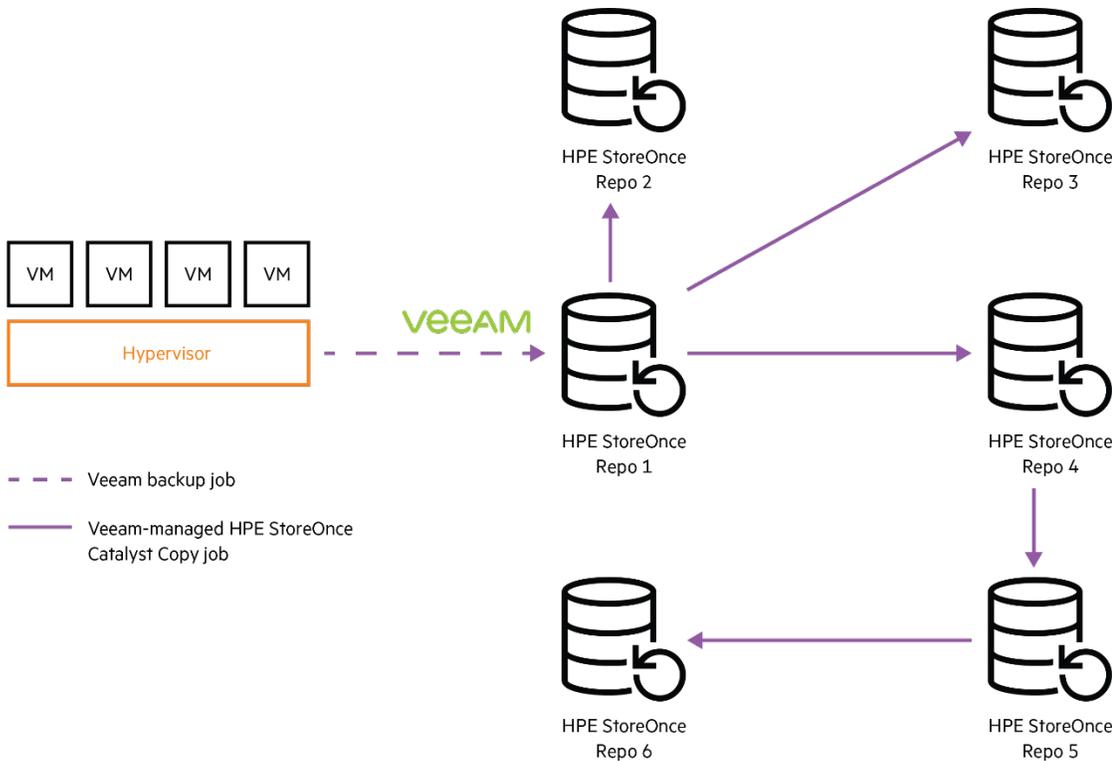


FIGURE 53. Replication model involving multiple parallel and cascading Veeam-managed HPE StoreOnce Catalyst Copy jobs

The following best practices are recommended when configuring a Veeam-managed HPE StoreOnce Catalyst Copy job:

- Either Catalyst Copy over Ethernet (CCoE) or Catalyst Copy over Fibre Channel (CCoFC) can be used.
- If CCoE is used, specify the HPE StoreOnce IP or FQDN, but not the short DN because it could lead to resolution problems.
- If encryption for data-in-transit is required, then, if possible, configure your WAN network devices for making the point-to-point encryption. If this encryption facility is not available, then it is possible to configure data-in-transit encryption on HPE StoreOnce (see the [Creating a Catalyst store for Veeam backups](#) section), but this increases the workload and could reduce the replication performance.

The replication protocol is defined by the target backup repository. If CCoE is desired, make sure to define the destination repository as CoE and not CoFC. Please note that it is supported to have the source repository accessed via CoFC, and the destination via CoE. If the destination needs to be accessed via CoFC for performing restore operations, but the replication must use CCoE, then it is necessary to implement the following workaround:

1. The production Veeam instance defines the target repository via CoE.
2. Install a new Veeam server to be used for restoring the destination replica via CoFC. The new Veeam server can be a VM.
3. On the new Veeam server, create a backup repository that points to the replicated copy via CoFC, which requires a server with FC connectivity. Both the production server and the new Veeam server can use the same proxy/gateway, so you may not need to install a new physical server just for restoring via CoFC.
4. Rescan the backup repository to import the restore point list to GUI.
5. The new Veeam server can be used for making a restore via CoFC from the Catalyst Copy replication HPE destination.



To initiate a Veeam-managed HPE StoreOnce Catalyst Copy job, select **Backup Copy**, followed by **HPE StoreOnce repository...**, shown in Figure 54.

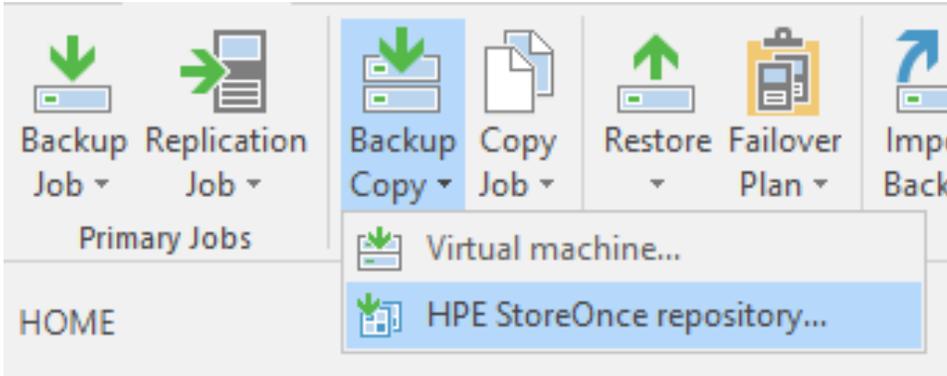


FIGURE 54. Initiating a Veeam-managed Catalyst Copy job

After naming the job, in the **Storage Mapping** section, click **Add...**. Designate the source Catalyst store and the target Catalyst store, illustrated in Figure 55.

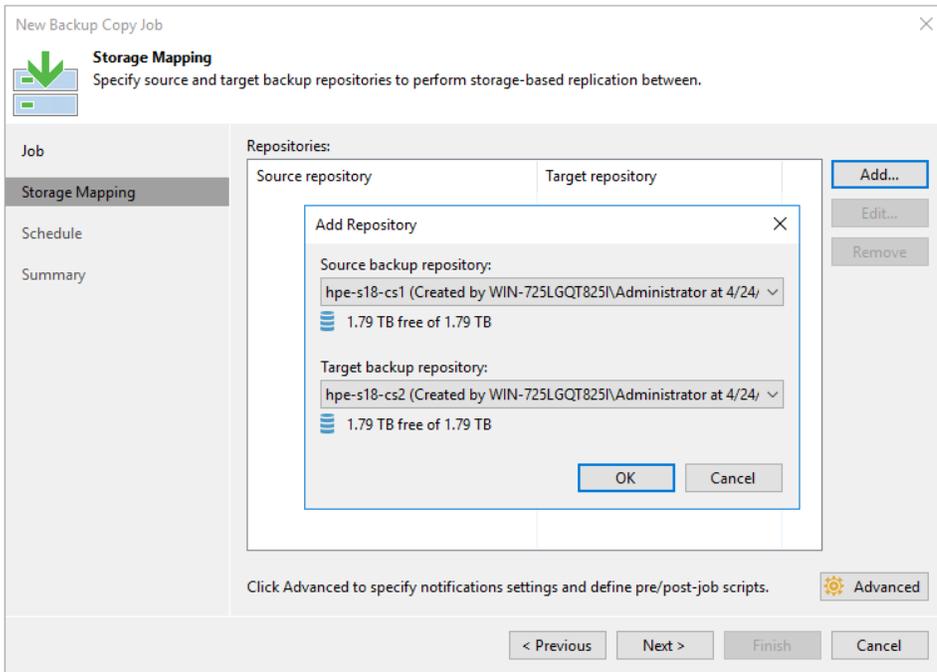


FIGURE 55. Selecting the source and target HPE StoreOnce Catalyst stores for a Veeam-managed HPE StoreOnce Catalyst copy job



For the advanced **Maintenance** settings shown in Figure 56, pay particular attention to the **Perform backup files health check** checkbox. When it is enabled, it can increase the workload of the target HPE StoreOnce system and the replication process can take longer to complete. HPE StoreOnce contains a sophisticated mechanism to ensure that the replicated file is the exact copy of the source file to reduce the risk of corruption in the presence of a network error.

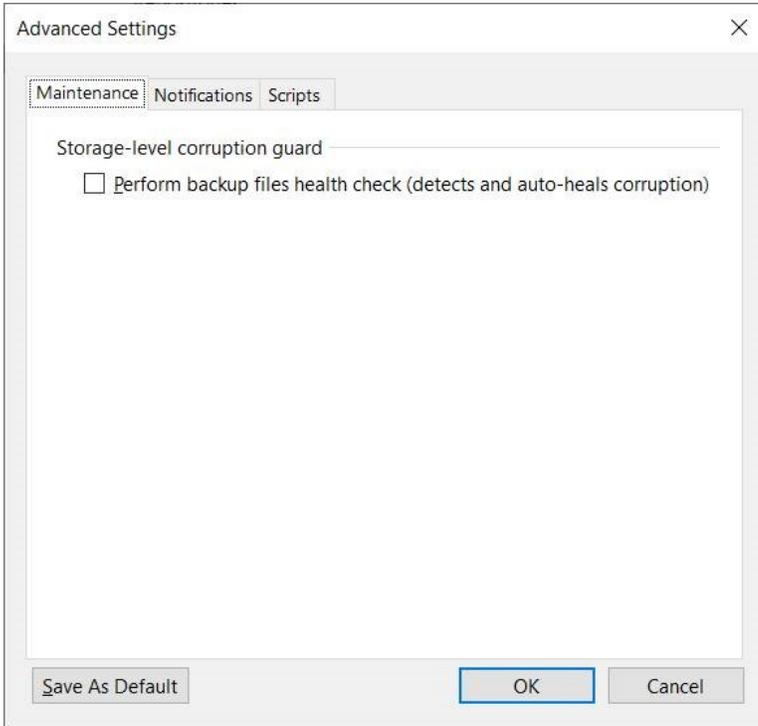


FIGURE 56. Veeam-managed Catalyst Copy job advanced "Maintenance" settings

Like other Backup Copy jobs, this job can be run on a schedule. This is configured in the **Schedule** section of the Backup Copy job, shown in Figure 57.

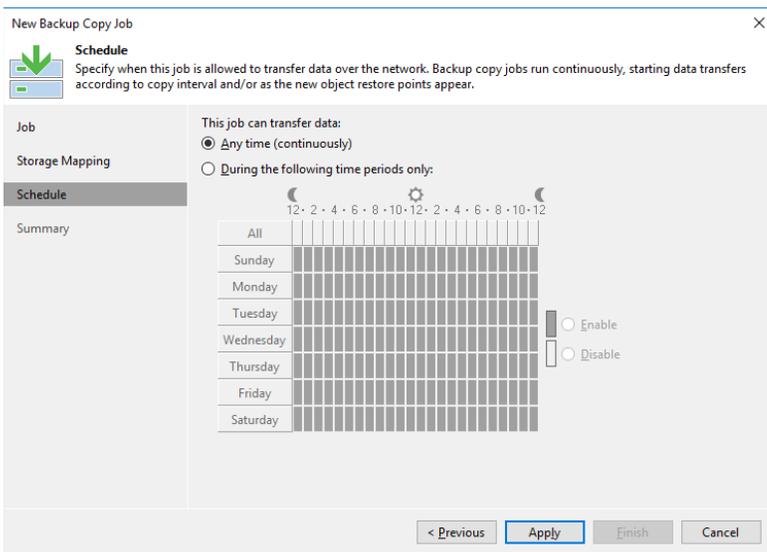


FIGURE 57. Veeam-managed Catalyst Copy jobs can be run on a schedule



Regarding parallel and cascade replication:

- To set up parallel replication from one backup repository to multiple destinations, it is necessary to set up an additional Backup Copy job for each replication.
- To set up cascade replication from *Source* to *Copy 1* and from *Copy 1* to *Copy 2*, there are two equally viable options:
 - Option 1: Add a line to the same Backup Copy job in the **Storage Mapping** section.
 - Option 2: Set up an additional Backup Copy job from *Copy 1* to *Copy 2*.
- There is no predefined limit on the cascade sequence length, the only warning is to avoid loops, such as *Source* to *Copy 1* and *Copy 1* to *Source*.

File replication behavior from a Catalyst Copy job (CCJ) can be described as follows:

- A CCJ makes a mirror of the source backup repository. If a file is deleted from the source, then the delete operation is replicated to the destination. In other words, it is not possible to define different retention on the source and destination repositories.
- A CCJ only replicates the backup files generated by a Backup Job for VMware, Hyper-V, Nutanix AHV, and Managed Agents.
- Files generated by a Backup Copy job are not processed by the CCJ.
- A CCJ replicates all the supported files inside the selected backup repository. If the intent is to replicate only the backup files from a specific job, then assign to the job an exclusive backup repository and create a CCJ that replicates that repository.
- To avoid creating too many Catalyst stores, it is possible to create multiple backup repositories inside the same Catalyst store. See [Creating multiple Veeam backup repositories inside the same HPE StoreOnce Catalyst store](#) for more details.
- When the Veeam Backup job is writing to a replicated backup repository that has GFS retention, then the GFS retention files are replicated too.

Delete behavior from a CCJ can be described as follows:

- Expired restore points on the source are deleted from replicas after a grace period (default=7 days). To change the grace-period length set the following registry key:

```
[HKEY_LOCAL_MACHINE\SOFTWARE\Veeam\Veeam Backup and Replication]
"CatalystCopyDeletionGuardDays"=dword:7
```

- If you remove the entire source backup repository or delete all the restore points from the source as "Delete from disk," then the replicated copies are not deleted.

CCJ limitations are as follows:

- It cannot process files generated by Backup Copy jobs.
- It cannot specify different retention for the replicated files.
- It does not support HPE StoreOnce Cloud Bank Storage.
- It does not support HPE StoreOnce Data Immutability.



Veeam Backup Copy Job with Catalyst write over WAN

Veeam backup copy jobs are designed to copy backups of selected VMs, or entire backup jobs with many VMs, to a remote location. They can also set simple or tiered GFS retention. The configuration shown in Figure 58 utilizes only one Veeam Gateway server located on the source side. It reads from the local repository and writes to the remote HPE StoreOnce appliance directly over the WAN.

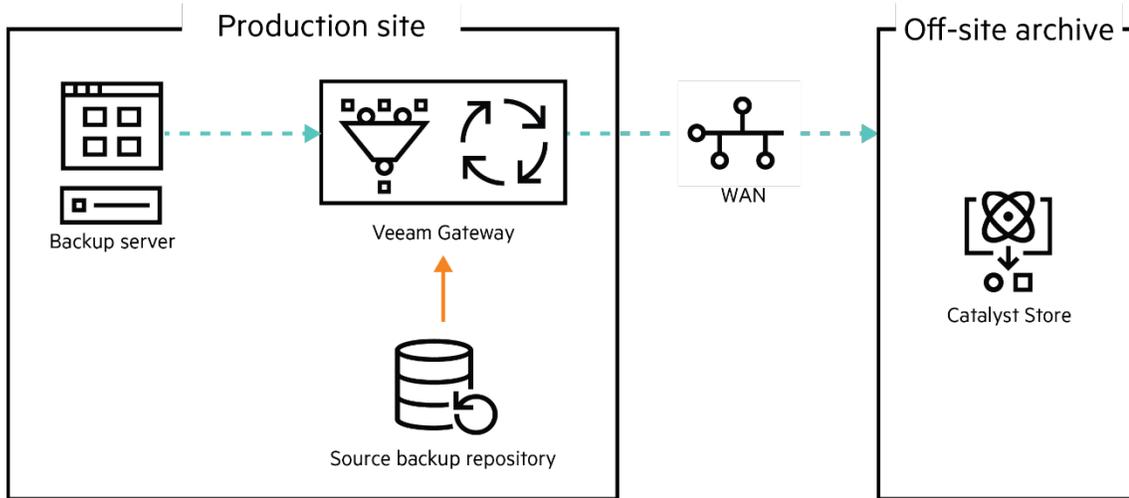


FIGURE 58. Utilizing a Veeam Gateway server to copy backups from a local HPE StoreOnce appliance to a remote HPE StoreOnce appliance

This configuration can be achieved by configuring on Veeam the same gateway servers for accessing the source backup repository and the BCJ target repository. In the example demonstrated in Figure 59, the gateway server “w2k16-12” is configured on both source and destination backup repositories.

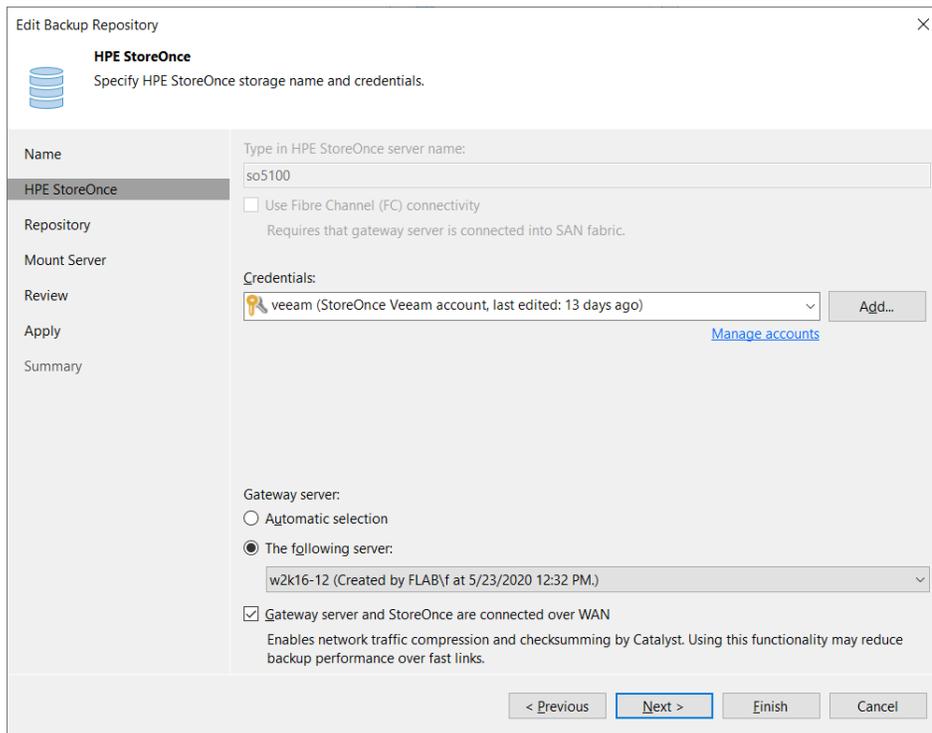


FIGURE 59. Configuring a gateway server for a Veeam backup repository



For the destination repository, enable **Gateway server and StoreOnce are connected over WAN**. For more information, see [Creating an HPE StoreOnce Catalyst-based Veeam backup repository](#) section. For more details on configuring a Veeam Gateway server, see [Deploying a Veeam Gateway server](#).

After configuring the same Veeam Gateway server for the source and destination backup repositories, the Backup Copy Job can now be configured. A BCJ is based on an “incremental forever” policy, and not on the forward incremental policy that works better with HPE StoreOnce. This policy is based on a single rolling full backup at the end of the chain.

The incremental forever causes a few limitations:

1. As long as incremental forever does not support multiple chains, then the daily retention is limited by the single chain length supported by the target HPE StoreOnce appliance. See helpcenter.veeam.com/docs/backup/vsphere/backup_copy_retention.html?ver=100 for more information.
2. It is necessary to set at least one weekly GFS. Otherwise, the rolling full backup chain will continue to grow forever.

As shown in Figure 60, the **Read the entire restore point from source backup instead of synthesizing it from increments** should not be selected, otherwise, the capacity utilization of the rolling full grows forever.

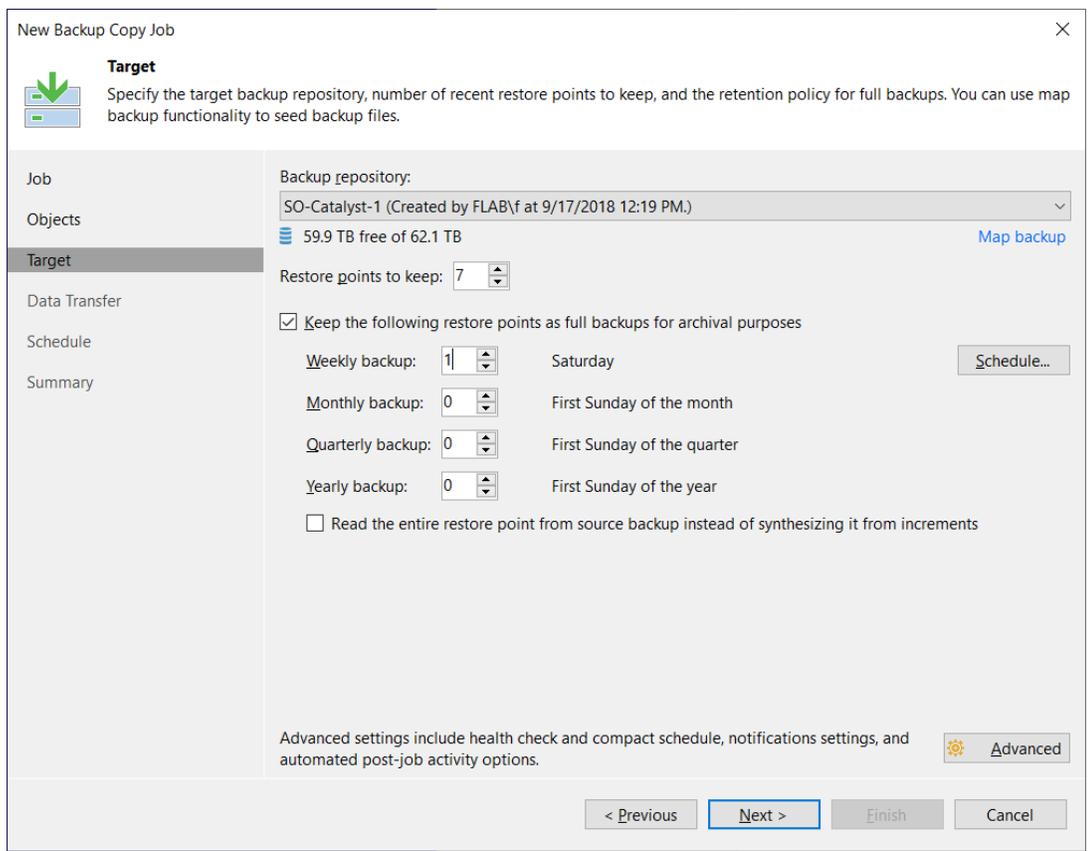


FIGURE 60. Backup Copy job “Target” details



The settings in the **Maintenance** tab shown in Figure 61 should be configured as follows:

Advanced Settings

Maintenance | Storage | Notifications | Scripts

Storage-level corruption guard

Perform backup files health check (detects and auto-heals corruption)

Monthly on: Last Saturday Months...

Weekly on selected days: Saturday Days...

Full backup file maintenance

Use these settings to defragment and compact full backup file periodically when the job schedule does not include periodic fulls.

Remove deleted items data after 30 days

Defragment and compact full backup file

Monthly on: Last Sunday Months...

Weekly on selected days: Sunday Days...

⚠ Maintenance is not required when periodic full backups are enabled.

Save As Default OK Cancel

FIGURE 61. Backup Copy job advanced maintenance settings

- Do not enable **Storage-level corruption guard**. While technically supported, storage-level corruption guard requires reading the entire backup file back to the gateway server, which in this architecture is typically over a WAN where link speed and reliability might affect its performance.

NOTE

Catalyst checksum data is sent to prevent transmission errors. HPE StoreOnce has additional data integrity algorithms to continuously verify its storage to prevent silent corruptions. Also, RAID 6 prevents data loss even after two concurrent disk failures.

- Do not enable **Full backup maintenance** for the same reasons mentioned above.



The settings in the **Storage** tab shown in Figure 62 should be configured as follows:

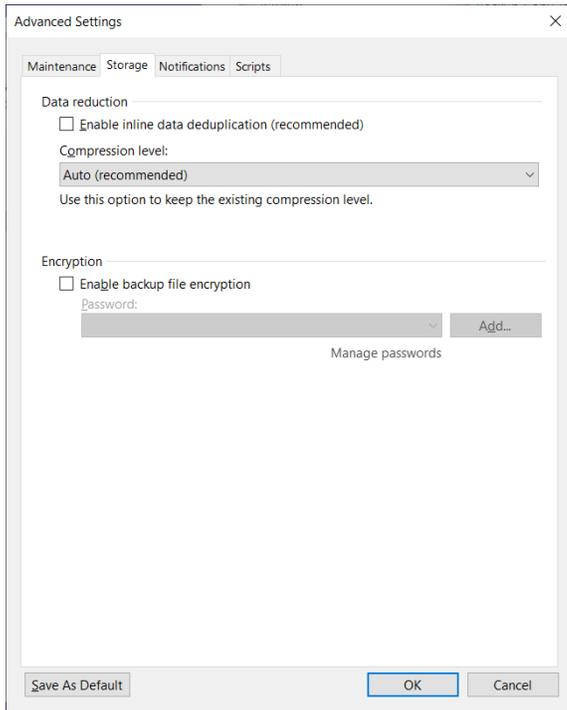


FIGURE 62. Backup Copy job advanced storage settings

- Do not check **Enable inline data deduplication** because deduplication is handled by HPE StoreOnce.
- Set the **Compression level** to **Auto** and make sure that the backup repository has **Decompress backup data blocks before storing**. See the [Creating an HPE StoreOnce Catalyst-based Veeam backup repository](#) section for more information.
- Do not select **Enable backup file encryption**. If the data on the source repository was already encrypted, the job will unencrypt before writing to the copy destination. If encryption for data-at-rest is required, use HPE StoreOnce encryption as described in the [Creating a Catalyst store for Veeam backups](#) section.



On the BCJ **Data Transfer** section, to send data directly from the source repository to the target repository, select **Direct** as shown in Figure 63.

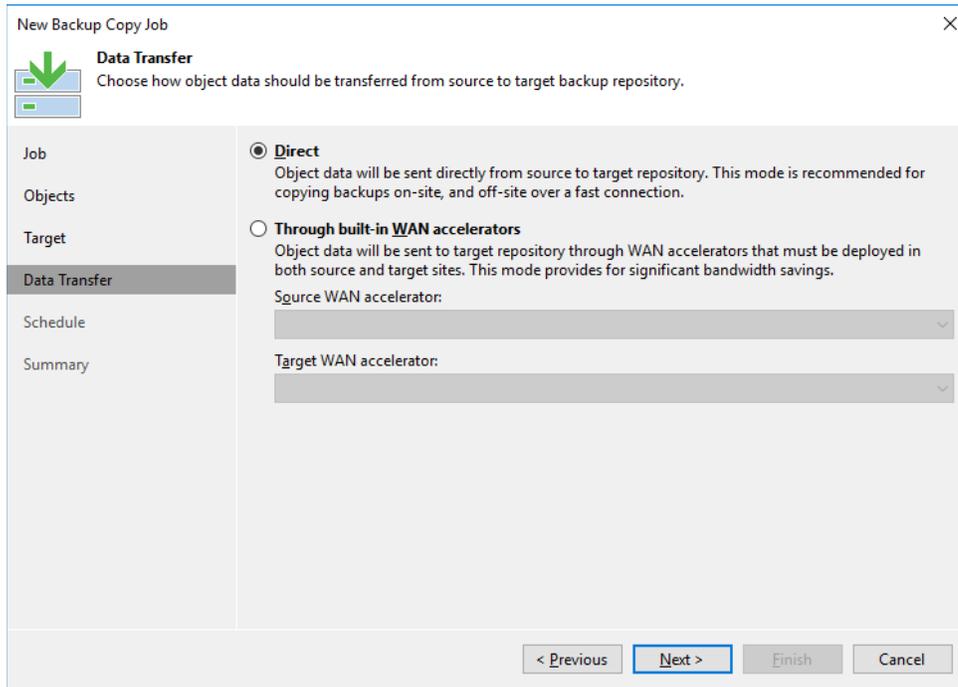


FIGURE 63. Specifying "Direct" as the Data Transfer type

Traditional Veeam configuration

The last option for copying remote backups involves the traditional process of backup and replication, shown in Figure 64. This setup comprises of two Veeam Gateway servers: one at the local site and one at the remote site. With this scenario, Veeam uses its own deduplication and compression engine with optional WAN accelerators. Compared to the options from the previous two sections, this deployment requires significantly higher bandwidth and requires a greater amount of hardware.

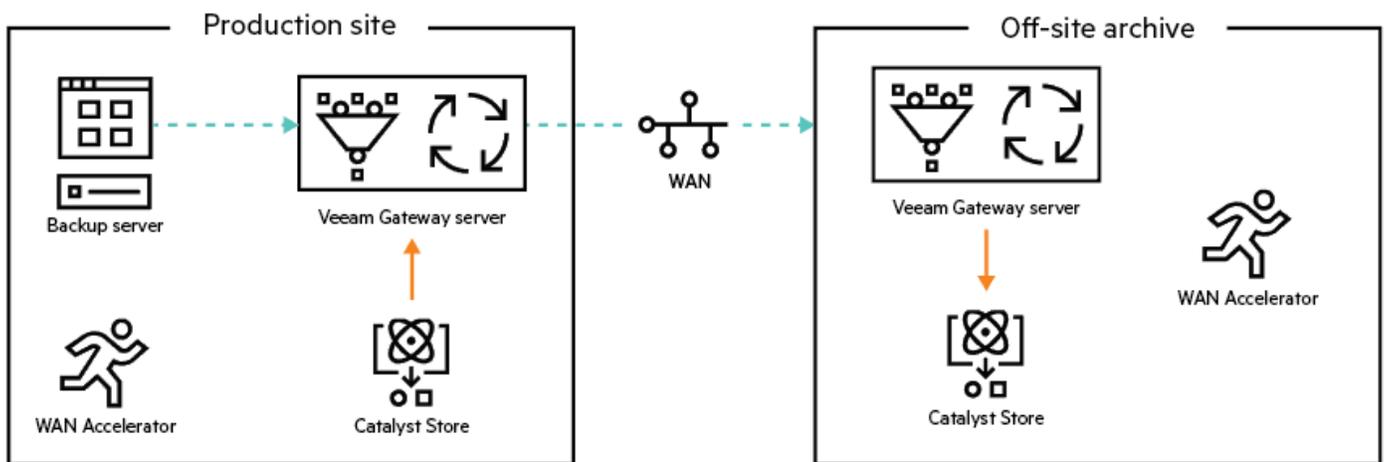


FIGURE 64. A traditional backup and replication setup with WAN accelerators and a Veeam Gateway server at both the local and remote sites



This configuration can be achieved by configuring two different gateway servers:

1. **Source GW** - This server is in the source site, close to the source repository.
2. **Destination GW** - This server is in the destination site, close to the BCJ target repository on an HPE StoreOnce Catalyst store.

For more details on configuring a Veeam Gateway server, see [Deploying a Veeam Gateway server](#). The Backup Copy job configuration steps are the same as in the [Veeam Backup Copy Job with Catalyst write over WAN](#) section.

Data transfer between the two gateway servers is not highly deduplicated as in the previous two solution designs. When the available bandwidth is very limited, it is also possible to configure the optional Veeam **Source WAN accelerator** and a **Target WAN accelerator**, as shown in [Figure 65](#). For more details on WAN accelerator configuration, see helpcenter.veeam.com/docs/backup/vsphere/wan_hiw.html?ver=100.

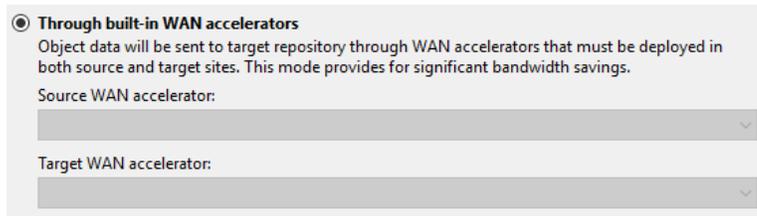


FIGURE 65. A Veeam Backup Copy job that uses WAN accelerators to transfer object data to the target backup repository

HPE Cloud Volume Backups with Veeam Backup & Replication

HPE Cloud Volumes are an enterprise-grade multi-cloud service for your applications in public clouds. This section provides an overview for using Veeam Backup & Replication with HPE Cloud Volume Backups (CVB). For more information on HPE Cloud Volumes, see cloudvolumes.hpe.com/welcome/block/overview/.

Perform the following steps to create a Veeam backup repository that utilizes CVB:

Step 1 – Create a CVB store

Log in to the CVB web portal, navigate to **Backup Stores**, shown in [Figure 66](#), and click **+ New Store**. Give the store a **Store Name** and **Store Description**, and then select the **Region** closest to your physical location. Click **Create Store**.

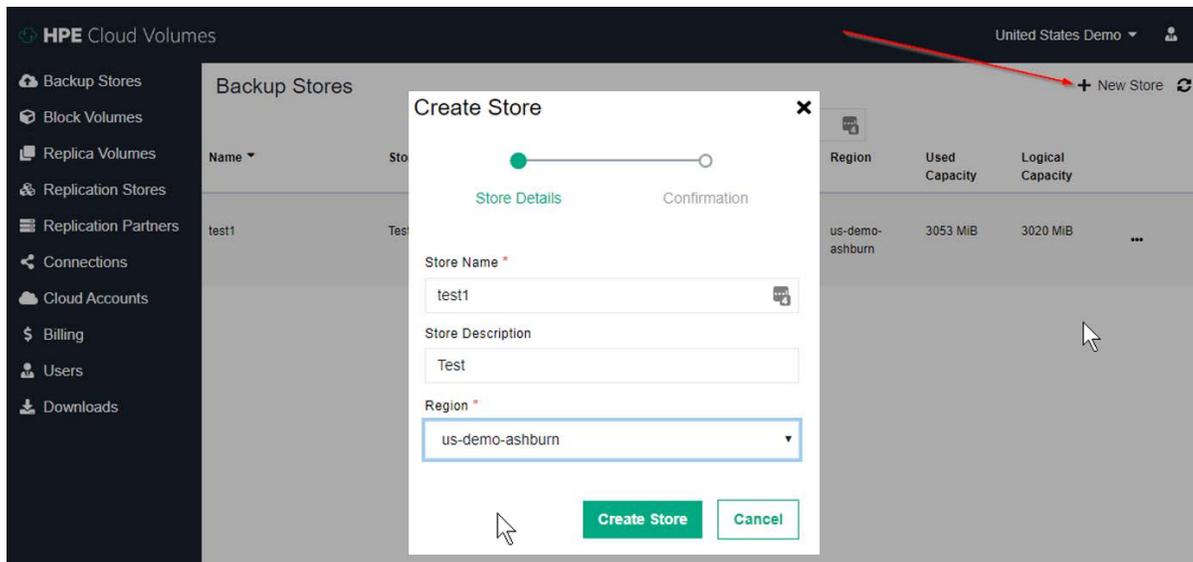


FIGURE 66. Creating a CVB store



NOTE

The **Store Name** must be a single word with less than 50 characters, be all lowercase, and contain only letters, numbers, and dashes.

Download the text file that contains your Username and Password, and then click **Close**. (It is possible to regenerate the credentials later by editing the Backup Store if needed) Then, download the Secure Client, as shown in Figure 67.

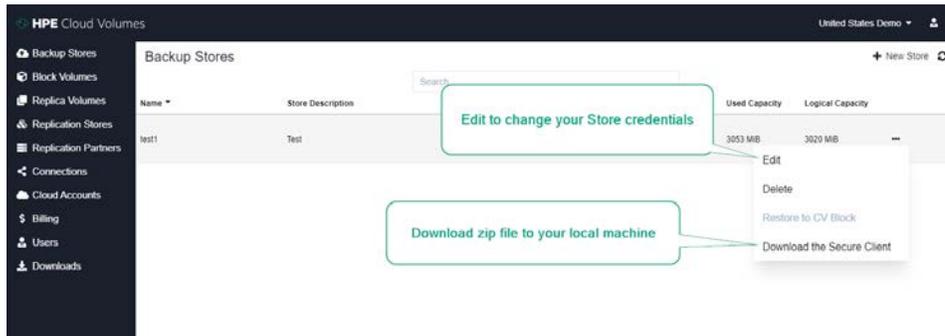


FIGURE 67. Editing CVB store credentials and downloading the Secure Client

Step 2 – Deploy Cloud Volumes Secure Client

The Secure Client is a Linux-based proxy server used to backup data from Veeam Backup & Replication to CVB. This is required to enforce end-to-end security over the open internet.

The following example uses CentOS 8 to deploy the Secure Client. The Linux commands can vary between Linux distributions.

The Secure Client is packaged as a ZIP file that consists of the following files:

- The Secure Client binary
- A YAML configuration file for the client
- A certificate for use by the client
- The private key for the certificate
- The Certificate Authority (CA) for the certificate
- A Readme file
- A license agreement file

The following are the prerequisites for running the Secure Client:

- Any Linux OS running in a physical or virtual machine
- A 2.0 GHz CPU with 4 Cores (minimum)
- 8 GB of RAM (minimum)
- Ports 9386, 9387, and 9388 are open in your firewall (outbound only)

To set up the Secure Client on the Linux machine:

1. Log in to Linux Machine and create the directory `/opt/cloudvolumes`.
2. Add a user named "secureclient" using the command: `sudo useradd secureclient`.
3. Upload the Secure Client to the Linux machine and extract the contents of the zip file to the `/opt/cloudvolumes` directory. An example command would be: `sudo unzip -j ./Downloads/hpe_cv_backup_secure_client.zip -d /opt/cloudvolumes/`.



4. Ensure that the Secure Client files have the proper file permissions by entering the following commands:
 - a. `sudo chmod 755 /opt/cloudvolumes/secure_client`
 - b. `sudo chmod 644 /opt/cloudvolumes/secure_client_config.yaml`
 - c. `sudo chmod 644 /opt/cloudvolumes/client.crt`
 - d. `sudo chmod 644 /opt/cloudvolumes/ca.crt`
 - e. `sudo chmod 600 /opt/cloudvolumes/client.key`
5. Grant ownership of the Secure Client folder and files to the “secureclient” user using the command:


```
sudo chown -R secureclient:secureclient /opt/cloudvolumes
```
6. Ensure ports 9386, 9387, and 9388 are open in your firewall outbound.
7. The Secure Client service must be started automatically on system boot. To do this, create a Linux system configuration file at `/etc/systemd/system/secure-client.service` with the following content:

```
[Service]
User=secureclient
Type=exec
ExecStart=/opt/cloudvolumes/secure_client -config /opt/cloudvolumes/secure_client_config.yaml
Restart=always
RestartSec=10s
LimitNOFILE=40000

[Install]
WantedBy=multi-user.target
```

8. Verify the YAML configuration file has the correct file path properties by modifying `/opt/cloudvolumes/secure_client_config.yaml` as follows:

```
# Certificate path for CDS signing authority
ca: /opt/cloudvolumes/ca.crt

# Client certificate issued by CDS to customer
ca: /opt/cloudvolumes/client.crt

# Client key issued by CDS to customer
key: /opt/cloudvolumes/client.key

# CBS public endpoint address
target1: demo-us-ashburn-1.cloudvolumes.hpe.com:9387
target2: demo-us-ashburn-1.cloudvolumes.hpe.com:9388

# Local ports to listen upon
source1: 0.0.0.0:9387
source2: 0.0.0.0:9388
```

The YAML file also contains the following properties that should not need to be modified. This includes:

- target1: The address and port to send traffic for HPE StoreOnce Catalyst command sessions
- target2: The address and port to send traffic for HPE StoreOnce Catalyst data sessions
- source1: The address and port to listen on for HPE StoreOnce Catalyst command sessions
- source2: The address and port to listen on for HPE StoreOnce Catalyst data sessions
- Generally, the address to listen on should always be 0.0.0.0.



9. Enable and start the Secure Client service by running

```
sudo systemctl enable secure-client
sudo systemctl start secure-client
```

Step 3 – Configure a Veeam Backup Repository

The final step is to create a Veeam Backup Repository that points to a CVB store. For this step, a CVB store is created using the same steps demonstrated in [Creating an HPE StoreOnce Catalyst-based Veeam backup repository](#). In this instance, use the IP of the Linux-based proxy server deployed in [Step 2 – Deploy Cloud Volumes Secure Client](#) for the IP of the HPE StoreOnce server name. For more about Veeam Backup & Replication support for HPE Cloud Volumes, see veeam.com/kb3143.

Recovering data from an HPE StoreOnce Catalyst repository

Veeam Backup & Replication handles backup files saved on backup repositories based on HPE StoreOnce Catalyst stores the same as repositories based on other storage types, such as local disks or NAS. HPE StoreOnce Catalyst store backups can be used for Instant VM Recovery (IVMR), Veeam On-Demand Sandbox including SureBackup, restoring guest VM files, restoring the entire VM, restoring virtual disks, and other restore options offered by Veeam Backup & Replication. For more information regarding restore performance, see the HPE Reference Configuration at h20195.www2.hp.com/v2/Getdocument.aspx?docname=a00064440enw.

Instant VM Recovery supports bulk processing so you can immediately restore multiple workloads at once. If you perform Instant VM Recovery for several workloads, Veeam Backup & Replication uses the resource scheduling mechanism to allocate and use optimal resources required for Instant VM Recovery.

Migrating Veeam Backup repositories to/from an HPE StoreOnce Catalyst store

There is an easy and effective methodology for migrating entire Veeam backup repositories to new storage platforms and even across different storage types. The migration process takes care of migrating all the backup files and the associated restore points. More information can be found at helpcenter.veeam.com/docs/backup/hyperv/sobr_evacuate.html?ver=100.

Examples of supported migrations include:

- Existing repository based on NAS/NTFS/XFS/ReFS → New HPE StoreOnce Catalyst-based backup repository
- Existing HPE StoreOnce Catalyst-based backup repository → New repository based on NAS/NTFS/ReFS
- Existing HPE StoreOnce Catalyst-based backup repository → New HPE StoreOnce Catalyst-based backup repository

Some limitations of this methodology include:

- The process migrates entire backup repositories. It is not possible to select data subsets.
- The process does not modify the data format. The existing data format must be compatible with the target repository. Specifically, compression, deduplication, encryption, backup policy, and chain length are not changed. Before starting the migration to HPE StoreOnce, make sure that the source format is compatible with the target.
- Make sure that the format of each Veeam Backup Job is compatible with the settings outlined in [Creating a Veeam backup job for VMs](#).



The following example demonstrates how to migrate a repository to an HPE StoreOnce Catalyst-based repository:

1. Navigate to **BACKUP INFRASTRUCTURE**, right-click **Scale-out Repositories**, and select **Add scale-out backup repository...**, demonstrated in Figure 68.

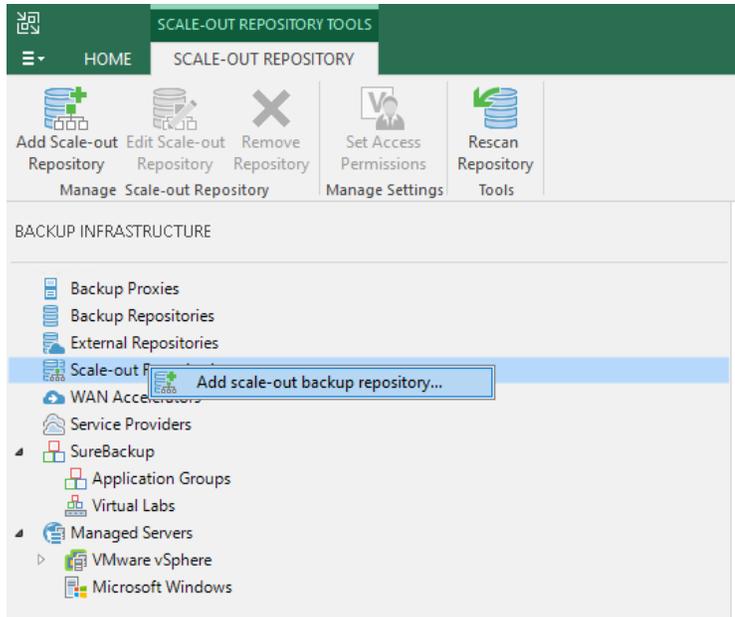


FIGURE 68. Adding a scale-out backup repository

2. Provide a **Name** and optional **Description** to the scale-out backup repository, illustrated in Figure 69.

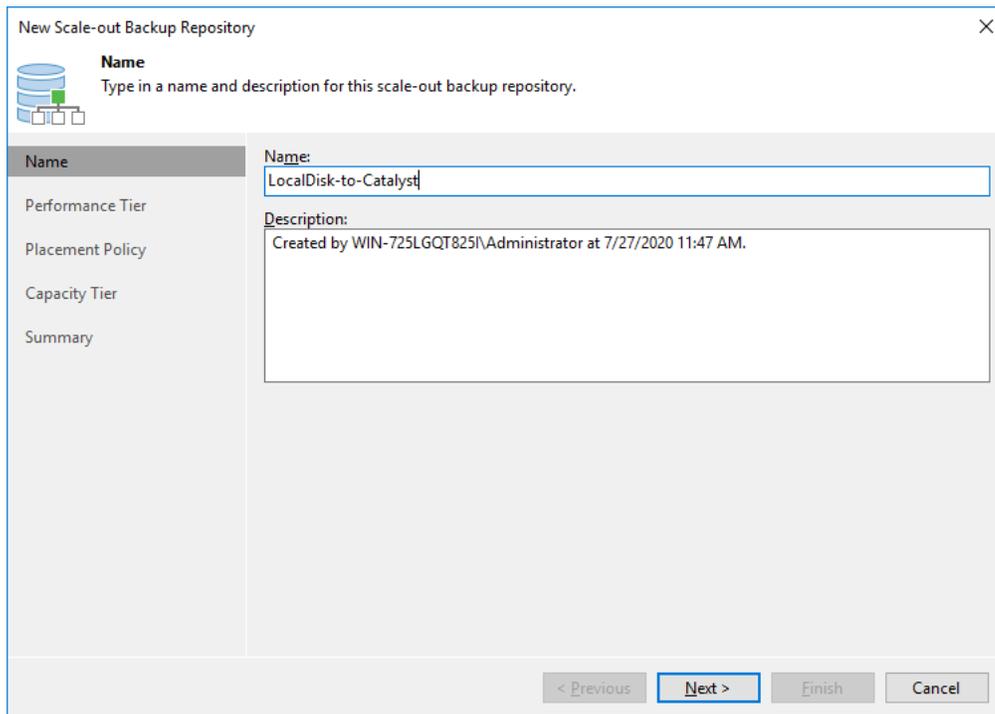


FIGURE 69. Providing a name and description for the scale-out backup repository



3. Add the source repository and the target repository, shown in Figure 70. The order that the repositories are added does not matter.

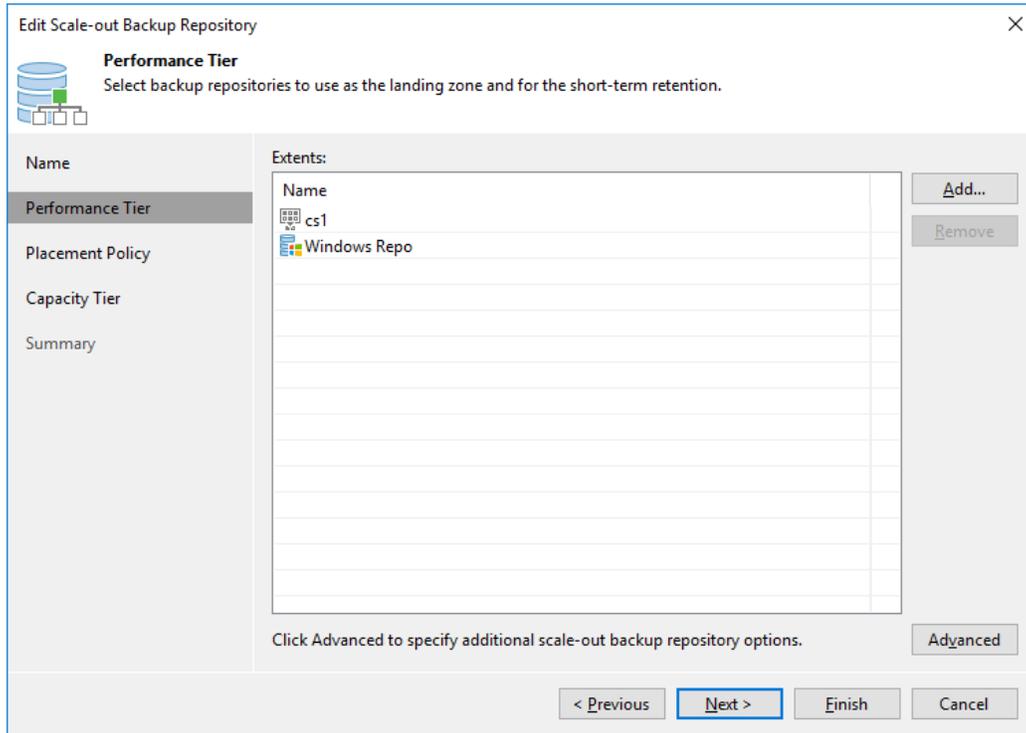


FIGURE 70. Selecting the "Performance Tier" for the scale-out backup repository

NOTE

Veeam automatically updates backup jobs utilizing the backup repositories added in Figure 70 to this newly created Scale-Out Backup Repository.

4. Make sure to set **Use per-VM backup files** (Figure 71) on the advanced settings of Figure 70.

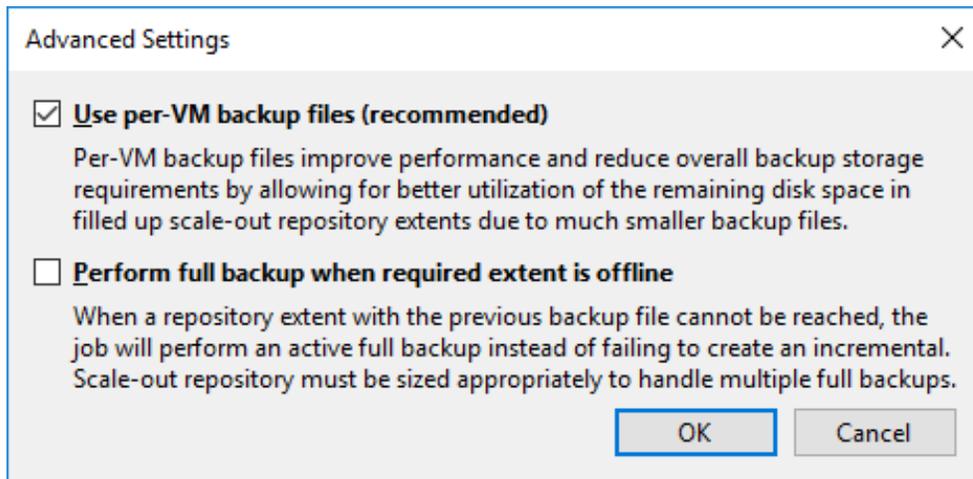


FIGURE 71. Advanced performance tier settings of a Veeam scale-out repository



5. Select **Data locality** to migrate all backup files to the target repository, demonstrated in Figure 72.

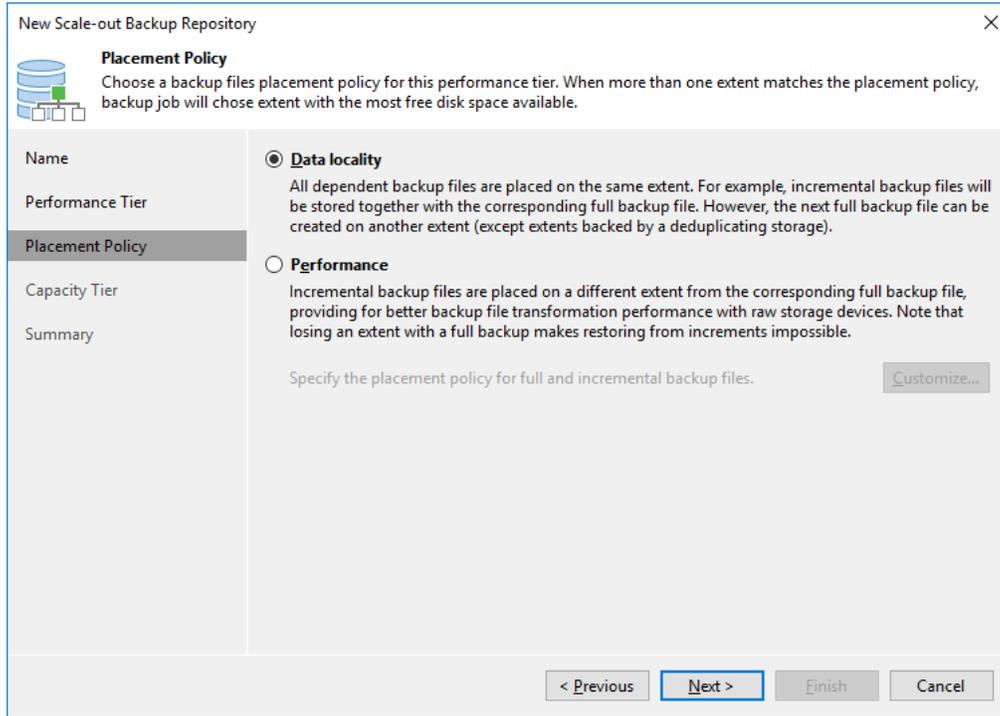


FIGURE 72. Choosing a "Placement Policy" for the scale-out backup repository

6. For the **Capacity Tier** shown in Figure 73, do not select **Extend scale-out backup repository capacity with object storage**.

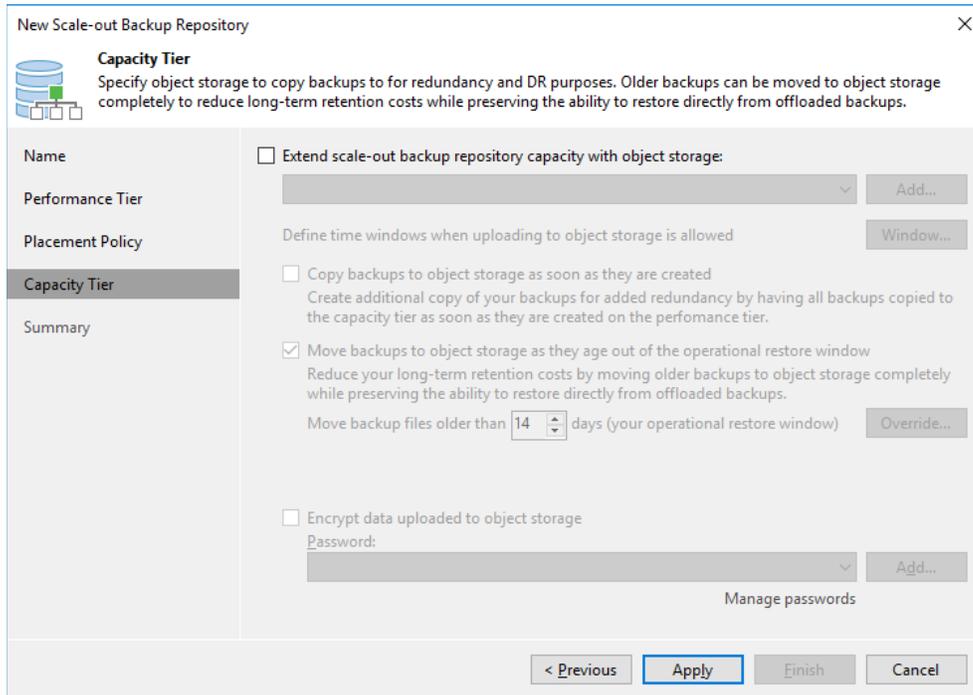


FIGURE 73. Optionally enabling object storage with a scale-out repository



7. Review the summary and click **Finish** (Figure 74).

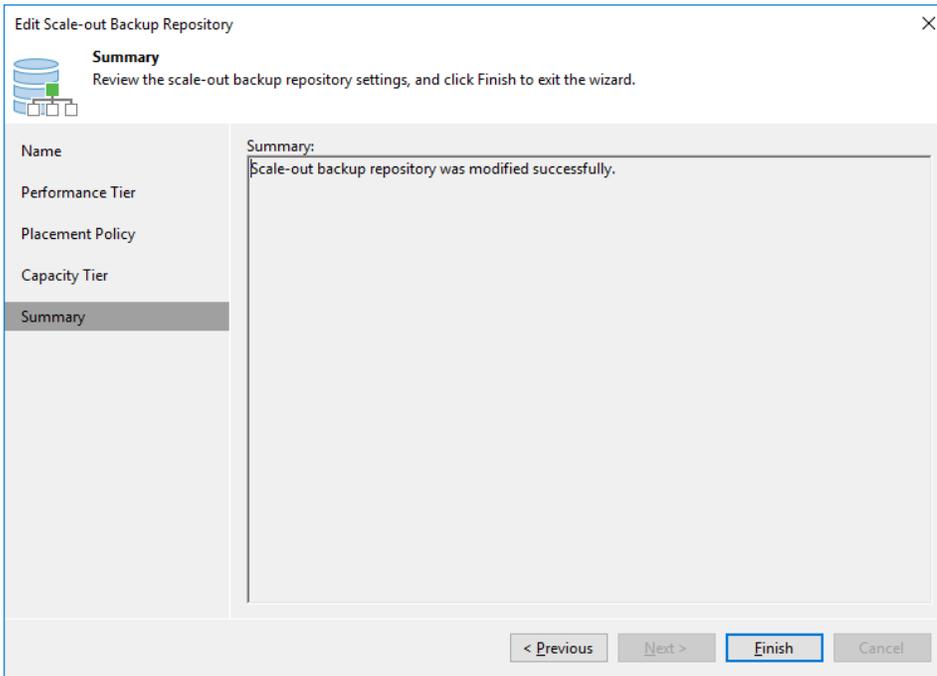


FIGURE 74. The summary page of the scale-out backup repository wizard

8. Referring to Figure 75, click the newly created scale-out repository containing the source repository and the target repository. In this scenario, “cs1” was the target repository, which was an HPE StoreOnce Catalyst store, and “Windows Repo” was the source repository, which was disk storage on a Windows server.

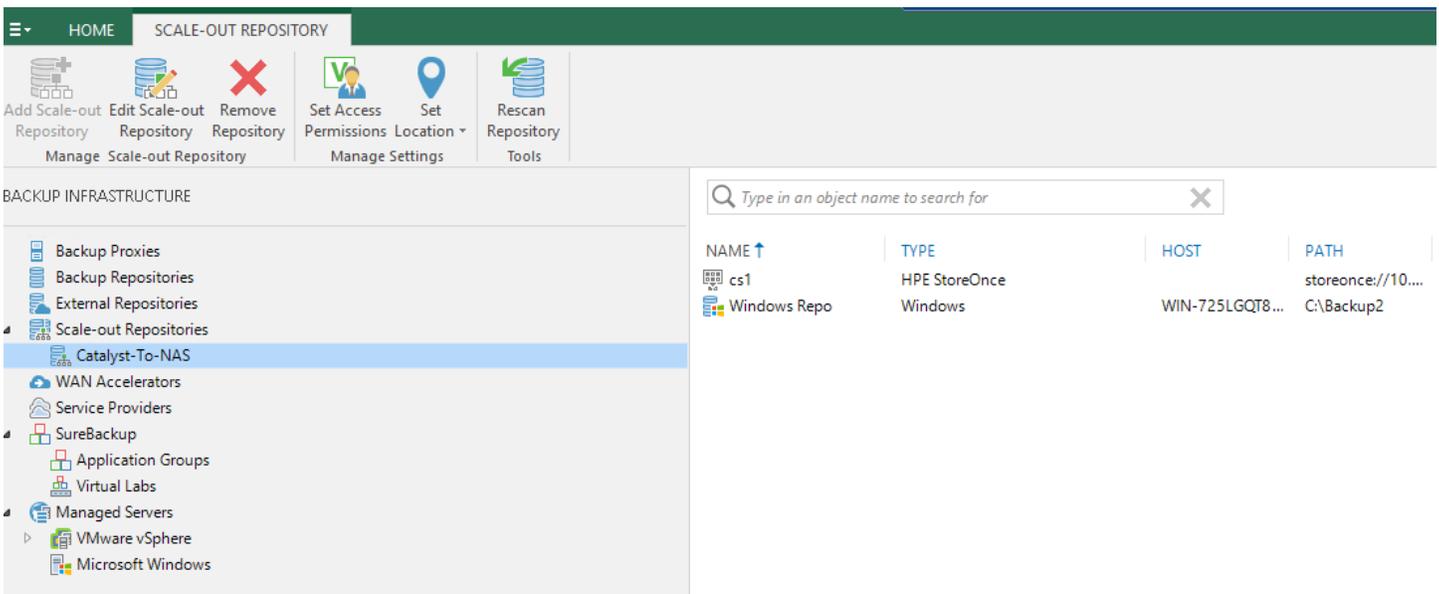


FIGURE 75. Selecting a scale-out repository



- 9. Right-click the source repository and put the source repository into **Maintenance mode**, shown in Figure 76. This is necessary to evacuate the backups from the source repository into the target repository.

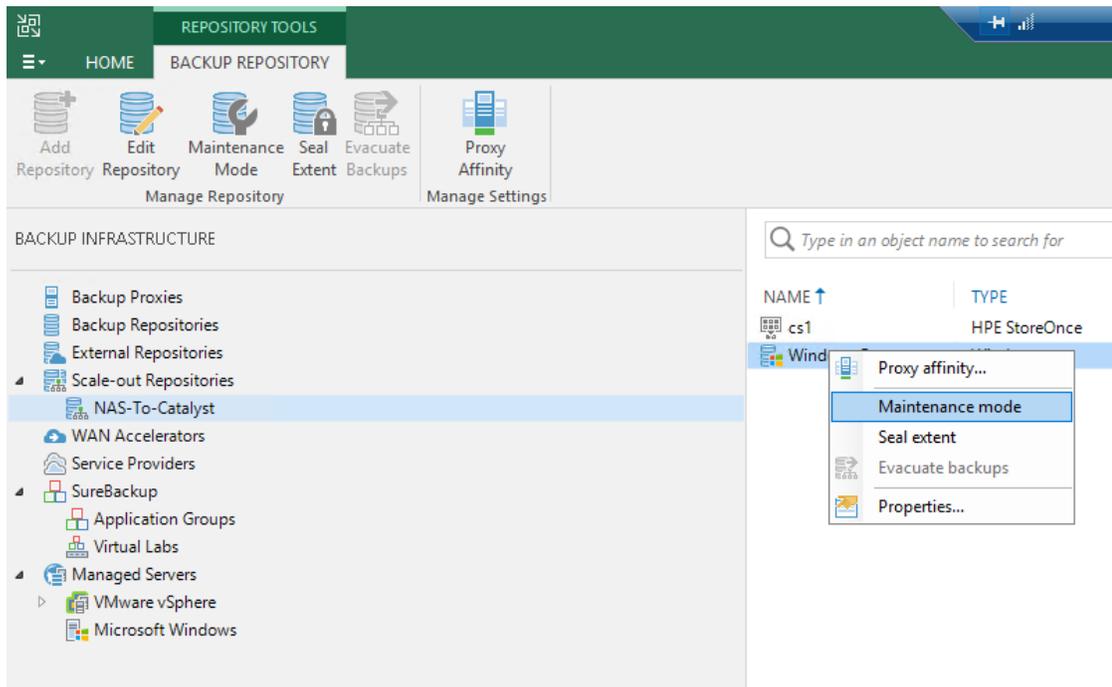


FIGURE 76. Placing a backup (source) repository into maintenance mode

- 10. Once in **Maintenance mode**, right-click the source repository and select **Evacuate backups**, demonstrated in Figure 77.

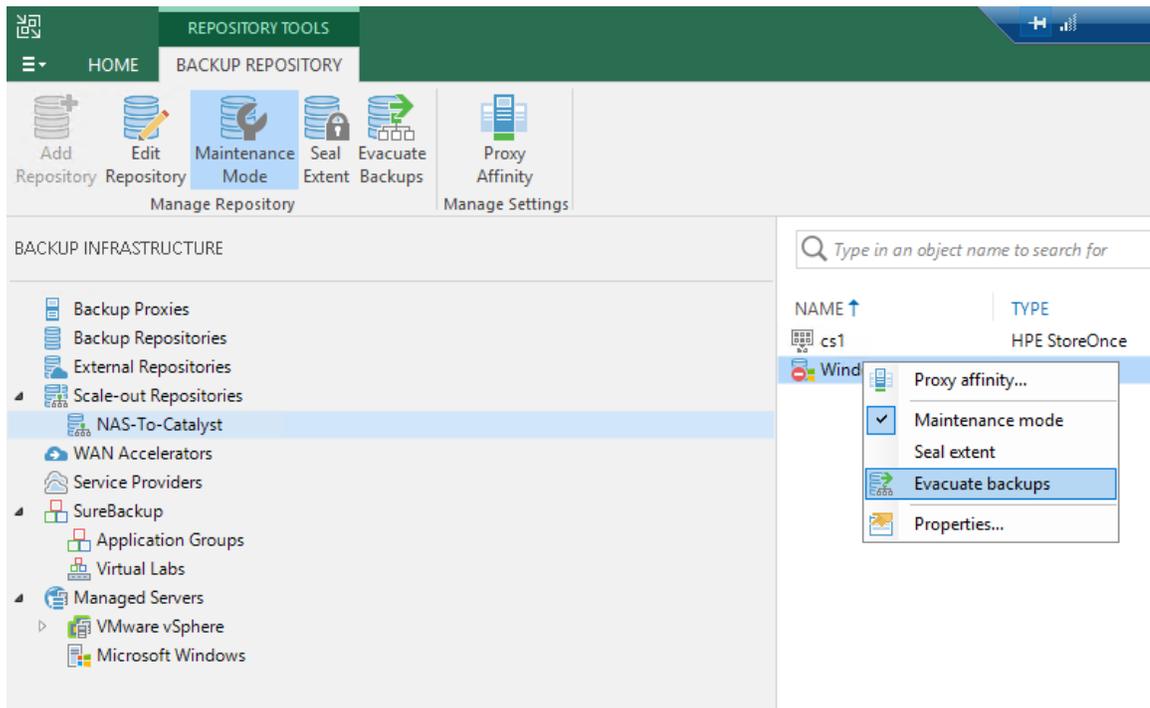


FIGURE 77. Evacuating the backups from a repository into a different Veeam repository



11. As shown in Figure 78, Veeam will now migrate the backups from the source repository to the target repository.

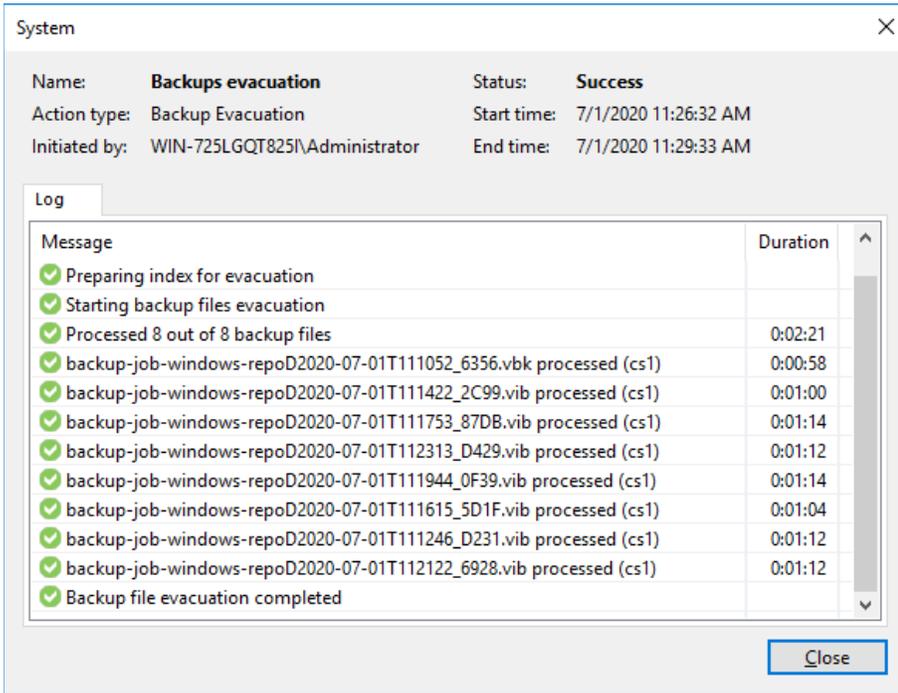


FIGURE 78. Veeam executing backup evacuation and migrating the backups to a separate repository

After the evacuation process is complete, it is possible to edit the SOBR and **Remove** the source extent, as shown in Figure 79.

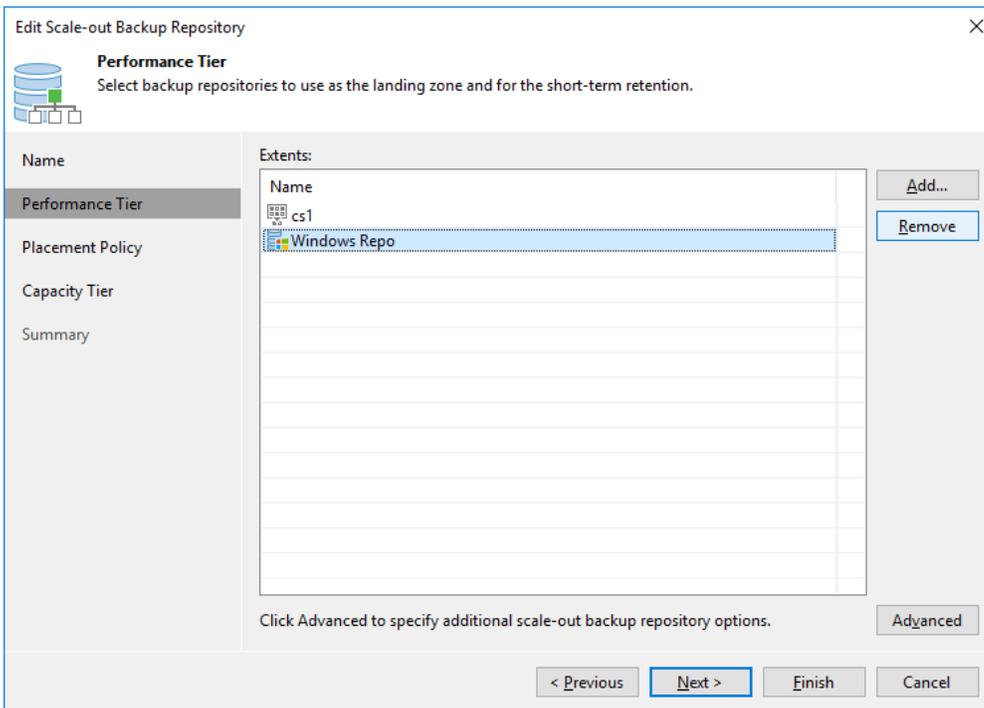


FIGURE 79. Removing the source repository from a Scale-Out Backup Repository



After the migration is completed, there is no need to convert the destination SOBR to a normal repository. The Backup Jobs that previously wrote to the migrated repository will now use the destination repository. All of the restore points are visible and available for restore operations, and the process for deleting expired restore points should work as expected.

NOTE

As demonstrated in Figure 80, the number of concurrent tasks should be limited when migrating a backup repository with a large number of restore points to avoid overwhelming the data-mover server.

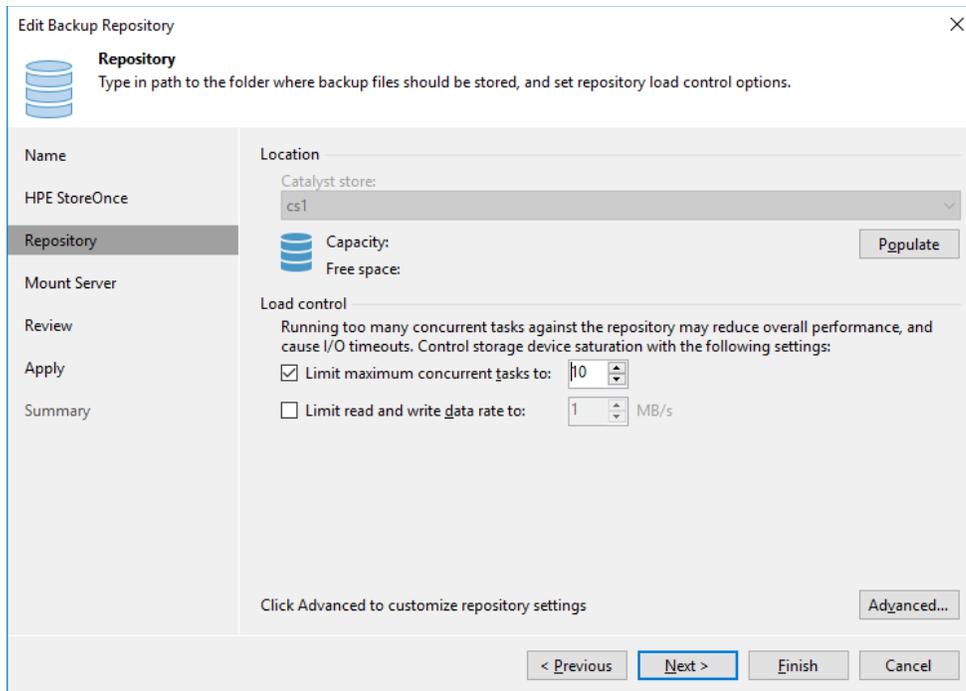


FIGURE 80. Limiting the maximum concurrent tasks for a Veeam repository

SUMMARY

Veeam Backup & Replication provides a solid software package for VM/physical backup, replication, and recovery; however, without effective storage infrastructure, meeting the demands of the business can be difficult. HPE StoreOnce systems offer powerful scale-out storage infrastructure that can be used in combination with Veeam Backup & Replication to significantly reduce the size of backup data and to streamline the virtual appliance data-protection foundation. With an automated backup and disaster recovery operation, together with secure data retention and built-in data encryption for data-at-rest, HPE StoreOnce is a powerful and reliable data protection tool.

With the HPE StoreOnce Catalyst backup solution, Hewlett Packard Enterprise, together with Veeam, offers an end-to-end data protection solution to recover virtual appliances rapidly, effectively, and with minimal disruption to business.



Reference Architecture

RESOURCES AND ADDITIONAL LINKS

HPE StoreOnce Data Protection Backup Appliances, hpe.com/storage/StoreOnce

HPE StoreOnce Compatibility Matrix, hpe.com/Storage/StoreOnceSupportMatrix

Protecting Data from Ransomware with HPE StoreOnce Catalyst, hpe.com/V2/getpdf.aspx/A00042003ENW.pdf

HPE StoreOnce (Veeam Help Center), helpcenter.veeam.com/docs/backup/vsphere/deduplicating_appliance_storeonce.html

Long-Term Retention Policy (Veeam Help Center), helpcenter.veeam.com/docs/backup/vsphere/gfs_retention_policy.html

HPE Reference Configuration for HPE 3PAR Storage and HPE StoreOnce with Veeam Availability Suite, hpe.com/v2/getdocument.aspx?docname=a00061747enw

HPE Reference Configuration for Veeam Availability Suite with HPE Nimble Storage, hpe.com/v2/getdocument.aspx?docname=a00079582enw

HPE Cloud Volumes Block Overview, cloudvolumes.hpe.com/welcome/block/overview/

HPE Reference Architectures, hpe.com/info/ra

HPE Servers, hpe.com/servers

HPE Storage, hpe.com/storage

HPE Networking, hpe.com/networking

HPE Technology Consulting Services, hpe.com/us/en/services/consulting.html

To help us improve our documents, please provide feedback at hpe.com/contact/feedback.

LEARN MORE AT

hpe.com/us/en/storage/storeonce.html

© Copyright 2020-2021 Hewlett Packard Enterprise Development LP. The information contained herein is subject to change without notice. The only warranties for Hewlett Packard Enterprise products and services are set forth in the express warranty statements accompanying such products and services. Nothing herein should be construed as constituting an additional warranty. Hewlett Packard Enterprise shall not be liable for technical or editorial errors or omissions contained herein.

Intel and Xeon are trademarks of Intel Corporation in the U.S. and other countries. Microsoft, Hyper-V, and Windows are registered trademarks or trademarks of Microsoft Corporation in the United States and/or other countries. VMware, VMware vCenter Server, VMware vSphere Hypervisor, VMware vSphere High Availability, VMware vSphere Distributed Resource Scheduler, and VMware vSphere Storage vMotion are registered trademarks or trademarks of VMware, Inc. and its subsidiaries in the United States and other jurisdictions. Veeam Backup & Replication and SureBackup are registered trademarks or trademarks of Veeam. Linux is the registered trademark of Linus Torvalds in the U.S. and other countries. All third-party trademarks are the property of their respective owners.

