

VEEAM



# Veeam Backup & Replication version 9.x with HPE StoreOnce Catalyst

Configuration guide

1+1=3  
HPE + Veeam  
Better Together

## Contents

<b>Configuration guide</b> .....	<b>1</b>
<b>Introduction</b> .....	<b>3</b>
Software and security requirements .....	3
Veeam & StoreOnce integration components .....	4
Supported topologies and use cases for Catalyst based repositories: .....	4
Configuring Catalyst Store Access Control .....	5
<b>Creating a Catalyst Store for Veeam Backups</b> .....	<b>6</b>
<b>How to connect to a StoreOnce Catalyst store via IP and Fibre Channel</b> .....	<b>8</b>
<b>Prerequisites for using Catalyst over Fibre Channel</b> .....	<b>9</b>
<b>Initial Fibre Channel configuration</b> .....	<b>9</b>
<b>Configuring of the Veeam Gateway server for Fibre Channel</b> .....	<b>11</b>
<b>Creating a StoreOnce Catalyst-based Veeam Backup Repository</b> .....	<b>12</b>
Creating a Veeam backup job .....	17
<b>Copying Veeam backups to a remote site</b> .....	<b>23</b>
Architecture "A" – Site to site through source and target Veeam gateway servers with local StoreOnce Catalyst stores ..	23
Architecture "B" – Site to site using only source gateway servers with remote Catalyst stores (suggested for bandwidth reduction and performance) .....	24
<b>Implementation steps for copying Veeam backups to a remote site</b> .....	<b>25</b>
Architecture "A" configuration .....	25
Architecture "B" restore configuration .....	35
<b>Utilizing backups</b> .....	<b>36</b>
<b>Resources</b> .....	<b>36</b>
<b>Appendix</b> .....	<b>37</b>
<b>About Veeam Software</b> .....	<b>38</b>
<b>About HPE StoreOnce</b> .....	<b>38</b>

## Introduction

This guide is for users responsible for setting up and managing data Availability infrastructures that include Veeam® Backup & Replication™ and Hewlett Packard Enterprise (HPE) StoreOnce backup appliances configured with Catalyst backup targets. It assumes the reader is familiar with Veeam Backup & Replication and HPE StoreOnce. Please refer to the Veeam Backup & Replication user guide for more detailed information. For information on HPE StoreOnce, please visit [www.hpe.com/storage/StoreOnce](http://www.hpe.com/storage/StoreOnce).

Veeam Backup & Replication integrates with StoreOnce Catalyst to enable source-side deduplication, virtual synthetic full backups and faster restores. This leads to more efficient backup data transfer over the network, reduced storage space and faster performance in comparison to backup targets with traditional NAS protocol (CIFS and NFS) connectivity.

This guide covers 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, two copy scenarios are compared. The first has a single local Veeam Gateway server and the second has two Veeam Gateway servers, one at the local site and one at the remote site.

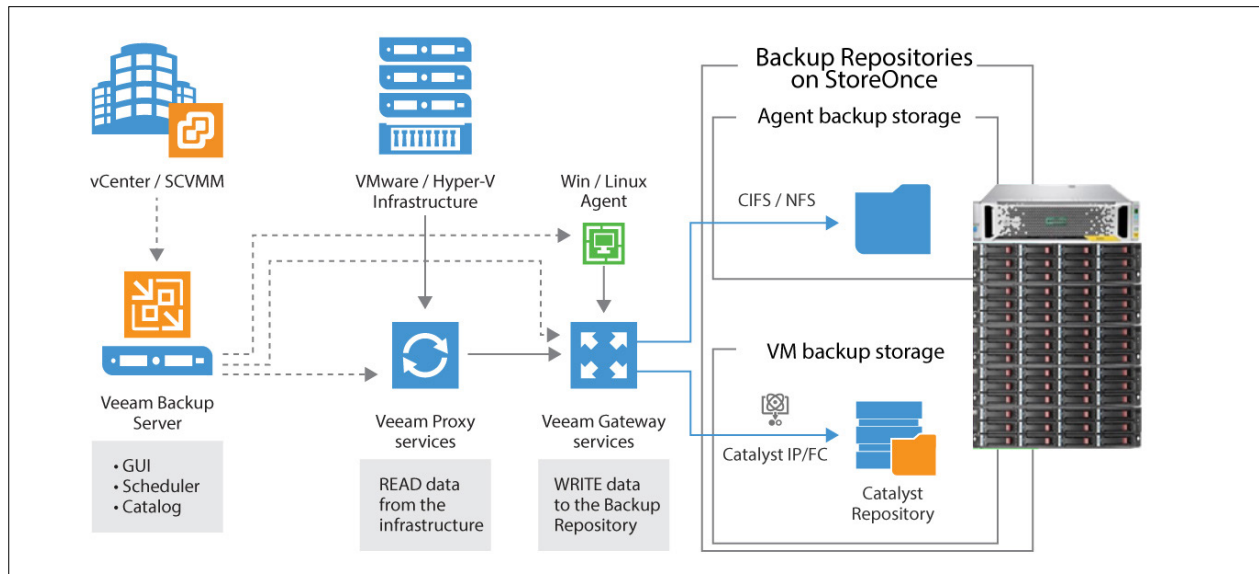
## Software and security requirements

- Veeam Backup & Replication Enterprise or Enterprise Plus version 9 Update 2 or later with StoreOnce 3.13.1 or later for general Catalyst integration
- Veeam Backup & Replication Enterprise or Enterprise Plus version 9.5 or later with StoreOnce 3.15.1 or later is required for Instant VM Recovery® and other vPower® NFS-based features. For full compatibility details please see the support matrix at [www.hpe.com/storage/DAPRCompatibility](http://www.hpe.com/storage/DAPRCompatibility)
- Administrator credentials for the HPE StoreOnce appliance
- Recommended –10 GbE, 8 Gb FC or 16 Gb FC connectivity from the Veeam Gateway server to the StoreOnce appliance

## Veeam & StoreOnce integration components

The figure below gives an overview of the main components in a Veeam Backup & Replication and StoreOnce deployment. This guide will focus on Veeam integration with StoreOnce Catalyst based backup repositories.

### Architecture overview



## 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.

### Backup over LAN

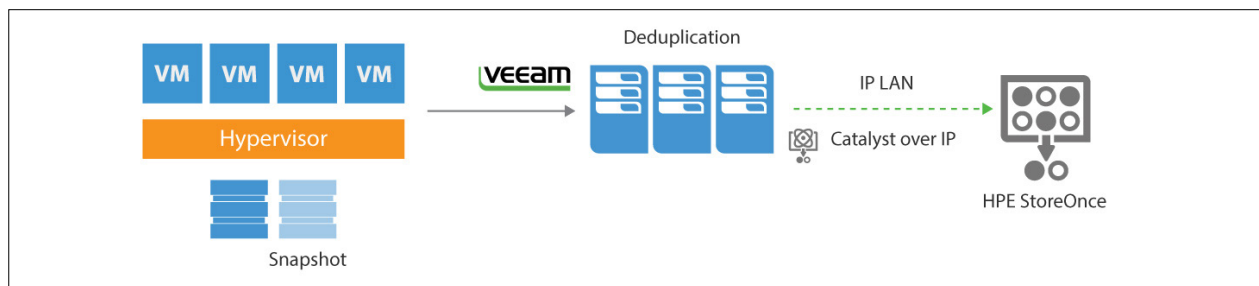


Figure 1: Veeam using Catalyst for source-side deduplication over LAN

## LAN-free Backup




Figure 2: Veeam using Catalyst for source-side deduplication over fibre channel

## Backup and backup copy over WAN

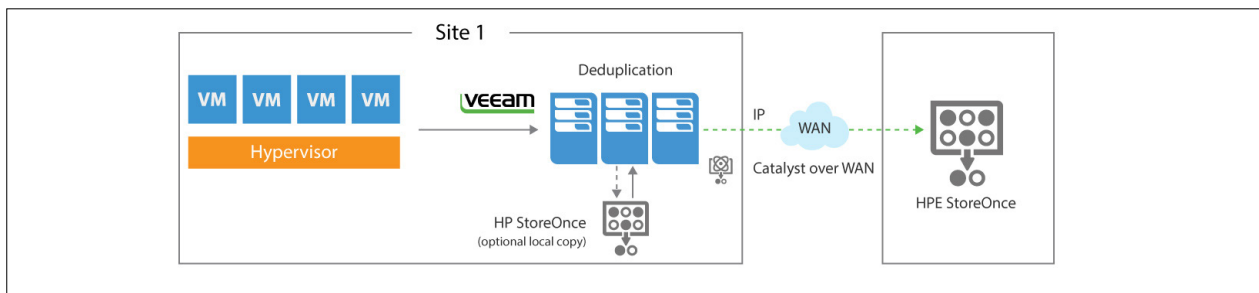


Figure 3: Veeam using Catalyst for source-side deduplication over a WAN for direct backup or backup copy to create a secondary copy of local backup data.

## Configuring Catalyst Store Access Control

It's a good practice to restrict access to the Catalyst stores only to the clients that should write and read backup data from the store. StoreOnce Catalyst Client access permission checking is used to control this access and is implemented by creating Catalyst client identifiers. Two modes can be set on StoreOnce Catalyst stores for access control:

- **No Client Access Permission Checking** allows all clients to have access to the StoreOnce Catalyst store(s)
- **Client Access Permission Checking** allows a selected group of client identifiers to access to the StoreOnce Catalyst store(s). To enable this mode, enable Client Access Permission Checking on the Settings tab of the StoreOnce GUI. To access this tab, navigate to *HPE StoreOnce>StoreOnce Catalyst>Settings*. When client access permissions are enabled, a Catalyst store must grant access to all clients that need access to the store. This is done using the Permissions tab of the configuration page for that StoreOnce Catalyst store. These credentials will be used in the Veeam Backup Repository creation wizard.

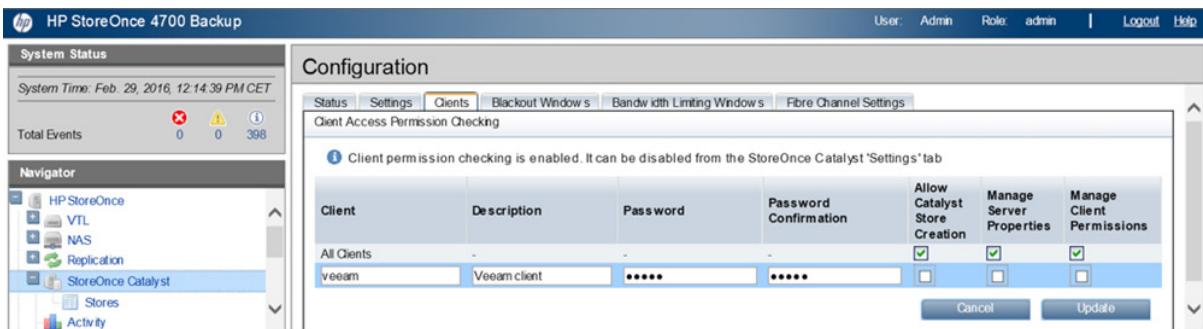


Figure 4: StoreOnce GUI for viewing and assigning permissions and a password to Catalyst client identifiers.

## Creating a Catalyst Store for Veeam Backups

Before configuring a Backup Repository, a StoreOnce Catalyst store(s) must be created on the StoreOnce appliance(s) that will store the backups and backup copies.

- Log onto the HPE StoreOnce Management Console (StoreOnce GUI) with administrator privileges.
- Navigate to *HPE StoreOnce>StoreOnce Catalyst>Stores*. Click *Create*.
- Change the store details as appropriate. If necessary, find the table describing the store details in the *StoreOnce Catalyst Stores* section of the HPE StoreOnce Backup system user guide for the product in use.
- Specify the Primary and Secondary Transfer Policies. There are two possible values, Low Bandwidth (LBW) or High Bandwidth (HBW). Specifying the value of the transfer policy determines how the backup data is transferred between the Veeam Server and the StoreOnce appliance. As shown in Figure 5, there is a primary and secondary transfer policy. The Primary (Default) Transfer Policy configures the default transfer policy. The Secondary Transfer Policy is set to the same value as the Primary (Default) Transfer Policy to enforce that policy or it is set to the other value to permit the Veeam server to use the secondary transfer policy if it calculates that is necessary.

**Recommendation:** The transfer policy configured between the Veeam server and the StoreOnce appliance determines if the backup data transfer will be low bandwidth (LBW) or high bandwidth (HBW). LBW transfer processes the backup data on the Veeam server to identify and only transfer unique backup data to the 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 transfer sends all the backup data from the Veeam server to the StoreOnce appliance. Duplicate data is identified and deduplicated on the StoreOnce appliance. This is called "target-side deduplication." Target-side deduplication does not reduce the backup data transfer over the network and should only be used with high throughput networks – 10 GbE, 8 Gb FC, or 16 Gb FC.

When testing or running a proof of concept, it is recommended to simulate as closely as possible the production setup. If a single backup job is used for testing, a HBW transfer policy may be faster than a LBW transfer policy. As more backup jobs are run in parallel, the speed advantage of a LBW transfer policy will increase. **LBW transfer (source-side deduplication) is recommended** for most use cases and should be used unless there is a reason not to use it.

**Note:** Do not create too many Catalyst stores. The maximum number of Catalyst Stores is reported in the StoreOnce user guide. There isn't a size limit for each Catalyst Store, but it's generally better to keep a store at less than ~100TB of used capacity (i.e.: 1 PB as logical capacity for a 10:1 deduplication).

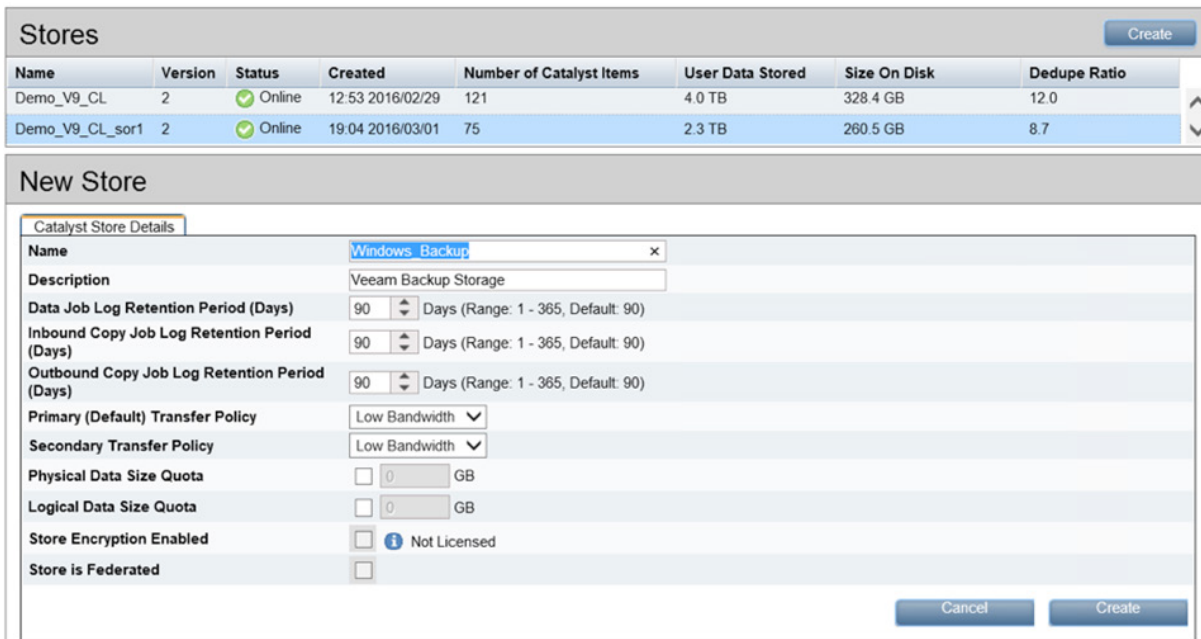


Figure 5: StoreOnce GUI showing a Catalyst store configured with a low-bandwidth primary and secondary transfer policy to reduce network bandwidth usage.

The recommended transfer policies are slightly different for Veeam Backup & Replication versions 9.0 and 9.5.

### Transfer Policy settings for Veeam Backup & Replication v9

- For source-side deduplication – recommended
  - Primary (Default) Transfer Policy = *Low bandwidth*
  - Secondary Transfer Policy = *Low bandwidth*
- For target-side deduplication
  - Primary Transfer Policy = High bandwidth
  - Secondary Transfer Policy = (any value)

### Transfer Policy settings for Veeam Backup & Replication 9.5

- For source-side deduplication – recommended
  - Primary (Default) Transfer Policy = Low bandwidth
  - Secondary Transfer Policy = (any value)
- For target-side deduplication
  - Primary Transfer Policy = High bandwidth
  - Secondary Transfer Policy = (any value)

By default, the new Catalyst store is accessible from "All Clients." Access can be restricted to named clients as shown in the screenshot below, as an option. This requires a client identifier to be created as described in the *Configuring Catalyst Store Access Control* section above.

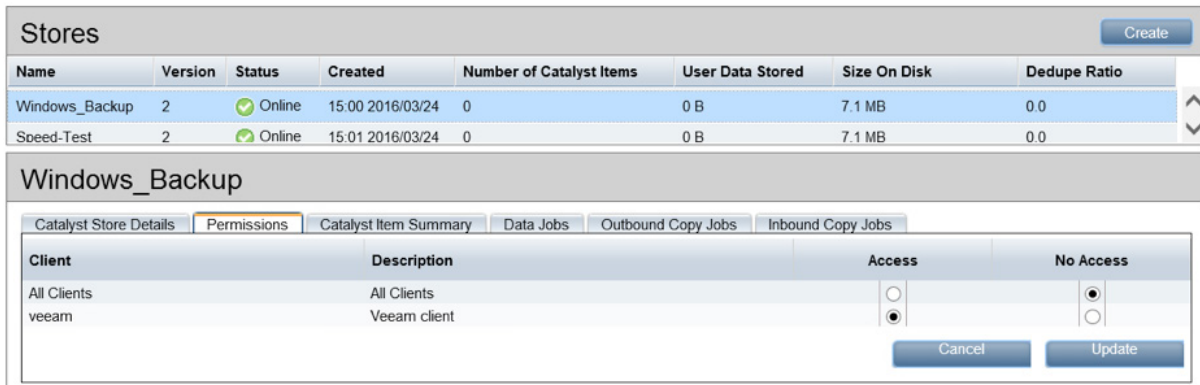


Figure 6: The permissions settings for the 'Windows Backup' store showing access restricted to the client "Veeam."

## How to connect to a StoreOnce Catalyst store via IP and Fibre Channel

HPE StoreOnce Catalyst provides source side deduplication and thus reduces the network load during backup and copy processes. 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.

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

- **IP/Ethernet:** StoreOnce Catalyst over Ethernet (CoE) is the most commonly used option and it is available with minimal configuration. The user just needs to make sure the Ethernet infrastructure has enough bandwidth to support the backup load, and that the Veeam Gateway server has access to this network with an appropriate IP address.
- **Fibre Channel (FC):** StoreOnce Catalyst over Fibre Channel (CoFC) functions the same way as StoreOnce CoE. However, some additional configuration is required to set up the backup and restore connections between the ports on the HPE StoreOnce Backup system and the ports on the client servers. This is done using the Fibre Channel Settings tab of the StoreOnce GUI. This GUI page is only available if CoFC is enabled on the StoreOnce appliance.

## Prerequisites for using Catalyst over Fibre Channel

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

**Note:** When using StoreOnce Catalyst over Fibre Channel (CoFC), administrator privileges are required to access OS-specific device files associated with StoreOnce CoFC devices. Fibre Channel users are advised to consult the StoreOnce support matrix (available at [www.hpe.com/storage/DAPRCompatibility](http://www.hpe.com/storage/DAPRCompatibility)) to ensure that Veeam server host bus adapters (HBAs), switches, Fibre Channel driver and firmware versions are supported.

## Initial Fibre Channel configuration

In the StoreOnce GUI, navigate to HPE StoreOnce>StoreOnce Catalyst>Fibre Channel Settings tab and locate the identifier at the top of the window as shown below. This is the Fibre Channel address of the StoreOnce that is used to identify the StoreOnce Backup system in Veeam. It is in the format COFC-<device-id> which will be inputted into Veeam where an Ethernet address would otherwise be provided.

The screenshot displays the 'Fibre Channel Settings' configuration page. At the top, there are navigation tabs: Status, Settings, Clients, Blackout Windows, Bandwidth Limiting Windows, and Fibre Channel Settings. The main content area is divided into sections:

- Identifier:** COFC-4p3ca82a1508001
- Identifier Alias:** COFC-4p3ca82a1508001
- Ports:** A section for configuring target ports.
- Target Ports:** A table listing four ports with their locations, status, speed, topology, and a checkbox for enabling them.
- Devices:** A section for configuring target devices.
- Target Devices:** A table listing four devices with their names, status, FC addresses, World Wide Node Names, World Wide Port Names, Number of Logins, and Number of Devices per Login.

Figure 7: The Catalyst Fibre Channel setting tab shows Fibre Channel settings and the number of target ports and FC devices.

- HPE recommends the default values in the target ports section. However, the speed can be edited.
- Locate the world-wide port name for each port on the StoreOnce Backup system. This is the information needed to zone the Veeam Gateway server Fibre Channel ports with the StoreOnce Fibre Channel ports.
- Locate the “Number of Logins” and “Number of Devices per Login” for each port. These values determine the number of concurrent backup and restore connections allowed on each Fibre Channel port on the HPE StoreOnce Backup system.
  - Number of logins defines the number of client-side ports that are zoned to connect to that port on the StoreOnce system and cannot be edited.
  - Devices per initiator port determines the number of backup and restore connections that are allowed from each Veeam Gateway server Fibre Channel port to the ports on the StoreOnce Backup system. This value should be increased to the maximum number of possible connections available in the GUI. If this value is increased, a device-file rescan on the Veeam server is needed before the change is recognized.

**Note:** *The number of possible paths between a Veeam Server Fibre Channel port and a StoreOnce Fibre Channel port is calculated as number of client ports zoned \* number of StoreOnce node ports zoned \* devices per initiator port.*

**Table:** Example relationships between device per initiator port and number of data sessions.

Veeam server	HPE StoreOnce	Number of devices per login	Maximum possible FC connections between Veeam Server FC ports and StoreOnce FC ports
1 port zoned to	4 ports	1	4
2 ports zoned to	4 ports	1	8
2 ports zoned to	2 ports	4	16
2 ports zoned to	4 ports	8	64

## Configuring of the Veeam Gateway server for Fibre Channel

Administrator permissions are required to run CoFC backups. 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 select "Scan for hardware changes" to detect the new devices.

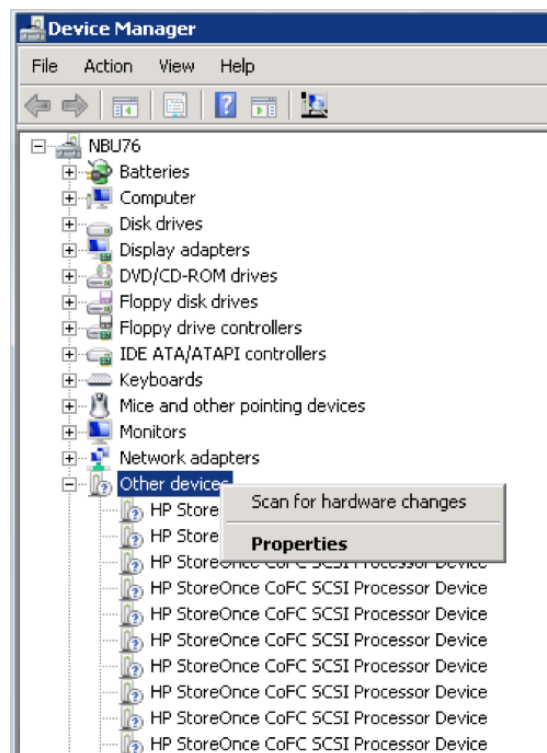


Figure 8: StoreOnce CoFC presents a device type of Processor and is shown under Other Devices in Windows Device Manager.

## Creating a StoreOnce Catalyst-based Veeam Backup Repository

Backup repositories are simply defined as storage for Veeam backup files created by Veeam Backup & Replication. A 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 a StoreOnce system. To create a Catalyst-enabled Veeam backup repository use the following steps:

- Click on "Backup Infrastructure" in Veeam Backup & Replication
- Right-click on "Backup Repositories" and choose "Add Backup Repository"

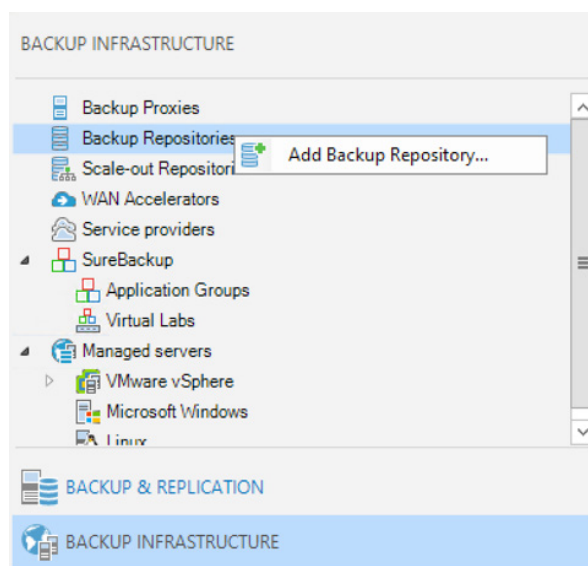


Figure 9: Veeam UI showing the creation of a new Backup Repository.

- **Repository name** – Give the repository a friendly name

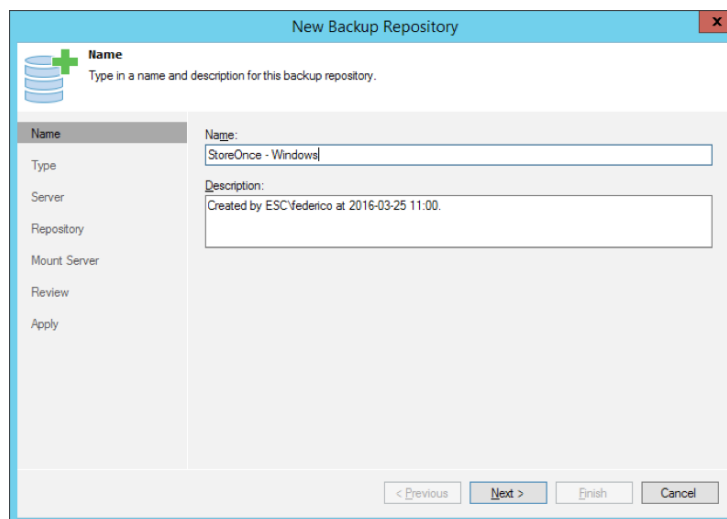


Figure 10: Defining the name of the backup repository.

- **Repository type** – Choose "Deduplicating storage appliance"

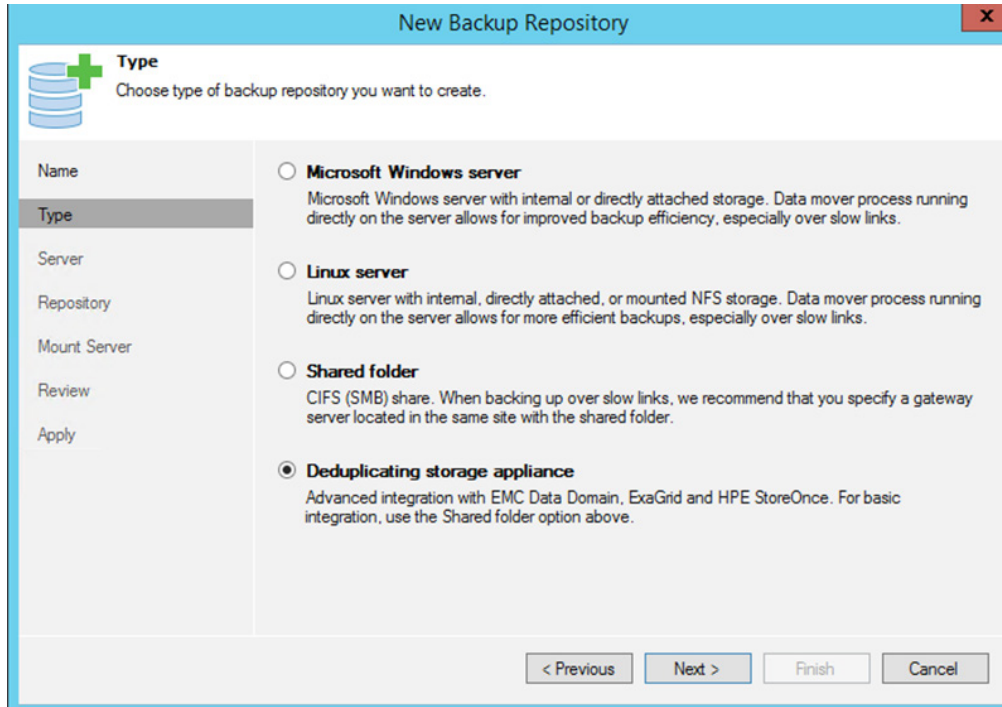


Figure 11: Specifying "Deduplicating storage appliance" as the type of backup repository.

- **Deduplicating storage type** – Choose "HPE StoreOnce"

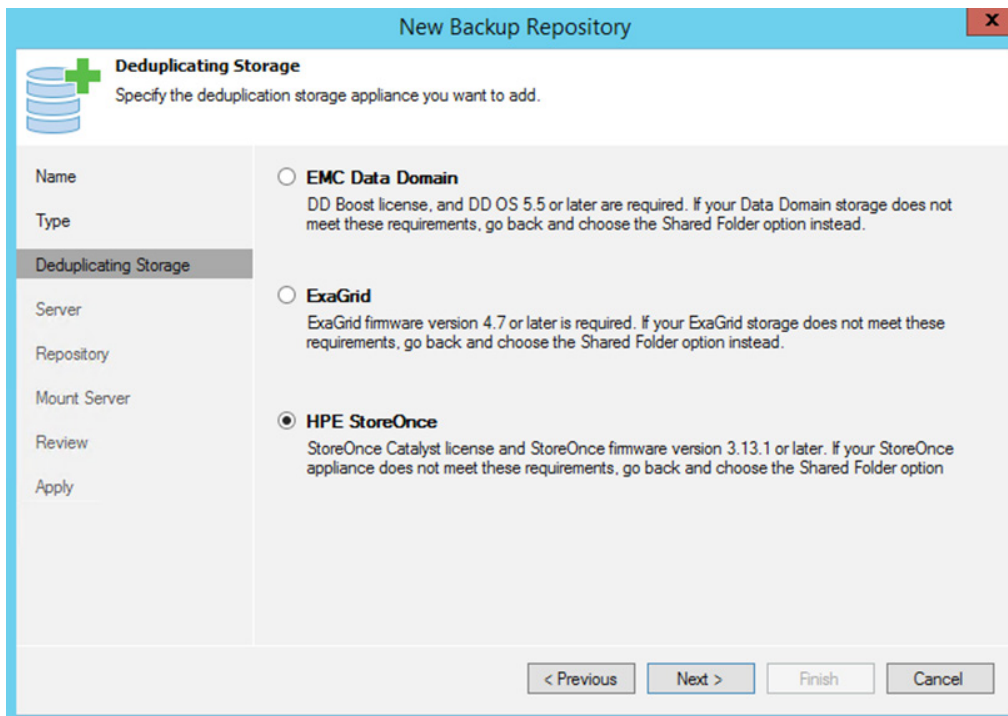


Figure 12: Specifying "HPE StoreOnce" as the type of backup repository being configured

- **Server name** – Enter the DNS name, IP address or choose Fibre Channel connectivity for the StoreOnce Appliance, along with credentials to access the Catalyst Stores. If “Client Access Permission Checking” has been enabled, those credentials will be used here. Otherwise, any user or admin credentials can be used. New credentials can be added by choosing “Add” and inputting the appropriate credentials that will be stored in the Veeam credentials database for future use

**Gateway server** –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 StoreOnce, where Veeam does not install a data mover on the StoreOnce appliance, Veeam uses a gateway server to host this service. The gateway server sits in front of the StoreOnce appliance 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 added to Veeam Backup & Replication that has access to the 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 in a way that gives it access to the StoreOnce.

**Note:** *Whenever possible it is suggested configuring both proxy and gateway roles to run on the same server. This eliminates an extra hop for backup traffic and reduces backup proxy CPU utilization.*

**Note:** *When using virtualized Veeam components in vSphere, ensure that the vNIC used for data traffic is set to VMXNET3 rather than the default E1000. This is particularly important for 10 GbE LAN.*

WAN settings (version 9.5 and newer) – Catalyst has been designed to work over LAN as well as WAN connections. When the communication between the Veeam Gateway and StoreOnce is across a WAN link, we recommend 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 the additional workload, it is recommended to not activate this functionality for on-site Catalyst repositories.

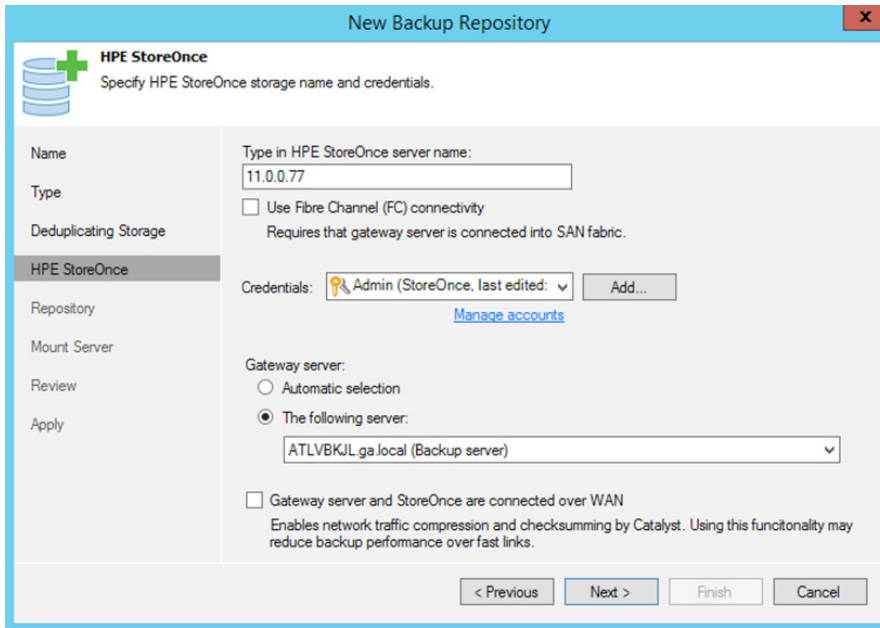


Figure 13: Specifying StoreOnce name/IP, credentials and Gateway Server options.

- **Catalyst Store** – Select the Catalyst Store to associate to this Veeam repository. The same Catalyst Store can be associated to more than one backup repository

**Load control** – It is possible to limit the system resources that Veeam uses for reading and writing to this backup repository. Veeam Backup & Replication has a built-in scheduler to ensure that the active data sessions do not exceed the StoreOnce appliance’s maximum number of data sessions, so this setting is not mandatory and can be left disabled.

**Note:** *Advanced settings should not be adjusted for StoreOnce-based backup repositories.*

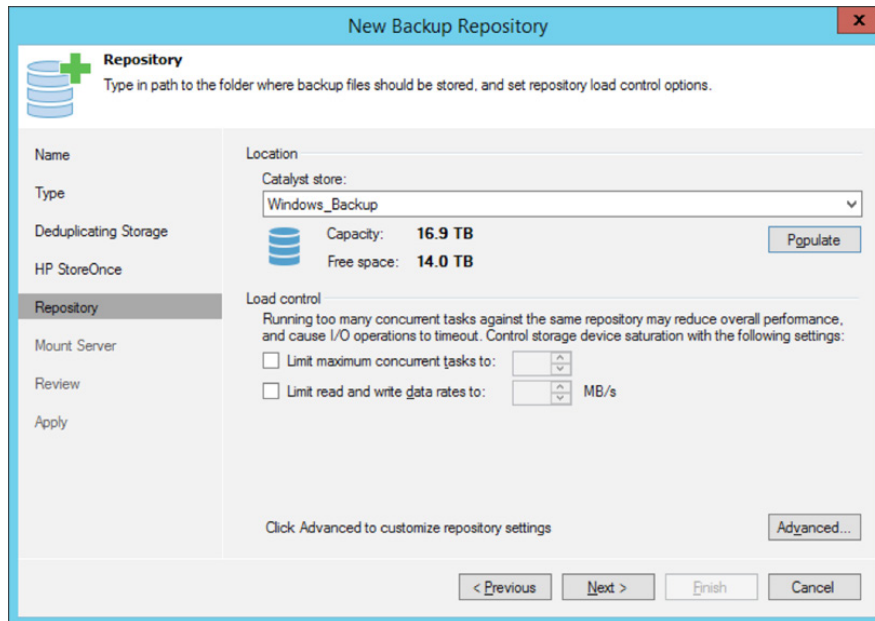


Figure 14: New backup repository wizard specifying Catalyst store and load control settings.

- **Mount server** – In the case of 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
- **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 on this page by default. It is best practice to have at least 15% of the largest VM(s) that could be recovered available at the specified path. For example, a 100 GB VM that may be restored would need 15 GB of free space at this location. It is also possible to write new data during a restore to a production datastore which can be chosen in the restore job itself

**Note:** In cases where the Gateway server selection is set to automatic, the Mount Server will also handle the functions of synthetic full backups and running backup copy agents.

**Note:** Veeam Backup & Replication Enterprise or Enterprise Plus 9.5 and StoreOnce v3.15.1 are required for vPower NFS support.

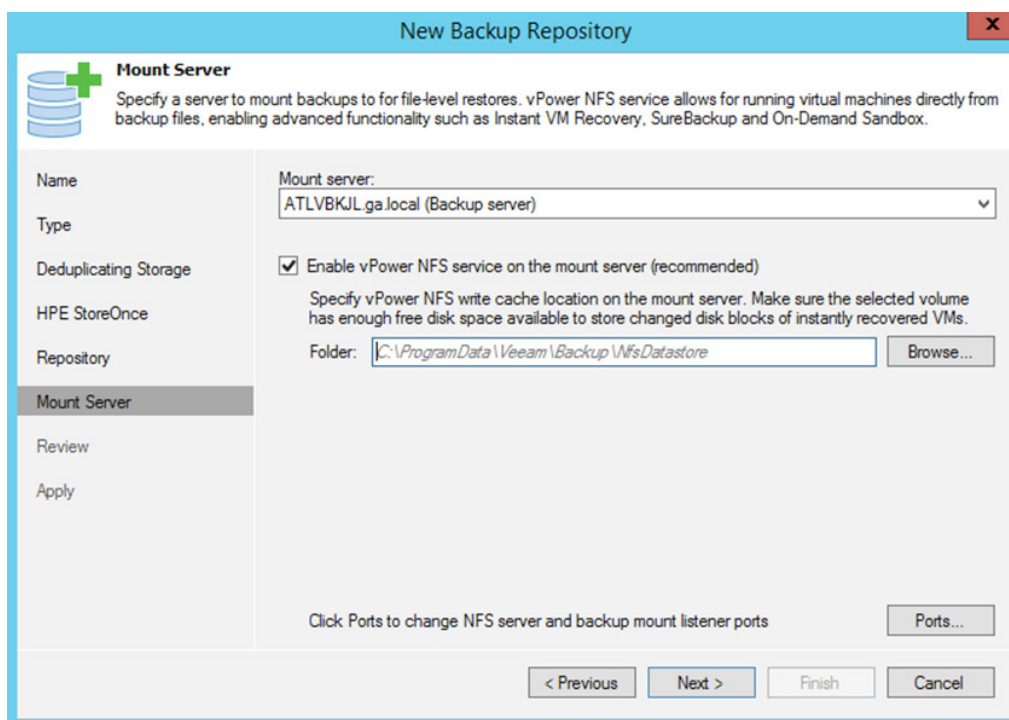


Figure 15: Specifying a mount server for restore operations in the New Backup Repository Wizard.

- Review all settings and click finish

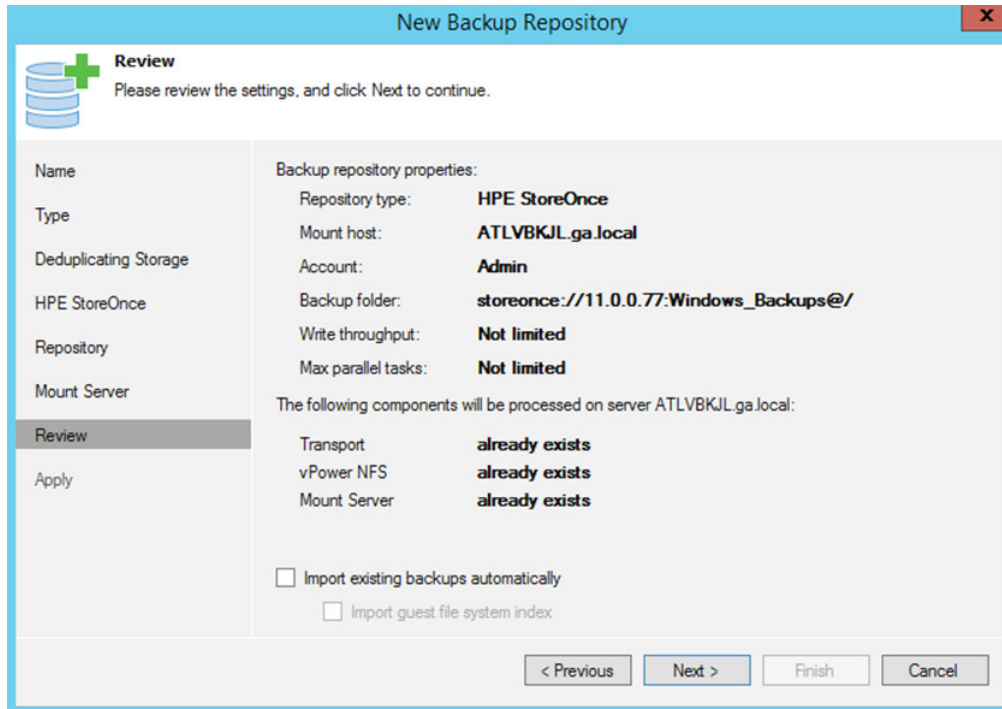


Figure 16: The summary page in the New Backup Repository Wizard.

## Creating a Veeam backup job

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 and application quiescence, in addition to policies for backup file handling. This section will review the recommended settings when creating a backup job writing to a StoreOnce Catalyst-based backup repository.

- Click on "Backup & Replication" in the left-hand menu bar
- Using the ribbon at the top of the user interface, choose "backup job"




Figure 17: Launching the New Backup Job wizard from the Veeam GUI.

- Give the job a friendly name
- Choose "Add" and select the VM(s) to add to this backup job

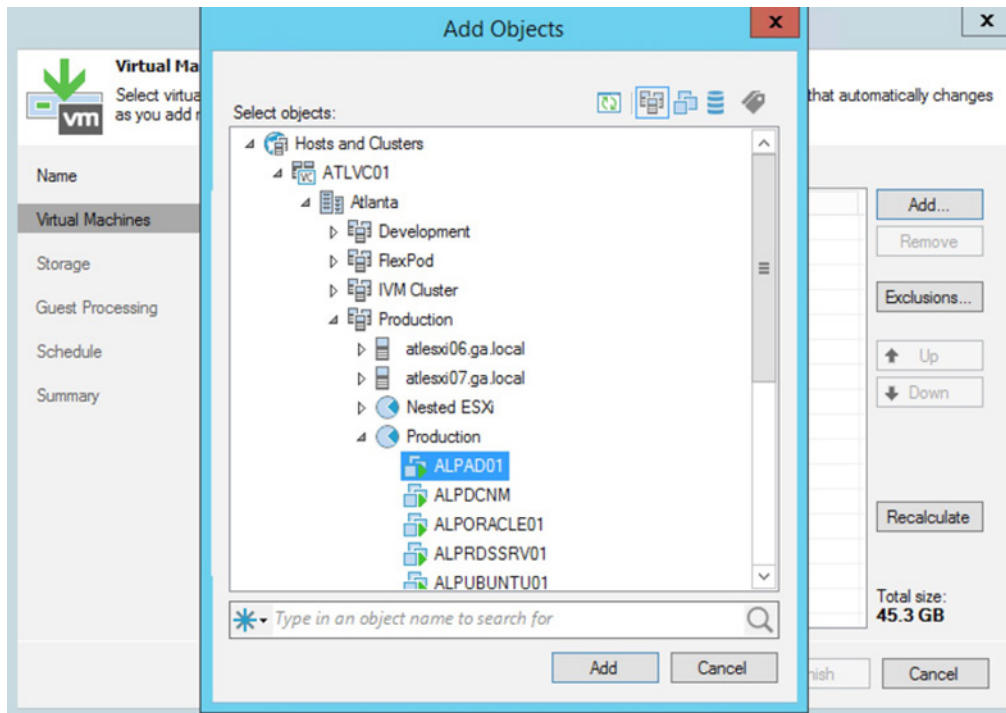


Figure 18: Launching the New Backup Job wizard from the Veeam GUI.

**Note:** If a hierarchy object such as "Resource Pool" or "Cluster" is chosen, Veeam will backup 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 removed from under the object in the infrastructure.

- **Backup proxy** – Veeam's intelligent load balancing between multiple backup proxies can be leveraged by leaving this as "Automatic." In the case of HPE StoreOnce, it is recommended that the backup proxy be the same as the gateway server whenever possible. This will reduce the number of "hops" in the data path and can reduce CPU utilization on the backup proxy
- **Backup repository** – Choose the StoreOnce Catalyst-based backup repository created in previous steps
- **Retention policy** – The number of restore points to keep on disk defines the retention for the VM backups in this job. To equate this to time on disk, schedule the job accordingly. For example, a daily backup schedule would mean 14 daily restore points as shown in Figure 19

**Note:** StoreOnce deduplication highly reduces the storage requirement for additional restore points. Usually a retention of 4 weeks does not require twice the capacity of two weeks, but only 30% more.

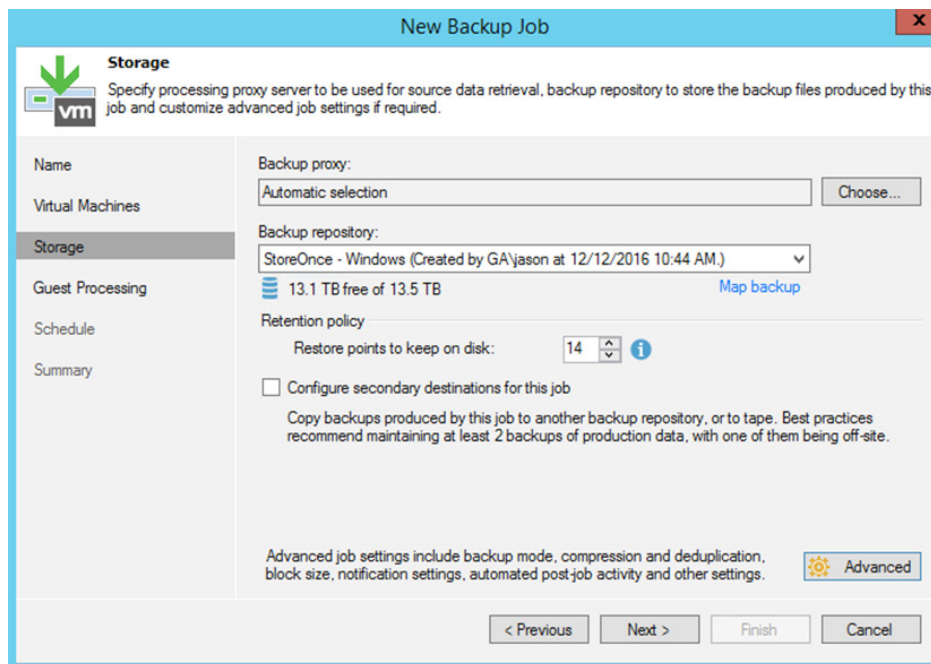


Figure 19: Specifying proxy, repository and retention parameters in the New Backup Job wizard.

- **Advanced settings** – For StoreOnce Catalyst-based backup repositories, the following settings are recommended. If different settings are selected, a notification will appear with an option to change them automatically.

## Backup Tab

### 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 new full backup. Veeam leverages the Catalyst API to offload (virtualize) this process from the Veeam servers to StoreOnce, 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 the workload on production storage and servers
- Veeam sends to StoreOnce the “instructions” for merging the previous full with all the subsequent Incremental backups to finally generate the new synthetic full
- 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.

Be aware that in Veeam Backup & Replication v9 Update 2 and 9.5, Veeam will open a new data session to the StoreOnce for each backup file within a backup chain that is being synthesized. For example, if a backup job has 10 virtual machines, each with seven restore points, Veeam will open up to 90 data sessions to StoreOnce (assuming a weekly synthetic). This could potentially utilize all available StoreOnce data session forcing other jobs to wait until the synthetic full frees up sessions. Veeam has a built-in scheduler to manage this process and to ensure that a StoreOnce's maximum data sessions aren't exceeded. A breakdown of the data sessions available on specific StoreOnce appliances can be found in the *Appendix* of this document.

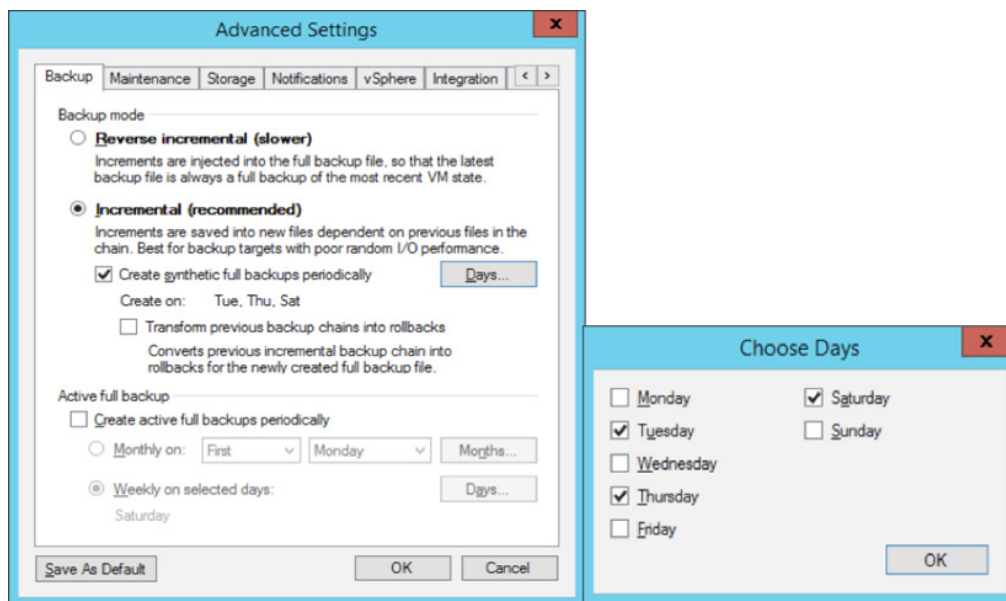


Figure 20: Specifying backup mode and synthetic full backups in the New Backup Job wizard.

**Note:** It is required that either a synthetic or active full backup be selected when backing up to HPE StoreOnce over Catalyst.

## Storage Tab

**Data deduplication** — Disabled

- **Compression level** — *Optimal*  
\*Be sure the default setting "decompress before writing" is selected in the Backup Repository advanced settings
- **Storage Optimization** — *Local target (16 TB+ backup files)*
- **Encryption** — *Disabled*

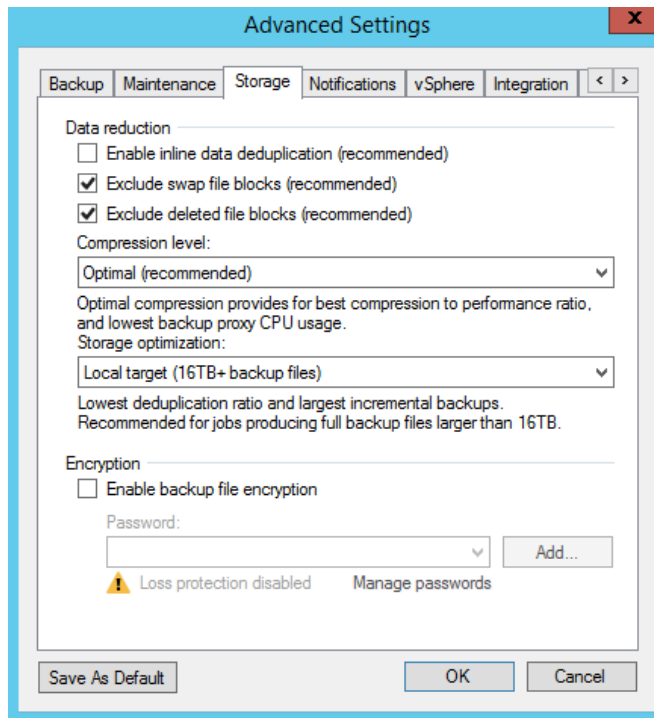


Figure 21: Specifying dedupe, compression, optimization and encryption options.

**Note:** If data-at-rest encryption is required, it is recommended to use StoreOnce embedded encryption rather than Veeam encryption because Veeam encryption will reduce StoreOnce data deduplication. To enable StoreOnce encryption, set "Store Encryption Enabled" at the time of creating the Catalyst store, as shown in Figure 22. Note StoreOnce encryption requires an additional license on all StoreOnce models except the StoreOnce VSA.

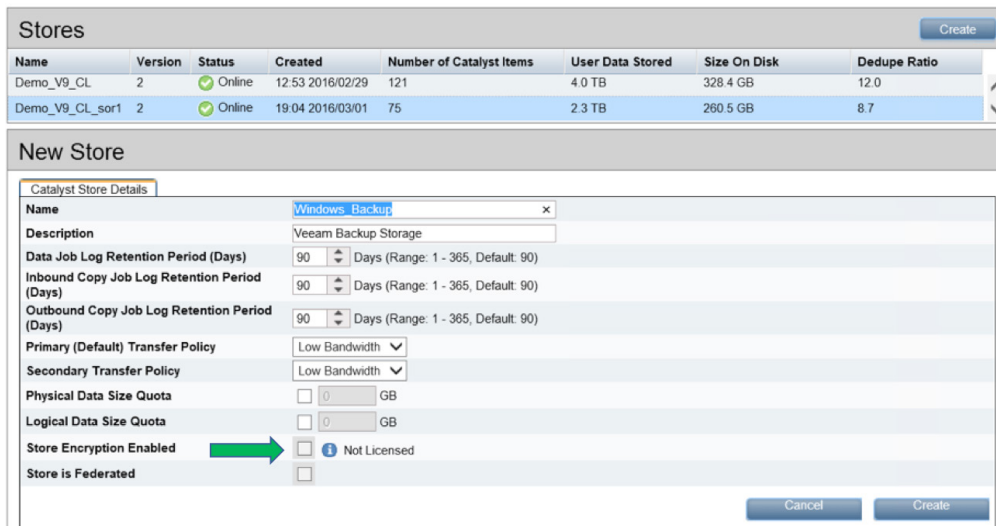


Figure 22: Enabling encryption in the StoreOnce Catalyst store GUI.

- **Guest Processing** – Guest processing enables transaction consistency for VM backups, indexing of VM guest file systems for quick searching through the optional Veeam Enterprise Manager, and the configuration of how logs of specific applications are handled. To enable these capabilities simply check the appropriate box and supply administrator credentials. Additional details on configuration of these features can be found in the [Veeam Backup & Replication user guide](#).

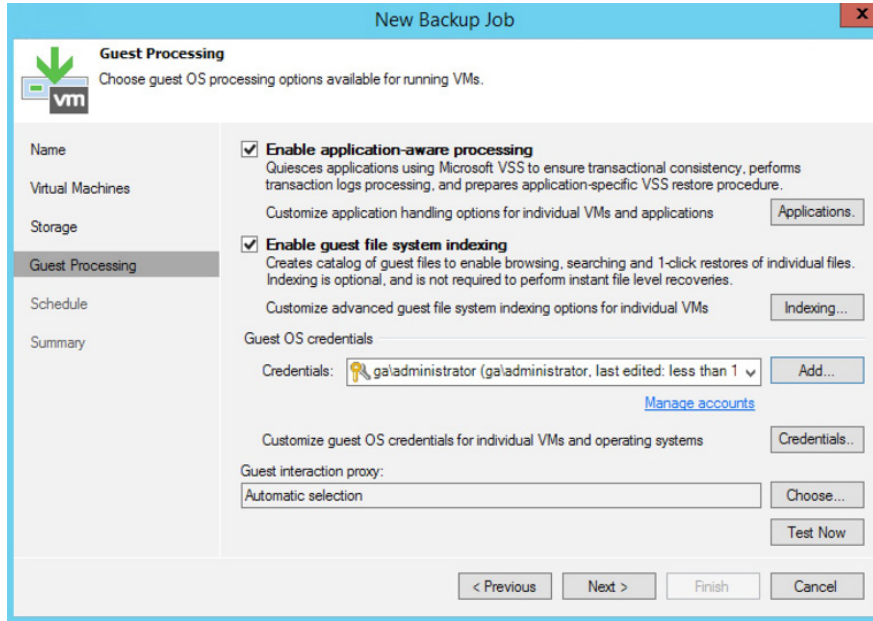


Figure 23: Setting Guest Processing and file indexing options in the New Backup Job wizard.

- **Job scheduling** – Usually backup jobs are run daily at night time. 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 needs to be maintained based on the StoreOnce model in use, then there will need to be a synthetic full twice per week.

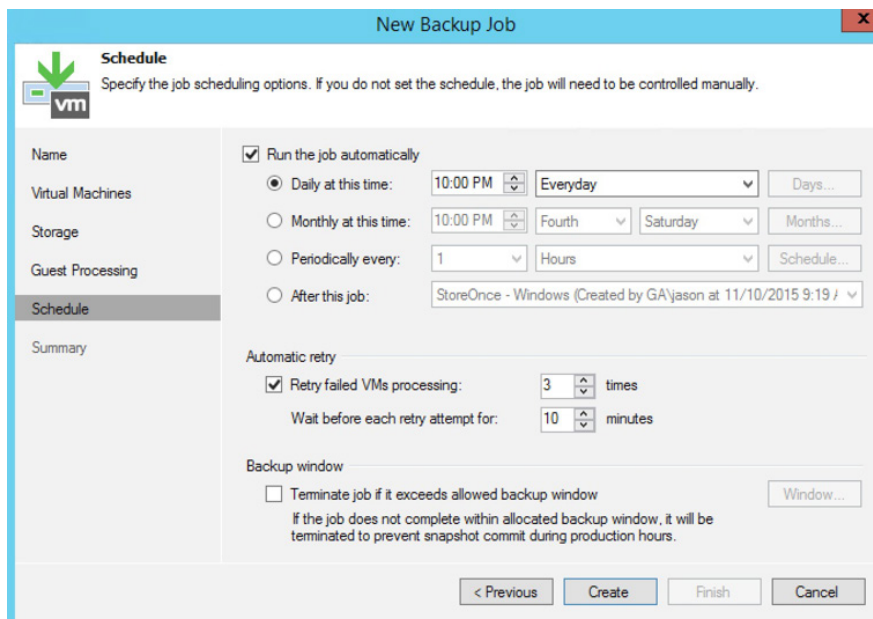


Figure 24: Set the job schedule to determine the RPO for the backup job.

## Copying Veeam backups to a remote site

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 StoreOnce appliances. Below we describe two architectures available to execute backup copies; both are based on the Veeam backup copy job.

### Architecture "A" – Site to site through source and target Veeam gateway servers with local StoreOnce Catalyst stores

#### Advantages:

1. Supports the Veeam health-check and heal functionality for backup copies
2. Facilitates the most efficient restores from and to the target site
3. Supports Veeam WAN acceleration

#### Disadvantages:

- Does not use StoreOnce Catalyst source-side deduplication for copy jobs
- Optional WAN accelerators are limited to a single stream, potentially becoming a throughput bottleneck

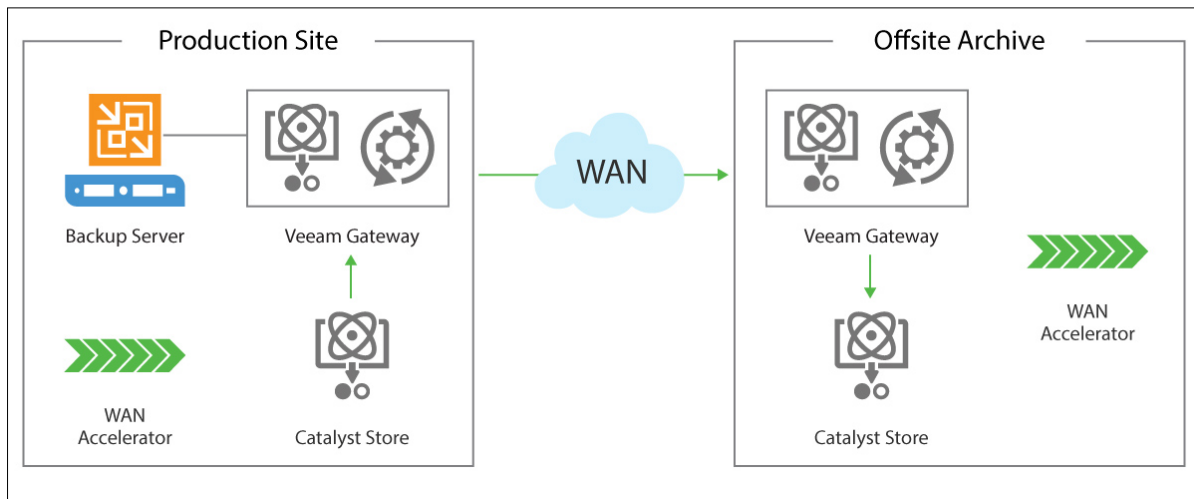


Figure 25: Backup Copy Job 'Architecture A' using Veeam WAN Accelerators to optimize the WAN utilization.

## Architecture "B" – Site to site using only source gateway servers with remote Catalyst stores (suggested for bandwidth reduction and performance)

### Advantages:

- Leverages StoreOnce Catalyst source-side deduplication which is reduce WAN bandwidth utilization
- Multiple concurrent streams within the same backup copy job
- Supports heavy workloads (1,000 Mb/s and higher) over high bandwidth WANs
- Does not require additional resources for running Veeam WAN accelerators

### Disadvantages:

- Veeam health-check and heal functionality for backup files is not recommended
- An additional Veeam gateway server at the remote site may be required to achieve the desired restore performance

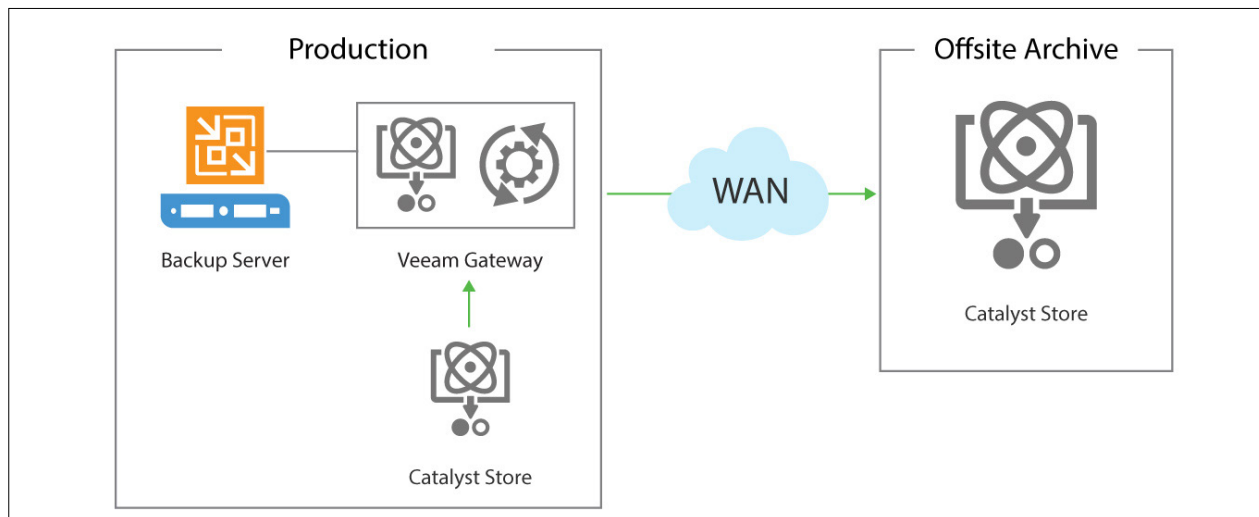


Figure 26: Backup Copy Job 'Architecture B' using StoreOnce Catalyst based deduplication and compression to highly optimize the WAN utilization.

**Note:** StoreOnce implements health-check processes for data integrity. Details can be found [here](#).

## Implementation steps for copying Veeam backups to a remote site

Veeam backup copy jobs are designed to copy backups of single VMs, or entire backup jobs with many VMs, to a remote location with simple or tiered GFS (Grandfather-Father-Son) retention scheme.

### Architecture "A" configuration

In this section, we are going to show the configuration of the architecture "A", where there is a Veeam gateway server at both locations.

- **Create a Catalyst Store on the remote StoreOnce appliance** — Create this Catalyst store as described in the "Creating a Catalyst Store for Veeam Backups" section above and shown in Figure 27.

Veeam_Backup_Copy					
Catalyst Store Details	Permissions	Catalyst Item Summary	Data Jobs	Outbound Copy Jobs	Inbound Copy Jobs
Name	Veeam_Backup_Copy				
Description	Offsite Copy of Backups				
Data Job Log Retention Period (Days)	90				
Inbound Copy Job Log Retention Period (Days)	90				
Outbound Copy Job Log Retention Period (Days)	90				
Primary (Default) Transfer Policy	Low Bandwidth				
Secondary Transfer Policy	Low Bandwidth				
Physical Data Size Quota	(No Limit)				
Logical Data Size Quota	(No Limit)				
Store Encryption Enabled	<input type="checkbox"/>				
Secure Erase Mode	None				
Store is Federated	<input type="checkbox"/>				

Figure 27: Creating a Catalyst Store called 'Veeam\_Backup\_Copy' using the StoreOnce GUI with a Low Bandwidth transfer policy.

- **Add a new Windows Server** — This will be the remote gateway server, so ensure that it has access to the Catalyst store created in the previous step.
  - Click on "Backup Infrastructure" in the Veeam Backup Console.
  - Right-click on "Managed Servers" and choose "Add Server," then choose "Microsoft Windows."

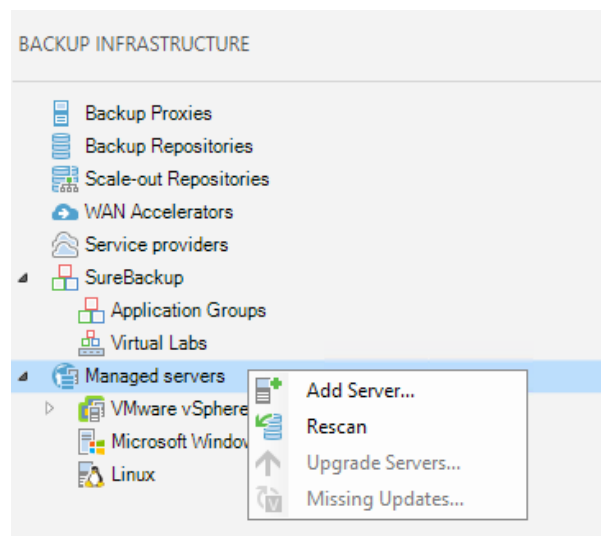


Figure 28: Adding a new server to act as a remote Gateway Server for StoreOnce.

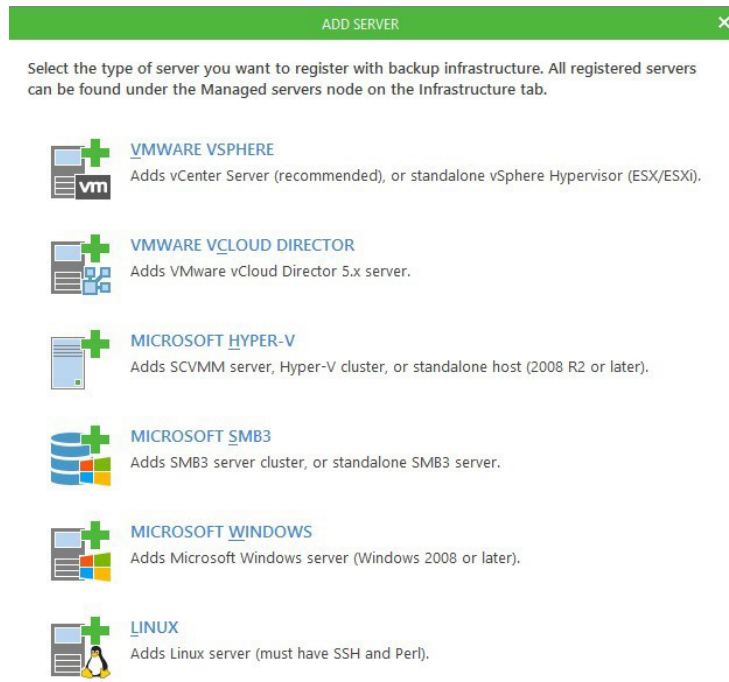


Figure 29: Selecting the Microsoft Windows option for the new Veeam Gateway server.

- Provide the server DNS name or IP address of the server selected as the gateway.




Figure 30: Setting IP or DNS name of the new Windows Server acting as a Veeam Gateway server.

- Provide local administrator credentials for the new server.

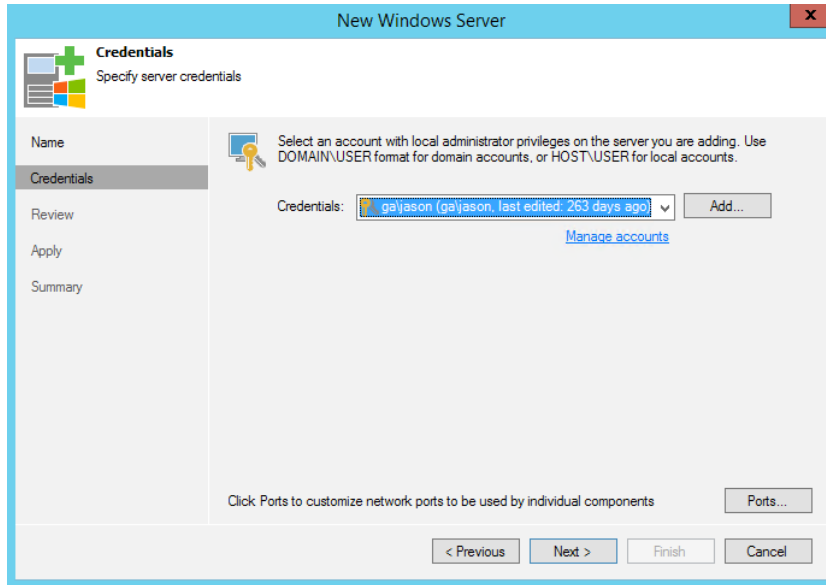


Figure 31: Providing domain or local administrator credentials.

- **Create a new backup repository**

Follow all the steps described in the “Creating a StoreOnce Catalyst-based Veeam Backup Repository” section above for creating a new repository. For the “Gateway server” choose the newly created Windows Server located at the remote site. This step is important because it defines that the backup- repository is accessible only by the specified gateway, thus every backup copy job targeting this backup repository will send data to this server before reaching the StoreOnce over Catalyst.

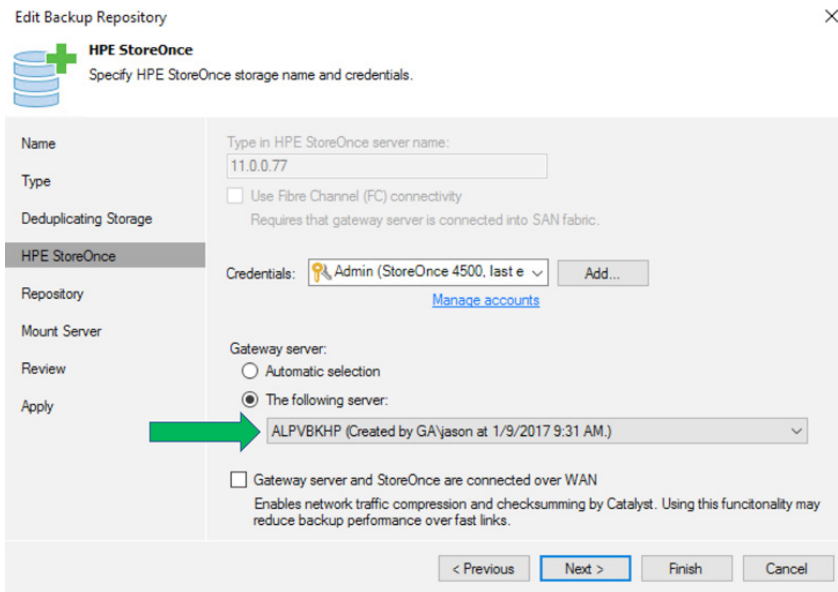


Figure 32: Specifying the remote Gateway server that will write to the remote StoreOnce over Catalyst.

- **Creating Veeam backup copy job**

In the Veeam GUI under “Backup & Replication” click on the “Backup Copy” toolbar button and choose the appropriate hypervisor for the VMs within the primary backup file.

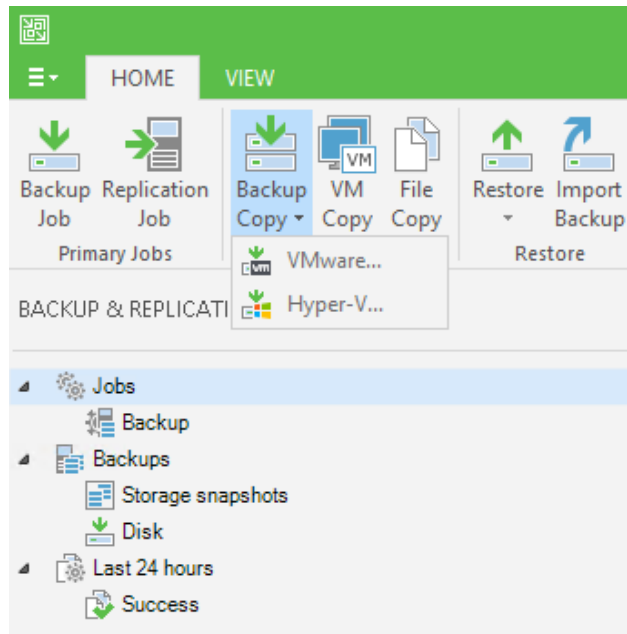


Figure 33: Creating a Veeam Backup Job.

- Give the backup copy job a name and choose the interval in which backup copies should be made – this can be per minute, per hour or per day. At the start of each copy interval, Veeam will look for new restore points on the primary backup repository to copy to the target location.

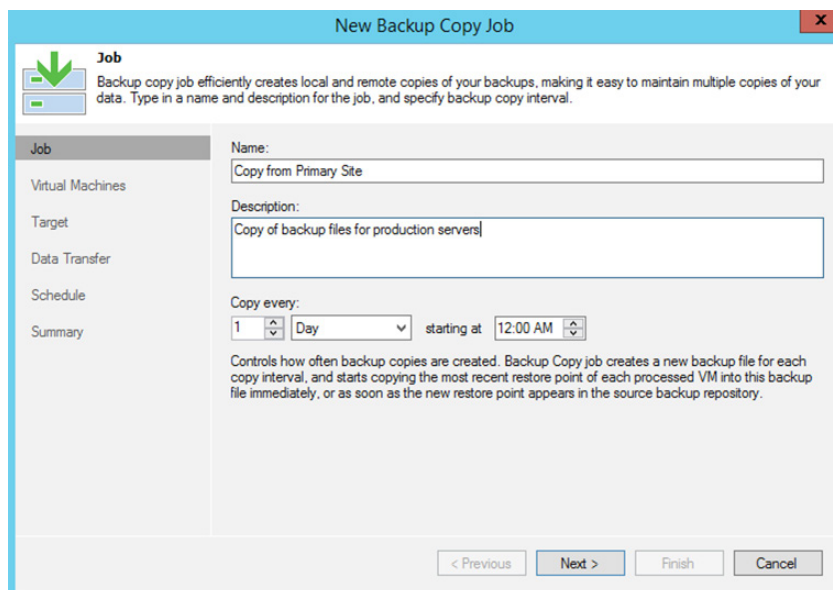


Figure 34: Defining a backup copy job name and copy interval.

- Choose the VM backups to be copied off site by selecting either backups, jobs or infrastructure. When choosing infrastructure or jobs, the copy job will wait until backups for the selected objects appear on the backup repository before transporting any data.

**Note:** The backup repository where the copies should be retrieved from can be found under Source.

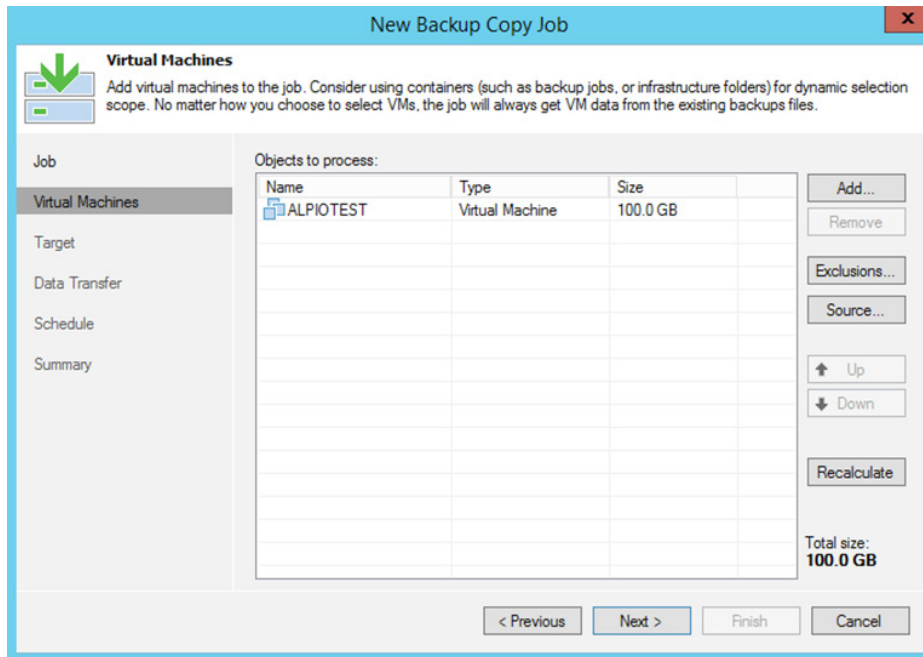


Figure 35: Selecting the source VM's or backups to be copied.

- **Target** – On this step, specify the Catalyst based target repository, the retention policies and the advanced parameters.
  - **Backup repository** – Choose the target StoreOnce Catalyst-based backup repository. Please note that every backup repository is dynamically or statically associated to a specific Veeam Gateway. The Veeam proxy that executes the backup copy job will read data from the source repository and will write data to the destination one communicating with the associated Gateway.
  - **Restore Points** – The number of restore points to keep in the primary retention chain. With Veeam 9.0 the maximum value is seven. With version 9.5, the maximum value depends on the StoreOnce model in use and can be as big as 42. For more details, refer to the table in the appendix and check the column "Max restore points per single backup chain." If a number higher than the maximum supported value is selected a warning message will appear.
  - **Keep the following restore points for archival purposes** – (Recommended). It is optional to choose a tiered retention schedule based on the GFS methodology. It is important to keep at least one weekly restore point. This setting forces the generation of a new full backup using the virtual synthetic full process. This operation reorganizes the full backup file reducing its capacity utilization.

For remote StoreOnce Catalyst repositories, leave the "Read the entire restore point from source backup" option un-checked.

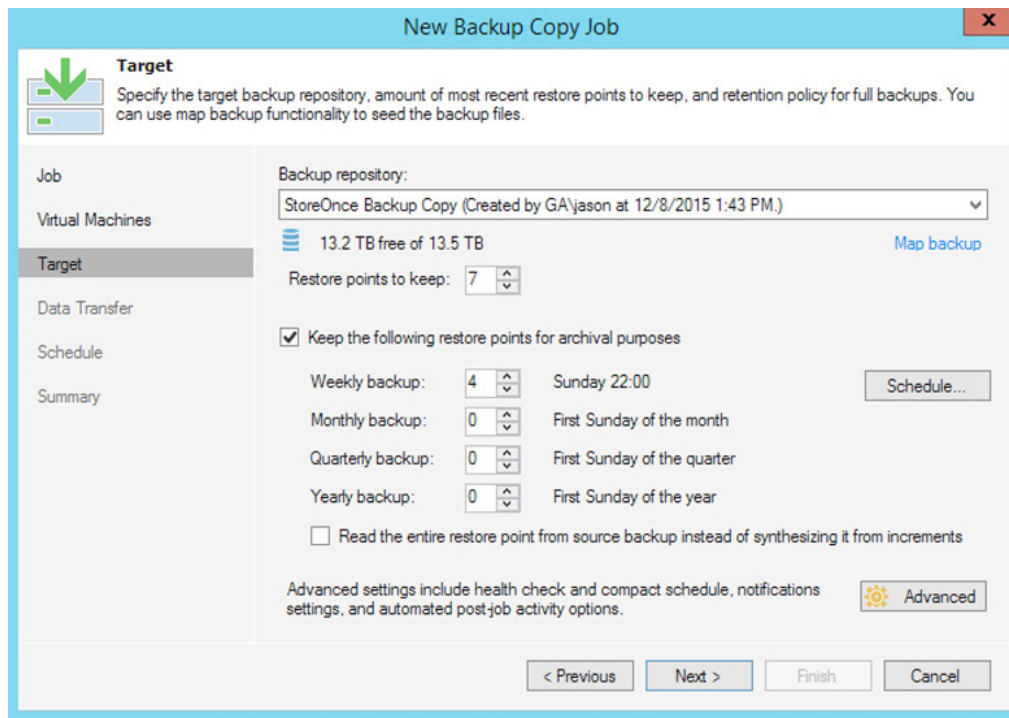


Figure 36: Defining the retention for the backup copy job.

- **Advanced** – Maintenance
  1. **Storage-level corruption guard** – Optional
  2. **Full backup file maintenance**
    - **Remove deleted VMs data** – Optional
    - **Defragment and compact full backup file** – Not supported
- **Advanced** – Storage
  1. **Data reduction**
    - **Deduplication** – Enabled
    - **Compression** – Auto

\*Be sure the default setting “decompress before writing” is selected in the Backup Repository advanced settings

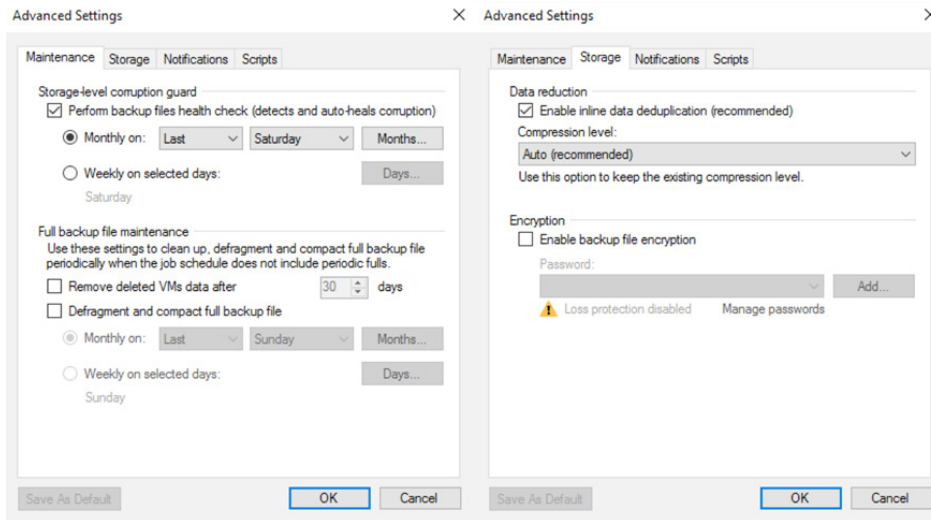


Figure 37: Adjusting backup copy job “advanced” settings for Architecture A.

- **Data Transfer** – Veeam Backup & Replication can send backup copies in a traditional direct copy over a WAN, or through WAN Accelerators.
  - **Direct** – Data is sent directly between the source and destination Veeam Gateway Servers, sending only unique incremental data based on Veeam’s tracking of changed data rather than Catalyst source-side deduplication.
  - **WAN accelerators** – Minimize Veeam backup copy traffic on WAN links by using source and target caches of data blocks already sent between the locations. Typically, this option is recommended on WAN connections of less than 100 Mb/s. A detailed overview of WAN accelerators and instructions for deployment can be found in the [Veeam Backup & Replication user guide](#).

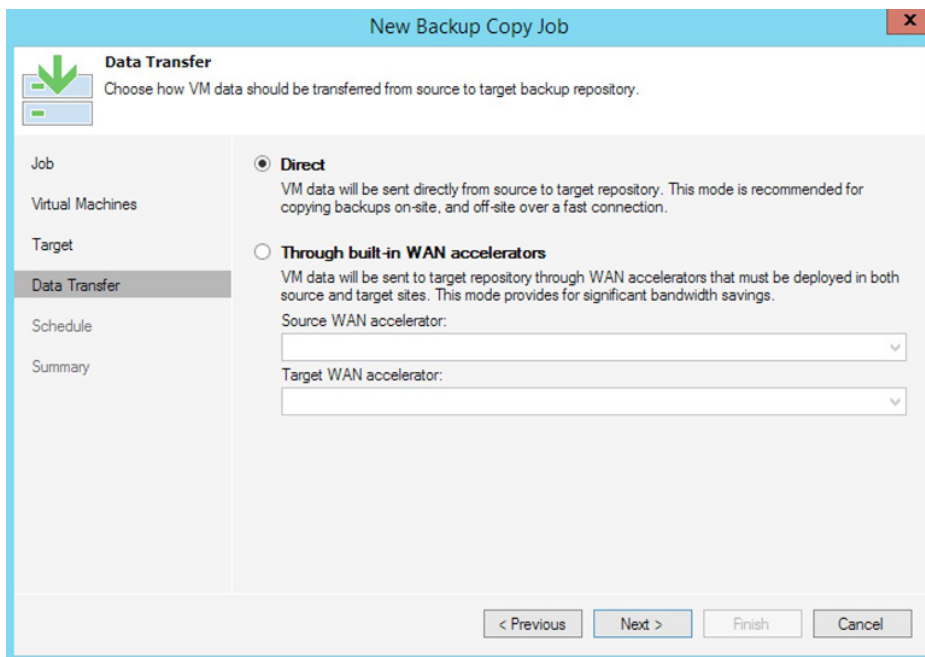


Figure 38: Specifying backup copy data transfer method.

- **Schedule** – By default, backup copy jobs are set to run continuously between copy intervals, waiting for new restore points that meet the interval's requirements. To manage network usage, data transfer can be paused during peak usage times, for example during hours of heavy infrastructure utilization, as shown in Figure 38.

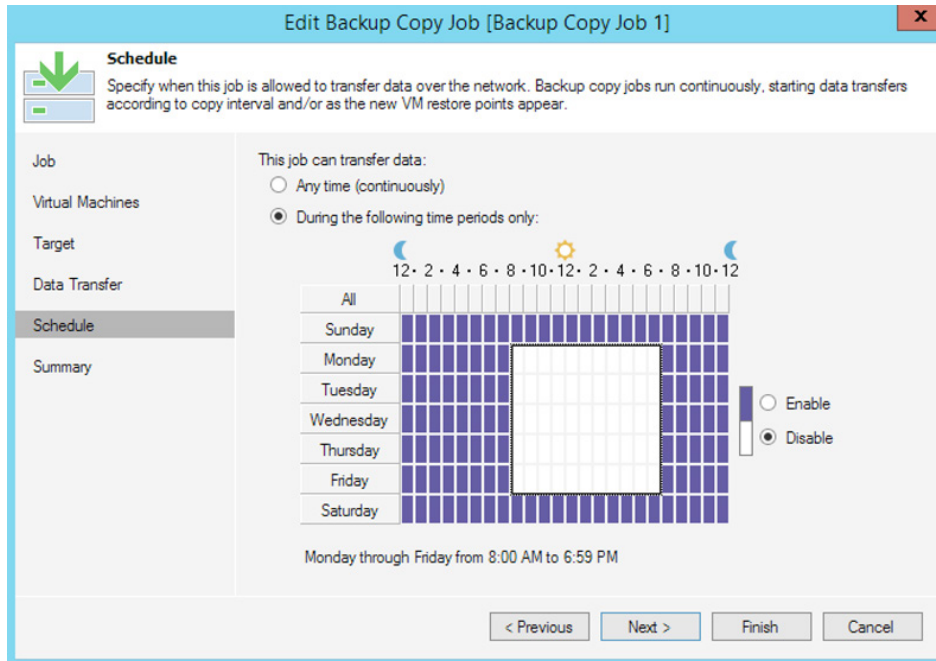


Figure 39: Specifying the backup copy schedule.

## Architecture "B" configuration

In this section, we show the configuration for using Veeam backup copy jobs between two locations with a Veeam gateway located on the source site writing directly over the WAN to the remote StoreOnce Catalyst-based backup repository. StoreOnce Catalyst makes source side deduplication and compression, and supports the high latency typical for WAN connectivity.

The advantage of this configuration comes from the high data reduction provided by Catalyst source-side deduplication and compression, which optimize the WAN utilization. This configuration does not require any additional hardware and supports a WAN connection of more than 1000Mb/s with multiple concurrent streams.

In comparison with "Architecture A", the communication over the WAN does not happen between a Veeam Gateway server at the source site and a Veeam Gateway server at the destination site, instead it happens between a Veeam Gateway server at the source site and the StoreOnce at the destination site.

- **Create a remote Catalyst Store**

Ensure that Low Bandwidth/Low Bandwidth transfer policies are selected to enable StoreOnce source-side deduplication.

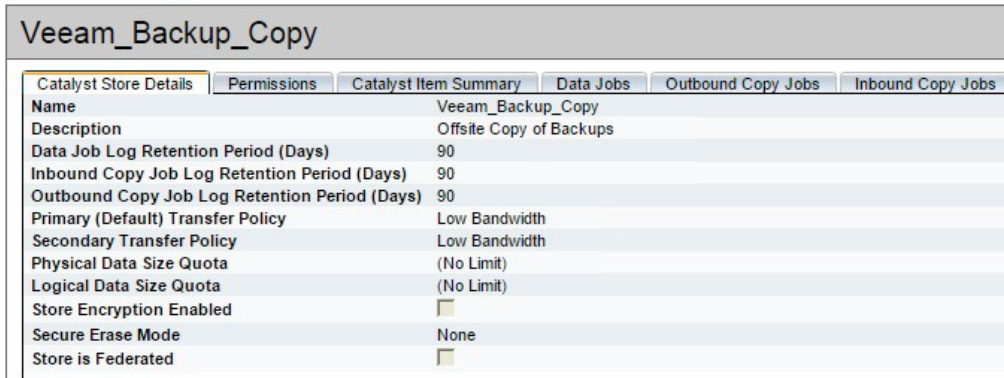


Figure 40: Creating a new remote Catalyst Store using the StoreOnce GUI.

- **Create a new backup repository**

Follow the steps described in the “Creating a StoreOnce Catalyst-based Veeam Backup Repository” section above to create a new backup repository. Choose a Gateway Server at the primary site and check the box labeled *Gateway server and StoreOnce are connected over WAN* which will enable additional compression and error correction using native StoreOnce checksum.

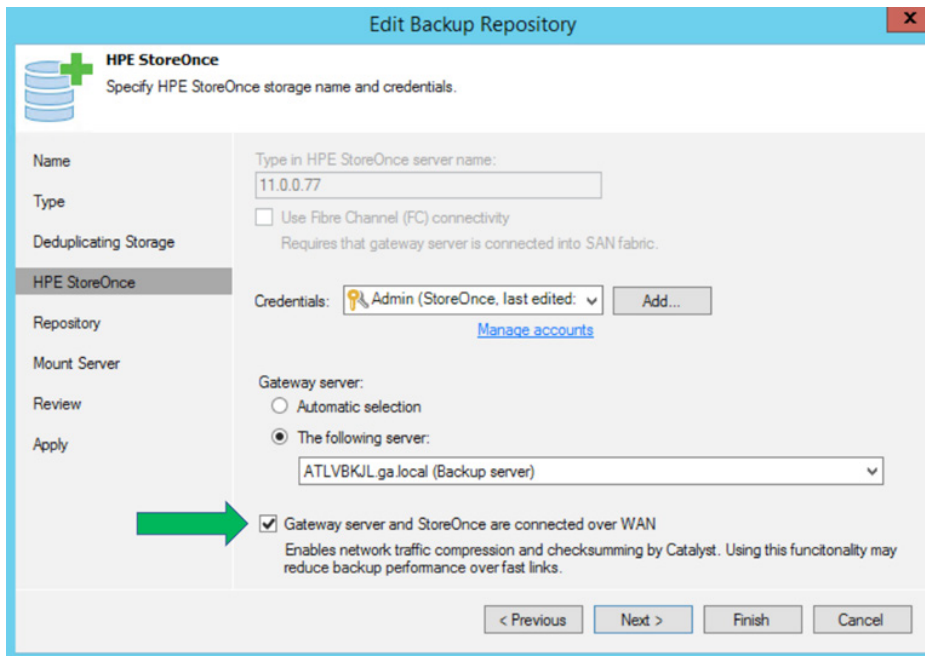


Figure 41: Enabling the “Gateway server and StoreOnce are connected over WAN” option for the remote backup repository.

- **Create a backup copy job**

Follow the same steps described in "Architecture A" except for the settings below.

- **Target** – In this step, specify the Catalyst-based target repository, the retention policies and the advanced parameters.
- **Advanced – Maintenance**
  - **Storage-level corruption guard** – Not recommended

**NOTE:** 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 may affect its performance.

- **Full backup file maintenance**
  - Remove deleted VMs data – Optional
  - Defragment and compact full backup file – Not supported
- **Advanced – Storage**
  - **Data Reduction**
    - Deduplication – Enabled
    - Compression – Auto

**NOTE:** Catalyst checksums data sent to prevent transmission errors. StoreOnce has additional data integrity algorithms to continuously verify its storage to prevent silent corruptions. In addition, RAID 6 prevents data loss even after two concurrent disk failures.

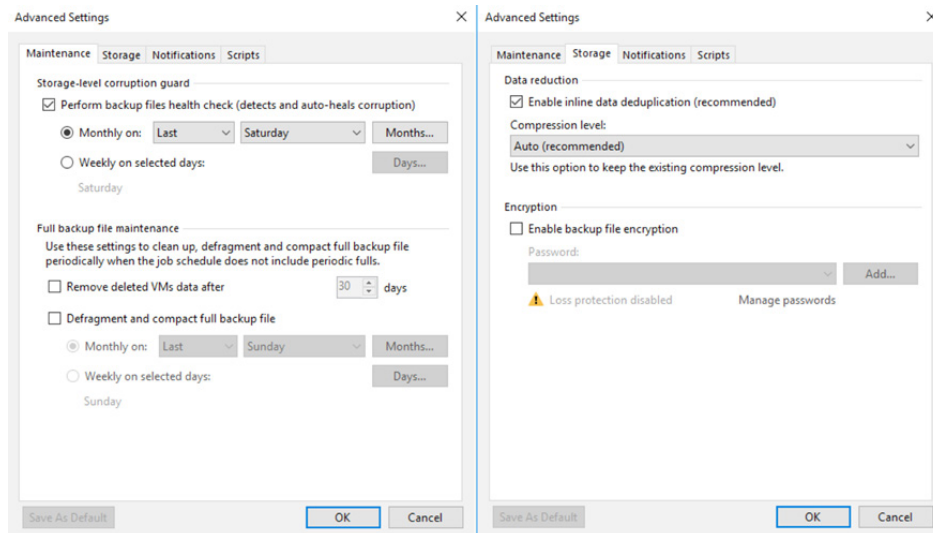


Figure 42: Adjusting backup copy job "advanced" settings for Architecture B.

- **Data Transfer** – Select “Direct” data transfer.

In this architecture, the “*Through built-in WAN accelerators*” is not supported and not required because the WAN is already optimized by Catalyst Source-side deduplication.

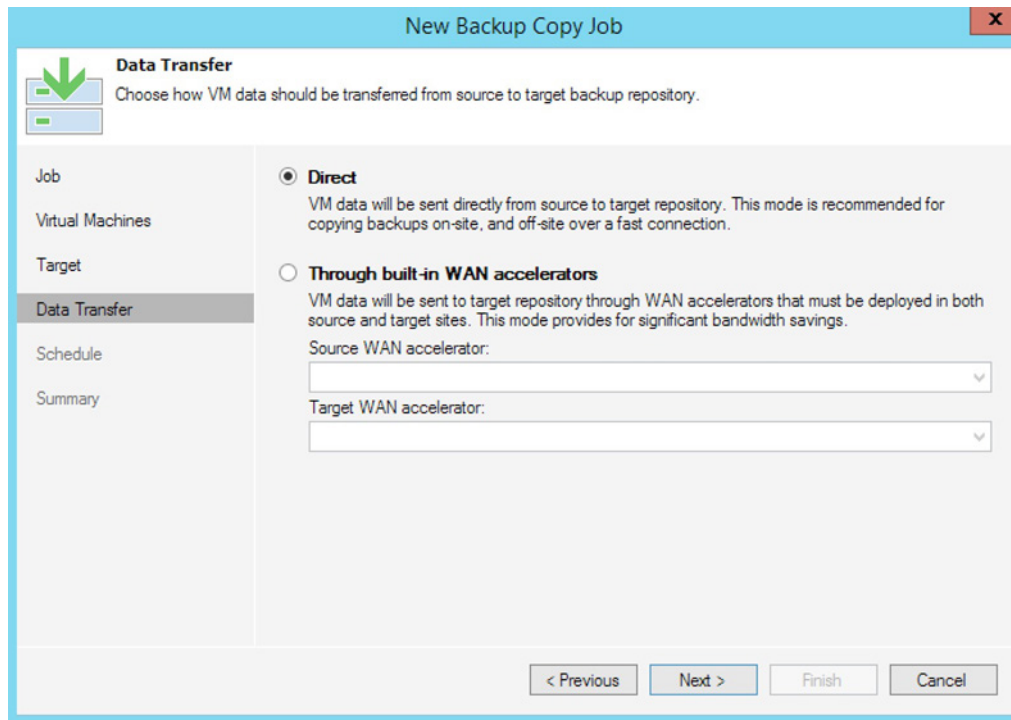


Figure 43: Choosing the direct transfer mode for the backup copy job.

## Architecture “B” restore configuration

Because there isn’t a gateway server at the remote location with the Architecture B and it is required that all data pass through a gateway server before restoring, it is recommended to add at least one standby gateway server that can be used on demand, specifically for this purpose. Many of the same steps described earlier in this guide apply for this configuration, but here is a summary:

1. **Add a Windows Server** – This should be located at the remote location
2. **Create a new Veeam Backup Repository pointing to the existing Catalyst Store**
  1. Use the off-site StoreOnce IP address, name or CoFC identifier and credentials
  2. Choose the same off-site Catalyst store as used in the setup of Architecture B for the backup repository
  3. Choose the newly added Windows Server as the gateway server and the mount server

Once added, this repository can be periodically rescanned and newly found backups will appear alongside “disk” and “disk (copy)” backups with a name of “disk (imported).” We recommend to restore from backups under *disk (imported)* as they will use the off-site gateway server and not the one on the primary site to retrieve restore data.

## Utilizing backups

Once primary and backup copy jobs have completed, backups for each job will be found under Backups>Disk for primary backups, Backups>Disk (copy) for backups created using backup copy jobs, or Backup>Disk (imported) for backups imported from backup repositories that aren't tied to a backup or backup copy job. Restores can be performed simply by right-clicking on the VMs for the restore and using the desired option.

**Note:** Application item restore options will not appear for VM backups containing supported applications if "Guest Aware Image Processing" isn't chosen in the backup job specification.

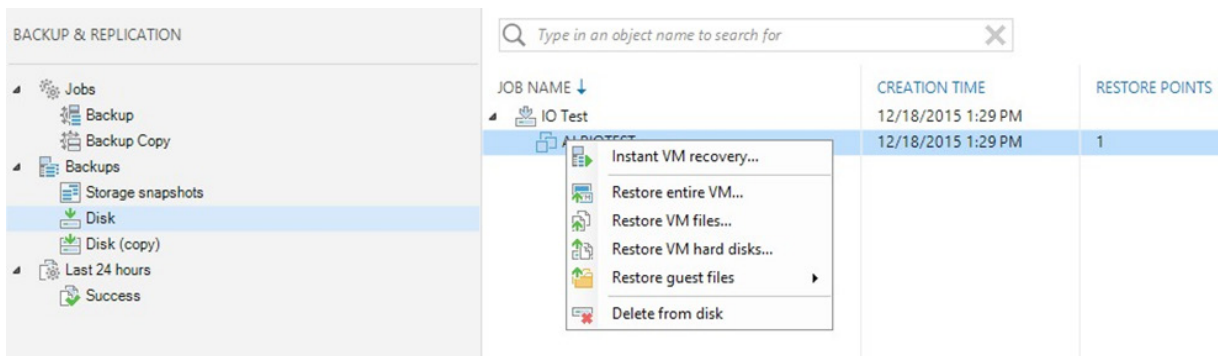


Figure 44: Restoring from disk-based backups in the Veeam GUI.

## Resources

- Veeam Backup & Replication user guide for Hyper-V or VMware
- StoreOnce support matrix
- Veeam with StoreOnce limitations
- Veeam Community Forums – Backup & Replication
- HPE StoreOnce page
- HPE-Veeam page
- HPE Technical Support Service for StoreOnce Catalyst

## Appendix

StoreOnce model	Max restore points per single backup chain (1)	Maximum concurrent backup sessions per node (2)
VSA	7	64
6200	14	192
2700	7	48
2900	14	96
4500	14	128
4700	14	192
4900	28	500
6500	28	512
3100	7	80
3500	14	192
5100	21	320
5500	35	1000
6600	42	1024

- This defines the maximum number of incremental backups associated with a full backup. If one backup per day is created on a StoreOnce 2700, then there needs to be at least one synthetic full per week. Please note, this limit does not influence the maximum number of restore points since it is possible to have multiple chains in retention.
- This number defines the limit for the StoreOnce node (server). Version 9.5 monitors the number of concurrent sessions and can pause jobs on "wait for resource" state until enough resources are available.

## About Veeam Software

Veeam® recognizes the new challenges companies across the globe face in enabling the Always-On Business™, a business that must operate 24.7.365. To address this, Veeam has pioneered a new market of *Availability for the Always-On Enterprise™* by helping organizations meet recovery time and point objectives (RTPO™) of less than 15 minutes for all applications and data, through a fundamentally new kind of solution that delivers high-speed recovery, data loss avoidance, verified protection, leveraged data and complete visibility. [Veeam Availability Suite™](#), which includes [Veeam Backup & Replication™](#), leverages virtualization, storage, and cloud technologies that enable the modern data center to help organizations save time, mitigate risks, and dramatically reduce capital and operational costs.

Founded in 2006, Veeam currently has 47,000 ProPartners and more than 242,000 customers worldwide. Veeam's global headquarters are located in Baar, Switzerland, and the company has offices throughout the world. To learn more, visit [www.veeam.com](http://www.veeam.com).

## About HPE StoreOnce

Hewlett Packard Enterprise was founded in 2015 from the split of Hewlett-Packard into HP Inc. and Hewlett Packard Enterprise. At HPE we help customers use technology to slash the time it takes to turn ideas into value. In turn, our customers transform industries, markets and lives.

HPE Storage delivers all flash and disk storage in software defined, hyper-converged, modular and traditional form factors. For secondary storage, HPE offers tape (StoreEver) and disk (StoreOnce) products. StoreOnce is available as purpose built backup appliances and as a virtual appliance. The StoreOnce appliances interoperate with Veeam and other leading backup software applications.

To learn more, visit <http://www.hpe.com> and [www.hpe.com/storage/StoreOnce](http://www.hpe.com/storage/StoreOnce).

AVAILABILITY for the Always-On Enterprise™

VEEAM

A nighttime photograph of a city skyline, likely Chicago, with several skyscrapers illuminated. The buildings are reflected in a body of water in the foreground. The overall color palette is dark with green and white highlights from the lights and text.

# Veeam makes the Fortune 500 Available.

# 24.7.365

To enable its **Digital Transformation**, 70% of the Fortune 500 rely on Veeam to ensure Availability of all data and applications. **24.7.365**