

Morro Data

Best Practice Guide

VM Backup to Cloud
with

Veeam Backup & Replication
Morro Data CloudNAS
Wasabi Hot Storage

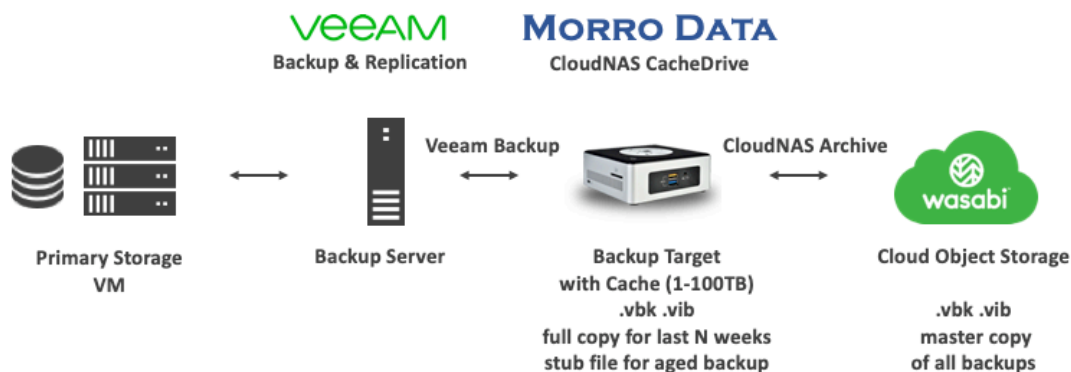
Introduction

VM backup and recovery is part of the critical IT operations to ensure business continuity. Traditionally IT has deployed an array of purpose-built backup appliances and applications to protect against server, infrastructure, and security failures. As VMs continue to spread in production, development, and verification environments, the never-ending challenge to expand VM backup repository has become a major challenge for system administrators.

Since VM backup footprint is usually quite large, cloud storage is increasingly being deployed for VM backup. However, cloud storage does not achieve the same performance level as on premises storage for recovery operation. For this reason, cloud storage has been used as tiered repository behind on premises storage.

In this best practice guide, we will show how Veeam Backup & Replication can work with Morro Data CloudNAS to keep the most recent backups on premises for fast recovery while archiving all backups in the retention window in the Wasabi Hot Storage cloud storage. CloudNAS caching not only provides buffer for most recent backup files but also simplifies the management of on premises storage and cloud storage as integral backup repository.

Concept for VM Backup and Archive



In the above diagram, the CloudNAS CacheDrive works as the Veeam backup target. CacheDrive's storage servers as the cache for the backup files and it uploads them to the cloud object storage for archive. This configuration has the following advantages:

- ✓ CacheDrive and cloud storage work as integral backup repository with unlimited capacity
- ✓ Duplicate backup copies in CacheDrive up to cache capacity
- ✓ Fast recovery of recent backups from CacheDrive

This Guide

For brevity, this guide assumes the reader is somewhat familiar with Veeam Backup & Replication, Morro Data CloudNAS, and Wasabi Hot Storage. We will focus on key discussions and major steps of configurations and skip some details.

This guide consists of the following parts:

- Part 1: Create the cloud storage bucket
- Part 2: Install and configure Morro Data CloudNAS
- Part 3: Configure Veeam backup repository using CacheDrive
- Part 4: Create the Veeam backup job
- Part 5: Run backup
- Part 6: Run recovery
- Summary

Part 1: Create the Cloud Storage Bucket

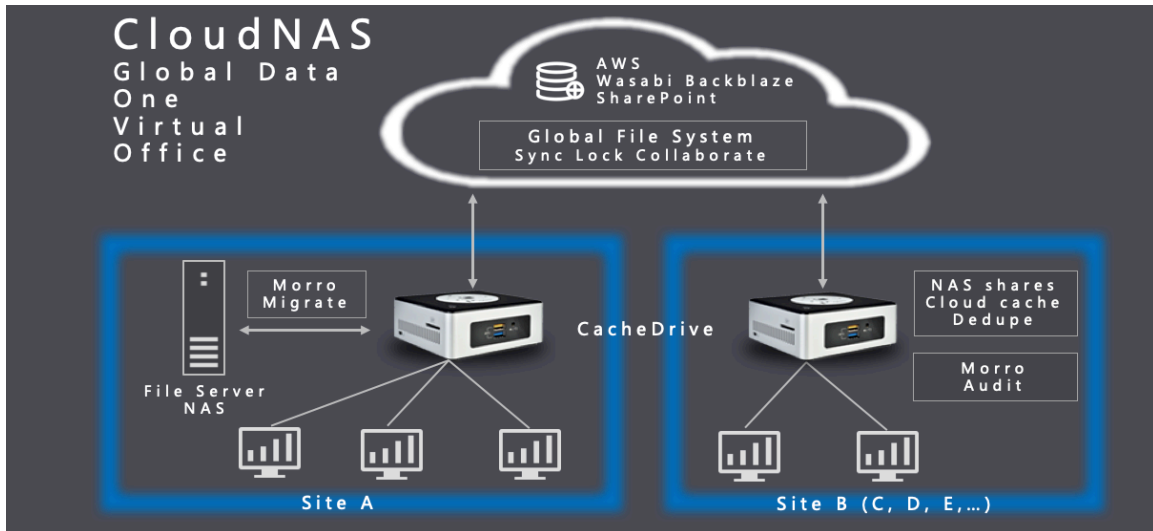
Creating a Wasabi Hot Storage account and storage bucket is quite straightforward. We assume the reader is familiar with this step. If you do not already have a Wasabi Hot Storage account, please go to the site to set up your account.

<https://www.wasabi.com>

Notes:

1. Make sure to set the bucket to Private.
2. Version setting can be left as default (keep all versions) although Veeam uses timestamp rather than version for recovery.

Part 2: Install and Configure Morro Data CloudNAS



Morro Data CloudNAS combines the high performance of a local NAS with the scalability and reliability of the cloud. Powered by the Morro Global File System and a hybrid cloud architecture, the CloudNAS Global File Services enable the following major applications:

- Sync among multiple sites
- Replicate to cloud and other sites
- Archive to cloud

Setting up CloudNAS is simple. The following are the steps we use for this guide:

1. Power up CacheDrive (physical or VM) and connect to Internet.
2. Sign up at <https://account.morrodatabackup.com>. After receiving the confirmation email, log in to Morro Cloud Manager (MCM).
3. In MCM > File System, add cloud storage by selecting the “Your Object Storage for Archive” option and choose Wasabi. Enter the necessary Wasabi Hot Storage information including bucket name, Key ID, and Secret Key.

4. Create a storage pool under the newly created cloud storage. Storage pool is a virtual layer between the physical cloud storage and the share.
5. Create a share under the new storage pool as our VM backup target. This share can map to any pathname in the B2 bucket. After creating the share, click the share icon to manage the share. In the share management panel example at the right, we create a share named "VeeamCacheShare". The share we create is of the share type Archive Share. It functions just like a regular share plus it uploads the files in the share to the B2 bucket pathname at scheduled intervals. Below is an example of the upload schedule and we will use it for this guide.

SCHEDULE

▼ Did you know ...

[SELECT ALL](#)
[CLEAR ALL](#)
[SYNC NOW](#)

S	M	T	W	T	F	S	
0:00	1:00	2:00	3:00	4:00	5:00	6:00	7:00
8:00	9:00	10:00	11:00	12:00	13:00	14:00	15:00
16:00	17:00	18:00	19:00	20:00	21:00	22:00	23:00

[CLOSE](#)
[OK](#)

Share ×

[MANAGE](#)
[ANALYZE](#)


▲


▼ Did you know ...

[Name*](#)
 VeeamCacheShare 15/45


[Comment](#)
 Veeam Backup 12/90

[Share Type](#)
 Archive

[User Permissions](#) 



[Source Gateway](#)
 VeeamCacheDrive

[Schedule](#) 

Continuous

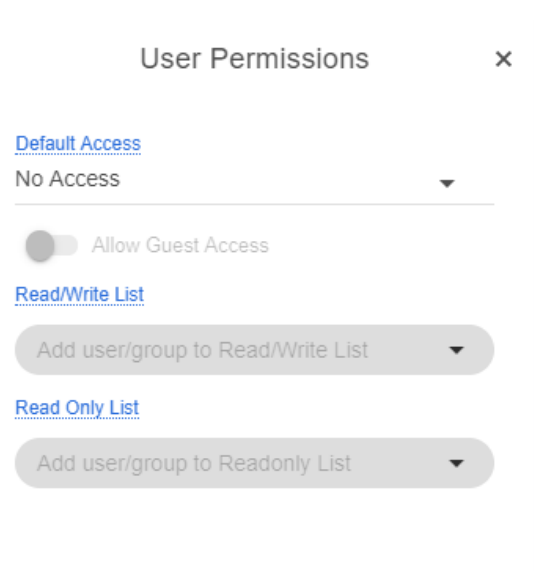
S
 M
 T
 W
 T
 F
 S

[Folder in Cloud Storage](#)
 /Veeam backup from Headquarter

About backup target share permission

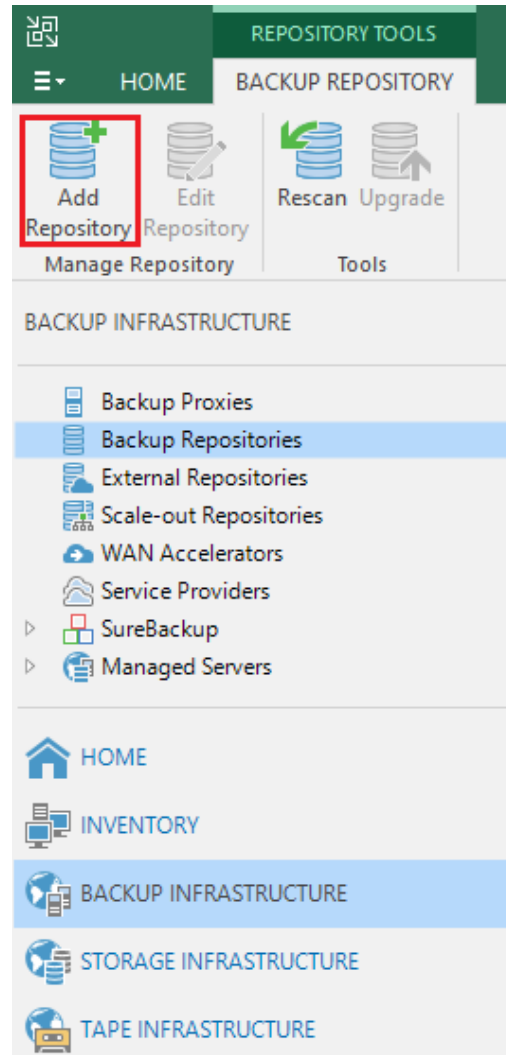
Strict permission should be used for the backup target share. We suggest to set share default access to “No Access” and do not give exceptions, as the example to the right shows. With this setting, only Local Administrator (admin) and your primary Domain Administrator will have administrator privileges.

Note: The default password for the Local Administrator (admin) is the same as the CloudNAS Business Administrator. However, the password of the Local Administrator can be separately set by changing it in the MCM > Team page.



Part 3: Configure Veeam Backup Repository using CacheDrive

CloudNAS CacheDrive functions as a Network Attached Storage and Veeam supports NAS as backup repository. In configuring Veeam, we use the CacheDrive as the backup target by clicking Backup Infrastructure > Backup Repositories and choose Add Repository. Click Network Attached Storage.



Add Backup Repository

Select the type of backup repository you want to add.



Direct attached storage

Microsoft Windows or Linux server with internal or direct attached storage. This configuration enables data movers to run directly on the server, allowing for fastest performance.



Network attached storage

Network share on a file server or a NAS system. When backing up to a remote share, we recommend that you select a gateway server located in the same site with the share.



Deduplicating storage appliance

Dell EMC Data Domain, ExaGrid, HPE StoreOnce or Quantum DXi. If you are unable to meet the requirements of advanced integration via native appliance API, use the network attached storage option instead.




Object storage

On-prem object storage system or a cloud object storage provider. Object storage based repositories can only be used for Capacity Tier of scale-out backup repositories, backing up directly to object storage is not currently supported.

We name this backup repository “CacheDriveStore”.

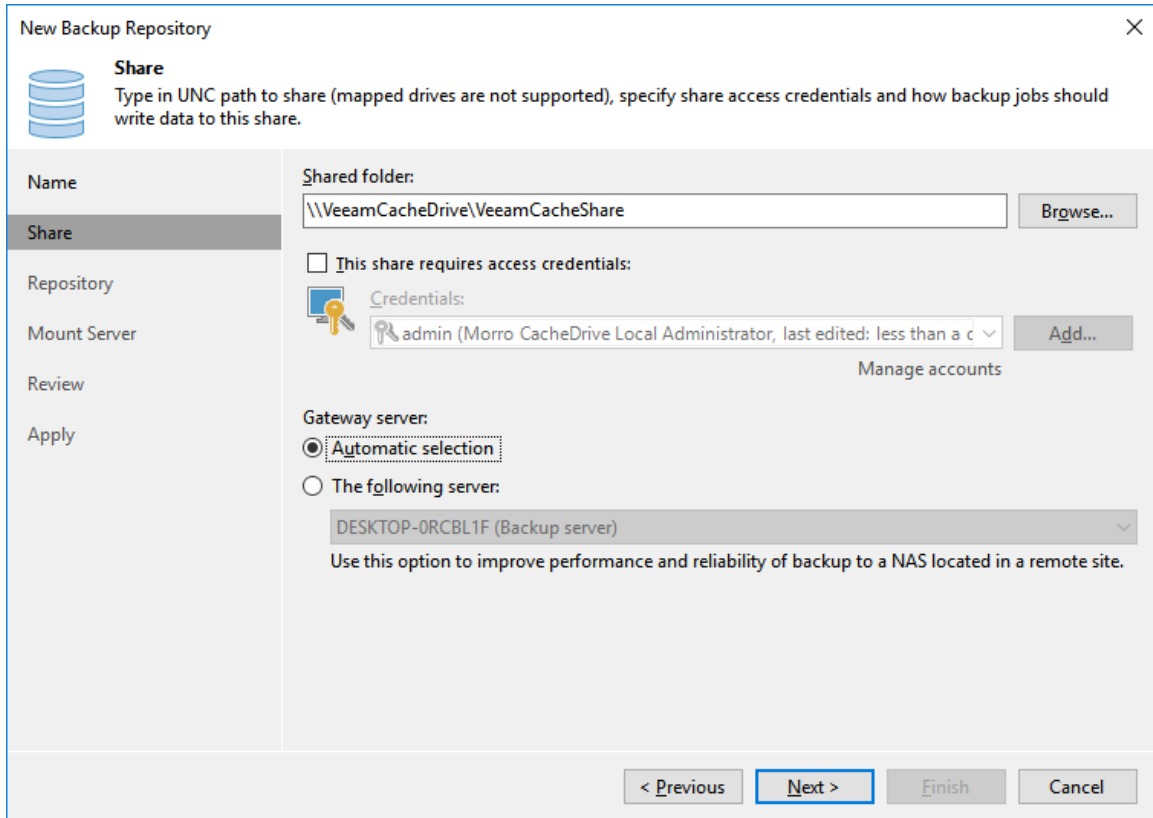
Edit Backup Repository

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

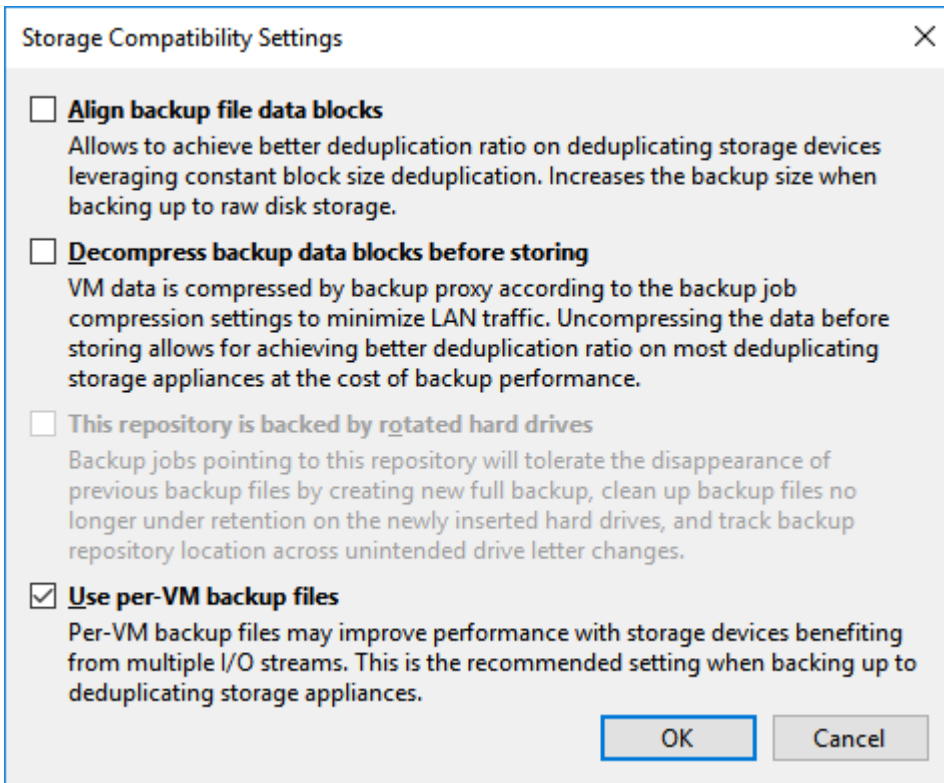
Name	<p>Name: CacheDriveStore</p>
Share	
Repository	<p>Description: Created by DESKTOP-0RCBL1F\hagi at 3/14/2019 2:31 PM.</p>
Mount Server	
Review	
Apply	

< Previous **Next >** Finish Cancel

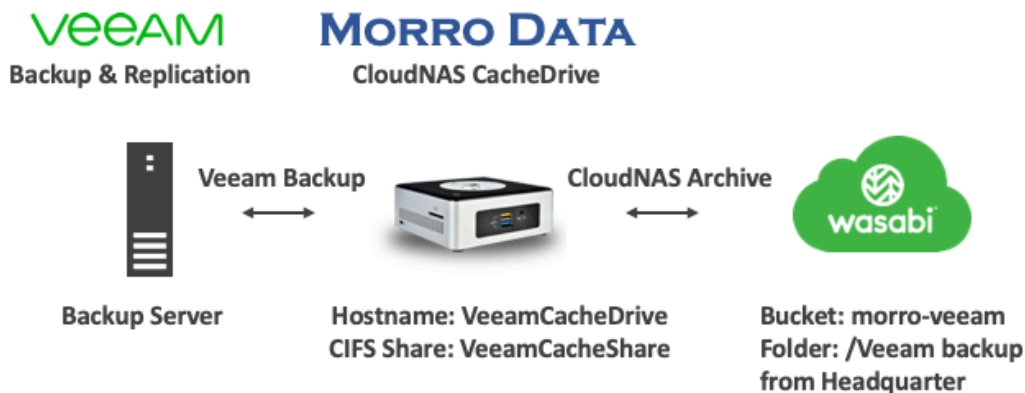
Next browse to select the backup target share “\\VeeamCacheDrive\VeeamCacheShare”. For share access credential, it is recommended to use the Local Administrator (‘admin’). If the CacheDrive is in a domain, we can also use the Primary Domain Administrator.



In Repository > Advanced Settings, we choose “Use per-VM backup files” to have more granularity for quickly restoring the particular VM when retrieving restore points from the cloud.



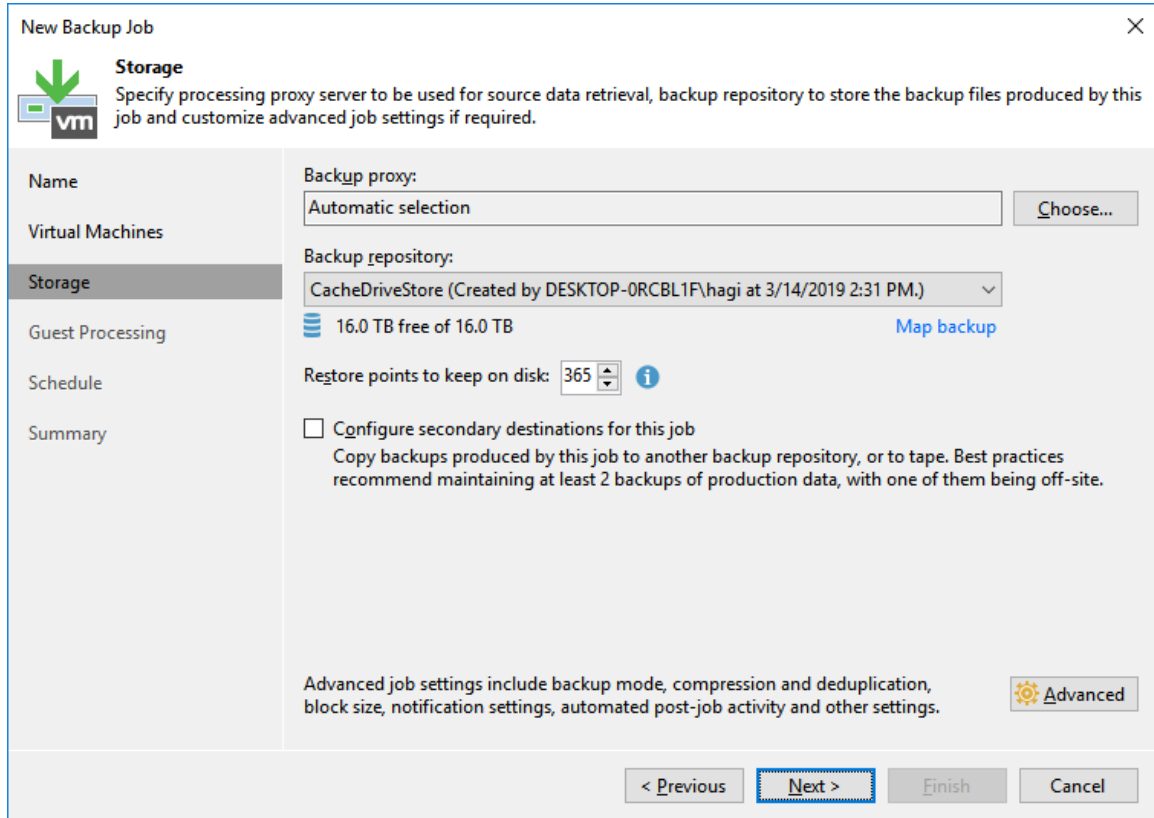
Review your configurations and click Apply for all changes. And now we have set up the CacheDrive as the VM backup target as below.



Part 4 – Create the Veeam Backup Job

This part is rather standard other than the part of setting the appropriate configuration for a backup repository including cloud storage to conserve the required upload bandwidth.

After adding the VMs for the backup job, select the backup repository “CacheDriveStore” that we created followed by defining the number of restore points that we require.

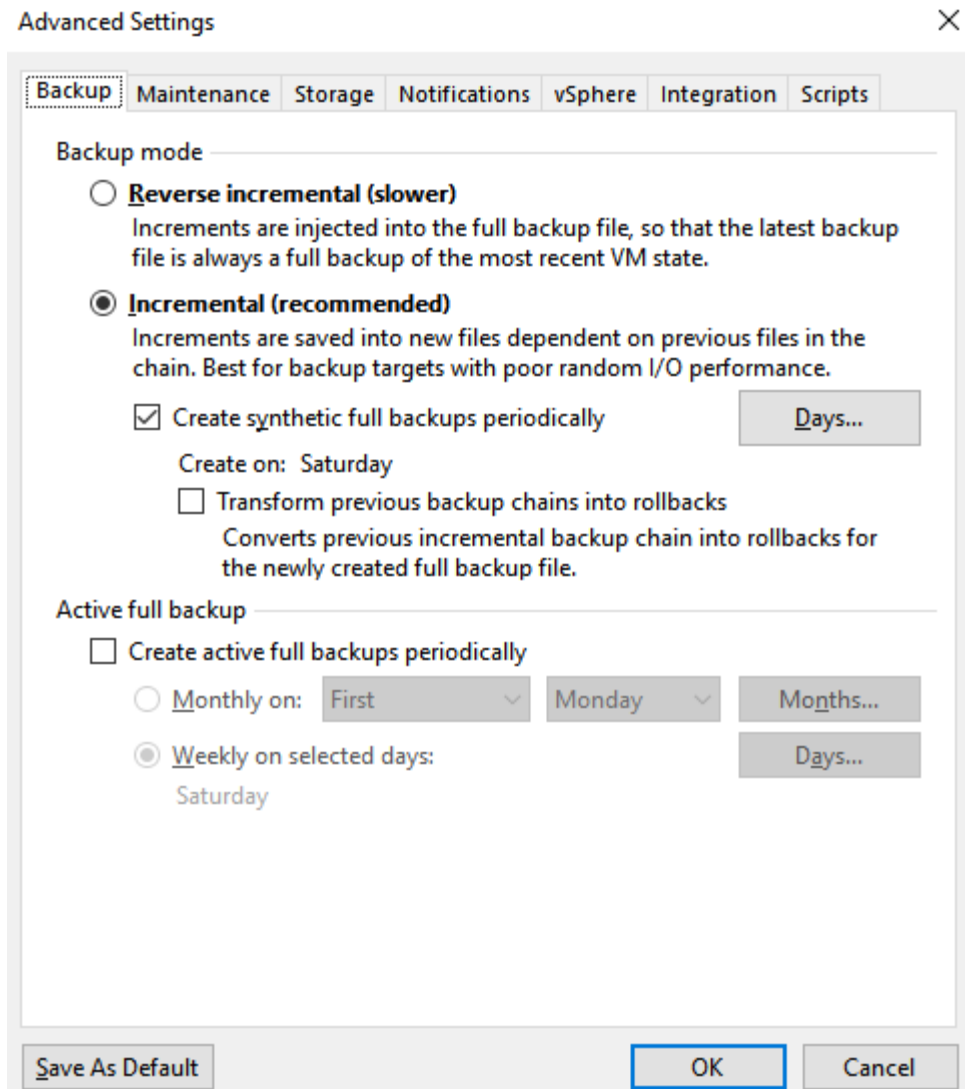


The screenshot shows the 'New Backup Job' wizard in the 'Storage' step. The interface includes a sidebar with navigation options: Name, Virtual Machines, Storage (selected), Guest Processing, Schedule, and Summary. The main area contains the following configuration options:

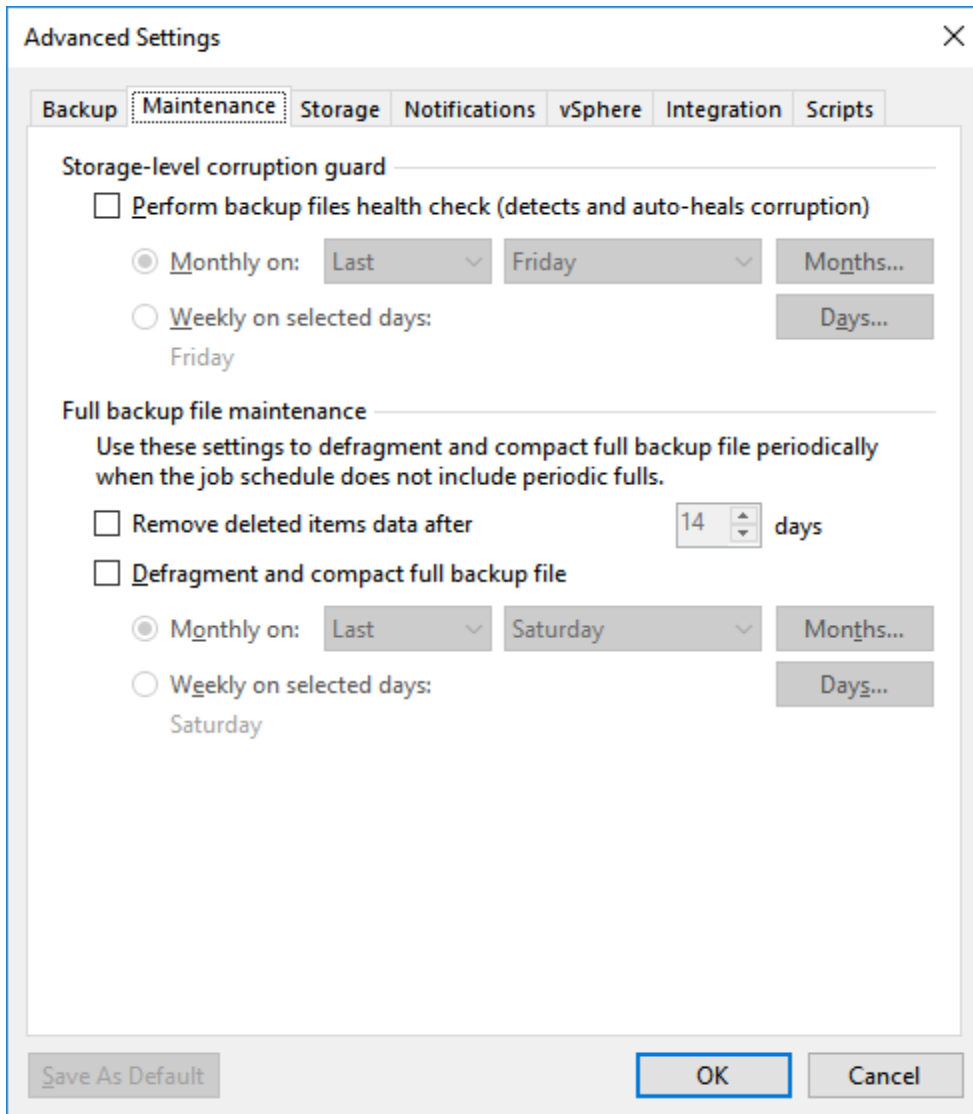
- Storage:** A green arrow icon and a 'vm' icon are shown next to the title. Below the title, it says: "Specify processing proxy server to be used for source data retrieval, backup repository to store the backup files produced by this job and customize advanced job settings if required."
- Backup proxy:** A text box containing "Automatic selection" and a "Choose..." button.
- Backup repository:** A dropdown menu showing "CacheDriveStore (Created by DESKTOP-0RCBL1F\hagi at 3/14/2019 2:31 PM.)" and a "Map backup" link.
- Free space:** A blue icon followed by "16.0 TB free of 16.0 TB".
- Restore points to keep on disk:** A spinner box set to "365" with an information icon.
- Secondary destinations:** An unchecked checkbox labeled "Configure secondary destinations for this job" with a sub-note: "Copy backups produced by this job to another backup repository, or to tape. Best practices recommend maintaining at least 2 backups of production data, with one of them being off-site."
- Advanced settings:** A note at the bottom states: "Advanced job settings include backup mode, compression and deduplication, block size, notification settings, automated post-job activity and other settings." Next to it is an "Advanced" button with a gear icon.

At the bottom of the wizard, there are four buttons: "< Previous", "Next >" (highlighted with a red box), "Finish", and "Cancel".

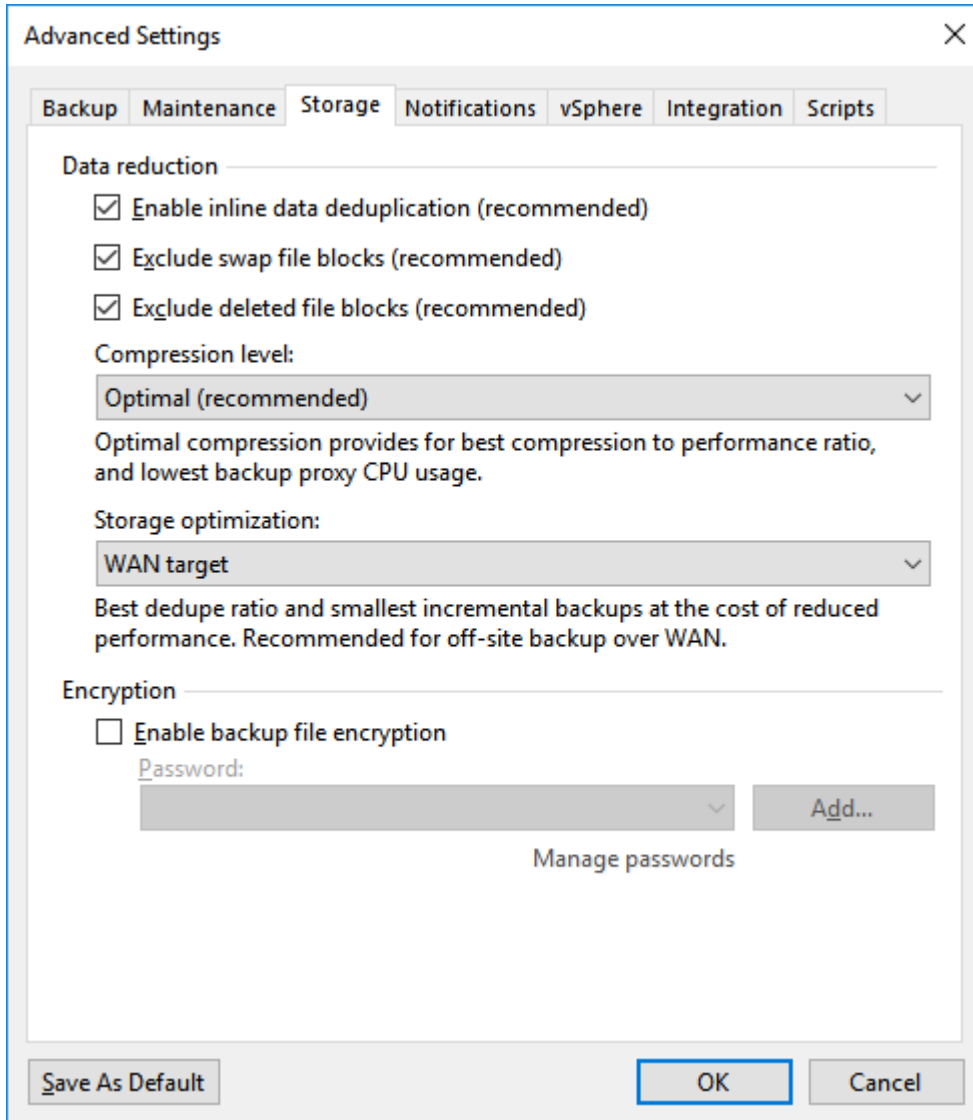
In Advanced Settings > Backup, select Incremental as recommended. Reverse Incremental would both create a new incremental backup file and cause the backup file .vbk to be partially updated for each incremental backup, needing more upload bandwidth.



In Advanced Settings > Maintenance tab, do not enable Perform backup files health check and do not enable Defragment and compact full backup file. We want to disable these settings to limit upload traffic as both operations result in creating new versions of old backup files.



In Advanced Settings > Storage tab, enable all the recommended data reduction options to reduce upload bandwidth requirements. Set compression level to Optimal or Extreme. Storage optimization should be set to WAN target again to reduce upload bandwidth.



Next we will configure the backup schedule to start the backup just before midnight every weekday. Assuming the backup job take less than 2 hours, we set VeeamCacheShare upload schedule at 2AM from Tuesday through Saturday.

Edit Backup Job [Backup ESX]

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

Name Run the job automatically

Virtual Machines Daily at this time: 11:59 PM On weekdays Days...

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

Guest Processing Periodically every: 1 Hours Schedule...

Schedule After this job: Backup CacheDrive (Created by DESKTOP-0RCBL1F,hagi at 3/14/2015)

Summary

Automatic retry

Retry failed items processing: 3 times

Wait before each retry attempt for: 10 minutes

Backup window

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

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

< Previous Apply Finish Cancel

Review the summary and click Finish to complete the backup job.

Part 5 – Run Backup

There are two parts of running the VM backup job – Veeam Backup and CloudNAS Archive. Veeam Backup refers to the backup operation performed by the Veeam backup proxy server. At the scheduled backup intervals, the proxy server reads from the VM datastore and compresses and writes to the CacheDrive. CloudNAS Archive refers to the snapshot and upload operations performed by the CacheDrive. At the scheduled upload intervals, CacheDrive takes a snapshot of the Archive Share and uploads the backup files to the cloud.

Local Backup Window and Cloud Upload Window


In most situations, the upload bandwidth is limited. So care should be taken to make sure upload is allotted enough time given the size of the backup files. After the first run, we can estimate the time required for a full backup. Based on that, we can set the upload schedule for the CloudNAS Archive share. Archive Share upload is based on snapshot so it is OK to schedule the next Veeam backup before the upload of the previous job completes.

User should schedule CloudNAS Archive so that it starts only after Veeam Backup completes. In the case that a large backup job requiring Veeam to continuously write backup files when CloudNAS Archive snapshot takes place, CloudNAS Archive will detect open files and will not upload partially completed backup files.

The following scenarios illustrate the relationship between time windows of Veeam Backup and CloudNAS Archive.

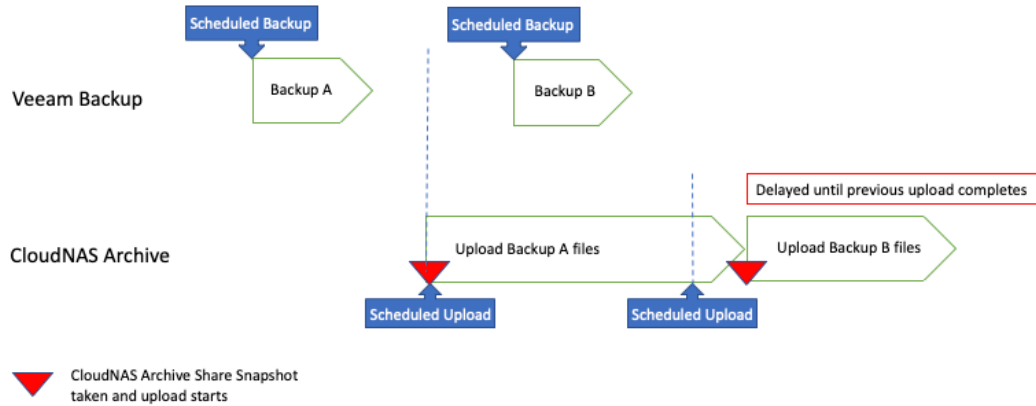
Veeam Backup to CloudNAS Archive Share

Normal Scenario



Veeam Backup to Archive Share

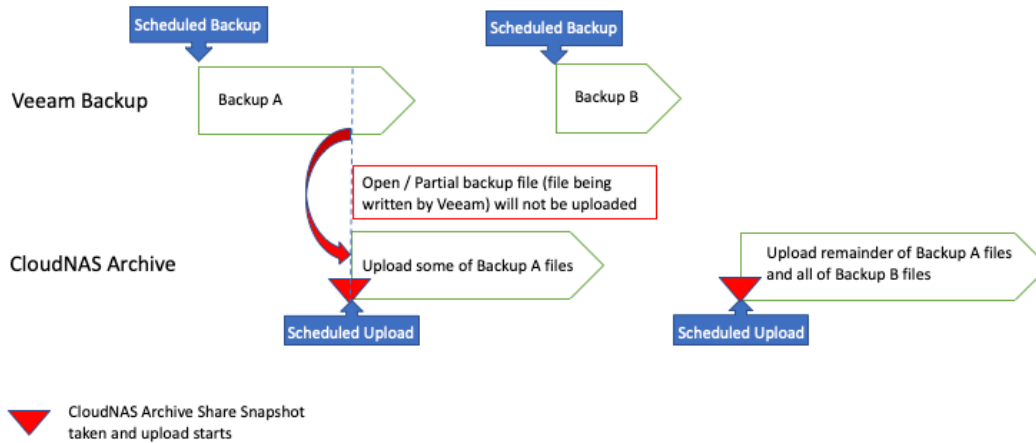
CloudNAS Archive upload takes longer than the next scheduled upload



Veeam Backup to Archive Share

CloudNAS Archive upload starts before Veeam Backup completes

(not recommended)




The next CloudNAS update is supposed to have CloudNAS Archive API to interface with Veeam post-freeze hook to automatically start upload as soon as Veeam Backup is completed. Please contact Morro Data support for details.

Files in Cloud Storage

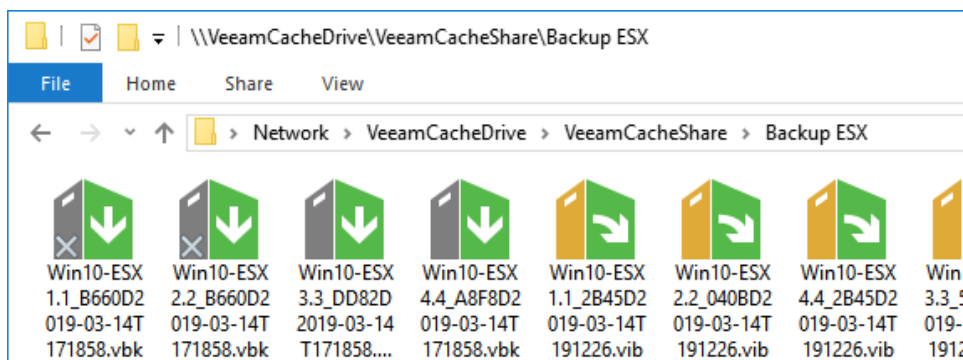
We can see the backup files in the Wasabi portal. There are three types of file extensions: “.vbm” is the xml file which contains the metadata for the backup job, “.vbk” is the fullback up file and “.vib” is the incremental backup. The “.vbm” file has multiple versions because it is updated with each job run.

File in CacheDrive Archive Share

Now let’s examine the files stored on the CacheDrive. During backup, CacheDrive will swap out the oldest files (already in cloud) to make room for new files. The cached-out files can still be seen in the CacheDrive. In Windows File Explorer, these cached-out files have an X on their file icons. The following screenshot shows the Windows File Explorer view of the files, and the first two “.vbk” files are cached-out files with an X on their file icons.

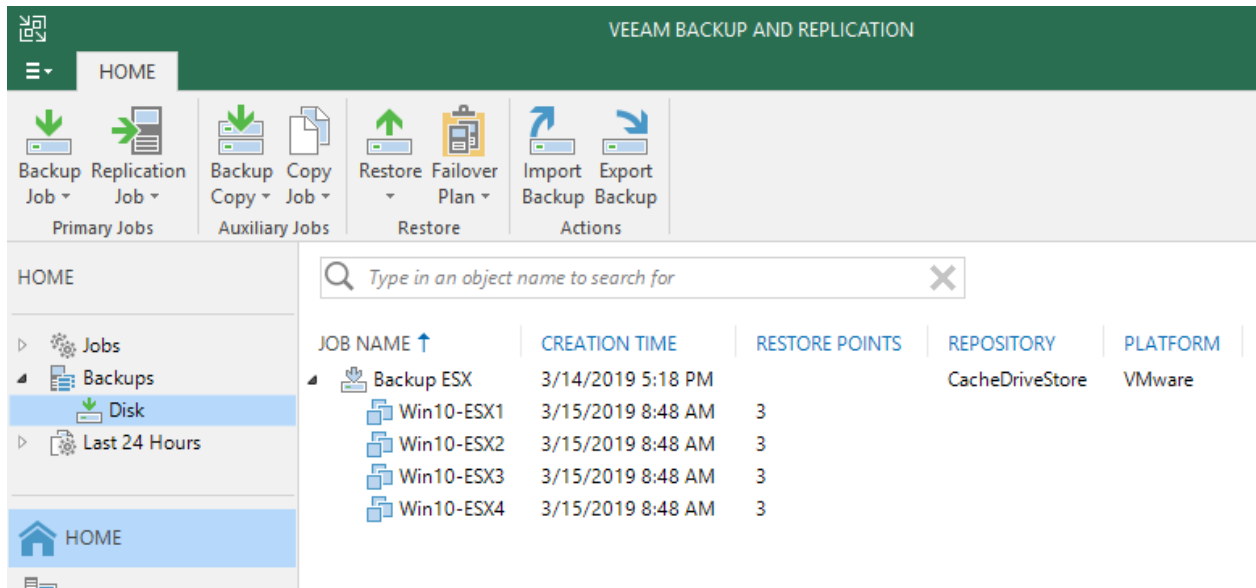


The following icon view shows the cache-out status more clearly.

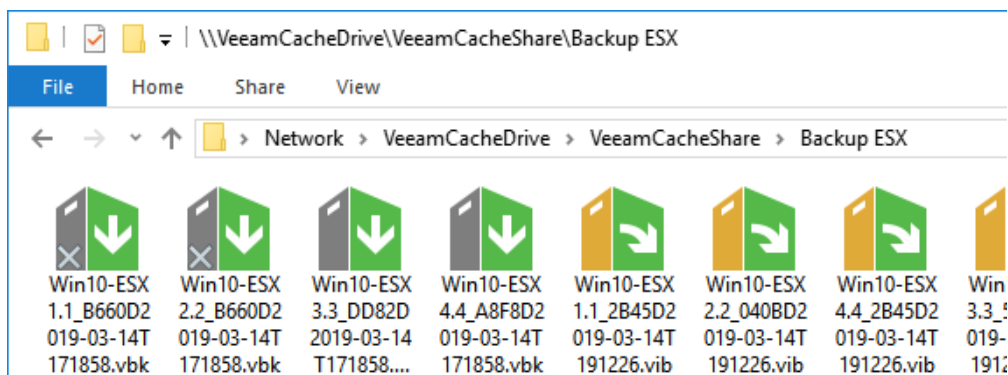


Part 6 – Run Recovery

A restore point is created after each backup job run. Each restore point corresponds to either a full backup or an incremental backup. The following screenshot shows four restore points available for recovery.



The following Windows File Browser view shows that the backup file for Win10-ESX1 is cached-out. We will attempt to restore using this cached-out file to show how recovery from cloud is like.



The recovery of a VM whose backup file is still in the CacheDrive is straightforward and fast. Here we try to restore “Win10-ESX1”, whose full backup file is already cached out. In the following recovery panel, we see 3 restore points. We pick one whose type is Incremental so to illustrate the recovery of backup files that are in the cloud only. In order to restore to this point, Veeam will need to retrieve the associated full backup file (.vbk) and the incremental file (.vib).

However, at this moment the content of the full backup file does not reside on the CacheDrive.

Full VM Restore

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

Virtual Machines to restore:

Type in a VM name for instant lookup

Name	Size	Restore point
Win10-ESX1	52.0 GB	less than a day ago (8:48 AM ...)

Add VM
Point...

Restore Points

Available restore points for Win10-ESX1:

Job	Type	Location
Backup ESX		
less than a day ago (8:48 AM Friday 3/15/2019)	Increment	CacheDriveStore
less than a day ago (7:13 PM Thursday 3/14/2019)	Increment	CacheDriveStore
less than a day ago (5:19 PM Thursday 3/14/2019)	Full	CacheDriveStore

Full VM Restore

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

VM Folder:

Name	New Name	Folder
Win10-ESX1	Win10-ESX1-Restore	vm

Select multiple VMs to apply settings change in bulk. Name... Folder...

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

< Previous Next > Finish Cancel

When Veeam tries to access the cached-out file, the CacheDrive will start downloading the file automatically. Currently CacheDrive has a built-in file read timeout limit of 90 seconds. In other words, if CacheDrive cannot provide the requested file in 90 seconds, it will return a timeout error to the application that requests the file. This time-out mechanism is to prevent some applications who cannot recover well from a long file read time. If the cached-out file can be downloaded before timeout, everything is normal. However backup files are typically large so it is expected to see timeout error after 90 seconds. We also suggest the use of per-VM backup so it is much faster to download from cloud in the case of restoring a single VM without downloading a much larger multi-VM backup file.

VM restore

VM name: **Win10-ESX1** Status: **Failed**
 Restore type: Full VM Restore Start time: 3/15/2019 9:32:54 AM
 Initiated by: DESKTOP-0RCBL1F\hagi End time: 3/15/2019 9:35:16 AM


Statistics Reason Parameters Log

Message	Duration
✓ Queued for processing at 3/15/2019 9:33:02 AM	
✓ Processing Win10-ESX1	0:02:13
✓ Required backup infrastructure resources have been assigned	
✓ 6 files to restore (52.0 GB)	
✓ Restoring [datastore] Win10-ESX1-Restore/Win10.vmx	
✓ Restoring file Win10.vmx (3.1 KB)	0:00:01
✓ Restoring file Win10.nvram (264.5 KB)	
✓ Registering restored VM on host: 172.18.2.104, pool: Resources, folder: vm, storag...	0:00:02
✓ Preparing for virtual disks restore	0:00:03
✓ Using proxy VMware Backup Proxy for restoring disk Hard disk 1	
✗ Restoring Hard disk 1 (52.0 GB) :	0:01:35
✗ Restore job failed Error: A device attached to the system is not functioning. Failed...	

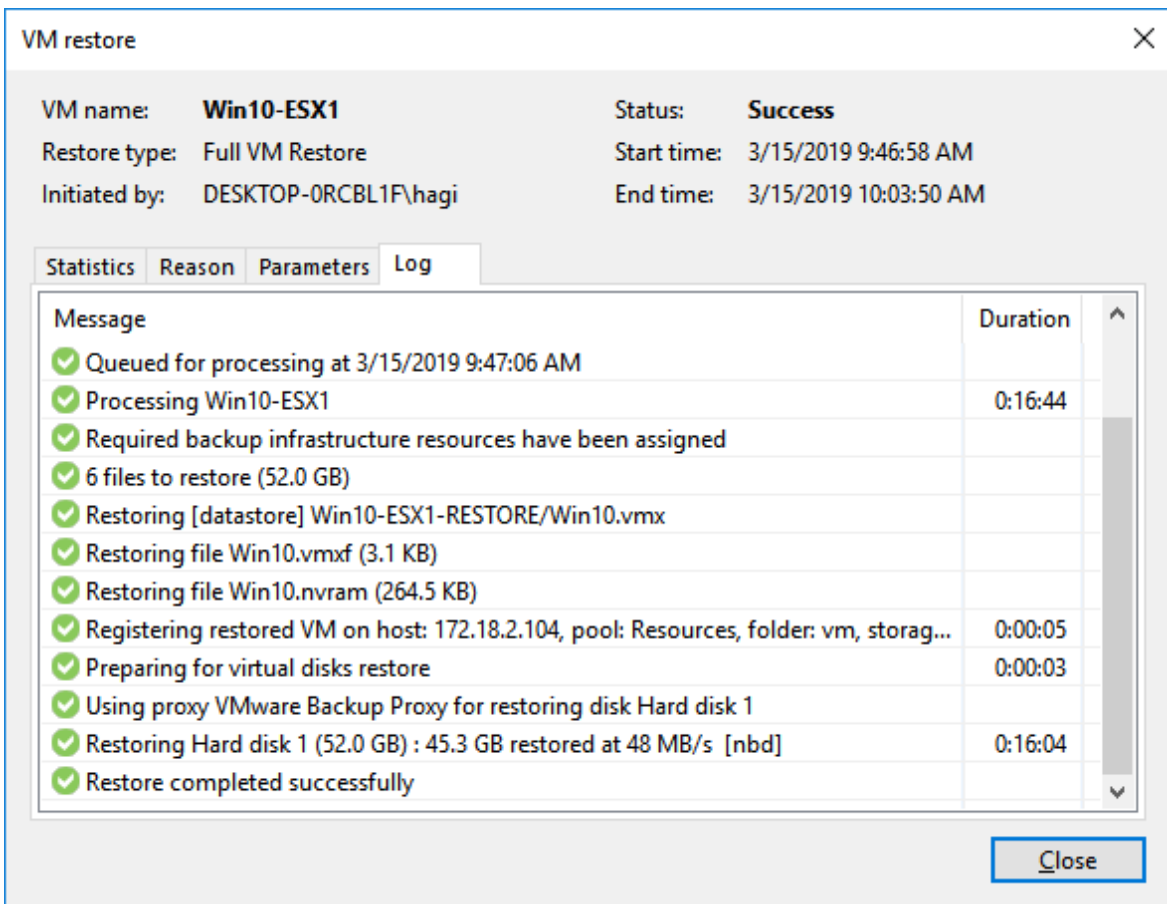
Restore job failed Error: A device attached to the system is not functioning. Failed to open file [\\VeeamCacheDrive\VeeamCacheShare\Backup ESX\Win10-ESX1.1_B660D2019-03-14T171858.vbk] in readonly mode. Failed to open storage for read access. Storage: [\\VeeamCacheDrive\VeeamCacheShare\Backup ESX\Win10-ESX1.1_B660D2019-03-14T171858.vbk]. Failed to upload disk. Shared memory connection was closed. Failed to download disk. Agent failed to process method {DataTransfer.SyncDisk}.

ST ENTERPRISE PLUS EDITION

When the above error “system is not functioning” appears, CacheDrive is already in the process of downloading the complete files. Every 100 mbps of download speed can download 1 GByte in 90 seconds. After the required download time, we can see that the X marks of the cached-out files are gone.



Now that the backup file is downloaded in the CacheDrive, we can restart the VM restore operation. The following shows a successful VM restore operation.



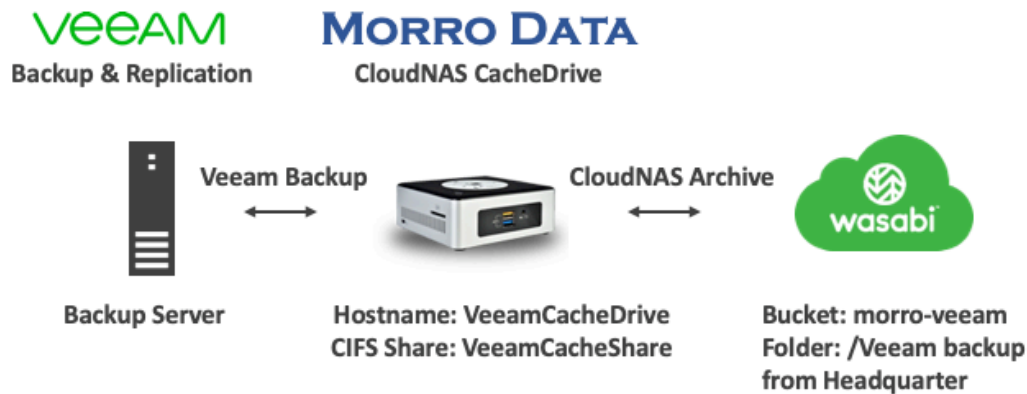
Disaster Recovery by downloading directly from Wasabi

In the unlikely event of a CacheDrive failure when VM needs to be restored, a replacement CacheDrive can quickly sync all cloud data to itself. After this initial sync, all files are in the cached-out state. When waiting for the replacement CacheDrive, all backup files can be accessed and downloaded from the Wasabi portal either directly or using a third-party tool.

Summary

In this guide, we have performed VM backup and recovery using Wasabi Hot Storage object storage with Morro Data CloudNAS. A CloudNAS CacheDrive was used as the backup target to keep recent backups on premises for fast recovery as well as upload all backups to B2 cloud storage. We present the following system diagram again for the summary.

Backup Repository Summary



In this guide we demonstrate that the combination of CacheDrive and Wasabi Hot Storage satisfies the following requirements for VM backup to cloud:

1. All backups are saved in the cloud for reliability and scalability
2. Recover recent backups from fast local storage
3. Simple IT for managing the backup repository