Best Practices Guide Veeam Data Platform

Table of Contents

CLOUDIAN®

Target Audience

This guide focuses on the configuration of Cloudian HyperStore and partly Veeam Backup and Replication and Backup for Windows 365. This paper is written for storage and network administrators who are familiar with Linux systems and tasked with maintaining or implementing a Cloudian HyperStore system in combination with Veeam Backup and Replication and Backup for Windows 365.

Solution Overview

Veeam Backup & Replication

Cloudian HyperStore enterprise storage complements Veeam Availability Suite with a fast, on-premises disk-based storage and built-in ransomware protection.

Scale out Backup Repository (SOBR)

Veeam Backup & Replication (VBR) uses Cloudian HyperStore as an on-premise capacity or performance tier via the S3 storage protocol using a scale-out repository. A scale-out backup repository (SOBR) contains both a primary backup repository as a performance tier and a capacity tier for long term storage. The capacity tier is simply an extension of the primary repository used for longer data retention periods and/or to store immutable backup copies.

Cloudian HyperStore in the Capacity Tier

Veeam Backup & Replication provides a policy-based engine that moves or copies the backup data from the performance tier to an S3 object storage target, such as Cloudian HyperStore. The backup meta-data is kept locally in the performance tier for intelligent search and browse capability. Cloudian HyperStore provides the ability to easily scale your archive storage capacity, performance and data durability options on-demand. Below is an illustration of a typical on-site Veeam Backup & Replication SOBR environment.

A few years ago, Veeam v9 update 4 introduced a few new features for leveraging S3 object storage and it was primarily as a repository for long term data retention for inactive backup chains. Back in early 2020, Veeam released Veeam B&R v10 which introduced the support “Copy Mode” for the capacity tier of a SOBR which allowed copies of active backup chains to be stored on the capacity tier. This also meant the data recovery and instant recovery could now be done directly from the capacity tier of a SOBR eliminating the requirement to re-inflate the backup data to the performance tier before a recovery operation can be performed.

In addition, one of the key features that was introduced with v10 R was the support of AWS S3 Object Lock API which provides the ability to create immutable copies of backups on s3 compatible Object Storage that also support this feature.

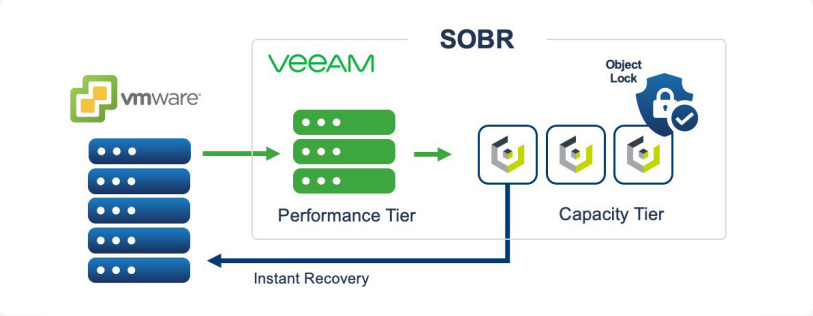

A Veeam Backup & Replication deployment typically consists of a single backup server, one or more backup proxies, one or more local or shared backup repositories and one or more source VMware vSphere or HyperV hosts with associated data stores that store virtual machines (VMs).

The Veeam Backup Server is responsible for defining and coordinating all activities in the backup domain. Backup proxy co-ordinates the data movement between the data sources, such as VMware, the backup repository and any object storage archive targets within a SOBR.

All backup infrastructure components engaged in the backup request make up a data pipe. Veeam treats VMs and other backup data as objects, not as a set of files. When backing up VMs, Veeam Backup & Replication copies a VM image as a whole, at a block level. Image-level backups can be used for different types of restore, including Instant VM Recovery, entire VM restore, VM file recovery, file-level recovery, and so on. A SOBR extends a performance tier, which is usually backups stored on local storage, to an S3 object storage pool. A SOBR policy is used to determine retention time in the performance tier before data is moved or copied to the object store for archival and or immutability. Veeam V10 introduced the ability to make copies of the backup data on the performance tier to the capacity in its SOBR policies. This provides the ability to create immutable copies of the backup data to an S3 compatible Object Storage target such as Cloudian HyperStore.

Now, with the release of Veeam Backup and Recovery version 12, Veeam extended their ability to leverage S3 Object Storage even further. They added features which include the ability to use S3 Object Storage as a standalone backup repository without the need of a primary tier of storage referred to as a "Performance Tier" in Veeam B&R, use s3 buckets as the performance tier and/or the capacity tier of a SOBR and Veeam B&R v12 now support using multiple s3 buckets extents as part of the same SOBR.

Cloudian HyperStore in the Performance Tier

Starting with Veeam Backup & Replication version 12 direct usage of Cloudian HyperStore in the Performance Tier is supported. One or multiple buckets can be configured and backups will be directly written to the bucket(s).

Veeam will load balance the the backups per VM over the multiple buckets. When you new buckets are added to the SOBR Performance Tier at a later point in time the existing VMs wont be moved to other buckets, but new VMs will be written to the newly added buckets.

In order to protect the data against ransomware immutability can be used in the Performance Tier, too.

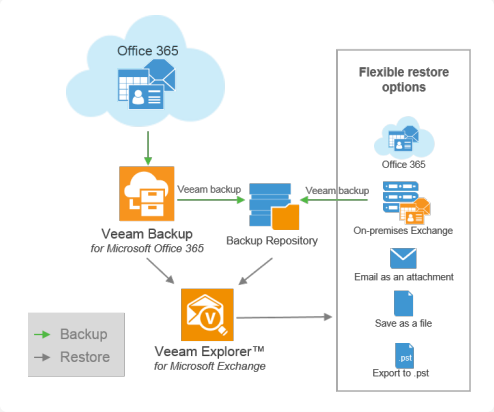

Veeam Backup for Windows 365

Microsoft Windows 365 enables your enterprise to work anywhere, anytime without the need to host your own email, files and content management infrastructure. But this does not replace your responsibility to backup business-critical Windows 365 data.

Microsoft Windows 365 offers geo-redundancy, which protects your data from site or device failure. But it does replace backup. With geo-redundancy alone, you have limited recovery options if your data is accidentally deleted or maliciously attacked.

Veeam Backup for Windows 365 (VB365) paired with Cloudian HyperStore provides a cost-effective solution to protect your Office 365 data. With Veeam and Cloudian, you have the access controls, data backups and data security that protect you from accidental deletion, rogue employees, and malware.

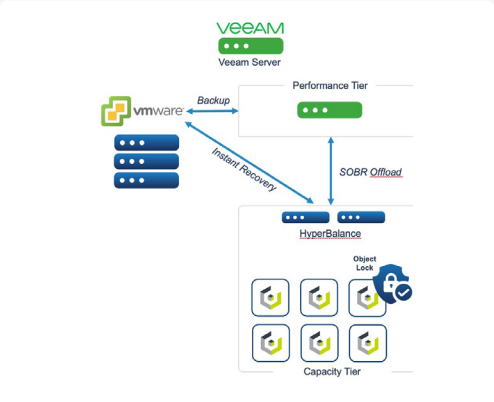

Solution Architecture

Veeam Backup & Replication

With Veeam Backup & Recovery Cloudian HyperStore is acting as the Capacity or Performance Tier of the Veeam SOBR. Depending on the configuration of the SOBR Capacity Tier the "SOBR Offload" is started once the Backup Job has loaded the data to the Performance Tier.

A usual HyperStore cluster consists of minimum 3 nodes which are placed behind a load balancer. The load balancer (e.g. Cloudian HyperBalance) is balancing the requests to the HyperStore nodes. Depending on the set up the load balancer can be virtualized or a physical pair of load balancer.

Amount and type of Cloudian HyperStore nodes depend on two factors. First of all the amount of nodes are calculated by the required storage amount. The storage amount depends on the backup and retention settings used. Additional the amount of nodes are calculated

by the required performance. This is the throughput required for reading and writing backups to the cluster and the performance required for deleting data.

The HyperStore cluster can be located in the same or in another data center. Backup and restore performance to and from the HyperStore cluster will depend on the connectivity between the data centers.

Cloudian HyperStore supports replication between data centers as well. If it is required to keep backup data in two or more data centers, the replication mechanism available with Cloudian Hyperstore can be used. With Cloudian Hyperstore, replication can be configured at the bucket level. For replication over a longer distances Cloudian HyperStore supports Cross Region Replication (CRR) as a mechanism to replicate the data. Depending on the situation it makes sense to use replication, CRR or a Veeam Backup Copy Job to replicate the data to the other data center(s).

Veeam Backup for Windows 365

Veeam Backup for Microsoft 365 is a comprehensive solution that allows you to back up and restore data of your Microsoft 365 organizations, including Microsoft Exchange, Microsoft SharePoint, Microsoft OneDrive for Business and Microsoft Teams data, as well as data of on-premises Microsoft Exchange and on-premises Microsoft SharePoint organizations.

When the backup job starts, the data first goes to the local cache, which is only stored in RAM memory, and afterwards is sent directly to the object storage. In the local cache Veeam stores the metadata to present the object directly via Veeam Explorer, but the content of the object itself (email body, attachment, appointment info, file contents…) is stored in the object storage.

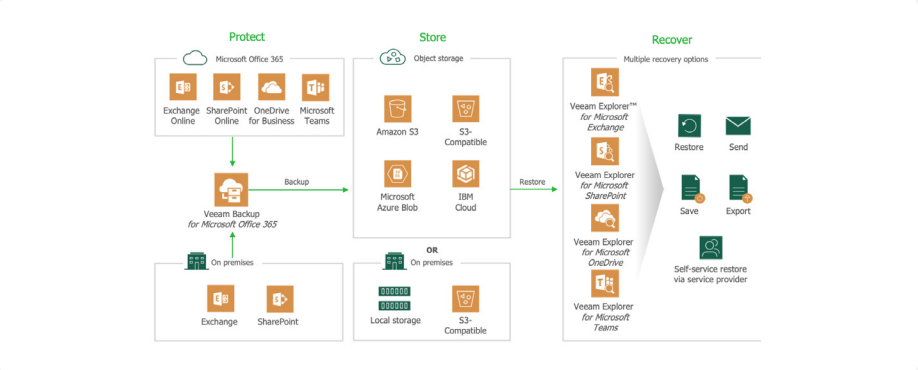

Object Sizes

Veeam Backup for Windows 365 does store data items in chunks to the object storage repository. These chunks are ,before compression, 5MB for Exchange data and 8MB for SharePoint and OneDrive data.

Additionally VB365 stores meta-data in own objects. These objects are small in size (100KB or less) and can make up for up to 50% of the overall number of objects in an object storage repository.

Cloudian Hyperstore Configuration

Recommended Version

Cloudian HyperStore 7.4.2 is the recommended version for Veeam deployments. Many optimization features were introduced in the Hyperstore 7.4 which will be beneficial to Veeam workloads.

HyperStore Secure Shell (HSH)

HyperStore Secure Shell (HSH) needs to be enabled to allow creation of Object Lock enabled buckets. Buckets with enabled Object Lock should only be used if there is a specific business (compliance) requirement to do so.

Important Note: Currently, Veeam Backup for Windows 365 does NOT support S3 Object Lock.

SSL Configuration

Veeam requires the endpoint to be connected via HTTPS only. A non encrypted connection is not supported by Veeam. A self signed certificate is supported by Veeam.

Bucket Storage Policy Configuration

Data Distribution Scheme

Cloudian recommends using a RF3, EC 4+2 or EC 6+3 storage policy for the s3 buckets. Overall less fragments are better.

Workload Type

Starting with Cloudian HyperStore Version 7.4 we implemented an easy way to optimize Cloudian for the Veeam workload. Cloudian HyperStore will detect the workload automatically and will optimize data handling for it. Inside the Storage Policy creation view you are now able to select a “Workload” type. In addition to using multiple buckets (hence multiple Veeam SOBRs in V11) to store Veeam data into.

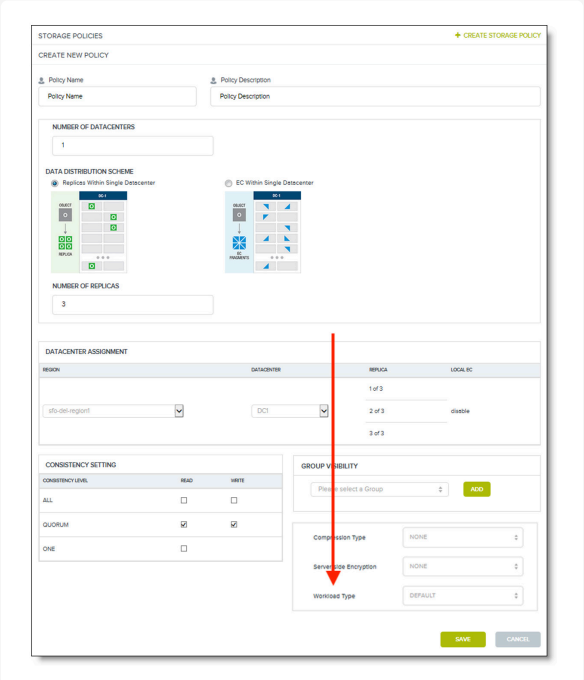

HyperStore 7.5.1 and newer

It’s recommended to choose the “DEFAULT” workload type as Cloudian HyperStore will automatically detect the workloads.

HyperStore 7.4

It’s recommended to choose the “VEEAM” workload type when using HyperStore buckets in the Capacity Tier and use “VEEAM12“ for buckets in the Veeam v12 Performance Tier.

For Veeam Backup for Windows 365 the workload type “VEEAMO365“ should be used.

HyperStore 7.4 does not display the “VEEAM12” or “VEEAMO365“ workload types out of the box on the Storage Policy Creation Panel and they have to be added manually:

Workload Type: VEEAM12

Get in touch with Cloudian Support to get the workload type added to your deployment.

{“bDep”:3,”bSuf”:0,”wlt”:[],”pRules”:[{“id”:0,”cond”:[{“n”:”P”,”o”:”START”,”v”:”Veeam/Backup/”},{“n”:”D”,”o”:”GE”,”v”:”6”}],”dep”:6,”suf”:0,”ud”:””,”st”:true}],”status”:true,”dw”:true}

Workload Type: VEEAMO365

Get in touch with Cloudian Support to get the workload type added to your deployment.

{“bDep”:3,”bSuf”:0,”wlt”:[“VEEAMO365”],”pRules”:[
{“id”:0,”cond”:[{“n”:”P”,”o”:”START”,”v”:”Veeam/
Backup365/”},{“n”:”D”,”o”:”GE”,”v”:”7”}],”dep”:7,”suf”:0,”ud”:””
“st”:true}],”status”:true,”dw”:true}

Older than HyperStore 7.4

In order to optimize queries and listing of objects in larger buckets, it is recommended to change the folder depth settings for the Veeam buckets inside Cloudian HyperStore by editing the mts.properties file before you create the Veeam bucket(s) . The process for doing this will be different depending if the Cloudian HyperStore deployment is used only for Veeam or if it is a mixed environment with other non Veeam buckets.

If you already have Veeam buckets and data in the buckets reach out to Cloudian Support to help on to change the bucket depth for existing buckets.

In the case you are using Cloudian HyperStore only for Veeam we recommend setting the overall bucket depth default to the new default. Once enabled all buckets which will be created afterwards will have a bucket depth of the new default.

1. Edit /etc/cloudian-{HS-VERSION}-puppet/modules/cloudians3/templates/mts.properties.erb

```
#vi /etc/cloudian-{HS-VERSION}-puppet/modules/cloudians3/templates/mts.properties.erb
```

With HSH:

```
# hspkg config -e mts.properties.erb
```

2. Add the property as below at the end of the file Veeam Capacity Tier:

```
cloudian.s3.folder.depth=7
```

Veeam Performance Tier :

```
cloudian.s3.folder.depth=6
```

Veeam Backup for Windows 365:

```
cloudian.s3.folder.depth=7
```

3. Push the configuration changes in puppet

```
# ./cloudianInstall.sh runpuppet=””
```

With HSH:

```
# hspkg install runpuppet=””
```

4. Restart S3 service (with or without HSH)

```
# ./cloudianService.sh -s s3 restart
```

With HSH:

```
# hspkg install
```

Important: All the buckets created after this will have a folder the same set folder depth!

When running mixed environments you can take back the changes after the required buckets were created by removing the added line, pushing the changes and restarting the S3 service again.

When running mixed environments you can take back the changes after the required buckets were created by removing the added line, pushing the changes and restarting the S3 service again.

Tune Cloudian for deletes (Version 7.4.x and newer)

Cloudian HyperStore can be used for many different types of workloads. In order to tune it for Veeam we recommend to adjust the handling of object deletion by the system.

cloudian.delete.dc.instances – Specifies the number of nodes within a datacenter to participate in the batch delete process. – DEFAULT IS 2

To distribute the load over all nodes in the cluster, in mts.properties change cloudian.delete.dc.instances from 2 to 0.

For Service Providers and environments not sized for Veeam workloads we recommend an additional change.

cloudian.s3.batch.delete.delay – delays the processing of each individual batch delete job by a number of milliseconds – DEFAULT IS 0

We recommend to additionally set cloudian.s3.batch.delete.delay from 0 to 50.

1. Edit /etc/cloudian-{HS-VERSION}-puppet/modules/cloudians3/templates/mts.properties.erb

```
#vi /etc/cloudian-{HS-VERSION}-puppet/modules/cloudians3/templates/mts.properties.erb
```

With HSH:

```
# hspkg config -e mts.properties.erb
```

2. Change the the property as below

```
cloudian.delete.dc.instances=0
```
(optional) `cloudian.s3.batch.delete.delay=50`

3. Push the configuration changes in puppet

```
# ./cloudianInstall.sh runpuppet=""
```

With HSH:

```
# hspkg install runpuppet=""
```

4. Restart S3 service (with or without HSH)

```
# ./cloudianService.sh -s s3 restart
```

With HSH:

```
# hspkg install
```

Bucket Properties

Veeam Backup & Replication

Veeam VBR Capacity Tier (S3 Offloading)

Before Veeam 12, only one bucket could be configured in the capacity tier. With Veeam VBR 12 multiple buckets are supported. For a better data management performance we recommend to use multiple buckets. As a rule of thumb create 1 bucket for every 100 VMs to be backed up.

Veeam VBR Performance Tier (Direct to Object)

When using Cloudian HyperStore in the performance tier multiple buckets should be created to optimize the data management. As a rule of thumb create 1 bucket for every 100 VMs to be backed up.

Veeam Backup for Windows 365

Only one bucket can be configured for each repository.

Versioning

Versioning SHOULD ABSOLUTELY NOT BE ENABLED on any bucket used for a Veeam SOBR unless it is used for backup immutability in combination with object lock.

Veeam Backup for Windows 365

Currently, Veeam Backup for Windows 365 does NOT support S3 Object Lock. Therefore, DO NOT use Object Lock nor versioning.

Life Cycle Policies

No lifecycle policies for tiering or expiring data are required or recommended. They absolutely should not be used.

Object Lock

Veeam Backup & Replication allows you to prohibit deletion or modification of data from object storage by making it immutable. Object Lock must be enabled when creating the bucket. Beside that, no other bucket level configuration is required. Don't configure any bucket level defaults. Veeam will manage object lock mode and retention by itself.

Veeam Backup for Windows 365

Currently, Veeam Backup for Windows 365 does NOT support S3 Object Lock. Therefore, DO NOT use Object Lock.

Tiering

Do not configure any tiering on object storage buckets used for Veeam Object Storage Repositories. This is not supported.

Tiering in object storages is based on object age. However, with Veeam's implementation even a very old block could still be relevant for the latest offloaded backup file when the block was not changed between the restore points. An object storage vendor can not know which blocks are still relevant and which not and thus can not make proper tiering decisions.

IAM and STS for Agent Backup

In Veeam VBR version 12, Veeam Agents for Windows / Linux / Mac will be able to leverage

the IAM/STS extensions of the S3 API to deliver secure multi-tenant backup within

a single S3 bucket. In order to allow access to these buckets Veeam will create IAM policies automatically.

This will only be used if the agents are standalone (not VBR managed).

Cloudian Configuration

Requirements

HyperStore Version: 7.5.1 or newer.

SSL needs to be configured for S3 and IAM/STS.

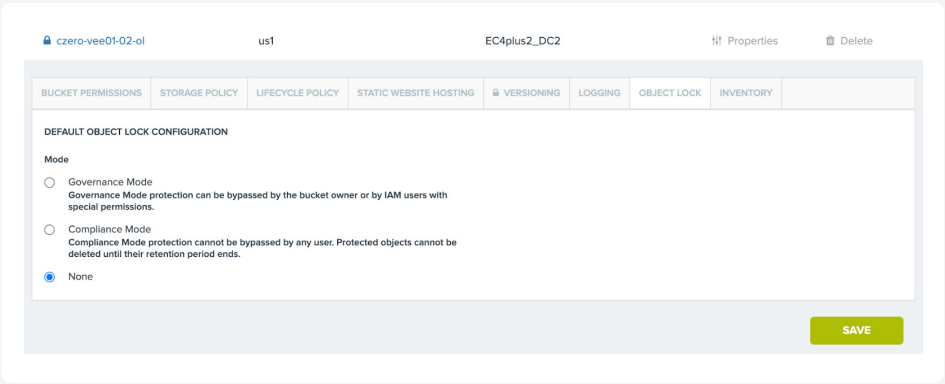

Configuration steps

1. Create a HyperStore user (e.g. “veeam_user“)
2. Login with the created user
3. Create an IAM user (e.g. “veeam_iam“)
4. Create new Access Key and Security Key
5. Create new IAM Policy by pasting below policy to the JSON tab:

```
{
“Version”: “2012-10-17”,
“Statement”: [
{
“Sid”: “VisualEditor0”,
“Effect”: “Allow”,
“Action”: “s3:”,
“Resource”: “”
},
{
“Sid”: “VisualEditor0”,
“Effect”: “Allow”,
“Action”: “sts:”,
“Resource”: “”
},
{
“Sid”: “VisualEditor0”,
“Effect”: “Allow”,
“Action”: “iam:”,
“Resource”: “”
}
]
}
```

6. Assign created IAM Policy to the user (e.g. “veeam_iam“)
7. Create a bucket and allow “Authenticated Users” access to the bucket.

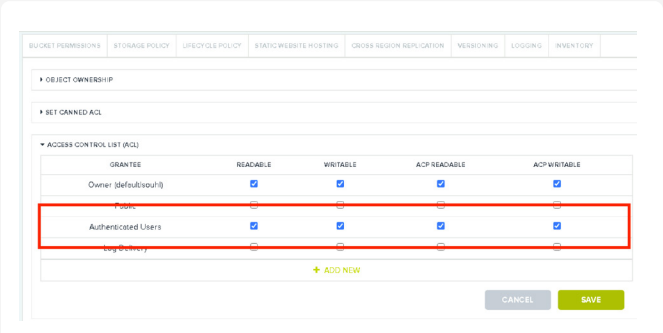

8. The access and secret key from the IAM user (e.g. “veeam_iam“) have to be used in the Veeam repository configuration, not the access and secret key of the HyperStore user (e.g. “veeam_user“)

Agent Configuration

- **Create Object Storage Repository**

In the Veeam console create a normal object storage repository (no SOBR).

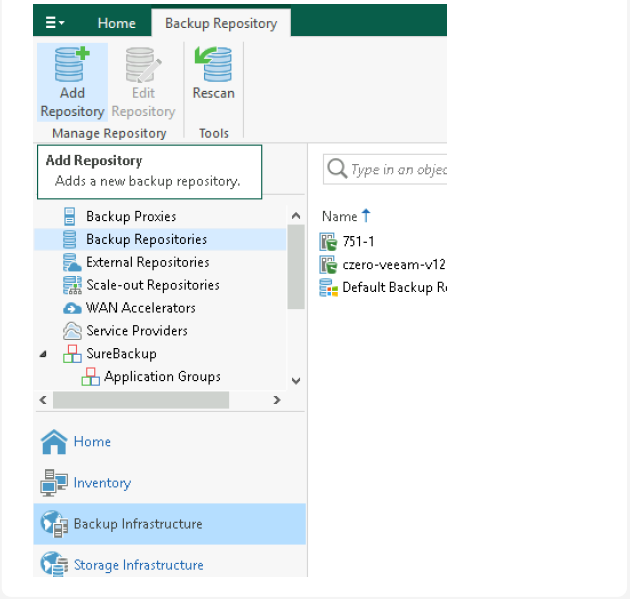

1. Add Repository
2. Choose “Object storage”
3. Choose “S3 Compatible”
4. Assign a name to the Object Storage Repository
5. Enter the S3 endpoint URL in the field “Service Point”
6. Add credentials and use the access and security key of the created IAM user
7. Choose the bucket you created and create folder inside the bucket
8. Finish the wizard with defaults

Configure Access Permissions of Object Storage Repository

1. Do a right click on the created repository and click on “Access Permissions”

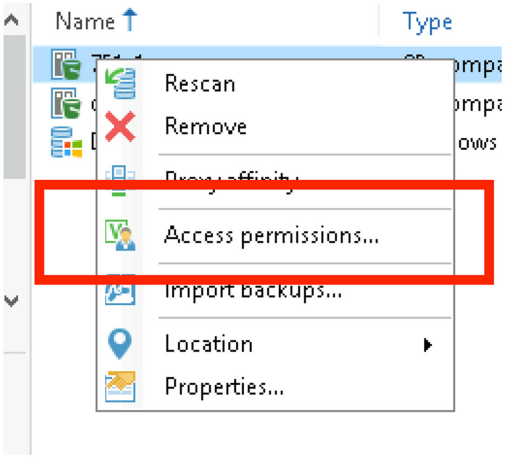

2. In the “Access Permissions” screen choose “Secure direct access” and provide your IAM and STS endpoints in the form: iam.cloudian.local:16443 (no leading “https://”)

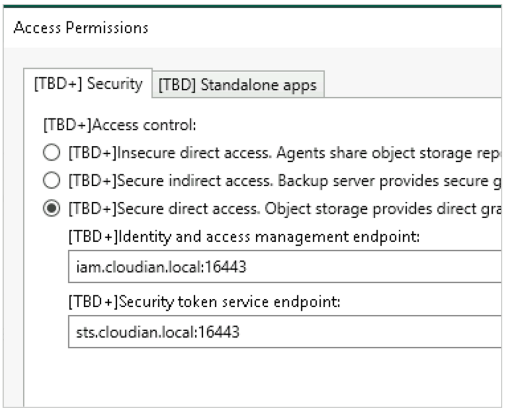

3. On the tab “Standalone apps” choose “Allow to everyone”.

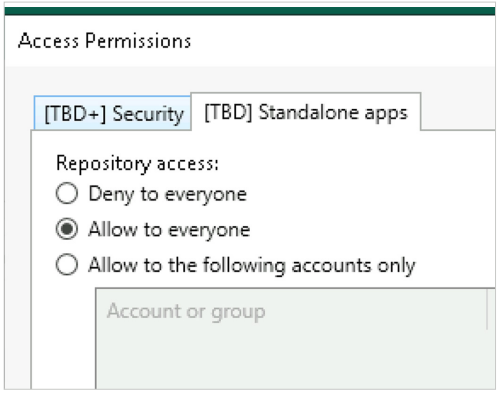

4. Confirm your settings by pressing “OK”.
5. Now you are ready to deploy your Veeam Agent on a Linux or Windows host.

Configure Veeam Agent (Windows/Linux)

Follow the steps in the Veeam Agent installation and configure the Agent to connect to your local VBR server.

When the agent asks to pick a destination choose the Veeam Backup Repository.

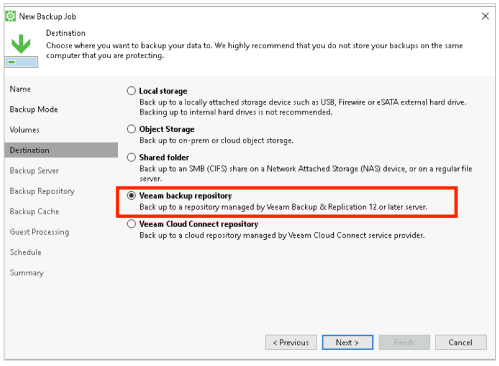

1. Choose the Repository you have created in the steps before.

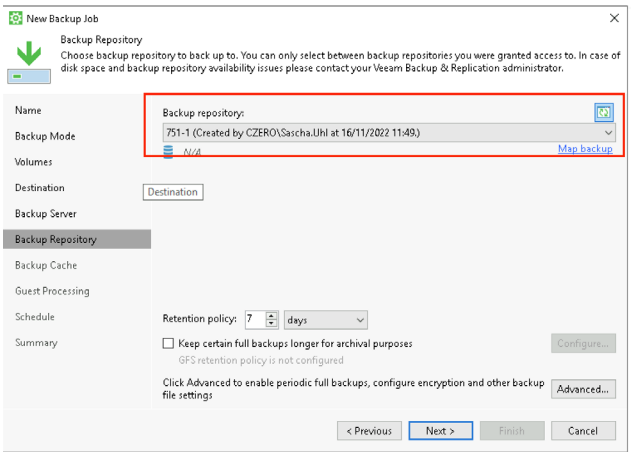

2. Click “Advanced” to configure “Local Target Large Blocks” in the storage options and confirm with “OK”.

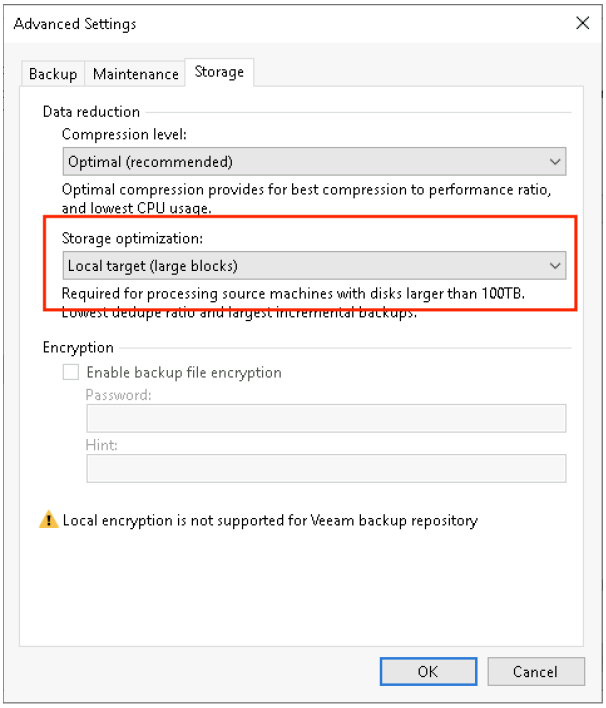

3. Finish the wizard with the defaults and start the job.

Veeam Backup And Replication Configuration

Veeam Backup & Replication

Veeam Version/Patches

Veeam 11

Minimum version has to be VBR version 11a CP4 (11.0.1.1261 20220302) which brings a special fix which is important when using object lock.

Veeam 12

Latest available

S3 Repository Configuration

Repository Tasks

To throttle the overall connections opened to the HyperStore cluster the total number of Repository Tasks across all S3 Repositories should be limited. As a rule of thumb you can calculate the total number of available Repository Tasks which can be used like this:

(100 x Number of Cloudian Nodes per site) / 64 = Task Limit

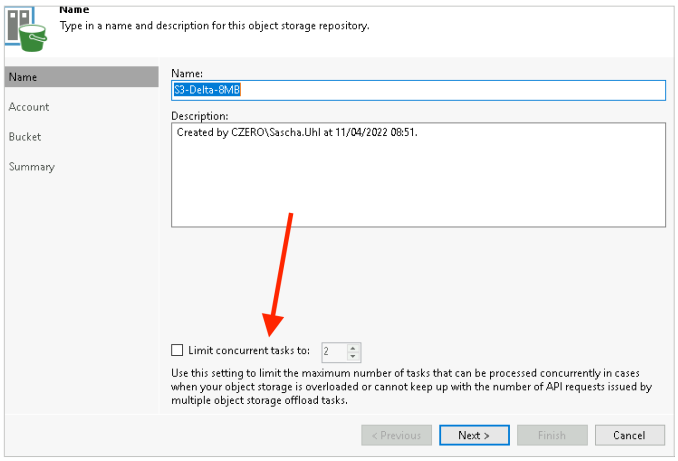

If more Tasks are required the amount threads used by Veeam Backup and Replication has to be reduced to e.g. 32. This can be done with a registry key:

`HKLM\SOFTWARE\Veeam\Veeam Backup and Replication\S3ConcurrentTaskLimit`

- On the Veeam server, open regedit in C:\Windows and locate the

HKEY_LOCAL_MACHINE > Software > Veeam > **Veeam Backup and Replication** entry.

- Add a new DWORD (32-bit) value entry with the name **S3ConcurrentTaskLimit**. When creating this entry, make sure that there are no blanks before or after the **S3ConcurrentTaskLimit** name.
- Right click on the new registry entry and click **Modify**. Ensure that the value is decimal.
- Close Backup & Replication and **reboot** the system to ensure that this registry entry is accepted.

When the registry key is applied on the VBR server it will be pushed out automatically to all Extents/GW.

With the new value a the available Tasks can be calculated: (100 x Number of Cloudian Nodes per site) / 32 = Task Limit.

Immutability Retention

Period length

Immutability retention period should be equal or shorter than the defined backup retention period.

Veeam's ransomware protection via object lock is limited to 90 days within the Veeam GUI. This is for a few purposeful reasons:

- Recovering from ransomware would rarely see a restore image older than 90 days
- If a customer "fat-fingered" a three digit number this would result in much longer than desired retention
- Remember, this works with AWS S3 and Service Providers. So accidentally submitted long periods would result in unwanted costs in those public cloud entities

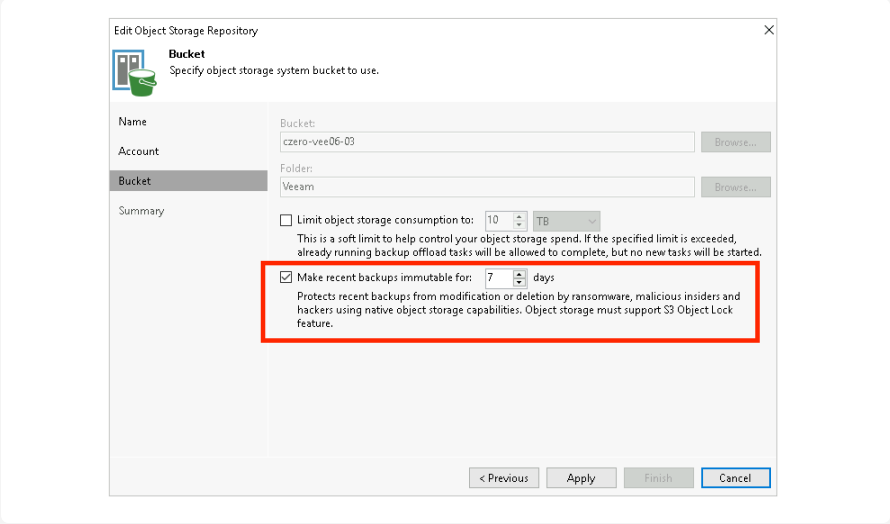

Cloudian recommends to stay in the 90 day retention period. The longer the retention period the more storage is required and will also generate excessive amounts of s3 requests in the future as metadata needs to be updated.

A longer retention period can be set via powershell. Consult with Veeam on how to do this.

Block Generation

To reduce I/O operations and associated costs, Veeam Backup & Replication will automatically add from 1 to 10 days to the immutability expiration date. This period is called Block Generation. You do not have to configure it, the Block Generation setting is applied automatically.

For example, if you set your immutability period to 30 days, Veeam Backup & Replication will add from 1 to 10 days to specific objects to reduce I/O operations with the storage over time. This will not change the retention and their effective immutability. It is a background optimization. Thus, if you need 30 days immutability period, set the period to 30 days.

More details can be found here: https://helpcenter.veeam.com/docs/backup/vsphere/block_gen.html?ver=110

SOBR Configuration

Veeam Backup and Replication Version 11

Veeam VBR 11 supports one object storage bucket per SOBR configuration. We recommend using multiple SOBRs pointing to different buckets. This will guarantee a better experience with the overall solution including better performance.

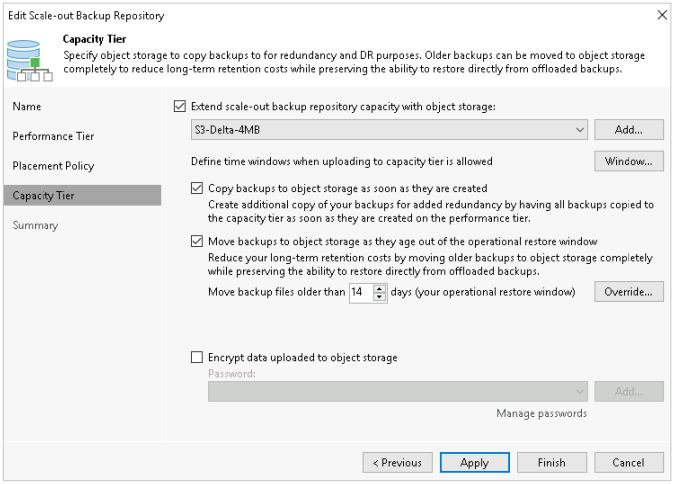

Specifically, deletes of large amounts of data at the same time can cause millions of S3 delete requests generated by Veeam. Using multiple SOBR allows it to distribute the requests which lowers the load on the S3 protocol.

Veeam Backup and Replication Version 12

Starting with Veeam V12 multiple buckets are supported in Performance and Capacity Tier. There is no limit to number of buckets which can be used. For customers upgrading to Veeam V12 or new customers who want to use the traditional SOBR configuration with Cloudian HyperStore in the Capacity Tier, multiple buckets should be configured. Veeam will balance the VMs automatically over the configured buckets. New buckets can be added afterwards, too, but Veeam would not rebalance existing VMs to the new buckets. Only new VM backups will go to the new buckets. Having multiple buckets in the Capacity Tier eliminates the need to have multiple SOBRs.

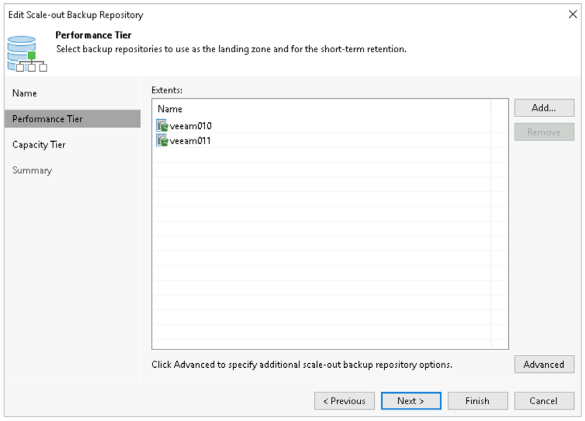

Important:

- New/separate bucket(s) should be used for every SOBR and workload (like VBO/Office365 Backup, Archive backup jobs).
- Different retention policies can be set per SOBR.
- In order to use a SOBR a Veeam Enterprise (up to 2 SOBR) or an Enterprise Plus (unlimited SOBRs) license is required.
- HyperStore has no limit on the amount of data in a bucket. But from Veeam performance point of view it can be beneficial to use multiple SOBR.

Backup Job Configuration

Storage Optimization Settings

Veeam allows you to configure a block size per Backup Job. The block size affects deduplication and incremental backup size. By default blocks are compressed and compression ratio is about 50%. The smaller the block size, the more calls are needed to the object storage to upload the data.

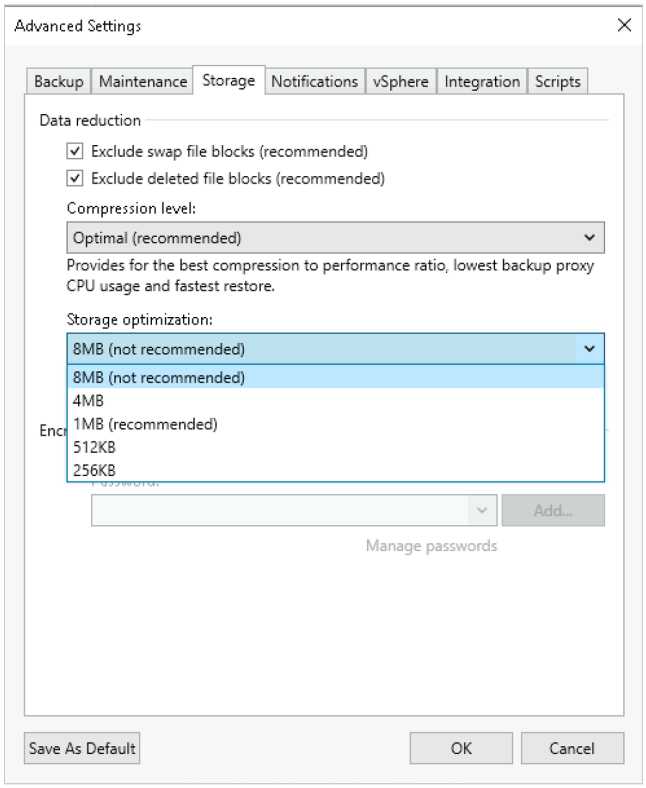

Additionally larger block size will increase the throughput to and from the Cloudian HyperStore system.

Cloudian recommends using “Local Target (large blocks)” (equal to 4MB or 2MB after compression object size) or even “Local Target (extra large blocks)” (equal to 8MB or 4MB after compression object size).

Changing the block size may result in larger storage consumption. Consult with your Cloudian SE on this aspect.

Block size changes only take effect after an "active full" backup is performed (Synthetic Full backup will not change block size).

How to configure in Veeam:

https://helpcenter.veeam.com/docs/backup/hyperv/backup_job_advanced_storage_hv.html?ver%20=100&ver=110

Enable "Extra Large Blocks" or "8MB" Blocks for Storage Optimization

To enable the "Extra Large Blocks" in the console GUI you have to enable this via a registry key:

1. Go to "Computer\HKEY_LOCAL_MACHINE\SOFTWARE\Veeam\Veeam Backup and Replication"
2. Add a value "UIShowLegacyBlockSize" with DWORD 32bit value "1".

Save as Default

In order to save the Storage configuration as a default for Backup Jobs created from now on you can save the new settings as a default. Simply click "Save as default" on the bottom of the Advanced Settings screen.

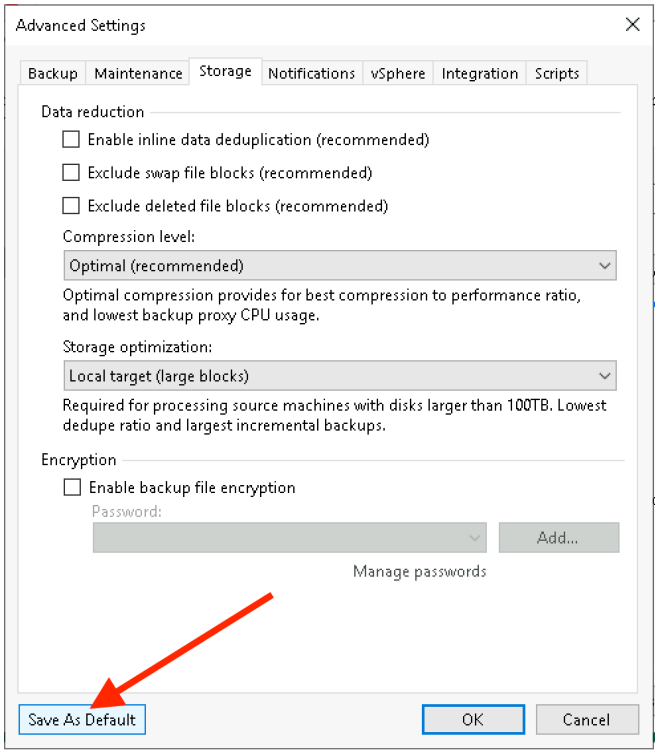

Veeam Backup For Windows 365

Cloudian HyperStore is the ideal target to store Windows 365 backups. Backups can be directly written to S3 with the need of large direct attached storage. To back up data to object storage, you can extend a backup repository with Cloudian HyperStore and map a backup job to such an extended repository.

Object Sizes

The backup data is grouped into quite large blocks. Exchange (365 Mail) is stored in approx. 5MB blocks, Sharepoint and OneDrive data are stored in approx. 8MB blocks.

Veeam Patches

Use the latest version of Veeam Backup for Windows 365.

Create own bucket/container for each repository with own user and restrictive ACL. Separate buckets/containers will also make sure to keep the failure domains smaller and allow for easier implementation of different retention settings.

Repository Configuration

Cloudian HyperStore is configured as a “S3 Compatible Object Storage” repository. Veeam has documented the setup process here:

https://helpcenter.veeam.com/docs/vbo365/guide/adding_s3compatible.html?ver=60

Backup Job Configuration

Veeam has documented the Backup Job configuration process here:

https://helpcenter.veeam.com/docs/vbo365/guide/vbo_new_backup_job.html?ver=60

Configuration Summary Table

Cloudian HyperStore Configuration

The table below provides a summary of the Best Practice recommendations. For a detailed procedure, see the previous section.

Cloudian HyperStore Settings	
Cloudian HyperStore Version	7.4.2 (7.5.1 for support of IAM/STS Agent Backup)
HyperStore Shell Enabled	Yes
Tune for deletes	Change `cloudian.s3.batch.delete.delay` and optional `cloudian.delete.dc.instances` in mts.properties
S3 over HTTPS	Yes
Storage Policies Settings	
Data Distribution Scheme	RF3, EC 4+2 or EC 6+3
Workload Type	“VEEAM” for Capacity Tier ”VEEAM12” for Performance Tier “VEEAMO365“ for Backup for Windows 365
New Storage Policy for every Veeam bucket	Yes, for optimal all around operational performance
Bucket Settings	
Versioning	Disabled
Life Cycle Policy	No
Tiering	No
Object Lock	Yes, when “Veeam immutability” is required

Veeam Configuration

The table below provides a summary of the Best Practice recommendations. For a detailed procedure, see the previous section.

Veeam Backup & Replication 12

Product Version	Latest
Backup Job - Storage Optimization Settings	"4 MB" or "8MB" (enable via registry key)
Immutability Retention Period	Should be < or = to backup retention period
S3 Repository Task Limit	The Task limit needs to be calculated based on amount of Cloudian HyperStore nodes. See corresponding chapter for details.
SOBR Configuration	Create a bucket for every 50 VMs (Performance or Capacity Tier, depending on the use case)

Veeam Backup & Replication 11

Product Version	11a CP4 (11.0.1.1261 20220302) or later
Backup Job - Storage Optimization Settings	"Local Target (large blocks)" or "Extra Local Target (large blocks)"
Immutability Retention Period	Should be < or = to backup retention period
S3 Repository Task Limit	The Task limit needs to be calculated based on amount of Cloudian HyperStore nodes. See corresponding chapter for details.
SOBR Configuration	Use a separate SOBR for every 200TB of Veeam Backup data.

Veeam Backup for Windows 365

Product Version	Latest
Immutability	No immutability support
Versioning	No versioning support

Conclusion

This guide detailed the key considerations for the deployment and performance tuning of Cloudian HyperStore when used as an S3 storage target within the Veeam architecture as well as detailed deployment examples to validate the overall solution functionality.

Note that Veeam and in particular, the scale-out repository performance improvement is very dependent on the overall environment and is sensitive to specific end-to-end changes. Further configuration modifications may be required that are outside the scope of this document. As discussed earlier, it is highly recommended to plan and execute a performance baseline with an appropriately size test source dataset before and after tuning in order to compare the results. For maximum benefits, tuning may require several iterations for best performance. Please consult with Veeam and/or Cloudian for further guidance and details.

Document Version History

Version	Description
1.0	Applying new format
2.0	Veeam VBR 12 and VBM365 added

Appendix A: Troubleshooting & Common Errors

Use immutability with existing Cloudian/Veeam deployment

The Veeam immutability feature requires the AWS S3 feature “Object Lock”. Cloudian HyperStore supports this feature with Version 7.2 or greater.

To use the object lock feature you need to enable it during bucket creation. It cannot be enabled on an existing bucket.

If you want to use immutability on an existing SOBR configuration you need to use the “seal” function. The bucket will be sealed and no new data will be written in it. You can add a new bucket to the SOBR configuration which has object lock enabled. New data will be written to this immutable bucket.

How to configure in Veeam V10 V11

Appendix B: How to create a bucket Cloudian HyperStore

1. Go to your Cloudian Management URL and login with user credentials offered by your system administrator

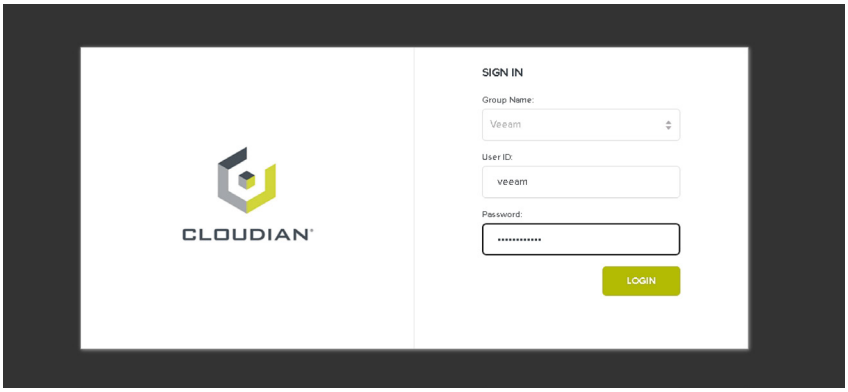

2. You will be prompted with an overview on the already created buckets. Click “Add new Bucket“ to start creating a new bucket.

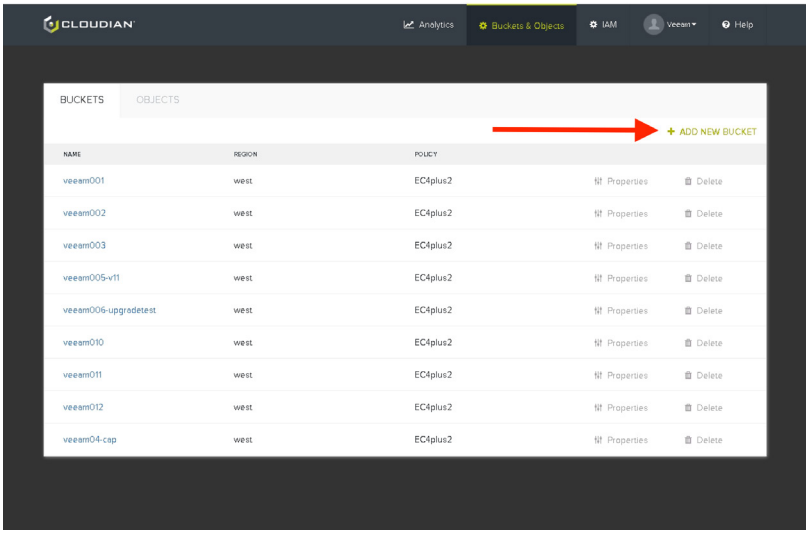

3. Enter the name of the bucket, the region and storage policy created for the Veeam workload.

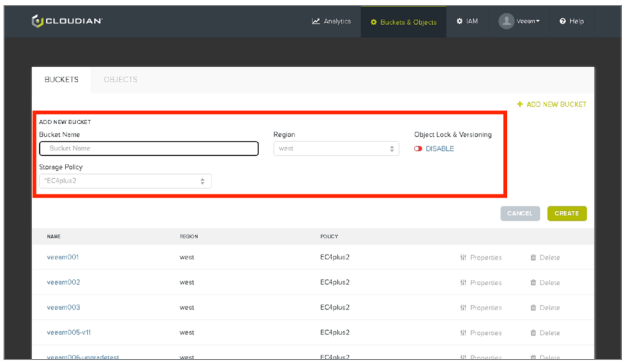

4. If you want to use immutability to protect the bucket against Ransomware you need to enable “Object Lock & Versioning”.

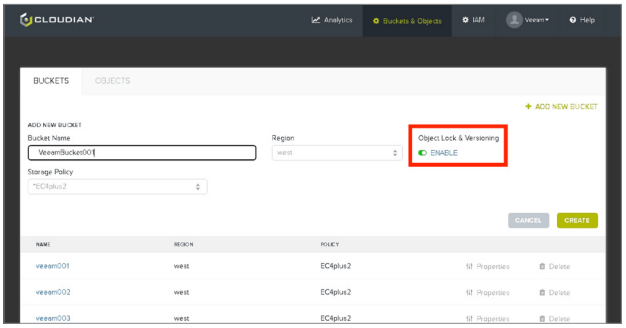

5. Once everything is configured properly press “Create” to create the bucket.

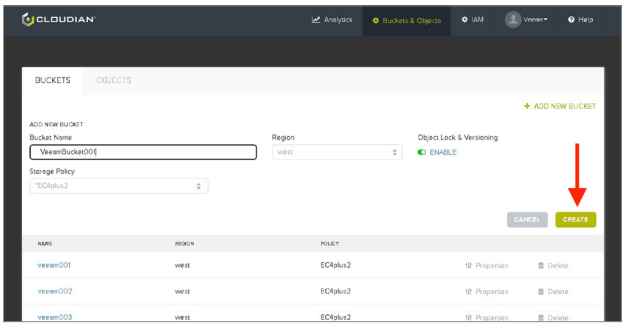

Appendix C: How to get S3 credentials and endpoint

1. After you logged in to the Cloudian Management Console you can access the Security section by click on the username.

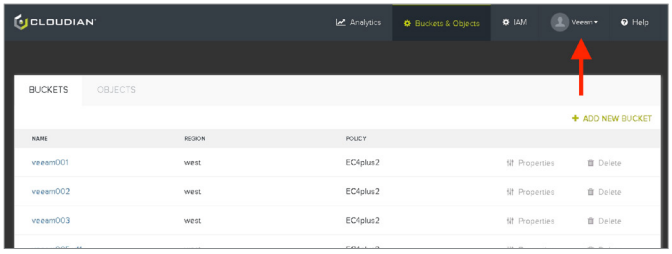

2. Click on “Security Credentials”

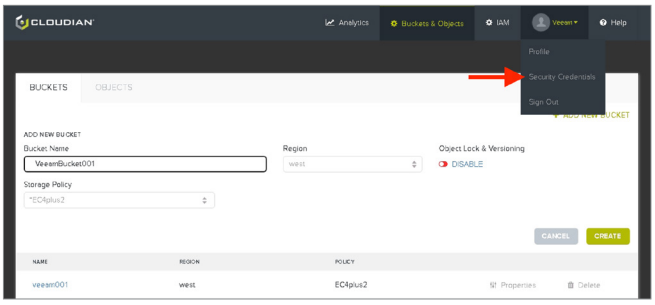

3. In the security screen you can see your “Access Key”, “S3 Endpoint (HTTPS)”.

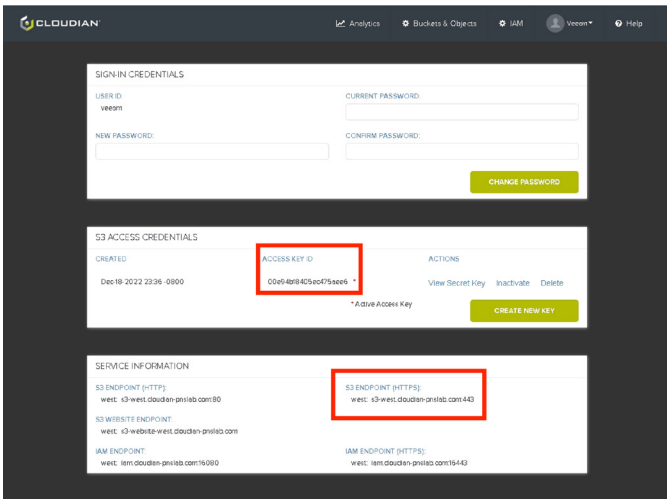

4. To see the “Security Key” you need to click on “View Secret Key”.

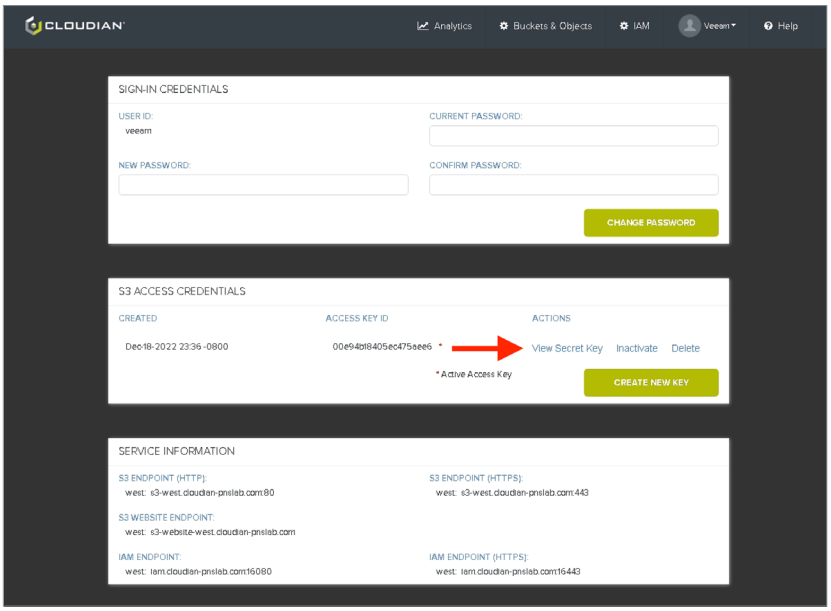

5. The Security Key will be shown in pop-up on top of the screen.

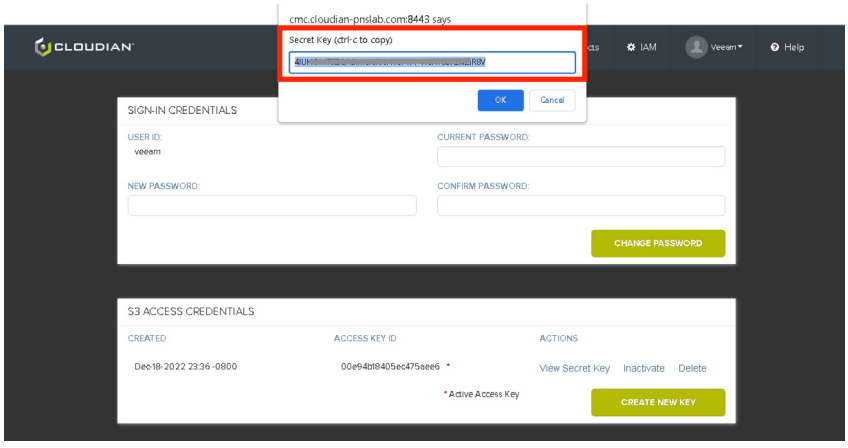

Appendix C: How to create a S3 repository in Veeam Console

1. Open the Veeam Console and go to “Backup Infrastructure”.

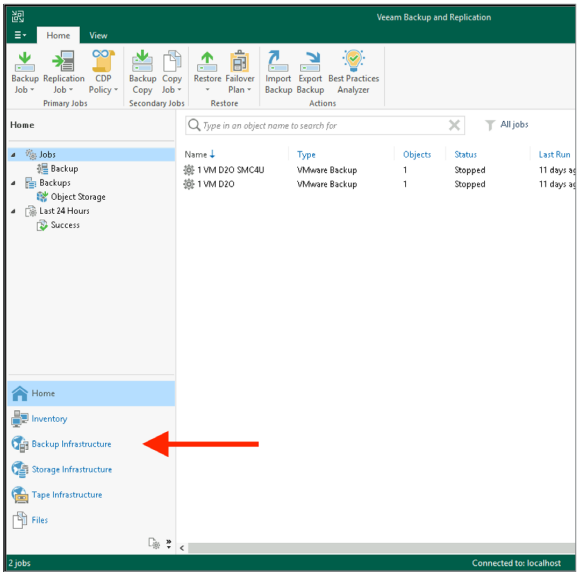

2. In “Backup Infrastructure” click on “Backup Repositories”

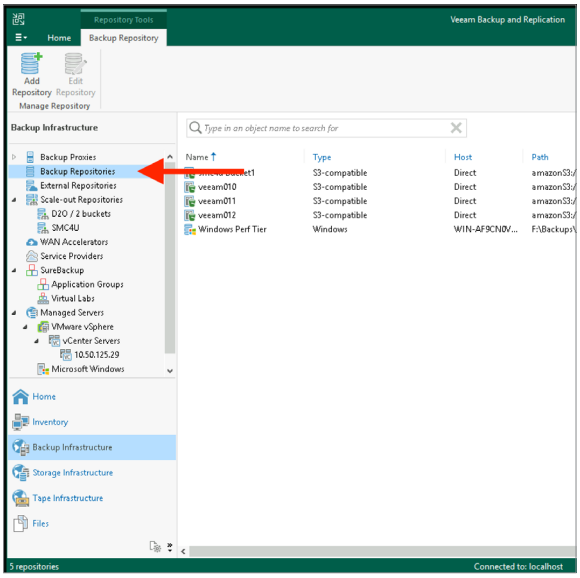

3. In “Backup Repositories” click on “Add Repository” to open the wizard.

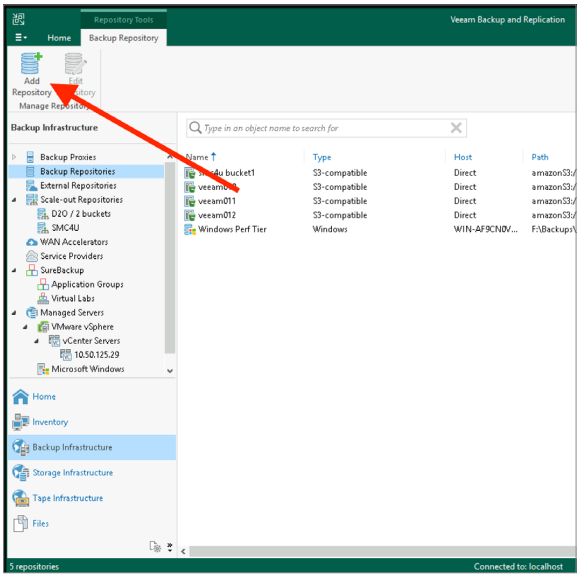

4. Click on “Object Storage”.

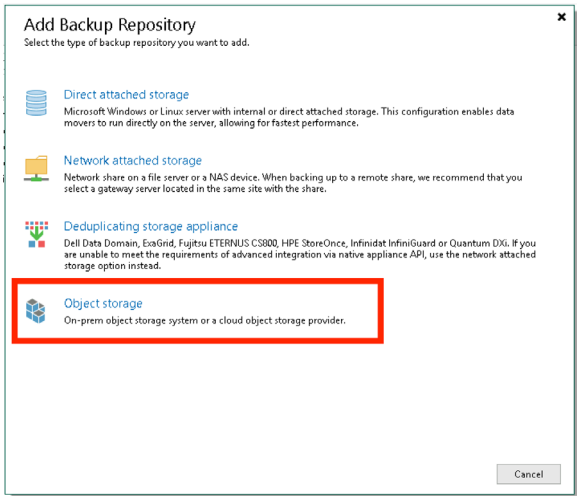

5. Click “S3 Compatible“ to open the wizard.

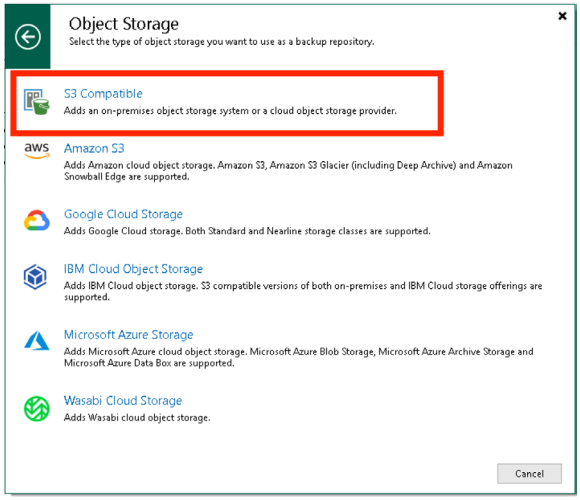

6. Enter a name for the bucket and set the “Limit concurrent tasks” value as described in the “S3 Repository Configuration“ section.

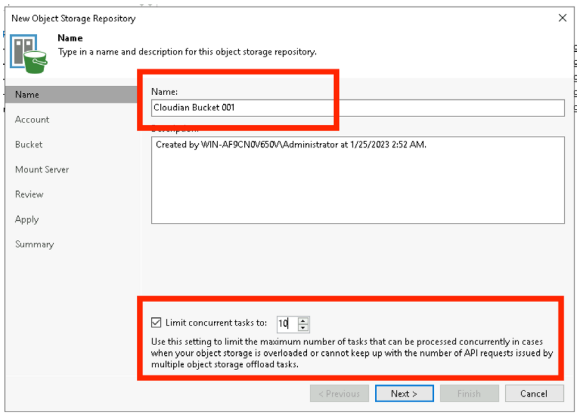

7. Enter your S3 Endpoint URL and ignore the region name.

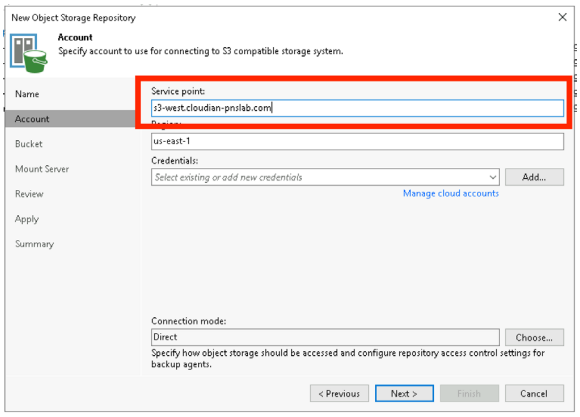

8. Add new credentials by click on the “Add…“ button.

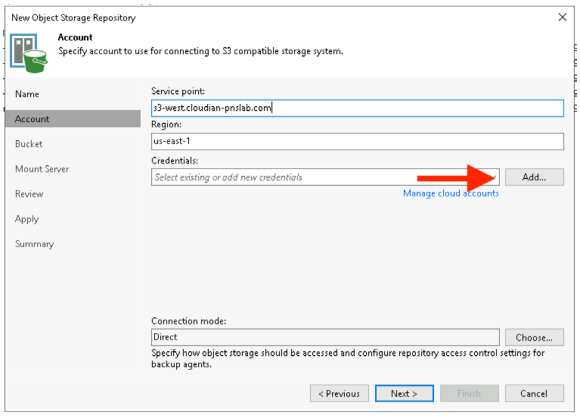

9. Enter your Access Key and Security Key.

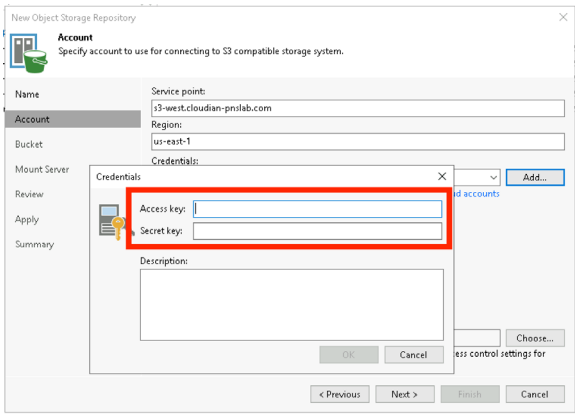

10. If you have already added another bucket you can pick the credentials from the dropdown.

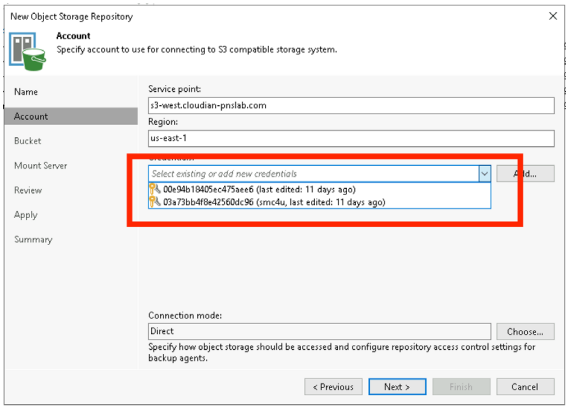

11. Ensure “Direct” connection mode is used and press “Next”

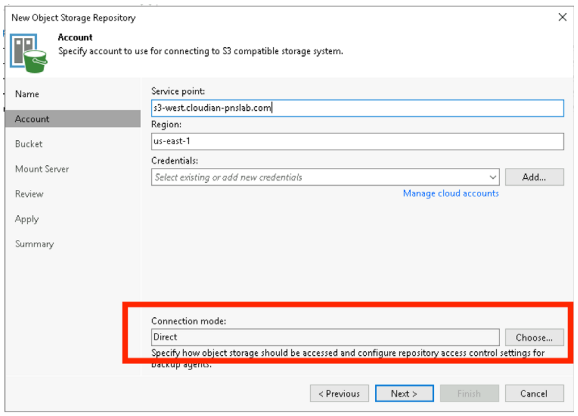

12. Press “Browse….” to choose a bucket.

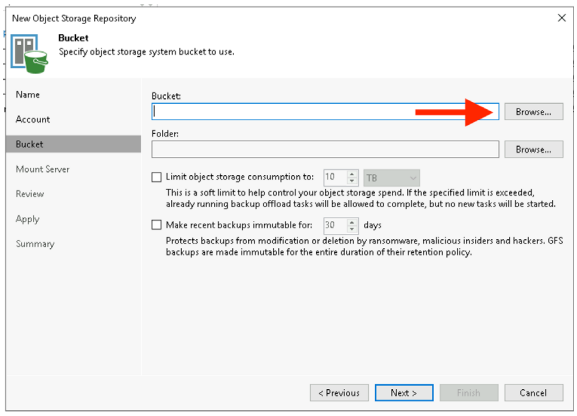

13. In the list of buckets pick the one you created for your Veeam workload and press “OK”.

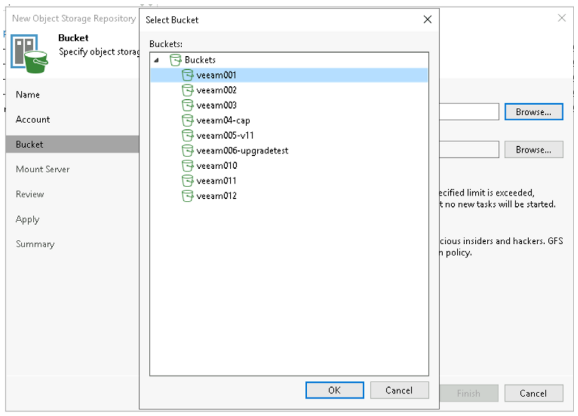

14. Press “Browse…” to select a folder into which the data should be stored.

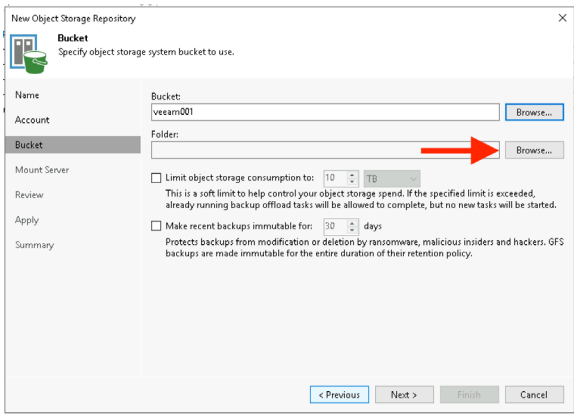

15. Select or create a folder by click on “New Folder” .

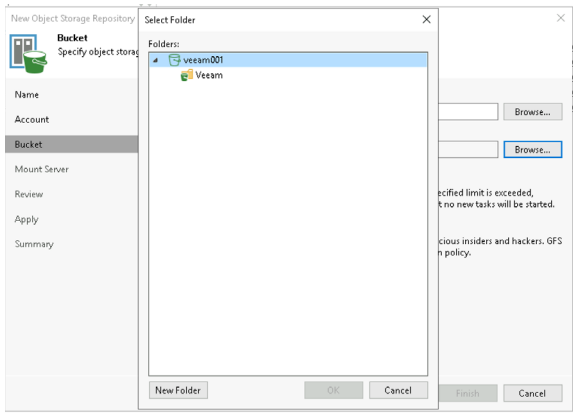

16. Optionally protect backups against ransomware attacks by setting a retention time.

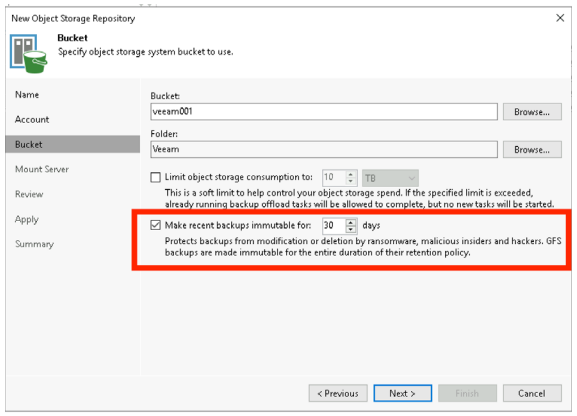

17. Skip the “Mount Server” page or adjust if required and press “Next”.

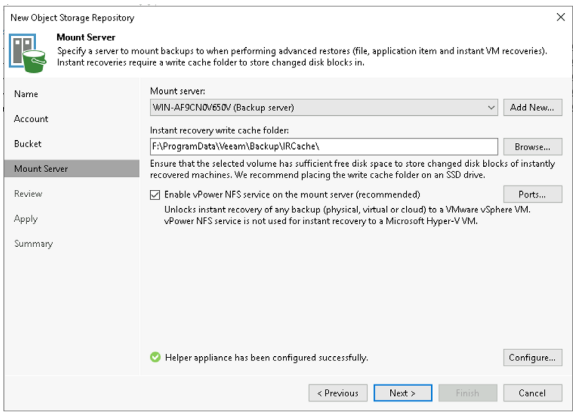

18. Review your config and press “Apply”.

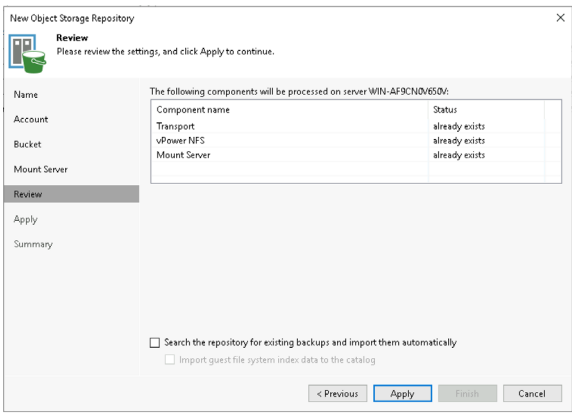

19. Repository will be created. Once done press “Next” to continue.

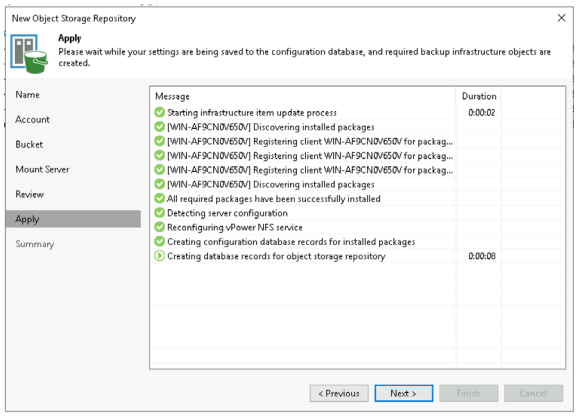

20. Press finish to end the wizard.

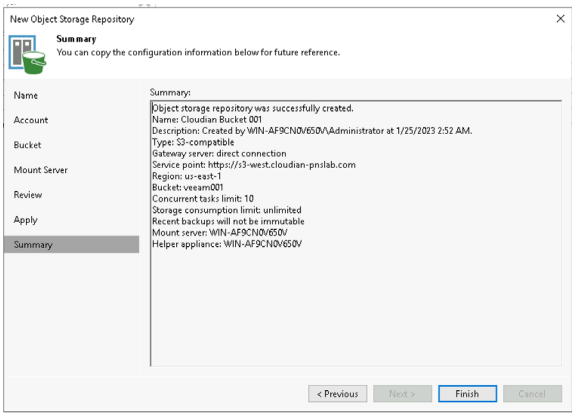

Appendix D: How to create a SOBR in Veeam Console

SOBR with bucket(s) in the capacity tier

1. Open the Veeam Console and go to “Backup Infrastructure”.

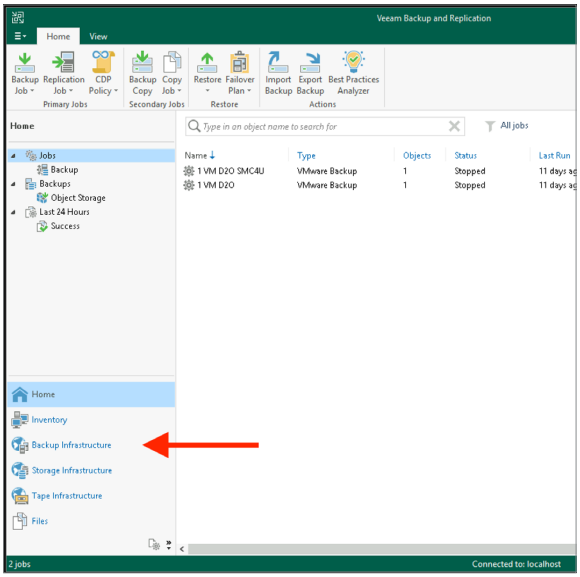

2. In “Backup Infrastructure” click on “Scale-Out Repositories”.

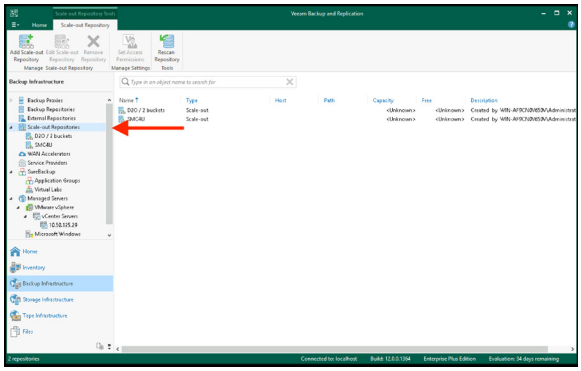

3. Click on “Add Scale-out Repository“ to open the wizard.

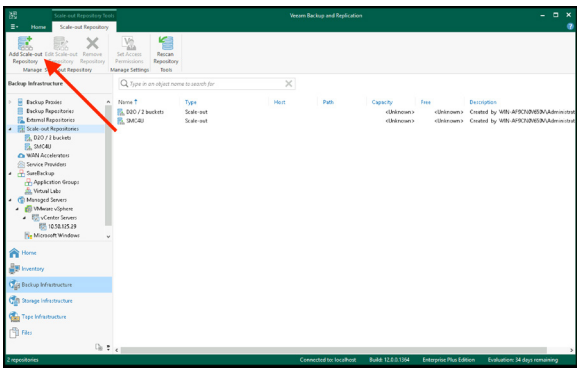

4. Enter a name for the Scale-out Repository you want to create and press “Next”.

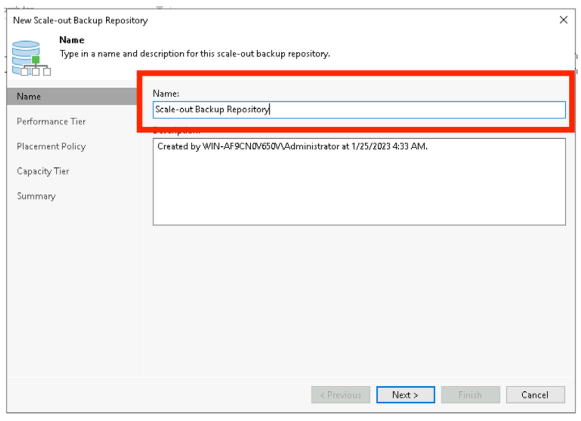

5. Click “Add…” to a repository to the Performance Tier.

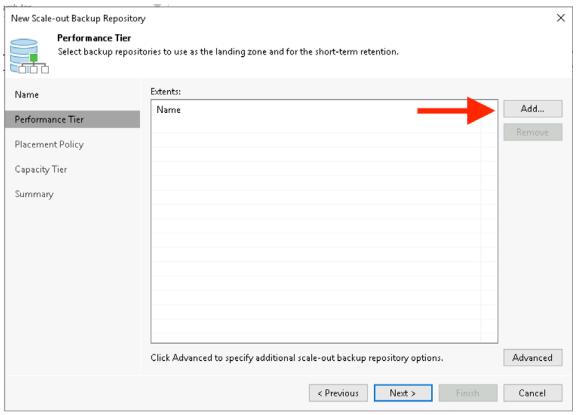

6. In a traditional SOBR the Performance Tier will be handled by direct attached (non-S3) repository. Click “OK” to confirm your selection.

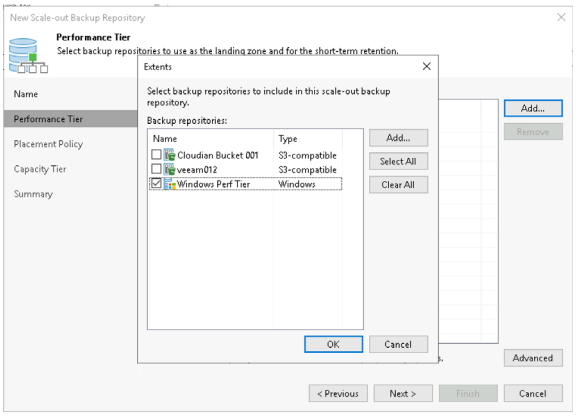

7. Click “Next” to proceed.

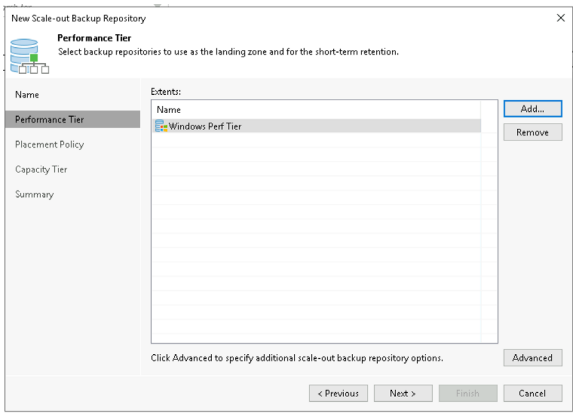

8. Enabling the offloading to S3 by checking the checkbox.

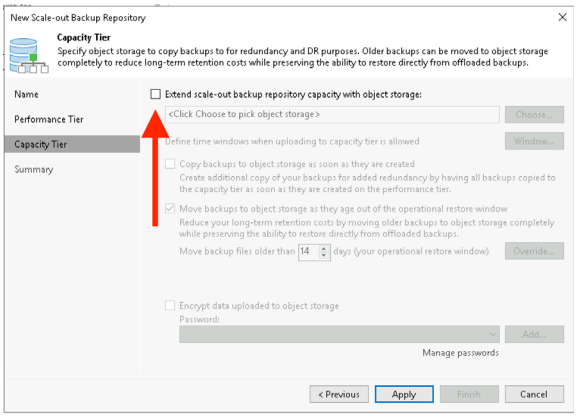

9. Click “Choose…” to select the target bucket(s).

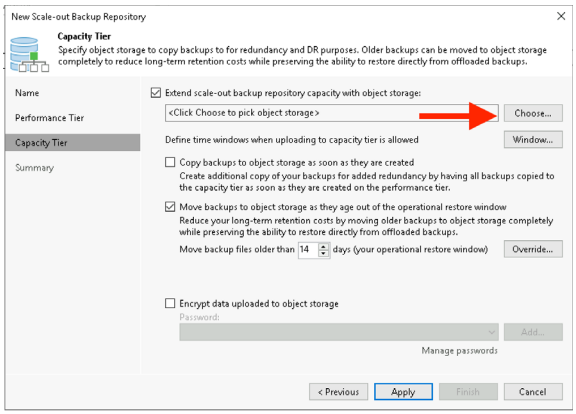

10. Select one or multiple buckets and confirm by clicking “OK”.

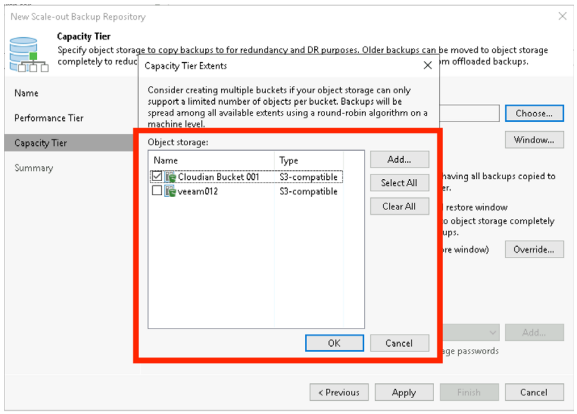

11. Check “Copy backups to object storage as soon as they are created” if you want to protect all backups against ransomware attacks by copying them to the Cloudian HyperStore system as soon as they are created. Press “Apply” to proceed.

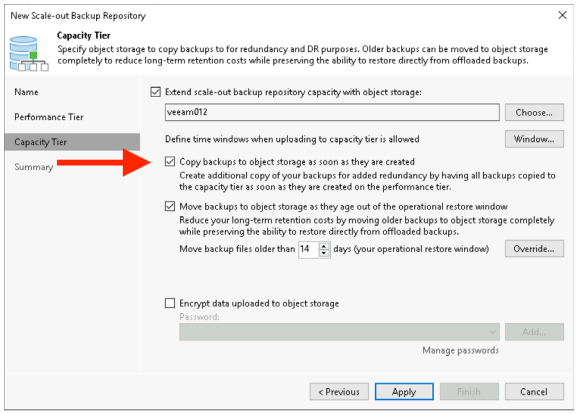

12. Press “Finish” to close the wizard.

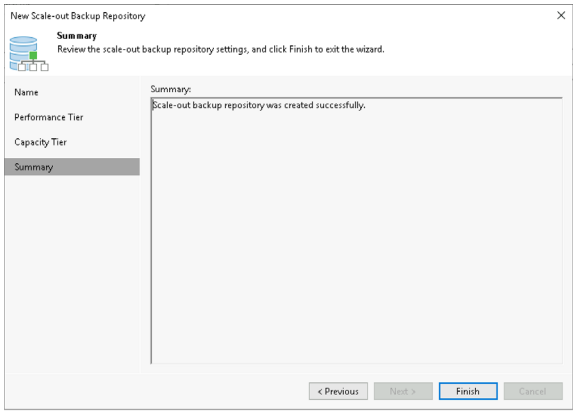

SOBR with bucket(s) in the performance tier

1. Open the Veeam Console and go to “Backup Infrastructure”.

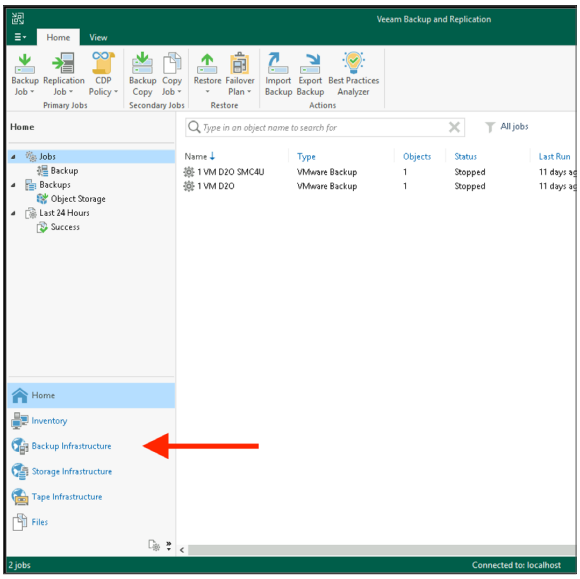

2. In “Backup Infrastructure” click on “Scale-Out Repositories”.

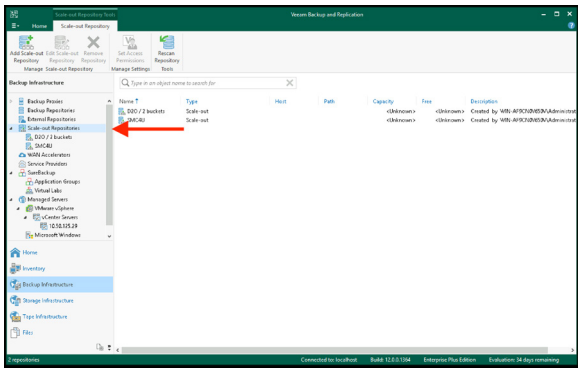

3. Click on “Add Scale-out Repository“ to open the wizard.

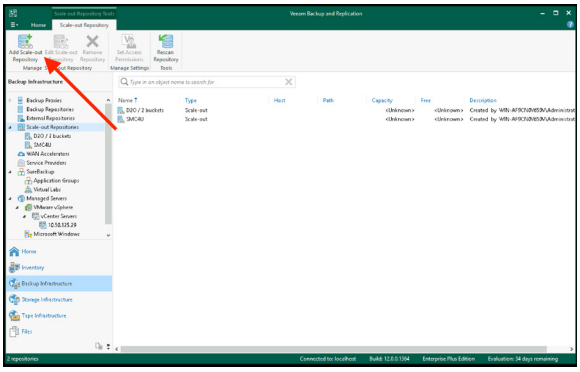

4. Enter a name for the Scale-out Repository you want to create and press “Next”.

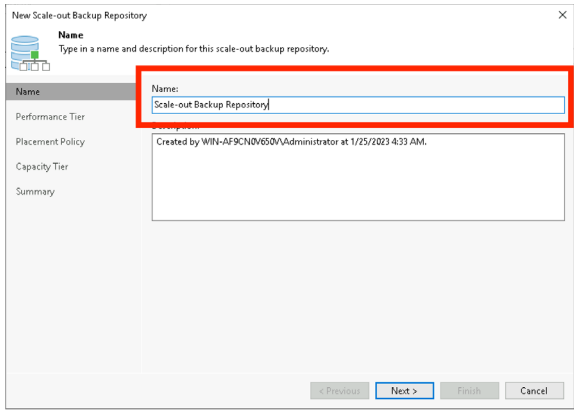

5. Click “Add…” to a repository to the Performance Tier.

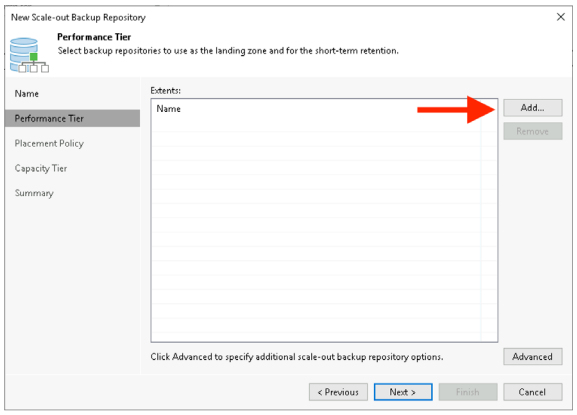

6. Select one or multiple buckets.

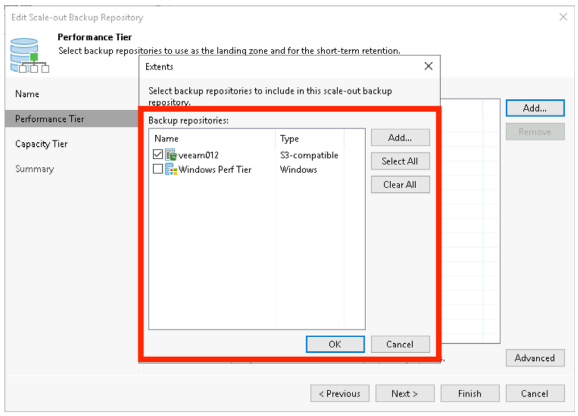

7. Click “Next” to proceed.

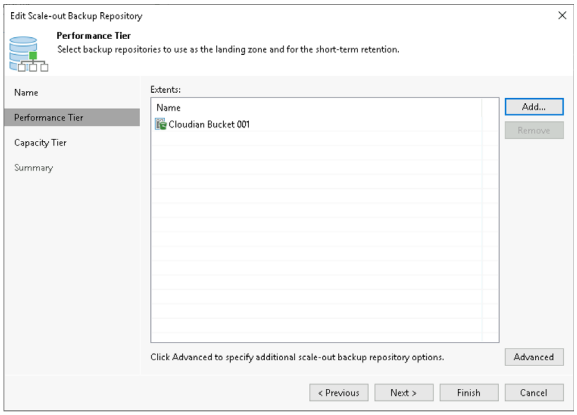

8. Skip the Capacity Tier config by clicking “Apply”.

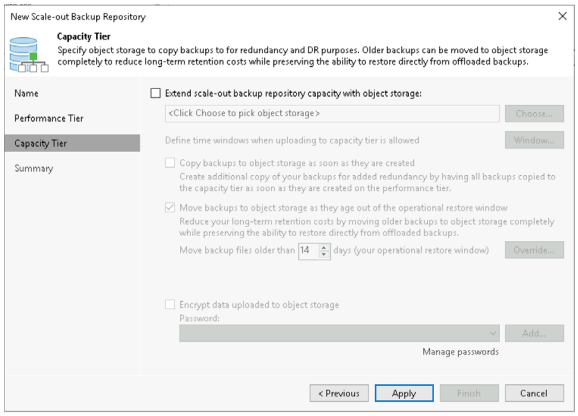

9. Press “Finish” to close the wizard.

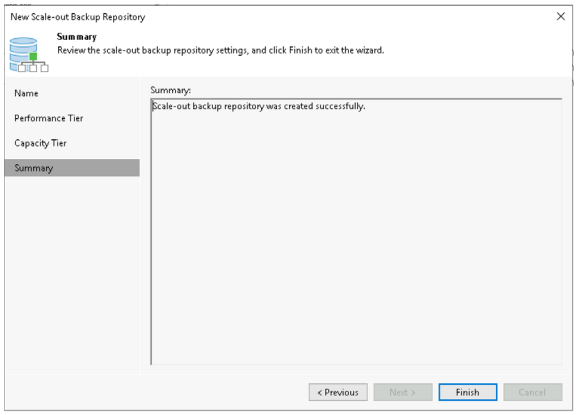

Appendix E: How to create a Backup Job in Veeam Console

1. Open the Veeam Console and go to “Home”.

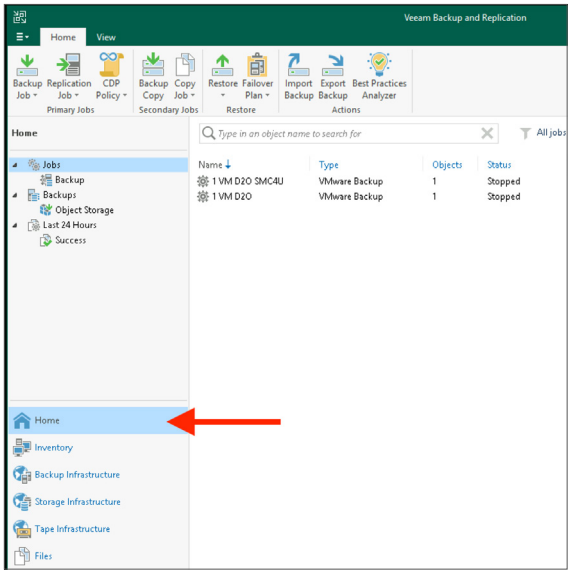

2. Click “Backup Job” and “Virtual Machine” to create a backup job to backup VMs.

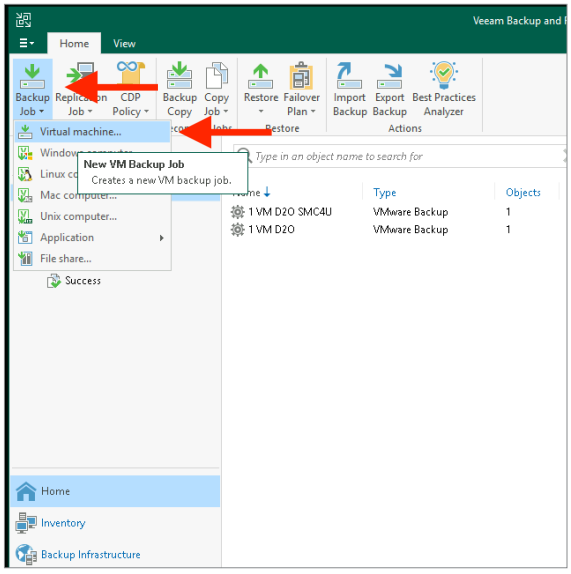

3. When the wizard has started enter a name for the backup job and click “Next“.

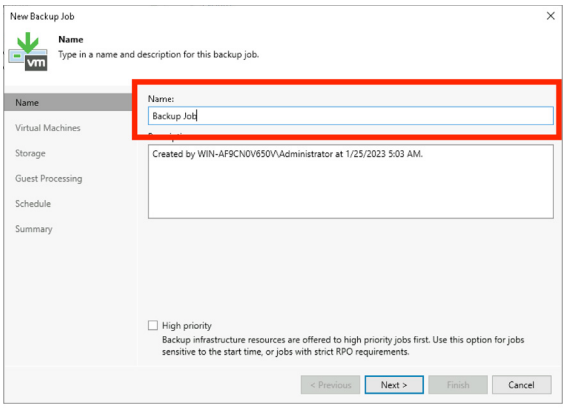

4. “Add..” to add virtual machines to the backup job.

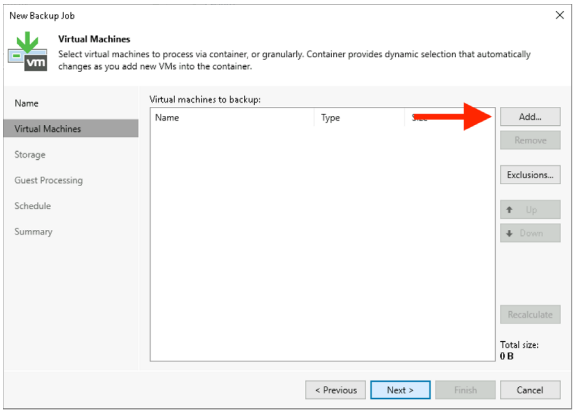

5. Select one or multiple VMs you want to include in the backup job and confirm your selection with “OK”.

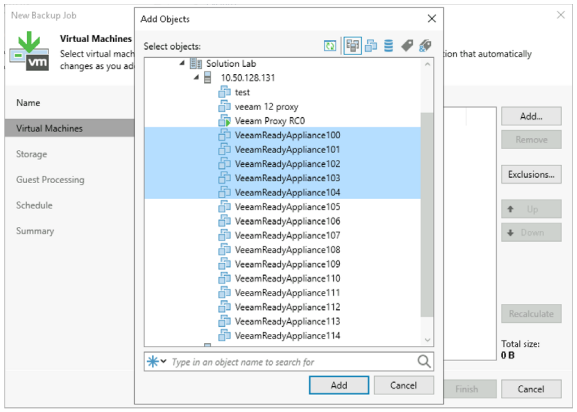

6. Press “Next” to continue.

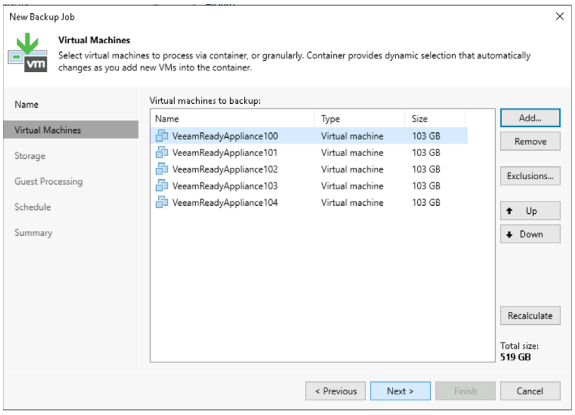

7. Select the “Backup repository” by choosing the created SOBR from the dropdown list.

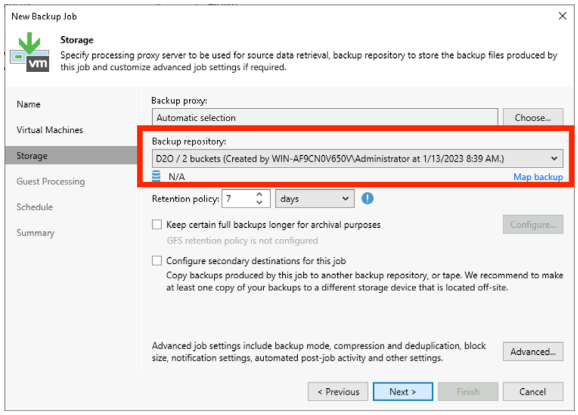

8. Press “Advanced…” to adjust the backup job settings.

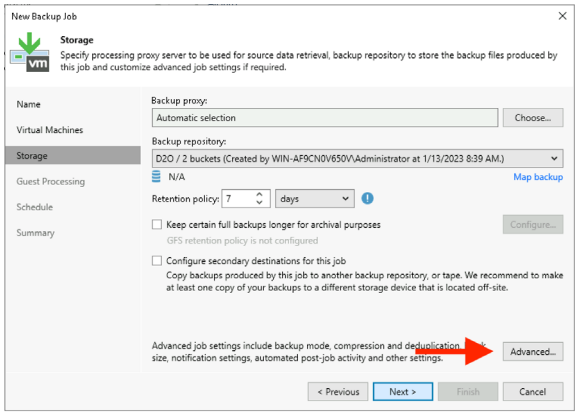

9. Click “Storage” to adjust the block size settings.

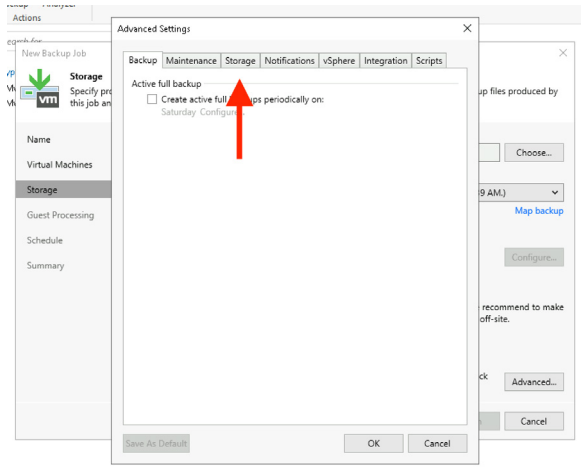

10. Change the “Storage optimization” settings to 4 or better 8 MB and confirm with clicking “OK”.

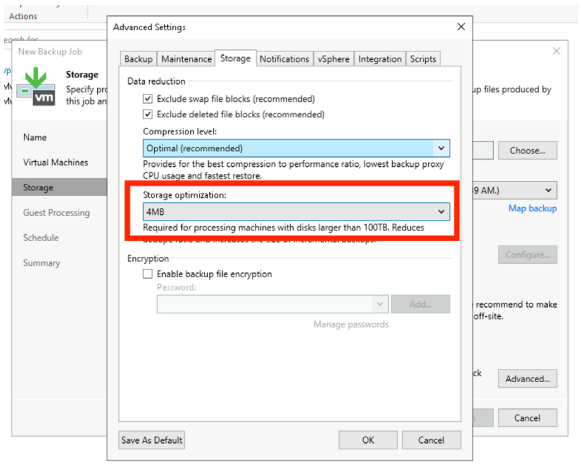

11. Press “Next” to continue.

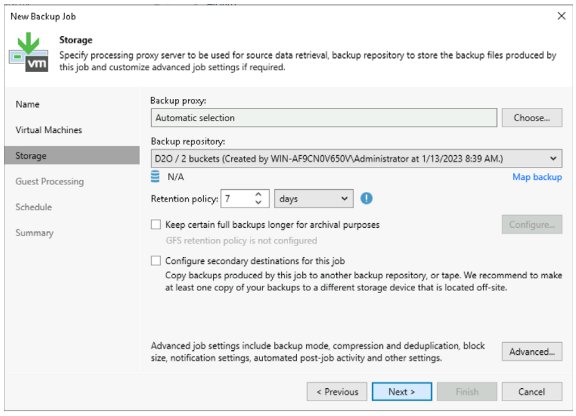

12. Change guest settings according to your requirements or just skip by clicking “Next”.

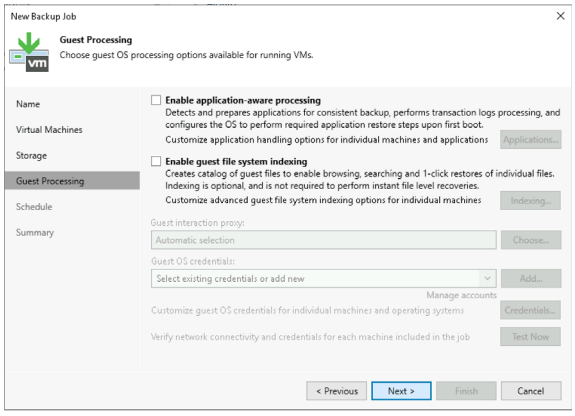

13. Configure your backup job schedule or skip by clicking “Next”.

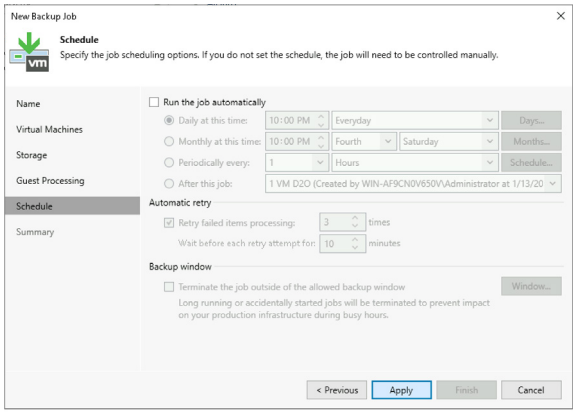

14. End the wizard by clicking “Finish”.

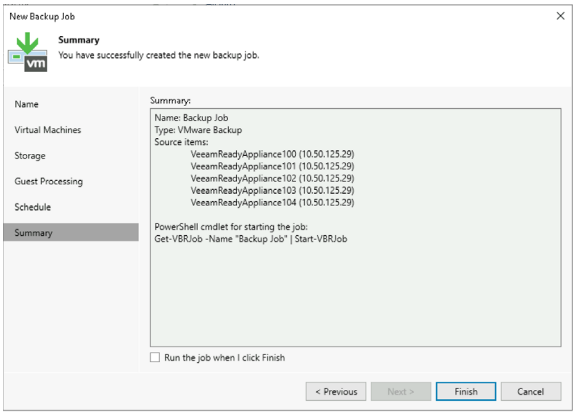

177 Bovet Road, Suite 450, San Mateo, California 94402
650.227.2380 | info@cloudian.com | cloudian.com