Veeam Backup & Replication Best Practices

V8 - 07/28/2021

Author: Sascha Uhl (sascha@cloudian.com)

This document explains best practices for Veeam deployments. Veeam is using a “SOBR” to tier data to Cloudian HyperStore.

CLOUDIAN, INC.
177 Bovet Road, Suite 450, San Mateo, CA 94402 Tel: 1.650.227.2380 | cloudian.com

©2020 Cloudian, Inc. Cloudian, the Cloudian logo, HyperFile, HyperScale, and HyperStore are registered trademarks or trademarks of Cloudian, Inc.
All other trademarks are property of their respective holders.

Changelog

Version	Changes
V6	Added change log Added information on number of concurrent connections in Veeam Added information on how to reduce immutability retention for POCs Updated Storage Policy considerations Updated Veeam Storage consumption
V7	Added note on pqlisting Added note on requirement of https New Storage Policy only required for large deployments Added notes on EC k+3 Max Bucket Size increased to 250TB Added PowerShell script to report SOBR offloading status Added reference to the Veeam Object Storage IO Sizing sheet Reworked Veeam Task Requirements Section
V8	Corrected reg key to reduce retention

CLOUDIAN, INC.
177 Bovet Road, Suite 450, San Mateo, CA 94402 Tel: 1.650.227.2380 | cloudian.com
©2020 Cloudian, Inc. Cloudian, the Cloudian logo, HyperFile, HyperScale, and HyperStore are registered trademarks or trademarks of Cloudian, Inc.
All other trademarks are property of their respective holders.

Useful Links

Cloudian HyperStore Deployment Guide with Veeam Backup & Replication Version 10 Immutability Feature:
https://data.cloudian.com/l/677273/2020-04-01/5ggyq/677273/63075/Cloudian_Veeam_Immutability_Feature_Deployment_Guide.pdf

Information on SOBR by Veeam:
Daily administration meets software-defined storage with the Scale-Out Backup Repository
https://www.veeam.com/blog/scale-out-backup-repository.html

Ultimate FAQ for Scale-out Backup Repository
https://www.veeam.com/blog/ultimate-faq-for-scale-out-backup-repository.html

Introducing the Veeam Scale-Out Backup Repository
https://www.youtube.com/watch?v=7ag6_Jm1uG4

Veeam Best Practices
https://www.veeambp.com/

CLOUDIAN, INC.
177 Bovet Road, Suite 450, San Mateo, CA 94402 Tel: 1.650.227.2380 | cloudian.com
©2020 Cloudian, Inc. Cloudian, the Cloudian logo, HyperFile, HyperScale, and HyperStore are registered trademarks or trademarks of Cloudian, Inc. All other trademarks are property of their respective holders.

Table of contents

Optimizations inside Veeam
- Summary
- Use multiple SOBRs
- Backup Job
  - Storage Settings
- SOBR Offloading
- V11 Task Throttling
- Number of concurrent connections (S3ConcurrentTaskLimit)
- Endpoint connection via HTTPS
- SOBR Offload Reporting
- Hints for Lab and POC
  - Archive Immutability Generation Days

Optimizations inside Cloudian HyperStore
- Summary
- Cloudian Version
- Storage Calculation
- Lifecycle Rules & Tiering
- Metadata Disk Size for Veeam Backup & Recovery
- Cloudian Sizing for Bulk Deletes
- Bucket Settings
  - Erasure Coding
  - Storage Policies
  - Change metadata protection
    - Replication Factor
    - (Replicated) Erasure Coding
    - Distributed Erasure Coding
  - Versioning
  - Bucket Folder Depth (Cassandra)
- PQLISTING

SECURITY
- Bucket usage
- Immutability

CLOUDIAN, INC.
177 Bovet Road, Suite 450, San Mateo, CA 94402 Tel: 1.650.227.2380 | cloudian.com ©2020 Cloudian, Inc. Cloudian, the Cloudian logo, HyperFile, HyperScale, and HyperStore are registered trademarks or trademarks of Cloudian, Inc. All other trademarks are property of their respective holders.

  - Hyperstore Secure Shell (HSH)
  - Retention Period
  - Use immutability with existing Cloudian/Veeam deployment

CLOUDIAN, INC.

177 Bovet Road, Suite 450, San Mateo, CA 94402 Tel: 1.650.227.2380 | cloudian.com ©2020 Cloudian, Inc. Cloudian, the Cloudian logo, HyperFile, HyperScale, and HyperStore are registered trademarks or trademarks of Cloudian, Inc. All other trademarks are property of their respective holders.

Optimizations inside Veeam

Summary

Here a summary of the topics below on the best practices which should be applied inside Veeam:

- Use multiple SOBRs
- Use “Local Target (Large Blocks)” for the “Backup Job” if possible
- Check “SOBR Offloading” settings
- Reduce thread count to lower CPU consumption in HyperStore
- Endpoint URL must be https

Use multiple SOBRs

Currently Veeam supports only 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 a better performance.

In order to use a SOBR a Veeam Enterprise (up to 2 SOBR) or an Enterprise Plus (unlimited SOBRs) license is required.

We recommend having a max bucket size of 250TB. Within the SOBR configuration option settings, it's possible to limit the max amount of space consumed by a capacity tier like S3 (on-prem like Cloudian HyperStore or cloud based).

Important:

- A new/separate bucket must be used for **every SOBR** and workload (like VBO/Office365 Backup, Archive backup jobs).
- Each SOBR requires its own performance tier (primary storage), whereby each SOBR could be a subdirectoy of the same primary storage
- Different retention policies can be set per SOBR.

CLOUDIAN, INC.
177 Bovet Road, Suite 450, San Mateo, CA 94402 Tel: 1.650.227.2380 | cloudian.com
©2020 Cloudian, Inc. Cloudian, the Cloudian logo, HyperFile, HyperScale, and HyperStore are registered trademarks or trademarks of Cloudian, Inc.
All other trademarks are property of their respective holders.

Backup Job

Storage Settings

Veeam allows 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. **Cloudian recommends to use "Local Target (large blocks)".** This is equal to a 4MB (=2MB after compression) object size. **Performance Tier storage requirements depend on this setting, too. Consult with Veeam or your partner before setting this.**

To consider when changing block size from 1MB "Local Target" to 4MB "Local Target (large blocks)":

- 4MB blocks have an up to 4x reduced object count compared to 1 MB blocks
- 4MB blocks may increase incremental backup disk required by 2x
- 4MB blocks increase time and reduces restore speed for random IO based restores. Example: SQL Explorer, Instant VM Recovery
- Block size changes only take effect after an active full is performed (Synthetic Full backup will not change block size)

How to configure in Veeam
V10:
https://helpcenter.veeam.com/docs/backup/hyperv/backup_job_advanced_storage_hv.html?ver=100
V11:
https://helpcenter.veeam.com/docs/backup/hyperv/backup_job_advanced_storage_hv.html?ver=110

SOBR Offloading

Veeam differentiates between repository and proxy tasks. For the SOBR offloading the aggregated number of concurrent tasks of all performance tier extents is used. If an extent has no limit on concurrent tasks, SOBR offload uses unlimited number of tasks (up to the # of backup chains).

©2020 Cloudian, Inc. Cloudian, the Cloudian logo, HyperFile, HyperScale, and HyperStore are registered trademarks or trademarks of Cloudian, Inc.
All other trademarks are property of their respective holders.
Each SOBR offload task will use up to 64 connections to the capacity tier. This is defined by

S3ConcurrentTaskLimit/AzureConcurrentTaskLimit in the Windows registry settings. Recommend approx 2,000 connections on-prem object storage (6k for public clouds). Since Veeam 11 additionally the number of tasks can be throttled on the Repository level.

To calculate the number of S3 operations that are processed in parallel (max) use this calculation: Available Repository Task Slots * S3ConcurrentTaskLimit = Max used S3 operations in parallel. Use the Veeam Object Storage IO calculator for the best setting.

SOBR Offload uses free repository task slots. As Backup and Restore jobs have higher precedence, SOBR offload will be paused until the task slot becomes available. Calculate 1 core and 4GB RAM per repository task slot and a recommended minimum is 2 cores/8GB RAM per repository. Details are explained in the Veeam Best Practices: https://bp.veeam.com/vbr/VBP/2_Design_Structures/D_Veeam_Components/D_backup_repositories/

When using “per job backups” one backup file per backup job is created and SOBR offload process will use one repository task per backup job chain.
In the case of “per vm backups” one backup file per virtual machine is created and the SOBR offload process will use one repository task per VM chain.

Settings vary a lot depending on the choosing architecture and size of the deployment.
It's strongly recommended to involve a Veeam SE for a proper sizing.

V11 Task Throttling

Veeam 11 introduced the capability to limit the maximum concurrent tasks at repository level. Each task slot allows processing of a single backup chain but is unchecked by default. The task throttling option should be preferably used to reduce load on an overloaded object storage system as it allows a single chain offload or download (at restore) to run at full speed. Repository setting overrides SOBR values!

Use the Veeam Object Storage IO calculator for the best setting.

Number of concurrent connections (S3ConcurrentTaskLimit)

The following Windows Registry key may need to be added to adjust the number of concurrent connections. By default it is not present. Depending on the value more or less threads are

©2020 Cloudian, Inc. Cloudian, the Cloudian logo, HyperFile, HyperScale, and HyperStore are registered trademarks or trademarks of Cloudian, Inc.
All other trademarks are property of their respective holders.

opened to write to the Cloudian HyperStore system. The value of this key is read only when the job starts.

Veeam provides the “Object Storage I/O Sizing Calculator” sheet which helps to find the correct values to set for the **S3ConcurrentTaskLimit**.

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

1. On the Veeam server, open regedit in C:\Windows and locate the HKEY_LOCAL_MACHINE > Software > Veeam > **Veeam Backup and Replication** entry.
2. 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.
3. Right click on the new registry entry and click **Modify**. Ensure that the value is decimal.
4. 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.

The default value of 64 (decimal) is applied to every worker process, meaning that the “SOBR Offloading” job overall may be using hundreds or thousands of concurrent connections, resulting in sub-optimal performance.

Try adjusting this to a low value initially, eg 2, and adjusting it upwards until optimal performance is achieved.

The performance will depend on the Veeam setup in the environment. If you have a single repository with 4 concurrent connections, and the above setting is left standard, you may have the ability to over power the object store. For example 4 x 64 connections (640 connections) This may result in high CPU usage.

When the registry key is applied on the VBR server it will be pushed out automatically to all Extents/GW.
When a Registry key is manually applied on the Extent/GW the registry key value on VBR server has a higher priority.

To verify the current SOBR offload job uses new setting check
%programdata%/Veeam/Backup/SOBR_Offload_name-of

©2020 Cloudian, Inc. Cloudian, the Cloudian logo, HyperFile, HyperScale, and HyperStore are registered trademarks or trademarks of Cloudian, Inc.
All other trademarks are property of their respective holders.

sobr/VMname/*.gate.log and search for S3ConcurrentTaskLimit.

Endpoint connection via HTTPS

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.

SOBR Offload Reporting

Veeam created a small PowerShell script to produce a report for SOBR offload jobs:

```
Connect-VBRServer -server "servername"
$sobrOffload = [Veeam.Backup.Model.EDbJobType]::ArchiveBackup
$jobs=[Veeam.Backup.Core.CBackupSession]::GetByTypeAndTimeInterval($s
obrOffload,(Get-Date).adddays(- 1), (Get-Date).adddays(1)) write
$jobs
```

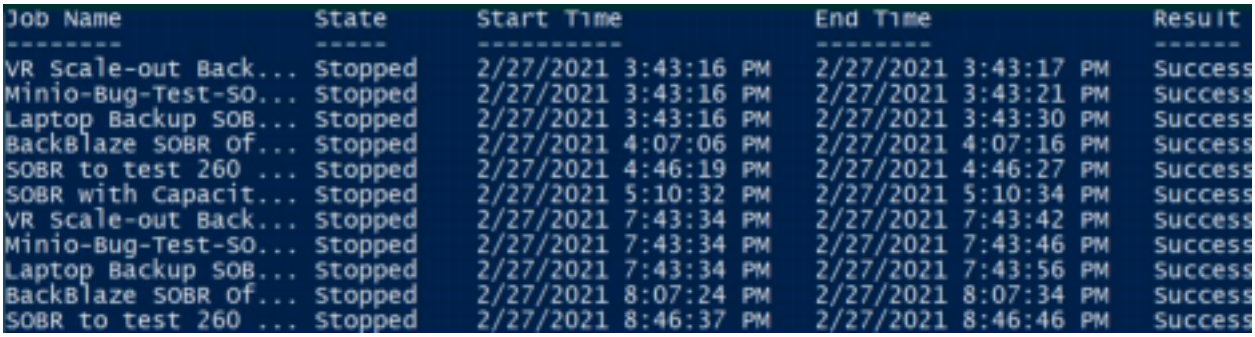

Job Name	State	Start Time	End Time	Result
VR Scale-out Back...	Stopped	2/27/2021 3:43:16 PM	2/27/2021 3:43:17 PM	Success
Minio-Bug-Test-SO...	Stopped	2/27/2021 3:43:16 PM	2/27/2021 3:43:21 PM	Success
Laptop Backup SOB...	Stopped	2/27/2021 3:43:16 PM	2/27/2021 3:43:30 PM	Success
BackBlaze SOBR Of...	Stopped	2/27/2021 4:07:06 PM	2/27/2021 4:07:16 PM	Success
SOBR to test 260 ...	Stopped	2/27/2021 4:46:19 PM	2/27/2021 4:46:27 PM	Success
SOBR with Capacit...	Stopped	2/27/2021 5:10:32 PM	2/27/2021 5:10:34 PM	Success
VR Scale-out Back...	Stopped	2/27/2021 7:43:34 PM	2/27/2021 7:43:42 PM	Success
Minio-Bug-Test-SO...	Stopped	2/27/2021 7:43:34 PM	2/27/2021 7:43:46 PM	Success
Laptop Backup SOB...	Stopped	2/27/2021 7:43:34 PM	2/27/2021 7:43:56 PM	Success
BackBlaze SOBR Of...	Stopped	2/27/2021 8:07:24 PM	2/27/2021 8:07:34 PM	Success
SOBR to test 260 ...	Stopped	2/27/2021 8:46:37 PM	2/27/2021 8:46:46 PM	Success

Hints for Lab and POC

Archive Immutability Generation Days

By default the number of archive immutability generations days will be set every 10 days by default. This means that an object is at least locked for 10 days by default. If a higher or lower frequency is required this can be configured by the following registry key.

Adding the following registry entry allows to set the value between 0-30.

©2020 Cloudian, Inc. Cloudian, the Cloudian logo, HyperFile, HyperScale, and HyperStore are registered trademarks or trademarks of Cloudian, Inc.
All other trademarks are property of their respective holders.

5. On the Veeam server, open regedit in C:\Windows and locate the HKEY_LOCAL_MACHINE > Software > Veeam > **Veeam Backup and Replication** entry.
6. Add a new DWORD (32-bit) value entry with the name
   For V10: SOBRArchiveImmutabilityGenerationDays
   For V11: SOBRArchiveImmutabilityGenerationDays
   When creating this entry, make sure that there are no blanks before or after the the name.
7. Right click on the new registry entry and click **Modify**. Ensure that the hexadecimal value to 0.
8. Close Backup & Replication and **reboot** the system to ensure that this registry entry is accepted.

CLOUDIAN, INC.
177 Bovet Road, Suite 450, San Mateo, CA 94402 Tel: 1.650.227.2380 | cloudian.com
©2020 Cloudian, Inc. Cloudian, the Cloudian logo, HyperFile, HyperScale, and HyperStore are registered trademarks or trademarks of Cloudian, Inc.

All other trademarks are property of their respective holders.

Optimizations inside Cloudian HyperStore

Summary

Here a summary of the topics below on the best practices which should be applied inside Cloudian HyperStore:

- Use Cloudian HyperStore 7.2.4.5 or later
- Ensure proper sizing of the Capacity Tier
- Ensure proper sizing for metadata disk size
- Use EC 4+2, 6+3 or RF3 when using “Local Target”
- No Lifecycle Rules & Tiering
- Change Bucket Folder Depth (Cassandra) to 7

Cloudian Version

We recommend to use at least Cloudian HyperStore 7.2.4.5 or later for Veeam deployments.

Storage Calculation

Veeam offer a sizing tool which is an online tool accessible here: https://vse.veeambp.com/#/

Overall the calculation of the required capacity can be very complex. The best approach is to get these information from Veeam.

Lifecycle Rules & Tiering

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

Tiering and lifecycle rules in object storages are 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.

CLOUDIAN, INC.
177 Bovet Road, Suite 450, San Mateo, CA 94402 Tel: 1.650.227.2380 | cloudian.com
©2020 Cloudian, Inc. Cloudian, the Cloudian logo, HyperFile, HyperScale, and HyperStore are registered trademarks or trademarks of Cloudian, Inc.

All other trademarks are property of their respective holders.

Metadata Disk Size for Veeam Backup & Recovery

By default block size is set to Local target, which is 1 MB before compression. Since compression ratio is very often around 2x, with this block size Veeam will write around 512 KB or less to the repository per each block by default.

Object metadata size is approx 290 Byte (without Object Lock) and 450 Byte (with Object Lock).

Because of the small objects it is possible we are running out of space on the metadata/OS disk. So, it is recommended to replace disks with larger drives (4 TB). Especially when using Cloudian HSA-4000/4100 appliances in most cases standard SSD size will be too small. Because of the many disks we do not recommend to use Cloudian HSX appliance for Veeam workloads at the moment.

Consult with the Cloudian Presales representative and ensure proper metadata sizing of the cluster.

Cloudian Sizing for Bulk Deletes

In the previous section we optimized the sizing for metadata capacity. Additionally to the sizing for metadata it is important to take performance of deletes into consideration, too. Veeam is doing bulk deletes frequently and we need to ensure that objects can be deleted in the required timeframe.

In some environments it might be better to use HSA-1600 appliances instead of HSA-4100 appliances.

Consult with the Cloudian Presales representative and ensure proper sizing of the cluster.

Bucket Settings

Erasure Coding

When using Erasure Coding for a Veeam bucket and the block size is smaller than “Local Target (Large Blocks)” use a RF3, EC 4+2 or EC 6+3 policy and avoid going higher. Overall less fragments are better.

CLOUDIAN, INC.
177 Bovet Road, Suite 450, San Mateo, CA 94402 Tel: 1.650.227.2380 | cloudian.com
©2020 Cloudian, Inc. Cloudian, the Cloudian logo, HyperFile, HyperScale, and HyperStore are registered trademarks or trademarks of Cloudian, Inc.
All other trademarks are property of their respective holders.

Storage Policies

In addition to using multiple buckets (hence multiple Veeam SOBRs) to store Veeam data into, it is Cloudian recommendation to create a new storage policy for every bucket which is used, even if the storage policies are all defined exactly the same way. For example, if in a deployment, one plans on creating four (4) s3 buckets to store Veeam data, it would be recommended to create four (4) storage policies and have each bucket use a different storage policy.

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.

Bucket Folder Depth (Cassandra)

In order to allow faster queries of the Cassandra database and larger bucket sizes it is recommended to change the folder depth of the Veeam buckets inside Cloudian HyperStore. Depending if Cloudian HyperStore deployment is only used for Veeam or if it is a mixed environment with other non Veeam buckets the process is different.

If you are using Cloudian HyperStore only for Veeam we recommend to set 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. Add the property as below at the end of the file
   ```
   cloudian.s3.folder.depth=7
   ```

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

   With HSH:
   ```
   # hspkg config -e mts.properties.erb
   ```

2. Push the configuration changes in puppet
   ```
   # ./cloudianInstall.sh runpuppet=""
   ```

   With HSH:

©2020 Cloudian, Inc. Cloudian, the Cloudian logo, HyperFile, HyperScale, and HyperStore are registered trademarks or trademarks of Cloudian, Inc.
All other trademarks are property of their respective holders.

```
# hspkg install runpuppet=””
```

3. Restart S3 service (with or without HSH)

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

With HSH:

```
# hspkg install
```

All the buckets created after this will have a folder depth of 7.

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 restart the S3 service again.

PQLISTING

Since Cloudian HyperStore 7.2.3 the PQListing is enabled by default. For older installations ensure it is enabled.

The value of the cloudian.s3.pqlisting parameter in the /etc/cloudian-7.2.4-puppet/modules/cloudians3/templates/mts.properties.erb Puppet configuration template has to be set to 'true'. After the change a puppet push and restart of the S3 service is required.

©2020 Cloudian, Inc. Cloudian, the Cloudian logo, HyperFile, HyperScale, and HyperStore are registered trademarks or trademarks of Cloudian, Inc.

All other trademarks are property of their respective holders.

SECURITY

Bucket usage

Purely out of security reasons it is recommended to create an own bucket and own user where possible for the Object Storage Repository and limit the user account to have only the required access on the object storage bucket.

Immutability

Veeam Backup & Replication allows you to prohibit deletion of data from object storage by making it immutable.

To make data immutable, Veeam Backup & Replication uses the Object Lock technology provided by Cloudian. Once imposed, the object lock prohibits deletion of data from object storage until the immutability expiration date is met. Inside the Veeam Capacity Tier settings it is specified how many days the object is immutable inside Cloudian object storage repository.

Note that Veeam Backup & Replication will automatically use Compliance object lock mode for each uploaded object.

The following operations cannot be performed on immutable data:

- Manual removal of data, as described in Removing Backups from Object Storage Repository.
- Removal of data by the retention policy, as described in Retention Policy.
- Removal of data using any cloud service provider tools.
- Removal of data by the cloud service provider technical support department. ● Removal of data by the Remove deleted items data after option, as described in Maintenance Settings.
- Immutable data is preserved as described in Block Generation.

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.

CLOUDIAN, INC.
177 Bovet Road, Suite 450, San Mateo, CA 94402 Tel: 1.650.227.2380 | cloudian.com
©2020 Cloudian, Inc. Cloudian, the Cloudian logo, HyperFile, HyperScale, and HyperStore are registered trademarks or trademarks of Cloudian, Inc.

All other trademarks are property of their respective holders.

Retention Period

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

Here an article from Veeam on the retention period:
V10: https://helpcenter.veeam.com/docs/backup/hyperv/block_gen.html?ver=100
V11: https://helpcenter.veeam.com/docs/backup/hyperv/block_gen.html?ver=110

How to configure immutable retention in Veeam
V10:
https://helpcenter.veeam.com/docs/backup/hyperv/compatible_storage_details.html?ver=100
V11:
https://helpcenter.veeam.com/docs/backup/hyperv/compatible_storage_details.html?ver=110

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:

©2020 Cloudian, Inc. Cloudian, the Cloudian logo, HyperFile, HyperScale, and HyperStore are registered trademarks or trademarks of Cloudian, Inc.
All other trademarks are property of their respective holders.

https://helpcenter.veeam.com/docs/backup/vsphere/osr_switch_to_seal.html?ver=100 V11: https://helpcenter.veeam.com/docs/backup/vsphere/osr_switch_to_seal.html?ver=110

CLOUDIAN, INC.
177 Bovet Road, Suite 450, San Mateo, CA 94402 Tel: 1.650.227.2380 | cloudian.com
©2020 Cloudian, Inc. Cloudian, the Cloudian logo, HyperFile, HyperScale, and HyperStore are registered trademarks or trademarks of Cloudian, Inc.

All other trademarks are property of their respective holders.