

# Optimizing & Securing Veeam Backups with Hitachi Content Platform for Cloud Scale

WHITE PAPER

## Solution Whitepaper: Veeam Ready Object with Immutability with HCP for Cloud Scale



By Hitachi Vantara

May 2022

# Contents

Executive Summary	3
Test Methodology	4
Test Environment	4
Hardware & Software Components	4
Solution Diagram	5
Configuring Hitachi Content Platform for Cloud Scale for Veeam	5
Test Cases	7
Test Results	8
Performance Tuning & HCPCS Configuration Guidelines	13
Additional Resources	13

## Executive Summary

Continued exponential growth of data in the industry has led to a never ending, increased demand on storage capacity and optimization of business-critical processes like backup and recovery. Aside from managing growth, are the ever so critical needs involving data security, risk mitigation, loss prevention, and compliance.

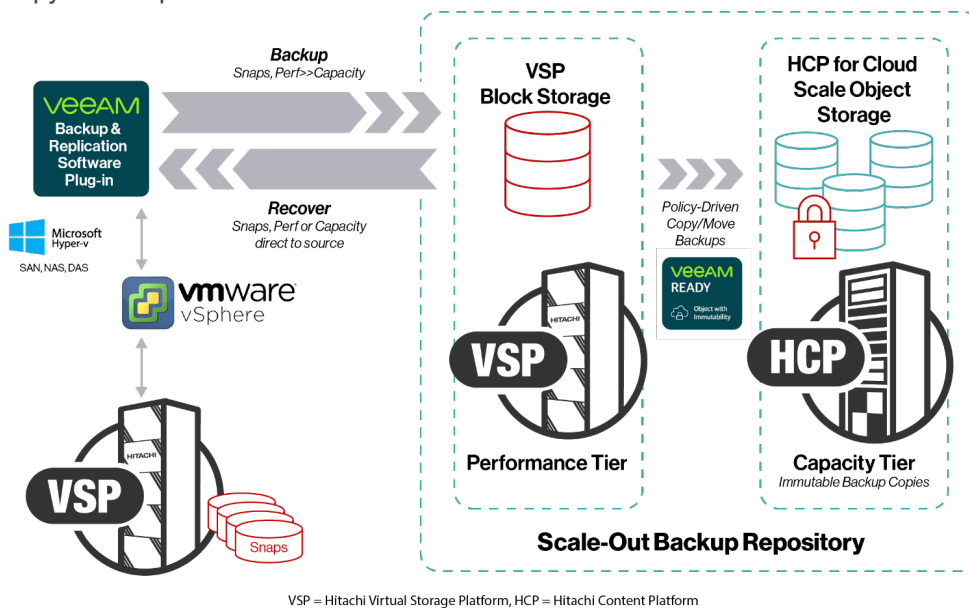
To help customers optimize their cloud data infrastructure, improve service levels for backups, and ensure immutability of critical recovery information, Hitachi Vantara and Veeam have partnered to deliver a “Veeam Object Ready – Immutable Backup” solution.

### A Massively Scalable and Cost-Effective Solution for Immutable Backups & Long-Term Data Retention

Hitachi Content Platform for Cloud Scale (HCP for Cloud Scale) with AWS S3 Object Lock is an on-premises cloud Object Based Storage (OBS) system that delivers a cost-effective, secure and robust data platform to power a myriad of enterprise use cases.

Veeam Backup & Replication, part of Veeam Availability Suite, delivers fast, flexible and reliable backup, recovery and replication for virtual, physical and cloud-based workloads. Veeam Cloud Tier, a feature of Veeam Backup & Replication Enterprise Edition, provides a Scale-Out Backup Repository for massively scalable capacity and long-term data retention using AWS S3 compliant OBS, such as HCP for Cloud Scale.

Primary backups are tiered to HCP for Cloud Scale based on aging or capacity policies and further made immutable by setting business required retention policies. In this manner, customers can improve backup and restore performance, reduce costs, and improve security in a coordinated fashion between the Veeam Backup & Replication application and the HCP for Cloud Scale platform. Due to immutable nature of the backups leveraging Object Lock, external threats to corrupt this “last resort to recover” will be denied and users can be assured that an “always recoverable data copy” will be preserved.



The purpose of this whitepaper is to document the test configuration, results and best practices.

The results captured in this document are specific to Veeam and Hitachi Vantara, including test plans, selected configurations, test methodologies, and processes used to measure results. Actual end user experiences may vary based on customer-specific environments and corresponding solution design. Please consult with your Veeam and Hitachi Vantara technical representatives prior to implementing this solution.

## Test Methodology

The goal of this testing was to validate integration, functionality and performance of Veeam Backup and Replication v11 with HCP for Cloud Scale v2.x using AWS S3 protocol. Testing was conducted in two phases as follows:

### 1. Basic Functional Run-Time Targets - 4.5 hours

- Multi-Part Upload
- File Level recovery (FLR)
- SOBR Offload Resiliency
- Large Offload to OBS
- Delete backup file on OBS

### 2. Immutability Tests with S3 Object Lock

- API Verification
- Verify Object Immutability
- SOBR Offload Resiliency

The test suite was developed and provided by Veeam. Configuration selection was at Hitachi's discretion. All results were submitted to and approved by Veeam. As such, HCP for Cloud Scale has received Veeam Object Ready "Immutable" designation and is listed on Veeam's Ready program site <https://www.veeam.com/ready.html>.

The corresponding Knowledge Base article can also be found at <https://www.veeam.com/kb3195>.

## Test Environment

This section describes the configuration that was tested, following Veeam Backup & Replication and HCP for Cloud Scale reference documentation.

The full Veeam documentation set can be accessed here at <https://www.veeam.com/documentation-guides-data-sheets.html>.

Veeam Community forum is here at <http://forums.veeam.com>.

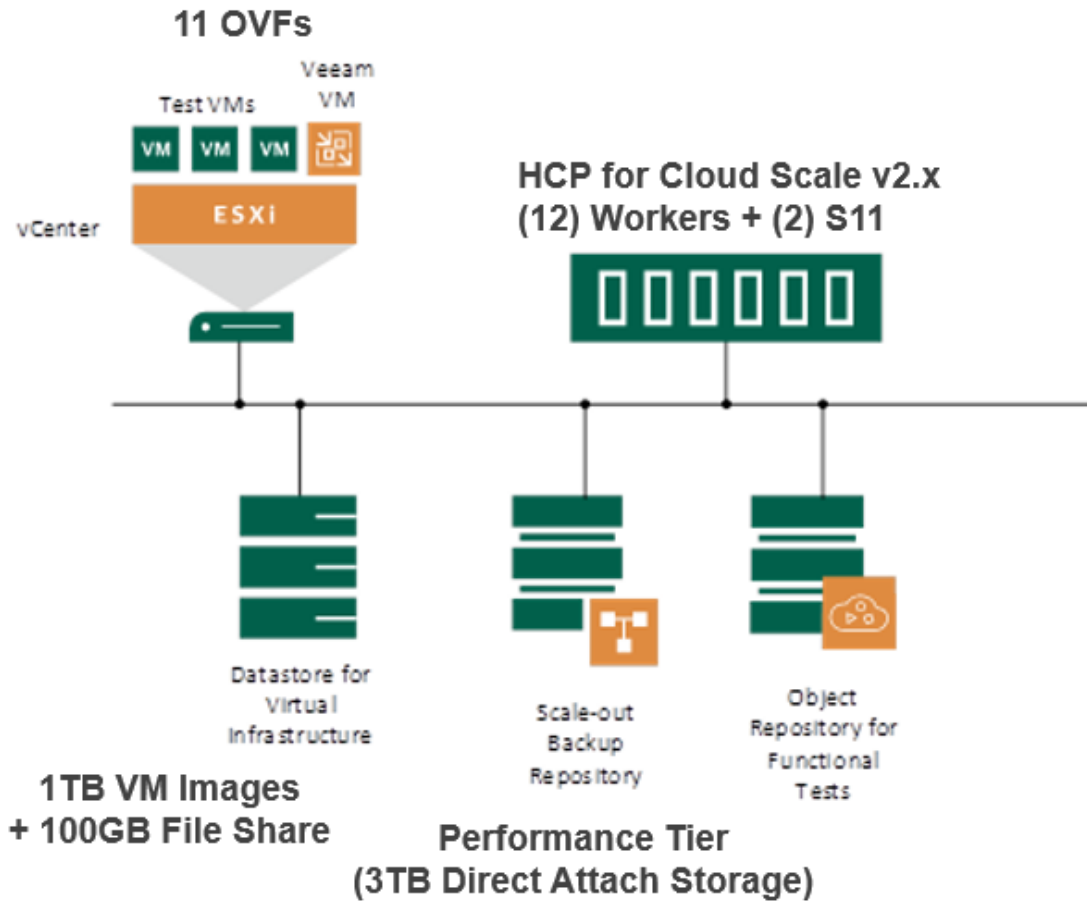
Hitachi Vantara Community site can be found here <https://community.hitachivantara.com/s/>.

## Hardware & Software Components

- (1) Veeam Backup and Restore Server
  - Two Intel Xeon E5 2650 v3 2.3 GHz, 10 Core, 128GB Ram per processor (256GB total), 2 × 10 GbE NIC & 17.9 TB local SSD storage
  - Veeam Backup and Restore v11
  - CentOS Linux Operating System 7.4.1708
  - Microsoft® Windows Server® Standard 2012 R2

- (1) Hitachi Content Platform Cloud Scale Cluster
  - (10) Worker nodes hosted on Quanta DS120 servers with 2 Intel Xeon Silver 4110 processors, 128 GB Ram per processor (256GB total), 1 x 10 GbE NIC and 12.5TB local storage.
  - (2) S11 S-nodes (100TBs minimum application usable storage)
  - HCPCS SoftwareV2.x

### Solution Diagram



### Configuring Hitachi Content Platform for Cloud Scale for Veeam

- HCPCS bucket setup

Performance settings on the HCP for CloudScale system (named: HCPCS) for Veeam Object Ready program:

- One bucket with versioning enabled and object lock disabled.
- 10 S3 Gateway services running, with memory settings max:12000m, heap:10000m.

- 10 Metadata Gateway services running, with memory max:40000m, heap:30000m.
- Memory of the Metadata Cache max:20000m, heap:18000m
- Memory of the Metadata Coordination to max:768m, heap:300m

The immutable bucket was created with the following AWS command line:

```
aws s3api create-bucket --endpoint-url <target> --bucket <bucket name> --region us-east-1 --object-lock-enabled-for-bucket
```

## Veeam Configuration

One Server running on Windows 2012 R2

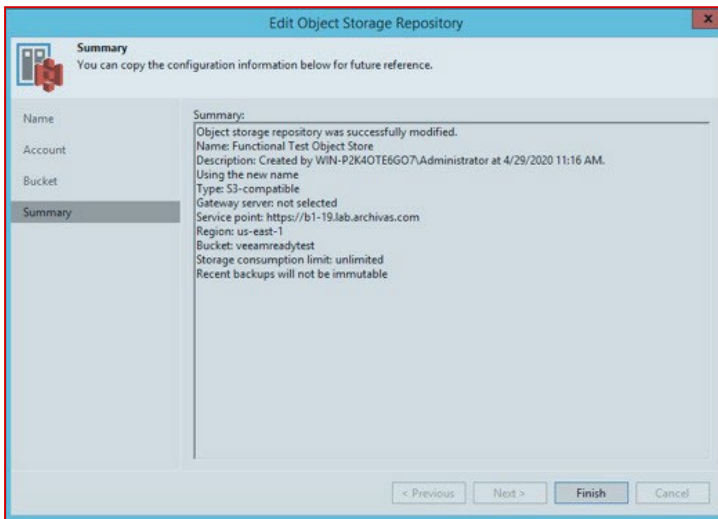
Memory: 32Gb of Ram

vCPUs: 24

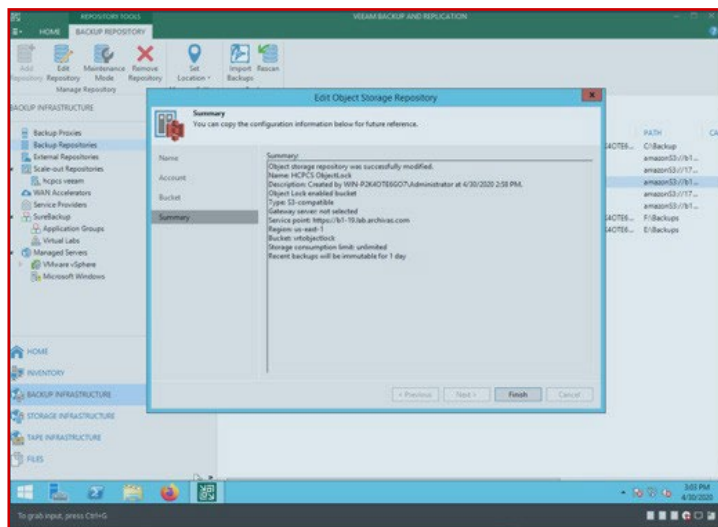
Drives: 500GB Drive for OS and programs

3TB Drive to be used for local backup

The Object Store Repository options:



The Immutable Object Store Repository options:



## Test Cases

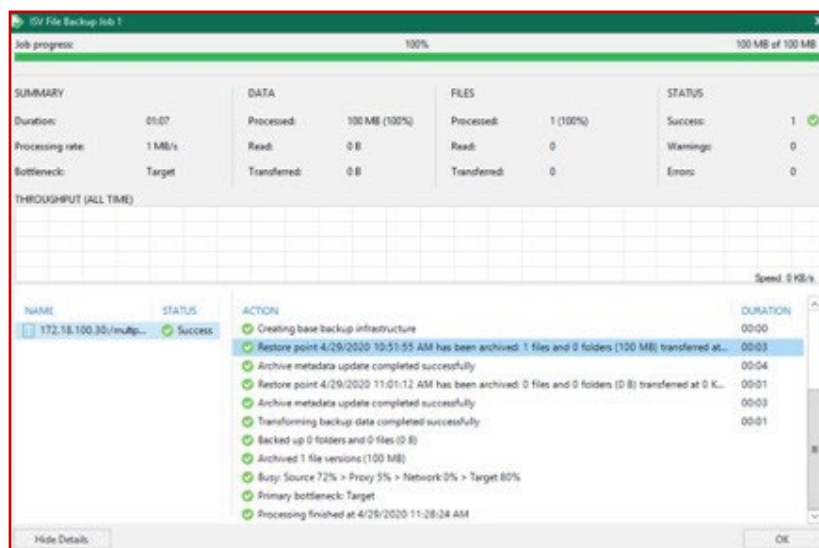
There were (8) specific test cases defined in the test plans, 5 for Veeam Ready Object and 3 for Veeam Ready Object Immutable as follow:

1. Multipart Upload Test – This is testing the default upload of a backup file from Veeam to HCPCS.
2. File Level Recovery Test – This test restores a file stored on HCPCS back to Veeam and finishing by having Veeam restore that file to the local system.
3. SOBR Offload Resiliency – This test two different recovery scenarios. The first scenario tests what happens if there is a network disruption for five minutes and then it is restored. The second test will have the backup fail due to a network outage the causes the backup to fail. After that, a manual offload will be started to confirm the remainder of the file is offloaded.
4. Perform a large offload to the Object Storage – Take 10 large, 100GB, VMs and use Veeam to offload to HCPCS. There is a performance piece to this test as well, the 1TB offload needs to be completed within four-and-a-half hours.
5. Delete a backup file on the S3 Compatible Storage – Confirm that deleting a file in Veeam will also be deleted on HCPCS. There is a performance piece to this test as well, the delete needs to be completed within four-and-a-half hours.
6. API Verification – Testing the results of S3 commands to get information on the buckets on the HCPCS system. This checks the results if Object Immutability is set as well as if it is not set.
7. Verification of Object Immutability – Run through a scenario testing the Object immutability on the HCPCS system. First a backup will be made with a 1-day immutable period. Then attempts to modify and delete it will be made and fail. After the 24-hour period the test runs an incremental backup to verify it can be modified. Finally, the object will be deleted to after the object's immutability period has passed to verify deletes work.
8. SOBR Offload Resiliency with Object Immutability - Take 10 large, 100GB, VMs and use Veeam to offload to HCPCS into a bucket with Object Immutability on. There is a performance piece to this test as well, the 1TB offload needs to be completed within four-and-a-half hours.

## Test Results

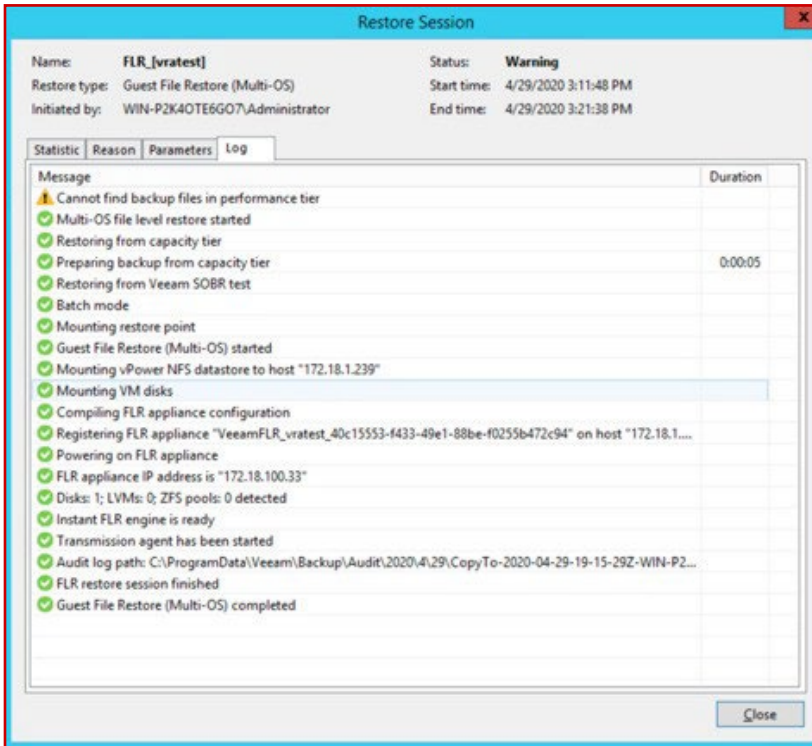
1. Multi-Part Upload Test

Backup a single file and then offload that backup to the Object Store.



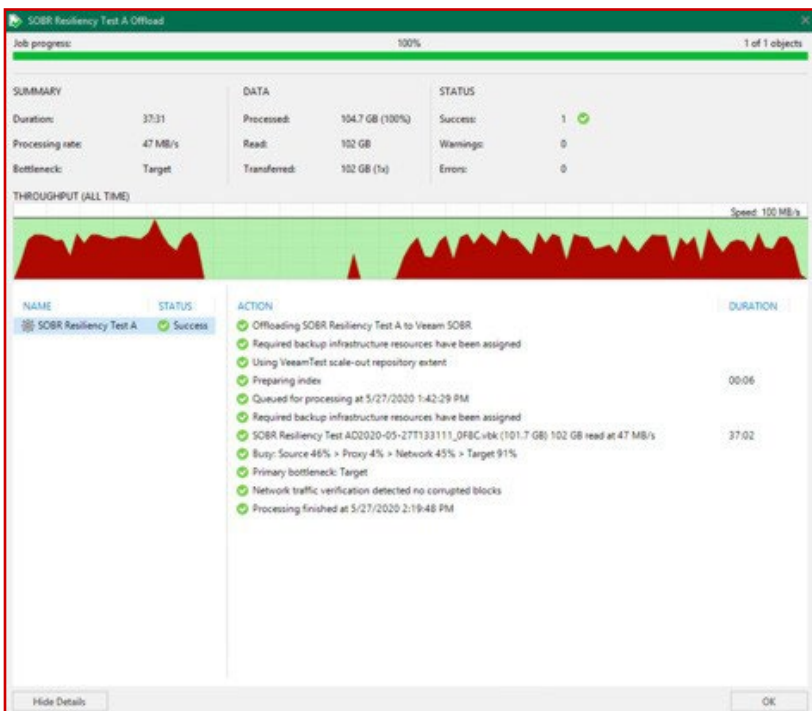
## 2. File Level Recovery Test

Restore the data from the Object Store.



## 3. SOBR Offload Resiliency Test

Run two backups. First, during an offload, have the network connection between the Veeam server and the Object store fail for five minutes. Then, restore the network connection and let the offload complete. Next, during the offload, have the network connection fail for over thirty minutes to cause the offload to time out. After the timeout, restore the network connection and manually run an offload to complete the originally failed offload.





**SUMMARY**

Duration:	48:21	Processed:	52.3 GB (100%)	Success:	0
Processing rate:	73 MB/s	Read:	51.5 GB	Warnings:	0
Bottleneck:	Target	Transferred:	51.5 GB (1x)	Errors:	1

**DATA**

**STATUS**

THROUGHPUT (ALL TIME) Speed: 110 MB/s

NAME	STATUS	ACTION	DURATION
SOBR Offload Resilien...	Failed	<ul style="list-style-type: none"> <li>Offloading SOBR Offload Resiliency 8 to Veeam SOBR</li> <li>Required backup infrastructure resources have been assigned</li> <li>Using VeeamTest scale-out repository extent</li> <li>Preparing index</li> <li>Queued for processing at 5/27/2020 11:23:51 AM</li> <li>Required backup infrastructure resources have been assigned</li> <li>SOBR Offload Resiliency 802020-05-27f111326_71f8.vbk (101.7 GB) 51.5 GB read at 73 MB/s</li> <li>WinHttpRequest: 12007: The server name or address could not be resolved</li> <li>Busy: Source 54% &gt; Proxy 4% &gt; Network 54% &gt; Target 89%</li> <li>Primary bottleneck: Target</li> <li>Network traffic verification detected no corrupted blocks</li> <li>Processing finished with errors at 5/27/2020 12:08:23 PM</li> </ul>	03:43

Hide Details OK

Aspos veeam Offload

Job progress: 100% 1 of 1 objects

**SUMMARY**

Duration:	13:46	Processed:	52.7 GB (100%)	Success:	1
Processing rate:	67 MB/s	Read:	50.8 GB	Warnings:	0
Bottleneck:	Target	Transferred:	50.8 GB (1x)	Errors:	0

**DATA**

**STATUS**

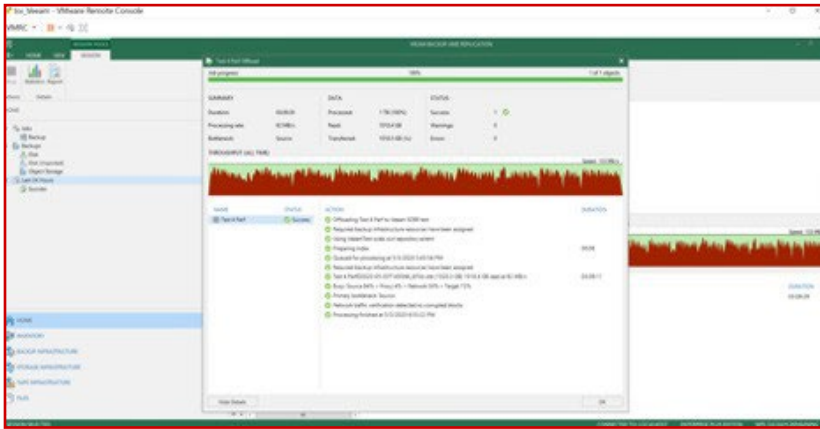
THROUGHPUT (ALL TIME) Speed: 105 MB/s

NAME	STATUS	ACTION	DURATION
SOBR Offload Resilien...	Success	<ul style="list-style-type: none"> <li>Offloading SOBR Offload Resiliency 8 to Veeam SOBR</li> <li>Required backup infrastructure resources have been assigned</li> <li>Preparing index</li> <li>Queued for processing at 5/27/2020 12:42:31 PM</li> <li>Required backup infrastructure resources have been assigned</li> <li>SOBR Offload Resiliency 802020-05-27f111326_71f8.vbk (101.7 GB) 50.8 GB read at 66 MB/s</li> <li>Busy: Source 54% &gt; Proxy 4% &gt; Network 53% &gt; Target 85%</li> <li>Primary bottleneck: Target</li> <li>Network traffic verification detected no corrupted blocks</li> <li>Processing finished at 5/27/2020 12:55:53 PM</li> </ul>	00:08

Visualization: Command Prom... OK

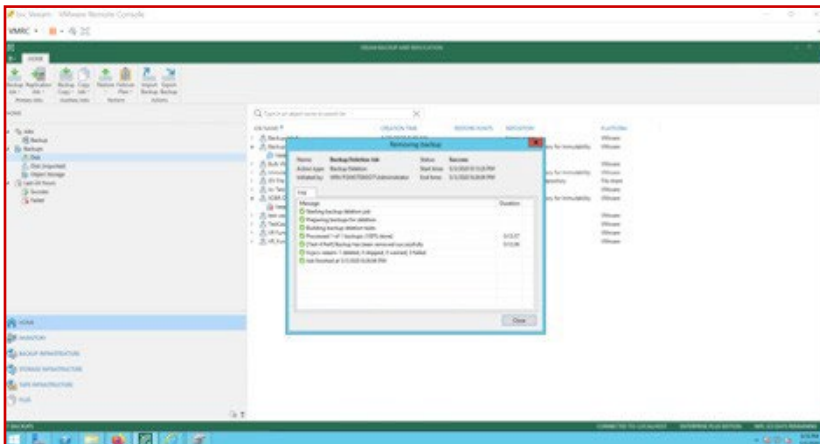
#### 4. Perform a large offload to the Object Store

Perform a backup consisting of ten VMs totaling a Terabyte in size. Measure the offload time to ensure it completes in under four-and-a-half hours.



#### 5. Delete a backup file from the S3 Compatible Storage

Delete the one Terabyte backup file and ensure it is removed from the Object Store in under 4.5 hours.



#### 6. API Verification

Run API calls against the Object Store to show the results based on whether the object lock flag is set. The first example is when the object lock flag is not set, the second when the object lock flag is set.

```
[mscerra@mscerra-dt Downloads]$ aws s3api get-object-lock-configuration --no-verify-ssl --endpoint-url https://bl-19.lab.archivas.com --bucket veenamobjectlockapi
/usr/local/aws-cli/v2/2.0.10/dist/urllib3/connectionpool.py:986: InsecureRequestWarning: Unverified HTTPS request is being made to host 'bl-19.lab.archivas.com'. Adding certificate verification is strongly advised. See: https://urllib3.readthedocs.io/en/latest/advanced-usage.html#ssl-warnings

An error occurred (ObjectLockConfigurationNotFoundError) when calling the GetObjectLockConfiguration operation: Object Lock configuration does not exist for this bucket
[mscerra@mscerra-dt Downloads]$
```

```
[mscerra@mscerra-dt Downloads]$ aws s3api get-object-lock-configuration --no-verify-ssl --endpoint-url https://bl-19.lab.archivas.com --bucket vrbobjectlock
/usr/local/aws-cli/v2/2.0.10/dist/urllib3/connectionpool.py:986: InsecureRequestWarning: Unverified HTTPS request is being made to host 'bl-19.lab.archivas.com'. Adding certificate verification is strongly advised. See: https://urllib3.readthedocs.io/en/latest/advanced-usage.html#ssl-warnings
{
  "ObjectLockConfiguration": {
    "ObjectLockEnabled": "Enabled"
  }
}
[mscerra@mscerra-dt Downloads]$
```

## 7. Verification of Object Immutability

Run tests sending the backup with the object immutable value set in both Veeam and the HCPCS system. First, backup the file. Second, after the object lock period (24 hours in this example) update the file to demonstrate that it can be updated.

Job progress: 100% 1 of 1 objects

SUMMARY		DATA		STATUS	
Duration:	01:00	Processed:	5 GB (100%)	Success:	1 ✓
Processing rate:	75 MB/s	Read:	2.2 GB	Warnings:	0
Bottleneck:	Target	Transferred:	2.2 GB (1x)	Errors:	0

THROUGHPUT (ALL TIME) Speed: 84 MB/s

NAME	STATUS	ACTION	DURATION
Immutable Test	Success	<ul style="list-style-type: none"> <li>Offloading Immutable Test to HCPCS ObjectLock</li> <li>Required backup infrastructure resources have been assigned</li> <li>Using VeeamReady-v10 (Backup server) scale-out repository extent</li> <li>Preparing index</li> <li>Queued for processing at 4/30/2020 3:36:27 PM</li> <li>Required backup infrastructure resources have been assigned</li> <li>Immutable Test02020-04-30T153355_F92C.vbk (2.2 GB) 2.2 GB read at 73 MB/s</li> <li>Busy: Source 36% &gt; Proxy 5% &gt; Network 34% &gt; Target 95%</li> <li>Primary bottleneck: Target</li> <li>Network traffic verification detected no corrupted blocks</li> <li>Processing finished at 4/30/2020 3:37:15 PM</li> </ul>	00:07

Job progress: 100% 1 of 1 objects

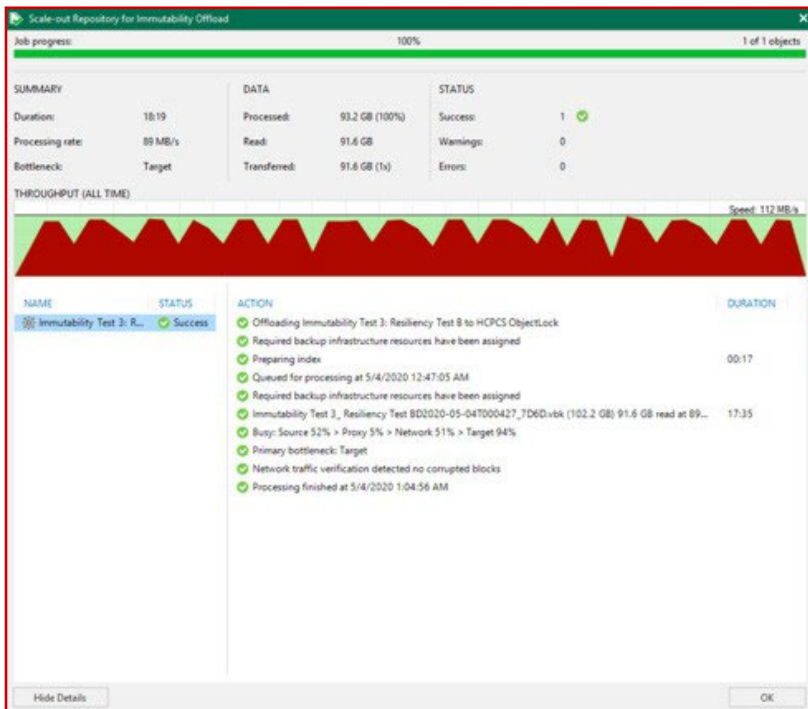
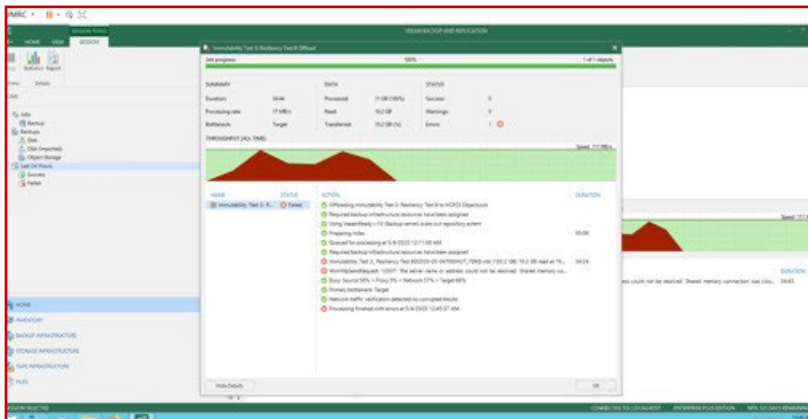
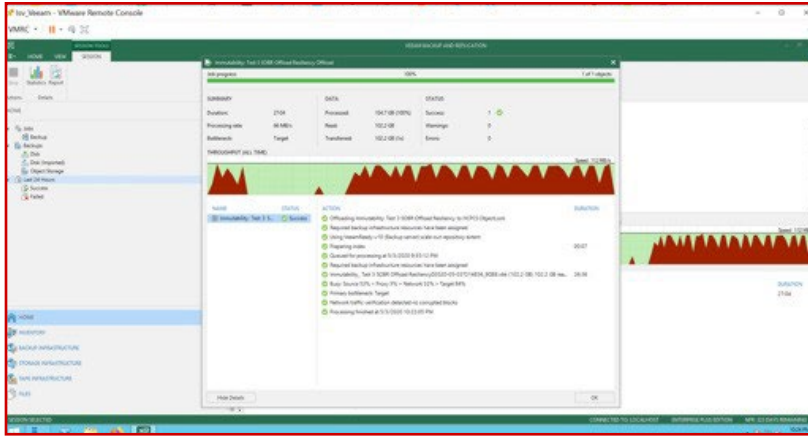
SUMMARY		DATA		STATUS	
Duration:	00:44	Processed:	57.1 MB (100%)	Success:	1 ✓
Processing rate:	2 MB/s	Read:	35.3 MB	Warnings:	0
Bottleneck:	Source	Transferred:	35.3 MB (1x)	Errors:	0

THROUGHPUT (ALL TIME) Speed: 0 KB/s

NAME	STATUS	ACTION	DURATION
Immutable Test	Success	<ul style="list-style-type: none"> <li>Offloading Immutable Test to HCPCS ObjectLock</li> <li>Required backup infrastructure resources have been assigned</li> <li>Preparing index</li> <li>Queued for processing at 4/30/2020 4:03:45 PM</li> <li>Required backup infrastructure resources have been assigned</li> <li>Immutable Test02020-04-30T160125_ABA6.vbk (2.2 GB) 35.3 MB read at 2 MB/s</li> <li>Busy: Source 92% &gt; Proxy 0% &gt; Network 32% &gt; Target 0%</li> <li>Primary bottleneck: Source</li> <li>Network traffic verification detected no corrupted blocks</li> <li>Processing finished at 4/30/2020 4:04:17 PM</li> </ul>	00:06

## 8. SOBR Offload Resiliency Object Immutability Test

Run two backups with the object immutable option chosen. First, during an offload, have the network connection between the Veeam server and the Object store fail for five minutes. After that, restore the connection and let the offload complete. Second, during the offload, have the network connection fail for over thirty minutes to cause the offload to time out. After the timeout, restore the network connection and manually run an offload to complete the originally failed offload.



## Performance Tuning & HCPCS Configuration Guidelines

### Simultaneous Connections

- Veeam has the option to set a number of threads for uploads. Setting one thread per S3 gateway produced the best results. In our example, we had ten S3 gateways and ten threads for Veeam.

The following are key questions to be answered, leading up to a proper system design:

1. **Identify if Object Immutability is required.** If so, the bucket will need to be created with the object immutable flag set on HCP for Cloud Scale as you cannot change the option once the bucket is created.
2. **Consider Retention Requirements.** In particular, whether deleting, freezing, or archiving of specific data sets will be necessary. Veeam Backup & Replication and HCP for Cloud Scale should be designed in conjunction to meet business compliance and data retention policies.

### Additional Resources

Check out the resources below to learn how Hitachi Vantara and Veeam work together to enable organizations to optimize their Veeam environments for reduced storage costs and greater efficiency.

[Learn how](#) Hitachi Content Platform Cloud Scale makes your data securely available anywhere, anytime.

[Read about](#) the power of the Hitachi Content Platform portfolio.

[Learn how](#) to use Veeam to safely backup and have your data available.

## Hitachi Vantara



Corporate Headquarters  
2535 Augustine Drive  
Santa Clara, CA 95054 USA  
[hitachivantara.com](http://hitachivantara.com) | [community.hitachivantara.com](http://community.hitachivantara.com)

Contact Information  
USA: 1-800-446-0744  
Global: 1-858-547-4526  
[hitachivantara.com/contact](http://hitachivantara.com/contact)

HITACHI is a registered trademark of Hitachi, Ltd. VSP is a trademark or registered trademark of Hitachi Vantara LLC. Microsoft, Azure and Windows are trademarks or registered trademarks of Microsoft Corporation. All other trademarks, service marks and company names are properties of their respective owners. **NOTE: Actual trademark text should reflect the real content. Use the trademark guidance in the Hitachi Vantara Brand Guide.**

WP-xxx-x Agency Initials Month 20XX