REFERENCE ARCHITECTURE

Quantum ActiveScale with Veeam Backup and Recovery

Abstract

This document defines a reference architecture for combined solution based on Quantum ActiveScale with Veeam Backup and Recovery.

Quantum

Table of Contents

Executive Summary

Data sets being managed by modern IT organizations continues to grow at exponential rates. The amount of data required for analytics and virtual infrastructures is staggering. And, as organizations become more dependent on their data, hackers are creating more sophisticated ways to attack that data for both destructive and monetary purposes. For these reasons and more, an organization's backup infrastructure is very much mission-critical.

As businesses strive to become more competitive, many are expected to do more with less staff. This means reliable, automated backup infrastructures that can be easily managed are critical for any successful business. For decades, Quantum has been a leader in providing successful businesses the solutions and knowledge needed to create their robust backup infrastructures.

This reference architecture describes one such solution, based on Quantum ActiveScale™ with Veeam Backup and Replication 11 (VBR). The reference architecture was verified by successfully completing the required testing defined by Veeam as part of the Veeam Ready program.

Objectives

The objectives of this paper are:

1. Articulate the importance of a robust backup infrastructure
2. Present a general, scalable backup infrastructure based on ActiveScale and VBR
3. Describe the components that make up the solution stack
4. Demonstrate that ActiveScale performs as expected when configured as Scale Out Backup Repository for VBR
5. Articulate a reference architecture that is suitable for a combined ActiveScale/VBR solution
6. Demonstrate the reference architecture is valid by performing various validation activities
7. Document additional guidelines, tips, and best practices to maximize success

The Need for Backup

A robust backup infrastructure is fundamental data protection for all organizations. Primary data is protected by creating and storing a copy of that data on a separate medium. The objective is to ensure the backup copy can be recovered in the event of a primary data failure.

There are numerous reasons primary data can fail. For example, data failures can be a result of hardware or software failures, data corruption, or a human-caused event, such as a malicious attack (virus, malware, ransomware, etc.), or simply accidental data deletion. Maintaining Backup copies allows an organization to restore a data set from an earlier point in time to quickly recover from a primary data failure.

Storing the copy of the data on a durable medium, safely and separate from the primary data, is critical. Object storage, with its extreme level of durability and favorable economics, has become the medium of

choice in a high-performance backup infrastructure. Quantum ActiveScale not only provides the durability and performance required, but advanced features such as Versioning and Object Lock for enhanced protection against malicious ransomware attacks.

Solution Overview

Data Centers are increasingly becoming a collection of virtualized servers. VBR seamlessly integrates with hypervisors to backup these virtual servers without the need to deploy backup client software on those servers. It does this by using the hypervisor API to create snapshots of the virtual server's disks, and then create the backup files from the snapshot. The primary advantage of this approach is that the backup is transparent to the virtualized servers, therefore, they do not experience backup processes overhead.

Quantum understand the importance of having a robust backup infrastructure for the virtualized data center. The Quantum Solutions Team has created, deployed, and tested this reference architecture to ensure overall reliability and integrity.

The solution described in this reference architecture is easily scalable, both up and out, without the need to rebalance storage. VBR has been certified by Quantum as being able to leverage ActiveScale's object-lock feature for enhanced ransomware protection. When backups are created, VBR will offload those backups from more expensive primary storage to ActiveScale based on user defined parameters. ActiveScale becomes the permanent storage repository for all but the most recent backups. In the event a restore is needed from these archived backups, whether it be an individual file or an entire data center, VBR will restore directly from ActiveScale eliminating the need to rehydrate the backups to the primary storage.

Technology Summary

The table below lists the technology components that make up the ActiveScale/VBR reference architecture outlined in this document. The paragraphs that follow the table provide more detail on the function of these components in the solution.

Technology	Version
Quantum ActiveScale Version	5.7 or higher
Quantum ActiveScale Model	P100E3 or larger
VMware ESXi Version	Although ESXi 6.5 was used in the deployment of this reference architecture, any hypervisor supported by Veeam is suitable.
Veeam Backup and Replication	11

About ActiveScale

With data quickly becoming the most valuable asset for many enterprises and critical for business survival, Quantum's ActiveScale is emerged as a preferred platform for long term secure storage in today's backup and archiving environments. The ActiveScale product line delivers flexible scalability, from a few hundred terabytes to hundreds of petabytes. ActiveScale's advanced features, such as Dynamic Data Placement and Dynamic Data Repair, simplify the deployment and management of the environment today and into the future.

About Veeam

Veeam Backup & Replication delivers Availability for ALL workloads — virtual, physical, and cloud — from a single management console, extending Veeam's leadership position from being the best for VMware vSphere and Microsoft Hyper-V to #1 Availability for any app, any data on any cloud. It allows customers to completely get rid of legacy backup forever and brings backup and replication together into a single software solution.

Reference Architecture

Veeam uses ActiveScale as a Scale Out Backup Repository (SOBR). When a backup is performed, it resides on a primary backup repository for a period of time. Rules are created that define how many and how long backups should reside in the primary repository. After meeting user-defined rules, these backup sets are copied to the scale out repository (ActiveScale). Once backups are safely written to ActiveScale, the VBR rules engine, based on user-defined parameters, will automatically purged from the primary repository either immediately, after a user-defined number of days, or the user can choose to not purge the backups from the primary repository at all.

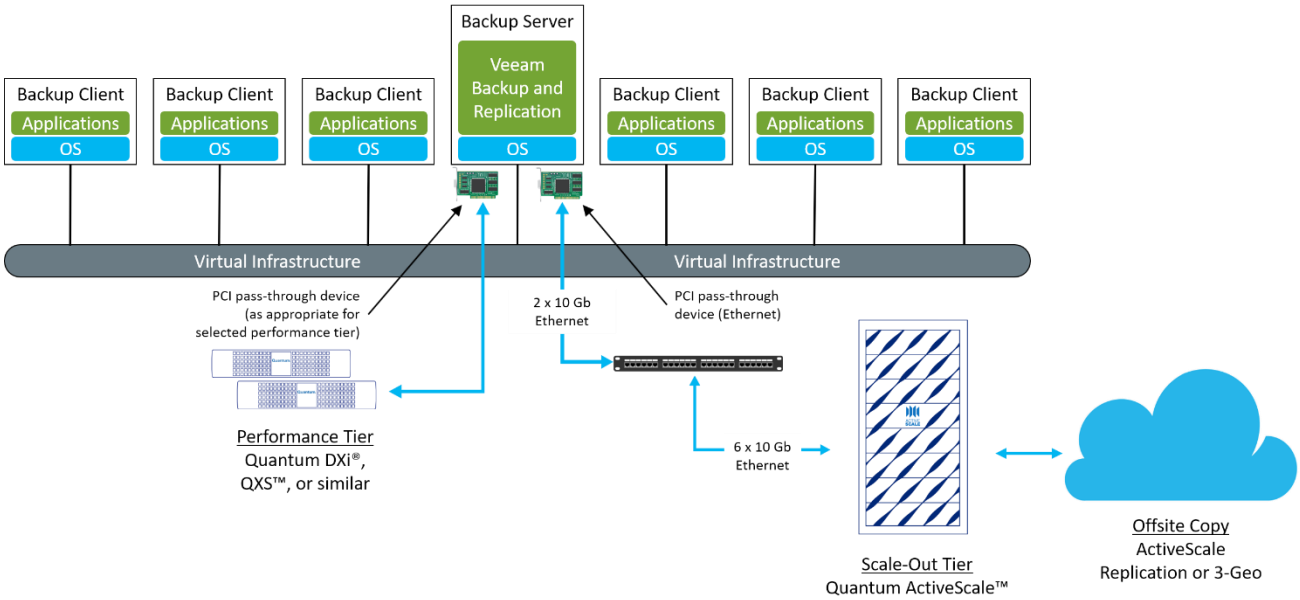

Network Requirements

Access to the ActiveScale S3 API is achieved by connecting to any one of several network interfaces. This is known as providing multiple end points and enables ActiveScale to support many thousands of parallel requests to the ActiveScale name spaces. Veeam however does not support multiple end points for any single object-based backup repository. For best performance, when configuring multiple object-based backup repositories, define them using different ActiveScale end points to minimize network congestion.

Quantum Configuration

ActiveScale Configuration

There are no special ActiveScale configuration considerations. Veeam however, will not create a bucket. The bucket(s) to be used by Veeam, must be created in advance using the ActiveScale GUI or any of a number of S3 tools that can create a bucket.

Note that when using VBR with immutability, object lock must be enabled on the bucket. When not using VBR immutability, object lock must not be enabled on the bucket.

Creating Buckets in ActiveScale

1. From ASView, select **Create Bucket**.

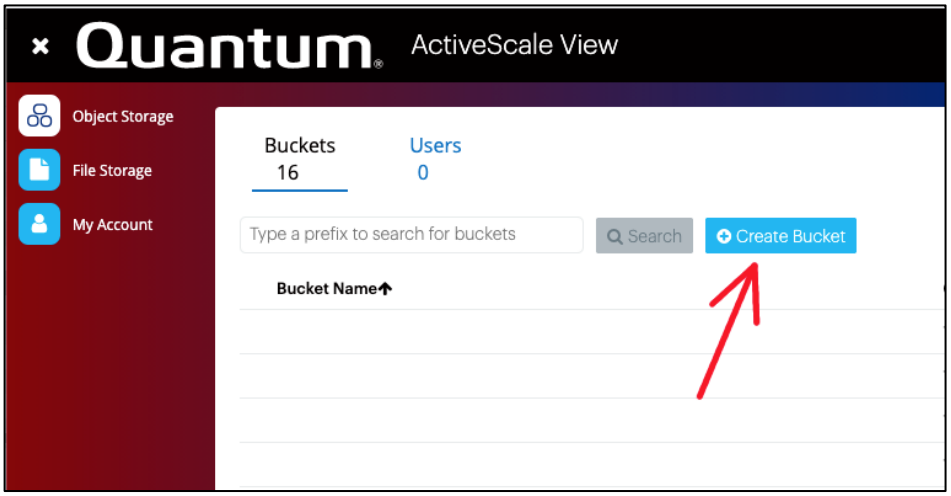

2. Enter the desired bucket name.

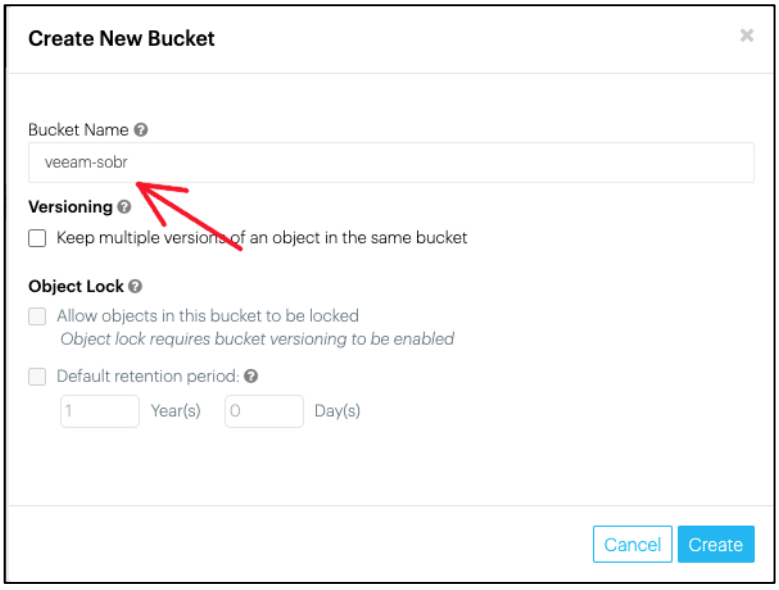

3. If the bucket is to be use with the VBR immutability feature
   a. Enable **Versioning**
   b. Enable **Object Lock**
   c. Ensure s default retention period is not specified

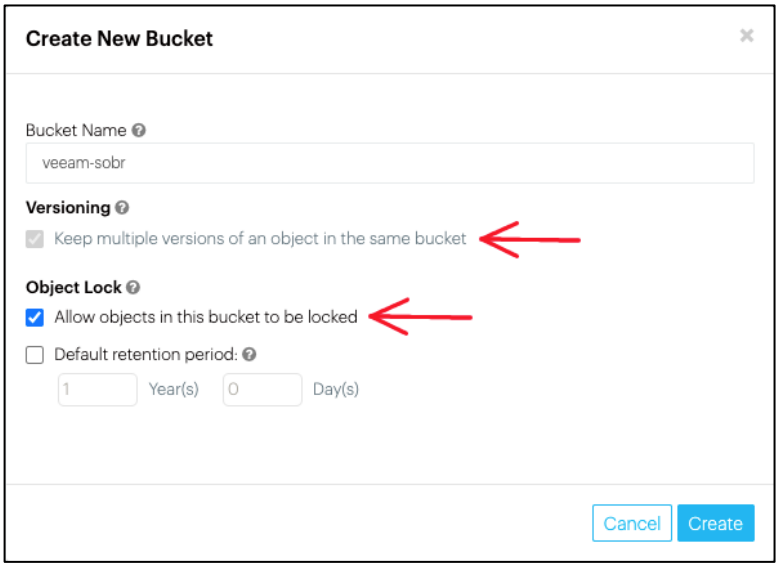

4. Select Create to complete the Create New Bucket wizard.

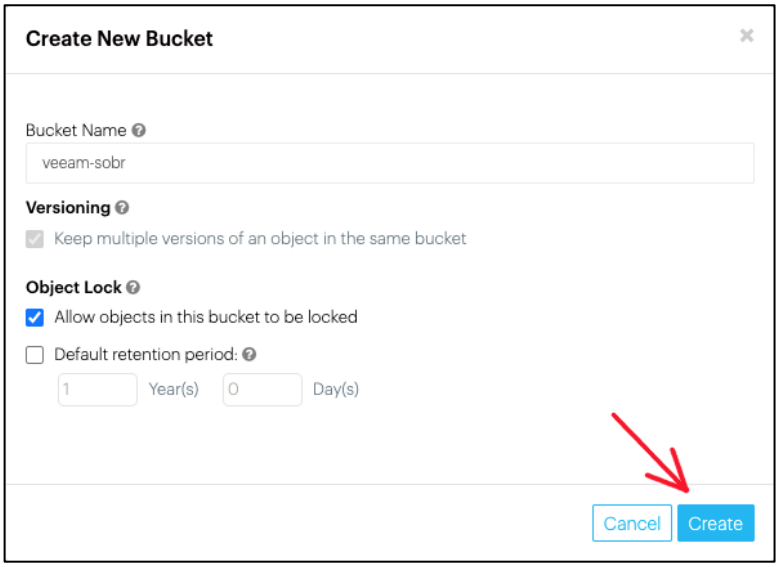

5. Confirm bucket properties by selecting the newly created bucket is ASView.

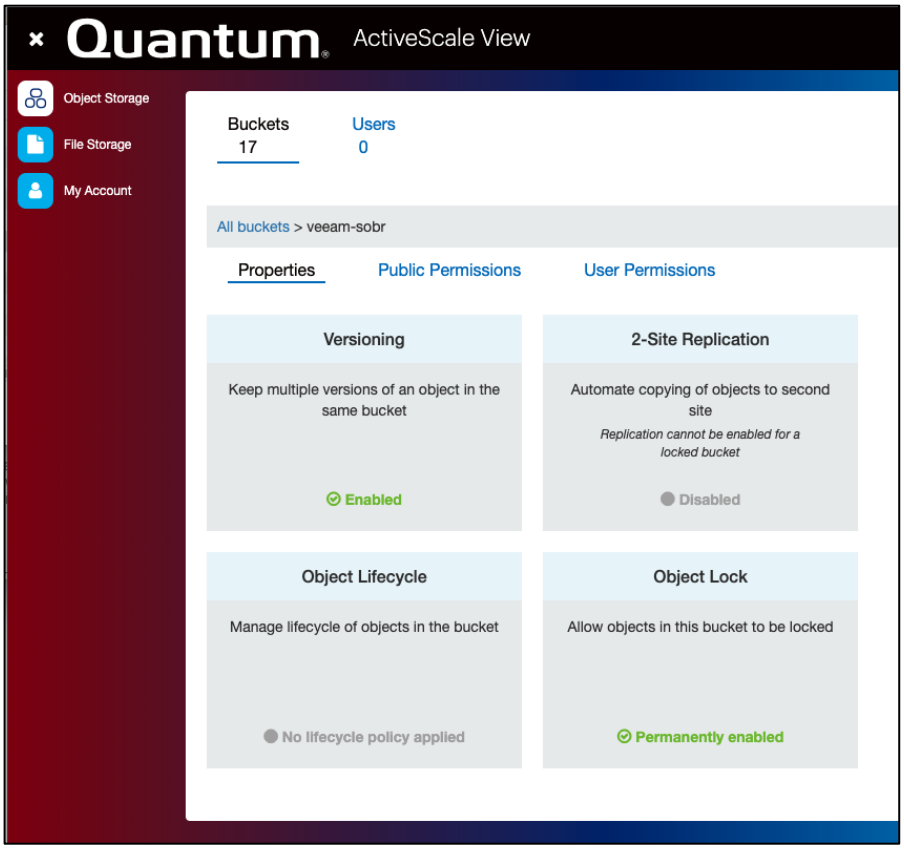

Veeam Backup and Replication Configuration

1. From the Veeam Console, Under **Backup Infrastructure** in Navigation Pane, right-click on **Backup Repository** and choose **Add Backup Repository.**

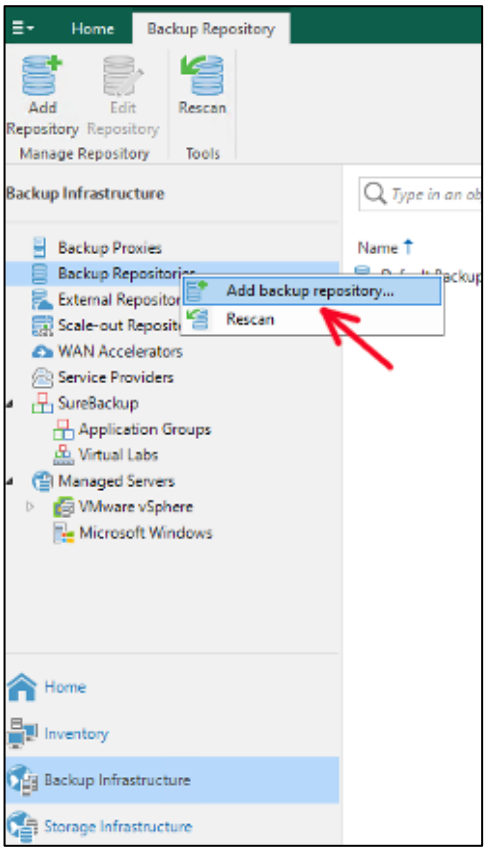

2. On the **Add Backup Repository** screen, select **Object Storage.**

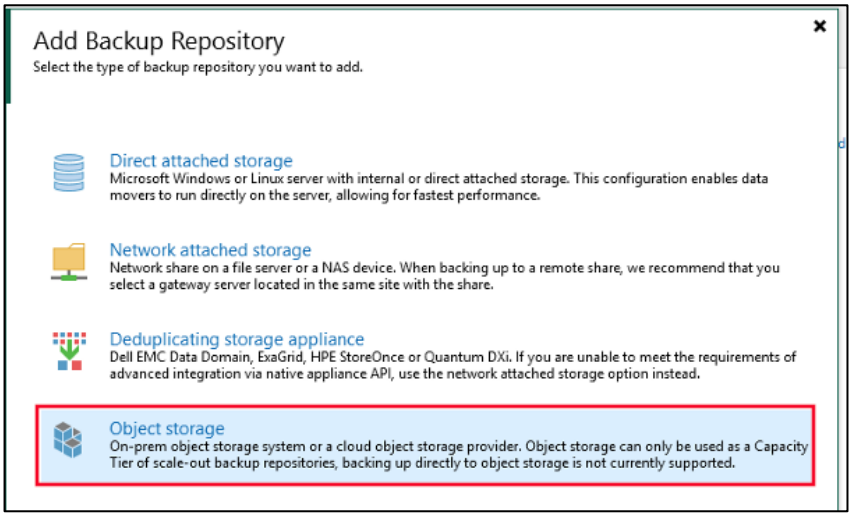

3. Choose S3 compatible object storage to start the Add Object Storage wizard.

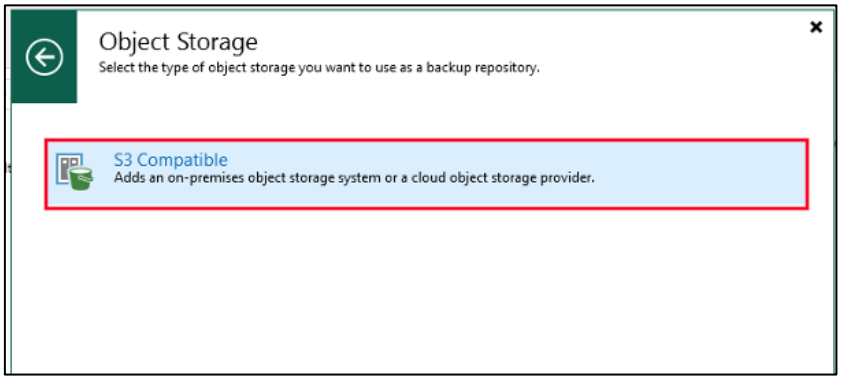

4. Enter in a name and description of the new archive repository. Click the **Next** button to continue.

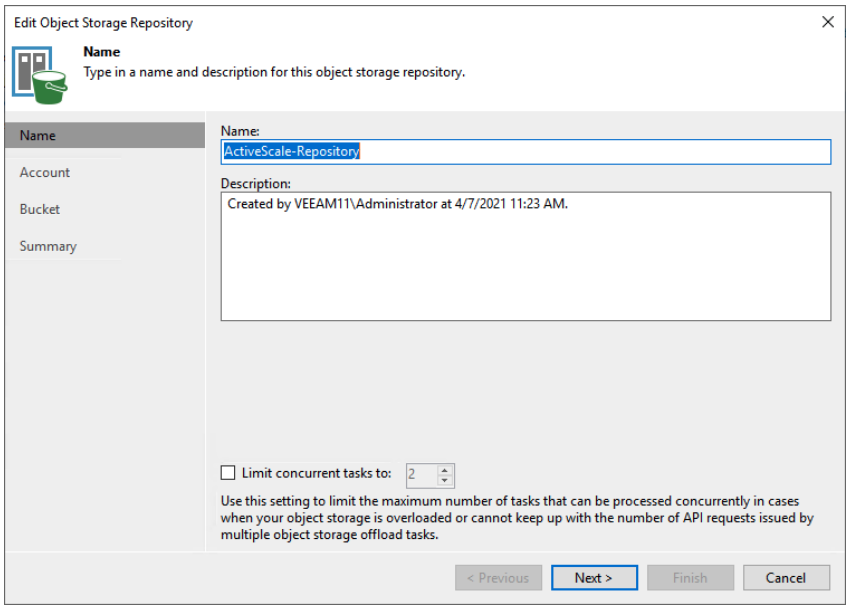

5. Enter your HTTPS endpoint address and region (be aware that Veeam only supports TLS/HTTPS encrypted targets). To add credentials, click on the Add button and enter the **Access key** and **Secret key** for your object storage and click **OK**.

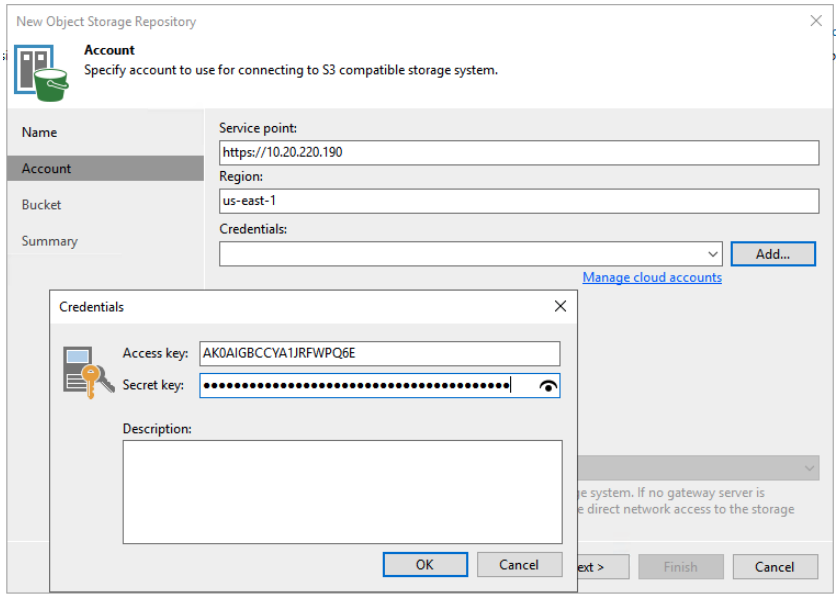

6. Leave the **Use gateway server** box unchecked and click **Next**. Accept any certificate warnings and **Continue**.

7. On the **Bucket** screen, identify the Bucket and Folder created for the repository. Click the **Next** button to proceed. Review the **Summary** screen and click **Finish** to create the object storage repository.

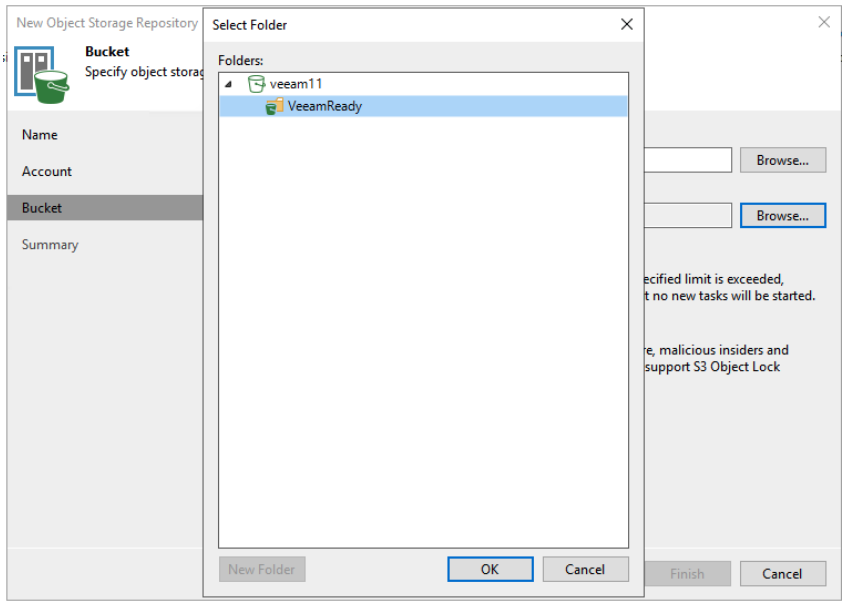

8. If Object Lock is too be used for this backup repository, check the **Make recent backups immutable for:** box and enter the desired value. Note that the ActiveScale bucket must have been created with the Object Lock feature enabled. Click **Apply** to continue

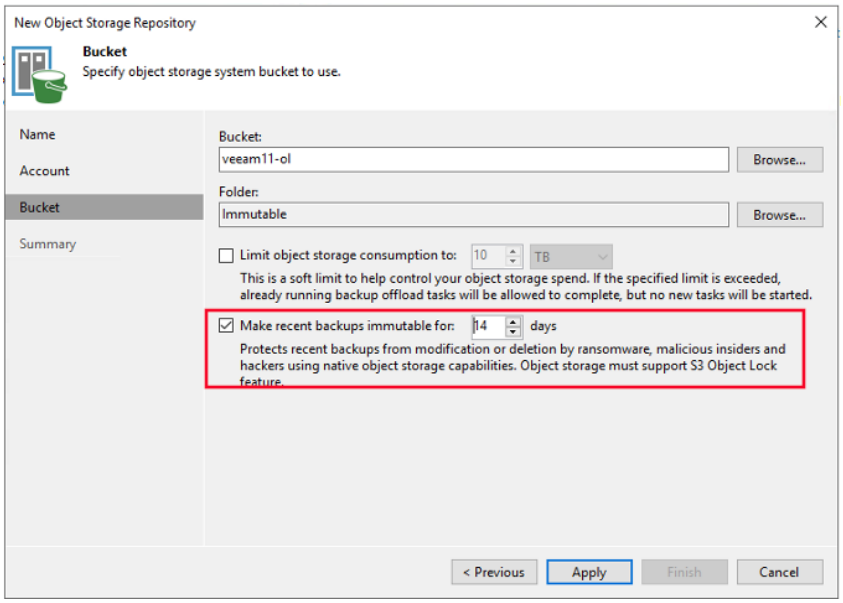

9. On the **Summary** screen, verify the storage repository was defined as intended. Notice it displays whether or not this backup repository is immutable. Click **Finish** to create the repository.

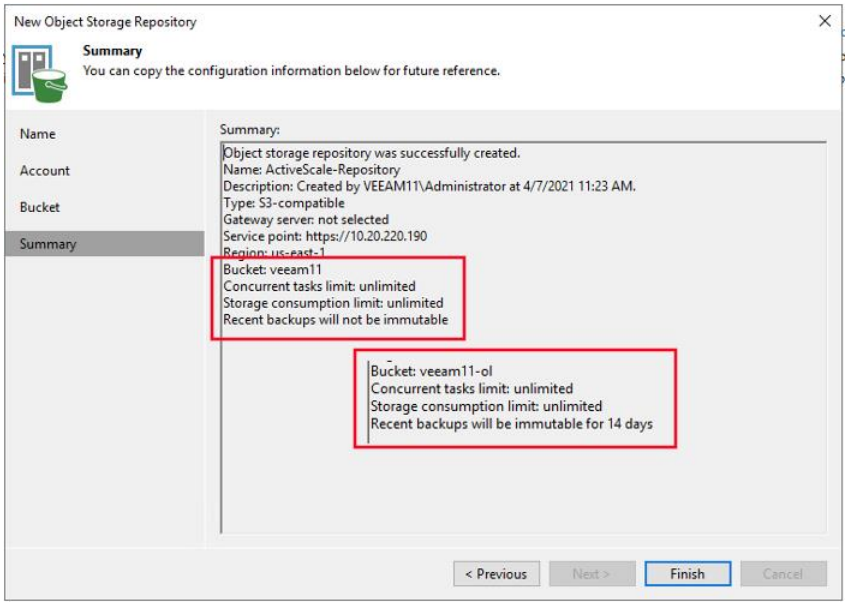

10. Under **Backup Infrastructure** pane, click on **Scale-out Repositories** and click the **Add Scale-out Repository** button on the top menu bar.

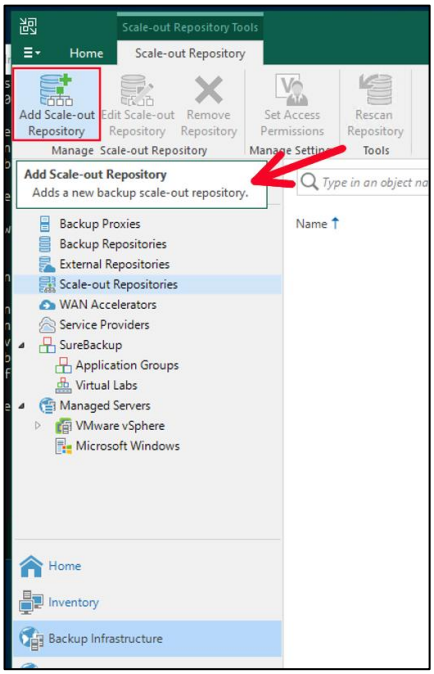

11. Type in a name and description for this new scale-out repository. Click **Next** to continue.

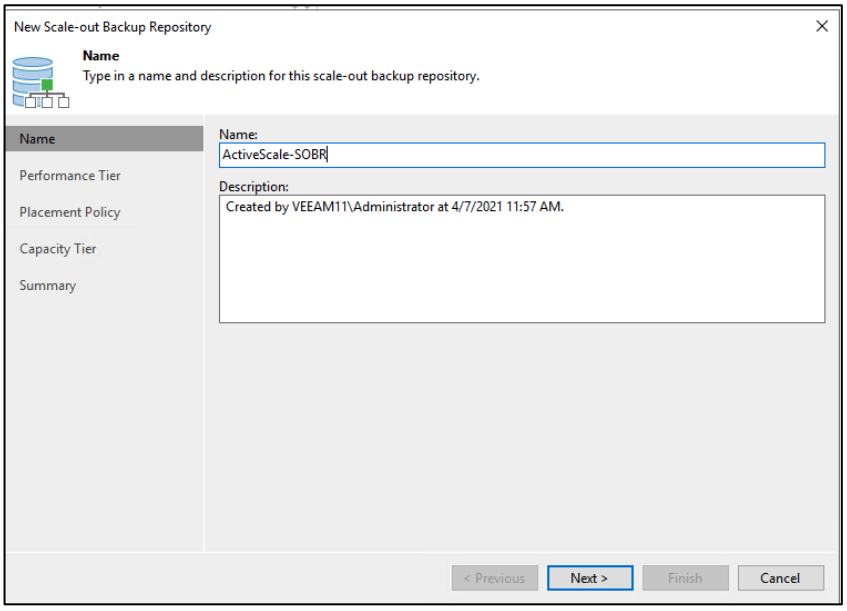

12. On the **Performance Tier** screen, click on the **Add** button to add a backup repository to the scale-out repository. Add in the primary repository that was created for your backups. Click **OK** to add the repository and click **Next** to continue.

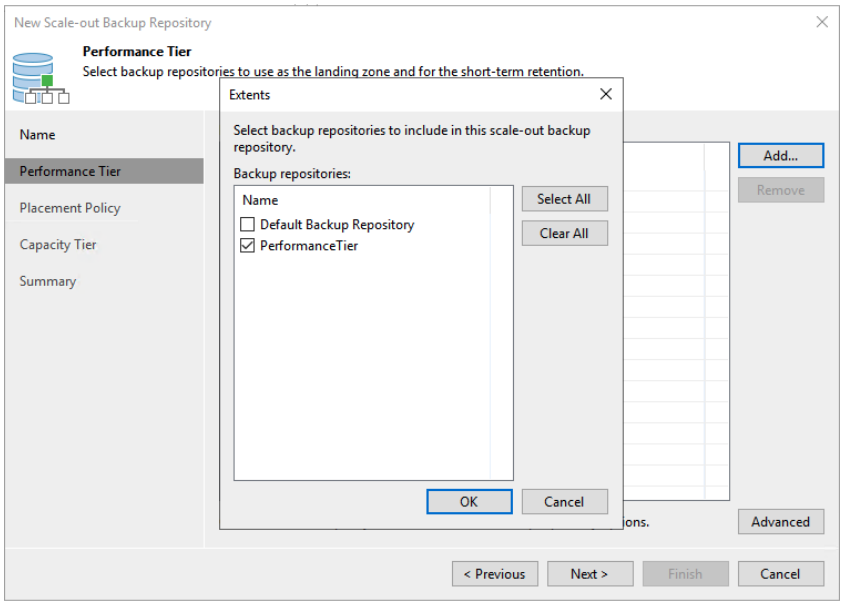

13. Click on the **Advanced** button to view advanced settings for the performance tier. Uncheck the **Use per-VM backup files** and click **OK**. Click **Next** to continue.

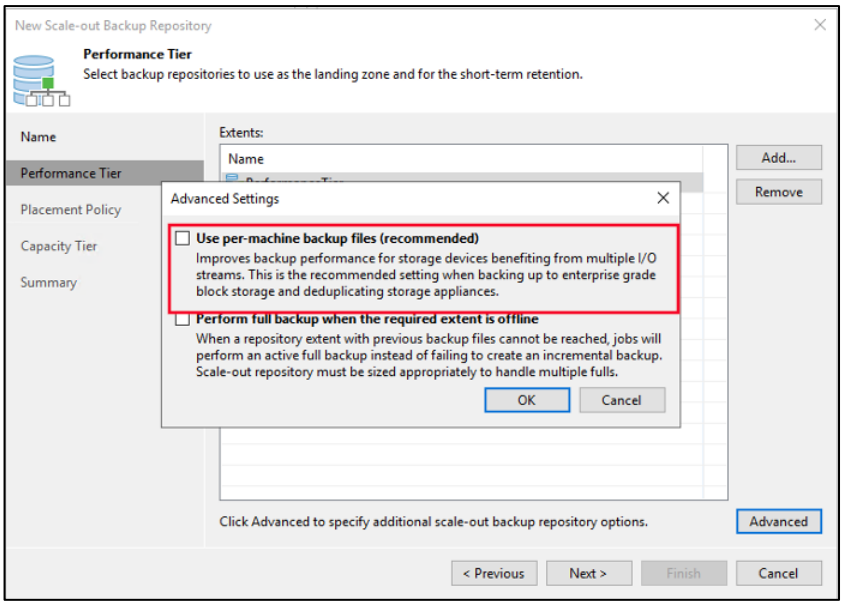

14. For the **Placement Policy** screen, keep the default **Data Locality** policy. Click Next to proceed.

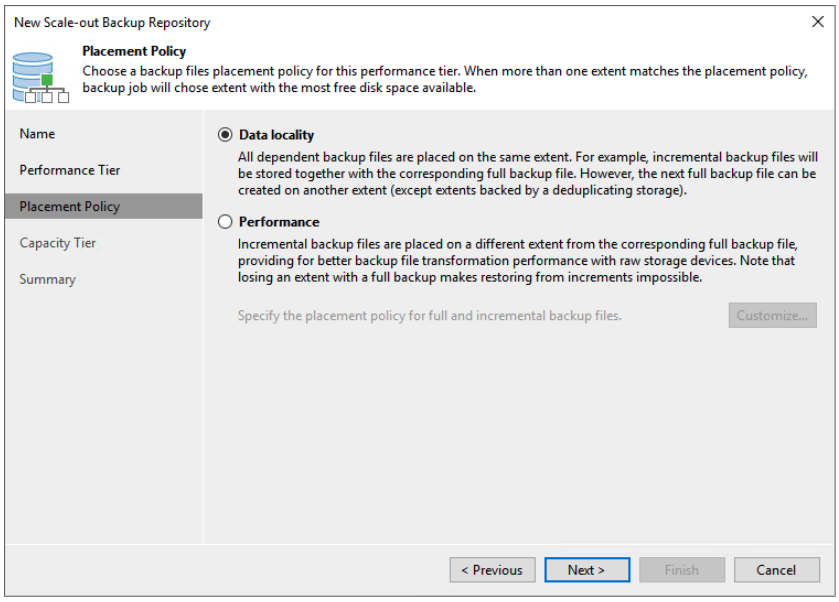

15. On the **Capacity Tier** screen, you have the option to either copy backups to object storage as they are created, or age backups to object storage after a defined period of time. Choose the option that best fits your desired use case. If using Veeam's immutability feature for protection against ransomware, it is recommended to copy backups to object storage as they are created. This ensure the backups are immediately protected.

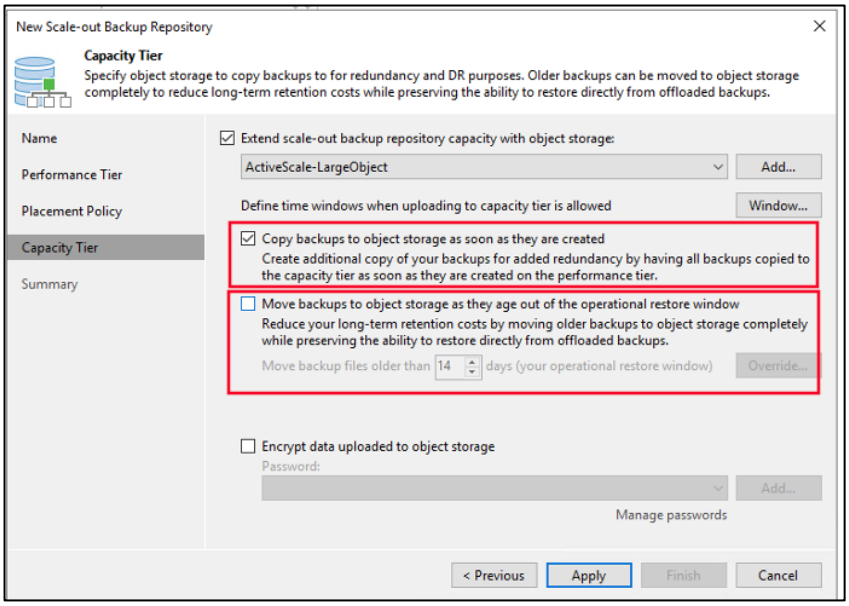

16. On the **Summary** screen, click **Finish** to create the Scale-out Repository

Validation Activities

Veeam uses a subset of the features provided by ActiveScale. The purpose of the integration testing is to ensure the S3 calls Veeam uses for its Scale Out Backup Repository feature are fully supported by ActiveScale. The following functionality was tested and confirmed successful to validate a Veeam/ActiveScale integration.

- File level backup to ActiveScale
- File level recovery from ActiveScale
- Offload backup image to ActiveScale
- Recover from ActiveScale
- Offload to ActiveScale and resume with intermittent network failure
- Offload to ActiveScale and restart after extended network outage
- Delete from ActiveScale
- Verify functionality when using Veeam managed encryption keys
- Validate Object Immutability Features

Considerations

Veeam performs backups by taking a snapshot of a VM backup client and mounting that snapshot on the Veeam server. Veeam is most efficient when the server is on the same ESX cluster as the backup clients. When the Veeam server is external to the backup clients, the snapshot is mounted as a remote NFS share which significantly reduces performance. Ensure the Veeam server has direct access to the data stores containing the VMs to be backed up.

Summary

Veeam Backup and Replication is a general-purpose application suited for a variety of use cases. There were no anomalous API level behaviors observed on either the Veeam or ActiveScale sides during validation of this reference architecture. A backup infrastructure based on Veeam and Quantum ActiveScale products is a viable and stable solution that can be deployed in virtually any environment.

References

The documents below were referenced to configure the software and systems for validation of this reference architecture.

Document Title	Download URL
ActiveScale S3 API Reference Guide	Link to Document
ActiveScale OS Admin Guide	Link to Document
Veeam User Guide for VMware vSphere	Link to Document

Document Feedback

For feedback, questions, and suggestions for improvements to this document contact the authors.

Contributor(s)

Version	Authors	Notes
1.0	Sherman Schorzman (Sherman.schorzman@quantum.com)	Initial Release – June 2020 ActiveScale 5.5 with Veeam 10
2.0	Sherman Schorzman (Sherman.schorzman@quantum.com) Tim Sherbak (timothy.sherbak@quantum.com)	July 2021 ActiveScale 5.7 with Veeam 11

