



REFERENCE ARCHITECTURE AND BEST PRACTICES

# Quantum ActiveScale with Veeam Backup and Recovery



## Abstract

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

# Table of Contents

Executive Summary.....	3
The Need for Backup.....	3
Solution Overview.....	3
Technology Summary.....	4
About Quantum ActiveScale.....	4
About Veeam Backup & Replication.....	4
About Storage Classes and Tiers.....	4
VBR Limitations .....	5
ActiveScale as a Backup Repository.....	5
ActiveScale as a Primary Repository.....	5
ActiveScale as a Scale-Out Backup Repository .....	5
SOBR Options .....	6
Reference Architecture .....	6
Basic ActiveScale/VBR Workflow Configuration .....	7
Considerations.....	8
Creating Buckets in ActiveScale.....	8
Create the SOBR – Local .....	10
Create the 3 Repositories to be Used in the SOBR.....	10
Create the SOBR for the Local Site .....	10
Create the SOBR – Remote.....	13
Create 2 Repositories to be Used in the SOBR for the Remote Site.....	13
Create the SOBR for the Remote Site.....	13
Create the Backup Job .....	14
Create the Backup Copy Job .....	16
Verify the Workflow .....	19
Unstructured Data .....	20
Performance Tweaks.....	21
Increase VBR Block Size .....	21
Load Balancing.....	21
DNS Round Robin.....	22
Summary .....	23
References .....	23

## Executive Summary

A good backup and recovery plan is essential for protecting your valuable data from unexpected events. This could be anything from a hardware failure to a cyberattack or even an accidental deletion. Having a plan in place ensures you have copies of your data stored securely and can restore them quickly if needed. This minimizes downtime, prevents potential financial losses, and reduces the stress of data loss. In short, a good backup and recovery plan gives you peace of mind knowing your data is safe.

This guide equips you with the best practices to build a robust backup environment based on [Quantum ActiveScale®](#) with [Veeam Backup and Replication 12.1](#) (VBR). We'll walk you through the setup of a typical deployment and address tuning parameters to ensure you get the best performance from both ActiveScale and VBR.

## The Need for Backup

The digital world is full of unforeseen storms that can threaten your data, and a backup and recovery plan acts as your umbrella and raincoat. Here's why it's crucial:

- **Safeguards Against Data Loss:** Hard drive crashes, accidental deletions, power surges—these events can wipe out your data in an instant. Backups create copies of your data, stored securely away from the original, so you can retrieve it if disaster strikes.
- **Combats Security Threats:** Cyberattacks like ransomware can lock you out of your data, demanding a ransom for its return. Backups ensure you have a clean, uninfected copy to restore from, minimizing damage and downtime.
- **Protects from Human Error:** We've all accidentally hit "delete." Backups provide a safety net, allowing you to recover files you've unintentionally erased.
- **Ensures Business Continuity:** Data loss can cripple a business. A robust recovery plan minimizes downtime, allowing you to get back up and running quickly, avoiding lost productivity and potential revenue.
- **Maintains Peace of Mind:** Knowing your data is safeguarded provides a sense of security. Backups eliminate the constant worry about losing critical information.

In essence, backups are your insurance policy against the unexpected. They offer a safety net, ensuring your valuable data is always protected and readily available, allowing you to focus on what matters most.

## Solution Overview

Data centers are increasingly becoming a collection of virtualized servers. Veeam's flagship product, *Veeam Backup & Replication (VBR)*, seamlessly integrates with hypervisors to back up 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 process overhead.

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

## Technology Summary

The table below lists the technology components that make up the ActiveScale/VBR solution stack 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	6.3 or higher
Quantum ActiveScale Deployment Model	Any ActiveScale appliance or supported software deployments
VMware ESXi	Although ESXi 7.0 was used in the deployment of this reference architecture, any hypervisor supported by Veeam is suitable.
Veeam Backup and Replication	12.1

### About Quantum ActiveScale

ActiveScale object storage provides a new, innovative approach to creating a simple, ‘always-on’ data repository that scales when and how you need it to—with the extreme data durability, accessibility, and security required of petabyte-scale growth. And ActiveScale Cold Storage reduces the cost of storing your cold data sets by up to 80% via an on-prem S3 Glacier tier.

### About Veeam Backup & Replication

VBR 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.

### About Storage Classes and Tiers

Both ActiveScale and VBR have the concept of storage classes or storage tiers. These classes and tiers are just a way to indicate the type of storage being used for backup jobs. With regards to how VBR interacts with object storage, classes and tiers can be mapped directly to an S3 storage class. The below table maps the ActiveScale Storage Classes and VBR Tiers referenced in this document to the appropriate S3 Storage Class.

ActiveScale Storage Class	VBR Tier	S3 Storage Class
Active Storage Class	Performance Tier	S3 Standard Class
Active Storage Class	Capacity Tier	S3 Standard Class
Cold Storage Class	Archive Tier	S3 Deep Glacier Class


Note: You can assign block, NAS, or object storage to a VBR Performance Tier. However, when ActiveScale is used for the Performance Tier, backup data will be PUT to ActiveScale’s Active Storage Class (S3 Standard Class).

## VBR Limitations

VBR has different features to protect different types of data. It also supports many types of storage to be used as backup repositories. But not all data can be protected in any repository. It is important to understand how VBR interacts with object storage, specifically ActiveScale, and be clear on certain VBR limitations with object-storage-based repositories. More information on object-storage-based repositories can be found here: [Object Storage Repositories](#)

## ActiveScale as a Backup Repository

ActiveScale is a storage target for VBR. When using ActiveScale as a storage target—either Active Storage Class or Cold Storage Class—VBR will use a base object inside a bucket as a root for the backup repository. Although technically incorrect, the VBR wizard refers to and displays that object as if it were a “folder”. It is the unique combination of bucket/base object (folder) that VBR considers a repository. Therefore, there are no restrictions on having multiple repositories in a single bucket. Similarly, if you choose, you could create multiple buckets—each with a single VBR repository—if that is more appropriate for your use case. From an ActiveScale perspective, the ActiveScale architecture is such that there is no performance benefit or penalty with either approach. Note, however, that the VBR repository setup wizard will create the repository inside the bucket but will not create the bucket itself. The bucket must be created in advance. The below images show the VBR repository wizard using both approaches.



## ActiveScale as a Primary Repository

The simplest type of repository in VBR is a single storage location used as a standalone repository for a backup job. It is referred to as a primary repository. Previous versions of VBR had more restrictions on storage types that could be used for a primary repository. Starting with version 12, VBR allows ActiveScale to be used as a primary backup repository.

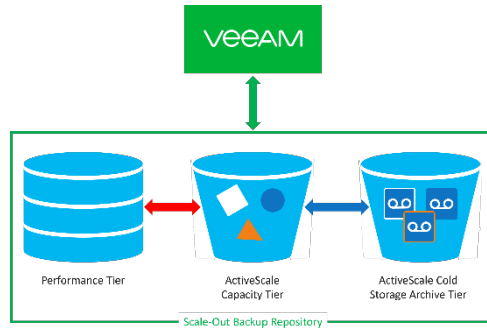
## ActiveScale as a Scale-Out Backup Repository

VBR has a special repository called a Scale-Out Backup Repository (SOBR). This is a logical construct that combines multiple backup repositories into a single repository with multiple extents. This allows for scalability, data storage efficiency, and performance optimization. VBR allows you to define how data is moved or copied to the various extents. The SOBR has a minimum of 2 extents but is capable of more. An SOBR will consist of 1 mandatory Performance Tier, 1 optional Capacity Tier, and 1 optional Archive Tier. Backup data will always be written to the Performance Tier. From there, the data will be copied, moved, and expired from the remaining extents based on VBR rules for that SOBR.


**IMPORTANT:** You may have the ability to move or copy data from one extent to another outside of VBR, for example using ActiveScale lifecycle policies or replication. Although this may seem tempting and may be even more efficient, it is unsupported. This will cause VBR to lose track of the data and you will have destroyed the integrity of the SOBR with no recovery path. All data movement must be done by VBR.

### SOBR Options

When creating your SOBR to leverage the VBR Archive Tier, you can choose to use 2 or 3 storage extents. The first option pictured here is an SOBR with 3 storage extents: a Performance Tier using any suitable storage such as direct attached storage or a NAS share, and the Capacity and Archive Tiers use a bucket in ActiveScale.



Another option is to use a 2-extent approach. Given the inherent durability and performance ActiveScale provides, this is a perfectly sound approach. The below image depicts an SOBR with 2 storage extents. Because VBR allows you to use ActiveScale’s Active Storage Class as the VBR Performance Tier, with this arrangement, backup data sets can be copied from the Performance Tier directly to the Archive Tier without the need for a Capacity Tier.



## Reference Architecture

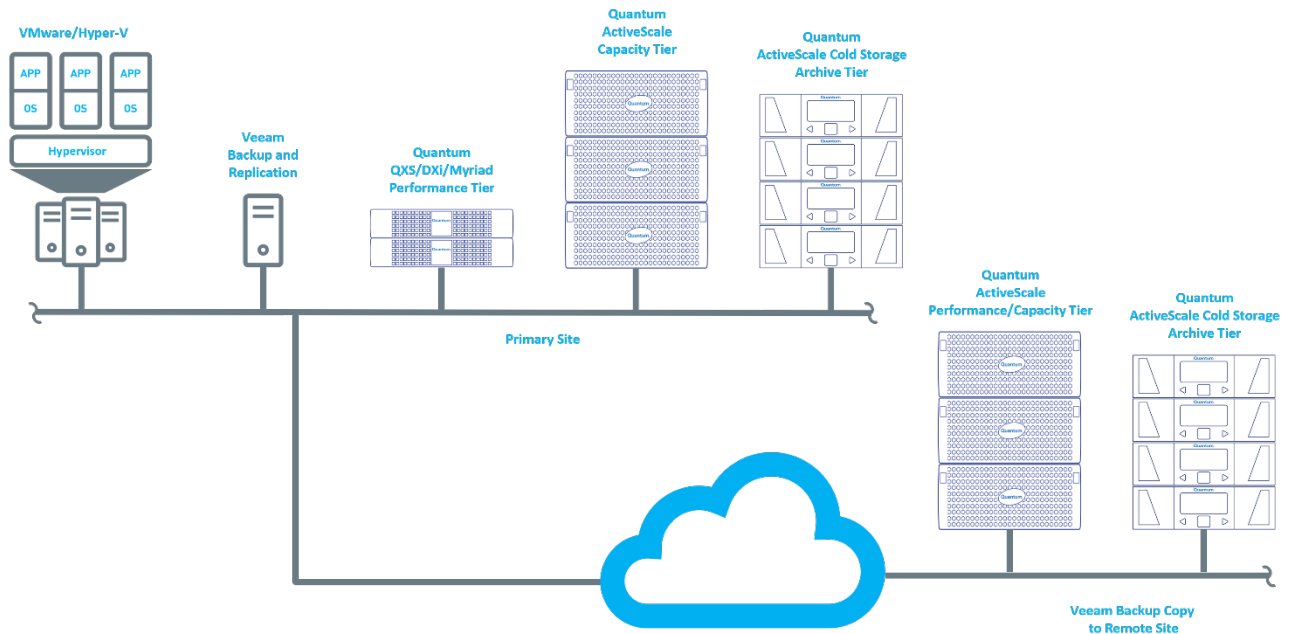
The reference architecture and setup described in this document are based on a typical environment with a primary location and remote site for disaster recovery. A combination of high-performance block storage (Quantum QXS™ or Quantum Myriad®) along with ActiveScale’s Active and Cold Storage Classes are used to ensure data is protected and rapidly available for the short to medium term and secured on low-cost Cold Storage for the long term.

At the primary location, for rapid backup and recovery, backup jobs should use an SOBR that uses high-performance block storage for the short-term Performance Tier to allow for rapid recovery of the most recent backup. The SOBR

should use ActiveScale’s Active Storage Class as the Capacity Tier for medium-term storage, and finally ActiveScale’s Cold Storage Class as the Archive Tier for long-term storage.

Because the remote site is intended for disaster recovery, high-performance block storage is not used. Rather, a 2-extent SOBR is used with ActiveScale’s Active Storage Class as the Performance Tier and its Cold Storage Class as the Archive Tier.

Based on user-defined parameters, Active Class and Cold Storage Class become 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.



## Basic ActiveScale/VBR Workflow Configuration

To create the VBR workflow described above, we will define an SOBR in the primary site that contains 3 storage extents or tiers: high-performance block for the Performance Tier, Active Storage Class for the Capacity Tier, and Cold Storage Class for the Archive Tier. At the remote site, an SOBR with 2 extents will be defined. This site will use Active Storage Class as the Performance Tier and Cold Storage Class as the Archive Tier. The remote SOBR, being used for disaster recovery only, will not use high-performance block storage.

The following steps will walk you through creating ActiveScale buckets, the SOBRs, and the VBR jobs required for the above workflow.



This document assumes the reader is familiar with VBR and knows how to create backup repositories and an SOBR. For more detailed instruction on how to create repositories and SOBRs, refer to the following How-To Guide:

[ActiveScale Cold Storage as a VBR Archive Tier](#)

Note: The above is hosted internally at Quantum. If required, your Quantum contact can provide you with the document.

## Considerations

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 S3 tool that can create buckets.


There are important VBR considerations when using ActiveScale as a VBR backup repository. VBR supports ActiveScale's Active Storage Class as a backup target. VBR also supports ActiveScale's Cold Storage Class as an Archive Tier repository. However, the Archive Tier has limitations; for example, the Archive Tier is not a standalone repository—it must be an extent of an SOBR. Also, not all VBR features support the Archive Tier; for example, VBR has a feature called Unstructured Data Backup designed to protect NAS storage. Because this feature has an instant recovery option, high-latency storage such as the VBR Archive Tier is not supported. More on unstructured data and the Archive Tier later.

VBR does support ActiveScale versioning and object-lock features, but it does not support ActiveScale lifecycle policies. To keep the VBR database in a consistent state, it must assign retention periods and be the entity that deletes expired objects. Therefore, VBR supports versioning on a bucket only when Object Lock is enabled as well.

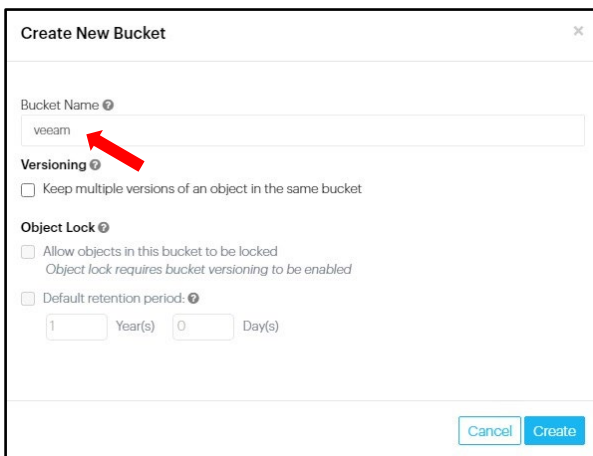
VBR does not support ActiveScale replication.

## Creating Buckets in ActiveScale

1. At the primary location, log into ASView and select **Create Bucket**.

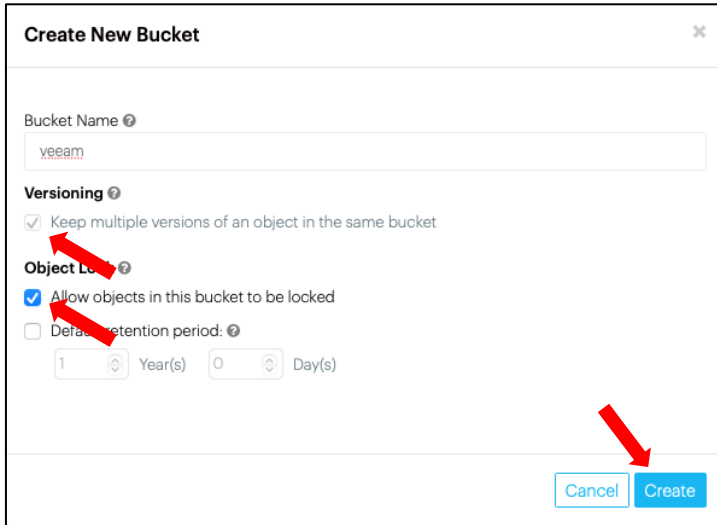


2. Enter the desired bucket name.



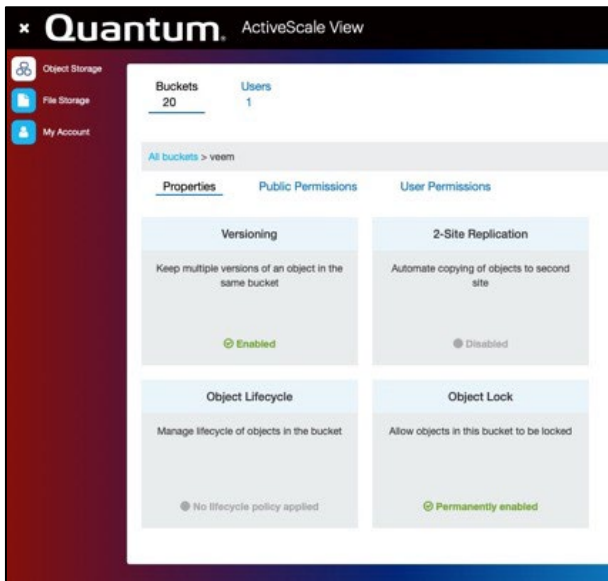


3. If the bucket is to be used with the VBR immutability feature:
  - a. Enable **Versioning**
  - b. Enable **Object Lock**
  - c. Ensure a default retention period is not specified
  - d. Click Create to complete the Bucket Creation Wizard



Note: VBR has an internal retention feature that leverages ActiveScale Object Lock. However, ActiveScale lifecycle policies are not supported. To keep the VBR database in a consistent state, it must assign retention periods and be the entity that deletes expired objects. Therefore, when enabling Object Lock, do not assign a default retention period. Also, keep in mind that VBR does not support the use of retention on a bucket that does not have Object Lock enabled.

4. Confirm bucket properties by selecting the newly created bucket in ASView.



5. Repeat this procedure on the ActiveScale system at the remote location.

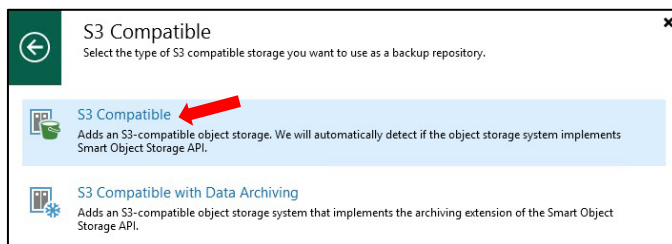
## Create the SOBR – Local

Provisioning Veeam entails creating the 3 backup repositories to be used as the 3 extents of your SOBR and creating the SOBR itself. There are a few things to watch for when creating the SOBR, so go through the following operations to correctly provision Veeam.

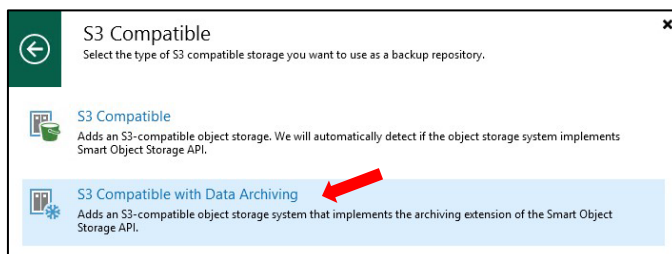
Note: When creating the ActiveScale repositories, you can use the same bucket or different buckets. There is no performance impact either way.

### Create the 3 Repositories to be Used in the SOBR

1. Create a repository to be used for the Performance Tier of the SOBR. Feel free to create a block based repository for the Performance Tier that is convenient for you.
2. Create a repository to be used as the Capacity Tier in the SOBR. This repository should be of type **S3 Compatible** (not S3 Compatible with Data Archiving). The repository creation wizard default options are OK for this repository.

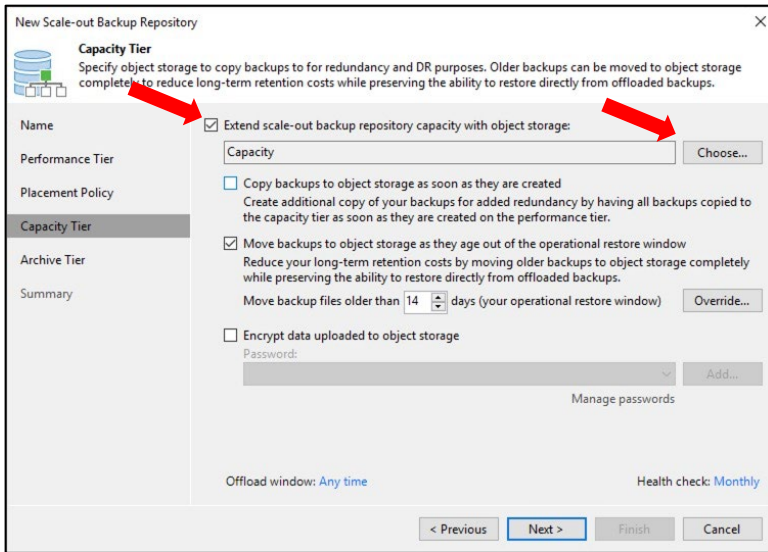


3. Create a repository to be used as the Archive Tier in the SOBR. This repository should be of type **S3 Compatible with Data Archiving**. The repository creation wizard default options are OK for this repository.



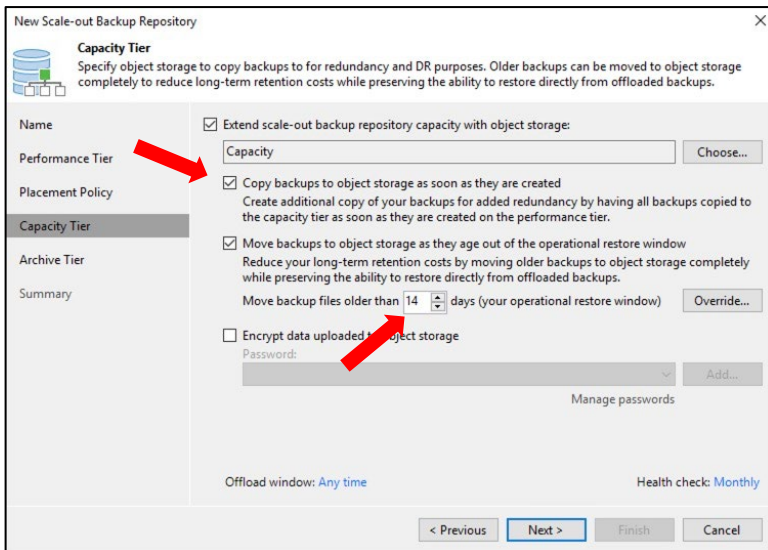
### Create the SOBR for the Local Site

1. Create the SOBR using the Add Scale-Out Backup Repository wizard.
2. Assign the Performance Tier using the Performance Tier repository created earlier, accepting the wizard defaults.
3. In the Capacity Tier section, you must indicate that you want to scale out to object storage by checking the *Extend scale-out backup repository capacity with object storage* checkbox.



4. You can choose when backups should get copied to the Capacity Tier (ActiveScale’s Active Storage Class). You can choose to age them out of the Performance Tier after some period of time or copy them to the Capacity Tier immediately. To get backups safely stored on ActiveScale as soon as possible, it is recommended you check the box indicating you want to *Copy the backups to Capacity Tier as soon as they are created*.


The following Capacity Tier section tells VBR to make a safe copy of backups to ActiveScale immediately, but not to purge them from the Performance Tier for 14 days. You can choose to leave a copy of backups in the Performance Tier for the number of days that is appropriate for your use case.



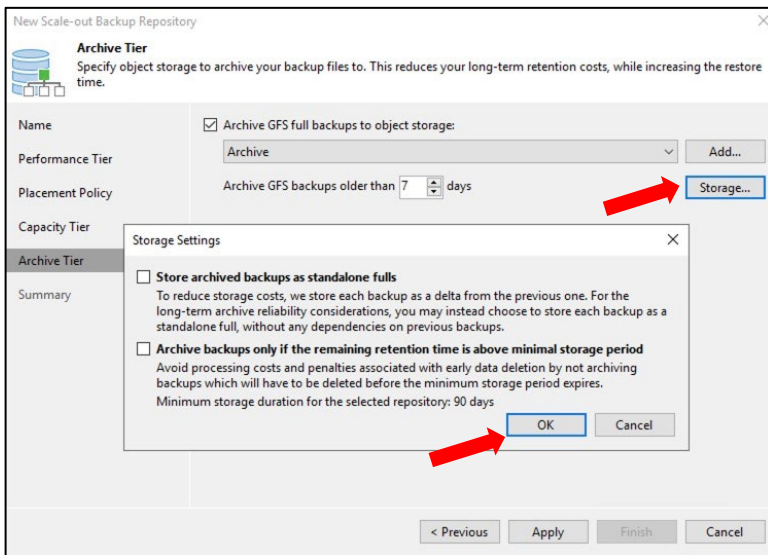
5. In the Archive Tier section, you enable offloading to ActiveScale Cold Storage by selecting the checkbox next to *Archive GFS full backups to object storage*. You then add a repository that was created with the type *S3 Compatible with Data Archiving*.

VBR does not simply copy objects from the Capacity Tier to the Archive Tier. Rather, it identifies backups that are eligible for archiving. From those backups, it creates Grandfather Father Son (GFS) restore points. It's the GFS restore points that get PUT to the Archive Tier. Eligible backups are closed backup chains that are older than the value assigned in the Archive Tier section of the wizard.

For more information on VBR GFS policies, refer to the following link: [VBR GFS Policies](#)



6. Before proceeding, there are options that can provide cost benefits when archiving to a public cloud provider. These options don't make sense when archiving to ActiveScale Cold Storage. While still in the Archive Tier section, click on the storage button. You will be presented with 2 storage setting options. Uncheck both options and click OK.



7. At this point you can complete the wizard, accepting the defaults.

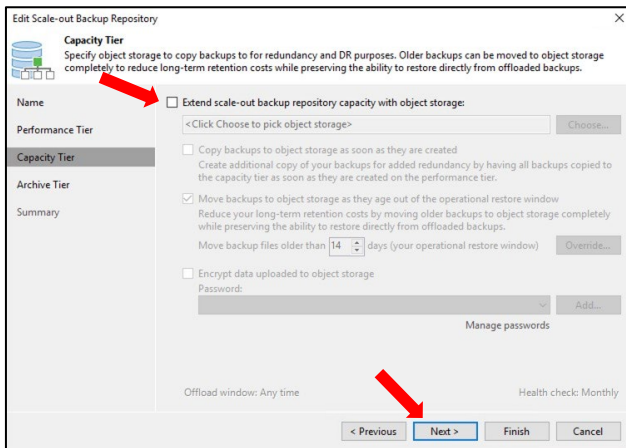
## Create the SOBR – Remote

### Create 2 Repositories to be Used in the SOBR for the Remote Site

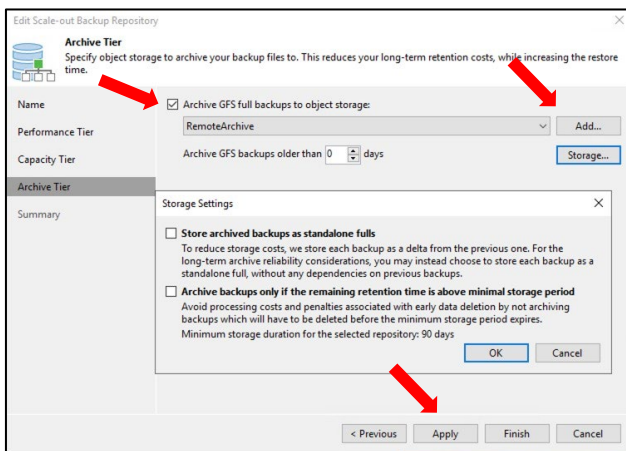
1. Create 2 repositories using bucket(s) on the ActiveScale system at the remote site: 1 repository of type **S3 Compatible** which will be used as the Performance Tier, and 1 repository of type **S3 Compatible with Data Archiving** to be used as the Archive Tier.

### Create the SOBR for the Remote Site

1. Create the SOBR using the Add Scale-out Backup Repository wizard.
2. Assign the Performance Tier using the Performance Tier repository created for the remote site.
3. In the Capacity Tier section, leave the *Extend scale-out backup repository capacity with object storage* checkbox unchecked. Click Next to move to the Archive Tier section.



4. In the Archive Tier section, assign the repository you created for the remote site Archive Tier, and set the GFS policy the same as was created for the local site.



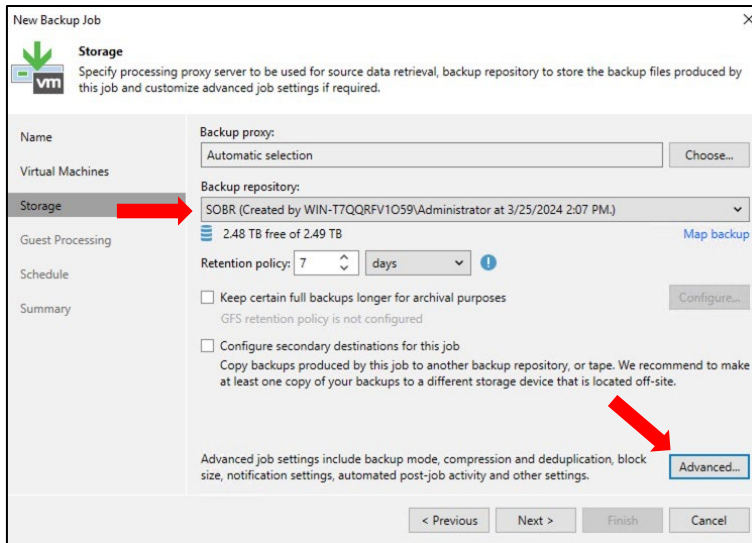
5. At this point you can complete the wizard, accepting the defaults.



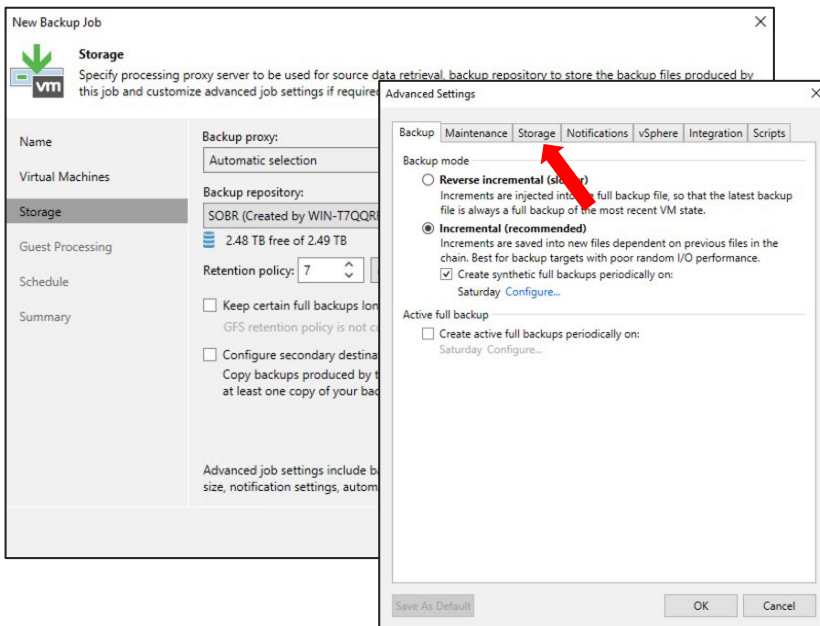
## Create the Backup Job

Creating a Backup Job that uses ActiveScale with the above SOBR is mostly straightforward. During the Backup Job creation, you can achieve significant performance benefits by increasing the block size Veeam uses when writing data to ActiveScale.

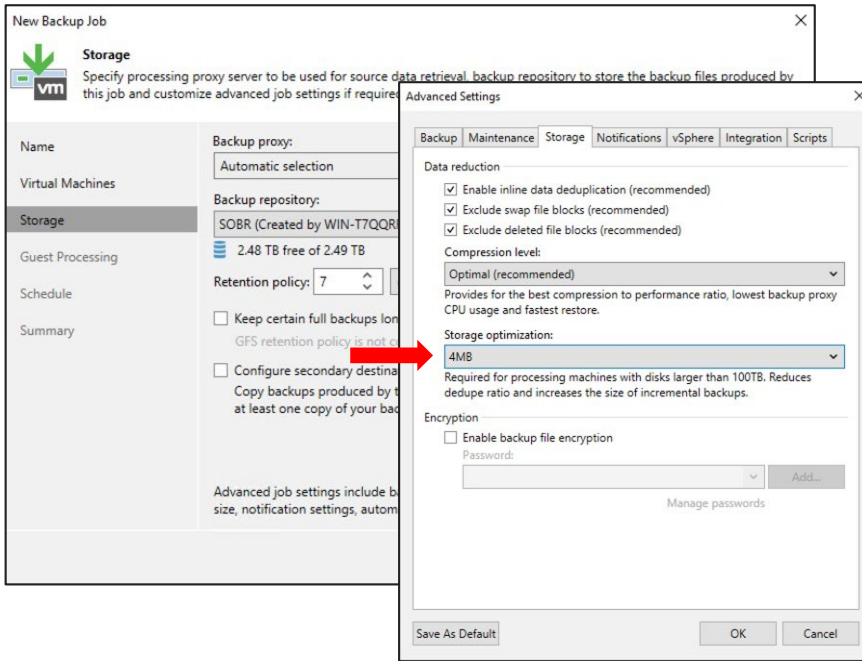
1. Start the Backup Job wizard by creating a new Backup Job.
2. Work through the wizard until you reach the Storage section. Select the SOBR that was created for the local site. At this point, click on Advanced.



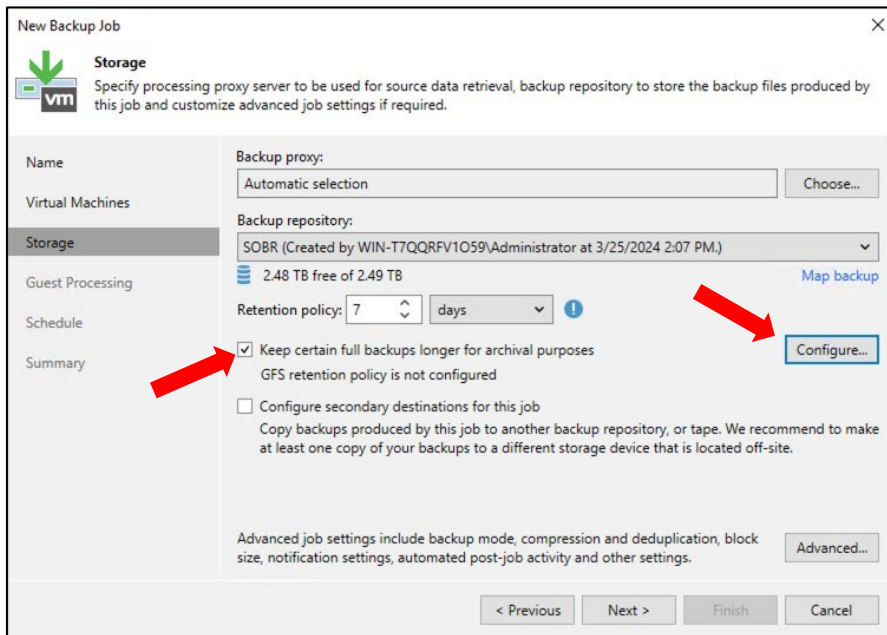
3. You will be presented with an Advanced Settings window. On the Advanced Settings window, select the Storage tab.



- 4. On the Storage tab, change the storage optimization block size from the default 1MB to 4MB. Click OK to save the block size setting.

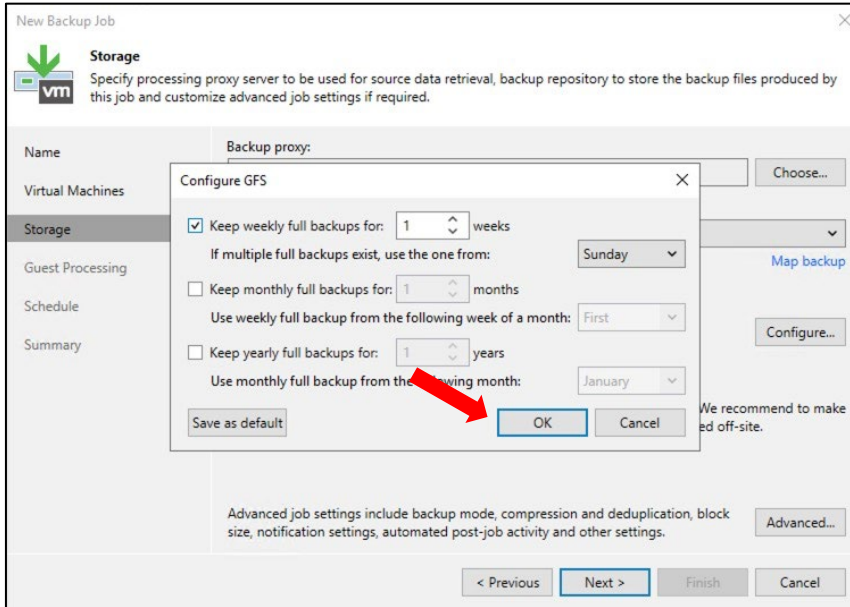


- 5. When using the Archive Tier, VBR creates GFS restore points and writes those restore points to ActiveScale's Cold Storage Class. To enable this feature, check the box next to *Keep certain full backups longer for archival purposes*. Then click *Configure* to define the GFS policy.





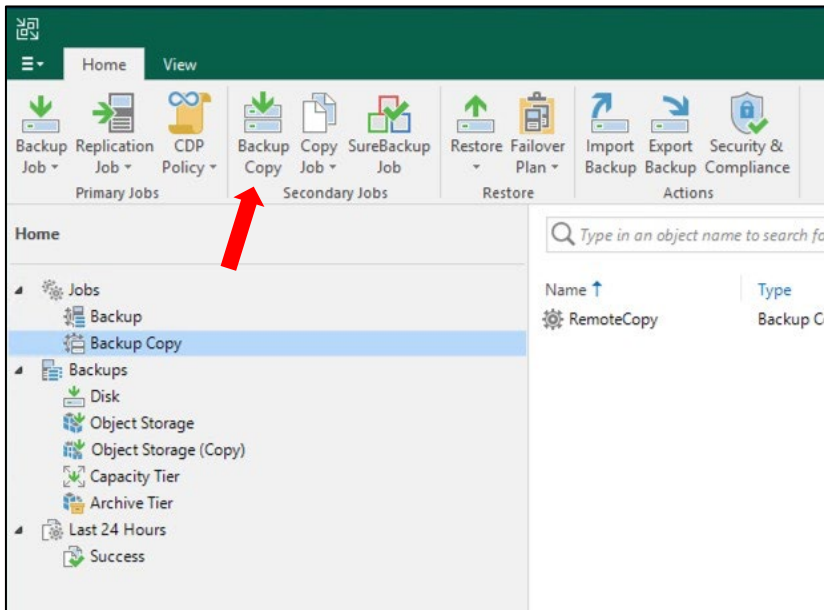
6. You are presented with a Configure GFS dialog. In this example, once a week on Sunday, VBR will look for eligible backups, create a GFS restore point from those backups, and write them to ActiveScale Cold Storage.



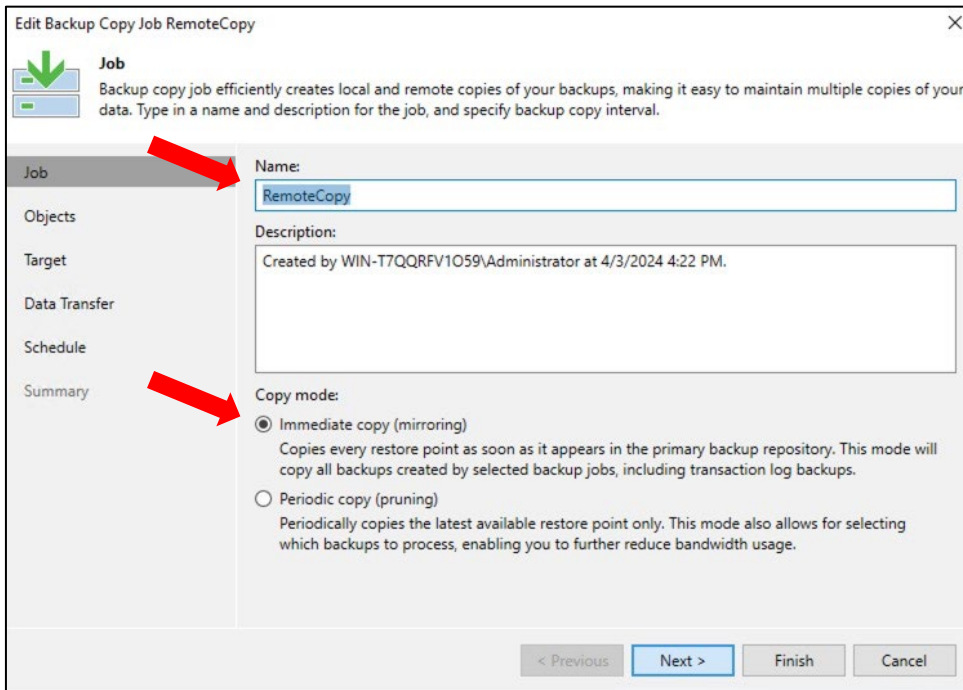
7. Continue with the wizard, accepting the defaults.

## Create the Backup Copy Job

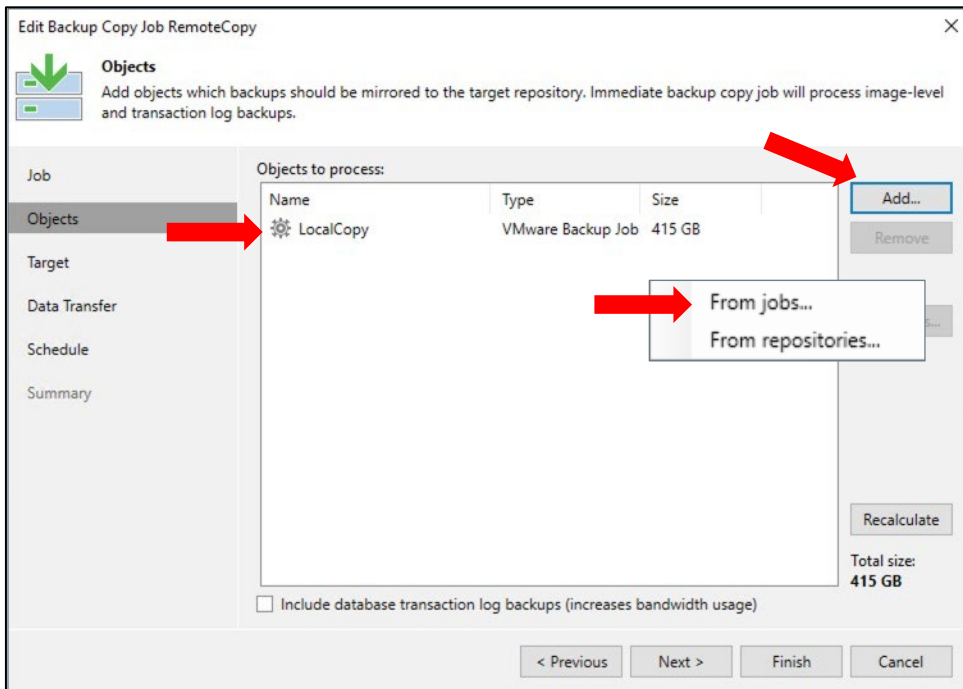
1. From the VBR console, start the Backup Copy Wizard.




2. Give the Job a name and select the copy mode that is appropriate for your use case. In this case, mirroring was selected. This will cause a copy to take place immediately after the original Backup Job completes.



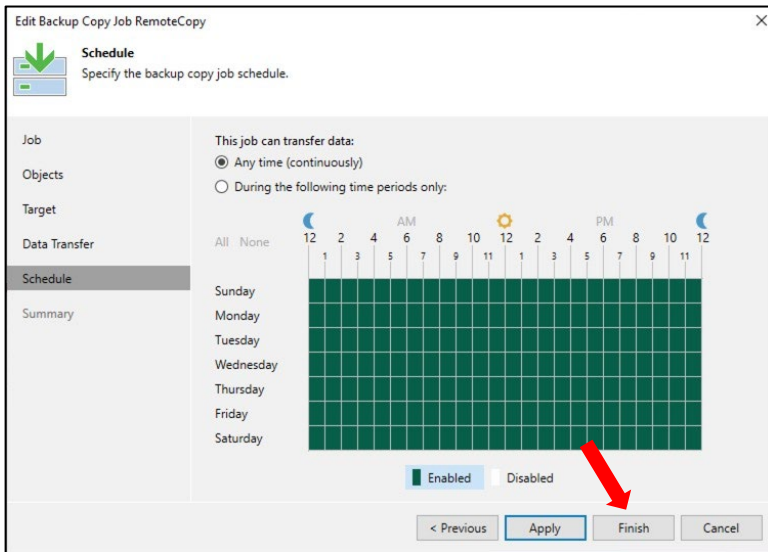
3. Select Add, then From jobs..., and then select the Backup Job used for the local site.



- In the Target section, select the SOBR that was created for the remote site. Define the retention and GFS policies as you did with the local site SOBR. Finally, check the box next to *Read the entire restore point from source backup instead of synthesizing it from increments*.



- In the Schedule section, define scheduling that makes sense for your use case.

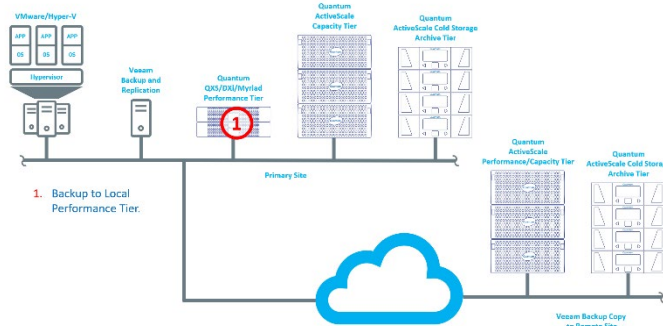


- Complete the wizard.

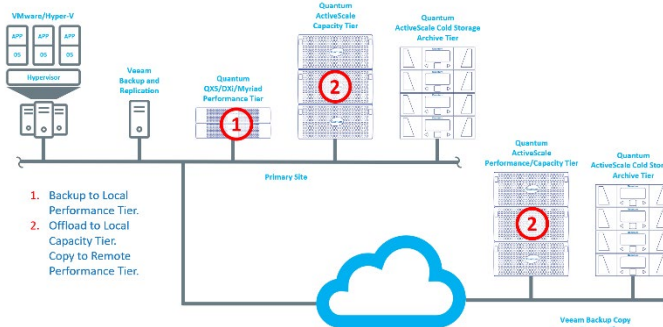
## Verify the Workflow

After setting up the Repositories, Backup Job, and Backup Copy Job, you can run the Backup Job and expect a workflow as described here:

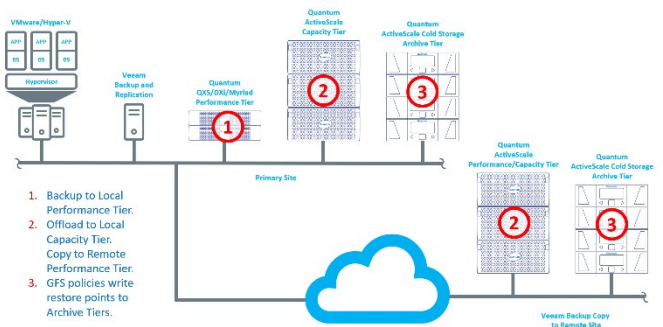
1. VBR creates a backup data set and writes it to the Performance Tier at the local site.



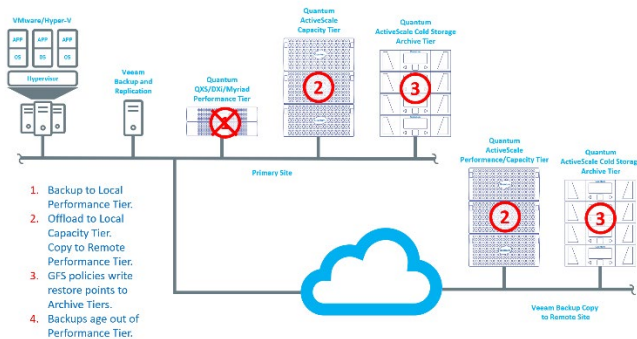
2. The Primary Site SOBR was created such that it will immediately make a copy of the backup data to the local Capacity Tier. At the same time, the Backup Copy Job was set to *Immediate Copy (mirroring)*, so the backup data in the local Performance Tier will get copied to the remote Performance Tier.



3. GFS policies will create restore points for eligible backups once weekly on Sunday. These restore points will be written to ActiveScale Cold Storage at both sites.



- Retention policies will age backups out of the Performance Tier after 14 days.



The above workflow is a result of how the example Repositories, Backup Job, and Backup Copy Job were created. Your retention and GFS policies should be adjusted for your use case.

## Unstructured Data

There is some confusion around how VBR works with unstructured data. Veeam defines unstructured data as: contents of SMB (CIFS) and NFS file shares, Windows and Linux-based file servers, NAS filers, and object storage systems. In other words, they are the individual files and/or objects that reside in a file or object storage system. Backing up unstructured data is different than backing up entire servers or virtual machines as was done in the examples above.

VBR has a feature called *Unstructured Data Backup*, described here:

### [Unstructured Data Backup](#)

This feature is designed around file-level backup and the ability to quickly restore files. Therefore, high-latency storage repositories, such as tape and the Archive Tier (ActiveScale’s Cold Storage Class), are not supported. However, the *Unstructured Data Backup* feature does support object storage when used as a Performance Tier or Capacity Tier.

There is another feature of VBR called *Unstructured Data Backup to Tape*, described here:

### [Unstructured Data Backup to Tape](#)

Despite similar naming, this is completely different than *Unstructured Data Backup*. This feature does allow the backup of files to tape but does not support the Archive Tier.

The reason for this is that VBR does not write individual files or objects to the Archive Tier. It writes GFS restore points. So, as of Version 12.1, unstructured data sources cannot take advantage of ActiveScale Cold Storage in a way that is supported by Veeam.

More information on working with unstructured data can be found here:

### [Working with Unstructured Data](#)



# Performance Tweaks

You can achieve a significant increase in both backup and restore performance by making a few changes from the default environment. In a nutshell, you want Backup Jobs to use VBRs largest block size and you want to spread VBR tasks over all available ActiveScale public interfaces. To take advantage of multiple ActiveScale interfaces, you can either create many different Backup Jobs or do some type of load balancing like DNS round robin.

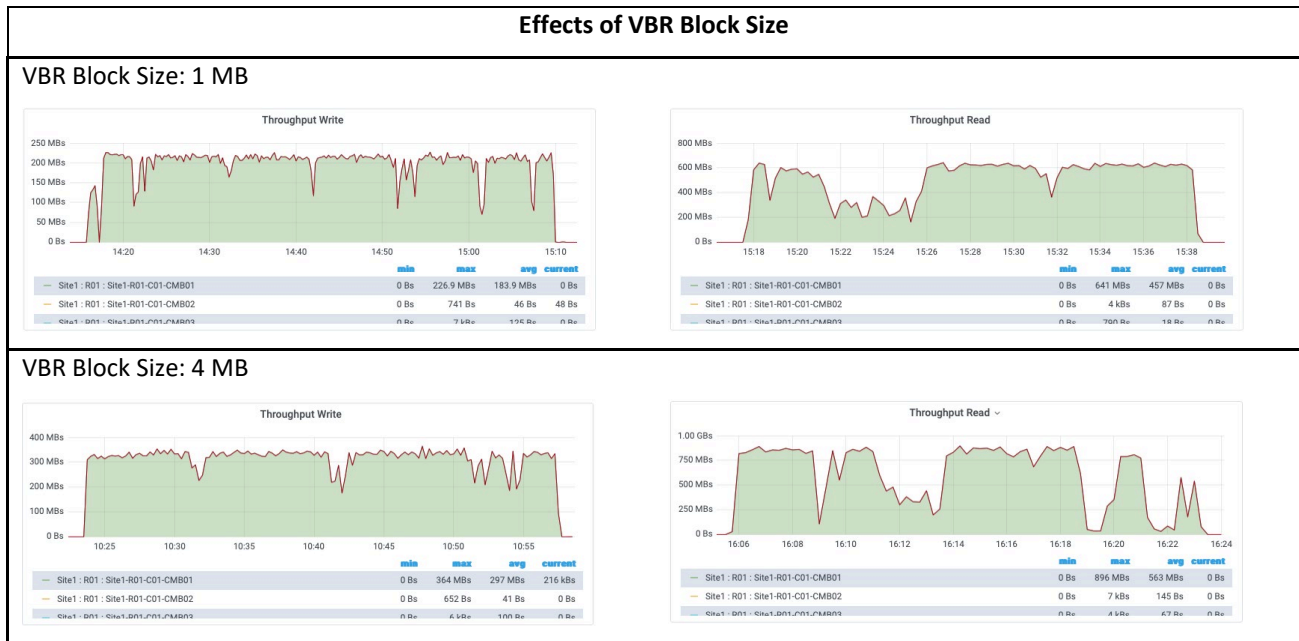
Note: The graphs in the below sections should not be considered performance benchmarks. Rather, the intent is to demonstrate the performance benefits when following the below recommendations.

## Increase VBR Block Size

Increasing the block size VBR uses is done during the Backup Job creation. The maximum block size VBR supports is 4 MB. This is explained in step 4 of the Create the Backup Job section of this document here: [Create the Backup Job](#)

The *Backup Copy Job* will use the same block size as the *Backup Job*.

The below graphs demonstrate the benefit of increasing the VBR block size.



## Load Balancing

ActiveScale does and always has supported strong consistency. Therefore, load balancing is a viable option to spread workload across ActiveScale’s public interfaces. If a production-worthy load balancer is available, it is the recommended way to maximize I/O bandwidth between VBR and ActiveScale. If a load balancer is not available, DNS round robin is an alternative.

## DNS Round Robin

VBR performs hostname resolution at the beginning of each task. So, when using DNS round robin, each task will resolve the DNS name for ActiveScale and use the address returned. That address will then be used for the entire task.

For example: A backup is a task. A Backup Job that backs up a single VM will resolve the hostname, get the address, and write the backup data to ActiveScale using that address and that address alone. A Backup Job that backs up 6 VMs, depending on how addresses are returned from the name resolution, may use up to 6 different addresses, resulting in a much better distribution across the ActiveScale interfaces.

Another consideration when using DNS round robin is Time to Live (TTL). This setting tells the local DNS cache how long a hostname address is valid. If the local cache has a valid IP for the hostname, it will not go back to the DNS server for name resolution, which has a negative impact on DNS round robin. Therefore, when using DNS round robin, it is recommended to disable DNS cache on the local VBR proxy.

The below graphs demonstrate task distribution across ActiveScale nodes with various DNS settings.





## Summary

By implementing the best practices outlined in this guide, you can harness the combined strengths of Quantum ActiveScale and Veeam Backup & Replication to create a robust and efficient data protection solution. This will ensure your critical data is always protected, minimize downtime in disaster scenarios, and provide peace of mind for your organization. Remember, these best practices are starting points, and you can adapt them to your specific needs and environment. Regularly review and update your backup strategy to ensure it continues to meet your evolving data protection requirements.

## 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	<a href="#">Link to Document</a>
ActiveScale OS Admin Guide	<a href="#">Link to Document</a>
Veeam User Guide for VMware vSphere	<a href="#">Link to Document</a>
Working with Unstructured Data	<a href="#">Link to Document</a>
Limitations for Scale-Out Backup Repositories	<a href="#">Link to Document</a>
Archive Tier Support	<a href="#">Link to Document</a>
Archive Tier Limitations	<a href="#">Link To Document</a>

Quantum delivers end-to-end data management solutions designed for the AI era. With over four decades of experience, our data platform has allowed customers to extract the maximum value from their unique, unstructured data. From high-performance ingest that powers AI applications and demanding data-intensive workloads, to massive, durable data lakes to fuel AI models, Quantum delivers the most comprehensive and cost-efficient solutions. Leading organizations in life sciences, government, media and entertainment, research, and industrial technology trust Quantum with their most valuable asset – their data. Quantum is listed on Nasdaq (QMCO). For more information visit [www.quantum.com](http://www.quantum.com).