

# RSTOR Space and Veeam

## Configuring Veeam Backup with RSTOR Space

November 2020 Prepared by: RSTOR Support



## **Table of Contents**

A	bstract	/Summary	. 3
1	Intr	oduction	. 4
	1.1	3 – 2 – 1 Rule	4
2	Vee	am Enterprise Features	. 5
	2.1	Scale Out Backup Repository (SOBR)	5
	2.2	Granular Recovery Functions	5
	2.3	Disaster Recovery	6
	2.4	Immutability and Legal Hold	6
3	Cor	figure Veeam to Use RSTOR Space	. 8



## **Abstract/Summary**

The purpose of this document is to provide a configuration guide and offer different use cases and methods for enhancing data protection with Veeam and RSTOR's Space S3 Object Platform. In this document the most common deployment scenarios and the benefits of each will be reviewed, finishing with the configuration steps.



#### **1** Introduction

While Veeam supports most S3 Object Storage platforms, most of the available options, such as AWS S3 or Azure Blob, actually create some unnecessary levels of frustration due to a lack of performance, and hidden, unpredictable, and variable costs. RSTOR Space addresses these common issues with a predictably priced, performant platform, built on 30+ patents, including patents for data transport and packaging, that set it apart as a leading S3 platform in an overall Veeam solution.

#### 1.1 3 – 2 – 1 Rule

The 3-2-1 rule is an industry standard practice which simply recommends that you have three copies of your backup data, on at least two different media types, with at least one copy off-site. Traditional legacy solutions tend to be backup to disk, possibly a replicated copy to more standard disk and then an archive off to tape. Usually the tape is then shipped off-site, maybe to a service provider like Iron Mountain and manually done by the team members to another datacenter or other location. While this method can satisfy the 3-2-1 rule, it typically includes unnecessary complexity, some inherent risk (what happens if there is a problem with the tapes), less control of your data, and higher costs.

More modern solutions are adopting object-based cloud storage for the archive tier, as opposed to tape. Object Based Storage (OBS) provides more flexibility, data control, resiliency, and can offer more economical options than traditional tape solutions. It also greatly reduces complexity, making a much simpler means of archiving and moving data as desired or, maybe, required. Most backup solutions now support some sort of direct integration support for S3 OBS, and Veeam has been an innovation leader in the backup space, with S3 cloud storage adoption being one of those features.

Using OBS helps satisfy the 3-2-1 rule by creating at least one copy of data, on different media (cloud OBS) and the copy is off-site. There are additional benefits of using OBS for archiving, including more resiliency with data replication options at the bucket level that can support DR and BC strategies, as well as other proactive functions, which will be reviewed later on in this document.

RSTOR Space works. There is simple, supported integration with Veeam, no complex deployment scenarios or unnecessarily complex configurations, with support for even some of the more complex features and functions Veeam has become known for. Adding RSTOR Space's unmatched performance and economics, it is a clear choice as the S3 platform in Veeam deployments.



#### 2 Veeam Enterprise Features

With RSTOR Space's innovative S3 OBS Platform, your data protection solution can be more complete than ever before. Offering exceptional performance, lower, predictable TCO, and including features such as native replication and object lock support, RSTOR Space is the perfect choice for any data storage needs. Combined with Veeam and the seamless integration, expanding the usability of your data sets has never been easier, and is more beneficial to your business.

While this article touches on many of the most common use cases for Veeam and RSTOR Space together, there are still many other ways this solution can resolve short comings of other solutions while providing a better user experience, lower TCO and more functionality. Reach out to your RSTOR Sales Engineer to find out more.

#### 2.1 Scale Out Backup Repository (SOBR)

One of Veeam's enterprise features and core components to their solution architecture, is their Scale Out Backup Repository (SOBR). SOBR is the ability to create logical repositories made up of multiple "extents", individual storage volumes combined to make up a single, logical repository for writing, storing and reading backup files from disks of varying types and locations. One of the types of Extents supported in a Veeam SOBR is a S3 bucket. This means you can have local, spinning disk volumes (block or file) combined with S3 OBS buckets to create a performant, yet cost effective, repository for your backup files with inherent tiering built in. With RSTOR Space's exceptional performance and flat, predictable cost, it is a clear choice to be used as extents in Veeam Scale Out Repositories.

#### 2.2 Granular Recovery Functions

Veeam is known for innovating many granular recovery functions, such as File Level Restore (FLR) without needing an agent installed on the guest. In most cases, it is not realistically feasible to perform these types of recovery functions from a backup file on a typical S3 OBS, due to performance limitations. While it may work eventually, the latency and time it takes to complete the task far outweighs the convenience. So, to overcome these limitations and irritations, Veeam will have to pull the data back to some form of primary disk before running the recovery task. Not only does this take time, but it also incurs expensive, unpredictable costs due to egress fees.

With RSTOR Space, these functions are seamless due to the unique performance capabilities, so recovery functions can be run directly from backup files stored on the RSTOR Space S3 platform. There is no need to take time and costly disk space to pull a



copy back before performing a desired restore function. This helps meet aggressive RTO and RPO's, keep over all costs down, and, if there was a need to pull a copy of a backup file for any other reason, it is accomplished simple and fast, with NO additional charges (egress, API requests, network transport). It is your data and you should have full control of it – anytime, anywhere. RSTOR Space keeps you in control of your data while lowering TCO.

#### 2.3 Disaster Recovery

Most of us have some sort of backup solution in place and have probably discussed, on many occasions, a Disaster Recovery (DR) strategy, but often there is no actual DR strategy or plan in place. Veeam has made DR very real and accessible to their customers, but it does add some layers of complexity, with additional associated costs to achieve. When you combine RSTOR Space with Veeam, DR becomes even easier, with less cost; RSTOR Space can provide native replication of the backup data sets to multiple locations, without needing to have the Veeam infrastructure in place to manage it. This saves costs up front, while still providing the building blocks (copies of the backup data files at different locations) to fully support a DR strategy and plan.

With RSTOR Space native replication, based on our patented network transport engine, large data sets can be replicated across the country, or even the globe, in a fraction of the time as other platforms. Because Veeam has a unique way of writing their backup data files, which includes a unique mechanism for maintaining the metadata, a new, default Veeam server could read those replicated backup files (assuming the proper security access has been met), create the catalogs and inventory and begin performing recovery tasks within minutes from any location – whether it be fully recovering servers, or simply restoring individual files... or any restore function in between.

When it comes to Disaster Recovery, performance and latency are the key factors that determine what RTO and RPOs are achievable, and with RSTOR Space and Veeam, you will be able to support more aggressive RTO and RPOs than ever before.

#### 2.4 Immutability and Legal Hold

More and more, the need for immutability and legal hold is becoming commonplace, and Veeam has introduced the function to do this when using S3 OBS as an archive tier for backups. However, not all S3 platforms support this important functionality. With RSTOR Space, immutability and legal hold are fully supported from within the Veeam application – allowing Veeam to fully manage retention policies and data pruning, on a S3 tier, as



required for compliance or any other reason. There are many reasons this feature support is essential to a complete solution, including Company policy, industry compliance regulations or even litigation, so if your solution cannot support these requirements, you could actually put your business in jeopardy.

Immutability can only be enforced if the S3 provider can support Object Locking. RSTOR Space provides full support in this use case by enabling S3 Bucket and object-locking. This needs to be enabled at the bucket level, allowing support for retention policies to be implemented at the bucket or individual object layer. Having a S3 Bucket with object locking enabled and full S3 API support allows the application to manage the policies to meet any immutability requirements,

Object Locking support consists of several variables, including the Mode (Governance or Compliance), which determines the severity of the lock, and retention variables to define the duration of the lock. There is also the need to support the associated S3 API requests to facilitate this functionality, which, again, RSTOR Space fully supports. With Veeam's ability to easily manage and automate these immutability requirements from within the management interfaces, on a supported platform, the ability to fully manage your data and adhere to compliance regulations is easier than ever.

## **RSTOR**

#### 3 Configure Veeam to Use RSTOR Space

This section describes step by step on how to configure Veeam to use RSTOR Space object storage platform.

- 1. First, we need to create a bucket to be used for the S3 repository. To do this, log in to the RSTOR web portal GUI and click on the create bucket button in the top right corner.
- 2. Next, we need to create a backup repository. In the Veeam GUI, under the "Backup Infrastructure" section, click on "Add Backup Repository" and select "Object Storage".



3. From the types of object storage list, select "S3 Compatible".

## **RSTOR**



4. Enter a name and description for the object storage repository, then click "Next".

New Object Storage Repository		×
Name Type in a name and c	lescription for this object storage repository.	
Name Account Bucket Summary	Name: rstor_nolock Description: Created by NBU01\Administrator at 8/11/2020 6:46 PM.	
	< Previous Next > Finish Cancel	



5. Enter the service point (this will be "s3.rstorcloud.io" for production accounts) and enter an arbitrary descriptor under region (RSTOR Space doesn't use regions, so this value will not be wrong).

New Object Storage Reposite	ory	×
Account Specify account	to use for connecting to S3 compatible storage system.	ł
Name	Service point:	
	s3.rstorcloud.io	
Account	Region:	
Bucket	west	
Summary	Credentials:	

6. For credentials, click "Add" to create and set the appropriate access keys for the object storage. Enter the access key and secret key to be used, and an optional description. Click "OK" once done.

\*Note: if you already have the proper object store credentials configured, simply select them from the drop-down menu available.

Credentia	als ×	
	Access key: Secret key:	
	Description: access keys for "noodlenate" account	
	OK Cancel	

7. Click "Next" after the credentials step has been completed.

\*Note: if the Veeam Backup Server does NOT have Internet access, this step may fail with a certificate validation error. In this case, please refer to <u>this Veeam KB article</u> for the workaround.

8. Select the bucket and the folder to be used as the backup repository, then click "Next".

## **RSTOR**



9. Review the summary presented, then click "Finish". In order to use your newly created object storage repository, it must be part of a Scale Out Repository, which will be described in the following steps.





- 10. To create a scale out repository, under the "Backup Infrastructure" section and "Scale Out Repositories", click to add a "New Scale-out Backup Repository".
- 11. Enter a name and description (if desired) and click "Next".

New Scale-out Backup Reposito	ry	×
Name		
Type in a name and	description for this scale-out backup repository.	
Name	Name:	
Performance Tier	Description	
Placement Policy	Created by NBU01\Administrator at 8/12/2020 7:15 PM.	
Capacity Tier		
Summary		
	< Previous Next > Finish C	Cancel

12. Under "Performance Tier", add the extent and click "Next".

Select backup	Tier repositories to use as the landing zone and for the short-term retention.	
Name	Extents:	
Porformanco Tior	Name	Add
	Repo1	Remove
Placement Policy		
Capacity Tier		
Summary		
	Click Advanced to specify additional scale-out backup repository options	Advance

13. Under "Placement Policy" select "Data Locality", then click "Next".

1	
New Scale-out Backup Repositor	y X
Placement Policy Choose a backup file backup job will chose	s placement policy for this performance tier. When more than one extent matches the placement policy, extent with the most free disk space available.
Name Performance Tier Placement Policy Capacity Tier Summary	<ul> <li><b>O bata locality</b>         All dependent backup files are placed on the same extent. For example, incremental backup files will be stored together with the corresponding full backup file. However, the next full backup file can be created on another extent (except extents backed by a deduplicating storage).     </li> <li><b>O Performance</b>         Incremental backup files are placed on a different extent from the corresponding full backup file, providing for better backup file transformation performance with raw storage devices. Note that losing an extent with a full backup makes restoring from increments impossible.     </li> <li>Specify the placement policy for full and incremental backup files.</li> </ul>
	< Previous Next > Finish Cancel

14. Under "Capacity Tier", select "Extend scale-out backup repository capacity with object storage", and select the RSTOR Space backup repository from the drop-down menu. Configure other variables as desired. When finished, click "Apply".

<sup>I</sup> New Scale-out Backup Reposito	ry	×
Capacity Tier Specify object stora completely to reduc	ge to copy backups to for redundancy and DR purposes. Older backups can be moved to ob e long-term retention costs while preserving the ability to restore directly from offloaded b	nject storage ackups.
Name	Extend scale-out backup repository capacity with object storage:	
Performance Tier	rstor_nolock ~	Add
Placement Policy	Define time windows when uploading to object storage is allowed	Window
Capacity Tier Summary	<ul> <li>Copy backups to object storage as soon as they are created</li> <li>Create additional copy of your backups for added redundancy by having all backu the capacity tier as soon as they are created on the performance tier.</li> <li>Move backups to object storage as they age out of the operational restore window Reduce your long-term retention costs by moving older backups to object storage while preserving the ability to restore directly from offloaded backups.</li> </ul>	ps copied to
	Move backup files older than b 🔄 days (your operational restore window)	Override
	Password:	Add
	Manage passwords	
	< Previous Apply Finish	Cancel



New Scale-out Backup Reposito	bry	×
Summary Review the scale-ou	ut backup repository settings, and click Finish to exit the wizard.	
Name Performance Tier Placement Policy Capacity Tier Summary	Summary: Scale- out backup repository was created successfully.	
	< Previous Next > Finish	Cancel

15. Click finish, and it is now all set up.