

# Explore Veeam 12's Unbreakable Backups

Roger Lund – Solutions Architect, Network Center Inc.  
[roger.lund@netcenter.net](mailto:roger.lund@netcenter.net)



# Table of Contents

- The Challenge : Ransomware is Increasing
- Discussion Point – Immutability?
- The 3-2-1 Rule
- Discussion Point- Why Wasabi?
- Discussion Point - Why Scale?
- Solution Review : Triple Play
- Discussion Point – Hardware / Lap Setup
- Setup – Linux Hardened Repository
- Setup – Scale Out Repository
- Demos
- Links



# The Challenge : Ransomware is Increasing



- According to the 2023 Data Protection Trends Report, 85%\* of organizations suffered at least one cyber attack in the preceding twelve months; an increase from 76% experienced in the prior year 56% of respondents report that their organization has suffered from a ransomware attack in the last 12 months Source Veeam Data Protection Trends Report 2023
- 60% of organizations need significant or complete overhauls between their backup and cyber teams Source Veeam Data Protection Trends Report 2023
- 87% of organizations have a risk management program that drives their security roadmap or strategy. That said, only 35% believe their program is working well, while 52% are seeking to improve their situation and the remaining 13% do not yet even have an established program. Source Veeam Data Protection Trends Report 2023
- Comparitech stated In 2020, 92 individual ransomware attacks affected over 600 separate clinics, hospitals, and organizations and more than 18 million patient records. We estimate the cost of these attacks to be almost \$21 billion. Source Comparitech.com



# Discussion Point – Immutability?



- What are Why Hardened , Immutable Repositories?
- Let's define what the Immutability Definition is:
  - Immutability ensures that once data is written, it cannot be modified or deleted for a specified period.
- Ransomware Protection:
  - In a world where ransomware attacks are on the rise, a hardened repository acts as a fortified defense, ensuring backup data remains uncompromised.
- Compliance:
  - Many industries mandate strict data retention policies. Immutability ensures that backup data cannot be tampered with, meeting stringent compliance requirements.
- Data Integrity:
  - Guarantees the authenticity of your data. You can trust that your backups are exactly as they were when initially saved.
- Operational Assurance:
  - In case of accidental deletions or modifications, immutable backups serve as a reliable point of recovery, ensuring business continuity.



# The 3-2-1 Rule



## Backup Strategy Rule to Live By

3



Different copies  
of data

2



Different media

1



Offsite copy

VEEAM

1



Of which is:  
offline  
air-gapped  
or immutable

0



No errors after  
backup recoverability  
verification

# The 3-2-1 Rule



- What are Why Hardened Immutable Repositories,?
- Let's define what the Immutability Definition is:
  - Immutability ensures that once data is written, it cannot be modified or deleted for a specified period.
- Ransomware Protection:
  - In a world where ransomware attacks are on the rise, a hardened repository acts as a fortified defense, ensuring backup data remains uncompromised.
- Compliance:
  - Many industries mandate strict data retention policies. Immutability ensures that backup data cannot be tampered with, meeting stringent compliance requirements.
- Data Integrity:
  - Guarantees the authenticity of your data. You can trust that your backups are exactly as they were when initially saved.
- Operational Assurance:
  - In case of accidental deletions or modifications, immutable backups serve as a reliable point of recovery, ensuring business continuity.





# Ultra-resilient backups



Scale-out Backup Repository™



## Performance Tier



DAS




NAS



Dedupe appliance



Hardened Repository

Immutability   
from start to finish

- Policy-based
- Transparent
- Space efficient
- Self-sufficient
- No extra costs

## Capacity Tier



S3-Compatible



IBM Object Storage



Google Cloud Platform



Amazon S3



Microsoft Azure Blob Storage



Amazon S3 Glacier



Microsoft Azure Archive Blob Storage



# Double-Play, Triple Play..

1

## Performance Tier



DAS



NAS



Dedupe appliance



Hardened Repository



Scale-out Backup Repository™

2

## Capacity Tier



S3-Compatible



IBM Object Storage



Google Cloud Platform



Amazon S3



Microsoft Azure Blob Storage

Jobs to tape



Immutability from start to finish

- Policy-based
- Transparent
- Space efficient
- Self-sufficient
- No extra costs

## Archive Tier



Amazon S3 Glacier



Microsoft Azure Archive Blob Storage



# Discussion Point- Why Wasabi? ○○○

## Wasabi and Veeam for the Most Resilient Data Protection Solution

### Key Features:

- Fast, reliable, and cost-effective backup & recovery
- Store more backups longer for increased cyber resilience
- Leverage low-cost cloud object storage with Veeam Scale Out Backup Repository
- Simplify the backup redundancies with direct-to-cloud primary and secondary copies
- Enhance ransomware protection with industry-leading immutability
- Seamless integration and operation
- Certified Veeam interoperability Veeam Ready Object and Veeam Ready Object with Object Lock

### Benefits:

- Lower overall IT operations and data recovery costs
- Unbreakable data protection from ransomware attacks
- No added costs to extend secure locks (API) or restore backups (egress)
- Improved risk mitigation with proven disaster recovery processes
- Increased confidence in the authenticity and availability of data backups

Source <https://wasabi.com/resource/veeam/>



# Discussion Point - Why Scale? ○○○

## Veeam Backup Architecture

- One or more source hosts
- One or more backup proxies
- Backup repository

## Veeam Scale out repository

- The main benefits of this feature are as follows:
  - It provides a convenient way of managing the backup storage.
  - The scale-out backup repository can be expanded at any moment: if the performance extents of your scale-out backup repository run out of space, you can add a new performance extent to the existing scale-out backup repository. It supports any backup target supported by Veeam: Windows or Linux servers with local or DAS storage, network shares, deduplicating storage appliances. All the features of any storage device or system are preserved.
  - It allows you to set up granular performance policy.
  - It provides practically unlimited cloud-based storage capacity: you can instruct Veeam Backup & Replication to offload data from extents to the cloud for long-term storage.

Source: [https://helpcenter.veeam.com/docs/backup/vsphere/backup\\_architecture.html?ver=110](https://helpcenter.veeam.com/docs/backup/vsphere/backup_architecture.html?ver=110)



# Solution Review : Triple Play



## Review:

- Veeam triple immutable solution
  - One performance repository
  - One capacity repository
  - One archive repository
  
- Step by Step
  - Setup each of the of the above
    - Linux Hardened Repository
    - Veeam Scale out repository



# Discussion Point : Hardware



- **Hardware Components**

- Server
  - What, works. What Doesn't
- NAS
  - Does it have to be a NAS?
- Networking
  - Speeds / Feeds that make sense.
- Cloud Storage
  - What makes sense?



# Setup – Lab Hardware



- **Hardware Components**

My Lab

- Server
  - Virtual
- NAS
  - IX Systems
    - Snapshots
- Networking
  - Speeds / Feeds
- Cloud Storage
  - S3
    - AWS S3
    - Wasabi
  - Archive tier
    - AWS Glacier





# Before we Start



Before we start:

- Best Practices for Hardening Veeam Backup Repositories based on Linux are:
  - K.I.S.S. design - Keep It Simple and Straightforward.
  - Make sure the servers are physical secured.
  - Create a dedicated repository account for Veeam, that can access the folder where you store backups.
  - Set permissions on the repository directory to only that account.
  - You do not need Root to use a Veeam Linux Repository. Also do not use SUDO.
  - Modify the Firewall, with dedicated rules for Veeam to allow access to specific ports.
  - Use Veeam encryption while storing backups on the repository.
- Source [https://bp.veeam.com/vbr/Security/hardening\\_backup\\_repository\\_linux.html](https://bp.veeam.com/vbr/Security/hardening_backup_repository_linux.html)



# Setup – Linux Hardened Repository



New Linux Server

**Name**  
Specify DNS name or IP address of Linux server. The server must have SSH and Perl installed.

**Name**  
DNS name or IP address:  
DNSNameHere

**Description:**  
Created by VEEAM01\Administrator at 10/8/2021 4:56 PM.

< Previous   **Next >**   Finish   Cancel

New Linux Server

**SSH Connection**  
Provide credentials for service console connection, and adjust secure shell (SSH) port number using advanced settings if required.

**Name**  
SSH Connection

**Credentials:**  
user001 (Single User Hardened Account, last edited: less than a day ago) [Add...](#) [Manage accounts](#)

Review  
Apply  
Summary

Customize advanced connection settings, such as SSH and data mover ports [Advanced...](#)

< Previous   Next >   Finish   Cancel

- Linux account...
- Linux private key...
- Single-use credentials for hardened repository...





Credentials

Username:

Password:

SSH port: 22

Non-root account

Elevate account privileges automatically

Add account to the sudoers file

Use "su" if "sudo" fails

Root password:

Description:

OK Cancel

Credentials

Username: locveeam

Password: ●●●●●●

SSH port: 22

Non-root account

Elevate account privileges automatically

Add account to the sudoers file

Use "su" if "sudo" fails

Root password:

Description:


one time hardened login

OK Cancel

Veeam Backup and Replication

192.168.168.15 SSH key fingerprint: ssh-rsa 3072 82:c2:1c:10:23:67:4c:f4:5c:6e:99:f3:11:af:d4:92  
Do you trust this server?

Yes No





### New Linux Server

**Review**  
Please review your settings and click Apply to continue.

Name

SSH Connection

**Review**

Apply

Summary

Due to these modifications the following components will be installed or removed on the target host:

Component name	Status
Transport	will be installed

After you click Apply missed components will be installed on the target host.

< Previous   **Apply**   Finish   Cancel

### New Linux Server

**Apply**  
Please wait while required operations are being performed, this may take a few minutes.

Name

SSH Connection

Review

**Apply**

Summary


Message	Duration
Starting infrastructure item update process	0:00:03
Checking if Veeam Data Mover service is supported by the Linux server	
Discovering existing Veeam Data Mover service	
Installing Veeam Data Mover service	0:00:13
Discovering existing Veeam Data Mover service	
Setting server certificate	0:00:01
Resolving server certificate thumbprint	
Setting client certificate	0:00:01
Configuring Veeam Data Mover service	
Restarting Veeam Data Mover service	
Testing Veeam Data Mover service connection	0:00:05
Collecting hardware info	0:00:03
Creating database records for server	0:00:09
Linux server saved successfully	

< Previous   **Next >**   Finish   Cancel





New Linux Server ×

 **Summary**  
You can copy the configuration information below for future reference.

Name	<b>Summary:</b> Linux host '192.168.168.15' was successfully created. <b>SSH options:</b> Credentials type: Single-use User: locveeam Elevate account privileges automatically: yes Add account to the sudoers file: no Use "su" if "sudo" fails: no <b>Hardware info:</b> Chassis type: Virtual (VMware) Cores count: 1 <b>Components:</b> Transport using port 6162
SSH Connection	
Review	
Apply	
<b>Summary</b>	

< Previous Next > Finish Cancel







## Add Backup Repository

Select the type of backup repository you want to add.



### Direct attached storage

Microsoft Windows or Linux server with internal or direct attached storage. This configuration enables data moves to run directly on the server, allowing for fastest performance.



### Network attached storage

Network share on a file server or a NAS device. When backing up to a remote share, we recommend that you select a gateway server located in the same site with the share.



### Deduplicating storage appliance

Dell EMC Data Domain, ExaGrid, HPE StoreOnce or Quantum DXi. If you are unable to meet the requirements of advanced integration via native appliance API, use the network attached storage option instead.



### Object storage

On-prem object storage system or a cloud object storage provider. Object storage can only be used as a Capacity Tier of scale-out backup repositories, backing up directly to object storage is not currently supported.

Cancel



## Direct Attached Storage

Select the operating system type of a server you want to use as a backup repository.



### Microsoft Windows

Adds local server storage presented as a regular volume or Storage Spaces. For better performance and storage efficiency, we recommend using ReFS.



### Linux

Adds local server storage, or locally mounted NFS share. The Linux server must use bash shell, and have SSH and Perl installed.

Cancel



New Backup Repository

**Name**  
Type in a name and description for this backup repository.

Name: Hardened Repository 1  
Description: Created by VEEAM01\Administrator at 10/8/2021 5:10 PM.

< Previous Next > Finish Cancel

New Backup Repository

**Server**  
Choose repository server. You can select server from the list of managed servers added to the console.

Repository server: 192.168.168.15 (Created by VEEAM01\Administrator at 10/8/2021 5:06 PM.) Add New...

Path	Capacity	Free	Populate
------	----------	------	----------

< Previous Next > Finish Cancel





### New Backup Repository

**Server**  
Choose repository server. You can select server from the list of managed servers added to the console.

Name	Repository server:
Server	192.168.168.15 (Created by VEEAM01\Administrator at 10/8/2021 5:06 PM.)
Repository	
Mount Server	
Review	
Apply	
Summary	

Path	Capacity	Free
/ (/dev/mapper/ubuntu--vg-ubuntu--lv)	48.5 GB	39.3 GB
/boot (/dev/sda2)	975.9 MB	802.7 MB
/dev (udev)	949.8 MB	949.8 MB
/dev/shm (tmpfs)	993.9 MB	993.9 MB
/mnt/veeamrepo (/dev/sdb)	1023.5 GB	1016.3 GB
/run (tmpfs)	198.8 MB	197.6 MB
/run/lock (tmpfs)	5 MB	5 MB
/run/user/1000 (tmpfs)	198.8 MB	198.8 MB
/run/user/1001 (tmpfs)	198.8 MB	198.8 MB
/snap/core18/1944 (/dev/loop0)	55.5 MB	0 B
/snap/core18/2128 (/dev/loop1)	55.5 MB	0 B
/snap/core20/1081 (/dev/loop2)	61.9 MB	0 B
/snap/core20/1169 (/dev/loop4)	61.9 MB	0 B
/snap/lxd/19188 (/dev/loop6)	70 MB	0 B
/snap/lxd/21545 (/dev/loop3)	67.3 MB	0 B

< Previous   **Next >**   Finish   Cancel

### New Backup Repository

**Repository**  
Type in path to the folder where backup files should be stored, and set repository load control options.

Name	Location
Server	Path to folder: /mnt/veeamrepo/backups
Repository	Capacity: <Unknown> Free space: <Unknown>
Mount Server	
Review	
Apply	
Summary	

Use fast cloning on XFS volumes (recommended)  
Reduces storage consumption and improves synthetic backup performance.

Make recent backups immutable for: 7 days  
Protects backups from modification or deletion by ransomware or hackers. GFS full backups are made immutable for the entire duration of their retention policy.

**Load control**  
Running too many concurrent tasks against the repository may reduce overall performance, and cause I/O timeouts. Control storage device saturation with the following settings:

Limit maximum concurrent tasks to: 4

Limit read and write data rate to: 1 MB/s

Click Advanced to customize repository settings. **Advanced**

< Previous   **Next >**   Finish   Cancel





### New Backup Repository

**Repository**  
Type in path to the folder where backup files should be stored, and set repository load control options.

Name	Location
Server	Path to folder: <input type="text" value="/mnt/veeamrepo/backups"/> <input <="" td="" type="button" value="Browse..."/>
<b>Repository</b>	<input type="button" value="Populate"/> Capacity: <b>1023.5 GB</b> Free space: <b>1016.3 GB</b>
Mount Server	<input checked="" type="checkbox"/> Use fast cloning on XFS volumes (recommended) Reduces storage consumption and improves synthetic backup performance.
Review	<input checked="" type="checkbox"/> Make recent backups immutable for: <input type="text" value="30"/> days Protects backups from modification or deletion by ransomware or hackers. GFS full backups are made immutable for the entire duration of their retention policy.
Apply	<b>Load control</b> Running too many concurrent tasks against the repository may reduce overall performance, and cause I/O timeouts. Control storage device saturation with the following settings:
Summary	<input checked="" type="checkbox"/> Limit maximum concurrent tasks to: <input type="text" value="4"/> <input type="checkbox"/> Limit read and write data rate to: <input type="text" value="1"/> MB/s

Click Advanced to customize repository settings.

< Previous **Next >** Finish Cancel

### New Backup Repository

**Mount Server**  
Specify a server to mount backups to when performing advanced restores (file, application item and instant VM recoveries). Instant recoveries require a write cache folder to store changed disk blocks in.

Name	Mount server:
Server	<input type="text" value="veeam01.lab02.vroger.com (Backup server)"/> <input <="" td="" type="button" value="Add New..."/>
Repository	Instant recovery write cache folder: <input type="text" value="C:\ProgramData\Veeam\Backup\IRCache\"/> <input <="" td="" type="button" value="Browse..."/>
<b>Mount Server</b>	Ensure that the selected volume has sufficient free disk space to store changed disk blocks of instantly recovered VMs. We recommend placing write cache on an SSD drive.
Review	<input checked="" type="checkbox"/> Enable vPower NFS service on the mount server (recommended) <input <="" td="" type="button" value="Ports..."/>
Apply	Unlocks instant recovery of any backup (physical, virtual or cloud) to a VMware vSphere VM. vPower NFS service is not used for instant recovery to a Microsoft Hyper-V VM.
Summary	

< Previous **Next >** Finish Cancel







### New Backup Repository



#### Review

Please review the settings, and click Apply to continue.

Name	The following components will be processed on server veeam01.lab02.vroger.com:	
Server	Component name	Status
Repository	Transport	already exists
Mount Server	vPower NFS	already exists
Review	Mount Server	already exists
Apply		
Summary		

Search the repository for existing backups and import them automatically  
 Import guest file system index data to the catalog

< Previous   **Apply**   Finish   Cancel

### New Backup Repository



#### Apply

Please wait while backup repository is created and saved in configuration, this may take a few minutes.

Name	Message	Duration
Server	Starting infrastructure item update process	0:00:03
Repository	[veeam01] Discovering installed packages	
Mount Server	[veeam01] Registering client VEEAM01 for package Transport	
Review	[veeam01] Registering client VEEAM01 for package vPower NFS	
Apply	[veeam01] Registering client VEEAM01 for package Mount Server	
Summary	[veeam01] Discovering installed packages	
	All required packages have been successfully installed	
	Detecting server configuration	
	Reconfiguring vPower NFS service	
	Creating configuration database records for installed packages	
	Collecting backup repository info	
	Creating database records for repository	0:00:05
	Backup repository has been added successfully	


< Previous   **Next >**   Finish   Cancel







New Backup Repository ✕

 **Summary**  
You can copy the configuration information below for future reference.

Name	Summary:
Server	Linux backup repository 'Hardened Repository 1' was successfully created.
Repository	Mount host: veeam01.lab02.vroger.com
Mount Server	Backup folder: /mnt/veeamrepo/backups
Review	Write throughput: unlimited
Apply	Max parallel tasks: 4
<b>Summary</b>	Fast cloning on XFS volumes: enabled
	Days of immutability: enabled for 30 days


< Previous Next > Finish Cancel



# Setup –Scale out Repository



New Scale-out Backup Repository ✕

 **Name**  
Type in a name and description for this scale-out backup repository.

**Name**

Performance Tier

Placement Policy

Capacity Tier

Summary

Name:  
Scale-out Backup Repository Hardened

Description:  
Created by VEEAM01\Administrator at 10/9/2023 3:55 PM.

< Previous **Next >** Finish Cancel







Extents ×

Select backup repositories to include in this scale-out backup repository.


Backup repositories:

Name	Type
<input type="checkbox"/> Backup Repository 1	Windows
<input checked="" type="checkbox"/> Hardened Repo	Hardened
<input type="checkbox"/> Object storage repositi...	S3-compatible





New Scale-out Backup Repository ×



**Placement Policy**

Choose a backup files placement policy for this performance tier. When more than one extent matches the placement policy, backup job will choose the extent with the most free disk space available.

Name	
Performance Tier	
Placement Policy	
Capacity Tier	
Summary	

**Data locality**

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).

**Performance**

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.

Specify the placement policy for full and incremental backup files. Customize...

**Strict placement policy enforcement**

By default, whenever we are unable to place a new backup according to the policy, we will violate it to ensure a backup could still be created. Select this option if you'd rather have backup jobs fail in such circumstances instead.


< Previous Next > Finish Cancel







New Scale-out Backup Repository ×

 **Capacity Tier**  
Specify object storage to copy backups to for redundancy and DR purposes. Older backups can be moved to object storage completely to reduce long-term retention costs while preserving the ability to restore directly from offloaded backups.

Name	<input checked="" type="checkbox"/> Extend scale-out backup repository capacity with object storage: <input data-bbox="996 601 1753 639" type="text" value=" &lt;Click Choose to pick object storage&gt; "/> <input data-bbox="1765 601 1895 639" type="button" value="Choose..."/>
Performance Tier	
Placement Policy	Define time windows when uploading to capacity tier is allowed <input data-bbox="1765 665 1895 701" type="button" value="Window..."/>
<b>Capacity Tier</b>	<input type="checkbox"/> Copy backups to object storage as soon as they are created Create additional copy of your backups for added redundancy by having all backups copied to the capacity tier as soon as they are created on the performance tier.
Summary	<input checked="" type="checkbox"/> 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 completely while preserving the ability to restore directly from offloaded backups. Move backup files older than <input data-bbox="1294 911 1370 946" type="text" value="14"/> days (your operational restore window) <input data-bbox="1765 911 1895 946" type="button" value="Override..."/>
	<input type="checkbox"/> Encrypt data uploaded to object storage Password: <input data-bbox="1026 1096 1760 1132" type="text"/> <input data-bbox="1765 1096 1895 1132" type="button" value="Add..."/> <input data-bbox="1574 1143 1753 1168" type="button" value="Manage passwords"/>






to copy backups to for redundancy and on purpose. Other backups can

### Capacity Tier Extents

Consider creating multiple buckets if your object storage can only support a limited number of objects per bucket. Backups will be spread among all available extents using a round-robin algorithm on a machine level.

Object storage:

Name	Type
<input type="checkbox"/>  Object storage reposi...	S3-compatible





## Add Object Storage Repository

Select the type of object storage you want to add.



### S3 Compatible

Adds an on-premises object storage system or a cloud object storage provider.



### Amazon S3

Adds Amazon cloud object storage. Amazon S3, Amazon S3 Glacier (including Deep Archive) and Amazon Snowball Edge are supported.



### Google Cloud Storage

Adds Google Cloud storage. Both Standard and Nearline storage classes are supported.



### IBM Cloud Object Storage

Adds IBM Cloud object storage. S3 compatible versions of both on-premises and IBM Cloud storage offerings are supported.



### Microsoft Azure Storage

Adds Microsoft Azure cloud object storage. Microsoft Azure Blob Storage, Microsoft Azure Archive Storage and Microsoft Azure Data Box are supported.



### Wasabi Cloud Storage

Adds Wasabi cloud object storage.

Cancel





## New Object Storage Repository



### Name

Type in a name and description for this object storage repository.

Name	Name: <input type="text" value="Wasabi"/>
Account	Description: <input type="text" value="Created by VEEAM01\Administrator at 10/9/2023 4:34 PM."/>
Bucket	
Mount Server	
Review	
Apply	
Summary	

Limit concurrent tasks to:

Use this setting to limit the maximum number of tasks that can be processed concurrently in cases when your object storage is overloaded or cannot keep up with the number of API requests issued by multiple object storage offload tasks.

< Previous   **Next >**   Finish   Cancel





## New Object Storage Repository



### Account

Specify an account to use for connecting to Wasabi Cloud Object Storage.


Name	Region: <input type="text" value="us-central"/>
Account	Credentials: <input type="text" value="PQC�3XSƒJƦAN6BHE4I5A (last edited: less than a day ago)"/> <input type="button" value="Add..."/> <a href="#">Manage cloud accounts</a>
Bucket	
Mount Server	
Review	
Apply	
Summary	
	Connection mode: <input type="text" value="Direct"/> <input type="button" value="Choose..."/> Specify how object storage should be accessed and configure repository access control settings for backup agents.

< Previous   Next >   Finish   Cancel





New Object Storage Repository ✕

 **Bucket**  
Specify object storage system bucket to use.

Name	Bucket:	<input type="text" value="demo-bucker-roger"/>	<input type="button" value="Browse..."/>
Account	Folder:	<input type="text" value="veeam"/>	<input type="button" value="Browse..."/>
<b>Bucket</b>	<input checked="" type="checkbox"/> Limit object storage consumption to:	<input type="text" value="1"/> <input type="button" value="TB"/>	
Mount Server	<p>This is a soft limit to help control your object storage spend. If the specified limit is exceeded, already running backup offload tasks will be allowed to complete, but no new tasks will be started.</p>		
Review	<input checked="" type="checkbox"/> Make recent backups immutable for:	<input type="text" value="30"/> days	
Apply	<p>Protects backups from modification or deletion by ransomware, malicious insiders and hackers. GFS backups are made immutable for the entire duration of their retention policy.</p>		
Summary			





New Object Storage Repository ✕



### Mount Server

Specify a server to mount backups to when performing advanced restores (file, application item and instant VM recoveries). Instant recoveries require a write cache folder to store changed disk blocks in.

Name	Mount server:
Account	veeam01 (Backup server) <span style="float: right;">Add New...</span>
Bucket	Instant recovery write cache folder: E:\ProgramData\Veeam\Backup\IRCache\ <span style="float: right;">Browse...</span>
<b>Mount Server</b>	Ensure that the selected volume has sufficient free disk space to store changed disk blocks of instantly recovered machines. We recommend placing the write cache folder on an SSD drive.
Review	<input checked="" type="checkbox"/> Enable vPower NFS service on the mount server (recommended) <span style="float: right;">Ports...</span>
Apply	Unlocks instant recovery of any backup (physical, virtual or cloud) to a VMware vSphere VM. vPower NFS service is not used for instant recovery to a Microsoft Hyper-V VM.
Summary	

✔ Helper appliance has been configured successfully. Configure...

< Previous Next > Finish Cancel







New Object Storage Repository ✕



### Mount Server

Specify a server to mount backups to when performing advanced restores (file, application item and instant VM recoveries). Instant recoveries require a write cache folder to store changed disk blocks in.

Name	Mount server:
Account	veeam01 (Backup server) <span style="float: right;">Add New...</span>
Bucket	Instant recovery write cache folder: E:\ProgramData\Veeam\Backup\IRCache\ <span style="float: right;">Browse...</span>
<b>Mount Server</b>	Ensure that the selected volume has sufficient free disk space to store changed disk blocks of instantly recovered machines. We recommend placing the write cache folder on an SSD drive.
Review	<input checked="" type="checkbox"/> Enable vPower NFS service on the mount server (recommended) <span style="float: right;">Ports...</span>
Apply	Unlocks instant recovery of any backup (physical, virtual or cloud) to a VMware vSphere VM. vPower NFS service is not used for instant recovery to a Microsoft Hyper-V VM.
Summary	

✔ Helper appliance has been configured successfully. Configure...

< Previous Next > Finish Cancel





Capacity Tier Extents

Consider creating multiple buckets if your object storage can only support a limited number of objects per bucket. Backups will be spread among all available extents using a round-robin algorithm on a machine level.

Object storage:


Name	Type
<input type="checkbox"/> Object storage reposi...	S3-compatible
<input checked="" type="checkbox"/> Wasabi	Wasabi

Buttons: Add..., Select All, Clear All, OK, Cancel





New Scale-out Backup Repository ✕

 **Capacity Tier**  
Specify object storage to copy backups to for redundancy and DR purposes. Older backups can be moved to object storage completely to reduce long-term retention costs while preserving the ability to restore directly from offloaded backups.

Name

Performance Tier

Placement Policy

**Capacity Tier**

Summary

Extend scale-out backup repository capacity with object storage:

Wasabi Choose...

Define time windows when uploading to capacity tier is allowed Window...

Copy backups to object storage as soon as they are created  
Create additional copy of your backups for added redundancy by having all backups copied to the capacity tier as soon as they are created on the performance tier.

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 completely while preserving the ability to restore directly from offloaded backups.

Move backup files older than 14 days (your operational restore window) Override...

Encrypt data uploaded to object storage  
Password:  
Created by VEEAM01\Administrator at 4/11/2020 7:53 PM. (Last edited: 1275 Add...)


[Manage passwords](#)

< Previous Apply Finish Cancel





New Scale-out Backup Repository ×

 **Capacity Tier**  
Specify object storage to copy backups to for redundancy and DR purposes. Older backups can be moved to object storage completely to reduce long-term retention costs while preserving the ability to restore directly from offloaded backups.

**Name**

**Performance Tier**

**Placement Policy**

**Capacity Tier**

**Summary**

Extend scale-out backup repository capacity with object storage:

Wasabi Choose...

Define time windows when uploading to capacity tier is allowed Window...

Copy backups to object storage as soon as they are created  
Create additional copy of your backups for added redundancy by having all backups copied to the capacity tier as soon as they are created on the performance tier.

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 completely while preserving the ability to restore directly from offloaded backups.

Move backup files older than 14 days (your operational restore window) Override...

Encrypt data uploaded to object storage  
Password:  
Created by VEEAM01\Administrator at 4/11/2020 7:53 PM. (Last edited: 1275 Add...)


[Manage passwords](#)

< Previous Apply Finish Cancel





New Scale-out Backup Repository ✕

 **Summary**  
Review the scale-out backup repository settings, and click Finish to exit the wizard.

Name	<b>Summary:</b> Scale-out backup repository was created successfully.
Performance Tier	
Placement Policy	
Capacity Tier	
<b>Summary</b>	

< Previous   Next >   **Finish**   Cancel



# Demo- Let's Make a Backup Job



Demo- Let's Recover a document





# Links



- Veeam backup and replication 12 what's new  
[https://www.veeam.com/veeam\\_backup\\_12\\_0\\_whats\\_new\\_wn.pdf](https://www.veeam.com/veeam_backup_12_0_whats_new_wn.pdf)
- Ransomware Prevention Kit
- <https://go.veeam.com/ransomware-prevention-kit>
- Protect against Ransomware with Immutable Backups: a Veeam Guide
- <https://www.veeam.com/wp-guide-protect-ransomware-immutable-backups.html>



