



Technical Report

# **NetApp StorageGRID with Veritas Enterprise Vault**

## Implementing NetApp StorageGRID as the primary storage for Enterprise Vault archives

### Configuration guide

Jonathan Wong, NetApp  
Rahul Patil, Veritas  
July 2022 | TR-4907

In partnership with



## **Abstract**

This configuration guide provides the steps to configure NetApp® StorageGRID® as a primary storage target with Veritas Enterprise Vault. It also describes how to configure StorageGRID for site failover in a disaster recovery (DR) scenario.

TABLE OF CONTENTS

**Basic configuration ..... 3**

- Prerequisites.....3
- Configure StorageGRID with Veritas Enterprise Vault.....4

**StorageGRID Object Lock configuration (optional) ..... 9**

- Prerequisites.....9
- Configure StorageGRID S3 Object Lock default bucket retention .....9
- Configure Enterprise Vault .....13

**StorageGRID site failover configuration (optional)..... 14**

- Prerequisites.....14
- Configure StorageGRID site failover .....14

**Where to find additional information ..... 16**

**Version history..... 17**

LIST OF FIGURES

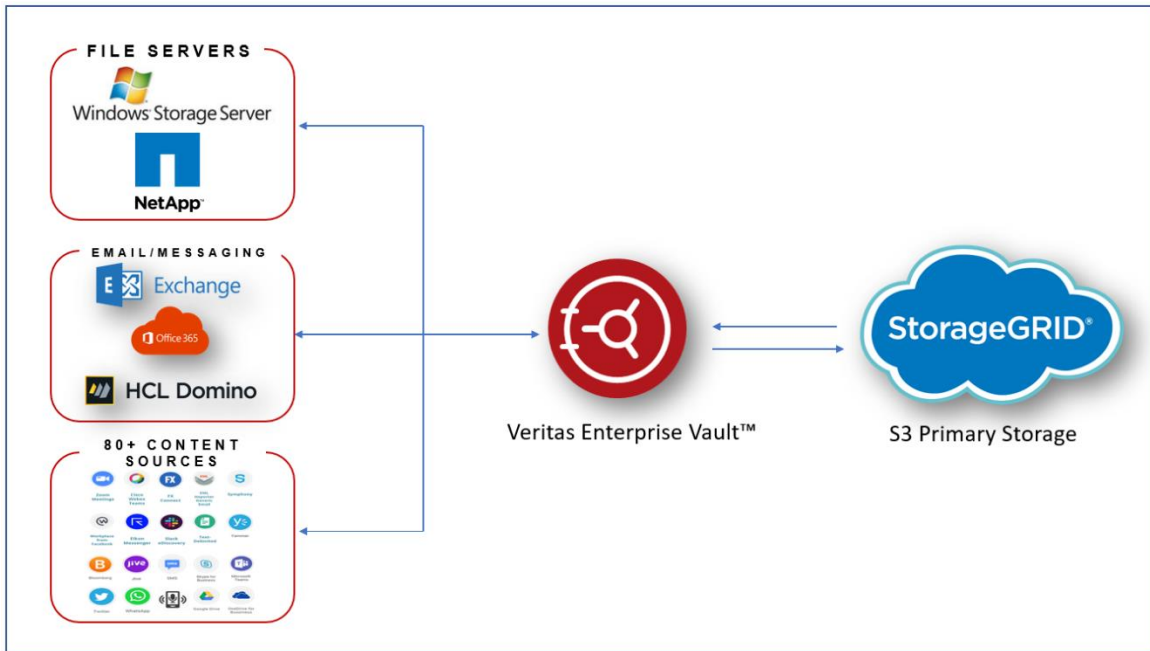
Figure 1) Veritas Enterprise Vault and StorageGRID architecture. ....3

## Reference architecture

StorageGRID provides an on-premises, S3-compatible cloud backup target for Veritas Enterprise Vault.

Figure 1 illustrates the Veritas Enterprise Vault and StorageGRID architecture.

Figure 1) Veritas Enterprise Vault and StorageGRID architecture.



## Basic configuration

This configuration guide is based on StorageGRID 11.5 and Enterprise Vault 14.1. For write once, read many (WORM) mode storage using S3 Object Lock, StorageGRID 11.6 and Enterprise Vault 14.2.2 was used. For more detailed information about these guidelines, see the [StorageGRID 11.6 Documentation](#) page or contact a StorageGRID expert.

### Prerequisites

Before you configure StorageGRID with Veritas Enterprise Vault, verify the following prerequisites:

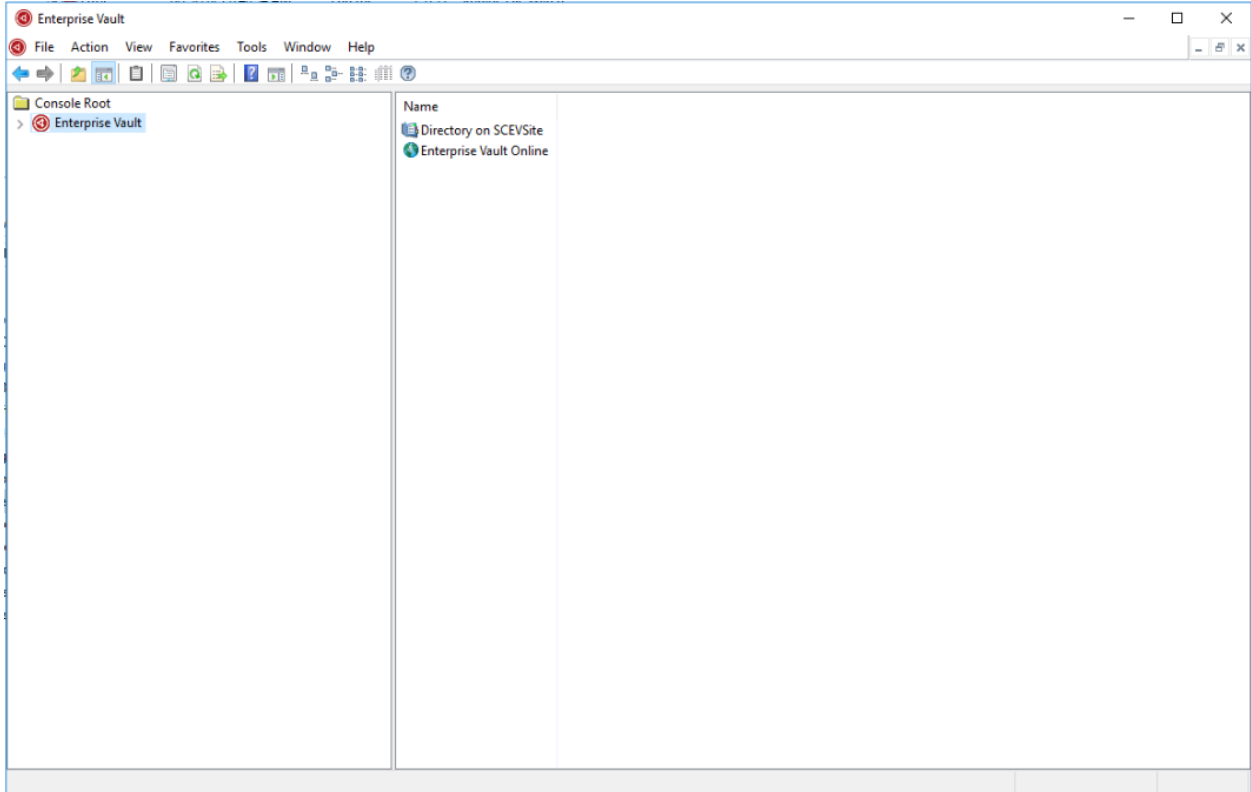
- StorageGRID 11.5+ is installed.  
**Note:** For WORM storage (Object Lock), StorageGRID 11.6 is required.
- Veritas Enterprise Vault 14.1+ is installed.  
**Note:** For WORM storage (Object Lock), Enterprise Vault version 14.2.2 is required.
- Vault store groups and a vault store has been created.  
For more information, see the Veritas Enterprise Vault Administration Guide.
- A StorageGRID tenant, access key, secret key and bucket have been created.
- A StorageGRID load balancer endpoint has been created (either HTTP or HTTPS).
- If using a self-signed certificate, add the StorageGRID self-signed CA certificate to the Enterprise Vault Servers. For more information, see this [Veritas Knowledge Base article](#).

- Update and apply the latest Enterprise Vault configuration file to enable supported storage solutions such as NetApp StorageGRID. For more information, see this [Veritas Knowledge Base article](#).

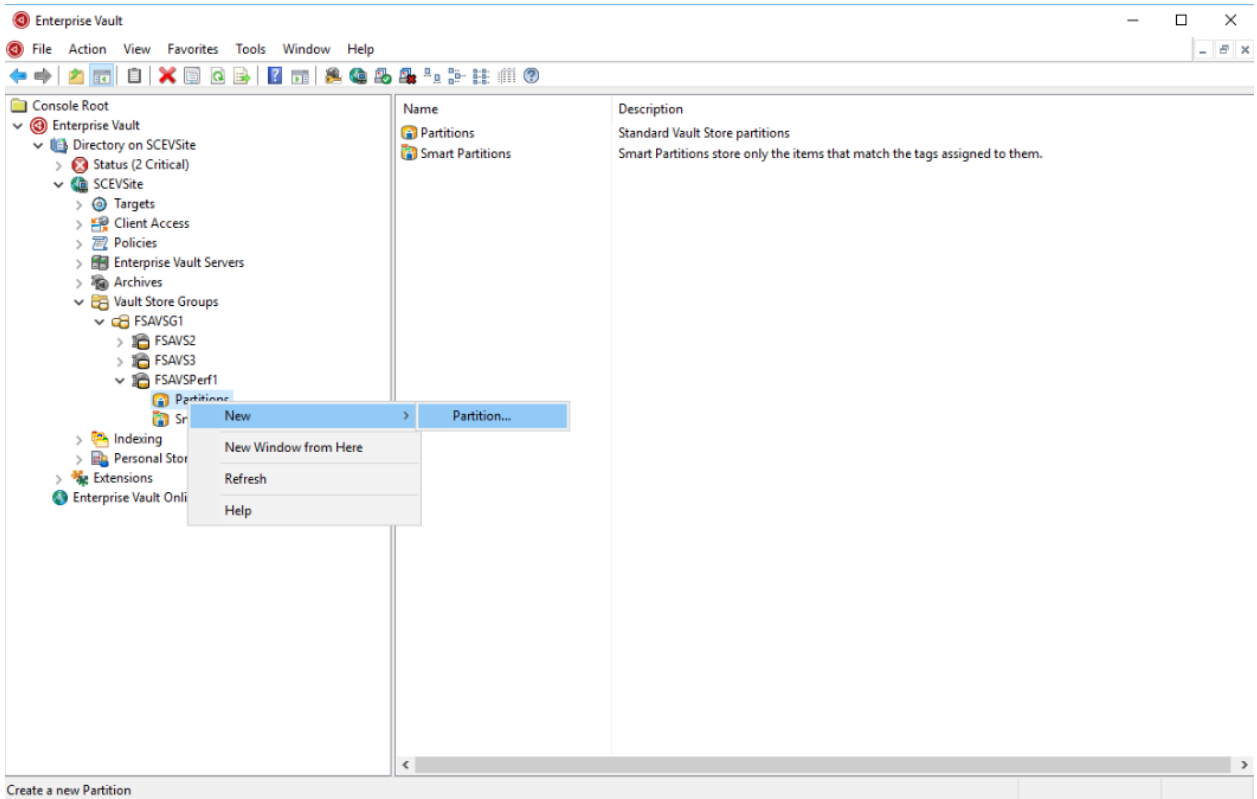
## Configure StorageGRID with Veritas Enterprise Vault

To configure StorageGRID with Veritas Enterprise Vault, complete the following steps:

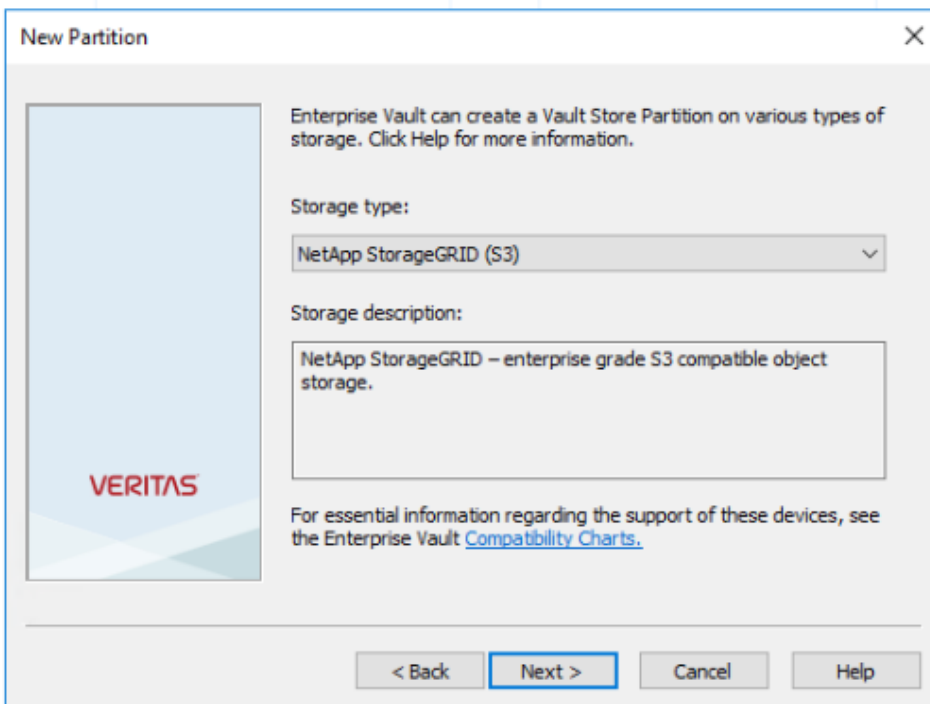
1. Launch the Enterprise Vault Administration console.



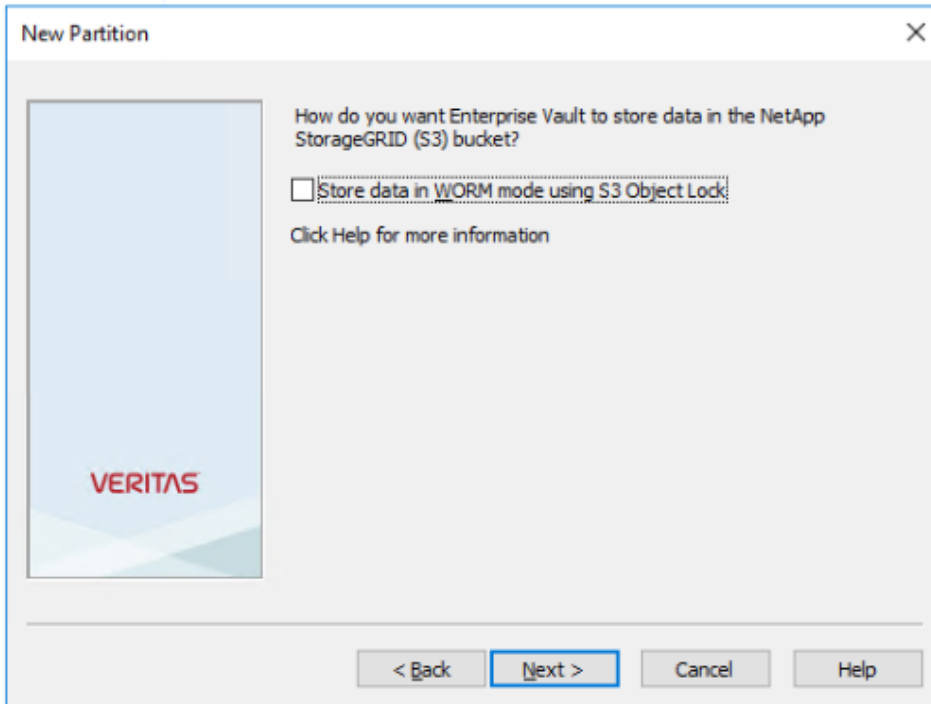
2. Create a new vault store partition in the appropriate vault store. Expand the Vault Store Groups folder and then the appropriate vault store. Right-click Partition and select New > Partition.



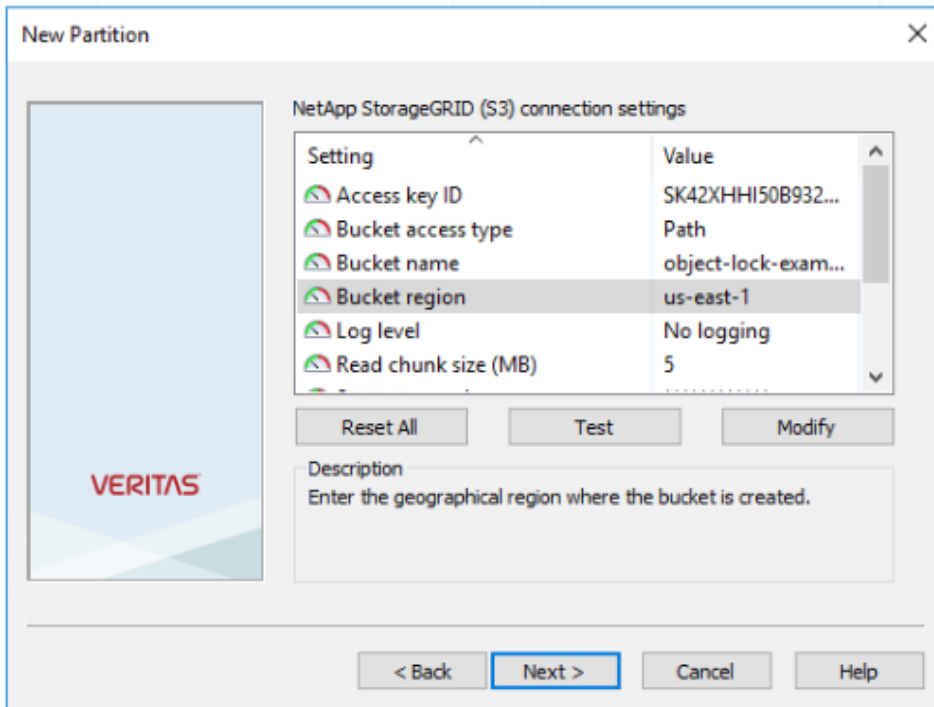
3. Follow the New Partition creation wizard. From the Storage Type drop-down menu, select NetApp StorageGRID (S3). Click Next.




4. Leave the Store Data in WORM Mode Using S3 Object Lock option unchecked. Click Next.



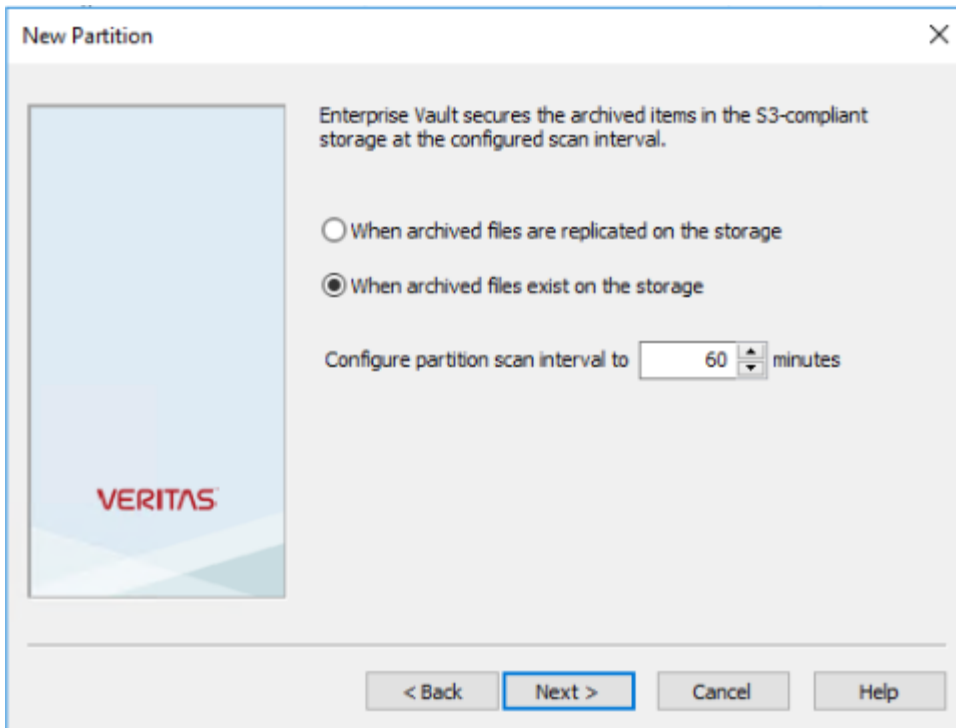
5. On the connection settings page, provide the following information:
- Access key ID
  - Secret access key
  - Service host name: Ensure to include the load balancer endpoint (LBE) port configured in StorageGRID (such as `https://<hostname>:<LBE_port>`).
  - Bucket name: Name of the precreated target bucket. Veritas Enterprise Vault does not create the bucket.
  - Bucket region: `us-east-1` is the default value.



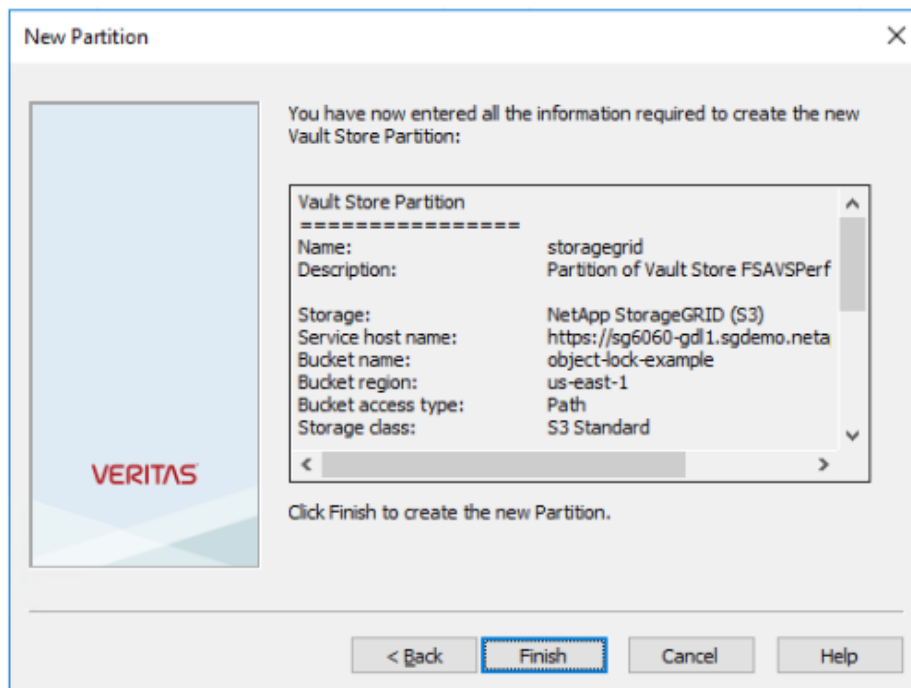
- To verify the connection to the StorageGRID bucket, click Test. Verify that the connection test was successful. Click OK and then Next.



- StorageGRID does not support the S3 replication parameter. To protect your objects, StorageGRID uses Information Lifecycle Management (ILM) rules to specify data protection schemes – multiple copies or erasure coding. Select the When Archived Files Exist on the Storage Option and click Next.

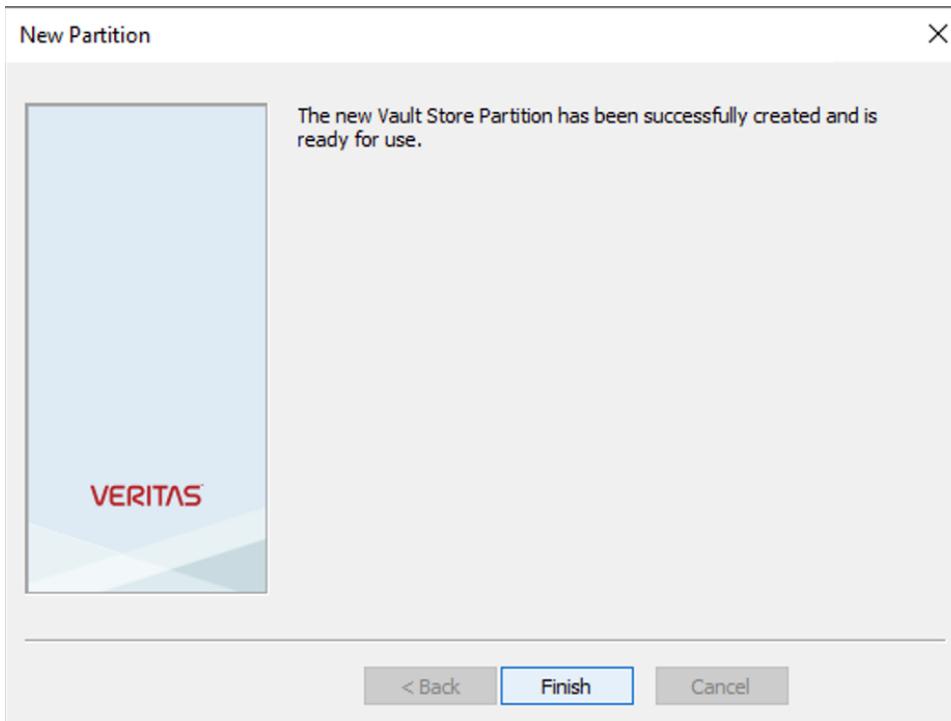


8. Verify the information on the summary page and click Finish.



9. After the new vault store partition has been successfully created, you can archive, restore, and search data in Enterprise Vault with StorageGRID as the primary storage.





## StorageGRID Object Lock configuration (optional)

### Prerequisites

For WORM storage, StorageGRID uses S3 Object Lock to retain objects for compliance. This requires StorageGRID 11.6, where S3 Object Lock default bucket retention was introduced. Enterprise Vault also requires version 14.2.2

### Configure StorageGRID S3 Object Lock default bucket retention

To configure the StorageGRID S3 Object Lock default bucket retention, complete the following steps:

1. In StorageGRID Tenant Manager, create a bucket and click Continue

# Create bucket

×

1 Enter details ————— 2 Manage object settings  
Optional

## Enter bucket details

Enter the bucket's name and select the bucket's region.

Bucket name ⓘ

Region ⓘ

Cancel Continue

2. Select the Enable S3 Object Lock option and click Create Bucket.

# Create bucket ✕

Enter details 2 Manage object settings Optional

## Manage object settings Optional

### Object versioning

Enable object versioning if you want to store every version of each object in this bucket. You can then retrieve previous versions of an object as needed.

*i* Object versioning has been enabled automatically because this bucket has S3 Object Lock enabled.

Enable object versioning

### S3 Object Lock

S3 Object Lock allows you to specify retention and legal hold settings for the objects ingested into a bucket. If you want to use S3 Object Lock, you must enable this setting when you create the bucket. You cannot add or disable S3 Object Lock after a bucket is created.

If S3 Object Lock is enabled, object versioning is enabled for the bucket automatically and cannot be suspended.

Enable S3 Object Lock

[Previous](#) [Create bucket](#)

3. After the bucket is created, select the bucket to view the bucket options. Expand the S3 Object Lock drop-down option.

### Overview

Name:	object-lock-example
Region:	us-east-1
S3 Object Lock:	Enabled
Date created:	2022-06-24 14:44:54 PDT

[View bucket contents in Experimental S3 Console](#)

**Bucket options** | [Bucket access](#) | [Platform services](#)

Consistency level	Read-after-new-write (default)
Last access time updates	Disabled
Object versioning	Enabled

### S3 Object Lock

Enabled

S3 Object Lock allows you to specify retention and legal hold settings for the objects ingested into a bucket. If you want to use S3 Object Lock, you must enable this setting when you create the bucket. You cannot enable or disable S3 Object Lock after a bucket is created.

After S3 Object Lock is enabled for a bucket, you can't disable it. You also can't suspend object versioning for the bucket.

S3 Object Lock

Enabled

Default retention

Disable

Enable

Save changes

- Under Default Retention, select Enable and set a default retention period of 1 day. Click Save Changes.

**S3 Object Lock** Enabled

S3 Object Lock allows you to specify retention and legal hold settings for the objects ingested into a bucket. If you want to use S3 Object Lock, you must enable this setting when you create the bucket. You cannot enable or disable S3 Object Lock after a bucket is created.

**Info** After S3 Object Lock is enabled for a bucket, you can't disable it. You also can't suspend object versioning for the bucket.

**S3 Object Lock**  
Enabled

**Default retention**

Disable

Enable

**Default retention mode**

**Compliance**  
No users can overwrite or delete protected object versions during the retention period.

**Default retention period**

1 Days

[Save changes](#)

The bucket is now ready to be used by Enterprise Vault to store WORM data.

## Configure Enterprise Vault

To configure Enterprise Vault, complete the following steps:

1. Repeat steps 1–3 in the “Basic configuration” section, but this time select the Store data in WORM Mode Using S3 Object Lock option. Click Next.

**New Partition** ✕

How do you want Enterprise Vault to store data in the NetApp StorageGRID (S3) bucket?

Store data in WORM mode using S3 Object Lock

[Click Help for more information](#)

**VERITAS**

2. When entering your S3 Bucket connection settings, make sure you are entering the name of an S3 bucket that has the S3 Object Lock Default retention enabled.
3. Test the connection to verify the settings.

## StorageGRID site failover configuration (optional)

It is a common for a StorageGRID architecture deployment to be multisite. Sites can either be active-active or active-passive for DR. In a DR scenario, make sure that Veritas Enterprise Vault can maintain connection to its primary storage (StorageGRID) and continue to ingest and retrieve data during a site failure. This section provides high-level configuration guidance for a two-site, active-passive deployment. For detailed information about these guidelines, see the [StorageGRID 11.6 Documentation](#) page or contact a StorageGRID expert.

### Prerequisites

Before you configure StorageGRID site failover, verify the following prerequisites:

- There is a two-site StorageGRID deployment; for example, SITE1 and SITE2.
- An admin node running the load balancer service or a gateway node, at each site, for load balancing has been created.
- A StorageGRID load balancer endpoint has been created.

### Configure StorageGRID site failover

To configure StorageGRID site failover, complete the following steps:

1. To ensure connectivity to StorageGRID during site failures, configure a high-availability (HA) group. From StorageGRID Grid Manager Interface (GMI), click Configuration, High Availability Groups, and + Create.

**Create High Availability Group**

**High Availability Group**

Name

Description

**Interfaces**

Select interfaces to include in the HA group. All interfaces must be in the same network subnet.

**Virtual IP Addresses**

Select interfaces before assigning virtual IP addresses.

2. Enter the required information. Click Select Interfaces and include both SITE1 and SITE2's network interfaces where SITE1 (the primary site) is the preferred master. Assign a virtual IP address within the same subnet. Click Save.

### Edit High Availability Group 'site1-HA'

**High Availability Group**

Name:

Description:

**Interfaces**

Select interfaces to include in the HA group. All interfaces must be in the same network subnet.

Node Name	Interface	IPv4 Subnet	Preferred Master
SITE1-ADM1	eth2	[REDACTED] 205.0/24	<input checked="" type="radio"/>
SITE2-ADM1	eth2	[REDACTED] 205.0/24	<input type="radio"/>

Displaying 2 interfaces.

**Virtual IP Addresses**

Virtual IP Subnet: 10.193.205.0/24. All virtual IP addresses must be within this subnet. There must be at least 1 and no more than 10 virtual IP addresses.

Virtual IP Address 1:

- This virtual IP (VIP) address should be associated to the S3 host name used during Veritas Enterprise Vault's partition configuration. The VIP address resolves traffic to SITE1—and during SITE1 failure, the VIP address transparently reroutes traffic to SITE2.
- Make sure the data is replicated to both SITE1 and SITE2. That way if SITE1 fails, the object data is still available from SITE2. This is done by first configuring the storage pools.  
From StorageGRID GMI, click ILM, Storage Pools, and then + Create. Follow the wizard to create two storage pools: one for SITE1 and another for SITE2.  
Storage pools are logical groupings of nodes used to define object placement

Storage Pool Details - site1

Number of Nodes: 4  
Storage Grade: All Storage Nodes

Node Name	Site Name	Used (%)
SITE1-S3	SITE1	0.449%
SITE1-S4	SITE1	0.401%
SITE1-S2	SITE1	0.393%
SITE1-S1	SITE1	0.312%

Storage Pool Details - site2		
Nodes Included <span style="float: right;">ILM Usage</span>		
Number of Nodes: 4		
Storage Grade: All Storage Nodes		
Node Name	Site Name	Used (%) <span style="float: right;">↑↓</span>
SITE2-S2	SITE2	0.382%
SITE2-S1	SITE2	0.417%
SITE2-S3	SITE2	0.434%
SITE2-S4	SITE2	0.323%

- From StorageGRID GMI, click ILM, Rules, and then + Create. Follow the wizard to create an ILM rule specifying one copy to be stored per site with an ingest behavior of Balanced.

**1 copy per site**

Description: 1 copy per site  
 Ingest Behavior: Balanced  
 Reference Time: Ingest Time  
 Filtering Criteria: Matches all objects.

Retention Diagram:

- Add the ILM rule into an ILM policy and activate the policy. This configuration results in the following outcome:
  - A virtual S3 endpoint IP where SITE1 is the primary and SITE2 is the secondary endpoint. If SITE1 fails, the VIP fails over to SITE2.
  - When archived data is sent from Veritas Enterprise Vault, StorageGRID ensures one copy is stored in SITE1 and another DR copy is stored in SITE2. If SITE1 fails, Enterprise Vault continues to ingest and retrieve from SITE2.

**Note:** Both of these configurations are transparent to Veritas Enterprise Vault. The S3 endpoint, bucket name, access keys, and so on are the same. There is no need to reconfigure the S3 connection settings on the Veritas Enterprise Vault partition.

## Where to find additional information

To learn more about the information that is described in this document, review the following documents and/or websites:

- NetApp StorageGRID Product Documentation <https://docs.netapp.com/us-en/storagegrid-116/>
- StorageGRID documentation resources <https://www.netapp.com/data-storage/storagegrid/documentation/>
- Veritas Enterprise Vault Product Documentation [https://www.veritas.com/support/en\\_US/article.100049763](https://www.veritas.com/support/en_US/article.100049763)



## Version history

Version	Date	Document version history
Version 1.0	September 2021	Initial release.
Version 1.1	December 2021	Added a KB article pertaining to self-signed certificates.
Version 1.2	July 2022	Added the Object Lock configuration for WORM storage.

Refer to the [Interoperability Matrix Tool \(IMT\)](#) on the NetApp Support site to validate that the exact product and feature versions described in this document are supported for your specific environment. The NetApp IMT defines the product components and versions that can be used to construct configurations that are supported by NetApp. Specific results depend on each customer's installation in accordance with published specifications.

### **Copyright information**

Copyright © 2022 NetApp, Inc. All Rights Reserved. Printed in the U.S. No part of this document covered by copyright may be reproduced in any form or by any means—graphic, electronic, or mechanical, including photocopying, recording, taping, or storage in an electronic retrieval system—without prior written permission of the copyright owner.

Software derived from copyrighted NetApp material is subject to the following license and disclaimer:

THIS SOFTWARE IS PROVIDED BY NETAPP "AS IS" AND WITHOUT ANY EXPRESS OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE, WHICH ARE HEREBY DISCLAIMED. IN NO EVENT SHALL NETAPP BE LIABLE FOR ANY DIRECT, INDIRECT, INCIDENTAL, SPECIAL, EXEMPLARY, OR CONSEQUENTIAL DAMAGES (INCLUDING, BUT NOT LIMITED TO, PROCUREMENT OF SUBSTITUTE GOODS OR SERVICES; LOSS OF USE, DATA, OR PROFITS; OR BUSINESS INTERRUPTION) HOWEVER CAUSED AND ON ANY THEORY OF LIABILITY, WHETHER IN CONTRACT, STRICT LIABILITY, OR TORT (INCLUDING NEGLIGENCE OR OTHERWISE) ARISING IN ANY WAY OUT OF THE USE OF THIS SOFTWARE, EVEN IF ADVISED OF THE POSSIBILITY OF SUCH DAMAGE.

NetApp reserves the right to change any products described herein at any time, and without notice. NetApp assumes no responsibility or liability arising from the use of products described herein, except as expressly agreed to in writing by NetApp. The use or purchase of this product does not convey a license under any patent rights, trademark rights, or any other intellectual property rights of NetApp.

The product described in this manual may be protected by one or more U.S. patents, foreign patents, or pending applications.

Data contained herein pertains to a commercial product and/or commercial service (as defined in FAR 2.101) and is proprietary to NetApp, Inc. The U.S. Government has a non-exclusive, non-transferrable, non-sublicensable, worldwide, limited irrevocable license to use the Data only in connection with and in support of the U.S. Government contract under which the Data was delivered. Except as provided herein, the Data may not be used, disclosed, reproduced, modified, performed, or displayed without the prior written approval of NetApp, Inc. United States Government license rights for the Department of Defense are limited to those rights identified in DFARS clause 252.227-7015(b).

### **Trademark information**

NETAPP, the NETAPP logo, and the marks listed at <http://www.netapp.com/TM> are trademarks of NetApp, Inc. Other company and product names may be trademarks of their respective owners.

TR-4907-0722