



Technical Report

# StorageGRID on NetApp HCI

## Solution Deployment Guide

Steven Pruchniewski and Amit Borulkar, NetApp  
November 2018 | TR-4734

### Abstract

This guide provides the recommended practices for deploying and configuring NetApp® StorageGRID® object storage on NetApp HCI. This solution enables an optimized, enterprise-grade object store capable of providing Amazon S3 compatible object storage and hybrid cloud services with the ability to seamlessly transition to a full-scale enterprise object solution.

## TABLE OF CONTENTS

<b>1</b>	<b>Solution Overview .....</b>	<b>3</b>
1.1	Use Case Summary .....	4
<b>2</b>	<b>Solution Technology .....</b>	<b>4</b>
2.1	NetApp StorageGRID .....	4
2.2	NetApp HCI .....	4
2.3	NetApp Deployment Engine .....	6
2.4	Technology Requirements .....	7
2.5	Solution Design .....	8
2.6	Deployment Procedures .....	9
<b>3</b>	<b>StorageGRID Configuration .....</b>	<b>19</b>
3.1	ILM and Data Management Rules .....	19
3.2	Storage Configuration .....	19
3.3	Customizing Storage Alerts .....	20
3.4	Getting Ready to Serve Data .....	20
3.5	Configuration Limit Testing .....	21
3.6	Realized Storage Efficiency .....	22
<b>4</b>	<b>Supporting Full-Scale Object Workloads .....</b>	<b>22</b>
4.1	Expanding for Object Space .....	22
4.2	Expanding for Metadata Space .....	22
<b>5</b>	<b>Data Protection Considerations .....</b>	<b>23</b>
	<b>Where to Find Additional Information .....</b>	<b>23</b>

## LIST OF TABLES

Table 1)	Hardware requirements .....	7
Table 2)	Software requirements .....	7
Table 3)	Deployment variables .....	9

## LIST OF FIGURES

Figure 1)	NetApp HCI minimum configuration .....	5
Figure 2)	Successful NDE deployment .....	6
Figure 3)	Solution design .....	8
Figure 4)	Storage Watermarks .....	19
Figure 6)	StorageGRID available storage .....	21

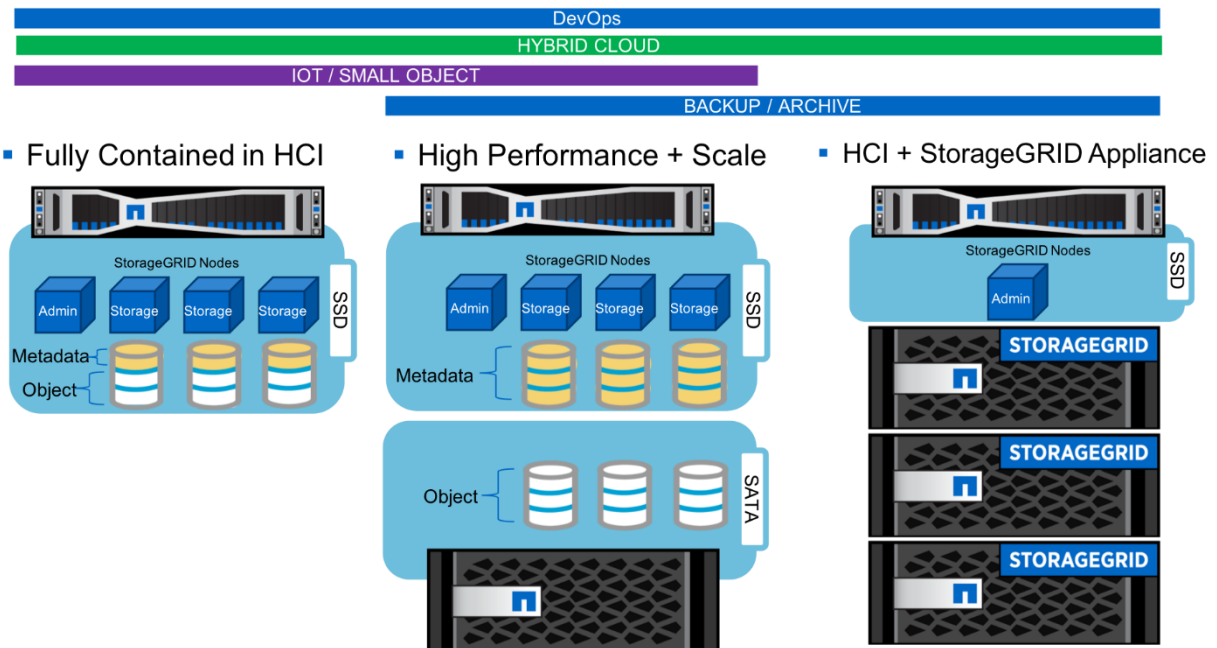
# 1 Solution Overview

The explosive growth of unstructured data has led to the rapid adoption of object storage for massive scale and efficiency, from both the storage capacity and operational perspective. The world is moving toward a model in which hot data is stored on flash technology and everything else on object stores. This technical report introduces a new solution: NetApp® StorageGRID® as an object store delivered on NetApp HCI architecture. Although several new object solutions are available in the market, StorageGRID has over a decade of production deployments with more than .5EB shipped. StorageGRID also integrates with leading cloud providers such as Amazon Web Services (AWS), enabling you to take advantage of hyperscaler resources with your own on-premises storage to enable hybrid cloud workflows. This solution provides an Amazon Simple Storage Service (S3) compatible object store with enterprise-grade features and the ability to seamlessly scale for massive scale and high-performance object workloads.

StorageGRID nodes are set up as virtual machines (VMs) on NetApp HCI, a system that is easy to deploy and manage. NetApp HCI provides a flexible architecture that allows StorageGRID to fit multiple use cases:

- **Fully contained in NetApp HCI.** This option reduces the footprint of object store deployment by using NetApp HCI resources.
- **High performance and scale.** This option uses the compute and flash storage of NetApp HCI for StorageGRID nodes, StorageGRID metadata and uses dense, lower-cost NetApp E-Series systems for object data.
- **NetApp HCI and StorageGRID appliance.** This option uses NetApp HCI to host the StorageGRID Admin Node and load balancer and uses the StorageGRID hardware appliance for Storage Nodes.

Figure 1) NetApp HCI deployment options.



This document describes the “fully contained in NetApp HCI” deployment option. The goal is to provide a fully functional enterprise-grade object store within the storage footprint of a single NetApp HCI chassis.

The solution is designed to use less than 10% of the compute and storage resources to provide up to 100 million objects and 1.65TB of object storage. Because more applications require S3 support and datasets are growing, this S3 target is expected to eventually serve a full-scale object workload requiring greater resources. This document describes how to configure StorageGRID to operate in a smaller footprint. It

also provides guidance for graduating this instance to a full-size object store that supports many billions of objects and petabytes of storage.

## 1.1 Use Case Summary

This solution applies to the following use cases:

- On-premises S3-compatible object storage
- Hybrid cloud workflows
- Shared image repositories that require simultaneous read/write capabilities—for example, an OpenShift container registry or Docker registry
- Workloads that ingest a large amount of unstructured data in which directory structures and file system usage can be cumbersome to manage (for example, Internet of Things [IoT] and deep learning workloads)
- Workloads that must comply with data governance and regulatory requirements yet maintain simple data access

To support other workloads that require more storage (such as backup), choose one of the other StorageGRID deployment options specified in the section “Solution Overview.”

## 2 Solution Technology

### 2.1 NetApp StorageGRID

StorageGRID deployment on NetApp HCI forms the foundation for an on-premises or hybrid cloud object store that has high performance, is easily extensible, and is software-defined.

StorageGRID scales performance and capacity by simply adding more StorageGRID Storage Nodes, an approach that closely matches with NetApp HCI's scaling model. Compute and storage resources can be scaled independently of one another, and you can expand your StorageGRID instance to support your workload. For example, you can add more StorageGRID nodes with less storage to support high-performance small objects, or you can have fewer StorageGRID nodes with more storage to achieve greater density.

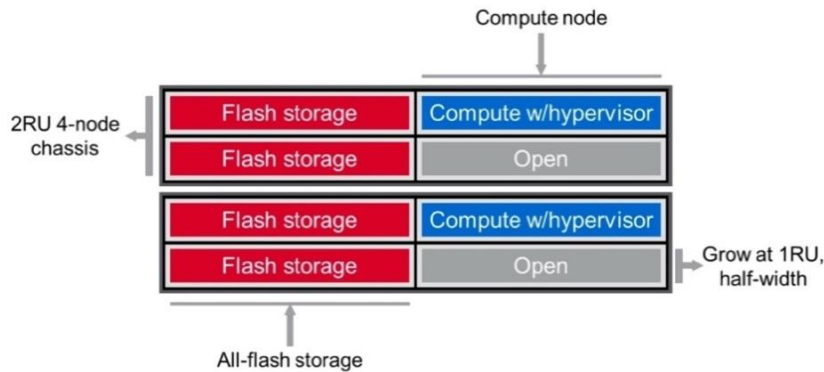
StorageGRID platform services are the foundation for realizing the promise of hybrid cloud, letting you tier and replicate to public or other S3-compatible clouds. StorageGRID can also use Amazon Simple Notification Service to invoke AWS Lambda functions, allowing you to store data in cost-effective on-premises objects while using compute features of the AWS marketplace. For more information about taking advantage of hybrid cloud features, see the [StorageGRID tenant administration documentation](#).

### 2.2 NetApp HCI

NetApp HCI is an enterprise-scale hybrid cloud infrastructure solution that delivers compute and storage resources in an agile, scalable, easy-to-manage two-rack-unit, four-node building block. It is based on:

- NetApp H-Series all-flash storage nodes running NetApp Element® software
- NetApp H-Series compute nodes running VMware ESXi
- The NetApp Deployment Engine (NDE) and NetApp Element vSphere plug-in, which enable deployment and management of NetApp HCI

Figure 2) NetApp HCI minimum configuration.



For details and technical specifications of compute and storage nodes in NetApp HCI, see the [NetApp HCI datasheet](#).

## NetApp HCI Design Principles

NetApp HCI enables you to run enterprise-class, virtualized, and containerized workloads in an accelerated manner by providing an agile turnkey infrastructure system. At its core, NetApp HCI is designed to provide predictable performance, linear scalability, and a simple deployment and management experience.

### Predictable

One of the biggest challenges in a multitenant environment is delivering predictable performance consistently for all workloads. Running multiple enterprise-grade workloads can result in resource contention: One workload can interfere with the performance of another. NetApp HCI alleviates this concern with quality of service (QoS) limits available natively with Element software. Element allows the granular control of every application and volume, eliminates “noisy neighbors,” and satisfies all performance SLAs. NetApp HCI’s multitenancy capabilities can help eliminate more than 90% of traditional performance-related problems.<sup>1</sup>

### Flexible

In the past, hyper converged infrastructure has required fixed resource ratios, limiting deployments to four to eight-node configurations. NetApp HCI, however, is a hybrid cloud infrastructure that scales compute and storage resources independently. Independent scaling avoids costly and inefficient overprovisioning, eliminates the 10% to 30% “HCI tax” from controller VM overhead, and simplifies capacity and performance planning.

With NetApp HCI, licensing costs are reduced (adding only compute or storage according to requirements). NetApp HCI is available in small, medium, and large storage and compute configurations that can be mixed. The architectural design choices enable you to confidently scale on your terms, making the infrastructure viable for core data center applications and platforms.

NetApp HCI is architected in building blocks at either the chassis or the node level. Each chassis can hold four nodes, made up of storage nodes, compute nodes, or both. A minimum configuration is two chassis

---

<sup>1</sup> <https://www.netapp.com/us/resources/esg-lab-report-quantifying-the-economic-value-of-a-solidfire-deployment>

with six nodes, consisting of four storage nodes and two compute nodes. Two more blank spots can be used for expansion. If you follow best practices, you can mix compute and storage nodes. You can scale resources nondisruptively through a simple GUI-driven process.

## Simple

An imperative in the IT community is to automate all routine tasks, eliminating the risk of user error while freeing up resources to focus on more interesting, higher-value projects. NetApp HCI allows IT departments to become more agile and responsive by simplifying deployment and ongoing management.

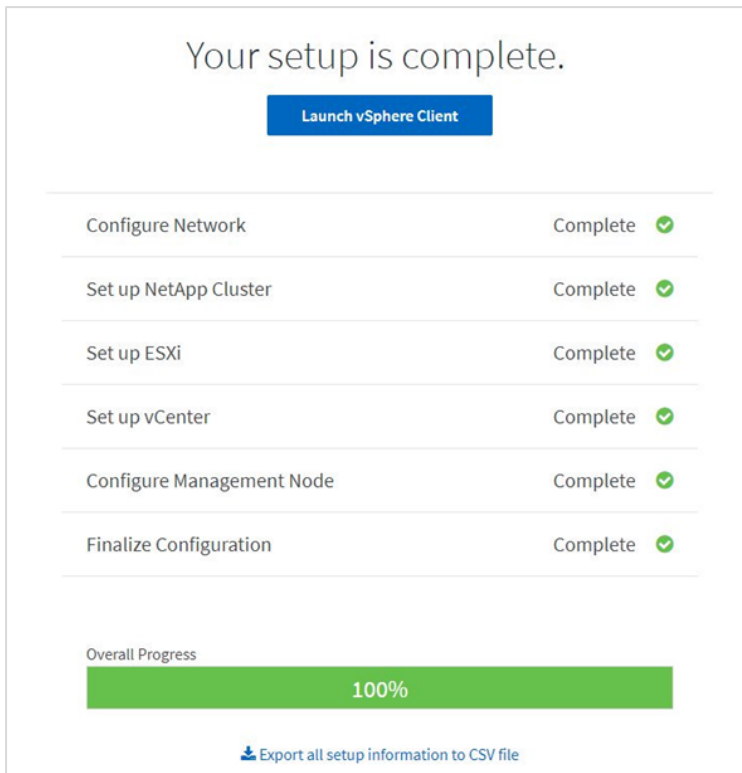
The NetApp Deployment Engine (NDE) eliminates most manual steps involved with deploying infrastructure, such as assigning names, network settings, and IP addresses and provisioning ESXi hosts and VMware datastores. You can expect the infrastructure to be functional in less than 30 minutes.

The VMware vCenter Plug-in simplifies management and makes it intuitive. NetApp HCI also works with a robust suite of APIs that enables integration into higher-level management, orchestration, backup, and disaster recovery tools.

## 2.3 NetApp Deployment Engine

NDE enables the quick deployment of NetApp HCI, including the NetApp Element software cluster and the VMware virtualized infrastructure. NDE simplifies Day 0 deployment by reducing the number of manual steps from over 300 to less than 30. Because NDE is intuitive and reuses data such as user name and password, you do not have to reenter information or set credentials at varying complexity levels. Likewise, assigning IP addresses is taken care of by NDE, allowing you to set a scheme and pool for all resources before actual configuration. Also, preinstallation checklists enable successful deployments because the system automatically checks for user errors, eliminating manual checks.

Figure 3) Successful NDE deployment.



As indicated in Figure 3, NDE optimally configures the data and management networks, configures the cluster, and sets up VMware ESXi, vCenter, and other required configurations to get your virtualized environment operating without risk.

For more information about NDE, see the [NetApp HCI Documentation Center](#).

For more information about deploying NetApp HCI, see the [NetApp HCI Deployment Guide](#).

## 2.4 Technology Requirements

This section covers the technology requirements for StorageGRID on NetApp HCI. Your requirements might vary. All the models of compute and storage listed in the NetApp [Interoperability Matrix Tool \(IMT\)](#) support the solution.

For more information about technical requirements and installation of NetApp HCI, review the [NetApp HCI Resources](#) page.

### Hardware

Table 1 lists the hardware components that were used to deploy the solution in the NetApp lab. The components might vary according to your organization's requirements.

Table 1) Hardware requirements.

Layer	Product Family	Quantity	Details
Compute	NetApp H500E	4	2 x Intel E5-2650v4; 12 cores; 2.2GHz 512GB RAM
Storage	NetApp H500S	4	6 x 960GB encrypting/non-encrypting

### Software

Table 2 lists the software components that were used to deploy the solution in the NetApp lab. The components might vary according to your organization's requirements.

Table 2) Software requirements.

Layer	Software	Version
Storage	NetApp Element	10.4
NetApp HCI engine	NetApp Deployment Engine	1.3.1
Hypervisor and above	Hypervisor	VMware vSphere ESXi 6.5 U2
	Hypervisor management system	VMware vCenter Server 6.5
	StorageGRID	11.1

**Note:** NetApp HCI is switch vendor agnostic and relies on standard enterprise-class data center switching features. The network design is described in section 2.5, "Solution Design."

### Object Store Limits

This optimized-footprint StorageGRID instance supports:

- 100 million objects
- 1.65TB of object data (protected by the two-copy policy)




As your workload approaches these limits, you can easily add more StorageGRID Storage Nodes to the system in a nondisruptive manner. StorageGRID licensing is based on capacity through either a perpetual or subscription-based license. NetApp recommends that you purchase a minimum 1TB license to enable support with this environment.

## 2.5 Solution Design

Figure 4 illustrates the solution design. The design trade-offs for the compute, network, and storage elements are described in the sections that follow.

Figure 4) Solution design.



### Compute

- The StorageGRID Storage Nodes are configured with VMware VM-VM anti-affinity rules. These rules enable the Storage Nodes to be provisioned on different NetApp HCI compute nodes, resulting in uniform resource distribution and high availability in physical node failure scenarios.
- Hot vMotion of StorageGRID Storage Nodes is not supported. Under certain circumstances, hot vMotion can cause Network Time Protocol (NTP) lock issues. Storage nodes should be powered off before vMotion.
- HAProxy is configured to load-balance the traffic among the three StorageGRID Storage Nodes.

### Network

- The solution implements a single network for all StorageGRID traffic. This network can be configured with a default gateway to reach AWS for StorageGRID platform services. A dedicated port group with VLAN tagging is created to carry the StorageGRID traffic.

**Note:** You might choose to implement separate networks to isolate client, grid, and Admin Network traffic. For detailed instructions about implementing additional networks, see the [StorageGRID Installation Guide](#).

### Storage

- A 3TB datastore is provisioned to host the StorageGRID VMs and object data.
- One 900GB VMDK is assigned to each StorageGRID Storage Node.
- The two-copy rule is configured on StorageGRID to achieve data availability in failure scenarios.



By default, Element software enables deduplication and compression. Therefore, greater storage efficiencies are realized without the overhead of storing an extra replica.

**Note:** The solution is designed to provide an enterprise-grade object store while consuming less than 10% of NetApp HCI resources.

## 2.6 Deployment Procedures

NetApp HCI deployment is outside the scope of this document. For deployment details, see the [NetApp HCI Documentation Center](#). Table 3 Table 3) Deployment variables.lists the variables used to deploy and configure StorageGRID.

**Table 3) Deployment variables.**

Variable	Value
<<grid_vlan>>	300
<<GridNetwork-Gateway>>	172.36.20.1
<<GridNetwork- Mask>>	255.255.255.0
<<GridNetwork-Primary-Admin-IP>>	172.36.20.100
<<GridNetwork-Primary-Admin-Name>>	sg-primary-admin
<<GridNetwork-Sg-sn01-IP>>	172.36.20.101
<<GridNetwork-Sg-sn01-Name>>	sg-sn-01
<<GridNetwork-Sg-sn02-IP>>	172.36.20.102
<<GridNetwork-Sg-sn02-Name>>	sg-sn-03
<<GridNetwork-Sg-sn03-IP>>	172.36.20.103
<<GridNetwork-Sg-sn03-Name>>	sg-sn-03
<<ntp-server>>	time.netapp.com
<<Grid-subnet>>	172.36.20.0/24
<<dns-server>>	172.36.20.1

**Note:** Ensure that all the components have been synchronized with an NTP server before proceeding with NetApp HCI deployment.

### Creating a StorageGRID Grid Network on the Switch Infrastructure


To set up the switches for StorageGRID on NetApp HCI, complete the following steps:

1. Log in to the 10/25GbE data switches.
2. Create a <<grid\_vlan>> VLAN for StorageGRID traffic on both switches.
3. Trunk the <<grid\_vlan>> VLAN onto the switch ports that are connected to NetApp HCI compute nodes.

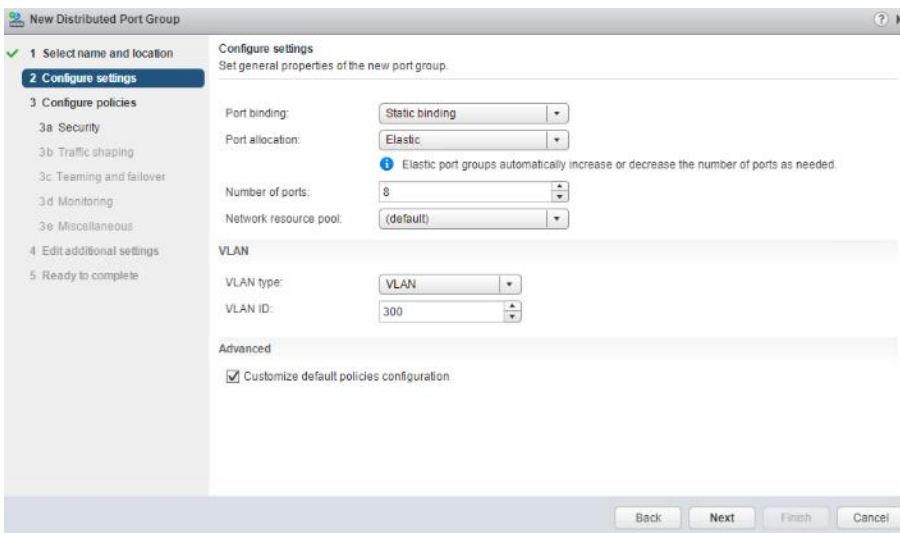
## Creating a StorageGRID Grid Network in vSphere

To create StorageGRID Grid Network in vSphere, complete the following steps:

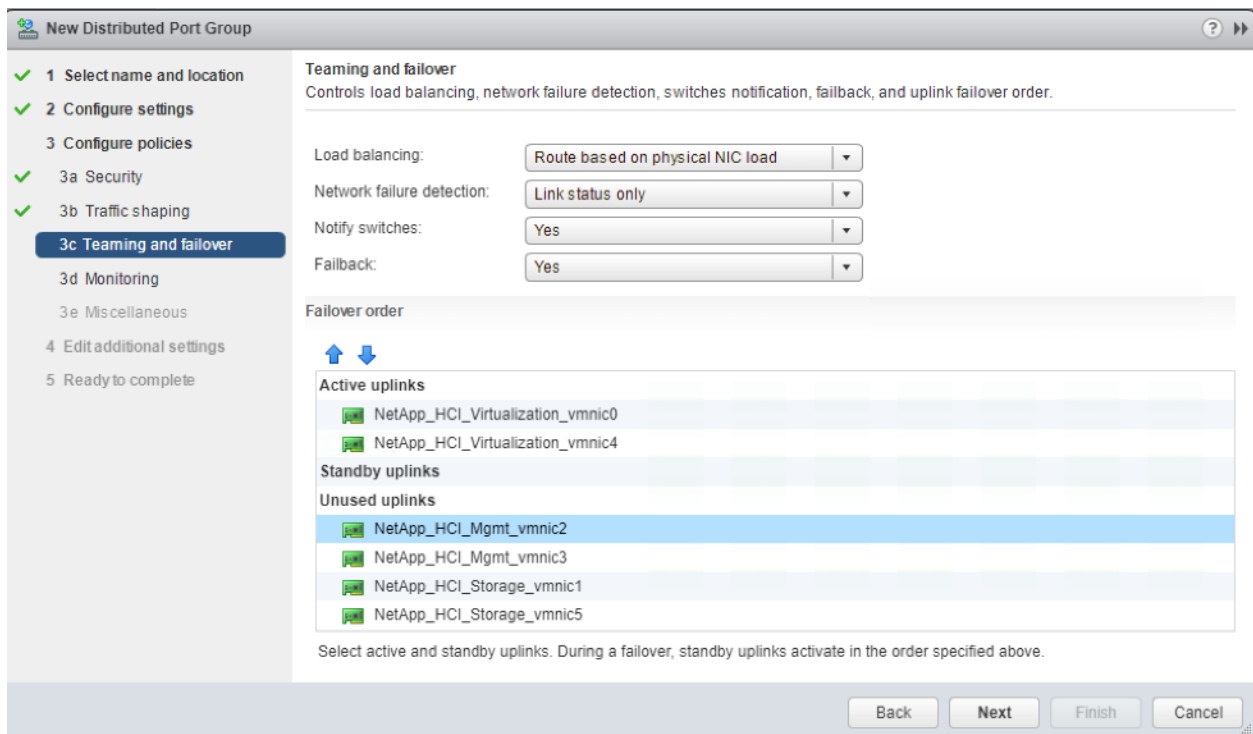
1. Create a distributed port group for the Grid Network.



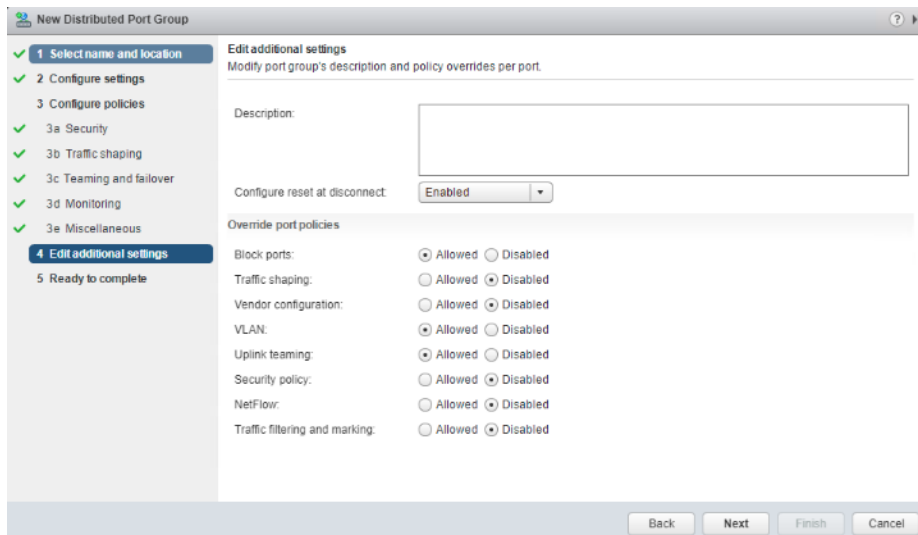
2. Configure <<grid\_vlan>> in the general properties and customize the default policies.



3. Click Next to configure security and traffic settings.
4. In the Teaming and Failover section, configure the `NetApp_HCI_Virtualization` uplinks.



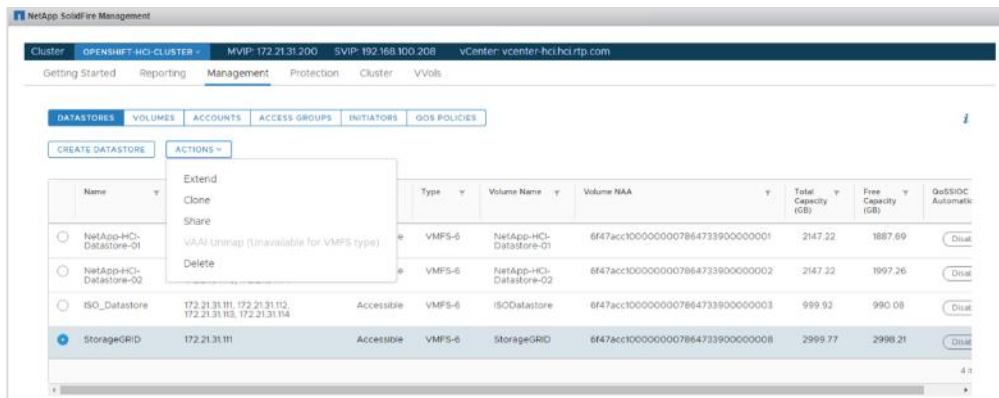
5. Allow block ports, VLAN, and uplink teaming, and complete the configuration.



## Creating a Datastore for StorageGRID Object Storage

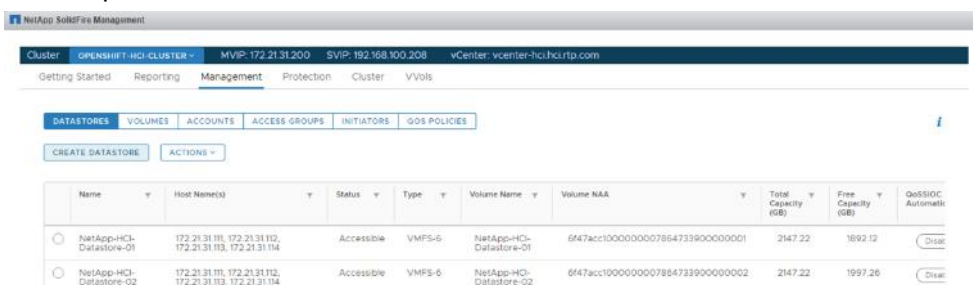
To create a datastore, complete the following steps:

1. Navigate to NetApp SolidFire® Management in vSphere.
2. Click the Management tab and then click Create Datastore.



3. Select one of the compute nodes in the NetApp HCI cluster.
4. Create a volume of size 3000GB and select the default custom settings for QoS.

5. Use the default NetApp-HCI access group and complete the datastore creation.
6. Select the datastore. From the Actions drop-down, click Share to mount the datastore to all the compute hosts in the ESXi cluster.



## Deploying the StorageGRID Admin Node and StorageGRID Storage Nodes

For detailed instructions for deploying StorageGRID in a VMware environment, see the [StorageGRID Documentation Center](#). This document only describes the modifications necessary for deploying a StorageGRID object store on NetApp HCI.

1. Download the latest StorageGRID installation suite from the [NetApp Support site](#).

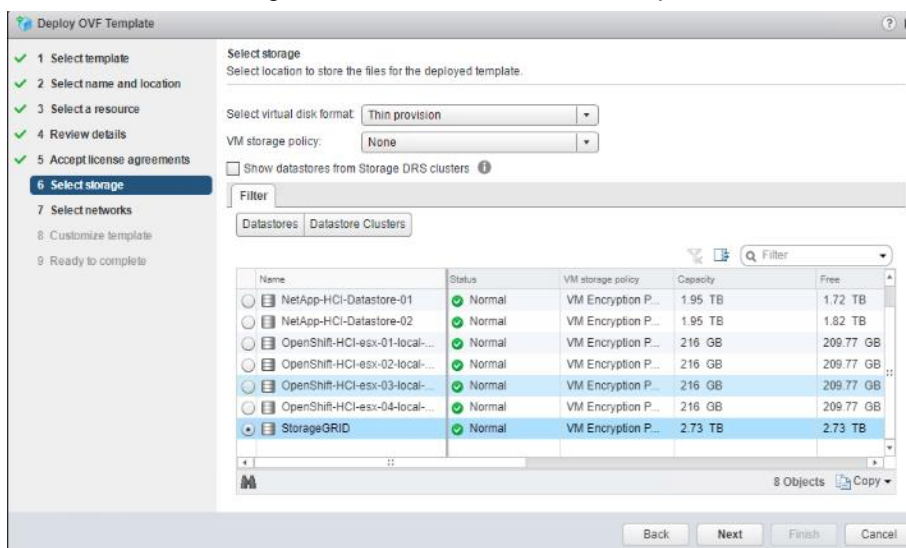
2. Because of this StorageGRID instance's smaller targeted footprint, each StorageGRID Storage Node has one 1TB disk assigned for storing the objects and metadata information. The default for StorageGRID is 3 x 4TB, and you can attach up to 16 LUNs per Storage Node.
3. Extract the downloaded StorageGRID artifacts.

```
tar -xvf StorageGRID-Webscale-11.1.0-VMware-20180619.1826.5d20160.tgz
```

4. Modify Disk-Section in the `vsphere-storage.ovf` file to reflect the new virtual disk capacity.

```
<DiskSection>
  <Info>Virtual disk information</Info>
  <Disk ovf:capacity="100" ovf:capacityAllocationUnits="byte * 2^30" ovf:diskId="sgroot"
  ovf:fileRef="sgroot-vmdk"
  ovf:format="http://www.vmware.com/interfaces/specifications/vmdk.html#streamOptimized"
  ovf:populatedSize="462553600" />
  <!-- @@DISK:ANCHOR@@
  <Disk ovf:capacity="@@DISK:CAPACITY@" ovf:capacityAllocationUnits="byte * 2^30"
  ovf:diskId="@@DISK:ID@"
  ovf:format="http://www.vmware.com/interfaces/specifications/vmdk.html#streamOptimized" />
  -->
  <!-- @@DISK:DEFAULT_START@@ -->
  <Disk ovf:capacity="900" ovf:capacityAllocationUnits="byte * 2^30" ovf:diskId="RangeDB disk
  1" ovf:format="http://www.vmware.com/interfaces/specifications/vmdk.html#streamOptimized" />
</DiskSection>
```

5. Deploy the StorageGRID primary Admin Node as described in the [StorageGRID Documentation Center](#).
  - a. Use the StorageGRID datastore created in the previous section with thin provisioning.



- b. Select the Grid Network created in the previous step and retain the Admin Network and Client Network at default settings.

Source Network	Destination Network
Grid Network	GridNetwork
Client Network	VM_Network
Admin Network	VM_Network

IP Allocation Settings  
IP protocol: IPv4      IP allocation: Static - Manual

- c. Configure <<GridNetwork-Primary-Admin-IP>>, <<GridNetwork-Gateway>>, <<GridNetwork-Mask>>, and <<GridNetwork-Primary-Admin-Name>>.

Customize template  
Customize the deployment properties of this software solution.

All properties have valid values

- Admin Network (eth1): 5 settings
- Client Network (eth2): 4 settings
- Grid Network (eth0): 3 settings
  - Grid network IP: The IP address of this node on the grid network. 172.36.20.100
  - Grid network gateway: The grid network gateway for this node (must be on the grid network subnet). Optional if the grid network consists of a single subnet, leave as 0.0.0.0 in that case. 172.36.20.1
  - Grid network mask: The netmask for this node's grid network subnet. 255.255.255.0
- Grid Node Parameters: 1 setting
  - Node name: Optional name for this node. Must be a valid hostname and consist only of letters, numbers and hyphens. sg-primary-admin

6. Deploy three StorageGRID Storage Nodes. For the deployment procedures, see the [StorageGRID Documentation Center](#).
  - a. Use the modified OVF template and select the StorageGRID datastore created previously.
  - b. Select the Grid Network created in the previous step and retain the Admin Network and Client Network at default settings.
  - c. Configure <<GridNetwork-Sg-sn01-IP>>, <<GridNetwork-Gateway>>, <<GridNetwork-Mask>>, and <<GridNetwork-Sg-sn01-Name>>.

**Deploy OVF Template**

- 1 Select template
- 2 Select name and location
- 3 Select a resource
- 4 Review details
- 5 Select storage
- 6 Select networks
- 7 Customize template**
- 8 Ready to complete

**Customize template**  
Customize the deployment properties of this software solution.

All properties have valid values [Show next...](#) [Collapse all...](#)

Client Network (eth2)	4 settings
Grid Network (eth0)	4 settings
Grid network IP	The IP address of this node on the grid network. 172.36.20.101
Grid network gateway	The grid network gateway for this node (must be on the grid network subnet). Optional if the grid network consists of a single subnet, leave as 0.0.0.0 in that case. 172.36.20.1
Grid network mask	The netmask for this node's grid network subnet. 255.255.255.0
Primary Admin IP	The grid network IP address of the primary admin node. Leave as 0.0.0.0 to attempt automatic discovery after deployment. 0.0.0.0
Grid Node Parameters	1 setting
Node name	Optional name for this node. Must be a valid hostname and consist only of letters, numbers and hyphens. sg-sn-01

Back Next Finish Cancel

- Provide the IP of the primary Admin Node. Optionally, the primary admin can be automatically discovered. This step might be required if you chose to deploy the Admin Node on a separate subnet.
- Repeat steps a, b, and c for the remaining two StorageGRID Storage Nodes.

## Creating VMware VM-VM Anti-Affinity Rules

The StorageGRID Storage Nodes are placed on different hosts, and VMware VM-VM anti-affinity rules prevent single physical hosts from being deprived of resources.

- Browse to the cluster.
- Click Configure > Configuration > VM/Host Rules > Add.

**VMware vSphere Web Client**

NetApp-HCI-Cluster-01

Configure > Configuration > VM/Host Rules > Add

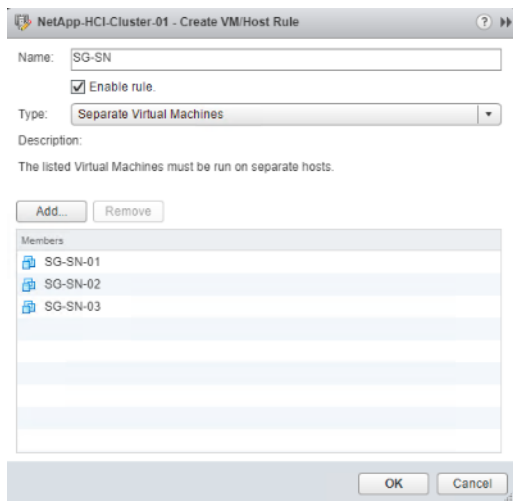
Name	Type	Enforced	Conflicts	Defined By
openshift-master	Separate Virtual Machines	Yes	0	User
openshift-infrastructure	Separate Virtual Machines	Yes	0	User

**VM/Host Rule Details**  
The listed 3 Virtual Machines must run on different hosts.

Rule Members	Conflicts
openshift-master01	0
openshift-master02	0
openshift-master03	0

- Create a rule to separate the StorageGRID Storage Node VMs.





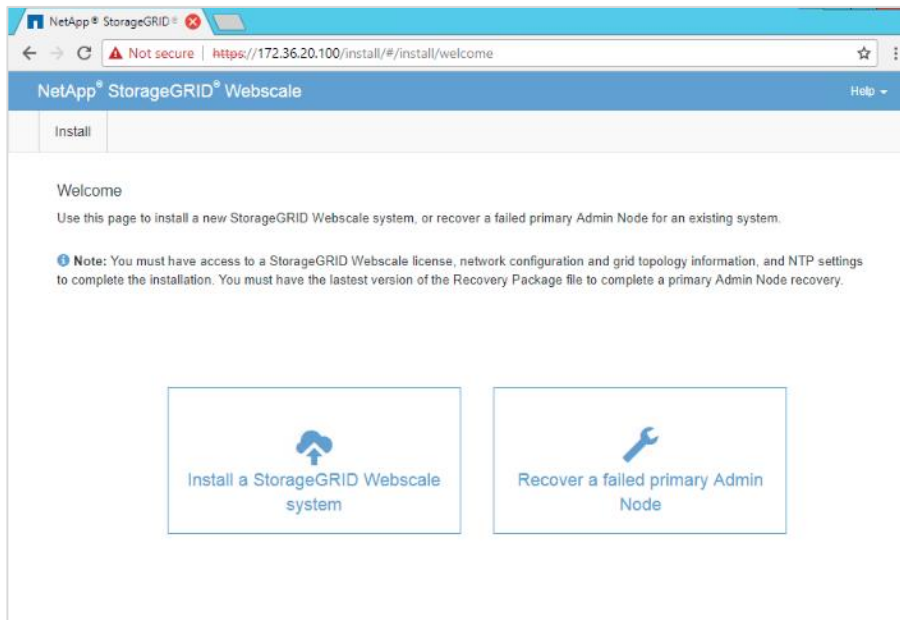
#### 4. Power on the StorageGRID Storage Nodes.

**Note:** Under certain circumstances, the live migration of Storage Nodes with vMotion can cause issues with NTP lock. Power off Storage Nodes before vMotion migration.

## Configuring the StorageGRID System and Completing the Installation

For detailed configuration instructions, see the [StorageGRID Documentation Center](#). This document only describes the modifications necessary for deploying a StorageGRID object store on NetApp HCI.

#### 1. Navigate to the primary admin URL in a web browser, and select Install a StorageGRID System.

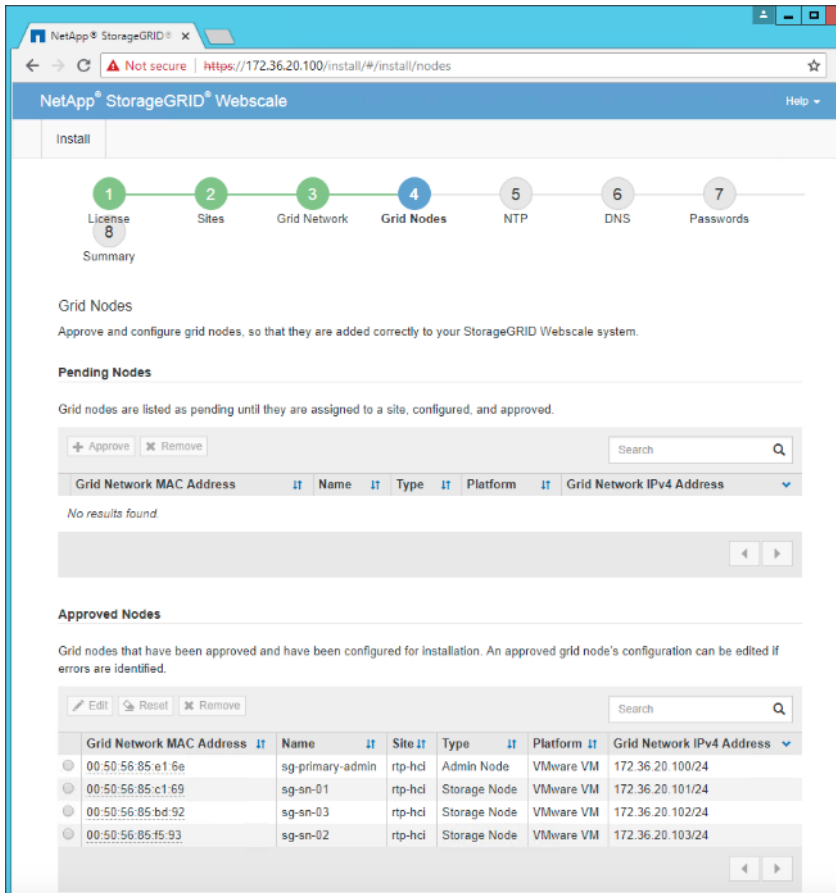


#### 2. Enter the grid name, license, and site name.

#### 3. Enter the <<Grid-subnet>> value.

The Storage Nodes are automatically discovered.

#### 4. Verify the details and approve the Storage and Admin Nodes.



5. Enter the <<ntp-server>> and <<dns-server>> values.
6. Configure the provisioning passwords and verify the information in the summary.
7. Download the recovery package and store it in a secure location.

## Configuring HAProxy to Load-Balance the StorageGRID Storage Nodes

HAProxy is used to load-balance the connections among the three StorageGRID Storage Nodes. StorageGRID also includes an optional load-balancer node that can be deployed as a VM or Docker container. For more information about the load-balancing options, see [TR-4626: StorageGRID Object Storage Load Balancer Options](#).

1. Deploy HAProxy on a Red Hat Enterprise Linux system.

```
yum install haproxy
chkconfig haproxy on
```

2. Make sure that appropriate rules are set for the firewall and SELinux.
3. Configure `/etc/haproxy/haproxy.cfg` to load-balance the StorageGRID Storage Nodes.

The following is an example of an HAProxy configuration. The highlighted sections are specific to StorageGRID.

```
# global parameters
global

    # Logging to syslog facility local0
    log /dev/log local0
# Proxy default configuration common for all frontend and backends
```

```

defaults

# passthrough any traffic via TCP
mode tcp

# apply log settings from the global section above to services
log global

# If sending a request to one server fails, try to send it to another, 3 times
retries 3

# Do not enforce session affinity (i.e., an HTTP session can be served by an
option redispatch

# Maximum number of simultaneous active connections from an upstream client
maxconn 5000

# Set the maximum time to wait for a connection attempt to a server to succeed
timeout connect 5s

# Set the maximum inactivity time on the client side.
timeout client 50s

# Set the maximum inactivity time on the server side.
timeout server 50s

# frontend specific configuration
frontend tcp-in

# bind to all network interfaces on port 443, restrict to specific IP if necessary
bind *:443

# for debugging purposes uncommenting the following option will enable basic
#option tcplog

# define a default backend
default_backend storagegrid

# backend specific configuration
backend storagegrid

# use all backup servers if primary servers are not available anymore
option allbackups

# balance connections using leastconn or roundrobin
balance leastconn

# define health check using HTTP OPTIONS call
option httpchk OPTIONS / HTTP/1.1

# declare backend servers
# check-ssl enables the health check using a SSL connection
# verify none disables certificate verifications, use
# verify required
# to enforce certificate verifications
server siteA-sn1 172.36.20.101:18082 check-ssl verify none
server siteA-sn2 172.36.20.102:18082 check-ssl verify none
server siteA-sn3 172.36.20.103:18082 check-ssl verify none

```

#### 4. Start the HAProxy service.

```
Systemctl start haproxy
```

**Note:** For this configuration, we chose to serve S3 on port 443. When deploying the default StorageGRID load balancer, S3 is served on port 8082.

## 3 StorageGRID Configuration

The configurations described in this section are specific to NetApp StorageGRID deployment on NetApp HCI. See the [StorageGRID administration documentation](#) for the configuration instructions to enable alerts through email, configuration of DNS, and SSL certificates.

### 3.1 ILM and Data Management Rules

StorageGRID uses powerful Information Lifecycle Management (ILM) data management rules to intelligently manage the placement of data. By default, the grid is configured with the Make 2 Copies rule. Each object copied into the grid is stored on separate StorageGRID Storage Nodes in two copies. Because the datastore being used by StorageGRID is deduplicated by NetApp Element, the space consumed is similar to a single copy.

StorageGRID also supports erasure coding for objects larger than 200KB. You can optionally configure a “2 plus 1” EC scheme in this configuration for greater efficiency and enable ILM rules to apply EC protection. Consult the [StorageGRID administration documentation](#) for details and additional options such as multisite protection.

### 3.2 Storage Configuration

#### Compression

NetApp Element software enables compression by default. NetApp recommends that you do not enable storage compression through StorageGRID.

#### Space Reservations

To keep StorageGRID from attempting to use more space than provided by the configured 3TB volume, set the storage values in the grid at a conservative reservation to stop taking new writes.

1. Select Configuration > Storage Options.
2. Click Configuration on the left, set the following values, and then click Apply Changes. The values are in bytes.
  - Storage Volume Read-Write Watermark: 200000000000
  - Storage Volume Soft Read-Only Watermark: 150000000000
  - Storage Volume Hard Read-Only Watermark: 100000000000
  - Metadata Reserved Space: 300000000000

With the custom settings in place, your storage watermarks appear as shown in Figure 5.

Figure 5) Storage watermarks.

#### Storage Watermarks

Description	Settings
Storage Volume Read-Write Watermark	200 GB
Storage Volume Soft Read-Only Watermark	150 GB
Storage Volume Hard Read-Only Watermark	100 GB
Metadata Reserved Space	300 GB


For definitions of storage volume watermark settings, see the [StorageGRID administration documentation](#).

### 3.3 Customizing Storage Alerts


#### Metadata Reserved Space

StorageGRID manages disk space for metadata and objects separately for greater flexibility and consistent performance.

To receive sufficient notice for adding more nodes or upgrading to a full-size StorageGRID instance, configure an alarm for 50% of available metadata reserved space. Due to the smaller allocation for metadata in this design, we selected a conservative setting that differs from the StorageGRID default of 3TB metadata space per StorageGRID Storage Node.

1. Click Configuration > Global Alarms.
2. Under Global Custom Alarms, click the  icon to create an alarm with the following values:
  - Service: DDS
  - Attribute: CDLP (Metadata Used Space (Percent))
  - Severity: Critical
  - Message: “Metadata space low - add more nodes”
  - Operator: >
  - Value: 50
3. Select Enabled and click Apply to save the alarm.











#### Object Storage Space

1. Click Configuration > Global Alarms.
2. Under Global Custom Alarms, click the  icon to create an alarm with the following values:
  - Service: LDR
  - Attribute: SAVP (Total Usable Space (Percent))
  - Severity: Major
  - Message: “Storage space low”
  - Operator: <
  - Value: 40

Your resulting custom alarms display as shown in Figure 7.

Figure 6) Global custom alarms.

Global Custom Alarms (2 Result(s))

Enabled	Service	Attribute	Severity	Message	Operator	Value	Additional Recipients	Actions
<input checked="" type="checkbox"/>	DDS	CDLP (Metadata Used Space (Percent))	 Critical	Metadata space low - add more nodes	>	50		   
<input checked="" type="checkbox"/>	LDR	SAVP (Total Usable Space (Percent))	 Major	Storage space low	<	40		   

### 3.4 Getting Ready to Serve Data

At this point, StorageGRID is ready to serve data. However, for compatibility with the widest array of applications and to eliminate common issues, follow the next steps.

## DNS Configuration

Configure your DNS with an alias for your service endpoint—for example, s3.company.com.

This alias will be associated with the IP configured on your HAProxy instance. You should also configure DNS to support a virtual host–style alias, in which the bucket name is part of the host name—for example, mybucket.s3.company.com.

You also need to configure an alias for wildcard queries: \*.s3.company.com.

## SSL Configuration

Install an SSL certificate configured for your DNS name and other Subject Alternative Names for virtual host support. For example:

CN = s3.company.com

DNS Name = \*.s3.company.com

NetApp recommends that you use a certificate from a trusted certificate authority. Many applications throw exceptions with untrusted certificates. The use of self-signed certificates is supported but might require more configuration on your application to enable trust.

For instructions on applying an SSL certificate and configuring support for multiple [S3 API endpoint domain names](#), see the [administration documentation](#).

## 3.5 Configuration Limit Testing

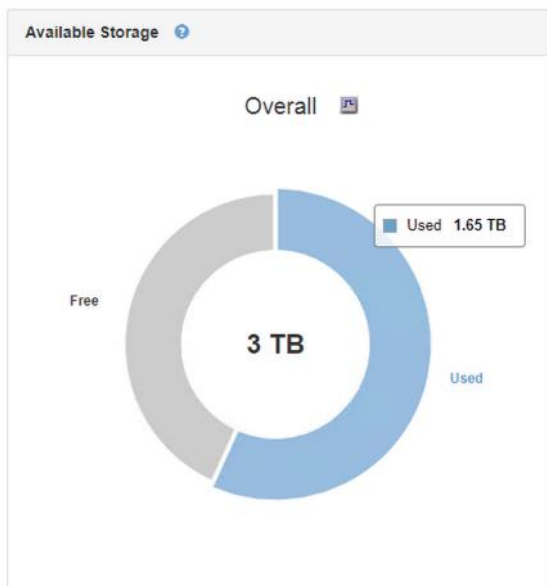
With the goal of creating a small footprint object store, we conducted the following tests:

- Consume 50% of metadata space.
- Consume object storage space to trigger space reservation limits.

As you approach the following limits, NetApp recommends expanding the grid:

- Object limits: 100 million
- Storage capacity: 1.65TB

Figure 7) StorageGRID available storage.



All testing was performed using S3tester: <https://github.com/s3tester/s3tester>.

### 3.6 Realized Storage Efficiency

In the NetApp Element management UI, examine the volume details for the StorageGRID datastore and view the Performance tab to see the used capacity. Depending on the storage efficiency achieved, you can provision a new datastore and add more VMDKs to the StorageGRID Storage Nodes.

## 4 Supporting Full-Scale Object Workloads

As your organization discovers more uses for object storage, the disk space and object count requirement might grow beyond the capacity of this configuration. To create a full-scale object storage system, you can add new nodes to StorageGRID and optionally decommission existing StorageGRID Storage Nodes. The NetApp HCI architecture is more than capable of supporting enterprise object storage workloads with hundreds of billions of objects. Object storage has traditionally focused on using lower-cost, ultra-dense storage. Therefore, you might choose to introduce other storage tiers such as NetApp E-Series systems or the NetApp StorageGRID appliance.

StorageGRID licensing is based on consumed storage, either on a perpetual or subscription basis. When you expand your grid to support more storage, you must reconcile the licensing. If the license capacity is exceeded, StorageGRID will continue serving data without limiting functionality. Licensing is required to enable Support.

### 4.1 Expanding for Object Space

If the system is approaching the predetermined limit of 1.65TB but the used metadata space remains low, you can simply add more disks to the StorageGRID Storage Nodes. Each Storage Node can have up to 16 disks dedicated to object storage. StorageGRID refers to these disks as RangeDB. The NetApp HCI installation has a single RangeDB volume, and you can add another 15 volumes by following the procedure outlined in the StorageGRID documentation, [VMware: Adding storage volumes to a Storage Node](#). After adding the disks, adjust the storage configuration settings as described in section 3.2. Note that the expansion of RangeDB volumes is not supported.

**Note:** With the storage efficiency savings realized, you can create additional datastores to house more VMDKs and attach them to the additional StorageGRID Storage Nodes.

### 4.2 Expanding for Metadata Space

After StorageGRID uses over 50% of its metadata space in this smaller configuration, NetApp recommends that you move to StorageGRID Storage Nodes that support the default object storage limit of 500 million per node.

Dedicated space for metadata is stored on the first RangeDB volume and, as stated previously, StorageGRID does not support the expansion of disks. To support an increased object count, you must add more StorageGRID Storage Nodes or move to full-sized StorageGRID Storage Nodes.

In the following example, three StorageGRID appliances are added and the existing VM StorageGRID Storage Nodes are decommissioned, retaining the Admin Node. The following steps are high level; for detailed instructions, see the [StorageGRID Expansion Guide](#).

1. Add three StorageGRID appliances with Administrative Domain Controller (ADC) service enabled.
2. Perform expansion steps through the StorageGRID UI.
3. After expansion is complete, decommission the three VM StorageGRID Storage Nodes as described in the [StorageGRID documentation](#). To maintain quorum, decommission the first two VM StorageGRID Storage Nodes and then the third.



With full-size Storage Nodes in place, the grid now supports 1.5 billion objects and the capacity of the appliances chosen. At this point, restore the storage configuration settings from section 3.2 to the defaults and remove the custom alarms. To continue to scale objects and capacity, continue adding StorageGRID Storage Nodes; each node adds support for 500 million objects.

## 5 Data Protection Considerations

Object storage is typically deployed such that all data stored in the system is considered protected, which is usually done by replicating a copy to a remote site. In this installation, all data is protected in the NetApp HCI system, but no remote copies are made. If site loss protection is required, you can expand this StorageGRID instance to more sites and change information lifecycle management (ILM) rules to distribute replicas or erasure code across many sites. For more details, see the [StorageGRID Expansion Guide](#).

Also, you can configure the StorageGRID CloudMirror service to replicate bucket contents to another StorageGRID instance or S3-compatible cloud. For instructions, see the [tenant administration documentation](#).

## Where to Find Additional Information

To learn more about the information described in this document, refer to the following documents and websites:

- NetApp HCI datasheet  
<https://www.netapp.com/us/media/ds-3881.pdf>
- NetApp HCI  
<https://www.netapp.com/us/products/converged-systems/hyper-converged-infrastructure.aspx>
- NetApp product documentation  
<http://docs.netapp.com>
- StorageGRID Documentation Center  
<http://docs.netapp.com/sgws-111/index.jsp>
- StorageGRID resources page  
<http://mysupport.netapp.com/storagegridwebscale/resources>

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 © 2018 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 item (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.