



Install StorageGRID nodes

StorageGRID software

NetApp
February 26, 2026

Table of Contents

- Install StorageGRID nodes 1
 - StorageGRID appliances 1
 - Install software-based nodes 1
 - Quick start for installing StorageGRID on a software-based node 1
 - Plan and prepare for installation on software-based nodes 2
 - Automate the software-based nodes installation 40
 - Deploy virtual grid nodes 56
 - Installation REST API 87
 - Troubleshoot installation issues 88
 - Script examples 90
- Configure grid and complete installation 94
 - Navigate to the Grid Manager 94
 - Specify the StorageGRID license information 95
 - Add sites 96
 - Specify Grid Network subnets 96
 - Approve pending grid nodes 97
 - Specify Network Time Protocol server information 101
 - Specify DNS server information 103
 - Specify the StorageGRID system passwords 103
 - Review your configuration and complete installation 105
 - Post-installation guidelines 106
- Where to go next 107
 - Required tasks 107
 - Optional tasks 107

Install StorageGRID nodes

StorageGRID appliances

Go to [StorageGRID Appliance Documentation](#) to learn how to install, configure, and maintain StorageGRID storage and services appliances.

Install software-based nodes

Quick start for installing StorageGRID on a software-based node

Follow these high-level steps to install a Linux or VMware StorageGRID node.



"Linux" refers to a RHEL, Ubuntu, or Debian deployment. For a list of supported versions, see the [NetApp Interoperability Matrix Tool \(IMT\)](#).

1

Preparation

- Learn about [StorageGRID architecture and network topology](#).
- Learn about the specifics of [StorageGRID networking](#).
- Gather and prepare the [Required information and materials](#).
- (VMware only) Install and configure [VMware vSphere Hypervisor, vCenter, and the ESX hosts](#).
- Prepare the required [CPU and RAM](#).
- Provide for [storage and performance requirements](#).
- (Linux only) [Prepare the Linux servers](#) that will host your StorageGRID nodes.

2

Deployment

Deploy grid nodes. When you deploy grid nodes, they are created as part of the StorageGRID system and connected to one or more networks.

- (Linux only) To deploy software-based grid nodes on the hosts you prepared in step 1, use the Linux command line and [node configuration files](#).
- (VMware only) Use the VMware vSphere Web Client, a .vmdk file, and a set of .ovf file templates to [deploy the software-based nodes as virtual machines \(VMs\)](#) on the servers you prepared in step 1.
- To deploy StorageGRID appliance nodes, follow the [Quick start for hardware installation](#).

3

Configuration

When all nodes have been deployed, use the Grid Manager to [configure the grid and complete the installation](#).

Automate the installation



"Linux" refers to a RHEL, Ubuntu, or Debian deployment. For a list of supported versions, see the [NetApp Interoperability Matrix Tool \(IMT\)](#).

Linux

To save time and provide consistency, you can automate the installation of the StorageGRID host service and the configuration of grid nodes.

- Use a standard orchestration framework such as Ansible, Puppet, or Chef to automate:
 - Installation of Linux
 - Configuration of networking and storage
 - Installation of the container engine and the StorageGRID host service
 - Deployment of virtual grid nodes

See [Automate the installation and configuration of the StorageGRID host service](#).

- After you deploy grid nodes, [automate the configuration of the StorageGRID system](#) using the Python configuration script provided in the installation archive.
- [Automate the installation and configuration of appliance grid nodes](#)
- If you are an advanced developer of StorageGRID deployments, automate the installation of grid nodes by using the [installation REST API](#).

VMware

To save time and provide consistency, you can automate the deployment and configuration of grid nodes and the configuration of the StorageGRID system.

- [Automate grid node deployment using VMware vSphere](#).
- After you deploy grid nodes, [automate the configuration of the StorageGRID system](#) using the Python configuration script provided in the installation archive.
- [Automate the installation and configuration of appliance grid nodes](#)
- If you are an advanced developer of StorageGRID deployments, automate the installation of grid nodes by using the [installation REST API](#).

Plan and prepare for installation on software-based nodes

Required information and materials

Before you install StorageGRID, gather and prepare the required information and materials.

Required information

Network plan

Which networks you intend to attach to each StorageGRID node. StorageGRID supports multiple networks for traffic separation, security, and administrative convenience.

See the StorageGRID [Networking guidelines](#).

Network information

IP addresses to assign to each grid node and the IP addresses of the DNS and NTP servers.

Servers for grid nodes

Identify a set of servers (physical, virtual, or both) that, in aggregate, provide sufficient resources to support the number and type of StorageGRID nodes you plan to deploy.



If your StorageGRID installation will not use StorageGRID appliance (hardware) Storage Nodes, you must use hardware RAID storage with battery-backed write cache (BBWC). StorageGRID does not support the use of virtual storage area networks (vSANs), software RAID, or no RAID protection.

Node migration (Ubuntu and Debian only, if needed)

Understand the [requirements for node migration](#), if you want to perform scheduled maintenance on physical hosts without any service interruption.

Related information

[NetApp Interoperability Matrix Tool](#)

Required materials

NetApp StorageGRID license

You must have a valid, digitally signed NetApp license.



A non-production license, which can be used for testing and proof of concept grids, is included in the StorageGRID installation archive.

StorageGRID installation archive

[Download the StorageGRID installation archive and extract the files.](#)

Service laptop

The StorageGRID system is installed through a service laptop.

The service laptop must have:

- Network port
- SSH client (for example, PuTTY)
- [Supported web browser](#)

StorageGRID documentation

- [Release notes](#)
- [Instructions for administering StorageGRID](#)

Download and extract the StorageGRID installation files

You must download the StorageGRID installation archive and extract the required files. Optionally, you can manually verify the files in the installation package.

Steps

1. Go to the [NetApp Downloads page for StorageGRID](#).

2. Select the button for downloading the latest release, or select another version from the drop-down menu and select **Go**.
3. Sign in with the username and password for your NetApp account.
4. If a Caution/MustRead statement appears, read it and select the checkbox.



You must apply any required hotfixes after you install the StorageGRID release. For more information, see the [hotfix procedure in the recovery and maintenance instructions](#)

5. Read the End User License Agreement, select the checkbox, and then select **Accept & Continue**.
6. In the **Install StorageGRID** column, select the .tgz or .zip installation archive for your software-based node type: RHEL, Ubuntu or Debian, or VMware.



Use the .zip file if you are running Windows on the service laptop.

7. Save the installation archive.
8. Code signature verification is manual on a Linux node. Optionally, if you need to verify the installation archive:
 - a. Download the StorageGRID code signature verification package. The file name for this package uses the format `StorageGRID_<version-number>_Code_Signature_Verification_Package.tar.gz`, where `<version-number>` is the StorageGRID software version.
 - b. Follow the steps to [manually verify the installation files](#).
9. Extract the files from the installation archive.
10. Choose the files you need.

The files you need depends on your planned grid topology and how you will deploy your StorageGRID system.



The paths listed in the table are relative to the top-level directory installed by the extracted installation archive.

RHEL

Path and file name	Description
<code>./rpms/README</code>	A text file that describes all of the files contained in the StorageGRID download file.
<code>./rpms/NLF000000.txt</code>	A free license that does not provide any support entitlement for the product.
<code>./rpms/StorageGRID-Webscale-Images-version-SHA.rpm</code>	RPM package for installing the StorageGRID node images on your RHEL hosts.
<code>./rpms/StorageGRID-Webscale-Service-version-SHA.rpm</code>	RPM package for installing the StorageGRID host service on your RHEL hosts.
Deployment scripting tool	Description
<code>./rpms/configure-storagegrid.py</code>	A Python script used to automate the configuration of a StorageGRID system.
<code>./rpms/configure-sga.py</code>	A Python script used to automate the configuration of StorageGRID appliances.
<code>./rpms/configure-storagegrid.sample.json</code>	An example configuration file for use with the <code>configure-storagegrid.py</code> script.
<code>./rpms/storagegrid-ssoauth.py</code>	An example Python script that you can use to sign in to the Grid Management API when single sign-on is enabled. You can also use this script for Ping Federate integration.
<code>./rpms/configure-storagegrid.blank.json</code>	A blank configuration file for use with the <code>configure-storagegrid.py</code> script.
<code>./rpms/extras/ansible</code>	Example Ansible role and playbook for configuring RHEL hosts for StorageGRID container deployment. You can customize the role or playbook as necessary.
<code>./rpms/storagegrid-ssoauth-azure.py</code>	An example Python script that you can use to sign in to the Grid Management API when single sign-on (SSO) is enabled using Active Directory or Ping Federate.
<code>./rpms/storagegrid-ssoauth-azure.js</code>	A helper script called by the companion <code>storagegrid-ssoauth-azure.py</code> Python script to perform SSO interactions with Azure.

Path and file name	Description
<code>./rpms/extras/api-schemas</code>	API schemas for StorageGRID. Note: Before you perform an upgrade, you can use these schemas to confirm that any code you have written to use StorageGRID management APIs will be compatible with the new StorageGRID release if you don't have a non-production StorageGRID environment for upgrade compatibility testing.

Ubuntu or Debian

Path and file name	Description
<code>./debs/README</code>	A text file that describes all of the files contained in the StorageGRID download file.
<code>./debs/NLF000000.txt</code>	A non-production NetApp License File that you can use for testing and proof of concept deployments.
<code>./debs/storagegrid-webscale-images-version-SHA.deb</code>	DEB package for installing the StorageGRID node images on Ubuntu or Debian hosts.
<code>./debs/storagegrid-webscale-images-version-SHA.deb.md5</code>	MD5 checksum for the file <code>/debs/storagegrid-webscale-images-version-SHA.deb</code> .
<code>./debs/storagegrid-webscale-service-version-SHA.deb</code>	DEB package for installing the StorageGRID host service on Ubuntu or Debian hosts.

Deployment scripting tool	Description
<code>./debs/configure-storagegrid.py</code>	A Python script used to automate the configuration of a StorageGRID system.
<code>./debs/configure-sga.py</code>	A Python script used to automate the configuration of StorageGRID appliances.
<code>./debs/storagegrid-ssoauth.py</code>	An example Python script that you can use to sign in to the Grid Management API when single sign-on is enabled. You can also use this script for Ping Federate integration.
<code>./debs/configure-storagegrid.sample.json</code>	An example configuration file for use with the <code>configure-storagegrid.py</code> script.
<code>./debs/configure-storagegrid.blank.json</code>	A blank configuration file for use with the <code>configure-storagegrid.py</code> script.

Path and file name	Description
<code>./debs/extras/ansible</code>	Example Ansible role and playbook for configuring Ubuntu or Debian hosts for StorageGRID container deployment. You can customize the role or playbook as necessary.
<code>./debs/storagegrid-ssoauth-azure.py</code>	An example Python script that you can use to sign in to the Grid Management API when single sign-on (SSO) is enabled using Active Directory or Ping Federate.
<code>./debs/storagegrid-ssoauth-azure.js</code>	A helper script called by the companion <code>storagegrid-ssoauth-azure.py</code> Python script to perform SSO interactions with Azure.
<code>./debs/extras/api-schemas</code>	API schemas for StorageGRID. Note: Before you perform an upgrade, you can use these schemas to confirm that any code you have written to use StorageGRID management APIs will be compatible with the new StorageGRID release if you don't have a non-production StorageGRID environment for upgrade compatibility testing.

VMware

Path and file name	Description
<code>./vsphere/README</code>	A text file that describes all of the files contained in the StorageGRID download file.
<code>./vsphere/NLF000000.txt</code>	A free license that does not provide any support entitlement for the product.
<code>./vsphere/NetApp-SG-version-SHA.vmdk</code>	The virtual machine disk file that is used as a template for creating grid node virtual machines.
<code>./vsphere/vsphere-primary-admin.ovf</code> <code>./vsphere/vsphere-primary-admin.mf</code>	The Open Virtualization Format template file (<code>.ovf</code>) and manifest file (<code>.mf</code>) for deploying the primary Admin Node.
<code>./vsphere/vsphere-non-primary-admin.ovf</code> <code>./vsphere/vsphere-non-primary-admin.mf</code>	The template file (<code>.ovf</code>) and manifest file (<code>.mf</code>) for deploying non-primary Admin Nodes.
<code>./vsphere/vsphere-gateway.ovf</code> <code>./vsphere/vsphere-gateway.mf</code>	The template file (<code>.ovf</code>) and manifest file (<code>.mf</code>) for deploying Gateway Nodes.

Path and file name	Description
./vsphere/vsphere-storage.ovf ./vsphere/vsphere-storage.mf	The template file (.ovf) and manifest file (.mf) for deploying virtual machine-based Storage Nodes.
Deployment scripting tool	Description
./vsphere/deploy-vsphere-ovftool.sh	A Bash shell script used to automate the deployment of virtual grid nodes.
./vsphere/deploy-vsphere-ovftool-sample.ini	An example configuration file for use with the deploy-vsphere-ovftool.sh script.
./vsphere/configure-storagegrid.py	A Python script used to automate the configuration of a StorageGRID system.
./vsphere/configure-sga.py	A Python script used to automate the configuration of StorageGRID appliances.
./vsphere/storagegrid-ssoauth.py	An example Python script that you can use to sign in to the Grid Management API when single sign-on (SSO) is enabled. You can also use this script for Ping Federate integration.
./vsphere/configure-storagegrid.sample.json	An example configuration file for use with the configure-storagegrid.py script.
./vsphere/configure-storagegrid.blank.json	A blank configuration file for use with the configure-storagegrid.py script.
./vsphere/storagegrid-ssoauth-azure.py	An example Python script that you can use to sign in to the Grid Management API when single sign-on (SSO) is enabled using Active Directory or Ping Federate.
./vsphere/storagegrid-ssoauth-azure.js	A helper script called by the companion storagegrid-ssoauth-azure.py Python script to perform SSO interactions with Azure.
./vsphere/extras/api-schemas	API schemas for StorageGRID. Note: Before you perform an upgrade, you can use these schemas to confirm that any code you have written to use StorageGRID management APIs will be compatible with the new StorageGRID release if you don't have a non-production StorageGRID environment for upgrade compatibility testing.

Manually verify installation files (optional)

If necessary, you can manually verify the files in the StorageGRID installation archive.

Before you begin

You have [downloaded the verification package](#) from the [NetApp Downloads page for StorageGRID](#).

Steps

1. Extract the artifacts from the verification package:

```
tar -xf StorageGRID_12.0.0_Code_Signature_Verification_Package.tar.gz
```

2. Ensure that these artifacts were extracted:

- Leaf certificate: Leaf-Cert.pem
- Certificate chain: CA-Int-Cert.pem
- Time stamp response chain: TS-Cert.pem
- Checksum file: sha256sum
- Checksum signature: sha256sum.sig
- Time stamp response file: sha256sum.sig.tsr

3. Use the chain to verify the leaf certificate is valid.

Example: `openssl verify -CAfile CA-Int-Cert.pem Leaf-Cert.pem`

Expected output: Leaf-Cert.pem: OK

4. If step 2 failed because of an expired leaf certificate, use the `tsr` file to verify.

Example: `openssl ts -CAfile CA-Int-Cert.pem -untrusted TS-Cert.pem -verify -data sha256sum.sig -in sha256sum.sig.tsr`

Expected output includes: Verification: OK

5. Create a public key file from the leaf certificate.

Example: `openssl x509 -pubkey -noout -in Leaf-Cert.pem > Leaf-Cert.pub`

Expected output: *none*

6. Use the public key to verify the `sha256sum` file against `sha256sum.sig`.

Example: `openssl dgst -sha256 -verify Leaf-Cert.pub -signature sha256sum.sig sha256sum`

Expected output: Verified OK

7. Verify the `sha256sum` file content against newly created checksums.

Example: `sha256sum -c sha256sum`

Expected output: *<filename>*: OK

<filename> is the name of the archive file you downloaded.

8. [Complete the remaining steps](#) to extract and choose the appropriate files from the installation archive.

Software requirements

You can use a virtual machine to host any type of StorageGRID node. You need one virtual machine for each grid node.

RHEL

To install StorageGRID on RHEL, you must install some third-party software packages. Some supported Linux distributions don't contain these packages by default. The software package versions that StorageGRID installations are tested on include those listed on this page.

If you select a Linux distribution and container runtime installation option that requires any of these packages, and they are not installed automatically by the Linux distribution, install one of the versions listed here if available from your provider or the supporting vendor for your Linux distribution. Otherwise, use the default package versions available from your vendor.

All installation options require either Podman or Docker. Do not install both packages. Install only the package required by your installation option.



Support for Docker as the container engine for software-only deployments is deprecated. Docker will be replaced with another container engine in a future release.

Python versions tested

- 3.5.2-2
- 3.6.8-2
- 3.6.8-38
- 3.6.9-1
- 3.7.3-1
- 3.8.10-0
- 3.9.2-1
- 3.9.10-2
- 3.9.16-1
- 3.10.6-1
- 3.11.2-6

Podman versions tested

- 3.2.3-0
- 3.4.4+ds1
- 4.1.1-7
- 4.2.0-11
- 4.3.1+ds1-8+b1
- 4.4.1-8
- 4.4.1-12

Docker versions tested



Docker support is deprecated and will be removed in a future release.

- Docker-CE 20.10.7

- Docker-CE 20.10.20-3
- Docker-CE 23.0.6-1
- Docker-CE 24.0.2-1
- Docker-CE 24.0.4-1
- Docker-CE 24.0.5-1
- Docker-CE 24.0.7-1
- 1.5-2

Ubuntu and Debian

To install StorageGRID on Ubuntu or Debian, you must install some third-party software packages. Some supported Linux distributions don't contain these packages by default. The software package versions that StorageGRID installations are tested on include those listed on this page.

If you select a Linux distribution and container runtime installation option that requires any of these packages, and they are not installed automatically by the Linux distribution, install one of the versions listed here if available from your provider or the supporting vendor for your Linux distribution. Otherwise, use the default package versions available from your vendor.

All installation options require either Podman or Docker. Do not install both packages. Install only the package required by your installation option.



Support for Docker as the container engine for software-only deployments is deprecated. Docker will be replaced with another container engine in a future release.

Python versions tested

- 3.5.2-2
- 3.6.8-2
- 3.6.8-38
- 3.6.9-1
- 3.7.3-1
- 3.8.10-0
- 3.9.2-1
- 3.9.10-2
- 3.9.16-1
- 3.10.6-1
- 3.11.2-6

Podman versions tested

- 3.2.3-0
- 3.4.4+ds1
- 4.1.1-7
- 4.2.0-11

- 4.3.1+ds1-8+b1
- 4.4.1-8
- 4.4.1-12

Docker versions tested



Docker support is deprecated and will be removed in a future release.

- Docker-CE 20.10.7
- Docker-CE 20.10.20-3
- Docker-CE 23.0.6-1
- Docker-CE 24.0.2-1
- Docker-CE 24.0.4-1
- Docker-CE 24.0.5-1
- Docker-CE 24.0.7-1
- 1.5-2

VMware

VMware vSphere Hypervisor

You must install VMware vSphere Hypervisor on a prepared physical server. The hardware must be configured correctly (including firmware versions and BIOS settings) before you install VMware software.

- Configure networking in the hypervisor as required to support networking for the StorageGRID system you are installing.

[Networking guidelines](#)

- Ensure that the datastore is large enough for the virtual machines and virtual disks that are required to host the grid nodes.
- If you create more than one datastore, name each datastore so that you can easily identify which datastore to use for each grid node when you create virtual machines.

ESX host configuration requirements



You must properly configure the network time protocol (NTP) on each ESX host. If the host time is incorrect, negative effects, including data loss, could occur.

VMware configuration requirements

You must install and configure VMware vSphere and vCenter before deploying StorageGRID nodes.

For supported versions of VMware vSphere Hypervisor and VMware vCenter Server software, see the [NetApp Interoperability Matrix Tool](#).

For the steps required to install these VMware products, see the VMware documentation.

CPU and RAM requirements

Before installing StorageGRID software, verify and configure the hardware so that it is ready to support the StorageGRID system.

Each StorageGRID node requires the following minimum resources:

- CPU cores: 8 per node
- RAM: Dependent on the total RAM available and the amount of non-StorageGRID software running on the system
 - Generally, at least 24 GB per node, and 2 to 16 GB less than the total system RAM
 - A minimum of 64 GB for each tenant that will have approximately 5,000 buckets

Software-based metadata-only node resources must match the existing Storage Nodes resources. For example:

- If the existing StorageGRID site is using SG6000 or SG6100 appliances, the software-based metadata-only nodes must meet the following minimum requirements:
 - 128 GB RAM
 - 8 core CPU
 - 8 TB SSD or equivalent storage for the Cassandra database (rangedb/0)
- If the existing StorageGRID site is using virtual Storage Nodes with 24 GB RAM, 8 core CPU, and 3 TB or 4TB of metadata storage, the software-based metadata-only nodes should use similar resources (24 GB RAM, 8 core CPU, and 4TB of metadata storage (rangedb/0)).

When adding a new StorageGRID site, the new site total metadata capacity should, at minimum, match existing StorageGRID sites and new site resources should match the Storage Nodes at existing StorageGRID sites.



"Linux" refers to a RHEL, Ubuntu, or Debian deployment. For a list of supported versions, see the [NetApp Interoperability Matrix Tool \(IMT\)](#).

Linux

Ensure that the number of StorageGRID nodes you plan to run on each physical or virtual host does not exceed the number of CPU cores or the physical RAM available. If the hosts aren't dedicated to running StorageGRID (not recommended), be sure to consider the resource requirements of the other applications.

VMware

VMware supports one node per virtual machine. Ensure that the StorageGRID node does not exceed the physical RAM available. Each virtual machine must be dedicated to running StorageGRID.



Monitor your CPU and memory usage regularly to ensure that these resources continue to accommodate your workload. For example, doubling the RAM and CPU allocation for virtual Storage Nodes would provide similar resources to those provided for StorageGRID appliance nodes. Additionally, if the amount of metadata per node exceeds 500 GB, consider increasing the RAM per node to 48 GB or more. For information about managing object metadata storage, increasing the Metadata Reserved Space setting, and monitoring CPU and memory usage, see the instructions for [administering](#), [monitoring](#), and [upgrading](#) StorageGRID.

If hyperthreading is enabled on the underlying physical hosts, you can provide 8 virtual cores (4 physical cores) per node. If hyperthreading is not enabled on the underlying physical hosts, you must provide 8 physical cores per node.

If you are using virtual machines as hosts and have control over the size and number of VMs, you should use a single VM for each StorageGRID node and size the VM accordingly.

(RHEL, Debian, and Ubuntu only) For production deployments, you should not run multiple Storage Nodes on the same physical storage hardware or virtual host. Each Storage Node in a single StorageGRID deployment should be in its own isolated failure domain. You can maximize the durability and availability of object data if you ensure that a single hardware failure can only impact a single Storage Node.

See also [Storage and performance requirements](#).

Storage and performance requirements

You must understand the storage requirements for StorageGRID nodes, so you can provide enough space to support the initial configuration and future storage expansion.

Storage and performance requirements vary based on your software-based node implementation.



"Linux" refers to a RHEL, Ubuntu, or Debian deployment. For a list of supported versions, see the [NetApp Interoperability Matrix Tool \(IMT\)](#).

Storage categories

StorageGRID nodes require three logical categories of storage:

- **Container pool** — Performance-tier (10K SAS or SSD) storage for the node containers, which will be assigned to the container engine storage driver when you install and configure the container engine on the hosts that will support your StorageGRID nodes.
- **System data** — Performance-tier (10K SAS or SSD) storage for per-node persistent storage of system data and transaction logs, which the StorageGRID host services will consume and map into individual nodes.
- **Object data** — Performance-tier (10K SAS or SSD) storage and capacity-tier (NL-SAS/SATA) bulk storage for the persistent storage of object data and object metadata.

You must use RAID-backed block devices for all storage categories. Non-redundant disks, SSDs, or JBODs aren't supported. You can use shared or local RAID storage for any of the storage categories; however, if you want to use the node migration capability in StorageGRID, you must store both system data and object data on shared storage. For more information, see [Node container migration requirements](#).

Performance requirements

The performance of the volumes used for the container pool, system data, and object metadata significantly impacts the overall performance of the system. You should use performance-tier (10K SAS or SSD) storage for these volumes to ensure adequate disk performance in terms of latency, input/output operations per second (IOPS), and throughput. You can use capacity-tier (NL-SAS/SATA) storage for the persistent storage of object data.

The volumes used for the container pool, system data, and object data must have write-back caching enabled. The cache must be on a protected or persistent media.

Requirements for hosts that use NetApp ONTAP storage

If the StorageGRID node uses storage assigned from a NetApp ONTAP system, confirm that the volume does not have a FabricPool tiering policy enabled. Disabling FabricPool tiering for volumes used with StorageGRID nodes simplifies troubleshooting and storage operations.



Never use FabricPool to tier any data related to StorageGRID back to StorageGRID itself. Tiering StorageGRID data back to StorageGRID increases troubleshooting and operational complexity.

Number of hosts required

Each StorageGRID site requires a minimum of three Storage Nodes.



In a production deployment, don't run more than one Storage Node on a single physical or virtual host. Using a dedicated host for each Storage Node provides an isolated failure domain.

Other types of nodes, such as Admin Nodes or Gateway Nodes, can be deployed on the same hosts, or they can be deployed on their own dedicated hosts as required.



Disk snapshots can't be used to restore grid nodes. Instead, refer to the [grid node recovery](#) procedures for each type of node.

Number of storage volumes for each node

The following table shows the number of storage volumes (LUNs) required for each host and the minimum size required for each LUN, based on which nodes will be deployed on that host.

The maximum tested LUN size is 39 TB.



These numbers are for each host, not for the entire grid.

LUN purpose	Storage category	Number of LUNs	Minimum size/LUN
Container engine storage pool	Container pool	1	Total number of nodes × 100 GB
/var/local volume	System data	1 for each node on this host	100 GB

LUN purpose	Storage category	Number of LUNs	Minimum size/LUN
Storage Node	Object data	3 for each Storage Node on this host Note: A Linux software-based Storage Node can have 1 to 48 storage volumes. A VMware software-based Storage Node can have 1 to 16 storage volumes. At least 3 storage volumes are recommended.	12 TB (4 TB/LUN, minimum) Maximum tested LUN size: 39 TB. See Storage requirements for Storage Nodes for more information.
Storage Node (metadata-only)	Object metadata	1	4 TB/LUN, minimum Maximum tested LUN size: 39 TB. See Storage requirements for Storage Nodes for more information. Note: Only one rangedb is required for metadata-only Storage Nodes.
Admin Node audit logs	System data	1 for each Admin Node on this host	200 GB
Admin Node tables	System data	1 for each Admin Node on this host	200 GB



Depending on the audit level configured, the size of user inputs such as S3 object key name, and how much audit log data you need to preserve, you might need to increase the size of the audit log LUN on each Admin Node. Generally, a grid generates approximately 1 KB of audit data per S3 operation, which would mean that a 200 GB LUN would support 70 million operations per day or 800 operations per second for two to three days.

Minimum storage space for a host

The following table shows the minimum storage space required for each type of node. You can use this table to determine the minimum amount of storage you must provide to the host in each storage category, based on which nodes will be deployed on that host.



Disk snapshots can't be used to restore grid nodes. Instead, refer to the [grid node recovery](#) procedures for each type of node.

Each node host requires a 100 GB LUN for the OS.

Type of node	Container pool	System data	Object data
Storage Node	100 GB	100 GB	4,000 GB
Admin Node	100 GB	500 GB (3 LUNs)	<i>not applicable</i>
Gateway Node	100 GB	100 GB	<i>not applicable</i>

Example: Calculating the storage requirements for a host or virtual machine

Suppose you plan to deploy three nodes on the same host or virtual machine: one Storage Node, one Admin Node, and one Gateway Node. You should provide a minimum of nine storage volumes to the host. You will need a minimum of 300 GB of performance-tier storage for the node containers, 700 GB of performance-tier storage for system data and transaction logs, and 12 TB of capacity-tier storage for object data.

Linux host example

Type of node	LUN purpose	Number of LUNs	LUN size
Storage Node	Container engine storage pool	1	300 GB (100 GB/node)
Storage Node	<code>/var/local</code> volume	1	100 GB
Storage Node	Object data	3	12 TB (4 TB/LUN)
Admin Node	<code>/var/local</code> volume	1	100 GB
Admin Node	Admin Node audit logs	1	200 GB
Admin Node	Admin Node tables	1	200 GB
Gateway Node	<code>/var/local</code> volume	1	100 GB
Total		9	Container pool: 300 GB System data: 700 GB Object data: 12,000 GB

VMware virtual machine example

Type of node	LUN purpose	Number of LUNs	LUN size
Storage Node	OS volume	1	100 GB
Storage Node	Object data	3	12 TB (4 TB/LUN)
Admin Node	OS volume	1	100 GB
Admin Node	Admin Node audit logs	1	200 GB
Admin Node	Admin Node tables	1	200 GB
Gateway Node	OS volume	1	100 GB
Total		8	System data: 700 GB Object data: 12,000 GB

Specific storage requirements for Storage Nodes


Linux and VMware have different storage requirements for storage nodes:

- A Linux software-based Storage Node can have 1 to 48 storage volumes
- A VMware software-based Storage Node can have 1 to 16 storage volumes
- Three or more storage volumes are recommended.
- Each storage volume should be 4 TB or larger.



An appliance Storage Node can also have up to 48 storage volumes.

As shown in the figure, StorageGRID reserves space for object metadata on storage volume 0 of each Storage Node. Any remaining space on storage volume 0 and any other storage volumes in the Storage Node are used exclusively for object data.



To provide redundancy and to protect object metadata from loss, StorageGRID stores three copies of the metadata for all objects in the system at each site. The three copies of object metadata are evenly distributed across all Storage Nodes at each site.

When installing a grid with metadata-only Storage Nodes, the grid must also contain a minimum number of nodes for object storage. See [Types of Storage Nodes](#) for more information about metadata-only Storage Nodes.

- For a single-site grid, at least two Storage Nodes are configured for objects and metadata.
- For a multi-site grid, at least one Storage Node per site are configured for objects and metadata.

When you assign space to volume 0 of a new Storage Node, you must ensure there is adequate space for that node's portion of all object metadata.

- At a minimum, you must assign at least 4 TB to volume 0.



If you use only one storage volume for a Storage Node and you assign 4 TB or less to the volume, the Storage Node might enter the storage read-only state on startup and store object metadata only.



If you assign less than 500 GB to volume 0 (non-production use only), 10% of the storage volume's capacity is reserved for metadata.

- Software-based metadata-only node resources must match the existing Storage Nodes resources. For example:
 - If the existing StorageGRID site is using SG6000 or SG6100 appliances, the software-based metadata-only nodes must meet the following minimum requirements:
 - 128 GB RAM
 - 8 core CPU
 - 8 TB SSD or equivalent storage for the Cassandra database (rangedb/0)
 - If the existing StorageGRID site is using virtual Storage Nodes with 24 GB RAM, 8 core CPU, and 3 TB or 4TB of metadata storage, the software-based metadata-only nodes should use similar resources (24 GB RAM, 8 core CPU, and 4TB of metadata storage (rangedb/0)).

When adding a new StorageGRID site, the new site total metadata capacity should, at minimum, match existing StorageGRID sites and new site resources should match the Storage Nodes at existing StorageGRID sites.

- If you are installing a new system (StorageGRID 11.6 or higher) and each Storage Node has 128 GB or more of RAM, assign 8 TB or more to volume 0. Using a larger value for volume 0 can increase the space allowed for metadata on each Storage Node.
- When configuring different Storage Nodes for a site, use the same setting for volume 0 if possible. If a site contains Storage Nodes of different sizes, the Storage Node with the smallest volume 0 will determine the metadata capacity of that site.

For details, go to [Manage object metadata storage](#).

Node container migration requirements (Linux)

The node migration feature allows you to manually move a node from one host to another. Typically, both hosts are in the same physical data center.



"Linux" refers to a RHEL, Ubuntu, or Debian deployment. For a list of supported versions, see the [NetApp Interoperability Matrix Tool \(IMT\)](#).

Node migration allows you to perform physical host maintenance without disrupting grid operations. You move all StorageGRID nodes, one at a time, to another host before taking the physical host offline. Migrating nodes requires only a short downtime for each node and should not affect operation or availability of grid services.

If you want to use the StorageGRID node migration feature, your deployment must meet additional requirements:

- Consistent network interface names across hosts in a single physical data center
- Shared storage for StorageGRID metadata and object repository volumes that is accessible by all hosts in a single physical data center. For example, you might use NetApp E-Series storage arrays.

If you are using virtual hosts and the underlying hypervisor layer supports VM migration, you might want to use this capability instead of the node migration feature in StorageGRID. In this case, you can ignore these additional requirements.

Before performing migration or hypervisor maintenance, shut down the nodes gracefully. See the instructions for [shutting down a grid node](#).

VMware Live Migration not supported

When performing bare-metal installation on VMware VMs, OpenStack Live Migration and VMware live vMotion cause the virtual machine clock time to jump and aren't supported for grid nodes of any type. Though rare, incorrect clock times can result in loss of data or configuration updates.

Cold migration is supported. In cold migration, you shut down the StorageGRID nodes before migrating them between hosts. See the instructions for [shutting down a grid node](#).

Consistent network interface names

To move a node from one host to another, the StorageGRID host service needs to have some confidence that the external network connectivity the node has at its current location can be duplicated at the new location. It gets this confidence through the use of consistent network interface names in the hosts.

Suppose, for example, that StorageGRID NodeA running on Host1 has been configured with the following interface mappings:

`eth0` → `bond0.1001`

`eth1` → `bond0.1002`

`eth2` → `bond0.1003`

The lefthand side of the arrows corresponds to the traditional interfaces as viewed from within a StorageGRID container (that is, the Grid, Admin, and Client Network interfaces, respectively). The righthand side of the arrows corresponds to the actual host interfaces providing these networks, which are three VLAN interfaces subordinate to the same physical interface bond.

Now, suppose you want to migrate NodeA to Host2. If Host2 also has interfaces named `bond0.1001`, `bond0.1002`, and `bond0.1003`, the system will allow the move, assuming that the like-named interfaces will provide the same connectivity on Host2 as they do on Host1. If Host2 does not have interfaces with the same names, the move will not be allowed.

There are many ways to achieve consistent network interface naming across multiple hosts; see [Configure the host network](#) for some examples.

Shared storage

To achieve rapid, low-overhead node migrations, the StorageGRID node migration feature does not physically move node data. Instead, node migration is performed as a pair of export and import operations, as follows:

- During the "node export" operation, a small amount of persistent state data is extracted from the node container running on HostA and cached on that node's system data volume. Then, the node container on HostA is deinstantiated.

- During the "node import" operation, the node container on HostB that uses the same network interface and block storage mappings that were in effect on HostA is instantiated. Then, the cached persistent state data is inserted into the new instance.

Given this mode of operation, all of the node's system data and object storage volumes must be accessible from both HostA and HostB for the migration to be allowed, and to work. In addition, they must have been mapped into the node using names that are guaranteed to refer to the same LUNs on HostA and HostB.

The following example shows one solution for block device mapping for a StorageGRID Storage Node, where DM multipathing is in use on the hosts, and the alias field has been used in `/etc/multipath.conf` to provide consistent, friendly block device names available on all hosts.

```
/var/local → /dev/mapper/sgws-sn1-var-local
rangedb0 → /dev/mapper/sgws-sn1-rangedb0
rangedb1 → /dev/mapper/sgws-sn1-rangedb1
rangedb2 → /dev/mapper/sgws-sn1-rangedb2
rangedb3 → /dev/mapper/sgws-sn1-rangedb3
```

Prepare the hosts (Linux)

How host-wide settings change during installation (Linux)

On bare metal systems, StorageGRID makes some changes to host-wide `sysctl` settings.



"Linux" refers to a RHEL, Ubuntu, or Debian deployment. For a list of supported versions, see the [NetApp Interoperability Matrix Tool \(IMT\)](#).

The following changes are made:

```
# Recommended Cassandra setting: CASSANDRA-3563, CASSANDRA-13008, DataStax
documentation
vm.max_map_count = 1048575

# core file customization
# Note: for cores generated by binaries running inside containers, this
# path is interpreted relative to the container filesystem namespace.
# External cores will go nowhere, unless /var/local/core also exists on
# the host.
kernel.core_pattern = /var/local/core/%e.core.%p

# Set the kernel minimum free memory to the greater of the current value
```

```
or
# 512MiB if the host has 48GiB or less of RAM or 1.83GiB if the host has
more than 48GiB of RTAM
vm.min_free_kbytes = 524288

# Enforce current default swappiness value to ensure the VM system has
some
# flexibility to garbage collect behind anonymous mappings. Bump
watermark_scale_factor
# to help avoid OOM conditions in the kernel during memory allocation
bursts. Bump
# dirty_ratio to 90 because we explicitly fsync data that needs to be
persistent, and
# so do not require the dirty_ratio safety net. A low dirty_ratio combined
with a large
# working set (nr_active_pages) can cause us to enter synchronous I/O mode
unnecessarily,
# with deleterious effects on performance.
vm.swappiness = 60
vm.watermark_scale_factor = 200
vm.dirty_ratio = 90

# Turn off slow start after idle
net.ipv4.tcp_slow_start_after_idle = 0

# Tune TCP window settings to improve throughput
net.core.rmem_max = 8388608
net.core.wmem_max = 8388608
net.ipv4.tcp_rmem = 4096 524288 8388608
net.ipv4.tcp_wmem = 4096 262144 8388608
net.core.netdev_max_backlog = 2500

# Turn on MTU probing
net.ipv4.tcp_mtu_probing = 1

# Be more liberal with firewall connection tracking
net.ipv4.netfilter.ip_conntrack_tcp_be_liberal = 1

# Reduce TCP keepalive time to reasonable levels to terminate dead
connections
net.ipv4.tcp_keepalive_time = 270
net.ipv4.tcp_keepalive_probes = 3
net.ipv4.tcp_keepalive_intvl = 30

# Increase the ARP cache size to tolerate being in a /16 subnet
net.ipv4.neigh.default.gc_thresh1 = 8192
```

```
net.ipv4.neigh.default.gc_thresh2 = 32768
net.ipv4.neigh.default.gc_thresh3 = 65536
net.ipv6.neigh.default.gc_thresh1 = 8192
net.ipv6.neigh.default.gc_thresh2 = 32768
net.ipv6.neigh.default.gc_thresh3 = 65536

# Disable IP forwarding, we are not a router
net.ipv4.ip_forward = 0

# Follow security best practices for ignoring broadcast ping requests
net.ipv4.icmp_echo_ignore_broadcasts = 1

# Increase the pending connection and accept backlog to handle larger
connection bursts.
net.core.somaxconn=4096
net.ipv4.tcp_max_syn_backlog=4096
```

Install Linux

You must install StorageGRID on all Linux grid hosts. For a list of supported versions, use the NetApp Interoperability Matrix Tool.

Before you begin

Ensure your operating system meets StorageGRID's minimum kernel version requirements, as listed below. Use the command `uname -r` to get your operating system's kernel version, or consult with your OS vendor.



"Linux" refers to a RHEL, Ubuntu, or Debian deployment. For a list of supported versions, see the [NetApp Interoperability Matrix Tool \(IMT\)](#).

RHEL

RHEL version	Minimum kernel version	Kernel package name
8.8 (deprecated)	4.18.0-477.10.1.el8_8.x86_64	kernel-4.18.0-477.10.1.el8_8.x86_64
8.10	4.18.0-553.el8_10.x86_64	kernel-4.18.0-553.el8_10.x86_64
9.0 (deprecated)	5.14.0-70.22.1.el9_0.x86_64	kernel-5.14.0-70.22.1.el9_0.x86_64
9.2 (deprecated)	5.14.0-284.11.1.el9_2.x86_64	kernel-5.14.0-284.11.1.el9_2.x86_64
9.4	5.14.0-427.18.1.el9_4.x86_64	kernel-5.14.0-427.18.1.el9_4.x86_64
9.6	5.14.0-570.18.1.el9_6.x86_64	kernel-5.14.0-570.18.1.el9_6.x86_64

Ubuntu

Note: Support for Ubuntu versions 18.04 and 20.04 have been deprecated and will be removed in a future release.

Ubuntu version	Minimum kernel version	Kernel package name
22.04.1	5.15.0-47-generic	linux-image-5.15.0-47-generic/jammy-updates,jammy-security,now 5.15.0-47.51
24.04	6.8.0-31-generic	linux-image-6.8.0-31-generic/noble,now 6.8.0-31.31

Debian

Note: Support for Debian version 11 has been deprecated and will be removed in a future release.

Debian version	Minimum kernel version	Kernel package name
11 (deprecated)	5.10.0-18-amd64	linux-image-5.10.0-18-amd64/stable,now 5.10.150-1
12	6.1.0-9-amd64	linux-image-6.1.0-9-amd64/stable,now 6.1.27-1

Steps

1. Install Linux on all physical or virtual grid hosts according to the distributor's instructions or your standard procedure.



Don't install any graphical desktop environments.

- If you are using the standard Linux installer when installing RHEL, select the "compute node" software configuration, if available, or "minimal install" base environment.
 - When installing Ubuntu, you must select **standard system utilities**. Selecting **OpenSSH server** is recommended to enable ssh access to your Ubuntu hosts. All other options can remain cleared.
2. Ensure that all hosts have access to package repositories, including the Extras channel for RHEL.
 3. If swap is enabled:
 - a. Run the following command: `$ sudo swapoff --all`
 - b. Remove all swap entries from `/etc/fstab` to persist the settings.



Failing to disable swap entirely can severely lower performance.

Understand AppArmor profile installation (Ubuntu and Debian)

If you are operating in a self-deployed Ubuntu environment and using the AppArmor mandatory access control system, the AppArmor profiles associated with packages you install on the base system might be blocked by the corresponding packages installed with StorageGRID.

By default, AppArmor profiles are installed for packages that you install on the base operating system. When you run these packages from the StorageGRID system container, the AppArmor profiles are blocked. The DHCP, MySQL, NTP, and tcdump base packages conflict with AppArmor, and other base packages might also conflict.

You have two choices for handling AppArmor profiles:

- Disable individual profiles for the packages installed on the base system that overlap with the packages in the StorageGRID system container. When you disable individual profiles, an entry appears in the StorageGRID log files indicating that AppArmor is enabled.

Use the following commands:

```
sudo ln -s /etc/apparmor.d/<profile.name> /etc/apparmor.d/disable/  
sudo apparmor_parser -R /etc/apparmor.d/<profile.name>
```

Example:

```
sudo ln -s /etc/apparmor.d/bin.ping /etc/apparmor.d/disable/  
sudo apparmor_parser -R /etc/apparmor.d/bin.ping
```

- Disable AppArmor altogether. For Ubuntu 9.10 or later, follow the instructions in the Ubuntu online community: [Disable AppArmor](#). Disabling AppArmor altogether might not be possible on newer Ubuntu versions.

After you disable AppArmor, no entries indicating that AppArmor is enabled will appear in the StorageGRID log files.

Configure the host network (Linux)

After completing the Linux installation on your hosts, you might need to perform some additional configuration to prepare a set of network interfaces on each host that are suitable for mapping into the StorageGRID nodes you will deploy later.



"Linux" refers to a RHEL, Ubuntu, or Debian deployment. For a list of supported versions, see the [NetApp Interoperability Matrix Tool \(IMT\)](#).

Before you begin

- You have reviewed the [StorageGRID networking guidelines](#).
- You have reviewed the information about [node container migration requirements](#).
- If you are using virtual hosts, you have read the [considerations and recommendations for MAC address cloning](#) before configuring the host network.



If you are using VMs as hosts, you should select VMXNET 3 as the virtual network adapter. The VMware E1000 network adapter has caused connectivity issues with StorageGRID containers deployed on certain distributions of Linux.

About this task

Grid nodes must be able to access the Grid Network and, optionally, the Admin and Client Networks. You provide this access by creating mappings that associate the host's physical interface to the virtual interfaces for each grid node. When creating host interfaces, use friendly names to facilitate deployment across all hosts, and to enable migration.

The same interface can be shared between the host and one or more nodes. For example, you might use the same interface for host access and node Admin Network access, to facilitate host and node maintenance. Although the same interface can be shared between the host and individual nodes, all must have different IP addresses. IP addresses can't be shared between nodes or between the host and any node.

You can use the same host network interface to provide the Grid Network interface for all StorageGRID nodes on the host; you can use a different host network interface for each node; or you can do something in between. However, you would not typically provide the same host network interface as both the Grid and Admin Network interfaces for a single node, or as the Grid Network interface for one node and the Client Network interface for another.

You can complete this task in many ways. For example, if your hosts are virtual machines and you are deploying one or two StorageGRID nodes for each host, you can create the correct number of network interfaces in the hypervisor, and use a 1-to-1 mapping. If you are deploying multiple nodes on bare metal hosts for production use, you can leverage the Linux networking stack's support for VLAN and LACP for fault tolerance and bandwidth sharing. The following sections provide detailed approaches for both of these examples. You don't need to use either of these examples; you can use any approach that meets your needs.



Don't use bond or bridge devices directly as the container network interface. Doing so could prevent node start-up caused by a kernel issue with the use of MACVLAN with bond and bridge devices in the container namespace. Instead, use a non-bond device, such as a VLAN or virtual Ethernet (veth) pair. Specify this device as the network interface in the node configuration file.

Considerations and recommendations for MAC address cloning

MAC address cloning causes the container to use the MAC address of the host, and the host to use the MAC address of either an address you specify or a randomly generated one. You should use MAC address cloning to avoid the use of promiscuous mode network configurations.

Enabling MAC cloning

In certain environments, security can be enhanced through MAC address cloning because it enables you to use a dedicated virtual NIC for the Admin Network, Grid Network, and Client Network. Having the container use the MAC address of the dedicated NIC on the host allows you to avoid using promiscuous mode network configurations.



MAC address cloning is intended to be used with virtual server installations and might not function properly with all physical appliance configurations.



If a node fails to start due to a MAC cloning targeted interface being busy, you might need to set the link to "down" before starting node. Additionally, it is possible that the virtual environment might prevent MAC cloning on a network interface while the link is up. If a node fails to set the MAC address and start due to an interface being busy, setting the link to "down" before starting the node might fix the issue.

MAC address cloning is disabled by default and must be set by node configuration keys. You should enable it when you install StorageGRID.

There is one key for each network:

- ADMIN_NETWORK_TARGET_TYPE_INTERFACE_CLONE_MAC
- GRID_NETWORK_TARGET_TYPE_INTERFACE_CLONE_MAC
- CLIENT_NETWORK_TARGET_TYPE_INTERFACE_CLONE_MAC

Setting the key to "true" causes the container to use the MAC address of the host's NIC. Additionally, the host will then use the MAC address of the specified container network. By default, the container address is a randomly generated address, but if you have set one using the `__NETWORK_MAC` node configuration key, that address is used instead. The host and container will always have different MAC addresses.



Enabling MAC cloning on a virtual host without also enabling promiscuous mode on the hypervisor might cause Linux host networking using the host's interface to stop working.

MAC cloning use cases

There are two use cases to consider with MAC cloning:

- **MAC cloning not enabled:** When the `__CLONE_MAC` key in the node configuration file is not set, or set to "false," the host will use the host NIC MAC and the container will have a StorageGRID-generated MAC unless a MAC is specified in the `__NETWORK_MAC` key. If an address is set in the `__NETWORK_MAC` key, the container will have the address specified in the `__NETWORK_MAC` key. This configuration of keys requires the use of promiscuous mode.
- **MAC cloning enabled:** When the `__CLONE_MAC` key in the node configuration file is set to "true," the container uses the host NIC MAC, and the host uses a StorageGRID-generated MAC unless a MAC is specified in the `__NETWORK_MAC` key. If an address is set in the `__NETWORK_MAC` key, the host uses the specified address instead of a generated one. In this configuration of keys, you should not use promiscuous mode.



If you don't want to use MAC address cloning and would rather allow all interfaces to receive and transmit data for MAC addresses other than the ones assigned by the hypervisor, ensure that the security properties at the virtual switch and port group levels are set to **Accept** for Promiscuous Mode, MAC Address Changes, and Forged Transmits. The values set on the virtual switch can be overridden by the values at the port group level, so ensure that settings are the same in both places.

To enable MAC cloning, see the [instructions for creating node configuration files](#).

MAC cloning example


Example of MAC cloning enabled with a host having MAC address of 11:22:33:44:55:66 for the interface ens256 and the following keys in the node configuration file:

- ADMIN_NETWORK_TARGET = ens256
- ADMIN_NETWORK_MAC = b2:9c:02:c2:27:10
- ADMIN_NETWORK_TARGET_TYPE_INTERFACE_CLONE_MAC = true

Result: the host MAC for ens256 is b2:9c:02:c2:27:10 and the Admin Network MAC is 11:22:33:44:55:66

Example 1: 1-to-1 mapping to physical or virtual NICs

Example 1 describes a simple physical interface mapping that requires little or no host-side configuration.



The Linux operating system creates the ensXYZ interfaces automatically during installation or boot, or when the interfaces are hot-added. No configuration is required other than ensuring that the interfaces are set to come up automatically after boot. You do have to determine which ensXYZ corresponds to which StorageGRID network (Grid, Admin, or Client) so you can provide the correct mappings later in the configuration process.

Note that the figure show multiple StorageGRID nodes; however, you would normally use this configuration for single-node VMs.

If Switch 1 is a physical switch, you should configure the ports connected to interfaces 10G1 through 10G3 for access mode, and place them on the appropriate VLANs.

Example 2: LACP bond carrying VLANs


Example 2 assumes you are familiar with bonding network interfaces and with creating VLAN interfaces on the Linux distribution you are using.

About this task

Example 2 describes a generic, flexible, VLAN-based scheme that facilitates the sharing of all available network bandwidth across all nodes on a single host. This example is particularly applicable to bare metal hosts.

To understand this example, suppose you have three separate subnets for the Grid, Admin, and Client Networks at each data center. The subnets are on separate VLANs (1001, 1002, and 1003) and are presented to the host on a LACP-bonded trunk port (bond0). You would configure three VLAN interfaces on the bond: bond0.1001, bond0.1002, and bond0.1003.

If you require separate VLANs and subnets for node networks on the same host, you can add VLAN interfaces on the bond and map them into the host (shown as bond0.1004 in the illustration).



Steps

1. Aggregate all physical network interfaces that will be used for StorageGRID network connectivity into a single LACP bond.

Use the same name for the bond on every host, for example, `bond0`.

2. Create VLAN interfaces that use this bond as their associated "physical device," using the standard VLAN interface naming convention `physdev-name.VLAN ID`.

Note that steps 1 and 2 require appropriate configuration on the edge switches terminating the other ends of the network links. The edge switch ports must also be aggregated into a LACP port channel, configured as a trunk, and allowed to pass all required VLANs.

Sample interface configuration files for this per-host networking configuration scheme are provided.

Related information

- [Example /etc/network/interfaces for Ubuntu and Debian](#)
- [Example /etc/sysconfig/network-scripts for RHEL](#)

Configure host storage (Linux)

You must allocate block storage volumes to each Linux host.

Before you begin

You have reviewed the following topics, which provide information you need to accomplish this task:

- [Storage and performance requirements](#)
- [Node container migration requirements](#)



"Linux" refers to a RHEL, Ubuntu, or Debian deployment. For a list of supported versions, see the [NetApp Interoperability Matrix Tool \(IMT\)](#).

About this task

When allocating block storage volumes (LUNs) to hosts, use the tables in "Storage requirements" to determine the following:

- Number of volumes required for each host (based on the number and types of nodes that will be deployed on that host)
- Storage category for each volume (that is, System Data or Object Data)
- Size of each volume

You will use this information as well as the persistent name assigned by Linux to each physical volume when you deploy StorageGRID nodes on the host.



You don't need to partition, format, or mount any of these volumes; you just need to ensure they are visible to the hosts.



Only one object-data LUN is required for metadata-only Storage Nodes.

Avoid using "raw" special device files (`/dev/sdb`, for example) as you compose your list of volume names. These files can change across reboots of the host, which will impact proper operation of the system. If you are using iSCSI LUNs and Device Mapper Multipathing, consider using multipath aliases in the `/dev/mapper` directory, especially if your SAN topology includes redundant network paths to the shared storage. Alternatively, you can use the system-created softlinks under `/dev/disk/by-path/` for your persistent

device names.

For example:

```
ls -l
$ ls -l /dev/disk/by-path/
total 0
lrwxrwxrwx 1 root root 9 Sep 19 18:53 pci-0000:00:07.1-ata-2 -> ../../sr0
lrwxrwxrwx 1 root root 9 Sep 19 18:53 pci-0000:03:00.0-scsi-0:0:0:0 ->
../../sda
lrwxrwxrwx 1 root root 10 Sep 19 18:53 pci-0000:03:00.0-scsi-0:0:0:0-part1
-> ../../sda1
lrwxrwxrwx 1 root root 10 Sep 19 18:53 pci-0000:03:00.0-scsi-0:0:0:0-part2
-> ../../sda2
lrwxrwxrwx 1 root root 9 Sep 19 18:53 pci-0000:03:00.0-scsi-0:0:1:0 ->
../../sdb
lrwxrwxrwx 1 root root 9 Sep 19 18:53 pci-0000:03:00.0-scsi-0:0:2:0 ->
../../sdc
lrwxrwxrwx 1 root root 9 Sep 19 18:53 pci-0000:03:00.0-scsi-0:0:3:0 ->
../../sdd
```

Results will differ for each installation.

Assign friendly names to each of these block storage volumes to simplify the initial StorageGRID installation and future maintenance procedures. If you are using the device mapper multipath driver for redundant access to shared storage volumes, you can use the `alias` field in your `/etc/multipath.conf` file.

For example:

```

multipaths {
    multipath {
        wwid 3600a09800059d6df00005df2573c2c30
        alias docker-storage-volume-hostA
    }
    multipath {
        wwid 3600a09800059d6df00005df3573c2c30
        alias sgws-adml-var-local
    }
    multipath {
        wwid 3600a09800059d6df00005df4573c2c30
        alias sgws-adml-audit-logs
    }
    multipath {
        wwid 3600a09800059d6df00005df5573c2c30
        alias sgws-adml-tables
    }
    multipath {
        wwid 3600a09800059d6df00005df6573c2c30
        alias sgws-gw1-var-local
    }
    multipath {
        wwid 3600a09800059d6df00005df7573c2c30
        alias sgws-sn1-var-local
    }
    multipath {
        wwid 3600a09800059d6df00005df7573c2c30
        alias sgws-sn1-rangedb-0
    }
    ...
}

```

Using the alias field in this way causes the aliases to appear as block devices in the `/dev/mapper` directory on the host, allowing you to specify a friendly, easily-validated name whenever a configuration or maintenance operation requires specifying a block storage volume.

If you are setting up shared storage to support StorageGRID node migration and using Device Mapper Multipathing, you can create and install a common `/etc/multipath.conf` on all co-located hosts. Just make sure to use a different container engine storage volume on each host. Using aliases and including the target hostname in the alias for each container engine storage volume LUN will make this easy to remember and is recommended.



Support for Docker as the container engine for software-only deployments is deprecated. Docker will be replaced with another container engine in a future release.

Related information

- [Configure container engine storage volume](#)

- [Storage and performance requirements](#)
- [Node container migration requirements](#)

Configure container engine storage volume (Linux)

Before installing the Docker or Podman container engine, you might need to format the storage volume and mount it.



Support for Docker as the container engine for software-only deployments is deprecated. Docker will be replaced with another container engine in a future release.



"Linux" refers to a RHEL, Ubuntu, or Debian deployment. For a list of supported versions, see the [NetApp Interoperability Matrix Tool \(IMT\)](#).

About this task

You can skip these steps if you plan to use the root volume for the Docker or Podman storage volume and have sufficient space available on the host partition containing:

- Podman: `/var/lib/containers`
- Docker: `/var/lib/docker`

Steps

1. Create a file system on the container engine storage volume:

RHEL

```
sudo mkfs.ext4 container-engine-storage-volume-device
```

Ubuntu or Debian

```
sudo mkfs.ext4 docker-storage-volume-device
```

2. Mount the container engine storage volume:

RHEL

- For Docker:

```
sudo mkdir -p /var/lib/docker
sudo mount container-storage-volume-device /var/lib/docker
```

- For Podman:

```
sudo mkdir -p /var/lib/containers
sudo mount container-storage-volume-device /var/lib/containers
```

Ubuntu or Debian

```
sudo mkdir -p /var/lib/docker
sudo mount docker-storage-volume-device /var/lib/docker
```

- For Podman:

```
sudo mkdir -p /var/lib/podman
sudo mount container-storage-volume-device /var/lib/podman
```

3. Add an entry for the container storage volume device to `/etc/fstab`.

- RHEL: `container-storage-volume-device`
- Ubuntu or Debian: `docker-storage-volume-device`

This step ensures that the storage volume will remount automatically after host reboots.

Install Docker

The StorageGRID system can run on Linux as a collection of containers.

- Before you can install StorageGRID for Ubuntu or Debian, you must install Docker.
- If you have chosen to use the Docker container engine, follow these steps to install Docker. Otherwise, [install Podman](#).



Support for Docker as the container engine for software-only deployments is deprecated. Docker will be replaced with another container engine in a future release.

Steps

1. Install Docker by following the instructions for your Linux distribution.



If Docker is not included with your Linux distribution, you can download it from the Docker website.

2. Ensure Docker has been enabled and started by running the following two commands:

```
sudo systemctl enable docker
```

```
sudo systemctl start docker
```

3. Confirm you have installed the expected version of Docker by entering the following:

```
sudo docker version
```

The Client and Server versions must be 1.11.0 or later.

Install Podman

The StorageGRID system runs as a collection of containers. If you have chosen to use the Podman container engine, follow these steps to install Podman. Otherwise, [install Docker](#).

Steps

1. Install Podman and Podman-Docker by following the instructions for your Linux distribution.



You must also install the Podman-Docker package when you install Podman.

2. Confirm you have installed the expected version of Podman and Podman-Docker by entering the following:

```
sudo docker version
```



The Podman-Docker package allows you to use Docker commands.

The Client and Server versions must be 3.2.3 or later.

```
Version: 3.2.3
API Version: 3.2.3
Go Version: go1.15.7
Built: Tue Jul 27 03:29:39 2021
OS/Arch: linux/amd64
```

Related information

[Configure host storage](#)

Install StorageGRID host services (Linux)

You use the StorageGRID package for your operating system type to install the StorageGRID host services.



"Linux" refers to a RHEL, Ubuntu, or Debian deployment. For a list of supported versions, see the [NetApp Interoperability Matrix Tool \(IMT\)](#).

RHEL

You use the StorageGRID RPM package to install the StorageGRID host services.

About this task

These instructions describe how to install the host services from the RPM packages. As an alternative, you can use the DNF repository metadata included in the installation archive to install the RPM packages remotely. See the DNF repository instructions for your Linux operating system.

Steps

1. Copy the StorageGRID RPM packages to each of your hosts, or make them available on shared storage.

For example, place them in the `/tmp` directory, so you can use the example command in the next step.

2. Log in to each host as root or using an account with sudo permission, and run the following commands in the order specified:

```
sudo dnf --nogpgcheck localinstall /tmp/StorageGRID-Webscale-Images-  
version-SHA.rpm
```

```
sudo dnf --nogpgcheck localinstall /tmp/StorageGRID-Webscale-  
Service-version-SHA.rpm
```



You must install the Images package first, and the Service package second.



If you placed the packages in a directory other than `/tmp`, modify the command to reflect the path you used.

Ubuntu or Debian

You use the StorageGRID DEB package to install the StorageGRID host services for Ubuntu, or Debian.

About this task

These instructions describe how to install the host services from the DEB packages. As an alternative, you can use the APT repository metadata included in the installation archive to install the DEB packages remotely. See the APT repository instructions for your Linux operating system.

Steps

1. Copy the StorageGRID DEB packages to each of your hosts, or make them available on shared storage.

For example, place them in the `/tmp` directory, so you can use the example command in the next step.

2. Log in to each host as root or using an account with sudo permission, and run the following commands.

You must install the `images` package first, and the `service` package second. If you placed the packages in a directory other than `/tmp`, modify the command to reflect the path you used.

```
sudo dpkg --install /tmp/storagegrid-webscale-images-version-SHA.deb
```

```
sudo dpkg --install /tmp/storagegrid-webscale-service-version-SHA.deb
```



Python 3 must already be installed before the StorageGRID packages can be installed. The `sudo dpkg --install /tmp/storagegrid-webscale-images-version-SHA.deb` command will fail until you have done so.

Automate the software-based nodes installation

Automate the installation (Linux)

You can automate the installation of the StorageGRID host service and the configuration of grid nodes.



About this task

"Linux" refers to a RHEL, Ubuntu, or Debian deployment. For a list of supported versions, see the [NetApp Interoperability Matrix Tool \(IMT\)](#).

Automating the deployment might be useful in any of the following cases:

- You already use a standard orchestration framework, such as Ansible, Puppet, or Chef, to deploy and configure physical or virtual hosts.
- You intend to deploy multiple StorageGRID instances.
- You are deploying a large, complex StorageGRID instance.

The StorageGRID host service is installed by a package and driven by configuration files. You can create the configuration files using one of these methods:

- [Create the configuration files](#) interactively during a manual installation.
- Prepare the configuration files ahead of time (or programmatically) to enable automated installation using standard orchestration frameworks, as describe in this article.

StorageGRID provides optional Python scripts for automating the configuration of StorageGRID appliances and the entire StorageGRID system (the "grid"). You can use these scripts directly, or you can inspect them to learn how to use the [StorageGRID installation REST API](#) in grid deployment and configuration tools you develop yourself.

Automate the installation and configuration of the StorageGRID host service

You can automate the installation of the StorageGRID host service using standard orchestration frameworks

such as Ansible, Puppet, Chef, Fabric, or SaltStack.

The StorageGRID host service is packaged in a DEB (Ubuntu or Debian) or an RPM (RHEL) and is driven by configuration files that you can prepare ahead of time (or programmatically) to enable automated installation. If you already use a standard orchestration framework to install and configure your Linux deployment, adding StorageGRID to your playbooks or recipes should be straightforward.

You can automate all of the steps for preparing the hosts and deploying virtual grid nodes.

Example Ansible role and playbook

Example Ansible role and playbook are supplied with the installation archive in the `/extras` folder. The Ansible playbook shows how the `storagegrid` role prepares the hosts and installs StorageGRID onto the target servers. You can customize the role or playbook as necessary.



The example playbook does not include the steps required to create network devices before starting the StorageGRID host service. Add these steps before finalizing and using the playbook.

RHEL

For RHEL, the installation tasks in the provided `storagegrid` role example use the `ansible.builtin.dnf` module to perform the installation from the local RPM files or a remote Yum repository. If the module is unavailable or not supported, you might need to edit the appropriate Ansible tasks in the following files to use the `yum` or `ansible.builtin.yum` module:

- `roles/storagegrid/tasks/rhel_install_from_repo.yml`
- `roles/storagegrid/tasks/rhel_install_from_local.yml`

Ubuntu or Debian

For Ubuntu or Debian, the installation tasks in the provided `storagegrid` role example use the `ansible.builtin.apt` module to perform the installation from the local DEB files or a remote apt repository. If the module is unavailable or not supported, you might need to edit the appropriate Ansible tasks in the following files to use the `ansible.builtin.apt` module:

- `roles/storagegrid/tasks/deb_install_from_repo.yml`
- `roles/storagegrid/tasks/deb_install_from_local.yml`

Automate the configuration of StorageGRID

After deploying the grid nodes, you can automate the configuration of the StorageGRID system.

Before you begin

- You know the location of the following files from the installation archive.

Filename	Description
<code>configure-storagegrid.py</code>	Python script used to automate the configuration

Filename	Description
<code>configure-storagegrid.sample.json</code>	Example configuration file for use with the script
<code>configure-storagegrid.blank.json</code>	Blank configuration file for use with the script

- You have created a `configure-storagegrid.json` configuration file. To create this file, you can modify the example configuration file (`configure-storagegrid.sample.json`) or the blank configuration file (`configure-storagegrid.blank.json`).



Store the management password and provisioning passphrase from the passwords section of the modified `configure-storagegrid.json` configuration file in a secure location. These passwords are required for installation, expansion, and maintenance procedures. You should also back up the modified `configure-storagegrid.json` configuration file and store it in a secure location.

About this task

You can use the `configure-storagegrid.py` Python script and the `configure-storagegrid.json` configuration file to automate the configuration of your StorageGRID system.



You can also configure the system using the Grid Manager or the Installation API.

Steps

- Log in to the Linux machine you are using to run the Python script.
- Change to the directory where you extracted the installation archive.

For example:

```
cd StorageGRID-Webscale-version/platform
```

where `platform` is `debs`, `rpms`, or `vsphere`.

- Run the Python script and use the configuration file you created.

For example:

```
./configure-storagegrid.py ./configure-storagegrid.json --start-install
```

Result

A Recovery Package `.zip` file is generated during the configuration process, and it is downloaded to the directory where you are running the installation and configuration process. You must back up the Recovery Package file so that you can recover the StorageGRID system if one or more grid nodes fails. For example, copy it to a secure, backed up network location and to a secure cloud storage location.



The Recovery Package file must be secured because it contains encryption keys and passwords that can be used to obtain data from the StorageGRID system.

If you specified that random passwords should be generated, open the `Passwords.txt` file and look for the passwords required to access your StorageGRID system.

```
#####  
##### The StorageGRID "Recovery Package" has been downloaded as: #####  
#####      ./sgws-recovery-package-994078-rev1.zip      #####  
#####   Safeguard this file as it will be needed in case of a   #####  
#####           StorageGRID node recovery.           #####  
#####
```

Your StorageGRID system is installed and configured when a confirmation message is displayed.

```
StorageGRID has been configured and installed.
```

Automate the installation (VMware)

You can use the VMware OVF Tool to automate the deployment of grid nodes. You can also automate the configuration of StorageGRID.

Automate grid node deployment

Use the VMware OVF Tool to automate the deployment of grid nodes.

Before you begin

- You have access to a Linux/Unix system with Bash 3.2 or later.
- You have VMware vSphere with vCenter
- You have VMware OVF Tool installed and correctly configured.
- You know the username and password to access VMware vSphere using the OVF Tool
- You have the sufficient permissions to deploy VMs from OVF files and power them on, and permissions to create additional volumes to attach to the VMs. See the `ovftool` documentation for details.
- You know the virtual infrastructure (VI) URL for the location in vSphere where you want to deploy the StorageGRID virtual machines. This URL will typically be a vApp, or Resource Pool. For example:
`vi://vcenter.example.com/vi/sgws`



You can use the VMware `ovftool` utility to determine this value (see the `ovftool` documentation for details).



If you are deploying to a vApp, the virtual machines will not start automatically the first time, and you must power them on manually.

- You have collected all the required information for the deployment configuration file. See [Collect information about your deployment environment](#) for information.
- You have access to the following files from the VMware installation archive for StorageGRID:

Filename	Description
NetApp-SG-version-SHA.vmdk	The virtual machine disk file that is used as a template for creating grid node virtual machines. Note: This file must be in the same folder as the .ovf and .mf files.
vsphere-primary-admin.ovf vsphere-primary-admin.mf	The Open Virtualization Format template file (.ovf) and manifest file (.mf) for deploying the primary Admin Node.
vsphere-non-primary-admin.ovf vsphere-non-primary-admin.mf	The template file (.ovf) and manifest file (.mf) for deploying non-primary Admin Nodes.
vsphere-gateway.ovf vsphere-gateway.mf	The template file (.ovf) and manifest file (.mf) for deploying Gateway Nodes.
vsphere-storage.ovf vsphere-storage.mf	The template file (.ovf) and manifest file (.mf) for deploying virtual machine-based Storage Nodes.
deploy-vsphere-ovftool.sh	The Bash shell script used to automate the deployment of virtual grid nodes.
deploy-vsphere-ovftool-sample.ini	The example configuration file for use with the deploy-vsphere-ovftool.sh script.

Define the configuration file for your deployment

You specify the information needed to deploy virtual grid nodes for StorageGRID in a configuration file, which is used by the `deploy-vsphere-ovftool.sh` Bash script. You can modify an example configuration file, so that you don't have to create the file from scratch.

Steps

1. Make a copy of the example configuration file (`deploy-vsphere-ovftool.sample.ini`). Save the new file as `deploy-vsphere-ovftool.ini` in the same directory as `deploy-vsphere-ovftool.sh`.
2. Open `deploy-vsphere-ovftool.ini`.
3. Enter all of the information required to deploy VMware virtual grid nodes.

See [Configuration file settings](#) for information.

4. When you have entered and verified all of the necessary information, save and close the file.

Configuration file settings

The `deploy-vsphere-ovftool.ini` configuration file contains the settings that are required to deploy virtual grid nodes.

The configuration file first lists global parameters, and then lists node-specific parameters in sections defined

by node name. When the file is used:

- *Global parameters* are applied to all grid nodes.
- *Node-specific parameters* override global parameters.

Global parameters

Global parameters are applied to all grid nodes, unless they are overridden by settings in individual sections. Place the parameters that apply to multiple nodes in the global parameter section, and then override these settings as necessary in the sections for individual nodes.

- **OVFTOOL_ARGUMENTS:** You can specify OVFTOOL_ARGUMENTS as global settings, or you can apply arguments individually to specific nodes. For example:

```
OVFTOOL_ARGUMENTS = --powerOn --noSSLVerify --diskMode=eagerZeroedThick
--datastore='datastore_name'
```

You can use the `--powerOffTarget` and `--overwrite` options to shut down and replace existing virtual machines.



You should deploy nodes to different datastores and specify OVFTOOL_ARGUMENTS for each node, instead of globally.

- **SOURCE:** The path to the StorageGRID virtual machine template (`.vmdk`) file and the `.ovf` and `.mf` files for individual grid nodes. This defaults to the current directory.

```
SOURCE = /downloads/StorageGRID-Webscale-version/vsphere
```

- **TARGET:** The VMware vSphere virtual infrastructure (vi) URL for the location where StorageGRID will be deployed. For example:

```
TARGET = vi://vcenter.example.com/vm/sgws
```

- **GRID_NETWORK_CONFIG:** The method used to acquire IP addresses, either `STATIC` or `DHCP`. The default is `STATIC`. If all or most of the nodes use the same method for acquiring IP addresses, you can specify that method here. You can then override the global setting by specifying different settings for one or more individual nodes. For example:

```
GRID_NETWORK_CONFIG = STATIC
```

- **GRID_NETWORK_TARGET:** The name of an existing VMware network to use for the Grid Network. If all or most of the nodes use the same network name, you can specify it here. You can then override the global setting by specifying different settings for one or more individual nodes. For example:

```
GRID_NETWORK_TARGET = SG Admin Network
```

- **GRID_NETWORK_MASK:** The network mask for the Grid Network. If all or most of the nodes use the same network mask, you can specify it here. You can then override the global setting by specifying different settings for one or more individual nodes. For example:

```
GRID_NETWORK_MASK = 255.255.255.0
```

- **GRID_NETWORK_GATEWAY:** The network gateway for the Grid Network. If all or most of the nodes use the same network gateway, you can specify it here. You can then override the global setting by specifying different settings for one or more individual nodes. For example:

```
GRID_NETWORK_GATEWAY = 10.1.0.1
```

- **GRID_NETWORK_MTU:** Optional. The maximum transmission unit (MTU) on the Grid Network. If specified, the value must be between 1280 and 9216. For example:

```
GRID_NETWORK_MTU = 9000
```

If omitted, 1400 is used.

If you want to use jumbo frames, set the MTU to a value suitable for jumbo frames, such as 9000. Otherwise, keep the default value.



The MTU value of the network must match the value configured on the virtual switch port in vSphere that the node is connected to. Otherwise, network performance issues or packet loss might occur.



For the best network performance, all nodes should be configured with similar MTU values on their Grid Network interfaces. The **Grid Network MTU mismatch** alert is triggered if there is a significant difference in MTU settings for the Grid Network on individual nodes. The MTU values don't have to be the same for all network types.

- **ADMIN_NETWORK_CONFIG:** The method used to acquire IP addresses, either DISABLED, STATIC, or DHCP. The default is DISABLED. If all or most of the nodes use the same method for acquiring IP addresses, you can specify that method here. You can then override the global setting by specifying different settings for one or more individual nodes. For example:

```
ADMIN_NETWORK_CONFIG = STATIC
```

- **ADMIN_NETWORK_TARGET:** The name of an existing VMware network to use for the Admin Network. This setting is required unless the Admin Network is disabled. If all or most of the nodes use the same network name, you can specify it here. Unlike the Grid Network, all nodes do not need to be connected to the same Admin Network. You can then override the global setting by specifying different settings for one or more individual nodes. For example:

```
ADMIN_NETWORK_TARGET = SG Admin Network
```

- **ADMIN_NETWORK_MASK:** The network mask for the Admin Network. This setting is required if you are using static IP addressing. If all or most of the nodes use the same network mask, you can specify it here. You can then override the global setting by specifying different settings for one or more individual nodes. For example:

```
ADMIN_NETWORK_MASK = 255.255.255.0
```

- **ADMIN_NETWORK_GATEWAY:** The network gateway for the Admin Network. This setting is required if you are using static IP addressing and you specify external subnets in the ADMIN_NETWORK_ESL setting. (That is, it is not required if ADMIN_NETWORK_ESL is empty.) If all or most of the nodes use the same network gateway, you can specify it here. You can then override the global setting by specifying different settings for one or more individual nodes. For example:

```
ADMIN_NETWORK_GATEWAY = 10.3.0.1
```

- **ADMIN_NETWORK_ESL:** The external subnet list (routes) for the Admin Network, specified as a comma-separated list of CIDR route destinations. If all or most of the nodes use the same external subnet list, you can specify it here. You can then override the global setting by specifying different settings for one or more individual nodes. For example:

```
ADMIN_NETWORK_ESL = 172.16.0.0/21,172.17.0.0/21
```

- **ADMIN_NETWORK_MTU:** Optional. The maximum transmission unit (MTU) on the Admin Network. Don't specify if ADMIN_NETWORK_CONFIG = DHCP. If specified, the value must be between 1280 and 9216. If omitted, 1400 is used. If you want to use jumbo frames, set the MTU to a value suitable for jumbo frames, such as 9000. Otherwise, keep the default value. If all or most of the nodes use the same MTU for the Admin Network, you can specify it here. You can then override the global setting by specifying different settings for one or more individual nodes. For example:

```
ADMIN_NETWORK_MTU = 8192
```

- **CLIENT_NETWORK_CONFIG:** The method used to acquire IP addresses, either DISABLED, STATIC, or DHCP. The default is DISABLED. If all or most of the nodes use the same method for acquiring IP addresses, you can specify that method here. You can then override the global setting by specifying different settings for one or more individual nodes. For example:

```
CLIENT_NETWORK_CONFIG = STATIC
```

- **CLIENT_NETWORK_TARGET:** The name of an existing VMware network to use for the Client Network. This setting is required unless the Client Network is disabled. If all or most of the nodes use the same network name, you can specify it here. Unlike the Grid Network, all nodes do not need to be connected to the same Client Network. You can then override the global setting by specifying different settings for one or more individual nodes. For example:

```
CLIENT_NETWORK_TARGET = SG Client Network
```

- **CLIENT_NETWORK_MASK:** The network mask for the Client Network. This setting is required if you are using static IP addressing. If all or most of the nodes use the same network mask, you can specify it here. You can then override the global setting by specifying different settings for one or more individual nodes. For example:

```
CLIENT_NETWORK_MASK = 255.255.255.0
```

- **CLIENT_NETWORK_GATEWAY:** The network gateway for the Client Network. This setting is required if you are using static IP addressing. If all or most of the nodes use the same network gateway, you can specify it here. You can then override the global setting by specifying different settings for one or more individual nodes. For example:

```
CLIENT_NETWORK_GATEWAY = 10.4.0.1
```

- **CLIENT_NETWORK_MTU:** Optional. The maximum transmission unit (MTU) on the Client Network. Don't specify if `CLIENT_NETWORK_CONFIG = DHCP`. If specified, the value must be between 1280 and 9216. If omitted, 1400 is used. If you want to use jumbo frames, set the MTU to a value suitable for jumbo frames, such as 9000. Otherwise, keep the default value. If all or most of the nodes use the same MTU for the Client Network, you can specify it here. You can then override the global setting by specifying different settings for one or more individual nodes. For example:

```
CLIENT_NETWORK_MTU = 8192
```

- **PORT_REMAP:** Remaps any port used by a node for internal grid node communications or external communications. Remapping ports is necessary if enterprise networking policies restrict one or more ports used by StorageGRID. For the list of ports used by StorageGRID, see internal grid node communications and external communications in [Networking guidelines](#).



Don't remap the ports you are planning to use to configure load balancer endpoints.



If only `PORT_REMAP` is set, the mapping that you specify is used for both inbound and outbound communications. If `PORT_REMAP_INBOUND` is also specified, `PORT_REMAP` applies only to outbound communications.

The format used is: *network type/protocol/default port used by grid node/new port*, where network type is grid, admin, or client, and protocol is tcp or udp.

For example:

```
PORT_REMAP = client/tcp/18082/443
```

If used alone, this example setting symmetrically maps both inbound and outbound communications for the grid node from port 18082 to port 443. If used in conjunction with `PORT_REMAP_INBOUND`, this example

setting maps outbound communications from port 18082 to port 443.

You can also remap multiple ports using a comma-separated list.

For example:

```
PORT_REMAP = client/tcp/18082/443, client/tcp/18083/80
```

- **PORT_REMAP_INBOUND:** Remaps inbound communications for the specified port. If you specify `PORT_REMAP_INBOUND` but don't specify a value for `PORT_REMAP`, outbound communications for the port are unchanged.



Don't remap the ports you are planning to use to configure load balancer endpoints.

The format used is: *network type/protocol/_default port used by grid node/new port*, where network type is grid, admin, or client, and protocol is tcp or udp.

For example:

```
PORT_REMAP_INBOUND = client/tcp/443/18082
```

This example takes traffic that is sent to port 443 to pass an internal firewall and directs it to port 18082, where the grid node is listening for S3 requests.

You can also remap multiple inbound ports using a comma-separated list.

For example:

```
PORT_REMAP_INBOUND = grid/tcp/3022/22, admin/tcp/3022/22
```

- **TEMPORARY_PASSWORD_TYPE:** The type of temporary installation password to be used when accessing the VM console or the StorageGRID Installation API, or using SSH, before the node joins the grid.



If all or most of the nodes use the same type of temporary installation password, specify the type in the global parameter section. Then, optionally use a different setting for an individual node. For example, if you select **Use Custom Password** globally, you can use **CUSTOM_TEMPORARY_PASSWORD=<password>** to set the password for each node.

TEMPORARY_PASSWORD_TYPE can be one of the following:

- **Use node name:** The node name is used as the temporary installation password and provides access to VM console, the StorageGRID Installation API, and SSH.
- **Disable password:** No temporary installation password will be used. If you need to access the VM to debug installation issues, see [Troubleshoot installation issues](#).
- **Use custom password:** The value provided with **CUSTOM_TEMPORARY_PASSWORD=<password>** is used as the temporary installation password and provides access to VM console, the StorageGRID Installation API, and SSH.



Optionally, you can omit the **TEMPORARY_PASSWORD_TYPE** parameter and only specify **CUSTOM_TEMPORARY_PASSWORD=<password>**.

- **CUSTOM_TEMPORARY_PASSWORD=<password>**

Optional. The temporary password to use during installation when accessing VM console, the StorageGRID Installation API, and SSH. Ignored if **TEMPORARY_PASSWORD_TYPE** is set to **Use node name** or **Disable password**.

Node-specific parameters

Each node is in its own section of the configuration file. Each node requires the following settings:

- The section head defines the node name that will be displayed in the Grid Manager. You can override that value by specifying the optional **NODE_NAME** parameter for the node.
- **NODE_TYPE**: VM_Admin_Node, VM_Storage_Node, or VM_API_Gateway_Node
- **STORAGE_TYPE**: combined, data, or metadata. This optional parameter for storage nodes defaults to combined (data and metadata) if it is not specified. For more information, see [Types of Storage Nodes](#).
- **GRID_NETWORK_IP**: The IP address for the node on the Grid Network.
- **ADMIN_NETWORK_IP**: The IP address for the node on the Admin Network. Required only if the node is attached to the Admin Network and **ADMIN_NETWORK_CONFIG** is set to **STATIC**.
- **CLIENT_NETWORK_IP**: The IP address for the node on the Client Network. Required only if the node is attached to the Client Network and **CLIENT_NETWORK_CONFIG** for this node is set to **STATIC**.
- **ADMIN_IP**: The IP address for the primary Admin node on the Grid Network. Use the value that you specify as the **GRID_NETWORK_IP** for the primary Admin Node. If you omit this parameter, the node attempts to discover the primary Admin Node IP using mDNS. For more information, see [How grid nodes discover the primary Admin Node](#).



The **ADMIN_IP** parameter is ignored for the primary Admin Node.

- Any parameters that were not set globally. For example, if a node is attached to the Admin Network and you did not specify **ADMIN_NETWORK** parameters globally, you must specify them for the node.

Primary Admin Node

The following additional settings are required for the primary Admin Node:

- **NODE_TYPE**: VM_Admin_Node
- **ADMIN_ROLE**: Primary

This example entry is for a primary Admin Node that is on all three networks:

```
[DC1-ADM1]
ADMIN_ROLE = Primary
NODE_TYPE = VM_Admin_Node
TEMPORARY_PASSWORD_TYPE = Use custom password
CUSTOM_TEMPORARY_PASSWORD = Passw0rd

GRID_NETWORK_IP = 10.1.0.2
ADMIN_NETWORK_IP = 10.3.0.2
CLIENT_NETWORK_IP = 10.4.0.2
```

The following additional setting is optional for the primary Admin Node:

- **DISK:** By default, Admin Nodes are assigned two additional 200 GB hard disks for audit and database use. You can increase these settings using the DISK parameter. For example:

```
DISK = INSTANCES=2, CAPACITY=300
```



For Admin nodes, INSTANCES must always equal 2.

Storage Node

The following additional setting is required for Storage Nodes:

- **NODE_TYPE:** VM_Storage_Node

This example entry is for a Storage Node that is on the Grid and Admin Networks, but not on the Client Network. This node uses the ADMIN_IP setting to specify the primary Admin Node's IP address on the Grid Network.

```
[DC1-S1]
NODE_TYPE = VM_Storage_Node

GRID_NETWORK_IP = 10.1.0.3
ADMIN_NETWORK_IP = 10.3.0.3

ADMIN_IP = 10.1.0.2
```

This second example entry is for a Storage Node on a Client Network where the customer's enterprise networking policy states that an S3 client application is only permitted to access the Storage Node using either port 80 or 443. The example configuration file uses PORT_REMAP to enable the Storage Node to send and receive S3 messages on port 443.

```
[DC2-S1]
NODE_TYPE = VM_Storage_Node

GRID_NETWORK_IP = 10.1.1.3
CLIENT_NETWORK_IP = 10.4.1.3
PORT_REMAP = client/tcp/18082/443

ADMIN_IP = 10.1.0.2
```

The last example creates a symmetric remapping for ssh traffic from port 22 to port 3022, but explicitly sets the values for both inbound and outbound traffic.

```
[DC1-S3]
NODE_TYPE = VM_Storage_Node

GRID_NETWORK_IP = 10.1.1.3

PORT_REMAP = grid/tcp/22/3022
PORT_REMAP_INBOUND = grid/tcp/3022/22

ADMIN_IP = 10.1.0.2
```

The following additional settings are optional for Storage Nodes:

- **DISK:** By default, Storage Nodes are assigned three 4 TB disks for RangeDB use. You can increase these settings with the DISK parameter. For example:

```
DISK = INSTANCES=16, CAPACITY=4096
```

- **STORAGE_TYPE:** By default, all new Storage Nodes are configured to store both object data and metadata, known as a *combined* Storage Node. You can change the Storage Node type to store only data or metadata with the STORAGE_TYPE parameter. For example:

```
STORAGE_TYPE = data
```

Gateway Node

The following additional setting is required for Gateway Nodes:

- **NODE_TYPE:** VM_API_Gateway

This example entry is for an example Gateway Node on all three networks. In this example, no Client Network parameters were specified in the global section of the configuration file, so they must be specified for the node:

```
[DC1-G1]
NODE_TYPE = VM_API_Gateway

GRID_NETWORK_IP = 10.1.0.5
ADMIN_NETWORK_IP = 10.3.0.5

CLIENT_NETWORK_CONFIG = STATIC
CLIENT_NETWORK_TARGET = SG Client Network
CLIENT_NETWORK_MASK = 255.255.255.0
CLIENT_NETWORK_GATEWAY = 10.4.0.1
CLIENT_NETWORK_IP = 10.4.0.5

ADMIN_IP = 10.1.0.2
```

Non-primary Admin Node

The following additional settings are required for non-primary Admin Nodes:

- **NODE_TYPE:** VM_Admin_Node
- **ADMIN_ROLE:** Non-Primary

This example entry is for a non-primary Admin Node that is not on the Client Network:

```
[DC2-ADM1]
ADMIN_ROLE = Non-Primary
NODE_TYPE = VM_Admin_Node

GRID_NETWORK_TARGET = SG Grid Network
GRID_NETWORK_IP = 10.1.0.6
ADMIN_NETWORK_IP = 10.3.0.6

ADMIN_IP = 10.1.0.2
```

The following additional setting is optional for non-primary Admin Nodes:

- **DISK:** By default, Admin Nodes are assigned two additional 200 GB hard disks for audit and database use. You can increase these settings using the DISK parameter. For example:

```
DISK = INSTANCES=2, CAPACITY=300
```



For Admin nodes, INSTANCES must always equal 2.

Run the Bash script

You can use the `deploy-vsphere-ovftool.sh` Bash script and the `deploy-vsphere-ovftool.ini` configuration

file you modified to automate the deployment of StorageGRID nodes in VMware vSphere.

Before you begin

You have created a `deploy-vsphere-ovftool.ini` configuration file for your environment.

You can use the help available with the Bash script by entering the help commands (`-h/--help`). For example:

```
./deploy-vsphere-ovftool.sh -h
```

or

```
./deploy-vsphere-ovftool.sh --help
```

Steps

1. Log in to the Linux machine you are using to run the Bash script.
2. Change to the directory where you extracted the installation archive.

For example:

```
cd StorageGRID-Webscale-version/vsphere
```

3. To deploy all grid nodes, run the Bash script with the appropriate options for your environment.

For example:

```
./deploy-vsphere-ovftool.sh --username=user --password=pwd ./deploy-vsphere-ovftool.ini
```

4. If a grid node failed to deploy because of an error, resolve the error and rerun the Bash script for only that node.

For example:

```
./deploy-vsphere-ovftool.sh --username=user --password=pwd --single -node="DC1-S3" ./deploy-vsphere-ovftool.ini
```

The deployment is complete when the status for each node is "Passed."

Deployment Summary

node	attempts	status
DC1-ADM1	1	Passed
DC1-G1	1	Passed
DC1-S1	1	Passed
DC1-S2	1	Passed
DC1-S3	1	Passed

Automate the configuration of StorageGRID

After deploying the grid nodes, you can automate the configuration of the StorageGRID system.

Before you begin

- You know the location of the following files from the installation archive.

Filename	Description
<code>configure-storagegrid.py</code>	Python script used to automate the configuration
<code>configure-storagegrid.sample.json</code>	Example configuration file for use with the script
<code>configure-storagegrid.blank.json</code>	Blank configuration file for use with the script

- You have created a `configure-storagegrid.json` configuration file. To create this file, you can modify the example configuration file (`configure-storagegrid.sample.json`) or the blank configuration file (`configure-storagegrid.blank.json`).



Store the management password and provisioning passphrase from the passwords section of the modified `configure-storagegrid.json` configuration file in a secure location. These passwords are required for installation, expansion, and maintenance procedures. You should also back up the modified `configure-storagegrid.json` configuration file and store it in a secure location.

About this task

You can use the `configure-storagegrid.py` Python script and the `configure-storagegrid.json` grid configuration file to automate the configuration of your StorageGRID system.



You can also configure the system using the Grid Manager or the Installation API.

Steps

- Log in to the Linux machine you are using to run the Python script.
- Change to the directory where you extracted the installation archive.

For example:

```
cd StorageGRID-Webscale-version/platform
```

where `platform` is `debs`, `rpms`, or `vsphere`.

3. Run the Python script and use the configuration file you created.

For example:

```
./configure-storagegrid.py ./configure-storagegrid.json --start-install
```

Result

A Recovery Package `.zip` file is generated during the configuration process, and it is downloaded to the directory where you are running the installation and configuration process. You must back up the Recovery Package file so that you can recover the StorageGRID system if one or more grid nodes fails. For example, copy it to a secure, backed up network location and to a secure cloud storage location.



The Recovery Package file must be secured because it contains encryption keys and passwords that can be used to obtain data from the StorageGRID system.

If you specified that random passwords should be generated, open the `Passwords.txt` file and look for the passwords required to access your StorageGRID system.

```
#####  
##### The StorageGRID "Recovery Package" has been downloaded as: #####  
#####      ./sgws-recovery-package-994078-rev1.zip      #####  
##### Safeguard this file as it will be needed in case of a #####  
#####      StorageGRID node recovery. #####  
#####
```

Your StorageGRID system is installed and configured when a confirmation message is displayed.

```
StorageGRID has been configured and installed.
```

Related information

- [Navigate to the Grid Manager](#)
- [Installation REST API](#)

Deploy virtual grid nodes

Collect information about your deployment environment (VMware)

Before deploying grid nodes, you must collect information about your network configuration and VMware environment.



It is more efficient to perform one single installation of all the nodes, rather than installing some nodes now and some nodes later.

VMware information

You must access the deployment environment and collect information about the VMware environment; the networks that were created for the Grid, Admin, and Client Networks; and the storage volume types you plan to use for Storage Nodes.

You must collect information about your VMware environment, including the following:

- The username and password for a VMware vSphere account that has appropriate permissions to complete the deployment.
- Host, datastore, and network configuration information for each StorageGRID node virtual machine.



VMware live vMotion causes the virtual machine clock time to jump and is not supported for grid nodes of any type. Though rare, incorrect clock times can result in loss of data or configuration updates.

Grid Network information

You must collect information about the VMware network created for the StorageGRID Grid Network (required), including:

- The network name.
- The method used to assign IP addresses, either static or DHCP.
 - If you are using static IP addresses, the required networking details for each grid node (IP address, gateway, network mask).
 - If you are using DHCP, the IP address of the primary Admin Node on the Grid Network. See [How grid nodes discover the primary Admin Node](#) for more information.

Admin Network information

For nodes that will be connected to the optional StorageGRID Admin Network, you must collect information about the VMware network created for this network, including:

- The network name.
- The method used to assign IP addresses, either static or DHCP.
 - If you are using static IP addresses, the required networking details for each grid node (IP address, gateway, network mask).
 - If you are using DHCP, the IP address of the primary Admin Node on the Grid Network. See [How grid nodes discover the primary Admin Node](#) for more information.
- The external subnet list (ESL) for the Admin Network.

Client Network information

For nodes that will be connected to the optional StorageGRID Client Network, you must collect information about the VMware network created for this network, including:

- The network name.

- The method used to assign IP addresses, either static or DHCP.
- If you are using static IP addresses, the required networking details for each grid node (IP address, gateway, network mask).

Information about additional interfaces

You can optionally add trunk or access interfaces to the VM in vCenter after you install the node. For example, you might want to add a trunk interface to an Admin or Gateway Node, so you can use VLAN interfaces to segregate the traffic belonging to different applications or tenants. Or, you might want to add an access interface to use in a high availability (HA) group.

The interfaces you add are displayed on the VLAN interfaces page and on the HA groups page in the Grid Manager.

- If you add a trunk interface, configure one or more VLAN interfaces for each new parent interface. See [configure VLAN interfaces](#).
- If you add an access interface, you must add it directly to HA groups. See [configure high availability groups](#).

Storage volumes for virtual Storage Nodes

You must collect the following information for virtual machine-based Storage Nodes:

- The number and size of storage volumes (storage LUNs) you plan to add. See [Storage and performance requirements](#).

Grid configuration information

You must collect information to configure your grid:

- Grid license
- Network Time Protocol (NTP) server IP addresses
- DNS server IP addresses

Create node configuration files for Linux deployments

Node configuration files are small text files that provide the information the StorageGRID host service needs to start a node and connect it to the appropriate network and block storage resources. Node configuration files are used for virtual nodes and aren't used for appliance nodes.



"Linux" refers to a RHEL, Ubuntu, or Debian deployment. For a list of supported versions, see the [NetApp Interoperability Matrix Tool \(IMT\)](#).

Location for node configuration files

Place the configuration file for each StorageGRID node in the `/etc/storagegrid/nodes` directory on the host where the node will run. For example, if you plan to run one Admin Node, one Gateway Node, and one Storage Node on HostA, you must place three node configuration files in `/etc/storagegrid/nodes` on HostA.

You can create the configuration files directly on each host using a text editor, such as vim or nano, or you can

create them elsewhere and move them to each host.

Naming of node configuration files

The names of the configuration files are significant. The format is `node-name.conf`, where `node-name` is a name you assign to the node. This name appears in the StorageGRID Installer and is used for node maintenance operations, such as node migration.

Node names must follow these rules:

- Must be unique
- Must start with a letter
- Can contain the characters A through Z and a through z
- Can contain the numbers 0 through 9
- Can contain one or more hyphens (-)
- Must be no more than 32 characters, not including the `.conf` extension

Any files in `/etc/storagegrid/nodes` that don't follow these naming conventions will not be parsed by the host service.

If you have a multi-site topology planned for your grid, a typical node naming scheme might be:

```
site-nodetype-nodenum.conf
```

For example, you might use `dc1-adm1.conf` for the first Admin Node in Data Center 1, and `dc2-sn3.conf` for the third Storage Node in Data Center 2. However, you can use any scheme you like, as long as all node names follow the naming rules.

Contents of a node configuration file

A configuration file contains key/value pairs, with one key and one value per line. For each key/value pair, follow these rules:

- The key and the value must be separated by an equal sign (=) and optional whitespace.
- The keys can contain no spaces.
- The values can contain embedded spaces.
- Any leading or trailing whitespace is ignored.

The following table defines the values for all supported keys. Each key has one of the following designations:

- **Required:** Required for every node or for the specified node types
- **Best practice:** Optional, although recommended
- **Optional:** Optional for all nodes

Admin Network keys

ADMIN_IP

Value	Designation
<p>Grid Network IPv4 address of the Admin Node you want to use to install the Linux-based node. For recovery, use the IP of the primary Admin Node if available; otherwise, use the IP of a non-primary Admin node. If you omit this parameter, the node attempts to discover a primary Admin Node using mDNS.</p> <p>How grid nodes discover the primary Admin Node</p> <p>Note: This value is ignored, and might be prohibited, on the primary Admin Node.</p>	Best practice

ADMIN_NETWORK_CONFIG

Value	Designation
DHCP, STATIC, or DISABLED	Optional

ADMIN_NETWORK_ESL

Value	Designation
<p>Comma-separated list of subnets in CIDR notation to which this node should communicate using the Admin Network gateway.</p> <p>Example: 172.16.0.0/21,172.17.0.0/21</p>	Optional

ADMIN_NETWORK_GATEWAY

Value	Designation
<p>IPv4 address of the local Admin Network gateway for this node. Must be on the subnet defined by ADMIN_NETWORK_IP and ADMIN_NETWORK_MASK. This value is ignored for DHCP-configured networks.</p> <p>Examples:</p> <p>1.1.1.1</p> <p>10.224.4.81</p>	Required if ADMIN_NETWORK_ESL is specified. Optional otherwise.

ADMIN_NETWORK_IP

Value	Designation
<p>IPv4 address of this node on the Admin Network. This key is only required when ADMIN_NETWORK_CONFIG = STATIC; don't specify it for other values.</p> <p>Examples:</p> <p>1.1.1.1</p> <p>10.224.4.81</p>	<p>Required when ADMIN_NETWORK_CONFIG = STATIC.</p> <p>Optional otherwise.</p>

ADMIN_NETWORK_MAC

Value	Designation
<p>The MAC address for the Admin Network interface in the container.</p> <p>This field is optional. If omitted, a MAC address will be generated automatically.</p> <p>Must be 6 pairs of hexadecimal digits separated by colons.</p> <p>Example: b2:9c:02:c2:27:10</p>	<p>Optional</p>

ADMIN_NETWORK_MASK

Value	Designation
<p>IPv4 netmask for this node, on the Admin Network. Specify this key when ADMIN_NETWORK_CONFIG = STATIC; don't specify it for other values.</p> <p>Examples:</p> <p>255.255.255.0</p> <p>255.255.248.0</p>	<p>Required if ADMIN_NETWORK_IP is specified and ADMIN_NETWORK_CONFIG = STATIC.</p> <p>Optional otherwise.</p>

ADMIN_NETWORK_MTU

Value	Designation
<p>The maximum transmission unit (MTU) for this node on the Admin Network. Don't specify if ADMIN_NETWORK_CONFIG = DHCP. If specified, the value must be between 1280 and 9216. If omitted, 1500 is used.</p> <p>If you want to use jumbo frames, set the MTU to a value suitable for jumbo frames, such as 9000. Otherwise, keep the default value.</p> <p>IMPORTANT: The MTU value of the network must match the value configured on the switch port the node is connected to. Otherwise, network performance issues or packet loss might occur.</p> <p>Examples:</p> <p>1500</p> <p>8192</p>	Optional

ADMIN_NETWORK_TARGET

Value	Designation
<p>Name of the host device that you will use for Admin Network access by the StorageGRID node. Only network interface names are supported. Typically, you use a different interface name than what was specified for GRID_NETWORK_TARGET or CLIENT_NETWORK_TARGET.</p> <p>Note: Don't use bond or bridge devices as the network target. Either configure a VLAN (or other virtual interface) on top of the bond device, or use a bridge and virtual Ethernet (veth) pair.</p> <p>Best practice: Specify a value even if this node will not initially have an Admin Network IP address. Then you can add an Admin Network IP address later, without having to reconfigure the node on the host.</p> <p>Examples:</p> <p>bond0.1002</p> <p>ens256</p>	Best practice

ADMIN_NETWORK_TARGET_TYPE

Value	Designation
Interface (This is the only supported value.)	Optional

ADMIN_NETWORK_TARGET_TYPE_INTERFACE_CLONE_MAC

Value	Designation
<p>True or False</p> <p>Set the key to "true" to cause the StorageGRID container use the MAC address of the host host target interface on the Admin Network.</p> <p>Best practice: In networks where promiscuous mode would be required, use the ADMIN_NETWORK_TARGET_TYPE_INTERFACE_CLONE_MAC key instead.</p> <p>For more details on MAC cloning for Linux, see Considerations and recommendations for MAC address cloning</p>	<p>Best practice</p>

ADMIN_ROLE

Value	Designation
<p>Primary or non-primary</p> <p>This key is only required when NODE_TYPE = VM_Admin_Node; don't specify it for other node types.</p>	<p>Required when NODE_TYPE = VM_Admin_Node</p> <p>Optional otherwise.</p>

Block device keys

BLOCK_DEVICE_AUDIT_LOGS

Value	Designation
<p>Path and name of the block device special file this node will use for persistent storage of audit logs.</p> <p>Examples:</p> <pre>/dev/disk/by-path/pci-0000:03:00.0-scsi-0:0:0:0</pre> <pre>/dev/disk/by-id/wwn-0x600a09800059d6df000060d757b475fd</pre> <pre>/dev/mapper/sgws-adm1-audit-logs</pre>	<p>Required for nodes with NODE_TYPE = VM_Admin_Node. Don't specify it for other node types.</p>

BLOCK_DEVICE_RANGEDB_nnn

Value	Designation
<p>Path and name of the block device special file this node will use for persistent object storage. This key is only required for nodes with <code>NODE_TYPE = VM_Storage_Node</code>; don't specify it for other node types.</p>	<p>Required: BLOCK_DEVICE_RANGEDB_000</p>
<p>Only BLOCK_DEVICE_RANGEDB_000 is required; the rest are optional. The block device specified for BLOCK_DEVICE_RANGEDB_000 must be at least 4 TB; the others can be smaller.</p>	<p>Optional: BLOCK_DEVICE_RANGEDB_001</p>
<p>Don't leave gaps. If you specify BLOCK_DEVICE_RANGEDB_005, you must also specify BLOCK_DEVICE_RANGEDB_004.</p>	<p>BLOCK_DEVICE_RANGEDB_002 BLOCK_DEVICE_RANGEDB_003</p>
<p>Note: For compatibility with existing deployments, two-digit keys are supported for upgraded nodes.</p>	<p>BLOCK_DEVICE_RANGEDB_004 BLOCK_DEVICE_RANGEDB_005</p>
<p>Examples:</p>	<p>BLOCK_DEVICE_RANGEDB_006</p>
<pre data-bbox="131 730 1036 762">/dev/disk/by-path/pci-0000:03:00.0-scsi-0:0:0:0</pre>	<p>BLOCK_DEVICE_RANGEDB_007</p>
<pre data-bbox="131 804 1036 877">/dev/disk/by-id/wwn-0x600a09800059d6df000060d757b475fd</pre>	<p>BLOCK_DEVICE_RANGEDB_008</p>
<pre data-bbox="131 909 1036 940">/dev/mapper/sgws-sn1-rangedb-000</pre>	<p>BLOCK_DEVICE_RANGEDB_009</p>
	<p>BLOCK_DEVICE_RANGEDB_010</p>
	<p>BLOCK_DEVICE_RANGEDB_011</p>
	<p>BLOCK_DEVICE_RANGEDB_012</p>
	<p>BLOCK_DEVICE_RANGEDB_013</p>
	<p>BLOCK_DEVICE_RANGEDB_014</p>
	<p>BLOCK_DEVICE_RANGEDB_015</p>

BLOCK_DEVICE_TABLES

Value	Designation
<p>Path and name of the block device special file this node will use for persistent storage of database tables. This key is only required for nodes with NODE_TYPE = VM_Admin_Node; don't specify it for other node types.</p> <p>Examples:</p> <pre>/dev/disk/by-path/pci-0000:03:00.0-scsi-0:0:0:0</pre> <pre>/dev/disk/by-id/wwn-0x600a09800059d6df000060d757b475fd</pre> <pre>/dev/mapper/sgws-adml-tables</pre>	Required

BLOCK_DEVICE_VAR_LOCAL

Value	Designation
<p>Path and name of the block device special file this node will use for its /var/local persistent storage.</p> <p>Examples:</p> <pre>/dev/disk/by-path/pci-0000:03:00.0-scsi-0:0:0:0</pre> <pre>/dev/disk/by-id/wwn-0x600a09800059d6df000060d757b475fd</pre> <pre>/dev/mapper/sgws-sn1-var-local</pre>	Required

Client Network keys

CLIENT_NETWORK_CONFIG

Value	Designation
DHCP, STATIC, or DISABLED	Optional

CLIENT_NETWORK_GATEWAY

Value	Designation

<p>IPv4 address of the local Client Network gateway for this node, which must be on the subnet defined by CLIENT_NETWORK_IP and CLIENT_NETWORK_MASK. This value is ignored for DHCP-configured networks.</p> <p>Examples:</p> <p>1.1.1.1</p> <p>10.224.4.81</p>	Optional
---	----------

CLIENT_NETWORK_IP

Value	Designation
<p>IPv4 address of this node on the Client Network.</p> <p>This key is only required when CLIENT_NETWORK_CONFIG = STATIC; don't specify it for other values.</p> <p>Examples:</p> <p>1.1.1.1</p> <p>10.224.4.81</p>	<p>Required when CLIENT_NETWORK_CONFIG = STATIC</p> <p>Optional otherwise.</p>

CLIENT_NETWORK_MAC

Value	Designation
<p>The MAC address for the Client Network interface in the container.</p> <p>This field is optional. If omitted, a MAC address will be generated automatically.</p> <p>Must be 6 pairs of hexadecimal digits separated by colons.</p> <p>Example: b2:9c:02:c2:27:20</p>	Optional

CLIENT_NETWORK_MASK

Value	Designation
<p>IPv4 netmask for this node on the Client Network.</p> <p>Specify this key when CLIENT_NETWORK_CONFIG = STATIC; don't specify it for other values.</p> <p>Examples:</p> <p>255.255.255.0</p> <p>255.255.248.0</p>	<p>Required if CLIENT_NETWORK_IP is specified and CLIENT_NETWORK_CONFIG = STATIC</p> <p>Optional otherwise.</p>

CLIENT_NETWORK_MTU

Value	Designation
<p>The maximum transmission unit (MTU) for this node on the Client Network. Don't specify if CLIENT_NETWORK_CONFIG = DHCP. If specified, the value must be between 1280 and 9216. If omitted, 1500 is used.</p> <p>If you want to use jumbo frames, set the MTU to a value suitable for jumbo frames, such as 9000. Otherwise, keep the default value.</p> <p>IMPORTANT: The MTU value of the network must match the value configured on the switch port the node is connected to. Otherwise, network performance issues or packet loss might occur.</p> <p>Examples:</p> <p>1500</p> <p>8192</p>	<p>Optional</p>

CLIENT_NETWORK_TARGET

Value	Designation
<p>Name of the host device that you will use for Client Network access by the StorageGRID node. Only network interface names are supported. Typically, you use a different interface name than what was specified for GRID_NETWORK_TARGET or ADMIN_NETWORK_TARGET.</p> <p>Note: Don't use bond or bridge devices as the network target. Either configure a VLAN (or other virtual interface) on top of the bond device, or use a bridge and virtual Ethernet (veth) pair.</p> <p>Best practice: Specify a value even if this node will not initially have a Client Network IP address. Then you can add a Client Network IP address later, without having to reconfigure the node on the host.</p> <p>Examples:</p> <pre>bond0.1003</pre> <pre>ens423</pre>	Best practice

CLIENT_NETWORK_TARGET_TYPE

Value	Designation
Interface (This is only supported value.)	Optional

CLIENT_NETWORK_TARGET_TYPE_INTERFACE_CLONE_MAC

Value	Designation
<p>True or False</p> <p>Set the key to "true" to cause the StorageGRID container to use the MAC address of the host target interface on the Client Network.</p> <p>Best practice: In networks where promiscuous mode would be required, use the CLIENT_NETWORK_TARGET_TYPE_INTERFACE_CLONE_MAC key instead.</p> <p>For more details on MAC cloning for Linux, see Considerations and recommendations for MAC address cloning</p>	Best practice

Grid Network keys

GRID_NETWORK_CONFIG

Value	Designation
STATIC or DHCP Defaults to STATIC if not specified.	Best practice

GRID_NETWORK_GATEWAY

Value	Designation
IPv4 address of the local Grid Network gateway for this node, which must be on the subnet defined by GRID_NETWORK_IP and GRID_NETWORK_MASK. This value is ignored for DHCP-configured networks. If the Grid Network is a single subnet with no gateway, use either the standard gateway address for the subnet (X.Y.Z.1) or this node's GRID_NETWORK_IP value; either value will simplify potential future Grid Network expansions.	Required

GRID_NETWORK_IP

Value	Designation
IPv4 address of this node on the Grid Network. This key is only required when GRID_NETWORK_CONFIG = STATIC; don't specify it for other values. Examples: 1.1.1.1 10.224.4.81	Required when GRID_NETWORK_CONFIG = STATIC Optional otherwise.

GRID_NETWORK_MAC

Value	Designation
The MAC address for the Grid Network interface in the container. Must be 6 pairs of hexadecimal digits separated by colons. Example: b2:9c:02:c2:27:30	Optional If omitted, a MAC address will be generated automatically.

GRID_NETWORK_MASK

Value	Designation
<p>IPv4 netmask for this node on the Grid Network. Specify this key when GRID_NETWORK_CONFIG = STATIC; don't specify it for other values.</p> <p>Examples:</p> <p>255.255.255.0</p> <p>255.255.248.0</p>	<p>Required when GRID_NETWORK_IP is specified and GRID_NETWORK_CONFIG = STATIC.</p> <p>Optional otherwise.</p>

GRID_NETWORK_MTU

Value	Designation
<p>The maximum transmission unit (MTU) for this node on the Grid Network. Don't specify if GRID_NETWORK_CONFIG = DHCP. If specified, the value must be between 1280 and 9216. If omitted, 1500 is used.</p> <p>If you want to use jumbo frames, set the MTU to a value suitable for jumbo frames, such as 9000. Otherwise, keep the default value.</p> <p>IMPORTANT: The MTU value of the network must match the value configured on the switch port the node is connected to. Otherwise, network performance issues or packet loss might occur.</p> <p>IMPORTANT: For the best network performance, all nodes should be configured with similar MTU values on their Grid Network interfaces. The Grid Network MTU mismatch alert is triggered if there is a significant difference in MTU settings for the Grid Network on individual nodes. The MTU values don't have to be the same for all network types.</p> <p>Examples:</p> <p>1500</p> <p>8192</p>	<p>Optional</p>

GRID_NETWORK_TARGET

Value	Designation
<p>Name of the host device that you will use for Grid Network access by the StorageGRID node. Only network interface names are supported. Typically, you use a different interface name than what was specified for ADMIN_NETWORK_TARGET or CLIENT_NETWORK_TARGET.</p> <p>Note: Don't use bond or bridge devices as the network target. Either configure a VLAN (or other virtual interface) on top of the bond device, or use a bridge and virtual Ethernet (veth) pair.</p> <p>Examples:</p> <pre>bond0.1001</pre> <pre>ens192</pre>	Required

GRID_NETWORK_TARGET_TYPE

Value	Designation
Interface (This is the only supported value.)	Optional

GRID_NETWORK_TARGET_TYPE_INTERFACE_CLONE_MAC

Value	Designation
<p>True or False</p> <p>Set the value of the key to "true" to cause the StorageGRID container to use the MAC address of the host target interface on the Grid Network.</p> <p>Best practice: In networks where promiscuous mode would be required, use the GRID_NETWORK_TARGET_TYPE_INTERFACE_CLONE_MAC key instead.</p> <p>For more details on MAC cloning for Linux, see Considerations and recommendations for MAC address cloning</p>	Best practice

Installation password key (temporary)

CUSTOM_TEMPORARY_PASSWORD_HASH

Value	Designation
<p>For the primary Admin Node, set a default temporary password for the StorageGRID Installation API during installation.</p> <p>Note: Set an installation password on the primary Admin Node only. If you attempt to set a password on another node type, validation of the node configuration file will fail.</p> <p>Setting this value has no effect when installation has completed.</p> <p>If this key is omitted, by default no temporary password is set. Alternatively, you can set a temporary password using the StorageGRID Installation API.</p> <p>Must be a <code>crypt()</code> SHA-512 password hash with format <code>\$6\$<salt>\$<password hash></code> for a password of at least 8 and no more than 32 characters.</p> <p>This hash can be generated using CLI tools, such as the <code>openssl passwd</code> command in SHA-512 mode.</p>	Best practice

Interfaces key

INTERFACE_TARGET_nnnn

Value	Designation
<p>Name and optional description for an extra interface you want to add to this node. You can add multiple extra interfaces to each node.</p> <p>For <i>nnnn</i>, specify a unique number for each INTERFACE_TARGET entry you are adding.</p> <p>For the value, specify the name of the physical interface on the bare-metal host. Then, optionally, add a comma and provide a description of the interface, which is displayed on the VLAN interfaces page and the HA groups page.</p> <p>Example: <code>INTERFACE_TARGET_0001=ens256, Trunk</code></p> <p>If you add a trunk interface, you must configure a VLAN interface in StorageGRID. If you add an access interface, you can add the interface directly to an HA group; you don't need to configure a VLAN interface.</p>	Optional

Maximum RAM key

MAXIMUM_RAM

Value	Designation
<p>The maximum amount of RAM that this node is allowed to consume. If this key is omitted, the node has no memory restrictions. When setting this field for a production-level node, specify a value that is at least 24 GB and 16 to 32 GB less than the total system RAM.</p> <p>Note: The RAM value affects a node's actual metadata reserved space. See the description of what Metadata Reserved Space is.</p> <p>The format for this field is <i>numberunit</i>, where <i>unit</i> can be b, k, m, or g.</p> <p>Examples:</p> <p>24g</p> <p>38654705664b</p> <p>Note: If you want to use this option, you must enable kernel support for memory cgroups.</p>	Optional

Node type keys

NODE_TYPE

Value	Designation
<p>Type of node:</p> <ul style="list-style-type: none"> • VM_Admin_Node • VM_Storage_Node • VM_Archive_Node • VM_API_Gateway 	Required

STORAGE_TYPE

Value	Designation
<p>Defines the type of objects a Storage Node contains. For more information, see Types of Storage Nodes. This key is only required for nodes with <code>NODE_TYPE = VM_Storage_Node</code>; don't specify it for other node types. Storage types:</p> <ul style="list-style-type: none"> • combined • data • metadata <p>Note: If the <code>STORAGE_TYPE</code> is not specified, the Storage Node type is set to combined (data and metadata) by default.</p>	Optional

Port remap keys



Support for port remapping is deprecated and will be removed in a future release. To remove remapped ports, refer to [Remove port remaps on bare metal hosts](#).

PORT_REMAP

Value	Designation
<p>Remaps any port used by a node for internal grid node communications or external communications. Remapping ports is necessary if enterprise networking policies restrict one or more ports used by StorageGRID, as described in Internal grid node communications or External communications.</p> <p>IMPORTANT: Don't remap the ports you are planning to use to configure load balancer endpoints.</p> <p>Note: If only <code>PORT_REMAP</code> is set, the mapping that you specify is used for both inbound and outbound communications. If <code>PORT_REMAP_INBOUND</code> is also specified, <code>PORT_REMAP</code> applies only to outbound communications.</p> <p>The format used is: <i>network type/protocol/default port used by grid node/new port</i>, where <i>network type</i> is grid, admin, or client, and <i>protocol</i> is tcp or udp.</p> <p>Example: <code>PORT_REMAP = client/tcp/18082/443</code></p> <p>You can also remap multiple ports using a comma-separated list.</p> <p>Example: <code>PORT_REMAP = client/tcp/18082/443, client/tcp/18083/80</code></p>	Optional

PORT_REMAP_INBOUND

Value	Designation
<p>Remaps inbound communications to the specified port. If you specify <code>PORT_REMAP_INBOUND</code> but don't specify a value for <code>PORT_REMAP</code>, outbound communications for the port are unchanged.</p> <p>IMPORTANT: Don't remap the ports you are planning to use to configure load balancer endpoints.</p> <p>The format used is: <i>network type/protocol/remapped port /default port used by grid node</i>, where <i>network type</i> is <code>grid</code>, <code>admin</code>, or <code>client</code>, and <i>protocol</i> is <code>tcp</code> or <code>udp</code>.</p> <p>Example: <code>PORT_REMAP_INBOUND = grid/tcp/3022/22</code></p> <p>You can also remap multiple inbound ports using a comma-separated list.</p> <p>Example: <code>PORT_REMAP_INBOUND = grid/tcp/3022/22, admin/tcp/3022/22</code></p>	Optional

How grid nodes discover the primary Admin Node

Grid nodes communicate with the primary Admin Node for configuration and management. Each grid node must know the IP address of the primary Admin Node on the Grid Network.

To ensure that a grid node can access the primary Admin Node, you can do either of the following when deploying the node:

- You can use the `ADMIN_IP` parameter to enter the primary Admin Node's IP address manually.
- You can omit the `ADMIN_IP` parameter to have the grid node discover the value automatically. Automatic discovery is especially useful when the Grid Network uses DHCP to assign the IP address to the primary Admin Node.

Automatic discovery of the primary Admin Node is accomplished using a multicast domain name system (mDNS). When the primary Admin Node first starts up, it publishes its IP address using mDNS. Other nodes on the same subnet can then query for the IP address and acquire it automatically. However, because multicast IP traffic is not normally routable across subnets, nodes on other subnets can't acquire the primary Admin Node's IP address directly.

If you use automatic discovery:



- You must include the `ADMIN_IP` setting for at least one grid node on any subnets that the primary Admin Node is not directly attached to. This grid node will then publish the primary Admin Node's IP address for other nodes on the subnet to discover with mDNS.
- Ensure that your network infrastructure supports passing multi-cast IP traffic within a subnet.

Deploy a StorageGRID node as a virtual machine (VMware)

You use VMware vSphere Web Client to deploy each grid node as a virtual machine. During deployment, each grid node is created and connected to one or more StorageGRID networks.

If you need to deploy any StorageGRID appliance Storage Nodes, see [Deploy appliance Storage Node](#).

Optionally, you can remap node ports or increase CPU or memory settings for the node before powering it on.

Before you begin

- You have reviewed how to [plan and prepare for installation](#), and you understand the requirements for software, CPU and RAM, and storage and performance.
- You are familiar with VMware vSphere Hypervisor and have experience deploying virtual machines in this environment.



The `open-vm-tools` package, an open-source implementation similar to VMware Tools, is included with the StorageGRID virtual machine. You don't need to install VMware Tools manually.

- You have downloaded and extracted the correct version of the StorageGRID installation archive for VMware.



If you are deploying the new node as part of an expansion or recovery operation, you must use the version of StorageGRID that is currently running on the grid.

- You have the StorageGRID Virtual Machine Disk (`.vmdk`) file:

```
NetApp-SG-version-SHA.vmdk
```

- You have the `.ovf` and `.mf` files for each type of grid node you are deploying:

Filename	Description
<code>vsphere-primary-admin.ovf</code> <code>vsphere-primary-admin.mf</code>	The template file and manifest file for the primary Admin Node.
<code>vsphere-non-primary-admin.ovf</code> <code>vsphere-non-primary-admin.mf</code>	The template file and manifest file for a non-primary Admin Node.
<code>vsphere-storage.ovf</code> <code>vsphere-storage.mf</code>	The template file and manifest file for a Storage Node.
<code>vsphere-gateway.ovf</code> <code>vsphere-gateway.mf</code>	The template file and manifest file for a Gateway Node.

- The `.vmdk`, `.ovf`, and `.mf` files are all in the same directory.

- You have a plan to minimize failure domains. For example, you should not deploy all Gateway Nodes on a single vSphere ESXi host.



In a production deployment, don't run more than one Storage Node on a single virtual machine. Do not run multiple virtual machines on the same ESXi host if that would create an unacceptable failure-domain issue.

- If you are deploying a node as part of an expansion or recovery operation, you have the [instructions for expanding a StorageGRID system](#) or the [recovery and maintenance instructions](#).
- If you are deploying a StorageGRID node as a virtual machine with storage assigned from a NetApp ONTAP system, you have confirmed that the volume does not have a FabricPool tiering policy enabled. For example, if a StorageGRID node is running as an virtual machine on a VMware host, ensure the volume backing the datastore for the node does not have a FabricPool tiering policy enabled. Disabling FabricPool tiering for volumes used with StorageGRID nodes simplifies troubleshooting and storage operations.



Never use FabricPool to tier any data related to StorageGRID back to StorageGRID itself. Tiering StorageGRID data back to StorageGRID increases troubleshooting and operational complexity.

About this task

Follow these instructions to initially deploy VMware nodes, add a new VMware node in an expansion, or replace a VMware node as part of a recovery operation. Except as noted in the steps, the node deployment procedure is the same for all node types, including Admin Nodes, Storage Nodes, and Gateway Nodes.

If you are installing a new StorageGRID system:

- You can deploy nodes in any order.
- You must ensure that each virtual machine can connect to the primary Admin Node over the Grid Network.
- You must deploy all grid nodes before configuring the grid.

If you are performing an expansion or recovery operation:

- You must ensure that the new virtual machine can connect to all other nodes over the Grid Network.

If you need to remap any of the node's ports, don't power on the new node until the port remap configuration is complete.



Support for port remapping is deprecated and will be removed in a future release. To remove remapped ports, refer to [Remove port remaps on bare metal hosts](#).

Steps

1. Using VCenter, deploy an OVF template.

If you specify a URL, point to a folder containing the following files. Otherwise, select each of these files from a local directory.

```
NetApp-SG-version-SHA.vmdk  
vsphere-node.ovf  
vsphere-node.mf
```

For example, if this is the first node you are deploying, use these files to deploy the primary Admin Node for your StorageGRID system:

```
NetApp-SG-version-SHA.vmdk
vsphere-primary-admin.ovf
vsphere-primary-admin.mf
```

2. Provide a name for the virtual machine.

The standard practice is to use the same name for both the virtual machine and the grid node.

3. Place the virtual machine in the appropriate vApp or resource pool.

4. If you are deploying the primary Admin Node, read and accept the End User License Agreement.

Depending on your version of vCenter, the order of the steps will vary for accepting the End User License Agreement, specifying the name of the virtual machine, and selecting a datastore.

5. Select storage for the virtual machine.

If you are deploying a node as part of recovery operation, perform the instructions in the [storage recovery step](#) to add new virtual disks, reattach virtual hard disks from the failed grid node, or both.

When deploying a Storage Node, use 3 or more storage volumes, with each storage volume being 4 TB or larger. You must assign at least 4 TB to volume 0.



The Storage Node .ovf file defines several VMDKs for storage. Unless these VMDKs meet your storage requirements, you should remove them and assign appropriate VMDKs or RDMs for storage before powering up the node. VMDKs are more commonly used in VMware environments and are easier to manage, while RDMs might provide better performance for workloads that use larger object sizes (for example, greater than 100 MB).



Some StorageGRID installations might use larger, more active storage volumes than typical virtualized workloads. You might need to tune some hypervisor parameters, such as `MaxAddressableSpaceTB`, to achieve optimal performance. If you encounter poor performance, contact your virtualization support resource to determine whether your environment could benefit from workload-specific configuration tuning.

6. Select networks.

Determine which StorageGRID networks the node will use by selecting a destination network for each source network.

- The Grid Network is required. You must select a destination network in the vSphere environment.
+
The Grid Network is used for all internal StorageGRID traffic. It provides connectivity among all nodes in the grid, across all sites and subnets. All nodes on the Grid Network must be able to communicate with all other nodes.
- If you use the Admin Network, select a different destination network in the vSphere environment. If you don't use the Admin Network, select the same destination you selected for the Grid Network.
- If you use the Client Network, select a different destination network in the vSphere environment. If you

don't use the Client Network, select the same destination you selected for the Grid Network.

- If you use an Admin or Client network, nodes do not have to be on the same Admin or Client networks.

7. For **Customize Template**, configure the required StorageGRID node properties.

a. Enter the **Node name**.



If you are recovering a grid node, you must enter the name of the node you are recovering.

b. Use the **Temporary installation password** drop-down to specify a temporary installation password, so that you can access the VM console or the StorageGRID Installation API, or use SSH, before the new node joins the grid.



The temporary installation password is only used during node installation. After a node has been added to the grid, you can access it using the [node console password](#), which is listed in the `Passwords.txt` file in the Recovery Package.

- **Use node name:** The value you provided for the **Node name** field is used as the temporary installation password.
- **Use custom password:** A custom password is used as the temporary installation password.
- **Disable password:** No temporary installation password will be used. If you need to access the VM to debug installation issues, see [Troubleshoot installation issues](#).

c. If you selected **Use custom password**, specify the temporary installation password you want to use in the **Custom password** field.

d. In the **Grid Network (eth0)** section, select STATIC or DHCP for the **Grid network IP configuration**.

- If you select STATIC, enter the **Grid network IP**, **Grid network mask**, **Grid network gateway**, and **Grid network MTU**.
- If you select DHCP, the **Grid network IP**, **Grid network mask**, and **Grid network gateway** are automatically assigned.

e. In the **Primary Admin IP** field, enter the IP address of the primary Admin Node for the Grid Network.



This step does not apply if the node you are deploying is the primary Admin Node.

If you omit the primary Admin Node IP address, the IP address will be automatically discovered if the primary Admin Node, or at least one other grid node with `ADMIN_IP` configured, is present on the same subnet. However, it is recommended to set the primary Admin Node IP address here.

f. In the **Admin Network (eth1)** section, select STATIC, DHCP, or DISABLED for the **Admin network IP configuration**.

- If you don't want to use the Admin Network, select DISABLED and enter **0.0.0.0** for the Admin Network IP. You can leave the other fields blank.
- If you select STATIC, enter the **Admin network IP**, **Admin network mask**, **Admin network gateway**, and **Admin network MTU**.
- If you select STATIC, enter the **Admin network external subnet list**. You must also configure a gateway.
- If you select DHCP, the **Admin network IP**, **Admin network mask**, and **Admin network gateway** are automatically assigned.

- g. In the **Client Network (eth2)** section, select STATIC, DHCP, or DISABLED for the **Client network IP configuration**.
 - If you don't want to use the Client Network, select DISABLED and enter **0.0.0.0** for the Client Network IP. You can leave the other fields blank.
 - If you select STATIC, enter the **Client network IP**, **Client network mask**, **Client network gateway**, and **Client network MTU**.
 - If you select DHCP, the **Client network IP**, **Client network mask**, and **Client network gateway** are automatically assigned.
8. Review the virtual machine configuration and make any changes necessary.
9. When you are ready to complete, select **Finish** to start the upload of the virtual machine.
10. If you deployed this node as part of recovery operation and this is not a full-node recovery, perform these steps after deployment is complete:
 - a. Right-click the virtual machine, and select **Edit Settings**.
 - b. Select each default virtual hard disk that has been designated for storage, and select **Remove**.
 - c. Depending on your data recovery circumstances, add new virtual disks according to your storage requirements, reattach any virtual hard disks preserved from the previously removed failed grid node, or both.

Note the following important guidelines:

- If you are adding new disks you should use the same type of storage device that was in use before node recovery.
 - The Storage Node .ovf file defines several VMDKs for storage. Unless these VMDKs meet your storage requirements, you should remove them and assign appropriate VMDKs or RDMs for storage before powering up the node. VMDKs are more commonly used in VMware environments and are easier to manage, while RDMs might provide better performance for workloads that use larger object sizes (for example, greater than 100 MB).
11. If you need to remap the ports used by this node, follow these steps.

You might need to remap a port if your enterprise networking policies restrict access to one or more ports that are used by StorageGRID. See the [networking guidelines](#) for the ports used by StorageGRID.



Don't remap the ports used in load balancer endpoints.

- a. Select the new VM.
- b. From the Configure tab, select **Settings > vApp Options**. The location of **vApp Options** depends on the version of vCenter.
- c. In the **Properties** table, locate PORT_REMAP_INBOUND and PORT_REMAP.
- d. To symmetrically map both inbound and outbound communications for a port, select **PORT_REMAP**.



Support for port remapping is deprecated and will be removed in a future release. To remove remapped ports, refer to [Remove port remaps on bare metal hosts](#).



If only PORT_REMAP is set, the mapping that you specify applies to both inbound and outbound communications. If PORT_REMAP_INBOUND is also specified, PORT_REMAP applies only to outbound communications.

i. Select **Set Value**.

ii. Enter the port mapping:

```
<network type>/<protocol>/<default port used by grid node>/<new port>
```

<network type> is grid, admin, or client, and <protocol> is tcp or udp.

For example, to remap ssh traffic from port 22 to port 3022, enter:

```
client/tcp/22/3022
```

You can remap multiple ports using a comma-separated list.

For example:

```
client/tcp/18082/443, client/tcp/18083/80
```

iii. Select **OK**.

e. To specify the port used for inbound communications to the node, select **PORT_REMAP_INBOUND**.



If you specify **PORT_REMAP_INBOUND** and don't specify a value for **PORT_REMAP**, outbound communications for the port are unchanged.

i. Select **Set Value**.

ii. Enter the port mapping:

```
<network type>/<protocol>/<remapped inbound port>/<default inbound port used by grid node>
```

<network type> is grid, admin, or client, and <protocol> is tcp or udp.

For example, to remap inbound SSH traffic that is sent to port 3022 so that it is received at port 22 by the grid node, enter the following:

```
client/tcp/3022/22
```

You can remap multiple inbound ports using a comma-separated list.

For example:

```
grid/tcp/3022/22, admin/tcp/3022/22
```

iii. Select **OK**

12. If you want to increase the CPU or memory for the node from the default settings:

a. Right-click the virtual machine, and select **Edit Settings**.

b. Change the number of CPUs or the amount of memory as required.

Set the **Memory Reservation** to the same size as the **Memory** allocated to the virtual machine.

c. Select **OK**.

13. Power on the virtual machine.

After you finish

If you deployed this node as part of an expansion or recovery procedure, return to those instructions to complete the procedure.

Example node configuration files (Linux)

You can use the example node configuration files to help set up the node configuration files for your StorageGRID system. The examples show node configuration files for all types of grid nodes.



"Linux" refers to a RHEL, Ubuntu, or Debian deployment. For a list of supported versions, see the [NetApp Interoperability Matrix Tool \(IMT\)](#).

For most nodes, you can add Admin and Client Network addressing information (IP, mask, gateway, and so on) when you configure the grid using the Grid Manager or the Installation API. The exception is the primary Admin Node. If you want to browse to the Admin Network IP of the primary Admin Node to complete grid configuration (because the Grid Network is not routed, for example), you must configure the Admin Network connection for the primary Admin Node in its node configuration file. This is shown in the example.



In the examples, the Client Network target has been configured as a best practice, even though the Client Network is disabled by default.

Example for primary Admin Node

Example file name: `/etc/storagegrid/nodes/dcl-adm1.conf`

Example file contents:

```

NODE_TYPE = VM_Admin_Node
ADMIN_ROLE = Primary
TEMPORARY_PASSWORD_TYPE = Use custom password
CUSTOM_TEMPORARY_PASSWORD = Passw0rd
BLOCK_DEVICE_VAR_LOCAL = /dev/mapper/dc1-adm1-var-local
BLOCK_DEVICE_AUDIT_LOGS = /dev/mapper/dc1-adm1-audit-logs
BLOCK_DEVICE_TABLES = /dev/mapper/dc1-adm1-tables
GRID_NETWORK_TARGET = bond0.1001
ADMIN_NETWORK_TARGET = bond0.1002
CLIENT_NETWORK_TARGET = bond0.1003

GRID_NETWORK_IP = 10.1.0.2
GRID_NETWORK_MASK = 255.255.255.0
GRID_NETWORK_GATEWAY = 10.1.0.1

ADMIN_NETWORK_CONFIG = STATIC
ADMIN_NETWORK_IP = 192.168.100.2
ADMIN_NETWORK_MASK = 255.255.248.0
ADMIN_NETWORK_GATEWAY = 192.168.100.1
ADMIN_NETWORK_ESL = 192.168.100.0/21,172.16.0.0/21,172.17.0.0/21

```

Example for Storage Node

Example file name: /etc/storagegrid/nodes/dc1-sn1.conf

Example file contents:

```

NODE_TYPE = VM_Storage_Node
ADMIN_IP = 10.1.0.2
BLOCK_DEVICE_VAR_LOCAL = /dev/mapper/dc1-sn1-var-local
BLOCK_DEVICE_RANGEDB_00 = /dev/mapper/dc1-sn1-rangedb-0
BLOCK_DEVICE_RANGEDB_01 = /dev/mapper/dc1-sn1-rangedb-1
BLOCK_DEVICE_RANGEDB_02 = /dev/mapper/dc1-sn1-rangedb-2
BLOCK_DEVICE_RANGEDB_03 = /dev/mapper/dc1-sn1-rangedb-3
GRID_NETWORK_TARGET = bond0.1001
ADMIN_NETWORK_TARGET = bond0.1002
CLIENT_NETWORK_TARGET = bond0.1003

GRID_NETWORK_IP = 10.1.0.3
GRID_NETWORK_MASK = 255.255.255.0
GRID_NETWORK_GATEWAY = 10.1.0.1

```

Example for Gateway Node

Example file name: /etc/storagegrid/nodes/dc1-gw1.conf

Example file contents:

```
NODE_TYPE = VM_API_Gateway
ADMIN_IP = 10.1.0.2
BLOCK_DEVICE_VAR_LOCAL = /dev/mapper/dc1-gw1-var-local
GRID_NETWORK_TARGET = bond0.1001
ADMIN_NETWORK_TARGET = bond0.1002
CLIENT_NETWORK_TARGET = bond0.1003
GRID_NETWORK_IP = 10.1.0.5
GRID_NETWORK_MASK = 255.255.255.0
GRID_NETWORK_GATEWAY = 10.1.0.1
```

Example for a non-primary Admin Node

Example file name: /etc/storagegrid/nodes/dc1-adm2.conf

Example file contents:

```
NODE_TYPE = VM_Admin_Node
ADMIN_ROLE = Non-Primary
ADMIN_IP = 10.1.0.2
BLOCK_DEVICE_VAR_LOCAL = /dev/mapper/dc1-adm2-var-local
BLOCK_DEVICE_AUDIT_LOGS = /dev/mapper/dc1-adm2-audit-logs
BLOCK_DEVICE_TABLES = /dev/mapper/dc1-adm2-tables
GRID_NETWORK_TARGET = bond0.1001
ADMIN_NETWORK_TARGET = bond0.1002
CLIENT_NETWORK_TARGET = bond0.1003

GRID_NETWORK_IP = 10.1.0.6
GRID_NETWORK_MASK = 255.255.255.0
GRID_NETWORK_GATEWAY = 10.1.0.1
```

Validate the StorageGRID configuration (Linux)

After creating configuration files in /etc/storagegrid/nodes for each of your StorageGRID nodes, you must validate the contents of those files.



"Linux" refers to a RHEL, Ubuntu, or Debian deployment. For a list of supported versions, see the [NetApp Interoperability Matrix Tool \(IMT\)](#).

To validate the contents of the configuration files, run the following command on each host:

```
sudo storagegrid node validate all
```

If the files are correct, the output shows **PASSED** for each configuration file, as shown in the example.



When using only one LUN on metadata-only nodes, you might receive a warning message that can be ignored.

```
Checking for misnamed node configuration files... PASSED
Checking configuration file for node dcl-adml... PASSED
Checking configuration file for node dcl-gw1... PASSED
Checking configuration file for node dcl-sn1... PASSED
Checking configuration file for node dcl-sn2... PASSED
Checking configuration file for node dcl-sn3... PASSED
Checking for duplication of unique values between nodes... PASSED
```



For an automated installation, you can suppress this output by using the `-q` or `--quiet` options in the `storagegrid` command (for example, `storagegrid --quiet...`). If you suppress the output, the command will have a non-zero exit value if any configuration warnings or errors were detected.

If the configuration files are incorrect, the issues are shown as **WARNING** and **ERROR**, as shown in the example. If any configuration errors are found, you must correct them before you continue with the installation.

```

Checking for misnamed node configuration files...
WARNING: ignoring /etc/storagegrid/nodes/dc1-adml
WARNING: ignoring /etc/storagegrid/nodes/dc1-sn2.conf.keep
WARNING: ignoring /etc/storagegrid/nodes/my-file.txt
Checking configuration file for node dc1-adml...
ERROR: NODE_TYPE = VM_Foo_Node
      VM_Foo_Node is not a valid node type.  See *.conf.sample
ERROR: ADMIN_ROLE = Foo
      Foo is not a valid admin role.  See *.conf.sample
ERROR: BLOCK_DEVICE_VAR_LOCAL = /dev/mapper/sgws-gw1-var-local
      /dev/mapper/sgws-gw1-var-local is not a valid block device
Checking configuration file for node dc1-gw1...
ERROR: GRID_NETWORK_TARGET = bond0.1001
      bond0.1001 is not a valid interface.  See `ip link show`
ERROR: GRID_NETWORK_IP = 10.1.3
      10.1.3 is not a valid IPv4 address
ERROR: GRID_NETWORK_MASK = 255.248.255.0
      255.248.255.0 is not a valid IPv4 subnet mask
Checking configuration file for node dc1-sn1...
ERROR: GRID_NETWORK_GATEWAY = 10.2.0.1
      10.2.0.1 is not on the local subnet
ERROR: ADMIN_NETWORK_ESL = 192.168.100.0/21,172.16.0foo
      Could not parse subnet list
Checking configuration file for node dc1-sn2... PASSED
Checking configuration file for node dc1-sn3... PASSED
Checking for duplication of unique values between nodes...
ERROR: GRID_NETWORK_IP = 10.1.0.4
      dc1-sn2 and dc1-sn3 have the same GRID_NETWORK_IP
ERROR: BLOCK_DEVICE_VAR_LOCAL = /dev/mapper/sgws-sn2-var-local
      dc1-sn2 and dc1-sn3 have the same BLOCK_DEVICE_VAR_LOCAL
ERROR: BLOCK_DEVICE_RANGEDB_00 = /dev/mapper/sgws-sn2-rangedb-0
      dc1-sn2 and dc1-sn3 have the same BLOCK_DEVICE_RANGEDB_00

```

Start the StorageGRID host service (Linux)

To start your StorageGRID nodes, and ensure they restart after a host reboot, you must enable and start the StorageGRID host service.



"Linux" refers to a RHEL, Ubuntu, or Debian deployment. For a list of supported versions, see the [NetApp Interoperability Matrix Tool \(IMT\)](#).

Steps

1. Run the following commands on each host:

```
sudo systemctl enable storagegrid
sudo systemctl start storagegrid
```

2. Run the following command to ensure the deployment is proceeding:

```
sudo storagegrid node status node-name
```

3. If any node returns a status of "Not Running" or "Stopped," run the following command:

```
sudo storagegrid node start node-name
```

4. If you have previously enabled and started the StorageGRID host service (or if you are unsure if the service has been enabled and started), also run the following command:

```
sudo systemctl reload-or-restart storagegrid
```

Installation REST API

StorageGRID provides the StorageGRID Installation API for performing installation tasks.

The API uses the Swagger open source API platform to provide the API documentation. Swagger allows both developers and non-developers to interact with the API in a user interface that illustrates how the API responds to parameters and options. This documentation assumes that you are familiar with standard web technologies and the JSON data format.



Any API operations you perform using the API Documentation webpage are live operations. Be careful not to create, update, or delete configuration data or other data by mistake.

Each REST API command includes the API's URL, an HTTP action, any required or optional URL parameters, and an expected API response.

StorageGRID Installation API

The StorageGRID Installation API is only available when you are initially configuring your StorageGRID system, and if you need to perform a primary Admin Node recovery. The Installation API can be accessed over HTTPS from the Grid Manager.

To access the API documentation, go to the installation web page on the primary Admin Node and select **Help > API documentation** from the menu bar.

The StorageGRID Installation API includes the following sections:

- **config** — Operations related to the product release and versions of the API. You can list the product release version and the major versions of the API supported by that release.
- **grid** — Grid-level configuration operations. You can get and update grid settings, including grid details, Grid

Network subnets, grid passwords, and NTP and DNS server IP addresses.

- **nodes** — Node-level configuration operations. You can retrieve a list of grid nodes, delete a grid node, configure a grid node, view a grid node, and reset a grid node's configuration.
- **provision** — Provisioning operations. You can start the provisioning operation and view the status of the provisioning operation.
- **recovery** — Primary Admin Node recovery operations. You can reset information, upload the Recover Package, start the recovery, and view the status of the recovery operation.
- **recovery-package** — Operations to download the Recovery Package.
- **sites** — Site-level configuration operations. You can create, view, delete, and modify a site.
- **temporary-password** — Operations on the temporary password to secure the mgmt-api during installation.

Troubleshoot installation issues

If any problems occur while installing your StorageGRID system, you can access the installation log files. Technical support might also need to use the installation log files to resolve issues.



"Linux" refers to a RHEL, Ubuntu, or Debian deployment. For a list of supported versions, see the [NetApp Interoperability Matrix Tool \(IMT\)](#).

Linux

The following installation log files are available from the container that is running each node:

- `/var/local/log/install.log` (found on all grid nodes)
- `/var/local/log/gdu-server.log` (found on the primary Admin Node)

The following installation log files are available from the host:

- `/var/log/storagegrid/daemon.log`
- `/var/log/storagegrid/nodes/node-name.log`

To learn how to access the log files, see [Collect log files and system data](#).

VMware

The following are the main installation log files, which technical support might need to resolve issues.

- `/var/local/log/install.log` (found on all grid nodes)
- `/var/local/log/gdu-server.log` (found on the primary Admin Node)

Virtual machine resource reservation requires adjustment

OVF files include a resource reservation designed to ensure that each grid node has sufficient RAM and CPU to operate efficiently. If you create virtual machines by deploying these OVF files on VMware and the predefined number of resources aren't available, the virtual machines will not start.

About this task

If you are certain that the VM host has sufficient resources for each grid node, manually adjust the resources allocated for each virtual machine, and then try starting the virtual machines.

Steps

1. In the VMware vSphere Hypervisor client tree, select the virtual machine that is not started.
2. Right-click the virtual machine, and select **Edit Settings**.
3. From the Virtual Machines Properties window, select the **Resources** tab.
4. Adjust the resources allocated to the virtual machine:
 - a. Select **CPU**, and then use the Reservation slider to adjust the MHz reserved for this virtual machine.
 - b. Select **Memory**, and then use the Reservation slider to adjust the MB reserved for this virtual machine.
5. Click **OK**.
6. Repeat as required for other virtual machines hosted on the same VM host.

Temporary installation password was disabled

When you deploy a VMware node, you can optionally specify a temporary installation password. You must have this password to access the VM console or use SSH before the new node joins the grid.

If you opted to disable the temporary installation password, you must perform additional steps to debug installation issues.

You can do either of the following:

- Redeploy the VM but specify a temporary installation password so you can access the console or use SSH to debug installation issues.
- Use vCenter to set the password:
 1. Power off the VM.
 2. Go to **VM**, select the **Configure** tab, and select **vApp Options**.
 3. Specify type of temporary installation password to set:
 - Select **CUSTOM_TEMPORARY_PASSWORD** to set a custom temporary password.
 - Select **TEMPORARY_PASSWORD_TYPE** to use the node name as the temporary password.
 4. Select **Set Value**.
 5. Set the temporary password:
 - Change **CUSTOM_TEMPORARY_PASSWORD** to a custom password value.
 - Update the **TEMPORARY_PASSWORD_TYPE** with the **Use node name** value.
 6. Restart the VM to apply the new password.

Related information

- To learn how to access the log files, see [Log files reference](#).
- [Troubleshoot a StorageGRID system](#)
- If you need additional help, contact [NetApp Support](#).

Script examples

Example `/etc/sysconfig/network-scripts` (RHEL)

You can use the example files to aggregate four Linux physical interfaces into a single LACP bond and then establish three VLAN interfaces subtending the bond for use as StorageGRID Grid, Admin, and Client Network interfaces.

Physical interfaces

Note that the switches at the other ends of the links must also treat the four ports as a single LACP trunk or port channel, and must pass at least the three referenced VLANs with tags.

`/etc/sysconfig/network-scripts/ifcfg-ens160`

```
TYPE=Ethernet
NAME=ens160
UUID=011b17dd-642a-4bb9-acae-d71f7e6c8720
DEVICE=ens160
ONBOOT=yes
MASTER=bond0
SLAVE=yes
```

/etc/sysconfig/network-scripts/ifcfg-ens192

```
TYPE=Ethernet
NAME=ens192
UUID=e28eb15f-76de-4e5f-9a01-c9200b58d19c
DEVICE=ens192
ONBOOT=yes
MASTER=bond0
SLAVE=yes
```

/etc/sysconfig/network-scripts/ifcfg-ens224

```
TYPE=Ethernet
NAME=ens224
UUID=b0e3d3ef-7472-4cde-902c-ef4f3248044b
DEVICE=ens224
ONBOOT=yes
MASTER=bond0
SLAVE=yes
```

/etc/sysconfig/network-scripts/ifcfg-ens256

```
TYPE=Ethernet
NAME=ens256
UUID=7cf7aabc-3e4b-43d0-809a-1e2378faa4cd
DEVICE=ens256
ONBOOT=yes
MASTER=bond0
SLAVE=yes
```

Bond interface

/etc/sysconfig/network-scripts/ifcfg-bond0

```
DEVICE=bond0
TYPE=Bond
BONDING_MASTER=yes
NAME=bond0
ONBOOT=yes
BONDING_OPTS=mode=802.3ad
```

VLAN interfaces

/etc/sysconfig/network-scripts/ifcfg-bond0.1001

```
VLAN=yes
TYPE=Vlan
DEVICE=bond0.1001
PHYSDEV=bond0
VLAN_ID=1001
REORDER_HDR=0
BOOTPROTO=none
UUID=296435de-8282-413b-8d33-c4dd40fca24a
ONBOOT=yes
```

/etc/sysconfig/network-scripts/ifcfg-bond0.1002

```
VLAN=yes
TYPE=Vlan
DEVICE=bond0.1002
PHYSDEV=bond0
VLAN_ID=1002
REORDER_HDR=0
BOOTPROTO=none
UUID=dbaaec72-0690-491c-973a-57b7dd00c581
ONBOOT=yes
```

/etc/sysconfig/network-scripts/ifcfg-bond0.1003

```
VLAN=yes
TYPE=Vlan
DEVICE=bond0.1003
PHYSDEV=bond0
VLAN_ID=1003
REORDER_HDR=0
BOOTPROTO=none
UUID=d1af4b30-32f5-40b4-8bb9-71a2fbf809a1
ONBOOT=yes
```

Example /etc/network/interfaces (Ubuntu and Debian)

The `/etc/network/interfaces` file includes three sections, which define the physical interfaces, bond interface, and VLAN interfaces. You can combine the three example sections into a single file, which will aggregate four Linux physical interfaces into a single LACP bond and then establish three VLAN interfaces subtending the bond for use as

StorageGRID Grid, Admin, and Client Network interfaces.

Physical interfaces

Note that the switches at the other ends of the links must also treat the four ports as a single LACP trunk or port channel, and must pass at least the three referenced VLANs with tags.

```
# loopback interface
auto lo
iface lo inet loopback

# ens160 interface
auto ens160
iface ens160 inet manual
    bond-master bond0
    bond-primary en160

# ens192 interface
auto ens192
iface ens192 inet manual
    bond-master bond0

# ens224 interface
auto ens224
iface ens224 inet manual
    bond-master bond0

# ens256 interface
auto ens256
iface ens256 inet manual
    bond-master bond0
```

Bond interface

```
# bond0 interface
auto bond0
iface bond0 inet manual
    bond-mode 4
    bond-miimon 100
    bond-slaves ens160 ens192 ens224 ens256
```

VLAN interfaces

```
# 1001 vlan
auto bond0.1001
iface bond0.1001 inet manual
vlan-raw-device bond0

# 1002 vlan
auto bond0.1002
iface bond0.1002 inet manual
vlan-raw-device bond0

# 1003 vlan
auto bond0.1003
iface bond0.1003 inet manual
vlan-raw-device bond0
```

Configure grid and complete installation

Navigate to the Grid Manager

You use the Grid Manager to define all of the information required to configure your StorageGRID system.

Before you begin

The primary Admin Node must be deployed and have completed the initial startup sequence.

Steps

1. Open your web browser and navigate to:

```
https://primary_admin_node_ip
```

Alternatively, you can access the Grid Manager on port 8443:

```
https://primary_admin_node_ip:8443
```

You can use the IP address for the primary Admin Node IP on the Grid Network or on the Admin Network, as appropriate for your network configuration. You might need to use the security/advanced option in your browser to navigate to an untrusted certificate.

2. Manage a temporary installer password as needed:
 - If a password has already been set using one of these methods, enter the password to proceed.
 - A user set the password while accessing the installer previously
 - For Linux, the password was automatically imported from the node config file at `/etc/storagegrid/nodes/<node_name>.conf`
 - For VMware, the SSH/console password was automatically imported from the OVF properties
 - If a password has not been set, optionally set a password to secure the StorageGRID installer.

3. Select **Install a StorageGRID system**.

The page used to configure a StorageGRID grid appears.

The screenshot shows the NetApp StorageGRID installation wizard interface. At the top, there is a blue header with "NetApp® StorageGRID®" and a "Help" dropdown. Below the header is a navigation bar with "Install" selected. A progress indicator shows eight steps: 1. License (highlighted in blue), 2. Sites, 3. Grid Network, 4. Grid Nodes, 5. NTP, 6. DNS, 7. Passwords, and 8. Summary. The main content area is titled "License" and contains the instruction: "Enter a grid name and upload the license file provided by NetApp for your StorageGRID system." There are two input fields: "Grid Name" with an empty text box, and "License File" with a "Browse" button.

Specify the StorageGRID license information

You must specify the name for your StorageGRID system and upload the license file provided by NetApp.

Steps

1. On the License page, enter a meaningful name for your StorageGRID system in the **Grid Name** field.

After installation, the name is displayed at the top of the Nodes menu.

2. Select **Browse**, locate the NetApp license file (*NLF-unique-id.txt*), and select **Open**.

The license file is validated, and the serial number is displayed.



The StorageGRID installation archive includes a free license that does not provide any support entitlement for the product. You can update to a license that offers support after installation.

This screenshot shows the same NetApp StorageGRID installation wizard interface as the previous one, but with the fields filled in. The "Grid Name" field contains "StorageGRID". The "License File" field has a "Browse" button and the filename "NLF-959007-Internal.txt" displayed next to it. The "License Serial Number" field contains "959007". The progress indicator remains the same, with step 1 highlighted.

3. Select **Next**.


Add sites

You must create at least one site when you are installing StorageGRID. You can create additional sites to increase the reliability and storage capacity of your StorageGRID system.

Steps

1. On the Sites page, enter the **Site Name**.
2. To add additional sites, click the plus sign next to the last site entry and enter the name in the new **Site Name** text box.

Add as many additional sites as required for your grid topology. You can add up to 16 sites.



Sites

In a single-site deployment, infrastructure and operations are centralized in one site.

In a multi-site deployment, infrastructure can be distributed asymmetrically across sites, and proportional to the needs of each site. Typically, sites are located in geographically different locations. Having multiple sites also allows the use of distributed replication and erasure coding for increased availability and resiliency.

Site Name 1	<input type="text" value="Raleigh"/>	✕
Site Name 2	<input type="text" value="Atlanta"/>	+ ✕

3. Click **Next**.

Specify Grid Network subnets

You must specify the subnets that are used on the Grid Network.

About this task

The subnet entries include the subnets for the Grid Network for each site in your StorageGRID system, along with any subnets that need to be reachable through the Grid Network.

If you have multiple grid subnets, the Grid Network gateway is required. All grid subnets specified must be reachable through this gateway.

Steps

1. Specify the CIDR network address for at least one Grid Network in the **Subnet 1** text box.
2. Click the plus sign next to the last entry to add an additional network entry. You must specify all subnets for all sites in the Grid Network.

- If you have already deployed at least one node, click **Discover Grid Networks Subnets** to automatically populate the Grid Network Subnet List with the subnets reported by grid nodes that have registered with the Grid Manager.
- You must manually add any subnets for NTP, DNS, LDAP, or other external servers accessed through the Grid Network gateway.

NetApp® StorageGRID® Help ▾

Install

1 License 2 Sites 3 **Grid Network** 4 Grid Nodes 5 NTP 6 DNS 7 Passwords 8 Summary

Grid Network

You must specify the subnets that are used on the Grid Network. These entries typically include the subnets for the Grid Network for each site in your StorageGRID system. Select Discover Grid Networks to automatically add subnets based on the network configuration of all registered nodes.

Note: You must manually add any subnets for NTP, DNS, LDAP, or other external servers accessed through the Grid Network gateway.

Subnet 1 +

3. Click **Next**.

Approve pending grid nodes

You must approve each grid node before it can join the StorageGRID system.

Before you begin

You have deployed all virtual and StorageGRID appliance grid nodes.



It is more efficient to perform one single installation of all the nodes, rather than installing some nodes now and some nodes later.

Steps

1. Review the Pending Nodes list, and confirm that it shows all of the grid nodes you deployed.



If a grid node is missing, confirm that it was deployed successfully and has the correct Grid Network IP of the primary admin node set for ADMIN_IP.

2. Select the radio button next to a pending node you want to approve.



Grid Nodes

Approve and configure grid nodes, so that they are added correctly to your StorageGRID system.

Pending Nodes

Grid nodes are listed as pending until they are assigned to a site, configured, and approved.

+ Approve		✗ Remove		Search		Q			
Grid Network MAC Address	↑↓	Name	↑↓	Type	↑↓	Platform	↑↓	Grid Network IPv4 Address	▼
<input checked="" type="radio"/>	50:6b:4b:42:d7:00	NetApp-SGA	Storage Node	StorageGRID Appliance	172.16.5.20/21				

Approved Nodes

Grid nodes that have been approved and have been configured for installation. An approved grid node's configuration can be edited if errors are identified.

✎ Edit		🔄 Reset		✗ Remove		Search		Q			
Grid Network MAC Address	↑↓	Name	↑↓	Site	↑↓	Type	↑↓	Platform	↑↓	Grid Network IPv4 Address	▼
<input type="radio"/>	00:50:56:87:42:ff	dc1-adm1	Raleigh	Admin Node	VMware VM	172.16.4.210/21					
<input type="radio"/>	00:50:56:87:c0:16	dc1-s1	Raleigh	Storage Node	VMware VM	172.16.4.211/21					
<input type="radio"/>	00:50:56:87:79:ee	dc1-s2	Raleigh	Storage Node	VMware VM	172.16.4.212/21					
<input type="radio"/>	00:50:56:87:db:9c	dc1-s3	Raleigh	Storage Node	VMware VM	172.16.4.213/21					
<input type="radio"/>	00:50:56:87:62:38	dc1-g1	Raleigh	API Gateway Node	VMware VM	172.16.4.214/21					

3. Click **Approve**.

4. In General Settings, modify settings for the following properties, as necessary:

- **Site:** The system name of the site for this grid node.
- **Name:** The system name for the node. The name defaults to the name you specified when you configured the node.

System names are required for internal StorageGRID operations and can't be changed after you complete the installation. However, during this step of the installation process, you can change system names as required.



For a VMware node, you can change the name here, but this action will not change the name of the virtual machine in vSphere.

- **NTP Role:** The Network Time Protocol (NTP) role of the grid node. The options are **Automatic**, **Primary**, and **Client**. Selecting **Automatic** assigns the Primary role to Admin Nodes, Storage Nodes with ADC services, Gateway Nodes, and any grid nodes that have non-static IP addresses. All other grid nodes are assigned the Client role.



Make sure that at least two nodes at each site can access at least four external NTP sources. If only one node at a site can reach the NTP sources, timing issues will occur if that node goes down. In addition, designating two nodes per site as primary NTP sources ensures accurate timing if a site is isolated from the rest of the grid.

- **Storage Type** (Storage Nodes only): Specify that a new Storage Node be used exclusively for data only, metadata only, or both. The options are **Data and metadata** ("combined"), **Data only**, and **Metadata only**.



See [Types of Storage Nodes](#) for information about requirements for these node types.

- **ADC service** (Storage Nodes only): Select **Automatic** to let the system determine whether the node requires the Administrative Domain Controller (ADC) service. The ADC service keeps track of the location and availability of grid services. At least three Storage Nodes at each site must include the ADC service. You can't add the ADC service to a node after it is deployed.

5. In Grid Network, modify settings for the following properties as necessary:

- **IPv4 Address (CIDR)**: The CIDR network address for the Grid Network interface (eth0 inside the container). For example: 192.168.1.234/21
- **Gateway**: The Grid Network gateway. For example: 192.168.0.1

The gateway is required if there are multiple grid subnets.



If you selected DHCP for the Grid Network configuration and you change the value here, the new value will be configured as a static address on the node. You must make sure the configured IP address is not within a DHCP address pool.

6. If you want to configure the Admin Network for the grid node, add or update the settings in the Admin Network section as necessary.

Enter the destination subnets of the routes out of this interface in the **Subnets (CIDR)** text box. If there are multiple Admin subnets, the Admin gateway is required.



If you selected DHCP for the Admin Network configuration and you change the value here, the new value will be configured as a static address on the node. You must make sure the configured IP address is not within a DHCP address pool.

Appliances: For a StorageGRID appliance, if the Admin Network was not configured during the initial installation using the StorageGRID Appliance Installer, it can't be configured in this Grid Manager dialog box. Instead, you must follow these steps:

- a. Reboot the appliance: In the Appliance Installer, select **Advanced > Reboot**.

Rebooting can take several minutes.

- b. Select **Configure Networking > Link Configuration** and enable the appropriate networks.
- c. Select **Configure Networking > IP Configuration** and configure the enabled networks.
- d. Return to the Home page and click **Start Installation**.
- e. In the Grid Manager: If the node is listed in the Approved Nodes table, remove the node.
- f. Remove the node from the Pending Nodes table.

- g. Wait for the node to reappear in the Pending Nodes list.
- h. Confirm that you can configure the appropriate networks. They should already be populated with the information you provided on the IP Configuration page of the Appliance Installer.

For additional information, see the [Quick start for hardware installation](#) to locate instructions for your appliance.

- 7. If you want to configure the Client Network for the grid node, add or update the settings in the Client Network section as necessary. If the Client Network is configured, the gateway is required, and it becomes the default gateway for the node after installation.



If you selected DHCP for the Client Network configuration and you change the value here, the new value will be configured as a static address on the node. You must make sure the configured IP address is not within a DHCP address pool.

Appliances: For a StorageGRID appliance, if the Client Network was not configured during the initial installation using the StorageGRID Appliance Installer, it can't be configured in this Grid Manager dialog box. Instead, you must follow these steps:

- a. Reboot the appliance: In the Appliance Installer, select **Advanced > Reboot**.


Rebooting can take several minutes.

- b. Select **Configure Networking > Link Configuration** and enable the appropriate networks.
- c. Select **Configure Networking > IP Configuration** and configure the enabled networks.
- d. Return to the Home page and click **Start Installation**.
- e. In the Grid Manager: If the node is listed in the Approved Nodes table, remove the node.
- f. Remove the node from the Pending Nodes table.
- g. Wait for the node to reappear in the Pending Nodes list.
- h. Confirm that you can configure the appropriate networks. They should already be populated with the information you provided on the IP Configuration page of the Appliance Installer.

For additional information, see the [Quick start for hardware installation](#) to locate instructions for your appliance.

- 8. Click **Save**.

The grid node entry moves to the Approved Nodes list.



Grid Nodes

Approve and configure grid nodes, so that they are added correctly to your StorageGRID system.

Pending Nodes

Grid nodes are listed as pending until they are assigned to a site, configured, and approved.

+ Approve
✖ Remove

Search Q

Grid Network MAC Address	Name	Type	Platform	Grid Network IPv4 Address
<i>No results found.</i>				

◀
▶

Approved Nodes

Grid nodes that have been approved and have been configured for installation. An approved grid node's configuration can be edited if errors are identified.

✎ Edit
🔄 Reset
✖ Remove

Search Q

	Grid Network MAC Address	Name	Site	Type	Platform	Grid Network IPv4 Address
<input type="radio"/>	00:50:56:87:42:ff	dc1-adm1	Raleigh	Admin Node	VMware VM	172.16.4.210/21
<input type="radio"/>	00:50:56:87:c0:16	dc1-s1	Raleigh	Storage Node	VMware VM	172.16.4.211/21
<input type="radio"/>	00:50:56:87:79:ee	dc1-s2	Raleigh	Storage Node	VMware VM	172.16.4.212/21
<input type="radio"/>	00:50:56:87:db:9c	dc1-s3	Raleigh	Storage Node	VMware VM	172.16.4.213/21
<input type="radio"/>	00:50:56:87:62:38	dc1-g1	Raleigh	API Gateway Node	VMware VM	172.16.4.214/21
<input type="radio"/>	50:6b:4b:42:d7:00	NetApp-SGA	Raleigh	Storage Node	StorageGRID Appliance	172.16.5.20/21

◀
▶

9. Repeat these steps for each pending grid node you want to approve.

You must approve all nodes that you want in the grid. However, you can return to this page at any time before you click **Install** on the Summary page. You can modify the properties of an approved grid node by selecting its radio button and clicking **Edit**.

10. When you are done approving grid nodes, click **Next**.

Specify Network Time Protocol server information

You must specify the Network Time Protocol (NTP) configuration information for the StorageGRID system, so that operations performed on separate servers can be kept synchronized.

About this task

You must specify IPv4 addresses for the NTP servers.

You must specify external NTP servers. The specified NTP servers must use the NTP protocol.

You must specify four NTP server references of Stratum 3 or better to prevent issues with time drift.



When specifying the external NTP source for a production-level StorageGRID installation, don't use the Windows Time (W32Time) service on a version of Windows earlier than Windows Server 2016. The time service on earlier versions of Windows is not sufficiently accurate and is not supported by Microsoft for use in high-accuracy environments, such as StorageGRID.

[Support boundary to configure the Windows Time service for high-accuracy environments](#)

The external NTP servers are used by the nodes to which you previously assigned Primary NTP roles.



Make sure that at least two nodes at each site can access at least four external NTP sources. If only one node at a site can reach the NTP sources, timing issues will occur if that node goes down. In addition, designating two nodes per site as primary NTP sources ensures accurate timing if a site is isolated from the rest of the grid.

(VMware only) Perform additional checks for VMware, such as ensuring that the hypervisor uses the same NTP source as the virtual machine, and using VMTools to disable the time sync between the hypervisor and StorageGRID virtual machines.

Steps

1. Specify the IPv4 addresses for at least four NTP servers in the **Server 1** to **Server 4** text boxes.
2. If necessary, select the plus sign next to the last entry to add additional server entries.

The screenshot shows the NetApp StorageGRID installation wizard interface. At the top, there is a blue header with "NetApp® StorageGRID®" and a "Help" dropdown. Below the header is a progress bar with eight steps: 1. License, 2. Sites, 3. Grid Network, 4. Grid Nodes, 5. NTP (highlighted in blue), 6. DNS, 7. Passwords, and 8. Summary. Below the progress bar, the "Network Time Protocol" section is visible. It contains the instruction: "Enter the IP addresses for at least four Network Time Protocol (NTP) servers, so that operations performed on separate servers are kept in sync." There are four input fields labeled "Server 1" through "Server 4". The values entered are: Server 1: 10.60.248.183, Server 2: 10.227.204.142, Server 3: 10.235.48.111, and Server 4: 0.0.0.0. A plus sign (+) is located to the right of the Server 4 input field.

3. Select **Next**.

Related information

[Networking guidelines](#)

Specify DNS server information

You must specify DNS information for your StorageGRID system, so that you can access external servers using hostnames instead of IP addresses.

About this task

Specifying [DNS server information](#) allows you to use Fully Qualified Domain Name (FQDN) hostnames rather than IP addresses for email notifications and AutoSupport.

To ensure proper operation, specify two or three DNS servers. If you specify more than three, it is possible that only three will be used because of known OS limitations on some platforms. If you have routing restrictions in your environment, you can [customize the DNS server list](#) for individual nodes (typically all nodes at a site) to use a different set of up to three DNS servers.

If possible, use DNS servers that each site can access locally to ensure that an islanded site can resolve the FQDNs for external destinations.

Steps

1. Specify the IPv4 address for at least one DNS server in the **Server 1** text box.
2. If necessary, select the plus sign next to the last entry to add additional server entries.

The screenshot shows the NetApp StorageGRID installation wizard interface. At the top, there is a blue header with the text "NetApp® StorageGRID®" and a "Help" dropdown menu. Below the header is a navigation bar with an "Install" button. A progress indicator consists of eight numbered circles: 1 (License), 2 (Sites), 3 (Grid Network), 4 (Grid Nodes), 5 (NTP), 6 (DNS), 7 (Passwords), and 8 (Summary). The "DNS" step (6) is currently active and highlighted in blue. Below the progress indicator, the section is titled "Domain Name Service". The instructions read: "Enter the IP address for at least one Domain Name System (DNS) server, so that server hostnames can be used instead of IP addresses. Specifying at least two DNS servers is recommended. Configuring DNS enables server connectivity, email notifications, and NetApp AutoSupport." There are two input fields for DNS servers. "Server 1" contains the IP address "10.224.223.130" and has a red "x" icon to its right. "Server 2" contains the IP address "10.224.223.136" and has a red "+ x" icon to its right, indicating that a new server can be added.

The best practice is to specify at least two DNS servers. You can specify up to six DNS servers.

3. Select **Next**.

Specify the StorageGRID system passwords

As part of installing your StorageGRID system, you need to enter the passwords to use to secure your system and perform maintenance tasks.

About this task

Use the Install passwords page to specify the provisioning passphrase and the grid management root user password.

- The provisioning passphrase is used as an encryption key and is not stored by the StorageGRID system.

- You must have the provisioning passphrase for installation, expansion, and maintenance procedures, including downloading the Recovery Package. Therefore, it is important that you store the provisioning passphrase in a secure location.
- You can change the provisioning passphrase from the Grid Manager if you have the current one.
- The grid management root user password can be changed using the Grid Manager.
- Randomly generated command line console and SSH passwords are stored in the `Passwords.txt` file in the Recovery Package.

Steps

1. In **Provisioning Passphrase**, enter the provisioning passphrase that will be required to make changes to the grid topology of your StorageGRID system.

Store the provisioning passphrase in a secure place.



If after the installation completes and you want to change the provisioning passphrase later, you can use the Grid Manager. Select **Configuration > Access control > Grid passwords**.

2. In **Confirm Provisioning Passphrase**, reenter the provisioning passphrase to confirm it.
3. In **Grid Management Root User Password**, enter the password to use to access the Grid Manager as the "root" user.



Store the provisioning passphrase in a secure location. It's required for installation, expansion, and maintenance procedures.

4. In **Confirm Root User Password**, reenter the Grid Manager password to confirm it.



Passwords

Enter secure passwords that meet your organization's security policies. A text file containing the command line passwords must be downloaded during the final installation step.

Provisioning Passphrase	<input type="password"/>
Confirm Provisioning Passphrase	<input type="password"/>
Grid Management Root User Password	<input type="password"/>
Confirm Root User Password	<input type="password"/>

Create random command line passwords.

- If you are installing a grid for proof of concept or demo purposes, optionally clear the **Create random command line passwords** checkbox.

For production deployments, random passwords should always be used for security reasons. Clear **Create random command line passwords** only for demo grids if you want to use default passwords to access grid nodes from the command line using the "root" or "admin" account.



You are prompted to download the Recovery Package file (`sgws-recovery-package-id-revision.zip`) after you click **Install** on the Summary page. You must [download this file](#) to complete the installation. The passwords required to access the system are stored in the `Passwords.txt` file, contained in the Recovery Package file.

- Click **Next**.

Review your configuration and complete installation

You must carefully review the configuration information you have entered to ensure that the installation completes successfully.

Steps

- View the **Summary** page.

NetApp® StorageGRID® Help ▾

Install

1 License 2 Sites 3 Grid Network 4 Grid Nodes 5 NTP 6 DNS 7 Passwords 8 **Summary**

Summary

Verify that all of the grid configuration information is correct, and then click Install. You can view the status of each grid node as it installs. Click the Modify links to go back and change the associated information.

General Settings

Grid Name	Grid1	Modify License
Passwords	Auto-generated random command line passwords	Modify Passwords

Networking

NTP	10.60.248.183 10.227.204.142 10.235.48.111	Modify NTP
DNS	10.224.223.130 10.224.223.136	Modify DNS
Grid Network	172.16.0.0/21	Modify Grid Network

Topology

Topology	Atlanta	Modify Sites	Modify Grid Nodes
	Raleigh		
	dc1-adm1 dc1-g1 dc1-s1 dc1-s2 dc1-s3 NetApp-SGA		

- Verify that all of the grid configuration information is correct. Use the Modify links on the Summary page to go back and correct any errors.

3. Click **Install**.



If a node is configured to use the Client Network, the default gateway for that node switches from the Grid Network to the Client Network when you click **Install**. If you lose connectivity, you must ensure that you are accessing the primary Admin Node through an accessible subnet. See [Networking guidelines](#) for details.

4. Click **Download Recovery Package**.

When the installation progresses to the point where the grid topology is defined, you are prompted to download the Recovery Package file (.zip), and confirm that you can successfully access the contents of this file. You must download the Recovery Package file so that you can recover the StorageGRID system if one or more grid nodes fail. The installation continues in the background, but you can't complete the installation and access the StorageGRID system until you download and verify this file.

5. Verify that you can extract the contents of the .zip file, and then save it in two safe, secure, and separate locations.



The Recovery Package file must be secured because it contains encryption keys and passwords that can be used to obtain data from the StorageGRID system.

6. Select the **I have successfully downloaded and verified the Recovery Package file** checkbox, and click **Next**.

If the installation is still in progress, the status page appears. This page indicates the progress of the installation for each grid node.

Installation Status

If necessary, you may [Download the Recovery Package file](#) again.

Name	Site	Grid Network IPv4 Address	Progress	Stage
dc1-adm1	Site1	172.16.4.215/21	<div style="width: 100%;"></div>	Starting services
dc1-g1	Site1	172.16.4.216/21	<div style="width: 100%;"></div>	Complete
dc1-s1	Site1	172.16.4.217/21	<div style="width: 75%;"></div>	Waiting for Dynamic IP Service peers
dc1-s2	Site1	172.16.4.218/21	<div style="width: 25%;"></div>	Downloading hotfix from primary Admin if needed
dc1-s3	Site1	172.16.4.219/21	<div style="width: 10%;"></div>	Downloading hotfix from primary Admin if needed

When the Complete stage is reached for all grid nodes, the sign-in page for the Grid Manager appears.

7. Sign in to the Grid Manager using the "root" user and the password you specified during the installation.

Post-installation guidelines

After completing grid node deployment and configuration, follow these guidelines for DHCP addressing and network configuration changes.

- If DHCP was used to assign IP addresses, configure a DHCP reservation for each IP address on the networks being used.

You can only set up DHCP during the deployment phase. You can't set up DHCP during configuration.



Nodes reboot when the Grid Network configuration is changed by DHCP, which can cause outages if a DHCP change affects multiple nodes at the same time.

- You must use the Change IP procedures if you want to change IP addresses, subnet masks, and default gateways for a grid node. See [Configure IP addresses](#).
- If you make networking configuration changes, including routing and gateway changes, client connectivity to the primary Admin Node and other grid nodes might be lost. Depending on the networking changes applied, you might need to reestablish these connections.

Where to go next

After completing an installation, perform the required integration and configuration tasks. You can perform the optional tasks as needed.

Required tasks

- (VMware only) Configure VMware vSphere Hypervisor for automatic restart.

You must configure the hypervisor to restart the virtual machines when the server restarts. Without an automatic restart, the virtual machines and grid nodes remain shut down after the server restarts. For details, see the VMware vSphere Hypervisor documentation.

- [Create a tenant account](#) for the S3 client protocol that will be used to store objects on your StorageGRID system.
- [Control system access](#) by configuring groups and user accounts. Optionally, you can [configure a federated identity source](#) (such as Active Directory or OpenLDAP), so you can import administration groups and users. Or, you can [create local groups and users](#).
- Integrate and test the [S3 API](#) client applications you will use to upload objects to your StorageGRID system.
- [Configure the information lifecycle management \(ILM\) rules and ILM policy](#) you want to use to protect object data.
- If your installation includes appliance Storage Nodes, use SANtricity OS to complete the following tasks:
 - Connect to each StorageGRID appliance.
 - Verify receipt of AutoSupport data.

See [Set up hardware](#).

- Review and follow the [StorageGRID system hardening guidelines](#) to eliminate security risks.
- [Configure email notifications for system alerts](#).

Optional tasks

- [Update grid node IP addresses](#) if they have changed since you planned your deployment and generated the recovery package.
- [Configure storage encryption](#), if required.
- [Configure storage compression](#) to reduce the size of stored objects, if required.
- [Configure VLAN interfaces](#) to isolate and partition network traffic, if required.

- [Configure high availability groups](#) to improve connection availability for the Grid Manager, Tenant Manager, and S3 clients, if required.
- [Configure load balancer endpoints](#) for S3 client connectivity, if required.

Copyright information

Copyright © 2026 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.

LIMITED RIGHTS LEGEND: Use, duplication, or disclosure by the government is subject to restrictions as set forth in subparagraph (b)(3) of the Rights in Technical Data -Noncommercial Items at DFARS 252.227-7013 (FEB 2014) and FAR 52.227-19 (DEC 2007).

Data contained herein pertains to a commercial product and/or commercial service (as defined in FAR 2.101) and is proprietary to NetApp, Inc. All NetApp technical data and computer software provided under this Agreement is commercial in nature and developed solely at private expense. The U.S. Government has a non-exclusive, non-transferrable, nonsublicensable, 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) (FEB 2014).

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.