



Install and upgrade

StorageGRID

NetApp
November 04, 2025

This PDF was generated from <https://docs.netapp.com/us-en/storagegrid-117/installconfig/index.html> on November 04, 2025. Always check docs.netapp.com for the latest.

Table of Contents

Install and upgrade	1
Install appliance hardware	1
Quick start for hardware installation	1
Hardware overviews	2
Prepare for installation	27
Automate appliance installation and configuration	72
Automate StorageGRID configuration	77
Overview of installation REST APIs	79
Install hardware	80
Set up hardware	120
Deploy appliance node	167
Install Red Hat Enterprise Linux or CentOS	202
Install Red Hat Enterprise Linux or CentOS: Overview	202
Plan and prepare for Red Hat or CentOS installation	203
Deploy virtual grid nodes (Red Hat or CentOS)	224
Configure the grid and complete installation (Red Hat or CentOS)	244
Automate the installation (Red Hat Enterprise Linux or CentOS)	258
Overview of the installation REST API	260
Where to go next	261
Troubleshoot installation issues	261
Example /etc/sysconfig/network-scripts	262
Install Ubuntu or Debian	264
Install Ubuntu or Debian: Overview	264
Plan and prepare for Ubuntu or Debian installation	265
Deploy virtual grid nodes (Ubuntu or Debian)	287
Configure grid and complete installation (Ubuntu or Debian)	307
Automate the installation (Ubuntu or Debian)	321
Overview of the installation REST API	323
Where to go next	324
Troubleshoot installation issues	325
Example /etc/network/interfaces	325
Install VMware	327
Install VMware: Overview	327
Plan and prepare for VMware installation	328
Deploy virtual machine grid nodes (VMware)	335
Configure the grid and complete installation (VMware)	343
Automate the installation (VMware)	357
Overview of the installation REST API	369
Where to go next	370
Troubleshoot installation issues	371
Upgrade StorageGRID software	372
Upgrade StorageGRID software: Overview	372
What's new in StorageGRID 11.7	372

Removed or deprecated features	376
Changes to the Grid Management API	378
Changes to the Tenant Management API	379
Plan and prepare for upgrade	379
Upgrade software	387
Troubleshoot upgrade issues	394

Install and upgrade

Install appliance hardware

Quick start for hardware installation

Follow these high-level steps to install and set up a StorageGRID appliance and deploy it as a node in your StorageGRID system.

1

Prepare for installation

1. Work with your NetApp Professional Services consultant to automate installation and configuration. See [Automate appliance installation and configuration](#).

This step is optional. However, streamlining and automating configuration steps can save time and provide consistency in the configuration of multiple appliances.

2. [Prepare site](#)
3. [Unpack boxes](#)
4. [Obtain additional equipment and tools](#)
5. [Review web browser requirements](#)
6. [Review appliance network connections](#)
7. [Gather installation information](#)

2

Install hardware

1. [Register hardware](#)
2. Install into cabinet or rack
 - [SGF6112](#)
 - [SG6000](#)
 - [SG5700](#)
 - [SG100 and SG1000](#)
3. Cable appliance
 - [SGF6112](#)
 - [SG6000](#)
 - [SG5700](#)
 - [SG100 and SG1000](#)
4. Connect power cords and apply power
 - [SGF6112](#)
 - [SG6000](#)
 - [SG5700](#)

- SG100 and SG1000

5. [View status indicators and codes](#)

3

Set up hardware

If you are configuring and deploying more than one appliance, use the NetApp ConfigBuilder tool to automate the following configuration and deployment steps. For guidance, contact your NetApp Professional Services consultant. See [Automate appliance installation and configuration](#).

1. Configure StorageGRID connections

- [Access StorageGRID Appliance Installer](#) and verify you are running most recent version
- Configure network links
- Configure StorageGRID IP addresses
- Verify network connections
- Verify port-level network connections

2. [Access and configure SANtricity System Manager \(SG6000 and SG5700\)](#)

3. [Configure the BMC interface \(SGF6112, SG6000, SG100, and SG1000\)](#)

4. Perform optional setup steps

- [Enable node encryption](#)
- [Change RAID mode \(SG6000 and SG5700\)](#)
- [Remap network ports](#)

4

Deploy appliance node

Deploy the appliance as a new node in your StorageGRID system.

- [Deploy appliance Storage Node](#)
- [Deploy services appliance node](#)

Hardware overviews

SGF6112 appliance: Overview

The StorageGRID SGF6112 appliance operates as a Storage Node in a StorageGRID system. The appliance can be used in a hybrid grid environment that combines appliance Storage Nodes and virtual (software-based) Storage Nodes.

The SGF6112 appliance provides the following features:

- 12 NVMe (nonvolatile memory express) SSD drives with integrated compute and storage controllers.
- Integrates the storage and computing elements for a StorageGRID Storage Node.
- Includes the StorageGRID Appliance Installer to simplify Storage Node deployment and configuration.
- Includes a baseboard management controller (BMC) for monitoring and diagnosing the hardware in the compute controller.

- Supports up to four 10-GbE or 25-GbE connections to the StorageGRID Grid Network and Client Network.

SGF6112 hardware description

The StorageGRID SGF6112 is an all-flash appliance that features a compact design with compute controller and storage controller integrated into a 1U chassis. The appliance supports 12 SSD NVMe drives with a storage capacity of up to 15.3 TB per drive.

Resilient object storage

The SGF6112 is designed with SSDs in a RAID that provides the following data protection features:

- Ability to function after the failure of a single SSD with no impact on object availability.
- Ability to function after multiple SSD failures with a minimum necessary reduction in object availability (based on the design of the underlying RAID scheme).
- Fully recoverable, while in service, from SSD failures that don't result in extreme damage to the RAID housing the node's root volume (the StorageGRID operating system).

SGF6112 hardware components

The SGF6112 appliance includes the following components:

Component	Description
Compute and storage platform	<p>A one-rack unit (1U) server that includes:</p> <ul style="list-style-type: none"> • Two 2.1/2.6 GHz 165 W processors providing 48 cores • 256 GB RAM • 2 × 1/10 GBase-T ports • 4 × 10/25 GbE Ethernet ports • 1 × 256 GB Internal Boot drive (includes StorageGRID software) • Baseboard management controller (BMC) that simplifies hardware management • Redundant power supplies and fans

SGF6112 diagrams


SGF6112 front view

This figure shows the front of the SGF6112 without the bezel. The appliance includes a 1U compute and storage platform that contains 12 SSD drives.



SGF6112 rear view

This figure shows the back of the SGF6112, including the ports, fans, and power supplies.



Callout	Port	Type	Use
1	Network ports 1-4	10/25-GbE, based on cable or SFP transceiver type (SFP28 and SFP+ modules are supported), switch speed, and configured link speed.	Connect to the Grid Network and the Client Network for StorageGRID.
2	BMC management port	1-GbE (RJ-45)	Connect to the appliance baseboard management controller.
3	Diagnostic and support ports	<ul style="list-style-type: none"> VGA USB Micro-USB console port Micro-SD slot module 	Reserved for technical support use.
4	Admin Network port 1	1/10-GbE (RJ-45)	Connect the appliance to the Admin Network for StorageGRID.
5	Admin Network port 2	1/10-GbE (RJ-45)	<p>Options:</p> <ul style="list-style-type: none"> Bond with Admin Network port 1 for a redundant connection to the Admin Network for StorageGRID. Leave disconnected and available for temporary local access (IP 169.254.0.1). During installation, use port 2 for IP configuration if DHCP-assigned IP addresses aren't available.

SG6060 and SG6060X appliances: Overview

The StorageGRID SG6060 and SG6060X appliances each include a compute controller

and a storage controller shelf that contains two storage controllers and 60 drives.

Optionally, 60-drive expansion shelves can be added to both appliances. There are no specification or functional differences between the SG6060 and SG6060X except for the location of the interconnect ports on the storage controller.

SG6060 and SG6060X components

The SG6060 and SG6060X appliances include the following components:


Component	Description
Compute controller	SG6000-CN controller, a one-rack unit (1U) server that includes: <ul style="list-style-type: none">• 40 cores (80 threads)• 192 GB RAM• Up to 4 × 25 Gbps aggregate Ethernet bandwidth• 4 × 16 Gbps Fibre Channel (FC) interconnect• Baseboard management controller (BMC) that simplifies hardware management• Redundant power supplies
Storage controller shelf	E-Series E2860 controller shelf (storage array), a 4U shelf that includes: <ul style="list-style-type: none">• Two E2800 series controllers (duplex configuration) to provide storage controller failover support<ul style="list-style-type: none">◦ The SG6060 contains E2800A storage controllers◦ The SG6060X contains E2800B storage controllers• Five-drawer drive shelf that holds sixty 3.5-inch drives (2 solid-state drives, or SSDs, and 58 NL-SAS drives)• Redundant power supplies and fans
Optional: Storage expansion shelves Note: Expansion shelves can be installed during initial deployment or added later.	E-Series DE460C enclosure, a 4U shelf that includes: <ul style="list-style-type: none">• Two input/output modules (IOMs)• Five drawers, each holding 12 NL-SAS drives, for a total of 60 drives• Redundant power supplies and fans Each SG6060 and SG6060X appliance can have one or two expansion shelves for a total of 180 drives (two of these drives are reserved for E-Series read cache).

SG6060 and SG6060X diagrams

The fronts of the SG6060 and SG6060X are identical.

SG6060 or SG6060X front view

This figure shows the front of the SG6060 or SG6060X, which includes a 1U compute controller and a 4U shelf containing two storage controllers and 60 drives in five drive drawers.




Callout	Description
1	SG6000-CN compute controller with front bezel
2	E2860 controller shelf with front bezel (optional expansion shelf appears identical)
3	SG6000-CN compute controller with front bezel removed

Callout	Description
4	E2860 controller shelf with front bezel removed (optional expansion shelf appears identical)

SG6060 rear view


This figure shows the back of the SG6060, including the compute and storage controllers, fans, and power supplies.



Callout	Description
1	Power supply (1 of 2) for SG6000-CN compute controller
2	Connectors for SG6000-CN compute controller
3	Fan (1 of 2) for E2860 controller shelf
4	E-Series E2800A storage controller (1 of 2) and connectors
5	Power supply (1 of 2) for E2860 controller shelf

SG6060X rear view


This figure shows the back of the SG6060X.



Callout	Description
1	Power supply (1 of 2) for SG6000-CN compute controller
2	Connectors for SG6000-CN compute controller
3	Fan (1 of 2) for E2860 controller shelf
4	E-Series E2800B storage controller (1 of 2) and connectors
5	Power supply (1 of 2) for E2860 controller shelf

Expansion shelf

This figure shows the back of the optional expansion shelf for the SG6060 and SG6060X, including the input/output modules (IOMs), fans, and power supplies. Each SG6060 can be installed with one or two expansion shelves, which can be included in the initial installation or added later.



Callout	Description
1	Fan (1 of 2) for expansion shelf
2	IOM (1 of 2) for expansion shelf
3	Power supply (1 of 2) for expansion shelf

SG6000 controllers

Each model of the StorageGRID SG6000 appliance includes an SG6000-CN compute controller in a 1U enclosure and duplex E-Series storage controllers in a 2U or 4U enclosure, depending on the model. Review the diagrams to learn more about each type of controller.

SG6000-CN compute controller


- Provides compute resources for the appliance.
- Includes the StorageGRID Appliance Installer.



StorageGRID software is not preinstalled on the appliance. This software is retrieved from the Admin Node when you deploy the appliance.

- Can connect to all three StorageGRID networks, including the Grid Network, the Admin Network, and the Client Network.
- Connects to the E-Series storage controllers and operates as the initiator.

SG6000-CN connectors




Callout	Port	Type	Use
1	Interconnect ports 1-4	16-Gb/s Fibre Channel (FC), with integrated optics	Connect the SG6000-CN controller to the E2800 controllers (two connections to each E2800).
2	Network ports 1-4	10-GbE or 25-GbE, based on cable or SFP transceiver type, switch speed, and configured link speed	Connect to the Grid Network and the Client Network for StorageGRID.
3	BMC management port	1-GbE (RJ-45)	Connect to the SG6000-CN baseboard management controller.
4	Diagnostic and support ports	<ul style="list-style-type: none"> VGA Serial, 115200 8-N-1 USB 	Reserved for technical support use.
5	Admin Network port 1	1-GbE (RJ-45)	Connect the SG6000-CN to the Admin Network for StorageGRID.
6	Admin Network port 2	1-GbE (RJ-45)	<p>Options:</p> <ul style="list-style-type: none"> Bond with management port 1 for a redundant connection to the Admin Network for StorageGRID. Leave unwired and available for temporary local access (IP 169.254.0.1). During installation, use port 2 for IP configuration if DHCP-assigned IP addresses aren't available.

SGF6024: EF570 storage controllers

- Two controllers for failover support.
- Manage the storage of data on the drives.
- Function as standard E-Series controllers in a duplex configuration.
- Include SANtricity OS Software (controller firmware).

- Include SANtricity System Manager for monitoring storage hardware and for managing alerts, the AutoSupport feature, and the Drive Security feature.
- Connect to the SG6000-CN controller and provide access to the flash storage.

EF570 connectors



Callout	Port	Type	Use
1	Interconnect ports 1 and 2	16-Gb/s FC optical SFP	Connect each of the EF570 controllers to the SG6000-CN controller. There are four connections to the SG6000-CN controller (two from each EF570).
2	Diagnostic and support ports	<ul style="list-style-type: none"> • RJ-45 serial port • Micro USB serial port • USB port 	Reserved for technical support use.
3	Drive expansion ports	12Gb/s SAS	Not used. The SGF6024 appliance does not support expansion drive shelves.
4	Management ports 1 and 2	1-Gb (RJ-45) Ethernet	<ul style="list-style-type: none"> • Port 1 connects to the network where you access SANtricity System Manager on a browser. • Port 2 is reserved for technical support use.

SG6060 and SG6060X: E2800 storage controllers

- Two controllers for failover support.
- Manage the storage of data on the drives.
- Function as standard E-Series controllers in a duplex configuration.
- Include SANtricity OS Software (controller firmware).
- Include SANtricity System Manager for monitoring storage hardware and for managing alerts, the AutoSupport feature, and the Drive Security feature.
- Connect to the SG6000-CN controller and provide access to the storage.

The SG6060 and SG6060X use E2800 storage controllers.


Appliance	Controller	Controller HIC
SG6060	Two E2800A storage controllers	None
SG6060X	Two E2800B storage controllers	Four-port HIC

The E2800A and the E2800B storage controllers are identical in specifications and function except for the location of the interconnect ports.



Don't use an E2800A and an E2800B in the same appliance.


E2800A connectors



Callout	Port	Type	Use
1	Interconnect ports 1 and 2	16-Gb/s FC optical SFP	<p>Connect each of the E2800A controllers to the SG6000-CN controller.</p> <p>There are four connections to the SG6000-CN controller (two from each E2800A).</p>

Callout	Port	Type	Use
2	Management ports 1 and 2	1-Gb (RJ-45) Ethernet	<ul style="list-style-type: none"> Port 1 Options: <ul style="list-style-type: none"> Connect to a management network to enable direct TCP/IP access to SANtricity System Manager Leave unwired to save a switch port and IP address. Access SANtricity System Manager using the Grid Manager or Storage Grid Appliance Installer UIs. <p>Note: some optional SANtricity functionality, such as NTP sync for accurate log timestamps, is not available when you choose to leave Port 1 unwired.</p> <p>Note: StorageGRID 11.5 or greater, and SANtricity 11.70 or greater, are required when you leave Port 1 unwired.</p> <ul style="list-style-type: none"> Port 2 is reserved for technical support use.
3	Diagnostic and support ports	<ul style="list-style-type: none"> RJ-45 serial port Micro USB serial port USB port 	Reserved for technical support use.
4	Drive expansion ports 1 and 2	12Gb/s SAS	Connect the ports to the drive expansion ports on the IOMs in the expansion shelf.

E2800B connectors




Callout	Port	Type	Use
1	Interconnect ports 1 and 2	16-Gb/s FC optical SFP	<p>Connect each of the E2800B controllers to the SG6000-CN controller.</p> <p>There are four connections to the SG6000-CN controller (two from each E2800B).</p>
2	Management ports 1 and 2	1-Gb (RJ-45) Ethernet	<ul style="list-style-type: none"> • Port 1 Options: <ul style="list-style-type: none"> ◦ Connect to a management network to enable direct TCP/IP access to SANtricity System Manager ◦ Leave unwired to save a switch port and IP address. Access SANtricity System Manager using the Grid Manager or Storage Grid Appliance Installer UIs. <p>Note: some optional SANtricity functionality, such as NTP sync for accurate log timestamps, is not available when you choose to leave Port 1 unwired.</p> <p>Note: StorageGRID 11.5 or greater, and SANtricity 11.70 or greater, are required when you leave Port 1 unwired.</p> <ul style="list-style-type: none"> • Port 2 is reserved for technical support use.
3	Diagnostic and support ports	<ul style="list-style-type: none"> • RJ-45 serial port • Micro USB serial port • USB port 	Reserved for technical support use.
4	Drive expansion ports 1 and 2	12Gb/s SAS	Connect the ports to the drive expansion ports on the IOMs in the expansion shelf.

SG6060 and SG6060X: IOMs for optional expansion shelves

The expansion shelf contains two input/output modules (IOMs) that connect to the storage controllers or to other expansion shelves.

IOM connectors



Callout	Port	Type	Use
1	Drive expansion ports 1-4	12Gb/s SAS	Connect each port to the storage controllers or additional expansion shelf (if any).

SG5700 appliance: Overview

The SG5700 StorageGRID appliance is an integrated storage and computing platform that operates as a Storage Node in a StorageGRID grid. The appliance can be used in a hybrid grid environment that combines appliance Storage Nodes and virtual (software-based) Storage Nodes.

StorageGRID SG5700 series appliance provides the following features:

- Integrate the storage and computing elements for a StorageGRID Storage Node.
- Include the StorageGRID Appliance Installer to simplify Storage Node deployment and configuration.
- Includes E-Series SANtricity System Manager for hardware management and monitoring.
- Support up to four 10-GbE or 25-GbE connections to the StorageGRID Grid Network and Client Network.
- Support Full Disk Encryption (FDE) drives or FIPS drives. When these drives are used with the Drive Security feature in SANtricity System Manager, unauthorized access to data is prevented.

The SG5700 appliance is available in four models: the SG5712 and SG5712X, and the SG5760 and SG5760X. There are no specification or functional differences between the SG5712 and SG5712X except for the location of the interconnect ports on the storage controller. Similarly, there are no specification or functional differences between the SG5760 and SG5760X except for the location of the interconnect ports on the storage controller.

SG5700 components

The SG5700 models include the following components:

Component	SG5712	SG5712X	SG5760	SG5760X
Compute controller	E5700SG controller	E5700SG controller	E5700SG controller	E5700SG controller
Storage controller	E2800A controller	E2800B controller	E2800A controller	E2800B controller

Component	SG5712	SG5712X	SG5760	SG5760X
Chassis	E-Series DE212C enclosure, a two rack-unit (2U) enclosure	E-Series DE212C enclosure, a two rack-unit (2U) enclosure	E-Series DE460C enclosure, a four rack-unit (4U) enclosure	E-Series DE460C enclosure, a four rack-unit (4U) enclosure
Drives	12 NL-SAS drives (3.5-inch)	12 NL-SAS drives (3.5-inch)	60 NL-SAS drives (3.5-inch)	60 NL-SAS drives (3.5-inch)
Redundant power supplies and fans	Two power-fan canisters	Two power-fan canisters	Two power canisters and two fan canisters	Two power canisters and two fan canisters

The maximum raw storage available in the StorageGRID appliance is fixed, based on the number of drives in each enclosure. You can't expand the available storage by adding a shelf with additional drives.

SG5700 diagrams


SG5712 front and rear views

The figures show the front and back of the SG5712, a 2U enclosure that holds 12 drives.



SG5712 components

The SG5712 includes two controllers and two power-fan canisters.



Callout	Description
1	E2800A controller (storage controller)
2	E5700SG controller (compute controller)
3	Power-fan canisters


SG5712X front and rear views

The figures show the front and back of the SG5712X, a 2U enclosure that holds 12 drives.



SG5712X components


The SG5712X includes two controllers and two power-fan canisters.



Callout	Description
1	E2800B controller (storage controller)
2	E5700SG controller (compute controller)
3	Power-fan canisters

SG5760 front and rear views

The figures show the front and back of the SG5760 model, a 4U enclosure that holds 60 drives in 5 drive drawers.





SG5760 components

The SG5760 includes two controllers, two fan canisters, and two power canisters.

Callout	Description
1	E2800A controller (storage controller)
2	E5700SG controller (compute controller)
3	Fan canister (1 of 2)
4	Power canister (1 of 2)

SG5760X front and rear views

The figures show the front and back of the SG5760X model, a 4U enclosure that holds 60 drives in 5 drive drawers.



SG5760X components

The SG5760X includes two controllers, two fan canisters, and two power canisters.

Callout	Description
1	E2800B controller (storage controller)
2	E5700SG controller (compute controller)
3	Fan canister (1 of 2)
4	Power canister (1 of 2)

Related information

[NetApp E-Series Systems Documentation Site](#)

SG5700 controllers

Both the 12-drive SG5712 and SG5712X and the 60-drive SG5760 and SG5760X models of the StorageGRID appliance include an E5700SG compute controller and an E-Series E2800 storage controller.

- The SG5712 and SG5760 use a E2800A controller.
- The SG5712X and the SG5760X use a E2800B controller.

The E2800A and E2800B controllers are identical in specification and function except for the location of the interconnect ports.

E5700SG compute controller


- Operates as the compute server for the appliance.
- Includes the StorageGRID Appliance Installer.



StorageGRID software is not preinstalled on the appliance. This software is accessed from the Admin Node when you deploy the appliance.

- Can connect to all three StorageGRID networks, including the Grid Network, the Admin Network, and the Client Network.
- Connects to the E2800 controller and operates as the initiator.

E5700SG connectors



Callout	Port	Type	Use
1	Interconnect ports 1 and 2	16Gb/s Fibre Channel (FC), optical SFP	Connect the E5700SG controller to the E2800 controller.
2	Diagnostic and support ports	<ul style="list-style-type: none">• RJ-45 serial port• Micro USB serial port• USB port	Reserved for technical support.
3	Drive expansion ports	12Gb/s SAS	Not used. StorageGRID appliances don't support expansion drive shelves.
4	Network ports 1-4	10-GbE or 25-GbE, based on SFP transceiver type, switch speed, and configured link speed	Connect to the Grid Network and the Client Network for StorageGRID.
5	Management port 1	1-Gb (RJ-45) Ethernet	Connect to the Admin Network for StorageGRID.

Callout	Port	Type	Use
6	Management port 2	1-Gb (RJ-45) Ethernet	<p>Options:</p> <ul style="list-style-type: none"> • Bond with management port 1 for a redundant connection to the Admin Network for StorageGRID. • Leave unwired and available for temporary local access (IP 169.254.0.1). • During installation, use port 2 for IP configuration if DHCP-assigned IP addresses aren't available.


E2800 storage controller

There are two versions of the E2800 storage controller used in the SG5700 appliances: E2800A and E2800B. The E2800A does not have a HIC, and the E2800B has a four-port HIC. The two controller versions have identical specifications and function except for the location of the interconnect ports.


The E2800 series storage controller has the following specifications:

- Operates as the storage controller for the appliance.
- Manages the storage of data on the drives.
- Functions as a standard E-Series controller in simplex mode.
- Includes SANtricity OS Software (controller firmware).
- Includes SANtricity System Manager for monitoring appliance hardware and for managing alerts, the AutoSupport feature, and the Drive Security feature.
- Connects to the E5700SG controller and operates as the target.

E2800A connectors



E2800B connectors



Callout	Port	Type	Use
1	Interconnect ports 1 and 2	16Gb/s FC optical SFP	Connect the E2800 controller to the E5700SG controller.
2	Management ports 1 and 2	1-Gb (RJ-45) Ethernet	<ul style="list-style-type: none"> Port 1 Options: <ul style="list-style-type: none"> Connect to a management network to enable direct TCP/IP access to SANtricity System Manager Leave unwired to save a switch port and IP address. Access SANtricity System Manager using the Grid Manager or Storage Grid Appliance Installer UIs. Note: some optional SANtricity functionality, such as NTP sync for accurate log timestamps, is not available when you choose to leave Port 1 unwired. Note: StorageGRID 11.5 or greater, and SANtricity 11.70 or greater, are required when you leave Port 1 unwired. Port 2 is reserved for technical support use.

Callout	Port	Type	Use
3	Diagnostic and support ports	<ul style="list-style-type: none"> • RJ-45 serial port • Micro USB serial port • USB port 	Reserved for technical support use.
4	Drive expansion ports.	12Gb/s SAS	Not used.

SG100 and SG1000 appliances: Overview

The StorageGRID SG100 services appliance and the SG1000 services appliance can operate as a Gateway Node and as an Admin Node to provide high availability load balancing services in a StorageGRID system. Both appliances can operate as Gateway Nodes and Admin Nodes (primary or non-primary) at the same time.

Appliance features

Both models of the services appliance provide the following features:

- Gateway Node or Admin Node functions for a StorageGRID system.
- The StorageGRID Appliance Installer to simplify node deployment and configuration.
- When deployed, can access StorageGRID software from an existing Admin Node or from software downloaded to a local drive. To further simplify the deployment process, a recent version of the software is preloaded onto the appliance during manufacturing.
- A baseboard management controller (BMC) for monitoring and diagnosing some of the appliance hardware.
- The ability to connect to all three StorageGRID networks, including the Grid Network, the Admin Network, and the Client Network:
 - The SG100 supports up to four 10- or 25-GbE connections to the Grid Network and Client Network.
 - The SG1000 supports up to four 10-, 25-, 40-, or 100-GbE connections to the Grid Network and Client Network.

SG100 and SG1000 diagrams

This figure shows the front of the SG100 and the SG1000 with the bezel removed. From the front, the two appliances are identical except for the product name on the bezel.




The two solid-state drives (SSDs), indicated by the orange outline, are used for storing the StorageGRID operating system and are mirrored using RAID 1 for redundancy. When the SG100 or SG1000 services appliance is configured as an Admin Node, these drives are used to store audit logs, metrics, and database tables.

The remaining drive slots are blank.

SG100 connectors


This figure shows the connectors on the back of the SG100.



Callout	Port	Type	Use
1	Network ports 1-4	10/25-GbE, based on cable or SFP transceiver type (SFP28 and SFP+ modules are supported), switch speed, and configured link speed	Connect to the Grid Network and the Client Network for StorageGRID.
2	BMC management port	1-GbE (RJ-45)	Connect to the appliance baseboard management controller.
3	Diagnostic and support ports	<ul style="list-style-type: none"> VGA Serial, 115200 8-N-1 USB 	Reserved for technical support use.
4	Admin Network port 1	1-GbE (RJ-45)	Connect the appliance to the Admin Network for StorageGRID.
5	Admin Network port 2	1-GbE (RJ-45)	<p>Options:</p> <ul style="list-style-type: none"> Bond with management port 1 for a redundant connection to the Admin Network for StorageGRID. Leave disconnected and available for temporary local access (IP 169.254.0.1). During installation, use port 2 for IP configuration if DHCP-assigned IP addresses aren't available.

SG1000 connectors

This figure shows the connectors on the back of the SG1000.



Callout	Port	Type	Use
1	Network ports 1-4	10/25/40/100-GbE, based on cable or transceiver type, switch speed, and configured link speed. QSFP28 and QSFP+ (40/100GbE) are supported natively and SFP28/SFP+ transceivers can be used with a QSA (sold separately) to use 10/25GbE speeds.	Connect to the Grid Network and the Client Network for StorageGRID.
2	BMC management port	1-GbE (RJ-45)	Connect to the appliance baseboard management controller.
3	Diagnostic and support ports	<ul style="list-style-type: none"> VGA Serial, 115200 8-N-1 USB 	Reserved for technical support use.
4	Admin Network port 1	1-GbE (RJ-45)	Connect the appliance to the Admin Network for StorageGRID.

Callout	Port	Type	Use
5	Admin Network port 2	1-GbE (RJ-45)	<p>Options:</p> <ul style="list-style-type: none"> • Bond with management port 1 for a redundant connection to the Admin Network for StorageGRID. • Leave disconnected and available for temporary local access (IP 169.254.0.1). • During installation, use port 2 for IP configuration if DHCP-assigned IP addresses aren't available.

SG100 and SG1000 applications

You can configure the StorageGRID services appliances in various ways to provide gateway services as well as redundancy of some grid administration services.

Appliances can be deployed in the following ways:

- Add to a new or existing grid as a Gateway Node
- Add to a new grid as a primary or non-primary Admin Node, or to an existing grid as a non-primary Admin Node
- Operate as a Gateway Node and Admin Node (primary or non-primary) at the same time

The appliance facilitates the use of high availability (HA) groups and intelligent load balancing for S3 or Swift data path connections.

The following examples describe how you can maximize the capabilities of the appliance:

- Use two SG100 or two SG1000 appliances to provide gateway services by configuring them as Gateway Nodes.



Don't deploy the SG100 and SG1000 service appliances in the same site. Unpredictable performance might result.

- Use two SG100 or two SG1000 appliances to provide redundancy of some grid administration services. Do this by configuring each appliance as Admin Nodes.
- Use two SG100 or two SG1000 appliances to provide highly available load balancing and traffic shaping services accessed through one or more virtual IP addresses. Do this by configuring the appliances as any combination of Admin Nodes or Gateway Nodes and adding both nodes to the same HA group.



If you use Admin Nodes and Gateway Nodes in the same HA group, Admin Node-only port will not fail over. See the instructions for [configuring HA groups](#).

When used with StorageGRID storage appliances, both the SG100 and the SG1000 services appliances enable deployment of appliance-only grids with no dependencies on external hypervisors or compute hardware.

Prepare for installation

Prepare site

Before installing the appliance, you must make sure that the site and the cabinet or rack you plan to use meet the specifications for a StorageGRID appliance.

Steps

1. Confirm that the site meets the requirements for temperature, humidity, altitude range, airflow, heat dissipation, wiring, power, and grounding. See the [NetApp Hardware Universe](#) for more information.
2. Confirm that your location provides the correct voltage of AC power:

Model	Requirement
SGF6112	100 to 240 volts AC
SG6060	240-volt AC
SGF6024	120-volt AC
SG5760	240-volt AC
SG100 and SG1000	120 to 240 volts AC

3. Obtain a 19-inch (48.3-cm) cabinet or rack to fit shelves of the following size (without cables).

SGF6112

Height	Width	Depth	Maximum weight
1.70 in. (4.31 cm)	18.98 in. (48.2 cm)	33.11 in. (84.1 cm)	43.83 lb. (19.88 kg)

SG6000

Type of shelf	Height	Width	Depth	Maximum weight
E2860 controller shelf (SG6060)	6.87 in. (17.46 cm)	17.66 in. (44.86 cm)	38.25 in. (97.16 cm)	250 lb. (113 kg)
Expansion shelf (SG6060) - Optional	6.87 in. (17.46 cm)	17.66 in. (44.86 cm)	38.25 in. (97.16 cm)	250 lb. (113 kg)
EF570 controller shelf (SGF6024)	3.35 in. (8.50 cm)	17.66 in. (44.86 cm)	19.00 in. (48.26 cm)	51.74 lb. (23.47 kg)
SG6000-CN compute controller	1.70 in. (4.32 cm)	17.32 in. (44.0 cm)	32.0 in. (81.3 cm)	39 lb. (17.7 kg)

SG5700

Appliance model	Height	Width	Depth	Maximum weight
SG5712 (12 drives)	3.41 in. (8.68 cm)	17.6 in. (44.7 cm)	21.1 in. (53.6 cm)	63.9 lb (29.0 kg)
SG5760 (60 drives)	6.87 in. (17.46 cm)	17.66 in. (44.86 cm)	38.25 in. (97.16 cm)	250 lb. (113 kg)

SG100 and SG1000

Height	Width	Depth	Maximum weight
1.70 in. (4.32 cm)	17.32 in. (44.0 cm)	32.0 in. (81.3 cm)	39 lb. (17.7 kg)

4. Decide where you are going to install the appliance.



When installing the E2860 controller shelf or optional expansion shelves, install hardware from the bottom to the top of the rack or cabinet to prevent the equipment from tipping over. To ensure that the heaviest equipment is at the bottom of the cabinet or rack, install the SG6000-CN controller above the E2860 controller shelf and expansion shelves.



Before committing to the installation, verify that the 0.5m optic cables shipped with an SG6000 appliance, or cables that you supply, are long enough for the planned layout.

5. Install any required network switches. See the [NetApp Interoperability Matrix Tool](#) for compatibility information.

Unpack boxes

Before installing your StorageGRID appliance, unpack all boxes and compare the contents to the items on the packing slip.

SGF6112 appliances

Hardware

Item	What it looks like
SGF6112	
Rail kit with instructions	
Front bezel	

Power cords

The shipment for an SGF6112 appliance includes the following power cords.




Your cabinet might have special power cords that you use instead of the power cords that ship with the appliance.

Item	What it looks like
Two power cords for your country	



SG6000 appliances


SG6060 hardware

Item	What it looks like
SG6000-CN controller	
E2860 controller shelf with no drives installed	
Two front bezels	
Two rail kits with instructions	
60 drives (2 SSD and 58 NL-SAS)	
Four handles	

Item	What it looks like
Back brackets and cage nuts for square-hole rack installation	

SG6060 expansion shelf

Item	What it looks like
Expansion shelf with no drives installed	
Front bezel	
60 NL-SAS drives	
One rail kit with instructions	
Four handles	

Item	What it looks like
Back brackets and cage nuts for square-hole rack installation	

SGF6024 hardware


Item	What it looks like
SG6000-CN controller	
EF570 flash array with 24 solid state (flash) drives installed	
Two front bezels	
Two rail kits with instructions	
Shelf endcaps	

Cables and connectors

The shipment for an SG6000 appliance includes the following cables and connectors.







Your cabinet might have special power cords that you use instead of the power cords that ship with the appliance.

Item	What it looks like
Four power cords for your country	
Optical cables and SFP transceivers	 <ul style="list-style-type: none"> Four optical cables for the FC interconnect ports Four SFP+ transceivers, which support 16-Gb/s FC
Optional: Two SAS cables for connecting each SG6060 expansion shelf	

SG5700 appliances

Hardware

Item	What it looks like
SG5712 appliance with 12 drives installed	
SG5760 appliance with no drives installed	
Front bezel for the appliance	

Item	What it looks like
Rail kit with instructions	
SG5760: Sixty drives	
SG5760: Handles	
SG5760: Back brackets and cage nuts for square-hole rack installation	


Cables and connectors

The shipment for an SG5700 appliance includes the following cables and connectors.



Your cabinet might have special power cords that you use instead of the power cords that ship with the appliance.

Item	What it looks like
Two power cords for your country	

Item	What it looks like
Optical cables and SFP transceivers	 <ul style="list-style-type: none"> • Two optical cables for the FC interconnect ports • Eight SFP+ transceivers, compatible with both the four 16Gb/s FC interconnect ports and the four 10-GbE network ports

SG100 and SG1000 appliances

Hardware

Item	What it looks like
SG100 or SG1000	
Rail kit with instructions	

Power cords

The shipment for an SG100 or SG1000 appliance includes the following power cords.



Your cabinet might have special power cords that you use instead of the power cords that ship with the appliance.


Item	What it looks like
Two power cords for your country	

Obtain additional equipment and tools

Before installing a StorageGRID appliance, confirm you have all of the additional equipment and tools that you need.


All appliances

You need the following equipment to install and configure all appliances.

Item	What it looks like
Screwdrivers	 <ul style="list-style-type: none"> • Phillips No. 2 screwdriver • Medium flat-blade screwdriver
ESD wrist strap	
Service laptop	 <ul style="list-style-type: none"> • Supported web browser • SSH client, such as PuTTY • 1-GbE (RJ-45) port <p> Some ports might not support 10/100 Ethernet speeds.</p>
Optional tools	 <ul style="list-style-type: none"> • Power drill with Phillips head bit • Flashlight

SGF6112

You need the following additional equipment to install and configure the SGF6112 hardware.

Item	What it looks like
Optical cables and transceivers	 <ul style="list-style-type: none">• One to four of either of these types of cables:<ul style="list-style-type: none">◦ TwinAx/Copper◦ Fibre/Optical• One to four of each of these transceivers/adapters based on link speed (mixed speeds aren't supported):<ul style="list-style-type: none">◦ 10-GbE SFP+◦ 25-GbE SFP28
RJ-45 (Cat5/Cat5e/Cat6/Cat6a) Ethernet cables	

SG6000

You need the following additional equipment to install and configure the SG6000 hardware.

Item	What it looks like
Optical cables and SFP transceivers	 <ul style="list-style-type: none">• One to four of either of these types of cables:<ul style="list-style-type: none">◦ TwinAx/Copper◦ Fibre/Optical• One to four of each of these transceivers/adapters, based on link speed (mixed speeds aren't supported):<ul style="list-style-type: none">◦ 10-GbE SFP+◦ 25-GbE SFP28
RJ-45 (Cat5/Cat5e/Cat6) Ethernet cables	

Item	What it looks like
Optional tools	Mechanized lift for 60-drive shelves

SG5700

You need the following additional equipment to install and configure the SG5700 hardware.

Item	What it looks like
Optical cables and SFP transceivers	  <ul style="list-style-type: none"> Optical cables for the 10/25-GbE ports you plan to use Optional: SFP28 transceivers if you want to use 25-GbE link speed
Ethernet cables	
Optional tools	Mechanized lift for SG5760

SG100 and SG1000

You need the following additional equipment to install and configure the SG100 and SG1000 hardware.

Item	What it looks like
Optical cables and transceivers	 <ul style="list-style-type: none"> One to four of either of these cable types: <ul style="list-style-type: none"> TwinAx/Copper Fibre/Optical One to four of each of these transceivers/adapters based on link speed (mixed speeds aren't supported): <ul style="list-style-type: none"> SG100: <ul style="list-style-type: none"> 10-GbE SFP+ 25-GbE SFP28 SG1000: <ul style="list-style-type: none"> 10-GbE QSFP-to-SFP adapter (QSA) and SFP+ 25-GbE QSFP-to-SFP adapter (QSA) and SFP28 40-GbE QSFP+ 100-GbE QFSP28
RJ-45 (Cat5/Cat5e/Cat6/Cat6a) Ethernet cables	

Web browser requirements

You must use a supported web browser.

Web browser	Minimum supported version
Google Chrome	107
Microsoft Edge	107
Mozilla Firefox	106

You should set the browser window to a recommended width.

Browser width	Pixels
Minimum	1024

Browser width	Pixels
Optimum	1280

Review appliance network connections

Review appliance network connections

Before installing the StorageGRID appliance, you should understand which networks can be connected to the appliance and how the ports on each controller are used.

StorageGRID network requirements are fully explained in the [Networking guidelines](#).

When you deploy a StorageGRID appliance as a node in a StorageGRID system, you can connect it to the following networks:

- **Grid Network for StorageGRID:** The Grid Network is used for all internal StorageGRID traffic. It provides connectivity between all nodes in the grid, across all sites and subnets. The Grid Network is required.
- **Admin Network for StorageGRID:** The Admin Network is a closed network used for system administration and maintenance. The Admin Network is typically a private network and does not need to be routable between sites. The Admin Network is optional.
- **Client Network for StorageGRID:** The Client Network is an open network used to provide access to client applications, including S3 and Swift. The Client Network provides client protocol access to the grid, so the Grid Network can be isolated and secured. You can configure the Client Network so that the appliance can be accessed over this network using only the ports you choose to open. The Client Network is optional.
- **Management network for SANtricity** (optional for storage appliances, not needed for the SGF6112): This network provides access to SANtricity System Manager, allowing you to monitor and manage the hardware components in the appliance and storage controller shelf. This management network can be the same as the Admin Network for StorageGRID, or it can be an independent management network.
- **BMC management network** (optional for SG100, SG1000, SG6000, and SGF6112): This network provides access to the baseboard management controller in the SG100, SG1000, SG6000, and SGF6112 appliances allowing you to monitor and manage the hardware components in the appliance. This management network can be the same as the Admin Network for StorageGRID, or it can be an independent management network.

If the optional BMC management network is not connected, some support and maintenance procedures will be more difficult to perform. You can leave the BMC management network unconnected except when needed for support purposes.




For detailed information about StorageGRID networks, see the [StorageGRID network types](#).

Network connections (SG5700)

When you install a SG5700 StorageGRID appliance, you connect the two controllers to each other and to the required networks.

The figure shows the two controllers in the SG5760, with the E2800 controller on the top and the E5700SG controller on the bottom. In the SG5712, the E2800 controller is to the left of the E5700SG controller.



Callout	Port	Type of port	Use
1	Two interconnect ports on each controller	16Gb/s FC optical SFP+	Connect the two controllers to each other.
2	Management port 1 on the E2800 controller	1-GbE (RJ-45)	Connects to the network where you access SANtricity System Manager. You can use the Admin Network for StorageGRID or an independent management network.
	Management port 2 on the E2800 controller	1-GbE (RJ-45)	Reserved for technical support.
3	Management port 1 on the E5700SG controller	1-GbE (RJ-45)	Connects the E5700SG controller to the Admin Network for StorageGRID.
	Management port 2 on the E5700SG controller	1-GbE (RJ-45)	<ul style="list-style-type: none"> Can be bonded with management port 1 if you want a redundant connection to the Admin Network. Can be left unwired and available for temporary local access (IP 169.254.0.1). During installation, can be used to connect the E5700SG controller to a service laptop if DHCP-assigned IP addresses aren't available.

Callout	Port	Type of port	Use
4	10/25-GbE ports 1-4 on the E5700SG controller	10-GbE or 25-GbE Note: The SFP+ transceivers included with the appliance support 10-GbE link speeds. If you want to use 25-GbE link speeds for the four network ports, you must provide SFP28 transceivers.	Connect to the Grid Network and the Client Network for StorageGRID. See Port bond modes (E5700SG controller) .

Port bond modes (SGF6112)


When [configuring network links](#) for the SGF6112 appliance, you can use port bonding for the ports that connect to the Grid Network and optional Client Network, and the 1/10-GbE management ports that connect to the optional Admin Network. Port bonding helps protect your data by providing redundant paths between StorageGRID networks and the appliance.

Network bond modes

The networking ports on the appliance support Fixed port bond mode or Aggregate port bond mode for the Grid Network and Client Network connections.

Fixed port bond mode

Fixed port bond mode is the default configuration for the networking ports.



Callout	Which ports are bonded
C	Ports 1 and 3 are bonded together for the Client Network, if this network is used.
G	Ports 2 and 4 are bonded together for the Grid Network.

When using Fixed port bond mode, the ports can be bonded using active-backup mode or Link Aggregation Control Protocol mode (LACP 802.3ad).


- In active-backup mode (default), only one port is active at a time. If the active port fails, its backup port automatically provides a failover connection. Port 4 provides a backup path for port 2 (Grid Network), and port 3 provides a backup path for port 1 (Client Network).
- In LACP mode, each pair of ports forms a logical channel between the appliance and the network, allowing

for higher throughput. If one port fails, the other port continues to provide the channel. Throughput is reduced, but connectivity is not impacted.

 If you don't need redundant connections, you can use only one port for each network. However, be aware that the **Storage appliance link down** alert might be triggered in the Grid Manager after StorageGRID is installed, indicating that a cable is unplugged. You can safely disable this alert rule.

Aggregate port bond mode

Aggregate port bond mode significantly increases the throughput for each StorageGRID network and provides additional failover paths.



Callout	Which ports are bonded
1	All connected ports are grouped in a single LACP bond, allowing all ports to be used for Grid Network and Client Network traffic.

If you plan to use aggregate port bond mode:


- You must use LACP network bond mode.
- You must specify a unique VLAN tag for each network. This VLAN tag will be added to each network packet to ensure that network traffic is routed to the correct network.
- The ports must be connected to switches that can support VLAN and LACP. If multiple switches are participating in the LACP bond, the switches must support multi-chassis link aggregation groups (MLAG), or equivalent.
- You understand how to configure the switches to use VLAN, LACP, and MLAG, or equivalent.

If you don't want to use all four ports, you can use one, two, or three ports. Using more than one port maximizes the chance that some network connectivity will remain available if one of the ports fails.

 If you choose to use fewer than four network ports, be aware that a **Services appliance link down** alert might be triggered in the Grid Manager after the appliance node is installed, indicating that a cable is unplugged. You can safely disable this alert rule for the triggered alert.

Network bond modes for management ports

For the two 1/10-GbE management ports, you can choose Independent network bond mode or Active-Backup network bond mode to connect to the optional Admin Network.



In Independent mode, only the management port on the left is connected to the Admin Network. This mode does not provide a redundant path. The management port on the right is unconnected and available for temporary local connections (uses IP address 169.254.0.1).

In Active-Backup mode, both management ports are connected to the Admin Network. Only one port is active at a time. If the active port fails, its backup port automatically provides a failover connection. Bonding these two physical ports into one logical management port provides a redundant path to the Admin Network.



If you need to make a temporary local connection to the appliance when the 1/10-GbE management ports are configured for Active-Backup mode, remove the cables from both management ports, plug your temporary cable into the management port on the right, and access the appliance using IP address 169.254.0.1.

Callout	Network bond mode
A	Active-Backup mode. Both management ports are bonded into one logical management port connected to the Admin Network.
I	Independent mode. The port on the left is connected to the Admin Network. The port on the right is available for temporary local connections (IP address 169.254.0.1).

Port bond modes (SG6000-CN controller)


When [configuring network links](#) for the SG6000-CN controller, you can use port bonding for the 10/25-GbE ports that connect to the Grid Network and optional Client Network, and the 1-GbE management ports that connect to the optional Admin Network. Port bonding helps protect your data by providing redundant paths between StorageGRID networks and the appliance.

Network bond modes for 10/25-GbE ports

The 10/25-GbE networking ports on the SG6000-CN controller support Fixed port bond mode or Aggregate port bond mode for the Grid Network and Client Network connections.

Fixed port bond mode

Fixed mode is the default configuration for the 10/25-GbE networking ports.



Callout	Which ports are bonded
C	Ports 1 and 3 are bonded together for the Client Network, if this network is used.
G	Ports 2 and 4 are bonded together for the Grid Network.

When using Fixed port bond mode, the ports can be bonded using active-backup mode or Link Aggregation Control Protocol mode (LACP 802.3ad).

- In active-backup mode (default), only one port is active at a time. If the active port fails, its backup port automatically provides a failover connection. Port 4 provides a backup path for port 2 (Grid Network), and port 3 provides a backup path for port 1 (Client Network).
- In LACP mode, each pair of ports forms a logical channel between the controller and the network, allowing for higher throughput. If one port fails, the other port continues to provide the channel. Throughput is reduced, but connectivity is not impacted.




If you don't need redundant connections, you can use only one port for each network. However, be aware that an alert will be triggered in the Grid Manager after StorageGRID is installed, indicating that the link is down. Because this port is disconnected on purpose, you can safely disable this alert.

From the Grid Manager, select **Alert > Rules**, select the rule, and click **Edit rule**. Then, uncheck the **Enabled** checkbox.

Aggregate port bond mode

Aggregate port bond mode significantly increases the throughput for each StorageGRID network and provides additional failover paths.



Callout	Which ports are bonded
1	All connected ports are grouped in a single LACP bond, allowing all ports to be used for Grid Network and Client Network traffic.

If you plan to use aggregate port bond mode:

- You must use LACP network bond mode.
- You must specify a unique VLAN tag for each network. This VLAN tag will be added to each network packet to ensure that network traffic is routed to the correct network.
- The ports must be connected to switches that can support VLAN and LACP. If multiple switches are participating in the LACP bond, the switches must support multi-chassis link aggregation groups (MLAG), or equivalent.
- You understand how to configure the switches to use VLAN, LACP, and MLAG, or equivalent.

If you don't want to use all four 10/25-GbE ports, you can use one, two, or three ports. Using more than one port maximizes the chance that some network connectivity will remain available if one of the 10/25-GbE ports fails.



If you choose to use fewer than four ports, be aware that one or more alarms will be raised in the Grid Manager after StorageGRID is installed, indicating that cables are unplugged. You can safely acknowledge the alarms to clear them.

Network bond modes for 1-GbE management ports

For the two 1-GbE management ports on the SG6000-CN controller, you can choose Independent network bond mode or Active-Backup network bond mode to connect to the optional Admin Network.

In Independent mode, only the management port on the left is connected to the Admin Network. This mode does not provide a redundant path. The management port on the right is unconnected and available for temporary local connections (uses IP address 169.254.0.1)

In Active-Backup mode, both management ports are connected to the Admin Network. Only one port is active at a time. If the active port fails, its backup port automatically provides a failover connection. Bonding these two physical ports into one logical management port provides a redundant path to the Admin Network.



If you need to make a temporary local connection to the SG6000-CN controller when the 1-GbE management ports are configured for Active-Backup mode, remove the cables from both management ports, plug your temporary cable into the management port on the right, and access the appliance using IP address 169.254.0.1.



Callout	Network bond mode
A	Both management ports are bonded into one logical management port connected to the Admin Network.
I	The port on the left is connected to the Admin Network. The port on the right is available for temporary local connections (IP address 169.254.0.1).

Port bond modes (E5700SG controller)


When [configuring network links](#) for the E5700SG controller, you can use port bonding for the 10/25-GbE ports that connect to the Grid Network and optional Client Network, and the 1-GbE management ports that connect to the optional Admin Network. Port bonding helps protect your data by providing redundant paths between StorageGRID networks and the appliance.

Network bond modes for 10/25-GbE ports

The 10/25-GbE networking ports on the E5700SG controller support Fixed port bond mode or Aggregate port bond mode for the Grid Network and Client Network connections.

Fixed port bond mode

Fixed mode is the default configuration for the 10/25-GbE networking ports.



Callout	Which ports are bonded
C	Ports 1 and 3 are bonded together for the Client Network, if this network is used.
G	Ports 2 and 4 are bonded together for the Grid Network.

When using Fixed port bond mode, you can use one of two network bond modes: Active-Backup or Link Aggregation Control Protocol (LACP).


- In Active-Backup mode (default), only one port is active at a time. If the active port fails, its backup port automatically provides a failover connection. Port 4 provides a backup path for port 2 (Grid Network), and port 3 provides a backup path for port 1 (Client Network).
- In LACP mode, each pair of ports forms a logical channel between the controller and the network, allowing for higher throughput. If one port fails, the other port continues to provide the channel. Throughput is reduced, but connectivity is not impacted.



If you don't need redundant connections, you can use only one port for each network. However, be aware that an alarm will be raised in the Grid Manager after StorageGRID is installed, indicating that a cable is unplugged. You can safely acknowledge this alarm to clear it.

Aggregate port bond mode

Aggregate port bond mode significantly increases the throughput for each StorageGRID network and provides additional failover paths.



Callout	Which ports are bonded
1	All connected ports are grouped in a single LACP bond, allowing all ports to be used for Grid Network and Client Network traffic.

If you plan to use Aggregate port bond mode:

- You must use LACP network bond mode.
- You must specify a unique VLAN tag for each network. This VLAN tag will be added to each network packet to ensure that network traffic is routed to the correct network.
- The ports must be connected to switches that can support VLAN and LACP. If multiple switches are participating in the LACP bond, the switches must support multi-chassis link aggregation groups (MLAG), or equivalent.
- You understand how to configure the switches to use VLAN, LACP, and MLAG, or equivalent.

If you don't want to use all four 10/25-GbE ports, you can use one, two, or three ports. Using more than one port maximizes the chance that some network connectivity will remain available if one of the 10/25-GbE ports fails.



If you choose to use fewer than four ports, be aware that one or more alarms will be raised in the Grid Manager after StorageGRID is installed, indicating that cables are unplugged. You can safely acknowledge the alarms to clear them.

Network bond modes for 1-GbE management ports


For the two 1-GbE management ports on the E5700SG controller, you can choose Independent network bond mode or Active-Backup network bond mode to connect to the optional Admin Network.

In Independent mode, only management port 1 is connected to the Admin Network. This mode does not provide a redundant path. Management port 2 is left unwired and available for temporary local connections (use IP address 169.254.0.1)

In Active-Backup mode, both management ports 1 and 2 are connected to the Admin Network. Only one port is active at a time. If the active port fails, its backup port automatically provides a failover connection. Bonding these two physical ports into one logical management port provides a redundant path to the Admin Network.



If you need to make a temporary local connection to the E5700SG controller when the 1-GbE management ports are configured for Active-Backup mode, remove the cables from both management ports, plug your temporary cable into management port 2, and access the appliance using IP address 169.254.0.1.



Port bond modes (SG100 and SG1000)

When configuring network links for the SG100 and SG1000 appliances, you can use port bonding for the ports that connect to the Grid Network and optional Client Network, and the 1-GbE management ports that connect to the optional Admin Network. Port bonding helps protect your data by providing redundant paths between StorageGRID networks and the appliance.


Network bond modes

The networking ports on the services appliance support Fixed port bond mode or Aggregate port bond mode for the Grid Network and Client Network connections.


Fixed port bond mode

Fixed port bond mode is the default configuration for the networking ports. The figures show how the network ports on the SG1000 or SG100 are bonded in fixed port bond mode.

SG100:



SG1000:



Callout	Which ports are bonded
C	Ports 1 and 3 are bonded together for the Client Network, if this network is used.
G	Ports 2 and 4 are bonded together for the Grid Network.

When using Fixed port bond mode, the ports can be bonded using active-backup mode or Link Aggregation

Control Protocol mode (LACP 802.3ad).

- In active-backup mode (default), only one port is active at a time. If the active port fails, its backup port automatically provides a failover connection. Port 4 provides a backup path for port 2 (Grid Network), and port 3 provides a backup path for port 1 (Client Network).
- In LACP mode, each pair of ports forms a logical channel between the services appliance and the network, allowing for higher throughput. If one port fails, the other port continues to provide the channel. Throughput is reduced, but connectivity is not impacted.




If you don't need redundant connections, you can use only one port for each network. However, be aware that the **Services appliance link down** alert might be triggered in the Grid Manager after StorageGRID is installed, indicating that a cable is unplugged. You can safely disable this alert rule.


Aggregate port bond mode

Aggregate port bond mode significantly increases the throughput for each StorageGRID network and provides additional failover paths. These figures show how the network ports are bonded in aggregate port bond mode.

SG100:



SG1000:



Callout	Which ports are bonded
1	All connected ports are grouped in a single LACP bond, allowing all ports to be used for Grid Network and Client Network traffic.

If you plan to use aggregate port bond mode:

- You must use LACP network bond mode.
- You must specify a unique VLAN tag for each network. This VLAN tag will be added to each network packet to ensure that network traffic is routed to the correct network.
- The ports must be connected to switches that can support VLAN and LACP. If multiple switches are participating in the LACP bond, the switches must support multi-chassis link aggregation groups (MLAG),

or equivalent.

- You understand how to configure the switches to use VLAN, LACP, and MLAG, or equivalent.

If you don't want to use all four ports, you can use one, two, or three ports. Using more than one port maximizes the chance that some network connectivity will remain available if one of the ports fails.



If you choose to use fewer than four network ports, be aware that a **Services appliance link down** alert might be triggered in the Grid Manager after the appliance node is installed, indicating that a cable is unplugged. You can safely disable this alert rule for the triggered alert.

Network bond modes for management ports

For the two 1-GbE management ports on the services appliance, you can choose Independent network bond mode or Active-Backup network bond mode to connect to the optional Admin Network. These figures show how the management ports on the appliances are bonded in network bond mode for the Admin Network.

SG100:



SG1000:



Callout	Network bond mode
A	Active-Backup mode. Both management ports are bonded into one logical management port connected to the Admin Network.
I	Independent mode. The port on the left is connected to the Admin Network. The port on the right is available for temporary local connections (IP address 169.254.0.1).

In Independent mode, only the management port on the left is connected to the Admin Network. This mode does not provide a redundant path. The management port on the right is unconnected and available for temporary local connections (uses IP address 169.254.0.1)

In Active-Backup mode, both management ports are connected to the Admin Network. Only one port is active at a time. If the active port fails, its backup port automatically provides a failover connection. Bonding these two physical ports into one logical management port provides a redundant path to the Admin Network.



If you need to make a temporary local connection to the services appliance when the 1-GbE management ports are configured for Active-Backup mode, remove the cables from both management ports, plug your temporary cable into the management port on the right, and access the appliance using IP address 169.254.0.1.

Gather installation information

Gather installation information: Overview

As you install and configure a StorageGRID appliance, you make decisions and gather information about Ethernet switch ports, IP addresses, and port and network bond modes.

Refer to the instructions for your appliance to determine what information you need:

- [SGF6112](#)
- [SG6000](#)
- [SG5700](#)
- [SG100 and SG1000](#)

Alternatively, you can work with your NetApp Professional Services consultant to use the NetApp ConfigBuilder tool to streamline and automate the configuration steps. See [Automate appliance installation and configuration](#).

Gather installation information (SGF6112)

Using the following tables, record the required information for each network you connect to the appliance. These values are required to install and configure the hardware.



Instead of using the tables, use the workbook provided with ConfigBuilder. Using the ConfigBuilder workbook allows you to upload your system information and generate a JSON file to automatically complete some configuration steps in the StorageGRID Appliance Installer. See [Automate appliance installation and configuration](#).

Check StorageGRID version

Before installing an SGF6112 appliance, confirm your StorageGRID system is using a required version of StorageGRID software.

Appliance	Required StorageGRID version
SGF6112	11.7 or later (latest hotfix recommended)

Administration and maintenance ports

The Admin Network for StorageGRID is an optional network, used for system administration and maintenance. The appliance connects to the Admin Network using the following ports on the appliance.

The following figure shows the RJ-45 ports on the SG6112 appliance.



Information needed	Your value
Admin Network enabled	Choose one: <ul style="list-style-type: none"> • No • Yes (default)
Network bond mode	Choose one: <ul style="list-style-type: none"> • Independent (default) • Active-Backup
Switch port for the left port circled in the diagram (default active port for Independent network bond mode)	
Switch port for the right port circled in the diagram (Active-Backup network bond mode only)	
MAC address for the Admin Network port Note: The MAC address label on the front of the appliance lists the MAC address for the BMC management port. To determine the MAC address for the Admin Network port, you must add 2 to the hexadecimal number on the label. For example, if the MAC address on the label ends in 09 , the MAC address for the Admin Port would end in 0B . If the MAC address on the label ends in (y)FF , the MAC address for the Admin Port would end in (y+1)01 . You can easily make this calculation by opening Calculator in Windows, setting it to Programmer mode, selecting Hex, typing the MAC address, then typing + 2 = .	
DHCP-assigned IP address for the Admin Network port, if available after power on Note: You can determine the DHCP-assigned IP address by using the MAC address to look up the assigned IP.	<ul style="list-style-type: none"> • IPv4 address (CIDR): • Gateway:
Static IP address you plan to use for the appliance node on the Admin Network Note: If your network does not have a gateway, specify the same static IPv4 address for the gateway.	<ul style="list-style-type: none"> • IPv4 address (CIDR): • Gateway:
Admin Network subnets (CIDR)	

Networking ports

The four networking ports on the appliance connect to the StorageGRID Grid Network and the optional Client Network.

Information needed	Your value
Link speed	For the SGF6112, choose one of the following: <ul style="list-style-type: none">• Auto (default)• 10 GbE• 25 GbE
Port bond mode	Choose one: <ul style="list-style-type: none">• Fixed (default)• Aggregate
Switch port for port 1 (Client Network for Fixed mode)	
Switch port for port 2 (Grid Network for Fixed mode)	
Switch port for port 3 (Client Network for Fixed mode)	
Switch port for port 4 (Grid Network for Fixed mode)	

Grid Network ports

The Grid Network for StorageGRID is a required network, used for all internal StorageGRID traffic. The appliance connects to the Grid Network using the four network ports.

Information needed	Your value
Network bond mode	Choose one: <ul style="list-style-type: none">• Active-Backup (default)• LACP (802.3ad)
VLAN tagging enabled	Choose one: <ul style="list-style-type: none">• No (default)• Yes
VLAN tag (if VLAN tagging is enabled)	Enter a value between 0 and 4095:

Information needed	Your value
DHCP-assigned IP address for the Grid Network, if available after power on	<ul style="list-style-type: none"> • IPv4 address (CIDR): • Gateway:
Static IP address you plan to use for the appliance node on the Grid Network Note: If your network does not have a gateway, specify the same static IPv4 address for the gateway.	<ul style="list-style-type: none"> • IPv4 address (CIDR): • Gateway:
Grid Network subnets (CIDRs)	
Maximum transmission unit (MTU) setting (optional). You can use the default value of 1500, or set the MTU to a value suitable for jumbo frames, such as 9000.	

Client Network ports

The Client Network for StorageGRID is an optional network, typically used to provide client protocol access to the grid. The appliance connects to the Client Network using the four network ports.

Information needed	Your value
Client Network enabled	Choose one: <ul style="list-style-type: none"> • No (default) • Yes
Network bond mode	Choose one: <ul style="list-style-type: none"> • Active-Backup (default) • LACP (802.3ad)
VLAN tagging enabled	Choose one: <ul style="list-style-type: none"> • No (default) • Yes
VLAN tag (If VLAN tagging is enabled)	Enter a value between 0 and 4095:
DHCP-assigned IP address for the Client Network, if available after power on	<ul style="list-style-type: none"> • IPv4 address (CIDR): • Gateway:

Information needed	Your value
<p>Static IP address you plan to use for the appliance node on the Client Network</p> <p>Note: If the Client Network is enabled, the default route on the appliance will use the gateway specified here.</p>	<ul style="list-style-type: none"> • IPv4 address (CIDR): • Gateway:

BMC management network ports

You can access the BMC interface on the appliance using the 1-GbE management port circled in the diagram. This port supports remote management of the controller hardware over Ethernet using the Intelligent Platform Management Interface (IPMI) standard.



You can enable or disable remote IPMI access for all appliances containing a BMC by using the management API private endpoint, PUT /private/bmc.

The following figure shows the BMC management port on the SG6112 appliance.



Information needed	Your value
Ethernet switch port you will connect to the BMC management port (circled in the diagram)	
DHCP-assigned IP address for the BMC management network, if available after power on	<ul style="list-style-type: none"> • IPv4 address (CIDR): • Gateway:
Static IP address you plan to use for the BMC management port	<ul style="list-style-type: none"> • IPv4 address (CIDR): • Gateway:

Related information

- [Cable appliance \(SGF6112\)](#)
- [Configure StorageGRID IP addresses](#)

Gather installation information (SG6000)

Using the tables, record the required information for each network you connect to the appliance. These values are required to install and configure the hardware.



Instead of using the tables, use the workbook provided with ConfigBuilder. Using the ConfigBuilder workbook allows you to upload your system information and generate a JSON file to automatically complete some configuration steps in the StorageGRID Appliance Installer. See [Automate appliance installation and configuration](#).

Information needed to connect to SANtricity System Manager on storage controllers

You connect both of the storage controllers in the appliance (either the E2800 series controllers or the EF570 controllers) to the management network you will use for SANtricity System Manager. The controllers are located in each appliance as follows:

- SG6060 and SG6060X: Controller A is on the top, and controller B is on the bottom.
- SGF6024: Controller A is on the left, and controller B is on the right.

Information needed	Your value for controller A	Your value for controller B
Ethernet switch port you will connect to management port 1 (labeled as P1 on the controller)		
MAC address for management port 1 (printed on a label near port P1)		
DHCP-assigned IP address for management port 1, if available after power on Note: If the network you will connect to the storage controller includes a DHCP server, the network administrator can use the MAC address to determine the IP address that was assigned by the DHCP server.		
Static IP address you plan to use for the appliance on the management network	For IPv4: <ul style="list-style-type: none"> IPv4 address: Subnet mask: Gateway: For IPv6: <ul style="list-style-type: none"> IPv6 address: Routable IP address: storage controller router IP address: 	For IPv4: <ul style="list-style-type: none"> IPv4 address: Subnet mask: Gateway: For IPv6: <ul style="list-style-type: none"> IPv6 address: Routable IP address: storage controller router IP address:
IP address format	Choose one: <ul style="list-style-type: none"> IPv4 IPv6 	Choose one: <ul style="list-style-type: none"> IPv4 IPv6

Information needed	Your value for controller A	Your value for controller B
Speed and duplex mode Note: You must make sure the Ethernet switch for the SANtricity System Manager management network is set to autonegotiate.	Must be: <ul style="list-style-type: none">Autonegotiate (default)	Must be: <ul style="list-style-type: none">Autonegotiate (default)

Information needed to connect SG6000-CN controller to Admin Network

The Admin Network for StorageGRID is an optional network, used for system administration and maintenance. The appliance connects to the Admin Network using the following 1-GbE management ports on the SG6000-CN controller.



Information needed	Your value
Admin Network enabled	Choose one: <ul style="list-style-type: none">NoYes (default)
Network bond mode	Choose one: <ul style="list-style-type: none">Independent (default)Active-Backup
Switch port for the left port in the red circle in the diagram (default active port for Independent network bond mode)	
Switch port for the right port in the red circle in the diagram (Active-Backup network bond mode only)	
MAC address for the Admin Network port Note: The MAC address label on the front of the SG6000-CN controller lists the MAC address for the BMC management port. To determine the MAC address for the Admin Network port, you must add 2 to the hexadecimal number on the label. For example, if the MAC address on the label ends in 09 , the MAC address for the Admin Port would end in 0B . If the MAC address on the label ends in (y)FF , the MAC address for the Admin Port would end in (y+1)01 . You can easily make this calculation by opening Calculator in Windows, setting it to Programmer mode, selecting Hex, typing the MAC address, then typing + 2 = .	

Information needed	Your value
<p>DHCP-assigned IP address for the Admin Network port, if available after power on</p> <p>Note: You can determine the DHCP-assigned IP address by using the MAC address to look up the assigned IP.</p>	<ul style="list-style-type: none"> • IPv4 address (CIDR): • Gateway:
<p>Static IP address you plan to use for the appliance Storage Node on the Admin Network</p> <p>Note: If your network does not have a gateway, specify the same static IPv4 address for the gateway.</p>	<ul style="list-style-type: none"> • IPv4 address (CIDR): • Gateway:
Admin Network subnets (CIDR)	

Information needed to connect and configure 10/25-GbE ports on SG6000-CN controller

The four 10/25-GbE ports on the SG6000-CN controller connect to the StorageGRID Grid Network and the optional Client Network.

Information needed	Your value
Link speed	<p>Choose one:</p> <ul style="list-style-type: none"> • Auto (default) • 10 GbE • 25 GbE
Port bond mode	<p>Choose one:</p> <ul style="list-style-type: none"> • Fixed (default) • Aggregate
Switch port for port 1 (Client Network for Fixed mode)	
Switch port for port 2 (Grid Network for Fixed mode)	
Switch port for port 3 (Client Network for Fixed mode)	
Switch port for port 4 (Grid Network for Fixed mode)	

Information needed to connect SG6000-CN controller to Grid Network

The Grid Network for StorageGRID is a required network, used for all internal StorageGRID traffic. The appliance connects to the Grid Network using the 10/25-GbE ports on the SG6000-CN controller.

Information needed	Your value
Network bond mode	Choose one: <ul style="list-style-type: none"> • Active-Backup (default) • LACP (802.3ad)
VLAN tagging enabled	Choose one: <ul style="list-style-type: none"> • No (default) • Yes
VLAN tag(if VLAN tagging is enabled)	Enter a value between 0 and 4095:
DHCP-assigned IP address for the Grid Network, if available after power on	<ul style="list-style-type: none"> • IPv4 address (CIDR): • Gateway:
Static IP address you plan to use for the appliance Storage Node on the Grid Network Note: If your network does not have a gateway, specify the same static IPv4 address for the gateway.	<ul style="list-style-type: none"> • IPv4 address (CIDR): • Gateway:
Grid Network subnets (CIDRs)	

Information needed to connect SG6000-CN controller to Client Network

The Client Network for StorageGRID is an optional network, typically used to provide client protocol access to the grid. The appliance connects to the Client Network using the 10/25-GbE ports on the SG6000-CN controller.

Information needed	Your value
Client Network enabled	Choose one: <ul style="list-style-type: none"> • No (default) • Yes
Network bond mode	Choose one: <ul style="list-style-type: none"> • Active-Backup (default) • LACP (802.3ad)
VLAN tagging enabled	Choose one: <ul style="list-style-type: none"> • No (default) • Yes

Information needed	Your value
VLAN tag (If VLAN tagging is enabled)	Enter a value between 0 and 4095:
DHCP-assigned IP address for the Client Network, if available after power on	<ul style="list-style-type: none"> • IPv4 address (CIDR): • Gateway:
<p>Static IP address you plan to use for the appliance Storage Node on the Client Network</p> <p>Note: If the Client Network is enabled, the default route on the controller will use the gateway specified here.</p>	<ul style="list-style-type: none"> • IPv4 address (CIDR): • Gateway:

Information needed to connect SG6000-CN controller to BMC management network

You can access the BMC interface on the SG6000-CN controller using the following 1-GbE management port. This port supports remote management of the controller hardware over Ethernet using the Intelligent Platform Management Interface (IPMI) standard.



You can enable or disable remote IPMI access for all appliances containing a BMC by using the management API private endpoint, PUT /private/bmc.

Information needed	Your value
Ethernet switch port you will connect to the BMC management port (circled in the diagram)	
DHCP-assigned IP address for the BMC management network, if available after power on	<ul style="list-style-type: none"> • IPv4 address (CIDR): • Gateway:
Static IP address you plan to use for the BMC management port	<ul style="list-style-type: none"> • IPv4 address (CIDR): • Gateway:

Related information

- [SG6000 controllers](#)
- [Review appliance network connections](#)
- [Port bond modes \(SG6000-CN controller\)](#)
- [Cable appliance \(SG6000\)](#)
- [Configure StorageGRID IP addresses](#)

Gather installation information (SG5700)

Using the tables, record the required information for each network you connect to the appliance. These values are required to install and configure the hardware.

 Instead of using the tables, use the workbook provided with ConfigBuilder. Using the ConfigBuilder workbook allows you to upload your system information and generate a JSON file to automatically complete some configuration steps in the StorageGRID Appliance Installer. See [Automate appliance installation and configuration](#).

Information needed to connect to SANtricity System Manager on E2800 controller

You connect the E2800 series controller to the management network you will use for SANtricity System Manager.

Information needed	Your value
Ethernet switch port you will connect to management port 1	
MAC address for management port 1 (printed on a label near port P1)	
DHCP-assigned IP address for management port 1, if available after power on	
Note: If the network you will connect to the E2800 controller includes a DHCP server, the network administrator can use the MAC address to determine the IP address that was assigned by the DHCP server.	
Speed and duplex mode	Must be: <ul style="list-style-type: none">Autonegotiate (default)
Note: You must make sure the Ethernet switch for the SANtricity System Manager management network is set to autonegotiate.	
IP address format	Choose one: <ul style="list-style-type: none">IPv4IPv6

Information needed	Your value
Static IP address you plan to use for the appliance on the management network	<p>For IPv4:</p> <ul style="list-style-type: none"> • IPv4 address: • Subnet mask: • Gateway: <p>For IPv6:</p> <ul style="list-style-type: none"> • IPv6 address: • Routable IP address: • E2800 controller router IP address:

Information needed to connect E5700SG controller to Admin Network

The Admin Network for StorageGRID is an optional network, used for system administration and maintenance. The appliance connects to the Admin Network using the 1-GbE management ports on the E5700SG controller.

Information needed	Your value
Admin Network enabled	<p>Choose one:</p> <ul style="list-style-type: none"> • No • Yes (default)
Network bond mode	<p>Choose one:</p> <ul style="list-style-type: none"> • Independent • Active-Backup
Switch port for port 1	
Switch port for port 2 (Active-Backup network bond mode only)	
DHCP-assigned IP address for management port 1, if available after power on Note: If the Admin Network includes a DHCP server, the E5700SG controller displays the DHCP-assigned IP address on its seven-segment display after it boots up. You can also determine the DHCP-assigned IP address by using the MAC address to look up the assigned IP.	<ul style="list-style-type: none"> • IPv4 address (CIDR): • Gateway:

Information needed	Your value
<p>Static IP address you plan to use for the appliance Storage Node on the Admin Network</p> <p>Note: If your network does not have a gateway, specify the same static IPv4 address for the gateway.</p>	<ul style="list-style-type: none"> • IPv4 address (CIDR): • Gateway:
Admin Network subnets (CIDR)	

Information needed to connect and configure 10/25-GbE ports on E5700SG controller

The four 10/25-GbE ports on the E5700SG controller connect to the StorageGRID Grid Network and Client Network.



See [Port bond modes \(E5700SG controller\)](#).

Information needed	Your value
<p>Link speed</p> <p>Note: If you select 25 GbE, install SPF28 transceivers. Autonegotiation is not supported, so you must also configure the ports and the connected switches for 25GbE.</p>	<p>Choose one:</p> <ul style="list-style-type: none"> • 10 GbE (default) • 25 GbE
<p>Port bond mode</p>	<p>Choose one:</p> <ul style="list-style-type: none"> • Fixed (default) • Aggregate
Switch port for port 1 (Client Network)	
Switch port for port 2 (Grid Network)	
Switch port for port 3 (Client Network)	
Switch port for port 4 (Grid Network)	

Information needed to connect E5700SG controller to Grid Network

The Grid Network for StorageGRID is a required network, used for all internal StorageGRID traffic. The appliance connects to the Grid Network using the 10/25-GbE ports on the E5700SG controller.



See [Port bond modes \(E5700SG controller\)](#).

Information needed	Your value
Network bond mode	Choose one: <ul style="list-style-type: none"> • Active-Backup (default) • LACP (802.3ad)
VLAN tagging enabled	Choose one: <ul style="list-style-type: none"> • No (default) • Yes
VLAN tag(if VLAN tagging is enabled)	Enter a value between 0 and 4095:
DHCP-assigned IP address for the Grid Network, if available after power on Note: If the Grid Network includes a DHCP server, the E5700SG controller displays the DHCP-assigned IP address for the Grid Network on its seven-segment display after it boots up.	<ul style="list-style-type: none"> • IPv4 address (CIDR): • Gateway:
Static IP address you plan to use for the appliance Storage Node on the Grid Network Note: If your network does not have a gateway, specify the same static IPv4 address for the gateway.	<ul style="list-style-type: none"> • IPv4 address (CIDR): • Gateway:
Grid Network subnets (CIDR) Note: If the Client Network is not enabled, the default route on the controller will use the gateway specified here.	

Information needed to connect E5700SG controller to Client Network

The Client Network for StorageGRID is an optional network, typically used to provide client protocol access to the grid. The appliance connects to the Client Network using the 10/25-GbE ports on the E5700SG controller.



See [Port bond modes \(E5700SG controller\)](#).

Information needed	Your value
Client Network enabled	Choose one: <ul style="list-style-type: none"> • No (default) • Yes

Information needed	Your value
Network bond mode	Choose one: <ul style="list-style-type: none"> • Active-Backup (default) • LACP (802.3ad)
VLAN tagging enabled	Choose one: <ul style="list-style-type: none"> • No (default) • Yes
VLAN tag (if VLAN tagging is enabled)	Enter a value between 0 and 4095:
DHCP-assigned IP address for the Client Network, if available after power on	<ul style="list-style-type: none"> • IPv4 address (CIDR): • Gateway:
Static IP address you plan to use for the appliance Storage Node on the Client Network Note: If the Client Network is enabled, the default route on the controller will use the gateway specified here.	<ul style="list-style-type: none"> • IPv4 address (CIDR): • Gateway:

Related information

- [Network connections \(SG5700\)](#)
- [Port bond modes \(E5700SG controller\)](#)
- [Configure hardware \(SG5700\)](#)

Gather installation information (SG100 and SG1000)

Using the tables, record the required information for each network you connect to the appliance. These values are required to install and configure the hardware.

 Instead of using the tables, use the workbook provided with ConfigBuilder. Using the ConfigBuilder workbook allows you to upload your system information and generate a JSON file to automatically complete some configuration steps in the StorageGRID Appliance Installer. See [Automate appliance installation and configuration](#).

Check StorageGRID version

Before installing an SG100 or SG1000 services appliance, confirm your StorageGRID system is using a required version of StorageGRID software.

Appliance	Required StorageGRID version
SG1000	11.3 or later (latest hotfix recommended)

Appliance	Required StorageGRID version
SG100	11.4 or later (latest hotfix recommended)

Administration and maintenance ports

The Admin Network for StorageGRID is an optional network, used for system administration and maintenance. The appliance connects to the Admin Network using the following 1-GbE management ports on the appliance.

SG100 RJ-45 ports:



SG1000 RJ-45 ports:



Information needed	Your value
Admin Network enabled	<p>Choose one:</p> <ul style="list-style-type: none"> • No • Yes (default)
Network bond mode	<p>Choose one:</p> <ul style="list-style-type: none"> • Independent (default) • Active-Backup
Switch port for the left port circled in the diagram (default active port for Independent network bond mode)	
Switch port for the right port circled in the diagram (Active-Backup network bond mode only)	

Information needed	Your value
<p>MAC address for the Admin Network port</p> <p>Note: The MAC address label on the front of the appliance lists the MAC address for the BMC management port. To determine the MAC address for the Admin Network port, add 2 to the hexadecimal number on the label. For example, if the MAC address on the label ends in 09, the MAC address for the Admin Port would end in 0B. If the MAC address on the label ends in (y)FF, the MAC address for the Admin Port would end in (y+1)01. You can easily make this calculation by opening Calculator in Windows, setting it to Programmer mode, selecting Hex, typing the MAC address, then typing + 2 =.</p>	
<p>DHCP-assigned IP address for the Admin Network port, if available after power on</p> <p>Note: You can determine the DHCP-assigned IP address by using the MAC address to look up the assigned IP.</p>	<ul style="list-style-type: none"> • IPv4 address (CIDR): • Gateway:
<p>Static IP address you plan to use for the appliance node on the Admin Network</p> <p>Note: If your network does not have a gateway, specify the same static IPv4 address for the gateway.</p>	<ul style="list-style-type: none"> • IPv4 address (CIDR): • Gateway:
Admin Network subnets (CIDR)	

Networking ports

The four networking ports on the appliance connect to the StorageGRID Grid Network and the optional Client Network.

Information needed	Your value
Link speed	<p>For the SG100, choose one of the following:</p> <ul style="list-style-type: none"> • Auto (default) • 10 GbE • 25 GbE <p>For the SG1000, choose one of the following:</p> <ul style="list-style-type: none"> • Auto (default) • 10 GbE • 25 GbE • 40 GbE • 100 GbE <p>Note: For the SG1000, 10- and 25-GbE speeds require the use of QSA adapters.</p>
Port bond mode	<p>Choose one:</p> <ul style="list-style-type: none"> • Fixed (default) • Aggregate
Switch port for port 1 (Client Network for Fixed mode)	
Switch port for port 2 (Grid Network for Fixed mode)	
Switch port for port 3 (Client Network for Fixed mode)	
Switch port for port 4 (Grid Network for Fixed mode)	

Grid Network ports

The Grid Network for StorageGRID is a required network, used for all internal StorageGRID traffic. The appliance connects to the Grid Network using the four network ports.

Information needed	Your value
Network bond mode	<p>Choose one:</p> <ul style="list-style-type: none"> • Active-Backup (default) • LACP (802.3ad)

Information needed	Your value
VLAN tagging enabled	Choose one: <ul style="list-style-type: none"> • No (default) • Yes
VLAN tag(if VLAN tagging is enabled)	Enter a value between 0 and 4095:
DHCP-assigned IP address for the Grid Network, if available after power on	<ul style="list-style-type: none"> • IPv4 address (CIDR): • Gateway:
Static IP address you plan to use for the appliance node on the Grid Network Note: If your network does not have a gateway, specify the same static IPv4 address for the gateway.	<ul style="list-style-type: none"> • IPv4 address (CIDR): • Gateway:
Grid Network subnets (CIDRs)	
Maximum transmission unit (MTU) setting (optional)You can use the default value of 1500, or set the MTU to a value suitable for jumbo frames, such as 9000.	

Client Network ports

The Client Network for StorageGRID is an optional network, typically used to provide client protocol access to the grid. The appliance connects to the Client Network using the four network ports.

Information needed	Your value
Client Network enabled	Choose one: <ul style="list-style-type: none"> • No (default) • Yes
Network bond mode	Choose one: <ul style="list-style-type: none"> • Active-Backup (default) • LACP (802.3ad)
VLAN tagging enabled	Choose one: <ul style="list-style-type: none"> • No (default) • Yes
VLAN tag (If VLAN tagging is enabled)	Enter a value between 0 and 4095:

Information needed	Your value
DHCP-assigned IP address for the Client Network, if available after power on	<ul style="list-style-type: none"> • IPv4 address (CIDR): • Gateway:
Static IP address you plan to use for the appliance node on the Client Network Note: If the Client Network is enabled, the default route on the appliance will use the gateway specified here.	<ul style="list-style-type: none"> • IPv4 address (CIDR): • Gateway:

BMC management network ports

You can access the BMC interface on the services appliance using the 1-GbE management port circled in the diagram. This port supports remote management of the controller hardware over Ethernet using the Intelligent Platform Management Interface (IPMI) standard.



You can enable or disable remote IPMI access for all appliances containing a BMC by using the management API private endpoint, PUT /private/bmc.

SG100 BMC management port:



SG1000 BMC management port:



Information needed	Your value
Ethernet switch port you will connect to the BMC management port (circled in the diagram)	
DHCP-assigned IP address for the BMC management network, if available after power on	<ul style="list-style-type: none"> • IPv4 address (CIDR): • Gateway:
Static IP address you plan to use for the BMC management port	<ul style="list-style-type: none"> • IPv4 address (CIDR): • Gateway:

Related information

- [Cable appliance \(SG100 and SG1000\)](#)
- [Configure StorageGRID IP addresses](#)

Automate appliance installation and configuration

Automating installation and configuration can be useful for deploying multiple StorageGRID instances or one large, complex StorageGRID instance.

Using NetApp StorageGRID tools, you can automate the installation and configuration of your StorageGRID appliances. After you install and configure the appliances, you can [automate the configuration of the entire StorageGRID system](#).

You can automate the configuration of the following:

- Grid Network, Admin Network, and Client Network IP addresses
- BMC interface
- Network links
 - Port bond mode
 - Network bond mode
 - Link speed

Automation options

To automate appliance installation and configuration, use one or more of the following options:

- Generate a JSON file that contains configuration details. Work with your NetApp Professional Services consultant to use the [NetApp ConfigBuilder tool](#) to complete these steps:

Step	Consult NetApp Professional Services	Use ConfigBuilder
1	Get sales order number	
2		Get workbook
3	Complete workbook	
4		Upload workbook
5		Generate JSON file
6	Upload JSON file to appliance	
7	Appliance ready for configuration. See Automate using Appliance Installer .	



You can use the same JSON file to configure more than one appliance.

Configuring your appliance using an uploaded JSON file is often more efficient than performing the configuration manually, especially if you have to configure many nodes. Performing the configuration manually requires using multiple pages in the StorageGRID Appliance Installer and applying the configuration file for each node one at a time.

- If you are an advanced user, you can use the following StorageGRID Python scripts to install and configure your system:
 - `configure-sga.py`: Automate the installation and configuration of your appliances. See [Automate appliance installation and configuration using `configure-sga.py` script](#).
 - `configure-storagegrid.py`: Configure other components of the entire StorageGRID system (the "grid"). See [Automate StorageGRID configuration](#).



You can use StorageGRID automation Python scripts directly, or you can use them as examples of how to use the StorageGRID Installation REST API in grid deployment and configuration tools you develop yourself. See the instructions for [downloading and extracting the StorageGRID installation files](#).

Automate appliance configuration using StorageGRID Appliance Installer

After you have generated a JSON file, you can automate the configuration of one or more appliances by using the StorageGRID Appliance Installer to upload the JSON file.

Before you begin

- The appliance has been installed in a rack, connected to your networks, and powered on.
- You have [generated the JSON file](#) with the guidance of your NetApp Professional Services consultant.
- Your appliance contains the latest firmware compatible with StorageGRID 11.5 or higher.
- You are connected to the StorageGRID Appliance Installer on the appliance you are configuring using a [supported web browser](#).

Steps

1. In the StorageGRID Appliance Installer, select **Advanced > Update Appliance Configuration**. The Update Appliance Configuration page appears.
2. Browse for and select the JSON file with the configuration you want to upload.

The file is uploaded and validated. When the validation process is complete, the file name is shown next to a green check mark.



You might lose connection to the appliance if the configuration from the JSON file includes sections for `link_config`, `networks`, or both. If you aren't reconnected within 1 minute, re-enter the appliance URL using one of the other IP addresses assigned to the appliance.

Upload JSON

JSON configuration	<input type="button" value="Browse"/>	appliances.orig.json
Node name	<input type="button" value="-- Select a node ▾"/>	
<input type="button" value="Apply JSON configuration"/>		

The **Node name** drop down is populated with the top-level node names defined in the JSON file.



If the file is not valid, the file name is shown in red and an error message is displayed in a yellow banner. The invalid file is not applied to the appliance. ConfigBuilder verifies that you have a valid JSON file.

3. Select a node from the list in the **Node name** drop down.

The **Apply JSON configuration** button becomes enabled.

4. Select **Apply JSON configuration**.

The configuration is applied to the selected node.

Automate appliance installation and configuration using `configure-sga.py` script

If you are an advanced user, you can use the `configure-sga.py` script to automate many of the installation and configuration tasks for StorageGRID appliance nodes, including installing and configuring a primary Admin Node. This script can be useful if you have a large number of appliances to configure.

You can also use the script to generate a JSON file that contains appliance configuration information. You can upload the JSON file to the StorageGRID Appliance Installer to configure all appliance nodes at the same time. You can also edit the JSON file, then upload it to apply a new configuration to one or more appliances.



This procedure is for advanced users with experience using command-line interfaces. Alternatively, you can [use the StorageGRID Appliance Installer to automate the configuration](#).

Before you begin

- The appliance has been installed in a rack, connected to your networks, and powered on.
- You have [generated the JSON file](#) with the guidance of your NetApp Professional Services consultant.
- Your appliance contains the latest firmware compatible with StorageGRID 11.5 or higher.
- You have configured the IP address of the Admin Network for the appliance.
- You have downloaded the `configure-sga.py` file. The file is included in the installation archive, or you can access it by clicking **Help > Appliance Installation Script** in the StorageGRID Appliance Installer.

Steps

1. Log in to the Linux machine you are using to run the Python script.
2. For general help with the script syntax and to see a list of the available parameters, enter the following:

```
configure-sga.py --help
```

The `configure-sga.py` script uses five subcommands:

- `advanced` for advanced StorageGRID appliance interactions, including BMC configuration and creating a JSON file containing the current configuration of the appliance
- `configure` for configuring the RAID mode, node name, and networking parameters
- `install` for starting a StorageGRID installation
- `monitor` for monitoring a StorageGRID installation

- reboot for rebooting the appliance

If you enter a subcommand (advanced, configure, install, monitor, or reboot) argument followed by the --help option you will get a different help text providing more detail on the options available within that subcommand:

```
configure-sga.py subcommand --help
```

If you will [back up the appliance configuration to a JSON file](#), ensure the node names follow these requirements:

- Each node name is unique if you want to automatically configure all appliance nodes using a JSON file.
- Must be a valid hostname containing at least 1 and no more than 32 characters.
- Can use letters, numbers, and hyphens.
- Can't start or end with a hyphen.
- Can't contain only numbers.

3. To apply the configuration from the JSON file to the appliance, enter the following, where *SGA-INSTALL-IP* is the Admin Network IP address for the appliance, *json-file-name* is the name of the JSON file, and *node-name-inside-json-file* is the name of the node with the configuration being applied:

```
configure-sga.py advanced --restore-file json-file-name --restore-node node-name-inside-json-file SGA-INSTALL-IP
```

4. To confirm the current configuration of the appliance node, enter the following where *SGA-INSTALL-IP* is the Admin Network IP address for the appliance:

```
configure-sga.py configure SGA-INSTALL-IP
```

The results show current IP information for the appliance, including the IP address of the primary Admin Node and information about the Admin, Grid, and Client Networks.

```
Connecting to +https://10.224.2.30:8443+ (Checking version and
connectivity.)
2021/02/25 16:25:11: Performing GET on /api/versions... Received 200
2021/02/25 16:25:11: Performing GET on /api/v2/system-info... Received
200
2021/02/25 16:25:11: Performing GET on /api/v2/admin-connection...
Received 200
2021/02/25 16:25:11: Performing GET on /api/v2/link-config... Received
200
2021/02/25 16:25:11: Performing GET on /api/v2/networks... Received 200
2021/02/25 16:25:11: Performing GET on /api/v2/system-config... Received
200

StorageGRID Appliance
  Name:          LAB-SGA-2-30
  Node type:    storage

StorageGRID primary Admin Node
  IP:           172.16.1.170
```

State: unknown
Message: Initializing...
Version: Unknown

Network Link Configuration

Link Status

Link	State	Speed (Gbps)
1	Up	10
2	Up	10
3	Up	10
4	Up	10
5	Up	1
6	Down	N/A

Link Settings

Port bond mode: FIXED
Link speed: 10GBE

Grid Network: ENABLED
Bonding mode: active-backup
VLAN: novlan
MAC Addresses: 00:a0:98:59:8e:8a 00:a0:98:59:8e:82

Admin Network: ENABLED
Bonding mode: no-bond
MAC Addresses: 00:80:e5:29:70:f4

Client Network: ENABLED
Bonding mode: active-backup
VLAN: novlan
MAC Addresses: 00:a0:98:59:8e:89 00:a0:98:59:8e:81

Grid Network

CIDR: 172.16.2.30/21 (Static)
MAC: 00:A0:98:59:8E:8A
Gateway: 172.16.0.1
Subnets: 172.17.0.0/21
172.18.0.0/21
192.168.0.0/21
MTU: 1500

Admin Network

CIDR: 10.224.2.30/21 (Static)
MAC: 00:80:E5:29:70:F4
Gateway: 10.224.0.1

```
Subnets: 10.0.0.0/8
          172.19.0.0/16
          172.21.0.0/16
MTU: 1500
```

```
Client Network
CIDR: 47.47.2.30/21 (Static)
MAC: 00:A0:98:59:8E:89
Gateway: 47.47.0.1
MTU: 2000
```

```
#####
##### If you are satisfied with this configuration, #####
##### execute the script with the "install" sub-command. #####
#####
```

5. If you need to change any of the values in the current configuration, use the `configure` subcommand to update them. For example, if you want to change the IP address that the appliance uses for connection to the primary Admin Node to 172.16.2.99, enter the following:

```
configure-sga.py configure --admin-ip 172.16.2.99 _SGA-INSTALL-IP_
```

6. If you want to back up the appliance configuration to a JSON file, use the `advanced` and `backup-file` subcommands. For example, if you want to back up the configuration of an appliance with IP address `SGA-INSTALL-IP` to a file named `appliance-SG1000.json`, enter the following:
`configure-sga.py advanced --backup-file appliance-SG1000.json SGA-INSTALL-IP`

The JSON file containing the configuration information is written to the same directory you executed the script from.



Check that the top-level node name in the generated JSON file matches the appliance name. Don't make any changes to this file unless you are an experienced user and have a thorough understanding of StorageGRID APIs.

7. When you are satisfied with the appliance configuration, use the `install` and `monitor` subcommands to install the appliance:

```
configure-sga.py install --monitor SGA-INSTALL-IP
```

8. If you want to reboot the appliance, enter the following:

```
configure-sga.py reboot SGA-INSTALL-IP
```

Automate StorageGRID configuration

After you have installed and configured 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
configure-storagegrid.py	Python script used to automate the configuration
configure-storagegrid.sample.json	Sample configuration file for use with the script
configure-storagegrid.blank.json	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 sample configuration file (`configure-storagegrid.sample.json`) or the blank configuration file (`configure-storagegrid.blank.json`).

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

1. Log in to the Linux machine you are using to run the Python script.
2. 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
```

After you finish

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, you need to extract 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.

Overview of installation REST APIs

StorageGRID provides two REST APIs for performing installation tasks: the StorageGRID Installation API and the StorageGRID Appliance Installer API.

Both APIs use 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 Docs 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.

StorageGRID Appliance Installer API

The StorageGRID Appliance Installer API can be accessed over HTTPS from *Controller_IP*:8443.

To access the API documentation, go to the StorageGRID Appliance Installer on the appliance and select **Help** > **API Docs** from the menu bar.

The StorageGRID Appliance Installer API includes the following sections:

- **clone**: Operations to configure and control node cloning.
- **encryption**: Operations to manage encryption and view encryption status.
- **hardware config**: Operations to configure system settings on attached hardware.
- **installation**: Operations for starting the appliance installation and for monitoring installation status.
- **networking**: Operations related to the Grid, Admin, and Client Network configuration for a StorageGRID appliance and appliance port settings.
- **setup**: Operations to help with initial appliance installation setup including requests to get information about the system and update the primary Admin Node IP.
- **support**: Operations for rebooting the controller and getting logs.
- **update-config**: Operations to update StorageGRID appliance configuration.
- **upgrade**: Operations related to upgrading appliance firmware.
- **uploadsg**: Operations for uploading StorageGRID installation files.

Install hardware

Register hardware

Registering the appliance hardware provides support benefits.

Steps

1. Locate the chassis serial number for the appliance. For SG6000 appliances the chassis serial number is on the storage controller shelf.

You can find the number on the packing slip, in your confirmation email, or on the appliance after you unpack it.



There are several serial numbers on the SG6000 storage appliance. The serial number on the storage controller shelf is the one that must be registered and used if you call for service or support on the SG6000 appliance.

2. Go to the NetApp Support Site at mysupport.netapp.com.

3. Determine whether you need to register the hardware:

If you are a...	Follow these steps...
Existing NetApp customer	<ol style="list-style-type: none">a. Sign in with your username and password.b. Select Products > My Products.c. Confirm that the new serial number is listed.d. If it is not, follow the instructions for new NetApp customers.
New NetApp customer	<ol style="list-style-type: none">a. Click Register Now, and create an account.b. Select Products > Register Products.c. Enter the product serial number and requested details. <p>After your registration is approved, you can download any required software. The approval process might take up to 24 hours.</p>

Install into cabinet or rack

Install into cabinet or rack (SGF6112)


You install a set of rails for the appliance in your cabinet or rack, and then slide the appliance onto the rails.

Before you begin

- You have reviewed the Safety Notices document included in the box, and understand the precautions for moving and installing hardware.
- You have the instructions packaged with the rail kit.


Steps

1. Carefully follow the instructions for the rail kit to install the rails in your cabinet or rack.
2. On the two rails installed in the cabinet or rack, extend the movable parts of the rails until you hear a click.



3. Insert the appliance into the rails.
4. Slide the appliance into the cabinet or rack.

When you can't move the appliance any further, pull the blue latches on both sides of the chassis to slide the appliance all the way in.



5. Tighten the captive screws on the appliance front panel to secure the appliance in the rack.



Don't attach the front bezel until after you power on the appliance.

SG6000

Install into cabinet or rack (SG6000)

For the SG6060 and SGF6024, you install rails in your cabinet or rack and slide the controller shelf, any expansion shelves, and the compute controller onto the rails. For the SG6060, don't install the drives in each shelf until the shelves are installed.

Model	Install	For information
SG6060	60-drive controller shelf and any 60-drive expansion shelves	Install 60-drive shelves
SG6060	60 drives into each shelf	Install drives

Model	Install	For information
SGF6024	24-drive controller shelf	Install 24-drive shelves
SG6060 and SGF6024	SG6000-CN compute controller	Install SG6000-CN controller

Install 60-drive shelves (SG6060)

You install a set of rails for the E2860 controller shelf in your cabinet or rack, and then slide the controller shelf onto the rails. If you are installing 60-drive expansion shelves, the same procedure applies.

Before you begin

- You have reviewed the Safety Notices document included in the box, and understand the precautions for moving and installing hardware.
- You have the instructions packaged with the rail kit.



Each 60-drive shelf weighs approximately 132 lb (60 kg) without drives installed. Four people or a mechanized lift are required to safely move the shelf.



To avoid damaging the hardware, never move the shelf if drives are installed. You must remove all drives before moving the shelf.



When installing the E2860 controller shelf or optional expansion shelves, install hardware from the bottom to the top of the rack or cabinet to prevent the equipment from tipping over. To ensure that the heaviest equipment is at the bottom of the cabinet or rack, install the SG6000-CN controller above the E2860 controller shelf and expansion shelves.



Before committing to the installation, verify that the 0.5m optic cables shipped with the appliance, or cables that you supply, are long enough for the planned layout.


Steps

1. Carefully follow the instructions for the rail kit to install the rails in your cabinet or rack.

For square hole cabinets, first install the provided cage nuts to secure the front and rear of the shelf with screws.

2. Remove the outer packing box for the appliance. Then, fold down the flaps on the inner box.
3. If you are lifting the appliance by hand, attach the four handles to the sides of the chassis.

Push up on each handle until it clicks into place.



4. Place the back of the shelf (the end with the connectors) on the rails.
5. Supporting the shelf from the bottom, slide it into the cabinet. If you are using the handles, use the thumb latches to detach one handle at a time as you slide the shelf in.


To remove the handles, pull back on the release latch, push down, then pull away from the shelf.

6. Secure the shelf to the front of the cabinet.

Insert screws into the first and third holes from the top of the shelf on both sides.

7. Secure the shelf to the rear of the cabinet.

Place two back brackets on each side of the upper rear section of the shelf. Insert screws into the first and third holes of each bracket.



8. Repeat these steps for any expansion shelves.

Install drives (SG6060)

After installing the 60-drive shelf into a cabinet or rack, install all 60 drives into the shelf. The shipment for the E2860 controller shelf includes two SSD drives, which you should install in the top drawer of the controller shelf. Each optional expansion shelf includes 60 HDD drives and no SSD drives.

Before you begin

You have installed the E2860 controller shelf or optional expansion shelves (one or two) in the cabinet or rack.



To avoid damaging the hardware, never move the shelf if drives are installed. You must remove all drives before moving the shelf.


Steps

1. Wrap the strap end of the ESD wristband around your wrist, and secure the clip end to a metal ground to prevent static discharge.
2. Remove the drives from their packaging.
3. Release the levers on the top drive drawer, and slide the drawer out using the levers.
4. Locate the two SSD drives.



Expansion shelves don't use SSD drives.

5. Raise each drive handle to a vertical position.
6. Install the two SSD drives in slots 0 and 1 (the first two slots along the lefthand side of the drawer).
7. Gently position each drive into its slot, and lower the raised drive handle until it clicks into place.



8. Install 10 HDD drives into the top drawer.

9. Slide the drawer back in by pushing on the center and closing both levers gently.



Stop pushing the drawer if you feel binding. Use the release levers at the front of the drawer to slide the drawer back out. Then, carefully reinsert the drawer into the slot.

10. Repeat these steps to install HDD drives into the other four drawers.



You must install all 60 drives to ensure correct operation.

11. Attach the front bezel to the shelf.

12. If you have expansion shelves, repeat these steps to install 12 HDD drives into each drawer of each expansion shelf.

13. Proceed to the instructions for installing the SG6000-CN into a cabinet or rack.

Install 24-drive shelves (SGF6024)

You install a set of rails for the EF570 controller shelf in your cabinet or rack, and then slide the array onto the rails.

Before you begin

- You have reviewed the Safety Notices document included in the box, and understand the precautions for moving and installing hardware.
- You have the instructions packaged with the rail kit.

Steps

1. Carefully follow the instructions for the rail kit to install the rails in your cabinet or rack.

For square hole cabinets, first install the provided cage nuts to secure the front and rear of the shelf with screws.

2. Remove the outer packing box for the appliance. Then, fold down the flaps on the inner box.
3. Place the back of the shelf (the end with the connectors) on the rails.



A fully loaded shelf weighs approximately 52 lb (24 kg). Two people are required to safely move the enclosure.

4. Carefully slide the enclosure all the way onto the rails.



You might need to adjust the rails to ensure that the enclosure slides all the way onto the rails.




Don't place additional equipment on the rails after you finish installing the enclosure. The rails aren't designed to bear additional weight.




If applicable, you might need to remove the shelf end caps or the system bezel to secure the enclosure to the rack post; if so, you need to replace the end caps or bezel when you are done.

5. Secure the enclosure to the front of the cabinet or rack and rails by inserting two M5 screws through the mounting brackets (preinstalled on either side of the front of the enclosure), the holes on the rack or system cabinet, and the holes on the front of rails.



6. Secure the enclosure to the back of the rails by inserting two M5 screws through the brackets at the enclosure and the rail kit bracket.
7. If applicable, replace the shelf end caps or the system bezel.



Install SG6000-CN controller (SG6060 and SG6024)


You install a set of rails for the SG6000-CN controller in your cabinet or rack, and then slide the controller onto the rails.

Before you begin

- You have reviewed the Safety Notices document included in the box, and understand the precautions for moving and installing hardware.
- You have the instructions packaged with the rail kit.
- You have installed the E2860 controller shelf and drives or the EF570 controller shelf.


Steps

1. Carefully follow the instructions for the rail kit to install the rails in your cabinet or rack.
2. On the two rails installed in the cabinet or rack, extend the movable parts of the rails until you hear a click.



3. Insert the SG6000-CN controller into the rails.
4. Slide the controller into the cabinet or rack.

When you can't move the controller any further, pull the blue latches on both sides of the chassis to slide the controller all the way in.



Don't attach the front bezel until after you power on the controller.

5. Tighten the captive screws on the controller front panel to secure the controller in the rack.



Install into cabinet or rack (SG5700)

You install a set of rails in your cabinet or rack and then slide the appliance onto the rails. If you have an SG5760, install the drives after installing the appliance.

Before you begin

- You have reviewed the Safety Notices document included in the box, and understand the precautions for moving and installing hardware.
- You have the instructions packaged with the rail kit.

Install SG5712

Follow these steps to install an SG5712 appliance into a rack or cabinet.



The SG5712 weighs approximately 64 lb (29 kg) when fully loaded with drives. Two people or a mechanized lift are required to safely move the SG5712.



Install hardware from the bottom of the rack or cabinet or rack up to prevent the equipment from tipping over.

Steps

1. Follow the instructions for the rail kit to install the rails.
2. Place the back of the appliance (the end with the connectors) on the rails.
3. Carefully slide the appliance all the way back into the cabinet or rack.
4. Secure the appliance to the cabinet or rack as directed in the rail kit instructions.
5. Attach the bezel to the front.

Install SG5760

Follow these steps to install an SG5760 appliance and any expansion shelves into a rack or cabinet.



Install hardware from the bottom of the rack or cabinet or rack up to prevent the equipment from tipping over.



The SG5760 weighs approximately 132 lb (60 kg) with no drives installed. Four people or a mechanized lift are required to safely move an empty SG5760.



To avoid damaging the hardware, never move an SG5760 if drives are installed. You must remove all drives before moving the shelf.

Steps

1. Follow the instructions for the rail kit to install the rails in your cabinet or rack.
2. Prepare to move the appliance:
 - a. Remove the outer packing box.
 - b. Fold down the flaps on the inner box.
 - c. If you are lifting the SG5760 by hand, attach the four handles to the sides of the chassis.

You remove these handles as you slide the appliance onto the rails.


3. If your cabinet has square hole, install the cage nuts so that you can secure the front and rear of the shelf with screws.
4. Place the back of the appliance (the end with the connectors) on the rails.
5. Supporting the appliance from the bottom, slide it into the rack or cabinet.

Use the thumb latches to detach the handles as you slide the appliance in.

6. Secure the appliance to the front of the rack by inserting two screws in the first and third holes (counting down from the top) on each side.
7. Secure the appliance to the rear of the rack or cabinet with the brackets.
8. Install 12 drives in each of the five drive drawers.

You must install all 60 drives to ensure correct operation.

- a. Put on the ESD wristband, and remove the drives from their packaging.
- b. Release the levers on the top drive drawer, and slide the drawer out using the levers.
- c. Raise the drive handle to vertical, and align the buttons on the drive with the notches on the drawer.



- d. Pressing gently on the top of the drive, rotate the drive handle down until the drive snaps into place.
- e. After installing the first 12 drives, slide the drawer back in by pushing on the center and closing both levers gently.
- f. Repeat these steps for the other four drawers.

9. Attach the front bezel.

Install into cabinet or rack (SG100 and SG1000)


You install a set of rails for the appliance in your cabinet or rack, and then slide the appliance onto the rails.

Before you begin

- You have reviewed the Safety Notices document included in the box, and understand the precautions for moving and installing hardware.
- You have the instructions packaged with the rail kit.


Steps

1. Carefully follow the instructions for the rail kit to install the rails in your cabinet or rack.
2. On the two rails installed in the cabinet or rack, extend the movable parts of the rails until you hear a click.



3. Insert the appliance into the rails.
4. Slide the appliance into the cabinet or rack.

When you can't move the appliance any further, pull the blue latches on both sides of the chassis to slide the appliance all the way in.



Don't attach the front bezel until after you power on the appliance.

Cable appliance

Cable appliance (SGF6112)

You connect the management port on the appliance to the service laptop and connect the network ports on the appliance to the Grid Network and optional Client Network for StorageGRID.

Before you begin


- You have an RJ-45 Ethernet cable for connecting the management port.
- You have one of the following options for the network ports. These items aren't provided with the appliance.
 - One to four TwinAx cables for connecting the four network ports.
 - One to four SFP+ or SFP28 transceivers if you plan to use optical cables for the ports.



Risk of exposure to laser radiation — Don't disassemble or remove any part of an SFP transceiver. You might be exposed to laser radiation.

About this task

The following figures show the ports on the back of the SGF6112.



Callout	Port	Type of port	Use
1	BMC management port on the appliance	1-GbE (RJ-45)	Connects to the network where you access the BMC interface.
2	Four 10/25-GbE network ports on the appliance		Connect to the Grid Network and the Client Network for StorageGRID.
3	Admin Network port on the appliance (labeled P1 in the figure)	1-GbE (RJ-45) Important: This port operates only at 1/10-GbE (RJ-45) and does not support 100-megabit speeds.	Connects the appliance to the Admin Network for StorageGRID.
	Rightmost RJ-45 port on the appliance	1-GbE (RJ-45) Important: This port operates only at 1/10-GbE (RJ-45) and does not support 100-megabit speeds.	<ul style="list-style-type: none"> Can be bonded with management port 1 if you want a redundant connection to the Admin Network. Can be left disconnected and available for temporary local access (IP 169.254.0.1). During installation, can be used to connect the appliance to a service laptop if DHCP-assigned IP addresses aren't available.

Steps

1. Connect the BMC management port on the appliance to the management network, using an Ethernet cable.

Although this connection is optional, it is recommended to facilitate support.

2. Connect the network ports on the appliance to the appropriate network switches, using TwinAx cables or optical cables and transceivers.

All four network ports must use the same link speed.



SGF6112 link speed (GbE)	Required equipment
10	SFP+ transceiver
25	SFP28 transceiver

- If you plan to use Fixed port bond mode (default), connect the ports to the StorageGRID Grid and Client Networks, as shown in the table.

Port	Connects to...
Port 1	Client Network (optional)
Port 2	Grid Network
Port 3	Client Network (optional)
Port 4	Grid Network

- If you plan to use the Aggregate port bond mode, connect one or more of the network ports to one or more switches. You should connect at least two of the four ports to avoid having a single point of failure. If you use more than one switch for a single LACP bond, the switches must support MLAG or equivalent.

3. If you plan to use the Admin Network for StorageGRID, connect the Admin Network port on the appliance to the Admin Network, using an Ethernet cable.

Cable appliance (SG6000)

You connect the storage controllers to the SG6000-CN controller, connect the management ports on all three controllers, and connect the network ports on the SG6000-CN controller to the Grid Network and optional Client Network for StorageGRID.

Before you begin

- You have the four optical cables provided with the appliance for connecting the two storage controllers to the SG6000-CN controller.
- You have RJ-45 Ethernet cables (four minimum) for connecting the management ports.
- You have one of the following options for the network ports. These items aren't provided with the appliance.
 - One to four TwinAx cables for connecting the four network ports.
 - One to four SFP+ or SFP28 transceivers if you plan to use optical cables for the ports.



Risk of exposure to laser radiation — Don't disassemble or remove any part of an SFP transceiver. You might be exposed to laser radiation.


About this task

The following figures show the three controllers in the SG6060 and SG6060X appliances, with the SG6000-CN compute controller on the top and the two E2800 storage controllers on the bottom. The SG6060 uses E2800A controllers, and the SG6060X uses E2800B controllers.


 Both versions of the E2800 controller have identical specifications and function except for the location of the interconnect ports.

 Don't use an E2800A and E2800B controller in the same appliance.

SG6060 connections:




SG6060X connections:



The following figure shows the three controllers in the SGF6024 appliance, with the SG6000-CN compute controller on the top and the two EF570 storage controllers side by side below the compute controller.

SGF6024 connections:



Callout	Port	Type of port	Use
1	BMC management port on the SG6000-CN controller	1-GbE (RJ-45)	Connects to the network where you access the BMC interface.

Callout	Port	Type of port	Use
2	FC connection ports: <ul style="list-style-type: none"> 4 on the SG6000-CN controller 2 on each storage controller 	16-Gb/s FC optical SFP+	Connect each storage controller to the SG6000-CN controller.
3	Four network ports on the SG6000-CN controller	10/25-GbE	Connect to the Grid Network and the Client Network for StorageGRID.
4	Admin Network port on the SG6000-CN controller (labeled P1 in the figure)	1-GbE (RJ-45) Important: This port operates only at 1000 baseT/full and does not support 10- or 100-megabit speeds.	Connects the SG6000-CN controller to the Admin Network for StorageGRID.
	Rightmost RJ-45 port on the SG6000-CN controller	1-GbE (RJ-45) Important: This port operates only at 1000 baseT/full and does not support 10- or 100-megabit speeds.	<ul style="list-style-type: none"> Can be bonded with management port 1 if you want a redundant connection to the Admin Network. Can be left unwired and available for temporary local access (IP 169.254.0.1). During installation, can be used to connect the SG6000-CN controller to a service laptop if DHCP-assigned IP addresses aren't available.
5	Management port 1 on each storage controller	1-GbE (RJ-45)	Connects to the network where you access SANtricity System Manager.
	Management port 2 on each storage controller	1-GbE (RJ-45)	Reserved for technical support.

Steps

1. Connect the BMC management port on the SG6000-CN controller to the management network, using an Ethernet cable.

Although this connection is optional, it is recommended to facilitate support.
2. Connect the two FC ports on each storage controller to the FC ports on the SG6000-CN controller, using four optical cables and four SFP+ transceivers for the storage controllers.
3. Connect the network ports on the SG6000-CN controller to the appropriate network switches, using TwinAx cables or optical cables and SFP+ or SFP28 transceivers.



The four network ports must use the same link speed. Install SFP+ transceivers if you plan to use 10-GbE link speeds. Install SFP28 transceivers if you plan to use 25-GbE link speeds.

- If you plan to use Fixed port bond mode (default), connect the ports to the StorageGRID Grid and Client Networks, as shown in the table.

Port	Connects to...
Port 1	Client Network (optional)
Port 2	Grid Network
Port 3	Client Network (optional)
Port 4	Grid Network

- If you plan to use the Aggregate port bond mode, connect one or more of the network ports to one or more switches. You should connect at least two of the four ports to avoid having a single point of failure. If you use more than one switch for a single LACP bond, the switches must support MLAG or equivalent.

4. If you plan to use the Admin Network for StorageGRID, connect the Admin Network port on the SG6000-CN controller to the Admin Network, using an Ethernet cable.
5. If you plan to use the management network for SANtricity System Manager, connect management port 1 (P1) on each storage controller (the RJ-45 port on the left) to the management network for SANtricity System Manager, using an Ethernet cable.

Don't use management port 2 (P2) on the storage controllers (the RJ-45 port on the right). This port is reserved for technical support.

Related information

[Port bond modes \(SG6000-CN controller\)](#)

Cable appliance (SG5700)

You connect the two controllers to each other, connect the management ports on each controller, and connect the 10/25-GbE ports on the E5700SG controller to the Grid Network and optional Client Network for StorageGRID.

Before you begin

- You have unpacked the following items, which are included with the appliance:
 - Two power cords.
 - Two optical cables for the FC interconnect ports on the controllers.
 - Eight SFP+ transceivers, which support either 10-GbE or 16-Gbps FC. The transceivers can be used with the two interconnect ports on both controllers and with the four 10/25-GbE network ports on the E5700SG controller, assuming you want the network ports to use a 10-GbE link speed.
- You have obtained the following items, which aren't included with the appliance:

- One to four optical cables for the 10/25-GbE ports you plan to use.
- One to four SFP28 transceivers, if you plan to use 25-GbE link speed.
- Ethernet cables for connecting the management ports.




Risk of exposure to laser radiation — Don't disassemble or remove any part of an SFP transceiver. You might be exposed to laser radiation.


About this task

The figures show the two controllers in the SG5760 and SG5760X, with the E2800 series storage controller on the top and the E5700SG controller on the bottom. In the SG5712 and SG5712X, the E2800 series storage controller is to the left of the E5700SG controller when viewed from the back.

SG5760 connections:



SG5760X connections:



Callout	Port	Type of port	Use
1	Two interconnect ports on each controller	16Gb/s FC optical SFP+	Connect the two controllers to each other.
2	Management port 1 on the E2800 series controller	1-GbE (RJ-45)	Connects to the network where you access SANtricity System Manager. You can use the Admin Network for StorageGRID or an independent management network.
2	Management port 2 on the E2800 series controller	1-GbE (RJ-45)	Reserved for technical support.
3	Management port 1 on the E5700SG controller	1-GbE (RJ-45)	Connects the E5700SG controller to the Admin Network for StorageGRID.
3	Management port 2 on the E5700SG controller	1-GbE (RJ-45)	<ul style="list-style-type: none"> Can be bonded with management port 1 if you want a redundant connection to the Admin Network. Can be left unwired and available for temporary local access (IP 169.254.0.1). During installation, can be used to connect the E5700SG controller to a service laptop if DHCP-assigned IP addresses aren't available.
4	10/25-GbE ports 1-4 on the E5700SG controller	10-GbE or 25-GbE Note: The SFP+ transceivers included with the appliance support 10-GbE link speeds. If you want to use 25-GbE link speeds for the four network ports, you must provide SFP28 transceivers.	Connect to the Grid Network and the Client Network for StorageGRID. See Port bond modes (E5700SG controller) .

Steps

1. Connect the E2800 controller to the E5700SG controller, using two optical cables and four of the eight SFP+ transceivers.

Connect this port...	To this port...
Interconnect port 1 on the E2800 controller	Interconnect port 1 on the E5700SG controller
Interconnect port 2 on the E2800 controller	Interconnect port 2 on the E5700SG controller

2. If you plan to use SANtricity System Manager, connect management port 1 (P1) on the E2800 controller (the RJ-45 port on the left) to the management network for SANtricity System Manager, using an Ethernet cable.

Don't use management port 2 (P2) on the E2800 controller (the RJ-45 port on the right). This port is reserved for technical support.

3. If you plan to use the Admin Network for StorageGRID, connect management port 1 on the E5700SG controller (the RJ-45 port on the left) to the Admin Network, using an Ethernet cable.

If you plan to use active-backup network bond mode for the Admin Network, connect management port 2 on the E5700SG controller (the RJ-45 port on the right) to the Admin Network, using an Ethernet cable.

4. Connect the 10/25-GbE ports on the E5700SG controller to the appropriate network switches, using optical cables and SFP+ or SFP28 transceivers.



All ports must use the same link speed. Install SFP+ transceivers if you plan to use 10-GbE link speeds. Install SFP28 transceivers if you plan to use 25-GbE link speeds.

- If you plan to use Fixed port bond mode (default), connect the ports to the StorageGRID Grid and Client Networks, as shown in the table.

Port	Connects to...
Port 1	Client Network (optional)
Port 2	Grid Network
Port 3	Client Network (optional)
Port 4	Grid Network

- If you plan to use the Aggregate port bond mode, connect one or more of the network ports to one or more switches. You should connect at least two of the four ports to avoid having a single point of failure. If you use more than one switch for a single LACP bond, the switches must support MLAG or equivalent.

Related information

[Access StorageGRID Appliance Installer](#)

Cable appliance (SG100 and SG1000)

You must connect the management port on the appliance to the service laptop and connect the network ports on the appliance to the Grid Network and optional Client

Network for StorageGRID.

Before you begin

- You have an RJ-45 Ethernet cable for connecting the management port.
- You have one of the following options for the network ports. These items aren't provided with the appliance.
 - One to four TwinAx cables for connecting the four network ports.
 - For the SG100, one to four SFP+ or SFP28 transceivers if you plan to use optical cables for the ports.
 - For the SG1000, one to four QSFP+ or QSFP28 transceivers if you plan to use optical cables for the ports.




Risk of exposure to laser radiation — Don't disassemble or remove any part of a SFP or QSFP transceiver. You might be exposed to laser radiation.


About this task

The following figures show the ports on the back of the appliance.

SG100 port connections:



SG1000 port connections:



Callout	Port	Type of port	Use
1	BMC management port on the appliance	1-GbE (RJ-45)	Connects to the network where you access the BMC interface.
2	Four network ports on the appliance	<ul style="list-style-type: none">• For the SG100: 10/25-GbE• For the SG1000: 10/25/40/100-GbE	Connect to the Grid Network and the Client Network for StorageGRID.

Callout	Port	Type of port	Use
3	Admin Network port on the appliance (labeled P1 in the figures)	1-GbE (RJ-45) Important: This port operates only at 1000 baseT/full and does not support 10- or 100-megabit speeds.	Connects the appliance to the Admin Network for StorageGRID.
	Rightmost RJ-45 port on the appliance	1-GbE (RJ-45) Important: This port operates only at 1000 baseT/full and does not support 10- or 100-megabit speeds.	<ul style="list-style-type: none"> • Can be bonded with management port 1 if you want a redundant connection to the Admin Network. • Can be left disconnected and available for temporary local access (IP 169.254.0.1). • During installation, can be used to connect the appliance to a service laptop if DHCP-assigned IP addresses aren't available.

Steps

1. Connect the BMC management port on the appliance to the management network, using an Ethernet cable.

Although this connection is optional, it is recommended to facilitate support.

2. Connect the network ports on the appliance to the appropriate network switches, using TwinAx cables or optical cables and transceivers.

All four network ports must use the same link speed. See the following table for the equipment required for your hardware and link speed.



SG100 link speed (GbE)	Required equipment
10	SFP+ transceiver
25	SFP28 transceiver
SG1000 link speed (GbE)	Required equipment
10	QSA and SFP+ transceiver
25	QSA and SFP28 transceiver
40	QSFP+ transceiver
100	QFSP28 transceiver

- If you plan to use Fixed port bond mode (default), connect the ports to the StorageGRID Grid and Client Networks, as shown in the table.

Port	Connects to...
Port 1	Client Network (optional)
Port 2	Grid Network
Port 3	Client Network (optional)
Port 4	Grid Network

- If you plan to use the Aggregate port bond mode, connect one or more of the network ports to one or more switches. You should connect at least two of the four ports to avoid having a single point of failure. If you use more than one switch for a single LACP bond, the switches must support MLAG or equivalent.

3. If you plan to use the Admin Network for StorageGRID, connect the Admin Network port on the appliance to the Admin Network, using an Ethernet cable.

Connect power cords and apply power

Connect power cords and apply power (SGF6112)

After connecting the network cables, you are ready to apply power to the appliance.

Steps

1. Connect a power cord to each of the two power supply units in the appliance.
2. Connect these two power cords to two different power distribution units (PDUs) in the cabinet or rack.
3. If the power button on the front of the appliance is not currently illuminated blue, press the button to turn on power to the appliance.

Don't press the power button again during the power-on process.

The LED on the power supply should be illuminated green without blinking.

4. If errors occur, correct any issues.
5. Attach the front bezel to the appliance if removed.

Related information

[View status indicators](#)

Connect power cords and apply power (SG6000)

After connecting the network cables, you are ready to apply power to the SG6000-CN controller and to the two storage controllers or optional expansion shelves.

Steps

1. Confirm that both controllers in the storage controller shelf are off.



Risk of electrical shock — Before connecting the power cords, make sure that the power switches for each of the two storage controllers are off.

2. If you have expansion shelves, confirm that both of the IOM power switches are off.



Risk of electrical shock — Before connecting the power cords, make sure that the two power switches for each of the expansion shelves are off.

3. Connect a power cord to each of the two power supply units in the SG6000-CN controller.
4. Connect these two power cords to two different power distribution units (PDUs) in the cabinet or rack.
5. Connect a power cord to each of the two power supply units in the storage controller shelf.
6. If you have expansion shelves, connect a power cord to each of the two power supply units in each expansion shelf.
7. Connect the two power cords in each storage shelf (including the optional expansion shelves) to two different PDUs in the cabinet or rack.
8. If the power button on the front of the SG6000-CN controller is not currently illuminated blue, press the button to turn on power to the controller.

Don't press the power button again during the power-on process.

9. Turn on the two power switches on the back of the storage controller shelf. If you have expansion shelves, turn on the two power switches for each shelf.
 - Don't turn off the power switches during the power-on process.
 - The fans in the storage controller shelf and optional expansion shelves might be very loud when they first start up. The loud noise during start-up is normal.
10. After the components have booted up, check their status.
 - Check the seven-segment display on the back of each storage controller. Refer to the article about viewing boot-up status codes for more information.
 - Verify that the power button on the front of the SG6000-CN controller is lit.
11. If errors occur, correct any issues.
12. Attach the front bezel to the SG6000-CN controller if removed.

Related information

- [View status indicators](#)
- [Reinstall SG6000-CN controller into cabinet or rack](#)

Connect power cords and apply power (SG5700)

When you apply power to the appliance, both controllers boot up.

Before you begin

Both appliance power switches must be off before connecting power.



Risk of electrical shock — Before connecting the power cords, make sure that the two power switches on the appliance are off.

Steps

1. Confirm that the two power switches on the appliance are off.
2. Connect the two power cords to the appliance.
3. Connect the two power cords to different power distribution units (PDUs) in the cabinet or rack.
4. Turn on the two power switches on the appliance.
 - Don't turn off the power switches during the power-on process.
 - The fans are very loud when they first start up. The loud noise during start-up is normal.
5. After the controllers have booted up, check their seven-segment displays.

Connect power cords and apply power (SG100 and SG1000)

After connecting the network cables, you are ready to apply power to the appliance.

Steps

1. Connect a power cord to each of the two power supply units in the appliance.
2. Connect these two power cords to two different power distribution units (PDUs) in the cabinet or rack.
3. If the power button on the front of the appliance is not currently illuminated blue, press the button to turn on power to the appliance.

Don't press the power button again during the power-on process.

4. If errors occur, correct any issues.
5. Attach the front bezel to the appliance if removed.

Related information

[View status indicators](#)

View status indicators and codes

The appliances and controllers include indicators that help you determine the status of the appliance components.

SGF6112

The appliance includes indicators that help you determine the status of the appliance controller and the SSDs:

- [Appliance indicators and buttons](#)
- [General boot-up codes](#)
- [SSD indicators](#)

Use this information to help [troubleshoot SGF6112 hardware installation](#).

Appliance indicators and buttons

The following figure shows indicators and buttons on the SG6112 appliance.



Callout	Display	State
1	Power button	<ul style="list-style-type: none">• Blue: the appliance is powered on.• Off: the appliance is powered off.
2	Reset button	Use this button to perform a hard reset of the controller.
3	Identify button	Using the BMC, this button can be set to blink, On (Solid), or Off. <ul style="list-style-type: none">• Blue, blinking: Identifies the appliance in the cabinet or rack.• Blue, solid: Identifies the appliance in the cabinet or rack.• Off: The appliance is not visually identifiable in the cabinet or rack.
4	Status LED	<ul style="list-style-type: none">• Amber, solid: An error has occurred. <p>Note: To view the boot-up and error codes, access the BMC interface.</p> <ul style="list-style-type: none">• Off: No errors are present.
5	PFR	This light is not used by the SGF6112 appliance and remains off.

General boot-up codes


During boot-up or after a hard reset of the appliance, the following occurs:

1. The baseboard management controller (BMC) logs codes for the boot-up sequence, including any errors that occur.
2. The power button lights up.
3. If any errors occur during boot-up, the alarm LED lights up.

To view the boot-up and error codes, [access the BMC interface](#).

SSD indicators

The following figure shows SSD indicators on the SG6112 appliance.



LED	Display	State
1	Drive status/fault	<ul style="list-style-type: none">• Blue (solid): drive is online• Amber (solid): drive failure• Off: slot is empty <p>Note: If a new working SSD is inserted into a working SGF6112 StorageGRID node, the LEDs on the SSD should blink initially, but stop blinking as soon as the system determines that the drive has enough capacity and is functional.</p>
2	Drive active	Blue (blinking): drive is being accessed

SG6000

The SG6000 appliance controllers include indicators that help you determine the status of the appliance controller:

- [Status indicators and buttons on SG6000-CN controller](#)
- [General boot-up codes](#)
- [Boot-up status codes for SG6000 storage controllers](#)

Use this information to help [troubleshoot SG6000 installation](#).

Status indicators and buttons on SG6000-CN controller

The SG6000-CN controller includes indicators that help you determine the status of the controller, including the following indicators and buttons.

The following figure shows status indicators and buttons on the SG6000-CN controller.



Callout	Display	Description
1	Power button	<ul style="list-style-type: none"> Blue: The controller is powered on. Off: The controller is powered off.
2	Reset button	<p><i>No indicator</i></p> <p>Use this button to perform a hard reset of the controller.</p>
3	Identify button	<ul style="list-style-type: none"> Blinking or solid blue: Identifies the controller in the cabinet or rack. Off: The controller is not visually identifiable in the cabinet or rack. <p>This button can be set to Blink, On (Solid), or Off.</p>
4	Alarm LED	<ul style="list-style-type: none"> Amber: An error has occurred. <p>Note: To view the boot-up and error codes, access the BMC interface.</p> <ul style="list-style-type: none"> Off: No errors are present.

General boot-up codes

During boot-up or after a hard reset of the SG6000-CN controller, the following occurs:

1. The baseboard management controller (BMC) logs codes for the boot-up sequence, including any errors that occur.
2. The power button lights up.
3. If any errors occur during boot-up, the alarm LED lights up.

To view the boot-up and error codes, [access the BMC interface](#).

Boot-up status codes for SG6000 storage controllers

Each storage controller has a seven-segment display that provides status codes as the controller powers up. The status codes are the same for both the E2800 controller and the EF570 controller.

For descriptions of these codes, see the E-Series system monitoring information for your storage controller type.

Steps

1. During boot-up, monitor progress by viewing the codes shown on the seven-segment display for each storage controller.

The seven-segment display on each storage controller shows the repeating sequence **OS, Sd, blank** to indicate that the controller is performing start-of-day processing.

2. After the controllers have booted up, confirm that each storage controller shows 99, which is the default ID for an E-Series controller shelf.

Make sure this value is displayed on both storage controllers, as shown in this example E2800 controller.



3. If one or both controllers show other values, see [Troubleshoot hardware installation \(SG6000 or SG5700\)](#) and confirm you completed the installation steps correctly. If you are unable to resolve the problem, contact technical support.

Related information

- [NetApp Support](#)
- [Power on SG6000-CN controller and verify operation](#)

SG5700

The appliance controllers include indicators that help you determine the status of the appliance controller:

- [SG5700 boot-up status codes](#)
- [Status indicators on E5700SG controller](#)
- [General boot-up codes](#)
- [E5700SG controller boot-up codes](#)
- [E5700SG controller error codes](#)

Use this information to help [troubleshoot SG5700 hardware installation](#).

SG5700 boot-up status codes

The seven-segment displays on each controller show status and error codes as the appliance powers up.


The E2800 controller and the E5700SG controller display different statuses and error codes.

To understand what these codes mean, see the following resources:

Controller	Reference
E2800 controller	<p><i>E5700 and E2800 System Monitoring Guide</i></p> <p>Note: The codes listed for the E-Series E5700 controller don't apply to the E5700SG controller in the appliance.</p>
E5700SG controller	"Status indicators on the E5700SG controller"

Steps

1. During boot-up, monitor progress by viewing the codes shown on the seven-segment displays.
 - The seven-segment display on the E2800 controller shows the repeating sequence **OS**, **Sd**, **blank** to indicate that it is performing start-of-day processing.
 - The seven-segment display on the E5700SG controller shows a sequence of codes, ending with **AA** and **FF**.
2. After the controllers have booted up, confirm the seven-segment displays show the following:



Controller	Seven-segment display
E2800 controller	Shows 99, which is the default ID for an E-Series controller shelf.

Controller	Seven-segment display
E5700SG controller	<p>Shows HO, followed by a repeating sequence of two numbers.</p> <div style="border: 1px solid #ccc; padding: 10px; margin-top: 10px;"> HO -- IP address for Admin Network -- IP address for Grid Network HO </div> <p>In the sequence, the first set of numbers is the DHCP-assigned IP address for the controller's management port 1. This address is used to connect the controller to the Admin Network for StorageGRID. The second set of numbers is the DHCP-assigned IP address used to connect the appliance to the Grid Network for StorageGRID.</p> <p>Note: If an IP address could not be assigned using DHCP, 0.0.0.0 is displayed.</p>


3. If the seven-segment displays show other values, see [Troubleshoot hardware installation \(SG6000 or SG5700\)](#) and confirm you completed the installation steps correctly. If you are unable to resolve the problem, contact technical support.

Status indicators on E5700SG controller

The seven-segment display and the LEDs on the E5700SG controller show status and error codes while the appliance powers up and while the hardware is initializing. You can use these displays to determine status and troubleshoot errors.

After the StorageGRID Appliance Installer has started, you should periodically review the status indicators on the E5700SG controller.

The following figure shows status indicators on the E5700SG controller.



Callout	Display	Description
1	Attention LED	Amber: The controller is faulty and requires operator attention, or the installation script was not found. Off: The controller is operating normally.
2	Seven-segment display	Shows a diagnostic code Seven-segment display sequences enable you to understand errors and the operational state of the appliance.
3	Expansion Port Attention LEDs	Amber: These LEDs are always amber (no link established) because the appliance does not use the expansion ports.
4	Host Port Link Status LEDs	Green: The link is up. Off: The link is down.
5	Ethernet Link State LEDs	Green: A link is established. Off: No link is established.
6	Ethernet Activity LEDs	Green: The link between the management port and the device to which it is connected (such as an Ethernet switch) is up. Off: There is no link between the controller and the connected device. Blinking Green: There is Ethernet activity.

General boot-up codes

During boot-up or after a hard reset of the appliance, the following occurs:

1. The seven-segment display on the E5700SG controller shows a general sequence of codes that is not specific to the controller. The general sequence ends with the codes AA and FF.
2. Boot-up codes that are specific to the E5700SG controller appear.

E5700SG controller boot-up codes

During a normal boot-up of the appliance, the seven-segment display on the E5700SG controller shows the following codes in the order listed:

Code	Indicates
HI	The master boot script has started.
PP	The system is checking to see if the FPGA needs to be updated.
HP	The system is checking to see if the 10/25-GbE controller firmware needs to be updated.
RB	The system is rebooting after applying firmware updates.
FP	The hardware subsystem firmware update checks have been completed. Inter-controller communication services are starting.
HE	The system is awaiting connectivity with the E2800 controller and synchronizing with the SANtricity operating system. Note: If this boot procedure does not progress past this stage, check the connections between the two controllers.
HC	The system is checking for existing StorageGRID installation data.
HO	The StorageGRID Appliance Installer is running.
HA	StorageGRID is running.

E5700SG controller error codes

These codes represent error conditions that might be shown on the E5700SG controller as the appliance boots up. Additional two-digit hexadecimal codes are displayed if specific low-level hardware errors occur. If any of these codes persists for more than a second or two, or if you are unable to resolve the error by following one of the prescribed troubleshooting procedures, contact technical support.

Code	Indicates
22	No master boot record found on any boot device.
23	The internal flash disk is not connected.
2A, 2B	Stuck bus, unable to read DIMM SPD data.
40	Invalid DIMMs.
41	Invalid DIMMs.
42	Memory test failed.

Code	Indicates
51	SPD reading failure.
92 to 96	PCI bus initialization.
A0 to A3	SATA drive initialization.
AB	Alternate boot code.
AE	Booting OS.
EA	DDR4 training failed.
E8	No memory installed.
EU	The installation script was not found.
EP	Installation or communication with the E2800 controller has failed.

Related information

- [NetApp Support](#)
- [E5700 and E2800 System Monitoring Guide](#)

SG100 and SG1000

The appliance includes indicators that help you determine the status of the appliance controller and the two SSDs:

- [Appliance indicators and buttons](#)
- [General boot-up codes](#)
- [SSD indicators](#)

Use this information to help [troubleshoot SG100 and SG1000 hardware installation](#).

Appliance indicators and buttons

The following figure shows status indicators and buttons on the SG100 and SG1000.



Callout	Display	State
1	Power button	<ul style="list-style-type: none"> • Blue: the appliance is powered on. • Off: the appliance is powered off.

Callout	Display	State
2	Reset button	Use this button to perform a hard reset of the controller.
3	Identify button	<p>This button can be set to Blink, On (Solid), or Off.</p> <ul style="list-style-type: none"> Blue, blinking: Identifies the appliance in the cabinet or rack. Blue, solid: Identifies the appliance in the cabinet or rack. Off: The appliance is not visually identifiable in the cabinet or rack.
4	Alarm LED	<ul style="list-style-type: none"> Amber, solid: An error has occurred. <p>Note: To view the boot-up and error codes, access the BMC interface.</p> <ul style="list-style-type: none"> Off: No errors are present.

General boot-up codes


During boot-up or after a hard reset of the appliance, the following occurs:

1. The baseboard management controller (BMC) logs codes for the boot-up sequence, including any errors that occur.
2. The power button lights up.
3. If any errors occur during boot-up, the alarm LED lights up.

To view the boot-up and error codes, [access the BMC interface](#).

SSD indicators

The following figure shows SSD indicators on the SG100 and SG1000.



LED	Display	State
1	Drive status/fault	<ul style="list-style-type: none"> Blue (solid): drive is online Amber (blinking): drive failure Off: slot is empty

LED	Display	State
2	Drive active	Blue (blinking): drive is being accessed

Set up hardware

Set up hardware: Overview

After applying power to the appliance, you configure the network connections that will be used by StorageGRID.

Configure required network connections

For all appliances, you perform several tasks to configure required network connections such as:

- Access the Appliance Installer
- Configure network links
- Verify port-level network connections

Additional configuration that might be required

Depending upon which appliance types you are configuring, additional hardware configuration might be required.

SANtricity System Manager

For SG6000 and SG5700, you configure SANtricity System Manager. The SANtricity software is used to monitor the hardware for these appliances.

BMC interface

The following appliances have a BMC interface that must be configured:

- SGF6112
- SG6000
- SG1000
- SG100

Optional configuration

- Storage appliances
 - Configure SANtricity System Manager (SG6000 and SG5700) the software you will use to monitor the hardware
 - Change the RAID mode
- Services appliances
 - Access the BMC interface for the SG100 and SG1000 and the SG6000-CN controller

Configure StorageGRID connections

Access StorageGRID Appliance Installer

You must access the StorageGRID Appliance Installer to verify the installer version and configure the connections between the appliance and the three StorageGRID networks: the Grid Network, the Admin Network (optional), and the Client Network (optional).

Before you begin

- You are using any management client that can connect to the StorageGRID Admin Network, or you have a service laptop.
- The client or service laptop has a [supported web browser](#).
- The services appliance or storage appliance controller is connected to all of the StorageGRID networks you plan to use.
- You know the IP address, gateway, and subnet for the services appliance or storage appliance controller on these networks.
- You have configured the network switches you plan to use.


About this task

To initially access the StorageGRID Appliance Installer, you can use the DHCP-assigned IP address for the Admin Network port on the services appliance or storage appliance controller (assuming it is connected to the Admin Network), or you can connect a service laptop directly to the services appliance or storage appliance controller.

Steps

1. If possible, use the DHCP address for the Admin Network port on the services appliance or storage appliance controller. The Admin Network port is highlighted in the following figure. (Use the IP address on the Grid Network if the Admin Network is not connected.)

SGF6112



SG6000-CN



E5700SG

For the E5700SG, you can do either of the following:

- Look at the seven-segment display on the E5700SG controller. If management port 1 and 10/25-GbE ports 2 and 4 on the E5700SG controller are connected to networks with DHCP servers, the controller attempts to obtain dynamically assigned IP addresses when you power on the enclosure. After the controller has completed the power-on process, its seven-segment display shows **HO**, followed by a repeating sequence of two numbers.

```
HO -- IP address for Admin Network -- IP address for Grid Network  
HO
```

In the sequence:

- The first set of numbers is the DHCP address for the appliance Storage Node on the Admin Network, if it is connected. This IP address is assigned to management port 1 on the E5700SG controller.
- The second set of numbers is the DHCP address for the appliance Storage Node on the Grid Network. This IP address is assigned to 10/25-GbE ports 2 and 4 when you first apply power to the appliance.



If an IP address could not be assigned using DHCP, 0.0.0.0 is displayed.

SG100



SG1000



- Locate the MAC address label on the front of the services appliance or storage appliance, and determine the MAC address for the Admin Network port.

The MAC address label lists the MAC address for the BMC management port.

To determine the MAC address for the Admin Network port, add **2** to the hexadecimal number on the label. For example, if the MAC address on the label ends in **09**, the MAC address for the Admin Port would end in **0B**. If the MAC address on the label ends in **(y)FF**, the MAC address for the Admin Port would end in **(y+1)01**. You can easily make this calculation by opening Calculator in Windows, setting it to Programmer mode, selecting Hex, typing the MAC address, then typing **+ 2 =**.

- b. Provide the MAC address to your network administrator, so they can look up the DHCP address for the appliance on the Admin Network.
- c. From the client, enter this URL for the StorageGRID Appliance Installer:

https://Appliance_IP:8443

For *Appliance_IP*, use the DHCP address (use the IP address for the Admin Network if you have it).

- d. If you are prompted with a security alert, view and install the certificate using the browser's installation wizard.

The alert will not appear the next time you access this URL.

The StorageGRID Appliance Installer Home page appears. The information and messages shown when you first access this page depend on how your appliance is currently connected to StorageGRID networks. Error messages might appear that will be resolved in later steps.

Home

 The installation is ready to be started. Review the settings below, and then click Start Installation.

This Node

Node type

Storage

Node name

MM-2-108-SGA-lab25

Cancel

Save

Primary Admin Node connection

Enable Admin Node discovery

Primary Admin Node IP

172.16.1.178

Connection state

Connection to 172.16.1.178 ready

Cancel

Save

Installation

Current state


Ready to start installation of MM-2-108-SGA-lab25 into grid with Admin Node 172.16.1.178 running StorageGRID 11.2.0, using StorageGRID software downloaded from the Admin Node.

Start Installation

2. If you can't obtain an IP address using DHCP, you can use a link-local connection.

SGF6112

Connect a service laptop directly to the rightmost RJ-45 port on the appliance, using an Ethernet cable.




SG6000-CN

Connect a service laptop directly to the rightmost RJ-45 port on the SG6000-CN controller, using an Ethernet cable.



E5700SG

Connect the service laptop to management port 2 on the E5700SG controller, using an Ethernet cable.



SG100

Connect a service laptop directly to the rightmost RJ-45 port on the services appliance, using an Ethernet cable.



SG1000

Connect a service laptop directly to the rightmost RJ-45 port on the services appliance, using an Ethernet cable.



- a. Open a web browser on the service laptop.
- b. Enter this URL for the StorageGRID Appliance Installer:
<https://169.254.0.1:8443>

The StorageGRID Appliance Installer Home page appears. The information and messages shown when you first access this page depend on how your appliance is currently connected to StorageGRID networks. Error messages might appear that will be resolved in later steps.



If you can't access the Home page over a link-local connection, configure the service laptop IP address as 169.254.0.2, and try again.

After you finish

After accessing the StorageGRID Appliance Installer:

- Verify that the StorageGRID Appliance Installer version on the appliance matches the software version installed on your StorageGRID system. Upgrade StorageGRID Appliance Installer, if necessary.

[Verify and upgrade StorageGRID Appliance Installer version](#)

- Review any messages displayed on the StorageGRID Appliance Installer Home page and configure the link configuration and the IP configuration, as required.

NetApp® StorageGRID® Appliance Installer

Home	Configure Networking	Configure Hardware	Monitor Installation	Advanced
Home				
This Node				
Node type	Gateway		<input type="button" value="Save"/>	
Node name	xlr8r-10		<input type="button" value="Save"/>	
Primary Admin Node connection				
Enable Admin Node discovery	<input type="checkbox"/>		<input type="button" value="Save"/>	
Primary Admin Node IP	192.168.7.44		<input type="button" value="Save"/>	
Connection state	Connection to 192.168.7.44 ready			
Installation				
Current state	Ready to start installation of xlr8r-10 into grid with Admin Node 192.168.7.44 running StorageGRID 11.8.0, using StorageGRID software downloaded from the Admin Node.			
<input type="button" value="Start Installation"/>				

Verify and upgrade StorageGRID Appliance Installer version

The StorageGRID Appliance Installer version on the appliance must match the software version installed on your StorageGRID system to ensure that all StorageGRID features are supported.

Before you begin

You have accessed the StorageGRID Appliance Installer.

About this task

StorageGRID appliances come from the factory preinstalled with the StorageGRID Appliance Installer. If you are adding an appliance to a recently upgraded StorageGRID system, you might need to manually upgrade the StorageGRID Appliance Installer before installing the appliance as a new node.

The StorageGRID Appliance Installer automatically upgrades when you upgrade to a new StorageGRID version. You don't need to upgrade the StorageGRID Appliance Installer on installed appliance nodes. This procedure is only required when you are installing an appliance that contains an earlier version of the StorageGRID Appliance Installer.

Steps

1. From the StorageGRID Appliance Installer, select **Advanced > Upgrade Firmware**.
2. Compare the Current Firmware version to the software version installed on your StorageGRID system. (From the top of the Grid Manager, select the help icon and select **About**.)

The second digit in the two versions should match. For example, if your StorageGRID system is running version 11.6.x.y, the StorageGRID Appliance Installer version should be 3.6.z.

3. If the appliance has a down-level version of the StorageGRID Appliance Installer, go to [NetApp Downloads: StorageGRID Appliance](#).

Sign in with the username and password for your NetApp account.

4. Download the appropriate version of the **Support file for StorageGRID Appliances** and the corresponding checksum file.

The Support file for StorageGRID appliances is a .zip archive that contains the current and previous firmware versions for all StorageGRID appliance models.

After downloading the Support file for StorageGRID appliances, extract the .zip archive and see the README file for important information about installing the StorageGRID Appliance Installer.

5. Follow the instructions on the Upgrade Firmware page of your StorageGRID Appliance Installer to perform these steps:
 - a. Upload the appropriate support file (firmware image) for your controller type. Some firmware versions also require uploading a checksum file. If you are prompted for a checksum file, it can also be found in the Support file for StorageGRID Appliances.
 - b. Upgrade the inactive partition.
 - c. Reboot and swap partitions.
 - d. Upload the appropriate support file (firmware image) again for your controller type. Some firmware versions also require uploading a checksum file. If you are prompted for a checksum file, it can also be found in the Support file for StorageGRID Appliances.

- e. Upgrade the second (inactive) partition.

Related information

[Accessing StorageGRID Appliance Installer](#)

Configure network links

You can configure network links for the ports used to connect the appliance to the Grid Network, the Client Network, and the Admin Network. You can set the link speed as well as the port and network bond modes.



If you are using ConfigBuilder to generate a JSON file, you can configure the network links automatically. See [Automate appliance installation and configuration](#).

Before you begin

- You have [obtained the additional equipment](#) required for your cable type and link speed.
- You have installed the correct transceivers in the ports, based on the link speed you plan to use.
- You have connected the network ports to switches that support your chosen speed.

If you plan to use Aggregate port bond mode, LACP network bond mode, or VLAN tagging:

- You have connected the network ports on the appliance to switches that can support VLAN and LACP.
- If multiple switches are participating in the LACP bond, the switches support multi-chassis link aggregation groups (MLAG), or equivalent.
- You understand how to configure the switches to use VLAN, LACP, and MLAG or equivalent.
- You know the unique VLAN tag to use for each network. This VLAN tag will be added to each network packet to ensure that network traffic is routed to the correct network.

About this task

You only need to configure the settings on the Link Configuration page if you want to use a non-default setting.



The LACP transmit hash policy is layer2+3.


The figures and tables summarize the options for the port bond mode and network bond mode for each appliance. See the following for more information:

- [Port bond modes \(SGF6112\)](#)
- [Port bond modes \(SG6000-CN\)](#)
- [Port bond modes \(E5700SG\)](#)
- [Port bond modes \(SG1000 and SG100\)](#)

SGF6112

Fixed port bond mode (default)

The figure shows how the four network ports are bonded in fixed port bond mode (default configuration).




Callout	Which ports are bonded
C	Ports 1 and 3 are bonded together for the Client Network, if this network is used.
G	Ports 2 and 4 are bonded together for the Grid Network.

The table summarizes the options for configuring the network ports. You only need to configure the settings on the Link Configuration page if you want to use a non-default setting.

Network bond mode	Client Network disabled (default)	Client Network enabled
Active-Backup (default)	<ul style="list-style-type: none">Ports 2 and 4 use an active-backup bond for the Grid Network.Ports 1 and 3 aren't used.A VLAN tag is optional.	<ul style="list-style-type: none">Ports 2 and 4 use an active-backup bond for the Grid Network.Ports 1 and 3 use an active-backup bond for the Client Network.VLAN tags can be specified for both networks for the convenience of the network administrator.
LACP (802.3ad)	<ul style="list-style-type: none">Ports 2 and 4 use an LACP bond for the Grid Network.Ports 1 and 3 aren't used.A VLAN tag is optional.	<ul style="list-style-type: none">Ports 2 and 4 use an LACP bond for the Grid Network.Ports 1 and 3 use an LACP bond for the Client Network.VLAN tags can be specified for both networks for the convenience of the network administrator.

Aggregate port bond mode

The figure shows how the four network ports are bonded in aggregate port bond mode.




Callout	Which ports are bonded
1	All four ports are grouped in a single LACP bond, allowing all ports to be used for Grid Network and Client Network traffic.

The table summarizes the options for configuring the network ports. You only need to configure the settings on the Link Configuration page if you want to use a non-default setting.

Network bond mode	Client Network disabled (default)	Client Network enabled
LACP (802.3ad) only	<ul style="list-style-type: none"> Ports 1-4 use a single LACP bond for the Grid Network. A single VLAN tag identifies Grid Network packets. 	<ul style="list-style-type: none"> Ports 1-4 use a single LACP bond for the Grid Network and the Client Network. Two VLAN tags allow Grid Network packets to be segregated from Client Network packets.

Active-Backup network bond mode for management ports


This figure shows how the two 1-GbE management ports on the SGF6112 are bonded in Active-Backup network bond mode for the Admin Network.



SG6000

Fixed port bond mode (default)

This figure shows how the four network ports are bonded in fixed port bond mode (default configuration)




Callout	Which ports are bonded
C	Ports 1 and 3 are bonded together for the Client Network, if this network is used.
G	Ports 2 and 4 are bonded together for the Grid Network.

The table summarizes the options for configuring the network ports. You only need to configure the settings on the Link Configuration page if you want to use a non-default setting.

Network bond mode	Client Network disabled (default)	Client Network enabled
Active-Backup (default)	<ul style="list-style-type: none"> Ports 2 and 4 use an active-backup bond for the Grid Network. Ports 1 and 3 aren't used. A VLAN tag is optional. 	<ul style="list-style-type: none"> Ports 2 and 4 use an active-backup bond for the Grid Network. Ports 1 and 3 use an active-backup bond for the Client Network. VLAN tags can be specified for both networks for the convenience of the network administrator.
LACP (802.3ad)	<ul style="list-style-type: none"> Ports 2 and 4 use an LACP bond for the Grid Network. Ports 1 and 3 aren't used. A VLAN tag is optional. 	<ul style="list-style-type: none"> Ports 2 and 4 use an LACP bond for the Grid Network. Ports 1 and 3 use an LACP bond for the Client Network. VLAN tags can be specified for both networks for the convenience of the network administrator.

Aggregate port bond mode

This figure shows how the four network ports are bonded in aggregate port bond mode.



Callout	Which ports are bonded
1	All four ports are grouped in a single LACP bond, allowing all ports to be used for Grid Network and Client Network traffic.

The table summarizes the options for configuring the network ports. You only need to configure the settings on the Link Configuration page if you want to use a non-default setting.

Network bond mode	Client Network disabled (default)	Client Network enabled
LACP (802.3ad) only	<ul style="list-style-type: none"> Ports 1-4 use a single LACP bond for the Grid Network. A single VLAN tag identifies Grid Network packets. 	<ul style="list-style-type: none"> Ports 1-4 use a single LACP bond for the Grid Network and the Client Network. Two VLAN tags allow Grid Network packets to be segregated from Client Network packets.

Active-Backup network bond mode for management ports


This figure shows how the two 1-GbE management ports on the SG6000-CN controller are bonded in Active-Backup network bond mode for the Admin Network.



SG5700

Fixed port bond mode (default)

This figure shows how the four 10/25-GbE ports are bonded in Fixed port bond mode (default configuration).




Callout	Which ports are bonded
C	Ports 1 and 3 are bonded together for the Client Network, if this network is used.
G	Ports 2 and 4 are bonded together for the Grid Network.

The table summarizes the options for configuring the four 10/25-GbE ports. You only need to configure the settings on the Link Configuration page if you want to use a non-default setting.

Network bond mode	Client Network disabled (default)	Client Network enabled
Active-Backup (default)	<ul style="list-style-type: none"> Ports 2 and 4 use an active-backup bond for the Grid Network. Ports 1 and 3 aren't used. A VLAN tag is optional. 	<ul style="list-style-type: none"> Ports 2 and 4 use an active-backup bond for the Grid Network. Ports 1 and 3 use an active-backup bond for the Client Network. VLAN tags can be specified for both networks for the convenience of the network administrator.
LACP (802.3ad)	<ul style="list-style-type: none"> Ports 2 and 4 use an LACP bond for the Grid Network. Ports 1 and 3 aren't used. A VLAN tag is optional. 	<ul style="list-style-type: none"> Ports 2 and 4 use an LACP bond for the Grid Network. Ports 1 and 3 use an LACP bond for the Client Network. VLAN tags can be specified for both networks for the convenience of the network administrator.

Aggregate port bond mode

This figure shows how the four 10/25-GbE ports are bonded in Aggregate port bond mode.



Callout	Which ports are bonded
1	All four ports are grouped in a single LACP bond, allowing all ports to be used for Grid Network and Client Network traffic.

The table summarizes the options for configuring the four 10/25-GbE ports. You only need to configure the settings on the Link Configuration page if you want to use a non-default setting.

Network bond mode	Client Network disabled (default)	Client Network enabled
LACP (802.3ad) only	<ul style="list-style-type: none"> Ports 1-4 use a single LACP bond for the Grid Network. A single VLAN tag identifies Grid Network packets. 	<ul style="list-style-type: none"> Ports 1-4 use a single LACP bond for the Grid Network and the Client Network. Two VLAN tags allow Grid Network packets to be segregated from Client Network packets.

Active-Backup network bond mode for management ports

This figure shows how the two 1-GbE management ports on the E5700SG controller are bonded in Active-Backup network bond mode for the Admin Network.




SG100 and SG1000


Fixed port bond mode (default)

The figures show how the four network ports on the SG1000 or SG100 are bonded in fixed port bond mode (default configuration).

SG1000:



SG100:



Callout	Which ports are bonded
C	Ports 1 and 3 are bonded together for the Client Network, if this network is used.
G	Ports 2 and 4 are bonded together for the Grid Network.


The table summarizes the options for configuring the four network ports. You only need to configure the settings on the Link Configuration page if you want to use a non-default setting.

Network bond mode	Client Network disabled (default)	Client Network enabled
Active-Backup (default)	<ul style="list-style-type: none"> Ports 2 and 4 use an active-backup bond for the Grid Network. Ports 1 and 3 aren't used. A VLAN tag is optional. 	<ul style="list-style-type: none"> Ports 2 and 4 use an active-backup bond for the Grid Network. Ports 1 and 3 use an active-backup bond for the Client Network. VLAN tags can be specified for both networks for the convenience of the network administrator.
LACP (802.3ad)	<ul style="list-style-type: none"> Ports 2 and 4 use an LACP bond for the Grid Network. Ports 1 and 3 aren't used. A VLAN tag is optional. 	<ul style="list-style-type: none"> Ports 2 and 4 use an LACP bond for the Grid Network. Ports 1 and 3 use an LACP bond for the Client Network. VLAN tags can be specified for both networks for the convenience of the network administrator.


Aggregate port bond mode

These figures show how the four network ports are bonded in aggregate port bond mode.

SG1000:



SG100:



Callout	Which ports are bonded
1	All four ports are grouped in a single LACP bond, allowing all ports to be used for Grid Network and Client Network traffic.

The table summarizes the options for configuring the four network ports. You only need to configure the settings on the Link Configuration page if you want to use a non-default setting.

Network bond mode	Client Network disabled (default)	Client Network enabled
LACP (802.3ad) only	<ul style="list-style-type: none"> Ports 1-4 use a single LACP bond for the Grid Network. A single VLAN tag identifies Grid Network packets. 	<ul style="list-style-type: none"> Ports 1-4 use a single LACP bond for the Grid Network and the Client Network. Two VLAN tags allow Grid Network packets to be segregated from Client Network packets.


Active-Backup network bond mode for management ports

These figures show how the two 1-GbE management ports on the appliances are bonded in Active-Backup network bond mode for the Admin Network.

SG1000:



SG100:



Steps

- From the menu bar of the StorageGRID Appliance Installer, click **Configure Networking > Link Configuration**.

The Network Link Configuration page displays a diagram of your appliance with the network and management ports numbered.

The Link Status table lists the link state, link speed, and other statistics of the numbered ports.

The first time you access this page:

- Link Speed is set to **Auto**.
- Port bond mode is set to **Fixed**.
- Network bond mode is set to **Active-Backup** for the Grid Network.
- The **Admin Network** is enabled, and the network bond mode is set to **Independent**.

- The **Client Network** is disabled.

2. Select the link speed for the network ports from the **Link speed** drop-down list.

The network switches you are using for the Grid Network and the Client Network must also support and be configured for this speed. You must use the appropriate adapters or transceivers for the configured link speed. Use Auto link speed when possible because this option negotiates both link speed and Forward Error Correction (FEC) mode with the link partner.

If you plan to use the 25-GbE link speed for the SG6000 or SG5700 network ports:

- Use SFP28 transceivers and SFP28 TwinAx cables or optical cables.
- For the SG6000, select **Auto** from the **Link speed** drop-down list.
- For the SG5700, select **25GbE** from the **Link speed** drop-down list.

3. Enable or disable the StorageGRID networks you plan to use.

The Grid Network is required. You can't disable this network.

- a. If the appliance is not connected to the Admin Network, clear the **Enable network** checkbox for the Admin Network.
- b. If the appliance is connected to the Client Network, select the **Enable network** checkbox for the Client Network.

The Client Network settings for the data NIC ports are now shown.

4. Refer to the table, and configure the port bond mode and the network bond mode.

This example shows:

- **Aggregate** and **LACP** selected for the Grid and the Client Networks. You must specify a unique VLAN tag for each network. You can select values between 0 and 4095.
- **Active-Backup** selected for the Admin Network.

Link Settings

Link speed

Port bond mode

Fixed Aggregate

Choose Fixed port bond mode if you want to use ports 2 and 4 for the Grid Network and ports 1 and 3 for the Client Network (if enabled). Choose Aggregate port bond mode if you want all connected ports to share a single LACP bond for both the Grid and Client Networks.

Grid Network

Enable network

Network bond mode

Active-Backup LACP (802.3ad)

If the port bond mode is Aggregate, all bonds must be in LACP (802.3ad) mode.

Enable VLAN (802.1q) tagging

VLAN (802.1q) tag

MAC Addresses

50:6b:4b:42:d7:00 50:6b:4b:42:d7:01 50:6b:4b:42:d7:24 50:6b:4b:42:d7:25

If you are using DHCP, it is recommended that you configure a permanent DHCP reservation. Use all of these MAC addresses in the reservation to assign one IP address to this network interface.

Admin Network

Enable network

Network bond mode

Independent Active-Backup

Connect the Admin Network to ports 5 and 6. If necessary, you can make a temporary direct Ethernet connection by disconnecting ports 5 and 6, then connecting to port 6 and using link-local IP address 169.254.0.1 for access.

MAC Addresses

d8:c4:97:2a:e4:95

If you are using DHCP, it is recommended that you configure a permanent DHCP reservation. Use all of these MAC addresses in the reservation to assign one IP address to this network interface.

Client Network

Enable network

Network bond mode

Active-Backup LACP (802.3ad)

If the port bond mode is Aggregate, all bonds must be in LACP (802.3ad) mode.

Enable VLAN (802.1q) tagging

VLAN (802.1q) tag

MAC Addresses

50:6b:4b:42:d7:00 50:6b:4b:42:d7:01 50:6b:4b:42:d7:24 50:6b:4b:42:d7:25

If you are using DHCP, it is recommended that you configure a permanent DHCP reservation. Use all of these MAC addresses in the reservation to assign one IP address to this network interface.

5. When you are satisfied with your selections, click **Save**.



You might lose your connection if you made changes to the network or link you are connected through. If you aren't reconnected within 1 minute, re-enter the URL for the StorageGRID Appliance Installer using one of the other IP addresses assigned to the appliance:

https://appliance_IP:8443

Configure StorageGRID IP addresses

You use the StorageGRID Appliance Installer to configure the IP addresses and routing information used for the services appliance or appliance Storage Node on the StorageGRID Grid, Admin, and Client Networks.

If you are using ConfigBuilder to generate a JSON file, you can configure IP addresses automatically. See [Automate appliance installation and configuration](#).

About this task

You must either assign a static IP for the appliance on each connected network or assign a permanent lease for the address on the DHCP server.

To change the link configuration, see the following instructions:

- [Change link configuration of the SGF6112 appliance](#)
- [Change link configuration of the SG6000-CN controller](#)
- [Change link configuration of the E5700SG controller](#)
- [Change link configuration of the SG100 or SG1000 services appliance](#)

Steps

1. In the StorageGRID Appliance Installer, select **Configure Networking > IP Configuration**.

The IP Configuration page appears.

2. To configure the Grid Network, select either **Static** or **DHCP** in the **Grid Network** section of the page.

3. If you selected **Static**, follow these steps to configure the Grid Network:

- a. Enter the static IPv4 address, using CIDR notation.
- b. Enter the gateway.

If your network does not have a gateway, re-enter the same static IPv4 address.

- c. If you want to use jumbo frames, change the MTU field to a value suitable for jumbo frames, such as 9000. Otherwise, keep the default value of 1500.



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.



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.

d. Click **Save**.

When you change the IP address, the gateway and list of subnets might also change.

If you lose your connection to the StorageGRID Appliance Installer, re-enter the URL using the new static IP address you just assigned. For example,

https://appliance_IP:8443

e. Confirm that the list of Grid Network subnets is correct.

If you have grid subnets, the Grid Network gateway is required. All grid subnets specified must be reachable through this gateway. These Grid Network subnets must also be defined in the Grid Network Subnet List on the primary Admin Node when you start StorageGRID installation.



The default route is not listed. If the Client Network is not enabled, the default route will use the Grid Network gateway.

- To add a subnet, click the insert icon  to the right of the last entry.
- To remove an unused subnet, click the delete icon .

f. Click **Save**.

4. If you selected **DHCP**, follow these steps to configure the Grid Network:

a. After you select the **DHCP** radio button, click **Save**.

The **IPv4 Address**, **Gateway**, and **Subnets** fields are automatically populated. If the DHCP server is set up to assign an MTU value, the **MTU** field is populated with that value, and the field becomes read-only.

Your web browser is automatically redirected to the new IP address for the StorageGRID Appliance Installer.

b. Confirm that the list of Grid Network subnets is correct.

If you have grid subnets, the Grid Network gateway is required. All grid subnets specified must be reachable through this gateway. These Grid Network subnets must also be defined in the Grid Network Subnet List on the primary Admin Node when you start StorageGRID installation.



The default route is not listed. If the Client Network is not enabled, the default route will use the Grid Network gateway.

- To add a subnet, click the insert icon  to the right of the last entry.
- To remove an unused subnet, click the delete icon .

c. If you want to use jumbo frames, change the MTU field to a value suitable for jumbo frames, such as 9000. Otherwise, keep the default value of 1500.



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.



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.

d. Click **Save**.

5. To configure the Admin Network, select either **Static** or **DHCP** in the **Admin Network** section of the page.



To configure the Admin Network, you enable the Admin Network on the Link Configuration page.

Admin Network

The Admin Network is a closed network used for system administration and maintenance. The Admin Network is typically a private network and does not need to be routable between sites.

IP Static DHCP

Assignment

IPv4 Address (CIDR)

Gateway

Subnets (CIDR) +

MTU ▼

Cancel Save

6. If you selected **Static**, follow these steps to configure the Admin Network:

a. Enter the static IPv4 address, using CIDR notation, for Management Port 1 on the appliance.

Management Port 1 is the left of the two 1-GbE RJ45 ports on the right end of the appliance.

b. Enter the gateway.

If your network does not have a gateway, re-enter the same static IPv4 address.

c. If you want to use jumbo frames, change the MTU field to a value suitable for jumbo frames, such as 9000. Otherwise, keep the default value of 1500.



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.

d. Click **Save**.

When you change the IP address, the gateway and list of subnets might also change.

If you lose your connection to the StorageGRID Appliance Installer, re-enter the URL using the new static IP address you just assigned. For example,

https://appliance:8443

e. Confirm that the list of Admin Network subnets is correct.

You must verify that all subnets can be reached using the gateway you provided.



The default route can't be made to use the Admin Network gateway.

- To add a subnet, click the insert icon  to the right of the last entry.
- To remove an unused subnet, click the delete icon .

f. Click **Save**.

7. If you selected **DHCP**, follow these steps to configure the Admin Network:

a. After you select the **DHCP** radio button, click **Save**.

The **IPv4 Address**, **Gateway**, and **Subnets** fields are automatically populated. If the DHCP server is set up to assign an MTU value, the **MTU** field is populated with that value, and the field becomes read-only.

Your web browser is automatically redirected to the new IP address for the StorageGRID Appliance Installer.

b. Confirm that the list of Admin Network subnets is correct.

You must verify that all subnets can be reached using the gateway you provided.



The default route can't be made to use the Admin Network gateway.

- To add a subnet, click the insert icon  to the right of the last entry.
- To remove an unused subnet, click the delete icon .

c. If you want to use jumbo frames, change the MTU field to a value suitable for jumbo frames, such as 9000. Otherwise, keep the default value of 1500.



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.

d. Click **Save**.

8. To configure the Client Network, select either **Static** or **DHCP** in the **Client Network** section of the page.



To configure the Client Network, you enable the Client Network on the Link Configuration page.

Client Network

The Client Network is an open network used to provide access to client applications, including S3 and Swift. The Client Network enables grid nodes to communicate with any subnet reachable through the Client Network gateway. The Client Network does not become operational until you complete the StorageGRID configuration steps.

IP Assignment

Static DHCP

IPv4 Address (CIDR)

Gateway

MTU

9. If you selected **Static**, follow these steps to configure the Client Network:

- Enter the static IPv4 address, using CIDR notation.
- Click **Save**.
- Confirm that the IP address for the Client Network gateway is correct.



If the Client Network is enabled, the default route is displayed. The default route uses the Client Network gateway and can't be moved to another interface while the Client Network is enabled.

- If you want to use jumbo frames, change the MTU field to a value suitable for jumbo frames, such as 9000. Otherwise, keep the default value of 1500.



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.

10. If you selected **DHCP**, follow these steps to configure the Client Network:

- After you select the **DHCP** radio button, click **Save**.

The **IPv4 Address** and **Gateway** fields are automatically populated. If the DHCP server is set up to assign an MTU value, the **MTU** field is populated with that value, and the field becomes read-only.

Your web browser is automatically redirected to the new IP address for the StorageGRID Appliance Installer.

b. Confirm that the gateway is correct.



If the Client Network is enabled, the default route is displayed. The default route uses the Client Network gateway and can't be moved to another interface while the Client Network is enabled.

c. If you want to use jumbo frames, change the MTU field to a value suitable for jumbo frames, such as 9000. Otherwise, keep the default value of 1500.



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.

Verify network connections

You should confirm you can access the StorageGRID networks you are using from the appliance. To validate routing through network gateways, you should test connectivity between the StorageGRID Appliance Installer and IP addresses on different subnets. You can also verify the MTU setting.

Steps

1. From the menu bar of the StorageGRID Appliance Installer, click **Configure Networking > Ping and MTU Test**.

The Ping and MTU Test page appears.

Ping and MTU Test

Use a ping request to check the appliance's connectivity to a remote host. Select the network you want to check connectivity through, and enter the IP address of the host you want to reach. To verify the MTU setting for the entire path through the network to the destination, select Test MTU.

Ping and MTU Test

Network	<input type="text" value="Grid"/>
Destination IPv4 Address or FQDN	<input type="text"/>
Test MTU	<input type="checkbox"/>
Test Connectivity	

2. From the **Network** drop-down box, select the network you want to test: Grid, Admin, or Client.
3. Enter the IPv4 address or fully qualified domain name (FQDN) for a host on that network.

For example, you might want to ping the gateway on the network or the primary Admin Node.

4. Optionally, select the **Test MTU** checkbox to verify the MTU setting for the entire path through the network to the destination.

For example, you can test the path between the appliance node and a node at a different site.

5. Click **Test Connectivity**.

If the network connection is valid, the "Ping test passed" message appears, with the ping command output listed.

Ping and MTU Test

Use a ping request to check the appliance's connectivity to a remote host. Select the network you want to check connectivity through, and enter the IP address of the host you want to reach. To verify the MTU setting for the entire path through the network to the destination, select Test MTU.

Ping and MTU Test

Network	Grid
Destination IPv4 Address or FQDN	10.96.104.223
Test MTU	<input checked="" type="checkbox"/>
Test Connectivity	
Ping test passed	
Ping command output	
<pre>PING 10.96.104.223 (10.96.104.223) 1472(1500) bytes of data. 1480 bytes from 10.96.104.223: icmp_seq=1 ttl=64 time=0.318 ms --- 10.96.104.223 ping statistics --- 1 packets transmitted, 1 received, 0% packet loss, time 0ms rtt min/avg/max/mdev = 0.318/0.318/0.318/0.000 ms Found MTU 1500 for 10.96.104.223 via br0</pre>	

Related information

- [Configure network links](#)
- [Change MTU setting](#)

Verify port-level network connections

To ensure that access between the StorageGRID Appliance Installer and other nodes is not obstructed by firewalls, confirm that the StorageGRID Appliance Installer can connect to a specific TCP port or set of ports at the specified IP address or range of addresses.

About this task

Using the list of ports provided in the StorageGRID Appliance Installer, you can test the connectivity between the appliance and the other nodes in your Grid Network.

Additionally, you can test connectivity on the Admin and Client Networks and on UDP ports, such as those used for external NFS or DNS servers. For a list of these ports, see the [network port reference](#).



The Grid Network ports listed in the port connectivity table are valid only for StorageGRID version 11.7.0. To verify which ports are correct for each node type, you should always consult the networking guidelines for your version of StorageGRID.

Steps

1. From the StorageGRID Appliance Installer, click **Configure Networking > Port Connectivity Test (nmap)**.

The Port Connectivity Test page appears.

The port connectivity table lists node types that require TCP connectivity on the Grid Network. For each node type, the table lists the Grid Network ports that should be accessible to your appliance.

You can test the connectivity between the appliance ports listed in the table and the other nodes in your Grid Network.

2. From the **Network** drop-down, select the network you want to test: **Grid, Admin, or Client**.
3. Specify a range of IPv4 addresses for the hosts on that network.

For example, you might want to probe the gateway on the network or the primary Admin Node.

Specify a range using a hyphen, as shown in the example.

4. Enter a TCP port number, a list of ports separated by commas, or a range of ports.

Port Connectivity Test

Network	<input type="text" value="Grid"/>
IPv4 Address Ranges	<input type="text" value="10.224.6.160-161"/>
Port Ranges	<input type="text" value="22,2022"/>
Protocol	<input checked="" type="radio"/> TCP <input type="radio"/> UDP
Test Connectivity	

5. Click **Test Connectivity**.

- If the selected port-level network connections are valid, the “Port connectivity test passed” message appears in a green banner. The nmap command output is listed below the banner.

Port connectivity test passed

Nmap command output. Note: Unreachable hosts will not appear in the output.

```
# Nmap 7.70 scan initiated Fri Nov 13 18:32:03 2020 as: /usr/bin/nmap -n -oN - -e br0 -p 22,2022 10.224.6.160-161
Nmap scan report for 10.224.6.160
Host is up (0.00072s latency).

PORT      STATE SERVICE
22/tcp    open  ssh
2022/tcp  open  down

Nmap scan report for 10.224.6.161
Host is up (0.00060s latency).

PORT      STATE SERVICE
22/tcp    open  ssh
2022/tcp  open  down

# Nmap done at Fri Nov 13 18:32:04 2020 -- 2 IP addresses (2 hosts up) scanned in 0.55 seconds
```

- If a port-level network connection is made to the remote host, but the host is not listening on one or more of the selected ports, the “Port connectivity test failed” message appears in a yellow banner. The nmap command output is listed below the banner.

Any remote port the host is not listening to has a state of “closed.” For example, you might see this yellow banner when the node you are trying to connect to is in a pre-installed state and the StorageGRID NMS service is not yet running on that node.

⚠ Port connectivity test failed

Connection not established. Services might not be listening on target ports.

Nmap command output. Note: Unreachable hosts will not appear in the output.

```
# Nmap 7.70 scan initiated Sat May 16 17:07:02 2020 as: /usr/bin/nmap -n -oN - -e br0 -p 22,80,443,1504,1505,1506,1508,7443,9999
Nmap scan report for 172.16.4.71
Host is up (0.00020s latency).

PORT      STATE SERVICE
22/tcp    open  ssh
80/tcp    open  http
443/tcp   open  https
1504/tcp  closed evb-elm
1505/tcp  open  funkproxy
1506/tcp  open  utcd
1508/tcp  open  diamond
7443/tcp  open  oracleas-https
9999/tcp  open  abyss
MAC Address: 00:50:56:87:39:AE (VMware)

# Nmap done at Sat May 16 17:07:03 2020 -- 1 IP address (1 host up) scanned in 0.59 seconds
```

- If a port-level network connection can’t be made for one or more selected ports, the “Port connectivity test failed” message appears in a red banner. The nmap command output is listed below the banner.

The red banner indicates that a TCP connection attempt to a port on the remote host was made, but nothing was returned to the sender. When no response is returned, the port has a state of “filtered” and is likely blocked by a firewall.



Ports with “closed” are also listed.

 Port connectivity test failed
Connection failed to one or more ports.

Nmap command output. Note: Unreachable hosts will not appear in the output.

```
# Nmap 7.70 scan initiated Sat May 16 17:11:01 2020 as: /usr/bin/nmap -n -oN - -e br0 -p 22,79,80,443,1504,1505,1506,1508,7443,9999 172.16.4.71
Nmap scan report for 172.16.4.71
Host is up (0.00029s latency).

PORT      STATE    SERVICE
22/tcp    open     ssh
79/tcp    filtered finger
80/tcp    open     http
443/tcp   open     https
1504/tcp  closed   evb-eim
1505/tcp  open     funkproxy
1506/tcp  open     utcd
1508/tcp  open     diagmond
7443/tcp  open     oracleas-https
9999/tcp  open     abyss
MAC Address: 00:50:56:87:39:AE (VMware)

# Nmap done at Sat May 16 17:11:02 2020 -- 1 IP address (1 host up) scanned in 1.60 seconds
```

Configure SANtricity System Manager (SG6000 and SG5700)

You can use SANtricity System Manager to monitor the status of the storage controllers, storage disks, and other hardware components in the storage controller shelf. You can also configure a proxy for E-Series AutoSupport that enables you to send AutoSupport messages from the appliance without the use of the management port.

Set up and access SANtricity System Manager

You might need to access SANtricity System Manager on the storage controller to monitor the hardware in the storage controller shelf or to configure E-Series AutoSupport.

Before you begin

- You are using a [supported web browser](#).
- To access SANtricity System Manager through Grid Manager, you have installed StorageGRID, and you have the Storage appliance administrator permission or Root access permission.
- To access SANtricity System Manager using the StorageGRID Appliance Installer, you have the SANtricity System Manager administrator username and password.
- To access SANtricity System Manager directly using a web browser, you have the SANtricity System Manager administrator username and password.



You must have SANtricity firmware 8.70 or higher to access SANtricity System Manager using the Grid Manager or the StorageGRID Appliance Installer. You can check your firmware version by using the StorageGRID Appliance Installer and selecting **Help > About**.



Accessing SANtricity System Manager from the Grid Manager or from the Appliance Installer is generally meant only for monitoring your hardware and configuring E-Series AutoSupport. Many features and operations within SANtricity System Manager such as upgrading firmware don't apply to monitoring your StorageGRID appliance. To avoid issues, always follow the hardware installation and maintenance instructions for your appliance.

About this task

There are three ways to access SANtricity System Manager, depending upon what stage of the installation and

configuration process you are in:

- If the appliance has not yet been deployed as a node in your StorageGRID system, you should use the Advanced tab in the StorageGRID Appliance Installer.



Once the node is deployed, you can no longer use the StorageGRID Appliance Installer to access SANtricity System Manager.

- If the appliance has been deployed as a node in your StorageGRID system, use the SANtricity System Manager tab on the Nodes page in Grid Manager.
- If you can't use the StorageGRID Appliance Installer or Grid Manager, you can access SANtricity System Manager directly using a web browser connected to the management port.

This procedure includes steps for your initial access to SANtricity System Manager. If you have already set up SANtricity System Manager, go to the [configure hardware alerts step](#).



Using either the Grid Manager or the StorageGRID Appliance Installer enables you to access SANtricity System Manager without having to configure or connect the management port of the appliance.

You use SANtricity System Manager to monitor the following:

- Performance data such as storage array level performance, I/O latency, CPU utilization, and throughput
- Hardware component status
- Support functions including viewing diagnostic data

You can use SANtricity System Manager to configure the following settings:

- Email alerts, SNMP alerts, or syslog alerts for the components in the storage controller shelf
- E-Series AutoSupport settings for the components in the storage controller shelf.

For additional details on E-Series AutoSupport, see the [NetApp E-Series Systems Documentation Site](#).

- Drive Security keys, which are needed to unlock secured drives (this step is required if the Drive Security feature is enabled)
- Administrator password for accessing SANtricity System Manager

Steps

1. Do one of the following:

- Use the StorageGRID Appliance Installer and select **Advanced > SANtricity System Manager**
- Use the Grid Manager and select **NODES > appliance Storage Node > SANtricity System Manager**



If these options aren't available or the login page does not appear, use the [IP addresses for the storage controllers](#). Access SANtricity System Manager by browsing to the storage controller IP.

2. Set or enter the administrator password.

SANtricity System Manager uses a single administrator password that is shared among all users.

Set Up SANtricity® System Manager



More (10 total) >

1 Welcome

2 Verify Hardware

3 Verify Hosts


4 Select Applications

5 Define Workloads

6 Ac

Welcome to the SANtricity® System Manager! With System Manager, you can...

- Configure your storage array and set up alerts.
- Monitor and troubleshoot any problems when they occur.
- Keep track of how your system is performing in real time.



3. Select **Cancel** to close the wizard.



Don't complete the Set Up wizard for a StorageGRID appliance.

4. Configure hardware alerts.

a. Select **Help** to access the online help for SANtricity System Manager.

b. Use the **Settings > Alerts** section of the online help to learn about alerts.

- c. Follow the “How To” instructions to set up email alerts, SNMP alerts, or syslog alerts.
5. Manage AutoSupport for the components in the storage controller shelf.
 - a. Select **Help** to access the online help for SANtricity System Manager.
 - b. Use the **SUPPORT > Support Center** section of the online help to learn about the AutoSupport feature.
 - c. Follow the “How To” instructions to manage AutoSupport.

For specific instructions on configuring a StorageGRID proxy for sending E-Series AutoSupport messages without using the management port, go to the [instructions for configuring storage proxy settings](#).

6. If the Drive Security feature is enabled for the appliance, create and manage the security key.
 - a. Select **Help** to access the online help for SANtricity System Manager.
 - b. Use the **Settings > System > Security key management** section of the online help to learn about Drive Security.
 - c. Follow the “How To” instructions to create and manage the security key.
7. Optionally, change the administrator password.
 - a. Select **Help** to access the online help for SANtricity System Manager.
 - b. Use the **Home > Storage array administration** section of the online help to learn about the administrator password.
 - c. Follow the “How To” instructions to change the password.

Review hardware status in SANtricity System Manager

You can use SANtricity System Manager to monitor and manage the individual hardware components in the storage controller shelf and to review hardware diagnostic and environmental information, such as component temperatures, as well as issues related to the drives.

Before you begin

- You are using a [supported web browser](#).
- To access SANtricity System Manager through Grid Manager, you have the Storage appliance administrator permission or Root access permission.
- To access SANtricity System Manager using the StorageGRID Appliance Installer, you have the SANtricity System Manager administrator username and password.
- To access SANtricity System Manager directly using a web browser, you have the SANtricity System Manager administrator username and password.



You must have SANtricity firmware 8.70 or higher to access SANtricity System Manager using the Grid Manager or the StorageGRID Appliance Installer.




Accessing SANtricity System Manager from the Grid Manager or from the Appliance Installer is generally meant only for monitoring your hardware and configuring E-Series AutoSupport. Many features and operations within SANtricity System Manager such as upgrading firmware don’t apply to monitoring your StorageGRID appliance. To avoid issues, always follow the hardware installation and maintenance instructions for your appliance.

Steps

1. Access SANtricity System Manager.
2. Enter the administrator username and password if required.
3. Click **Cancel** to close the Set Up wizard and to display the SANtricity System Manager home page.

The SANtricity System Manager home page appears. In SANtricity System Manager, the controller shelf is referred to as a storage array.



4. Review the information displayed for appliance hardware and confirm that all hardware components have a status of Optimal.
 - a. Click the **Hardware** tab.
 - b. Click **Show back of shelf**.

HARDWARE

[Learn More >](#)

Legend ▾

Show status icon details 

Controller Shelf 99 ▾     

[Show front of shelf](#)



From the back of the shelf, you can view both storage controllers, the battery in each storage controller, the two power canisters, the two fan canisters, and expansion shelves (if any). You can also view component temperatures.

- c. To see the settings for each storage controller, select the controller, and select **View settings** from the context menu.
- d. To see the settings for other components in the back of the shelf, select the component you want to view.
- e. Click **Show front of shelf**, and select the component you want to view.

From the front of the shelf, you can view the drives and the drive drawers for the storage controller shelf or the expansion shelves (if any).

If the status of any component is Needs Attention, follow the steps in the Recovery Guru to resolve the issue or contact technical support.

Set IP addresses for storage controllers using StorageGRID Appliance Installer

Management port 1 on each storage controller connects the appliance to the management network for SANtricity System Manager. If you can't access SANtricity System Manager from the StorageGRID Appliance Installer, set a static IP address for each storage controller to ensure that you don't lose your management connection to the hardware and the controller firmware in the controller shelf.

Before you begin

- You are using any management client that can connect to the StorageGRID Admin Network, or you have a service laptop.
- The client or service laptop has a supported web browser.

About this task

DHCP-assigned addresses can change at any time. Assign static IP addresses to the controllers to ensure consistent accessibility.



Follow this procedure only if you don't have access to SANtricity System Manager from the StorageGRID Appliance Installer (**Advanced > SANtricity System Manager**) or Grid Manager (**NODES > SANtricity System Manager**).

Steps

1. From the client, enter the URL for the StorageGRID Appliance Installer:

`https://Appliance_Controller_IP:8443`

For `Appliance_Controller_IP`, use the IP address for the appliance on any StorageGRID network.

The StorageGRID Appliance Installer Home page appears.

2. Select **Configure Hardware > Storage Controller Network Configuration**.

The Storage Controller Network Configuration page appears.

3. Depending on your network configuration, select **Enabled** for IPv4, IPv6, or both.

4. Make a note of the IPv4 address that is automatically displayed.

DHCP is the default method for assigning an IP address to the storage controller management port.



It might take a few minutes for the DHCP values to appear.

IPv4 Address Assignment	<input type="radio"/> Static <input checked="" type="radio"/> DHCP
IPv4 Address (CIDR)	10.224.5.166/21
Default Gateway	10.224.0.1

5. Optionally, set a static IP address for the storage controller management port.



You should either assign a static IP for the management port or assign a permanent lease for the address on the DHCP server.

- a. Select **Static**.
- b. Enter the IPv4 address, using CIDR notation.
- c. Enter the default gateway.

IPv4 Address Assignment	<input checked="" type="radio"/> Static <input type="radio"/> DHCP
IPv4 Address (CIDR)	10.224.2.200/21
Default Gateway	10.224.0.1

- d. Click **Save**.

It might take a few minutes for your changes to be applied.

When you connect to SANtricity System Manager, you will use the new static IP address as the URL:
https://Storage_Controller_IP

Configure BMC interface (SGF6112, SG6000, SG100, and SG1000)

BMC interface: Overview (SGF6112, SG6000, SG100, and SG1000)

The user interface for the baseboard management controller (BMC) on the SGF6112, SG6000, or services appliance provides status information about the hardware and allows you to configure SNMP settings and other options for the appliances.

Use the following procedures in this section to configure the BMC when you install the appliance:

- [Change admin or root password for BMC interface](#)
- [Set IP address for BMC management port](#)
- [Access BMC interface](#)
- [Configure SNMP settings](#)
- [Set up email notifications for BMC alerts](#)

If the appliance has already been installed into a grid and is running StorageGRID software, use the following procedures:



- [Place the appliance into maintenance mode](#) to access the StorageGRID appliance installer.
- See [Set IP address for BMC management port](#) for information about accessing the BMC interface using the StorageGRID Appliance Installer.

Change admin or root password for BMC interface

For security, you must change the password for the BMC's admin or root user.

Before you begin

The management client is using a [supported web browser](#).

About this task

When you first install the appliance, the BMC uses a default password for the admin or root user. You must change the password for the admin or root user to secure your system.

The default user depends on when you installed your StorageGRID appliance. The default user is **admin** for new installations and **root** for older installations.


Steps

1. From the client, enter the URL for the StorageGRID Appliance Installer:
https://Appliance_IP:8443

For *Appliance_IP*, use the IP address for the appliance on any StorageGRID network.

The StorageGRID Appliance Installer Home page appears.

2. Select **Configure Hardware > BMC Configuration**.



The screenshot shows the 'NetApp® StorageGRID® Appliance Installer' interface. The top navigation bar includes 'Home', 'Configure Networking', 'Configure Hardware' (which is currently selected and highlighted with a yellow oval), and 'Monitor Installation'. Below the navigation bar, there are two main links: 'Home' and 'BMC Configuration' (which is also highlighted with a yellow oval), and 'Storage Controller Network Configuration'.

The Baseboard Management Controller Configuration page appears.

3. Enter a new password for the admin or root account in the two fields provided.

4. Click **Save**.


Set IP address for BMC management port

Before you can access the BMC interface, configure the IP address for the BMC management port on the SGF6112, SG6000-CN controller, or services appliances.

If you are using ConfigBuilder to generate a JSON file, you can configure IP addresses automatically. See [Automate appliance installation and configuration](#).

Before you begin

- The management client is using a [supported web browser](#).
- You are using any management client that can connect to a StorageGRID network.
- The BMC management port is connected to the management network you plan to use.



About this task

For support purposes, the BMC management port allows low-level hardware access.



You should only connect this port to a secure, trusted, internal management network. If no such network is available, leave the BMC port unconnected or blocked, unless a BMC connection is requested by technical support.

Steps


1. From the client, enter the URL for the StorageGRID Appliance Installer:

`https://Appliance_IP:8443`

For Appliance_IP, use the IP address for the appliance on any StorageGRID network.

The StorageGRID Appliance Installer Home page appears.

2. Select **Configure Hardware > BMC Configuration**.



The Baseboard Management Controller Configuration page appears.

3. Make a note of the IPv4 address that is automatically displayed.

DHCP is the default method for assigning an IP address to this port.



It might take a few minutes for the DHCP values to appear.

Baseboard Management Controller Configuration

LAN IP Settings

IP Assignment	<input type="radio"/> Static <input checked="" type="radio"/> DHCP
MAC Address	d8:c4:97:28:50:62
IPv4 Address (CIDR)	10.224.3.225/21
Default gateway	10.224.0.1

Cancel **Save**

4. Optionally, set a static IP address for the BMC management port.



You should either assign a static IP for the BMC management port or assign a permanent lease for the address on the DHCP server.

- a. Select **Static**.
- b. Enter the IPv4 address, using CIDR notation.
- c. Enter the default gateway.

Baseboard Management Controller Configuration

LAN IP Settings

IP Assignment Static DHCP

MAC Address d8:c4:97:28:50:62

IPv4 Address (CIDR) 10.224.3.225/21

Default gateway 10.224.0.1

Cancel Save

- d. Click **Save**.

It might take a few minutes for your changes to be applied.

Access BMC interface


You can access the BMC interface using the DHCP or static IP address for the BMC management port on the following appliance models:

- SGF6112
- SG6000
- SG1000
- SG100

Before you begin

- The management client is using a [supported web browser](#).
- The BMC management port on the appliance is connected to the management network you plan to use.

SGF6112



SG6000



SG1000



SG100



Steps

1. Enter the URL for the BMC interface:

https://BMC_Port_IP

For *BMC_Port_IP*, use the DHCP or static IP address for the BMC management port.

The BMC sign-in page appears.




If you haven't yet configured *BMC_Port_IP*, follow the instructions in [Configure BMC interface](#). If you are unable to follow that procedure due to a hardware problem, and have not yet configured a BMC IP address, you might still be able to access the BMC. By default, the BMC obtains an IP address using DHCP. If DHCP is enabled on the BMC network, your network administrator can provide the IP address assigned to the BMC MAC, which is printed on the label on the front of the appliance. If DHCP is not enabled on the BMC network, the BMC will not respond after a few minutes and assign itself the default static IP 192.168.0.120. You might need to connect your laptop directly to the BMC port, and change the networking setting to assign your laptop an IP such as 192.168.0.200/24, in order to browse to 192.168.0.120.

2. Enter the admin or root username and password, using the password you set when you [changed the default root password](#):



The default user depends on when you installed your StorageGRID appliance. The default user is **admin** for new installations and **root** for older installations.

3. Select **Sign me in**.



4. Optionally, create additional users by selecting **Settings > User Management** and clicking on any "disabled" user.



When users sign in for the first time, they might be prompted to change their password for increased security.

Configure SNMP settings for BMC

If you are familiar with configuring SNMP for hardware, you can use the BMC interface to configure the SNMP settings for the SGF6112, SG6000, and services appliances. You can provide secure community strings, enable SNMP Trap, and specify up to five SNMP destinations.

Before you begin

- You know how to access the BMC dashboard.
- You have experience in configuring SNMP settings for SNMPv1-v2c equipment.



BMC settings made by this procedure might not be preserved if the appliance fails and has to be replaced. Make sure you have a record of all settings you have applied, so they can be easily reapplied after a hardware replacement if necessary.

Steps

1. From the BMC dashboard, select **Settings > SNMP Settings**.
2. On the SNMP Settings page, select **Enable SNMP V1/V2**, and then provide a Read-Only Community String and a Read-Write Community String.

The Read-Only Community String is like a user ID or password. You should change this value to prevent intruders from getting information about your network setup. The Read-Write Community String protects the device against unauthorized changes.

3. Optionally, select **Enable Trap**, and enter the required information.



Enter the Destination IP for each SNMP trap using an IP address. Fully qualified domain names aren't supported.

Enable traps if you want the appliance to send immediate notifications to an SNMP console when it is in an unusual state. Depending on the device, traps might indicate hardware failures of various components, link up/down conditions, temperature thresholds being exceeded, or high traffic.

4. Optionally, click **Send Test Trap** to test your settings.

5. If the settings are correct, click **Save**.

Set up email notifications for BMC alerts

If you want email notifications to be sent when alerts occur, use the BMC interface to configure SMTP settings, users, LAN destinations, alert policies, and event filters.




BMC settings made by this procedure might not be preserved if the SG6000-CN controller or services appliance fails and has to be replaced. Make sure you have a record of all settings you have applied, so they can be easily reapplied after a hardware replacement if necessary.

Before you begin

You know how to access the BMC dashboard.

About this task

In the BMC interface, you use the **SMTP Settings**, **User Management**, and **Platform Event Filters** options on the Settings page to configure email notifications.



Steps

1. Configure SNMP settings for BMC.

- a. Select **Settings > SMTP Settings**.
- b. For Sender Email ID, enter a valid email address.

This email address is provided as the From address when the BMC sends email.

2. Set up users to receive alerts.

a. From the BMC dashboard, select **Settings > User Management**.

b. Add at least one user to receive alert notifications.

The email address you configure for a user is the address the BMC sends alert notifications to. For example, you could add a generic user, such as “notification-user,” and use the email address of a technical support team email distribution list.

3. Configure the LAN destination for alerts.

a. Select **Settings > Platform Event Filters > LAN Destinations**.

b. Configure at least one LAN destination.

- Select **Email** as the Destination Type.

- For BMC Username, select a user name that you added earlier.

- If you added multiple users and want all of them to receive notification emails, add a LAN Destination for each user.

c. Send a test alert.

4. Configure alert policies so you can define when and where the BMC sends alerts.

a. Select **Settings > Platform Event Filters > Alert Policies**.

b. Configure at least one alert policy for each LAN destination.

- For Policy Group Number, select **1**.

- For Policy Action, select **Always send alert to this destination**.

- For LAN Channel, select **1**.

- In the Destination Selector, select the LAN destination for the policy.

5. Configure event filters to direct alerts for different event types to the appropriate users.

a. Select **Settings > Platform Event Filters > Event Filters**.

b. For Alert Policy Group Number, enter **1**.

c. Create filters for every event you want the Alert Policy Group to be notified about.

- You can create event filters for power actions, specific sensor events, or all events.

- If you are uncertain which events to monitor, select **All Sensors** for Sensor Type and **All Events** for Event Options. If you receive unwanted notifications, you can change your selections later.

Optional: Enable node encryption

If you enable node encryption, the disks in your appliance can be protected by secure key management server (KMS) encryption against physical loss or removal from the site. You must select and enable node encryption during appliance installation. You can't disable node encryption after the KMS encryption process starts.

If you are using ConfigBuilder to generate a JSON file, you can enable node encryption automatically. See [Automate appliance installation and configuration](#).

Before you begin

Review the information about [configuring KMS](#).

About this task

An appliance that has node encryption enabled connects to the external key management server (KMS) that is configured for the StorageGRID site. Each KMS (or KMS cluster) manages the encryption keys for all appliance nodes at the site. These keys encrypt and decrypt the data on each disk in an appliance that has node encryption enabled.

A KMS can be set up in Grid Manager before or after the appliance is installed in StorageGRID. See the information about KMS and appliance configuration in the instructions for administering StorageGRID for additional details.

- If a KMS is set up before installing the appliance, KMS-controlled encryption begins when you enable node encryption on the appliance and add it to a StorageGRID site where KMS is configured.
- If a KMS is not set up before you install the appliance, KMS-controlled encryption is performed on each appliance that has node encryption enabled as soon as a KMS is configured and available for the site that contains the appliance node.

 When an appliance is installed with node encryption enabled, a temporary key is assigned. The data on the appliance is not protected until the appliance is connected to the Key Management System (KMS) and a KMS security key is set. See the [KMS appliance configuration overview](#) for additional information.

Without the KMS key needed to decrypt the disk, data on the appliance can't be retrieved and is effectively lost. This is the case whenever the decryption key can't be retrieved from the KMS. The key becomes inaccessible if a customer clears the KMS configuration, a KMS key expires, connection to the KMS is lost, or the appliance is removed from the StorageGRID system where its KMS keys are installed.

Steps

1. Open a browser, and enter one of the IP addresses for the appliance's compute controller.

`https://Controller_IP:8443`

`Controller_IP` is the IP address of the compute controller (not the storage controller) on any of the three StorageGRID networks.

The StorageGRID Appliance Installer Home page appears.



After the appliance has been encrypted with a KMS key, the appliance disks can't be decrypted without using the same KMS key.

2. Select **Configure Hardware > Node Encryption**.

NetApp® StorageGRID® Appliance Installer

Help ▾

Home	Configure Networking ▾	Configure Hardware ▾	Monitor Installation	Advanced ▾	
------	------------------------	----------------------	----------------------	------------	--

Node Encryption

Node encryption allows you to use an external key management server (KMS) to encrypt all StorageGRID data on this appliance. If node encryption is enabled for the appliance and a KMS is configured for the site, you cannot access any data on the appliance unless the appliance can communicate with the KMS.

Encryption Status

 You can only enable node encryption for an appliance during installation. You cannot enable or disable the node encryption setting after the appliance is installed.

Enable node encryption

Save

Key Management Server Details

3. Select **Enable node encryption**.

Before appliance installation, you can clear **Enable node encryption** without risk of data loss. When the installation begins, the appliance node accesses the KMS encryption keys in your StorageGRID system and begins disk encryption. You can't disable node encryption after the appliance is installed.



After you add an appliance that has node encryption enabled to a StorageGRID site that has a KMS, you can't stop using KMS encryption for the node.

4. Select **Save**.

5. Deploy the appliance as a node in your StorageGRID system.

KMS-controlled encryption begins when the appliance accesses the KMS keys configured for your StorageGRID site. The installer displays progress messages during the KMS encryption process, which might take a few minutes depending on the number of disk volumes in the appliance.



Appliances are initially configured with a random non-KMS encryption key assigned to each disk volume. The disks are encrypted using this temporary encryption key, that is not secure, until the appliance that has node encryption enabled accesses the KMS keys configured for your StorageGRID site.

After you finish

You can view node-encryption status, KMS details, and the certificates in use when the appliance node is in maintenance mode. See [Monitor node encryption in maintenance mode](#) for information.

Optional: Change RAID mode

On some appliance models, you can change to a different RAID mode on the appliance to accommodate your storage and recovery requirements. You can only change the mode before deploying the appliance Storage Node.

If you are using ConfigBuilder to generate a JSON file, you can change the RAID mode automatically. See [Automate appliance installation and configuration](#).

About this task

If supported by your appliance, you can choose one of the following volume configuration options:

- **Dynamic Disk Pools (DDP):** This mode uses two parity drives for every eight data drives. This is the default and recommended mode for all appliances. When compared to RAID 6, DDP delivers better system performance, reduced rebuild times after drive failures, and ease of management. DDP also provides drawer loss protection in SG5760 appliances.



DDP does not provide drawer loss protection in SG6060 appliances because of the two SSDs. Drawer loss protection is effective in any expansion shelves that are added to an SG6060.

- **DDP16:** This mode uses two parity drives for every 16 data drives, which results in higher storage efficiency compared to DDP. When compared to RAID 6, DDP16 delivers better system performance, reduced rebuild times after drive failures, ease of management, and comparable storage efficiency. To use DDP16 mode, your configuration must contain at least 20 drives. DDP16 does not provide drawer loss protection.
- **RAID6:** This mode uses two parity drives for every 16 or more data drives. It is a hardware protection scheme that uses parity stripes on each disk, and allows for two disk failures within the RAID set before any data is lost. To use RAID 6 mode, your configuration must contain at least 20 drives. Although RAID 6 can increase storage efficiency of the appliance when compared to DDP, it is not recommended for most StorageGRID environments.



If any volumes have already been configured or if StorageGRID was previously installed, changing the RAID mode causes the volumes to be removed and replaced. Any data on those volumes will be lost.

SG6000

Before you begin

- You are using any client that can connect to StorageGRID.
- The client has a [supported web browser](#).

Steps

1. Open a browser, and enter one of the IP addresses for the appliance's compute controller.

`https://Controller_IP:8443`

`Controller_IP` is the IP address of the compute controller (not the storage controller) on any of the three StorageGRID networks.

The StorageGRID Appliance Installer Home page appears.

2. Select **Advanced > RAID Mode**.
3. On the **Configure RAID Mode** page, select the desired RAID mode from the Mode drop-down list.
4. Click **Save**.

SG5760

Before you begin

- You have an SG5760 with 60 drives. If you have an SG5712, you must use the default DDP mode.
- You are using any client that can connect to StorageGRID.
- The client has a [supported web browser](#).

Steps

1. Using the service laptop, open a web browser and access the StorageGRID Appliance Installer:
`https://E5700SG_Controller_IP:8443`

Where `E5700SG_Controller_IP` is any of the IP addresses for the E5700SG controller.

2. Select **Advanced > RAID Mode**.
3. On the **Configure RAID Mode** page, select the desired RAID mode from the Mode drop-down list.
4. Click **Save**.

Related information

[NetApp E-Series Systems Documentation Site](#)

Optional: Remap network ports for appliance

You can optionally remap the internal ports on an appliance node to different external ports. For example, you might need to remap ports because of a firewall issue.

Before you begin

- You have previously accessed the StorageGRID Appliance Installer.

About this task

You can't use remapped ports for load balancer endpoints. If you need to remove a remapped port, follow the steps in [Remove port remaps](#).

Steps

1. From the StorageGRID Appliance Installer, select **Configure Networking > Remap Ports**.

The Remap Port page appears.

2. From the **Network** drop-down box, select the network for the port you want to remap: Grid, Admin, or Client.
3. From the **Protocol** drop-down box, select the IP protocol: TCP or UDP.
4. From the **Remap Direction** drop-down box, select which traffic direction you want to remap for this port: Inbound, Outbound, or Bi-directional.
5. For **Original Port**, enter the number of the port you want to remap.
6. For **Mapped-To Port**, enter the number of the port you want to use instead.
7. Select **Add Rule**.

The new port mapping is added to the table, and the remapping takes effect immediately.

8. To remove a port mapping, select the radio button for the rule you want to remove, and select **Remove Selected Rule**.

Deploy appliance node

Deploy appliance Storage Node

After installing and configuring the storage appliance, you can deploy it as a Storage Node in a StorageGRID system. When you deploy an appliance as a Storage Node, you use the StorageGRID Appliance Installer included on the appliance.

Before you begin

- If you are cloning an appliance node, continue following the [appliance node cloning](#) process.
- The appliance has been installed in a rack or cabinet, connected to your networks, and powered on.
- Network links, IP addresses, and port remapping (if necessary) have been configured for the appliance using the StorageGRID Appliance Installer.
- You know one of the IP addresses assigned to the appliance's compute controller. You can use the IP address for any attached StorageGRID network.
- The primary Admin Node for the StorageGRID system has been deployed.
- All Grid Network subnets listed on the IP Configuration page of the StorageGRID Appliance Installer have been defined in the Grid Network Subnet List on the primary Admin Node.
- You have a service laptop with a supported web browser.

About this task

Each storage appliance functions as a single Storage Node. Any appliance can connect to the Grid Network, the Admin Network, and the Client Network

To deploy an appliance Storage Node in a StorageGRID system, you access the StorageGRID Appliance Installer and perform the following steps:

- You specify or confirm the IP address of the primary Admin Node and the name of the Storage Node.
- You start the deployment and wait as volumes are configured and the software is installed.
- When the installation pauses partway through the appliance installation tasks, you resume the installation by signing into the Grid Manager, approving all grid nodes, and completing the StorageGRID installation and deployment processes.



If you need to deploy multiple appliance nodes at one time, you can automate the installation process by using the `configure-sga.py` Appliance Installation script.

- If you are performing an expansion or recovery operation, follow the appropriate instructions:
 - To add an appliance Storage Node to an existing StorageGRID system, see the instructions for [adding grid nodes](#).
 - To deploy an appliance Storage Node as part of a recovery operation, see instructions [recovering an appliance Storage Node](#).

Steps

1. Open a browser, and enter one of the IP addresses for the appliance's compute controller.

`https://Controller_IP:8443`

The StorageGRID Appliance Installer Home page appears.

NetApp® StorageGRID® Appliance Installer

[Home](#)[Configure Networking](#)[Configure Hardware](#)[Monitor Installation](#)[Advanced](#)

Home

 The installation is ready to be started. Review the settings below, and then click Start Installation.

Primary Admin Node connection

Enable Admin Node discovery

Primary Admin Node IP

Connection state Connection to 172.16.4.210 ready

[Cancel](#)[Save](#)

Node name

Node name

[Cancel](#)[Save](#)

Installation

Current state Ready to start installation of NetApp-SGA into grid with Admin Node 172.16.4.210.

[Start Installation](#)

2. In the **Primary Admin Node connection** section, determine whether you need to specify the IP address for the primary Admin Node.

If you have previously installed other nodes in this data center, the StorageGRID Appliance Installer can discover this IP address automatically, assuming the primary Admin Node, or at least one other grid node with ADMIN_IP configured, is present on the same subnet.

3. If this IP address is not shown or you need to change it, specify the address:

Option	Description
Manual IP entry	<ol style="list-style-type: none"> Clear the Enable Admin Node discovery checkbox. Enter the IP address manually. Click Save. Wait for the connection state for the new IP address to become ready.
Automatic discovery of all connected primary Admin Nodes	<ol style="list-style-type: none"> Select the Enable Admin Node discovery checkbox. Wait for the list of discovered IP addresses to be displayed. Select the primary Admin Node for the grid where this appliance Storage Node will be deployed. Click Save. Wait for the connection state for the new IP address to become ready.

4. In the **Node name** field, provide the system name you want to use for this appliance node, and click **Save**.

The name that appears here will be the appliance node's system name. System names are required for internal StorageGRID operations and can't be changed.

5. In the **Installation** section, confirm that the current state is "Ready to start installation of *node_name* into grid with primary Admin Node *admin_ip*" and that the **Start Installation** button is enabled.

If the **Start Installation** button is not enabled, you might need to change the network configuration or port settings. For instructions, see the maintenance instructions for your appliance.



If you are deploying the Storage Node appliance as a node cloning target, stop the deployment process here and continue the [node cloning procedure](#).

6. From the StorageGRID Appliance Installer home page, click **Start Installation**.

The Current state changes to "Installation is in progress," and the Monitor Installation page is displayed.



If you need to access the Monitor Installation page manually, click **Monitor Installation**.

7. If your grid includes multiple appliance Storage Nodes, repeat these steps for each appliance.



If you need to deploy multiple appliance Storage Nodes at one time, you can automate the installation process by using the `configure-sga.py` Appliance Installation script.

Deploy services appliance node

You can deploy a services appliance as a primary Admin Node, a non-primary Admin Node, or a Gateway Node. Both the SG100 and the SG1000 appliances can operate as Gateway Nodes and Admin Nodes (primary or non-primary) at the same time.

Deploy services appliance as primary Admin Node

When you deploy a services appliance as a primary Admin Node, you use the StorageGRID Appliance Installer included on the appliance to install the StorageGRID software, or you upload the software version you want to install. You must install and configure the primary Admin Node before you install any other appliance node types. A primary Admin Node can connect to the Grid Network, and to the optional Admin Network and Client Network, if one or both are configured.

Before you begin

- The appliance has been installed in a rack or cabinet, connected to your networks, and powered on.
- Network links, IP addresses, and port remapping (if necessary) have been configured for the appliance using the StorageGRID Appliance Installer.
- You have a service laptop with a [supported web browser](#).
- You know one of the IP addresses assigned to the appliance. You can use the IP address for any attached StorageGRID network.

About this task

To install StorageGRID on an appliance primary Admin Node:

- You use the StorageGRID Appliance Installer to install the StorageGRID software. If you want to install a different version of the software, you first upload it using the StorageGRID Appliance Installer.
- You wait as the software is installed.
- When the software has been installed, the appliance is rebooted automatically.

Steps

1. Open a browser, and enter the IP address for the appliance.

https://services_appliance_IP:8443

The StorageGRID Appliance Installer Home page appears.

2. In the **This Node** section, select **Primary Admin**.
3. In the **Node name** field, enter the name you want to use for this appliance node, and click **Save**.

The node name is assigned to this appliance node in the StorageGRID system. It is shown on the Grid Nodes page in the Grid Manager.

4. Optionally, to install a different version of the StorageGRID software, follow these steps:
 - a. Download the installation archive:
[NetApp Downloads: StorageGRID Appliance](#)
 - b. Extract the archive.
 - c. From the StorageGRID Appliance Installer, select **Advanced > Upload StorageGRID Software**.
 - d. Click **Remove** to remove the current software package.

NetApp® StorageGRID® Appliance Installer

	Home	Configure Networking ▾	Configure Hardware ▾	Monitor Installation	Advanced ▾
--	------	------------------------	----------------------	----------------------	------------

Upload StorageGRID Software

If this node is the primary Admin Node of a new deployment, you must use this page to upload the StorageGRID software installation package, unless the version of the software you want to install has already been uploaded. If you are adding this node to an existing deployment, you can avoid network traffic by uploading the installation package that matches the software version running on the existing grid. If you do not upload the correct package, the node obtains the software from the grid's primary Admin Node during installation.

Current StorageGRID Installation Software

Version 11.3.0

Package Name storagegrid-webscale-images-11-3-0_11.3.0-20190806.1731.4064510_amd64.deb

[Remove](#)

e. Click **Browse** for the software package you downloaded and extracted, and then click **Browse** for the checksum file.

NetApp® StorageGRID® Appliance Installer

	Home	Configure Networking ▾	Configure Hardware ▾	Monitor Installation	Advanced ▾
--	------	------------------------	----------------------	----------------------	------------

Upload StorageGRID Software

If this node is the primary Admin Node of a new deployment, you must use this page to upload the StorageGRID software installation package, unless the version of the software you want to install has already been uploaded. If you are adding this node to an existing deployment, you can avoid network traffic by uploading the installation package that matches the software version running on the existing grid. If you do not upload the correct package, the node obtains the software from the grid's primary Admin Node during installation.

Current StorageGRID Installation Software

Version None

Package Name None

Upload StorageGRID Installation Software

Software Package [Browse](#)

Checksum File [Browse](#)

f. Select **Home** to return to the Home page.

5. Confirm that the current state is “Ready to start installation of primary Admin Node name with software version x.y” and that the **Start Installation** button is enabled.



If you are deploying the Admin Node appliance as a node cloning target, stop the deployment process here and continue the [node cloning procedure](#).

6. From the StorageGRID Appliance Installer home page, click **Start Installation**.

Home

1 The installation is ready to be started. Review the settings below, and then click Start Installation.

This Node

Node type	Primary Admin (with Load Balancer)
Node name	xlr8r-8
<input type="button" value="Cancel"/> <input type="button" value="Save"/>	

Installation

Current state	Ready to start installation of xlr8r-8 as primary Admin Node of a new grid running StorageGRID 11.6.0.
<input type="button" value="Start Installation"/>	

The Current state changes to “Installation is in progress,” and the Monitor Installation page is displayed.



If you need to access the Monitor Installation page manually, click **Monitor Installation** from the menu bar.

Deploy services appliance as Gateway or non-primary Admin Node

When you deploy a services appliance as a Gateway Node or non-primary Admin Node, you use the StorageGRID Appliance Installer included on the appliance.

Before you begin

- The appliance has been installed in a rack or cabinet, connected to your networks, and powered on.
- Network links, IP addresses, and port remapping (if necessary) have been configured for the appliance using the StorageGRID Appliance Installer.
- The primary Admin Node for the StorageGRID system has been deployed.
- All Grid Network subnets listed on the IP Configuration page of the StorageGRID Appliance Installer have been defined in the Grid Network Subnet List on the primary Admin Node.
- You have a service laptop with a [supported web browser](#).
- You know the IP address assigned to the appliance. You can use the IP address for any attached StorageGRID network.

About this task

To install StorageGRID on a services appliance node:

- You specify or confirm the IP address of the primary Admin Node and the name of the appliance node.
- You start the installation and wait as the software is installed.

Partway through the appliance Gateway Node installation tasks, the installation pauses. To resume the

installation, you sign into the Grid Manager, approve all grid nodes, and complete the StorageGRID installation process. The installation of a non-primary Admin Node does not require your approval.



Don't deploy the SG100 and SG1000 service appliances in the same site. Unpredictable performance might result.



If you need to deploy multiple appliance nodes at one time, you can automate the installation process. See [Automate appliance installation and configuration](#).

Steps

1. Open a browser, and enter the IP address for the appliance.

https://Controller_IP:8443

The StorageGRID Appliance Installer Home page appears.

2. In the Primary Admin Node connection section, determine whether you need to specify the IP address for the primary Admin Node.

If you have previously installed other nodes in this data center, the StorageGRID Appliance Installer can discover this IP address automatically, assuming the primary Admin Node, or at least one other grid node with ADMIN_IP configured, is present on the same subnet.

3. If this IP address is not shown or you need to change it, specify the address:

Option	Description
Manual IP entry	<ol style="list-style-type: none">Clear the Enable Admin Node discovery checkbox.Enter the IP address manually.Click Save.Wait for the connection state for the new IP address to become ready.
Automatic discovery of all connected primary Admin Nodes	<ol style="list-style-type: none">Select the Enable Admin Node discovery checkbox.Wait for the list of discovered IP addresses to be displayed.Select the primary Admin Node for the grid where this appliance Storage Node will be deployed.Click Save.Wait for the connection state for the new IP address to become ready.

4. In the **Node name** field, provide the system name you want to use for this appliance node, and click **Save**.


The name that appears here will be the appliance node's system name. System names are required for internal StorageGRID operations and can't be changed.

5. Optionally, to install a different version of the StorageGRID software, follow these steps:

- Download the installation archive:

NetApp Downloads: StorageGRID Appliance

- b. Extract the archive.
- c. From the StorageGRID Appliance Installer, select **Advanced > Upload StorageGRID Software**.
- d. Click **Remove** to remove the current software package.



NetApp® StorageGRID® Appliance Installer

Home Configure Networking ▾ Configure Hardware ▾ Monitor Installation Advanced ▾

Upload StorageGRID Software

If this node is the primary Admin Node of a new deployment, you must use this page to upload the StorageGRID software installation package, unless the version of the software you want to install has already been uploaded. If you are adding this node to an existing deployment, you can avoid network traffic by uploading the installation package that matches the software version running on the existing grid. If you do not upload the correct package, the node obtains the software from the grid's primary Admin Node during installation.

Current StorageGRID Installation Software

Version	11.3.0
Package Name	storagegrid-webscale-images-11-3-0_11.3.0-20190806.1731.4064510_amd64.deb

Remove

- e. Click **Browse** for the software package you downloaded and extracted, and then click **Browse** for the checksum file.



NetApp® StorageGRID® Appliance Installer

Home Configure Networking ▾ Configure Hardware ▾ Monitor Installation Advanced ▾

Upload StorageGRID Software

If this node is the primary Admin Node of a new deployment, you must use this page to upload the StorageGRID software installation package, unless the version of the software you want to install has already been uploaded. If you are adding this node to an existing deployment, you can avoid network traffic by uploading the installation package that matches the software version running on the existing grid. If you do not upload the correct package, the node obtains the software from the grid's primary Admin Node during installation.

Current StorageGRID Installation Software

Version	None
Package Name	None



Upload StorageGRID Installation Software

Software Package	Browse
Checksum File	Browse

- f. Select **Home** to return to the Home page.
6. In the Installation section, confirm that the current state is "Ready to start installation of *node_name* into grid with primary Admin Node *admin_ip*" and that the **Start Installation** button is enabled.

If the **Start Installation** button is not enabled, you might need to change the network configuration or port settings. For instructions, see the maintenance instructions for your appliance.

7. From the StorageGRID Appliance Installer home page, click **Start Installation**.

The installation is ready to be started. Review the settings below, and then click Start Installation.

This Node

Node type: Non-primary Admin (with Load Balancer)

Node name: GW-SG1000-003-074

Primary Admin Node connection

Enable Admin Node discovery

Primary Admin Node IP: 172.16.6.32

Connection state: Connection to 172.16.6.32 ready

Installation

Current state: Ready to start installation of GW-SG1000-003-074 into grid with Admin Node 172.16.6.32 running StorageGRID 11.6.0, using StorageGRID software downloaded from the Admin Node.

The Current state changes to “Installation is in progress,” and the Monitor Installation page is displayed.



If you need to access the Monitor Installation page manually, click **Monitor Installation** from the menu bar.

8. If your grid includes multiple appliance nodes, repeat the previous steps for each appliance.

Monitor appliance installation

The StorageGRID Appliance Installer provides status until installation is complete. When the software installation is complete, the appliance is rebooted.

Example 1. Steps

Storage appliance

1. To monitor the installation progress, click **Monitor Installation**.

The Monitor Installation page shows the installation progress.

Monitor Installation

1. Configure storage		
Step	Progress	Status
Connect to storage controller	<div style="width: 100%; background-color: green;"></div>	Complete
Clear existing configuration	<div style="width: 100%; background-color: green;"></div>	Complete
Configure volumes	<div style="width: 25%; background-color: blue;"></div>	Creating volume StorageGRID-obj-00
Configure host settings	<div style="width: 0%; background-color: lightgray;"></div>	Pending

2. Install OS		
Step	Progress	Status
3. Install StorageGRID	<div style="width: 0%; background-color: lightgray;"></div>	Pending
4. Finalize installation	<div style="width: 0%; background-color: lightgray;"></div>	Pending

The blue status bar indicates which task is currently in progress. Green status bars indicate tasks that have completed successfully.



The installer ensures that tasks completed in a previous install aren't re-run. If you are re-running an installation, any tasks that don't need to be re-run are shown with a green status bar and a status of "Skipped."

2. Review the progress of the first two installation stages.

1. Configure storage

During this stage, the installer connects to the storage controller, clears any existing configuration, creates RAIDs according to the configured RAID Mode, allocates volumes for StorageGRID software and object data storage, and configures host settings.

2. Install OS

During this stage, the installer copies the base operating system image for StorageGRID to the appliance.

3. Continue monitoring the installation progress until the **Install StorageGRID** stage pauses and a message appears on the embedded console, prompting you to approve this node on the Admin Node using the Grid Manager. Go to the next step.

Monitor Installation

1. Configure storage	Complete
2. Install OS	Complete
3. Install StorageGRID	Running
4. Finalize installation	Pending

Connected (unencrypted) to: QEMU

```
/platform.type=: Device or resource busy
[2017-07-31T22:09:12.362566]    INFO -- [INSG] NOTICE: seeding /var/local with c
ontainer data
[2017-07-31T22:09:12.366205]    INFO -- [INSG] Fixing permissions
[2017-07-31T22:09:12.369633]    INFO -- [INSG] Enabling syslog
[2017-07-31T22:09:12.511533]    INFO -- [INSG] Stopping system logging: syslog-n
g.
[2017-07-31T22:09:12.570096]    INFO -- [INSG] Starting system logging: syslog-n
g.
[2017-07-31T22:09:12.576360]    INFO -- [INSG] Beginning negotiation for downloa
d of node configuration
[2017-07-31T22:09:12.581363]    INFO -- [INSG]
[2017-07-31T22:09:12.585066]    INFO -- [INSG]
[2017-07-31T22:09:12.588314]    INFO -- [INSG]
[2017-07-31T22:09:12.591851]    INFO -- [INSG]
[2017-07-31T22:09:12.594886]    INFO -- [INSG]
[2017-07-31T22:09:12.598360]    INFO -- [INSG]
[2017-07-31T22:09:12.601324]    INFO -- [INSG]
[2017-07-31T22:09:12.604759]    INFO -- [INSG]
[2017-07-31T22:09:12.607800]    INFO -- [INSG]
[2017-07-31T22:09:12.610985]    INFO -- [INSG]
[2017-07-31T22:09:12.614597]    INFO -- [INSG]
[2017-07-31T22:09:12.618282]    INFO -- [INSG] Please approve this node on the A
dmin Node GMI to proceed...
_
```

4. Go to the Grid Manager of the Primary Admin node, approve the pending storage node, and complete the StorageGRID installation process.

When you click **Install** from the Grid Manager, Stage 3 completes and stage 4, **Finalize Installation**, begins. When stage 4 completes, the controller is rebooted.

Services appliance

1. To monitor the installation progress, click **Monitor Installation** from the menu bar.

The Monitor Installation page shows the installation progress.

Monitor Installation

1. Configure storage	Complete												
2. Install OS	Running												
<table><thead><tr><th>Step</th><th>Progress</th><th>Status</th></tr></thead><tbody><tr><td>Obtain installer binaries</td><td><div style="width: 100%; background-color: green;"></div></td><td>Complete</td></tr><tr><td>Configure installer</td><td><div style="width: 100%; background-color: green;"></div></td><td>Complete</td></tr><tr><td>Install OS</td><td><div style="width: 50%; background-color: #0072BD;"></div></td><td>Installer VM running</td></tr></tbody></table>		Step	Progress	Status	Obtain installer binaries	<div style="width: 100%; background-color: green;"></div>	Complete	Configure installer	<div style="width: 100%; background-color: green;"></div>	Complete	Install OS	<div style="width: 50%; background-color: #0072BD;"></div>	Installer VM running
Step	Progress	Status											
Obtain installer binaries	<div style="width: 100%; background-color: green;"></div>	Complete											
Configure installer	<div style="width: 100%; background-color: green;"></div>	Complete											
Install OS	<div style="width: 50%; background-color: #0072BD;"></div>	Installer VM running											
3. Install StorageGRID	Pending												
4. Finalize installation	Pending												

The blue status bar indicates which task is currently in progress. Green status bars indicate tasks that have completed successfully.



The installer ensures that tasks completed in a previous install aren't re-run. If you are re-running an installation, any tasks that don't need to be re-run are shown with a green status bar and a status of "Skipped."

2. Review the progress of first two installation stages.

◦ 1. Configure storage

During this stage, the installer clears any existing configuration from the drives in the appliance, and configures host settings.

◦ 2. Install OS

During this stage, the installer copies the base operating system image for StorageGRID to the appliance.

3. Continue monitoring the installation progress until one of the following processes occurs:

- For all appliance nodes except the primary Admin Node, the Install StorageGRID stage pauses and a message appears on the embedded console, prompting you to approve this node on the Admin Node using the Grid Manager. Go to the next step.
- For appliance primary Admin Node installation, you don't need to approve the node. The appliance is rebooted. You can skip the next step.



During installation of an appliance primary Admin Node, a fifth phase appears (see the example screen shot showing four phases). If the fifth phase is in progress for more than 10 minutes, refresh the web page manually.

Monitor Installation

1. Configure storage	Complete
2. Install OS	Complete
3. Install StorageGRID	Running
4. Finalize installation	Pending

Connected (unencrypted) to: QEMU

```
/platform.type■: Device or resource busy
[2017-07-31T22:09:12.362566]    INFO -- [INSG] NOTICE: seeding /var/local with c
ontainer data
[2017-07-31T22:09:12.366205]    INFO -- [INSG] Fixing permissions
[2017-07-31T22:09:12.369633]    INFO -- [INSG] Enabling syslog
[2017-07-31T22:09:12.511533]    INFO -- [INSG] Stopping system logging: syslog-n
g.
[2017-07-31T22:09:12.570096]    INFO -- [INSG] Starting system logging: syslog-n
g.
[2017-07-31T22:09:12.576360]    INFO -- [INSG] Beginning negotiation for downlo
ad of node configuration
[2017-07-31T22:09:12.581363]    INFO -- [INSG]
[2017-07-31T22:09:12.585066]    INFO -- [INSG]
[2017-07-31T22:09:12.588314]    INFO -- [INSG]
[2017-07-31T22:09:12.591851]    INFO -- [INSG]
[2017-07-31T22:09:12.594886]    INFO -- [INSG]
[2017-07-31T22:09:12.598360]    INFO -- [INSG]
[2017-07-31T22:09:12.601324]    INFO -- [INSG]
[2017-07-31T22:09:12.604759]    INFO -- [INSG]
[2017-07-31T22:09:12.607800]    INFO -- [INSG]
[2017-07-31T22:09:12.610985]    INFO -- [INSG]
[2017-07-31T22:09:12.614597]    INFO -- [INSG]
[2017-07-31T22:09:12.618282]    INFO -- [INSG] Please approve this node on the A
dmin Node GMI to proceed...
```

4. Go to the Grid Manager of the Primary Admin node, approve the pending grid node, and complete the StorageGRID installation process.

When you click **Install** from the Grid Manager, Stage 3 completes and stage 4, **Finalize Installation**, begins. When stage 4 completes, the appliance is rebooted.

Reboot appliance while StorageGRID Appliance Installer is running

You might need to reboot the appliance while the StorageGRID Appliance Installer is running. For example, you might need to reboot the appliance if the installation fails.

About this task

This procedure only applies when the appliance is running the StorageGRID Appliance Installer. Once the installation is completed, this step no longer works because the StorageGRID Appliance Installer is no longer available.

Steps

1. From the StorageGRID Appliance Installer, click **Advanced > Reboot Controller**, and then select one of these options:
 - Select **Reboot into StorageGRID** to reboot the controller with the node rejoining the grid. Select this option if you are done working in maintenance mode and are ready to return the node to normal operation.
 - Select **Reboot into Maintenance Mode** to reboot the controller with the node remaining in maintenance mode. (This option is available only when the controller is in maintenance mode.) Select this option if there are additional maintenance operations you need to perform on the node before rejoining the grid.



The appliance is rebooted.

Troubleshoot hardware installation (SGF6112)

If you encounter issues during the installation, you might find it helpful to review troubleshooting information related to hardware setup and connectivity issues.

View boot-up codes (SGF6112)

When you apply power to the appliance, the BMC logs a series of boot-up codes. You can view these codes on a graphical console that is connected to the BMC management port.

Before you begin

- You know how to access the BMC dashboard.
- If you want to use serial-over-LAN (SOL), you have experience using IPMI SOL console applications.

Steps

1. Select one of the following methods for viewing the boot-up codes for the appliance controller, and gather the required equipment.

Method	Required equipment
VGA console	<ul style="list-style-type: none"> • VGA-capable monitor • VGA cable
KVM	<ul style="list-style-type: none"> • RJ-45 cable
Serial port	<ul style="list-style-type: none"> • DB-9 serial cable • Virtual serial terminal
SOL	<ul style="list-style-type: none"> • Virtual serial terminal

2. If you are using a VGA console, perform these steps:
 - a. Connect a VGA-capable monitor to the VGA port on the back of the appliance.
 - b. View the codes displayed on the monitor.
3. If you are using BMC KVM, perform these steps:
 - a. Connect to the BMC management port and log in to the BMC web interface.
 - b. Select **Remote Control**.
 - c. Launch the KVM.
 - d. View the codes on the virtual monitor.
4. If you are using a serial port and terminal, perform these steps:
 - a. Connect to the serial USB port on the back of the appliance.
 - b. Use settings 115200 8-N-1.
 - c. View the codes printed over the serial terminal.
5. If you are using SOL, perform these steps:
 - a. Connect to the IPMI SOL using the BMC IP address and login credentials.

```
ipmitool -I lanplus -H BMC_Port_IP -U admin -P Password sol activate
```

 - b. View the codes on the virtual serial terminal.
6. Use the table to look up the codes for your appliance.

Code	Indicates
HI	The master boot script has started.
HP	The system is checking to see if the network interface card (NIC) firmware needs to be updated.
RB	The system is rebooting after applying firmware updates.
FP	The hardware subsystem firmware update checks have been completed. Inter-controller communication services are starting.

Code	Indicates
HC	The system is checking for existing StorageGRID installation data.
HO	The StorageGRID appliance is running.
HA	StorageGRID is running.

Related information

[Access BMC interface](#)

View error codes (SGF6112)

If a hardware error occurs when the appliance is booting up, the BMC logs an error code. As required, you can view these error codes using the BMC interface, and then work with technical support to resolve the issue.

Before you begin

- You know how to access the BMC dashboard.

Steps

1. From the BMC dashboard, select **BIOS POST Code**.
2. Review the information displayed for Current Code and the Previous Code.

If any of the following error codes are shown, work with technical support to resolve the issue.

Code	Indicates
0x0E	Microcode not found
0x0F	Microcode not loaded
0x50	Memory initialization error. Invalid memory type or incompatible memory speed.
0x51	Memory initialization error. SPD reading has failed.
0x52	Memory initialization error. Invalid memory size or memory modules don't match.
0x53	Memory initialization error. No usable memory detected.
0x54	Unspecified memory initialization error
0x55	Memory not installed

Code	Indicates
0x56	Invalid CPU type or speed
0x57	CPU mismatch
0x58	CPU self-test failed, or possible CPU cache error
0x59	CPU micro-code is not found, or micro-code update failed
0x5A	Internal CPU error
0x5B	Reset PPI is not available
0x5C	PEI phase BMC self-test failure
0xD0	CPU initialization error
0xD1	North bridge initialization error
0xD2	South bridge initialization error
0xD3	Some architectural protocols aren't available
0xD4	PCI resource allocation error. Out of resources.
0xD5	No space for legacy option ROM
0xD6	No console output devices are found
0xD7	No console input devices are found
0xD8	Invalid password
0xD9	Error loading boot option (LoadImage returned error)
0xDA	Boot option failed (StartImage returned error)
0xDB	Flash update failed
0xDC	Reset protocol is not available
0xDD	DXE phase BMC self-test failure
0xE8	MRC: ERR_NO_MEMORY

Code	Indicates
0xE9	MRC: ERR_LT_LOCK
0xEA	MRC: ERR_DDR_INIT
0xEB	MRC: ERR_MEM_TEST
0xEC	MRC: ERR_VENDOR_SPECIFIC
0xED	MRC: ERR_DIMM_COMPAT
0xEE	MRC: ERR_MRC_COMPATIBILITY
0xEF	MRC: ERR_MRC_STRUCT
0xF0	MRC: ERR_SET_VDD
0xF1	MRC: ERR_IOT_MEM_BUFFER
0xF2	MRC: ERR_RC_INTERNAL
0xF3	MRC: ERR_INVALID_REG_ACCESS
0xF4	MRC: ERR_SET_MC_FREQ
0xF5	MRC: ERR_READ_MC_FREQ
0x70	MRC: ERR_DIMM_CHANNEL
0x74	MRC: ERR_BIST_CHECK
0xF6	MRC: ERR_SMBUS
0xF7	MRC: ERR_PCU
0xF8	MRC: ERR_NGN
0xF9	MRC: ERR_INTERLEAVE_FAILURE

Hardware setup appears to hang (SGF6112)

The StorageGRID Appliance Installer might not be available if hardware faults or cabling errors prevent the appliance from completing its boot-up processing.

Steps

1. Review the LEDs on the appliance and the boot-up and error codes displayed in the BMC.
2. If you need help resolving an issue, contact technical support.

Related information

- [View boot-up codes \(SGF6112\)](#)
- [View error codes \(SGF6112\)](#)

Troubleshoot connection issues (SGF6112)

If you encounter connection issues during the StorageGRID appliance installation, you should perform the corrective action steps listed.

Unable to connect to appliance

If you can't connect to the services appliance, there might be a network issue, or the hardware installation might not have been completed successfully.

Steps

1. Try to ping the appliance using the IP address for the appliance :

`ping appliance_IP`

2. If you receive no response from the ping, confirm you are using the correct IP address.

You can use the IP address of the appliance on the Grid Network, the Admin Network, or the Client Network.

3. If the IP address is correct, check appliance cabling, QSFP or SFP transceivers, and the network setup.
4. If physical access to the appliance is available, you can use a direct connection to the permanent link-local IP `169.254.0.1` to check controller networking configuration and update if necessary. For detailed instructions, see step 2 in [Access StorageGRID Appliance Installer](#).

If that does not resolve the issue, contact technical support.

5. If the ping was successful, open a web browser.
6. Enter the URL for the StorageGRID Appliance Installer:

`https://appliances_controller_IP:8443`

The Home page appears.

Troubleshoot hardware installation (SG6000 or SG5700)

If you encounter issues during the installation, you might find it helpful to review troubleshooting information related to hardware setup and connectivity issues.

[View boot-up codes \(SG6000-CN controller\)](#)

When you apply power to the appliance, the BMC logs a series of boot-up codes for the SG6000-CN controller. You can view these codes in several ways.

Before you begin

- You know how to access the BMC dashboard.
- If you want to use serial-over-LAN (SOL), you have experience using IPMI SOL console applications.

Steps

1. Select one of the following methods for viewing the boot-up codes for the appliance controller, and gather the required equipment.

Method	Required equipment
VGA console	<ul style="list-style-type: none"> • VGA-capable monitor • VGA cable
KVM	<ul style="list-style-type: none"> • RJ-45 cable
Serial port	<ul style="list-style-type: none"> • DB-9 serial cable • Virtual serial terminal
SOL	<ul style="list-style-type: none"> • Virtual serial terminal

2. If you are using a VGA console, perform these steps:
 - a. Connect a VGA-capable monitor to the VGA port on the back of the appliance.
 - b. View the codes displayed on the monitor.
3. If you are using BMC KVM, perform these steps:
 - a. Connect to the BMC management port and log in to the BMC web interface.
 - b. Select **Remote Control**.
 - c. Launch the KVM.
 - d. View the codes on the virtual monitor.
4. If you are using a serial port and terminal, perform these steps:
 - a. Connect to the DB-9 serial port on the back of the appliance.
 - b. Use settings 115200 8-N-1.
 - c. View the codes printed over the serial terminal.
5. If you are using SOL, perform these steps:
 - a. Connect to the IPMI SOL using the BMC IP address and login credentials.

```
ipmitool -I lanplus -H BMC_Port_IP -U admin -P Password sol activate
```

 In some cases the default username might be `root` instead of `admin`.

 - b. View the codes on the virtual serial terminal.
6. Use the table to look up the codes for your appliance.

Code	Indicates
HI	The master boot script has started.
HP	The system is checking to see if the network interface card (NIC) firmware needs to be updated.
RB	The system is rebooting after applying firmware updates.
FP	The hardware subsystem firmware update checks have been completed. Inter-controller communication services are starting.
HE	<p>For an appliance Storage Node only:</p> <p>The system is awaiting connectivity with the storage controllers and synchronizing with the SANtricity operating system.</p> <p>Note: If the boot-up procedure does not progress past this stage, perform these steps:</p> <ol style="list-style-type: none"> <li data-bbox="523 825 1383 889">Confirm that the four interconnect cables between the SG6000-CN controller and the two storage controllers are securely connected. <li data-bbox="523 903 1323 935">As required, replace one or more of the cables, and try again. <li data-bbox="523 952 1295 984">If this does not resolve the issue, contact technical support.
HC	The system is checking for existing StorageGRID installation data.
HO	The StorageGRID Appliance Installer is running.
HA	StorageGRID is running.

View error codes (SG6000-CN controller)

If a hardware error occurs when the SG6000-CN controller is booting up, the BMC logs an error code. As required, you can view these error codes using the BMC interface, and then work with technical support to resolve the issue.

Before you begin

- You know how to access the BMC dashboard.

Steps

1. From the BMC dashboard, select **BIOS POST Code**.
2. Review the information displayed for Current Code and the Previous Code.

If any of the following error codes are shown, work with technical support to resolve the issue.

Code	Indicates
0x0E	Microcode not found
0x0F	Microcode not loaded
0x50	Memory initialization error. Invalid memory type or incompatible memory speed.
0x51	Memory initialization error. SPD reading has failed.
0x52	Memory initialization error. Invalid memory size or memory modules don't match.
0x53	Memory initialization error. No usable memory detected.
0x54	Unspecified memory initialization error
0x55	Memory not installed
0x56	Invalid CPU type or speed
0x57	CPU mismatch
0x58	CPU self-test failed, or possible CPU cache error
0x59	CPU micro-code is not found, or micro-code update failed
0x5A	Internal CPU error
0x5B	Reset PPI is not available
0x5C	PEI phase BMC self-test failure
0xD0	CPU initialization error
0xD1	North bridge initialization error
0xD2	South bridge initialization error
0xD3	Some architectural protocols aren't available
0xD4	PCI resource allocation error. Out of resources.
0xD5	No space for legacy option ROM

Code	Indicates
0xD6	No console output devices are found
0xD7	No console input devices are found
0xD8	Invalid password
0xD9	Error loading boot option (LoadImage returned error)
0xDA	Boot option failed (StartImage returned error)
0xDB	Flash update failed
0xDC	Reset protocol is not available
0xDD	DXE phase BMC self-test failure
0xE8	MRC: ERR_NO_MEMORY
0xE9	MRC: ERR_LT_LOCK
0xEA	MRC: ERR_DDR_INIT
0xEB	MRC: ERR_MEM_TEST
0xEC	MRC: ERR_VENDOR_SPECIFIC
0xED	MRC: ERR_DIMM_COMPAT
0xEE	MRC: ERR_MRC_COMPATIBILITY
0xEF	MRC: ERR_MRC_STRUCT
0xF0	MRC: ERR_SET_VDD
0xF1	MRC: ERR_IOT_MEM_BUFFER
0xF2	MRC: ERR_RC_INTERNAL
0xF3	MRC: ERR_INVALID_REG_ACCESS
0xF4	MRC: ERR_SET_MC_FREQ
0xF5	MRC: ERR_READ_MC_FREQ

Code	Indicates
0x70	MRC: ERR_DIMM_CHANNEL
0x74	MRC: ERR_BIST_CHECK
0xF6	MRC: ERR_SMBUS
0xF7	MRC: ERR_PCU
0xF8	MRC: ERR_NGN
0xF9	MRC: ERR_INTERLEAVE_FAILURE

Hardware setup appears to hang (SG6000 or SG5700)

The StorageGRID Appliance Installer might not be available if hardware faults or cabling errors prevent the storage controllers or the appliance controller from completing their boot-up processing.

Example 2. Steps

SG6000

1. For the storage controllers, watch the codes on the seven-segment displays.

While the hardware is initializing during power up, the two seven-segment displays show a sequence of codes. When the hardware boots successfully, both seven-segment displays show 99.

2. Review the LEDs on the SG6000-CN controller and the boot-up and error codes displayed in the BMC.
3. If you need help resolving an issue, contact technical support.

SG5700

1. Watch the codes on the seven-segment displays.

While the hardware is initializing during power up, the two seven-segment displays show a sequence of codes. When the hardware boots successfully, the seven-segment displays show different codes for each controller.

2. Review the codes on the seven-segment display for the E5700SG controller.



The installation and provisioning take time. Some installation phases don't report updates to the StorageGRID Appliance Installer for several minutes.

If an error occurs, the seven-segment display flashes a sequence, such as HE.

3. To understand what these codes mean, see the following resources:

Controller	Reference
E5700SG controller	<ul style="list-style-type: none">“Status indicators on the E5700SG controller”“HE error: Error synchronizing with SANtricity OS Software”
E2800 controller	<p><i>E5700 and E2800 System Monitoring Guide</i></p> <p>Note: The codes described for the E-Series E5700 controller don't apply to the E5700SG controller in the appliance.</p>

4. If this does not resolve the issue, contact technical support.

Related information

- [View status indicators](#)
- [NetApp E-Series Systems Documentation Site](#)
- [HE error: Error synchronizing with SANtricity OS Software](#)
- [E5700 and E2800 System Monitoring Guide](#)
- [View boot-up codes \(SG6000-CN controller\)](#)
- [View error codes \(SG6000-CN controller\)](#)

Troubleshoot connection issues (SG6000 or SG5700)

If you encounter connection issues during the StorageGRID appliance installation, you should perform the corrective action steps listed.

Unable to connect to SG6000 appliance

If you can't connect to the appliance, there might be a network issue, or the hardware installation might not have been completed successfully.

Steps

1. If you are unable to connect to SANtricity System Manager:
 - a. Try to ping the appliance using the IP address for either storage controller on the management network for SANtricity System Manager:
ping *Storage_Controller_IP*
 - b. If you receive no response from the ping, confirm you are using the correct IP address.

Use the IP address for management port 1 on either storage controller.
 - c. If the IP address is correct, check appliance cabling and the network setup.

If that does not resolve the issue, contact technical support.
 - d. If the ping was successful, open a web browser.
 - e. Enter the URL for SANtricity System Manager:
https://*Storage_Controller_IP*

The log in page for SANtricity System Manager appears.
2. If you are unable to connect to the SG6000-CN controller:
 - a. Try to ping the appliance using the IP address for the SG6000-CN controller:
ping *SG6000-CN_Controller_IP*
 - b. If you receive no response from the ping, confirm you are using the correct IP address.

You can use the IP address of the appliance on the Grid Network, the Admin Network, or the Client Network.
 - c. If the IP address is correct, check appliance cabling, SFP transceivers, and the network setup.
 - d. If physical access to the SG6000-CN is available, you can use a direct connection to the permanent link-local IP 169.254.0.1 to check controller networking configuration and update if necessary. For detailed instructions, see step 2 in [Accessing StorageGRID Appliance Installer](#).

If that does not resolve the issue, contact technical support.
 - e. If the ping was successful, open a web browser.
 - f. Enter the URL for the StorageGRID Appliance Installer:
https://*SG6000-CN_Controller_IP*:8443

The Home page appears.

SG6060 expansion shelves don't appear in Appliance Installer

If you have installed expansion shelves for the SG6060 and they don't appear in the StorageGRID Appliance Installer, you should verify that the shelves have been completely installed and powered on.

About this task

You can verify that the expansion shelves are connected to the appliance by viewing the following information in the StorageGRID Appliance Installer:

- The **Home** page contains a message about expansion shelves.

 The storage system contains 2 expansion shelves.

- The **Advanced > RAID Mode** page indicates by number of drives whether or not the appliance includes expansion shelves. For example, in the following screen shot, two SSDs and 178 HDDs are shown. An SG6060 with two expansion shelves contains 180 total drives.

Configure RAID Mode

This appliance contains the following drives.

Type	Size	Number of drives
SSD	800 GB	2
HDD	11.8 TB	178

If the StorageGRID Appliance Installer pages don't indicate that expansion shelves are present, follow this procedure.

Steps

1. Verify that all required cables have been firmly connected. See [Cable appliance \(SG6000\)](#).
2. Verify that you have powered on the expansion shelves. See [Connect power cords and apply power \(SG6000\)](#).
3. If you need help resolving an issue, contact technical support.

Unable to connect to SG5700 appliance

If you can't connect to the appliance, there might be a network issue, or the hardware installation might not have been completed successfully.

Steps

1. If you are unable to connect to SANtricity System Manager:
 - a. Try to ping the appliance using the IP address for the E2800 controller on the management network for SANtricity System Manager:
ping E2800_Controller_IP
 - b. If you receive no response from the ping, confirm you are using the correct IP address.

Use the IP address for management port 1 on the E2800 controller.
- c. If the IP address is correct, check appliance cabling and the network setup.

If that does not resolve the issue, contact technical support.

- d. If the ping was successful, open a web browser.
- e. Enter the URL for SANtricity System Manager:
`https://E2800_Controller_IP`

The log in page for SANtricity System Manager appears.

2. If you are unable to connect to the E5700SG controller:

- a. Try to ping the appliance using the IP address for the E5700SG controller:
`ping E5700SG_Controller_IP`
- b. If you receive no response from the ping, confirm you are using the correct IP address.

You can use the IP address of the appliance on the Grid Network, the Admin Network, or the Client Network.
- c. If the IP address is correct, check appliance cabling, SFP transceivers, and the network setup.

If that does not resolve the issue, contact technical support.

- d. If the ping was successful, open a web browser.
- e. Enter the URL for the StorageGRID Appliance Installer:
`https://E5700SG_Controller_IP:8443`

The Home page appears.

Related information

[View status indicators](#)

HE error: Error synchronizing with SANtricity OS Software (SG5700)

The seven-segment display on the compute controller shows an HE error code if the StorageGRID Appliance Installer can't synchronize with SANtricity OS Software.

About this task

If an HE error code is displayed, perform this corrective action.

Steps

1. Check the integrity of the two SAS interconnect cables, and confirm they are securely connected.
2. As required, replace one or both of the cables, and try again.
3. If this does not resolve the issue, contact technical support.

Troubleshoot hardware installation (SG100 and SG1000)

If you encounter issues during the installation, you might find it helpful to review troubleshooting information related to hardware setup and connectivity issues.

View boot-up codes (SG100 and SG1000)

When you apply power to the appliance, the BMC logs a series of boot-up codes. You can view these codes on a graphical console that is connected to the BMC management port.

Before you begin

- You know how to access the BMC dashboard.
- If you want to use serial-over-LAN (SOL), you have experience using IPMI SOL console applications.

Steps

1. Select one of the following methods for viewing the boot-up codes for the appliance controller, and gather the required equipment.

Method	Required equipment
VGA console	<ul style="list-style-type: none">• VGA-capable monitor• VGA cable
KVM	<ul style="list-style-type: none">• RJ-45 cable
Serial port	<ul style="list-style-type: none">• DB-9 serial cable• Virtual serial terminal
SOL	<ul style="list-style-type: none">• Virtual serial terminal

2. If you are using a VGA console, perform these steps:
 - a. Connect a VGA-capable monitor to the VGA port on the back of the appliance.
 - b. View the codes displayed on the monitor.
3. If you are using BMC KVM, perform these steps:
 - a. Connect to the BMC management port and log in to the BMC web interface.
 - b. Select **Remote Control**.
 - c. Launch the KVM.
 - d. View the codes on the virtual monitor.
4. If you are using a serial port and terminal, perform these steps:
 - a. Connect to the DB-9 serial port on the back of the appliance.
 - b. Use settings 115200 8-N-1.
 - c. View the codes printed over the serial terminal.
5. If you are using SOL, perform these steps:
 - a. Connect to the IPMI SOL using the BMC IP address and login credentials.

```
ipmitool -I lanplus -H BMC_Port_IP -U admin -P Password sol activate
```



In some cases the default username might be `root` instead of `admin`.

b. View the codes on the virtual serial terminal.

6. Use the table to look up the codes for your appliance.

Code	Indicates
HI	The master boot script has started.
HP	The system is checking to see if the network interface card (NIC) firmware needs to be updated.
RB	The system is rebooting after applying firmware updates.
FP	The hardware subsystem firmware update checks have been completed. Inter-controller communication services are starting.
HC	The system is checking for existing StorageGRID installation data.
HO	The StorageGRID appliance is running.
HA	StorageGRID is running.

Related information

[Access BMC interface](#)

View error codes (SG100 and SG1000)

If a hardware error occurs when the appliance is booting up, the BMC logs an error code. As required, you can view these error codes using the BMC interface, and then work with technical support to resolve the issue.

Before you begin

- You know how to access the BMC dashboard.

Steps

1. From the BMC dashboard, select **BIOS POST Code**.
2. Review the information displayed for Current Code and the Previous Code.

If any of the following error codes are shown, work with technical support to resolve the issue.

Code	Indicates
0x0E	Microcode not found
0x0F	Microcode not loaded

Code	Indicates
0x50	Memory initialization error. Invalid memory type or incompatible memory speed.
0x51	Memory initialization error. SPD reading has failed.
0x52	Memory initialization error. Invalid memory size or memory modules don't match.
0x53	Memory initialization error. No usable memory detected.
0x54	Unspecified memory initialization error
0x55	Memory not installed
0x56	Invalid CPU type or speed
0x57	CPU mismatch
0x58	CPU self-test failed, or possible CPU cache error
0x59	CPU micro-code is not found, or micro-code update failed
0x5A	Internal CPU error
0x5B	Reset PPI is not available
0x5C	PEI phase BMC self-test failure
0xD0	CPU initialization error
0xD1	North bridge initialization error
0xD2	South bridge initialization error
0xD3	Some architectural protocols aren't available
0xD4	PCI resource allocation error. Out of resources.
0xD5	No space for legacy option ROM
0xD6	No console output devices are found
0xD7	No console input devices are found

Code	Indicates
0xD8	Invalid password
0xD9	Error loading boot option (LoadImage returned error)
0xDA	Boot option failed (StartImage returned error)
0xDB	Flash update failed
0xDC	Reset protocol is not available
0xDD	DXE phase BMC self-test failure
0xE8	MRC: ERR_NO_MEMORY
0xE9	MRC: ERR_LT_LOCK
0xEA	MRC: ERR_DDR_INIT
0xEB	MRC: ERR_MEM_TEST
0xEC	MRC: ERR_VENDOR_SPECIFIC
0xED	MRC: ERR_DIMM_COMPAT
0xEE	MRC: ERR_MRC_COMPATIBILITY
0xEF	MRC: ERR_MRC_STRUCT
0xF0	MRC: ERR_SET_VDD
0xF1	MRC: ERR_IOT_MEM_BUFFER
0xF2	MRC: ERR_RC_INTERNAL
0xF3	MRC: ERR_INVALID_REG_ACCESS
0xF4	MRC: ERR_SET_MC_FREQ
0xF5	MRC: ERR_READ_MC_FREQ
0x70	MRC: ERR_DIMM_CHANNEL
0x74	MRC: ERR_BIST_CHECK

Code	Indicates
0xF6	MRC: ERR_SMBUS
0xF7	MRC: ERR_PCU
0xF8	MRC: ERR_NGN
0xF9	MRC: ERR_INTERLEAVE_FAILURE

Hardware setup appears to hang (SG100 and SG1000)

The StorageGRID Appliance Installer might not be available if hardware faults or cabling errors prevent the appliance from completing its boot-up processing.

Steps

1. Review the LEDs on the appliance and the boot-up and error codes displayed in the BMC.
2. If you need help resolving an issue, contact technical support.

Related information

- [View boot-up codes \(SG100 and SG1000\)](#)
- [View error codes \(SG100 and SG1000\)](#)

Troubleshoot connection issues (SG100 and SG1000)

If you encounter connection issues during the StorageGRID appliance installation, you should perform the corrective action steps listed.

Unable to connect to appliance

If you can't connect to the services appliance, there might be a network issue, or the hardware installation might not have been completed successfully.

Steps

1. Try to ping the appliance using the IP address for the appliance :
ping services_appliance_IP
2. If you receive no response from the ping, confirm you are using the correct IP address.

You can use the IP address of the appliance on the Grid Network, the Admin Network, or the Client Network.

3. If the IP address is correct, check appliance cabling, QSFP or SFP transceivers, and the network setup.
4. If physical access to the appliance is available, you can use a direct connection to the permanent link-local IP 169.254.0.1 to check controller networking configuration and update if necessary. For detailed instructions, see step 2 in [Access StorageGRID Appliance Installer](#).

If that does not resolve the issue, contact technical support.

5. If the ping was successful, open a web browser.
6. Enter the URL for the StorageGRID Appliance Installer:
https://appliances_controller_IP:8443

The Home page appears.

Install Red Hat Enterprise Linux or CentOS

Install Red Hat Enterprise Linux or CentOS: Overview

Installing a StorageGRID system in a Red Hat Enterprise Linux (RHEL) or CentOS Linux environment includes three primary steps.

1. **Preparation:** During planning and preparation, you perform the following tasks:
 - Learn about the hardware and storage requirements for StorageGRID.
 - Learn about the specifics of [StorageGRID networking](#) so you can configure your network appropriately.
 - Identify and prepare the physical or virtual servers you plan to use to host your StorageGRID grid nodes.
 - On the servers you have prepared:
 - Install Linux
 - Configure the host network
 - Configure host storage
 - Install container engine
 - Install the StorageGRID host services
2. **Deployment:** Deploy grid nodes using the appropriate user interface. When you deploy grid nodes, they are created as part of the StorageGRID system and connected to one or more networks.
 - a. Use the Linux command line and node configuration files to deploy software-based grid nodes on the hosts you prepared in step 1.
 - b. Use the StorageGRID Appliance Installer to deploy StorageGRID appliance nodes.



Hardware-specific installation and integration instructions aren't included in the StorageGRID installation procedure. To learn how to install StorageGRID appliances, see the [Quick start for hardware installation](#) to locate instructions for your appliance.

3. **Configuration:** When all nodes have been deployed, use the Grid Manager to configure the grid and complete the installation.

These instructions recommend a standard approach for deploying and configuring a StorageGRID system. See also the information about the following alternative approaches:

- Use a standard orchestration framework such as Ansible, Puppet, or Chef to install RHEL or CentOS, configure networking and storage, install the container engine and the StorageGRID host service, and deploy virtual grid nodes.
- Automate the deployment and configuration of the StorageGRID system using a Python configuration script (provided in the installation archive).

- Automate the deployment and configuration of appliance grid nodes with a Python configuration script (available from the installation archive or from the StorageGRID Appliance Installer).
- If you are an advanced developer of StorageGRID deployments, use the installation REST APIs to automate the installation of StorageGRID grid nodes.

Plan and prepare for Red Hat or CentOS installation

Before you install (Red Hat or CentOS)

Before deploying grid nodes and configuring StorageGRID, you must be familiar with the steps and requirements for completing the procedure.

The StorageGRID deployment and configuration procedures assume that you are familiar with the architecture and operation of the StorageGRID system.

You can deploy a single site or multiple sites at one time; however, all sites must meet the minimum requirement of having at least three Storage Nodes.

Before starting a StorageGRID installation, you must:

- Understand the compute requirements, including the minimum CPU and RAM requirements for each node.
- Understand how StorageGRID supports multiple networks for traffic separation, security, and administrative convenience, and have a plan for which networks you intend to attach to each StorageGRID node.

See the [StorageGRID Networking guidelines](#).

- Understand the storage and performance requirements of each type of grid node.
- 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.
- Understand the [requirements for node migration](#), if you want to perform scheduled maintenance on physical hosts without any service interruption.
- Gather all networking information in advance. Unless you are using DHCP, gather the IP addresses to assign to each grid node, and the IP addresses of the DNS and NTP servers that will be used.
- Install, connect, and configure all required hardware, including any StorageGRID appliances, to specifications.



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.



Hardware-specific installation and integration instructions aren't included in the StorageGRID installation procedure. To learn how to install StorageGRID appliances, see [Install appliance hardware](#).

- Decide which of the available deployment and configuration tools you want to use.

Required materials

Before you install StorageGRID, you must gather and prepare required materials.

Item	Notes
NetApp StorageGRID license	<p>You must have a valid, digitally signed NetApp license.</p> <p>Note: A non-production license, which can be used for testing and proof of concept grids, is included in the StorageGRID installation archive.</p>
StorageGRID installation archive	<p>You must download the StorageGRID installation archive and extract the files.</p>
Service laptop	<p>The StorageGRID system is installed through a service laptop.</p> <p>The service laptop must have:</p> <ul style="list-style-type: none"> • Network port • SSH client (for example, PuTTY) • Supported web browser
StorageGRID documentation	<ul style="list-style-type: none"> • Release notes • Instructions for administering StorageGRID

Related information

[NetApp Interoperability Matrix Tool](#)

Download and extract the StorageGRID installation files

You must download the StorageGRID installation archive and extract the required files.

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 file for Red Hat Enterprise Linux or CentOS.
7. Save and extract the archive file.
8. Choose the files you need from the following list.

The files you need depend 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

Path and file name	Description
./rpms/README	A text file that describes all of the files contained in the StorageGRID download file.
./rpms/NLF000000.txt	A free license that does not provide any support entitlement for the product.
./rpms/StorageGRID-Webscale-Images-version-SHA.rpm	RPM package for installing the StorageGRID node images on your RHEL or CentOS hosts.
./rpms/StorageGRID-Webscale-Service-version-SHA.rpm	RPM package for installing the StorageGRID host service on your RHEL or CentOS hosts.
Deployment scripting tool	Description
./rpms/configure-storagegrid.py	A Python script used to automate the configuration of a StorageGRID system.
./rpms/configure-sga.py	A Python script used to automate the configuration of StorageGRID appliances.
./rpms/configure-storagegrid.sample.json	An example configuration file for use with the configure-storagegrid.py script.
./rpms/storagegrid-ssoauth.py	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.
./rpms/configure-storagegrid.blank.json	A blank configuration file for use with the configure-storagegrid.py script.
./rpms/extras/ansible	Example Ansible role and playbook for configuring RHEL or CentOS hosts for StorageGRID container deployment. You can customize the role or playbook as necessary.
./rpms/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.

Path and file name	Description
./rpms/storagegrid-ssoauth-azure.js	A helper script called by the companion <code>storagegrid-ssoauth-azure.py</code> Python script to perform SSO interactions with Azure.
./rpms/extras/api-schemas	<p>API schemas for StorageGRID.</p> <p>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.</p>

CPU and RAM requirements

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

For information about supported servers, see the [NetApp Interoperability Matrix Tool](#).

Each StorageGRID node requires the following minimum resources:

- CPU cores: 8 per node
- RAM: At least 24 GB per node, and 2 to 16 GB less than the total system RAM, depending on the total RAM available and the amount of non-StorageGRID software running on the system

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.

 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.

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.

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.

Number of storage volumes for each host

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	90 GB
Storage Node	Object data	3 for each Storage Node on this host Note: A software-based Storage Node can have 1 to 16 storage volumes; at least 3 storage volumes are recommended.	12 TB (4 TB/LUN) See Storage requirements for Storage Nodes for more information.
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.

Type of node	Container pool	System data	Object data
Storage Node	100 GB	90 GB	4,000 GB

Type of node	Container pool	System data	Object data
Admin Node	100 GB	490 GB (3 LUNs)	<i>not applicable</i>
Gateway Node	100 GB	90 GB	<i>not applicable</i>
Archive Node	100 GB	90 GB	<i>not applicable</i>

Example: Calculating the storage requirements for a host

Suppose you plan to deploy three nodes on the same host: 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, 670 GB of performance-tier storage for system data and transaction logs, and 12 TB of capacity-tier storage for object data.

Type of node	LUN purpose	Number of LUNs	LUN size
Storage Node	Container engine storage pool	1	300 GB (100 GB/node)
Storage Node	/var/local volume	1	90 GB
Storage Node	Object data	3	12 TB (4 TB/LUN)
Admin Node	/var/local volume	1	90 GB
Admin Node	Admin Node audit logs	1	200 GB
Admin Node	Admin Node tables	1	200 GB
Gateway Node	/var/local volume	1	90 GB
Total		9	Container pool: 300 GB System data: 670 GB Object data: 12,000 GB


Storage requirements for Storage Nodes

A software-based Storage Node can have 1 to 16 storage volumes—3 or more storage volumes are recommended. Each storage volume should be 4 TB or larger.



An appliance Storage Node can 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 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.

- 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

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.

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

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 [Configuring](#)

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:

1. 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.
2. 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`

Deployment tools

You might benefit from automating all or part of the StorageGRID installation.

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 that can be created interactively during a manual installation, or prepared ahead of time (or programmatically) to enable automated installation using standard orchestration frameworks. StorageGRID provides optional Python scripts for automating the configuration of StorageGRID appliances, and the whole 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.

If you are interested in automating all or part of your StorageGRID deployment, review [Automate the installation](#) before beginning the installation process.

Prepare the hosts (Red Hat or CentOS)

How host-wide settings change during installation

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

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 Linux on all grid hosts. Use the [NetApp Interoperability Matrix Tool](#) to get a list of supported versions.



Ensure that your operating system is upgraded to Linux kernel 4.15 or higher.

Steps

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



If you are using the standard Linux installer, NetApp recommends selecting the “compute node” software configuration, if available, or “minimal install” base environment. Don’t install any graphical desktop environments.

2. Ensure that all hosts have access to package repositories, including the Extras channel.

You might need these additional packages later in this installation procedure.

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.

Configure the host network (Red Hat Enterprise Linux or CentOS)

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.

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.

Related information

[Creating node configuration files](#)

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 shows 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


About this task

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

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/sysconfig/network-scripts`

Configure host storage

You must allocate block storage volumes to each 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

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.

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-adm1-var-local
    }
    multipath {
        wwid 3600a09800059d6df00005df4573c2c30
        alias sgws-adm1-audit-logs
    }
    multipath {
        wwid 3600a09800059d6df00005df5573c2c30
        alias sgws-adm1-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
    }
    ...
}
```

This will cause 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.

Related information

[Configure container engine storage volume](#)

Configure container engine storage volume

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

About this task

You can skip these steps if you plan to use local storage for the Docker or Podman storage volume and have sufficient space available on the host partition containing `/var/lib/docker` for Docker and `/var/lib/containers` for Podman.



Podman is supported only on Red Hat Enterprise Linux (RHEL).

Steps

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

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

2. Mount the container engine storage volume:

- 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
```

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

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

Install Docker

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

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 on Red Hat Enterprise Linux as a collection of containers. If you have chosen to use the Podman container engine, follow these steps to install Podman. Otherwise, [install Docker](#).



Podman is supported only on Red Hat Enterprise Linux (RHEL).

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
```

Install StorageGRID host services

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 Yum repository metadata included in the installation archive to install the RPM packages remotely. See the Yum 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 yum --nogpgcheck localinstall /tmp/StorageGRID-Webscale-Images-  
version-SHA.rpm
```

```
sudo yum --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.

Deploy virtual grid nodes (Red Hat or CentOS)

Create node configuration files for Red Hat Enterprise Linux or CentOS 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.

Where do I put the node configuration files?

You must 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.

What do I name the 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-nodenumber.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.

What is in a node configuration file?

The configuration files contain key/value pairs, with one key and one value per line. For each key/value pair, you must 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.

Some keys are required for every node, while others are optional or only required for certain node types.

The table defines the acceptable values for all supported keys. In the middle column:

R: required

BP: best practice

O: optional

Key	R, BP, or O?	Value
ADMIN_IP	BP	<p>Grid Network IPv4 address of the primary Admin Node for the grid to which this node belongs. Use the same value you specified for GRID_NETWORK_IP for the grid node with NODE_TYPE = VM_Admin_Node and ADMIN_ROLE = Primary. 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>
ADMIN_NETWORK_CONFIG	O	DHCP, STATIC, or DISABLED
ADMIN_NETWORK_ESL	O	<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>
ADMIN_NETWORK_GATEWAY	O (R)	<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>Note: This parameter is required if ADMIN_NETWORK_ESL is specified.</p> <p>Examples:</p> <p>1.1.1.1</p> <p>10.224.4.81</p>
ADMIN_NETWORK_IP	O	<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>

Key	R, BP, or O?	Value
ADMIN_NETWORK_MAC	O	<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>
ADMIN_NETWORK_MASK	O	<p>IPv4 netmask for 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>255.255.255.0</p> <p>255.255.248.0</p>
ADMIN_NETWORK_MTU	O	<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>

Key	R, BP, or O?	Value
ADMIN_NETWORK_TARGET	BP	<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>
ADMIN_NETWORK_TARGET_TYPE	O	<p>Interface</p> <p>(This is the only supported value.)</p>
ADMIN_NETWORK_TARGET_TYPE_INTERFACE_CLONE_MAC	BP	<p>True or False</p> <p>Set the key to "true" to cause the StorageGRID container use the MAC address of the 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:</p> <p>Considerations and recommendations for MAC address cloning (Red Hat Enterprise Linux or CentOS)</p> <p>Considerations and recommendations for MAC address cloning (Ubuntu or Debian)</p>
ADMIN_ROLE	R	<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>

Key	R, BP, or O?	Value
BLOCK_DEVICE_AUDIT_LOGS	R	<p>Path and name of the block device special file this node will use for persistent storage of audit logs. 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 /dev/disk/by-id/wwn-0x600a09800059d6df000060d757b475fd /dev/mapper/sgws-adm1-audit-logs</pre>

Key	R, BP, or O?	Value
BLOCK_DEVICE_RANGEDB_000	R	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 NODE_TYPE = VM_Storage_Node; don't specify it for other node types.
BLOCK_DEVICE_RANGEDB_001		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.
BLOCK_DEVICE_RANGEDB_002		
BLOCK_DEVICE_RANGEDB_003		
BLOCK_DEVICE_RANGEDB_004		
BLOCK_DEVICE_RANGEDB_005		
BLOCK_DEVICE_RANGEDB_006		
BLOCK_DEVICE_RANGEDB_007		
BLOCK_DEVICE_RANGEDB_008		
BLOCK_DEVICE_RANGEDB_009		
BLOCK_DEVICE_RANGEDB_010		
BLOCK_DEVICE_RANGEDB_011		
BLOCK_DEVICE_RANGEDB_012		
BLOCK_DEVICE_RANGEDB_013		
BLOCK_DEVICE_RANGEDB_014		
BLOCK_DEVICE_RANGEDB_015		

Key	R, BP, or O?	Value
BLOCK_DEVICE_TABLES	R	<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 /dev/disk/by-id/wwn-0x600a09800059d6df000060d757b475fd /dev/mapper/sgws-adm1-tables</pre>
BLOCK_DEVICE_VAR_LOCAL	R	<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 /dev/disk/by-id/wwn-0x600a09800059d6df000060d757b475fd /dev/mapper/sgws-sn1-var-local</pre>
CLIENT_NETWORK_CONFIG	O	DHCP, STATIC, or DISABLED
CLIENT_NETWORK_GATEWAY	O	<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> <pre>1.1.1.1 10.224.4.81</pre>

Key	R, BP, or O?	Value
CLIENT_NETWORK_IP	O	<p>IPv4 address of this node on the Client Network. 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>
CLIENT_NETWORK_MAC	O	<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>
CLIENT_NETWORK_MASK	O	<p>IPv4 netmask for this node on the Client Network. This key is only required 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>
CLIENT_NETWORK_MTU	O	<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>

Key	R, BP, or O?	Value
CLIENT_NETWORK_TARGET	BP	<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> <p>bond0.1003</p> <p>ens423</p>
CLIENT_NETWORK_TARGET_TYPE	O	<p>Interface</p> <p>(This is only supported value.)</p>
CLIENT_NETWORK_TARGET_TYPE_INTERFACE_CLONE_MAC	BP	<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:</p> <p>Considerations and recommendations for MAC address cloning (Red Hat Enterprise Linux or CentOS)</p> <p>Considerations and recommendations for MAC address cloning (Ubuntu or Debian)</p>
GRID_NETWORK_CONFIG	BP	<p>STATIC or DHCP</p> <p>(Defaults to STATIC if not specified.)</p>

Key	R, BP, or O?	Value
GRID_NETWORK_GATEWAY	R	<p>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.</p> <p>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.</p>
GRID_NETWORK_IP	R	<p>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.</p> <p>Examples:</p> <p>1.1.1.1</p> <p>10.224.4.81</p>
GRID_NETWORK_MAC	O	<p>The MAC address for the Grid 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:c2:27:30</p>
GRID_NETWORK_MASK	O	<p>IPv4 netmask for this node on the Grid Network. This key is only required 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>

Key	R, BP, or O?	Value
GRID_NETWORK_MTU	O	<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 8192</p>
GRID_NETWORK_TARGET	R	<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> <p>bond0.1001</p> <p>ens192</p>
GRID_NETWORK_TARGET_TYPE	O	<p>Interface</p> <p>(This is the only supported value.)</p>

Key	R, BP, or O?	Value
GRID_NETWORK_TARGET_TYPE_INTERFACE_CLONE_MAC	BP	<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:</p> <p>Considerations and recommendations for MAC address cloning (Red Hat Enterprise Linux or CentOS)</p> <p>Considerations and recommendations for MAC address cloning (Ubuntu or Debian)</p>
INTERFACE_TARGET_nnnn	O	<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>For example: INTERFACE_TARGET_0001=ens256, Trunk</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>

Key	R, BP, or O?	Value
MAXIMUM_RAM	O	<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 <number><unit>, where <unit> 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>
NODE_TYPE	R	<p>Type of node:</p> <p>VM_Admin_Node VM_Storage_Node VM_Archive_Node VM_API_Gateway</p>
PORT_REMAP	O	<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 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.</p> <p>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.</p> <p>For example:</p> <pre>PORT_REMAP = client/tcp/18082/443</pre>

Key	R, BP, or O?	Value
PORT_REMAP_INBOUND	O	<p>Remaps inbound communications to 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.</p> <p>IMPORTANT: Don't remap the ports you are planning to use to configure load balancer endpoints.</p> <p>The format used is: <network type>/<protocol:>/<remapped port >/<default port used by grid node>, where <network type> is grid, admin, or client, and protocol is tcp or udp.</p> <p>For example:</p> <pre>PORT_REMAP_INBOUND = grid/tcp/3022/22</pre>

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.

Example node configuration files

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.

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/dc1-adm1.conf

Example file contents:

```
NODE_TYPE = VM_Admin_Node
ADMIN_ROLE = Primary
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 Archive Node

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

Example file contents:

```

NODE_TYPE = VM_Archive_Node
ADMIN_IP = 10.1.0.2
BLOCK_DEVICE_VAR_LOCAL = /dev/mapper/dc1-arcl-var-local
GRID_NETWORK_TARGET = bond0.1001
ADMIN_NETWORK_TARGET = bond0.1002
CLIENT_NETWORK_TARGET = bond0.1003

GRID_NETWORK_IP = 10.1.0.4
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

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

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.

```
Checking for misnamed node configuration files... PASSED
Checking configuration file for node dc1-adm1... PASSED
Checking configuration file for node dc1-gw1... PASSED
Checking configuration file for node dc1-sn1... PASSED
Checking configuration file for node dc1-sn2... PASSED
Checking configuration file for node dc1-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-adm1
  WARNING: ignoring /etc/storagegrid/nodes/dc1-sn2.conf.keep
  WARNING: ignoring /etc/storagegrid/nodes/my-file.txt
Checking configuration file for node dc1-adm1...
  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

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

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
```

Configure the grid and complete installation (Red Hat or CentOS)

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 one of the following addresses:

`https://primary_admin_node_ip`

`client_network_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.

2. Select **Install a StorageGRID system**.

The page used to configure a StorageGRID system appears.

NetApp® StorageGRID®

Help ▾

Install



License

Enter a grid name and upload the license file provided by NetApp for your StorageGRID system.

Grid Name

License File

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.

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

License

Enter a grid name and upload the license file provided by NetApp for your StorageGRID system.

Grid Name: StorageGRID

License File: NLF-959007-Internal.txt

License Serial Number: 959007

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.



NetApp® StorageGRID®

Help ▾

Install

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

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	Raleigh	×
Site Name 2	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.

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.



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.

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.

						Search
	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.

						Search
	Grid Network MAC Address	Name	Site	Type	Platform	Grid Network IPv4 Address
<input checked="" type="radio"/>	00:50:56:87:42:ff	dc1-adm1	Raleigh	Admin Node	VMware VM	172.16.4.210/21
<input checked="" type="radio"/>	00:50:56:87:c0:16	dc1-s1	Raleigh	Storage Node	VMware VM	172.16.4.211/21
<input checked="" type="radio"/>	00:50:56:87:79:ee	dc1-s2	Raleigh	Storage Node	VMware VM	172.16.4.212/21
<input checked="" type="radio"/>	00:50:56:87:db:9c	dc1-s3	Raleigh	Storage Node	VMware VM	172.16.4.213/21
<input checked="" 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:

Storage Node Configuration

General Settings

Site	<input type="text" value="Raleigh"/>
Name	<input type="text" value="NetApp-SGA"/>
NTP Role	<input type="text" value="Automatic"/>
ADC Service	<input type="text" value="Automatic"/>

Grid Network

Configuration	STATIC
IPv4 Address (CIDR)	<input type="text" value="172.16.5.20/21"/>
Gateway	<input type="text" value="172.16.5.20"/>

Admin Network

Configuration	STATIC
IPv4 Address (CIDR)	<input type="text" value="10.224.5.20/21"/>
Gateway	<input type="text" value="10.224.0.1"/>
Subnets (CIDR)	<input type="text" value="10.0.0.0/8"/> ×
	<input type="text" value="172.19.0.0/16"/> ×
	<input type="text" value="172.21.0.0/16"/> + ×

Client Network

Configuration	STATIC
IPv4 Address (CIDR)	<input type="text" value="47.47.5.20/21"/>
Gateway	<input type="text" value="47.47.0.1"/>

- **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.

- **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.

- **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 resulting 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 resulting 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 installation instructions for your appliance model.

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 resulting 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 installation 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	Search
	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	Search
	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.

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.

NetApp® StorageGRID®

Help ▾

Install

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

Network Time Protocol

Enter the IP addresses for at least four Network Time Protocol (NTP) servers, so that operations performed on separate servers are kept in sync.

Server 1	10.60.248.183
Server 2	10.227.204.142
Server 3	10.235.48.111
Server 4	0.0.0.0

+

3. Select **Next**.

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.

If the DNS server information is omitted or incorrectly configured, a DNST alarm is triggered on each grid node's SSM service. The alarm clears when DNS is configured correctly and the new server information has reached all grid nodes.

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.

NetApp® StorageGRID®

Help ▾

Install

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

Domain Name Service

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.

Server 1 10.224.223.130 X

Server 2 10.224.223.136 + X

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 password in a secure place.

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

NetApp® StorageGRID®

Help ▾

Install

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

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
Confirm Provisioning Passphrase
Grid Management Root User Password
Confirm Root User Password

Create random command line passwords.

5. 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.

6. 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

1. 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	

2. 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 style="width: 50%;"></div></div>	Starting services
dc1-q1	Site1	172.16.4.216/21	<div style="width: 100%;"><div style="width: 100%;"></div></div>	Complete
dc1-s1	Site1	172.16.4.217/21	<div style="width: 100%;"><div style="width: 50%;"></div></div>	Waiting for Dynamic IP Service peers
dc1-s2	Site1	172.16.4.218/21	<div style="width: 100%;"><div style="width: 10%;"></div></div>	Downloading hotfix from primary Admin if needed
dc1-s3	Site1	172.16.4.219/21	<div style="width: 100%;"><div style="width: 10%;"></div></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 their IP addresses change, which can cause outages if a DHCP address 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.

Automate the installation (Red Hat Enterprise Linux or CentOS)

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

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 described 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 an RPM 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 RHEL or CentOS, adding StorageGRID to your playbooks or recipes should be straightforward.

See the example Ansible role and playbook in the `/extras` folder supplied with the installation archive. The Ansible playbook shows how the `storagegrid` role prepares the host 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.

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

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
configure-storagegrid.py	Python script used to automate the configuration
configure-storagegrid.sample.json	Sample configuration file for use with the script
configure-storagegrid.blank.json	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 sample configuration file (`configure-storagegrid.sample.json`) or the blank configuration file (`configure-storagegrid.blank.json`).

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

1. Log in to the Linux machine you are using to run the Python script.
2. 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 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

[Overview of the installation REST API](#)

Overview of the 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 Docs 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 Recovery Package, start the recovery, and view the status of the recovery operation.
- **recovery-package** — Operations to download the Recovery Package.
- **schemas** — API schemas for advanced deployments
- **sites** — Site-level configuration operations. You can create, view, delete, and modify a site.

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

- [Create a tenant account](#) for each client protocol (Swift or S3) 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](#) or [Swift 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](#).
- If your StorageGRID system includes any Archive Nodes (deprecated), configure the Archive Node's connection to the target external archival storage system.

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 access to the system for auditing purposes](#) through an NFS file share.

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.

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](#).

Related information

[Troubleshoot a StorageGRID system](#)

Example `/etc/sysconfig/network-scripts`

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
```

Install Ubuntu or Debian

Install Ubuntu or Debian: Overview

Installing a StorageGRID system in an Ubuntu or Debian environment includes three primary steps.

1. **Preparation:** During planning and preparation, you perform the following tasks:

- Learn about the hardware and storage requirements for StorageGRID.
- Learn about the specifics of [StorageGRID networking](#) so you can configure your network appropriately.
- Identify and prepare the physical or virtual servers you plan to use to host your StorageGRID grid nodes.
- On the servers you have prepared:
 - Install Linux
 - Configure the host network
 - Configure host storage
 - Install Docker
 - Install the StorageGRID host services

2. **Deployment:** Deploy grid nodes using the appropriate user interface. When you deploy grid nodes, they are created as part of the StorageGRID system and connected to one or more networks.

- a. Use the Linux command line and node configuration files to deploy virtual grid nodes on the hosts you prepared in step 1.
- b. Use the StorageGRID Appliance Installer to deploy StorageGRID appliance nodes.



Hardware-specific installation and integration instructions aren't included in the StorageGRID installation procedure. To learn how to install StorageGRID appliances, see the [Quick start for hardware installation](#) to locate instructions for your appliance.

3. **Configuration:** When all nodes have been deployed, use the Grid Manager to configure the grid and complete the installation.

These instructions recommend a standard approach for deploying and configuring a StorageGRID system in an Ubuntu or Debian environment. See also the information about the following alternative approaches:

- Use a standard orchestration framework such as Ansible, Puppet, or Chef to install Ubuntu or Debian, configure networking and storage, install Docker and the StorageGRID host service, and deploy virtual grid nodes.
- Automate the deployment and configuration of the StorageGRID system using a Python configuration script (provided in the installation archive).
- Automate the deployment and configuration of appliance grid nodes with a Python configuration script (available from the installation archive or from the StorageGRID Appliance Installer).
- If you are an advanced developer of StorageGRID deployments, use the installation REST APIs to automate the installation of StorageGRID grid nodes.

Plan and prepare for Ubuntu or Debian installation

Before you install (Ubuntu or Debian)

Before deploying grid nodes and configuring StorageGRID, you must be familiar with the steps and requirements for completing the procedure.

The StorageGRID deployment and configuration procedures assume that you are familiar with the architecture and operation of the StorageGRID system.

You can deploy a single site or multiple sites at one time; however, all sites must meet the minimum

requirement of having at least three Storage Nodes.

Before starting a StorageGRID installation, you must:

- Understand the compute requirements, including the minimum CPU and RAM requirements for each node.
- Understand how StorageGRID supports multiple networks for traffic separation, security, and administrative convenience, and have a plan for which networks you intend to attach to each StorageGRID node.

See the [StorageGRID Networking guidelines](#).

- Understand the storage and performance requirements of each type of grid node.
- 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.
- Understand the [requirements for node migration](#), if you want to perform scheduled maintenance on physical hosts without any service interruption.
- Gather all networking information in advance. Unless you are using DHCP, gather the IP addresses to assign to each grid node, and the IP addresses of the DNS and NTP servers that will be used.
- Install, connect, and configure all required hardware, including any StorageGRID appliances, to specifications.



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.



Hardware-specific installation and integration instructions aren't included in the StorageGRID installation procedure. To learn how to install StorageGRID appliances, see [Install appliance hardware](#).

- Decide which of the available deployment and configuration tools you want to use.

Required materials

Before you install StorageGRID, you must gather and prepare required materials.

Item	Notes
NetApp StorageGRID license	You must have a valid, digitally signed NetApp license. Note: 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	You must download the StorageGRID installation archive and extract the files.

Item	Notes
Service laptop	<p>The StorageGRID system is installed through a service laptop.</p> <p>The service laptop must have:</p> <ul style="list-style-type: none"> • Network port • SSH client (for example, PuTTY) • Supported web browser
StorageGRID documentation	<ul style="list-style-type: none"> • Release notes • Instructions for administering StorageGRID

Related information

[NetApp Interoperability Matrix Tool](#)

Download and extract the StorageGRID installation files

You must download the StorageGRID installation archive and extract the required files.

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**.

The downloads page for the version you selected appears. The page contains three columns:

6. In the **Install StorageGRID** column, select the .tgz or .zip file for Ubuntu or Debian.



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

7. Save and extract the archive file.
8. Choose the files you need from the following list.

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



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

Path and file name	Description
./debs/README	A text file that describes all of the files contained in the StorageGRID download file.
./debs/NLF000000.txt	A non-production NetApp License File that you can use for testing and proof of concept deployments.
./debs/storagegrid-webscale-images-version-SHA.deb	DEB package for installing the StorageGRID node images on Ubuntu or Debian hosts.
./debs/storagegrid-webscale-images-version-SHA.deb.md5	MD5 checksum for the file /debs/storagegrid-webscale-images-version-SHA.deb.
./debs/storagegrid-webscale-service-version-SHA.deb	DEB package for installing the StorageGRID host service on Ubuntu or Debian hosts.
Deployment scripting tool	Description
./debs/configure-storagegrid.py	A Python script used to automate the configuration of a StorageGRID system.
./debs/configure-sga.py	A Python script used to automate the configuration of StorageGRID appliances.
./debs/storagegrid-ssoauth.py	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.
./debs/configure-storagegrid.sample.json	An example configuration file for use with the configure-storagegrid.py script.
./debs/configure-storagegrid.blank.json	A blank configuration file for use with the configure-storagegrid.py script.
./debs/extras/ansible	Example Ansible role and playbook for configuring Ubuntu or Debian hosts for StorageGRID container deployment. You can customize the role or playbook as necessary.
./debs/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.
./debs/storagegrid-ssoauth-azure.js	A helper script called by the companion storagegrid-ssoauth-azure.py Python script to perform SSO interactions with Azure.

Path and file name	Description
./debs/extras/api-schemas	<p>API schemas for StorageGRID.</p> <p>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.</p>

CPU and RAM requirements

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

For information about supported servers, see the [NetApp Interoperability Matrix Tool](#).

Each StorageGRID node requires the following minimum resources:

- CPU cores: 8 per node
- RAM: At least 24 GB per node, and 2 to 16 GB less than the total system RAM, depending on the total RAM available and the amount of non-StorageGRID software running on the system

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.

 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.

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.

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 Docker storage driver when you install and configure Docker 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.

Number of storage volumes for each host

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	90 GB
Storage Node	Object data	3 for each Storage Node on this host Note: A software-based Storage Node can have 1 to 16 storage volumes; at least 3 storage volumes are recommended.	12 TB (4 TB/LUN) See storage requirements for Storage Nodes for more information.
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.

Type of node	Container pool	System data	Object data
Storage Node	100 GB	90 GB	4,000 GB

Type of node	Container pool	System data	Object data
Admin Node	100 GB	490 GB (3 LUNs)	<i>not applicable</i>
Gateway Node	100 GB	90 GB	<i>not applicable</i>
Archive Node	100 GB	90 GB	<i>not applicable</i>

Example: Calculating the storage requirements for a host

Suppose you plan to deploy three nodes on the same host: 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, 670 GB of performance-tier storage for system data and transaction logs, and 12 TB of capacity-tier storage for object data.

Type of node	LUN purpose	Number of LUNs	LUN size
Storage Node	Docker storage pool	1	300 GB (100 GB/node)
Storage Node	/var/local volume	1	90 GB
Storage Node	Object data	3	12 TB (4 TB/LUN)
Admin Node	/var/local volume	1	90 GB
Admin Node	Admin Node audit logs	1	200 GB
Admin Node	Admin Node tables	1	200 GB
Gateway Node	/var/local volume	1	90 GB
Total		9	Container pool: 300 GB System data: 670 GB Object data: 12,000 GB


Storage requirements for Storage Nodes

A software-based Storage Node can have 1 to 16 storage volumes—3 or more storage volumes are recommended. Each storage volume should be 4 TB or larger.



An appliance Storage Node can 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 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.

- 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

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.

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

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:

Steps

1. 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.
2. 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`

Deployment tools

You might benefit from automating all or part of the StorageGRID installation.

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 that can be created interactively during a manual installation, or prepared ahead of time (or programmatically) to enable automated installation using standard orchestration frameworks. StorageGRID provides optional Python scripts for automating the configuration of StorageGRID appliances, and the whole 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.

If you are interested in automating all or part of your StorageGRID deployment, review [Automate the installation](#) before beginning the installation process.

Prepare the hosts (Ubuntu or Debian)

How host-wide settings change during installation

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

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 Linux on all grid hosts. Use the [NetApp Interoperability Matrix Tool \(IMT\)](#) to get a list of supported versions.



Ensure that your operating system is upgraded to Linux kernel 4.15 or higher.

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. 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 Ubuntu or Debian package repositories.

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

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.

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

Configure the host network (Ubuntu or Debian)

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.

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 shows 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 $10G_1$ through $10G_3$ 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](#)

Configure host storage

You must allocate block storage volumes to each 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](#)

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.

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-adm1-var-local
    }
    multipath {
        wwid 3600a09800059d6df00005df4573c2c30
        alias sgws-adm1-audit-logs
    }
    multipath {
        wwid 3600a09800059d6df00005df5573c2c30
        alias sgws-adm1-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
    }
    ...
}

```

This will cause 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 Docker storage volume on each host. Using aliases and including the target hostname in the alias for each Docker storage volume LUN will make this easy to remember and is recommended.

Related information

[Storage and performance requirements](#)

[Node container migration requirements](#)

Configure the Docker storage volume

Before installing Docker, you might need to format the Docker storage volume and mount it on `/var/lib/docker`.

About this task

You can skip these steps if you plan to use local storage for the Docker storage volume and have sufficient space available on the host partition containing `/var/lib`.

Steps

1. Create a file system on the Docker storage volume:

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

2. Mount the Docker storage volume:

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

3. Add an entry for `docker-storage-volume-device` to `/etc/fstab`.

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

Install Docker

The StorageGRID system runs on Linux as a collection of Docker containers. Before you can install StorageGRID, you must install Docker.

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.

Related information

[Configure host storage](#)

Install StorageGRID host services

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

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 2.7 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.

Deploy virtual grid nodes (Ubuntu or Debian)

Create node configuration files for Ubuntu or Debian 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.

Where do I put the node configuration files?

You must 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.

What do I name the 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-nodenumber.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.

What is in a node configuration file?

The configuration files contain key/value pairs, with one key and one value per line. For each key/value pair, you must 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.

Some keys are required for every node, while others are optional or only required for certain node types.

The table defines the acceptable values for all supported keys. In the middle column:

R: required

BP: best practice

O: optional

Key	R, BP, or O?	Value
ADMIN_IP	BP	<p>Grid Network IPv4 address of the primary Admin Node for the grid to which this node belongs. Use the same value you specified for GRID_NETWORK_IP for the grid node with NODE_TYPE = VM_Admin_Node and ADMIN_ROLE = Primary. 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>
ADMIN_NETWORK_CONFIG	O	DHCP, STATIC, or DISABLED
ADMIN_NETWORK_ESL	O	<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>
ADMIN_NETWORK_GATEWAY	O (R)	<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>Note: This parameter is required if ADMIN_NETWORK_ESL is specified.</p> <p>Examples:</p> <p>1.1.1.1</p> <p>10.224.4.81</p>
ADMIN_NETWORK_IP	O	<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>

Key	R, BP, or O?	Value
ADMIN_NETWORK_MAC	O	<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>
ADMIN_NETWORK_MASK	O	<p>IPv4 netmask for 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>255.255.255.0</p> <p>255.255.248.0</p>
ADMIN_NETWORK_MTU	O	<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>

Key	R, BP, or O?	Value
ADMIN_NETWORK_TARGET	BP	<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>
ADMIN_NETWORK_TARGET_TYPE	O	<p>Interface</p> <p>(This is the only supported value.)</p>
ADMIN_NETWORK_TARGET_TYPE_INTERFACE_CLONE_MAC	BP	<p>True or False</p> <p>Set the key to "true" to cause the StorageGRID container use the MAC address of the 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:</p> <p>Considerations and recommendations for MAC address cloning (Red Hat Enterprise Linux or CentOS)</p> <p>Considerations and recommendations for MAC address cloning (Ubuntu or Debian)</p>
ADMIN_ROLE	R	<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>

Key	R, BP, or O?	Value
BLOCK_DEVICE_AUDIT_LOGS	R	<p>Path and name of the block device special file this node will use for persistent storage of audit logs. 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 /dev/disk/by-id/wwn-0x600a09800059d6df000060d757b475fd /dev/mapper/sgws-adm1-audit-logs</pre>

Key	R, BP, or O?	Value
BLOCK_DEVICE_RANGEDB_000	R	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 NODE_TYPE = VM_Storage_Node; don't specify it for other node types.
BLOCK_DEVICE_RANGEDB_001		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.
BLOCK_DEVICE_RANGEDB_002		
BLOCK_DEVICE_RANGEDB_003		
BLOCK_DEVICE_RANGEDB_004		
BLOCK_DEVICE_RANGEDB_005		
BLOCK_DEVICE_RANGEDB_006		
BLOCK_DEVICE_RANGEDB_007		
BLOCK_DEVICE_RANGEDB_008		
BLOCK_DEVICE_RANGEDB_009		
BLOCK_DEVICE_RANGEDB_010		
BLOCK_DEVICE_RANGEDB_011		
BLOCK_DEVICE_RANGEDB_012		
BLOCK_DEVICE_RANGEDB_013		
BLOCK_DEVICE_RANGEDB_014		
BLOCK_DEVICE_RANGEDB_015		

Key	R, BP, or O?	Value
BLOCK_DEVICE_TABLES	R	<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 /dev/disk/by-id/wwn-0x600a09800059d6df000060d757b475fd /dev/mapper/sgws-adm1-tables</pre>
BLOCK_DEVICE_VAR_LOCAL	R	<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 /dev/disk/by-id/wwn-0x600a09800059d6df000060d757b475fd /dev/mapper/sgws-sn1-var-local</pre>
CLIENT_NETWORK_CONFIG	O	DHCP, STATIC, or DISABLED
CLIENT_NETWORK_GATEWAY	O	<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> <pre>1.1.1.1 10.224.4.81</pre>

Key	R, BP, or O?	Value
CLIENT_NETWORK_IP	O	<p>IPv4 address of this node on the Client Network. 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>
CLIENT_NETWORK_MAC	O	<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>
CLIENT_NETWORK_MASK	O	<p>IPv4 netmask for this node on the Client Network. This key is only required 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>
CLIENT_NETWORK_MTU	O	<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>

Key	R, BP, or O?	Value
CLIENT_NETWORK_TARGET	BP	<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> <p>bond0.1003</p> <p>ens423</p>
CLIENT_NETWORK_TARGET_TYPE	O	<p>Interface</p> <p>(This is only supported value.)</p>
CLIENT_NETWORK_TARGET_TYPE_INTERFACE_CLONE_MAC	BP	<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:</p> <p>Considerations and recommendations for MAC address cloning (Red Hat Enterprise Linux or CentOS)</p> <p>Considerations and recommendations for MAC address cloning (Ubuntu or Debian)</p>
GRID_NETWORK_CONFIG	BP	<p>STATIC or DHCP</p> <p>(Defaults to STATIC if not specified.)</p>

Key	R, BP, or O?	Value
GRID_NETWORK_GATEWAY	R	<p>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.</p> <p>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.</p>
GRID_NETWORK_IP	R	<p>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.</p> <p>Examples:</p> <p>1.1.1.1</p> <p>10.224.4.81</p>
GRID_NETWORK_MAC	O	<p>The MAC address for the Grid 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:c2:27:30</p>
GRID_NETWORK_MASK	O	<p>IPv4 netmask for this node on the Grid Network. This key is only required 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>

Key	R, BP, or O?	Value
GRID_NETWORK_MTU	O	<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 8192</p>
GRID_NETWORK_TARGET	R	<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> <p>bond0.1001</p> <p>ens192</p>
GRID_NETWORK_TARGET_TYPE	O	<p>Interface</p> <p>(This is the only supported value.)</p>

Key	R, BP, or O?	Value
GRID_NETWORK_TARGET_TYPE_INTERFACE_CLONE_MAC	BP	<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:</p> <p>Considerations and recommendations for MAC address cloning (Red Hat Enterprise Linux or CentOS)</p> <p>Considerations and recommendations for MAC address cloning (Ubuntu or Debian)</p>
INTERFACE_TARGET_nnnn	O	<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>For example: INTERFACE_TARGET_0001=ens256, Trunk</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>

Key	R, BP, or O?	Value
MAXIMUM_RAM	O	<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 <number><unit>, where <unit> 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>
NODE_TYPE	R	<p>Type of node:</p> <p>VM_Admin_Node VM_Storage_Node VM_Archive_Node VM_API_Gateway</p>
PORT_REMAP	O	<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 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.</p> <p>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.</p> <p>For example:</p> <pre>PORT_REMAP = client/tcp/18082/443</pre>

Key	R, BP, or O?	Value
PORT_REMAP_INBOUND	O	<p>Remaps inbound communications to 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.</p> <p>IMPORTANT: Don't remap the ports you are planning to use to configure load balancer endpoints.</p> <p>The format used is: <network type>/<protocol:>/<remapped port >/<default port used by grid node>, where <network type> is grid, admin, or client, and protocol is tcp or udp.</p> <p>For example:</p> <pre>PORT_REMAP_INBOUND = grid/tcp/3022/22</pre>

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.

Example node configuration files

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.

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/dc1-adm1.conf

Example file contents:

```
NODE_TYPE = VM_Admin_Node
ADMIN_ROLE = Primary
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 Archive Node

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

Example file contents:

```

NODE_TYPE = VM_Archive_Node
ADMIN_IP = 10.1.0.2
BLOCK_DEVICE_VAR_LOCAL = /dev/mapper/dc1-arcl-var-local
GRID_NETWORK_TARGET = bond0.1001
ADMIN_NETWORK_TARGET = bond0.1002
CLIENT_NETWORK_TARGET = bond0.1003

GRID_NETWORK_IP = 10.1.0.4
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

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

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.

```
Checking for misnamed node configuration files... PASSED
Checking configuration file for node dc1-adm1... PASSED
Checking configuration file for node dc1-gw1... PASSED
Checking configuration file for node dc1-sn1... PASSED
Checking configuration file for node dc1-sn2... PASSED
Checking configuration file for node dc1-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-adm1
  WARNING: ignoring /etc/storagegrid/nodes/dc1-sn2.conf.keep
  WARNING: ignoring /etc/storagegrid/nodes/my-file.txt
Checking configuration file for node dc1-adm1...
  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

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

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
```

Configure grid and complete installation (Ubuntu or Debian)

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 one of the following addresses:

`https://primary_admin_node_ip`

`client_network_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.

2. Select **Install a StorageGRID system**.

The page used to configure a StorageGRID system appears.

NetApp® StorageGRID®

Help ▾

Install



License

Enter a grid name and upload the license file provided by NetApp for your StorageGRID system.

Grid Name

License File

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.

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

License

Enter a grid name and upload the license file provided by NetApp for your StorageGRID system.

Grid Name: StorageGRID

License File: NLF-959007-Internal.txt

License Serial Number: 959007


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.

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.



NetApp® StorageGRID®

Help ▾

Install

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

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	Raleigh	×	
Site Name 2	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.

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.



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	172.16.0.0/21	+
Discover Grid Network subnets		

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.
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.

					<input type="button" value="Approve"/>	<input type="button" value="Remove"/>	Search <input type="text"/>
	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	<input type="button" value="<"/>	<input type="button" value=">"/>

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.

					<input type="button" value="Edit"/>	<input type="button" value="Reset"/>	<input type="button" value="Remove"/>	Search <input type="text"/>
	Grid Network MAC Address	Name	Site	Type	Platform	Grid Network IPv4 Address		
<input checked="" type="radio"/>	00:50:56:87:42:ff	dc1-adm1	Raleigh	Admin Node	VMware VM	172.16.4.210/21	<input type="button" value="<"/>	
<input checked="" type="radio"/>	00:50:56:87:c0:16	dc1-s1	Raleigh	Storage Node	VMware VM	172.16.4.211/21	<input type="button" value=">"/>	
<input checked="" type="radio"/>	00:50:56:87:79:ee	dc1-s2	Raleigh	Storage Node	VMware VM	172.16.4.212/21		
<input checked="" type="radio"/>	00:50:56:87:db:9c	dc1-s3	Raleigh	Storage Node	VMware VM	172.16.4.213/21		
<input checked="" 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:

Storage Node Configuration

General Settings

Site	<input type="text" value="Raleigh"/>
Name	<input type="text" value="NetApp-SGA"/>
NTP Role	<input type="text" value="Automatic"/>
ADC Service	<input type="text" value="Automatic"/>

Grid Network

Configuration	STATIC
IPv4 Address (CIDR)	<input type="text" value="172.16.5.20/21"/>
Gateway	<input type="text" value="172.16.5.20"/>

Admin Network

Configuration	STATIC
IPv4 Address (CIDR)	<input type="text" value="10.224.5.20/21"/>
Gateway	<input type="text" value="10.224.0.1"/>
Subnets (CIDR)	<input type="text" value="10.0.0.0/8"/> ×
	<input type="text" value="172.19.0.0/16"/> ×
	<input type="text" value="172.21.0.0/16"/> + ×

Client Network

Configuration	STATIC
IPv4 Address (CIDR)	<input type="text" value="47.47.5.20/21"/>
Gateway	<input type="text" value="47.47.0.1"/>

- **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.

- **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.

- **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 resulting 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 resulting 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 resulting 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.

To learn how to install StorageGRID appliances, 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.

		<input type="button" value="Approve"/>	<input type="button" value="Remove"/>			Search	<input type="button" value=""/>				
	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.

		<input type="button" value="Edit"/>	<input type="button" value="Reset"/>	<input type="button" value="Remove"/>			Search	<input type="button" value=""/>				
	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.

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.

NetApp® StorageGRID®

Help ▾

Install

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

Network Time Protocol

Enter the IP addresses for at least four Network Time Protocol (NTP) servers, so that operations performed on separate servers are kept in sync.

Server 1	10.60.248.183
Server 2	10.227.204.142
Server 3	10.235.48.111
Server 4	0.0.0.0

+

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.

If the DNS server information is omitted or incorrectly configured, a DNST alarm is triggered on each grid node's SSM service. The alarm clears when DNS is configured correctly and the new server information has reached all grid nodes.

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 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 password in a secure place.

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

NetApp® StorageGRID®

Help ▾

Install

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

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	••••••••
Confirm Provisioning Passphrase	••••••••
Grid Management Root User Password	••••••••
Confirm Root User Password	••••••••

Create random command line passwords.

5. 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.

6. 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

1. 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

2. 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-admin1	Site1	172.16.4.215/21	<div style="width: 80%;"><div style="width: 100%;"></div></div>	Starting services
dc1-g1	Site1	172.16.4.216/21	<div style="width: 100%;"><div style="width: 100%;"></div></div>	Complete
dc1-s1	Site1	172.16.4.217/21	<div style="width: 70%;"><div style="width: 100%;"></div></div>	Waiting for Dynamic IP Service peers
dc1-s2	Site1	172.16.4.218/21	<div style="width: 10%;"><div style="width: 100%;"></div></div>	Downloading hotfix from primary Admin if needed
dc1-s3	Site1	172.16.4.219/21	<div style="width: 10%;"><div style="width: 100%;"></div></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 their IP addresses change, which can cause outages if a DHCP address 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.

Automate the installation (Ubuntu or Debian)

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

About this task

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 that can be created interactively during a manual installation, or prepared ahead of time (or programmatically) to enable automated installation using standard orchestration frameworks. StorageGRID provides optional Python scripts for automating the configuration of StorageGRID appliances, and the whole 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 and is driven by configuration files that can be prepared ahead of time (or programmatically) to enable automated installation. If you already use a standard orchestration framework to install and configure Ubuntu or Debian, adding StorageGRID to your playbooks or recipes should be straightforward.

You can automate these tasks:

1. Installing Linux
2. Configuring Linux
3. Configuring host network interfaces to meet StorageGRID requirements
4. Configuring host storage to meet StorageGRID requirements
5. Installing Docker
6. Installing the StorageGRID host service
7. Creating StorageGRID node configuration files in `/etc/storagegrid/nodes`
8. Validating StorageGRID node configuration files
9. Starting the StorageGRID host service

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.

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>	Sample 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 sample configuration file (`configure-storagegrid.sample.json`) or the blank configuration file (`configure-storagegrid.blank.json`).

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

1. Log in to the Linux machine you are using to run the Python script.
2. 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

[Overview of the installation REST API](#)

Overview of the 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 Docs 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.
- **schemas** — API schemas for advanced deployments
- **sites** — Site-level configuration operations. You can create, view, delete, and modify a site.

Related information

[Automating the installation](#)

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

- [Create a tenant account](#) for each client protocol (Swift or S3) 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](#) or [Swift 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](#).
- If your StorageGRID system includes any Archive Nodes (deprecated), configure the Archive Node's connection to the target external archival storage system.

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 access to the system for auditing purposes](#) through an NFS file share.

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.

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](#).

Related information

[Troubleshoot a StorageGRID system](#)

Example /etc/network/interfaces

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
```

Install VMware

Install VMware: Overview

Installing a StorageGRID system in a VMware environment includes three primary steps.

1. Preparation: During planning and preparation, you perform the following tasks:

- Learn about the hardware, software, virtual machine, storage, and performance requirements for StorageGRID.
- Learn about the specifics of [StorageGRID networking](#) so you can configure your network appropriately.
- Identify and prepare the physical servers you plan to use to host your StorageGRID grid nodes.
- On the servers you have prepared:
 - Install VMware vSphere Hypervisor
 - Configure the ESX hosts
 - Install and configure VMware vSphere and vCenter

2. Deployment: Deploy grid nodes using the VMware vSphere Web Client. When you deploy grid nodes, they are created as part of the StorageGRID system and connected to one or more networks.

- a. 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.
- b. Use the StorageGRID Appliance Installer to deploy StorageGRID appliance nodes.



Hardware-specific installation and integration instructions aren't included in the StorageGRID installation procedure. To learn how to install StorageGRID appliances, see the [Quick start for hardware installation](#) to locate instructions for your appliance.

3. Configuration: When all nodes have been deployed, use the Grid Manager to configure the grid and complete the installation.

These instructions recommend a standard approach for deploying and configuring a StorageGRID system in a

VMware environment. See also the information about the following alternative approaches:

- Use the `deploy-vsphere-ovftool.sh` Bash script (available from the installation archive) to deploy grid nodes in VMware vSphere.
- Automate the deployment and configuration of the StorageGRID system using a Python configuration script (provided in the installation archive).
- Automate the deployment and configuration of appliance grid nodes with a Python configuration script (available from the installation archive or from the StorageGRID Appliance Installer).
- If you are an advanced developer of StorageGRID deployments, use the installation REST APIs to automate the installation of StorageGRID grid nodes.

Plan and prepare for VMware installation

Before you install (VMware)

Before deploying grid nodes and configuring the StorageGRID grid, you must be familiar with the steps and requirements for completing the procedure.

The StorageGRID deployment and configuration procedures assume that you are familiar with the architecture and operational functionality of the StorageGRID system.

You can deploy a single site or multiple sites at one time; however, all sites must meet the minimum requirement of having at least three Storage Nodes.

Before starting the node deployment and grid configuration procedure, you must:

- Plan the StorageGRID deployment.
- Install, connect, and configure all required hardware, including any StorageGRID appliances, to specifications.



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.



Hardware-specific installation and integration instructions aren't included in the StorageGRID installation procedure. To learn how to install StorageGRID appliances, see [Install appliance hardware](#).

- Understand the [available network options](#) and how each network option should be implemented on grid nodes.
- Gather all networking information in advance. Unless you are using DHCP, gather the IP addresses to assign to each grid node, and the IP addresses of the DNS and NTP servers that will be used.
- Decide which of the available deployment and configuration tools you want to use.

Required materials

Before you install StorageGRID, you must gather and prepare required materials.

Item	Notes
NetApp StorageGRID license	<p>You must have a valid, digitally signed NetApp license.</p> <p>Note: The StorageGRID installation archive includes a free license that does not provide any support entitlement for the product.</p>
StorageGRID installation archive	<p>You must download the StorageGRID installation archive and extract the files.</p>
VMware software and documentation	<p>During installation, you use VMware vSphere Web Client to deploy virtual grid nodes on virtual machines. For supported versions, see the NetApp Interoperability Matrix Tool.</p>
Service laptop	<p>The StorageGRID system is installed through a service laptop. The service laptop must have:</p> <ul style="list-style-type: none"> • Network port • SSH client (for example, PuTTY) • Supported web browser
StorageGRID documentation	<ul style="list-style-type: none"> • Release notes • Instructions for administering StorageGRID

Download and extract the StorageGRID installation files

You must download the StorageGRID installation archives and extract the files..

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 file for VMware.



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

7. Save and extract the archive file.
8. Choose the files you need from the following list.

The files you need depend 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.

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

Path and file name	Description
./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.
./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	<p>API schemas for StorageGRID.</p> <p>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.</p>

Software requirements

You can use a virtual machine to host any type of StorageGRID grid node. One virtual machine is required for each grid node installed on the VMware server.

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 grid 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.

For information about supported servers, see the [NetApp Interoperability Matrix Tool](#).

Each StorageGRID node requires the following minimum resources:

- CPU cores: 8 per node
- RAM: At least 24 GB per node, and 2 to 16 GB less than the total system RAM, depending on the total RAM available and the amount of non-StorageGRID software running on the system

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.



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.

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 and performance requirements for StorageGRID nodes hosted by virtual machines, so you can provide enough space to support the initial configuration and future storage expansion.

Performance requirements

The performance of the OS volume and of the first storage volume significantly impacts the overall performance of the system. Ensure that these provide adequate disk performance in terms of latency, input/output operations per second (IOPS), and throughput.

All StorageGRID nodes require that the OS drive and all storage volumes have write-back caching enabled. The cache must be on a protected or persistent media.

Requirements for virtual machines that use NetApp ONTAP storage

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 a 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.

Number of virtual machines 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 virtual machine server. Using a dedicated virtual machine 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 virtual machine host, or they can be deployed on their own dedicated virtual machine hosts as required. However, if you have multiple nodes of the same type (two Gateway Nodes, for example), don't install all instances on the same virtual machine host.

Storage requirements by node type

In a production environment, the virtual machines for StorageGRID grid nodes must meet different requirements, depending on the types of nodes.

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

Node Type	Storage
Admin Node	100 GB LUN for OS 200 GB LUN for Admin Node tables 200 GB LUN for Admin Node audit log
Storage Node	100 GB LUN for OS 3 LUNs for each Storage Node on this host Note: A Storage Node can have 1 to 16 storage LUNs; at least 3 storage LUNs are recommended. Minimum size per LUN: 4 TB Maximum tested LUN size: 39 TB.
Gateway Node	100 GB LUN for OS
Archive Node	100 GB LUN for OS

 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.


Storage requirements for Storage Nodes

A software-based Storage Node can have 1 to 16 storage volumes—3 or more storage volumes are recommended. Each storage volume should be 4 TB or larger.



An appliance Storage Node can 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 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.

- 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](#).

Deploy virtual machine grid nodes (VMware)

Collect information about your deployment environment

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 grid 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.
- If you aren't using DHCP, the required networking details for each grid node (IP address, gateway, and network mask).
- If you aren't 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).
- 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

Related information

[How grid nodes discover the primary Admin Node](#)

[Storage and performance requirements](#)

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

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
vsphere-primary-admin.ovf	The template file and manifest file for the primary Admin Node.
vsphere-primary-admin.mf	
vsphere-non-primary-admin.ovf	The template file and manifest file for a non-primary Admin Node.
vsphere-non-primary-admin.mf	

Filename	Description
vsphere-archive.ovf	The template file and manifest file for an Archive Node.
vsphere-archive.mf	
vsphere-gateway.ovf	The template file and manifest file for a Gateway Node.
vsphere-gateway.mf	
vsphere-storage.ovf	The template file and manifest file for a Storage Node.
vsphere-storage.mf	

- The `.vdmk`, `.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 virtual machine server.



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

- 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 a 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, Gateway Nodes, and Archive Nodes.

If you are installing a new StorageGRID system:

- You must deploy the primary Admin Node before you deploy any other grid node.
- 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 the primary Admin Node 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.

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  
sphere-primary-admin.ovf  
sphere-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.
- 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.

7. Under **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. 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.

c. 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.

d. 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.

e. 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**.



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. Scroll back to the top of the table, and select **Edit**.
- ii. On the Type tab, select **User configurable**, and select **Save**.
- iii. Select **Set Value**.
- iv. 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

v. 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. Scroll back to the top of the table, and select **Edit**.

ii. On the Type tab, select **User configurable**, and select **Save**.

iii. Select **Set Value**.

iv. 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

v. 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.

Configure the grid and complete installation (VMware)

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 one of the following addresses:

`https://primary_admin_node_ip`

`https://client_network_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. Select **Install a StorageGRID system**.

The page used to configure a StorageGRID grid appears.

NetApp® StorageGRID®

Help ▾

Install

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

License

Enter a grid name and upload the license file provided by NetApp for your StorageGRID system.

Grid Name

License File

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.

License

Enter a grid name and upload the license file provided by NetApp for your StorageGRID system.

Grid Name: StorageGRID

License File: NLF-959007-Internal.txt

License Serial Number: 959007

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.

NetApp® StorageGRID®

Help ▾

Install

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	Raleigh	<input type="button" value="X"/>
Site Name 2	Atlanta	<input type="button" value="+ X"/>

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.

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.



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.

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.

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.

						Search	🔍			
	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.

						Search	🔍					
	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:

Storage Node Configuration

General Settings

Site	Raleigh
Name	NetApp-SGA
NTP Role	Automatic
ADC Service	Automatic

Grid Network

Configuration	STATIC
IPv4 Address (CIDR)	172.16.5.20/21
Gateway	172.16.5.20

Admin Network

Configuration	STATIC
IPv4 Address (CIDR)	10.224.5.20/21
Gateway	10.224.0.1
Subnets (CIDR)	10.0.0.0/8
	172.19.0.0/16
	172.21.0.0/16

+ ×

Client Network

Configuration	STATIC
IPv4 Address (CIDR)	47.47.5.20/21
Gateway	47.47.0.1

- **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.

- **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: 192.168.1.1234/21
- **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 resulting 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 resulting 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 resulting 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	Search
	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	Search
	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.

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.

NetApp® StorageGRID®

Help ▾

Install

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

Network Time Protocol

Enter the IP addresses for at least four Network Time Protocol (NTP) servers, so that operations performed on separate servers are kept in sync.

Server 1	10.60.248.183
Server 2	10.227.204.142
Server 3	10.235.48.111
Server 4	0.0.0.0

+

3. Select **Next**.

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.

If the DNS server information is omitted or incorrectly configured, a DNST alarm is triggered on each grid node's SSM service. The alarm clears when DNS is configured correctly and the new server information has reached all grid nodes.

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 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 password in a secure place.

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

NetApp® StorageGRID®

Help ▾

Install

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

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
Confirm Provisioning Passphrase
Grid Management Root User Password
Confirm Root User Password

Create random command line passwords.

5. 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.

6. 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

1. View the **Summary** page.

2. 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	Search	Filter
dc1-admin1	Site1	172.16.4.215/21	<div style="width: 80%;"><div style="width: 100%;"></div></div>	Starting services		
dc1-g1	Site1	172.16.4.216/21	<div style="width: 100%;"><div style="width: 100%;"></div></div>	Complete		
dc1-s1	Site1	172.16.4.217/21	<div style="width: 70%;"><div style="width: 100%;"></div></div>	Waiting for Dynamic IP Service peers		
dc1-s2	Site1	172.16.4.218/21	<div style="width: 10%;"><div style="width: 100%;"></div></div>	Downloading hotfix from primary Admin if needed		
dc1-s3	Site1	172.16.4.219/21	<div style="width: 10%;"><div style="width: 100%;"></div></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 their IP addresses change, which can cause outages if a DHCP address 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.

Automate the installation (VMware)

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

Automate grid node deployment

Use VMware vSphere 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 OVF Tool 4.1 installed and correctly configured.
- You know the username and password required to access VMware vSphere using the OVF Tool.
- 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 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 <code>.ovf</code> and <code>.mf</code> files.
vsphere-primary-admin.ovf vsphere-primary-admin.mf	The Open Virtualization Format template file (<code>.ovf</code>) and manifest file (<code>.mf</code>) for deploying the primary Admin Node.
vsphere-non-primary-admin.ovf vsphere-non-primary-admin.mf	The template file (<code>.ovf</code>) and manifest file (<code>.mf</code>) for deploying non-primary Admin Nodes.

Filename	Description
vsphere-archive.ovf	The template file (.ovf) and manifest file (.mf) for deploying Archive Nodes.
vsphere-archive.mf	
vsphere-gateway.ovf	The template file (.ovf) and manifest file (.mf) for deploying Gateway Nodes.
vsphere-gateway.mf	
vsphere-storage.ovf	The template file (.ovf) and manifest file (.mf) for deploying virtual machine-based Storage Nodes.
vsphere-storage.mf	
deploy-vsphere-ovftool.sh	The Bash shell script used to automate the deployment of virtual grid nodes.
deploy-vsphere-ovftool-sample.ini	The sample 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 a sample configuration file, so that you don't have to create the file from scratch.

Steps

1. Make a copy of the sample 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 = DHCP
```

- **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 = 8192
```

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 switch port 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. 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. 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.

- **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.

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, VM_Archive_Node, or VM_API_Gateway_Node
- **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

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 setting is 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
```

Archive Node

The following additional setting is required for Archive Nodes:

- **NODE_TYPE**: VM_Archive_Node

This example entry is for an Archive Node that is on the Grid and Admin Networks, but not on the Client Network.

```
[DC1-ARC1]
NODE_TYPE = VM_Archive_Node

GRID_NETWORK_IP = 10.1.0.4
ADMIN_NETWORK_IP = 10.3.0.4

ADMIN_IP = 10.1.0.2
```

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 grid 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
configure-storagegrid.py	Python script used to automate the configuration
configure-storagegrid.sample.json	Sample configuration file for use with the script
configure-storagegrid.blank.json	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 sample configuration file (`configure-storagegrid.sample.json`) or the blank configuration file (`configure-storagegrid.blank.json`).

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

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

For example:

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

where `platform` is `debs`, `rpm`s, 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](#)

[Overview of the installation REST API](#)

Overview of the 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 Docs 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.
- **schemas** — API schemas for advanced deployments
- **sites** — Site-level configuration operations. You can create, view, delete, and modify a site.

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

- 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 each client protocol (Swift or S3) 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](#) or [Swift 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](#).
- If your StorageGRID system includes any Archive Nodes (deprecated), configure the Archive Node's connection to the target external archival storage system.

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 access to the system for auditing purposes](#) through an NFS file share.

Troubleshoot installation issues

If any problems occur while installing your StorageGRID system, you can access the installation log files.

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)

Related information

To learn how to access the log files, see [Log files reference](#).

If you need additional help, contact [NetApp Support](#).

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.

Upgrade StorageGRID software

Upgrade StorageGRID software: Overview

Use these instructions to upgrade a StorageGRID system to a new release.

About these instructions

These instructions describe what's new in StorageGRID 11.7 and provide step-by-step instructions for upgrading all nodes in your StorageGRID system to the new release.

Before you begin

Review these topics to learn about the new features and enhancements in StorageGRID 11.7, determine whether any features have been deprecated or removed, and find out about changes to StorageGRID APIs.

- [What's new in StorageGRID 11.7](#)
- [Removed or deprecated features](#)
- [Changes to the Grid Management API](#)
- [Changes to the Tenant Management API](#)

What's new in StorageGRID 11.7

This release of StorageGRID introduces the following features and functional changes.

New features

Grid federation

You can configure a grid federation connection between two StorageGRID systems to clone tenant account information and replicate bucket objects between the grids for disaster recovery. See [What is grid federation?](#), [What is account clone](#), and [What is cross-grid replication](#).

Improved read availability

The read-after-new-write (default) consistency control was improved to be more available. GET/HEAD requests for non-existent objects will succeed with up to one Storage Node offline at each site. Buckets are no longer required to be set to the Available consistency control for this scenario. For example, applications that check existence of an object before creation will properly function with read-after-new-write even during software upgrade when one Storage Node is offline.

Rename grid, sites, and nodes

A new maintenance procedure lets you change the display names that are shown throughout the Grid Manager. You can update display names safely and whenever you need. See [Rename grid, sites, and nodes](#).

FabricPool and S3 setup wizard

The FabricPool and S3 setup wizard guides you through each step of configuring StorageGRID for use with ONTAP FabricPool or other S3 client application and produces a file you can use when entering required values in the other application. See [Use FabricPool setup wizard](#) and [Use S3 setup wizard](#).

Related to this change, a banner is now displayed on the dashboard to remind new users to configure [S3 endpoint domain names](#) for S3 virtual-hosted-style requests and set up [email notifications for alerts](#).

Firewall controls

The Firewall control page enables you to manage the external access of ports on nodes in your grid, and to define host addresses and IP subnets that are allowed access to closed ports. The new page also includes the untrusted Client Network settings, which now allow you to select additional ports you want open when untrusted Client Network is configured. See [Configure internal firewall](#).

Enhanced security policies

You can now determine which protocols and ciphers are used to establish secure TLS connections with client applications and secure SSH connections to internal StorageGRID services. See [Manage the TLS and SSH policy](#).

Load balancer endpoint changes

When [configuring load balancer endpoints](#), you can now:

- Allow all tenants to access the endpoint (default), or specify a list of allowed or blocked tenants to provide better security isolation between tenants and their endpoints.
- Use the **Node Type** binding mode to require clients to use the IP address (or corresponding FQDN) of any Admin Node or the IP address of any Gateway Node, based on the type of node you select.

SGF6112 all-flash appliance

The new StorageGRID SGF6112 storage appliance features a compact design with compute controller and storage controller integrated into a 1U chassis. The appliance supports 12 SSD NVMe drives with a storage capacity of up to 15.3 TB per drive. The SSD drives are in a RAID that provides resilient object storage. See [SGF6112 appliance: Overview](#).

Other Grid Manager enhancements

ILM enhancements

The improved ILM wizard makes it easier to specify filters, enter time periods and placements, and view retention diagrams. Erasure-coding profiles are created automatically when you select a storage pool and an EC scheme for a placement. For new StorageGRID 11.7 installations (not upgrades), a storage pool is automatically created for each site and the new **1 Copy Per Site** default rule ensures that new multi-site installations will have site-loss protection by default. See [Manage objects with ILM](#).

Customizable dashboard

You can now configure custom dashboards for the Grid Manager. See [View and manage the dashboard](#).

Volume restoration UI

Storage volume restoration lets you restore object data if a storage volume fails. For StorageGRID 11.7, you

can start volume restoration from Grid Manager in addition to the existing method of entering commands manually. Using Grid Manager is now the preferred method for restoring object data. See [Restore object data using Grid Manager](#).

Upgrade and hotfix UI

When you upgrade to StorageGRID 11.7, you can apply the latest 11.7 hotfix at the same time. The StorageGRID upgrade page shows the recommended upgrade path and links directly to the correct download pages. See [Perform upgrade](#).

Units for storage values

You can now select base 10 or base 2 units for the storage values displayed in the Grid Manager and Tenant Manager. Select the user drop-down in the upper right of the Grid Manager or Tenant Manager, then select [User preferences](#).

Access MIB from Grid Manager

You can now access SNMP-compliant MIB files from the Grid Manager using the SNMP agent page. See [Access MIB files](#).

Custom storage grades for new nodes

When you perform an expansion to add a new site or new Storage Nodes, you can now assign a custom storage grade to each new node. See [Perform expansion](#).

Tenant Manager updates

Cross-grid replication

Tenant accounts that have permission to use a [grid federation connection](#) can clone tenant groups, users, and S3 keys from one grid to another and use cross-grid replication to replicate bucket objects between two grids. See [Clone tenant groups and users](#) and [Manage cross-grid replication](#).

Delete all objects from bucket

Tenant Manager users can now delete all objects in a bucket, so the bucket can be deleted. See [Delete objects in bucket](#).

S3 Object Lock default retention

Tenant Manager users can now enable and configure default retention when creating S3 Object Lock buckets. See [Create an S3 bucket](#).

S3 updates

S3 Object Lock governance mode

When specifying the S3 Object Lock settings for an object or the default retention settings for a bucket, you can now use governance mode. This retention mode allows users with special permission to bypass certain retention settings. See [Use S3 Object Lock to retain objects](#) and [Use S3 REST API to configure S3 Object Lock](#).

S3 group policy for ransomware mitigation

When added as the group policy for an S3 tenant account, the sample policy helps mitigate ransomware attacks. It prevents older object versions from being permanently deleted. See [Create groups for an S3 tenant](#).

NewerNoncurrentVersions threshold for S3 buckets

The `NewerNoncurrentVersions` action in the bucket lifecycle configuration specifies the number of noncurrent versions retained in a versioned S3 bucket. This threshold overrides lifecycle rules provided by ILM. See [How objects are deleted](#).

S3 Select updates

S3 `SelectObjectContent` now offers support for Parquet objects. In addition, you can now use S3 Select with Admin and Gateway load balancer endpoints that are bare metal nodes running a kernel with cgroup v2 enabled. See [S3 SelectObjectContent](#).

Other enhancements

Certificate subject optional

The certificate subject field is now optional. If this field is left blank, the generated certificate uses the first domain name or IP address as the subject common name (CN). See [Manage security certificates](#).

ILM audit message category and new messages

An audit message category was added for ILM operations and includes the IDEL, LKCU, and ORLM messages. This new category is set to **Normal**. See [ILM operations audit messages](#).

In addition, new audit messages were added to support new 11.7 functionality:

- [BROR: Bucket Read Only Request](#)
- [CGRR: Cross-Grid Replication Request](#)
- [EBDL: Empty Bucket Delete](#)
- [EBKR: Empty Bucket Request](#)
- [S3SL: S3 Select Request](#)

New alerts

The following new alerts were added for StorageGRID 11.7:

- Appliance DAS drive fault detected
- Appliance DAS drive rebuilding
- Appliance fan fault detected
- Appliance NIC fault detected
- Appliance SSD critical warning
- AutoSupport message failed to send
- Cassandra oversize write error
- Cross-grid replication permanent request failure

- Cross-grid replication resources unavailable
- Debug performance impact
- Expiration of grid federation certificate
- FabricPool bucket has unsupported bucket consistency setting
- Firewall configuration failure
- Grid federation connection failure
- Storage appliance fan fault detected
- Storage Node not in desired storage state
- Storage volume needs attention
- Storage volume needs to be restored
- Storage volume offline
- Trace configuration enabled
- Volume Restoration failed to start replicated data repair

Documentation changes

- A new quick reference summarizes how StorageGRID supports Amazon Simple Storage Service (S3) APIs. See [Quick reference: Supported S3 API requests](#).
- The new [StorageGRID quick start](#) lists the high-level steps for configuring and using a StorageGRID system and provides links to the relevant instructions.
- The appliance hardware installation instructions were combined and consolidated for ease of use. A quick start was added as a high-level guide to hardware installation. See [Quick start for hardware installation](#).
- The maintenance instructions common to all appliance models were combined, consolidated, and moved to the maintenance section of the doc site. See [Common node maintenance: Overview](#).
- The maintenance instructions specific to each appliance model were also moved to the maintenance section:

[Maintain SGF6112 hardware](#)

[Maintain SG6000 hardware](#)

[Maintain SG5700 hardware](#)

[Maintain SG100 and SG1000 hardware](#)

Removed or deprecated features

Some features were removed or deprecated in this release. Review these items to understand whether you need to update client applications or modify your configuration before you upgrade.

Connection Load Balancer (CLB) service removed

The Connection Load Balancer (CLB) service on Gateway Nodes was deprecated in StorageGRID 11.4 and has now been completely removed from the software. To distribute incoming network connections from client applications to Storage Nodes, you can configure load balancer endpoints for the Load Balancer service, which

is included on all Admin Nodes and Gateway Nodes, or you can integrate a third-party load balancer. See [Considerations for load balancing](#).

If custom certificates were set up for the S3 or Swift API in the existing StorageGRID version, the CLB ports 8082, 8083, 8084, and 8085 will be automatically converted to load balancer endpoints during the upgrade to StorageGRID 11.7.

SG5600 appliance is End of Support

The SG5600 appliance has reached End Of Support. Contact your NetApp Sales Representative for hardware refresh options.

If you need to perform maintenance procedures on SG5600 hardware, use the [StorageGRID 11.6 instructions](#).

Swift support deprecated

As of the StorageGRID 11.7 release, support for Swift client applications has been deprecated. The user interface and APIs that support Swift client applications will be removed in a future release.

Archive Node support deprecated

Support for Archive Nodes (for both archiving to the cloud using the S3 API and archiving to tape using TSM middleware) is deprecated and will be removed in a future release. Moving objects from an Archive Node to an external archival storage system has been replaced by ILM Cloud Storage Pools, which offer more functionality.

See:

- [Migrate objects to a Cloud Storage Pool](#)
- [Use Cloud Storage Pools](#)

In addition, you should remove Archive Nodes from the active ILM policy in StorageGRID 11.7 or earlier. Removing object data stored on Archive Nodes will simplify future upgrades. See [Working with ILM rules and ILM policies](#).

Audit export through CIFS/Samba removed

Audit export through CIFS/Samba was deprecated in StorageGRID Webscale 11.1 and has now been removed. As required, you can [use an external syslog server](#) or [configure audit client access for NFS](#).

Option to specify a storage pool as a temporary location removed

Previously, when you created an ILM rule with an object placement that includes a single storage pool, you were prompted to specify a second storage pool to use as a temporary location. Starting with StorageGRID 11.7, this option has been removed.

Grid Manager options moved or removed

Several Grid Manager options were moved or removed.

- The [Compress stored objects](#) option was moved to **CONFIGURATION > System > Object compression**.
- The **Network Transfer Encryption** internal connection setting was removed and replaced by the [TLS and SSH policies](#) tab on the new **CONFIGURATION > Security > Security settings** page.



The AES256-SHA option was the default in StorageGRID 11.6 and is the only setting available in StorageGRID 11.7. The AES128-SHA value is ignored in the Grid Management API. During the StorageGRID 11.7 upgrade, the network transfer encryption algorithm is set to AES256-SHA.

- The **Stored object encryption**, **Prevent client modification**, and **Enable HTTP for Storage Node connections** options were moved to the [Network and objects tab](#) on the new **CONFIGURATION > Security > Security settings** page.
- The **Browser inactivity timeout** option was moved to the new **CONFIGURATION > Security > Security settings** page.
- The **Link cost** option was moved to **SUPPORT > Other > Link cost**.
- The list of NMS entities was moved to **SUPPORT > Other > NMS entities**.
- The **Stored Object Hashing** option was removed. The **SHA-1** and **SHA-256** settings are no longer used for internal background verification because they require additional CPU resources over MD5 and packet CRC32 check.
- The **Preferred sender** option was removed. If your StorageGRID deployment includes multiple Admin Nodes, the primary Admin Node is the preferred sender for alert notifications, AutoSupport messages, SNMP traps and informs, and legacy alarm notifications. If the primary Admin Node becomes unavailable, notifications are temporarily sent by other Admin Nodes. See [What is an Admin Node?](#).
- The [Untrusted Client Network settings](#) were moved to **CONFIGURATION > Firewall control**.

S3 endpoint domain name format restrictions

The use of IP addresses as endpoint domain names is unsupported. Future releases will prevent the configuration. If you need to use IP addresses for endpoint domain names, contact technical support. See [S3 endpoint domain names](#).

User initiated Volume Lost command removed

The `proc/CMSI/Volume_Lost` has been removed. Use the `repair-data start-replicated-volume-repair` command to restore replicated data for a volume.

Changes to the Grid Management API

StorageGRID 11.7 uses version 3 of the Grid Management API. Version 3 deprecates version 2; however, version 1 and version 2 are still supported.



You can continue to use version 1 and version 2 of the management API with StorageGRID 11.7; however, support for these versions of the API will be removed in a future release of StorageGRID. After upgrading to StorageGRID 11.7, the deprecated v1 and v2 APIs can be deactivated using the `PUT /grid/config/management API`.

To learn more, go to [Use the Grid Management API](#).

Display names now included in responses to node-health requests

Related to the new [Rename grid, sites, and nodes procedure](#), after you rename a site or node, the item's name (its system name) and its display name are both returned by the **node-health** API.

Can create bucket and access keys for new S3 tenant

New `s3Bucket` and `s3AccessKey` options were added to the **accounts** API. When you create an S3 tenant account using the Grid Management API, you can optionally create a bucket for that tenant as well as the access key ID and secret key for the tenant's root user.

Can change storage state for Storage Node

You can use the new **node-storage-state** API endpoints to determine and change the state of the storage in a Storage Node (online, offline, read-only).

Changes to the Tenant Management API

StorageGRID 11.7 uses version 3 of the Tenant Management API. Version 3 deprecates version 2; however, version 1 and version 2 are still supported.

 You can continue to use version 1 and version 2 of the management API with StorageGRID 11.7; however, support for these versions of the API will be removed in a future release of StorageGRID. After upgrading to StorageGRID 11.7, the deprecated v1 and v2 APIs can be deactivated using the `PUT /grid/config/management` API.

New endpoints for grid federation

You can use the **grid-federation-connections** API endpoints to list grid federation connections for the current tenant and to clear the last cross-grid replication error for the current tenant and selected grid federation connection.

To learn more, go to [Understand the Tenant Management API](#).

Plan and prepare for upgrade

Estimate the time to complete an upgrade

When planning an upgrade to StorageGRID 11.7, you must consider when to upgrade, based on how long the upgrade might take. You must also be aware of which operations you can and can't perform during each stage of the upgrade.

About this task

The time required to complete a StorageGRID upgrade depends on a variety of factors such as client load and hardware performance.

The table summarizes the main upgrade tasks and lists the approximate time required for each task. The steps after the table provide instructions you can use to estimate the upgrade time for your system.

Upgrade task	Description	Approximate time required	During this task
Run prechecks and upgrade primary Admin Node	The upgrade prechecks are run, and the primary Admin Node is stopped, upgraded, and restarted.	30 minutes to 1 hour, with SG100 and SG1000 appliance nodes requiring the most time. Unresolved precheck errors will increase this time.	You can't access the primary Admin Node. Connection errors might be reported, which you can ignore. Running the upgrade prechecks before starting the upgrade lets you resolve any errors before the scheduled upgrade maintenance window.
Start upgrade service	The software file is distributed, and the upgrade service is started.	3 minutes per grid node	
Upgrade other grid nodes	The software on all other grid nodes is upgraded, in the order in which you approve the nodes. Every node in your system will be brought down one at a time.	15 minutes to 1 hour per node, with appliance nodes requiring the most time Note: For appliance nodes, the StorageGRID Appliance Installer is automatically updated to the latest release.	<ul style="list-style-type: none"> Don't change the grid configuration. Don't change the audit level configuration. Don't update the ILM configuration. You are prevented from performing other maintenance procedures, such as hotfix, decommission, or expansion. Note: If you need to perform a recovery, contact technical support.
Enable features	The new features for the new version are enabled.	Less than 5 minutes	<ul style="list-style-type: none"> Don't change the grid configuration. Don't change the audit level configuration. Don't update the ILM configuration. You can't perform another maintenance procedure.
Upgrade database	The upgrade process checks each node to verify that the Cassandra database does not need to be updated.	10 seconds per node or a few minutes for the entire grid	<p>The upgrade from StorageGRID 11.6 to 11.7 does not require a Cassandra database upgrade; however, the Cassandra service will be stopped and restarted on each Storage Node.</p> <p>For future StorageGRID feature releases, the Cassandra database update step might take several days to complete.</p>
Final upgrade steps	Temporary files are removed and the upgrade to the new release completes.	5 minutes	When the Final upgrade steps task completes, you can perform all maintenance procedures.

Steps

1. Estimate the time required to upgrade all grid nodes.
 - a. Multiply the number of nodes in your StorageGRID system by 1 hour/node.

As a general rule, appliance nodes take longer to upgrade than software-based nodes.
 - b. Add 1 hour to this time to account for the time required to download the `.upgrade` file, run precheck validations, and complete the final upgrade steps.
2. If you have Linux nodes, add 15 minutes for each node to account for the time required to download and install the RPM or DEB package.
3. Calculate the total estimated time for the upgrade by adding the results of steps 1 and 2.

Example: Estimated time to upgrade to StorageGRID 11.7

Suppose your system has 14 grid nodes, of which 8 are Linux nodes.

1. Multiply 14 by 1 hour/node.
2. Add 1 hour to account for the download, precheck, and final steps.

The estimated time to upgrade all nodes is 15 hours.

3. Multiply 8 by 15 minutes/node to account for the time to install the RPM or DEB package on the Linux nodes.

The estimated time for this step is 2 hours.

4. Add the values together.

You should allow up to 17 hours to complete the upgrade of your system to StorageGRID 11.7.0.

 As required, you can split the maintenance window into smaller windows by approving subsets of grid nodes to upgrade in multiple sessions. For example, you might prefer to upgrade the nodes at site A in one session and then upgrade the nodes at site B in a later session. If you choose to perform the upgrade in more than one session, be aware that you can't start using the new features until all nodes have been upgraded.

How your system is affected during the upgrade

You must understand how your StorageGRID system will be affected during the upgrade.

StorageGRID upgrades are non-disruptive

The StorageGRID system can ingest and retrieve data from client applications throughout the upgrade process. If you approve all nodes of the same type to upgrade (for example, Storage Nodes), the nodes are brought down one at a time, so there is no time when all grid nodes or all grid nodes of a certain type are unavailable.

To allow for continued availability, ensure that your ILM policy contains rules that specify storing multiple copies of each object. You must also ensure that all external S3 or Swift clients are configured to send requests to one of the following:

- A high availability (HA) group virtual IP address

- A high availability third-party load balancer
- Multiple Gateway Nodes for each client
- Multiple Storage Nodes for each client

Appliance firmware is upgraded

During the StorageGRID 11.7 upgrade:

- All StorageGRID appliance nodes are automatically upgraded to StorageGRID Appliance Installer firmware version 3.7.
- SG6060 and SGF6024 appliances are automatically upgraded to BIOS firmware version 3B07.EX and BMC firmware version 3.97.07.
- SG100 and SG1000 appliances are automatically upgraded to BIOS firmware version 3B12.EC and BMC firmware version 4.71.07.

Alerts might be triggered

Alerts might be triggered when services start and stop and when the StorageGRID system is operating as a mixed-version environment (some grid nodes running an earlier version, while others have been upgraded to a later version). Other alerts might be triggered after the upgrade completes.

For example, you might see the **Unable to communicate with node** alert when services are stopped, or you might see the **Cassandra communication error** alert when some nodes have been upgraded to StorageGRID 11.7 but other nodes are still running StorageGRID 11.6. In general, these alerts will clear when the upgrade completes.

The **ILM placement unachievable** alert might be triggered when Storage Nodes are stopped during the upgrade to StorageGRID 11.7. This alert might persist for 1 day after the upgrade completes.

After the upgrade completes, you can review any upgrade-related alerts by selecting **Recently resolved alerts** or **Current alerts** from the Grid Manager dashboard.

Many SNMP notifications are generated

Be aware that a large number of SNMP notifications might be generated when grid nodes are stopped and restarted during the upgrade. To avoid excessive notifications, clear the **Enable SNMP Agent Notifications** checkbox (**CONFIGURATION > Monitoring > SNMP agent**) to disable SNMP notifications before you start the upgrade. Then, re-enable notifications after the upgrade is complete.

Configuration changes are restricted



This list applies specifically to upgrades from StorageGRID 11.6 to StorageGRID 11.7. If you're upgrading to another StorageGRID release, refer to the list of restricted changes in the upgrade instructions for that release.

Until the **Enable New Feature** task completes:

- Don't make any grid configuration changes.
- Don't enable or disable any new features. For example, don't attempt to create a grid federation connection until both StorageGRID systems have been updated to StorageGRID 11.7.
- Don't update the ILM configuration. Otherwise, you might experience inconsistent and unexpected ILM behavior.

- Don't apply a hotfix or recover a grid node.



Contact technical support if you need to recover a node during upgrade.

- You should not manage HA groups, VLAN interfaces, or load balancer endpoints while you're upgrading to StorageGRID 11.7.
- Don't delete any HA groups until the upgrade to StorageGRID 11.7 is complete. Virtual IP addresses in other HA groups might become inaccessible.

Until the **Final Upgrade Steps** task completes:

- Don't perform an expansion procedure.
- Don't perform a decommission procedure.

You can't view bucket details or manage buckets from the Tenant Manager

During the upgrade to StorageGRID 11.7 (that is, while the system is operating as a mixed-version environment), you can't view bucket details or manage buckets using the Tenant Manager. One of the following errors appears on the Buckets page in Tenant Manager:

- You can't use this API while you're upgrading to 11.7.
- You can't view bucket versioning details in the Tenant Manager while you're upgrading to 11.7.

This error will resolve after the upgrade to 11.7 is complete.

Workaround

While the 11.7 upgrade is in progress, use the following tools to view bucket details or manage buckets, instead of using the Tenant Manager:

- To perform standard S3 operations on a bucket, use either the [S3 REST API](#) or the [Tenant Management API](#).
- To perform StorageGRID custom operations on a bucket (for example, viewing and modifying the bucket consistency level, enabling or disabling last access time updates, or configuring search integration), use the Tenant Management API.

TLS ciphers or SSH configurations might change

If TLS ciphers or SSH configurations have been manually changed or are inconsistent across nodes, all nodes will be overwritten to be either Legacy Compatibility or Modern Compatibility after upgrade. If you used `fips-ciphers.sh` in StorageGRID 11.6, the Common Criteria policy is applied to all nodes. Otherwise, the Legacy Compatibility policy is applied. If you require Common Criteria validated configurations, you must use the Common Criteria policy or the FIPS strict policy. If you did not use `fips-ciphers.sh`, you should use the new Modern Compatibility setting after upgrade completes. To configure ciphers, go to **CONFIGURATION > Security > Security settings** and select **TLS and SSH policies**.

CLB ports might be converted to load balancer endpoints

The legacy Connection Load Balancer (CLB) service has been removed in StorageGRID 11.7. If CLB configuration is detected during upgrade prechecks, the **Legacy CLB load balancer activity detected** alert will be triggered. If custom certificates were set up for the S3 or Swift API in the existing StorageGRID version, the CLB ports 8082, 8083, 8084, and 8085 will be converted to load balancer endpoints during upgrade to StorageGRID 11.7.

See also [Considerations for load balancing](#).

Impact of an upgrade on groups and user accounts

You must understand the impact of the StorageGRID upgrade, so that you can update groups and user accounts appropriately after the upgrade is complete.

Changes to group permissions and options

After upgrading to StorageGRID 11.7, optionally assign the following new permission to tenant user groups.

Permission	Description
Tenant Manager > Manage objects with S3 Console	<p>When combined with the Manage all buckets permission, this permission allows users to access the Experimental S3 Console from the Buckets page.</p> <p>Users who have this permission but who don't have the Manage all buckets permission can still navigate directly to the Experimental S3 Console.</p>

See [Tenant management permissions](#).

Verify the installed version of StorageGRID

Before starting the upgrade, you must verify that the previous version of StorageGRID is currently installed with the latest available hotfix applied.

About this task

Before you upgrade to StorageGRID 11.7, your grid must have StorageGRID 11.6 installed. If you are currently using a previous version of StorageGRID, you must install all previous upgrade files along with their latest hotfixes (strongly recommended) until your grid's current version is StorageGRID 11.6.x.y.

One possible upgrade path is shown in the [example](#).



NetApp strongly recommends that you apply the latest hotfix for each StorageGRID version before upgrading to the next version and that you also apply the latest hotfix for each new version you install. In some cases, you must apply a hotfix to avoid the risk of data loss. See [NetApp Downloads: StorageGRID](#) and the release notes for each hotfix to learn more.

Note that you can run a script to update from 11.3.0.13+ to 11.4.0.y in one step and from 11.4.0.7+ to 11.5.0.y in one step. See [NetApp Knowledge Base: How to run combined major upgrade and hotfix script for StorageGRID](#).

Steps

1. Sign in to the Grid Manager using a [supported web browser](#).
2. From the top of the Grid Manager, select **Help > About**.
3. Verify that **Version** is 11.6.x.y.

In the StorageGRID 11.6.x.y version number:

- The **major release** has an *x* value of 0 (11.6.0).
- A **hotfix**, if one has been applied, has a *y* value (for example, 11.6.0.1).

4. If **Version** is not 11.6.x.y, go to [NetApp Downloads: StorageGRID](#) to download the files for each previous release, including the latest hotfix for each release.
5. Obtain the the upgrade instructions for each release you downloaded. Then, perform the software upgrade procedure for that release, and apply the latest hotfix for that release (strongly recommended).

See the [StorageGRID hotfix procedure](#).

Example: Upgrade to StorageGRID 11.6 from version 11.3.0.8

The following example shows the steps to upgrade from StorageGRID version 11.3.0.8 to version 11.6 in preparation for a StorageGRID 11.7 upgrade.



Optionally, you can run a script to combine steps 2 and 3 (update from 11.3.0.13+ to 11.4.0.y) and to combine steps 4 and 5 (update from 11.4.0.7+ to 11.5.0.y). See [NetApp Knowledge Base: How to run combined major upgrade and hotfix script for StorageGRID](#).

Download and install software in the following sequence to prepare your system for upgrade:

1. Apply the latest StorageGRID 11.3.0.y hotfix.
2. Upgrade to the StorageGRID 11.4.0 major release.
3. Apply the latest StorageGRID 11.4.0.y hotfix.
4. Upgrade to the StorageGRID 11.5.0 major release.
5. Apply the latest StorageGRID 11.5.0.y hotfix.
6. Upgrade to the StorageGRID 11.6.0 major release.
7. Apply the latest StorageGRID 11.6.0.y hotfix.

Obtain the required materials for a software upgrade

Before you begin the software upgrade, you must obtain all required materials so you can complete the upgrade successfully.

Item	Notes
Service laptop	The service laptop must have: <ul style="list-style-type: none"> • Network port • SSH client (for example, PuTTY)
Supported web browser	Browser support typically changes for each StorageGRID release. Make sure your browser is compatible with the new StorageGRID version.
Provisioning passphrase	The passphrase is created and documented when the StorageGRID system is first installed. The provisioning passphrase is not listed in the <code>Passwords.txt</code> file.

Item	Notes
Linux RPM or DEB archive	<p>If any nodes are deployed on Linux hosts, you must download and install the RPM or DEB package on all hosts before you start the upgrade.</p> <p>Important: Ensure that your operating system is upgraded to Linux kernel 4.15 or higher.</p>
StorageGRID documentation	<ul style="list-style-type: none"> • Release notes for StorageGRID 11.7 (sign in required). Be sure to read these carefully before starting the upgrade. • StorageGRID software upgrade resolution guide for the major version you are upgrading to (sign in required) • Other StorageGRID 11.7 documentation, as required.

Check the system's condition

Before upgrading a StorageGRID system, you must verify the system is ready to accommodate the upgrade. You must ensure that the system is running normally and that all grid nodes are operational.

Steps

1. Sign in to the Grid Manager using a [supported web browser](#).
2. Check for and resolve any active alerts.
3. Confirm that no conflicting grid tasks are active or pending.
 - a. Select **SUPPORT > Tools > Grid topology**.
 - b. Select **site > primary Admin Node > CMN > Grid Tasks > Configuration**.

Information lifecycle management evaluation (ILME) tasks are the only grid tasks that can run concurrently with the software upgrade.

- c. If any other grid tasks are active or pending, wait for them to finish or release their lock.



Contact technical support if a task does not finish or release its lock.

4. Refer to [Internal grid node communications](#) and [External communications](#) to ensure that all required ports for StorageGRID 11.7 are opened before you upgrade.

The following [internal ports](#) must be open before you upgrade to StorageGRID 11.7:

Port	Description
1055	Used for firewall controls knocking protocol.
8011	Before upgrading, confirm this port is open between all grid nodes on the Grid Network.
10342	

Port	Description
18086	<p>TCP port used for new LDR service.</p> <p>Before upgrading, confirm this port is open from all grid nodes to all Storage Nodes.</p>




If you have opened any custom firewall ports, you are notified during the upgrade precheck. You must contact technical support before proceeding with the upgrade.

Upgrade software

Upgrade workflow

Before starting the upgrade, review the general workflow. The StorageGRID Upgrade page guides you through each upgrade step.



1. If any StorageGRID nodes are deployed on Linux hosts, [install the RPM or DEB package on each host](#) before you start the upgrade.
2. From the primary Admin Node, access the StorageGRID Upgrade page and upload the upgrade file and the hotfix file, if required.
3. Download the current Recovery Package.
4. Run upgrade prechecks to detect and resolve any issues before you start the actual upgrade.
5. Start the upgrade, which runs prechecks and upgrades the primary Admin Node automatically. You can't access the Grid Manager while the primary Admin Node is being upgraded. Audit logs will also be unavailable. This upgrade can take up to 30 minutes.
6. After the primary Admin Node has been upgraded, download a new Recovery Package.

7. Approve the grid nodes. You can approve individual grid nodes, groups of grid nodes, or all grid nodes.



Don't approve the upgrade for a grid node unless you are sure that node is ready to be stopped and rebooted.

8. Resume operations. When all grid nodes have been upgraded, new features are enabled and you can resume operations. You must wait to perform a decommission or expansion procedure until the background **Upgrade database** task and the **Final upgrade steps** task have completed.

Related information

[Estimate the time to complete an upgrade](#)

Linux: Download and install the RPM or DEB package on all hosts

If any StorageGRID nodes are deployed on Linux hosts, you must download and install an additional RPM or DEB package on each of these hosts before you start the upgrade.

Download upgrade, Linux, and hotfix files

When you perform a StorageGRID upgrade from the Grid Manager, you are prompted to download the upgrade archive and any required hotfix as the first step. However, if you need to download files to upgrade Linux hosts, you can save time by downloading all required files in advance.

Steps

1. Go to [NetApp Downloads: StorageGRID](#).
2. Select the button for downloading the latest release, or select another version from the drop-down menu and select **Go**.
- StorageGRID software versions have this format: 11.x.y. StorageGRID hotfixes have this format: 11.x.y.z.
3. Sign in with the username and password for your NetApp account.
4. If a Caution/MustRead notice appears, make note of the hotfix number, and select the checkbox.
5. Read the End User License Agreement, select the checkbox, and then select **Accept & Continue**.

The downloads page for the version you selected appears. The page contains three columns.

6. From the second column (**Upgrade StorageGRID**), download two files:
 - The upgrade archive for the latest release (this is the file in the section labeled **VMware, SG1000, or SG100 Primary Admin Node**). While this file is not needed until you perform the upgrade, downloading it now will save time.
 - An RPM or DEB archive in either .tgz or .zip format. Select the .zip file if you are running Windows on the service laptop.
 - Red Hat Enterprise Linux or CentOS
StorageGRID-Webscale-version-RPM-uniqueID.zip
StorageGRID-Webscale-version-RPM-uniqueID.tgz
 - Ubuntu or Debian
StorageGRID-Webscale-version-DEB-uniqueID.zip
StorageGRID-Webscale-version-DEB-uniqueID.tgz
7. If you needed to agree to a Caution/MustRead notice because of a required hotfix, download the hotfix:

- a. Go back to [NetApp Downloads: StorageGRID](#).
- b. Select the hotfix number from the drop-down.
- c. Agree to the Caution notice and EULA again.
- d. Download and save the hotfix and its README.

You will be prompted to upload the hotfix file on the StorageGRID Upgrade page when you start the upgrade.

Install archive on all Linux hosts

Perform these steps before upgrading StorageGRID software.

Steps

1. Extract the RPM or DEB packages from the installation file.
2. Install the RPM or DEB packages on all Linux hosts.

See the steps for installing StorageGRID host services in the installation instructions:

- [Red Hat Enterprise Linux or CentOS: Install StorageGRID host services](#)
- [Ubuntu or Debian: Install StorageGRID host services](#)

The new packages are installed as additional packages. Don't remove the existing packages.

Perform the upgrade

You can upgrade to StorageGRID 11.7 and apply the latest hotfix for that release at the same time. The StorageGRID upgrade page provides the recommended upgrade path and links directly to the correct download pages.

Before you begin

You have reviewed all of the considerations and completed all of the planning and preparation steps.

Access StorageGRID Upgrade page

As a first step, access the StorageGRID Upgrade page in the Grid Manager.

Steps

1. Sign in to the Grid Manager using a [supported web browser](#).
2. Select **MAINTENANCE > System > Software update**.
3. From the StorageGRID upgrade tile, select **Upgrade**.

Select files

The update path on the StorageGRID Upgrade page indicates which major versions (for example, 11.7.0) and hotfixes (for example, 11.7.0.1) you must install to get to the latest StorageGRID release. You should install the recommended versions and hotfixes in the order shown.



If no update path is shown, your browser might not be able to access the NetApp Support Site, or the **Check for software updates** checkbox on the AutoSupport page (**SUPPORT > Tools > AutoSupport**) might be disabled.

Steps

1. For the **Select files** step, review the update path.
2. From the Download files section, select each **Download** link to download the required files from the NetApp Support Site.

If no update path is shown, go to the [NetApp Downloads: StorageGRID](#) to determine if a new version or hotfix is available and to download the files you need.



If you needed to download and install an RPM or DEB package on all Linux hosts, you might already have the StorageGRID upgrade and hotfix files listed in the update path.

3. Select **Browse** to upload the version upgrade file to StorageGRID:
`NetApp_StorageGRID_11.7.0_Software_uniqueID.upgrade`

When the upload and validation process is done, a green check mark appears next to the file name.

4. If you downloaded a hotfix file, select **Browse** to upload that file. The hotfix will be automatically applied as part of the version upgrade.
5. Select **Continue**.

Run prechecks

Running prechecks allows you to detect and resolve any upgrade issues before you start upgrading your grid.

Steps

1. For the **Run prechecks** step, start by entering the provisioning passphrase for your grid.
2. Select **Download recovery package**.

You should download the current copy of the Recovery Package file before you upgrade the primary Admin Node. The Recovery Package file allows you to restore the system if a failure occurs.

3. When the file is downloaded, confirm that you can access the contents, including the `Passwords.txt` file.
4. Copy the downloaded file (`.zip`) to 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.

5. Select **Run prechecks**, and wait for the prechecks to complete.
6. Review the details for each reported precheck and resolve any reported errors. See the [StorageGRID software upgrade resolution guide](#) for the StorageGRID 11.7 release.

You must resolve all precheck *errors* before you can upgrade your system. However, you don't need to address precheck *warnings* before upgrading.



If you have opened any custom firewall ports, you are notified during the precheck validation. You must contact technical support before proceeding with the upgrade.

7. If you made any configuration changes to resolve the reported issues, select **Run prechecks** again to get updated results.

If all errors have been resolved, you are prompted to start the upgrade.

Start upgrade and upgrade primary Admin Node

When you start the upgrade, the upgrade prechecks are run again, and the primary Admin Node is automatically upgraded. This part of the upgrade can take up to 30 minutes.



You will not be able to access any other Grid Manager pages while the primary Admin Node is being upgraded. Audit logs will also be unavailable.

Steps

1. Select **Start upgrade**.

A warning appears to remind you will temporarily lose access to the Grid Manager.

2. Select **OK** to acknowledge the warning and start the upgrade.
3. Wait for the upgrade prechecks to be performed and for the primary Admin Node to be upgraded.



If any precheck errors are reported, resolve them and select **Start upgrade** again.

If the grid has another Admin Node that is online and ready, you can use it to monitor the status of the primary Admin Node. As soon as the primary Admin Node is upgraded, you can approve the other grid nodes.

4. As required, select **Continue** to access the **Upgrade other nodes** step.

Upgrade other nodes

You must upgrade all grid nodes, but you can perform multiple upgrade sessions and customize the upgrade sequence. For example, you might prefer to upgrade the nodes at site A in one session and then upgrade the nodes at site B in a later session. If you choose to perform the upgrade in more than one session, be aware that you can't start using the new features until all nodes have been upgraded.

If the order in which nodes are upgraded is important, approve nodes or groups of nodes one at a time and wait until the upgrade is complete on each node before approving the next node or group of nodes.



When the upgrade starts on a grid node, the services on that node are stopped. Later, the grid node is rebooted. To avoid service interruptions for client applications that are communicating with the node, don't approve the upgrade for a node unless you are sure that node is ready to be stopped and rebooted. As required, schedule a maintenance window or notify customers.

Steps

1. For the **Upgrade other nodes** step, review the Summary, which provides the start time for the upgrade as a whole and the status for each major upgrade task.
 - **Start upgrade service** is the first upgrade task. During this task, the software file is distributed to the grid nodes, and the upgrade service is started on each node.
 - When the **Start upgrade service** task is complete, the **Upgrade other grid nodes** task starts, and you are prompted to download a new copy of the Recovery Package.

2. When prompted, enter your provisioning passphrase and download a new copy of the Recovery Package.



You should download a new copy of the Recovery Package file after the primary Admin Node is upgraded. The Recovery Package file allows you to restore the system if a failure occurs.

3. Review the status tables for each type of node. There are tables for non-primary Admin Nodes, Gateway Nodes, Storage Nodes, and Archive Nodes.

A grid node can be in one of these stages when the tables first appear:

- Unpacking the upgrade
- Downloading
- Waiting to be approved

4. When you are ready to select grid nodes for upgrade (or if you need to unapprove selected nodes), use these instructions:

Task	Instruction
Search for specific nodes to approve, such as all nodes at a particular site	Enter the search string in the Search field
Select all nodes for upgrade	Select Approve all nodes
Select all nodes of the same type for upgrade (for example, all Storage Nodes)	Select the Approve all button for the node type If you approve more than one node of the same type, the nodes will be upgraded one at a time.
Select an individual node for upgrade	Select the Approve button for the node
Postpone the upgrade on all selected nodes	Select Unapprove all nodes
Postpone the upgrade on all selected nodes of the same type	Select the Unapprove all button for the node type
Postpone the upgrade on an individual node	Select the Unapprove button for the node

5. Wait for the approved nodes to proceed through these upgrade stages:

- Approved and waiting to be upgraded
- Stopping services



You can't remove a node when its Stage reaches **Stopping services**. The **Unapprove** button is disabled.

- Stopping container
- Cleaning up Docker images

- Upgrading base OS packages



When an appliance node reaches this stage, the StorageGRID Appliance Installer software on the appliance is updated. This automated process ensures that the StorageGRID Appliance Installer version remains in sync with the StorageGRID software version.

- Rebooting



Some appliance models might reboot multiple times to upgrade the firmware and BIOS.

- Performing steps after reboot
- Starting services
- Done

6. Repeat the [approval step](#) as many times as needed until all grid nodes have been upgraded.

Complete upgrade

When all grid nodes have completed the upgrade stages, the **Upgrade other grid nodes** task is shown as Completed. The remaining upgrade tasks are performed automatically in the background.

Steps

1. As soon as the **Enable features** task is complete (which occurs quickly), you can start using the [new features](#) in the upgraded StorageGRID version.
2. During the **Upgrade database** task, the upgrade process checks each node to verify that the Cassandra database does not need to be updated.



The upgrade from StorageGRID 11.6 to 11.7 does not require a Cassandra database upgrade; however, the Cassandra service will be stopped and restarted on each Storage Node. For future StorageGRID feature releases, the Cassandra database update step might take several days to complete.

3. When the **Upgrade database** task has completed, wait a few minutes for the **Final upgrade steps** to complete.
4. When the **Final upgrade steps** have completed, the upgrade is done. The first step, **Select files**, is redisplayed with a green success banner.
5. Verify that grid operations have returned to normal:
 - a. Check that the services are operating normally and that there are no unexpected alerts.
 - b. Confirm that client connections to the StorageGRID system are operating as expected.

Troubleshoot upgrade issues

If something goes wrong when you perform an upgrade, you might be able to resolve the issue yourself. If you can't resolve an issue, gather as much information as you can and then contact technical support.

Upgrade does not complete

The following sections describe how to recover from situations where the upgrade has partially failed.

Upgrade precheck errors

To detect and resolve issues, you can manually run the upgrade prechecks before starting the actual upgrade. Most precheck errors provide information about how to resolve the issue.

Provisioning failures

If the automatic provisioning process fails, contact technical support.

Grid node crashes or fails to start

If a grid node crashes during the upgrade process or fails to start successfully after the upgrade finishes, contact technical support to investigate and to correct any underlying issues.

Ingest or data retrieval is interrupted

If data ingest or retrieval is unexpectedly interrupted when you aren't upgrading a grid node, contact technical support.

Database upgrade errors

If the database upgrade fails with an error, retry the upgrade. If it fails again, contact technical support.

Related information

[Checking the system's condition before upgrading software](#)

User interface issues

You might experience issues with the Grid Manager or the Tenant Manager during or after the upgrade.

Grid Manager displays multiple error messages during upgrade

If you refresh your browser or navigate to another Grid Manager page while the primary Admin Node is being upgraded, you might see multiple “503: Service Unavailable” and “Problem connecting to the server” messages. You can safely ignore these messages—they will stop appearing soon as the node is upgraded.

If these messages appear for more than an hour after you started the upgrade, something might have occurred that prevented the primary Admin Node from being upgraded. If you are unable to resolve the issue on your own, contact technical support.

Web interface does not respond as expected

The Grid Manager or the Tenant Manager might not respond as expected after StorageGRID software is upgraded.

If you experience issues with the web interface:

- Make sure you are using a [supported web browser](#).



Browser support typically changes for each StorageGRID release.

- Clear your web browser cache.

Clearing the cache removes outdated resources used by the previous version of StorageGRID software, and permits the user interface to operate correctly again. For instructions, see the documentation for your web browser.

“Docker image availability check” error messages

When attempting to start the upgrade process, you might receive an error message that states “The following issues were identified by the Docker image availability check validation suite.” All issues must be resolved before you can complete the upgrade.

Contact technical support if you are unsure of the changes required to resolve the identified issues.

Message	Cause	Solution
Unable to determine upgrade version. Upgrade version info file {file_path} did not match the expected format.	The upgrade package is corrupt.	Re-upload the upgrade package, and try again. If the problem persists, contact technical support.
Upgrade version info file {file_path} was not found. Unable to determine upgrade version.	The upgrade package is corrupt.	Re-upload the upgrade package, and try again. If the problem persists, contact technical support.
Unable to determine currently installed release version on {node_name}.	A critical file on the node is corrupt.	Contact technical support.
Connection error while attempting to list versions on {node_name}	The node is offline or the connection was interrupted.	Check to make sure that all nodes are online and reachable from the primary Admin Node, and try again.
The host for node {node_name} does not have StorageGRID {upgrade_version} image loaded. Images and services must be installed on the host before the upgrade can proceed.	The RPM or DEB packages for the upgrade have not been installed on the host where the node is running, or the images are still in the process of being imported. Note: This error only applies to nodes that are running as containers on Linux.	Check to make sure that the RPM or DEB packages have been installed on all Linux hosts where nodes are running. Make sure the version is correct for both the service and the images file. Wait a few minutes, and try again. See Linux: Install RPM or DEB package on all hosts .
Error while checking node {node_name}	An unexpected error occurred.	Wait a few minutes, and try again.
Uncaught error while running prechecks. {error_string}	An unexpected error occurred.	Wait a few minutes, and try again.

Copyright information

Copyright © 2025 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.