



StorageGRID® 11.3

Installation Guide

For Red Hat® Enterprise Linux® or CentOS Deployments

January 2020 | 215-14195_2020-01_en-us
doccomments@netapp.com



Contents

Installation overview	5
Planning and preparation	7
Required materials	8
Downloading and extracting the StorageGRID installation files	8
Hardware requirements	10
Networking requirements	10
Network model	11
Networking guidelines	12
Network installation and provisioning	14
Host network configuration	16
Deployment tools	17
Internal grid node communications	17
External communications	19
Networking and ports for platform services and Cloud Storage Pools	23
Storage requirements	23
Node container migration requirements	26
Web browser requirements	28
Preparing the hosts	28
Installing Linux	28
Configuring the host network	29
Configuring host storage	31
Installing Docker	33
Installing StorageGRID host services	35
Deploying virtual grid nodes	36
Creating node configuration files	36
How grid nodes discover the primary Admin Node	46
Example node configuration files	46
Validating the StorageGRID configuration	48
Starting the StorageGRID host service	49
Configuring the grid and completing installation	51
Navigating to the Grid Manager	51
Specifying the StorageGRID license information	52
Adding sites	52
Specifying Grid Network subnets	53
Approving pending grid nodes	54
Specifying Network Time Protocol server information	58
Specifying Domain Name System server information	59
Specifying the StorageGRID system passwords	59
Reviewing your configuration and completing installation	60
Automating the installation	63
Automating the installation and configuration of the StorageGRID host service	63

Automating the configuration and installation of appliance Storage Nodes	63
Automating the configuration of StorageGRID	65
Overview of installation REST APIs	68
Where to go next	70
Troubleshooting installation issues	72
Sample /etc/sysconfig/network-scripts	73
Copyright	75
Trademark information	76
How to send comments about documentation and receive update notifications	77

Installation 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 CPU, network, and storage requirements for StorageGRID.
 - 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. You use the Linux command line to deploy virtual grid nodes on the hosts you prepared in Step 1.
 - b. You use the StorageGRID Appliance Installer to deploy StorageGRID appliance nodes.
- 3. Configuration:** When all nodes have been deployed, you use the StorageGRID Grid Manager to configure the grid and complete the installation.

This document recommends a standard approach for deploying and configuring a StorageGRID system. It also provides 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 Docker and the StorageGRID host service, and deploy virtual grid nodes.
- Configure the StorageGRID system using a Python configuration script (provided in the installation archive).
- Deploy and configure appliance grid nodes with a second Python configuration script (available from the installation archive or from the StorageGRID Appliance Installer).
- Use the installation REST APIs to automate the installation of StorageGRID virtual nodes and appliance nodes.

Related concepts

[Planning and preparation](#) on page 7

[Overview of installation REST APIs](#) on page 68

Related tasks

[Deploying virtual grid nodes](#) on page 36

Configuring the grid and completing installation on page 51

Automating the installation on page 63

Planning and preparation

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 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 StorageGRID's 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.
- 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 are 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 domain name system (DNS) and network time protocol (NTP) servers that will be used.
- Install, connect, and configure all required hardware, including any StorageGRID appliances, to specifications.

Note: Hardware-specific installation and integration instructions are not included in the StorageGRID installation procedure. To learn how to install StorageGRID appliances, see the installation and maintenance instructions for your appliance.

- Decide which of the available deployment and configuration tools you want to use.

Steps

1. [Required materials](#) on page 8
2. [Downloading and extracting the StorageGRID installation files](#) on page 8
3. [Hardware requirements](#) on page 10
4. [Networking requirements](#) on page 10
5. [Storage requirements](#) on page 23
6. [Node container migration requirements](#) on page 26
7. [Web browser requirements](#) on page 28
8. [Preparing the hosts](#) on page 28

Related information

[SG1000 appliance installation and maintenance](#)
[SG6000 appliance installation and maintenance](#)
[SG5700 appliance installation and maintenance](#)

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 one of the following StorageGRID installation archives and extract the files to your service laptop. <ul style="list-style-type: none"> StorageGRID-Webscale-version-RPM-uniqueID.zip StorageGRID-Webscale-version-RPM-uniqueID.tgz
Service laptop	The StorageGRID system is installed through a service laptop. The service laptop must have: <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 tasks

[Downloading and extracting the StorageGRID installation files](#) on page 8

Related references

[Web browser requirements](#) on page 28

Related information

[Administering StorageGRID](#)

[StorageGRID release notes](#)

Downloading and extracting 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.

[NetApp Downloads: StorageGRID](#)

2. Click the button for downloading the latest release, or select another version from the pull-down menu and click **Go**.
3. Sign in using the username and password for your NetApp account.
4. Read and accept the End User License Agreement.

The downloads page for the version you selected appears. The page contains columns for new installation files, upgrade files, and NAS Bridge.

5. In the **New install files** column, click the button for the appropriate software.

Download the .tgz or .zip archive file for your platform.

- `StorageGRID-Webscale-version-RPM-uniqueID.zip`
- `StorageGRID-Webscale-version-RPM-uniqueID.tgz`

The compressed files contain the RPM files and scripts for Red Hat Enterprise Linux or CentOS.

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

6. Save and extract the archive file.
7. 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 grid.

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

Table 1: Files for Red Hat Enterprise Linux or CentOS Linux

Path and file name	Description
<code>/rpms/README</code>	A text file that describes all of the files contained in the StorageGRID download file.
<code>/rpms/NLF000000.txt</code>	A free license that does not provide any support entitlement for the product.
<code>/rpms/StorageGRID-Webscale-Images-version-SHA.rpm</code>	RPM package for installing the StorageGRID node images on your RHEL or CentOS hosts.
<code>/rpms/StorageGRID-Webscale-Service-version-SHA.rpm</code>	RPM package for installing the StorageGRID host service on your RHEL or CentOS hosts.
Deployment scripting tools	
<code>/rpms/configure-storagegrid.py</code>	A Python script used to automate the configuration of a StorageGRID system.
<code>/rpms/configure-storagegrid.sample.json</code>	A sample configuration file for use with the <code>configure-storagegrid.py</code> script.
<code>/rpms/configure-storagegrid.blank.json</code>	A blank configuration file for use with the <code>configure-storagegrid.py</code> script.
<code>/rpms/configure-sga.py</code>	A Python script used to automate the configuration of StorageGRID appliances.

Path and file name	Description
/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.

Hardware requirements

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

The following table lists the supported minimum resource requirements for each StorageGRID node. Use these values to 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 are not dedicated to running StorageGRID (not recommended), be sure to consider the resource requirements of the other applications.

Type of node	CPU cores	RAM
Admin	8	24 GB
Storage	8	24 GB
Gateway	8	24 GB
Archive	8	24 GB

Note: 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.

Note: 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 according to the table.

Note: For production deployments, you should not run multiple Storage Nodes on the same physical 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.

Networking requirements

You must verify that the networking infrastructure and configuration is in place to support your StorageGRID system.

For more information on networking configuration and supported network topologies, see the *Grid primer*.

Related information

[Grid primer](#)

Network model

You can configure three networks for use with the StorageGRID system. Each network type must be on a separate subnet with no overlap.

To understand how these three networks are used, consider the three types of network traffic that are processed by nodes in a StorageGRID system:

- Grid traffic: The internal StorageGRID traffic that travels between all nodes in the grid
- Admin traffic: The traffic used for system administration and maintenance
- Client traffic: The traffic that travels between external client applications and the grid, including all object storage requests from S3 and Swift clients

To allow you more precise control and security, you can configure one, two, or three networks to manage these three types of traffic.

Grid Network

The Grid Network is required. It is used for all internal StorageGRID traffic. The Grid Network provides connectivity between all nodes in the grid, across all sites and subnets. All hosts on the Grid Network must be able to talk to all other hosts. The Grid Network can consist of multiple subnets. Networks containing critical grid services, such as NTP, can also be added as Grid subnets.

When the Grid Network is the only StorageGRID network, it is also used for all admin traffic and all client traffic. The Grid Network gateway is the node default gateway unless the node has the Client Network configured.

Attention: When configuring the Grid Network, you must ensure that the network is secured from untrusted clients, such as those on the open internet.

The Grid Network is mapped using `GRID_NETWORK_TARGET` (eth0 inside the container).

IP/ Mask	Gateway	Static routes	Default route (0.0.0.0)
CIDR for static IP	The Grid Network gateway must be configured if there are multiple grid subnets. The Grid Network gateway is the node default gateway until grid configuration is complete.	Static routes are generated automatically for all nodes to all subnets configured in the global Grid Network Subnet List.	The Grid Network Gateway IP is the default gateway. If a Client Network is added, the default gateway switches from the Grid Network gateway to the Client Network gateway when grid configuration is complete.

Admin Network

The Admin Network is optional. It 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.

Using the Admin Network for administrative access allows the Grid Network to be isolated and secure. Typical uses of the Admin Network include access to the Grid Manager, access to critical services, such as NTP and DNS, access to audit logs on Admin Nodes, and SSH access to all nodes for maintenance and support. The Admin Network is never used for internal grid traffic. An Admin Network gateway is provided and allows the Admin Network to span multiple subnets. However, the Admin Network gateway is never used as the node default gateway.

The Admin Network is mapped using `ADMIN_NETWORK_TARGET` (eth1 inside the container).

IP/Mask	Gateway	Static routes	Default route (0.0.0.0/0)
CIDR for static IP	The Admin Network gateway is required if multiple admin subnets are defined.	Static routes are generated automatically to each subnet configured in the node's Admin Network Subnet List.	N/A

Client Network

The Client Network is also optional. It is an open network used to provide access to grid services for client applications such as S3 and Swift. The Client Network enables grid nodes to communicate with any subnet reachable through the Client Network gateway. Optionally, you can configure the Client Network so that the appliance can be accessed over this network using only the ports that you choose to open. The Client Network does not become operational until you complete the StorageGRID configuration steps.

You can use the Client Network to provide client access to the grid, so you can isolate and secure the Grid Network. The following nodes are often configured with a Client Network:

- Gateway Nodes and Storage Nodes, because these nodes provide S3 and Swift protocol access to the grid.
- Admin Nodes, because these nodes provide access to the Tenant Manager.

When a Client Network is configured, the Client Network gateway is required and becomes the node default gateway after the grid has been configured.

The Client Network is mapped using `CLIENT_NETWORK_TARGET` (eth2 inside the container).

IP/Mask	Gateway	Static routes	Default route (0.0.0.0/0)
CIDR for static IP	The Client Network gateway is required if the Client Network is configured. The Client Network gateway becomes the default route for the grid node when grid configuration is complete.	N/A	Added if a Client Network Gateway IP is configured

Networking guidelines

Follow these guidelines when configuring StorageGRID networks.

Grid Network guidelines

- You configure the Grid Network Subnet List using the Grid Manager to enable static route generation between subnets on the Grid Network.
- Each node must be attached to the Grid Network and must be able to communicate with the primary Admin Node using the networking configuration you specify when deploying the node.

Subnet guidelines

Note: The following restrictions are enforced by the Grid Manager during deployment. They are provided here to assist in pre-deployment network planning.

- The subnet mask for any network IP address cannot be 255.255.255.254 or 255.255.255.255 (/31 or /32 in CIDR notation).

- The subnets defined by a given node's IP/mask (CIDR) values for all networks it uses cannot overlap each other.
- The subnet defined by a given node's Admin Network IP/mask (CIDR) value cannot overlap any subnet in the Grid Network Subnet List.
- The subnet defined by a given node's Client Network IP/mask (CIDR) value cannot overlap any subnet in the Grid Network Subnet List, nor any subnet present in that node's Admin Network External Subnet List (A-ESL).
- The subnet defined by the Grid Network IP/mask (CIDR) value of any given node in the grid must be contained within at least one subnet present in the Grid Network Subnet List.
- No subnet in the Grid Network Subnet List can overlap with any subnet in any node's Admin Network External Subnet List (A-ESL).

Gateway guidelines

- If set, the gateway for a given network must lie within the subnet defined by the node's IP/mask (CIDR) value for that network.
- If you configure an interface using static addressing, you must specify a gateway address other than 0.0.0.0.
- If you do not have a gateway, best practice is to set the gateway address to be equal to the Grid Network IP address or to the .1 address on the subnet.

General guidelines

- At least one NTP server must be reachable by the primary Admin Node, using the networking configuration you specified when deploying the primary Admin Node.
- If you are not ready to configure the optional Admin and Client Networks during deployment, you can configure these networks when you approve grid nodes during the configuration steps.
- Admin Nodes must always be secured from untrusted clients, such as those on the open internet. You must ensure that no untrusted client can access any Admin Node on the Grid Network, the Admin Network, or the Client Network.
If you are using a Client Network, you can help secure StorageGRID from hostile attacks by accepting inbound client traffic only on explicitly configured endpoints. See the information about managing untrusted Client Networks in the instructions for administering StorageGRID.
- Admin Nodes and Gateway Nodes that you intend to add to high availability groups must be configured with a static IP address.

Docker containerized deployments

Docker-related network configuration is not required in a StorageGRID system.

However, bond or bridge network devices cannot be used directly as the container network interface, so use a non-bond device as the network interface in the node configuration file. See the instructions for creating node configuration files for more information.

Using network address translation (NAT)

You can use network address translation (NAT) between external clients and grid nodes, such as to provide a public IP address for a Gateway Node. However, you must not use NAT on the Grid Network between grid nodes or between StorageGRID sites. When you use private IPv4 addresses for the Grid Network, those addresses must be directly routable from every grid node at every site. Using NAT to bridge a public network segment is supported only when you employ a tunneling

application that is transparent to all nodes in the grid, meaning the grid nodes require no knowledge of public IP addresses.

Post-configuration guidelines

After completing configuration:

- If DHCP was used to assign IP addresses, you should configure a DHCP reservation for each IP address on these networks. You can only set up DHCP during the deployment phase. You cannot set up DHCP during configuration.
 - Attention:** Because nodes reboot when their IP addresses change, outages can occur if a DHCP address change affects multiple nodes at once.
- You must use the IP address change procedures if you want to change IP addresses, subnet masks, and default gateways for a grid node. See information about configuring IP addresses in the recovery and maintenance instructions.
- 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 re-establish these connections.

For more information about StorageGRID system architecture and topology, review the networking topics in the *Grid primer*.

Related tasks

[Approving pending grid nodes](#) on page 54

Related information

[Recovery and maintenance](#)

[Grid primer](#)

Network installation and provisioning

You must understand how the Grid Network and the optional Admin and Client Networks are used during node deployment and grid configuration.

When you first deploy a node, you must attach the node to the Grid Network and ensure it has access to the primary Admin Node. If the Grid Network is isolated, you can configure the Admin Network on the primary Admin Node for configuration and installation access from outside the Grid Network. A Grid Network with a gateway configured becomes the default gateway for a node during deployment. The default gateway allows grid nodes on separate subnets to communicate with the primary Admin Node before the grid has been configured.

After the nodes are deployed, they register themselves with the primary Admin Node using the Grid Network. You can then use the Grid Manager, the `configure-storagegrid.py` Python script, or the Installation API to configure the grid and approve the registered nodes. During grid configuration, you can configure multiple grid subnets. Static routes to these subnets through the Grid Network gateway will be created on each node when you complete grid configuration. If necessary, subnets containing NTP servers or requiring access to the Grid Manager or API can also be configured as grid subnets.

If you want to disable the Admin Network or Client Network, you can remove the configuration from them during the node approval process.

Note: When using the Client Network, keep in mind that a node's default gateway will switch from the Grid Network to the Client Network when you complete the grid configuration steps. For all nodes, you must ensure that the node does not lose access to external NTP servers when the

gateway switches. For Admin Nodes, you must also ensure that browsers or API clients do not lose access to the Grid Manager. To maintain access, perform one of the following steps:

- When configuring the node, route Grid Manager traffic (Admin Nodes only) and NTP traffic through the Admin Network.
- Add subnets to the Grid Network Subnet List (GNSL) that include the IPs of remote clients and servers that should communicate with the grid over the Grid Network.
- Ensure that both the Grid and Client Network gateways can route traffic to and from the external NTP servers and browsers or other Grid Manager API clients.

If you are creating...	Behavior	Recommended configuration
Grid Network only	All Grid, Admin, and Client traffic flows over the Grid Network. The Grid Network gateway is the node default gateway.	
Grid Network and Admin Network	Grid and Client traffic flows over the Grid Network. Administrative traffic flows over the Admin Network. The Grid Network gateway is the node default gateway.	
Grid Network and Client Network (no Admin Network)	When a node is deployed, the Grid Network gateway is the node default gateway. Subnets providing access to the Grid Manager and NTP servers should be included as Grid Network subnets during configuration. When you complete the grid configuration steps, the Client Network gateway becomes the node default gateway.	Allow NTP and installer client access through both the Grid and Client Network gateways. <i>or</i> Add the NTP or installer Client subnets, or both, as Grid Networks.
All three networks (Grid, Admin, and Client)	When a node is deployed, the Grid Network gateway is the node default gateway. Subnets providing access to the Grid Manager and NTP servers should be included on the Grid Network subnets or as Admin Network subnets during configuration. When you complete the grid configuration steps, the Client Network gateway becomes the node default gateway.	Allow NTP and installer client access through both the Grid and Client Network gateways. <i>or</i> Add the NTP or installer client subnets, or both, as Grid Networks (so explicit routes will be created). <i>or</i> Add NTP and installer client subnets to the Admin Network External Subnet List (AESL).

If you are creating...	Behavior	Recommended configuration
Client Network, but at a later time	Subnets providing access to the Grid Manager and NTP servers should be included as Grid Networks or as Admin subnets. The Client Network gateway will become the node default gateway	Allow NTP and installer client access through both the Grid and Client Network gateways. <i>or</i> Add the NTP or installer client subnets, or both, as Grid Networks (so explicit routes will be created). <i>or</i> Add NTP and installer client subnets to the AESL.

Host network configuration

Before starting your StorageGRID deployment, determine which networks (Grid, Admin, Client) each node will use. You must ensure that all networks are extended to the appropriate physical or virtual hosts and that each host has sufficient bandwidth.

If you are using physical hosts to support grid nodes, extend all networks and subnets used by each node at the site to all physical hosts at that site. This strategy simplifies host configuration and enables future node migration.

Note: When configuring host network connections, you must also obtain an IP address for the physical host itself. You must also open the required ports to the host.

Note: A physical interface on the host can be used by the host itself and one or more nodes running on the host. Any IP addresses assigned to the host or nodes using this interface must be unique.

The following table provides the minimum bandwidth recommendations for each type of StorageGRID node and each type of network. You must provision each physical or virtual host with sufficient network bandwidth to meet the aggregate minimum bandwidth requirements for the total number and type of StorageGRID nodes you plan to run on that host.

Type of node	Type of network		
	Grid	Admin	Client
Admin	10 Gbps	1 Gbps	1 Gbps
Gateway	10 Gbps	1 Gbps	10 Gbps
Storage	10 Gbps	1 Gbps	10 Gbps
Archive	10 Gbps	1 Gbps	10 Gbps

Note: This table does not include SAN bandwidth, which is required for access to shared storage. If you are using shared storage accessed over Ethernet (iSCSI or FCoE), you should provision separate physical interfaces on each host to provide sufficient SAN bandwidth. To avoid introducing a bottleneck, SAN bandwidth for a given host should roughly match aggregate Storage Node network bandwidth for all Storage Nodes running on that host.

Use the table to determine the minimum number of network interfaces to provision on each host, based on the number and type of StorageGRID nodes you plan to run on that host.

For example, suppose you want to use the following configuration:

- Run one Admin Node, one Gateway Node, and one Storage Node on a single host

- Connect the Grid and Admin Networks on the Admin Node (requires $10 + 1 = 11$ Gbps)
- Connect the Grid and Client Networks on the Gateway Node (requires $10 + 10 = 20$ Gbps)
- Connect the Grid Network on the Storage Node (requires 10 Gbps)

In this scenario, you should provide a minimum of $11 + 20 + 10 = 41$ Gbps of network bandwidth, which could be met by two 40 Gbps interfaces or five 10 Gbps interfaces, potentially aggregated into trunks and then shared by the three or more VLANs carrying the Grid, Admin, and Client subnets local to the physical data center containing the host.

For some recommended ways of configuring physical and network resources on the hosts in your StorageGRID cluster to prepare for your StorageGRID grid deployment, see “Configuring the host network.”

Related tasks

[Configuring the host network](#) on page 29

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 “Automating the installation” before beginning the installation process.

Related concepts

[Overview of installation REST APIs](#) on page 68

Related tasks

[Automating the installation](#) on page 63

Internal grid node communications

The following ports must be accessible to grid nodes on the Grid Network. Ensure that the required ports for the grid node type are open on the server.

The StorageGRID internal firewall only allows incoming connections to these ports (excluding 22, 80, 123, and 443) on the Grid Network.

If enterprise networking policies restrict the availability of any of these ports, you can remap ports using a configuration file setting.

Note: NetApp recommends that you enable Internet Control Message Protocol (ICMP) traffic between grid nodes. Allowing ICMP traffic can improve failover performance when a grid node cannot be reached.

In addition to ICMP and the ports listed in the table, StorageGRID uses the Virtual Router Redundancy Protocol (VRRP). VRRP is an internet protocol that uses IP number 112. StorageGRID uses VRRP in unicast mode only, as supported by the Keepalived software. VRRP is required only if high-availability (HA) groups are configured. For more information about Keepalived, reference the Keepalived documentation.

Port	TCP or UDP	From	To	Details
22	TCP	Primary Admin Node	All nodes	For maintenance procedures, the primary Admin Node must be able to communicate with all other nodes using SSH on port 22. Allowing SSH traffic from other nodes is optional.
80	TCP	Appliances	Primary Admin Node	Used by StorageGRID appliances to communicate with the primary Admin Node to start the installation.
123	UDP	All nodes	All nodes	Network time protocol service. Every node synchronizes the time with every other node using NTP.
443	TCP	All nodes	Primary Admin Node	Used for communicating status to the primary Admin Node during installation and other maintenance procedures.
1139	TCP	Storage Nodes	Storage Nodes	Internal traffic between Storage Nodes.
1501	TCP	All nodes	Storage Nodes with ADC	Reporting, auditing, and configuration internal traffic.
1502	TCP	All nodes	Storage Nodes	S3- and Swift-related internal traffic.
1504	TCP	All nodes	Admin Nodes	NMS service reporting and configuration internal traffic.
1505	TCP	All nodes	Admin Nodes	AMS service internal traffic.
1506	TCP	All nodes	All nodes	Server status internal traffic.
1507	TCP	All nodes	Gateway Nodes	Load balancer internal traffic.
1508	TCP	All nodes	Primary Admin Node	Configuration management internal traffic.
1509	TCP	All nodes	Archive Nodes	Archive Node internal traffic.
1511	TCP	All nodes	Storage Nodes	Metadata internal traffic.
5353	UDP	All nodes	All nodes	Optionally used for full-grid IP changes and for primary Admin Node discovery during installation, expansion, and recovery.

Port	TCP or UDP	From	To	Details
7001	TCP	Storage Nodes	Storage Nodes	Cassandra TLS inter-node cluster communication.
7443	TCP	All Nodes	Admin Nodes	Internal traffic for maintenance procedures and error reporting.
9042	TCP	Storage Nodes	Storage Nodes	Cassandra client port.
9999	TCP	All nodes	All nodes	Internal traffic for multiple services. Includes maintenance procedures, metrics, and networking updates.
11139	TCP	Archive/Storage Nodes	Archive/Storage Nodes	Internal traffic between Storage Nodes and Archive Nodes.
18000	TCP	Admin/Storage Nodes	Storage Nodes with ADC	Account service internal traffic.
18001	TCP	Admin/Storage Nodes	Storage Nodes with ADC	Identity Federation internal traffic.
18002	TCP	Admin/Storage Nodes	Storage Nodes	Internal API traffic related to object protocols.
18003	TCP	Admin/Storage Nodes	Storage Nodes with ADC	Platform services internal traffic.
18017	TCP	Admin/Storage Nodes	Storage Nodes	Data Mover service internal traffic for Cloud Storage Pools.
18082	TCP	Admin/Storage Nodes	Storage Nodes	S3-related internal traffic.
18083	TCP	All nodes	Storage Nodes	Swift-related internal traffic.
18200	TCP	Admin/Storage Nodes	Storage Nodes	Additional statistics about client requests.
19000	TCP	Admin/Storage Nodes	Storage Nodes with ADC	Keystone service internal traffic.

External communications

Clients need to communicate with grid nodes and, by extension, the servers that host them in order to ingest and retrieve content. The ports used depends on the protocols chosen to ingest and retrieve content.

If enterprise networking policies restrict the availability of any of the ports used for traffic into or out of the nodes, you can remap ports when deploying nodes.

The following table shows the ports used for traffic into the nodes.

Note: This list does not include ports that might be configured as load balancer endpoints. For more information, see the section on configuring load balancer endpoints.

Port	TCP or UDP	Protocol	From	To	Details
22	TCP	SSH	Service laptop	All nodes	SSH or console access is required for procedures with console steps. Optionally, you can use port 2022 instead of 22.
25	TCP	SMTP	Admin Nodes	Email server	Used for alerts and email-based AutoSupport. You can override the default port setting of 25 using the Email Servers page.
53	TCP/UDP	DNS	All nodes	DNS servers	Used for domain name system.
67	UDP	DHCP	All nodes	DHCP service	Optionally used to support DHCP-based network configuration. The dhclient service does not run for statically-configured grids.
68	UDP	DHCP	DHCP service	All nodes	Optionally used to support DHCP-based network configuration. The dhclient service does not run for grids that use static IP addresses.
80	TCP	HTTP	Browser	Admin Nodes	Port 80 redirects to port 443 for the Admin Node user interface.
			Browser	Appliances	Port 80 redirects to port 8443 for the StorageGRID appliance's install/maintenance user interface.
			Storage Nodes with ADC	AWS	Used for platform services messages sent to AWS or other external services that use HTTP. Tenants can override the default HTTP port setting of 80 when creating an endpoint.
			Storage Nodes	AWS	Cloud Storage Pools requests sent to AWS targets that use HTTP. Grid administrators can override the default HTTP port setting of 80 when configuring a Cloud Storage Pool.
111	TCP/UDP	RPCBind	NFS client	Admin Nodes	Used by NFS-based audit export (portmap).
123	UDP	NTP	Primary NTP nodes	External NTP	Network time protocol service. Nodes selected as primary NTP sources also synchronize clock times with the external NTP time sources.
137	UDP	NetBIOS	SMB client	Admin Nodes	Used by SMB-based audit export for clients that require NetBIOS support.

Port	TCP or UDP	Protocol	From	To	Details
138	UDP	NetBIOS	SMB client	Admin Nodes	Used by SMB-based audit export for clients that require NetBIOS support.
139	TCP	SMB	SMB client	Admin Nodes	Used by SMB-based audit export for clients that require NetBIOS support.
161	TCP/UDP	SNMP	SNMP client	All nodes	Used for SNMP polling. All nodes provide basic information; Admin Nodes additionally provide alarm data. Defaults to UDP port 161 when configured. Note: For information about using SNMP with StorageGRID, contact your NetApp account representative.
162	TCP/UDP	SNMP Notifications	All nodes	Notification destinations	Outbound SNMP notifications/traps default to UDP port 162. Note: For information about using SNMP with StorageGRID, contact your NetApp account representative.
389	TCP/UDP	LDAP	Storage Nodes with ADC	Active Directory/ LDAP	Used for connecting to an Active Directory or LDAP server for Identity Federation.
443	TCP	HTTPS	Browser	Admin Nodes	Used by web browsers and management API clients for accessing the Grid Manager and Tenant Manager.
			Admin Nodes	Active Directory	Used by Admin Nodes connecting to Active Directory if single sign-on (SSO) is enabled.
			Archive Nodes	AWS S3	Used for accessing AWS S3 from Archive Nodes.
			Storage Nodes with ADC	AWS	Used for platform services messages sent to AWS or other external services that use HTTPS. Tenants can override the default HTTP port setting of 443 when creating an endpoint.
			Storage Nodes	AWS	Cloud Storage Pools requests sent to AWS targets that use HTTPS. Grid administrators can override the default HTTPS port setting of 443 when configuring a Cloud Storage Pool.
445	TCP	SMB	SMB client	Admin Nodes	Used by SMB-based audit export.
903	TCP	NFS	NFS client	Admin Nodes	Used by NFS-based audit export (<code>rpc.mountd</code>).

Port	TCP or UDP	Protocol	From	To	Details
2022	TCP	SSH	Service laptop	All nodes	SSH or console access is required for procedures with console steps. Optionally, you can use port 2022 can instead of 22.
2049	TCP	NFS	NFS client	Admin Nodes	Used by NFS-based audit export (nfs).
8022	TCP	SSH	Service laptop	All nodes	SSH on port 8022 grants access to the base operating system for support and troubleshooting. This port is not required to be accessible between grid nodes or during normal operations.
8082	TCP	HTTPS	S3 clients	Gateway Nodes	S3-related external traffic to Gateway Nodes (HTTPS).
8083	TCP	HTTPS	Swift clients	Gateway Nodes	Swift-related external traffic to Gateway Nodes (HTTPS).
8084	TCP	HTTP	S3 clients	Gateway Nodes	S3-related external traffic to Gateway Nodes (HTTP).
8085	TCP	HTTP	Swift clients	Gateway Nodes	Swift-related external traffic to Gateway Nodes (HTTP).
8443	TCP	HTTPS	Browser	Admin Nodes	Optional. Used by web browsers and management API clients for accessing the Grid Manager. Can be used to separate Grid Manager and Tenant Manager communications.
9022	TCP	SSH	Service laptop	Appliances	Grants access to StorageGRID appliances in pre-configuration mode for support and troubleshooting. This port is not required to be accessible between grid nodes or during normal operations.
9443	TCP	HTTPS	Browser	Admin Nodes	Optional. Used by web browsers and management API clients for accessing the Tenant Manager. Can be used to separate Grid Manager and Tenant Manager communications.
18082	TCP	HTTPS	S3 clients	Storage Nodes	S3-related external traffic to Storage Nodes (HTTPS).
18083	TCP	HTTPS	Swift clients	Storage Nodes	Swift-related external traffic to Storage Nodes (HTTPS).
18084	TCP	HTTP	S3 clients	Storage Nodes	S3-related external traffic to Storage Nodes (HTTP).
18085	TCP	HTTP	Swift clients	Storage Nodes	Swift-related external traffic to Storage Nodes (HTTP).

Networking and ports for platform services and Cloud Storage Pools

If you plan to use platform services or Cloud Storage Pools, you must configure grid networking and ports to ensure that the destination endpoints can be reached.

To ensure that platform services messages can be delivered to the configured endpoints, you must configure the network or networks containing ADC Storage Nodes. Similarly, to ensure that ILM rules can move objects to and from a specified Cloud Storage Pool, you must configure the network or networks that contain your grid's Storage Nodes.

By default, platform services and Cloud Storage Pool communications use the following ports:

- **80**: For endpoint URIs that begin with `http`
- **443**: For endpoint URIs that begin with `https`
- **8082**: For external HTTPS communications received by a Gateway Node
- **8084**: For external HTTP communications received by a Gateway Node

A different port can be specified when the endpoint is created or edited.

If you use a non-transparent proxy server, you must also configure proxy settings to allow messages to be sent to external endpoints, such as an endpoint on the internet. See *administering StorageGRID* to learn how to configure proxy settings.

Related information

[Administering StorageGRID](#)

[Using tenant accounts](#)

Storage requirements

You must understand the storage requirements for StorageGRID nodes, so you can provide the appropriate type and amount of storage for each host.

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 metadata** – Performance-tier (10K SAS or SSD) storage for per-node persistent storage of system metadata 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 are not supported. You can use shared or local RAID storage for any of the storage categories; however, if you want to use StorageGRID's node migration capability, you must store both system metadata and object data on shared storage.

Performance requirements

The performance of the volumes used for the container pool, system metadata, 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 metadata, 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 AFF storage

If the host for the StorageGRID node will use storage assigned from a NetApp AFF 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.

Attention: 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.

Attention: In a production deployment, do not 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.

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

LUN purpose	Storage category	Number of LUNs	Minimum size/LUN
Docker storage pool	Container pool	1	Total number of nodes × 100 GB
/var/local volume	System metadata	1 for each node on this host	90 GB
Storage Node	Object data	3 for each Storage Node on this host Note: A Storage Node can have 1 to 16 storage volumes; at least 3 storage volumes are recommended.	4,000 GB See “Storage requirements for Storage Nodes” below for more information.
Admin Node audit logs	System metadata	1 for each Admin Node on this host	200 GB
Admin Node tables	System metadata	1 for each Admin Node on this host	200 GB

Attention: 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. As a rule of thumb, 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.

Type of node	Container pool	System metadata	Object data
Storage Node	100 GB	90 GB	4,000 GB
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 metadata and transaction logs, and 12 TB of capacity-tier storage for object data.

Type of node	LUN purpose	Number of LUNs	LUN size
	Docker storage pool	1	300 GB (100 GB/node)
Storage Node	/var/local volume	1	90 GB
	Object data	3	4,000 GB
Admin Node	/var/local volume	1	90 GB
	Admin Node audit logs	1	200 GB
	Admin Node tables	1	200 GB
Gateway Node	/var/local volume	1	90 GB
Total		9	Container pool: 300 GB System metadata: 670 GB Object data: 12,000 GB

Storage requirements for Storage Nodes

A Storage Node can have 1 to 16 storage volumes; however, 3 or more storage volumes are recommended. Each storage volume should be 4 TB or larger.

You must assign at least 4 TB to volume 0 of each Storage Node. StorageGRID reserves up to 4 TB of space on volume 0 (the Metadata Reserved Space, or CAWM) for object metadata and for essential database operations, such as compaction and repair.

If you assign additional space to volume 0, that space is used for object storage. Assigning more than 4 TB to volume 0 is recommended.

Note: 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.

Note: 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.

Related concepts

[Node container migration requirements](#) on page 26

Related information

[Recovery and maintenance](#)

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 simply 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 StorageGRID's node migration feature. In this case, you can ignore these additional requirements.

Before performing migration or hypervisor maintenance, shut down the nodes gracefully. See the recovery and maintenance 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 are not 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 procedure for shutting down a grid node in the recovery and maintenance instructions.

Consistent network interface names

In order 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

In order 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 metadata 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 metadata 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
```

Related tasks

[Configuring the host network](#) on page 29

Related information

[Recovery and maintenance](#)

Web browser requirements

You must use a supported web browser.

Web browser	Minimum supported version
Google Chrome	74
Microsoft Internet Explorer	11 (Native Mode)
Mozilla Firefox	67

You should set the browser window to a recommended width.

Browser width	Pixels
Minimum	1024
Optimum	1280

Preparing the hosts

You must complete the following steps to prepare your physical or virtual hosts for StorageGRID. Note that you can automate many or all of these steps using standard server configuration frameworks such as Ansible, Puppet, or Chef.

Steps

1. [Installing Linux](#) on page 28
2. [Configuring the host network](#) on page 29
3. [Configuring host storage](#) on page 31
4. [Installing Docker](#) on page 33
5. [Installing StorageGRID host services](#) on page 35

Related concepts

[Automating the installation and configuration of the StorageGRID host service](#) on page 63

Installing Linux

You must install Red Hat Enterprise Linux or CentOS Linux on all grid hosts. Use the NetApp Interoperability Matrix Tool to get a list of supported versions.

Steps

1. Install Linux on all physical or virtual grid hosts according to the distributor's instructions or your standard procedure.
 - Note:** If you are using the standard Linux installer, NetApp recommends selecting the "compute node" software configuration, if available, or "minimal install" base environment. Do not 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.

Attention: Failing to disable swap entirely can severely lower performance.

Related information

[NetApp Interoperability Matrix Tool](#)

Configuring the host network

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 following topics, which provide information you need to accomplish this task:

[Networking requirements](#) on page 10

[Node container migration requirements](#) on page 26

Attention: If you are using VMs as hosts, you must allow all interfaces to receive and transmit data for MAC addresses other than the ones assigned by the hypervisor. To do this, 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.

Note: 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 cannot be shared between nodes or between the host and any node.

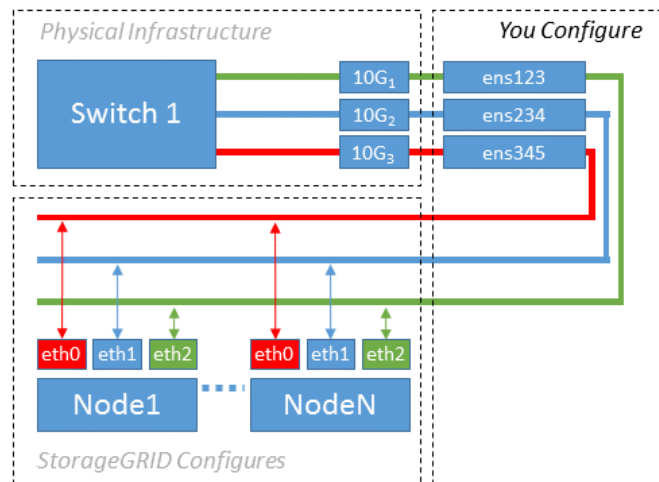
You can use the same host network interface to provide the Grid Network interface for all StorageGRID nodes on the host; you can use a different host network interface for each node; or you can do something in between. However, you would not typically provide the same host network interface as both the 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 simply 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 do not need to use either of these examples; you can use any approach that meets your needs.

Note: Do not use bond or bridge devices as the network interface specified in the node configuration file, as this could prevent node startup due to an issue with networking in the container namespace. Either put a VLAN (or other virtual interface) on top of the bond device, or use a bridge and virtual Ethernet (veth) pair. .

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

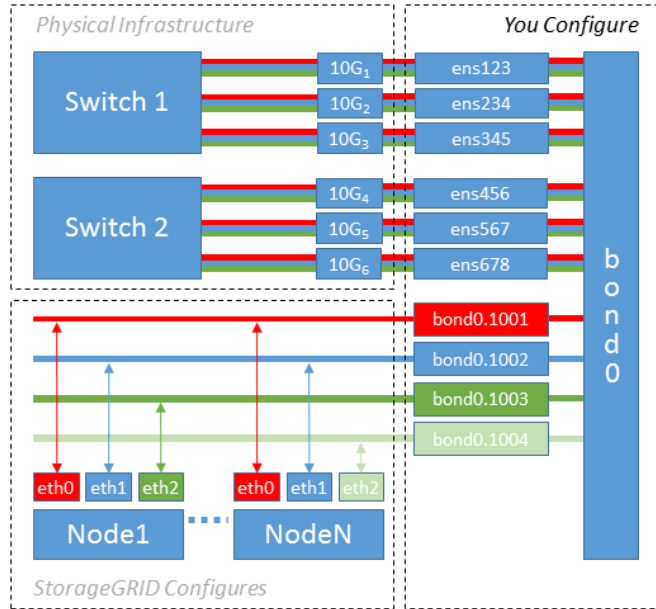
Example 2 assumes you are familiar with bonding network interfaces and with creating VLAN interfaces on the Linux distribution you are using.

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 references

[Sample /etc/sysconfig/network-scripts](#) on page 73

Configuring 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 requirements](#) on page 23

[Node container migration requirements](#) on page 26

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 Metadata 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.

Note: You do not 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.

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

For example:

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



```
}
...
```

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.

Note: 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 tasks

[Installing Docker](#) on page 33

Installing Docker

The StorageGRID system runs on Red Hat Enterprise Linux or CentOS as a collection of Docker containers. Before you can install StorageGRID, you must install Docker.

About this task

When using Docker with Red Hat Enterprise Linux, you must use either the overlay2 storage driver or the devicemapper storage driver.

- If the Red Hat Enterprise Linux version is 7.5 or later, the overlay2 storage driver is the recommended storage driver.
- If the Red Hat Enterprise Linux version is earlier than 7.5, the devicemapper storage driver is the recommended storage driver. On production systems using devicemapper, you must use `direct-lvm` mode.

For details about the overlay2 and devicemapper storage drivers, see the Docker documentation and your installation instructions for Red Hat Enterprise Linux.

Steps

1. If you are installing Docker on Red Hat Enterprise Linux, activate the RHEL Extras channel.

```
sudo subscription-manager repos --enable rhel-7-server-extras-rpms
```

Note: The CentOS Extras repository is enabled by default.

2. Install Docker.

```
sudo yum install 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 1.10.3 or later.

```
Client:
Version: 1.10.3
API version: 1.22
Package version: docker-common-1.10.3-46.el7.14.x86_64
```

```

Go version: go1.6.2
Git commit: 5206701-unsupported
Built: Mon Aug 29 14:00:01 2016
OS/Arch: linux/amd64

Server:
Version: 1.10.3
API version: 1.22
Package version: docker-common-1.10.3-46.el7.14.x86_64
Go version: go1.6.2
Git commit: 5206701-unsupported
Built: Mon Aug 29 14:00:01 2016
OS/Arch: linux/amd64

```

4. Determine which storage driver you are using:

- If the Red Hat Enterprise Linux version is 7.5 or later, you have the overlay2 storage driver, which does not need to be configured. Go to step 5.
- If the Red Hat Enterprise Linux version is earlier than 7.5, you have the devicemapper storage driver. By default, this storage driver is configured in `loopback-lvm` mode, which is not suitable for production. Follow these steps to configure the storage driver in `direct-lvm` mode.

- a. After completing the Docker installation on a host, enter the command:

```
sudo docker info
```

If the output of this command shows that the Data loop file and Metadata loop file are underneath `/var/lib/docker/devicemapper/devicemapper`, then Docker is configured in `loopback-lvm` mode.

- b. Reconfigure Docker in `direct-lvm` mode before continuing with this installation process.

You will use the Docker storage pool volume you allocated for each host when configuring host storage as the block device backing your `dm.thinpooldev` container storage pool. The referenced instructions use the example block device name `/dev/xvdf`. You must replace `/dev/xvdf` with the persistent name of the Docker storage pool volume allocated to the host on which you are reconfiguring docker.

Attention: If you use shared storage for the Docker storage pool volumes, you must use a different physical volume for each host. System data will be corrupted if two hosts attempt to use the same Docker storage pool volume for container storage.

5. Run the following two commands to enable and start Docker:

```
sudo systemctl enable docker
```

```
sudo systemctl start docker
```

Related tasks

[Configuring host storage](#) on page 31

Related information

[Docker Documentation: Docker and the OverlayFS storage driver](#)

[Docker Documentation: Docker and the Device Mapper storage driver](#)

Installing 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
```

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

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

Deploying virtual grid nodes

To deploy virtual grid nodes on Red Hat Enterprise Linux or CentOS hosts, you create node configuration files for all nodes, validate the files, and start the StorageGRID host service, which starts the nodes. If you need to deploy any StorageGRID appliance Storage Nodes, see the installation and maintenance instructions for the appliance after you have deployed all virtual nodes.

Steps

1. [Creating node configuration files](#) on page 36
2. [Validating the StorageGRID configuration](#) on page 48
3. [Starting the StorageGRID host service](#) on page 49

Related information

[SG1000 appliance installation and maintenance](#)

[SG5600 appliance installation and maintenance](#)

[SG5700 appliance installation and maintenance](#)

[SG6000 appliance installation and maintenance](#)

Creating node configuration files

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 are not 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 do not 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>-<node-type><node-number>.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. The middle column indicates whether the key is Required (R), a Best Practice (BP), or Optional (O).

Key	R, BP, or O?	Value
ADMIN_IP	BP	Grid Network IPv4 address of the primary Admin Node for the grid to which this node belongs. Use the same value you specified for <code>GRID_NETWORK_IP</code> for the grid node with <code>NODE_TYPE = VM_Admin_Node</code> and <code>ADMIN_ROLE = Primary</code> . If you omit this parameter, the node attempts to discover a primary Admin Node using mDNS. See “How grid nodes discover the primary Admin Node.” Note: This value is ignored, and might be prohibited, on the primary Admin Node.
<code>ADMIN_NETWORK_CONFIG</code>	O	DHCP, STATIC, or DISABLED (Defaults to DISABLED if not specified.)
<code>ADMIN_NETWORK_ESL</code>	O	Comma-separated list of subnets in CIDR notation to which this node should communicate via the Admin Network gateway. Example: <code>172.16.0.0/21,172.17.0.0/21</code>

Key	R, BP, or O?	Value
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> <ul style="list-style-type: none"> • 1.1.1.1 • 10.224.4.81
ADMIN_NETWORK_IP	O	<p>IPv4 address of this node on the Admin Network. This key is only required when ADMIN_NETWORK_CONFIG = STATIC; do not specify it for other values.</p> <p>Examples:</p> <ul style="list-style-type: none"> • 1.1.1.1 • 10.224.4.81
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. 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; do not specify it for other values.</p> <p>Examples:</p> <ul style="list-style-type: none"> • 255.255.255.0 • 255.255.248.0
ADMIN_NETWORK_MTU	O	<p>The maximum transmission unit (MTU) for this node on the Admin Network. Do not specify if ADMIN_NETWORK_CONFIG = DHCP. If specified, the value must be between 68 and 65535. If omitted, 1400 is used.</p> <p>Examples:</p> <ul style="list-style-type: none"> • 1500 • 8192

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 <code>GRID_NETWORK_TARGET</code> or <code>CLIENT_NETWORK_TARGET</code>.</p> <p>Note: Do not 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> <ul style="list-style-type: none"> • bond0.1002 • ens256
ADMIN_NETWORK_TARGET_TYPE	O	<p>Interface (This is the only supported value.)</p>
ADMIN_ROLE	R	<p>Primary or Non-Primary</p> <p>This key is only required when <code>NODE_TYPE = VM_Admin_Node</code>; do not specify it for other node types.</p>
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 <code>NODE_TYPE = VM_Admin_Node</code>; do not specify it for other node types.</p> <p>Examples:</p> <ul style="list-style-type: none"> • /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

Key	R, BP, or O?	Value
<p>BLOCK_DEVICE_RANGEDB_00 BLOCK_DEVICE_RANGEDB_01 BLOCK_DEVICE_RANGEDB_02 BLOCK_DEVICE_RANGEDB_03 BLOCK_DEVICE_RANGEDB_04 BLOCK_DEVICE_RANGEDB_05 BLOCK_DEVICE_RANGEDB_06 BLOCK_DEVICE_RANGEDB_07 BLOCK_DEVICE_RANGEDB_08 BLOCK_DEVICE_RANGEDB_09 BLOCK_DEVICE_RANGEDB_10 BLOCK_DEVICE_RANGEDB_11 BLOCK_DEVICE_RANGEDB_12 BLOCK_DEVICE_RANGEDB_13 BLOCK_DEVICE_RANGEDB_14 BLOCK_DEVICE_RANGEDB_15</p>	R	<p>Path and name of the block device special file this node will use for persistent object storage. This key is only required for nodes with <code>NODE_TYPE = VM_Storage_Node</code>; do not specify it for other node types.</p> <p>Only <code>BLOCK_DEVICE_RANGEDB_00</code> is required; the rest are optional. The block device specified for <code>BLOCK_DEVICE_RANGEDB_00</code> must be at least 4 TB; the others can be smaller.</p> <p>Note: Do not leave gaps. If you specify <code>BLOCK_DEVICE_RANGEDB_05</code>, you must also specify <code>BLOCK_DEVICE_RANGEDB_04</code>.</p> <p>Examples:</p> <ul style="list-style-type: none"> • <code>/dev/disk/by-path/ pci-0000:03:00.0-scsi-0:0:0:0</code> • <code>/dev/disk/by-id/ wwn-0x600a09800059d6df000060d7 57b475fd</code> • <code>/dev/mapper/sgws-snl-rangedb-0</code>
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 <code>NODE_TYPE = VM_Admin_Node</code>; do not specify it for other node types.</p> <p>Examples:</p> <ul style="list-style-type: none"> • <code>/dev/disk/by-path/ pci-0000:03:00.0-scsi-0:0:0:0</code> • <code>/dev/disk/by-id/ wwn-0x600a09800059d6df000060d7 57b475fd</code> • <code>/dev/mapper/sgws-adml-tables</code>

Key	R, BP, or O?	Value
BLOCK_DEVICE_VAR_LOCAL	R	Path and name of the block device special file this node will use for its <code>/var/local</code> persistent storage. Examples: <ul style="list-style-type: none"> • <code>/dev/disk/by-path/pci-0000:03:00.0-scsi-0:0:0:0</code> • <code>/dev/disk/by-id/wwn-0x600a09800059d6df000060d757b475fd</code> • <code>/dev/mapper/sgws-snl-var-local</code>
CLIENT_NETWORK_CONFIG	O	DHCP, STATIC, or DISABLED (Defaults to DISABLED if not specified.)
CLIENT_NETWORK_GATEWAY	O	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. Examples: <ul style="list-style-type: none"> • 1.1.1.1 • 10.224.4.81
CLIENT_NETWORK_IP	O	IPv4 address of this node on the Client Network. This key is only required when CLIENT_NETWORK_CONFIG = STATIC; do not specify it for other values. Examples: <ul style="list-style-type: none"> • 1.1.1.1 • 10.224.4.81
CLIENT_NETWORK_MAC	O	The MAC address for the Client Network interface in the container. This field is optional. If omitted, a MAC address will be generated automatically. Must be 6 pairs of hexadecimal digits separated by colons. Example: <code>b2:9c:02:c2:27:20</code>

Key	R, BP, or O?	Value
CLIENT_NETWORK_MASK	O	IPv4 netmask for this node on the Client Network. This key is only required when CLIENT_NETWORK_CONFIG = STATIC; do not specify it for other values. Examples: <ul style="list-style-type: none"> • 255.255.255.0 • 255.255.248.0
CLIENT_NETWORK_MTU	O	The maximum transmission unit (MTU) for this node on the Client Network. Do not specify if CLIENT_NETWORK_CONFIG = DHCP. If specified, the value must be between 68 and 65535. If omitted, 1400 is used. Examples: <ul style="list-style-type: none"> • 1500 • 8192
CLIENT_NETWORK_TARGET	BP	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. Note: Do not 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. 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. Examples: <ul style="list-style-type: none"> • bond0.1003 • ens423
CLIENT_NETWORK_TARGET_TYPE	O	Interface (This is only supported value.)
GRID_NETWORK_CONFIG	BP	STATIC or DHCP (Defaults to STATIC if not specified.)

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; do not specify it for other values.</p> <p>Examples:</p> <ul style="list-style-type: none"> • 1.1.1.1 • 10.224.4.81
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. 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; do not specify it for other values.</p> <p>Examples:</p> <ul style="list-style-type: none"> • 255.255.255.0 • 255.255.248.0
GRID_NETWORK_MTU	O	<p>The maximum transmission unit (MTU) for this node on the Grid Network. Do not specify if GRID_NETWORK_CONFIG = DHCP. If specified, the value must be between 68 and 65535. If omitted, 1400 is used.</p> <p>Examples:</p> <ul style="list-style-type: none"> • 1500 • 8192

Key	R, BP, or O?	Value
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 <code>ADMIN_NETWORK_TARGET</code> or <code>CLIENT_NETWORK_TARGET</code>.</p> <p>Note: Do not 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> <ul style="list-style-type: none"> • bond0.1001 • ens192
GRID_NETWORK_TARGET_TYPE	O	<p>Interface (This is the only supported value.)</p>
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, do not specify a value less than 24g (24 GB).</p> <p>The format for this field is <code><number><unit></code>, where <code><unit></code> can be b, k, m, or g.</p> <p>Examples:</p> <p>24g 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> <ul style="list-style-type: none"> • VM_Admin_Node • VM_Storage_Node • VM_Archive_Node • VM_API_Gateway

Key	R, BP, or O?	Value
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>Attention: Do not 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: <i><network type>/<protocol>/<default port used by grid node>/<new port></i>, 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>
PORT_REMAP_INBOUND	O	<p>Remaps inbound communications to the specified port. If you specify PORT_REMAP_INBOUND but do not specify a value for PORT_REMAP, outbound communications for the port are unchanged.</p> <p>Attention: Do not remap the ports you are planning to use to configure load balancer endpoints.</p> <p>The format used is: <i><network type>/<protocol:>/<remapped port >/<default port used by grid node></i>, where 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>

Related concepts

[Internal grid node communications](#) on page 17

[External communications](#) on page 19

Related references

[How grid nodes discover the primary Admin Node](#) on page 46

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 cannot acquire the primary Admin Node's IP address directly.

Attention: 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.

Note: 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-adml.conf`

Example file contents:

```

NODE_TYPE = VM_Admin_Node
ADMIN_ROLE = Primary
BLOCK_DEVICE_VAR_LOCAL = /dev/mapper/dc1-adml-var-local
BLOCK_DEVICE_AUDIT_LOGS = /dev/mapper/dc1-adml-audit-logs

```

```

BLOCK_DEVICE_TABLES = /dev/mapper/dcl-adml-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/dcl-snl.conf

Example file contents:

```

NODE_TYPE = VM_Storage_Node
ADMIN_IP = 10.1.0.2
BLOCK_DEVICE_VAR_LOCAL = /dev/mapper/dcl-snl-var-local
BLOCK_DEVICE_RANGEDB_00 = /dev/mapper/dcl-snl-rangedb-0
BLOCK_DEVICE_RANGEDB_01 = /dev/mapper/dcl-snl-rangedb-1
BLOCK_DEVICE_RANGEDB_02 = /dev/mapper/dcl-snl-rangedb-2
BLOCK_DEVICE_RANGEDB_03 = /dev/mapper/dcl-snl-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/dcl-arcl.conf

Example file contents:

```

NODE_TYPE = VM_Archive_Node
ADMIN_IP = 10.1.0.2
BLOCK_DEVICE_VAR_LOCAL = /dev/mapper/dcl-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/dcl-gw1.conf

Example file contents:

```

NODE_TYPE = VM_API_Gateway
ADMIN_IP = 10.1.0.2
BLOCK_DEVICE_VAR_LOCAL = /dev/mapper/dcl-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
```

Validating 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
```

Tip: 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/dcl-adml
WARNING: ignoring /etc/storagegrid/nodes/dcl-sn2.conf.keep
WARNING: ignoring /etc/storagegrid/nodes/my-file.txt
Checking configuration file for node dcl-adml...
ERROR: NODE_TYPE = VM_Foo_Node
      VM_Foo_Node is not a valid node type.  See *.conf.sample
ERROR: ADMIN_ROLE = Foo
      Foo is not a valid admin role.  See *.conf.sample
ERROR: BLOCK_DEVICE_VAR_LOCAL = /dev/mapper/sgws-gw1-var-local
      /dev/mapper/sgws-gw1-var-local is not a valid block device
Checking configuration file for node dcl-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 dcl-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 dcl-sn2... PASSED
Checking configuration file for node dcl-sn3... PASSED
Checking for duplication of unique values between nodes...
ERROR: GRID_NETWORK_IP = 10.1.0.4
      dcl-sn2 and dcl-sn3 have the same GRID_NETWORK_IP
ERROR: BLOCK_DEVICE_VAR_LOCAL = /dev/mapper/sgws-sn2-var-local
      dcl-sn2 and dcl-sn3 have the same BLOCK_DEVICE_VAR_LOCAL
ERROR: BLOCK_DEVICE_RANGEDB_00 = /dev/mapper/sgws-sn2-rangedb-0
      dcl-sn2 and dcl-sn3 have the same BLOCK_DEVICE_RANGEDB_00

```

Starting 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

```

For any node that returns a status of “Not-Running” or “ Stopped”, run the following command:

```

sudo storagegrid node start node-name

```

3. 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
```

Configuring the grid and completing installation

You complete installation by configuring the StorageGRID grid from the Grid Manager on the primary Admin Node.

Steps

1. [Navigating to the Grid Manager](#) on page 51
2. [Specifying the StorageGRID license information](#) on page 52
3. [Adding sites](#) on page 52
4. [Specifying Grid Network subnets](#) on page 53
5. [Approving pending grid nodes](#) on page 54
6. [Specifying Network Time Protocol server information](#) on page 58
7. [Specifying Domain Name System server information](#) on page 59
8. [Specifying the StorageGRID system passwords](#) on page 59
9. [Reviewing your configuration and completing installation](#) on page 60

Navigating 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 the following address:

`https://primary_admin_node_ip`

Note: 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. Click **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

Specifying 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 **Grid Name**.
The name is displayed as the top level in the grid topology tree after installation.

2. Click **Browse**, locate the NetApp License File (`NLFunique_id.txt`), and click **Open**.

The license file is validated, and the serial number and licensed storage capacity are displayed.

Note: 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.

The screenshot shows the NetApp StorageGRID installation wizard interface. At the top, there is a blue header with the text "NetApp® StorageGRID®" and a "Help" link. Below the header is a navigation bar with an "Install" button. A progress indicator shows eight steps: 1. License (highlighted), 2. Sites, 3. Grid Network, 4. Grid Nodes, 5. NTP, 6. DNS, 7. Passwords, and 8. Summary. The main content area is titled "License" and contains the instruction: "Enter a grid name and upload the license file provided by NetApp for your StorageGRID system." Below this instruction are four input fields: "Grid Name" with the value "Grid1", "New License File" with a "Browse" button, "License Serial Number" with the value "950719", and "Storage Capacity (TB)" with the value "240".

3. Click **Next**.

Adding sites

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

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 ✕

Site Name 2 + ✕

3. Click **Next**.

Specifying 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 via the Grid Network (for example, the subnets hosting your NTP servers).

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.

NetApp® StorageGRID® Help ▾

Install

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

Grid Network

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

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

Subnet 1 +

3. Click **Next**.

Approving pending grid nodes

You must approve each grid node before it joins the StorageGRID grid.

Before you begin

All virtual and StorageGRID appliance grid nodes must have been deployed.

Steps

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

Note: 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.

Grid Network MAC Address	Name	Type	Platform	Grid Network IPv4 Address
<input checked="" type="radio"/> 50:5b: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.

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	<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"/> x
	<input type="text" value="172.19.0.0/16"/> x
	<input type="text" value="172.21.0.0/16"/> + x

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 name of the site with which this grid node will be associated.
- **Name:** The host name that will be assigned to the node, and the name that will be displayed in the Grid Manager. The name defaults to the name you specified during node deployment, but you can change the name here 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.

Attention: 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 cannot 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

Note: The gateway is required if there are multiple grid subnets.

Note: 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.

Note: 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 cannot be configured in this Grid Manager dialog box. Instead, you must follow these steps:

- Reboot the appliance: In the Appliance Installer, select **Advanced > Reboot**.
Rebooting can take several minutes.
- Select **Configure Networking > Link Configuration** and enable the appropriate networks.
- Select **Configure Networking > IP Configuration** and configure the enabled networks.
- Return to the Home page and click **Start Installation**.
- In the Grid Manager: If the node is listed in the **Approved Nodes** table, reset the node.
- Remove the node from the **Pending Nodes** table.
- Wait for the node to reappear in the **Pending Nodes** list.
- Confirm that you can configure the appropriate networks. They should already be populated with the information you provided on the **IP Configuration** page.

For additional information, see the installation and maintenance 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.

Note: 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 cannot 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, reset 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.

For additional information, see the installation and maintenance instructions for your appliance model.

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.

Grid Network MAC Address	Name	Type	Platform	Grid Network IPv4 Address
No results found.				

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.

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

- 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**.

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

Specifying 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 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.

Note: When specifying the external NTP source for a production-level StorageGRID installation, do not 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.

Attention: 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

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

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

3. Click **Next**.

Specifying Domain Name System server information

You must specify Domain Name System (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. Specifying at least two DNS servers is recommended.

Attention: Provide two to six IP addresses for DNS servers. You should select DNS servers that each site can access locally in the event of network islanding. This is to ensure an islanded site continues to have access to the DNS service. After configuring the grid-wide DNS server list, you can further customize the DNS server list for each node. For details, see information about modifying the DNS configuration in the recovery and maintenance instructions.

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

The screenshot shows the NetApp StorageGRID installation wizard interface. At the top, there is a blue header with "NetApp StorageGRID" and a "Help" dropdown. Below the header is a progress bar with eight steps: 1. License, 2. Sites, 3. Grid Network, 4. Grid Nodes, 5. NTP, 6. DNS (highlighted in blue), 7. Passwords, and 8. Summary. Below the progress bar, the "Domain Name Service" section is visible. It contains the following text: "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." Below this text are two input fields for DNS servers. The first field is labeled "Server 1" and contains the IP address "10.224.223.130". To the right of this field is a red "x" icon. The second field is labeled "Server 2" and contains the IP address "10.224.223.136". To the right of this field are a red "+" icon and a red "x" icon.

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

3. Click **Next**.

Specifying the StorageGRID system passwords

You need to enter the passwords to use to secure your StorageGRID system.

Steps

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

You should record this password in a secure place.

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

The screenshot shows the NetApp StorageGRID installation wizard interface. At the top, there is a blue header with the text "NetApp® StorageGRID®" and a "Help" dropdown menu. Below the header is a navigation bar with an "Install" button. A progress indicator shows eight steps: 1. License, 2. Sites, 3. Grid Network, 4. Grid Nodes, 5. NTP, 6. DNS, 7. Passwords (highlighted in blue), and 8. Summary. The main content area is titled "Passwords" and contains the following text: "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." Below this text are four password input fields: "Provisioning Passphrase", "Confirm Provisioning Passphrase", "Grid Management Root User Password", and "Confirm Root User Password". Each field contains a series of dots representing masked characters. At the bottom of the form, there is a checkbox labeled "Create random command line passwords." which is currently checked.

5. If you are installing a grid for proof of concept or demo purposes, optionally deselect the **Create random command line passwords** check box.

For production deployments, random passwords should always be used for security reasons. Deselect **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.

Attention: 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**.

Reviewing your configuration and completing 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**.

Note: 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 “Network installation and provisioning” 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 cannot 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.

Attention: 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** check box, and click **Next**.

Download Recovery Package

Before proceeding, you must download the Recovery Package file. This file is necessary to recover the StorageGRID system if a failure occurs.

When the download completes, open the .zip file and confirm it includes a "gpt-backup" directory and a second .zip file. Then, extract this inner .zip file and confirm you can open the passwords.txt file.

After you have verified the contents, copy the Recovery Package file 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.

i The Recovery Package is required for recovery procedures and must be stored in a secure location.

Download Recovery Package

I have successfully downloaded and verified the Recovery Package file.

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%; background-color: #0070c0;"></div>	Starting services
dc1-g1	Site1	172.16.4.216/21	<div style="width: 100%; background-color: #0070c0;"></div>	Complete
dc1-s1	Site1	172.16.4.217/21	<div style="width: 80%; background-color: #0070c0;"></div>	Waiting for Dynamic IP Service peers
dc1-s2	Site1	172.16.4.218/21	<div style="width: 20%; background-color: #0070c0;"></div>	Downloading hotfix from primary Admin if needed
dc1-s3	Site1	172.16.4.219/21	<div style="width: 20%; background-color: #0070c0;"></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.

Related concepts

[Network installation and provisioning](#) on page 14

Automating the installation

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

Steps

1. [Automating the installation and configuration of the StorageGRID host service](#) on page 63
2. [Automating the configuration and installation of appliance Storage Nodes](#) on page 63
3. [Automating the configuration of StorageGRID](#) on page 65

Automating 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 can be prepared 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.

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 host and installs StorageGRID onto the target servers. You can customize the role or playbook as necessary.

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

Related tasks

[Preparing the hosts](#) on page 28

[Deploying virtual grid nodes](#) on page 36

Automating the configuration and installation of appliance Storage Nodes

You can use the `configure-sga.py` script to configure and install StorageGRID appliance Storage Nodes.

Before you begin

- The appliance has been installed in a rack, connected to your networks, and powered on.
- Network links and IP addresses have been configured for the appliance using the StorageGRID Appliance Installer.
- The primary Admin Node for the StorageGRID grid has been deployed, and you know its IP address.
- 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.

See the installation and maintenance instructions for your appliance.

- You know the location of the `configure-sga.py` file. The file is included in the installation archive, or you can access it by clicking **Help > Appliance Installation Script** from the StorageGRID Appliance Installer.

Steps

1. Log in to the Linux machine you are using to run the Python script.
2. For 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 four subcommands: `configure`, `install`, `monitor`, and `reboot`. To get help for each subcommand, enter the following:

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

3. To confirm the current configuration of the appliance Storage Node, enter the following:

```
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://10.224.2.32:8443 (Checking version and
connectivity)
2017/07/30 21:56:34: Performing GET on /api/versions... Received 200
2017/07/30 21:56:34: Performing GET on /api/v2/system-info...
Received 200
2017/07/30 21:56:34: Performing GET on /api/v2/admin-connection...
Received 200
2017/07/30 21:56:34: Performing GET on /api/v2/link-config...
Received 200
2017/07/30 21:56:34: Performing GET on /api/v2/networks... Received
200
2017/07/30 21:56:35: Performing GET on /api/v2/system-config...
Received 200
```

```
StorageGRID Webscale Appliance
Name: MS-SGA1
```

```
StorageGRID primary Admin Node
IP: 172.16.1.170
State: ready
Message: Connection validated at 2017-07-31T01:55:52.982161
```

Bond and VLAN configuration

```
Port bond mode: FIXED
Link speed: 10GBE

Admin Network: ENABLED
Bonding mode: no-bond
VLAN: novlan

Grid Network: ENABLED
Bonding mode: lacp
VLAN: novlan

Client Network: ENABLED
Bonding mode: lacp
VLAN: novlan
```

```
Admin Network
CIDR: 10.224.2.32/21 (Static)
```



```

MAC:          00:80:E5:43:AE:BC
Gateway:     10.224.0.1
Subnets:    192.168.8.0/24
              172.19.0.0/16
              10.0.0.0/8
              192.168.9.0/24
              172.21.0.0/16

Grid Network
CIDR:        172.16.2.32/21 (Static)
MAC:         00:A0:98:A1:85:AE
Gateway:     172.16.0.1
Subnets:    172.18.0.0/21
              172.17.0.0/21
              192.168.0.0/21

Client Network
CIDR:        47.47.2.32/21 (Static)
MAC:         00:A0:98:A1:85:AD
Gateway:     47.47.0.1

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

```

4. 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 for the primary Admin Node to 10.224.2.99, enter the following:

```
configure-sga.py configure --admin-ip 10.224.2.99 SGA-INSTALL-IP
```

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

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

6. If you are experiencing difficulties trying to configure the appliance and want to try rebooting it, enter the following:

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

Related information

[SG6000 appliance installation and maintenance](#)

[SG5700 appliance installation and maintenance](#)

[SG5600 appliance installation and maintenance](#)

Automating the configuration of StorageGRID

After deploying virtual grid nodes and appliance 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.

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

Steps

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

For example:

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

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

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

For example:

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

Result

A Recovery Package `.zip` file is generated during the configuration process, and it is downloaded to the directory from which 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.

Attention: 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.
```

Related concepts

[Overview of installation REST APIs](#) on page 68

Related tasks

[Configuring the grid and completing installation](#) on page 51

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 (JavaScript Object Notation) data format.

Attention: 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 in the event that 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:

- **hardware configuration** – 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.
- **upgrade** – Operations related to upgrading appliance firmware.
- **uploadsg** – Operations for uploading StorageGRID installation files.

Where to go next

After completing an installation, you must perform a series of integration and configuration steps. Some steps are required; others are optional.

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 or Swift API client applications you will use to upload objects to your StorageGRID system.
- When you are ready, configure the information lifecycle management (ILM) rules and ILM policy you want to use to protect object data.

Note: When you install StorageGRID, the default ILM policy, Baseline 2 Copies Policy, is active. This policy includes the stock ILM rule (Make 2 Copies), and it applies if no other policy has been activated.

- If your installation includes appliance Storage Nodes, use SANtricity software to complete the following tasks:
 - Connect to each StorageGRID appliance.
 - Verify receipt of AutoSupport data.
- If your StorageGRID system includes any Archive Nodes, configure the Archive Node's connection to the target external archival storage system.

Note: If any Archive Nodes will use Tivoli Storage Manager as the external archival storage system, you must also configure Tivoli Storage Manager.
- Review and follow the StorageGRID system hardening guidelines to eliminate security risks.

Optional tasks

- Configure mailing lists and contents of notification emails for system alarms and service status changes.
- Update grid node IP addresses if they have changed since you planned your deployment and generated the Recovery Package. See information about changing IP addresses in the recovery and maintenance instructions.
- Configure storage encryption, if required.
- Configure storage compression to reduce the size of stored objects, if required.
- Configure audit client access. You can configure access to the system for auditing purposes through an NFS or a CIFS file share. See the instructions for administering StorageGRID.

Note: Audit export through CIFS/Samba has been deprecated and will be removed in a future StorageGRID release.

Related information

[Administering StorageGRID](#)

[Grid primer](#)

[Implementing S3 client applications](#)

[Implementing Swift client applications](#)

[Monitoring and troubleshooting StorageGRID](#)

[Recovery and maintenance](#)

[SG1000 appliance installation and maintenance](#)

[SG5600 appliance installation and maintenance](#)

[SG5700 appliance installation and maintenance](#)

[SG6000 appliance installation and maintenance](#)

[StorageGRID release notes](#)

[StorageGRID system hardening](#)

[Understanding audit messages](#)

[Upgrading StorageGRID](#)

Troubleshooting 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 the instructions for monitoring and troubleshooting StorageGRID. For help troubleshooting appliance installation issues, see the installation and maintenance instructions for your appliances. If you need additional help, contact technical support.

Related information

[*Monitoring and troubleshooting StorageGRID*](#)
[*SG1000 appliance installation and maintenance*](#)
[*SG6000 appliance installation and maintenance*](#)
[*SG5700 appliance installation and maintenance*](#)
[*SG5600 appliance installation and maintenance*](#)
[*NetApp Support*](#)

Sample /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.

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.

Physical interfaces

/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
```

Copyright

Copyright © 2019 NetApp, Inc. All rights reserved. Printed in the U.S.

No part of this document covered by copyright may be reproduced in any form or by any means—graphic, electronic, or mechanical, including photocopying, recording, taping, or storage in an electronic retrieval system—without prior written permission of the copyright owner.

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

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

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

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

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

Trademark information

NETAPP, the NETAPP logo, and the marks listed on the NetApp Trademarks page are trademarks of NetApp, Inc. Other company and product names may be trademarks of their respective owners.

<http://www.netapp.com/us/legal/netapptmlist.aspx>

How to send comments about documentation and receive update notifications

You can help us to improve the quality of our documentation by sending us your feedback. You can receive automatic notification when production-level (GA/FCS) documentation is initially released or important changes are made to existing production-level documents.

If you have suggestions for improving this document, send us your comments by email.

[*doccomments@netapp.com*](mailto:doccomments@netapp.com)

To help us direct your comments to the correct division, include in the subject line the product name, version, and operating system.

If you want to be notified automatically when production-level documentation is released or important changes are made to existing production-level documents, follow Twitter account @NetAppDoc.

You can also contact us in the following ways:

- NetApp, Inc., 1395 Crossman Ave., Sunnyvale, CA 94089 U.S.
- Telephone: +1 (408) 822-6000
- Fax: +1 (408) 822-4501
- Support telephone: +1 (888) 463-8277