



Networking guidelines

StorageGRID 11.7

NetApp
April 12, 2024

Table of Contents

- Networking guidelines 1
 - Networking guidelines: Overview 1
 - StorageGRID network types 2
 - Network topology examples 5
 - Networking requirements 13
 - Network-specific requirements 14
 - Deployment-specific networking considerations 16
 - Network installation and provisioning 19
 - Post-installation guidelines 19
 - Network port reference 20

Networking guidelines

Networking guidelines: Overview

Use these guidelines to learn about StorageGRID architecture and networking topologies and to learn the requirements for network configuration and provisioning.

About these instructions

These guidelines provide information you can use to create the StorageGRID networking infrastructure before deploying and configuring StorageGRID nodes. Use these guidelines to help ensure that communication can occur among all the nodes in the grid and between the grid and external clients and services.

External clients and external services need to connect to StorageGRID networks to perform functions such as the following:

- Store and retrieve object data
- Receive email notifications
- Access the StorageGRID management interface (the Grid Manager and Tenant Manager)
- Access the audit share (optional)
- Provide services such as:
 - Network Time Protocol (NTP)
 - Domain name system (DNS)
 - Key Management Server (KMS)

StorageGRID networking must be configured appropriately to handle the traffic for these functions and more.

Before you begin

Configuring the networking for a StorageGRID system requires a high level of experience with Ethernet switching, TCP/IP networking, subnets, network routing, and firewalls.

Before you configure networking, become familiar with StorageGRID architecture as described in [Learn about StorageGRID](#).

After you determine which StorageGRID networks you want to use and how those networks will be configured, you can install and configure the StorageGRID nodes by following the appropriate instructions.

Install software-based nodes

- [Install Red Hat Enterprise Linux or CentOS](#)
- [Install Ubuntu or Debian](#)
- [Install VMware](#)

Install appliance nodes

- [Install appliance hardware](#)

Configure and administer StorageGRID software

- [Administer StorageGRID](#)
- [Release notes](#)

StorageGRID network types

The grid nodes in a StorageGRID system process *grid traffic*, *admin traffic*, and *client traffic*. You must configure the networking appropriately to manage these three types of traffic and to provide control and security.

Traffic types

Traffic type	Description	Network type
Grid traffic	The internal StorageGRID traffic that travels between all nodes in the grid. All grid nodes must be able to communicate with all other grid nodes over this network.	Grid Network (required)
Admin traffic	The traffic used for system administration and maintenance.	Admin Network (optional), VLAN network (optional)
Client traffic	The traffic that travels between external client applications and the grid, including all object storage requests from S3 and Swift clients.	Client Network (optional), VLAN network (optional)

You can configure networking in the following ways:

- Grid Network only
- Grid and Admin Networks
- Grid and Client Networks
- Grid, Admin, and Client Networks

The Grid Network is mandatory and can manage all grid traffic. The Admin and Client Networks can be included at the time of installation or added later to adapt to changes in requirements. Although the Admin Network and Client Network are optional, when you use these networks to handle administrative and client traffic, the Grid Network can be made isolated and secure.

Internal ports are only accessible over the Grid Network. External ports are accessible from all network types. This flexibility provides multiple options for designing a StorageGRID deployment and setting up external IP and port filtering in switches and firewalls. See [internal grid node communications](#) and [external communications](#).

Network interfaces

StorageGRID nodes are connected to each network using the following specific interfaces:

Network	Interface name
Grid Network (required)	eth0

Network	Interface name
Admin Network (optional)	eth1
Client Network (optional)	eth2

For details about mapping virtual or physical ports to node network interfaces, see the installation instructions:

Software-based nodes

- [Install Red Hat Enterprise Linux or CentOS](#)
- [Install Ubuntu or Debian](#)
- [Install VMware](#)

Appliance nodes

- [SGF6112 storage appliance](#)
- [SG6000 storage appliance](#)
- [SG5700 storage appliance](#)
- [SG100 and SG1000 services appliances](#)

Network information for each node

You must configure the following for each network you enable on a node:

- IP address
- Subnet mask
- Gateway IP address

You can only configure one IP address/mask/gateway combination for each of the three networks on each grid node. If you don't want to configure a gateway for a network, you should use the IP address as the gateway address.

High availability groups

High availability (HA) groups provide the ability to add virtual IP (VIP) addresses to the Grid or Client Network interface. For more information, see [Manage high availability groups](#).

Grid Network

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



StorageGRID does not support network address translation (NAT) between nodes.

The Grid Network can be used for all admin traffic and all client traffic, even if the Admin Network and Client Network are configured. The Grid Network gateway is the node default gateway unless the node has the Client Network configured.



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

Note the following requirements and details for the Grid Network gateway:

- 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.
- 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. When configured, it can be used for system administration and maintenance traffic. The Admin Network is typically a private network and does not need to be routable between nodes.

You can choose which grid nodes should have the Admin Network enabled on them.

When you use the Admin Network, administrative and maintenance traffic does not need to travel across the Grid Network. Typical uses of the Admin Network include the following:

- Access to the Grid Manager and Tenant Manager user interfaces.
- Access to critical services such as NTP servers, DNS servers, external key management servers (KMS), and Lightweight Directory Access Protocol (LDAP) servers.
- Access to audit logs on Admin Nodes.
- Secure Shell Protocol (SSH) access 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 communicate with multiple external subnets. However, the Admin Network gateway is never used as the node default gateway.

Note the following requirements and details for the Admin Network gateway:

- The Admin Network gateway is required if connections will be made from outside of the Admin Network subnet or if multiple Admin Network subnets are configured.
- Static routes are created for each subnet configured in the node's Admin Network Subnet List.

Client Network

The Client Network is optional. When configured, it is used to provide access to grid services for client applications such as S3 and Swift. If you plan to make StorageGRID data accessible to an external resource (for example, a Cloud Storage Pool or the StorageGRID CloudMirror replication service), the external resource can also use the Client Network. Grid nodes can communicate with any subnet reachable through the Client Network gateway.

You can choose which grid nodes should have the Client Network enabled on them. All nodes don't have to be on the same Client Network, and nodes will never communicate with each other over the Client Network. The Client Network does not become operational until grid installation is complete.

For added security, you can specify that a node's Client Network interface be untrusted so that the Client

Network will be more restrictive of which connections are allowed. If a node's Client Network interface is untrusted, the interface accepts outbound connections such as those used by CloudMirror replication, but only accepts inbound connections on ports that have been explicitly configured as load balancer endpoints. See [Manage firewall controls](#) and [Configure load balancer endpoints](#).

When you use a Client Network, client traffic does not need to travel across the Grid Network. Grid Network traffic can be separated onto a secure, non-routable network. The following node types are often configured with a Client Network:

- Gateway Nodes, because these nodes provide access to the StorageGRID Load Balancer service and S3 and Swift client access to the grid.
- Storage Nodes, because these nodes provide access to the S3 and Swift protocols and to Cloud Storage Pools and the CloudMirror replication service.
- Admin Nodes, to ensure that tenant users can connect to the Tenant Manager without needing to use the Admin Network.

Note the following for the Client Network gateway:

- 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.

Optional VLAN networks

As required, you can optionally use virtual LAN (VLAN) networks for client traffic and for some types of admin traffic. Grid traffic, however, can't use a VLAN interface. The internal StorageGRID traffic between nodes must always use the Grid Network on eth0.

To support the use VLANs, you must configure one or more interfaces on a node as trunk interfaces at the switch. You can configure the Grid Network interface (eth0) or the Client Network interface (eth2) to be a trunk, or you can add trunk interfaces to the node.

If eth0 is configured as a trunk, Grid Network traffic flows over the trunk native interface, as configured on the switch. Similarly, if eth2 is configured as a trunk, and the Client Network is also configured on the same node, the Client Network uses the trunk port's native VLAN as configured on the switch.

Only inbound admin traffic, such as used for SSH, Grid Manager, or Tenant Manager traffic, is supported over VLAN networks. Outbound traffic, such as used for NTP, DNS, LDAP, KMS, and Cloud Storage Pools, is not supported over VLAN networks.



VLAN interfaces can be added to Admin Nodes and Gateway Nodes only. You can't use a VLAN interface for client or admin access to Storage Nodes or Archive Nodes.

See [Configure VLAN interfaces](#) for instructions and guidelines.

VLAN interfaces are only used in HA groups and are assigned VIP addresses on the active node. See [Manage high availability groups](#) for instructions and guidelines.

Network topology examples

Grid Network topology

The simplest network topology is created by configuring the Grid Network only.

When you configure the Grid Network, you establish the host IP address, subnet mask, and Gateway IP address for the eth0 interface for each grid node.

During configuration, you must add all Grid Network subnets to the Grid Network Subnet List (GNSL). This list includes all subnets for all sites, and might also include external subnets that provide access to critical services such as NTP, DNS, or LDAP.


At installation, the Grid Network interface applies static routes for all subnets in the GNSL and sets the node's default route to the Grid Network gateway if one is configured. The GNSL is not required if there is no Client Network and the Grid Network gateway is the node's default route. Host routes to all other nodes in the grid are also generated.

In this example, all traffic shares the same network, including traffic related to S3 and Swift client requests and administrative and maintenance functions.



This topology is appropriate for single-site deployments that aren't externally available, proof-of-concept or test deployments, or when a third-party load balancer acts as the client access boundary. When possible, the Grid Network should be used exclusively for internal traffic. Both the Admin Network and the Client Network have additional firewall restrictions that block external traffic to internal services. Using the Grid Network for external client traffic is supported, but this use offers fewer layers of protection.

Topology example: Grid Network only



GNSL → 200.200.200.0/24

Grid Network		
Nodes	IP/mask	Gateway
Admin	200.200.200.32/24	200.200.200.1
Storage	200.200.200.33/24	200.200.200.1
Storage	200.200.200.34/24	200.200.200.1
Storage	200.200.200.35/24	200.200.200.1
Storage	200.200.200.36/24	200.200.200.1
Gateway	200.200.200.37/24	200.200.200.1

System Generated

Nodes	Routes	Type	From
All	0.0.0.0/0 → 200.200.200.1	Default	Grid Network gateway
	200.200.200.0/24 → eth0	Link	Interface IP/mask

Admin Network topology


Having an Admin Network is optional. One way that you can use an Admin Network and a Grid Network is to configure a routable Grid Network and a bounded Admin Network for each node.

When you configure the Admin Network, you establish the host IP address, subnet mask, and Gateway IP address for the eth1 interface for each grid node.

The Admin Network can be unique to each node and can consist of multiple subnets. Each node can be configured with an Admin External Subnet List (AESL). The AESL lists the subnets reachable over the Admin Network for each node. The AESL must also include the subnets of any services the grid will access over the Admin Network, such as NTP, DNS, KMS, and LDAP. Static routes are applied for each subnet in the AESL.

In this example, the Grid Network is used for traffic related to S3 and Swift client requests and object management. while the Admin Network is used for administrative functions.

Topology example: Grid and Admin Networks



GNSL → 172.16.0.0/16

AESL (all) → 10.10.1.0/24 10.10.2.0/24 10.10.3.0/24

Nodes	Grid Network		Admin Network	
	IP/mask	Gateway	IP/mask	Gateway
Admin	172.16.200.32/24	172.16.200.1	10.10.1.10/24	10.10.1.1
Storage 1	172.16.200.33/24	172.16.200.1	10.10.1.11/24	10.10.1.1
Storage 2	172.16.200.34/24	172.16.200.1	10.10.3.65/24	10.10.3.1
Storage 3	172.16.200.35/24	172.16.200.1	10.10.1.12/24	10.10.1.1
Storage 4	172.16.200.36/24	172.16.200.1	10.10.1.13/24	10.10.1.1
Gateway	172.16.200.37/24	172.16.200.1	10.10.3.66/24	10.10.3.1

System Generated					
Nodes	Routes			Type	From
All	0.0.0.0/0	→	172.16.200.1	Default	Grid Network gateway
Admin, Storage 1, 3, and 4	172.16.0.0/16	→	eth0	Static	GNSL
	10.10.1.0/24	→	eth1	Link	Interface IP/mask
	10.10.2.0/24	→	10.10.1.1	Static	AESL
	10.10.3.0/24	→	10.10.1.1	Static	AESL
Storage 2, Gateway	172.16.0.0/16	→	eth0	Static	GNSL
	10.10.1.0/24	→	10.10.3.1	Static	AESL
	10.10.2.0/24	→	10.10.3.1	Static	AESL
	10.10.3.0/24	→	eth1	Link	Interface IP/mask

Client Network topology


Having a Client Network is optional. Using a Client Network allows client network traffic (for example, S3 and Swift) to be separated from grid internal traffic, which allows grid networking to be more secure. Administrative traffic can be handled by either the Client or Grid Network when the Admin Network is not configured.

When you configure the Client Network, you establish the host IP address, subnet mask, and Gateway IP address for the eth2 interface for the configured node. Each node's Client Network can be independent of the Client Network on any other node.

If you configure a Client Network for a node during installation, the node's default gateway switches from the Grid Network gateway to the Client Network gateway when installation is complete. If a Client Network is added later, the node's default gateway switches in the same way.

In this example, the Client Network is used for S3 and Swift client requests and for administrative functions, while the Grid Network is dedicated to internal object management operations.

Topology example: Grid and Client Networks



GNSL → 172.16.0.0/16

Nodes	Grid Network	Client Network	
	IP/mask	IP/mask	Gateway
Admin	172.16.200.32/24	37.5.5.10/24	37.5.5.1
Storage	172.16.200.33/24	37.5.5.11/24	37.5.5.1
Storage	172.16.200.34/24	37.5.5.12/24	37.5.5.1
Storage	172.16.200.35/24	37.5.5.13/24	37.5.5.1
Storage	172.16.200.36/24	37.5.5.14/24	37.5.5.1
Gateway	172.16.200.37/24	37.5.5.15/24	37.5.5.1

System Generated

Nodes	Routes		Type	From
All	0.0.0.0/0	→ 37.5.5.1	Default	Client Network gateway
	172.16.0.0/16	→ eth0	Link	Interface IP/mask
	37.5.5.0/24	→ eth2	Link	Interface IP/mask


Topology for all three networks

You can configure all three networks into a network topology consisting of a private Grid Network, bounded site-specific Admin Networks, and open Client Networks. Using load balancer endpoints and untrusted Client Networks can provide additional security if needed.

In this example:

- The Grid Network is used for network traffic related to internal object management operations.
- The Admin Network is used for traffic related to administrative functions.
- The Client Network is used for traffic related to S3 and Swift client requests.

Topology example: Grid, Admin, and Client Networks



Networking requirements

You must verify that the current networking infrastructure and configuration can support the planned StorageGRID network design.

General networking requirements

All StorageGRID deployments must be able to support the following connections.

These connections can occur through the Grid, Admin, or Client Networks, or the combinations of these networks as illustrated in the network topology examples.

- **Management connections:** Inbound connections from an administrator to the node, usually through SSH. Web browser access to the Grid Manager, the Tenant Manager, and the StorageGRID Appliance Installer.
- **NTP server connections:** Outbound UDP connection that receives an inbound UDP response.

At least one NTP server must be reachable by the primary Admin Node.

- **DNS server connections:** Outbound UDP connection that receives an inbound UDP response.
- **LDAP/Active Directory server connections:** Outbound TCP connection from the Identity service on Storage Nodes.
- **AutoSupport:** Outbound TCP connection from the Admin Nodes to either `support.netapp.com` or a customer-configured proxy.
- **External key management server:** Outbound TCP connection from each appliance node with node encryption enabled.
- Inbound TCP connections from S3 and Swift clients.
- Outbound requests from StorageGRID platform services such as CloudMirror replication or from Cloud Storage Pools.

If StorageGRID is unable to contact any of the provisioned NTP or DNS servers using the default routing rules, it will automatically attempt contact on all networks (Grid, Admin, and Client) as long as the IP addresses of the DNS and NTP servers are specified. If the NTP or DNS servers can be reached on any network, StorageGRID will automatically create additional routing rules to ensure that network is used for all future attempts to connect to it.



Although you can use these automatically discovered host routes, in general you should manually configure the DNS and NTP routes to ensure connectivity in case automatic discovery fails.

If you aren't 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. Additionally, you can configure these networks after installation, using the Change IP tool (see [Configure IP addresses](#)).

Only S3 and Swift client connections and SSH, Grid Manager, and Tenant Manager administrative connections are supported over VLAN interfaces. Outbound connections, such as to NTP, DNS, LDAP, AutoSupport, and KMS servers, must go over the Client, Admin, or Grid Network interfaces directly. If the interface is configured as a trunk to support VLAN interfaces, this traffic will flow over the interface's native VLAN, as configured at the switch.

Wide Area Networks (WANs) for multiple sites

When configuring a StorageGRID system with multiple sites, the WAN connection between sites must have a minimum bandwidth of 25 Mbit/second in each direction before accounting for client traffic. Data replication or erasure coding between sites, node or site expansion, node recovery, and other operations or configurations will require additional bandwidth.

Actual minimum WAN bandwidth requirements depend on client activity and the ILM protection scheme. For assistance estimating the minimum WAN bandwidth requirements, contact your NetApp Professional Services consultant.

Connections for Admin Nodes and Gateway Nodes

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.

Admin Nodes and Gateway Nodes that you plan to add to high availability groups must be configured with a static IP address. For more information, see [Manage high availability groups](#).

Using network address translation (NAT)

Don't use network address translation (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. As required, however, you can use NAT between external clients and grid nodes, such as to provide a public IP address for a Gateway Node. 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.

Network-specific requirements

Follow the requirements for each StorageGRID network type.

Network gateways and routers

- If set, the gateway for a given network must be within the specific network's subnet.
- If you configure an interface using static addressing, you must specify a gateway address other than 0.0.0.0.
- If you don't have a gateway, the best practice is to set the gateway address to be the IP address of the network interface.

Subnets



Each network must be connected to its own subnet that does not overlap with any other network on the node.

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 can't be 255.255.255.254 or 255.255.255.255 (/31 or /32 in CIDR notation).

- The subnet defined by a network interface IP address and subnet mask (CIDR) can't overlap the subnet of any other interface configured on the same node.
- The Grid Network subnet for each node must be included in the GNSL.
- The Admin Network subnet can't overlap the Grid Network subnet, the Client Network subnet, or any subnet in the GNSL.
- The subnets in the AESL can't overlap with any subnets in the GNSL.
- The Client Network subnet can't overlap the Grid Network subnet, the Admin Network subnet, any subnet in the GNSL, or any subnet in the AESL.

Grid Network

- At deployment time, each grid 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.
- During normal grid operations, each grid node must be able to communicate with all other grid nodes over the Grid Network.



The Grid Network must be directly routable between each node. Network address translation (NAT) between nodes is not supported.

- If the Grid Network consists of multiple subnets, add them to the Grid Network Subnet List (GNSL). Static routes are created on all nodes for each subnet in the GNSL.
- If the Grid Network interface is configured as a trunk to support VLAN interfaces, the trunk native VLAN must be the VLAN used for Grid Network traffic. All grid nodes must be accessible over the trunk native VLAN.

Admin Network

The Admin Network is optional. If you plan to configure an Admin Network, follow these requirements and guidelines.

Typical uses of the Admin Network include management connections, AutoSupport, KMS, and connections to critical servers such as NTP, DNS, and LDAP if these connections aren't provided through the Grid Network or Client Network.



The Admin Network and AESL can be unique to each node, as long as the desired network services and clients are reachable.



You must define at least one subnet on the Admin Network to enable inbound connections from external subnets. Static routes are automatically generated on each node for each subnet in the AESL.

Client Network

The Client Network is optional. If you plan to configure a Client Network, note the following considerations.

- The Client Network is designed to support traffic from S3 and Swift clients. If configured, the Client Network gateway becomes the node's default gateway.
- If you use a Client Network, you can help secure StorageGRID from hostile attacks by accepting inbound

client traffic only on explicitly configured load balancer endpoints. See [Configure load balancer endpoints](#).

- If the Client Network interface is configured as a trunk to support VLAN interfaces, consider whether configuring the Client Network interface (eth2) is necessary. If configured, Client Network traffic will flow over the trunk native VLAN, as configured in the switch.

Deployment-specific networking considerations

Linux deployments

For efficiency, reliability, and security, the StorageGRID system runs on Linux as a collection of container engines. Container engine-related network configuration is not required in a StorageGRID system.

Use a non-bond device, such as a VLAN or virtual Ethernet (veth) pair, for the container network interface. Specify this device as the network interface in the node configuration file.



Don't use bond or bridge devices directly as the container network interface. Doing so could prevent node start-up because of a kernel issue with the use of macvlan with bond and bridge devices in the container namespace.

See the installation instructions for [Red Hat Enterprise Linux or CentOS](#) or [Ubuntu or Debian](#) deployments.

Host network configuration for container engine deployments

Before starting your StorageGRID deployment on a container engine platform, determine which networks (Grid, Admin, Client) each node will use. You must ensure that each node's network interface is configured on the correct virtual or physical host interface, and that each network has sufficient bandwidth.

Physical hosts

If you are using physical hosts to support grid nodes:

- Make sure all hosts use the same host interface for each node interface. This strategy simplifies host configuration and enables future node migration.
- Obtain an IP address for the physical host itself.



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 host and the node can't share IP addresses.

- Open the required ports to the host.
- If you intend to use VLAN interfaces in StorageGRID, the host must have one or more trunk interfaces that provide access to the desired VLANs. These interfaces can be passed into the node container as eth0, eth2, or as additional interfaces. To add trunk or access interfaces, see the following:
 - **RHEL or CentOS (before installing the node):** [Create node configuration files](#)
 - **Ubuntu or Debian (before installing the node):** [Create node configuration files](#)
 - **RHEL, CentOS, Ubuntu, or Debian (after installing the node):** [Linux: Add trunk or access interfaces to a node](#)

Minimum bandwidth recommendations

The following table provides the minimum LAN 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
	Minimum LAN bandwidth		
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



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 the 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, to 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 deployment, see the following:

- [Configure the host network \(Red Hat Enterprise Linux or CentOS\)](#)
- [Configure the host network \(Ubuntu or Debian\)](#)

Networking and ports for platform services and Cloud Storage Pools

If you plan to use StorageGRID platform services or Cloud Storage Pools, you must

configure grid networking and firewalls to ensure that the destination endpoints can be reached.

Networking for platform services

As described in [Manage platform services for tenants](#) and [What are platform services?](#), platform services include external services that provide search integration, event notification, and CloudMirror replication.

Platform services require access from Storage Nodes that host the StorageGRID ADC service to the external service endpoints. Examples for providing access include:

- On the Storage Nodes with ADC services, configure unique Admin Networks with AESL entries that route to the target endpoints.
- Rely on the default route provided by a Client Network. If you use the default route, you can use the [untrusted Client Network feature](#) to restrict inbound connections.

Networking for Cloud Storage Pools

Cloud Storage Pools also require access from Storage Nodes to the endpoints provided by the external service used, such as Amazon S3 Glacier or Microsoft Azure Blob storage. For information, see [What is a Cloud Storage Pool?](#).

Ports for platform services and Cloud Storage Pools

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`

A different port can be specified when the endpoint is created or edited. See [Network port reference](#).

If you use a non-transparent proxy server, you must also [configure storage proxy settings](#) to allow messages to be sent to external endpoints, such as an endpoint on the internet.

VLANs and platform services and Cloud Storage Pools

You can't use VLAN networks for platform services or Cloud Storage Pools. The destination endpoints must be reachable over the Grid, Admin, or Client Network.

Appliance nodes

You can configure the network ports on StorageGRID appliances to use the port bond modes that meet your requirements for throughput, redundancy, and failover.

The 10/25-GbE ports on the StorageGRID appliances can be configured in Fixed or Aggregate bond mode for connections to the Grid Network and Client Network.

The 1-GbE Admin Network ports can be configured in Independent or Active-Backup mode for connections to the Admin Network.

See the information about port bond modes for your appliance:

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

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.

Initial deployment of a node

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.

If necessary, subnets containing NTP servers or requiring access to the Grid Manager or API can also be configured as grid subnets.

Automatic node registration with primary Admin Node

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.

Disabling the Admin Network or Client Network

If you want to disable the Admin Network or Client Network, you can remove the configuration from them during the node approval process, or you can use the Change IP tool after installation is complete (see [Configure IP addresses](#)).

Post-installation guidelines

After completing grid node deployment and configuration, follow these guidelines for DHCP addressing and network configuration changes.

- If DHCP was used to assign IP addresses, configure a DHCP reservation for each IP address on the networks being used.

You can only set up DHCP during the deployment phase. You can't set up DHCP during configuration.



Nodes reboot when their IP addresses change, which can cause outages if a DHCP address change affects multiple nodes at the same time.

- You must use the Change IP procedures if you want to change IP addresses, subnet masks, and default gateways for a grid node. See [Configure IP addresses](#).
- If you make networking configuration changes, including routing and gateway changes, client connectivity to the primary Admin Node and other grid nodes might be lost. Depending on the networking changes applied, you might need to reestablish these connections.

Network port reference

You must ensure the network infrastructure can provide internal and external communication between nodes within the grid and to external clients and services. You might need access across internal and external firewalls, switching systems, and routing systems.

Use the details provided for [Internal grid node communications](#) and [External communications](#) to determine how to configure each required port.

Internal grid node communications

The StorageGRID internal firewall allows incoming connections to specific ports on the Grid Network. Connections are also accepted on ports defined by load balancer endpoints.



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 can't 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 protocol number 112. StorageGRID uses VRRP in unicast mode only. VRRP is required only if [high availability groups](#) are configured.

Guidelines for Linux-based nodes

If enterprise networking policies restrict access to any of these ports, you can remap ports at deployment time using a deployment configuration parameter. For more information about port remapping and deployment configuration parameters, see:

- [Install Red Hat Enterprise Linux or CentOS](#)
- [Install Ubuntu or Debian](#)

Guidelines for VMware-based nodes

Configure the following ports only if you need to define firewall restrictions that are external to VMware networking.

If enterprise networking policies restrict access to any of these ports, you can remap ports when you deploy nodes using the VMware vSphere Web Client, or by using a configuration file setting when automating grid node deployment. For more information about port remapping and deployment configuration parameters, see [Install VMware](#).

Guidelines for appliance nodes

If enterprise networking policies restrict access to any of these ports, you can remap ports using the StorageGRID Appliance Installer. See [Optional: Remap network ports for appliance](#).

StorageGRID internal ports

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 its 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.
1055	TCP	All nodes	Primary Admin Node	Internal traffic for installation, expansion, recovery, 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.

Port	TCP or UDP	From	To	Details
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.
7001	TCP	Storage Nodes	Storage Nodes	Cassandra TLS inter-node cluster communication.
7443	TCP	All nodes	Primary Admin Node	Internal traffic for installation, expansion, recovery, other maintenance procedures, and error reporting.
8011	TCP	All nodes	Primary Admin Node	Internal traffic for installation, expansion, recovery, and other maintenance procedures.
8443	TCP	Primary Admin Node	Appliance nodes	Internal traffic related to the maintenance mode procedure.
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.
10226	TCP	Storage Nodes	Primary Admin Node	Used by StorageGRID appliances for forwarding AutoSupport messages from E-Series SANtricity System Manager to the primary Admin Node.
10342	TCP	All nodes	Primary Admin Node	Internal traffic for installation, expansion, recovery, and other maintenance procedures.
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.

Port	TCP or UDP	From	To	Details
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.
18019	TCP	Storage Nodes	Storage Nodes	Chunk service internal traffic for erasure coding.
18082	TCP	Admin/Storage Nodes	Storage Nodes	S3-related internal traffic.
18083	TCP	All nodes	Storage Nodes	Swift-related internal traffic.
18086	TCP	All grid nodes	All Storage Nodes	Internal traffic related to LDR service.
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.

Related information

[External communications](#)

External communications

Clients need to communicate with grid nodes to ingest and retrieve content. The ports used depends on the object storage protocols chosen. These ports need to be accessible to the client.

Restricted access to ports

If enterprise networking policies restrict access to any of the ports, you can use [load balancer endpoints](#) to allow access on user-defined ports. You can then use [untrusted Client Networks](#) to allow access only on load balancer endpoint ports.

Port remapping

To use systems and protocols such as SMTP, DNS, SSH, or DHCP, you must remap ports when deploying nodes. However, you should not remap load balancer endpoints. For information about port remapping, see the installation instructions:

- [Install Red Hat Enterprise Linux or CentOS](#)
- [Install Ubuntu or Debian](#)
- [Install VMware](#)
- [Optional: Remap network ports for appliance](#)

Ports used for external communications

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



This list does not include ports that might be configured as [load balancer endpoints](#) or used for [syslog servers](#).

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 DNS.
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.
80	TCP	HTTP	Browser	Appliances	Port 80 redirects to port 8443 for the StorageGRID Appliance Installer.

Port	TCP or UDP	Protocol	From	To	Details
80	TCP	HTTP	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.
80	TCP	HTTP	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). Note: This port is required only if NFS-based audit export is enabled.
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. Note: This port is required only if SMB-based audit export is enabled.
138	UDP	NetBIOS	SMB client	Admin Nodes	Used by SMB-based audit export for clients that require NetBIOS support. Note: This port is required only if SMB-based audit export is enabled.
139	TCP	SMB	SMB client	Admin Nodes	Used by SMB-based audit export for clients that require NetBIOS support. Note: This port is required only if SMB-based audit export is enabled.

Port	TCP or UDP	Protocol	From	To	Details
161	TCP/ UDP	SNMP	SNMP client	All nodes	<p>Used for SNMP polling. All nodes provide basic information; Admin Nodes also provide alert and alarm data. Defaults to UDP port 161 when configured.</p> <p>Note: This port is only required, and is only opened on the node firewall if SNMP is configured. If you plan to use SNMP, you can configure alternate ports.</p> <p>Note: For information about using SNMP with StorageGRID, contact your NetApp account representative.</p>
162	TCP/ UDP	SNMP Notifications	All nodes	Notification destinations	<p>Outbound SNMP notifications and traps default to UDP port 162.</p> <p>Note: This port is only required if SNMP is enabled and notification destinations are configured. If you plan to use SNMP, you can configure alternate ports.</p> <p>Note: For information about using SNMP with StorageGRID, contact your NetApp account representative.</p>
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	<p>Used by web browsers and management API clients for accessing the Grid Manager and Tenant Manager.</p> <p>Note: If you close Grid Manager ports 443 or 8443, any users currently connected on a blocked port, including you, will lose access to Grid Manager unless their IP address has been added to the Privileged address list. See Configure firewall controls to configure privileged IP addresses.</p>
443	TCP	HTTPS	Admin Nodes	Active Directory	Used by Admin Nodes connecting to Active Directory if single sign-on (SSO) is enabled.
443	TCP	HTTPS	Archive Nodes	Amazon S3	Used for accessing Amazon S3 from Archive Nodes.

Port	TCP or UDP	Protocol	From	To	Details
443	TCP	HTTPS	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.
443	TCP	HTTPS	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. Note: This port is required only if SMB-based audit export is enabled.
903	TCP	NFS	NFS client	Admin Nodes	Used by NFS-based audit export (<code>rpc.mountd</code>). Note: This port is required only if NFS-based audit export is enabled.
2022	TCP	SSH	Service laptop	All nodes	SSH or console access is required for procedures with console steps. Optionally, you can use port 22 instead of 2022.
2049	TCP	NFS	NFS client	Admin Nodes	Used by NFS-based audit export (nfs). Note: This port is required only if NFS-based audit export is enabled.
5353	UDP	mDNS	All nodes	All nodes	Provides the multicast DNS (mDNS) service that is used for full-grid IP changes and for primary Admin Node discovery during installation, expansion, and recovery.
5696	TCP	KMIP	Appliance	KMS	Key Management Interoperability Protocol (KMIP) external traffic from appliances configured for node encryption to the Key Management Server (KMS), unless a different port is specified on the KMS configuration page of the StorageGRID Appliance Installer.

Port	TCP or UDP	Protocol	From	To	Details
8022	TCP	SSH	Service laptop	All nodes	SSH on port 8022 grants access to the base operating system on appliance and virtual node platforms for support and troubleshooting. This port is not used for Linux-based (bare metal) nodes and is not required to be accessible between grid nodes or during normal operations.
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. Note: If you close Grid Manager ports 443 or 8443, any users currently connected on a blocked port, including you, will lose access to Grid Manager unless their IP address has been added to the Privileged address list. See Configure firewall controls to configure privileged IP addresses.
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.
9091	TCP	HTTPS	External Grafana service	Admin Nodes	Used by external Grafana services for secure access to the StorageGRID Prometheus service. Note: This port is required only if certificate-based Prometheus access is enabled.
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 client traffic directly to Storage Nodes (HTTPS).
18083	TCP	HTTPS	Swift clients	Storage Nodes	Swift client traffic directly to Storage Nodes (HTTPS).
18084	TCP	HTTP	S3 clients	Storage Nodes	S3 client traffic directly to Storage Nodes (HTTP).

Port	TCP or UDP	Protocol	From	To	Details
18085	TCP	HTTP	Swift clients	Storage Nodes	Swift client traffic directly to Storage Nodes (HTTP).
23000-23999	TCP	HTTPS	All nodes on the source grid for cross-grid replication	Admin Nodes and Gateway Nodes on the destination grid for cross-grid replication	This range of ports is reserved for grid federation connections. Both grids in a given connection use the same port.

Copyright information

Copyright © 2024 NetApp, Inc. All Rights Reserved. Printed in the U.S. No part of this document covered by copyright may be reproduced in any form or by any means—graphic, electronic, or mechanical, including photocopying, recording, taping, or storage in an electronic retrieval system—without prior written permission of the copyright owner.

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

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

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

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

LIMITED RIGHTS LEGEND: Use, duplication, or disclosure by the government is subject to restrictions as set forth in subparagraph (b)(3) of the Rights in Technical Data -Noncommercial Items at DFARS 252.227-7013 (FEB 2014) and FAR 52.227-19 (DEC 2007).

Data contained herein pertains to a commercial product and/or commercial service (as defined in FAR 2.101) and is proprietary to NetApp, Inc. All NetApp technical data and computer software provided under this Agreement is commercial in nature and developed solely at private expense. The U.S. Government has a non-exclusive, non-transferrable, nonsublicensable, worldwide, limited irrevocable license to use the Data only in connection with and in support of the U.S. Government contract under which the Data was delivered. Except as provided herein, the Data may not be used, disclosed, reproduced, modified, performed, or displayed without the prior written approval of NetApp, Inc. United States Government license rights for the Department of Defense are limited to those rights identified in DFARS clause 252.227-7015(b) (FEB 2014).

Trademark information

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