



Configure S3 and Swift client connections

StorageGRID

NetApp
June 10, 2022

Table of Contents

- Configure S3 and Swift client connections 1
 - About S3 and Swift client connections 1
 - Summary: IP addresses and ports for client connections 1
- Configure VLAN interfaces 4
- Manage high availability groups 7
- Manage load balancing 19
- Configure S3 API endpoint domain names 30
- Enable HTTP for client communications 32
- Control which client operations are permitted 33

Configure S3 and Swift client connections

About S3 and Swift client connections

As a grid administrator, you manage the configuration options that control how S3 and Swift tenants can connect client applications to your StorageGRID system to store and retrieve data. There are a number of different options to meet different client and tenant requirements.

Client applications can store or retrieve objects by connecting to any of the following:

- The Load Balancer service on Admin Nodes or Gateway Nodes, or optionally, the virtual IP address of a high availability (HA) group of Admin Nodes or Gateway Nodes
- The CLB service on Gateway Nodes, or optionally, the virtual IP address of a high availability group of Gateway Nodes



The CLB service is deprecated. Clients configured before the StorageGRID 11.3 release can continue to use the CLB service on Gateway Nodes. All other client applications that depend on StorageGRID to provide load balancing should connect using the Load Balancer service.

- Storage Nodes, with or without an external load balancer

You can optionally configure the following features on your StorageGRID system:

- **VLAN interfaces:** You can create virtual LAN (VLAN) interfaces on Admin Nodes and Gateway Nodes to isolate and partition client and tenant traffic for security, flexibility, and performance. After creating a VLAN interface, you add it to a high availability (HA) group.
- **High availability groups:** You can create an HA group of the interfaces for Gateway Nodes or Admin Nodes to create an active-backup configuration, or you can use round-robin DNS or a third-party load balancer and multiple HA groups to achieve an active-active configuration. Client connections are made using the virtual IP addresses of HA groups.
- **Load Balancer service:** You can enable clients to use the Load Balancer service by creating load balancer endpoints for client connections. When creating a load balancer endpoint, you specify a port number, whether the endpoint accepts HTTP or HTTPS connections, the type of client (S3 or Swift) that will use the endpoint, and the certificate to be used for HTTPS connections (if applicable).
- **Untrusted Client Network:** You can make the Client Network more secure by configuring it as untrusted. When the Client Network is untrusted, clients can only connect using load balancer endpoints.

You can also enable the use of HTTP for clients that connect to StorageGRID either directly to Storage Nodes or using the CLB service (deprecated), and you can configure S3 API endpoint domain names for S3 clients.

Summary: IP addresses and ports for client connections

Client applications can connect to StorageGRID using the IP address of a grid node and the port number of a service on that node. If high availability (HA) groups are configured, client applications can connect using the virtual IP address of the HA group.

About this task

This table summarizes the different ways that clients can connect to StorageGRID and the IP addresses and ports that are used for each type of connection. The instructions describe how to find this information in the Grid Manager if load balancer endpoints and high availability (HA) groups are already configured.

Where connection is made	Service that client connects to	IP address	Port
HA group	Load Balancer	Virtual IP address of an HA group	<ul style="list-style-type: none"> • Load balancer endpoint port
HA group	CLB Note: The CLB service is deprecated.	Virtual IP address of an HA group	Default S3 ports: <ul style="list-style-type: none"> • HTTPS: 8082 • HTTP: 8084 Default Swift ports: <ul style="list-style-type: none"> • HTTPS:8083 • HTTP:8085
Admin Node	Load Balancer	IP address of the Admin Node	<ul style="list-style-type: none"> • Load balancer endpoint port
Gateway Node	Load Balancer	IP address of the Gateway Node	<ul style="list-style-type: none"> • Load balancer endpoint port
Gateway Node	CLB Note: The CLB service is deprecated.	IP address of the Gateway Node Note: By default, HTTP ports for CLB and LDR are not enabled.	Default S3 ports: <ul style="list-style-type: none"> • HTTPS: 8082 • HTTP: 8084 Default Swift ports: <ul style="list-style-type: none"> • HTTPS:8083 • HTTP:8085
Storage Node	LDR	IP address of Storage Node	Default S3 ports: <ul style="list-style-type: none"> • HTTPS: 18082 • HTTP: 18084 Default Swift ports: <ul style="list-style-type: none"> • HTTPS: 18083 • HTTP:18085

Examples

To connect an S3 client to the Load Balancer endpoint of an HA group of Gateway Nodes, use a URL structured as shown below:

- `https://VIP-of-HA-group:LB-endpoint-port`

For example, if the virtual IP address of the HA group is 192.0.2.5 and the port number of an S3 Load Balancer endpoint is 10443, then an S3 client could use the following URL to connect to StorageGRID:

- `https://192.0.2.5:10443`

To connect a Swift client to the Load Balancer endpoint of an HA group of Gateway Nodes, use a URL structured as shown below:

- `https://VIP-of-HA-group:LB-endpoint-port`

For example, if the virtual IP address of the HA group is 192.0.2.6 and the port number of a Swift Load Balancer endpoint is 10444, then a Swift client could use the following URL to connect to StorageGRID:

- `https://192.0.2.6:10444`

It is possible to configure a DNS name for the IP address that clients use to connect to StorageGRID. Contact your local network administrator.

Steps

1. Sign in to the Grid Manager using a [supported web browser](#).
2. To find the IP address of a grid node:
 - a. Select **NODES**.
 - b. Select the Admin Node, Gateway Node, or Storage Node to which you want to connect.
 - c. Select the **Overview** tab.
 - d. In the Node Information section, note the IP addresses for the node.
 - e. Select **Show more** to view IPv6 addresses and interface mappings.

You can establish connections from client applications to any of the IP addresses in the list:

- **eth0**: Grid Network
- **eth1**: Admin Network (optional)
- **eth2**: Client Network (optional)



If you are viewing an Admin Node or a Gateway Node and it is the active node in a high availability group, the virtual IP address of the HA group is shown on eth2.

3. To find the virtual IP address of a high availability group:
 - a. Select **CONFIGURATION > Network > High availability groups**.
 - b. In the table, note the virtual IP address of the HA group.
4. To find the port number of a Load Balancer endpoint:
 - a. Select **CONFIGURATION > Network > Load balancer endpoints**.

The Load Balancer Endpoints page appears, showing the list of endpoints that have already been configured.

- b. Select an endpoint, and select **Edit endpoint**.

The Edit Endpoint window opens and displays additional details about the endpoint.

- c. Confirm that the endpoint you have selected is configured for use with the correct protocol (S3 or Swift), then select **Cancel**.
- d. Note the port number for the endpoint that you want to use for a client connection.



If the port number is 80 or 443, the endpoint is configured only on Gateway Nodes, since those ports are reserved on Admin Nodes. All other ports are configured on both Gateway Nodes and Admin Nodes.

Configure VLAN interfaces

You can create virtual LAN (VLAN) interfaces on Admin Nodes and Gateway Nodes and use them in HA groups and load balancer endpoints to isolate and partition traffic for security, flexibility, and performance.

Considerations for VLAN interfaces

- You create a VLAN interface by entering a VLAN ID and choosing a parent interface on one or more nodes.
- A parent interface must be configured as a trunk interface at the switch.
- A parent interface can be the Grid Network (eth0), the Client Network (eth2), or an additional trunk interface for the VM or bare-metal host (for example, ens256).
- For each VLAN interface, you can select only one parent interface for a given node. For example, you cannot use both the Grid Network interface and the Client Network interface on the same Gateway Node as the parent interface for the same VLAN.
- If the VLAN interface is for Admin Node traffic, which includes traffic related to the Grid Manager and the Tenant Manager, select interfaces on Admin Nodes only.
- If the VLAN interface is for S3 or Swift client traffic, select interfaces on either Admin Nodes or Gateway Nodes.
- If you need to add trunk interfaces, see the following for details:
 - **VMware (after installing the node):** [VMware: Add trunk or access interfaces to a node](#)
 - **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](#)

Create a VLAN interface

What you'll need

- You are signed in to the Grid Manager using a [supported web browser](#).
- You have the Root access permission.
- A trunk interface has been configured in the network and attached to the VM or Linux node. You know the name of the trunk interface.
- You know the ID of the VLAN you are configuring.

About this task

Your network administrator might have configured one or more trunk interfaces and one or more VLANs to segregate the client or admin traffic belonging to different applications or tenants. Each VLAN is identified by a numeric ID or tag. For example, your network might use VLAN 100 for FabricPool traffic and VLAN 200 for an archive application.

You can use the Grid Manager to create VLAN interfaces that allow clients to access StorageGRID on a specific VLAN. When you create VLAN interfaces, you specify the VLAN ID and select parent (trunk) interfaces on one or more nodes.

Access the wizard

1. Select **CONFIGURATION > Network > VLAN interfaces**.
2. Select **Create**.

Enter details for the VLAN interfaces

1. Specify the ID of the VLAN in your network. You can enter any value between 1 and 4094.

VLAN IDs do not need to be unique. For example, you might use VLAN ID 200 for admin traffic at one site and the same VLAN ID for client traffic at another site. You can create separate VLAN interfaces with different sets of parent interfaces at each site. However, two VLAN interfaces with the same ID cannot share the same interface on a node.

If you specify an ID that has already been used, a message appears. You can continue creating another VLAN interface for the same VLAN ID, or you can select **Cancel** and then edit the existing ID.

2. Optionally, enter a short description for the VLAN interface.

VLAN details

VLAN ID ?

203

Description (optional) ?

VLAN for S3 tenants. Uses Admin and Gateway Nodes at site 1.

60/64

Cancel Continue

3. Select **Continue**.

Choose parent interfaces

The table lists the available interfaces for all Admin Nodes and Gateway Nodes at each site in your grid. Admin

Network (eth1) interfaces cannot be used as parent interfaces and are not shown.

1. Select one or more parent interfaces to attach this VLAN to.

For example, you might want to attach a VLAN to the Client Network (eth2) interface for a Gateway Node and an Admin Node.

Parent interfaces

Select one or more parent interfaces for this VLAN interface. You can only select one parent interface on each node for each VLAN interface.


	Site ?	Node name ?	Interface ?	Description ?	Node type ?	Attached VLANs ?
<input type="checkbox"/>	Data Center 2	DC2-ADM1	eth0	Grid Network	Non-primary Admin	—
<input checked="" type="checkbox"/>	Data Center 2	DC2-ADM1	eth2	Client Network	Non-primary Admin	—
<input type="checkbox"/>	Data Center 1	DC1-G1	eth0	Grid Network	Gateway	—
<input checked="" type="checkbox"/>	Data Center 1	DC1-G1	eth2	Client Network	Gateway	—
<input type="checkbox"/>	Data Center 1	DC1-ADM1	eth0	Grid Network	Primary Admin	—

2 interfaces are selected.

[Previous](#) [Continue](#)

2. Select **Continue**.

Confirm the settings

1. Review the configuration and make any changes.
 - If you need to change the VLAN ID or description, select **Enter VLAN details** at the top of the page.
 - If you need to change a parent interface, select **Choose parent interfaces** at the top of the page or select **Previous**.
 - If you need to remove a parent interface, select the trash can .
2. Select **Save**.
3. Wait up to 5 minutes for the new interface to appear as a selection on the High availability groups page and to be listed in the **Network interfaces** table for the node (**NODES > parent interface node > Network**).

Edit a VLAN interface

When you edit a VLAN interface, you can make the following types of changes:

- Change the VLAN ID or description.
- Add or remove parent interfaces.

For example, you might want to remove a parent interface from a VLAN interface if you plan to decommission the associated node.

Note the following:

- You can't change a VLAN ID if the VLAN interface is used in an HA group.
- You can't remove a parent interface if that parent interface is used in an HA group.

For example, suppose VLAN 200 is attached to parent interfaces on Nodes A and B. If an HA group uses the VLAN 200 interface for Node A and the eth2 interface for Node B, you can remove the unused parent interface for Node B, but you can't remove the used parent interface for Node A.

Steps

1. Select **CONFIGURATION > Network > VLAN interfaces**.
2. Select the check box for the VLAN interface you want to edit. Then, select **Actions > Edit**.
3. Optionally, update the VLAN ID or the description. Then, select **Continue**.

You can't update a VLAN ID if the VLAN is used in an HA group.

4. Optionally, select or unselect the check boxes to add parent interfaces or to remove unused interfaces. Then, select **Continue**.
5. Review the configuration and make any changes.
6. Select **Save**.

Remove a VLAN interface

You can remove one or more VLAN interfaces.

You can't remove a VLAN interface if it is currently used in an HA group. You must remove the VLAN interface from the HA group before you can remove it.

To avoid any disruptions in client traffic, consider doing one of the following:

- Add a new VLAN interface to the HA group before removing this VLAN interface.
- Create a new HA group that does not use this VLAN interface.
- If the VLAN interface you want to remove is currently the active interface, edit the HA group. Move the VLAN interface you want to remove to the bottom of the priority list. Wait until communication is established on the new primary interface and then remove the old interface from the HA group. Finally, delete the VLAN interface on that node.

Steps

1. Select **CONFIGURATION > Network > VLAN interfaces**.
2. Select the check box for each VLAN interface you want to remove. Then, select **Actions > Delete**.
3. Select **Yes** to confirm your selection.

All VLAN interfaces you selected are removed. A green success banner appears on the VLAN interfaces page.

Manage high availability groups

Manage high availability (HA) groups: Overview

You can group the network interfaces of multiple Admin and Gateway Nodes into a high availability (HA) group. If the active interface in the HA group fails, a backup interface can manage the workload.

What is an HA group?

You can use high availability (HA) groups to provide highly available data connections for S3 and Swift clients or to provide highly available connections to the Grid Manager and the Tenant Manager.

Each HA group provides access to the shared services on the selected nodes.

- HA groups that include Gateway Nodes, Admin Nodes, or both provide highly available data connections for S3 and Swift clients.
- HA groups that include only Admin Nodes provide highly available connections to the Grid Manager and the Tenant Manager.
- An HA group that includes only SG100 or SG1000 appliances and VMware-based software nodes can provide highly available connections for [S3 tenants that use S3 Select](#). HA groups are recommended when using S3 Select, but not required.

How do you create an HA group?

1. You select a network interface for one or more Admin Nodes or Gateway Nodes. You can use a Grid Network (eth0) interface, Client Network (eth2) interface, VLAN interface, or an access interface you have added to the node.



You cannot add an interface to an HA group if it has a DHCP-assigned IP address.

2. You specify one interface to be the Primary interface. The Primary interface is the active interface unless a failure occurs.
3. You determine the priority order for any Backup interfaces.
4. You assign one to 10 virtual IP (VIP) addresses to the group. Clients applications can use any of these VIP addresses to connect to StorageGRID.

For instructions, see [Configure high availability groups](#).

What is the active interface?

During normal operation, all of the VIP addresses for the HA group are added to the Primary interface, which is the first interface in the priority order. As long as the Primary interface remains available, it is used when clients connect to any VIP address for the group. That is, during normal operation, the Primary interface is the “active” interface for the group.

Similarly, during normal operation, any lower priority interfaces for the HA group act as “backup” interfaces. These backup interfaces are not used unless the Primary (currently active) interface becomes unavailable.

View the current HA group status of a node

To see if a node is assigned to an HA group and determine its current status, select **NODES > node**.

If the **Overview** tab includes an entry for **HA groups**, the node is assigned to the HA groups listed. The value

after the group name is the current status of the node in the HA group:

- **Active:** The HA group is currently being hosted on this node.
- **Backup:** The HA group is not currently using this node; this is a backup interface.
- **Stopped:** The HA group cannot be hosted on this node because the High Availability (keepalived) service has been stopped manually.
- **Fault:** The HA group cannot be hosted on this node because of one or more of the following:
 - The Load Balancer (nginx-gw) service is not running on the node.
 - The node's eth0 or VIP interface is down.
 - The node is down.

In this example, the primary Admin Node has been added to two HA groups. This node is currently the active interface for the Admin clients group and a backup interface for the FabricPool clients group.

DC1-ADM1 (Primary Admin Node) [🔗](#)

Overview Hardware Network Storage Load balancer Tasks

Node information [?](#)

Name:	DC1-ADM1
Type:	Primary Admin Node
ID:	ce00d9c8-8a79-4742-bdef-c9c658db5315
Connection state:	✔ Connected
Software version:	11.6.0 (build 20211207.1804.614bc17)
HA groups:	Admin clients (Active) FabricPool clients (Backup)
IP addresses:	172.16.1.225 - eth0 (Grid Network) 10.224.1.225 - eth1 (Admin Network) 47.47.0.2, 47.47.1.225 - eth2 (Client Network) Show additional IP addresses ▼

What happens when the active interface fails?

The interface that currently hosts the VIP addresses is the active interface. If the HA group includes more than one interface and the active interface fails, the VIP addresses move to the first available backup interface in

the priority order. If that interface fails, the VIP addresses move to the next available backup interface, and so on.

Failover can be triggered for any of these reasons:

- The node on which the interface is configured goes down.
- The node on which the interface is configured loses connectivity to all other nodes for at least 2 minutes.
- The active interface goes down.
- The Load Balancer service stops.
- The High Availability service stops.



Failover might not be triggered by network failures external to the node that hosts the active interface. Similarly, failover is not triggered by the failure of the CLB service (deprecated) or services for the Grid Manager or the Tenant Manager.

The failover process generally takes only a few seconds and is fast enough that client applications should experience little impact and can rely on normal retry behaviors to continue operation.

When failure is resolved and a higher priority interface becomes available again, the VIP addresses are automatically moved to the highest priority interface that is available.

How are HA groups used?

You can use high availability (HA) groups to provide highly available connections to StorageGRID for object data and for administrative use.

- An HA group can provide highly available administrative connections to the Grid Manager or the Tenant Manager.
- An HA group can provide highly available data connections for S3 and Swift clients.
- An HA group that contains only one interface allows you to provide many VIP addresses and to explicitly set IPv6 addresses.

An HA group can provide high availability only if all nodes included in the group provide the same services. When you create an HA group, add interfaces from the types of nodes that provide the services you require.

- **Admin Nodes:** Include the Load Balancer service and enable access to the Grid Manager or the Tenant Manager.
- **Gateway Nodes:** Include the Load Balancer service and the CLB service (deprecated).

Purpose of HA group	Add nodes of this type to the HA group
Access to Grid Manager	<ul style="list-style-type: none"> • Primary Admin Node (Primary) • Non-primary Admin Nodes <p>Note: The primary Admin Node must be the Primary interface. Some maintenance procedures can only be performed from the primary Admin Node.</p>
Access to Tenant Manager only	<ul style="list-style-type: none"> • Primary or non-primary Admin Nodes

Purpose of HA group	Add nodes of this type to the HA group
S3 or Swift client access — Load Balancer service	<ul style="list-style-type: none"> • Admin Nodes • Gateway Nodes
S3 client access for S3 Select	<ul style="list-style-type: none"> • SG100 or SG1000 appliances • VMware-based software nodes <p>Note: HA groups are recommended when using S3 Select, but not required.</p>
S3 or Swift client access — CLB service Note: The CLB service is deprecated.	<ul style="list-style-type: none"> • Gateway Nodes

Limitations of using HA groups with Grid Manager or Tenant Manager

If a Grid Manager or Tenant Manager service fails, HA group failover is not triggered.

If you are signed in to the Grid Manager or the Tenant Manager when failover occurs, you are signed out and must sign in again to resume your task.

Some maintenance procedures cannot be performed when the primary Admin Node is unavailable. During failover, you can use the Grid Manager to monitor your StorageGRID system.

Limitations of using HA groups with the CLB service

The failure of the CLB service does not trigger failover within the HA group.

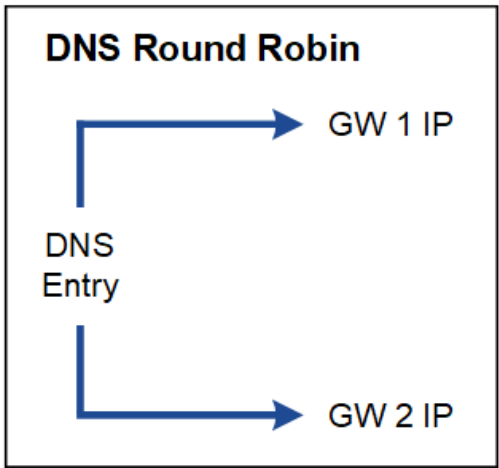



The CLB service is deprecated.


Configuration options for HA groups

The following diagrams provide examples of different ways you can configure HA groups. Each option has advantages and disadvantages.

In the diagrams, blue indicates the primary interface in the HA group and yellow indicates the backup interface in the HA group.



GW = Gateway Node
VIP = Virtual IP address



The table summarizes the benefits of each HA configuration shown in the diagram.

Configuration	Advantages	Disadvantages
Active-Backup HA	<ul style="list-style-type: none"> Managed by StorageGRID with no external dependencies. Fast failover. 	<ul style="list-style-type: none"> Only one node in an HA group is active. At least one node per HA group will be idle.
DNS Round Robin	<ul style="list-style-type: none"> Increased aggregate throughput. No idle hosts. 	<ul style="list-style-type: none"> Slow failover, which could depend on client behavior. Requires configuration of hardware outside of StorageGRID. Needs a customer-implemented health check.

Configuration	Advantages	Disadvantages
Active-Active HA	<ul style="list-style-type: none"> • Traffic is distributed across multiple HA groups. • High aggregate throughput that scales with the number of HA groups. • Fast failover. 	<ul style="list-style-type: none"> • More complex to configure. • Requires configuration of hardware outside of StorageGRID. • Needs a customer-implemented health check.

Configure high availability groups

You can configure high availability (HA) groups to provide highly available access to the services on Admin Nodes or Gateway Nodes.

What you'll need

- You are signed in to the Grid Manager using a [supported web browser](#).
- You have the Root access permission.
- If you plan to use a VLAN interface in an HA group, you have created the VLAN interface. See [Configure VLAN interfaces](#).
- If you plan to use an access interface for a node in an HA group, you have created the interface:
 - **Red Hat Enterprise Linux or CentOS (before installing the node):** [Create node configuration files](#)
 - **Ubuntu or Debian (before installing the node):** [Create node configuration files](#)
 - **Linux (after installing the node):** [Linux: Add trunk or access interfaces to a node](#)
 - **VMware (after installing the node):** [VMware: Add trunk or access interfaces to a node](#)

Create a high availability group

When you create a high availability group, you select one or more interfaces and organize them in priority order. Then, you assign one or more VIP addresses to the group.

An interface must be for a Gateway Node or an Admin Node to be included in an HA group. An HA group can only use one interface for any given node; however, other interfaces for the same node can be used in other HA groups.

Access the wizard

1. Select **CONFIGURATION > Network > High availability groups**.
2. Select **Create**.

Enter details for the HA group

1. Provide a unique name for the HA group.

Create a high availability group ✕

1 Enter details — 2 Add interfaces — 3 Prioritize interfaces — 4 Enter IP addresses

Enter details for the HA group

HA group name

Description (optional)

2. Optionally, enter a description for the HA group.
3. Select **Continue**.

Add interfaces to the HA group

1. Select one or more interfaces to add to this HA group.

Use the column headers to sort the rows, or enter a search term to locate interfaces more quickly.

Add interfaces to the HA group

Select one or more interfaces for this HA group. You can select only one interface for each node.

Search... Total interface count: 4

Node	Interface	Site	IPv4 subnet	Node type
<input type="checkbox"/> DC1-ADM1-104-96	eth0	DC1	10.96.104.0/22	Primary Admin Node
<input type="checkbox"/> DC1-ADM1-104-96	eth2	DC1	—	Primary Admin Node
<input type="checkbox"/> DC2-ADM1-104-103	eth0	DC2	10.96.104.0/22	Admin Node
<input type="checkbox"/> DC2-ADM1-104-103	eth2	DC2	—	Admin Node


0 interfaces selected



After creating a VLAN interface, wait up to 5 minutes for the new interface to appear in the table.

Guidelines for selecting interfaces

- You must select at least one interface.
- You can select only one interface for a node.
- If the HA group is for HA protection of Admin Node services, which include the Grid Manager and the Tenant Manager, select interfaces on Admin Nodes only.
- If the HA group is for HA protection of S3 or Swift client traffic, select interfaces on Admin Nodes, Gateway Nodes, or both.
- If the HA group is for HA protection of the deprecated CLB service, select interfaces on Gateway Nodes only.
- If you select interfaces on different types of nodes, an informational note appears. You are reminded that if a failover occurs, services provided by the previously active node might not be available on the newly active node. For example, a backup Gateway Node cannot provide HA protection of Admin Node services. Similarly, a backup Admin Node cannot perform all of the maintenance procedures that the primary Admin Node can provide.
- If you cannot select an interface, its check box is disabled. The tool tip provides more information.



- You cannot select an interface if its subnet value or gateway conflicts with another selected interface.
- You cannot select a configured interface if it does not have a static IP address.

2. Select **Continue**.

Determine the priority order

1. Determine the Primary interface and any Backup (failover) interfaces for this HA group.

Drag and drop rows to change the values in the **Priority order** column.

Determine the priority order

Determine the primary interface and the backup (failover) interfaces for this HA group. Drag and drop rows or select the arrows.

Priority order 	Node	Interface 	Node type 
1 (Primary interface)	 DC1-ADM1-104-96 	eth2	Primary Admin Node
2	 DC2-ADM1-104-103 	eth2	Admin Node



If the HA group provides access to the Grid Manager, you must select an interface on the primary Admin Node to be the Primary interface. Some maintenance procedures can only be performed from the primary Admin Node.

The first interface in the list is the Primary interface. The Primary interface is the active interface unless a failure occurs.

If the HA group includes more than one interface and the Primary interface fails, the VIP addresses move to the highest priority interface that is available. If that interface fails, the VIP addresses move to the next highest priority interface that is available, and so on.

2. Select **Continue**.

Enter IP addresses


1. In the **Subnet CIDR** field, specify the VIP subnet in CIDR notation—an IPv4 address followed by a slash and the subnet length (0-32).

The network address must not have any host bits set. For example, 192.16.0.0/22.




If you use a 32-bit prefix, the VIP network address also serves as the gateway address and the VIP address.

Enter details for the HA group


Subnet CIDR 

Specify the subnet in CIDR notation. The optional gateway IP and all VIPs must be in this subnet.

IPv4 address followed by a slash and the subnet length (0-32)

Gateway IP address (optional) 

Optionally specify the IP address of the gateway, which must be in the subnet. If the subnet address length is 32, the gateway IP address is automatically set to the subnet IP.

Virtual IP address 

Specify at least 1 and no more than 10 virtual IPs for the HA group. All virtual IPs must be in the same subnet. If the subnet length is 32, only one VIP is allowed, which is automatically set to the subnet/gateway IP.

[Add another IP address](#)

- Optionally, if any S3, Swift, administrative or tenant clients will access these VIP addresses from a different subnet, enter the **Gateway IP address**. The gateway address must be within the VIP subnet.

Client and admin users will use this gateway to access the virtual IP addresses.

- Enter one or more **virtual IP addresses** for the HA group. You can add up to 10 IP addresses. All VIPs must be within the VIP subnet.

You must provide at least one IPv4 address. Optionally, you can specify additional IPv4 and IPv6 addresses.

- Select **Create HA group** and select **Finish**.

The HA Group is created, and you can now use the configured virtual IP addresses.



Wait up to 15 minutes for changes to an HA group to be applied to all nodes.

Next steps

If you will use this HA group for load balancing, create a load balancer endpoint to determine the port and network protocol and to attach any required certificates. See [Configure load balancer endpoints](#).

Edit a high availability group

You can edit a high availability (HA) group to change its name and description, add or remove interfaces, change the priority order, or add or update virtual IP addresses.

For example, you might need to edit an HA group if you want to remove the node associated with a selected interface in a site or node decommission procedure.

Steps

1. Select **CONFIGURATION > Network > High availability groups**.

The High availability groups page shows all existing HA groups.

High availability groups [Learn more about HA groups](#)

You can group the network interfaces of multiple Admin and Gateway Nodes into a high availability (HA) group. If the active interface in the group fails, a backup interface can manage the workload.

Each HA group provides access to the shared services on the selected nodes. Select Gateway Nodes, Admin Nodes, or both for load balancing. Select Admin Nodes for management services. All interfaces in a group must be in the same subnet. You assign one or more virtual IP addresses (VIPs) to each group. Clients use these VIPs to connect to StorageGRID.

- You cannot select an interface if it has a DHCP-assigned IP address.
- Wait up to 15 minutes for changes to an HA group to be applied to all nodes.

Create Actions Search... Total HA groups count: 2

<input type="checkbox"/>	Name ?	Description ?	Virtual IP address ?	Interfaces (in priority order) ?
<input type="checkbox"/>	FabricPool	Use for FabricPool client access	10.96.104.5 10.96.104.6	DC1-ADM1-104-96:eth2 (active) DC2-ADM1-104-103:eth2
<input type="checkbox"/>	S3 Clients	use for S3 client access	10.96.104.10	DC1-ADM1-104-96:eth0 DC2-ADM1-104-103:eth0

< Previous 1 Next >

2. Select the check box for the HA group you want to edit.
3. Do one of the following, based on what you want to update:
 - Select **Actions > Edit virtual IP address** to add or remove VIP addresses.
 - Select **Actions > Edit HA group** to update the group's name or description, add or remove interfaces, change the priority order, or add or remove VIP addresses.
4. If you selected **Edit virtual IP address**:
 - a. Update the virtual IP addresses for the HA group.
 - b. Select **Save**.
 - c. Select **Finish**.
5. If you selected **Edit HA group**:
 - a. Optionally, update the group's name or description.
 - b. Optionally, select or unselect the check boxes to add or remove interfaces.



If the HA group provides access to the Grid Manager, you must select an interface on the primary Admin Node to be the Primary interface. Some maintenance procedures can only be performed from the primary Admin Node

- c. Optionally, drag and drop rows to change the priority order of the Primary interface and any Backup interfaces for this HA group.
- d. Optionally, update the virtual IP addresses.
- e. Select **Save** and then select **Finish**.



Wait up to 15 minutes for changes to an HA group to be applied to all nodes.

Remove a high availability group

You can remove one or more high availability (HA) groups at a time. However, you cannot remove an HA group if it is bound to one or more load balancer endpoints.

To prevent client disruptions, update any affected S3 or Swift client applications before you remove an HA group. Update each client to connect using another IP address, for example, the virtual IP address of a different HA group or the IP address that was configured for an interface during installation.

Steps

1. Select **CONFIGURATION > Network > High availability groups**.
2. Select the check box for each HA group you want to remove. Then, select **Actions > Remove HA group**.
3. Review the message and select **Delete HA group** to confirm your selection.

All HA groups you selected are removed. A green success banner appears on the High availability groups page.

Manage load balancing

Manage load balancing: Overview

You can use the StorageGRID load balancing functions to handle ingest and retrieval workloads from S3 and Swift clients. Load balancing maximizes speed and connection capacity by distributing the workloads and connections across multiple Storage Nodes.

You can load balance client workloads in the following ways:

- Use the Load Balancer service, which is installed on Admin Nodes and Gateway Nodes. The Load Balancer service provides Layer 7 load balancing and performs TLS termination of client requests, inspects the requests, and establishes new secure connections to the Storage Nodes. This is the recommended load balancing mechanism.

See [How load balancing works - Load Balancer service](#).

- Use the deprecated Connection Load Balancer (CLB) service, which is installed on Gateway Nodes only. The CLB service provides Layer 4 load balancing and supports link costs.

See [How load balancing works - CLB service \(deprecated\)](#).

- Integrate a third-party load balancer. Contact your NetApp account representative for details.

How load balancing works - Load Balancer service

The Load Balancer service distributes incoming network connections from client applications to Storage Nodes. To enable load balancing, you must configure load balancer endpoints using the Grid Manager.

You can configure load balancer endpoints only for Admin Nodes or Gateway Nodes, since these node types contain the Load Balancer service. You cannot configure endpoints for Storage Nodes or Archive Nodes.

Each load balancer endpoint specifies a port, a network protocol (HTTP or HTTPS), a client type (S3 or Swift), and a binding mode. HTTPS endpoints require a server certificate. Binding modes allow you to restrict the accessibility of endpoint ports to:

- The virtual IP addresses (VIPs) of specific high availability (HA) groups
- Specific network interfaces of specific Admin and Gateway Nodes

Port considerations

Clients can access any of the endpoints you configure on any node running the Load Balancer service, with two exceptions: ports 80 and 443 are reserved on Admin Nodes, so endpoints configured on these ports support load balancing operations only on Gateway Nodes.

If you have remapped any ports, you cannot use the same ports to configure load balancer endpoints. You can create endpoints using remapped ports, but those endpoints will be remapped to the original CLB ports and service, not the Load Balancer service. Follow the steps in [Remove port remaps](#).



The CLB service is deprecated.

CPU availability

The Load Balancer service on each Admin Node and Gateway Node operates independently when forwarding S3 or Swift traffic to the Storage Nodes. Through a weighting process, the Load Balancer service routes more requests to Storage Nodes with higher CPU availability. Node CPU load information is updated every few minutes, but weighting might be updated more frequently. All Storage Nodes are assigned a minimal base weight value, even if a node reports 100% utilization or fails to report its utilization.

In some cases, information about CPU availability is limited to the site where the Load Balancer service is located.

Configure load balancer endpoints

Load balancer endpoints determine the ports and network protocols S3 and Swift clients can use when connecting to the StorageGRID load balancer on Gateway and Admin Nodes.

What you'll need

- You are signed in to the Grid Manager using a [supported web browser](#).
- You have the Root access permission.

- If you previously remapped a port you intend to use for the load balancer endpoint, you have [removed the port remap](#).
- You have created any high availability (HA) groups you plan to use. HA groups are recommended, but not required. See [Manage high availability groups](#).
- If the load balancer endpoint will be used by [S3 tenants for S3 Select](#), it must not use the IP addresses or FQDNs of any bare-metal nodes. Only SG100 or SG1000 appliances and VMware-based software nodes are allowed for the load balancer endpoints used for S3 Select.
- You have configured any VLAN interfaces you plan to use. See [Configure VLAN interfaces](#).
- If you are creating an HTTPS endpoint (recommended), you have the information for the server certificate.



Changes to an endpoint certificate can take up to 15 minutes to be applied to all nodes.

- To upload a certificate, you need the server certificate, the certificate private key, and optionally, a CA bundle.
- To generate a certificate, you need all of the domain names and IP addresses that S3 or Swift clients will use to access the endpoint. You must also know the subject (Distinguished Name).
- If you want to use the StorageGRID S3 and Swift API certificate (which can also be used for connections directly to Storage Nodes), you have already replaced the default certificate with a custom certificate signed by an external certificate authority. See [Configure S3 and Swift API certificates](#).

The certificate can use wildcards to represent the fully qualified domain names of all Admin Nodes and Gateway Nodes running the Load Balancer service. For example, `*.storagegrid.example.com` uses the `*` wildcard to represent `adm1.storagegrid.example.com` and `gn1.storagegrid.example.com`. See [Configure S3 API endpoint domain names](#).

Create a load balancer endpoint

Each load balancer endpoint specifies a port, a client type (S3 or Swift), and a network protocol (HTTP or HTTPS).

Access the wizard

1. Select **CONFIGURATION > Network > Load balancer endpoints**.
2. Select **Create**.

Enter endpoint details

1. Enter details for the endpoint.

Create a load balancer endpoint ✕

1 Enter endpoint details
 2 Select binding mode
 3 Attach certificate

Endpoint details

Name ?

Port ?

Enter an unused port or accept the suggested port.

Client type ?

Select the type of client application that will use this endpoint.

S3
 Swift

Network protocol ?

Select the network protocol clients will use with this endpoint. If you select HTTPS, attach the security certificate before saving the endpoint.

HTTPS (recommended)
 HTTP

Cancel
Continue

Field	Description
Name	A descriptive name for the endpoint, which will appear in the table on the Load balancer endpoints page.
Port	<p>The port clients will use to connect to the Load Balancer service on Admin Nodes and Gateway Nodes.</p> <p>Accept the suggested port number or enter any external port that is not used by another grid service. Enter a value between 1 and 65535.</p> <p>If you enter 80 or 443, the endpoint is configured only on Gateway Nodes. These ports are reserved on Admin Nodes.</p> <p>See the Networking guidelines for information about external ports.</p>
Client type	The type of client application that will use this endpoint, either S3 or Swift .

Field	Description
Network protocol	<p>The network protocol that clients will use when connecting to this endpoint.</p> <ul style="list-style-type: none"> • Select HTTPS for secure, TLS encrypted communication (recommended). You must attach a security certificate before you can save the endpoint. • Select HTTP for less secure, unencrypted communication. Use HTTP only for a non-production grid.

2. Select **Continue**.

Select the binding mode

1. Select a binding mode for the endpoint to control how the endpoint is accessed.

Option	Description
Global (default)	<p>Clients can access the endpoint using a fully qualified domain name (FQDN), the IP address of any Gateway Node or Admin Node, or the virtual IP address of any HA group on any network.</p> <p>Use the Global setting (default) unless you need to restrict the accessibility of this endpoint.</p>
Node interfaces	<p>Clients must use the IP address of a selected node and network interface to access this endpoint.</p>
Virtual IPs of HA groups	<p>Clients must use a virtual IP address of an HA group to access this endpoint.</p> <p>Endpoints with this binding mode can all use the same port number, as long as the HA groups you select for the endpoints do not overlap.</p> <p>Endpoints with this mode can all use the same port number as long as the interfaces you select for the endpoints do not overlap.</p>



If you use the same port for more than one endpoint, an endpoint using **Virtual IPs of HA groups** mode overrides an endpoint using **Node interfaces** mode, which overrides an endpoint using **Global** mode.

2. If you selected **Node interfaces**, select one or more node interfaces for each Admin Node or Gateway Node that you want to associate with this endpoint.

Binding mode ?

Select a binding mode if you plan to monitor or limit the use of this endpoint with a traffic classification policy.

The binding mode controls how the endpoint is accessed—using any IP address or using specific IP addresses and network interfaces.

Global Node interfaces Virtual IPs of HA groups

If you use the same port for more than one endpoint, an endpoint bound to HA groups overrides an endpoint bound to Node interfaces, which overrides a Global endpoint. If this behavior does not meet your requirements, consider using a different port number for each endpoint.

Search... Total interface count: 3

<input type="checkbox"/>	Node ?	Node interface ?	Site ?	IP address ?	Node type ?
<input type="checkbox"/>	DC1-ADM1	eth0 ?	Data Center 1	172.16.3.246 and 2 more	Primary Admin Node
<input type="checkbox"/>	DC1-ADM1	eth1 ?	Data Center 1	10.224.3.246 and 5 more	Primary Admin Node
<input type="checkbox"/>	DC1-ADM1	eth2 ?	Data Center 1	47.47.3.246 and 3 more	Primary Admin Node

3. If you selected **Virtual IPs of HA groups**, select one or more HA groups.

Binding mode ?

Select a binding mode if you plan to monitor or limit the use of this endpoint with a traffic classification policy.

The binding mode controls how the endpoint is accessed—using any IP address or using specific IP addresses and network interfaces.

Global Node interfaces Virtual IPs of HA groups

If you use the same port for more than one endpoint, an endpoint bound to HA groups overrides an endpoint bound to Node interfaces, which overrides a Global endpoint. If this behavior does not meet your requirements, consider using a different port number for each endpoint.

Search... Total interface count: 2

<input type="checkbox"/>	Name ?	Description ?	Virtual IP address ?	Interfaces (in priority order) ?
<input type="checkbox"/>	FabricPool	Use for FabricPool client access	10.96.104.5 10.96.104.6	DC1-ADM1-104-96:eth2 (active) DC2-ADM1-104-103:eth2
<input type="checkbox"/>	S3 Clients	use for S3 client access	10.96.104.10	DC1-ADM1-104-96:eth0 DC2-ADM1-104-103:eth0

4. If you are creating an **HTTP** endpoint, you do not need to attach a certificate. Select **Create** to add the new load balancer endpoint. Then, go to [After you finish](#). Otherwise, select **Continue** to attach the certificate.

Attach certificate

1. If you are creating an **HTTPS** endpoint, select the type of security certificate you want to attach to the endpoint.

The certificate secures the connections between S3 and Swift clients and the Load Balancer service on Admin Node or Gateway Nodes.

- **Upload certificate.** Select this option if you have custom certificates to upload.
- **Generate certificate.** Select this option if you have the values needed to generate a custom certificate.
- **Use StorageGRID S3 and Swift certificate.** Select this option if you want to use the global S3 and Swift API certificate, which can also be used for connections directly to Storage Nodes.

You cannot select this option unless you have replaced the default S3 and Swift API certificate, which is signed by the grid CA, with a custom certificate signed by an external certificate authority. See [Configure S3 and Swift API certificates](#).

2. If you are not using the StorageGRID S3 and Swift certificate, upload or generate the certificate.

Upload certificate

- a. Select **Upload certificate**.
- b. Upload the required server certificate files:
 - **Server certificate**: The custom server certificate file in PEM encoding.
 - **Certificate private key**: The custom server certificate private key file (.key).



EC private keys must be 224 bits or larger. RSA private keys must be 2048 bits or larger.

- **CA bundle**: A single optional file containing the certificates from each intermediate issuing certificate authority (CA). The file should contain each of the PEM-encoded CA certificate files, concatenated in certificate chain order.
- c. Expand **Certificate details** to see the metadata for each certificate you uploaded. If you uploaded an optional CA bundle, each certificate displays on its own tab.

- Select **Download certificate** to save the certificate file or select **Download CA bundle** to save the certificate bundle.

Specify the certificate file name and download location. Save the file with the extension .pem.

For example: storagegrid_certificate.pem

- Select **Copy certificate PEM** or **Copy CA bundle PEM** to copy the certificate contents for pasting elsewhere.
- d. Select **Create**.

The load balancer endpoint is created. The custom certificate is used for all subsequent new connections between S3 and Swift clients and the endpoint.

Generate certificate

- a. Select **Generate certificate**.
- b. Specify the certificate information:
 - **Domain name**: One or more fully qualified domain names to include in the certificate. Use an * as a wildcard to represent multiple domain names.
 - **IP**: One or more IP addresses to include in the certificate.
 - **Subject**: X.509 subject or distinguished name (DN) of the certificate owner.
 - **Days valid**: Number of days after creation that the certificate expires.

- c. Select **Generate**.

- d. Select **Certificate details** to see the metadata for the generated certificate.

- Select **Download certificate** to save the certificate file.

Specify the certificate file name and download location. Save the file with the extension .pem.

For example: storagegrid_certificate.pem

- Select **Copy certificate PEM** to copy the certificate contents for pasting elsewhere.
- e. Select **Create**.

The load balancer endpoint is created. The custom certificate is used for all subsequent new connections between S3 and Swift clients and this endpoint.

After you finish

1. If you use a domain name system (DNS), ensure that the DNS includes a record to associate the StorageGRID fully qualified domain name to each IP address that clients will use to make connections.

The IP address you enter in the DNS record depends on whether you are using an HA group of load-balancing nodes:

- If you have configured a HA group, clients will connect to the virtual IP addresses of that HA group.
- If you are not using a HA group, clients will connect to the StorageGRID Load Balancer service using the IP address of any Gateway Node or Admin Node.

You must also ensure that the DNS record references all required endpoint domain names, including any wildcard names.

2. Provide S3 and Swift clients with the information needed to connect to the endpoint:


- Port number
- Fully qualified domain name or IP address
- Any required certificate details






View and edit load balancer endpoints

You can view details for existing load balancer endpoints, including the certificate metadata for a secured endpoint. You can also change an endpoint's name or binding mode and update any associated certificates.

You cannot change the service type (S3 or Swift), the port, or the protocol (HTTP or HTTPS).

- To view basic information for all load balancer endpoints, review the table on the Load balancer endpoints page.



<input type="checkbox"/>	Name 	Port 	Network protocol 	Binding mode 	Certificate expiration 
<input type="checkbox"/>	FabricPool endpoint	10443	HTTPS	Global	Oct 19th, 2022

- To view all details about a specific endpoint, including certificate metadata, select the endpoint's name in the table.

FabricPool endpoint

Port: 10443
 Client type: S3
 Network protocol: HTTPS
 Binding mode: Global
 Endpoint ID: c2b6feb3-c567-449d-b717-4fed98c4a411

Remove

Binding Mode

Certificate

You can select a different binding mode or change IP addresses for the current binding mode.

Edit binding mode

Binding mode: Global




This endpoint uses the Global binding mode. Unless there are one or more overriding endpoints for the same port, clients can access this endpoint using the IP address of any Gateway Node, any Admin Node, or the virtual IP of any HA group on any network.

- To edit an endpoint, use the **Actions** menu on the Load balancer endpoints page or the details page for a specific endpoint.



After editing an endpoint, you might need to wait up to 15 minutes for your changes to be applied to all nodes.

Task	Actions menu	Details page
Edit endpoint name	<ol style="list-style-type: none"> Select the check box for the endpoint. Select Actions > Edit endpoint name. Enter the new name. Select Save. 	<ol style="list-style-type: none"> Select the endpoint name to display the details. Select the edit icon . Enter the new name. Select Save.
Edit endpoint binding mode	<ol style="list-style-type: none"> Select the check box for the endpoint. Select Actions > Edit endpoint binding mode. Update the binding mode as required. Select Save changes. 	<ol style="list-style-type: none"> Select the endpoint name to display the details. Select Edit binding mode. Update the binding mode as required. Select Save changes.

Task	Actions menu	Details page
Edit endpoint certificate	<ol style="list-style-type: none"> Select the check box for the endpoint. Select Actions > Edit endpoint certificate. Upload or generate a new custom certificate or begin using the global S3 and Swift certificate, as required. Select Save changes. 	<ol style="list-style-type: none"> Select the endpoint name to display the details. Select the Certificate tab. Select Edit certificate. Upload or generate a new custom certificate or begin using the global S3 and Swift certificate, as required. Select Save changes.

Remove load balancer endpoints

You can remove one or more endpoints using the **Actions** menu, or you can remove a single endpoint from the details page.



To prevent client disruptions, update any affected S3 or Swift client applications before you remove a load balancer endpoint. Update each client to connect using a port assigned to another load balancer endpoint. Be sure to update any required certificate information as well.

- To remove one or more endpoints:
 - From the Load balancer page, select the check box for each endpoint you want to remove.
 - Select **Actions > Remove**.
 - Select **OK**.
- To remove one endpoint from the details page:
 - From the Load balancer page, select the endpoint name.
 - Select **Remove** on the details page.
 - Select **OK**.

How load balancing works - CLB service (deprecated)

The Connection Load Balancer (CLB) service on Gateway Nodes is deprecated. The Load Balancer service is now the recommended load balancing mechanism.

The CLB service uses Layer 4 load balancing to distribute incoming TCP network connections from client applications to the optimal Storage Node based on availability, system load, and the administrator-configured link cost. When the optimal Storage Node is chosen, the CLB service establishes a two-way network connection and forwards the traffic to and from the chosen node. The CLB does not consider the Grid Network configuration when directing incoming network connections.

To view information about the CLB service, select **SUPPORT > Tools > Grid topology**, and then expand a Gateway Node until you can select **CLB** and the options below it.

The screenshot shows the Grid Manager interface. On the left is a tree view of the StorageGRID Webscale Deployment topology. On the right is the 'Overview: Summary - DC1-G1-98-161' page, which includes a 'Storage Capacity' table.

Storage Capacity	
Storage Nodes Installed:	N/A
Storage Nodes Readable:	N/A
Storage Nodes Writable:	N/A
Installed Storage Capacity:	N/A
Used Storage Capacity:	N/A
Used Storage Capacity for Data:	N/A
Used Storage Capacity for Metadata:	N/A
Usable Storage Capacity:	N/A

If you choose to use the CLB service, you should consider configuring link costs for your StorageGRID system.

- [What link costs are](#)
- [Update link costs](#)

Configure S3 API endpoint domain names

To support S3 virtual hosted-style requests, you must use the Grid Manager to configure the list of endpoint domain names that S3 clients connect to.

What you'll need

- You are signed in to the Grid Manager using a [supported web browser](#).
- You have specific access permissions.
- You have confirmed that a grid upgrade is not in progress.



Do not make any changes to the domain name configuration when a grid upgrade is in progress.

About this task

To enable clients to use S3 endpoint domain names, you must do all of the following:

- Use the Grid Manager to add the S3 endpoint domain names to the StorageGRID system.
- Ensure that the certificate the client uses for HTTPS connections to StorageGRID is signed for all domain names that the client requires.

For example, if the endpoint is `s3.company.com`, you must ensure that the certificate used for HTTPS connections includes the `s3.company.com` endpoint and the endpoint's wildcard Subject Alternative Name (SAN): `*.s3.company.com`.

- Configure the DNS server used by the client. Include DNS records for the IP addresses that clients use to make connections, and ensure that the records reference all required endpoint domain names, including any wildcard names.



Clients can connect to StorageGRID using the IP address of a Gateway Node, an Admin Node, or a Storage Node, or by connecting to the virtual IP address of a high availability group. You should understand how client applications connect to the grid so you include the correct IP addresses in the DNS records.

Clients that use HTTPS connections (recommended) to the grid can use either of these certificates:

- Clients that connect to a load balancer endpoint can use a custom certificate for that endpoint. Each load balancer endpoint can be configured to recognize different endpoint domain names.
- Clients that connect to a load balancer endpoint, directly to a Storage Node, or directly to the deprecated CLB service on a Gateway Node can customize the global S3 and Swift API certificate to include all required endpoint domain names.

Steps

1. Select **CONFIGURATION** > **Network** > **Domain names**.

The Endpoint Domain Names page appears.

Endpoint Domain Names

Virtual Hosted-Style Requests

Enable support of S3 virtual hosted-style requests by specifying API endpoint domain names. Support is disabled if this list is empty. Examples: s3.example.com, s3.example.co.uk, s3-east.example.com

Endpoint 1	<input type="text" value="s3.example.com"/>	✕
Endpoint 2	<input type="text"/>	+ ✕

2. Enter the list of S3 API endpoint domain names in the **Endpoint** fields. Use the **+** icon to add additional fields.

If this list is empty, support for S3 virtual hosted-style requests is disabled.

3. Select **Save**.
4. Ensure that the server certificates that clients use match the required endpoint domain names.
 - If clients connect to a load balancer endpoint that uses its own certificate, update the certificate associated with the endpoint.
 - If clients connect to a load balancer endpoint that uses the global S3 and Swift API certificate, directly to Storage Nodes, or to the CLB service on Gateway Nodes, update the global S3 and Swift API certificate.
5. Add the DNS records required to ensure that endpoint domain name requests can be resolved.

Result

Now, when clients use the endpoint `bucket.s3.company.com`, the DNS server resolves to the correct endpoint and the certificate authenticates the endpoint as expected.

Related information

- [Use S3](#)

- [View IP addresses](#)
- [Configure high availability groups](#)
- [Configure S3 and Swift API certificates](#)
- [Configure load balancer endpoints](#)

Enable HTTP for client communications

By default, client applications use the HTTPS network protocol for all connections to Storage Nodes or to the deprecated CLB service on Gateway Nodes. You can optionally enable HTTP for these connections, for example, when testing a non-production grid.

What you'll need

- You are signed in to the Grid Manager using a [supported web browser](#).
- You have specific access permissions.

About this task

Complete this task only if S3 and Swift clients need to make HTTP connections directly to Storage Nodes or to the deprecated CLB service on Gateway Nodes.

You do not need to complete this task for clients that only use HTTPS connections or for clients that connect to the Load Balancer service (because you can configure each Load Balancer endpoint to use either HTTP or HTTPS). See the information on configuring load balancer endpoints for more information.

See [Summary: IP addresses and ports for client connections](#) to learn which ports S3 and Swift clients use when connecting to Storage Nodes or to the deprecated CLB service using HTTP or HTTPS



Be careful when enabling HTTP for a production grid because requests will be sent unencrypted.

Steps

1. Select **CONFIGURATION > System > Grid options**.
2. In the Network Options section, select the **Enable HTTP Connection** check box.

Network Options



3. Select **Save**.

Related information

- [Configure load balancer endpoints](#)
- [Use S3](#)

- [Use Swift](#)

Control which client operations are permitted

You can select the Prevent Client Modification grid option to deny specific HTTP client operations.

What you'll need

- You are signed in to the Grid Manager using a [supported web browser](#).
- You have specific access permissions.

About this task

Prevent Client Modification is a system wide setting. When the Prevent Client Modification option is selected, the following requests are denied:

• S3 REST API

- Delete Bucket requests
- Any requests to modify an existing object's data, user-defined metadata, or S3 object tagging



This setting does not apply to buckets with versioning enabled. Versioning already prevents modifications to object data, user-defined metadata, and object tagging.

• Swift REST API

- Delete Container requests
- Requests to modify any existing object. For example, the following operations are denied: Put Overwrite, Delete, Metadata Update, and so on.

Steps

1. Select **CONFIGURATION > System > Grid options**.
2. In the Network Options section, select the **Prevent Client Modification** check box.

Network Options

The screenshot shows the 'Network Options' configuration section. It contains three settings:

- Prevent Client Modification**: A checkbox that is checked, highlighted with a yellow rounded rectangle. A blue question mark icon is located below the text.
- Enable HTTP Connection**: An unchecked checkbox. A blue question mark icon is located below the text.
- Network Transfer Encryption**: A radio button selection. The 'AES256-SHA' option is selected (indicated by a filled circle), while 'AES128-SHA' is unselected (indicated by an empty circle). A blue question mark icon is located below the text.

3. Select **Save**.

Copyright Information

Copyright © 2022 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.

RESTRICTED RIGHTS LEGEND: Use, duplication, or disclosure by the government is subject to restrictions as set forth in subparagraph (c)(1)(ii) of the Rights in Technical Data and Computer Software clause at DFARS 252.277-7103 (October 1988) and FAR 52-227-19 (June 1987).

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.