Maintain

StorageGRID 11.7

NetApp
April 12, 2024

Table of Contents

Maintain

Expand your grid

Expand your grid: Overview

Use these instructions to expand the capacity or capabilities of your StorageGRID system without interrupting system operations.

About these instructions

These instructions describe how to perform a StorageGRID expansion to add storage volumes to Storage Nodes, new grid nodes to an existing site, or an entire new site.

These instructions are for technical personnel who are responsible for configuring and supporting the StorageGRID system after it has been installed.

Expansion workflow

The reason you are performing the expansion determines how many new nodes of each type you must add and the location of those new nodes. For example, there are different node requirements if you are performing an expansion to increase storage capacity, add metadata capacity, or add redundancy or new capabilities.

As shown in the workflow, the steps for performing an expansion depend on whether you are adding storage volumes to a Storage Node, adding new nodes to an existing site, or adding a new site. In all cases, you can perform the expansion without interrupting the operation of your current system.

The steps for adding nodes also depend on whether you are adding StorageGRID appliances or hosts running VMware or Linux.

> ⓘ “Linux” refers to a Red Hat® Enterprise Linux®, Ubuntu®, CentOS, or Debian® deployment. Use the NetApp Interoperability Matrix Tool (IMT) to get a list of supported versions.

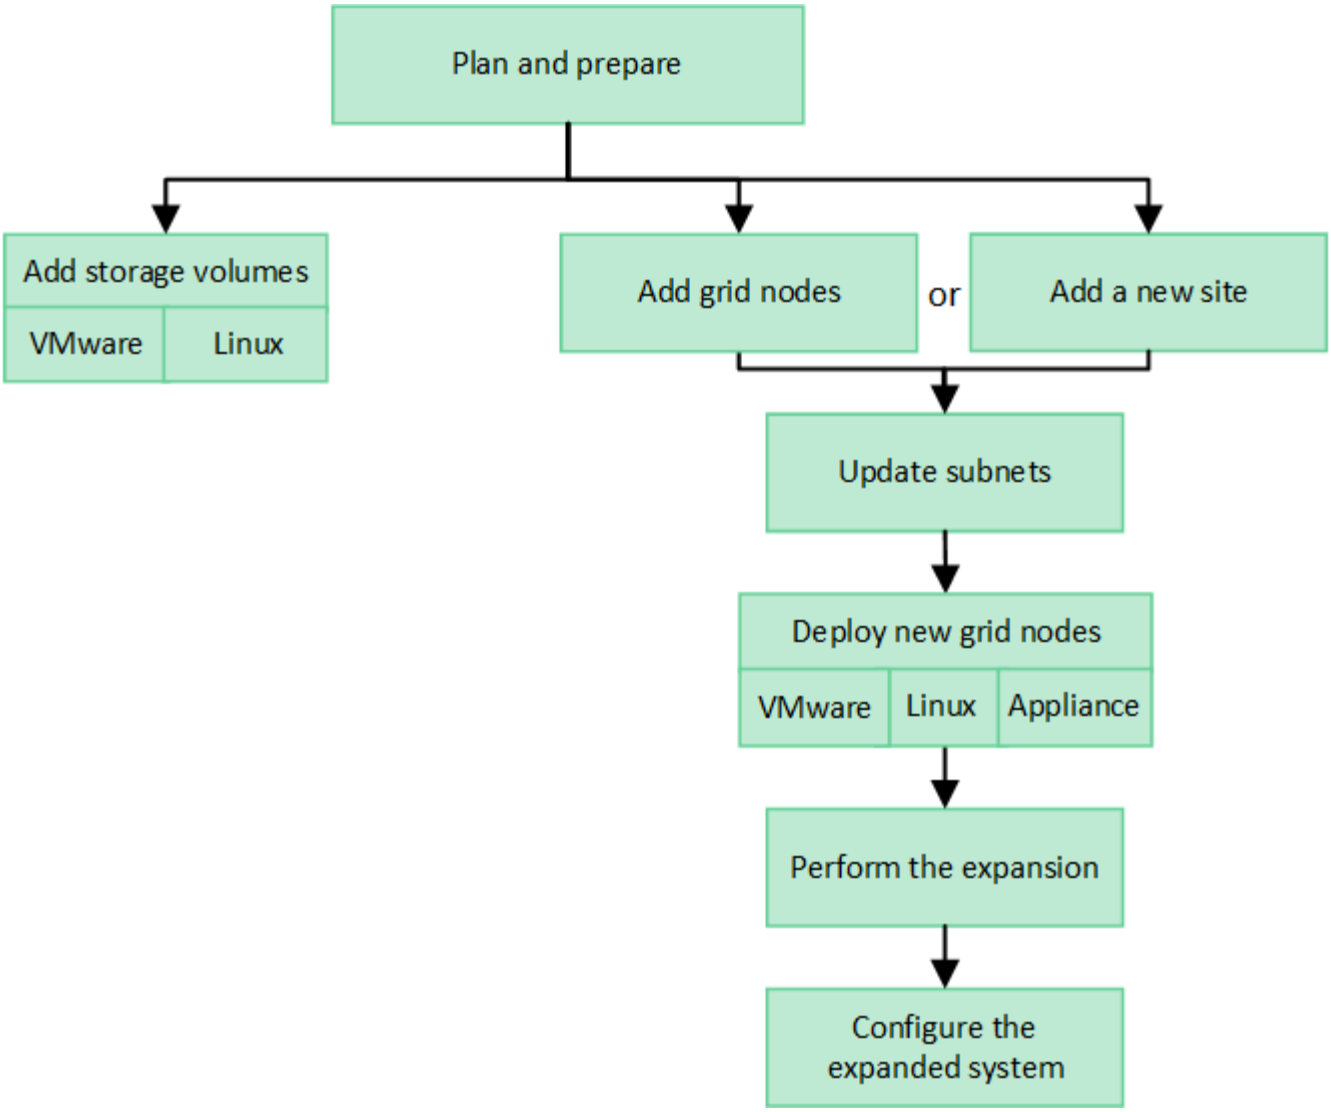

Plan StorageGRID expansion

Add storage capacity

Guidelines for adding object capacity

You can expand the object storage capacity of your StorageGRID system by adding storage volumes to existing Storage Nodes or by adding new Storage Nodes to existing sites. You must add storage capacity in a way that meets the requirements of your information lifecycle management (ILM) policy.

Guidelines for adding storage volumes

Before adding storage volumes to existing Storage Nodes, review the following guidelines and limitations:

- You must examine your current ILM rules to determine where and when to add storage volumes to increase the storage available for replicated objects or erasure-coded objects.
- You can't increase the metadata capacity of your system by adding storage volumes because object metadata is stored only on volume 0.

- Each software-based Storage Node can support a maximum of 16 storage volumes. If you need to add capacity beyond that, you must add new Storage Nodes.
- You can add one or two expansion shelves to each SG6060 appliance. Each expansion shelf adds 16 storage volumes. With both expansion shelves installed, the SG6060 can support a total of 48 storage volumes.
- You can't add storage volumes to any other storage appliance.
- You can't increase the size of an existing storage volume.
- You can't add storage volumes to a Storage Node at the same time you are performing a system upgrade, recovery operation, or another expansion.

After you have decided to add storage volumes and have determined which Storage Nodes you must expand to satisfy your ILM policy, follow the instructions for your type of Storage Node:

- To add one or two expansion shelves to an SG6060 storage appliance, go to Add expansion shelf to deployed SG6060.
- For a software-based node, follow the instructions for adding storage volumes to Storage Nodes.

Guidelines for adding Storage Nodes

Before adding Storage Nodes to existing sites, review the following guidelines and limitations:

- You must examine your current ILM rules to determine where and when to add Storage Nodes to increase the storage available for replicated objects or erasure-coded objects.
- You should not add more than 10 Storage Nodes in a single expansion procedure.
- You can add Storage Nodes to more than one site in a single expansion procedure.
- You can add Storage Nodes and other types of nodes in a single expansion procedure.
- Before starting the expansion procedure, you must confirm that all data-repair operations performed as part of a recovery are complete. See Check data repair jobs.
- If you need to remove Storage Nodes before or after performing an expansion, you should not decommission more than 10 Storage Nodes in a single Decommission Node procedure.

Guidelines for ADC service on Storage Nodes

When configuring the expansion, you must choose whether to include the Administrative Domain Controller (ADC) service on each new Storage Node. The ADC service keeps track of the location and availability of grid services.

- The StorageGRID system requires a quorum of ADC services to be available at each site and at all times.
- At least three Storage Nodes at each site must include the ADC service.
- Adding the ADC service to every Storage Node is not recommended. Including too many ADC services can cause slowdowns due to the increased amount of communication between nodes.
- A single grid should not have more than 48 Storage Nodes with the ADC service. This is equivalent to 16 sites with three ADC services at each site.
- In general, when you select the **ADC Service** setting for a new node, you should select **Automatic**. Select **Yes** only if the new node will replace another Storage Node that includes the ADC service. Because you can't decommission a Storage Node if too few ADC services would remain, this ensures that a new ADC service is available before the old service is removed.

- You can’t add the ADC service to a node after it is deployed.

Add storage capacity for replicated objects

If the information lifecycle management (ILM) policy for your deployment includes a rule that creates replicated copies of objects, you must consider how much storage to add and where to add the new storage volumes or Storage Nodes.

For guidance on where to add additional storage, examine the ILM rules that create replicated copies. If ILM rules create two or more object copies, plan to add storage in each location where object copies are made. As a simple example, if you have a two-site grid and an ILM rule that creates one object copy at each site, you must add storage to each site to increase the overall object capacity of the grid. For information about object replication, see What is replication?.

For performance reasons, you should attempt to keep storage capacity and compute power balanced across sites. So, for this example, you should add the same number of Storage Nodes to each site or additional storage volumes at each site.

If you have a more complex ILM policy that includes rules that place objects in different locations based on criteria such as bucket name, or rules that change object locations over time, your analysis of where storage is required for the expansion will be similar, but more complex.

Charting how quickly overall storage capacity is being consumed can help you understand how much storage to add in the expansion, and when the additional storage space will be required. You can use the Grid Manager to monitor and chart storage capacity.

When planning the timing of an expansion, remember to consider how long it might take to procure and install additional storage.

Add storage capacity for erasure-coded objects

If your ILM policy includes a rule that makes erasure-coded copies, you must plan where to add new storage and when to add new storage. The amount of storage you add and the timing of the addition can affect the grid’s usable storage capacity.

The first step in planning a storage expansion is to examine the rules in your ILM policy that create erasure-coded objects. Because StorageGRID creates *k*+*m* fragments for every erasure-coded object and stores each fragment on a different Storage Node, you must ensure that at least *k*+*m* Storage Nodes have space for new erasure-coded data after the expansion. If the erasure coding profile provides site-loss protection, you must add storage to each site. See What are erasure coding schemes? for information about erasure coding profiles.

The number of nodes you need to add also depends on how full the existing nodes are when you perform the expansion.

General recommendation for adding storage capacity for erasure-coded objects

If you want to avoid detailed calculations, you can add two Storage Nodes per site when existing Storage Nodes reach 70% capacity.

This general recommendation provides reasonable results across a wide range of erasure-coding schemes for both single-site grids and for grids where erasure coding provides site-loss protection.

To better understand the factors that led to this recommendation or to develop a more precise plan for your

site, see Considerations for rebalancing erasure-coded data. For a custom recommendation optimized for your situation, contact your NetApp Professional Services consultant.

Considerations for rebalancing erasure-coded data

If you are performing an expansion to add Storage Nodes and you use ILM rules to erasure code data, you might need to perform the EC rebalance procedure if you can't add enough Storage Nodes for the erasure-coding scheme you are using.

After reviewing these considerations, perform the expansion, and then go to Rebalance erasure-coded data after adding Storage Nodes to run the procedure.

What is EC rebalancing?

EC rebalancing is a StorageGRID procedure that might be required after a Storage Node expansion. The procedure is run as a command-line script from the primary Admin Node. When you run the EC rebalance procedure, StorageGRID redistributes erasure-coded fragments among the existing and the newly added Storage Nodes at a site.

The EC rebalance procedure:

- Only moves erasure-coded object data. It does not move replicated object data.
- Redistributes the data within a site. It does not move data between sites.
- Redistributes data among all Storage Nodes at a site. It does not redistribute data within storage volumes.
- Does not consider the replicated data usage on each Storage Node when determining where to move erasure-coded data.
- Redistributes erasure-coded data evenly between Storage Nodes without considering the relative capacities of each node.
- Will not distribute erasure-coded data to Storage Nodes that are more than 70% full.
- Might decrease the performance of ILM operations and S3 and Swift client operations when it runs—additional resources are required to redistribute the erasure-coding fragments.
- Might need to be run more than once for systems with very large numbers of erasure-coded objects. To limit resource usage, the maximum number of moves is capped for each job.

When the EC rebalance procedure is complete:

- Erasure-coded data will have moved from Storage Nodes with less available space to Storage Nodes with more available space.
- The data protection of erasure-coded objects will be unchanged.
- Used (%) values might be different between Storage Nodes for two reasons:
  - Replicated object copies will continue to consume space on the existing nodes—the EC rebalance procedure does not move replicated data.
  - Larger-capacity nodes will be relatively less full than smaller-capacity nodes, even though all nodes will end up with approximately the same amount of erasure-coded data.

    For example, suppose three 200-TB nodes are each filled to 80% (200 × 0.8 = 160 TB on each node, or 480 TB for the site). If you add a 400-TB node and run the rebalance procedure, all nodes will now have approximately the same amount of erasure-code data (480/4 = 120 TB). However, the Used (%) for the larger node will be less than the Used (%) for the smaller nodes.

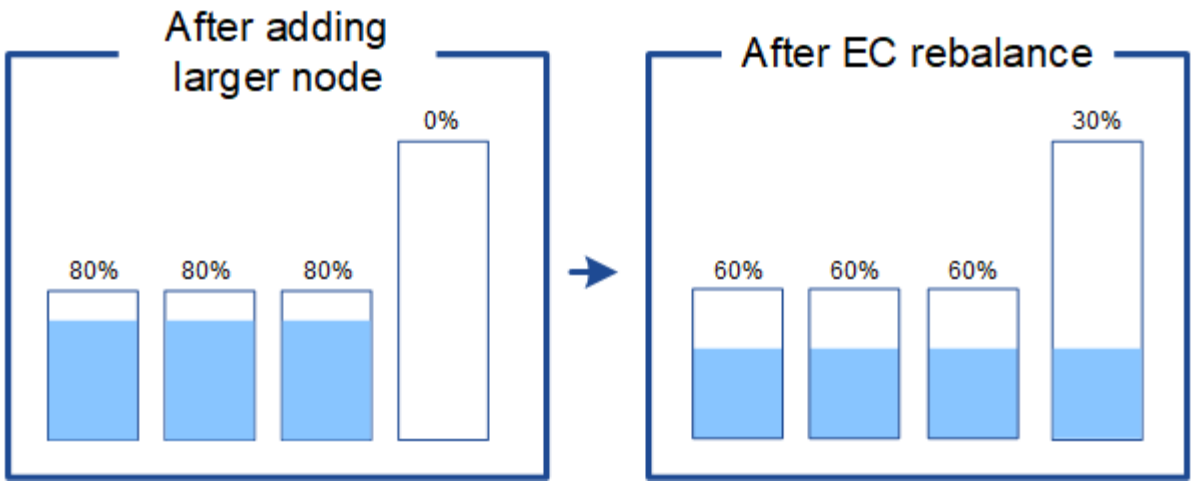

When to rebalance erasure-coded data

Consider the following scenario:

- StorageGRID is running at a single site, which contains three Storage Nodes.
- The ILM policy uses a 2+1 erasure-coding rule for all objects larger than 1.0 MB and a 2-copy replication rule for smaller objects.
- All Storage Nodes have become completely full. The **Low Object Storage** alert has been triggered at the major severity level.

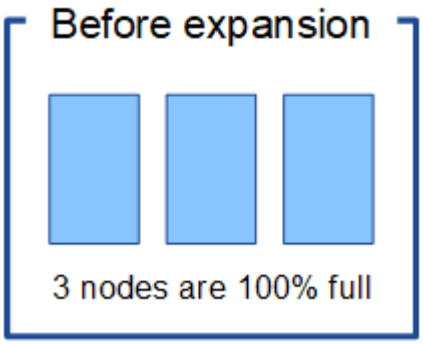

Rebalance is not required if you add enough nodes

To understand when EC rebalance is not required, suppose you added three (or more) new Storage Nodes. In this case, you don't need to perform EC rebalance. The original Storage Nodes will remain full, but new objects will now use the three new nodes for 2+1 erasure coding—the two data fragments and the one parity fragment can each be stored on a different node.

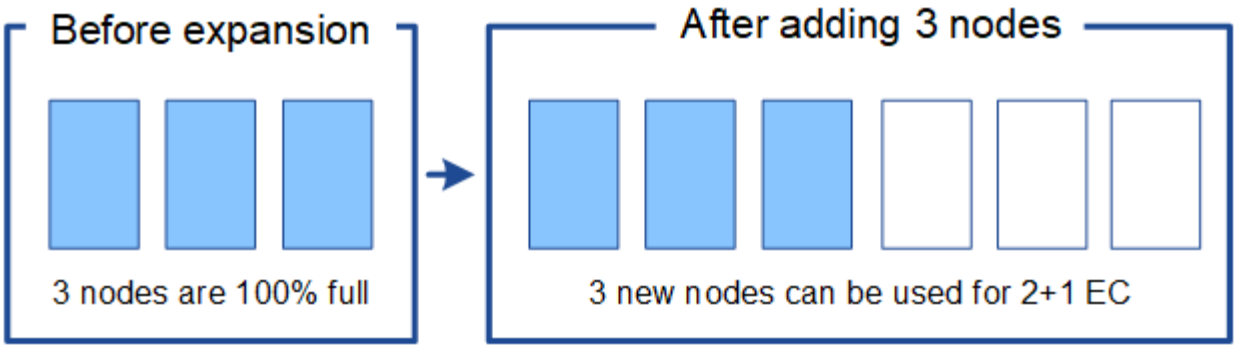

While you can run the EC rebalance procedure in this case, moving the existing erasure-coded data will temporarily decrease the grid's performance, which might impact client operations.

Rebalance is required if you can't add enough nodes

To understand when EC rebalance is required, suppose you can only add two Storage Nodes, instead of three. Because the 2+1 scheme requires at least three Storage Nodes to have space available, the empty nodes can't be used for new erasure-coded data.

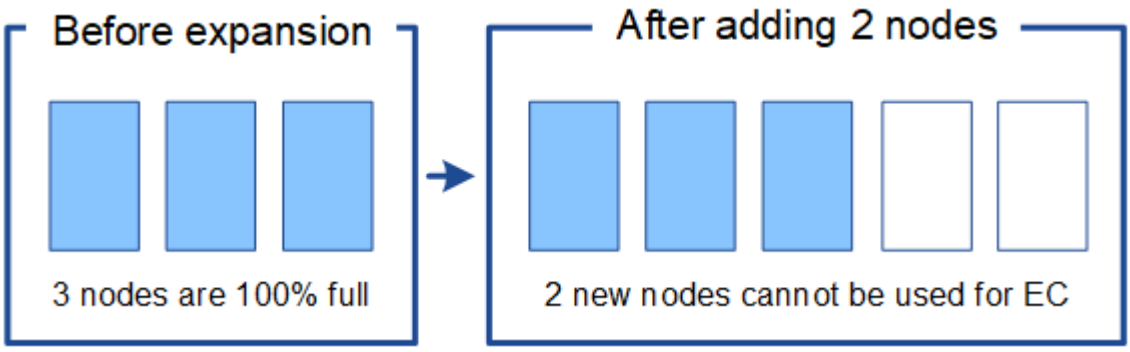

To make use of the new Storage Nodes, you should run the EC rebalance procedure. When this procedure runs, StorageGRID redistributes existing erasure-coded data and parity fragments among all Storage Nodes at the site. In this example, when the EC rebalance procedure is complete, all five nodes are now only 60% full, and objects can continue to be ingested into the 2+1 erasure coding scheme on all Storage Nodes.

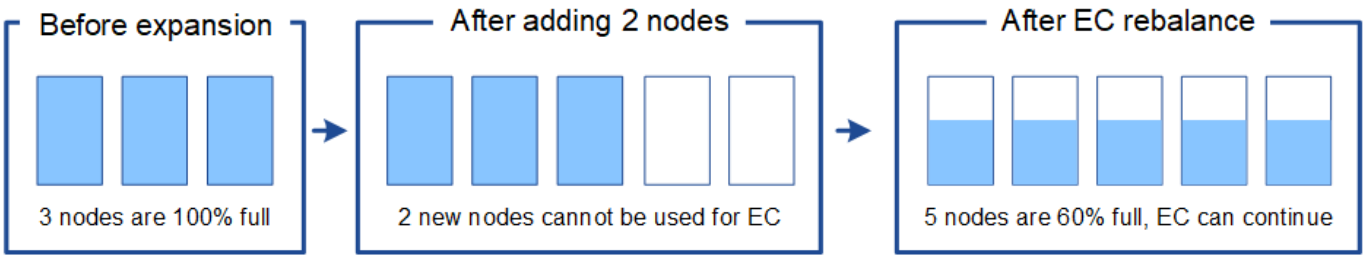

Recommendations for EC rebalancing

NetApp requires EC rebalancing if *all* of the following statements are true:

- You use erasure coding for your object data.
- The **Low Object Storage** alert has been triggered for one or more Storage Nodes at a site, indicating that the nodes are 80% or more full.

- You are unable to add enough new Storage Nodes for the erasure-coding scheme in use. See Add storage capacity for erasure-coded objects.
- Your S3 and Swift clients can tolerate lower performance for their write and read operations while the EC rebalance procedure is running.

You can optionally run the EC rebalance procedure if you prefer Storage Nodes to be filled to similar levels and your S3 and Swift clients can tolerate lower performance for their write and read operations while the EC rebalance procedure is running.

How EC rebalance procedure interacts with other maintenance tasks

You can't perform certain maintenance procedures at the same time you are running the EC rebalance procedure.

Procedure	Allowed during EC rebalance procedure?
Additional EC rebalance procedures	No. You can only run one EC rebalance procedure at a time.
Decommission procedure EC data repair job	No. • You are prevented from starting a decommission procedure or an EC data repair while the EC rebalance procedure is running. • You are prevented from starting the EC rebalance procedure while a Storage Node decommission procedure or an EC data repair is running.
Expansion procedure	No. If you need to add new Storage Nodes in an expansion, run the EC rebalance procedure after adding all new nodes.
Upgrade procedure	No. If you need to upgrade StorageGRID software, perform the upgrade procedure before or after running the EC rebalance procedure. As required, you can terminate the EC rebalance procedure to perform a software upgrade.
Appliance node clone procedure	No. If you need to clone an appliance Storage Node, run the EC rebalance procedure after adding the new node.
Hotfix procedure	Yes. You can apply a StorageGRID hotfix while the EC rebalance procedure is running.

Procedure	Allowed during EC rebalance procedure?
Other maintenance procedures	No. You must terminate the EC rebalance procedure before running other maintenance procedures.

How EC rebalance procedure interacts with ILM

While the EC rebalance procedure is running, avoid making ILM changes that might change the location of existing erasure-coded objects. For example, don't start using an ILM rule that has a different erasure coding profile. If you need to make such ILM changes, you should terminate the EC rebalance procedure.

Add metadata capacity

To ensure that adequate space is available for object metadata, you might need to perform an expansion procedure to add new Storage Nodes at each site.

StorageGRID reserves space for object metadata on volume 0 of each Storage Node. Three copies of all object metadata are maintained at each site, evenly distributed across all Storage Nodes.

You can use the Grid Manager to monitor the metadata capacity of Storage Nodes and to estimate how quickly metadata capacity is being consumed. In addition, the **Low metadata storage** alert is triggered for a Storage Node when the used metadata space reaches certain thresholds.

Note that a grid's object metadata capacity might be consumed faster than its object storage capacity, depending on how you use the grid. For example, if you typically ingest large numbers of small objects or add large quantities of user metadata or tags to objects, you might need to add Storage Nodes to increase metadata capacity even though sufficient object storage capacity remains.

For more information, see the following:

- Manage object metadata storage
- Monitor object metadata capacity for each Storage Node

Guidelines for increasing metadata capacity

Before adding Storage Nodes to increase metadata capacity, review the following guidelines and limitations:

- Assuming sufficient object storage capacity is available, having more space available for object metadata increases the number of objects you can store in your StorageGRID system.
- You can increase a grid's metadata capacity by adding one or more Storage Nodes to each site.
- The actual space reserved for object metadata on any given Storage Node depends on the Metadata Reserved Space storage option (system-wide setting), the amount of RAM allocated to the node, and the size of the node's volume 0.
- You can't increase metadata capacity by adding storage volumes to existing Storage Nodes, because metadata is stored only on volume 0.
- You can't increase metadata capacity by adding a new site.
- StorageGRID keeps three copies of all object metadata at every site. For this reason, the metadata capacity for your system is limited by the metadata capacity of your smallest site.

- When adding metadata capacity, you should add the same number of Storage Nodes to each site.

See the description of what Metadata Reserved Space is.

How metadata is redistributed when you add Storage Nodes

When you add Storage Nodes in an expansion, StorageGRID redistributes the existing object metadata to the new nodes at each site, which increases the overall metadata capacity of the grid. No user action is required.

The following figure shows how StorageGRID redistributes object metadata when you add Storage Nodes in an expansion. The left side of the figure represents volume 0 of three Storage Nodes before an expansion. Metadata is consuming a relatively large portion of each node's available metadata space, and the **Low metadata storage** alert has been triggered.

The right side of the figure shows how the existing metadata is redistributed after two Storage Nodes are added to the site. The amount of metadata on each node has decreased, the **Low metadata storage** alert is no longer triggered, and the space available for metadata has increased.

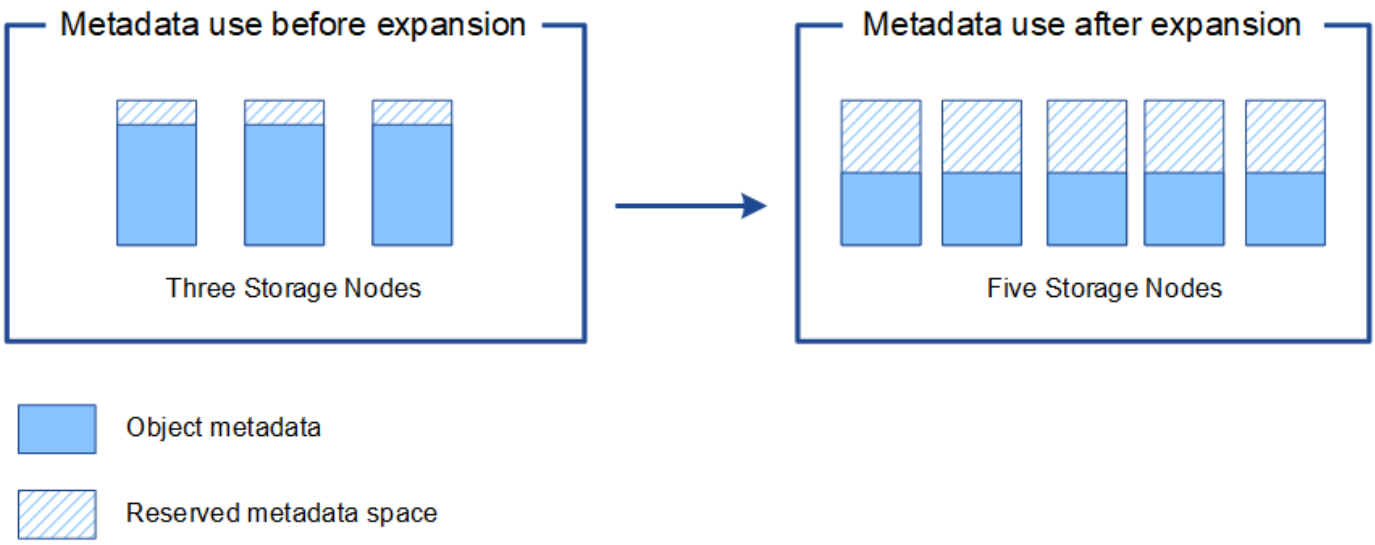

Add grid nodes to add capabilities to your system

You can add redundancy or additional capabilities to a StorageGRID system by adding new grid nodes to existing sites.

For example, you might choose to add Gateway Nodes to use in a high availability (HA) group, or you might add an Admin Node at a remote site to permit monitoring using a local node.

You can add one or more of the following types of nodes to one or more existing sites in a single expansion operation:

- Non-primary Admin Nodes
- Storage Nodes
- Gateway Nodes

When preparing to add grid nodes, be aware of the following limitations:

- The primary Admin Node is deployed during the initial installation. You can't add a primary Admin Node

during an expansion.

- You can add Storage Nodes and other types of nodes in the same expansion.
- When adding Storage Nodes, you must carefully plan the number and location of the new nodes. See Guidelines for adding object capacity.
- If the **Set new node default** option is **Untrusted** on the Untrusted Client Networks tab on the Firewall control page, client applications that connect to expansion nodes using the Client Network must connect using a load balancer endpoint port (**CONFIGURATION** > **Security** > **Firewall control**). See the instructions to change the security setting for the new node and to configure load balancer endpoints.

Add a new site

You can expand your StorageGRID system by adding a new site.

Guidelines for adding a site

Before adding a site, review the following requirements and limitations:

- You can only add one site per expansion operation.
- You can't add grid nodes to an existing site as part of the same expansion.
- All sites must include at least three Storage Nodes.
- Adding a new site does not automatically increase the number of objects you can store. The total object capacity of a grid depends on the amount of available storage, the ILM policy, and the metadata capacity at each site.
- When sizing a new site, you must ensure that it includes enough metadata capacity.

  StorageGRID keeps a copy of all object metadata at every site. When you add a new site, you must ensure that it includes enough metadata capacity for the existing object metadata and enough metadata capacity for growth.

  For more information, see the following:

  - Manage object metadata storage
  - Monitor object metadata capacity for each Storage Node
- You must consider the available network bandwidth between sites, and the level of network latency. Metadata updates are continually replicated between sites even if all objects are stored only at the site where they are ingested.
- Because your StorageGRID system remains operational during the expansion, you must review ILM rules before starting the expansion procedure. You must ensure that object copies aren't stored to the new site until the expansion procedure is complete.

  For example, before you begin the expansion, determine if any rules use the default storage pool (All Storage Nodes). If they do, you must create a new storage pool that contains the existing Storage Nodes and update your ILM rules to use the new storage pool. Otherwise, objects will be copied to the new site as soon as the first node at that site becomes active.

  For more information about changing ILM when adding a new site, see the example for changing an ILM policy.

Gather required materials

Before performing an expansion operation, gather the materials and install and configure any new hardware and networks.

Item	Notes
StorageGRID installation archive	If you are adding new grid nodes or a new site, you must download and extract the StorageGRID installation archive. You must use the same version that is currently running on the grid. For details, see the instructions for downloading and extracting the StorageGRID installation files. **Note:** You don't need to download files if you are adding new storage volumes to existing Storage Nodes or installing a new StorageGRID appliance.
Service laptop	The service laptop has the following: • Network port • SSH client (for example, PuTTY) • Supported web browser
`Passwords.txt` file	Contains the passwords required to access grid nodes on the command line. Included in the Recovery Package.
Provisioning passphrase	The passphrase is created and documented when the StorageGRID system is first installed. The provisioning passphrase is not in the `Passwords.txt` file.

StorageGRID documentation	• Administer StorageGRID
	• Release notes
	• Installation instructions for your platform
	  ◦ Install Red Hat Enterprise Linux or CentOS
	  ◦ Install Ubuntu or Debian
	  ◦ Install VMware

Current documentation for your platform	For supported versions, see the Interoperability Matrix Tool (IMT).

Download and extract the StorageGRID installation files

Before you can add new grid nodes or a new site, you must download the appropriate StorageGRID installation archive and extract the files.

About this task

You must perform expansion operations using the version of StorageGRID that is currently running on the grid.

Steps

1. Go to NetApp Downloads: StorageGRID.
2. Select the version of StorageGRID that is currently running on the grid.
3. Sign in with the username and password for your NetApp account.
4. Read the End User License Agreement, select the checkbox, and then select **Accept & Continue**.
5. In the **Install StorageGRID** column of the download page, select the `.tgz` or `.zip` file for your platform.

   The version shown in the installation archive file must match the version of the software that is currently installed.

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

Platform	Installation archive

Red Hat Enterprise Linux or CentOS	`StorageGRID-Webscale-`*`version`*`-RPM-`*`uniqueID`*`.zip`
	
	`StorageGRID-Webscale-`*`version`*`-RPM-`*`uniqueID`*`.tgz`

Ubuntu or Debian or Appliances	`StorageGRID-Webscale-`*`version`*`-DEB-`*`uniqueID`*`.zip`
	
	`StorageGRID-Webscale-`*`version`*`-DEB-`*`uniqueID`*`.tgz`

VMware	`StorageGRID-Webscale-`*`version`*`-VMware-`*`uniqueID`*`.zip`
	
	`StorageGRID-Webscale-`*`version`*`-VMware-`*`uniqueID`*`.tgz`

OpenStack/other Hypervisor	To expand an existing deployment on OpenStack, you must deploy a virtual machine running one of the supported Linux distributions listed above and follow the appropriate instructions for Linux.

6. Download and extract the archive file.
7. Follow the appropriate step for your platform to choose the files you need, based on your platform, planned grid topology, and how you will expand your StorageGRID system.

   The paths listed in the step for each platform are relative to the top-level directory installed by the archive file.

8. If you are expanding a Red Hat Enterprise Linux or CentOS system, select the appropriate files.

Path and file name	Description
`./rpms/README`	A text file that describes all of the files contained in the StorageGRID download file.
`./rpms/NLF000000.txt`	A free license that does not provide any support entitlement for the product.
`./rpms/StorageGRID-Webscale-Images-`*`version`*`-SHA.rpm`	RPM package for installing the StorageGRID node images on your RHEL or CentOS hosts.

Path and file name	Description
`./rpms/StorageGRID-Webscale-Service-version-SHA.rpm`	RPM package for installing the StorageGRID host service on your RHEL or CentOS hosts.
Deployment scripting tool	**Description**
`./rpms/configure-storagegrid.py`	A Python script used to automate the configuration of a StorageGRID system.
`./rpms/configure-sga.py`	A Python script used to automate the configuration of StorageGRID appliances.
`./rpms/configure-storagegrid.sample.json`	An example configuration file for use with the `configure-storagegrid.py` script.
`./rpms/storagegrid-ssoauth.py`	An example Python script that you can use to sign in to the Grid Management API when single sign-on is enabled. You can also use this script for Ping Federate.
`./rpms/configure-storagegrid.blank.json`	A blank configuration file for use with the `configure-storagegrid.py` script.
`./rpms/extras/ansible`	Example Ansible role and playbook for configuring RHEL or CentOS hosts for StorageGRID container deployment. You can customize the role or playbook as necessary.
`./rpms/storagegrid-ssoauth-azure.py`	An example Python script that you can use to sign in to the Grid Management API when single sign-on (SSO) is enabled using Active Directory or Ping Federate.
`./rpms/storagegrid-ssoauth-azure.js`	A helper script called by the companion `storagegrid-ssoauth-azure.py` Python script to perform SSO interactions with Azure.
`./rpms/extras/api-schemas`	API schemas for StorageGRID. **Note**: Before you perform an upgrade, you can use these schemas to confirm that any code you have written to use StorageGRID management APIs will be compatible with the new StorageGRID release if you don’t have a non-production StorageGRID environment for upgrade compatibility testing.

9. If you are expanding an Ubuntu or Debian system, select the appropriate files.

Path and file name	Description
`./debs/README`	A text file that describes all of the files contained in the StorageGRID download file.
`./debs/NLF000000.txt`	A non-production NetApp License File that you can use for testing and proof of concept deployments.
`./debs/storagegrid-webscale-images-version-SHA.deb`	DEB package for installing the StorageGRID node images on Ubuntu or Debian hosts.
`./debs/storagegrid-webscale-images-version-SHA.deb.md5`	MD5 checksum for the file `/debs/storagegrid-webscale-images-version-SHA.deb`.
`./debs/storagegrid-webscale-service-version-SHA.deb`	DEB package for installing the StorageGRID host service on Ubuntu or Debian hosts.
Deployment scripting tool	**Description**
`./debs/configure-storagegrid.py`	A Python script used to automate the configuration of a StorageGRID system.
`./debs/configure-sga.py`	A Python script used to automate the configuration of StorageGRID appliances.
`./debs/storagegrid-ssoauth.py`	An example Python script that you can use to sign in to the Grid Management API when single sign-on is enabled. You can also use this script for Ping Federate.
`./debs/configure-storagegrid.sample.json`	An example configuration file for use with the `configure-storagegrid.py` script.
`./debs/configure-storagegrid.blank.json`	A blank configuration file for use with the `configure-storagegrid.py` script.
`./debs/extras/ansible`	Example Ansible role and playbook for configuring Ubuntu or Debian hosts for StorageGRID container deployment. You can customize the role or playbook as necessary.
`./debs/storagegrid-ssoauth-azure.py`	An example Python script that you can use to sign in to the Grid Management API when single sign-on (SSO) is enabled using Active Directory or Ping Federate.
`./debs/storagegrid-ssoauth-azure.js`	A helper script called by the companion `storagegrid-ssoauth-azure.py` Python script to perform SSO interactions with Azure.

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

10. If you are expanding a VMware system, select the appropriate files.

Path and file name	Description
`./vsphere/README`	A text file that describes all of the files contained in the StorageGRID download file.
`./vsphere/NLF000000.txt`	A free license that does not provide any support entitlement for the product.
`./vsphere/NetApp-SG-version-SHA.vmdk`	The virtual machine disk file that is used as a template for creating grid node virtual machines.

`./vsphere/vsphere-primary-admin.ovf`
 	
`./vsphere/vsphere-primary-admin.mf`
	The Open Virtualization Format template file (`.ovf`) and manifest file (`.mf`) for deploying the primary Admin Node.

`./vsphere/vsphere-non-primary-admin.ovf`
 	
`./vsphere/vsphere-non-primary-admin.mf`
	The template file (`.ovf`) and manifest file (`.mf`) for deploying non-primary Admin Nodes.

`./vsphere/vsphere-archive.ovf`
 	
`./vsphere/vsphere-archive.mf`
	The template file (`.ovf`) and manifest file (`.mf`) for deploying Archive Nodes.

`./vsphere/vsphere-gateway.ovf`
 	
`./vsphere/vsphere-gateway.mf`
	The template file (`.ovf`) and manifest file (`.mf`) for deploying Gateway Nodes.

`./vsphere/vsphere-storage.ovf`
 	
`./vsphere/vsphere-storage.mf`
	The template file (`.ovf`) and manifest file (`.mf`) for deploying virtual machine-based Storage Nodes.

Deployment scripting tool	**Description**
`./vsphere/deploy-vsphere-ovftool.sh`	A Bash shell script used to automate the deployment of virtual grid nodes.
`./vsphere/deploy-vsphere-ovftool-sample.ini`	An example configuration file for use with the `deploy-vsphere-ovftool.sh` script.

Path and file name	Description
`./vsphere/configure-storagegrid.py`	A Python script used to automate the configuration of a StorageGRID system.
`./vsphere/configure-sga.py`	A Python script used to automate the configuration of StorageGRID appliances.
`./vsphere/storagegrid-ssoauth.py`	An example Python script that you can use to sign in to the Grid Management API when single sign-on (SSO) is enabled. You can also use this script for Ping Federate.
`./vsphere/configure-storagegrid.sample.json`	An example configuration file for use with the `configure-storagegrid.py` script.
`./vsphere/configure-storagegrid.blank.json`	A blank configuration file for use with the `configure-storagegrid.py` script.
`./vsphere/storagegrid-ssoauth-azure.py`	An example Python script that you can use to sign in to the Grid Management API when single sign-on (SSO) is enabled using Active Directory or Ping Federate.
`./vsphere/storagegrid-ssoauth-azure.js`	A helper script called by the companion `storagegrid-ssoauth-azure.py` Python script to perform SSO interactions with Azure.
`./vsphere/extras/api-schemas`	API schemas for StorageGRID. **Note**: Before you perform an upgrade, you can use these schemas to confirm that any code you have written to use StorageGRID management APIs will be compatible with the new StorageGRID release if you don’t have a non-production StorageGRID environment for upgrade compatibility testing.

11. If you are expanding a StorageGRID appliance-based system, select the appropriate files.

Path and file name	Description
`./debs/storagegrid-webscale-images-version-SHA.deb`	DEB package for installing the StorageGRID node images on your appliances.
`./debs/storagegrid-webscale-images-version-SHA.deb.md5`	MD5 checksum for the file `/debs/storagegridwebscale- images-version-SHA.deb`.

For appliance installation, these files are only required if you need to avoid network traffic. The appliance can download the required files from the primary Admin Node.

Verify hardware and networking

Before beginning the expansion of your StorageGRID system, ensure the following:

- The hardware needed to support the new grid nodes or new site has been installed and configured.
- All new nodes have bidirectional communication paths to all existing and new nodes (a requirement for the Grid Network). In particular, confirm that the following TCP ports are open between the new nodes you are adding in the expansion and the primary Admin Node:
  - 1055
  - 7443
  - 8011
  - 10342

  See Internal grid node communications.
- The primary Admin Node can communicate with all expansion servers that are intended to host the StorageGRID system.
- If any of the new nodes has a Grid Network IP address on a subnet not previously used, you have already added the new subnet to the Grid Network subnet list. Otherwise, you will have to cancel the expansion, add the new subnet, and start the procedure again.
- You aren't using 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. Using NAT to bridge the Grid Network across a public network segment is supported only if you use a tunneling application that is transparent to all nodes in the grid, meaning the grid nodes require no knowledge of public IP addresses.

  This NAT restriction is specific to grid nodes and the Grid Network. As required, you can use NAT between external clients and grid nodes, such as to provide a public IP address for a Gateway Node.

Add storage volumes

Add storage volumes to Storage Nodes

You can expand the storage capacity of Storage Nodes that have 16 or fewer storage volumes by adding additional storage volumes. You might need to add storage volumes to more than one Storage Node to satisfy ILM requirements for replicated or erasure-coded copies.

Before you begin

Before adding storage volumes, review the guidelines for adding object capacity to ensure that you know where to add volumes to meet the requirements of your ILM policy.

These instructions apply to software-based Storage Nodes only. See Add expansion shelf to deployed SG6060 to learn how to add storage volumes to SG6060 by installing expansion shelves. Other appliance Storage Nodes can't be expanded.

About this task

The underlying storage of a Storage Node is divided into storage volumes. Storage volumes are block-based storage devices that are formatted by the StorageGRID system and mounted to store objects. Each Storage Node can support up to 16 storage volumes, which are called *object stores* in the Grid Manager.

Object metadata is always stored in object store 0.

Each object store is mounted on a volume that corresponds to its ID. For example, the object store with an ID of 0000 corresponds to the `/var/local/rangedb/0` mount point.

Before adding new storage volumes, use the Grid Manager to view the current object stores for each Storage Node as well as the corresponding mount points. You can use this information when adding storage volumes.

Steps

1. Select **NODES** > ***site*** > ***Storage Node*** > **Storage**.
2. Scroll down to view the amounts of available storage for each volume and object store.

   For appliance Storage Nodes, the Worldwide Name for each disk matches the volume world-wide identifier (WWID) that appears when you view standard volume properties in SANtricity OS (the management software connected to the appliance's storage controller).

   To help you interpret disk read and write statistics related to volume mount points, the first portion of the name shown in the **Name** column of the Disk Devices table (that is, *sdc*, *sdd*, *sde*, and so on) matches the value shown in the **Device** column of the Volumes table.

Disk devices

Name	World Wide Name	I/O load	Read rate	Write rate
sdc(8:16,sdb)	N/A	0.05%	0 bytes/s	4 KB/s
sde(8:48,sdd)	N/A	0.00%	0 bytes/s	82 bytes/s
sdf(8:64,sde)	N/A	0.00%	0 bytes/s	82 bytes/s
sdg(8:80,sdf)	N/A	0.00%	0 bytes/s	82 bytes/s
sdd(8:32,sdc)	N/A	0.00%	0 bytes/s	82 bytes/s
croot(8:1,sda1)	N/A	0.04%	0 bytes/s	4 KB/s
cvloc(8:2,sda2)	N/A	0.95%	0 bytes/s	52 KB/s

Volumes

Mount point	Device	Status	Size	Available	Write cache status
/	croot	Online	21.00 GB	14.73 GB	Unknown
/var/local	cvloc	Online	85.86 GB	80.94 GB	Unknown
/var/local/rangedb/0	sdc	Online	107.32 GB	107.17 GB	Enabled
/var/local/rangedb/1	sdd	Online	107.32 GB	107.18 GB	Enabled
/var/local/rangedb/2	sde	Online	107.32 GB	107.18 GB	Enabled
/var/local/rangedb/3	sdf	Online	107.32 GB	107.18 GB	Enabled
/var/local/rangedb/4	sdg	Online	107.32 GB	107.18 GB	Enabled

Object stores

ID	Size	Available	Replicated data	EC data	Object data (%)	Health
0000	107.32 GB	96.44 GB	1.55 MB	0 bytes	0.00%	No Errors
0001	107.32 GB	107.18 GB	0 bytes	0 bytes	0.00%	No Errors
0002	107.32 GB	107.18 GB	0 bytes	0 bytes	0.00%	No Errors
0003	107.32 GB	107.18 GB	0 bytes	0 bytes	0.00%	No Errors
0004	107.32 GB	107.18 GB	0 bytes	0 bytes	0.00%	No Errors

3. Follow the instructions for your platform to add new storage volumes to the Storage Node.
   - VMware: Add storage volumes to Storage Node
   - Linux: Add direct-attached or SAN volumes to Storage Node

VMware: Add storage volumes to Storage Node

If a Storage Node includes fewer than 16 storage volumes, you can increase its capacity by using VMware vSphere to add volumes.

Before you begin

- You have access to the instructions for installing StorageGRID for VMware deployments.
  - Install VMware
- You have the `Passwords.txt` file.
- You have specific access permissions.

> Don't attempt to add storage volumes to a Storage Node while a software upgrade, recovery procedure, or another expansion procedure is active.

About this task

The Storage Node is unavailable for a brief time when you add storage volumes. You should perform this procedure on one Storage Node at a time to avoid impacting client-facing grid services.

Steps

1. If necessary, install new storage hardware and create new VMware datastores.
2. Add one or more hard disks to the virtual machine for use as storage (object stores).
   a. Open VMware vSphere Client.
   b. Edit the virtual machine settings to add one or more additional hard disks.

   The hard disks are typically configured as Virtual Machine Disks (VMDKs). VMDKs are more commonly used and are easier to manage, while RDMs might provide better performance for workloads that use larger object sizes (for example, greater than 100 MB). For more information about adding hard disks to virtual machines, see the VMware vSphere documentation.
3. Restart the virtual machine by using the **Restart Guest OS** option in the VMware vSphere Client, or by entering the following command in an ssh session to the virtual machine:`sudo reboot`

   > Don't use **Power Off** or **Reset** to restart the virtual machine.
4. Configure the new storage for use by the Storage Node:
   a. Log in to the grid node:
      i. Enter the following command: `ssh admin@grid_node_IP`
      ii. Enter the password listed in the `Passwords.txt` file.
      iii. Enter the following command to switch to root: `su -`
      iv. Enter the password listed in the `Passwords.txt` file. When you are logged in as root, the prompt changes from `$` to `#`.

b. Configure the new storage volumes:

```
sudo add_rangedbs.rb
```

This script finds any new storage volumes and prompts you to format them.

c. Enter **y** to accept the formatting.

d. If any of the volumes have previously been formatted, decide if you want to reformat them.

- Enter **y** to reformat.
- Enter **n** to skip reformatting.

The `setup_rangedbs.sh` script runs automatically.

5. Check that the services start correctly:

a. View a list of the status of all services on the server:

```
sudo storagegrid-status
```

The status is updated automatically.

b. Wait until all services are Running or Verified.

c. Exit the status screen:

```
Ctrl+C
```

6. Verify that the Storage Node is online:

a. Sign in to the Grid Manager using a supported web browser.

b. Select **SUPPORT** > **Tools** > **Grid topology**.

c. Select ***site*** > ***Storage Node*** > **LDR** > **Storage**.

d. Select the **Configuration** tab and then the **Main** tab.

e. If the **Storage State - Desired** drop-down list is set to Read-only or Offline, select **Online**.

f. Select **Apply Changes**.

7. To see the new object stores:

a. Select **NODES** > ***site*** > ***Storage Node*** > **Storage**.

b. View the details in the **Object Stores** table.

Result

You can use the expanded capacity of the Storage Nodes to save object data.

Linux: Add direct-attached or SAN volumes to Storage Node

If a Storage Node includes fewer than 16 storage volumes, you can increase its capacity by adding new block storage devices, making them visible to the Linux hosts, and adding the new block device mappings to the StorageGRID configuration file used for the Storage Node.

Before you begin

- You have access to the instructions for installing StorageGRID for your Linux platform.
  - Install Red Hat Enterprise Linux or CentOS
  - Install Ubuntu or Debian
- You have the `Passwords.txt` file.
- You have specific access permissions.

> Don't attempt to add storage volumes to a Storage Node while a software upgrade, recovery procedure, or another expansion procedure is active.

About this task

The Storage Node is unavailable for a brief time when you add storage volumes. You should perform this procedure on one Storage Node at a time to avoid impacting client-facing grid services.

Steps

1. Install the new storage hardware.

   For more information, see the documentation provided by your hardware vendor.

2. Create new block storage volumes of the desired sizes.
   - Attach the new drives and update the RAID controller configuration as needed, or allocate the new SAN LUNs on the shared storage arrays and allow the Linux host to access them.
   - Use the same persistent naming scheme you used for the storage volumes on the existing Storage Node.
   - If you use the StorageGRID node migration feature, make the new volumes visible to other Linux hosts that are migration targets for this Storage Node. For more information, see the instructions for installing StorageGRID for your Linux platform.
3. Log in to the Linux host supporting the Storage Node as root or with an account that has sudo permission.
4. Confirm that the new storage volumes are visible on the Linux host.

   You might have to rescan for devices.

5. Run the following command to temporarily disable the Storage Node:

   ```
   sudo storagegrid node stop <node-name>
   ```

6. Using a text editor such as vim or pico, edit the node configuration file for the Storage Node, which can be found at `/etc/storagegrid/nodes/<node-name>.conf`.
7. Locate the section of the node configuration file that contains the existing object storage block device mappings.

   In the example, `BLOCK_DEVICE_RANGEDB_00` to `BLOCK_DEVICE_RANGEDB_03` are the existing object storage block device mappings.

```
NODE_TYPE = VM_Storage_Node
ADMIN_IP = 10.1.0.2
BLOCK_DEVICE_VAR_LOCAL = /dev/mapper/sgws-sn1-var-local
BLOCK_DEVICE_RANGEDB_00 = /dev/mapper/sgws-sn1-rangedb-0
BLOCK_DEVICE_RANGEDB_01 = /dev/mapper/sgws-sn1-rangedb-1
BLOCK_DEVICE_RANGEDB_02 = /dev/mapper/sgws-sn1-rangedb-2
BLOCK_DEVICE_RANGEDB_03 = /dev/mapper/sgws-sn1-rangedb-3
GRID_NETWORK_TARGET = bond0.1001
ADMIN_NETWORK_TARGET = bond0.1002
CLIENT_NETWORK_TARGET = bond0.1003
GRID_NETWORK_IP = 10.1.0.3
GRID_NETWORK_MASK = 255.255.255.0
GRID_NETWORK_GATEWAY = 10.1.0.1
```

8. Add new object storage block device mappings corresponding to the block storage volumes you added for this Storage Node.

   Make sure to start at the next `BLOCK_DEVICE_RANGEDB_nn`. Don't leave a gap.

   - Based on the example above, start at `BLOCK_DEVICE_RANGEDB_04`.
   - In the example below, four new block storage volumes have been added to the node: `BLOCK_DEVICE_RANGEDB_04` to `BLOCK_DEVICE_RANGEDB_07`.

```
NODE_TYPE = VM_Storage_Node
ADMIN_IP = 10.1.0.2
BLOCK_DEVICE_VAR_LOCAL = /dev/mapper/sgws-sn1-var-local
BLOCK_DEVICE_RANGEDB_00 = /dev/mapper/sgws-sn1-rangedb-0
BLOCK_DEVICE_RANGEDB_01 = /dev/mapper/sgws-sn1-rangedb-1
BLOCK_DEVICE_RANGEDB_02 = /dev/mapper/sgws-sn1-rangedb-2
BLOCK_DEVICE_RANGEDB_03 = /dev/mapper/sgws-sn1-rangedb-3
BLOCK_DEVICE_RANGEDB_04 = /dev/mapper/sgws-sn1-rangedb-4
BLOCK_DEVICE_RANGEDB_05 = /dev/mapper/sgws-sn1-rangedb-5
BLOCK_DEVICE_RANGEDB_06 = /dev/mapper/sgws-sn1-rangedb-6
BLOCK_DEVICE_RANGEDB_07 = /dev/mapper/sgws-sn1-rangedb-7
GRID_NETWORK_TARGET = bond0.1001
ADMIN_NETWORK_TARGET = bond0.1002
CLIENT_NETWORK_TARGET = bond0.1003
GRID_NETWORK_IP = 10.1.0.3
GRID_NETWORK_MASK = 255.255.255.0
GRID_NETWORK_GATEWAY = 10.1.0.1
```

9. Run the following command to validate your changes to the node configuration file for the Storage Node:

   ```
   sudo storagegrid node validate <node-name>
   ```

   Address any errors or warnings before proceeding to the next step.

If you observe an error similar to the following, it means that the node configuration file is attempting to map the block device used by `<node-name>` for `<PURPOSE>` to the given `<path-name>` in the Linux file system, but there is not a valid block device special file (or softlink to a block device special file) at that location.

```
Checking configuration file for node <node-name>…
ERROR: BLOCK_DEVICE_<PURPOSE> = <path-name>
<path-name> is not a valid block device
```

Verify that you entered the correct `<path-name>`.

10. Run the following command to restart the node with the new block device mappings in place:

    ```
    sudo storagegrid node start <node-name>
    ```

11. Log in to the Storage Node as admin using the password listed in the `Passwords.txt` file.
12. Check that the services start correctly:
    a. View a list of the status of all services on the server:
    `sudo storagegrid-status`

    The status is updated automatically.

    b. Wait until all services are Running or Verified.
    c. Exit the status screen:

    `Ctrl+C`

13. Configure the new storage for use by the Storage Node:
    a. Configure the new storage volumes:

    `sudo add_rangedbs.rb`

    This script finds any new storage volumes and prompts you to format them.

    b. Enter **y** to format the storage volumes.
    c. If any of the volumes have previously been formatted, decide if you want to reformat them.
    - Enter **y** to reformat.
    - Enter **n** to skip reformatting.

    The `setup_rangedbs.sh` script runs automatically.

14. Verify that the Storage Node is online:
    a. Sign in to the Grid Manager using a supported web browser.
    b. Select **SUPPORT** > **Tools** > **Grid topology**.
    c. Select ***site*** > ***Storage Node*** > **LDR** > **Storage**.
    d. Select the **Configuration** tab and then the **Main** tab.

e. If the **Storage State - Desired** drop-down list is set to Read-only or Offline, select **Online**.

f. Click **Apply Changes**.

15. To see the new object stores:

a. Select **NODES** > ***site*** > ***Storage Node*** > **Storage**.

b. View the details in the **Object Stores** table.

Result

You can now use the expanded capacity of the Storage Nodes to save object data.

Add grid nodes or site

Add grid nodes to existing site or add new site

You can follow this procedure to add grid nodes to existing sites or to add a new site, but you can't perform both types of expansion at the same time.

Before you begin

- You have the Root access or Maintenance permission.
- All existing nodes in the grid are up and running across all sites.
- Any previous expansion, upgrade, decommissioning, or recovery procedures are complete.

> ⓘ You are prevented from starting an expansion while another expansion, upgrade, recovery, or active decommission procedure is in progress. However, if necessary, you can pause a decommission procedure to start an expansion.

Steps

1. Update subnets for Grid Network.
2. Deploy new grid nodes.
3. Perform expansion.

Update subnets for Grid Network

When you add grid nodes or a new site in an expansion, you might need to update or add subnets to the Grid Network.

StorageGRID maintains a list of the network subnets used to communicate between grid nodes on the Grid Network (eth0). These entries include the subnets used for the Grid Network by each site in your StorageGRID system as well as any subnets used for NTP, DNS, LDAP, or other external servers accessed through the Grid Network gateway.

Before you begin

- You are signed in to the Grid Manager using a supported web browser.
- You have the Maintenance or Root access permission.
- You have the provisioning passphrase.
- You have the network addresses, in CIDR notation, of the subnets you want to configure.

About this task

If any of the new nodes has a Grid Network IP address on a subnet not previously used, you must add the new subnet to the Grid Network subnet list before starting the expansion. Otherwise, you will have to cancel the expansion, add the new subnet, and start the procedure again.

Steps

1. Select **MAINTENANCE** > **Network** > **Grid Network**.
2. Select **Add another subnet** to add a new subnet in CIDR notation.

   For example, enter `10.96.104.0/22`.

3. Enter the provisioning passphrase, and select **Save**.
4. Wait until the changes are applied, then download a new Recovery Package.
   a. Select **MAINTENANCE** > **System** > **Recovery package**.
   b. Enter the **Provisioning Passphrase**.

> The Recovery Package file must be secured because it contains encryption keys and passwords that can be used to obtain data from the StorageGRID system. It is also used to recover the primary Admin Node.

The subnets you have specified are configured automatically for your StorageGRID system.

Deploy new grid nodes

The steps for deploying new grid nodes in an expansion are the same as the steps used when the grid was first installed. You must deploy all new grid nodes before you can perform the expansion.

When you expand the grid, the nodes you add don't have to match the existing node types. You can add VMware nodes, Linux container-based nodes, or appliance nodes.

VMware: Deploy grid nodes

You must deploy a virtual machine in VMware vSphere for each VMware node you want to add in the expansion.

Steps

1. Deploy the new node as virtual machine and connect it to one or more StorageGRID networks.

   When you deploy the node, you can optionally remap node ports or increase CPU or memory settings.

2. After you have deployed all new VMware nodes, perform the expansion procedure.

Linux: Deploy grid nodes

You can deploy grid nodes on new Linux hosts or on existing Linux hosts. If you need additional Linux hosts to support the CPU, RAM, and storage requirements of the StorageGRID nodes you want to add to your grid, you prepare them in the same way you prepared the hosts when you first installed them. Then, you deploy the expansion nodes in the same way you deployed grid nodes during installation.

Before you begin

- You have the instructions for installing StorageGRID for your version of Linux, and you have reviewed the hardware and storage requirements.
  - Install Red Hat Enterprise Linux or CentOS
  - Install Ubuntu or Debian
- If you plan to deploy new grid nodes on existing hosts, you have confirmed the existing hosts have enough CPU, RAM, and storage capacity for the additional nodes.
- You have a plan to minimize failure domains. For example, you should not deploy all Gateway Nodes on a single physical host.

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

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

Steps

1. If you are adding new hosts, access the installation instructions for deploying StorageGRID nodes.
2. To deploy the new hosts, follow the instructions for preparing the hosts.
3. To create node configuration files and to validate the StorageGRID configuration, follow the instructions for deploying grid nodes.
4. If you are adding nodes to a new Linux host, start the StorageGRID host service.
5. If you are adding nodes to an existing Linux host, start the new nodes using the storagegrid host service CLI:`sudo storagegrid node start [<node name\>]`

After you finish

After deploying all new grid nodes, you can perform the expansion.

Appliances: Deploying Storage, Gateway, or non-primary Admin Nodes

To install the StorageGRID software on an appliance node, you use the StorageGRID Appliance Installer, which is included on the appliance. In an expansion, each storage appliance functions as a single Storage Node, and each services appliance functions as a single Gateway Node or non-primary Admin Node. Any appliance can connect to the Grid Network, the Admin Network, and the Client Network.

Before you begin

- The appliance has been installed in a rack or cabinet, connected to your networks, and powered on.
- You have completed the Set up hardware steps.

  Setting up appliance hardware includes the required steps for configuring StorageGRID connections (network links and IP addresses) as well the optional steps for enabling node encryption, changing the RAID mode, and remapping network ports.

- All Grid Network subnets listed on the IP Configuration page of the StorageGRID Appliance Installer have been defined in the Grid Network Subnet List on the primary Admin Node.
- The StorageGRID Appliance Installer firmware on the replacement appliance is compatible with the StorageGRID software version currently running on your grid. If the versions aren't compatible, you must upgrade the StorageGRID Appliance Installer firmware.

- You have a service laptop with a supported web browser.
- You know one of the IP addresses assigned to the appliance's compute controller. You can use the IP address for any attached StorageGRID network.

About this task

The process of installing StorageGRID on an appliance node has the following phases:

- You specify or confirm the IP address of the primary Admin Node and the name of the appliance node.
- You start the installation and wait as volumes are configured and the software is installed.

  Partway through appliance installation tasks, the installation pauses. To resume the installation, you sign into the Grid Manager, approve all grid nodes, and complete the StorageGRID installation process.

> If you need to deploy multiple appliance nodes at one time, you can automate the installation process by using the `configure-sga.py` Appliance Installation Script.

Steps

1. Open a browser, and enter one of the IP addresses for the appliance's compute controller.

   `https://Controller_IP:8443`

   The StorageGRID Appliance Installer Home page appears.

2. In the **Primary Admin Node** connection section, determine whether you need to specify the IP address for the primary Admin Node.

   If you have previously installed other nodes in this data center, the StorageGRID Appliance Installer can discover this IP address automatically, assuming the primary Admin Node, or at least one other grid node with ADMIN_IP configured, is present on the same subnet.

3. If this IP address is not shown or you need to change it, specify the address:

Option	Description
Manual IP entry	a. Clear the **Enable Admin Node discovery** checkbox. b. Enter the IP address manually. c. Click **Save**. d. Wait for the connection state for the new IP address to become ready.
Automatic discovery of all connected primary Admin Nodes	a. Select the **Enable Admin Node discovery** checkbox. b. Wait for the list of discovered IP addresses to be displayed. c. Select the primary Admin Node for the grid where this appliance Storage Node will be deployed. d. Click **Save**. e. Wait for the connection state for the new IP address to become ready.

4. In the **Node name** field, enter the name you want to use for this appliance node, and select **Save**.

   The node name is assigned to this appliance node in the StorageGRID system. It is shown on the Nodes page (Overview tab) in the Grid Manager. If required, you can change the name when you approve the node.

5. In the **Installation** section, confirm that the current state is "Ready to start installation of *node name* into grid with primary Admin Node *admin_ip*" and that the **Start Installation** button is enabled.

   If the **Start Installation** button is not enabled, you might need to change the network configuration or port settings. For instructions, see the maintenance instructions for your appliance.

6. From the StorageGRID Appliance Installer home page, select **Start Installation**.

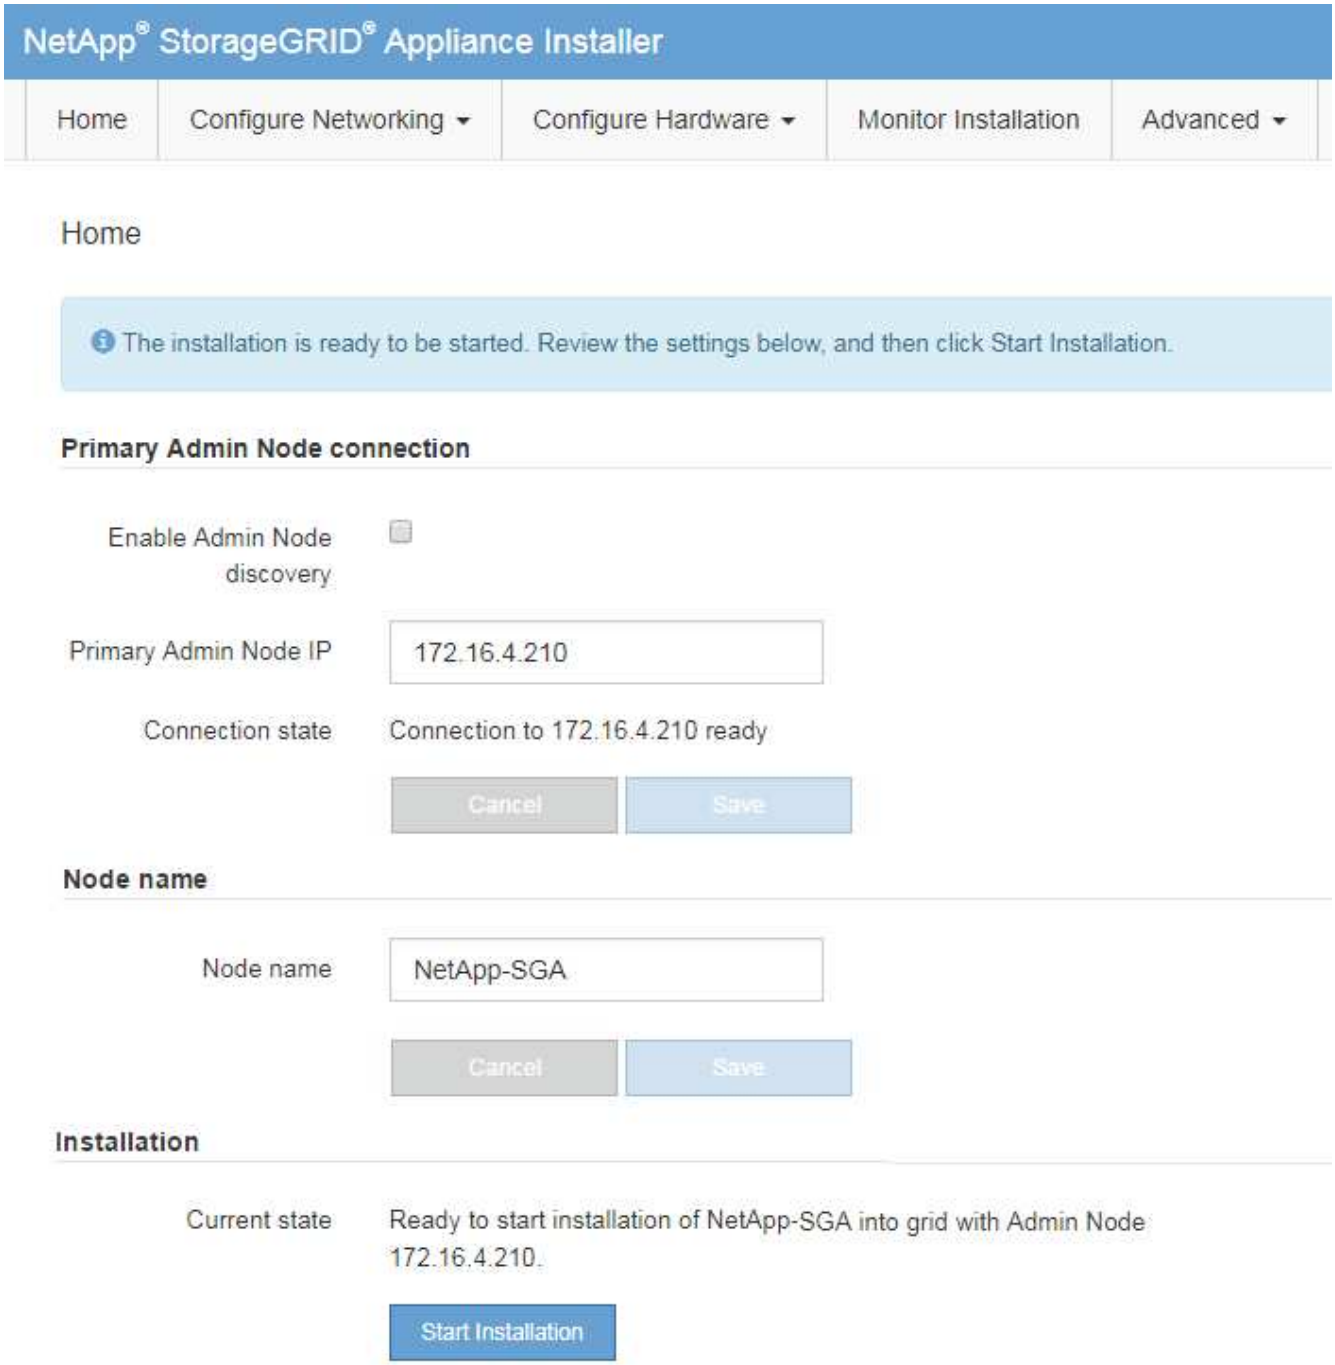

The Current state changes to “Installation is in progress,” and the Monitor Installation page is displayed.

7. If your expansion includes multiple appliance nodes, repeat the previous steps for each appliance.

   If you need to deploy multiple appliance Storage Nodes at one time, you can automate the installation process by using the configure-sga.py appliance installation script.

8. If you need to manually access the Monitor Installation page, select **Monitor Installation** from the menu bar.

   The Monitor Installation page shows the installation progress.

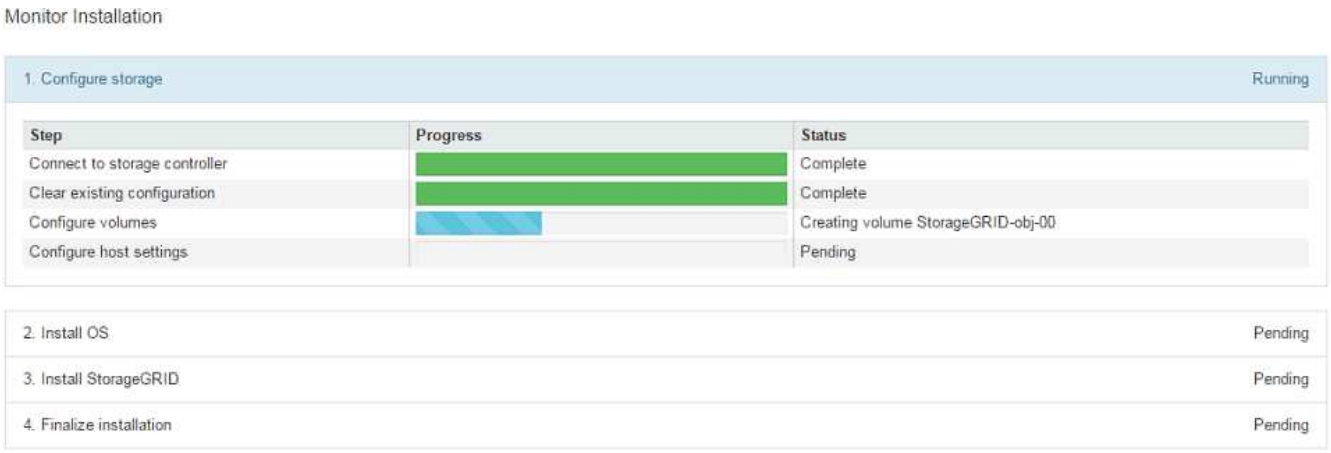

   The blue status bar indicates which task is currently in progress. Green status bars indicate tasks that have completed successfully.

   The installer ensures that tasks completed in a previous install aren’t re-run. If you are re-running an installation, any tasks that don’t need to be re-run are shown with a green status bar and a status of “Skipped.”

9. Review the progress of first two installation stages.

   1. Configure appliance

   During this stage, one of the following processes occurs:

   - For a storage appliance, the installer connects to the storage controller, clears any existing configuration, communicates with SANtricity OS to configure volumes, and configures host settings.
   - For a services appliance, the installer clears any existing configuration from the drives in the compute controller, and configures host settings.

     2. Install OS

     During this stage, the installer copies the base operating system image for StorageGRID to the appliance.

10. Continue monitoring the installation progress until a message appears in the console window, prompting you to use the Grid Manager to approve the node.

Wait until all nodes you added in this expansion are ready for approval before going to the Grid Manager to approve the nodes.

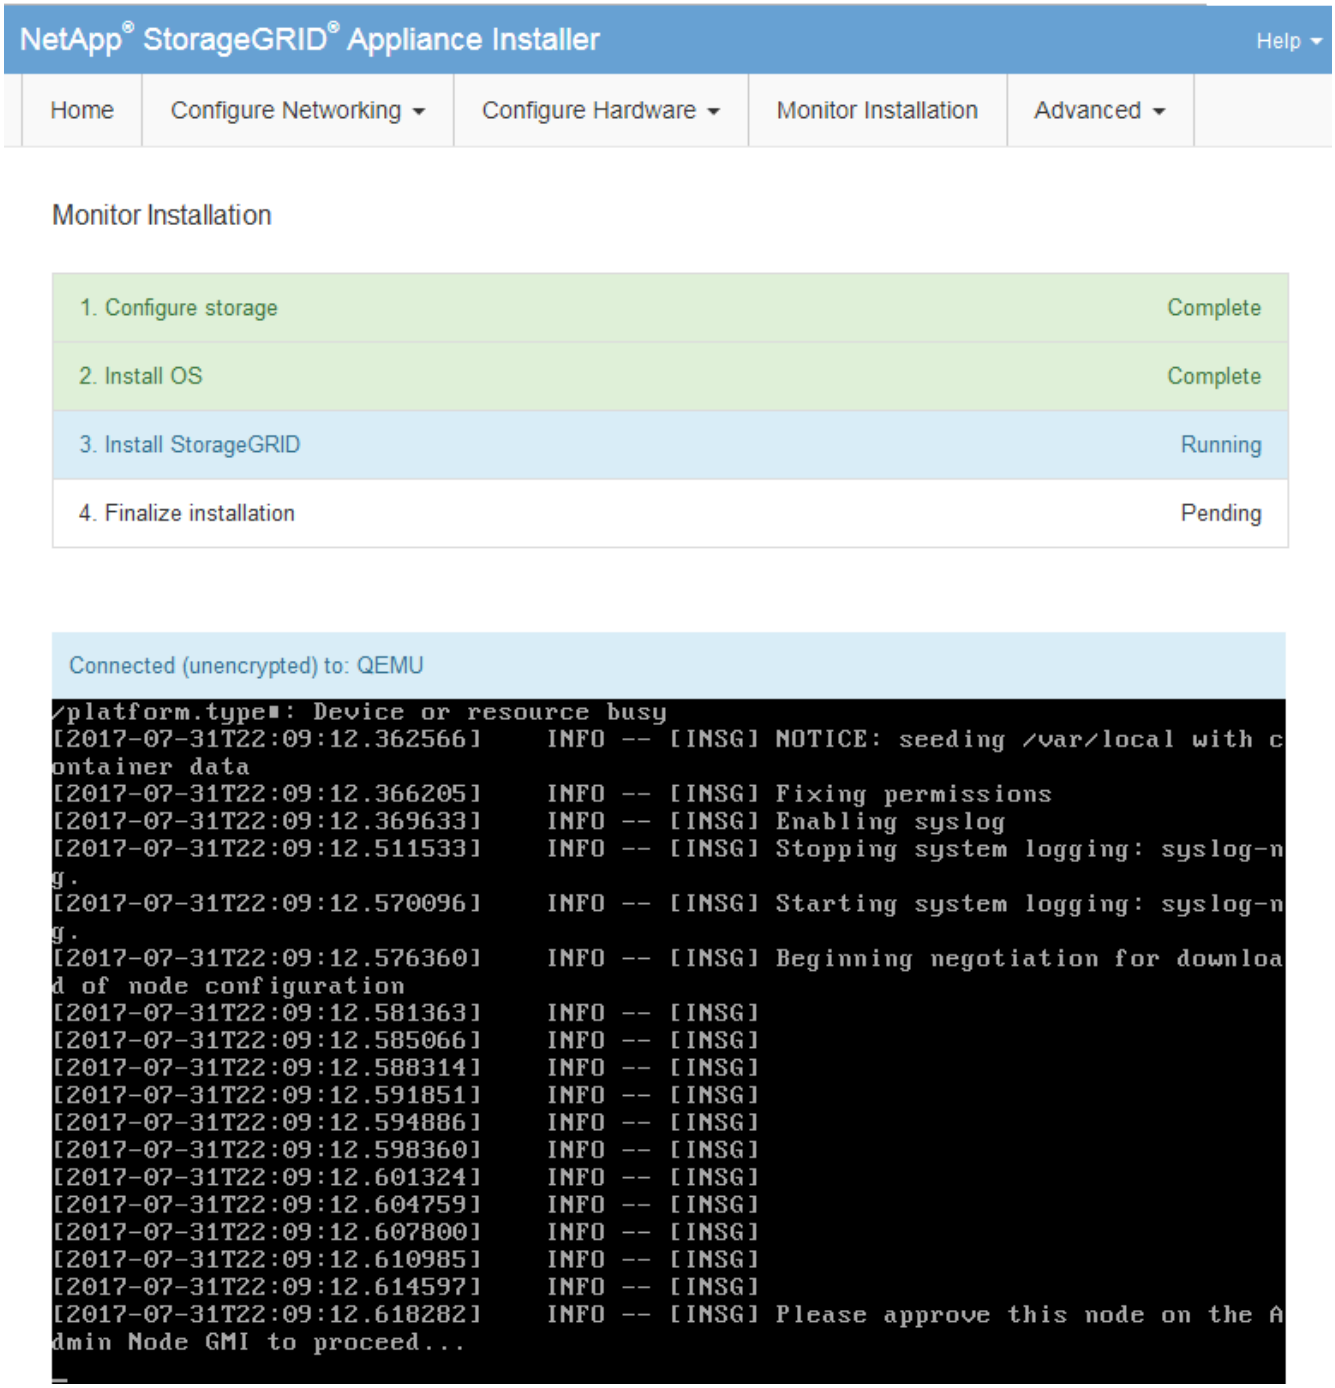

Perform expansion

When you perform the expansion, the new grid nodes are added to your existing StorageGRID deployment.

Before you begin

- You are signed in to the Grid Manager using a supported web browser.
- You have the provisioning passphrase.

- You have deployed all of the grid nodes that are being added in this expansion.
- You have the Maintenance or Root access permission.
- If you are adding Storage Nodes, you have confirmed that all data-repair operations performed as part of a recovery are complete. See Check data repair jobs.
- If you are adding Storage Nodes and you want to assign a custom storage grade to those nodes, you have already created the custom storage grade. You also have either the Root access permission or both the Maintenance and ILM permissions.
- If you are adding a new site, you have reviewed and updated ILM rules. You must ensure that object copies aren't stored to the new site until after the expansion is complete. For example, if a rule uses the default storage pool (**All Storage Nodes**), you must create a new storage pool that contains only the existing Storage Nodes and update ILM rules and the ILM policy to use that new storage pool. Otherwise, objects will be copied to the new site as soon as the first node at that site becomes active.

About this task

Performing the expansion includes these main user tasks:

1. Configure the expansion.
2. Start the expansion.
3. Download a new Recovery Package file.
4. Monitor the expansion steps and stages until all new nodes are installed and configured and all services have started.

> Some expansion steps and stages might take a significant amount of time to run on a large grid. For example, streaming Cassandra to a new Storage Node might take only a few minutes if the Cassandra database is empty. However, if the Cassandra database includes a large amount of object metadata, this stage might take several hours or longer. Don't reboot any Storage Nodes during the either the "Expanding the Cassandra cluster" or "Starting Cassandra and streaming data" stages.

Steps

1. Select **MAINTENANCE** > **Tasks** > **Expansion**.

   The Grid Expansion page appears. The Pending Nodes section lists the nodes that are ready to be added.

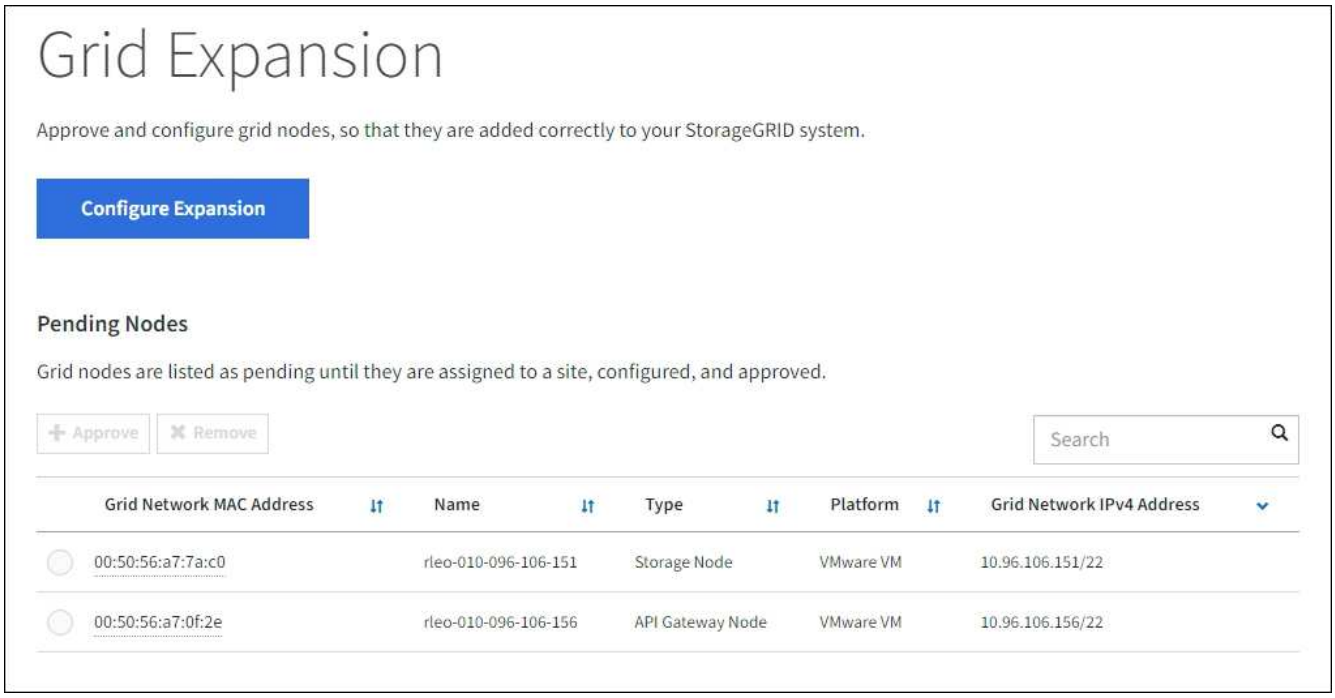

2. Select **Configure Expansion**.

   The Site Selection dialog box appears.

3. Select the type of expansion you are starting:
   - If you are adding a new site, select **New**, and enter the name of the new site.
   - If you are adding one or more nodes to an existing site, select **Existing**.
4. Select **Save**.
5. Review the **Pending Nodes** list, and confirm that it shows all of the grid nodes you deployed.

   As required, you can position your cursor over a node's **Grid Network MAC Address** to see details about that node.

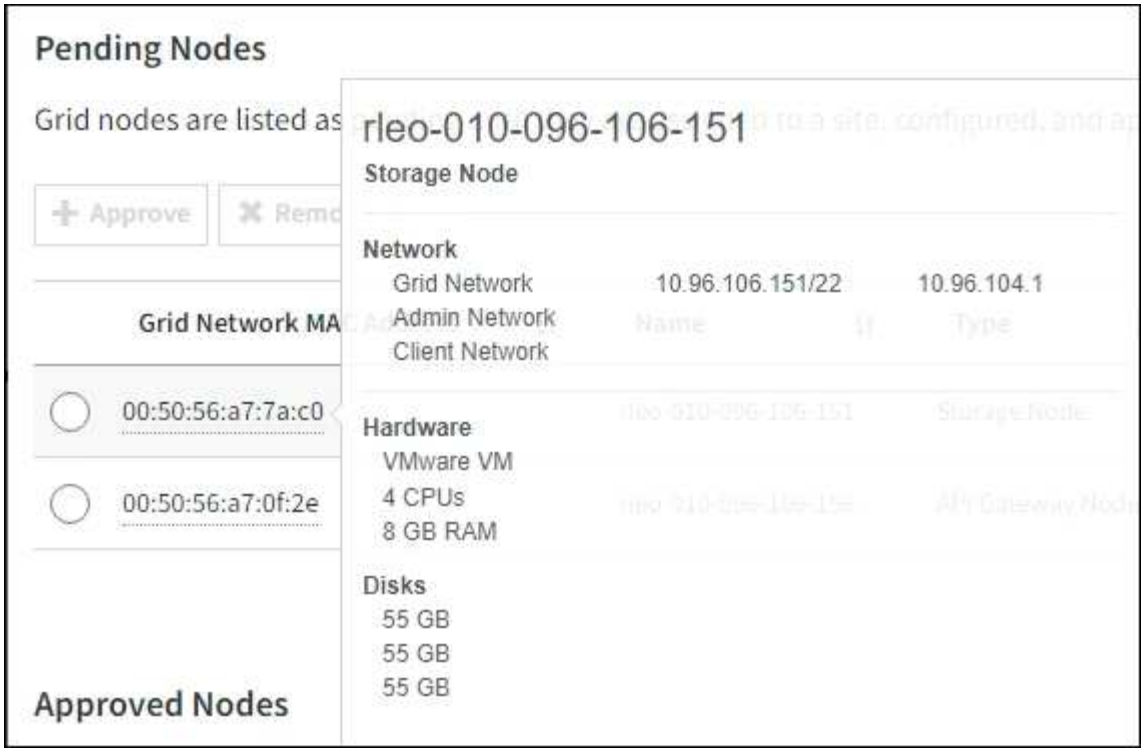

If a node is missing, confirm that it was deployed successfully.

6. From the list of pending nodes, approve the nodes you want to add in this expansion.
   a. Select the radio button next to the first pending grid node you want to approve.
   b. Select **Approve**.

      The grid node configuration form appears.

   c. As required, modify the general settings:

Field	Description
Site	The name of the site the grid node will be associated with. If you are adding multiple nodes, be sure to select the correct site for each node. If you are adding a new site, all nodes are added to the new site.
Name	The system name for the node. System names are required for internal StorageGRID operations and can’t be changed.

Field	Description
NTP Role	The Network Time Protocol (NTP) role of the grid node: • Select **Automatic** (default) to automatically assign the NTP role to the node. The Primary role will be assigned to Admin Nodes, Storage Nodes with ADC services, Gateway Nodes, and any grid nodes that have non-static IP addresses. The Client role will be assigned to all other grid nodes. • Select **Primary** to manually assign the Primary NTP role to the node. At least two nodes at each site should have the Primary role to provide redundant system access to external timing sources. • Select **Client** to manually assign the Client NTP role to the node.
ADC Service (Storage Nodes only)	Whether this Storage Node will run the Administrative Domain Controller (ADC) service. The ADC service keeps track of the location and availability of grid services. At least three Storage Nodes at each site must include the ADC service. You can't add the ADC service to a node after it is deployed. • Select **Yes** if the Storage Node you are replacing includes the ADC service. Because you can't decommission a Storage Node if too few ADC services would remain, this ensures that a new ADC service is available before the old service is removed. • Select **Automatic** to let the system determine whether this node requires the ADC service. Learn about the ADC quorum.
Storage Grade (Storage Nodes only)	Use the **Default** storage grade, or select the custom storage grade you want to assign to this new node. Storage grades are used by ILM storage pools, so your selection can affect which objects will be placed on the Storage Node.

d. As required, modify the settings for the Grid Network, Admin Network, and Client Network.

- **IPv4 Address (CIDR)**: The CIDR network address for the network interface. For example: 172.16.10.100/24

> If you discover that nodes have duplicate IP addresses on the Grid Network while you are approving nodes, you must cancel the expansion, redeploy the virtual machines or appliances with a non-duplicate IP, and restart the expansion.

- **Gateway**: The default gateway of the grid node. For example: 172.16.10.1
- **Subnets (CIDR)**: One or more subnetworks for the Admin Network.

e. Select **Save**.

The approved grid node moves to the Approved Nodes list.

- To modify the properties of an approved grid node, select its radio button, and select **Edit**.
- To move an approved grid node back to the Pending Nodes list, select its radio button, and select **Reset**.
- To permanently remove an approved grid node, power the node off. Then, select its radio button, and select **Remove**.

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

> If possible, you should approve all pending grid notes and perform a single expansion. More time will be required if you perform multiple small expansions.

7. When you have approved all grid nodes, enter the **Provisioning Passphrase**, and select **Expand**.

After a few minutes, this page updates to display the status of the expansion procedure. When tasks that affect individual grid nodes are in progress, the Grid Node Status section lists the current status for each grid node.

> During the “Installing grid nodes” step for a new appliance, the StorageGRID Appliance Installer shows installation moving from Stage 3 to Stage 4, Finalize Installation. When Stage 4 completes, the controller is rebooted.

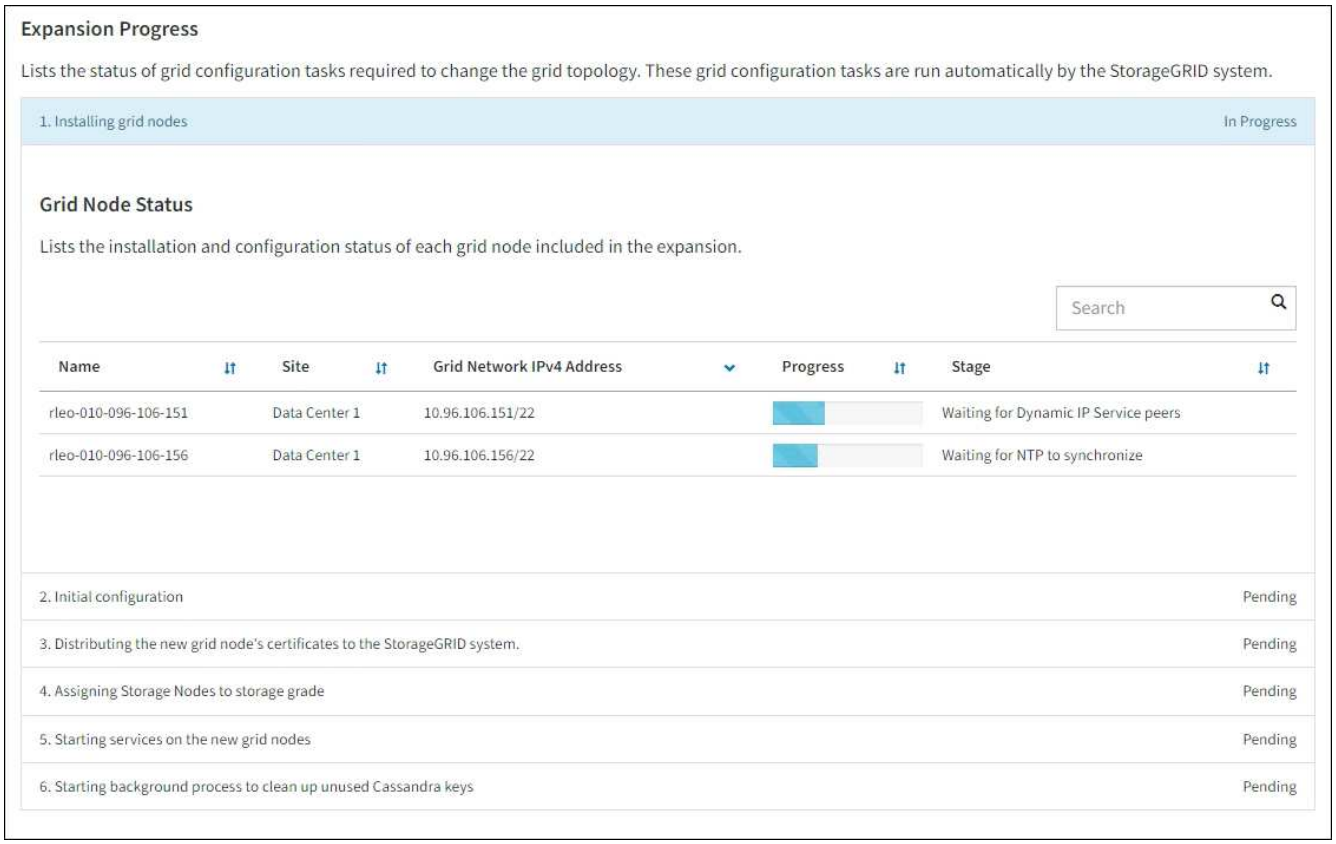

> A site expansion includes an additional task to configure Cassandra for the new site.

8. As soon as the **Download Recovery Package** link appears, download the Recovery Package file.

You must download an updated copy of the Recovery Package file as soon as possible after making grid

topology changes to the StorageGRID system. The Recovery Package file allows you to restore the system if a failure occurs.

a. Select the download link.

b. Enter the provisioning passphrase, and select **Start Download**.

c. When the download completes, open the `.zip` file and confirm that you can access the contents, including the `Passwords.txt` file.

d. Copy the downloaded Recovery Package file (`.zip`) to two safe, secure, and separate locations.

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

9. If you are adding Storage Nodes to an existing site or adding a site, monitor the Cassandra stages, which occur when services are started on the new grid nodes.

> Don’t reboot any Storage Nodes during either the "Expanding the Cassandra cluster" or "Starting Cassandra and streaming data" stages. These stages might take many hours to complete for each new Storage Node, especially if existing Storage Nodes contain a large amount of object metadata.

Adding Storage Nodes

If you are adding Storage Nodes to an existing site, review the percentage shown in the “Starting Cassandra and streaming data” status message.

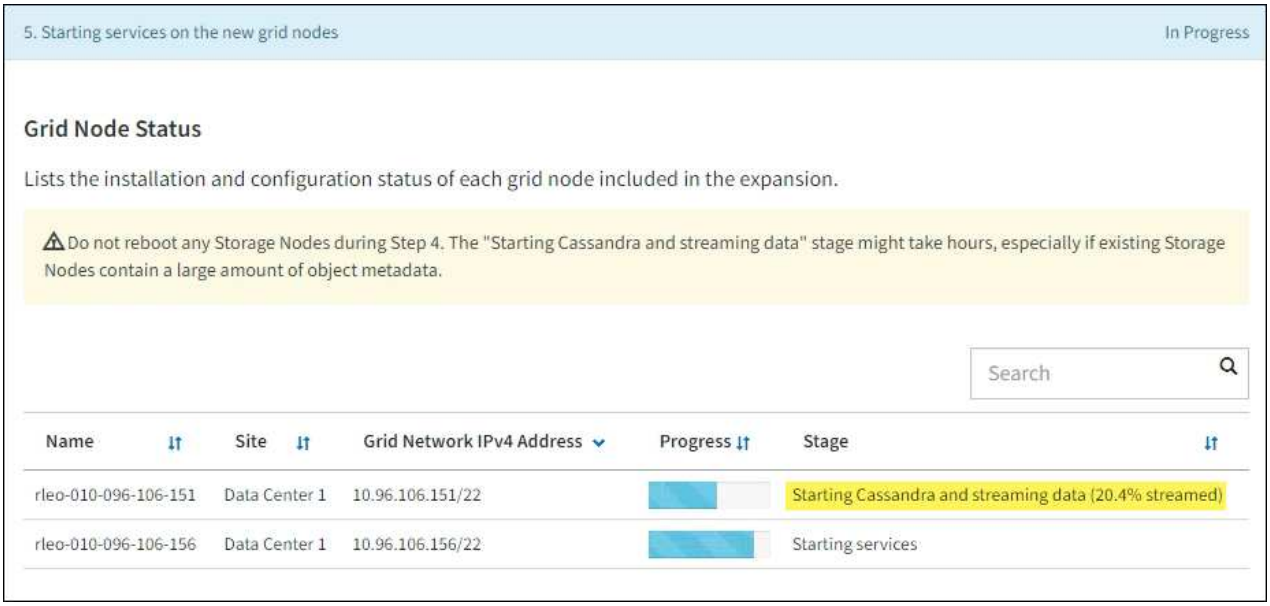

This percentage estimates how complete the Cassandra streaming operation is, based on the total amount of Cassandra data available and the amount that has already been written to the new node.

Adding site

If you are adding a new site, use `nodetool status` to monitor the progress of Cassandra streaming and to see how much metadata has been copied to the new site during the "Expanding the Cassandra cluster" stage. The total Data Load on the new site should be within about 20% of the total of a current site.

10. Continue monitoring the expansion until all tasks are complete and the **Configure Expansion** button reappears.

After you finish

Depending on which types of grid nodes you added, perform additional integration and configuration steps. See Configuration steps after expansion.

Configure expanded system

Configuration steps after expansion

After completing an expansion, you must perform additional integration and configuration steps.

About this task

You must complete the configuration tasks listed below for the grid nodes or sites you are adding in your expansion. Some tasks might be optional, depending on the options selected when installing and administering your system, and how you want to configure the nodes and sites added during the expansion.

Steps

1. If you added a site:
   - Create a storage pool for the site and each storage grade you selected for the new Storage Nodes.
   - Confirm that the ILM policy meets the new requirements. If rule changes are required, create new rules and update the ILM policy. If the rules are already correct, activate a new policy with no rule changes to ensure StorageGRID uses the new nodes.
   - Confirm that Network Time Protocol (NTP) servers are accessible from that site. See Manage NTP servers.

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

2. If you added one or more Storage Nodes to an existing site:
   - View storage pool details to confirm that each node you added is included in the expected storage pools and used in the expected ILM rules.
   - Confirm that the ILM policy meets the new requirements. If rule changes are required, create new rules and update the ILM policy. If the rules are already correct, activate a new policy with no rule changes to ensure StorageGRID uses the new nodes.
   - Verify that the Storage Node is active and able to ingest objects.
   - If you were unable to add the recommended number of Storage Nodes, rebalance erasure-coded data. See Rebalance erasure-coded data after adding Storage Nodes.
3. If you added a Gateway Node:
   - If high availability (HA) groups are used for client connections, optionally add the Gateway Node to an HA group. Select **CONFIGURATION** > **Network** > **High availability groups** to review the list of existing HA groups and to add the new node. See Configure high availability groups.
4. If you added an Admin Node:
   a. If single sign-on (SSO) is enabled for your StorageGRID system, create a relying party trust for the new Admin Node. You can't sign in to the node until you create this relying party trust. See Configure single sign-on.
   b. If you plan to use the Load Balancer service on Admin Nodes, optionally add the new Admin Node to an HA group. Select **CONFIGURATION** > **Network** > **High availability groups** to review the list of existing HA groups and to add the new node. See Configure high availability groups.
   c. Optionally, copy the Admin Node database from the primary Admin Node to the expansion Admin Node if you want to keep the attribute and audit information consistent on each Admin Node. See Copy the Admin Node database.
   d. Optionally, copy the Prometheus database from the primary Admin Node to the expansion Admin Node if you want to keep the historical metrics consistent on each Admin Node. See Copy Prometheus metrics.
   e. Optionally, copy the existing audit logs from the primary Admin Node to the expansion Admin Node if you want to keep the historical log information consistent on each Admin Node. See Copy audit logs.
   f. Optionally, configure access to the system for auditing purposes through an NFS file share. See Configure audit client access for NFS.
5. To check if expansion nodes were added with an untrusted Client Network or to change whether a node's Client Network is untrusted or trusted, go to **CONFIGURATION** > **Security** > **Firewall control**.

If the Client Network on the expansion node is untrusted, then connections to the node on the Client Network must be made using a load balancer endpoint. See Manage firewall controls.

6. Configure the DNS.

   If you have been specifying DNS settings separately for each grid node, you must add custom per-node DNS settings for the new nodes. See Modify DNS configuration for single grid node.

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

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

Verify that Storage Node is active

After an expansion operation that adds new Storage Nodes completes, the StorageGRID system should automatically start using the new Storage Nodes. You must use the StorageGRID system to verify that the new Storage Node is active.

Steps

1. Sign in to the Grid Manager using a supported web browser.
2. Select **NODES** > ***Expansion Storage Node*** > **Storage**.
3. Position your cursor over the **Storage Used - Object Data** graph to view the value for **Used**, which is the amount of the Total usable space that has been used for object data.
4. Verify that the value of **Used** is increasing as you move your cursor to the right on the graph.

Copy Admin Node database

When adding Admin Nodes through an expansion procedure, you can optionally copy the database from the primary Admin Node to the new Admin Node. Copying the database allows you to retain historical information about attributes, alerts, and alerts.

Before you begin

- You have completed the required expansion steps to add an Admin Node.
- You have the `Passwords.txt` file.
- You have the provisioning passphrase.

About this task

The StorageGRID software activation process creates an empty database for the NMS service on the expansion Admin Node. When the NMS service starts on the expansion Admin Node, it records information for servers and services that are currently part of the system or added later. This Admin Node database includes the following information:

- Alert history
- Alarm history
- Historical attribute data, which is used in the charts and text reports available from the **SUPPORT** > **Tools**

> **Grid topology** page

To ensure that the Admin Node database is consistent between nodes, you can copy the database from the primary Admin Node to the expansion Admin Node.

> Copying the database from the primary Admin Node (the*source Admin Node*) to an expansion Admin Node can take up to several hours to complete. During this period, the Grid Manager is inaccessible.

Use these steps to stop the MI service and the Management API service on both the primary Admin Node and the expansion Admin Node before copying the database.

Steps

1. Complete the following steps on the primary Admin Node:
   a. Log in to the Admin Node:
      i. Enter the following command: `ssh admin@grid_node_IP`
      ii. Enter the password listed in the `Passwords.txt` file.
      iii. Enter the following command to switch to root: `su -`
      iv. Enter the password listed in the `Passwords.txt` file.
   b. Run the following command: `recover-access-points`
   c. Enter the provisioning passphrase.
   d. Stop the MI service: `service mi stop`
   e. Stop the Management Application Program Interface (mgmt-api) service: `service mgmt-api stop`
2. Complete the following steps on the expansion Admin Node:
   a. Log in to the expansion Admin Node:
      i. Enter the following command: `ssh admin@grid_node_IP`
      ii. Enter the password listed in the `Passwords.txt` file.
      iii. Enter the following command to switch to root: `su -`
      iv. Enter the password listed in the `Passwords.txt` file.
   b. Stop the MI service: `service mi stop`
   c. Stop the mgmt-api service: `service mgmt-api stop`
   d. Add the SSH private key to the SSH agent. Enter:`ssh-add`
   e. Enter the SSH Access Password listed in the `Passwords.txt` file.
   f. Copy the database from the source Admin Node to the expansion Admin Node: `/usr/local/mi/bin/mi-clone-db.sh Source_Admin_Node_IP`
   g. When prompted, confirm that you want to overwrite the MI database on the expansion Admin Node.

      The database and its historical data are copied to the expansion Admin Node. When the copy operation is done, the script starts the expansion Admin Node.

   h. When you no longer require passwordless access to other servers, remove the private key from the SSH agent. Enter:`ssh-add -D`

3. Restart the services on the primary Admin Node: `service servermanager start`

Copy Prometheus metrics

After adding a new Admin Node, you can optionally copy the historical metrics maintained by Prometheus from the primary Admin Node to the new Admin Node. Copying the metrics ensures that historical metrics are consistent between Admin Nodes.

Before you begin

- The new Admin Node is installed and running.
- You have the `Passwords.txt` file.
- You have the provisioning passphrase.

About this task

When you add an Admin Node, the software installation process creates a new Prometheus database. You can keep the historical metrics consistent between nodes by copying the Prometheus database from the primary Admin Node (the *source Admin Node*) to the new Admin Node.

> Copying the Prometheus database might take an hour or more. Some Grid Manager features will be unavailable while services are stopped on the source Admin Node.

Steps

1. Log in to the source Admin Node:
   a. Enter the following command: `ssh admin@grid_node_IP`
   b. Enter the password listed in the `Passwords.txt` file.
   c. Enter the following command to switch to root: `su -`
   d. Enter the password listed in the `Passwords.txt` file.
2. From the source Admin Node, stop the Prometheus service: `service prometheus stop`
3. Complete the following steps on the new Admin Node:
   a. Log in to the new Admin Node:
      i. Enter the following command: `ssh admin@grid_node_IP`
      ii. Enter the password listed in the `Passwords.txt` file.
      iii. Enter the following command to switch to root: `su -`
      iv. Enter the password listed in the `Passwords.txt` file.
   b. Stop the Prometheus service: `service prometheus stop`
   c. Add the SSH private key to the SSH agent. Enter:`ssh-add`
   d. Enter the SSH Access Password listed in the `Passwords.txt` file.
   e. Copy the Prometheus database from the source Admin Node to the new Admin Node: `/usr/local/prometheus/bin/prometheus-clone-db.sh Source_Admin_Node_IP`
   f. When prompted, press **Enter** to confirm that you want to destroy the new Prometheus database on the new Admin Node.

The original Prometheus database and its historical data are copied to the new Admin Node. When the copy operation is done, the script starts the new Admin Node. The following status appears:

```
Database cloned, starting services
```

g. When you no longer require passwordless access to other servers, remove the private key from the SSH agent. Enter:

```
ssh-add -D
```

4. Restart the Prometheus service on the source Admin Node.

```
service prometheus start
```

Copy audit logs

When you add a new Admin Node through an expansion procedure, its AMS service only logs events and actions that occur after it joins the system. As required, you can copy audit logs from a previously installed Admin Node to the new expansion Admin Node so that it is in sync with the rest of the StorageGRID system.

Before you begin

- You have completed the required expansion steps to add an Admin Node.
- You have the `Passwords.txt` file.

About this task

To make historical audit messages available on a new Admin Node, you must copy the audit log files manually from an existing Admin Node to the expansion Admin Node.

> By default, audit information is sent to the audit log on Admin Nodes. You can skip these steps if either of the following applies:
>
> - You configured an external syslog server and audit logs are now being sent to the syslog server instead of to Admin Nodes.
> - You explicitly specified that audit messages should be saved only on the local nodes that generated them.
>
> See Configure audit messages and log destinations for details.

Steps

1. Log in to the primary Admin Node:
   a. Enter the following command: `ssh admin@_primary_Admin_Node_IP`
   b. Enter the password listed in the `Passwords.txt` file.
   c. Enter the following command to switch to root: `su -`
   d. Enter the password listed in the `Passwords.txt` file.

   When you are logged in as root, the prompt changes from `$` to `#`.

2. Stop the AMS service to prevent it from creating a new file: `service ams stop`
3. Rename the `audit.log` file to ensure that it does not overwrite the file on the expansion Admin Node you are copying it to:

   ```
   cd /var/local/audit/export
   ls -l
   mv audit.log new_name.txt
   ```

4. Copy all audit log files to the expansion Admin Node:

   ```
   scp -p * IP_address:/var/local/audit/export
   ```

5. If prompted for the passphrase for `/root/.ssh/id_rsa`, enter the SSH Access Password for the Primary Admin Node listed in the `Passwords.txt` file.
6. Restore the original `audit.log` file:

   ```
   mv new_name.txt audit.log
   ```

7. Start the AMS service:

   ```
   service ams start
   ```

8. Log out from the server:

   ```
   exit
   ```

9. Log in to the expansion Admin Node:
   a. Enter the following command: `ssh admin@expansion_Admin_Node_IP`
   b. Enter the password listed in the `Passwords.txt` file.
   c. Enter the following command to switch to root: `su -`
   d. Enter the password listed in the `Passwords.txt` file.

      When you are logged in as root, the prompt changes from `$` to `#`.

10. Update the user and group settings for the audit log files:

    ```
    cd /var/local/audit/export
    chown ams-user:bycast *
    ```

11. Log out from the server:

    ```
    exit
    ```

Rebalance erasure-coded data after adding Storage Nodes

After you add Storage Nodes, you can use the EC rebalance procedure to redistribute erasure-coded fragments among the existing and new Storage Nodes.

Before you begin

- You have completed the expansion steps to add the new Storage Nodes.
- You have reviewed the considerations for rebalancing erasure-coded data.
- You understand that replicated object data will not be moved by this procedure and that the EC rebalance procedure does not consider the replicated data usage on each Storage Node when determining where to move erasure-coded data.
- You have the `Passwords.txt` file.

What happens when this procedure runs

Before starting the procedure, note the following:

- The EC rebalance procedure will not start if one or more volumes are offline (unmounted) or if they are online (mounted) but in an error state.
- The EC rebalance procedure temporarily reserves a large amount of storage. Storage alerts might be triggered, but will resolve when the rebalance is complete. If there is not enough storage for the reservation, the EC rebalance procedure will fail. Storage reservations are released when the EC rebalance procedure completes, whether the procedure failed or succeeded.
- If a volume goes offline or experiences an error while an EC rebalance is in process, the rebalance process terminates partially completed with no loss of data. The EC rebalance procedure can be resumed at the point where it terminated when all volumes are online without error.
- When the EC rebalance procedure is running, the performance of ILM operations and S3 and Swift client operations might be impacted.

> ⓘ S3 and Swift API operations to upload objects (or object parts) might fail during the EC rebalancing procedure if they require more than 24 hours to complete. Long-duration PUT operations will fail if the applicable ILM rule uses Balanced or Strict placement on ingest. The following error will be reported: `500 Internal Server Error`.

Steps

1. Review the current object storage details for the site you plan to rebalance.
   a. Select **NODES**.
   b. Select the first Storage Node at the site.
   c. Select the **Storage** tab.
   d. Position your cursor over the Storage Used - Object Data chart to see the current amount of replicated data and erasure-coded data on the Storage Node.
   e. Repeat these steps to view the other Storage Nodes at the site.
2. Log in to the primary Admin Node:
   a. Enter the following command: `ssh admin@primary_Admin_Node_IP`
   b. Enter the password listed in the `Passwords.txt` file.
   c. Enter the following command to switch to root: `su -`
   d. Enter the password listed in the `Passwords.txt` file.

      When you are logged in as root, the prompt changes from `$` to `#`.
3. Start the procedure:

```
rebalance-data start --site "site-name"
```

For "*site-name*", specify the first site where you added new Storage Node or nodes. Enclose `site-name` in quotes.

The EC rebalance procedure starts, and a job ID is returned.

4. Copy the job ID.
5. Monitor the status of the EC rebalance procedure.
   - To view the status of a single EC rebalance procedure:

     ```
     rebalance-data status --job-id job-id
     ```

     For *job-id*, specify the ID that was returned when you started the procedure.

   - To view the status of the current EC rebalance procedure and any previously completed procedures:

     ```
     rebalance-data status
     ```

     > To get help on the rebalance-data command:
     >
     > ```
     > rebalance-data --help
     > ```

6. Perform additional steps, based on the status returned:
   - If the status is `In progress`, the EC rebalance operation is still running. You should periodically monitor the procedure until it completes.

     To view the estimated time to completion and the completion percentage for the current job:

     a. Select **SUPPORT** > **Tools** > **Metrics**.
     b. Select **EC Overview** in the Grafana section.
     c. Look at the **Grid EC Job Estimated Time to Completion** and **Grid EC Job Percentage Completed** dashboards.

   - If the status is `Success`, optionally review object storage to see the updated details for the site.

     Erasure-coded data should now be more balanced among the Storage Nodes at the site.

     > If the following message appears, run the EC rebalance procedure again until all erasure-coded data has been rebalanced:
     >
     > ```
     > The moves in this rebalance job have been limited. To rebalance
     > additional data, start EC rebalance again for the same site.
     > ```

   - If the status is `Failure`:

     a. Confirm that all Storage Nodes at the site are connected to the grid.
     b. Check for and resolve any alerts that might be affecting these Storage Nodes.
     c. Restart the EC rebalance procedure:

        ```
        rebalance-data start --job-id job-id
        ```

d. If the status of the EC rebalance procedure is still `Failure`, contact technical support.

7. If the EC rebalance procedure is generating too much load (for example, ingest operations are affected), pause the procedure.

   `rebalance-data pause --job-id job-id`

8. If you need to terminate the EC rebalance procedure (for example, so you can perform a StorageGRID software upgrade), enter the following:

   `rebalance-data terminate --job-id job-id`

   > When you terminate an EC rebalance procedure, any data fragments that have already been moved remain in the new location. Data is not moved back to the original location.

9. If you are using erasure coding at more than one site, run this procedure for all other affected sites.

Contact technical support

If you encounter errors during the grid expansion process that you are unable to resolve, or if a grid task fails, contact technical support.

About this task

When you contact technical support, you must provide the required log files to assist in troubleshooting the errors you are encountering.

Steps

1. Connect to the expansion node that has experienced failures:

   a. Enter the following command:`ssh -p 8022 admin@grid_node_IP`

   > Port 8022 is the SSH port of the base OS, while port 22 is the SSH port of the container engine running StorageGRID.

   b. Enter the password listed in the `Passwords.txt` file.

   c. Enter the following command to switch to root: `su -`

   d. Enter the password listed in the `Passwords.txt` file.

   Once you are logged in as root, the prompt changes from `$` to `#`.

2. Depending on the stage the installation reached, retrieve any of the following logs that are available on the grid node:

Platform	Logs
VMware	• `/var/log/daemon.log` • `/var/log/storagegrid/daemon.log` • `/var/log/storagegrid/nodes/<node-name>.log`

Platform	Logs
Linux	• `/var/log/storagegrid/daemon.log` • `/etc/storagegrid/nodes/<node-name>.conf` (for each failed node) • `/var/log/storagegrid/nodes/<node-name>.log` (for each failed node; might not exist)

Recover nodes and maintain your grid

Recover and maintain: Overview

Use these instructions to maintain your StorageGRID system and to recover from failures.

About these instructions

These instructions describe how to apply a software hotfix, recover grid nodes, recover a failed site, decommission grid nodes or an entire site, perform network maintenance, perform host-level and middleware maintenance procedures, and perform grid node procedures.

> In these instructions, “Linux” refers to a Red Hat® Enterprise Linux®, Ubuntu®, CentOS, or Debian® deployment. Use the NetApp Interoperability Matrix Tool to get a list of supported versions.

Before you begin

- You have a broad understanding of the StorageGRID system.
- You have reviewed your StorageGRID system’s topology and you understand the grid configuration.
- You understand that you must follow all instructions exactly and heed all warnings.
- You understand that maintenance procedures not described aren’t supported or require a services engagement.

Maintenance procedures for appliances

For hardware procedures, see the maintenance instructions for your StorageGRID appliance.

- Maintain SGF6112 appliance
- Maintain SG100 and SG1000 appliances
- Maintain SG6000 appliance
- Maintain SG5700 appliance

Download Recovery Package

The Recovery Package file allows you to restore the StorageGRID system if a failure occurs.

Before you begin

- From the primary Admin Node, you are signed in to the Grid Manager using a supported web browser.
- You have the provisioning passphrase.
- You have specific access permissions.

Download the current Recovery Package file before making grid topology changes to the StorageGRID system or before upgrading software. Then, download a new copy of the Recovery Package after making grid topology changes or after upgrading software.

Steps

1. Select **MAINTENANCE** > **System** > **Recovery package**.
2. Enter the provisioning passphrase, and select **Start Download**.

   The download starts immediately.

3. When the download completes, open the `.zip` file and confirm that you can access the contents, including the `Passwords.txt` file.
4. Copy the downloaded Recovery Package file (`.zip`) to two safe, secure, and separate locations.

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

StorageGRID hotfix procedure

StorageGRID hotfix procedure: Overview

You might need to apply a hotfix to your StorageGRID system if issues with the software are detected and resolved between feature releases.

StorageGRID hotfixes contain software changes that are made available outside of a feature or patch release. The same changes are included in a future release. In addition, each hotfix release contains a roll-up of all previous hotfixes within the feature or patch release.

Considerations for applying a hotfix

You can't apply a StorageGRID hotfix when another maintenance procedure is running. For example, you can't apply a hotfix while a decommission, expansion, or recovery procedure is running.

> If a node or site decommission procedure is paused, you can safely apply a hotfix. In addition, you might be able to apply a hotfix during the final stages of a StorageGRID upgrade procedure. See the instructions for upgrading StorageGRID software for details.

After you upload the hotfix in the Grid Manager, the hotfix is applied automatically to the primary Admin Node. Then, you can approve the application of the hotfix to the rest of the nodes in your StorageGRID system.

If a hotfix fails to be applied to one or more nodes, the reason for the failure appears in the Details column of the hotfix progress table. You must resolve whatever issues caused the failures and then retry the entire process. Nodes with a previously successful application of the hotfix will be skipped in subsequent applications. You can safely retry the hotfix process as many times as required until all nodes have been updated. The hotfix must be successfully installed on all grid nodes in order for the application to be complete.

While grid nodes are updated with the new hotfix version, the actual changes in a hotfix might only affect specific services on specific types of nodes. For example, a hotfix might only affect the LDR service on Storage Nodes.

How hotfixes are applied for recovery and expansion

After a hotfix has been applied to your grid, the primary Admin Node automatically installs the same hotfix version to any nodes restored by recovery operations or added in an expansion.

However, if you need to recover the primary Admin Node, you must manually install the correct StorageGRID release and then apply the hotfix. The final StorageGRID version of the primary Admin Node must match the version of the other nodes in the grid.

The following example illustrates how to apply a hotfix when recovering the primary Admin Node:

1. Assume the grid is running a StorageGRID 11.*A.B* version with the latest hotfix. The “grid version” is 11.*A.B.y*.
2. The primary Admin Node fails.
3. You redeploy the primary Admin Node using StorageGRID 11.*A.B*, and perform the recovery procedure.

   As required to match the grid version, you can use a minor release when deploying the node; you don’t need to deploy the major release first.

4. You then apply hotfix 11.*A.B.y* to the primary Admin Node.

For more information, see Configure replacement primary Admin Node.

How your system is affected when you apply a hotfix

You must understand how your StorageGRID system will be affected when you apply a hotfix.

Client applications might experience short-term disruptions

The StorageGRID system can ingest and retrieve data from client applications throughout the hotfix process; however, client connections to individual Gateway Nodes or Storage Nodes might be disrupted temporarily if the hotfix needs to restart services on those nodes. Connectivity will be restored after the hotfix process completes and services resume on the individual nodes.

You might need to schedule downtime to apply a hotfix if loss of connectivity for a short period is not acceptable. You can use selective approval to schedule when certain nodes are updated.

You can use multiple gateways and high availability (HA) groups to provide automatic failover during the hotfix process. See the instructions for configuring high availability groups.

Alerts and SNMP notifications might be triggered

Alerts and SNMP notifications might be triggered when services are restarted and when the StorageGRID system is operating as a mixed-version environment (some grid nodes running an earlier version, while others have been upgraded to a later version). In general, these alerts and notifications will clear when the hotfix completes.

Configuration changes are restricted

When applying a hotfix to StorageGRID:

- Don't make any grid configuration changes (for example, specifying Grid Network subnets or approving pending grid nodes) until the hotfix has been applied to all nodes.
- Don't update the ILM configuration until the hotfix has been applied to all nodes.

Obtain required materials for hotfix

Before applying a hotfix, you must obtain all required materials.

Item	Notes
StorageGRID hotfix file	You must download the StorageGRID hotfix file.
• Network port • Supported web browser • SSH client (for example, PuTTY)	
Recovery Package (`.zip`) file	Before applying a hotfix, download the most recent Recovery Package file in case any problems occur during the hotfix.Then, after the hotfix has been applied, download a new copy of the Recovery Package file and save it in a safe location. The updated Recovery Package file allows you to restore the system if a failure occurs.
`Passwords.txt file`	Optional and used only if you are applying a hotfix manually using the SSH client. The `Passwords.txt` file is part of the Recovery Package `.zip` file.
Provisioning passphrase	The passphrase is created and documented when the StorageGRID system is first installed. The provisioning passphrase is not listed in the `Passwords.txt` file.
Related documentation	`readme.txt` file for the hotfix. This file is included on the hotfix download page. Be sure to review the `readme` file carefully before applying the hotfix.

Download hotfix file

You must download the hotfix file before you can apply the hotfix.

Steps

1. Go to NetApp Downloads: StorageGRID.
2. Select the down arrow under **Available Software** to see a list of hotfixes that are available to download.

   Hotfix file versions have the form: 11.4.*x.y*.

3. Review the changes that are included in the update.

   > If you have just recovered the primary Admin Node and you need to apply a hotfix, select the same hotfix version that is installed on the other grid nodes.

   a. Select the hotfix version you want to download, and select **Go**.
   b. Sign in using the username and password for your NetApp account.
   c. Read and accept the End User License Agreement.

      The download page for the version you selected appears.

   d. Download the hotfix `readme.txt` file to view a summary of the changes included in the hotfix.
4. Select the download button for the hotfix, and save the file.

   > Don't change the name of this file.

   > If you are using a macOS device, the hotfix file might be automatically saved as a `.txt` file. If it is, you must rename the file without the `.txt` extension.

5. Select a location for the download, and select **Save**.

Check system's condition before applying hotfix

You must verify the system is ready to accommodate the hotfix.

1. Sign in to the Grid Manager using a supported web browser.
2. If possible, ensure that the system is running normally and that all grid nodes are connected to the grid.

   Connected nodes have green check marks on the Nodes page.

3. Check for and resolve any current alerts, if possible.
4. Ensure no other maintenance procedures are in progress, such as an upgrade, recovery, expansion, or decommission procedure.

   You should wait for any active maintenance procedures to complete before applying a hotfix.

   You can't apply a StorageGRID hotfix when another maintenance procedure is running. For example, you can't apply a hotfix while a decommission, expansion, or recovery procedure is running.

   > If a node or site decommission procedure is paused, you can safely apply a hotfix. In addition, you might be able to apply a hotfix during the final stages of a StorageGRID upgrade procedure. See the instructions for upgrading StorageGRID software.

Apply hotfix

The hotfix is first applied automatically to the primary Admin Node. Then, you must approve the application of the hotfix to other grid nodes until all nodes are running the same software version. You can customize the approval sequence by selecting to

approve individual grid nodes, groups of grid nodes, or all grid nodes.

Before you begin

- You have reviewed the considerations for applying a hotfix.
- You have the provisioning passphrase.
- You have Root access or the Maintenance permission.

About this task

- You can delay applying a hotfix to a node, but the hotfix process is not complete until you apply the hotfix to all nodes.
- You can't perform a StorageGRID software upgrade or a SANtricity OS update until you have completed the hotfix process.

Steps

1. Sign in to the Grid Manager using a supported web browser.
2. Select **MAINTENANCE** > **System** > **Software update**.

   The Software Update page appears.

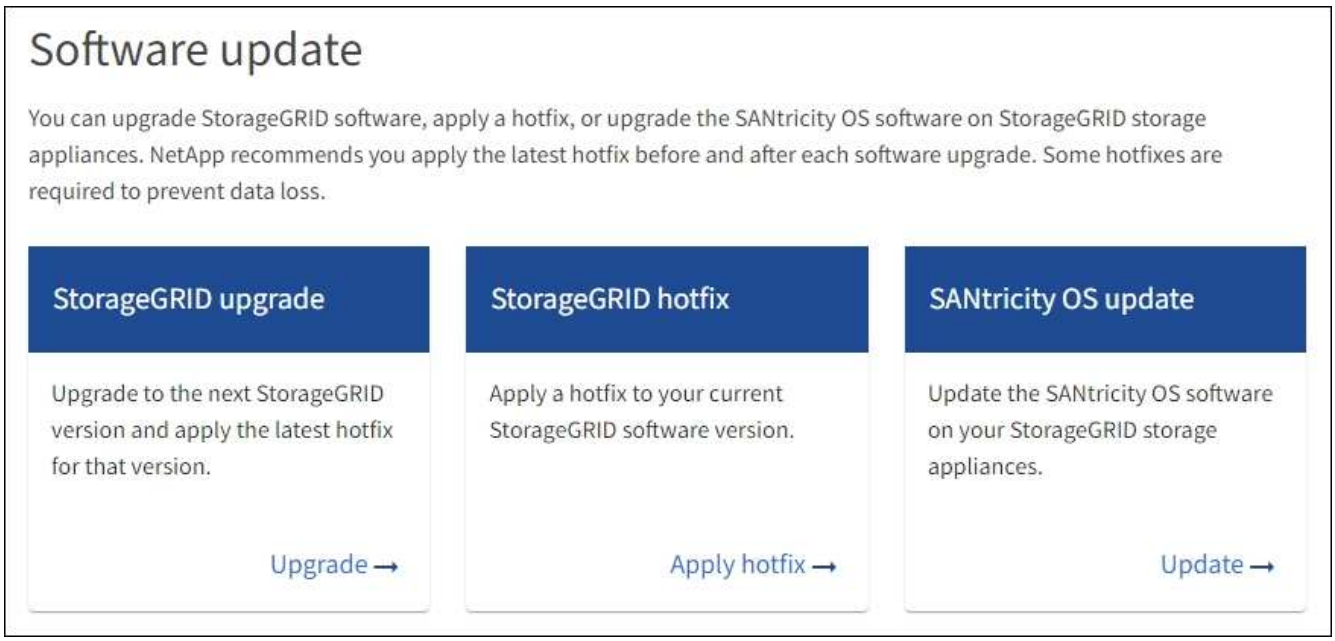

3. Select **Apply hotfix**.

   The StorageGRID Hotfix page appears.

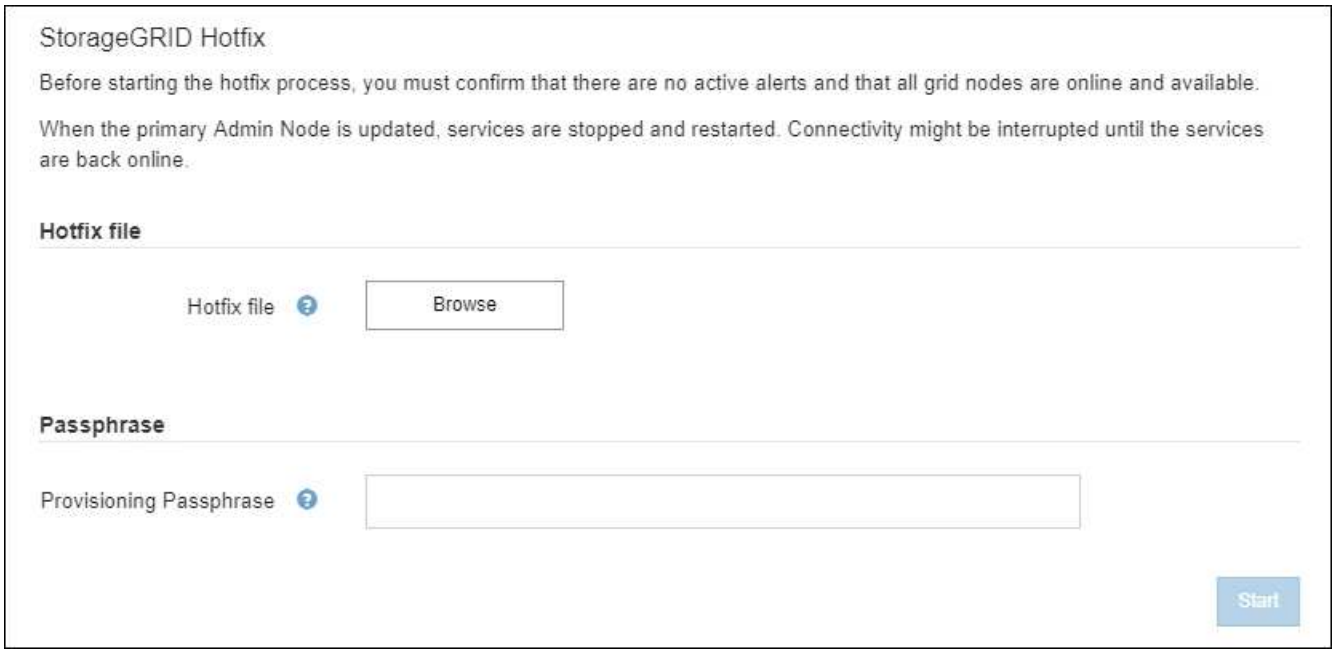

4. Select the hotfix file you downloaded from the NetApp Support Site.
   a. Select **Browse**.
   b. Locate and select the file.

      `hotfix-install-version`

   c. Select **Open**.

      The file is uploaded. When the upload is finished, the file name is shown in the Details field.

      Don't change the file name because it is part of the verification process.

5. Enter the provisioning passphrase in the text box.

   The **Start** button becomes enabled.

6. Select **Start**.

   A warning appears stating that your browser's connection might be lost temporarily as services on the primary Admin Node are restarted.

7. Select **OK** to start applying the hotfix to the primary Admin Node.

   When the hotfix starts:

   a. The hotfix validations are run.

      If any errors are reported, resolve them, re-upload the hotfix file, and select **Start** again.

   b. The hotfix installation progress table appears.

      This table shows all nodes in your grid and the current stage of the hotfix installation for each node. The nodes in the table are grouped by type (Admin Nodes, Gateway Nodes, Storage Nodes, and

Archive Nodes).

c. The progress bar reaches completion, and then the primary Admin Node is shown as “Complete.”

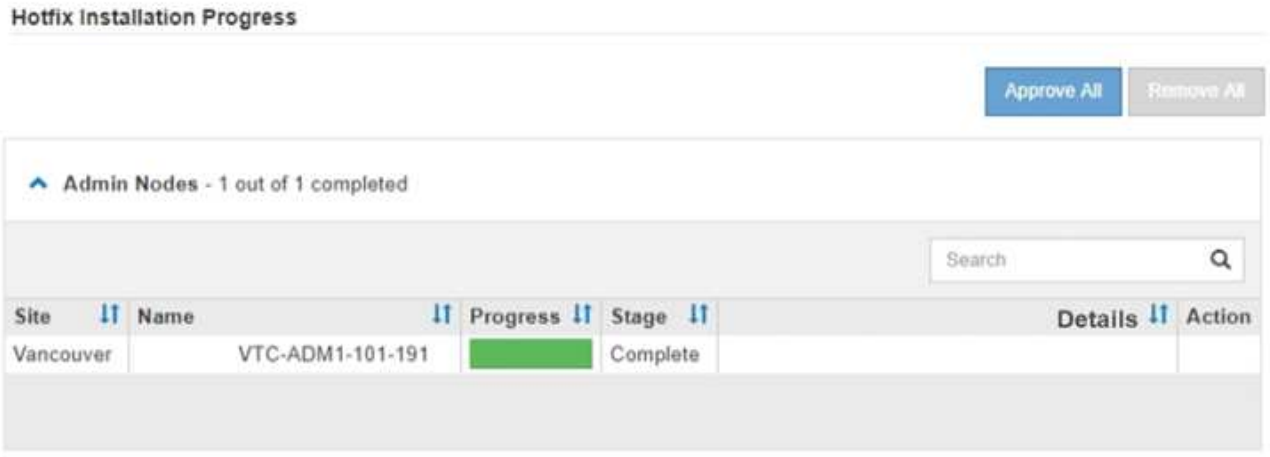

8. Optionally, sort the lists of nodes in each grouping in ascending or descending order by **Site**, **Name**, **Progress**, **Stage**, or **Details**. Or, enter a term in the **Search** box to search for specific nodes.
9. Approve the grid nodes that are ready to be updated. Approved nodes of the same type are upgraded one at a time.

> Don’t approve the hotfix for a node unless you are sure the node is ready to be updated. When the hotfix is applied to a grid node, some services on that node might be restarted. These operations might cause service interruptions for clients that are communicating with the node.

- Select one or more **Approve** buttons to add one or more individual nodes to the hotfix queue.
- Select the **Approve All** button within each grouping to add all nodes of the same type to the hotfix queue. If you have entered search criteria in the **Search** box, the **Approve All** button applies to all the nodes selected by the search criteria.

> The **Approve All** button at the top of the page approves all nodes listed on the page, while the **Approve All** button at the top of a table grouping only approves all nodes in that group. If the order in which nodes are upgraded is important, approve nodes or groups of nodes one at a time and wait until the upgrade is complete on each node before approving the next node(s).

- Select the top-level **Approve All** button at the top of the page to add all nodes in the grid to the hotfix queue.

> You must complete the StorageGRID hotfix before you can start a different software update. If you are unable to complete the hotfix, contact technical support.

- Select **Remove** or **Remove All** to remove a node or all nodes from the hotfix queue.

  When the Stage progresses beyond “Queued,” the **Remove** button is hidden and you can no longer remove the node from the hotfix process.

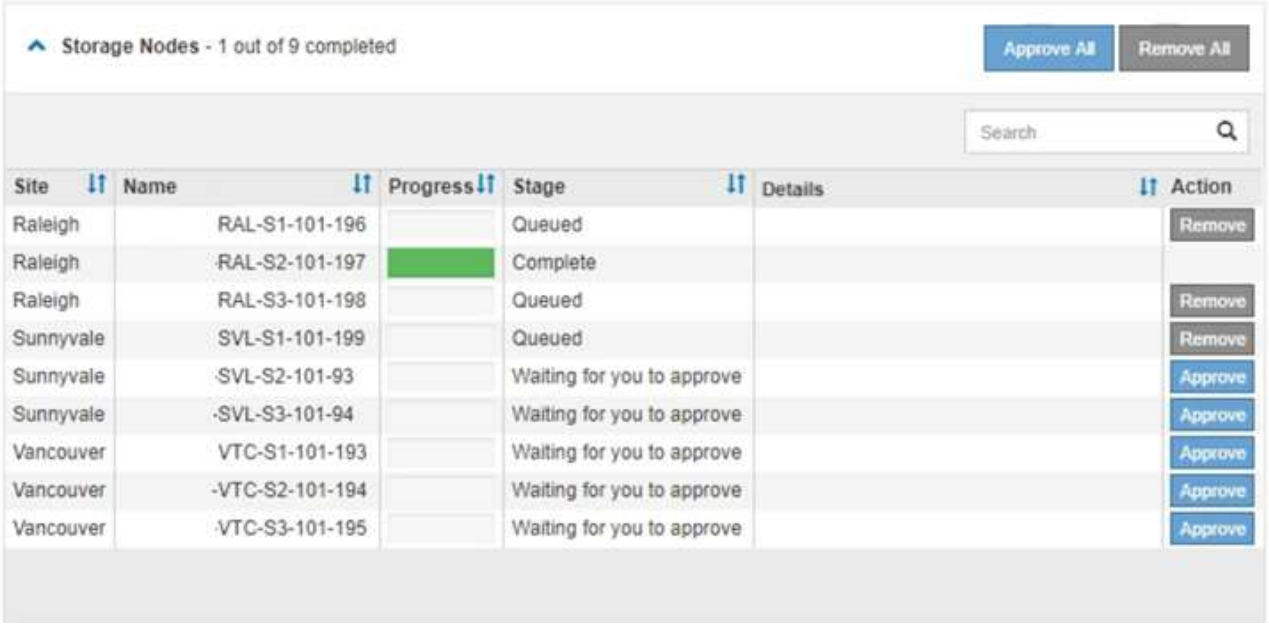

10. Wait while the hotfix is applied to each approved grid node.

    When the hotfix has been successfully installed on all nodes, the Hotfix Installation Progress table closes. A green banner shows the date and time the hotfix was completed.

11. If the hotfix could not be applied to any nodes, review the error for each node, resolve the issue, and repeat these steps.

    The procedure is not complete until the hotfix is successfully applied to all nodes. You can safely retry the hotfix process as many times as required until it is complete.

Grid node recovery procedures

Grid node recovery procedures: Overview

If a grid node fails, you can recover it by replacing the failed physical or virtual server, reinstalling StorageGRID software, and restoring recoverable data.

Grid nodes can fail if a hardware, virtualization, operating system, or software fault renders the node inoperable or unreliable. There are many kinds of failure that can trigger the need to recover a grid node.

The steps to recover a grid node vary, depending on the platform where the grid node is hosted and on the type of grid node. Each type of grid node has a specific recovery procedure, which you must follow exactly.

Generally, you try to preserve data from the failed grid node where possible, repair or replace the failed node, use the Grid Manager to configure the replacement node, and restore the node's data.

> If an entire StorageGRID site has failed, contact technical support. Technical support will work with you to develop and execute a site recovery plan that maximizes the amount of data that is recovered, and meets your business objectives. See How site recovery is performed by technical support.

Warnings and considerations for grid node recovery

If a grid node fails, you must recover it as soon as possible. You must review all warnings and considerations for node recovery before you begin.

> StorageGRID is a distributed system composed of multiple nodes working with each other. Don't use disk snapshots to restore grid nodes. Instead, refer to the recovery and maintenance procedures for each type of node.

Some of the reasons for recovering a failed grid node as soon as possible include the following:

- A failed grid node can reduce the redundancy of system and object data, leaving you vulnerable to the risk of permanent data loss if another node fails.
- A failed grid node can impact the efficiency of day-to-day operations.
- A failed grid node can reduce your ability to monitor system operations.
- A failed grid node can cause a 500 internal server error if strict ILM rules are in place.
- If a grid node is not recovered promptly, recovery times might increase. For example, queues might develop that need to be cleared before recovery is complete.

Always follow the recovery procedure for the specific type of grid node you are recovering. Recovery procedures vary for primary or non-primary Admin Nodes, Gateway Nodes, Archive Nodes, appliance nodes, and Storage Nodes.

Preconditions for recovering grid nodes

All of the following conditions are assumed when recovering grid nodes:

- The failed physical or virtual hardware has been replaced and configured.
- The StorageGRID Appliance Installer version on the replacement appliance matches the software version of your StorageGRID system, as described in Verify and upgrade StorageGRID Appliance Installer version.
- If you are recovering a grid node other than the primary Admin Node, there is connectivity between the grid node being recovered and the primary Admin Node.

Order of node recovery if a server hosting more than one grid node fails

If a server that is hosting more than one grid node fails, you can recover the nodes in any order. However, if the failed server is hosting the primary Admin Node, you must recover that node first. Recovering the primary Admin Node first prevents other node recoveries from halting as they wait to contact the primary Admin Node.

IP addresses for recovered nodes

Don't attempt to recover a node using an IP address that is currently assigned to any other node. When you deploy the new node, use the failed node's current IP address or an unused IP address.

If you use a new IP address to deploy the new node and then recover the node, the new IP address will continue to be used for the recovered node. If you want to revert to the original IP address, use the Change IP tool after the recovery is complete.

Gather required materials for grid node recovery

Before performing maintenance procedures, you must ensure you have the necessary

materials to recover a failed grid node.

Item	Notes
StorageGRID installation archive	If you need to recover a grid node, you need to download the StorageGRID installation files for your platform. **Note:** You don't need to download files if you are recovering failed storage volumes on a Storage Node.
Service laptop	The service laptop must have the following: • Network port • SSH client (for example, PuTTY) • Supported web browser
Recovery Package `.zip` file	Obtain a copy of the most recent Recovery Package `.zip` file: `sgws-recovery-package-id-revision.zip` The contents of the `.zip` file are updated each time the system is modified. You are directed to store the most recent version of the Recovery Package in a secure location after making such changes. Use the most recent copy to recover from grid failures. If the primary Admin Node is operating normally, you can download the Recovery Package from the Grid Manager. Select **MAINTENANCE** > **System** > **Recovery package**. If you can't access the Grid Manager, you can find encrypted copies of the Recovery Package on some Storage Nodes that contain the ADC service. On each Storage Node, examine this location for the Recovery Package: `/var/local/install/sgws-recovery-package-grid-id-revision.zip.gpg` Use the Recovery Package with the highest revision number.
`Passwords.txt` file	Contains the passwords required to access grid nodes on the command line. Included in the Recovery Package.
Provisioning passphrase	The passphrase is created and documented when the StorageGRID system is first installed. The provisioning passphrase is not in the `Passwords.txt` file.
Current documentation for your platform	Go to the platform vendor's website for documentation. For the current supported versions of your platform, see the NetApp Interoperability Matrix Tool.

Download and extract StorageGRID installation files

Download the software and extract the files, unless you are recovering failed storage volumes on a Storage Node.

You must use the version of StorageGRID that is currently running on the grid.

Steps

1. Determine which version of the software is currently installed. From the top of the Grid Manager, select the help icon and select **About**.
2. Go to the NetApp Downloads page for StorageGRID.
3. Select the version of StorageGRID that is currently running on the grid.

   StorageGRID software versions have this format: `11.x.y`.

4. Sign in with the username and password for your NetApp account.
5. Read the End User License Agreement, select the checkbox, and then select **Accept & Continue**.
6. In the **Install StorageGRID** column of the download page, select the `.tgz` or `.zip` file for your platform.

   The version shown in the installation archive file must match the version of the software that is currently installed.

   Use the `.zip` file if you are running Windows.

Platform	Installation archive

Red Hat Enterprise Linux or CentOS	`StorageGRID-Webscale-`*`version`*`-RPM-`*`uniqueID`*`.zip`
	
	`StorageGRID-Webscale-`*`version`*`-RPM-`*`uniqueID`*`.tgz`

Ubuntu or Debian or Appliances	`StorageGRID-Webscale-`*`version`*`-DEB-`*`uniqueID`*`.zip`
	
	`StorageGRID-Webscale-`*`version`*`-DEB-`*`uniqueID`*`.tgz`

VMware	`StorageGRID-Webscale-`*`version`*`-VMware-`*`uniqueID`*`.zip`
	
	`StorageGRID-Webscale-`*`version`*`-VMware-`*`uniqueID`*`.tgz`

7. Download and extract the archive file.
8. Follow the appropriate step for your platform to choose the files you need, based on your platform and which grid nodes you need to recover.

   The paths listed in the step for each platform are relative to the top-level directory installed by the archive file.

9. If you are recovering a Red Hat Enterprise Linux or CentOS system, select the appropriate files.

Path and file name	Description
`./rpms/README`	A text file that describes all of the files contained in the StorageGRID download file.
`./rpms/NLF000000.txt`	A free license that does not provide any support entitlement for the product.

Path and file name	Description
`./rpms/StorageGRID-Webscale-Images-version-SHA.rpm`	RPM package for installing the StorageGRID node images on your RHEL or CentOS hosts.
`./rpms/StorageGRID-Webscale-Service-version-SHA.rpm`	RPM package for installing the StorageGRID host service on your RHEL or CentOS hosts.
Deployment scripting tool	**Description**
`./rpms/configure-storagegrid.py`	A Python script used to automate the configuration of a StorageGRID system.
`./rpms/configure-sga.py`	A Python script used to automate the configuration of StorageGRID appliances.
`./rpms/configure-storagegrid.sample.json`	An example configuration file for use with the `configure-storagegrid.py` script.
`./rpms/storagegrid-ssoauth.py`	An example Python script that you can use to sign in to the Grid Management API when single sign-on is enabled. You can also use this script for Ping Federate.
`./rpms/configure-storagegrid.blank.json`	A blank configuration file for use with the `configure-storagegrid.py` script.
`./rpms/extras/ansible`	Example Ansible role and playbook for configuring RHEL or CentOS hosts for StorageGRID container deployment. You can customize the role or playbook as necessary.
`./rpms/storagegrid-ssoauth-azure.py`	An example Python script that you can use to sign in to the Grid Management API when single sign-on (SSO) is enabled using Active Directory or Ping Federate.
`./rpms/storagegrid-ssoauth-azure.js`	A helper script called by the companion `storagegrid-ssoauth-azure.py` Python script to perform SSO interactions with Azure.
`./rpms/extras/api-schemas`	API schemas for StorageGRID. **Note**: Before you perform an upgrade, you can use these schemas to confirm that any code you have written to use StorageGRID management APIs will be compatible with the new StorageGRID release if you don’t have a non-production StorageGRID environment for upgrade compatibility testing.

10. If you are recovering an Ubuntu or Debian system, select the appropriate files.

Path and file name	Description
`./debs/README`	A text file that describes all of the files contained in the StorageGRID download file.
`./debs/NLF000000.txt`	A non-production NetApp License File that you can use for testing and proof of concept deployments.
`./debs/storagegrid-webscale-images-version-SHA.deb`	DEB package for installing the StorageGRID node images on Ubuntu or Debian hosts.
`./debs/storagegrid-webscale-images-version-SHA.deb.md5`	MD5 checksum for the file `/debs/storagegrid-webscale-images-version-SHA.deb`.
`./debs/storagegrid-webscale-service-version-SHA.deb`	DEB package for installing the StorageGRID host service on Ubuntu or Debian hosts.
Deployment scripting tool	**Description**
`./debs/configure-storagegrid.py`	A Python script used to automate the configuration of a StorageGRID system.
`./debs/configure-sga.py`	A Python script used to automate the configuration of StorageGRID appliances.
`./debs/storagegrid-ssoauth.py`	An example Python script that you can use to sign in to the Grid Management API when single sign-on is enabled. You can also use this script for Ping Federate.
`./debs/configure-storagegrid.sample.json`	An example configuration file for use with the `configure-storagegrid.py` script.
`./debs/configure-storagegrid.blank.json`	A blank configuration file for use with the `configure-storagegrid.py` script.
`./debs/extras/ansible`	Example Ansible role and playbook for configuring Ubuntu or Debian hosts for StorageGRID container deployment. You can customize the role or playbook as necessary.
`./debs/storagegrid-ssoauth-azure.py`	An example Python script that you can use to sign in to the Grid Management API when single sign-on (SSO) is enabled using Active Directory or Ping Federate.

Path and file name	Description
`./debs/storagegrid-ssoauth-azure.js`	A helper script called by the companion `storagegrid-ssoauth-azure.py` Python script to perform SSO interactions with Azure.
`./debs/extras/api-schemas`	API schemas for StorageGRID. **Note**: Before you perform an upgrade, you can use these schemas to confirm that any code you have written to use StorageGRID management APIs will be compatible with the new StorageGRID release if you don't have a non-production StorageGRID environment for upgrade compatibility testing.

11. If you are recovering a VMware system, select the appropriate files.

Path and file name	Description
`./vsphere/README`	A text file that describes all of the files contained in the StorageGRID download file.
`./vsphere/NLF000000.txt`	A free license that does not provide any support entitlement for the product.
`./vsphere/NetApp-SG-version-SHA.vmdk`	The virtual machine disk file that is used as a template for creating grid node virtual machines.

`./vsphere/vsphere-primary-admin.ovf`
 	
`./vsphere/vsphere-primary-admin.mf`
	The Open Virtualization Format template file (`.ovf`) and manifest file (`.mf`) for deploying the primary Admin Node.

`./vsphere/vsphere-non-primary-admin.ovf`
 	
`./vsphere/vsphere-non-primary-admin.mf`
	The template file (`.ovf`) and manifest file (`.mf`) for deploying non-primary Admin Nodes.

`./vsphere/vsphere-archive.ovf`
 	
`./vsphere/vsphere-archive.mf`
	The template file (`.ovf`) and manifest file (`.mf`) for deploying Archive Nodes.

`./vsphere/vsphere-gateway.ovf`
 	
`./vsphere/vsphere-gateway.mf`
	The template file (`.ovf`) and manifest file (`.mf`) for deploying Gateway Nodes.

`./vsphere/vsphere-storage.ovf`
 	
`./vsphere/vsphere-storage.mf`
	The template file (`.ovf`) and manifest file (`.mf`) for deploying virtual machine-based Storage Nodes.

Deployment scripting tool	**Description**
`./vsphere/deploy-vsphere-ovftool.sh`	A Bash shell script used to automate the deployment of virtual grid nodes.

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

12. If you are recovering a StorageGRID appliance-based system, select the appropriate files.

Path and file name	Description
`./debs/storagegrid-webscale-images-version-SHA.deb`	DEB package for installing the StorageGRID node images on your appliances.

Path and file name	Description
`./debs/storagegrid-webscale-images-version-SHA.deb.md5`	MD5 checksum for the file `/debs/storagegridwebscale- images-version-SHA.deb`.

For appliance installation, these files are only required if you need to avoid network traffic. The appliance can download the required files from the primary Admin Node.

Select node recovery procedure

You must select the correct recovery procedure for the type of node that has failed.

Grid node	Recovery procedure
More than one Storage Node	Contact technical support. If more than one Storage Node has failed, technical support must assist with recovery to prevent database inconsistencies that could lead to data loss. A site recovery procedure might be required. How site recovery is performed by technical support
A single Storage Node	The Storage Node recovery procedure depends on the type and duration of the failure. Recover from Storage Node failures
Admin Node	The Admin Node procedure depends on whether you need to recover the primary Admin Node or a non-primary Admin Node. Recover from Admin Node failures
Gateway Node	Recover from Gateway Node failures.
Archive Node	Recover from Archive Node failures.

If a server that is hosting more than one grid node fails, you can recover the nodes in any order. However, if the failed server is hosting the primary Admin Node, you must recover that node first. Recovering the primary Admin Node first prevents other node recoveries from halting as they wait to contact the primary Admin Node.

Recover from Storage Node failures

Recover from Storage Node failures: Overview

The procedure for recovering a failed Storage Node depends on the type of failure and the type of Storage Node that has failed.

Use this table to select the recovery procedure for a failed Storage Node.

Issue	Action	Notes
• More than one Storage Node has failed. • A second Storage Node has failed less than 15 days after a Storage Node failure or recovery. This includes the case where a Storage Node fails while recovery of another Storage Node is still in progress.	Contact technical support.	Recovering more than one Storage Node (or more than one Storage Node within 15 days) might affect the integrity of the Cassandra database, which can cause data loss. Technical support can determine when it is safe to begin recovery of a second Storage Node. **Note**: If more than one Storage Node that contains the ADC service fails at a site, you lose any pending platform service requests for that site.
More than one Storage Node at a site has failed or an entire site has failed.	Contact technical support. It might be necessary to perform a site recovery procedure.	Technical support will assess your situation and develop a recovery plan. See How site recovery is performed by technical support.
A Storage Node has been offline for more than 15 days.	Recover Storage Node down more than 15 days	This procedure is required to ensure Cassandra database integrity.
An appliance Storage Node has failed.	Recover appliance Storage Node	The recovery procedure for appliance Storage Nodes is the same for all failures.
One or more storage volumes have failed, but the system drive is intact	Recover from storage volume failure where system drive is intact	This procedure is used for software-based Storage Nodes.
The system drive has failed.	Recover from system drive failure	The node replacement procedure depends on the deployment platform and on whether any storage volumes have also failed.

Some StorageGRID recovery procedures use Reaper to handle Cassandra repairs. Repairs occur automatically as soon as the related or required services have started. You might notice script output that mentions “reaper” or “Cassandra repair.” If you see an error message indicating the repair has failed, run the command indicated in the error message.

Recover Storage Node down more than 15 days

If a single Storage Node has been offline and not connected to other Storage Nodes for more than 15 days, you must rebuild Cassandra on the node.

Before you begin

- You have checked that a Storage Node decommissioning is not in progress, or you have paused the node decommission procedure. (In the Grid Manager, select **MAINTENANCE** > **Tasks** > **Decommission**.)
- You have checked that an expansion is not in progress. (In the Grid Manager, select **MAINTENANCE** > **Tasks** > **Expansion**.)

About this task

Storage Nodes have a Cassandra database that includes object metadata. If a Storage Node has not been able to communicate with other Storage Nodes for more than 15 days, StorageGRID assumes that node's Cassandra database is stale. The Storage Node can't rejoin the grid until Cassandra has been rebuilt using information from other Storage Nodes.

Use this procedure to rebuild Cassandra only if a single Storage Node is down. Contact technical support if additional Storage Nodes are offline or if Cassandra has been rebuilt on another Storage Node within the last 15 days; for example, Cassandra might have been rebuilt as part of the procedures to recover failed storage volumes or to recover a failed Storage Node.

If more than one Storage Node has failed (or is offline), contact technical support. Don't perform the following recovery procedure. Data loss could occur.

If this is the second Storage Node failure in less than 15 days after a Storage Node failure or recovery, contact technical support. Don't perform the following recovery procedure. Data loss could occur.

If more than one Storage Node at a site has failed, a site recovery procedure might be required. See How site recovery is performed by technical support.

Steps

1. If necessary, power on the Storage Node that needs to be recovered.
2. Log in to the grid node:
   a. Enter the following command: `ssh admin@grid_node_IP`
   b. Enter the password listed in the `Passwords.txt` file.
   c. Enter the following command to switch to root: `su -`
   d. Enter the password listed in the `Passwords.txt` file.

      When you are logged in as root, the prompt changes from `$` to `#`.+

      

      If you are unable to log in to the grid node, the system disk might not be intact. Go to the procedure for recovering from system drive failure.

3. Perform the following checks on the Storage Node:
   a. Issue this command: `nodetool status`

      The output should be `Connection refused`

   b. In the Grid Manager, select **SUPPORT** > **Tools** > **Grid topology**.
   c. Select ***Site*** > **Storage Node** > **SSM** > **Services**. Verify that the Cassandra service displays `Not Running`.

d. Select **Storage Node** > **SSM** > **Resources**. Verify that there is no error status in the Volumes section.

e. Issue this command: `grep -i Cassandra /var/local/log/servermanager.log`

You should see the following message in the output:

```
Cassandra not started because it has been offline for more than 15
day grace period - rebuild Cassandra
```

4. Issue this command, and monitor the script output: `check-cassandra-rebuild`
   - If the Cassandra service depending on volume 0 is running, you will be prompted to stop it. Enter: **y**

     If the Cassandra service is already stopped, you aren't prompted. The Cassandra service is stopped only for volume 0.

   - Review the warnings in the script. If none of them apply, confirm that you want to rebuild Cassandra. Enter: **y**

     Some StorageGRID recovery procedures use Reaper to handle Cassandra repairs. Repairs occur automatically as soon as the related or required services have started. You might notice script output that mentions "reaper" or "Cassandra repair." If you see an error message indicating the repair has failed, run the command indicated in the error message.

5. After the rebuild completes, perform the following checks:
   a. In the Grid Manager, select **SUPPORT** > **Tools** > **Grid topology**.
   b. Select ***Site*** > **recovered Storage Node** > **SSM** > **Services**.
   c. Confirm that all services are running.
   d. Select **DDS** > **Data Store**.
   e. Confirm that the **Data Store Status** is "Up" and the **Data Store State** is "Normal."

Recover appliance Storage Node

Recover appliance Storage Node: Workflow

The procedure for recovering a failed StorageGRID appliance Storage Node is the same whether you are recovering from the loss of the system drive or from the loss of storage volumes only.

As shown in the workflow diagram, you must prepare the appliance and reinstall software, configure the node to rejoin the grid, reformat storage, and restore object data.

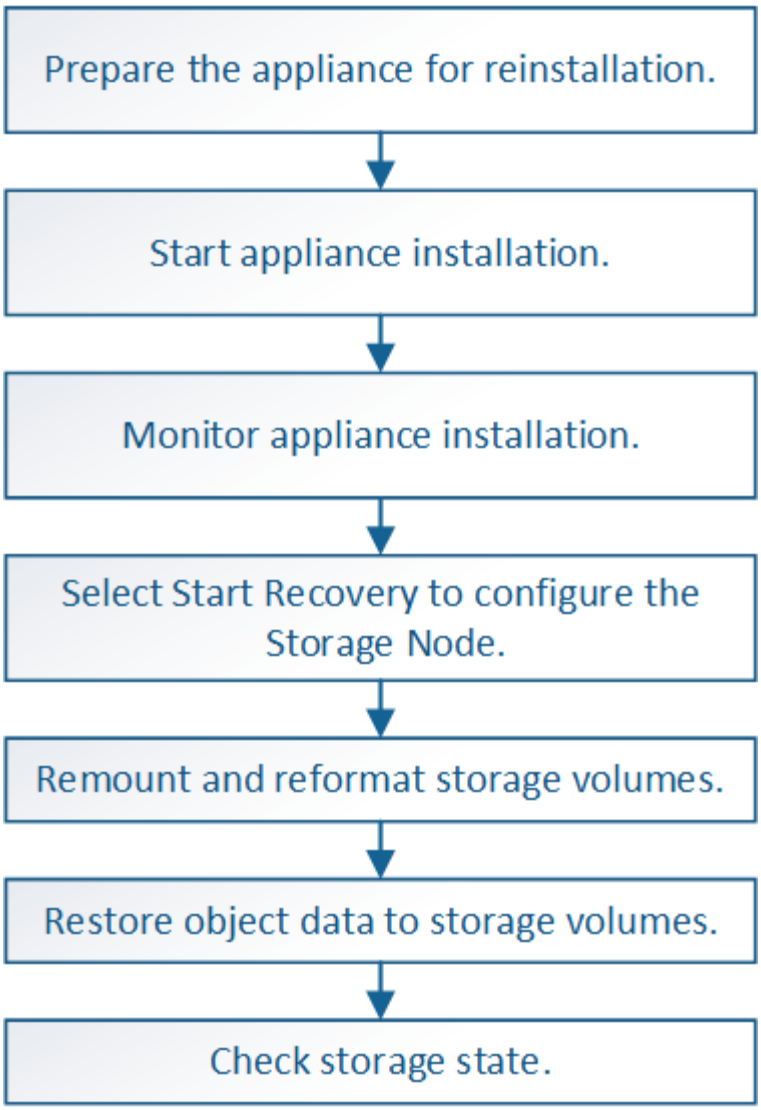

Cautions and notes for recovering appliance Storage Nodes

If more than one Storage Node has failed (or is offline), contact technical support. Don't perform the following recovery procedure. Data loss could occur.

If this is the second Storage Node failure in less than 15 days after a Storage Node failure or recovery, contact technical support. Rebuilding Cassandra on two or more Storage Nodes within 15 days can result in data loss.

If more than one Storage Node at a site has failed, a site recovery procedure might be required. See How site recovery is performed by technical support.

If ILM rules are configured to store only one replicated copy and the copy exists on a storage volume that has failed, you will not be able to recover the object.

If you encounter a Services: Status - Cassandra (SVST) alarm during recovery, see Recover failed storage volumes and rebuild Cassandra database. After Cassandra is rebuilt, alarms should clear. If alarms don't clear, contact technical support.

For hardware maintenance procedures, such as instructions for replacing a controller or reinstalling SANtricity OS, see the maintenance instructions for your storage appliance.

- SGF6112 storage appliances
- SG6000 storage appliances
- SG5700 storage appliances

Prepare appliance Storage Node for reinstallation

When recovering an appliance Storage Node, you must first prepare the appliance for reinstallation of StorageGRID software.

Steps

1. Log in to the failed Storage Node:
   a. Enter the following command: `ssh admin@grid_node_IP`
   b. Enter the password listed in the `Passwords.txt` file.
   c. Enter the following command to switch to root: `su -`
   d. Enter the password listed in the `Passwords.txt` file.

   When you are logged in as root, the prompt changes from `$` to `#`.
2. Prepare the appliance Storage Node for the installation of StorageGRID software. `sgareinstall`
3. When prompted to continue, enter: `y`

   The appliance reboots, and your SSH session ends. It usually takes about 5 minutes for the StorageGRID Appliance Installer to become available, although in some cases you might need to wait up to 30 minutes.

   Don't attempt to accelerate the reboot by cycling power or otherwise resetting the appliance. You might interrupt automatic BIOS, BMC, or other firmware upgrades.

   The StorageGRID appliance Storage Node is reset, and data on the Storage Node is no longer accessible. IP addresses configured during the original installation process should remain intact; however, it is recommended that you confirm this when the procedure completes.

   After executing the `sgareinstall` command, all StorageGRID-provisioned accounts, passwords, and SSH keys are removed, and new host keys are generated.

Start StorageGRID appliance installation

To install StorageGRID on an appliance Storage Node, you use the StorageGRID Appliance Installer, which is included on the appliance.

Before you begin

- The appliance has been installed in a rack, connected to your networks, and powered on.
- Network links and IP addresses have been configured for the appliance using the StorageGRID Appliance Installer.

- You know the IP address of the primary Admin Node for the StorageGRID grid.
- All Grid Network subnets listed on the IP Configuration page of the StorageGRID Appliance Installer have been defined in the Grid Network Subnet List on the primary Admin Node.
- You have completed these prerequisite tasks by following the installation instructions for your storage appliance. See Quick start for hardware installation.
- You are using a supported web browser.
- You know one of the IP addresses assigned to the compute controller in the appliance. You can use the IP address for the Admin Network (management port 1 on the controller), the Grid Network, or the Client Network.

About this task

To install StorageGRID on an appliance Storage Node:

- You specify or confirm the IP address of the primary Admin Node and the hostname (system name) of the node.
- You start the installation and wait as volumes are configured and the software is installed.
- Partway through the process, the installation pauses. To resume the installation, you must sign into the Grid Manager and configure the pending Storage Node as a replacement for the failed node.
- After you have configured the node, the appliance installation process completes, and the appliance is rebooted.

Steps

1. Open a browser and enter one of the IP addresses for the compute controller in the appliance.

   `https://Controller_IP:8443`

   The StorageGRID Appliance Installer Home page appears.

2. In the Primary Admin Node connection section, determine whether you need to specify the IP address for the primary Admin Node.

   The StorageGRID Appliance Installer can discover this IP address automatically, assuming the primary Admin Node, or at least one other grid node with ADMIN_IP configured, is present on the same subnet.

3. If this IP address is not shown or you need to change it, specify the address:

Option	Steps
Manual IP entry	a. Clear the **Enable Admin Node discovery** checkbox. b. Enter the IP address manually. c. Click **Save**. d. Wait while the connection state for the new IP address becomes “ready.”

Option	Steps
Automatic discovery of all connected primary Admin Nodes	a. Select the **Enable Admin Node discovery** checkbox. b. From the list of discovered IP addresses, select the primary Admin Node for the grid where this appliance Storage Node will be deployed. c. Click **Save**. d. Wait while the connection state for the new IP address becomes “ready.”

4. In the **Node Name** field, enter the same hostname (system name) that was used for the node you are recovering, and click **Save**.
5. In the Installation section, confirm that the current state is “Ready to start installation of *node name* into grid with Primary Admin Node *admin_ip*” and that the **Start Installation** button is enabled.

   If the **Start Installation** button is not enabled, you might need to change the network configuration or port settings. For instructions, see the maintenance instructions for your appliance.
6. From the StorageGRID Appliance Installer home page, click **Start Installation**.

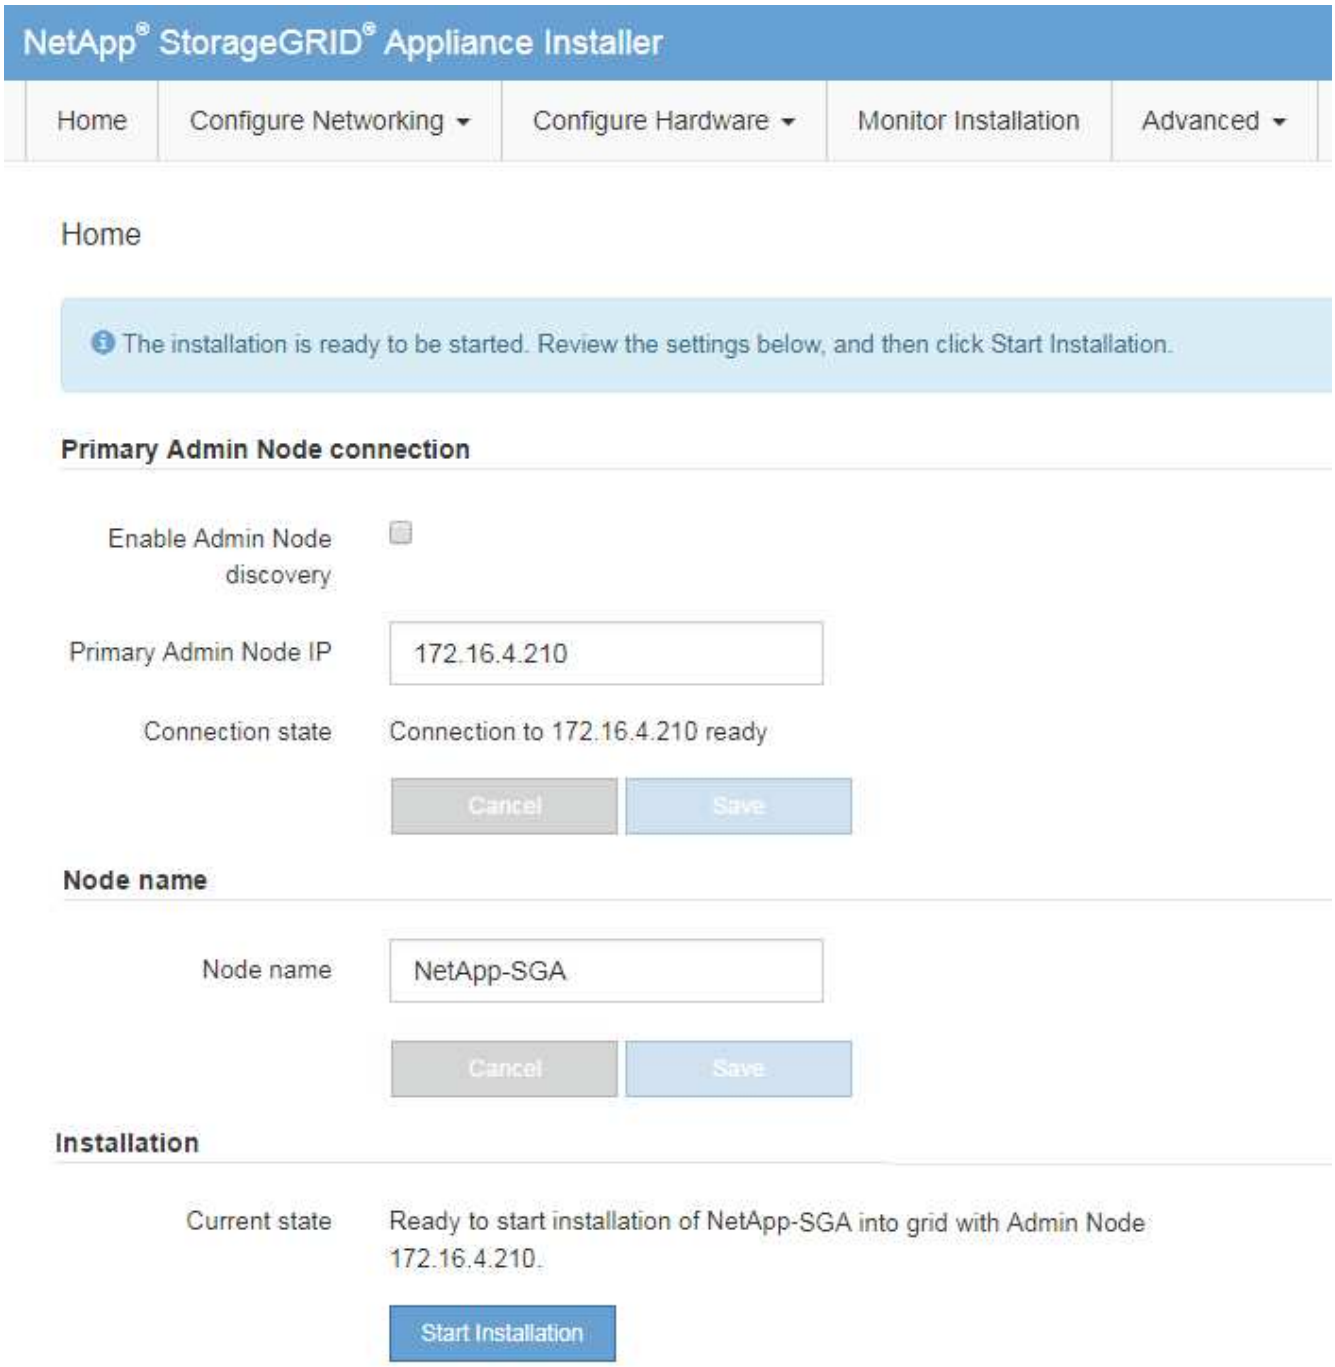

The Current state changes to “Installation is in progress,” and the Monitor Installation page is displayed.

If you need to access the Monitor Installation page manually, click **Monitor Installation** from the menu bar. See Monitor appliance installation.

Monitor StorageGRID appliance installation

The StorageGRID Appliance Installer provides status until installation is complete. When the software installation is complete, the appliance is rebooted.

Steps

1. To monitor the installation progress, click **Monitor Installation** from the menu bar.

The Monitor Installation page shows the installation progress.

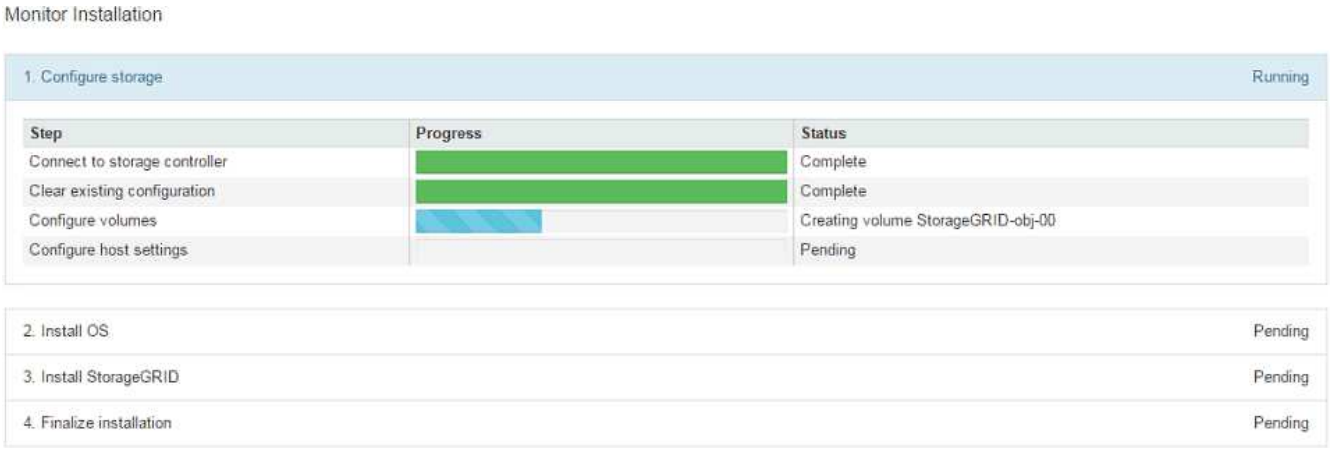

The blue status bar indicates which task is currently in progress. Green status bars indicate tasks that have completed successfully.

> The installer ensures that tasks completed in a previous install aren't re-run. If you are re-running an installation, any tasks that don't need to be re-run are shown with a green status bar and a status of "Skipped."

2. Review the progress of first two installation stages.
   - **1. Configure storage**

     During this stage, the installer connects to the storage controller, clears any existing configuration, communicates with SANtricity OS to configure volumes, and configures host settings.

   - **2. Install OS**

     During this stage, the installer copies the base operating system image for StorageGRID to the appliance.
3. Continue monitoring the installation progress until the **Install StorageGRID** stage pauses and a message appears on the embedded console prompting you to approve this node on the Admin Node using the Grid Manager.

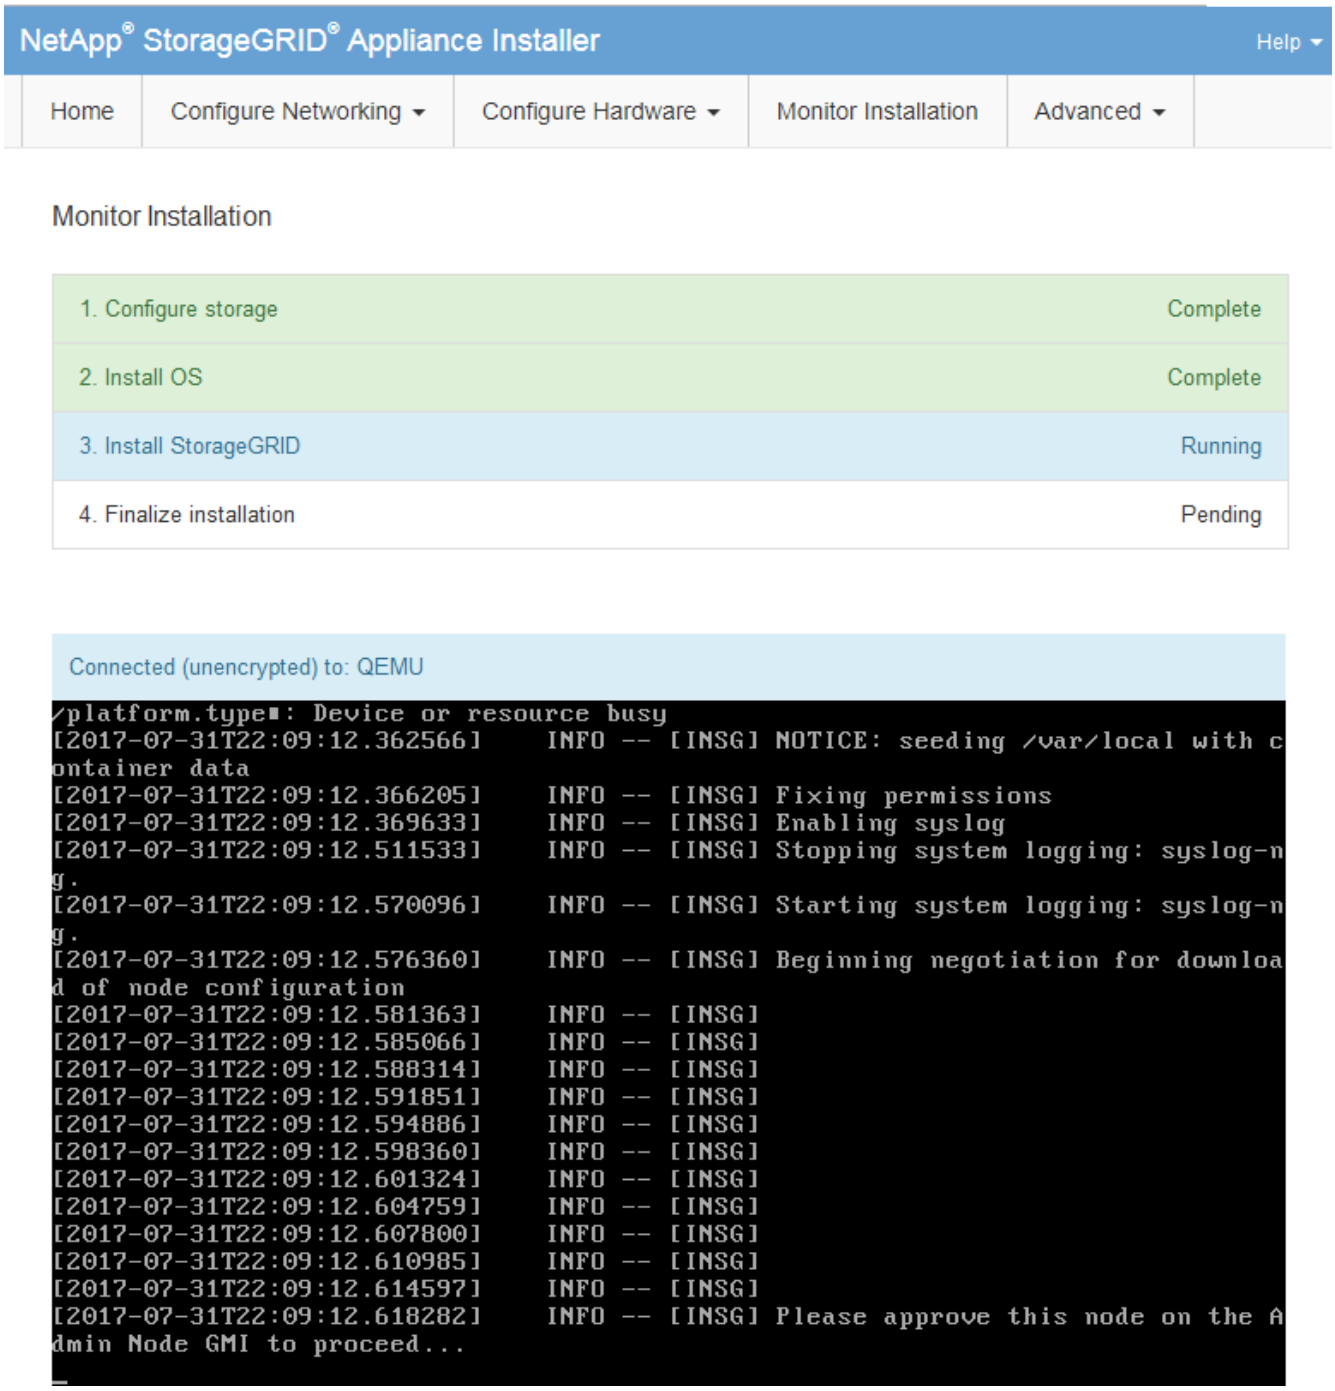

4. Go to Select Start Recovery to configure appliance Storage Node.

Select Start Recovery to configure appliance Storage Node

You must select Start Recovery in the Grid Manager to configure an appliance Storage Node as a replacement for the failed node.

Before you begin

- You are signed in to the Grid Manager using a supported web browser.
- You have the Maintenance or Root access permission.
- You have the provisioning passphrase.

- You have deployed a recovery appliance Storage Node.
- You have the start date of any repair jobs for erasure-coded data.
- You have verified that the Storage Node has not been rebuilt within the last 15 days.

Steps

1. From the Grid Manager, select **MAINTENANCE** > **Tasks** > **Recovery**.
2. Select the grid node you want to recover in the Pending Nodes list.

   Nodes appear in the list after they fail, but you can't select a node until it has been reinstalled and is ready for recovery.

3. Enter the **Provisioning Passphrase**.
4. Click **Start Recovery**.

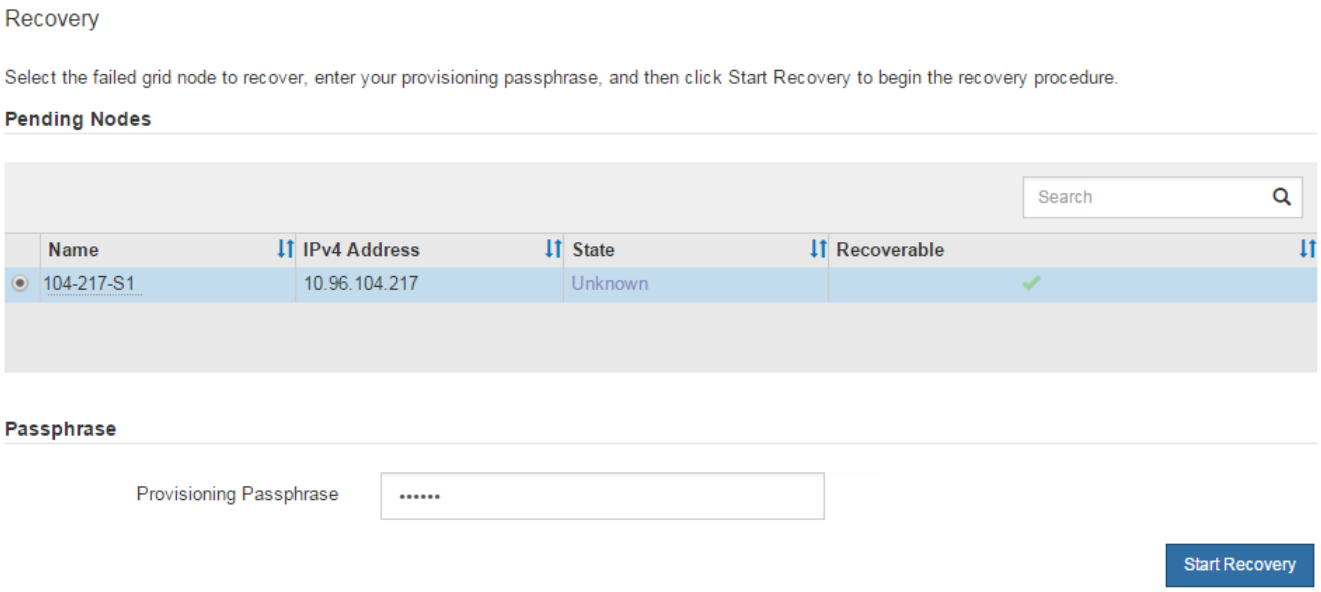

5. Monitor the progress of the recovery in the Recovering Grid Node table.

   When the grid node reaches the "Waiting for Manual Steps" stage, go to the next topic and perform the manual steps to remount and reformat appliance storage volumes.

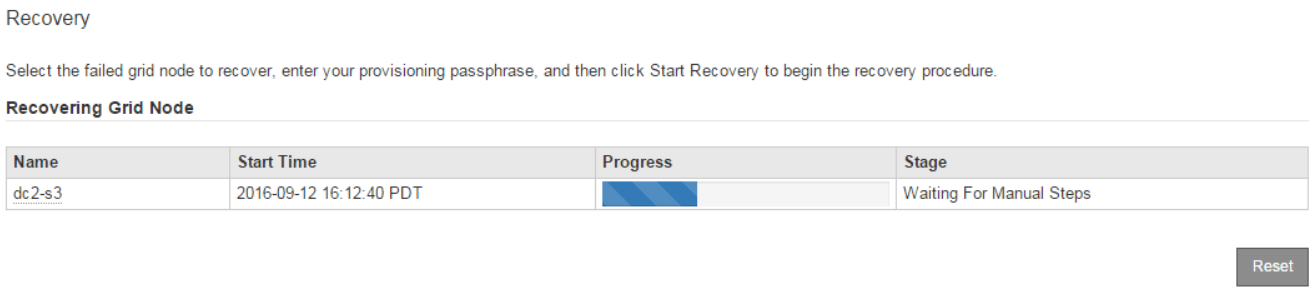

At any point during the recovery, you can click **Reset** to start a new recovery. A dialog box appears, indicating that the node will be left in an indeterminate state if you reset the procedure.

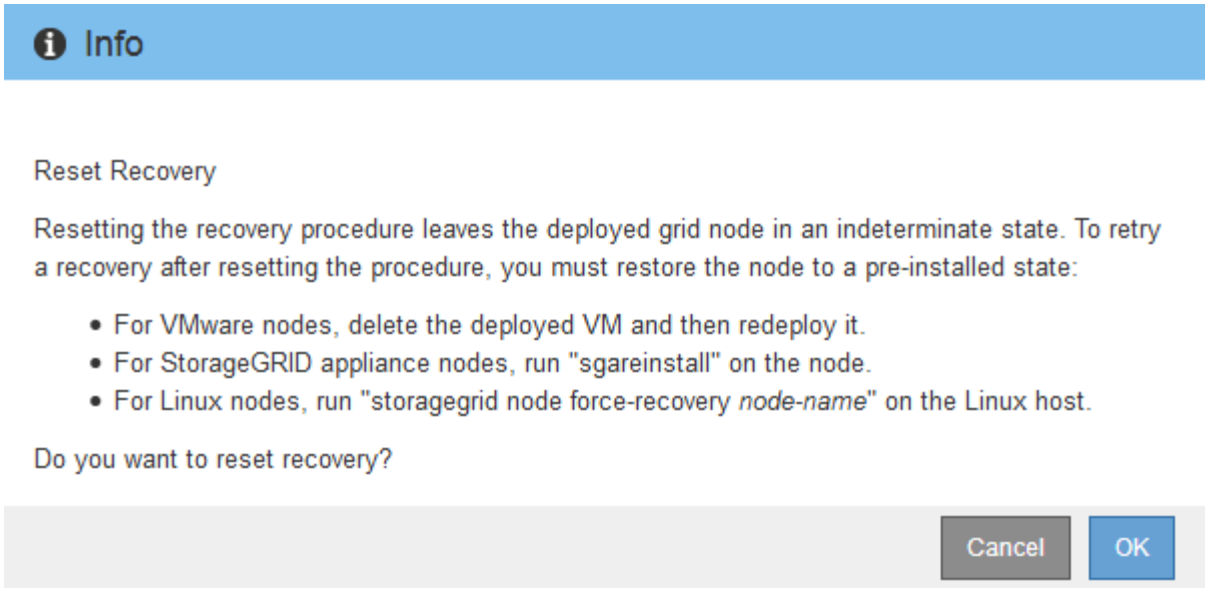

If you want to retry the recovery after resetting the procedure, you must restore the appliance node to a pre-installed state by running `sgareinstall` on the node.

Remount and reformat appliance storage volumes (manual steps)

You must manually run two scripts to remount preserved storage volumes and reformat any failed storage volumes. The first script remounts volumes that are properly formatted as StorageGRID storage volumes. The second script reformats any unmounted volumes, rebuilds the Cassandra database, if needed, and starts services.

Before you begin

- You have already replaced the hardware for any failed storage volumes that you know require replacement.

  Running the `sn-remount-volumes` script might help you identify additional failed storage volumes.

- You have checked that a Storage Node decommissioning is not in progress, or you have paused the node decommission procedure. (In the Grid Manager, select **MAINTENANCE** > **Tasks** > **Decommission**.)
- You have checked that an expansion is not in progress. (In the Grid Manager, select **MAINTENANCE** > **Tasks** > **Expansion**.)

> Contact technical support if more than one Storage Node is offline or if a Storage Node in this grid has been rebuilt in the last 15 days. Don't run the `sn-recovery-postinstall.sh` script. Rebuilding Cassandra on two or more Storage Nodes within 15 days of each other might result in data loss.

About this task

To complete this procedure, you perform these high-level tasks:

- Log in to the recovered Storage Node.
- Run the `sn-remount-volumes` script to remount properly formatted storage volumes. When this script runs, it does the following:

- Mounts and unmounts each storage volume to replay the XFS journal.
- Performs an XFS file consistency check.
- If the file system is consistent, determines if the storage volume is a properly formatted StorageGRID storage volume.
- If the storage volume is properly formatted, remounts the storage volume. Any existing data on the volume remains intact.

- Review the script output and resolve any issues.
- Run the `sn-recovery-postinstall.sh` script. When this script runs, it does the following.

> Don't reboot a Storage Node during recovery before running `sn-recovery-postinstall.sh` (step 4) to reformat the failed storage volumes and restore object metadata. Rebooting the Storage Node before `sn-recovery-postinstall.sh` completes causes errors for services that attempt to start and causes StorageGRID appliance nodes to exit maintenance mode.

  - Reformats any storage volumes that the `sn-remount-volumes` script could not mount or that were found to be improperly formatted.

    > If a storage volume is reformatted, any data on that volume is lost. You must perform an additional procedure to restore object data from other locations in the grid, assuming that ILM rules were configured to store more than one object copy.

  - Rebuilds the Cassandra database on the node, if needed.
  - Starts the services on the Storage Node.

Steps

1. Log in to the recovered Storage Node:
   a. Enter the following command: `ssh admin@grid_node_IP`
   b. Enter the password listed in the `Passwords.txt` file.
   c. Enter the following command to switch to root: `su -`
   d. Enter the password listed in the `Passwords.txt` file.

   When you are logged in as root, the prompt changes from `$` to `#`.

2. Run the first script to remount any properly formatted storage volumes.

   > If all storage volumes are new and need to be formatted, or if all storage volumes have failed, you can skip this step and run the second script to reformat all unmounted storage volumes.

   a. Run the script: `sn-remount-volumes`

      This script might take hours to run on storage volumes that contain data.

   b. As the script runs, review the output and answer any prompts.

As required, you can use the `tail -f` command to monitor the contents of the script's log file (`/var/local/log/sn-remount-volumes.log`) . The log file contains more detailed information than the command line output.

```
root@SG:~ # sn-remount-volumes
The configured LDR noid is 12632740

====== Device /dev/sdb ======
Mount and unmount device /dev/sdb and checking file system
consistency:
The device is consistent.
Check rangedb structure on device /dev/sdb:
Mount device /dev/sdb to /tmp/sdb-654321 with rangedb mount options
This device has all rangedb directories.
Found LDR node id 12632740, volume number 0 in the volID file
Attempting to remount /dev/sdb
Device /dev/sdb remounted successfully

====== Device /dev/sdc ======
Mount and unmount device /dev/sdc and checking file system
consistency:
Error: File system consistency check retry failed on device /dev/sdc.
You can see the diagnosis information in the /var/local/log/sn-
remount-volumes.log.

This volume could be new or damaged. If you run sn-recovery-
postinstall.sh, this volume and any data on this volume will be
deleted. If you only had two copies of object data, you will
temporarily have only a single copy.
StorageGRID Webscale will attempt to restore data redundancy by
making additional replicated copies or EC fragments, according to the
rules in the active ILM policy.

Don't continue to the next step if you believe that the data
remaining on this volume can't be rebuilt from elsewhere in the grid
(for example, if your ILM policy uses a rule that makes only one copy
or if volumes have failed on multiple nodes). Instead, contact
support to determine how to recover your data.

====== Device /dev/sdd ======
Mount and unmount device /dev/sdd and checking file system
consistency:
Failed to mount device /dev/sdd
This device could be an uninitialized disk or has corrupted
superblock.
File system check might take a long time. Do you want to continue? (y
```

```
or n) [y/N]? y

Error: File system consistency check retry failed on device /dev/sdd.
You can see the diagnosis information in the /var/local/log/sn-
remount-volumes.log.

This volume could be new or damaged. If you run sn-recovery-
postinstall.sh, this volume and any data on this volume will be
deleted. If you only had two copies of object data, you will
temporarily have only a single copy.
StorageGRID Webscale will attempt to restore data redundancy by
making additional replicated copies or EC fragments, according to the
rules in the active ILM policy.

Don't continue to the next step if you believe that the data
remaining on this volume can't be rebuilt from elsewhere in the grid
(for example, if your ILM policy uses a rule that makes only one copy
or if volumes have failed on multiple nodes). Instead, contact
support to determine how to recover your data.

====== Device /dev/sde ======
Mount and unmount device /dev/sde and checking file system
consistency:
The device is consistent.
Check rangedb structure on device /dev/sde:
Mount device /dev/sde to /tmp/sde-654321 with rangedb mount options
This device has all rangedb directories.
Found LDR node id 12000078, volume number 9 in the volID file
Error: This volume does not belong to this node. Fix the attached
volume and re-run this script.
```

In the example output, one storage volume was remounted successfully and three storage volumes had errors.

- `/dev/sdb` passed the XFS file system consistency check and had a valid volume structure, so it was remounted successfully. Data on devices that are remounted by the script is preserved.
- `/dev/sdc` failed the XFS file system consistency check because the storage volume was new or corrupt.
- `/dev/sdd` could not be mounted because the disk was not initialized or the disk's superblock was corrupted. When the script can't mount a storage volume, it asks if you want to run the file system consistency check.
  - If the storage volume is attached to a new disk, answer **N** to the prompt. You don't need check the file system on a new disk.
  - If the storage volume is attached to an existing disk, answer **Y** to the prompt. You can use the results of the file system check to determine the source of the corruption. The results are saved in the `/var/local/log/sn-remount-volumes.log` log file.

- `/dev/sde` passed the XFS file system consistency check and had a valid volume structure; however, the LDR node ID in the `volID` file did not match the ID for this Storage Node (the `configured LDR noid` displayed at the top). This message indicates that this volume belongs to another Storage Node.

3. Review the script output and resolve any issues.

> If a storage volume failed the XFS file system consistency check or could not be mounted, carefully review the error messages in the output. You must understand the implications of running the `sn-recovery-postinstall.sh` script on these volumes.

a. Check to make sure that the results include an entry for all of the volumes you expected. If any volumes aren't listed, rerun the script.

b. Review the messages for all mounted devices. Make sure there are no errors indicating that a storage volume does not belong to this Storage Node.

In the example, the output for /dev/sde includes the following error message:

```
Error: This volume does not belong to this node. Fix the attached
volume and re-run this script.
```

> If a storage volume is reported as belonging to another Storage Node, contact technical support. If you run the `sn-recovery-postinstall.sh` script, the storage volume will be reformatted, which might cause data loss.

c. If any storage devices could not be mounted, make a note of the device name, and repair or replace the device.

> You must repair or replace any storage devices that could not be mounted.

You will use the device name to look up the volume ID, which is required input when you run the `repair-data` script to restore object data to the volume (the next procedure).

d. After repairing or replacing all unmountable devices, run the `sn-remount-volumes` script again to confirm that all storage volumes that can be remounted have been remounted.

> If a storage volume can't be mounted or is improperly formatted, and you continue to the next step, the volume and any data on the volume will be deleted. If you had two copies of object data, you will have only a single copy until you complete the next procedure (restoring object data).

> Don't run the `sn-recovery-postinstall.sh` script if you believe that the data remaining on a failed storage volume can't be rebuilt from elsewhere in the grid (for example, if your ILM policy uses a rule that makes only one copy or if volumes have failed on multiple nodes). Instead, contact technical support to determine how to recover your data.

4. Run the `sn-recovery-postinstall.sh` script: `sn-recovery-postinstall.sh`

This script reformats any storage volumes that could not be mounted or that were found to be improperly

formatted; rebuilds the Cassandra database on the node, if needed; and starts the services on the Storage Node.

Be aware of the following:

- The script might take hours to run.
- In general, you should leave the SSH session alone while the script is running.
- Don't press **Ctrl+C** while the SSH session is active.
- The script will run in the background if a network disruption occurs and terminates the SSH session, but you can view the progress from the Recovery page.
- If the Storage Node uses the RSM service, the script might appear to stall for 5 minutes as node services are restarted. This 5-minute delay is expected whenever the RSM service boots for the first time.

> The RSM service is present on Storage Nodes that include the ADC service.

> Some StorageGRID recovery procedures use Reaper to handle Cassandra repairs. Repairs occur automatically as soon as the related or required services have started. You might notice script output that mentions "reaper" or "Cassandra repair." If you see an error message indicating the repair has failed, run the command indicated in the error message.

5. As the `sn-recovery-postinstall.sh` script runs, monitor the Recovery page in the Grid Manager.

   The Progress bar and the Stage column on the Recovery page provide a high-level status of the `sn-recovery-postinstall.sh` script.

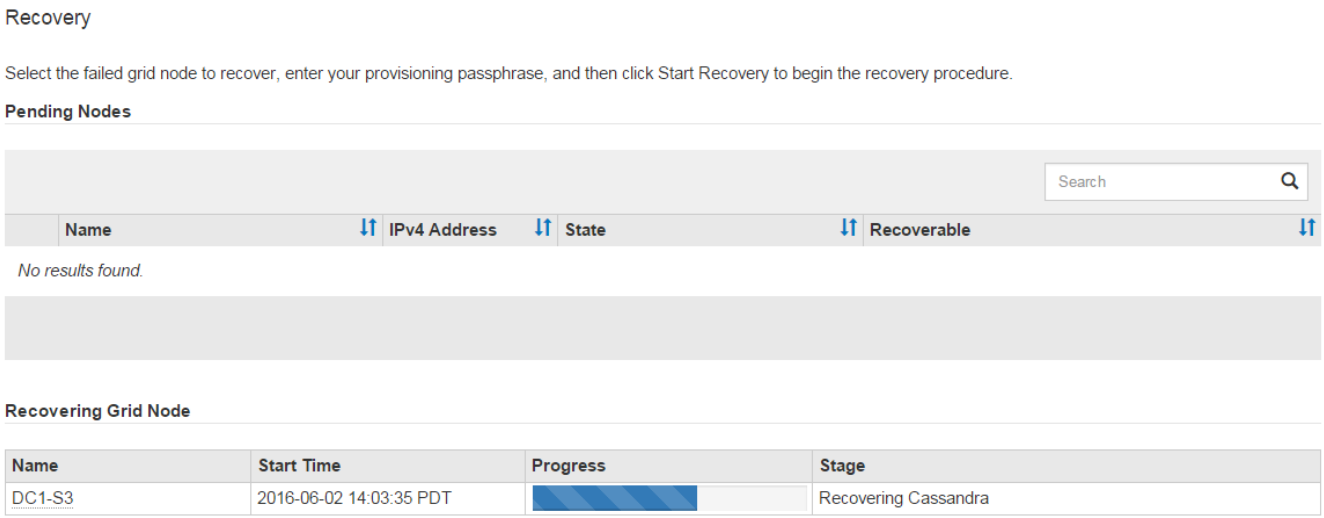

6. After the `sn-recovery-postinstall.sh` script has started services on the node, you can restore object data to any storage volumes that were formatted by the script.

   The script asks if you want to restore object data manually.

   - In most cases, you should restore object data using Grid Manager. Answer `n` to use the Grid Manager.
   - In rare cases, such as when instructed by technical support, or when you know that the replacement node has fewer volumes available for object storage than the original node, you must restore object data manually using the `repair-data` script. If one of these cases applies, answer `y`.

If you answer `y` to restore object data manually:

- You aren't able to restore object data using Grid Manager.
- You can monitor the progress of manual restoration jobs using Grid Manager.

Restore object data to storage volume for appliance

After recovering storage volumes for the appliance Storage Node, you can restore the replicated or erasure-coded object data that was lost when the Storage Node failed.

Which procedure should I use?

Whenever possible, restore object data using the **Volume restoration** page in the Grid Manager.

- If the volumes are listed at **MAINTENANCE** > **Volume restoration** > **Nodes to restore**, restore object data using the Volume restoration page in the Grid Manager.
- If the volumes aren't listed at **MAINTENANCE** > **Volume restoration** > **Nodes to restore**, follow the steps below for using the `repair-data` script to restore object data.

  If the recovered Storage Node contains fewer volumes than the node it is replacing, you must use the `repair-data` script.

Use the `repair-data` script to restore object data

Before you begin

- You must have confirmed that the recovered Storage Node has a Connection State of **Connected** on the **NODES** > **Overview** tab in the Grid Manager.

About this task

Object data can be restored from other Storage Nodes, an Archive Node, or a Cloud Storage Pool, assuming that the grid's ILM rules were configured such that object copies are available.

Note the following:

- If an ILM rule was configured to store only one replicated copy and that copy existed on a storage volume that failed, you will not be able to recover the object.
- If the only remaining copy of an object is in a Cloud Storage Pool, StorageGRID must issue multiple requests to the Cloud Storage Pool endpoint to restore object data. Before performing this procedure, contact technical support for help in estimating the recovery time frame and the associated costs.
- If the only remaining copy of an object is on an Archive Node, object data is retrieved from the Archive Node. Restoring object data to a Storage Node from an Archive Node takes longer than restoring copies from other Storage Nodes because of the latency associated with retrievals from external archival storage systems.

About the `repair-data` script

To restore object data, you run the `repair-data` script. This script begins the process of restoring object data and works with ILM scanning to ensure that ILM rules are met.

Select **Replicated data** or **Erasure-coded (EC) data** below to learn the different options for the `repair-data` script, based on whether you are restoring replicated data or erasure-coded data. If you need to restore both types of data, you must run both sets of commands.

For more information about the `repair-data` script, enter `repair-data --help` from the command line of the primary Admin Node.

Replicated data

Two commands are available for restoring replicated data, based on whether you need to repair the entire node or only certain volumes on the node:

```
repair-data start-replicated-node-repair

repair-data start-replicated-volume-repair
```

You can track repairs of replicated data with this command:

```
repair-data show-replicated-repair-status
```

Erasure coded (EC) data

Two commands are available for restoring erasure-coded data, based on whether you need to repair the entire node or only certain volumes on the node:

```
repair-data start-ec-node-repair

repair-data start-ec-volume-repair
```

You can track repairs of erasure-coded data with this command:

```
repair-data show-ec-repair-status
```

Repairs of erasure-coded data can begin while some Storage Nodes are offline. However, if all erasure-coded data can't be accounted for, the repair can't be completed. Repair will complete after all nodes are available.

The EC repair job temporarily reserves a large amount of storage. Storage alerts might be triggered, but will resolve when the repair is complete. If there is not enough storage for the reservation, the EC repair job will fail. Storage reservations are released when the EC repair job completes, whether the job failed or succeeded.

Find hostname for Storage Node

1. Log in to the primary Admin Node:
   a. Enter the following command: `ssh admin@primary_Admin_Node_IP`
   b. Enter the password listed in the `Passwords.txt` file.
   c. Enter the following command to switch to root: `su -`
   d. Enter the password listed in the `Passwords.txt` file.

When you are logged in as root, the prompt changes from `$` to `#`.

2. Use the `/etc/hosts` file to find the hostname of the Storage Node for the restored storage volumes. To see a list of all nodes in the grid, enter the following: `cat /etc/hosts`.

Repair data if all volumes have failed

If all storage volumes have failed, repair the entire node. Follow the instructions for **replicated data**, **erasure-coded (EC) data**, or both, based on whether you use replicated data, erasure-coded (EC) data, or both.

If only some volumes have failed, go to Repair data if only some volumes have failed.

> You can't run `repair-data` operations for more than one node at the same time. To recover multiple nodes, contact technical support.

Replicated data

If your grid includes replicated data, use the `repair-data start-replicated-node-repair` command with the `--nodes` option, where `--nodes` is the hostname (system name), to repair the entire Storage Node.

This command repairs the replicated data on a Storage Node named SG-DC-SN3:

```
repair-data start-replicated-node-repair --nodes SG-DC-SN3
```

> As object data is restored, the **Objects Lost** alert is triggered if the StorageGRID system can't locate replicated object data. Alerts might be triggered on Storage Nodes throughout the system. You should determine the cause of the loss and if recovery is possible. See Investigate lost objects.

Erasure coded (EC) data

If your grid contains erasure-coded data, use the `repair-data start-ec-node-repair` command with the `--nodes` option, where `--nodes` is the hostname (system name), to repair the entire Storage Node.

This command repairs the erasure-coded data on a Storage Node named SG-DC-SN3:

```
repair-data start-ec-node-repair --nodes SG-DC-SN3
```

The operation returns a unique `repair ID` that identifies this `repair_data` operation. Use this `repair ID` to track the progress and result of the `repair_data` operation. No other feedback is returned as the recovery process completes.

> Repairs of erasure-coded data can begin while some Storage Nodes are offline. Repair will complete after all nodes are available.

Repair data if only some volumes have failed

If only some of the volumes have failed, repair the affected volumes. Follow the instructions for **replicated data**, **erasure-coded (EC) data**, or both, based on whether you use replicated data, erasure-coded (EC) data, or both.

If all volumes have failed, go to Repair data if all volumes have failed.

Enter the volume IDs in hexadecimal. For example, `0000` is the first volume and `000F` is the sixteenth volume. You can specify one volume, a range of volumes, or multiple volumes that aren't in a sequence.

All the volumes must be on the same Storage Node. If you need to restore volumes for more than one Storage Node, contact technical support.

Replicated data

If your grid contains replicated data, use the `start-replicated-volume-repair` command with the `--nodes` option to identify the node (where `--nodes` is the hostname of the node). Then add either the `--volumes` or `--volume-range` option, as shown in the following examples.

Single volume: This command restores replicated data to volume `0002` on a Storage Node named SG-DC-SN3:

```
repair-data start-replicated-volume-repair --nodes SG-DC-SN3 --volumes 0002
```

Range of volumes: This command restores replicated data to all volumes in the range `0003` to `0009` on a Storage Node named SG-DC-SN3:

```
repair-data start-replicated-volume-repair --nodes SG-DC-SN3 --volume-range
0003,0009
```

Multiple volumes not in a sequence: This command restores replicated data to volumes `0001`, `0005`, and `0008` on a Storage Node named SG-DC-SN3:

```
repair-data start-replicated-volume-repair --nodes SG-DC-SN3 --volumes
0001,0005,0008
```

> As object data is restored, the **Objects Lost** alert is triggered if the StorageGRID system can't locate replicated object data. Alerts might be triggered on Storage Nodes throughout the system. Note the alert description and recommended actions to determine the cause of the loss and if recovery is possible.

Erasure coded (EC) data

If your grid contains erasure-coded data, use the `start-ec-volume-repair` command with the `--nodes` option to identify the node (where `--nodes` is the hostname of the node). Then add either the `--volumes` or `--volume-range` option, as shown in the following examples.

Single volume: This command restores erasure-coded data to volume `0007` on a Storage Node named SG-DC-SN3:

```
repair-data start-ec-volume-repair --nodes SG-DC-SN3 --volumes 0007
```

Range of volumes: This command restores erasure-coded data to all volumes in the range `0004` to `0006` on a Storage Node named SG-DC-SN3:

```
repair-data start-ec-volume-repair --nodes SG-DC-SN3 --volume-range 0004,0006
```

Multiple volumes not in a sequence: This command restores erasure-coded data to volumes `000A`, `000C`, and `000E` on a Storage Node named SG-DC-SN3:

```
repair-data start-ec-volume-repair --nodes SG-DC-SN3 --volumes 000A,000C,000E
```

The `repair-data` operation returns a unique `repair ID` that identifies this `repair_data` operation. Use this `repair ID` to track the progress and result of the `repair_data` operation. No other feedback is returned as the recovery process completes.

Repairs of erasure-coded data can begin while some Storage Nodes are offline. Repair will complete after all nodes are available.

Monitor repairs

Monitor the status of the repair jobs, based on whether you use **replicated data**, **erasure-coded (EC) data**, or both.

You can also monitor the status of volume restoration jobs in process and view a history of restoration jobs completed in Grid Manager.

Replicated data

- To get an estimated percent completion for the replicated repair, add the `show-replicated-repair-status` option to the repair-data command.

  ```
  repair-data show-replicated-repair-status
  ```

- To determine if repairs are complete:
  1. Select **NODES** > ***Storage Node being repaired*** > **ILM**.
  2. Review the attributes in the Evaluation section. When repairs are complete, the **Awaiting - All** attribute indicates 0 objects.
- To monitor the repair in more detail:
  1. Select **SUPPORT** > **Tools** > **Grid topology**.
  2. Select ***grid*** > ***Storage Node being repaired*** > **LDR** > **Data Store**.
  3. Use a combination of the following attributes to determine, as well as possible, if replicated repairs are complete.

     Cassandra inconsistencies might be present, and failed repairs aren't tracked.

     - **Repairs Attempted (XRPA)**: Use this attribute to track the progress of replicated repairs. This attribute increases each time a Storage Node tries to repair a high-risk object. When this attribute does not increase for a period longer than the current scan period (provided by the **Scan Period — Estimated** attribute), it means that ILM scanning found no high-risk objects that need to be repaired on any nodes.

       High-risk objects are objects that are at risk of being completely lost. This does not include objects that don't satisfy their ILM configuration.

     - **Scan Period — Estimated (XSCM)**: Use this attribute to estimate when a policy change will be applied to previously ingested objects. If the **Repairs Attempted** attribute does not increase for a period longer than the current scan period, it is probable that replicated repairs are done. Note that the scan period can change. The **Scan Period — Estimated (XSCM)** attribute applies to the entire grid and is the maximum of all node scan periods. You can query the **Scan Period — Estimated** attribute history for the grid to determine an appropriate time frame.

Erasure coded (EC) data

To monitor the repair of erasure-coded data and retry any requests that might have failed:

1. Determine the status of erasure-coded data repairs:
   - Select **SUPPORT** > **Tools** > **Metrics** to view the estimated time to completion and the completion percentage for the current job. Then, select **EC Overview** in the Grafana section. Look at the **Grid EC Job Estimated Time to Completion** and **Grid EC Job Percentage Completed** dashboards.
   - Use this command to see the status of a specific `repair-data` operation:

     ```
     repair-data show-ec-repair-status --repair-id repair ID
     ```

   - Use this command to list all repairs:

```
repair-data show-ec-repair-status
```

The output lists information, including `repair ID`, for all previously and currently running repairs.

2. If the output shows that the repair operation failed, use the `--repair-id` option to retry the repair.

   This command retries a failed node repair, using the repair ID 6949309319275667690:

   ```
   repair-data start-ec-node-repair --repair-id 6949309319275667690
   ```

   This command retries a failed volume repair, using the repair ID 6949309319275667690:

   ```
   repair-data start-ec-volume-repair --repair-id 6949309319275667690
   ```

Check storage state after recovering appliance Storage Node

After recovering an appliance Storage Node, you must verify that the desired state of the appliance Storage Node is set to online and ensure that the state will be online by default whenever the Storage Node server is restarted.

Before you begin

- You are signed in to the Grid Manager using a supported web browser.
- The Storage Node has been recovered, and data recovery is complete.

Steps

1. Select **SUPPORT** > **Tools** > **Grid topology**.
2. Check the values of **Recovered Storage Node** > **LDR** > **Storage** > **Storage State — Desired** and **Storage State — Current**.

   The value of both attributes should be Online.

3. If the Storage State — Desired is set to Read-only, complete the following steps:
   a. Click the **Configuration** tab.
   b. From the **Storage State — Desired** drop-down list, select **Online**.
   c. Click **Apply Changes**.
   d. Click the **Overview** tab and confirm that the values of **Storage State — Desired** and **Storage State — Current** are updated to Online.

Recover from storage volume failure where system drive is intact

Recover from storage volume failure where system drive is intact: Workflow

You must complete a series of tasks to recover a software-based Storage Node where one or more storage volumes on the Storage Node have failed, but the system drive is intact. If only storage volumes have failed, the Storage Node is still available to the StorageGRID system.

This recovery procedure applies to software-based Storage Nodes only. If storage volumes have failed on an appliance Storage Node, use the appliance procedure instead: Recover appliance Storage Node.

As shown in the workflow diagram, you must identify and unmount failed storage volumes, recover the volumes, rebuild the Cassandra database, and restore object data.

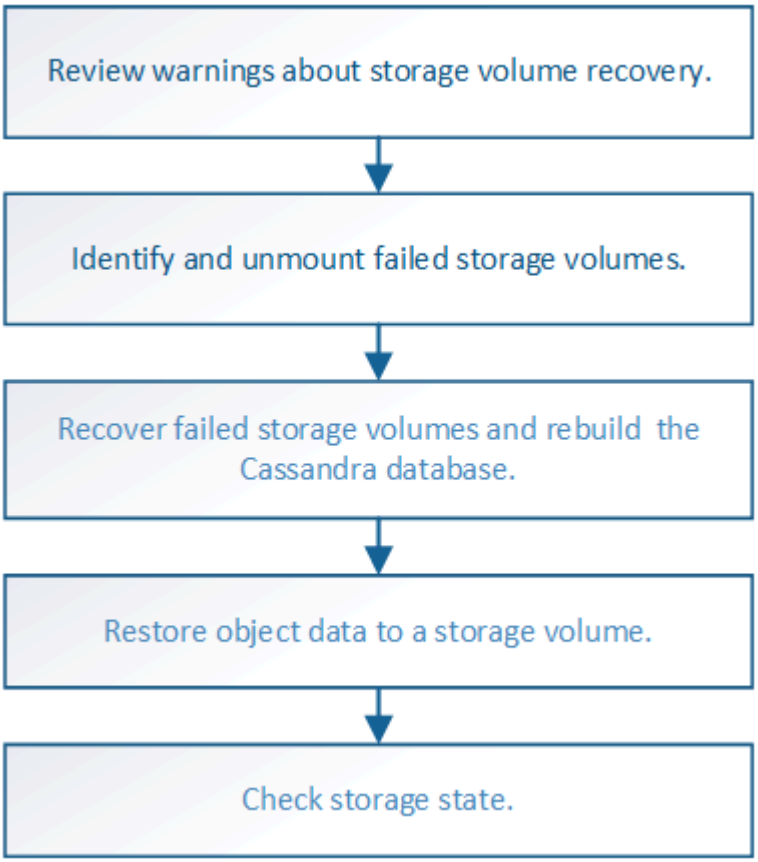

Review warnings about storage volume recovery

Before recovering failed storage volumes for a Storage Node, you must review the following warnings.

The storage volumes (or rangedbs) in a Storage Node are identified by a hexadecimal number, which is known as the volume ID. For example, 0000 is the first volume and 000F is the sixteenth volume. The first object store (volume 0) on each Storage Node uses up to 4 TB of space for object metadata and Cassandra database operations; any remaining space on that volume is used for object data. All other storage volumes are used exclusively for object data.

If volume 0 fails and needs to be recovered, the Cassandra database might be rebuilt as part of the volume recovery procedure. Cassandra might also be rebuilt in the following circumstances:

- A Storage Node is brought back online after having been offline for more than 15 days.
- The system drive and one or more storage volumes fails and is recovered.

When Cassandra is rebuilt, the system uses information from other Storage Nodes. If too many Storage Nodes are offline, some Cassandra data might not be available. If Cassandra has been rebuilt recently, Cassandra data might not yet be consistent across the grid. Data loss can occur if Cassandra is rebuilt when too many

Storage Nodes are offline or if two or more Storage Nodes are rebuilt within 15 days of each other.

> If more than one Storage Node has failed (or is offline), contact technical support. Don't perform the following recovery procedure. Data loss could occur.

> If this is the second Storage Node failure in less than 15 days after a Storage Node failure or recovery, contact technical support. Rebuilding Cassandra on two or more Storage Nodes within 15 days can result in data loss.

> If more than one Storage Node at a site has failed, a site recovery procedure might be required. See How site recovery is performed by technical support.

> If ILM rules are configured to store only one replicated copy and the copy exists on a storage volume that has failed, you will not be able to recover the object.

> If you encounter a Services: Status - Cassandra (SVST) alarm during recovery, see Recover failed storage volumes and rebuild Cassandra database. After Cassandra is rebuilt, alarms should clear. If alarms don't clear, contact technical support.

Related information

Warnings and considerations for grid node recovery

Identify and unmount failed storage volumes

When recovering a Storage Node with failed storage volumes, you must identify and unmount the failed volumes. You must verify that only the failed storage volumes are reformatted as part of the recovery procedure.

Before you begin

You are signed in to the Grid Manager using a supported web browser.

About this task

You should recover failed storage volumes as soon as possible.

The first step of the recovery process is to detect volumes that have become detached, need to be unmounted, or have I/O errors. If failed volumes are still attached but have a randomly corrupted file system, the system might not detect any corruption in unused or unallocated parts of the disk.

> You must finish this procedure before performing manual steps to recover the volumes, such as adding or re-attaching the disks, stopping the node, starting the node, or rebooting. Otherwise, when you run the `reformat_storage_block_devices.rb` script, you might encounter a file system error that causes the script to hang or fail.

> Repair the hardware and properly attach the disks before running the `reboot` command.

> Identify failed storage volumes carefully. You will use this information to verify which volumes must be reformatted. Once a volume has been reformatted, data on the volume can't be recovered.

To correctly recover failed storage volumes, you need to know both the device names of the failed storage volumes and their volume IDs.

At installation, each storage device is assigned a file system universal unique identifier (UUID) and is mounted to a rangedb directory on the Storage Node using that assigned file system UUID. The file system UUID and the rangedb directory are listed in the `/etc/fstab` file. The device name, rangedb directory, and the size of the mounted volume are displayed in the Grid Manager.

In the following example, device `/dev/sdc` has a volume size of 4 TB, is mounted to `/var/local/rangedb/0`, using the device name `/dev/disk/by-uuid/822b0547-3b2b-472e-ad5e-e1cf1809faba` in the `/etc/fstab` file:

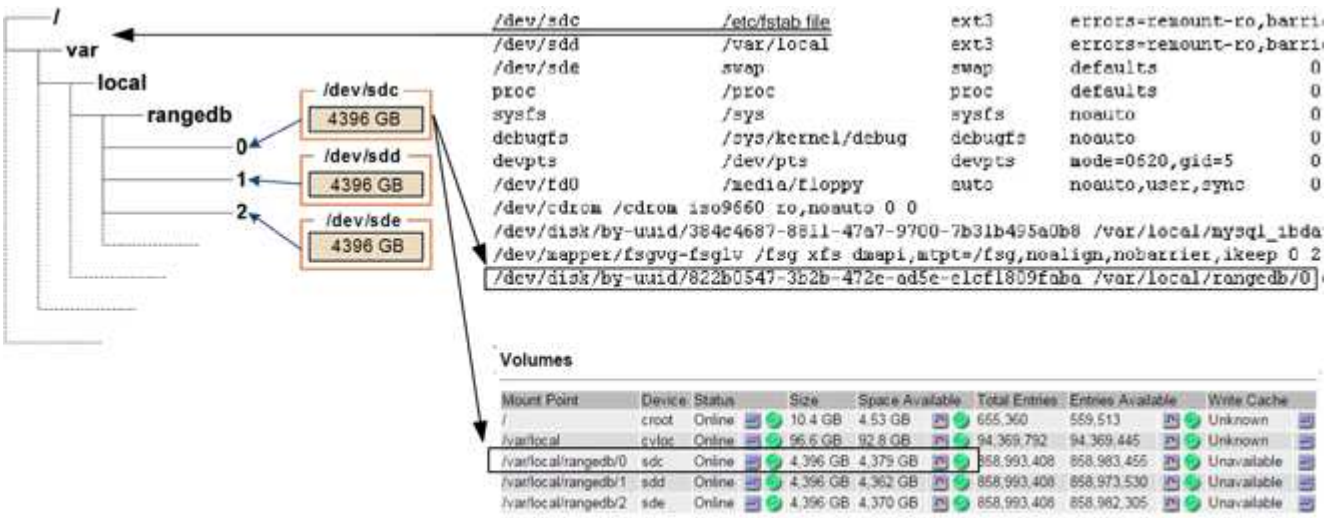

Steps

1. Complete the following steps to record the failed storage volumes and their device names:
   a. Select **SUPPORT** > **Tools** > **Grid topology**.
   b. Select **site** > **failed Storage Node** > **LDR** > **Storage** > **Overview** > **Main**, and look for object stores with alarms.

Object Stores

ID	Total	Available	Stored Data	Stored (%)	Health
0000	96.6 GB	96.6 GB	823 KB	0.001 %	Error
0001	107 GB	107 GB	0 B	0 %	No Errors
0002	107 GB	107 GB	0 B	0 %	No Errors

   c. Select **site** > **failed Storage Node** > **SSM** > **Resources** > **Overview** > **Main**. Determine the mount point and volume size of each failed storage volume identified in the previous step.

      Object stores are numbered in hex notation. For example, 0000 is the first volume and 000F is the sixteenth volume. In the example, the object store with an ID of 0000 corresponds to `/var/local/rangedb/0` with device name sdc and a size of 107 GB.

Volumes

Mount Point	Device	Status	Size	Space Available	Total Entries	Entries Available	Write Cache
/	croot	Online	10.4 GB	4.17 GB	655,360	554,806	Unknown
/var/local	cvloc	Online	96.6 GB	96.1 GB	94,369,792	94,369,423	Unknown
/var/local/rangedb/0	sdc	Online	107 GB	107 GB	104,857,600	104,856,202	Enabled
/var/local/rangedb/1	sdd	Online	107 GB	107 GB	104,857,600	104,856,536	Enabled
/var/local/rangedb/2	sde	Online	107 GB	107 GB	104,857,600	104,856,536	Enabled

2. Log in to the failed Storage Node:

   a. Enter the following command: `ssh admin@grid_node_IP`

   b. Enter the password listed in the `Passwords.txt` file.

   c. Enter the following command to switch to root: `su -`

   d. Enter the password listed in the `Passwords.txt` file.

   When you are logged in as root, the prompt changes from `$` to `#`.

3. Run the following script to unmount a failed storage volume:

   ```
   sn-unmount-volume object_store_ID
   ```

   The `object_store_ID` is the ID of the failed storage volume. For example, specify `0` in the command for an object store with ID 0000.

4. If prompted, press **y** to stop the Cassandra service depending on storage volume 0.

   If the Cassandra service is already stopped, you aren't prompted. The Cassandra service is stopped only for volume 0.

   ```
   root@Storage-180:~/var/local/tmp/storage~ # sn-unmount-volume 0
   Services depending on storage volume 0 (cassandra) aren't down.
   Services depending on storage volume 0 must be stopped before running
   this script.
   Stop services that require storage volume 0 [y/N]? y
   Shutting down services that require storage volume 0.
   Services requiring storage volume 0 stopped.
   Unmounting /var/local/rangedb/0
   /var/local/rangedb/0 is unmounted.
   ```

   In a few seconds, the volume is unmounted. Messages appear indicating each step of the process. The final message indicates that the volume is unmounted.

5. If the unmount fails because the volume is busy, you can force an unmount using the `--use-umountof` option:

   Forcing an unmount using the `--use-umountof` option might cause processes or services using the volume to behave unexpectedly or crash.

```
root@Storage-180:~ # sn-unmount-volume --use-umountof
/var/local/rangedb/2
Unmounting /var/local/rangedb/2 using umountof
/var/local/rangedb/2 is unmounted.
Informing LDR service of changes to storage volumes
```

Recover failed storage volumes and rebuild Cassandra database

You must run a script that reformats and remounts storage on failed storage volumes, and rebuilds the Cassandra database on the Storage Node if the system determines that it is necessary.

Before you begin

- You have the `Passwords.txt` file.
- The system drives on the server are intact.
- The cause of the failure has been identified and, if necessary, replacement storage hardware has already been acquired.
- The total size of the replacement storage is the same as the original.
- You have checked that a Storage Node decommissioning is not in progress, or you have paused the node decommission procedure. (In the Grid Manager, select **MAINTENANCE** > **Tasks** > **Decommission**.)
- You have checked that an expansion is not in progress. (In the Grid Manager, select **MAINTENANCE** > **Tasks** > **Expansion**.)
- You have reviewed the warnings about storage volume recovery.

Steps

1. As needed, replace failed physical or virtual storage associated with the failed storage volumes that you identified and unmounted earlier.

   Don't remount the volumes in this step. The storage is remounted and added to `/etc/fstab` in a later step.

2. In the Grid Manager, go to **NODES** > **`appliance Storage Node`** > **Hardware**. In the StorageGRID Appliance section of the page, verify that the Storage RAID mode is healthy.
3. Log in to the failed Storage Node:
   a. Enter the following command: `ssh admin@`*`grid_node_IP`*
   b. Enter the password listed in the `Passwords.txt` file.
   c. Enter the following command to switch to root: `su -`
   d. Enter the password listed in the `Passwords.txt` file.

      When you are logged in as root, the prompt changes from `$` to `#`.

4. Use a text editor (vi or vim) to delete failed volumes from the `/etc/fstab` file and then save the file.

Commenting out a failed volume in the `/etc/fstab` file is insufficient. The volume must be deleted from `fstab` as the recovery process verifies that all lines in the `fstab` file match the mounted file systems.

5. Reformat any failed storage volumes and rebuild the Cassandra database if it is necessary. Enter: `reformat_storage_block_devices.rb`
   - When storage volume 0 is unmounted, prompts and messages will indicate that the Cassandra service is being stopped.
   - You will be prompted to rebuild the Cassandra database if it is necessary.
     - Review the warnings. If none of them apply, rebuild the Cassandra database. Enter: **y**
     - If more than one Storage Node is offline or if another Storage Node has been rebuilt in the last 15 days. Enter: **n**

       The script will exit without rebuilding Cassandra. Contact technical support.
   - For each rangedb drive on the Storage Node, when you are asked: `Reformat the rangedb drive <name> (device <major number>:<minor number>)? [y/n]?`, enter one of the following responses:
     - **y** to reformat a drive that had errors. This reformats the storage volume and adds the reformatted storage volume to the `/etc/fstab` file.
     - **n** if the drive contains no errors, and you don't want to reformat it.

       Selecting **n** exits the script. Either mount the drive (if you think the data on the drive should be retained and the drive was unmounted in error) or remove the drive. Then, run the `reformat_storage_block_devices.rb` command again.

       Some StorageGRID recovery procedures use Reaper to handle Cassandra repairs. Repairs occur automatically as soon as the related or required services have started. You might notice script output that mentions "reaper" or "Cassandra repair." If you see an error message indicating the repair has failed, run the command indicated in the error message.

     In the following example output, the drive `/dev/sdf` must be reformatted, and Cassandra did not need to be rebuilt:

```
root@DC1-S1:~ # reformat_storage_block_devices.rb
Formatting devices that are not in use...
Skipping in use device /dev/sdc
Skipping in use device /dev/sdd
Skipping in use device /dev/sde
Reformat the rangedb drive /dev/sdf (device 8:64)? [Y/n]? y
Successfully formatted /dev/sdf with UUID b951bfcb-4804-41ad-b490-
805dfd8df16c
All devices processed
Running: /usr/local/ldr/setup_rangedb.sh 12368435
Cassandra does not need rebuilding.
Starting services.
Informing storage services of new volume

Reformatting done.  Now do manual steps to
restore copies of data.
```

After the storage volumes are reformatted and remounted and necessary Cassandra operations are complete, you can restore object data using Grid Manager.

Restore object data to storage volume where system drive is intact

After recovering a storage volume on a Storage Node where the system drive is intact, you can restore the replicated or erasure-coded object data that was lost when the storage volume failed.

Which procedure should I use?

Whenever possible, restore object data using the **Volume restoration** page in the Grid Manager.

- If the volumes are listed at **MAINTENANCE** > **Volume restoration** > **Nodes to restore**, restore object data using the Volume restoration page in the Grid Manager.
- If the volumes aren't listed at **MAINTENANCE** > **Volume restoration** > **Nodes to restore**, follow the steps below for using the `repair-data` script to restore object data.

  If the recovered Storage Node contains fewer volumes than the node it is replacing, you must use the `repair-data` script.

Use the `repair-data` script to restore object data

Before you begin

- You must have confirmed that the recovered Storage Node has a Connection State of **Connected** on the **NODES** > **Overview** tab in the Grid Manager.

About this task

Object data can be restored from other Storage Nodes, an Archive Node, or a Cloud Storage Pool, assuming

that the grid's ILM rules were configured such that object copies are available.

Note the following:

- If an ILM rule was configured to store only one replicated copy and that copy existed on a storage volume that failed, you will not be able to recover the object.
- If the only remaining copy of an object is in a Cloud Storage Pool, StorageGRID must issue multiple requests to the Cloud Storage Pool endpoint to restore object data. Before performing this procedure, contact technical support for help in estimating the recovery time frame and the associated costs.
- If the only remaining copy of an object is on an Archive Node, object data is retrieved from the Archive Node. Restoring object data to a Storage Node from an Archive Node takes longer than restoring copies from other Storage Nodes because of the latency associated with retrievals from external archival storage systems.

About the `repair-data` script

To restore object data, you run the `repair-data` script. This script begins the process of restoring object data and works with ILM scanning to ensure that ILM rules are met.

Select **Replicated data** or **Erasure-coded (EC) data** below to learn the different options for the `repair-data` script, based on whether you are restoring replicated data or erasure-coded data. If you need to restore both types of data, you must run both sets of commands.

> For more information about the `repair-data` script, enter `repair-data --help` from the command line of the primary Admin Node.

Replicated data

Two commands are available for restoring replicated data, based on whether you need to repair the entire node or only certain volumes on the node:

```
repair-data start-replicated-node-repair

repair-data start-replicated-volume-repair
```

You can track repairs of replicated data with this command:

```
repair-data show-replicated-repair-status
```

Erasure coded (EC) data

Two commands are available for restoring erasure-coded data, based on whether you need to repair the entire node or only certain volumes on the node:

```
repair-data start-ec-node-repair

repair-data start-ec-volume-repair
```

You can track repairs of erasure-coded data with this command:

```
repair-data show-ec-repair-status
```

> Repairs of erasure-coded data can begin while some Storage Nodes are offline. However, if all erasure-coded data can't be accounted for, the repair can't be completed. Repair will complete after all nodes are available.

> The EC repair job temporarily reserves a large amount of storage. Storage alerts might be triggered, but will resolve when the repair is complete. If there is not enough storage for the reservation, the EC repair job will fail. Storage reservations are released when the EC repair job completes, whether the job failed or succeeded.

Find hostname for Storage Node

1. Log in to the primary Admin Node:
   a. Enter the following command: `ssh admin@primary_Admin_Node_IP`
   b. Enter the password listed in the `Passwords.txt` file.
   c. Enter the following command to switch to root: `su -`
   d. Enter the password listed in the `Passwords.txt` file.

   When you are logged in as root, the prompt changes from `$` to `#`.

2. Use the `/etc/hosts` file to find the hostname of the Storage Node for the restored storage volumes. To see a list of all nodes in the grid, enter the following: `cat /etc/hosts`.

Repair data if all volumes have failed

If all storage volumes have failed, repair the entire node. Follow the instructions for **replicated data**, **erasure-coded (EC) data**, or both, based on whether you use replicated data, erasure-coded (EC) data, or both.

If only some volumes have failed, go to Repair data if only some volumes have failed.

> You can't run `repair-data` operations for more than one node at the same time. To recover multiple nodes, contact technical support.

Replicated data

If your grid includes replicated data, use the `repair-data start-replicated-node-repair` command with the `--nodes` option, where `--nodes` is the hostname (system name), to repair the entire Storage Node.

This command repairs the replicated data on a Storage Node named SG-DC-SN3:

```
repair-data start-replicated-node-repair --nodes SG-DC-SN3
```

> As object data is restored, the **Objects Lost** alert is triggered if the StorageGRID system can't locate replicated object data. Alerts might be triggered on Storage Nodes throughout the system. You should determine the cause of the loss and if recovery is possible. See Investigate lost objects.

Erasure coded (EC) data

If your grid contains erasure-coded data, use the `repair-data start-ec-node-repair` command with the `--nodes` option, where `--nodes` is the hostname (system name), to repair the entire Storage Node.

This command repairs the erasure-coded data on a Storage Node named SG-DC-SN3:

```
repair-data start-ec-node-repair --nodes SG-DC-SN3
```

The operation returns a unique `repair ID` that identifies this `repair_data` operation. Use this `repair ID` to track the progress and result of the `repair_data` operation. No other feedback is returned as the recovery process completes.

> Repairs of erasure-coded data can begin while some Storage Nodes are offline. Repair will complete after all nodes are available.

Repair data if only some volumes have failed

If only some of the volumes have failed, repair the affected volumes. Follow the instructions for **replicated data**, **erasure-coded (EC) data**, or both, based on whether you use replicated data, erasure-coded (EC) data, or both.

If all volumes have failed, go to Repair data if all volumes have failed.

Enter the volume IDs in hexadecimal. For example, `0000` is the first volume and `000F` is the sixteenth volume. You can specify one volume, a range of volumes, or multiple volumes that aren't in a sequence.

All the volumes must be on the same Storage Node. If you need to restore volumes for more than one Storage Node, contact technical support.

Replicated data

If your grid contains replicated data, use the `start-replicated-volume-repair` command with the `--nodes` option to identify the node (where `--nodes` is the hostname of the node). Then add either the `--volumes` or `--volume-range` option, as shown in the following examples.

Single volume: This command restores replicated data to volume `0002` on a Storage Node named SG-DC-SN3:

```
repair-data start-replicated-volume-repair --nodes SG-DC-SN3 --volumes 0002
```

Range of volumes: This command restores replicated data to all volumes in the range `0003` to `0009` on a Storage Node named SG-DC-SN3:

```
repair-data start-replicated-volume-repair --nodes SG-DC-SN3 --volume-range
0003,0009
```

Multiple volumes not in a sequence: This command restores replicated data to volumes `0001`, `0005`, and `0008` on a Storage Node named SG-DC-SN3:

```
repair-data start-replicated-volume-repair --nodes SG-DC-SN3 --volumes
0001,0005,0008
```

> As object data is restored, the **Objects Lost** alert is triggered if the StorageGRID system can't locate replicated object data. Alerts might be triggered on Storage Nodes throughout the system. Note the alert description and recommended actions to determine the cause of the loss and if recovery is possible.

Erasure coded (EC) data

If your grid contains erasure-coded data, use the `start-ec-volume-repair` command with the `--nodes` option to identify the node (where `--nodes` is the hostname of the node). Then add either the `--volumes` or `--volume-range` option, as shown in the following examples.

Single volume: This command restores erasure-coded data to volume `0007` on a Storage Node named SG-DC-SN3:

```
repair-data start-ec-volume-repair --nodes SG-DC-SN3 --volumes 0007
```

Range of volumes: This command restores erasure-coded data to all volumes in the range `0004` to `0006` on a Storage Node named SG-DC-SN3:

```
repair-data start-ec-volume-repair --nodes SG-DC-SN3 --volume-range 0004,0006
```

Multiple volumes not in a sequence: This command restores erasure-coded data to volumes `000A`, `000C`, and `000E` on a Storage Node named SG-DC-SN3:

```
repair-data start-ec-volume-repair --nodes SG-DC-SN3 --volumes 000A,000C,000E
```

The `repair-data` operation returns a unique `repair ID` that identifies this `repair_data` operation. Use this `repair ID` to track the progress and result of the `repair_data` operation. No other feedback is returned as the recovery process completes.

Repairs of erasure-coded data can begin while some Storage Nodes are offline. Repair will complete after all nodes are available.

Monitor repairs

Monitor the status of the repair jobs, based on whether you use **replicated data**, **erasure-coded (EC) data**, or both.

You can also monitor the status of volume restoration jobs in process and view a history of restoration jobs completed in Grid Manager.

Replicated data

- To get an estimated percent completion for the replicated repair, add the `show-replicated-repair-status` option to the repair-data command.

  ```
  repair-data show-replicated-repair-status
  ```

- To determine if repairs are complete:
  1. Select **NODES** > ***Storage Node being repaired*** > **ILM**.
  2. Review the attributes in the Evaluation section. When repairs are complete, the **Awaiting - All** attribute indicates 0 objects.
- To monitor the repair in more detail:
  1. Select **SUPPORT** > **Tools** > **Grid topology**.
  2. Select ***grid*** > ***Storage Node being repaired*** > **LDR** > **Data Store**.
  3. Use a combination of the following attributes to determine, as well as possible, if replicated repairs are complete.

     Cassandra inconsistencies might be present, and failed repairs aren't tracked.

     - **Repairs Attempted (XRPA)**: Use this attribute to track the progress of replicated repairs. This attribute increases each time a Storage Node tries to repair a high-risk object. When this attribute does not increase for a period longer than the current scan period (provided by the **Scan Period — Estimated** attribute), it means that ILM scanning found no high-risk objects that need to be repaired on any nodes.

       High-risk objects are objects that are at risk of being completely lost. This does not include objects that don't satisfy their ILM configuration.

     - **Scan Period — Estimated (XSCM)**: Use this attribute to estimate when a policy change will be applied to previously ingested objects. If the **Repairs Attempted** attribute does not increase for a period longer than the current scan period, it is probable that replicated repairs are done. Note that the scan period can change. The **Scan Period — Estimated (XSCM)** attribute applies to the entire grid and is the maximum of all node scan periods. You can query the **Scan Period — Estimated** attribute history for the grid to determine an appropriate time frame.

Erasure coded (EC) data

To monitor the repair of erasure-coded data and retry any requests that might have failed:

1. Determine the status of erasure-coded data repairs:
   - Select **SUPPORT** > **Tools** > **Metrics** to view the estimated time to completion and the completion percentage for the current job. Then, select **EC Overview** in the Grafana section. Look at the **Grid EC Job Estimated Time to Completion** and **Grid EC Job Percentage Completed** dashboards.
   - Use this command to see the status of a specific `repair-data` operation:

     ```
     repair-data show-ec-repair-status --repair-id repair ID
     ```

   - Use this command to list all repairs:

```
repair-data show-ec-repair-status
```

The output lists information, including `repair ID`, for all previously and currently running repairs.

2. If the output shows that the repair operation failed, use the `--repair-id` option to retry the repair.

   This command retries a failed node repair, using the repair ID 6949309319275667690:

   ```
   repair-data start-ec-node-repair --repair-id 6949309319275667690
   ```

   This command retries a failed volume repair, using the repair ID 6949309319275667690:

   ```
   repair-data start-ec-volume-repair --repair-id 6949309319275667690
   ```

Check storage state after recovering storage volumes

After recovering storage volumes, you must verify that the desired state of the Storage Node is set to online and ensure that the state will be online by default whenever the Storage Node server is restarted.

Before you begin

- You are signed in to the Grid Manager using a supported web browser.
- The Storage Node has been recovered, and data recovery is complete.

Steps

1. Select **SUPPORT** > **Tools** > **Grid topology**.
2. Check the values of **Recovered Storage Node** > **LDR** > **Storage** > **Storage State — Desired** and **Storage State — Current**.

   The value of both attributes should be Online.

3. If the Storage State — Desired is set to Read-only, complete the following steps:
   a. Click the **Configuration** tab.
   b. From the **Storage State — Desired** drop-down list, select **Online**.
   c. Click **Apply Changes**.
   d. Click the **Overview** tab and confirm that the values of **Storage State — Desired** and **Storage State — Current** are updated to Online.

Recover from system drive failure

Recover from system drive failure: Workflow

If the system drive on a software-based Storage Node has failed, the Storage Node is not available to the StorageGRID system. You must complete a specific set of tasks to recover from a system drive failure.

Use this procedure to recover from a system drive failure on a software-based Storage Node. This procedure includes the steps to follow if any storage volumes also failed or can't be remounted.

This procedure applies to software-based Storage Nodes only. You must follow a different procedure to recover an appliance Storage Node.

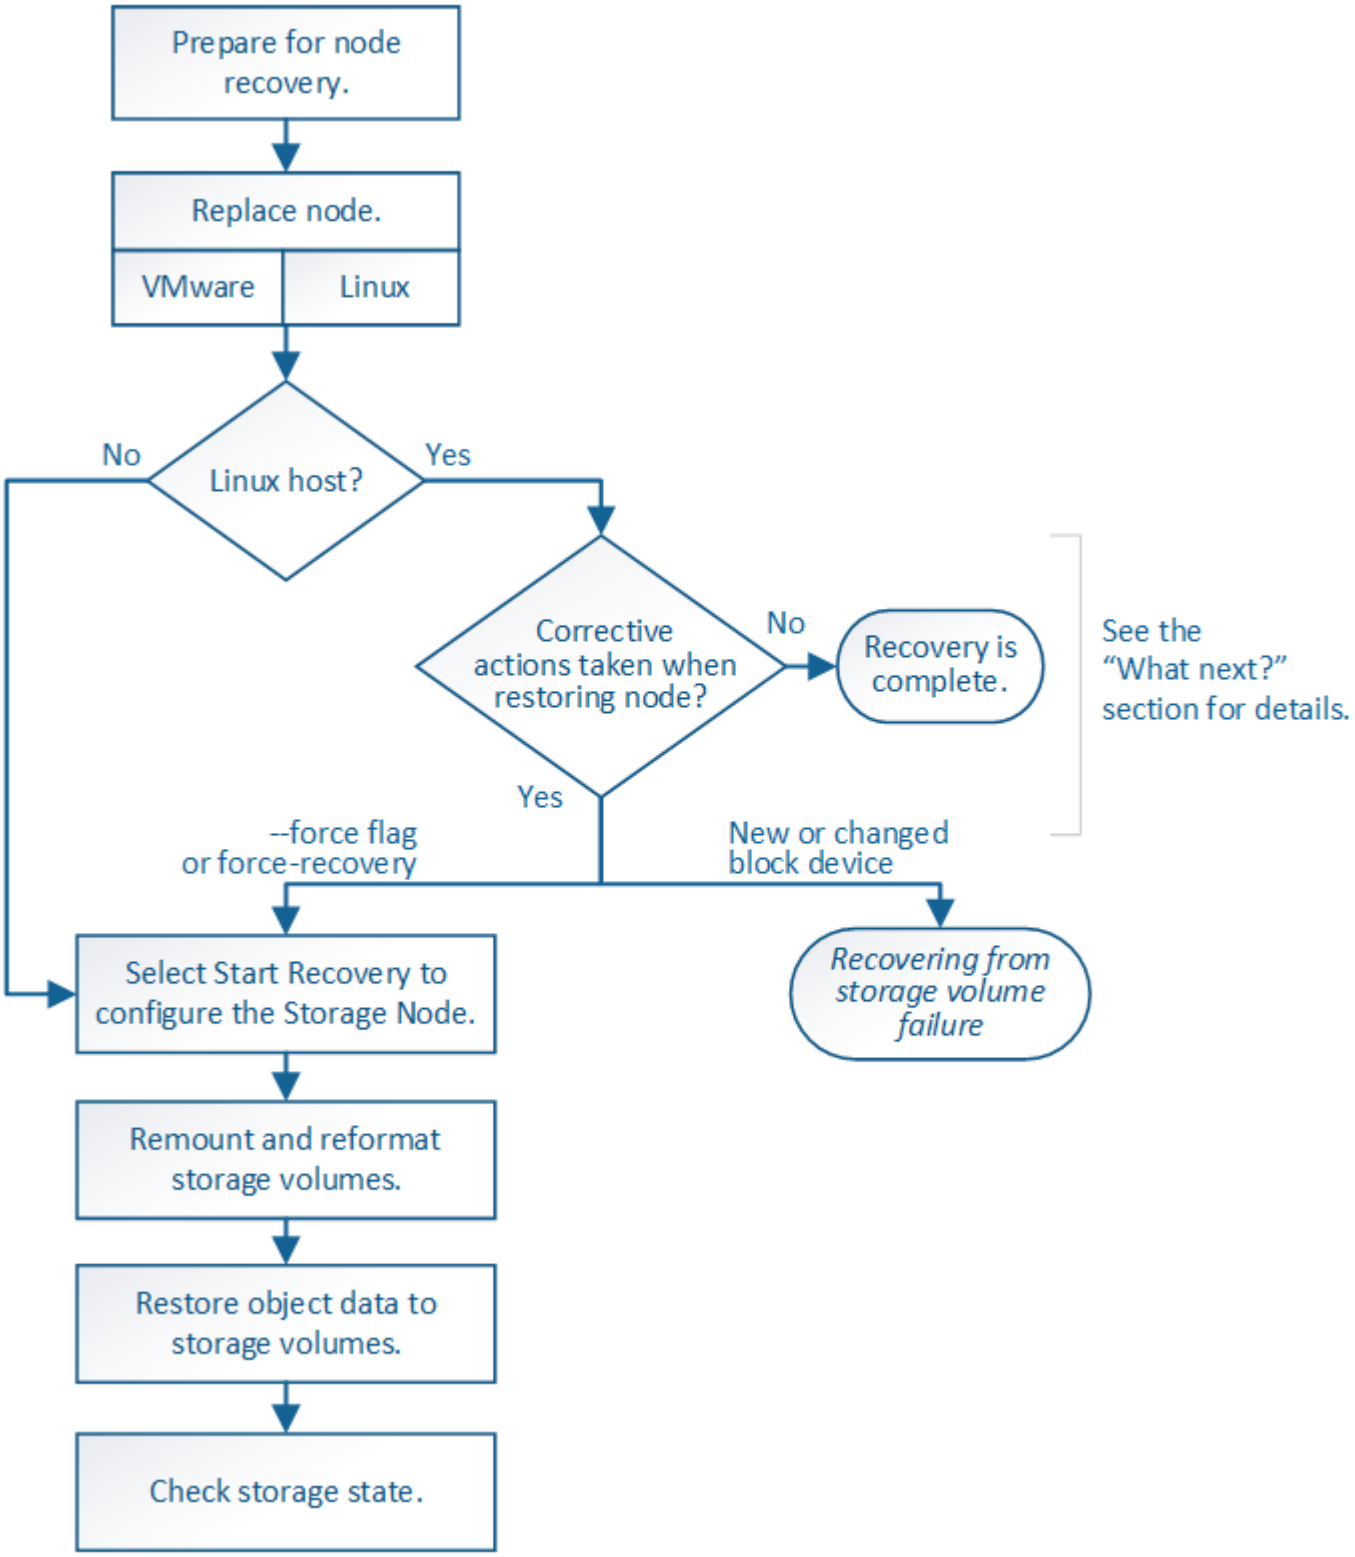

Review warnings for Storage Node system drive recovery

Before recovering a failed system drive of a Storage Node, review the general warnings and considerations for grid node recovery and the following specific warnings.

Storage Nodes have a Cassandra database that includes object metadata. The Cassandra database might be rebuilt in the following circumstances:

- A Storage Node is brought back online after having been offline for more than 15 days.
- A storage volume has failed and been recovered.
- The system drive and one or more storage volumes fails and is recovered.

When Cassandra is rebuilt, the system uses information from other Storage Nodes. If too many Storage Nodes are offline, some Cassandra data might not be available. If Cassandra has been rebuilt recently, Cassandra data might not yet be consistent across the grid. Data loss can occur if Cassandra is rebuilt when too many Storage Nodes are offline or if two or more Storage Nodes are rebuilt within 15 days of each other.

If more than one Storage Node has failed (or is offline), contact technical support. Don't perform the following recovery procedure. Data loss could occur.

If this is the second Storage Node failure in less than 15 days after a Storage Node failure or recovery, contact technical support. Rebuilding Cassandra on two or more Storage Nodes within 15 days can result in data loss.

If more than one Storage Node at a site has failed, a site recovery procedure might be required. See How site recovery is performed by technical support.

If this Storage Node is in read-only maintenance mode to allow for the retrieval of objects by another Storage Node with failed storage volumes, recover volumes on the Storage Node with failed storage volumes before recovering this failed Storage Node. See the instructions to recover from storage volume failure where system drive is intact.

If ILM rules are configured to store only one replicated copy and the copy exists on a storage volume that has failed, you will not be able to recover the object.

If you encounter a Services: Status - Cassandra (SVST) alarm during recovery, see Recover failed storage volumes and rebuild Cassandra database. After Cassandra is rebuilt, alarms should clear. If alarms don't clear, contact technical support.

Replace the Storage Node

If the system drive has failed, you must first replace the Storage Node.

You must select the node replacement procedure for your platform. The steps to replace a node are the same for all types of grid nodes.

This procedure applies to software-based Storage Nodes only. You must follow a different procedure to recover an appliance Storage Node.

Linux: If you aren't sure if your system drive has failed, follow the instructions to replace the node to determine which recovery steps are required.

Platform	Procedure
VMware	Replace a VMware node
Linux	Replace a Linux node
OpenStack	NetApp-provided virtual machine disk files and scripts for OpenStack are no longer supported for recovery operations. If you need to recover a node running in an OpenStack deployment, download the files for your Linux operating system. Then, follow the procedure for replacing a Linux node.

Select Start Recovery to configure Storage Node

After replacing a Storage Node, you must select Start Recovery in the Grid Manager to configure the new node as a replacement for the failed node.

Before you begin

- You are signed in to the Grid Manager using a supported web browser.
- You have the Maintenance or Root access permission.
- You have the provisioning passphrase.
- You have deployed and configured the replacement node.
- You have the start date of any repair jobs for erasure-coded data.
- You have verified that the Storage Node has not been rebuilt within the last 15 days.

About this task

If the Storage Node is installed as a container on a Linux host, you must perform this step only if one of these is true:

- You had to use the `--force` flag to import the node, or you issued `storagegrid node force-recovery node-name`
- You had to do a full node reinstall, or you needed to restore /var/local.

Steps

1. From the Grid Manager, select **MAINTENANCE** > **Tasks** > **Recovery**.
2. Select the grid node you want to recover in the Pending Nodes list.

   Nodes appear in the list after they fail, but you can't select a node until it has been reinstalled and is ready for recovery.

3. Enter the **Provisioning Passphrase**.
4. Click **Start Recovery**.

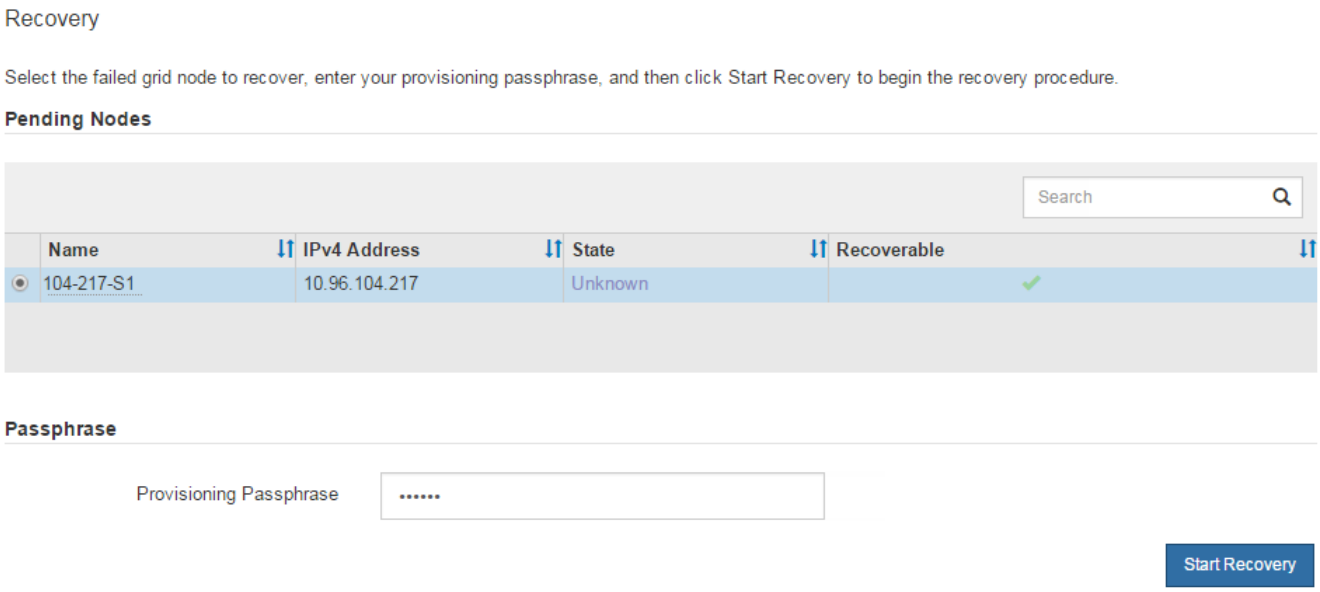

5. Monitor the progress of the recovery in the Recovering Grid Node table.

> While the recovery procedure is running, you can click **Reset** to start a new recovery. A dialog box appears, indicating that the node will be left in an indeterminate state if you reset the procedure.

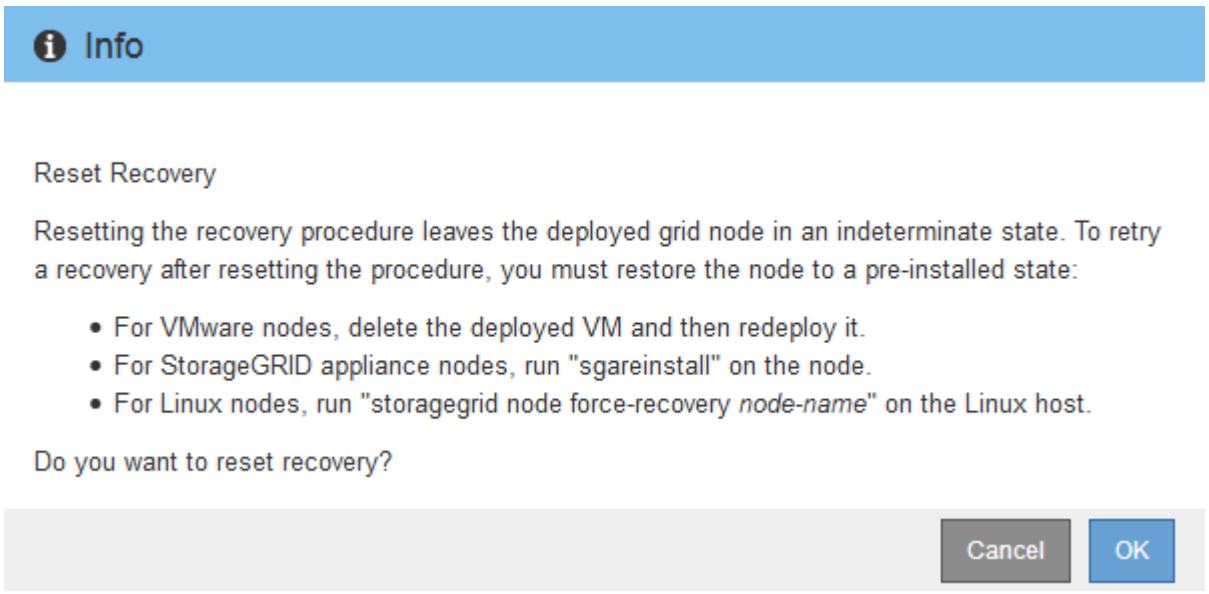

If you want to retry the recovery after resetting the procedure, you must restore the node to a pre-installed state, as follows:

- **VMware**: Delete the deployed virtual grid node. Then, when you are ready to restart the recovery, redeploy the node.
- **Linux**: Restart the node by running this command on the Linux host: `storagegrid node force-recovery node-name`

6. When the Storage Node reaches the “Waiting for Manual Steps” stage, go to Remount and reformat storage volumes (manual steps).

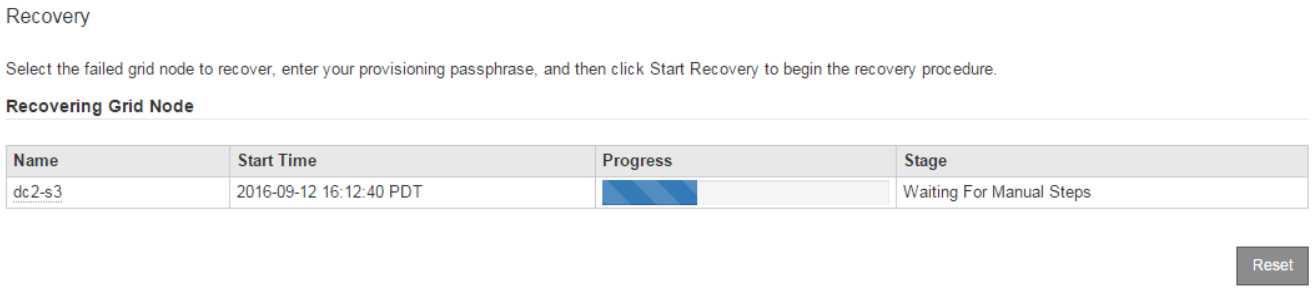

Remount and reformat storage volumes (manual steps)

You must manually run two scripts to remount preserved storage volumes and to reformat any failed storage volumes. The first script remounts volumes that are properly formatted as StorageGRID storage volumes. The second script reformats any unmounted volumes, rebuilds Cassandra, if needed, and starts services.

Before you begin

- You have already replaced the hardware for any failed storage volumes that you know require replacement.

  Running the `sn-remount-volumes` script might help you identify additional failed storage volumes.

- You have checked that a Storage Node decommissioning is not in progress, or you have paused the node decommission procedure. (In the Grid Manager, select **MAINTENANCE** > **Tasks** > **Decommission**.)
- You have checked that an expansion is not in progress. (In the Grid Manager, select **MAINTENANCE** > **Tasks** > **Expansion**.)
- You have reviewed the warnings for Storage Node system drive recovery.

> Contact technical support if more than one Storage Node is offline or if a Storage Node in this grid has been rebuilt in the last 15 days. Don't run the `sn-recovery-postinstall.sh` script. Rebuilding Cassandra on two or more Storage Nodes within 15 days of each other might result in data loss.

About this task

To complete this procedure, you perform these high-level tasks:

- Log in to the recovered Storage Node.
- Run the `sn-remount-volumes` script to remount properly formatted storage volumes. When this script runs, it does the following:
  - Mounts and unmounts each storage volume to replay the XFS journal.
  - Performs an XFS file consistency check.
  - If the file system is consistent, determines if the storage volume is a properly formatted StorageGRID storage volume.
  - If the storage volume is properly formatted, remounts the storage volume. Any existing data on the volume remains intact.
- Review the script output and resolve any issues.

- Run the `sn-recovery-postinstall.sh` script. When this script runs, it does the following.

> Don't reboot a Storage Node during recovery before running `sn-recovery-postinstall.sh` to reformat the failed storage volumes and restore object metadata. Rebooting the Storage Node before `sn-recovery-postinstall.sh` completes causes errors for services that attempt to start and causes StorageGRID appliance nodes to exit maintenance mode. See the step for post-install script.

  - Reformats any storage volumes that the `sn-remount-volumes` script could not mount or that were found to be improperly formatted.

> If a storage volume is reformatted, any data on that volume is lost. You must perform an additional procedure to restore object data from other locations in the grid, assuming that ILM rules were configured to store more than one object copy.

  - Rebuilds the Cassandra database on the node, if needed.
  - Starts the services on the Storage Node.

Steps

1. Log in to the recovered Storage Node:
   a. Enter the following command: `ssh admin@grid_node_IP`
   b. Enter the password listed in the `Passwords.txt` file.
   c. Enter the following command to switch to root: `su -`
   d. Enter the password listed in the `Passwords.txt` file.

   When you are logged in as root, the prompt changes from `$` to `#`.

2. Run the first script to remount any properly formatted storage volumes.

> If all storage volumes are new and need to be formatted, or if all storage volumes have failed, you can skip this step and run the second script to reformat all unmounted storage volumes.

   a. Run the script: `sn-remount-volumes`

      This script might take hours to run on storage volumes that contain data.

   b. As the script runs, review the output and answer any prompts.

> As required, you can use the `tail -f` command to monitor the contents of the script's log file (`/var/local/log/sn-remount-volumes.log`) . The log file contains more detailed information than the command line output.

```
root@SG:~ # sn-remount-volumes
The configured LDR noid is 12632740

====== Device /dev/sdb ======
```

```
Mount and unmount device /dev/sdb and checking file system
consistency:
The device is consistent.
Check rangedb structure on device /dev/sdb:
Mount device /dev/sdb to /tmp/sdb-654321 with rangedb mount options
This device has all rangedb directories.
Found LDR node id 12632740, volume number 0 in the volID file
Attempting to remount /dev/sdb
Device /dev/sdb remounted successfully

====== Device /dev/sdc ======
Mount and unmount device /dev/sdc and checking file system
consistency:
Error: File system consistency check retry failed on device /dev/sdc.
You can see the diagnosis information in the /var/local/log/sn-
remount-volumes.log.

This volume could be new or damaged. If you run sn-recovery-
postinstall.sh,
this volume and any data on this volume will be deleted. If you only
had two
copies of object data, you will temporarily have only a single copy.
StorageGRID Webscale will attempt to restore data redundancy by
making
additional replicated copies or EC fragments, according to the rules
in
the active ILM policy.

Don't continue to the next step if you believe that the data
remaining on
this volume can't be rebuilt from elsewhere in the grid (for example,
if
your ILM policy uses a rule that makes only one copy or if volumes
have
failed on multiple nodes). Instead, contact support to determine how
to
recover your data.

====== Device /dev/sdd ======
Mount and unmount device /dev/sdd and checking file system
consistency:
Failed to mount device /dev/sdd
This device could be an uninitialized disk or has corrupted
superblock.
File system check might take a long time. Do you want to continue? (y
or n) [y/N]? y
```

```
Error: File system consistency check retry failed on device /dev/sdd.
You can see the diagnosis information in the /var/local/log/sn-
remount-volumes.log.

This volume could be new or damaged. If you run sn-recovery-
postinstall.sh,
this volume and any data on this volume will be deleted. If you only
had two
copies of object data, you will temporarily have only a single copy.
StorageGRID Webscale will attempt to restore data redundancy by
making
additional replicated copies or EC fragments, according to the rules
in
the active ILM policy.

Don't continue to the next step if you believe that the data
remaining on
this volume can't be rebuilt from elsewhere in the grid (for example,
if
your ILM policy uses a rule that makes only one copy or if volumes
have
failed on multiple nodes). Instead, contact support to determine how
to
recover your data.

====== Device /dev/sde ======
Mount and unmount device /dev/sde and checking file system
consistency:
The device is consistent.
Check rangedb structure on device /dev/sde:
Mount device /dev/sde to /tmp/sde-654321 with rangedb mount options
This device has all rangedb directories.
Found LDR node id 12000078, volume number 9 in the volID file
Error: This volume does not belong to this node. Fix the attached
volume and re-run this script.
```

In the example output, one storage volume was remounted successfully and three storage volumes had errors.

- `/dev/sdb` passed the XFS file system consistency check and had a valid volume structure, so it was remounted successfully. Data on devices that are remounted by the script is preserved.
- `/dev/sdc` failed the XFS file system consistency check because the storage volume was new or corrupt.
- `/dev/sdd` could not be mounted because the disk was not initialized or the disk's superblock was corrupted. When the script can't mount a storage volume, it asks if you want to run the file system consistency check.

- If the storage volume is attached to a new disk, answer **N** to the prompt. You don't need check the file system on a new disk.
- If the storage volume is attached to an existing disk, answer **Y** to the prompt. You can use the results of the file system check to determine the source of the corruption. The results are saved in the `/var/local/log/sn-remount-volumes.log` log file.

- `/dev/sde` passed the XFS file system consistency check and had a valid volume structure; however, the LDR node ID in the volID file did not match the ID for this Storage Node (the `configured LDR noid` displayed at the top). This message indicates that this volume belongs to another Storage Node.

3. Review the script output and resolve any issues.

> If a storage volume failed the XFS file system consistency check or could not be mounted, carefully review the error messages in the output. You must understand the implications of running the `sn-recovery-postinstall.sh` script on these volumes.

a. Check to make sure that the results include an entry for all of the volumes you expected. If any volumes aren't listed, rerun the script.

b. Review the messages for all mounted devices. Make sure there are no errors indicating that a storage volume does not belong to this Storage Node.

In the example, the output for `/dev/sde` includes the following error message:

```
Error: This volume does not belong to this node. Fix the attached
volume and re-run this script.
```

> If a storage volume is reported as belonging to another Storage Node, contact technical support. If you run the `sn-recovery-postinstall.sh` script, the storage volume will be reformatted, which might cause data loss.

c. If any storage devices could not be mounted, make a note of the device name, and repair or replace the device.

> You must repair or replace any storage devices that could not be mounted.

You will use the device name to look up the volume ID, which is required input when you run the `repair-data` script to restore object data to the volume (the next procedure).

d. After repairing or replacing all unmountable devices, run the `sn-remount-volumes` script again to confirm that all storage volumes that can be remounted have been remounted.

> If a storage volume can't be mounted or is improperly formatted, and you continue to the next step, the volume and any data on the volume will be deleted. If you had two copies of object data, you will have only a single copy until you complete the next procedure (restoring object data).

Don't run the `sn-recovery-postinstall.sh` script if you believe that the data remaining on a failed storage volume can't be rebuilt from elsewhere in the grid (for example, if your ILM policy uses a rule that makes only one copy or if volumes have failed on multiple nodes). Instead, contact technical support to determine how to recover your data.

4. Run the `sn-recovery-postinstall.sh` script: `sn-recovery-postinstall.sh`

   This script reformats any storage volumes that could not be mounted or that were found to be improperly formatted; rebuilds the Cassandra database on the node, if needed; and starts the services on the Storage Node.

   Be aware of the following:

   - The script might take hours to run.
   - In general, you should leave the SSH session alone while the script is running.
   - Don't press **Ctrl+C** while the SSH session is active.
   - The script will run in the background if a network disruption occurs and terminates the SSH session, but you can view the progress from the Recovery page.
   - If the Storage Node uses the RSM service, the script might appear to stall for 5 minutes as node services are restarted. This 5-minute delay is expected whenever the RSM service boots for the first time.

     The RSM service is present on Storage Nodes that include the ADC service.

   Some StorageGRID recovery procedures use Reaper to handle Cassandra repairs. Repairs occur automatically as soon as the related or required services have started. You might notice script output that mentions "reaper" or "Cassandra repair." If you see an error message indicating the repair has failed, run the command indicated in the error message.

5. As the `sn-recovery-postinstall.sh` script runs, monitor the Recovery page in the Grid Manager.

   The Progress bar and the Stage column on the Recovery page provide a high-level status of the `sn-recovery-postinstall.sh` script.

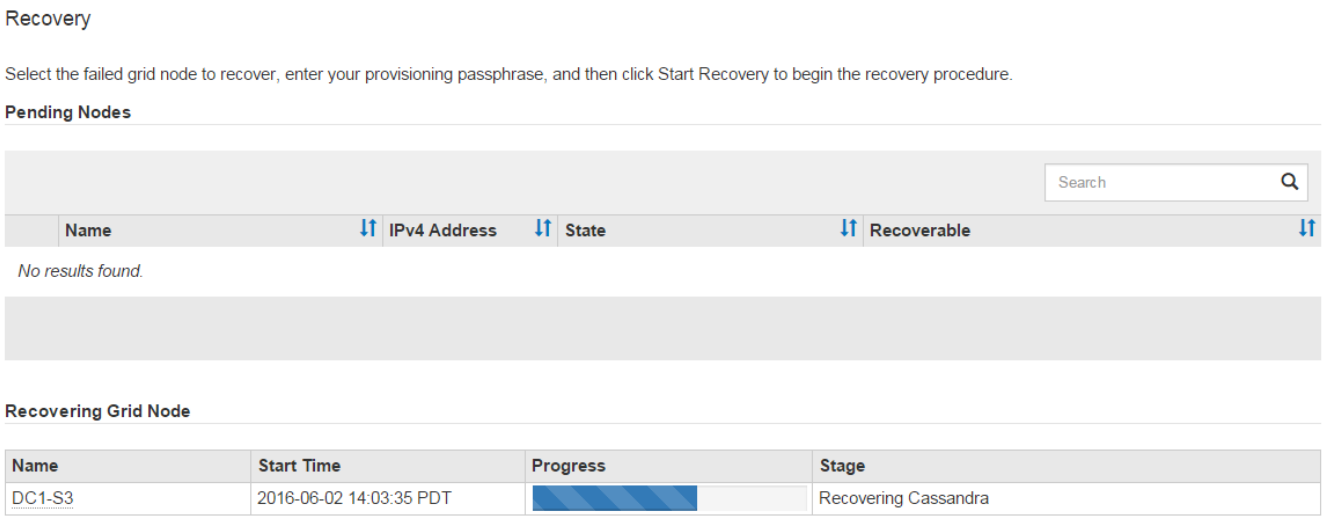

6. After the `sn-recovery-postinstall.sh` script has started services on the node, you can restore object data to any storage volumes that were formatted by the script.

   The script asks if you want to restore object data manually.

   - In most cases, you should restore object data using Grid Manager. Answer `n` to use the Grid Manager.
   - In rare cases, such as when instructed by technical support, or when you know that the replacement node has fewer volumes available for object storage than the original node, you must restore object data manually using the `repair-data` script. If one of these cases applies, answer `y`.

   > If you answer `y` to restore object data manually:
   >
   > - You aren't able to restore object data using Grid Manager.
   > - You can monitor the progress of manual restoration jobs using Grid Manager.

Restore object data to storage volume (system drive failure)

After recovering storage volumes for a non-appliance Storage Node, you can restore the replicated or erasure-coded object data that was lost when the Storage Node failed.

Which procedure should I use?

Whenever possible, restore object data using the **Volume restoration** page in the Grid Manager.

- If the volumes are listed at **MAINTENANCE** > **Volume restoration** > **Nodes to restore**, restore object data using the Volume restoration page in the Grid Manager.
- If the volumes aren't listed at **MAINTENANCE** > **Volume restoration** > **Nodes to restore**, follow the steps below for using the `repair-data` script to restore object data.

  If the recovered Storage Node contains fewer volumes than the node it is replacing, you must use the `repair-data` script.

Use the `repair-data` script to restore object data

Before you begin

- You must have confirmed that the recovered Storage Node has a Connection State of **Connected** on the **NODES** > **Overview** tab in the Grid Manager.

About this task

Object data can be restored from other Storage Nodes, an Archive Node, or a Cloud Storage Pool, assuming that the grid's ILM rules were configured such that object copies are available.

Note the following:

- If an ILM rule was configured to store only one replicated copy and that copy existed on a storage volume that failed, you will not be able to recover the object.
- If the only remaining copy of an object is in a Cloud Storage Pool, StorageGRID must issue multiple requests to the Cloud Storage Pool endpoint to restore object data. Before performing this procedure, contact technical support for help in estimating the recovery time frame and the associated costs.

- If the only remaining copy of an object is on an Archive Node, object data is retrieved from the Archive Node. Restoring object data to a Storage Node from an Archive Node takes longer than restoring copies from other Storage Nodes because of the latency associated with retrievals from external archival storage systems.

About the `repair-data` script

To restore object data, you run the `repair-data` script. This script begins the process of restoring object data and works with ILM scanning to ensure that ILM rules are met.

Select **Replicated data** or **Erasure-coded (EC) data** below to learn the different options for the `repair-data` script, based on whether you are restoring replicated data or erasure-coded data. If you need to restore both types of data, you must run both sets of commands.

For more information about the `repair-data` script, enter `repair-data --help` from the command line of the primary Admin Node.

Replicated data

Two commands are available for restoring replicated data, based on whether you need to repair the entire node or only certain volumes on the node:

```
repair-data start-replicated-node-repair

repair-data start-replicated-volume-repair
```

You can track repairs of replicated data with this command:

```
repair-data show-replicated-repair-status
```

Erasure coded (EC) data

Two commands are available for restoring erasure-coded data, based on whether you need to repair the entire node or only certain volumes on the node:

```
repair-data start-ec-node-repair

repair-data start-ec-volume-repair
```

You can track repairs of erasure-coded data with this command:

```
repair-data show-ec-repair-status
```

Repairs of erasure-coded data can begin while some Storage Nodes are offline. However, if all erasure-coded data can't be accounted for, the repair can't be completed. Repair will complete after all nodes are available.

The EC repair job temporarily reserves a large amount of storage. Storage alerts might be triggered, but will resolve when the repair is complete. If there is not enough storage for the reservation, the EC repair job will fail. Storage reservations are released when the EC repair job completes, whether the job failed or succeeded.

Find hostname for Storage Node

1. Log in to the primary Admin Node:
   a. Enter the following command: `ssh admin@primary_Admin_Node_IP`
   b. Enter the password listed in the `Passwords.txt` file.
   c. Enter the following command to switch to root: `su -`
   d. Enter the password listed in the `Passwords.txt` file.

      When you are logged in as root, the prompt changes from `$` to `#`.

2. Use the `/etc/hosts` file to find the hostname of the Storage Node for the restored storage volumes. To see a list of all nodes in the grid, enter the following: `cat /etc/hosts`.

Repair data if all volumes have failed

If all storage volumes have failed, repair the entire node. Follow the instructions for **replicated data**, **erasure-coded (EC) data**, or both, based on whether you use replicated data, erasure-coded (EC) data, or both.

If only some volumes have failed, go to Repair data if only some volumes have failed.

> You can't run `repair-data` operations for more than one node at the same time. To recover multiple nodes, contact technical support.

Replicated data

If your grid includes replicated data, use the `repair-data start-replicated-node-repair` command with the `--nodes` option, where `--nodes` is the hostname (system name), to repair the entire Storage Node.

This command repairs the replicated data on a Storage Node named SG-DC-SN3:

```
repair-data start-replicated-node-repair --nodes SG-DC-SN3
```

> As object data is restored, the **Objects Lost** alert is triggered if the StorageGRID system can't locate replicated object data. Alerts might be triggered on Storage Nodes throughout the system. You should determine the cause of the loss and if recovery is possible. See Investigate lost objects.

Erasure coded (EC) data

If your grid contains erasure-coded data, use the `repair-data start-ec-node-repair` command with the `--nodes` option, where `--nodes` is the hostname (system name), to repair the entire Storage Node.

This command repairs the erasure-coded data on a Storage Node named SG-DC-SN3:

```
repair-data start-ec-node-repair --nodes SG-DC-SN3
```

The operation returns a unique `repair ID` that identifies this `repair_data` operation. Use this `repair ID` to track the progress and result of the `repair_data` operation. No other feedback is returned as the recovery process completes.

> Repairs of erasure-coded data can begin while some Storage Nodes are offline. Repair will complete after all nodes are available.

Repair data if only some volumes have failed

If only some of the volumes have failed, repair the affected volumes. Follow the instructions for **replicated data**, **erasure-coded (EC) data**, or both, based on whether you use replicated data, erasure-coded (EC) data, or both.

If all volumes have failed, go to Repair data if all volumes have failed.

Enter the volume IDs in hexadecimal. For example, `0000` is the first volume and `000F` is the sixteenth volume. You can specify one volume, a range of volumes, or multiple volumes that aren't in a sequence.

All the volumes must be on the same Storage Node. If you need to restore volumes for more than one Storage Node, contact technical support.

Replicated data

If your grid contains replicated data, use the `start-replicated-volume-repair` command with the `--nodes` option to identify the node (where `--nodes` is the hostname of the node). Then add either the `--volumes` or `--volume-range` option, as shown in the following examples.

Single volume: This command restores replicated data to volume `0002` on a Storage Node named SG-DC-SN3:

```
repair-data start-replicated-volume-repair --nodes SG-DC-SN3 --volumes 0002
```

Range of volumes: This command restores replicated data to all volumes in the range `0003` to `0009` on a Storage Node named SG-DC-SN3:

```
repair-data start-replicated-volume-repair --nodes SG-DC-SN3 --volume-range
0003,0009
```

Multiple volumes not in a sequence: This command restores replicated data to volumes `0001`, `0005`, and `0008` on a Storage Node named SG-DC-SN3:

```
repair-data start-replicated-volume-repair --nodes SG-DC-SN3 --volumes
0001,0005,0008
```

> As object data is restored, the **Objects Lost** alert is triggered if the StorageGRID system can't locate replicated object data. Alerts might be triggered on Storage Nodes throughout the system. Note the alert description and recommended actions to determine the cause of the loss and if recovery is possible.

Erasure coded (EC) data

If your grid contains erasure-coded data, use the `start-ec-volume-repair` command with the `--nodes` option to identify the node (where `--nodes` is the hostname of the node). Then add either the `--volumes` or `--volume-range` option, as shown in the following examples.

Single volume: This command restores erasure-coded data to volume `0007` on a Storage Node named SG-DC-SN3:

```
repair-data start-ec-volume-repair --nodes SG-DC-SN3 --volumes 0007
```

Range of volumes: This command restores erasure-coded data to all volumes in the range `0004` to `0006` on a Storage Node named SG-DC-SN3:

```
repair-data start-ec-volume-repair --nodes SG-DC-SN3 --volume-range 0004,0006
```

Multiple volumes not in a sequence: This command restores erasure-coded data to volumes `000A`, `000C`, and `000E` on a Storage Node named SG-DC-SN3:

```
repair-data start-ec-volume-repair --nodes SG-DC-SN3 --volumes 000A,000C,000E
```

The `repair-data` operation returns a unique `repair ID` that identifies this `repair_data` operation. Use this `repair ID` to track the progress and result of the `repair_data` operation. No other feedback is returned as the recovery process completes.

Repairs of erasure-coded data can begin while some Storage Nodes are offline. Repair will complete after all nodes are available.

Monitor repairs

Monitor the status of the repair jobs, based on whether you use **replicated data**, **erasure-coded (EC) data**, or both.

You can also monitor the status of volume restoration jobs in process and view a history of restoration jobs completed in Grid Manager.

Replicated data

- To get an estimated percent completion for the replicated repair, add the `show-replicated-repair-status` option to the repair-data command.

  ```
  repair-data show-replicated-repair-status
  ```

- To determine if repairs are complete:
  1. Select **NODES** > ***Storage Node being repaired*** > **ILM**.
  2. Review the attributes in the Evaluation section. When repairs are complete, the **Awaiting - All** attribute indicates 0 objects.
- To monitor the repair in more detail:
  1. Select **SUPPORT** > **Tools** > **Grid topology**.
  2. Select ***grid*** > ***Storage Node being repaired*** > **LDR** > **Data Store**.
  3. Use a combination of the following attributes to determine, as well as possible, if replicated repairs are complete.

     Cassandra inconsistencies might be present, and failed repairs aren't tracked.

     - **Repairs Attempted (XRPA)**: Use this attribute to track the progress of replicated repairs. This attribute increases each time a Storage Node tries to repair a high-risk object. When this attribute does not increase for a period longer than the current scan period (provided by the **Scan Period — Estimated** attribute), it means that ILM scanning found no high-risk objects that need to be repaired on any nodes.

       High-risk objects are objects that are at risk of being completely lost. This does not include objects that don't satisfy their ILM configuration.

     - **Scan Period — Estimated (XSCM)**: Use this attribute to estimate when a policy change will be applied to previously ingested objects. If the **Repairs Attempted** attribute does not increase for a period longer than the current scan period, it is probable that replicated repairs are done. Note that the scan period can change. The **Scan Period — Estimated (XSCM)** attribute applies to the entire grid and is the maximum of all node scan periods. You can query the **Scan Period — Estimated** attribute history for the grid to determine an appropriate time frame.

Erasure coded (EC) data

To monitor the repair of erasure-coded data and retry any requests that might have failed:

1. Determine the status of erasure-coded data repairs:
   - Select **SUPPORT** > **Tools** > **Metrics** to view the estimated time to completion and the completion percentage for the current job. Then, select **EC Overview** in the Grafana section. Look at the **Grid EC Job Estimated Time to Completion** and **Grid EC Job Percentage Completed** dashboards.
   - Use this command to see the status of a specific `repair-data` operation:

     ```
     repair-data show-ec-repair-status --repair-id repair ID
     ```

   - Use this command to list all repairs:

```
repair-data show-ec-repair-status
```

The output lists information, including `repair ID`, for all previously and currently running repairs.

2. If the output shows that the repair operation failed, use the `--repair-id` option to retry the repair.

   This command retries a failed node repair, using the repair ID 6949309319275667690:

   ```
   repair-data start-ec-node-repair --repair-id 6949309319275667690
   ```

   This command retries a failed volume repair, using the repair ID 6949309319275667690:

   ```
   repair-data start-ec-volume-repair --repair-id 6949309319275667690
   ```

Check storage state after recovering Storage Node system drive

After recovering the system drive for a Storage Node, you must verify that the desired state of the Storage Node is set to online and ensure that the state will be online by default whenever the Storage Node server is restarted.

Before you begin

- You are signed in to the Grid Manager using a supported web browser.
- The Storage Node has been recovered, and data recovery is complete.

Steps

1. Select **SUPPORT** > **Tools** > **Grid topology**.
2. Check the values of **Recovered Storage Node** > **LDR** > **Storage** > **Storage State — Desired** and **Storage State — Current**.

   The value of both attributes should be Online.

3. If the Storage State — Desired is set to Read-only, complete the following steps:
   a. Click the **Configuration** tab.
   b. From the **Storage State — Desired** drop-down list, select **Online**.
   c. Click **Apply Changes**.
   d. Click the **Overview** tab and confirm that the values of **Storage State — Desired** and **Storage State — Current** are updated to Online.

Restore object data using Grid Manager

You can restore object data for a failed storage volume or Storage Node using Grid Manager. You can also use Grid Manager to monitor restoration processes in progress and display a restoration history.

Before you begin

- You have completed either of these procedures to format failed volumes:
  - Remount and reformat appliance storage volumes (manual steps)

    - Remount and reformat storage volumes (manual steps)
- You have confirmed that the Storage Node where you are restoring objects has a Connection State of **Connected** on the **NODES** > **Overview** tab in the Grid Manager.
- You have confirmed the following:
    - A grid expansion to add a Storage Node is not in process.
    - A Storage Node decommission is not in process or failed.
    - A recovery of a failed storage volume is not in process.
    - A recovery of a Storage Node with a failed system drive is not in process.
    - An EC rebalance job is not in process.
    - Appliance node cloning is not in process.

About this task

After you have replaced the drives and performed the manual steps to format the volumes, Grid Manager displays the volumes as candidates for restoration on the **MAINTENANCE** > **Volume restoration** > **Nodes to restore** tab.

Whenever possible, restore object data using the **Volume restoration** page in the Grid Manager. Follow these guidelines:

- If the volumes are listed at **MAINTENANCE** > **Volume restoration** > **Nodes to restore**, restore object data as described in the steps below. The volumes will be listed if:
    - Some, but not all, storage volumes in a node have failed
    - All storage volumes in a node have failed and are being replaced with the same number of volumes or more volumes

  The Volume restoration page in the Grid Manager also allows you to monitor the volume restoration process and view restoration history.
- If the volumes aren't listed in the Grid Manager as candidates for restoration, follow the appropriate steps for using the `repair-data` script to restore object data:
    - Restoring object data to storage volume (system drive failure)
    - Restore object data to storage volume where system drive is intact
    - Restore object data to storage volume for appliance

  If the recovered Storage Node contains fewer volumes than the node it is replacing, you must use the `repair-data` script.

You can restore two types of object data:

- Replicated data objects are restored from other locations, assuming that the grid's ILM rules were configured to make object copies available.
    - If an ILM rule was configured to store only one replicated copy and that copy existed on a storage volume that failed, you will not be able to recover the object.
    - If the only remaining copy of an object is in a Cloud Storage Pool, StorageGRID must issue multiple requests to the Cloud Storage Pool endpoint to restore object data.
    - If the only remaining copy of an object is on an Archive Node, object data is retrieved from the Archive

Node. Restoring object data to a Storage Node from an Archive Node takes longer than restoring object copies from other Storage Nodes.

- Erasure coded (EC) data objects are restored by reassembling the stored fragments. Corrupt or lost fragments are recreated by the erasure-coding algorithm from the remaining data and parity fragments.

> Volume restoration is dependent on the availability of resources where object copies are stored. Progress of volume restoration is nonlinear and might take days or weeks to complete.

Restore failed volume or node

Follow these steps to restore a failed volume or node.

Steps

1. In Grid Manager go to **MAINTENANCE** > **Volume restoration**.
2. Select the **Nodes to restore** tab.

   The number on the tab indicates the number of nodes with volumes requiring restoration.

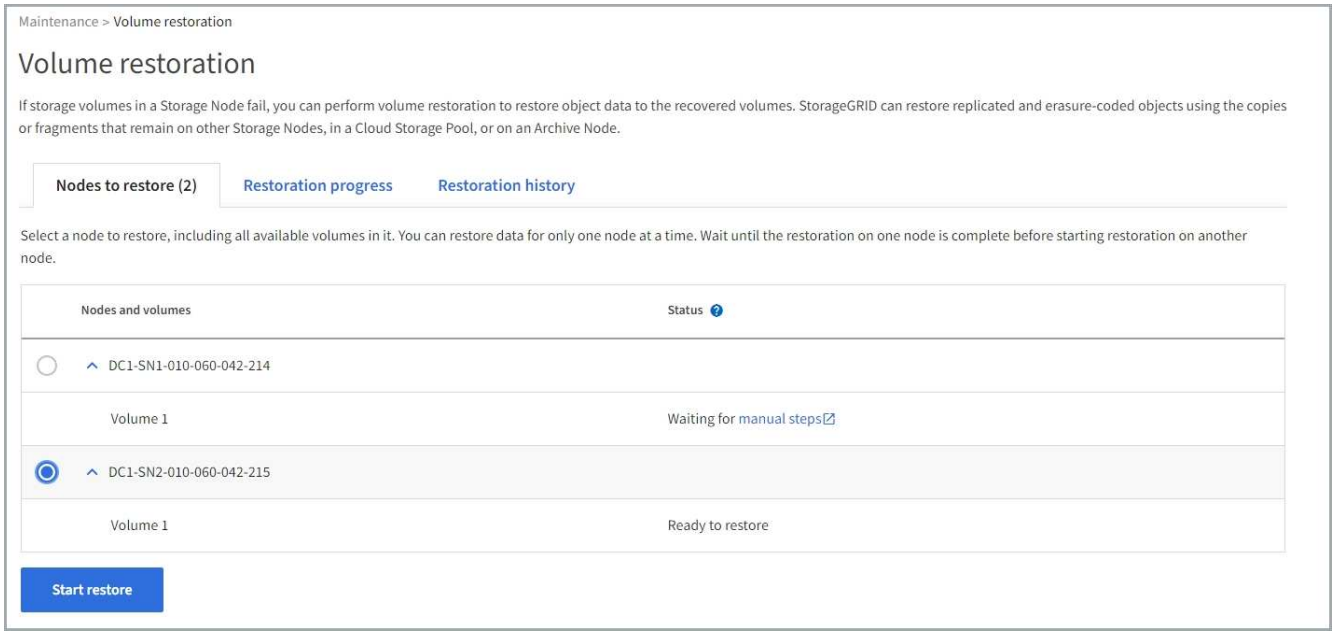

3. Expand each node to see the volumes in it that need restoration and their status.
4. Correct any issues preventing restoration of each volume that are indicated when you select Waiting for manual steps, if it displays as the volume status.
5. Select a node to restore where all the volumes indicate a Ready to restore status.

   You can only restore the volumes for one node at a time.

   Each volume in the node must indicate that it is ready to restore.

6. Select **Start restore**.
7. Address any warnings that might appear or select **Start anyway** to ignore the warnings and start the restoration.

Nodes are moved from the **Nodes to restore** tab to the **Restoration progress** tab when the restoration starts.

If a volume restoration can't be started, the node returns to the **Nodes to restore** tab.

View restoration progress

The **Restoration progress** tab shows the status of the volume restoration process and information about the volumes for a node being restored.

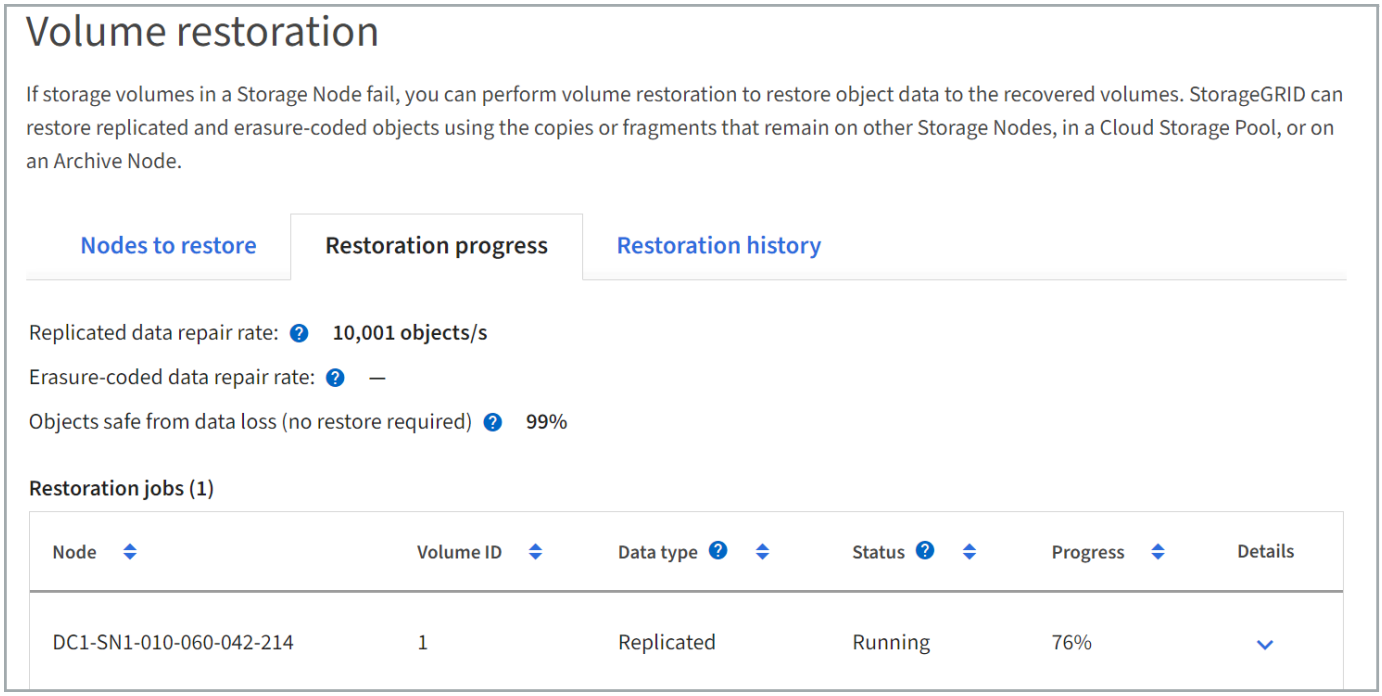

Data repair rates for replicated and erasure-coded objects in all volumes are averages summarizing all restorations in process, including those restorations initiated using the `repair-data` script. The percentage of objects in those volumes that are intact and don't require restoration is also indicated.

> Replicated data restoration is dependent on the availability of resources where the replicated copies are stored. Progress of replicated data restoration is nonlinear and might take days or weeks to complete.

The Restoration jobs section displays information about volume restorations started from Grid Manager.

- The number in the Restoration jobs section heading indicates the number of volumes that are either being restored or queued for restoration.
- The table displays information about each volume in a node being restored and its progress.
  - The progress for each node displays the percentage for each job.
  - Expand the Details column to display the restoration start time and job ID.
- If a volume restoration fails:
  - The Status column indicates failed.
  - An error appears, indicating the cause of failure.

    Correct the issues indicated in the error. Then select **Retry** to re-initiate the volume restoration.

    If multiple restoration jobs have failed, selecting **Retry** starts the most recently failed job.

View restoration history

The **Restoration history** tab shows information about all volume restorations that have successfully completed.

> Sizes aren't applicable for replicated objects and display only for restorations that contain erasure coded (EC) data objects.

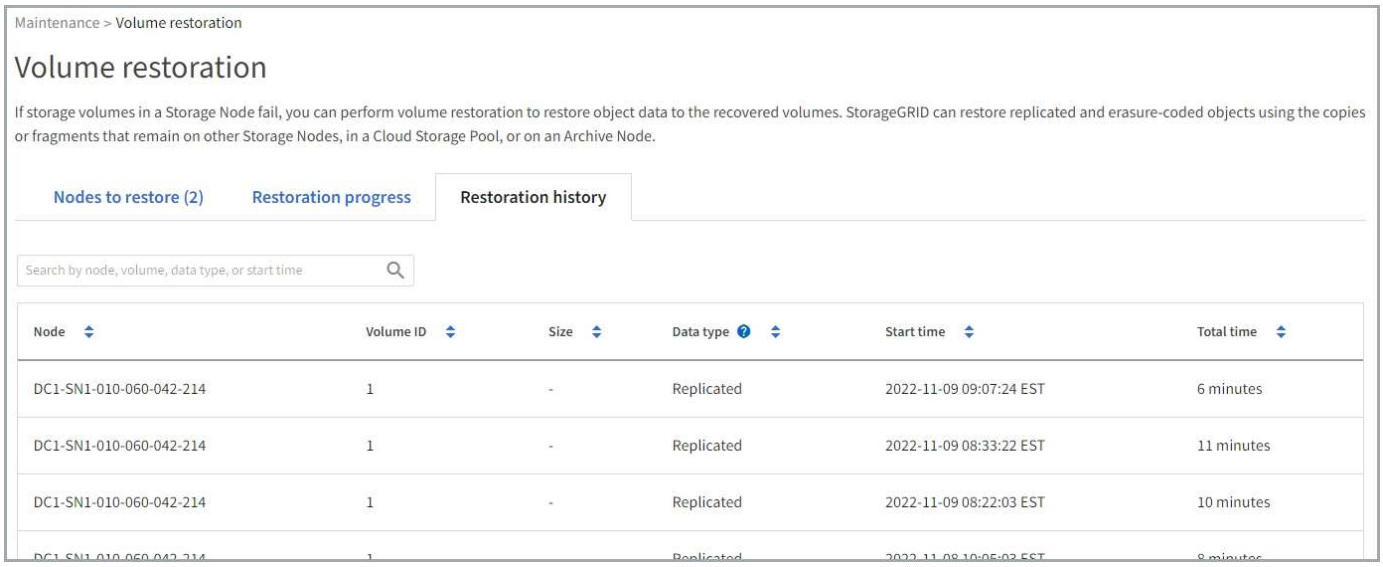

Monitor repair-data jobs

You can monitor the status of repair jobs by using the `repair-data` script from the command line.

These include jobs that you initiated manually, or jobs that StorageGRID initiated automatically as part of a decommission procedure.

> If you are running volume restoration jobs, monitor the progress and view a history of those jobs in the Grid Manager instead.

Monitor the status of `repair-data` jobs based on whether you use **replicated data**, **erasure-coded (EC) data**, or both.

Replicated data

- To get an estimated percent completion for the replicated repair, add the `show-replicated-repair-status` option to the repair-data command.

  ```
  repair-data show-replicated-repair-status
  ```

- To determine if repairs are complete:
  1. Select **NODES** > ***Storage Node being repaired*** > **ILM**.
  2. Review the attributes in the Evaluation section. When repairs are complete, the **Awaiting - All** attribute indicates 0 objects.
- To monitor the repair in more detail:
  1. Select **SUPPORT** > **Tools** > **Grid topology**.
  2. Select ***grid*** > ***Storage Node being repaired*** > **LDR** > **Data Store**.
  3. Use a combination of the following attributes to determine, as well as possible, if replicated repairs are complete.

     Cassandra inconsistencies might be present, and failed repairs aren't tracked.

     - **Repairs Attempted (XRPA)**: Use this attribute to track the progress of replicated repairs. This attribute increases each time a Storage Node tries to repair a high-risk object. When this attribute does not increase for a period longer than the current scan period (provided by the **Scan Period — Estimated** attribute), it means that ILM scanning found no high-risk objects that need to be repaired on any nodes.

       High-risk objects are objects that are at risk of being completely lost. This does not include objects that don't satisfy their ILM configuration.

     - **Scan Period — Estimated (XSCM)**: Use this attribute to estimate when a policy change will be applied to previously ingested objects. If the **Repairs Attempted** attribute does not increase for a period longer than the current scan period, it is probable that replicated repairs are done. Note that the scan period can change. The **Scan Period — Estimated (XSCM)** attribute applies to the entire grid and is the maximum of all node scan periods. You can query the **Scan Period — Estimated** attribute history for the grid to determine an appropriate time frame.

Erasure coded (EC) data

To monitor the repair of erasure-coded data and retry any requests that might have failed:

1. Determine the status of erasure-coded data repairs:
   - Select **SUPPORT** > **Tools** > **Metrics** to view the estimated time to completion and the completion percentage for the current job. Then, select **EC Overview** in the Grafana section. Look at the **Grid EC Job Estimated Time to Completion** and **Grid EC Job Percentage Completed** dashboards.
   - Use this command to see the status of a specific `repair-data` operation:

     ```
     repair-data show-ec-repair-status --repair-id repair ID
     ```

   - Use this command to list all repairs:

```
repair-data show-ec-repair-status
```

The output lists information, including `repair ID`, for all previously and currently running repairs.

2. If the output shows that the repair operation failed, use the `--repair-id` option to retry the repair.

   This command retries a failed node repair, using the repair ID 6949309319275667690:

   ```
   repair-data start-ec-node-repair --repair-id 6949309319275667690
   ```

   This command retries a failed volume repair, using the repair ID 6949309319275667690:

   ```
   repair-data start-ec-volume-repair --repair-id 6949309319275667690
   ```

Recover from Admin Node failures

Recover from Admin Node failures: Workflow

The recovery process for an Admin Node depends on whether it is the primary Admin Node or a non-primary Admin Node.

The high-level steps for recovering a primary or non-primary Admin Node are the same, although the details of the steps differ.

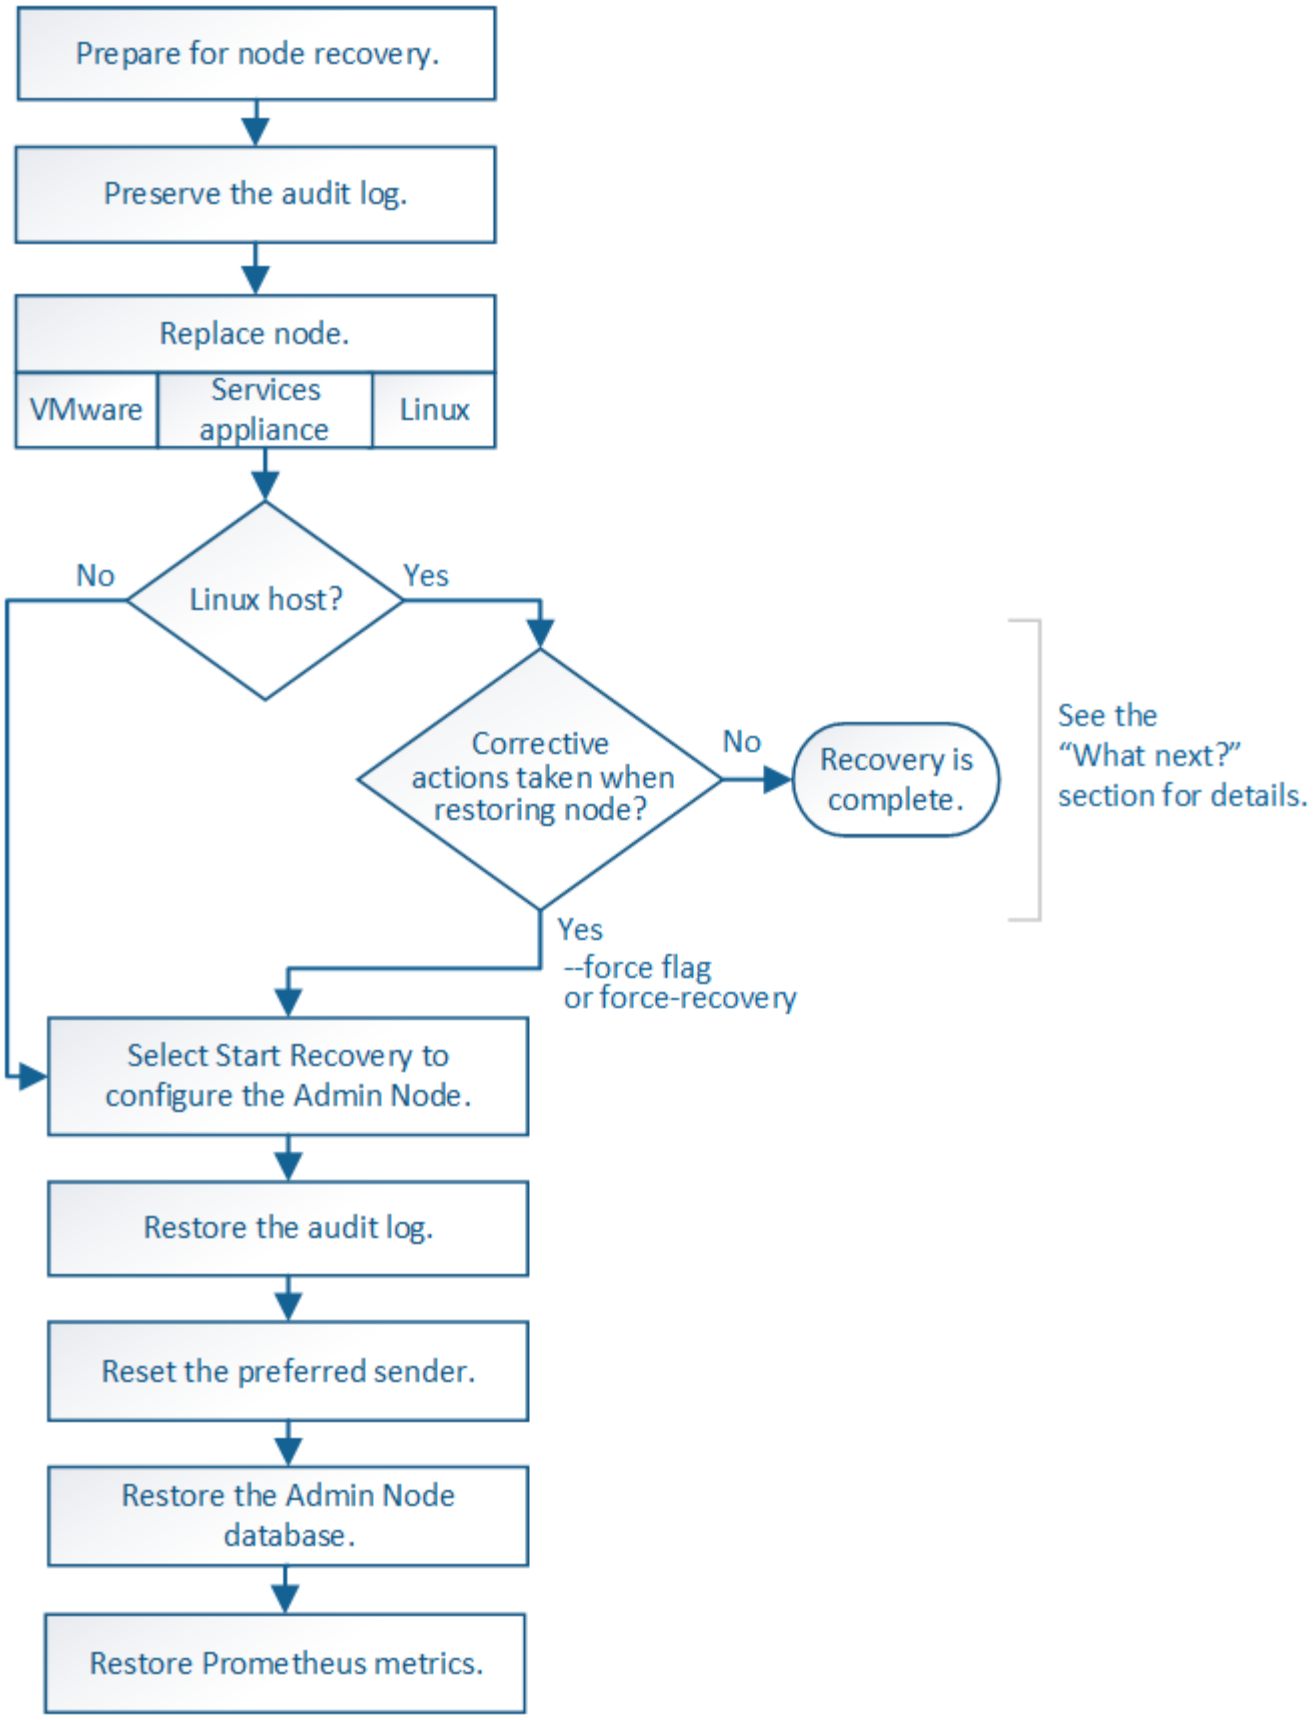

Always follow the correct recovery procedure for the Admin Node you are recovering. The procedures look the same at a high level, but differ in the details.

Choices

- Recover from primary Admin Node failures
- Recover from non-primary Admin Node failures

Recover from primary Admin Node failures

Recover from primary Admin Node failures: Overview

You must complete a specific set of tasks to recover from a primary Admin Node failure. The primary Admin Node hosts the Configuration Management Node (CMN) service for the grid.

A failed primary Admin Node should be replaced promptly. The Configuration Management Node (CMN) service on the primary Admin Node is responsible for issuing blocks of object identifiers for the grid. These identifiers are assigned to objects as they are ingested. New objects can't be ingested unless there are identifiers available. Object ingest can continue while the CMN is unavailable because approximately one month's supply of identifiers is cached in the grid. However, after cached identifiers are exhausted, no new objects can be added.

> ⓘ You must repair or replace a failed primary Admin Node within approximately a month or the grid might lose its ability to ingest new objects. The exact time period depends on your rate of object ingest: if you need a more accurate assessment of the time frame for your grid, contact technical support.

Copy audit logs from failed primary Admin Node

If you are able to copy audit logs from the failed primary Admin Node, you should preserve them to maintain the grid's record of system activity and usage. You can restore the preserved audit logs to the recovered primary Admin Node after it is up and running.

About this task

This procedure copies the audit log files from the failed Admin Node to a temporary location on a separate grid node. These preserved audit logs can then be copied to the replacement Admin Node. Audit logs aren't automatically copied to the new Admin Node.

Depending on the type of failure, you might not be able to copy audit logs from a failed Admin Node. If the deployment has only one Admin Node, the recovered Admin Node starts recording events to the audit log in a new empty file and previously recorded data is lost. If the deployment includes more than one Admin Node, you can recover the audit logs from another Admin Node.

> ⓘ If the audit logs aren't accessible on the failed Admin Node now, you might be able to access them later, for example, after host recovery.

Steps

1. Log in to the failed Admin Node if possible. Otherwise, log in to the primary Admin Node or another Admin Node, if available.
   a. Enter the following command: `ssh admin@grid_node_IP`
   b. Enter the password listed in the `Passwords.txt` file.
   c. Enter the following command to switch to root: `su -`
   d. Enter the password listed in the `Passwords.txt` file.

When you are logged in as root, the prompt changes from `$` to `#`.

2. Stop the AMS service to prevent it from creating a new log file:`service ams stop`
3. Rename the audit.log file so that it does not overwrite the existing file when you copy it to the recovered Admin Node.

   Rename audit.log to a unique numbered file name. For example, rename the audit.log file to `2023-10-25.txt.1`.

   ```
   cd /var/local/audit/export
   ls -l
   mv audit.log 2023-10-25.txt.1
   ```

4. Restart the AMS service: `service ams start`
5. Create the directory to copy all audit log files to a temporary location on a separate grid node: `ssh admin@grid_node_IP mkdir -p /var/local/tmp/saved-audit-logs`

   When prompted, enter the password for admin.

6. Copy all audit log files: `scp -p * admin@grid_node_IP:/var/local/tmp/saved-audit-logs`

   When prompted, enter the password for admin.

7. Log out as root: `exit`

Replace primary Admin Node

To recover a primary Admin Node, you must first replace the physical or virtual hardware.

You can replace a failed primary Admin Node with a primary Admin Node running on the same platform, or you can replace a primary Admin Node running on VMware or a Linux host with a primary Admin Node hosted on a services appliance.

Use the procedure that matches the replacement platform you select for the node. After you complete the node replacement procedure (which is suitable for all node types), that procedure will direct you to the next step for primary Admin Node recovery.

Replacement platform	Procedure
VMware	Replace a VMware node
Linux	Replace a Linux node
SG100 and SG1000 services appliances	Replace a services appliance

Replacement platform	Procedure
OpenStack	NetApp-provided virtual machine disk files and scripts for OpenStack are no longer supported for recovery operations. If you need to recover a node running in an OpenStack deployment, download the files for your Linux operating system. Then, follow the procedure for replacing a Linux node.

Configure replacement primary Admin Node

The replacement node must be configured as the primary Admin Node for your StorageGRID system.

Before you begin

- For primary Admin Nodes hosted on virtual machines, the virtual machine has been deployed, powered on, and initialized.
- For primary Admin Nodes hosted on a services appliance, you have replaced the appliance and have installed software. See the installation instructions for your appliance.
- You have the latest backup of the Recovery Package file (`sgws-recovery-package-id-revision.zip`).
- You have the provisioning passphrase.

Steps

1. Open your web browser and navigate to `https://primary_admin_node_ip`.

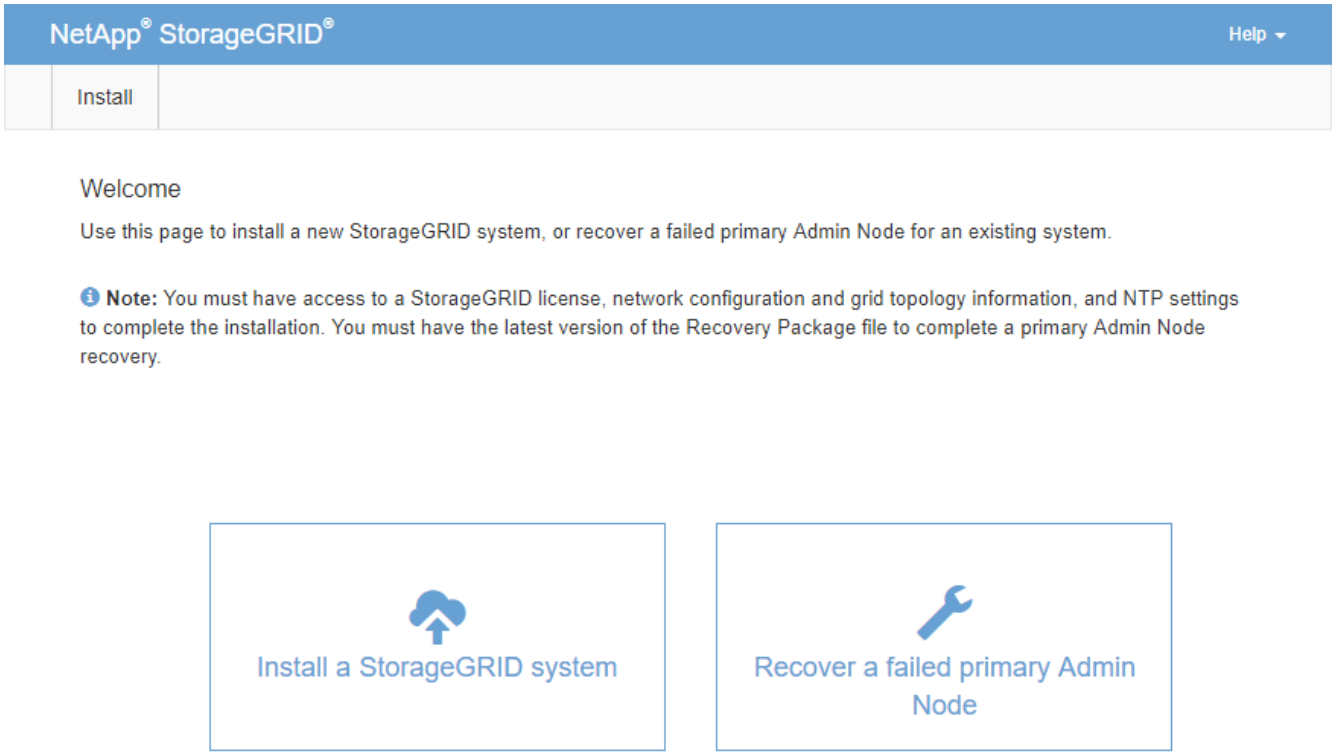

2. Click **Recover a failed primary Admin Node**.

3. Upload the most recent backup of the Recovery Package:
   a. Click **Browse**.
   b. Locate the most recent Recovery Package file for your StorageGRID system, and click **Open**.
4. Enter the provisioning passphrase.
5. Click **Start Recovery**.

   The recovery process begins. The Grid Manager might become unavailable for a few minutes as the required services start. When the recovery is complete, the sign in page is displayed.

6. If single sign-on (SSO) is enabled for your StorageGRID system and the relying party trust for the Admin Node you recovered was configured to use the default management interface certificate, update (or delete and recreate) the node's relying party trust in Active Directory Federation Services (AD FS). Use the new default server certificate that was generated during the Admin Node recovery process.

   > To configure a relying party trust, see Configure single sign-on. To access the default server certificate, log in to the command shell of the Admin Node. Go to the `/var/local/mgmt-api` directory, and select the `server.crt` file.

7. Determine if you need to apply a hotfix.
   a. Sign in to the Grid Manager using a supported web browser.
   b. Select **NODES**.
   c. From the list on the left, select the primary Admin Node.
   d. On the Overview tab, note the version displayed in the **Software Version** field.
   e. Select any other grid node.
   f. On the Overview tab, note the version displayed in the **Software Version** field.
      - If the versions displayed in the **Software Version** fields are the same, you don't need to apply a hotfix.
      - If the versions displayed in the **Software Version** fields are different, you must apply a hotfix to update the recovered primary Admin Node to the same version.

Restore audit log on recovered primary Admin Node

If you were able to preserve the audit log from the failed primary Admin Node, you can copy it to the primary Admin Node you are recovering.

Before you begin

- The recovered Admin Node is installed and running.
- You have copied the audit logs to another location after the original Admin Node failed.

About this task

If an Admin Node fails, audit logs saved to that Admin Node are potentially lost. It might be possible to preserve data from loss by copying audit logs from the failed Admin Node and then restoring these audit logs to the recovered Admin Node. Depending on the failure, it might not be possible to copy audit logs from the failed Admin Node. In that case, if the deployment has more than one Admin Node, you can recover audit logs from another Admin Node as audit logs are replicated to all Admin Nodes.

If there is only one Admin Node and the audit log can't be copied from the failed node, the recovered Admin

Node starts recording events to the audit log as if the installation is new.

You must recover an Admin Node as soon as possible to restore logging functionality.

> By default, audit information is sent to the audit log on Admin Nodes. You can skip these steps if either of the following applies:
>
> - You configured an external syslog server and audit logs are now being sent to the syslog server instead of to Admin Nodes.
> - You explicitly specified that audit messages should be saved only on the local nodes that generated them.
>
> See Configure audit messages and log destinations for details.

Steps

1. Log in to the recovered Admin Node:
   a. Enter the following command: `ssh admin@recovery_Admin_Node_IP`
   b. Enter the password listed in the `Passwords.txt` file.
   c. Enter the following command to switch to root: `su -`
   d. Enter the password listed in the `Passwords.txt` file.

   After you are logged in as root, the prompt changes from `$` to `#`.
2. Check which audit files have been preserved: `cd /var/local/audit/export`
3. Copy the preserved audit log files to the recovered Admin Node: `scp admin@grid_node_IP:/var/local/tmp/saved-audit-logs/YYYY* .`

   When prompted, enter the password for admin.
4. For security, delete the audit logs from the failed grid node after verifying that they have been copied successfully to the recovered Admin Node.
5. Update the user and group settings of the audit log files on the recovered Admin Node: `chown ams-user: bycast *`
6. Log out as root: `exit`

You must also restore any pre-existing client access to the audit share. For more information, see Configure audit client access.

Restore Admin Node database when recovering primary Admin Node

If you want to retain the historical information about attributes, alarms, and alerts on a primary Admin Node that has failed, you can restore the Admin Node database. You can only restore this database if your StorageGRID system includes another Admin Node.

Before you begin

- The recovered Admin Node is installed and running.
- The StorageGRID system includes at least two Admin Nodes.

- You have the `Passwords.txt` file.
- You have the provisioning passphrase.

About this task

If an Admin Node fails, the historical information stored in its Admin Node database is lost. This database includes the following information:

- Alert history
- Alarm history
- Historical attribute data, which is used in the charts and text reports available from the **SUPPORT** > **Tools** > **Grid topology** page.

When you recover an Admin Node, the software installation process creates an empty Admin Node database on the recovered node. However, the new database only includes information for servers and services that are currently part of the system or added later.

If you restored a primary Admin Node and your StorageGRID system has another Admin Node, you can restore the historical information by copying the Admin Node database from a non-primary Admin Node (the *source Admin Node*) to the recovered primary Admin Node. If your system has only a primary Admin Node, you can't restore the Admin Node database.

> Copying the Admin Node database might take several hours. Some Grid Manager features will be unavailable while services are stopped on the source Admin Node.

Steps

1. Log in to the source Admin Node:
   a. Enter the following command: `ssh admin@grid_node_IP`
   b. Enter the password listed in the `Passwords.txt` file.
   c. Enter the following command to switch to root: `su -`
   d. Enter the password listed in the `Passwords.txt` file.
2. From the source Admin Node, stop the MI service: `service mi stop`
3. From the source Admin Node, stop the Management Application Program Interface (mgmt-api) service: `service mgmt-api stop`
4. Complete the following steps on the recovered Admin Node:
   a. Log in to the recovered Admin Node:
      i. Enter the following command: `ssh admin@grid_node_IP`
      ii. Enter the password listed in the `Passwords.txt` file.
      iii. Enter the following command to switch to root: `su -`
      iv. Enter the password listed in the `Passwords.txt` file.
   b. Stop the MI service: `service mi stop`
   c. Stop the mgmt-api service: `service mgmt-api stop`
   d. Add the SSH private key to the SSH agent. Enter:`ssh-add`

e. Enter the SSH Access Password listed in the `Passwords.txt` file.

f. Copy the database from the source Admin Node to the recovered Admin Node: `/usr/local/mi/bin/mi-clone-db.sh Source_Admin_Node_IP`

g. When prompted, confirm that you want to overwrite the MI database on the recovered Admin Node.

   The database and its historical data are copied to the recovered Admin Node. When the copy operation is done, the script starts the recovered Admin Node.

h. When you no longer require passwordless access to other servers, remove the private key from the SSH agent. Enter:`ssh-add -D`

5. Restart the services on the source Admin Node: `service servermanager start`

Restore Prometheus metrics when recovering primary Admin Node

Optionally, you can retain the historical metrics maintained by Prometheus on a primary Admin Node that has failed. The Prometheus metrics can only be restored if your StorageGRID system includes another Admin Node.

Before you begin

- The recovered Admin Node is installed and running.
- The StorageGRID system includes at least two Admin Nodes.
- You have the `Passwords.txt` file.
- You have the provisioning passphrase.

About this task

If an Admin Node fails, the metrics maintained in the Prometheus database on the Admin Node are lost. When you recover the Admin Node, the software installation process creates a new Prometheus database. After the recovered Admin Node is started, it records metrics as if you had performed a new installation of the StorageGRID system.

If you restored a primary Admin Node and your StorageGRID system has another Admin Node, you can restore the historical metrics by copying the Prometheus database from a non-primary Admin Node (the *source Admin Node*) to the recovered primary Admin Node. If your system has only a primary Admin Node, you can't restore the Prometheus database.

> Copying the Prometheus database might take an hour or more. Some Grid Manager features will be unavailable while services are stopped on the source Admin Node.

Steps

1. Log in to the source Admin Node:
   a. Enter the following command: `ssh admin@grid_node_IP`
   b. Enter the password listed in the `Passwords.txt` file.
   c. Enter the following command to switch to root: `su -`
   d. Enter the password listed in the `Passwords.txt` file.
2. From the source Admin Node, stop the Prometheus service: `service prometheus stop`

3. Complete the following steps on the recovered Admin Node:
   a. Log in to the recovered Admin Node:
      i. Enter the following command: `ssh admin@grid_node_IP`
      ii. Enter the password listed in the `Passwords.txt` file.
      iii. Enter the following command to switch to root: `su -`
      iv. Enter the password listed in the `Passwords.txt` file.
   b. Stop the Prometheus service: `service prometheus stop`
   c. Add the SSH private key to the SSH agent. Enter:`ssh-add`
   d. Enter the SSH Access Password listed in the `Passwords.txt` file.
   e. Copy the Prometheus database from the source Admin Node to the recovered Admin Node: `/usr/local/prometheus/bin/prometheus-clone-db.sh Source_Admin_Node_IP`
   f. When prompted, press **Enter** to confirm that you want to destroy the new Prometheus database on the recovered Admin Node.

      The original Prometheus database and its historical data are copied to the recovered Admin Node. When the copy operation is done, the script starts the recovered Admin Node. The following status appears:

      Database cloned, starting services

   g. When you no longer require passwordless access to other servers, remove the private key from the SSH agent. Enter:`ssh-add -D`
4. Restart the Prometheus service on the source Admin Node.`service prometheus start`

Recover from non-primary Admin Node failures

Recover from non-primary Admin Node failures: Overview

You must complete the following tasks to recover from a non-primary Admin Node failure. One Admin Node hosts the Configuration Management Node (CMN) service and is known as the primary Admin Node. Although you can have multiple Admin Nodes, each StorageGRID system includes only one primary Admin Node. All other Admin Nodes are non-primary Admin Nodes.

Copy audit logs from failed non-primary Admin Node

If you are able to copy audit logs from the failed Admin Node, you should preserve them to maintain the grid's record of system activity and usage. You can restore the preserved audit logs to the recovered non-primary Admin Node after it is up and running.

This procedure copies the audit log files from the failed Admin Node to a temporary location on a separate grid node. These preserved audit logs can then be copied to the replacement Admin Node. Audit logs aren't automatically copied to the new Admin Node.

Depending on the type of failure, you might not be able to copy audit logs from a failed Admin Node. If the deployment has only one Admin Node, the recovered Admin Node starts recording events to the audit log in a

new empty file and previously recorded data is lost. If the deployment includes more than one Admin Node, you can recover the audit logs from another Admin Node.

> If the audit logs aren't accessible on the failed Admin Node now, you might be able to access them later, for example, after host recovery.

1. Log in to the failed Admin Node if possible. Otherwise, log in to the primary Admin Node or another Admin Node, if available.
   a. Enter the following command: `ssh admin@grid_node_IP`
   b. Enter the password listed in the `Passwords.txt` file.
   c. Enter the following command to switch to root: `su -`
   d. Enter the password listed in the `Passwords.txt` file.

   When you are logged in as root, the prompt changes from `$` to `#`.
2. Stop the AMS service to prevent it from creating a new log file:`service ams stop`
3. Rename the audit.log file so that it does not overwrite the existing file when you copy it to the recovered Admin Node.

   Rename audit.log to a unique numbered file name. For example, rename the audit.log file to `2023-10-25.txt.1`.

   ```
   cd /var/local/audit/export
   ls -l
   mv audit.log 2023-10-25.txt.1
   ```
4. Restart the AMS service: `service ams start`
5. Create the directory to copy all audit log files to a temporary location on a separate grid node: `ssh admin@grid_node_IP mkdir -p /var/local/tmp/saved-audit-logs`

   When prompted, enter the password for admin.
6. Copy all audit log files: `scp -p * admin@grid_node_IP:/var/local/tmp/saved-audit-logs`

   When prompted, enter the password for admin.
7. Log out as root: `exit`

Replace non-primary Admin Node

To recover a non-primary Admin Node, you first must replace the physical or virtual hardware.

You can replace a failed non-primary Admin Node with a non-primary Admin Node running on the same platform, or you can replace a non-primary Admin Node running on VMware or a Linux host with a non-primary Admin Node hosted on a services appliance.

Use the procedure that matches the replacement platform you select for the node. After you complete the node replacement procedure (which is suitable for all node types), that procedure will direct you to the next step for non-primary Admin Node recovery.

Replacement platform	Procedure
VMware	Replace a VMware node
Linux	Replace a Linux node
SG100 and SG1000 services appliances	Replace a services appliance
OpenStack	NetApp-provided virtual machine disk files and scripts for OpenStack are no longer supported for recovery operations. If you need to recover a node running in an OpenStack deployment, download the files for your Linux operating system. Then, follow the procedure for replacing a Linux node.

Select Start Recovery to configure non-primary Admin Node

After replacing a non-primary Admin Node, you must select Start Recovery in the Grid Manager to configure the new node as a replacement for the failed node.

Before you begin

- You are signed in to the Grid Manager using a supported web browser.
- You have the Maintenance or Root access permission.
- You have the provisioning passphrase.
- You have deployed and configured the replacement node.

Steps

1. From the Grid Manager, select **MAINTENANCE** > **Tasks** > **Recovery**.
2. Select the grid node you want to recover in the Pending Nodes list.

   Nodes appear in the list after they fail, but you can't select a node until it has been reinstalled and is ready for recovery.

3. Enter the **Provisioning Passphrase**.
4. Click **Start Recovery**.

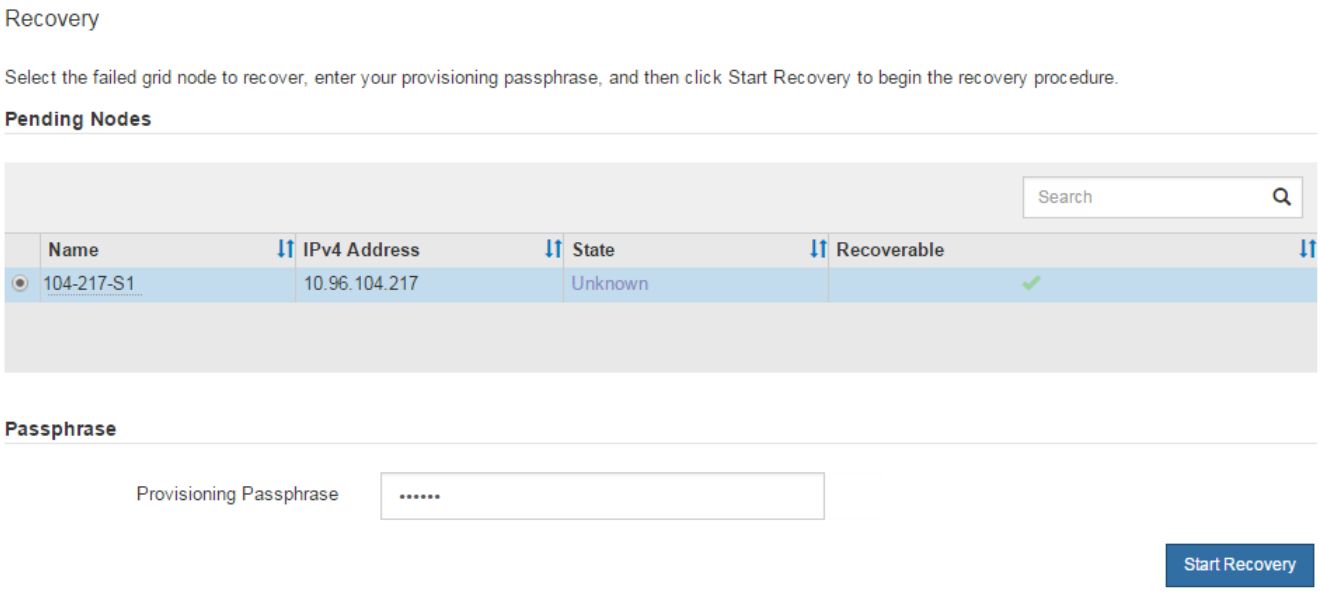

5. Monitor the progress of the recovery in the Recovering Grid Node table.

> While the recovery procedure is running, you can click **Reset** to start a new recovery. A dialog box appears, indicating that the node will be left in an indeterminate state if you reset the procedure.

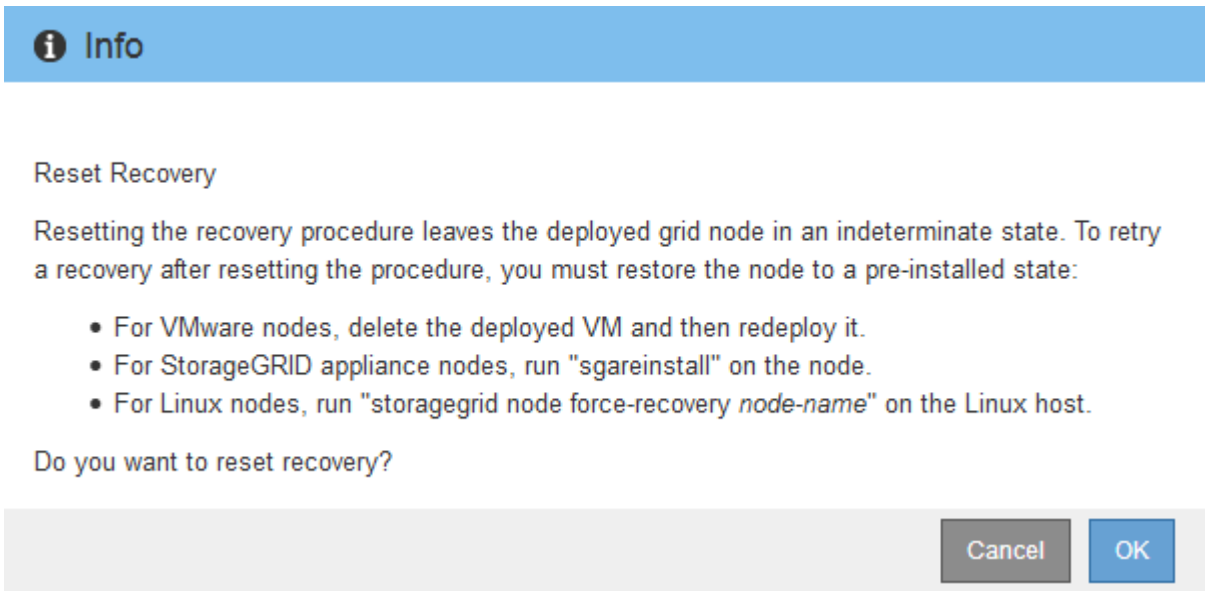

If you want to retry the recovery after resetting the procedure, you must restore the node to a pre-installed state, as follows:

- **VMware**: Delete the deployed virtual grid node. Then, when you are ready to restart the recovery, redeploy the node.
- **Linux**: Restart the node by running this command on the Linux host: `storagegrid node force-recovery node-name`
- **Appliance**: If you want to retry the recovery after resetting the procedure, you must restore the appliance node to a pre-installed state by running `sgareinstall` on the node. See Prepare appliance for reinstallation (platform replacement only).

6. If single sign-on (SSO) is enabled for your StorageGRID system and the relying party trust for the Admin Node you recovered was configured to use the default management interface certificate, update (or delete and recreate) the node's relying party trust in Active Directory Federation Services (AD FS). Use the new default server certificate that was generated during the Admin Node recovery process.

> To configure a relying party trust, see Configure single sign-on. To access the default server certificate, log in to the command shell of the Admin Node. Go to the `/var/local/mgmt-api` directory, and select the `server.crt` file.

Restore audit log on recovered non-primary Admin Node

If you were able to preserve the audit log from the failed non-primary Admin Node, so that historical audit log information is retained, you can copy it to the non-primary Admin Node you are recovering.

Before you begin

- The recovered Admin Node is installed and running.
- You have copied the audit logs to another location after the original Admin Node failed.

About this task

If an Admin Node fails, audit logs saved to that Admin Node are potentially lost. It might be possible to preserve data from loss by copying audit logs from the failed Admin Node and then restoring these audit logs to the recovered Admin Node. Depending on the failure, it might not be possible to copy audit logs from the failed Admin Node. In that case, if the deployment has more than one Admin Node, you can recover audit logs from another Admin Node as audit logs are replicated to all Admin Nodes.

If there is only one Admin Node and the audit log can't be copied from the failed node, the recovered Admin Node starts recording events to the audit log as if the installation is new.

You must recover an Admin Node as soon as possible to restore logging functionality.

> By default, audit information is sent to the audit log on Admin Nodes. You can skip these steps if either of the following applies:
>
> - You configured an external syslog server and audit logs are now being sent to the syslog server instead of to Admin Nodes.
> - You explicitly specified that audit messages should be saved only on the local nodes that generated them.
>
> See Configure audit messages and log destinations for details.

Steps

1. Log in to the recovered Admin Node:
   a. Enter the following command: + `ssh admin@recovery_Admin_Node_IP`
   b. Enter the password listed in the `Passwords.txt` file.
   c. Enter the following command to switch to root: `su -`
   d. Enter the password listed in the `Passwords.txt` file.

After you are logged in as root, the prompt changes from `$` to `#`.

2. Check which audit files have been preserved:

   ```
   cd /var/local/audit/export
   ```

3. Copy the preserved audit log files to the recovered Admin Node:

   ```
   scp admin@grid_node_IP:/var/local/tmp/saved-audit-logs/YYYY*
   ```

   When prompted, enter the password for admin.

4. For security, delete the audit logs from the failed grid node after verifying that they have been copied successfully to the recovered Admin Node.
5. Update the user and group settings of the audit log files on the recovered Admin Node:

   ```
   chown ams-user:bycast *
   ```

6. Log out as root: `exit`

You must also restore any pre-existing client access to the audit share. For more information, see Configure audit client access.

Restore Admin Node database when recovering non-primary Admin Node

If you want to retain the historical information about attributes, alarms, and alerts on a non-primary Admin Node that has failed, you can restore the Admin Node database from the primary Admin Node.

Before you begin

- The recovered Admin Node is installed and running.
- The StorageGRID system includes at least two Admin Nodes.
- You have the `Passwords.txt` file.
- You have the provisioning passphrase.

About this task

If an Admin Node fails, the historical information stored in its Admin Node database is lost. This database includes the following information:

- Alert history
- Alarm history
- Historical attribute data, which is used in the charts and text reports available from the **SUPPORT** > **Tools** > **Grid topology** page.

When you recover an Admin Node, the software installation process creates an empty Admin Node database on the recovered node. However, the new database only includes information for servers and services that are currently part of the system or added later.

If you restored a non-primary Admin Node, you can restore the historical information by copying the Admin Node database from the primary Admin Node (the *source Admin Node*) to the recovered node.

Copying the Admin Node database might take several hours. Some Grid Manager features will be unavailable while services are stopped on the source node.

Steps

1. Log in to the source Admin Node:
   a. Enter the following command: `ssh admin@grid_node_IP`
   b. Enter the password listed in the `Passwords.txt` file.
   c. Enter the following command to switch to root: `su -`
   d. Enter the password listed in the `Passwords.txt` file.
2. Run the following command from the source Admin Node. Then, enter the provisioning passphrase if prompted. `recover-access-points`
3. From the source Admin Node, stop the MI service: `service mi stop`
4. From the source Admin Node, stop the Management Application Program Interface (mgmt-api) service: `service mgmt-api stop`
5. Complete the following steps on the recovered Admin Node:
   a. Log in to the recovered Admin Node:
      i. Enter the following command: `ssh admin@grid_node_IP`
      ii. Enter the password listed in the `Passwords.txt` file.
      iii. Enter the following command to switch to root: `su -`
      iv. Enter the password listed in the `Passwords.txt` file.
   b. Stop the MI service: `service mi stop`
   c. Stop the mgmt-api service: `service mgmt-api stop`
   d. Add the SSH private key to the SSH agent. Enter:`ssh-add`
   e. Enter the SSH Access Password listed in the `Passwords.txt` file.
   f. Copy the database from the source Admin Node to the recovered Admin Node: `/usr/local/mi/bin/mi-clone-db.sh Source_Admin_Node_IP`
   g. When prompted, confirm that you want to overwrite the MI database on the recovered Admin Node.

      The database and its historical data are copied to the recovered Admin Node. When the copy operation is done, the script starts the recovered Admin Node.

   h. When you no longer require passwordless access to other servers, remove the private key from the SSH agent. Enter:`ssh-add -D`
6. Restart the services on the source Admin Node: `service servermanager start`

Restore Prometheus metrics when recovering non-primary Admin Node

Optionally, you can retain the historical metrics maintained by Prometheus on a non-primary Admin Node that has failed.

Before you begin

- The recovered Admin Node is installed and running.
- The StorageGRID system includes at least two Admin Nodes.
- You have the `Passwords.txt` file.
- You have the provisioning passphrase.

About this task

If an Admin Node fails, the metrics maintained in the Prometheus database on the Admin Node are lost. When you recover the Admin Node, the software installation process creates a new Prometheus database. After the recovered Admin Node is started, it records metrics as if you had performed a new installation of the StorageGRID system.

If you restored a non-primary Admin Node, you can restore the historical metrics by copying the Prometheus database from the primary Admin Node (the *source Admin Node*) to the recovered Admin Node.

> Copying the Prometheus database might take an hour or more. Some Grid Manager features will be unavailable while services are stopped on the source Admin Node.

Steps

1. Log in to the source Admin Node:
   a. Enter the following command: `ssh admin@grid_node_IP`
   b. Enter the password listed in the `Passwords.txt` file.
   c. Enter the following command to switch to root: `su -`
   d. Enter the password listed in the `Passwords.txt` file.
2. From the source Admin Node, stop the Prometheus service: `service prometheus stop`
3. Complete the following steps on the recovered Admin Node:
   a. Log in to the recovered Admin Node:
      i. Enter the following command: `ssh admin@grid_node_IP`
      ii. Enter the password listed in the `Passwords.txt` file.
      iii. Enter the following command to switch to root: `su -`
      iv. Enter the password listed in the `Passwords.txt` file.
   b. Stop the Prometheus service: `service prometheus stop`
   c. Add the SSH private key to the SSH agent. Enter:`ssh-add`
   d. Enter the SSH Access Password listed in the `Passwords.txt` file.
   e. Copy the Prometheus database from the source Admin Node to the recovered Admin Node: `/usr/local/prometheus/bin/prometheus-clone-db.sh Source_Admin_Node_IP`
   f. When prompted, press **Enter** to confirm that you want to destroy the new Prometheus database on the recovered Admin Node.

      The original Prometheus database and its historical data are copied to the recovered Admin Node. When the copy operation is done, the script starts the recovered Admin Node. The following status appears:

      Database cloned, starting services

g. When you no longer require passwordless access to other servers, remove the private key from the SSH agent. Enter:`ssh-add -D`

4. Restart the Prometheus service on the source Admin Node.`service prometheus start`

Recover from Gateway Node failures

Recover from Gateway Node failures: Workflow

You must complete a sequence of tasks in exact order to recover from a Gateway Node failure.

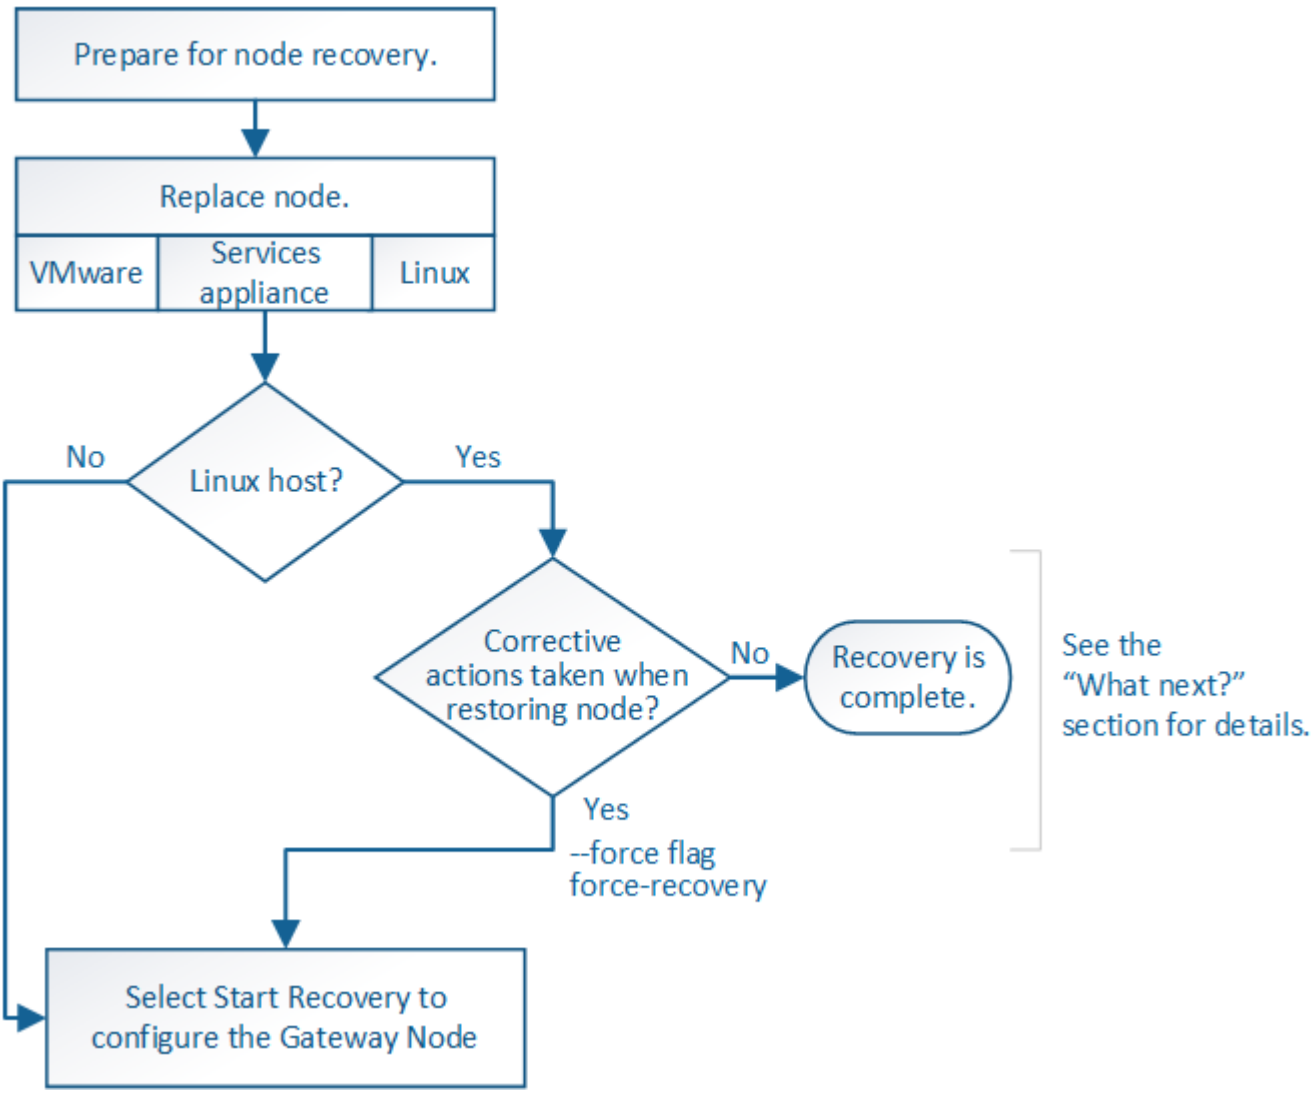

Replace Gateway Node

You can replace a failed Gateway Node with a Gateway Node running on the same physical or virtual hardware, or you can replace a Gateway Node running on VMware or a Linux host with a Gateway Node hosted on a services appliance.

The node replacement procedure you must follow depends on which platform will be used by the replacement node. After you complete the node replacement procedure (which is suitable for all node types), that procedure will direct you to the next step for Gateway Node recovery.

Replacement platform	Procedure
VMware	Replace a VMware node
Linux	Replace a Linux node
SG100 and SG1000 services appliances	Replace a services appliance
OpenStack	NetApp-provided virtual machine disk files and scripts for OpenStack are no longer supported for recovery operations. If you need to recover a node running in an OpenStack deployment, download the files for your Linux operating system. Then, follow the procedure for replacing a Linux node.

Select Start Recovery to configure Gateway Node

After replacing a Gateway Node, you must select Start Recovery in the Grid Manager to configure the new node as a replacement for the failed node.

Before you begin

- You are signed in to the Grid Manager using a supported web browser.
- You have the Maintenance or Root access permission.
- You have the provisioning passphrase.
- You have deployed and configured the replacement node.

Steps

1. From the Grid Manager, select **MAINTENANCE** > **Tasks** > **Recovery**.
2. Select the grid node you want to recover in the Pending Nodes list.

   Nodes appear in the list after they fail, but you can't select a node until it has been reinstalled and is ready for recovery.

3. Enter the **Provisioning Passphrase**.
4. Click **Start Recovery**.

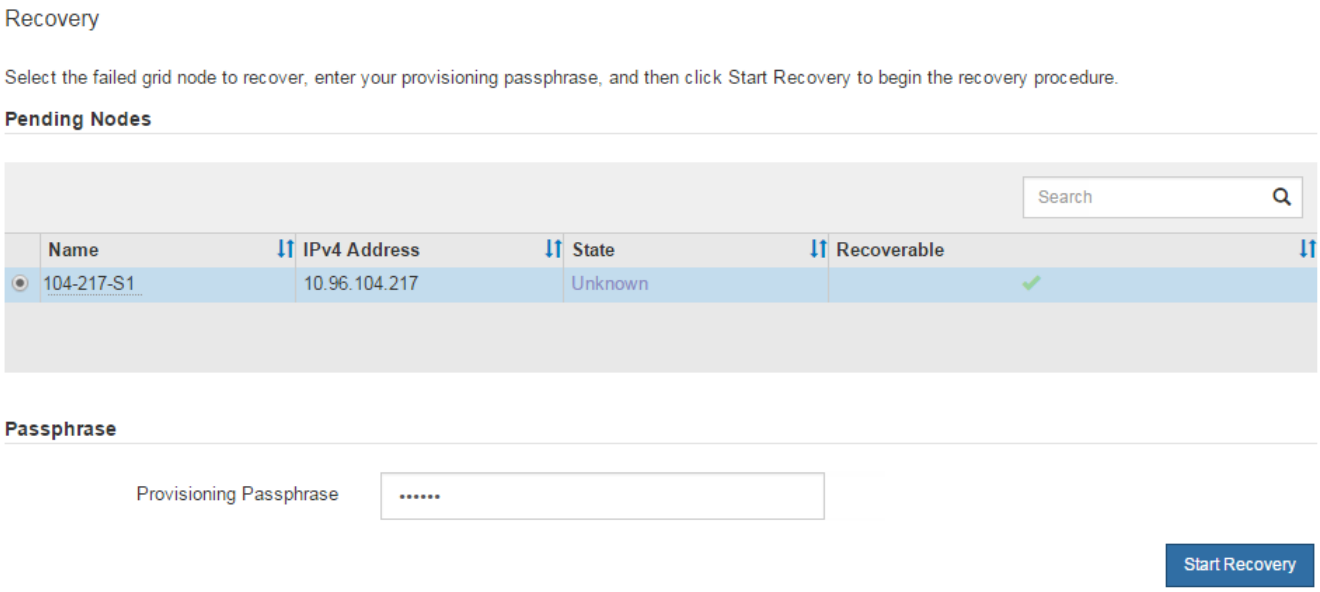

5. Monitor the progress of the recovery in the Recovering Grid Node table.

> While the recovery procedure is running, you can click **Reset** to start a new recovery. A dialog box appears, indicating that the node will be left in an indeterminate state if you reset the procedure.

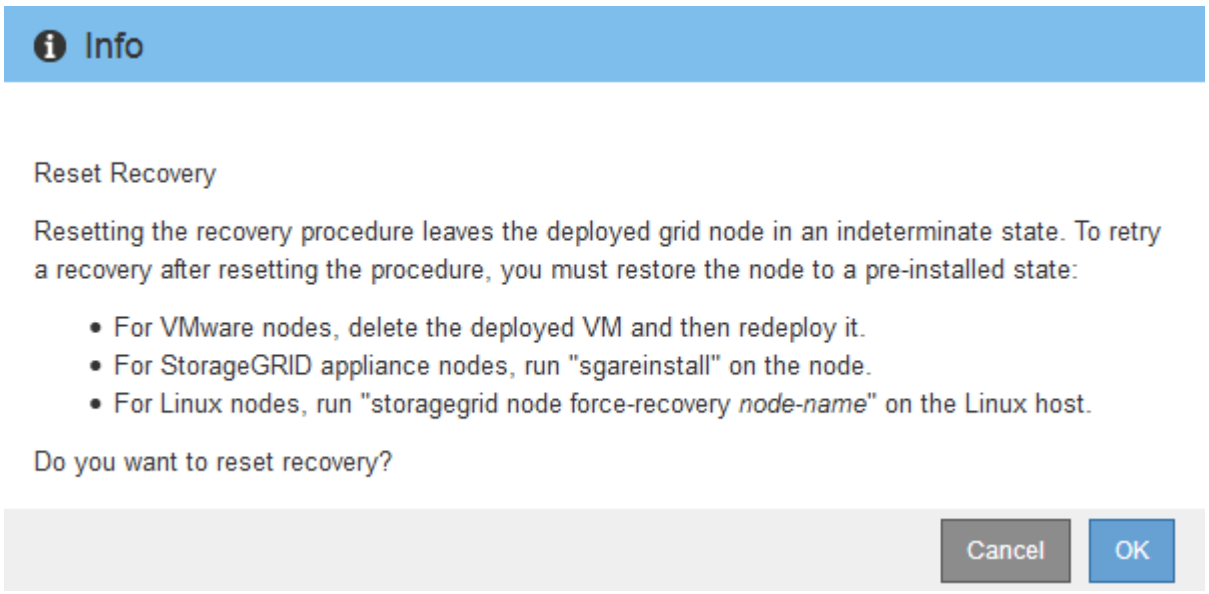

If you want to retry the recovery after resetting the procedure, you must restore the node to a pre-installed state, as follows:

- **VMware**: Delete the deployed virtual grid node. Then, when you are ready to restart the recovery, redeploy the node.
- **Linux**: Restart the node by running this command on the Linux host: `storagegrid node force-recovery node-name`
- **Appliance**: If you want to retry the recovery after resetting the procedure, you must restore the appliance node to a pre-installed state by running `sgareinstall` on the node. See Prepare appliance for reinstallation (platform replacement only).

Recover from Archive Node failures

Recover from Archive Node failures: Workflow

You must complete a sequence of tasks in exact order to recover from an Archive Node failure.

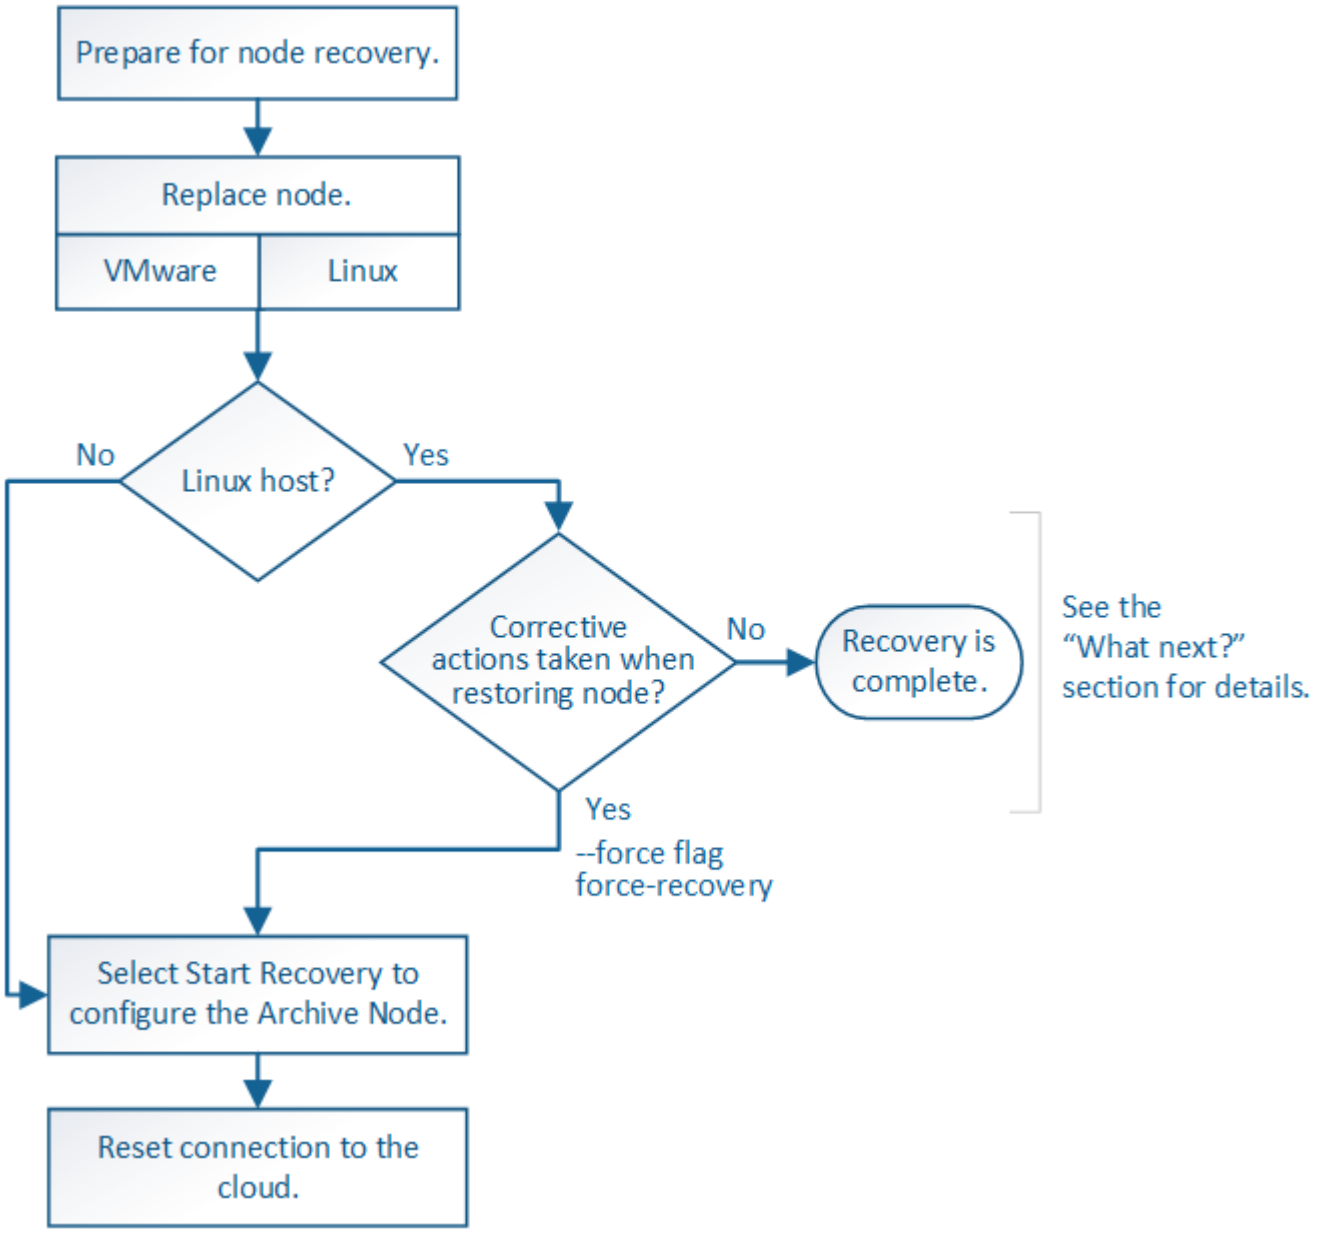

Archive Node recovery is affected by the following issues:

- If the ILM policy is configured to replicate a single copy.

  In a StorageGRID system that is configured to make a single copy of objects, an Archive Node failure might result in an unrecoverable loss of data. If there is a failure, all such objects are lost; however, you must still perform recovery procedures to "clean up" your StorageGRID system and purge lost object information from the database.

- If an Archive Node failure occurs during Storage Node recovery.

If the Archive Node fails while processing bulk retrievals as part of a Storage Node recovery, you must repeat the procedure to recover copies of object data to the Storage Node from the beginning to ensure that all object data retrieved from the Archive Node is restored to the Storage Node.

Replace Archive Node

To recover an Archive Node, you must first replace the node.

You must select the node replacement procedure for your platform. The steps to replace a node are the same for all types of grid nodes.

Platform	Procedure
VMware	Replace a VMware node
Linux	Replace a Linux node
OpenStack	NetApp-provided virtual machine disk files and scripts for OpenStack are no longer supported for recovery operations. If you need to recover a node running in an OpenStack deployment, download the files for your Linux operating system. Then, follow the procedure for replacing a Linux node.

Select Start Recovery to configure Archive Node

After replacing an Archive Node, you must select Start Recovery in the Grid Manager to configure the new node as a replacement for the failed node.

Before you begin

- You are signed in to the Grid Manager using a supported web browser.
- You have the Maintenance or Root access permission.
- You have the provisioning passphrase.
- You have deployed and configured the replacement node.

Steps

1. From the Grid Manager, select **MAINTENANCE** > **Tasks** > **Recovery**.
2. Select the grid node you want to recover in the Pending Nodes list.

   Nodes appear in the list after they fail, but you can't select a node until it has been reinstalled and is ready for recovery.

3. Enter the **Provisioning Passphrase**.
4. Click **Start Recovery**.

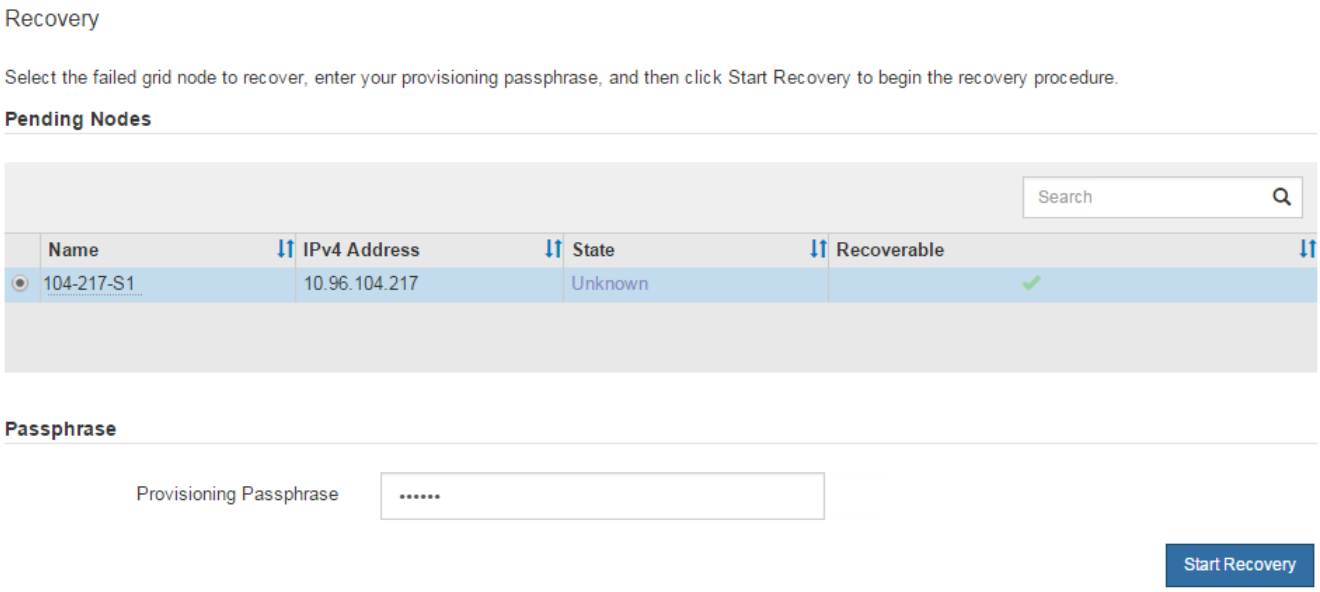

5. Monitor the progress of the recovery in the Recovering Grid Node table.

> While the recovery procedure is running, you can click **Reset** to start a new recovery. A dialog box appears, indicating that the node will be left in an indeterminate state if you reset the procedure.

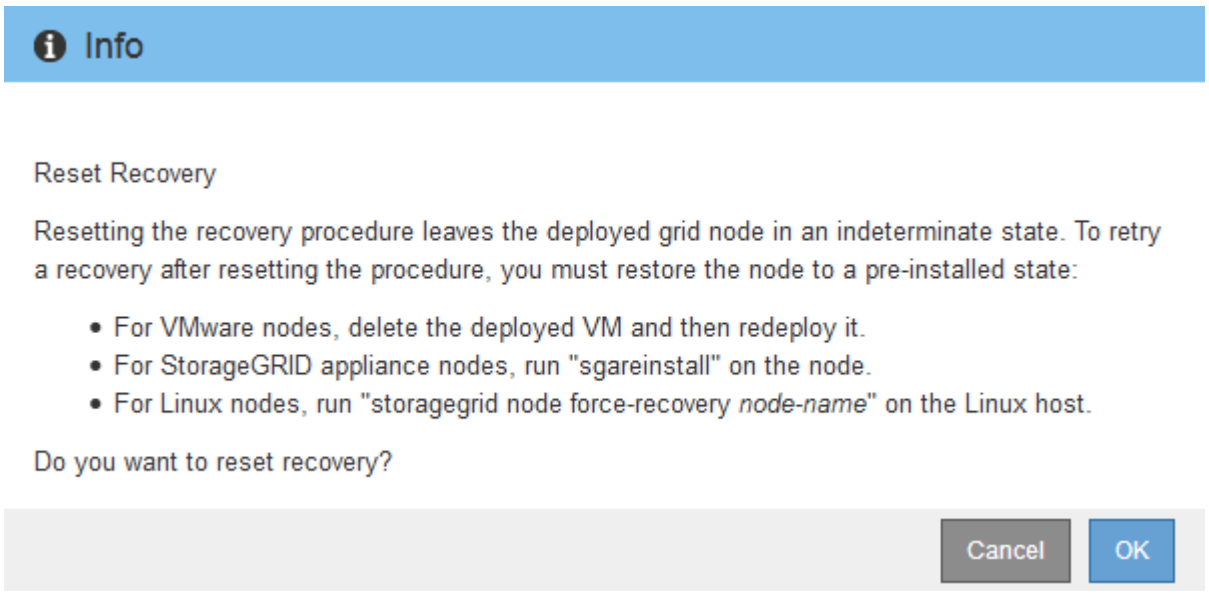

If you want to retry the recovery after resetting the procedure, you must restore the node to a pre-installed state, as follows:

- **VMware**: Delete the deployed virtual grid node. Then, when you are ready to restart the recovery, redeploy the node.
- **Linux**: Restart the node by running this command on the Linux host: `storagegrid node force-recovery node-name`

Reset Archive Node connection to the cloud

After you recover an Archive Node that targets the cloud through the S3 API, you need to modify configuration settings to reset connections. An Outbound Replication Status (ORSU) alarm is triggered if the Archive Node is unable to retrieve object data.

> If your Archive Node connects to external storage through TSM middleware, then the node resets itself automatically and you don't need to reconfigure.

Before you begin

You are signed in to the Grid Manager using a supported web browser.

Steps

1. Select **SUPPORT** > **Tools** > **Grid topology**.
2. Select **Archive Node** > **ARC** > **Target**.
3. Edit the **Access Key** field by entering an incorrect value and click **Apply Changes**.
4. Edit the **Access Key** field by entering the correct value and click **Apply Changes**.

All grid node types: Replace VMware node

When you recover a failed StorageGRID node that was hosted on VMware, you remove the failed node and deploy a recovery node.

Before you begin

You have determined that the virtual machine can't be restored and must be replaced.

About this task

You use the VMware vSphere Web Client to first remove the virtual machine associated with the failed grid node. Then, you can deploy a new virtual machine.

This procedure is only one step in the grid node recovery process. The node removal and deployment procedure is the same for all VMware nodes, including Admin Nodes, Storage Nodes, Gateway Nodes, and Archive Nodes.

Steps

1. Log in to VMware vSphere Web Client.
2. Navigate to the failed grid node virtual machine.
3. Make a note of all of the information required to deploy the recovery node.
   a. Right-click the virtual machine, select the **Edit Settings** tab, and note the settings in use.
   b. Select the **vApp Options** tab to view and record the grid node network settings.
4. If the failed grid node is a Storage Node, determine if any of the virtual hard disks used for data storage are undamaged and preserve them for reattachment to the recovered grid node.
5. Power off the virtual machine.
6. Select **Actions** > **All vCenter Actions** > **Delete from Disk** to delete the virtual machine.
7. Deploy a new virtual machine to be the replacement node, and connect it to one or more StorageGRID networks. For instructions see Deploying a StorageGRID node as a virtual machine.

When you deploy the node, you can optionally remap node ports or increase CPU or memory settings.

> After deploying the new node, you can add new virtual disks according to your storage requirements, reattach any virtual hard disks preserved from the previously removed failed grid node, or both.

8. Complete the node recovery procedure, based on the type of node you are recovering.

Type of node	Go to
Primary Admin Node	Configure replacement primary Admin Node
Non-primary Admin Node	Select Start Recovery to configure non-primary Admin Node
Gateway Node	Select Start Recovery to configure Gateway Node
Storage Node	Select Start Recovery to configure Storage Node
Archive Node	Select Start Recovery to configure Archive Node

All grid node types: Replace Linux node

All grid node types: Replace Linux node

If a failure requires that you deploy one or more new physical or virtual hosts or reinstall Linux on an existing host, deploy and configure the replacement host before you can recover the grid node. This procedure is one step of the grid node recovery process for all types of grid nodes.

"Linux" refers to a Red Hat® Enterprise Linux®, Ubuntu®, CentOS, or Debian® deployment. Use the NetApp Interoperability Matrix Tool (IMT) to get a list of supported versions.

This procedure is only performed as one step in the process of recovering software-based Storage Nodes, primary or non-primary Admin Nodes, Gateway Nodes, or Archive Nodes. The steps are identical regardless of the type of grid node you are recovering.

If more than one grid node is hosted on a physical or virtual Linux host, you can recover the grid nodes in any order. However, recovering a primary Admin Node first, if present, prevents the recovery of other grid nodes from stalling as they try to contact the primary Admin Node to register for recovery.

Deploy new Linux hosts

With a few exceptions, you prepare the new hosts as you did during the initial installation process.

To deploy new or reinstalled physical or virtual Linux hosts, follow the procedure for preparing the hosts in the StorageGRID installation instructions for your Linux operating system:

- Install Linux (Red Hat Enterprise Linux or CentOS)

- Install Linux (Ubuntu or Debian)

This procedure includes steps to accomplish the following tasks:

1. Install Linux.
2. Configure the host network.
3. Configure host storage.
4. Install the container engine.
5. Install the StorageGRID host service.

> Stop after you complete the “Install StorageGRID host service” task in the installation instructions. Don’t start the “Deploying grid nodes” task.

As you perform these steps, note the following important guidelines:

- Be sure to use the same host interface names you used on the original host.
- If you use shared storage to support your StorageGRID nodes, or you have moved some or all of the drives or SSDs from the failed to the replacement nodes, you must reestablish the same storage mappings that were present on the original host. For example, if you used WWIDs and aliases in `/etc/multipath.conf` as recommended in the installation instructions, be sure to use the same alias/WWID pairs in `/etc/multipath.conf` on the replacement host.
- If the StorageGRID node uses storage assigned from a NetApp ONTAP system, confirm that the volume does not have a FabricPool tiering policy enabled. Disabling FabricPool tiering for volumes used with StorageGRID nodes simplifies troubleshooting and storage operations.

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

Restore grid nodes to the host

To restore a failed grid node to a new Linux host, you perform these steps to restore the node configuration file.

1. Restore and validate the node by restoring the node configuration file. For a new install, you create a node configuration file for each grid node to be installed on a host. When restoring a grid node to a replacement host, you restore or replace the node configuration file for any failed grid nodes.
2. Start the StorageGRID host service.
3. As needed, recover any nodes that fail to start.

If any block storage volumes were preserved from the previous host, you might have to perform additional recovery procedures. The commands in this section help you determine which additional procedures are required.

Restore and validate grid nodes

You must restore the grid configuration files for any failed grid nodes, and then validate the grid configuration files and resolve any errors.

About this task

You can import any grid node that should be present on the host, as long as its `/var/local` volume was not lost as a result of the failure of the previous host. For example, the `/var/local` volume might still exist if you used shared storage for StorageGRID system data volumes, as described in the StorageGRID installation instructions for your Linux operating system. Importing the node restores its node configuration file to the host.

If it is not possible to import missing nodes, you must re-create their grid configuration files.

You must then validate the grid configuration file, and resolve any networking or storage issues that might occur before going on to restart StorageGRID. When you re-create the configuration file for a node, you must use the same name for the replacement node that was used for the node you are recovering.

See the installation instructions for more information about the location of the `/var/local` volume for a node.

- Install Red Hat Enterprise Linux or CentOS
- Install Ubuntu or Debian

Steps

1. At the command line of the recovered host, list all currently configured StorageGRID grid nodes:`sudo storagegrid node list`

   If no grid nodes are configured, there will be no output. If some grid nodes are configured, expect output in the following format:

   ```
   Name               Metadata-Volume
   ================================================================
   dc1-adm1           /dev/mapper/sgws-adm1-var-local
   dc1-gw1            /dev/mapper/sgws-gw1-var-local
   dc1-sn1            /dev/mapper/sgws-sn1-var-local
   dc1-arc1           /dev/mapper/sgws-arc1-var-local
   ```

   If some or all of the grid nodes that should be configured on the host aren't listed, you need to restore the missing grid nodes.

2. To import grid nodes that have a `/var/local` volume:

   a. Run the following command for each node you want to import:`sudo storagegrid node import node-var-local-volume-path`

      The `storagegrid node import` command succeeds only if the target node was shut down cleanly on the host on which it last ran. If that is not the case, you will observe an error similar to the following:

      ```
      This node (node-name) appears to be owned by another host (UUID host-uuid).

      Use the --force flag if you are sure import is safe.
      ```

   b. If you see the error about the node being owned by another host, run the command again with the `--force` flag to complete the import:`sudo storagegrid --force node import node-var-local-volume-path`

Any nodes imported with the `--force` flag will require additional recovery steps before they can rejoin the grid, as described in What's next: Perform additional recovery steps, if required.

3. For grid nodes that don't have a `/var/local` volume, re-create the node's configuration file to restore it to the host. For instructions, see:
   - Create node configuration files for Red Hat Enterprise Linux or CentOS
   - Create node configuration files for Ubuntu or Debian

   When you re-create the configuration file for a node, you must use the same name for the replacement node that was used for the node you are recovering. For Linux deployments, ensure that the configuration file name contains the node name. You should use the same network interfaces, block device mappings, and IP addresses when possible. This practice minimizes the amount of data that needs to be copied to the node during recovery, which could make the recovery significantly faster (in some cases, minutes rather than weeks).

   If you use any new block devices (devices that the StorageGRID node did not use previously) as values for any of the configuration variables that start with `BLOCK_DEVICE_` when you are re-creating the configuration file for a node, follow the guidelines in Fix missing block device errors.

4. Run the following command on the recovered host to list all StorageGRID nodes.

   ```
   sudo storagegrid node list
   ```

5. Validate the node configuration file for each grid node whose name was shown in the storagegrid node list output:

   ```
   sudo storagegrid node validate node-name
   ```

   You must address any errors or warnings before starting the StorageGRID host service. The following sections give more detail on errors that might have special significance during recovery.

Fix missing network interface errors

If the host network is not configured correctly or a name is misspelled, an error occurs when StorageGRID checks the mapping specified in the `/etc/storagegrid/nodes/node-name.conf` file.

You might see an error or warning matching this pattern:

```
Checking configuration file /etc/storagegrid/nodes/<node-name>.conf for
node <node-name>...
ERROR: <node-name>: GRID_NETWORK_TARGET = <host-interface-name>
       <node-name>: Interface <host-interface-name>' does not exist
```

The error could be reported for the Grid Network, the Admin Network, or the Client Network. This error means that the `/etc/storagegrid/nodes/node-name.conf` file maps the indicated StorageGRID network to the host interface named *host-interface-name*, but there is no interface with that name on the current host.

If you receive this error, verify that you completed the steps in Deploy new Linux hosts. Use the same names for all host interfaces as were used on the original host.

If you are unable to name the host interfaces to match the node configuration file, you can edit the node configuration file and change the value of the GRID_NETWORK_TARGET, the ADMIN_NETWORK_TARGET, or the CLIENT_NETWORK_TARGET to match an existing host interface.

Make sure the host interface provides access to the appropriate physical network port or VLAN, and that the interface does not directly reference a bond or bridge device. You must either configure a VLAN (or other virtual interface) on top of the bond device on the host, or use a bridge and virtual Ethernet (veth) pair.

Fix missing block device errors

The system checks that each recovered node maps to a valid block device special file or a valid softlink to a block device special file. If StorageGRID finds invalid mapping in the `/etc/storagegrid/nodes/node-name.conf` file, a missing block device error displays.

If you observe an error matching this pattern:

```
Checking configuration file /etc/storagegrid/nodes/<node-name>.conf for
node <node-name>...
ERROR: <node-name>: BLOCK_DEVICE_PURPOSE = <path-name>
       <node-name>: <path-name> does not exist
```

It means that `/etc/storagegrid/nodes/node-name.conf` maps the block device used by *node-name* for `PURPOSE` to the given path-name in the Linux file system, but there is not a valid block device special file, or softlink to a block device special file, at that location.

Verify that you completed the steps in Deploy new Linux hosts. Use the same persistent device names for all block devices as were used on the original host.

If you are unable to restore or re-create the missing block device special file, you can allocate a new block device of the appropriate size and storage category and edit the node configuration file to change the value of `BLOCK_DEVICE_PURPOSE` to point to the new block device special file.

Determine the appropriate size and storage category using the tables for your Linux operating system:

- Storage and performance requirements for Red Hat Enterprise Linux or CentOS
- Storage and performance requirements for Ubuntu or Debian

Review the recommendations for configuring host storage before proceeding with the block device replacement:

- Configure host storage for Red Hat Enterprise Linux or CentOS
- Configure host storage for Ubuntu or Debian

> If you must provide a new block storage device for any of the configuration file variables starting with `BLOCK_DEVICE_` because the original block device was lost with the failed host, ensure the new block device is unformatted before attempting further recovery procedures. The new block device will be unformatted if you are using shared storage and have created a new volume. If you are unsure, run the following command against any new block storage device special files.

Run the following command only for new block storage devices. Don't run this command if you believe the block storage still contains valid data for the node being recovered, as any data on the device will be lost.

```
sudo dd if=/dev/zero of=/dev/mapper/my-block-device-name bs=1G count=1
```

Start StorageGRID host service

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

Steps

1. Run the following commands on each host:

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

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

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

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

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

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

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

Recover nodes that fail to start normally

If a StorageGRID node doesn't rejoin the grid normally and doesn't show up as recoverable, it might be corrupted. You can force the node into recovery mode.

Steps

1. Confirm that the node's network configuration is correct.

   The node might have failed to rejoin the grid because of incorrect network interface mappings or an incorrect Grid Network IP address or gateway.

2. If the network configuration is correct, issue the `force-recovery` command:

   `sudo storagegrid node force-recovery` *`node-name`*

3. Perform the additional recovery steps for the node. See What's next: Perform additional recovery steps, if required.

What's next: Perform additional recovery steps, if required

Depending on the specific actions you took to get the StorageGRID nodes running on the replacement host, you might need to perform additional recovery steps for each node.

Node recovery is complete if you did not need to take any corrective actions while you replaced the Linux host or restored the failed grid node to the new host.

Corrective actions and next steps

During node replacement, you might have needed to take one of these corrective actions:

- You had to use the `--force` flag to import the node.
- For any `<PURPOSE>`, the value of the `BLOCK_DEVICE_<PURPOSE>` configuration file variable refers to a block device that does not contain the same data it did before the host failure.
- You issued `storagegrid node force-recovery` *`node-name`* for the node.
- You added a new block device.

If you took **any** of these corrective actions, you must perform additional recovery steps.

Type of recovery	Next step
Primary Admin Node	Configure replacement primary Admin Node
Non-primary Admin Node	Select Start Recovery to configure non-primary Admin Node
Gateway Node	Select Start Recovery to configure Gateway Node
Archive Node	Select Start Recovery to configure Archive Node
Storage Node (software-based): • If you had to use the `--force` flag to import the node, or you issued `storagegrid node force-recovery` *`node-name`* • If you had to do a full node reinstall, or you needed to restore /var/local	Select Start Recovery to configure Storage Node

Type of recovery	Next step
Storage Node (software-based): • If you added a new block device. • If, for any `<PURPOSE>`, the value of the `BLOCK_DEVICE_<PURPOSE>` configuration file variable refers to a block device that does not contain the same data it did before the host failure.	Recover from storage volume failure where system drive is intact

Replace failed node with services appliance

Replace failed node with services appliance: Overview

You can use an SG100 or SG1000 services appliance to recover a failed Gateway Node, a failed non-primary Admin Node, or a failed primary Admin Node that was hosted on VMware, a Linux host, or a services appliance. This procedure is one step of the grid node recovery procedure.

Before you begin

- You have determined that one of the following situations is true:
  - The virtual machine hosting the node can't be restored.
  - The physical or virtual Linux host for the grid node has failed, and must be replaced.
  - The services appliance hosting the grid node must be replaced.
- You have confirmed that the StorageGRID Appliance Installer version on the services appliance matches the software version of your StorageGRID system. See Verify and upgrade StorageGRID Appliance Installer version.

> Don't deploy both an SG100 and an SG1000 service appliance in the same site. Unpredictable performance might result.

About this task

You can use an SG100 or SG1000 services appliance to recover a failed grid node in the following cases:

- The failed node was hosted on VMware or Linux (platform change)
- The failed node was hosted on a services appliance (platform replacement)

Install services appliance (platform change only)

When you are recovering a failed grid node that was hosted on VMware or a Linux host and you are using a services appliance for the replacement node, you must first install the new appliance hardware using the same node name (system name) as the failed node.

Before you begin

You have the following information about the failed node:

- **Node name**: You must install the services appliance using the same node name as the failed node. The

node name is the hostname (system name).

- **IP addresses**: You can assign the services appliance the same IP addresses as the failed node, which is the preferred option, or you can select a new unused IP address on each network.

About this task

Perform this procedure only if you are recovering a failed node that was hosted on VMware or Linux and are replacing it with a node hosted on a services appliance.

Steps

1. Follow the instructions for installing a new SG100 or SG1000 services appliance. See Quick start for hardware installation.
2. When prompted for a node name, use the node name of the failed node.

Prepare appliance for reinstallation (platform replacement only)

When recovering a grid node that was hosted on a services appliance, you must first prepare the appliance for reinstallation of StorageGRID software.

Perform this procedure only if you are replacing a failed node that was hosted on a services appliance. Don't follow these steps if the failed node was originally hosted on VMware or a Linux host.

Steps

1. Log in to the failed grid node:

   a. Enter the following command: `ssh admin@grid_node_IP`

   b. Enter the password listed in the `Passwords.txt` file.

   c. Enter the following command to switch to root: `su -`

   d. Enter the password listed in the `Passwords.txt` file.

   When you are logged in as root, the prompt changes from `$` to `#`.

2. Prepare the appliance for the installation of StorageGRID software. Enter: `sgareinstall`
3. When prompted to continue, enter: `y`

   The appliance reboots, and your SSH session ends. It usually takes about 5 minutes for the StorageGRID Appliance Installer to become available, although in some cases you might need to wait up to 30 minutes.

   The services appliance is reset, and data on the grid node is no longer accessible. IP addresses configured during the original installation process should remain intact; however, it is recommended that you confirm this when the procedure completes.

   After executing the `sgareinstall` command, all StorageGRID-provisioned accounts, passwords, and SSH keys are removed, and new host keys are generated.

Start software installation on services appliance

To install a Gateway Node or Admin Node on an SG100 or SG1000 services appliance, you use the StorageGRID Appliance Installer, which is included on the appliance.

Before you begin

- The appliance is installed in a rack, connected to your networks, and powered on.
- Network links and IP addresses are configured for the appliance using the StorageGRID Appliance Installer.
- If you are installing a Gateway Node or non-primary Admin Node, you know the IP address of the primary Admin Node for the StorageGRID grid.
- All Grid Network subnets listed on the IP Configuration page of the StorageGRID Appliance Installer are defined in the Grid Network Subnet List on the primary Admin Node.

  See Quick start for hardware installation.

- You are using a supported web browser.
- You have one of the IP addresses assigned to the appliance. You can use the IP address for the Admin Network, the Grid Network, or the Client Network.
- If you are installing a primary Admin Node, you have the Ubuntu or Debian install files for this version of StorageGRID available.

> A recent version of StorageGRID software is preloaded onto the services appliance during manufacturing. If the preloaded version of software matches the version being used in your StorageGRID deployment, you don't need the installation files.

About this task

To install StorageGRID software on an SG100 or SG1000 services appliance:

- For a primary Admin Node, you specify the name of the node and then upload the appropriate software packages (if required).
- For a non-primary Admin Node or a Gateway Node, you specify or confirm the IP address of the primary Admin Node and the name of the node.
- You start the installation and wait as volumes are configured and the software is installed.
- Partway through the process, the installation pauses. To resume the installation, you must sign into the Grid Manager and configure the pending node as a replacement for the failed node.
- After you have configured the node, the appliance installation process completes, and the appliance is rebooted.

Steps

1. Open a browser and enter one of the IP addresses for the SG100 or SG1000 services appliance.

   `https://Controller_IP:8443`

   The StorageGRID Appliance Installer Home page appears.

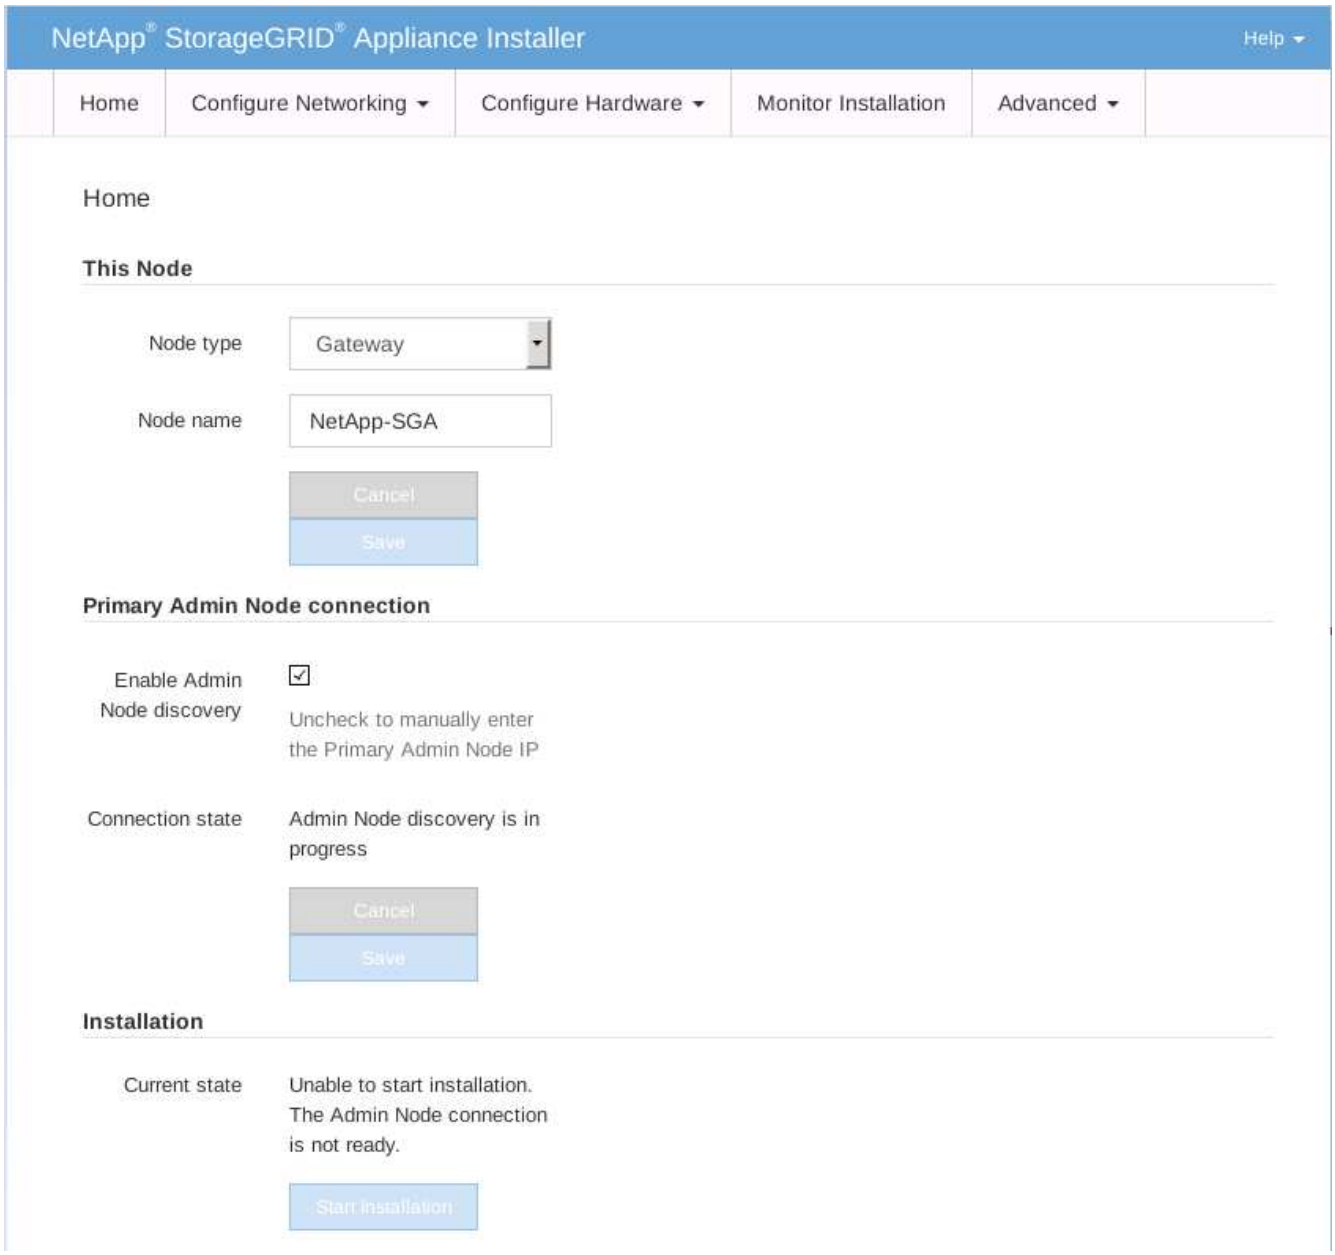

2. To install a Primary Admin Node:
   a. In the This Node section, for **Node Type**, select **Primary Admin**.
   b. In the **Node Name** field, enter the same name that was used for the node you are recovering, and click **Save**.
   c. In the Installation section, check the software version listed under Current state

      If the version of software that is ready to install is correct, skip ahead to the Installation step.

   d. If you need to upload a different version of software, under the **Advanced** menu, select **Upload StorageGRID Software**.

      TheUpload StorageGRID Software page appears.

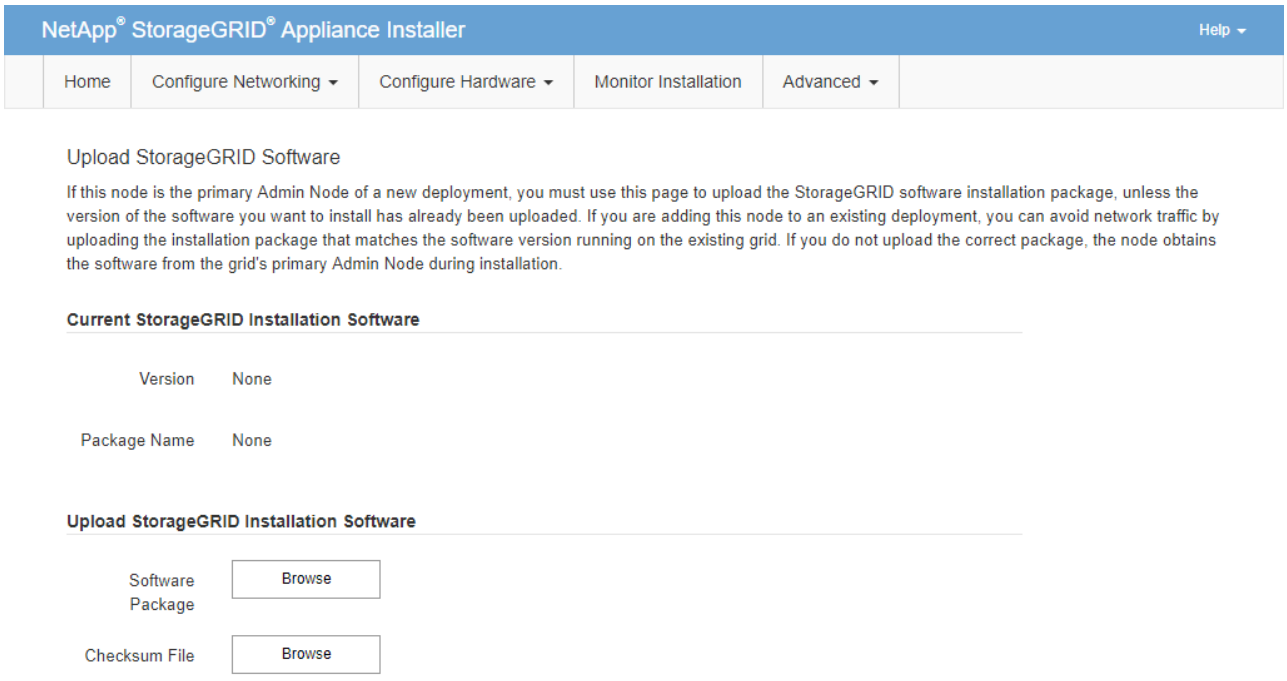

e. Click **Browse** to upload the **Software Package** and **Checksum File** for StorageGRID software.

   The files are automatically uploaded after you select them.

f. Click **Home** to return to the StorageGRID Appliance Installer Home page.

3. To install a Gateway Node or non-Primary Admin Node:

   a. In the This Node section, for **Node Type**, select **Gateway** or **Non-Primary Admin**, depending on the type of node you are restoring.

   b. In the **Node Name** field, enter the same name that was used for the node you are recovering, and click **Save**.

   c. In the Primary Admin Node connection section, determine whether you need to specify the IP address for the primary Admin Node.

      The StorageGRID Appliance Installer can discover this IP address automatically, assuming the primary Admin Node, or at least one other grid node with ADMIN_IP configured, is present on the same subnet.

   d. If this IP address is not shown or you need to change it, specify the address:

Option	Description
Manual IP entry	a. Clear the **Enable Admin Node discovery** checkbox. b. Enter the IP address manually. c. Click **Save**. d. Wait while the connection state for the new IP address becomes “ready.”

Option	Description
Automatic discovery of all connected primary Admin Nodes	a. Select the **Enable Admin Node discovery** checkbox. b. From the list of discovered IP addresses, select the primary Admin Node for the grid where this services appliance will be deployed. c. Click **Save**. d. Wait while the connection state for the new IP address becomes “ready.”

4. In the Installation section, confirm that the current state is Ready to start installation of node name and that the **Start Installation** button is enabled.

   If the **Start Installation** button is not enabled, you might need to change the network configuration or port settings. For instructions, see the maintenance instructions for your appliance.

5. From the StorageGRID Appliance Installer home page, click **Start Installation**.

   The Current state changes to “Installation is in progress,” and the Monitor Installation page is displayed.

   > If you need to access the Monitor Installation page manually, click **Monitor Installation** from the menu bar.

Related information

Maintain SG100 and SG1000 appliances

Monitor services appliance installation

The StorageGRID Appliance Installer provides status until installation is complete. When the software installation is complete, the appliance is rebooted.

Steps

1. To monitor the installation progress, click **Monitor Installation** from the menu bar.

   The Monitor Installation page shows the installation progress.

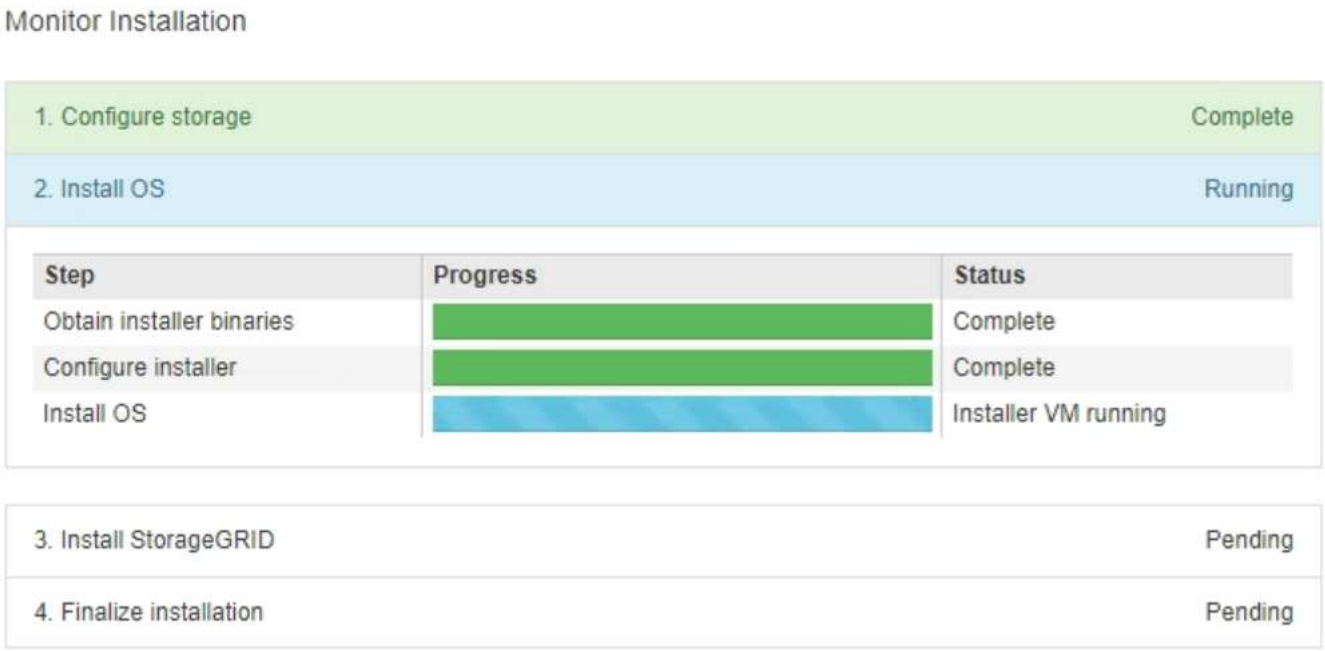

The blue status bar indicates which task is currently in progress. Green status bars indicate tasks that have completed successfully.

> The installer ensures that tasks completed in a previous install aren't re-run. If you are re-running an installation, any tasks that don't need to be re-run are shown with a green status bar and a status of "Skipped."

2. Review the progress of first two installation stages.
   - **1. Configure storage**

     During this stage, the installer clears any existing configuration from the drives, and configures host settings.

   - **2. Install OS**

     During this stage, the installer copies the base operating system image for StorageGRID from the primary Admin Node to the appliance or installs the base operating system from the installation package for the primary Admin Node.

3. Continue monitoring the installation progress until one of the following occurs:
   - For appliance Gateway Nodes or non-primary appliance Admin Nodes, the **Install StorageGRID** stage pauses and a message appears on the embedded console, prompting you to approve this node on the Admin Node using the Grid Manager.

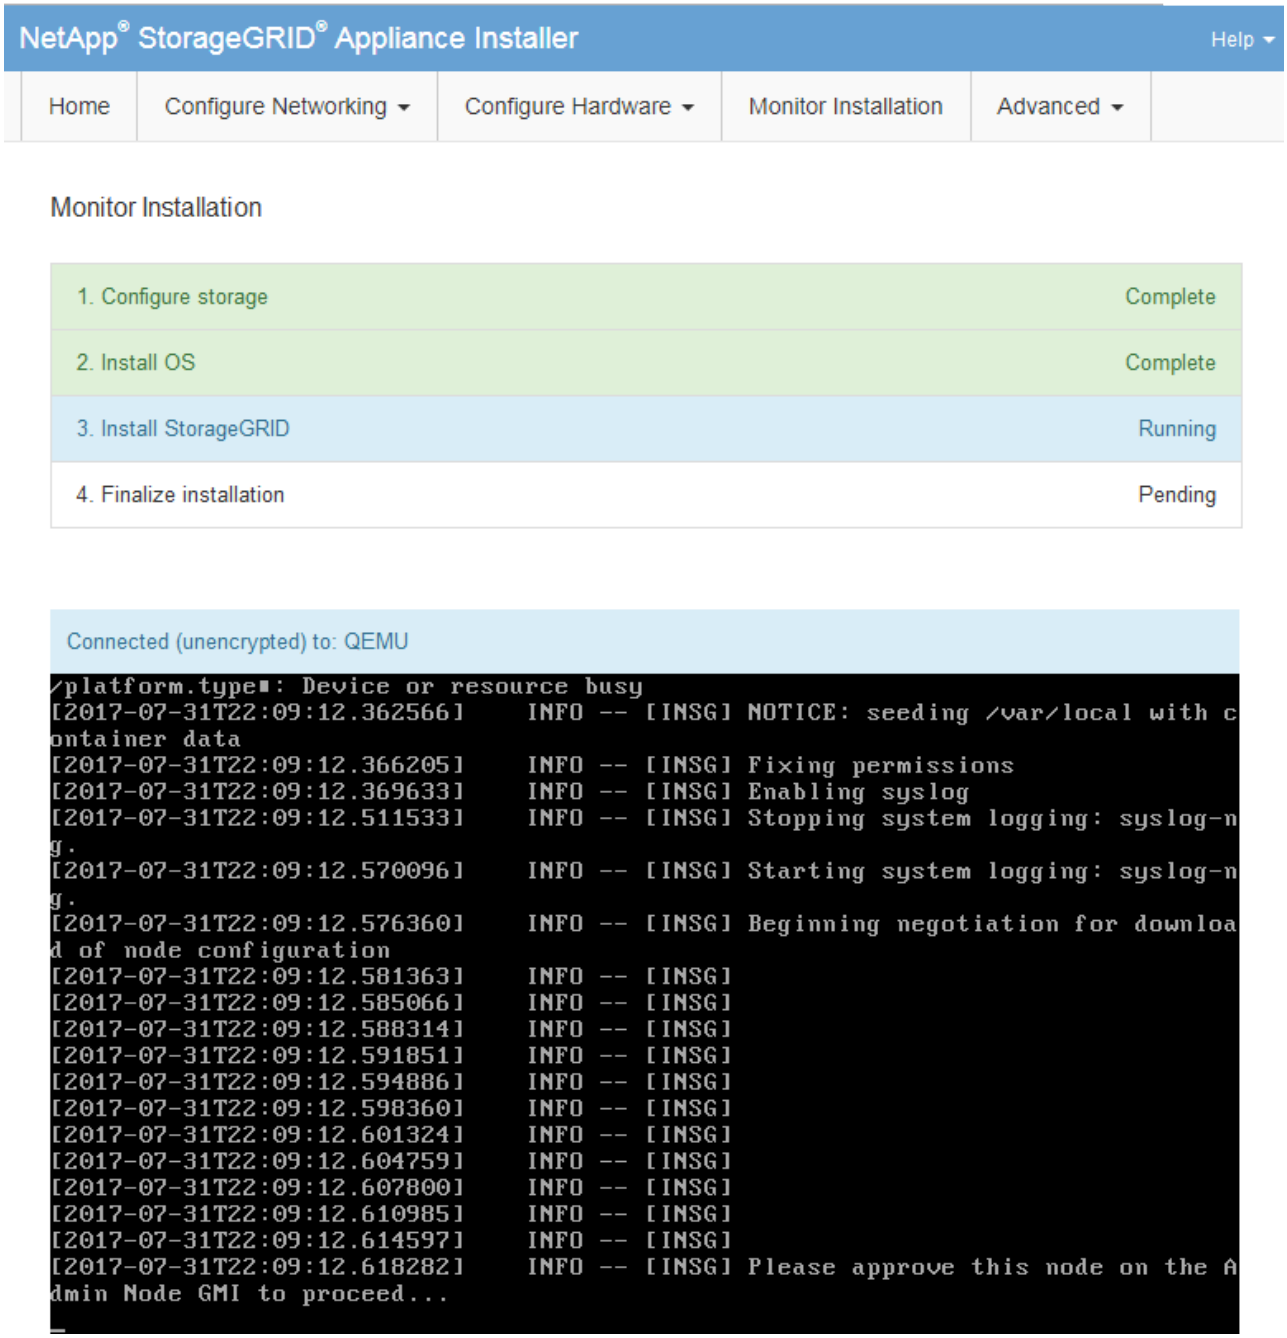

- For appliance primary Admin Nodes, a fifth phase (Load StorageGRID Installer) appears. If the fifth phase is in progress for more than 10 minutes, refresh the page manually.

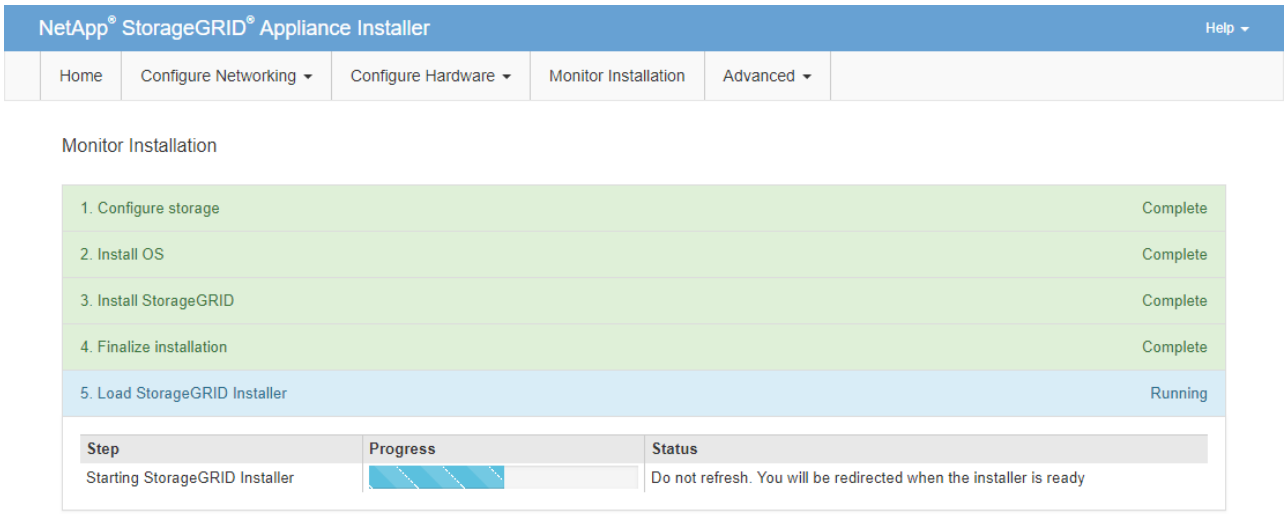

4. Go to the next step of the recovery process for the type of appliance grid node that you are recovering.

Type of recovery	Reference
Gateway Node	Select Start Recovery to configure Gateway Node
Non-primary Admin Node	Select Start Recovery to configure non-primary Admin Node
Primary Admin Node	Configure replacement primary Admin Node

How site recovery is performed by technical support

If an entire StorageGRID site fails or if multiple Storage Nodes fail, you must contact technical support. Technical support will assess your situation, develop a recovery plan, and then recover the failed nodes or site in a way that meets your business objectives, optimizes recovery time, and prevents unnecessary data loss.

Site recovery can only be performed by technical support.

StorageGRID systems are resilient to a wide variety of failures, and you can successfully perform many recovery and maintenance procedures yourself. However, it is difficult to create a simple, generalized site recovery procedure because the detailed steps depend on factors that are specific to your situation. For example:

- **Your business objectives**: After the complete loss of a StorageGRID site, you should evaluate how best to meet your business objectives. For example, do you want to rebuild the lost site in-place? Do you want to replace the lost StorageGRID site in a new location? Every customer's situation is different, and your recovery plan must be designed to address your priorities.
- **Exact nature of the failure**: Before beginning a site recovery, establish if any nodes at the failed site are intact or if any Storage Nodes contain recoverable objects. If you rebuild nodes or storage volumes that contain valid data, unnecessary data loss could occur.
- **Active ILM policy**: The number, type, and location of object copies in your grid is controlled by your active ILM policy. The specifics of your ILM policy can affect the amount of recoverable data, as well as the specific techniques required for recovery.

If a site contains the only copy of an object and the site is lost, the object is lost.

- **Bucket (or container) consistency**: The consistency level applied to a bucket (or container) affects whether StorageGRID fully replicates object metadata to all nodes and sites before telling a client that object ingest was successful. If your consistency level allows for eventual consistency, some object metadata might have been lost in the site failure. This can affect the amount of recoverable data and potentially the details of the recovery procedure.
- **History of recent changes**: The details of your recovery procedure can be affected by whether any maintenance procedures were in progress at the time of the failure or whether any recent changes were made to your ILM policy. Technical support must assess the recent history of your grid as well as its current situation before beginning a site recovery.

Overview of site recovery

This is a general overview of the process that technical support uses to recover a failed site.

Site recovery can only be performed by technical support.

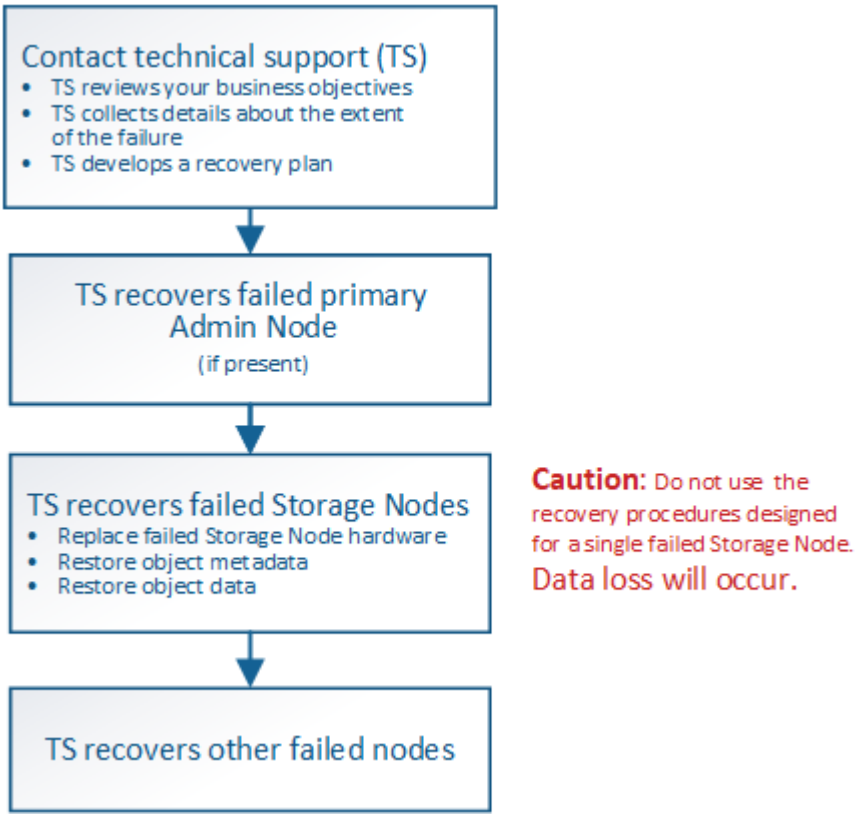

1. Contact technical support.

   Technical support does a detailed assessment of the failure and works with you to review your business objectives. Based on this information, technical support develops a recovery plan tailored for your situation.

2. Technical support recovers the primary Admin Node if it has failed.
3. Technical support recovers all Storage Nodes, following this outline:
   a. Replace Storage Node hardware or virtual machines as required.
   b. Restore object metadata to the failed site.

c. Restore object data to the recovered Storage Nodes.

> Data loss will occur if the recovery procedures for a single failed Storage Node are used.

> When an entire site has failed, specialized commands are required to successfully restore objects and object metadata.

4. Technical support recovers other failed nodes.

   After object metadata and data have been recovered, failed Gateway Nodes, non-primary Admin Nodes, or Archive Nodes can be recovered using standard procedures.

Related information

Site decommission

Decommission procedure

Decommission procedure: Overview

You can perform a decommission procedure to permanently remove grid nodes or an entire site from the StorageGRID system.

To remove a grid node or a site, you perform one of the following decommission procedures:

- Perform a grid node decommission to remove one or more nodes, which can be at one or more sites. The nodes you remove can be online and connected to the StorageGRID system, or they can be offline and disconnected.
- Perform a site decommission to remove a site. You perform a **connected site decommission** if all nodes are connected to StorageGRID. You perform a **disconnected site decommission** if all nodes are disconnected from StorageGRID.

> Before performing a disconnected site decommission, you must contact your NetApp account representative. NetApp will review your requirements before enabling all steps in the Decommission Site wizard. You should not attempt a disconnected site decommission if you believe it might be possible to recover the site or to recover object data from the site.

If a site contains a mixture of connected ( ) and disconnected nodes (or), you must bring all offline nodes back online.

> If you need to perform a second maintenance procedure, you can pause the decommission procedure while the Storage Nodes are being removed. The **Pause** button is enabled only when the ILM evaluation or erasure-coded data decommissioning stages are reached; however, ILM evaluation (data migration) will continue to run in the background. After the second maintenance procedure is complete, you can resume decommissioning.

Grid node decommission

Grid node decommission: Overview

You can use the node decommission procedure to remove one or more Storage Nodes, Gateway Nodes, or non-primary Admin Nodes at one or more sites. You can’t decommission the primary Admin Node or an Archive Node.

In general, you should decommission grid nodes only while they are connected to the StorageGRID system and all nodes are in normal health (have green icons on the **NODES** pages and on the **Decommission Nodes** page). However, if required, you can decommission a grid node that is disconnected. Before removing a disconnected node, make sure you understand the implications and restrictions of that process.

Use the node decommission procedure when any of the following are true:

- You have added a larger Storage Node to the system and you want to remove one or more smaller Storage Nodes, while at the same time preserving objects.
- You require less total storage.
- You no longer require a Gateway Node.
- You no longer require a non-primary Admin Node.
- Your grid includes a disconnected node that you can’t recover or bring back online.

The flowchart shows the high-level steps for decommissioning grid nodes.

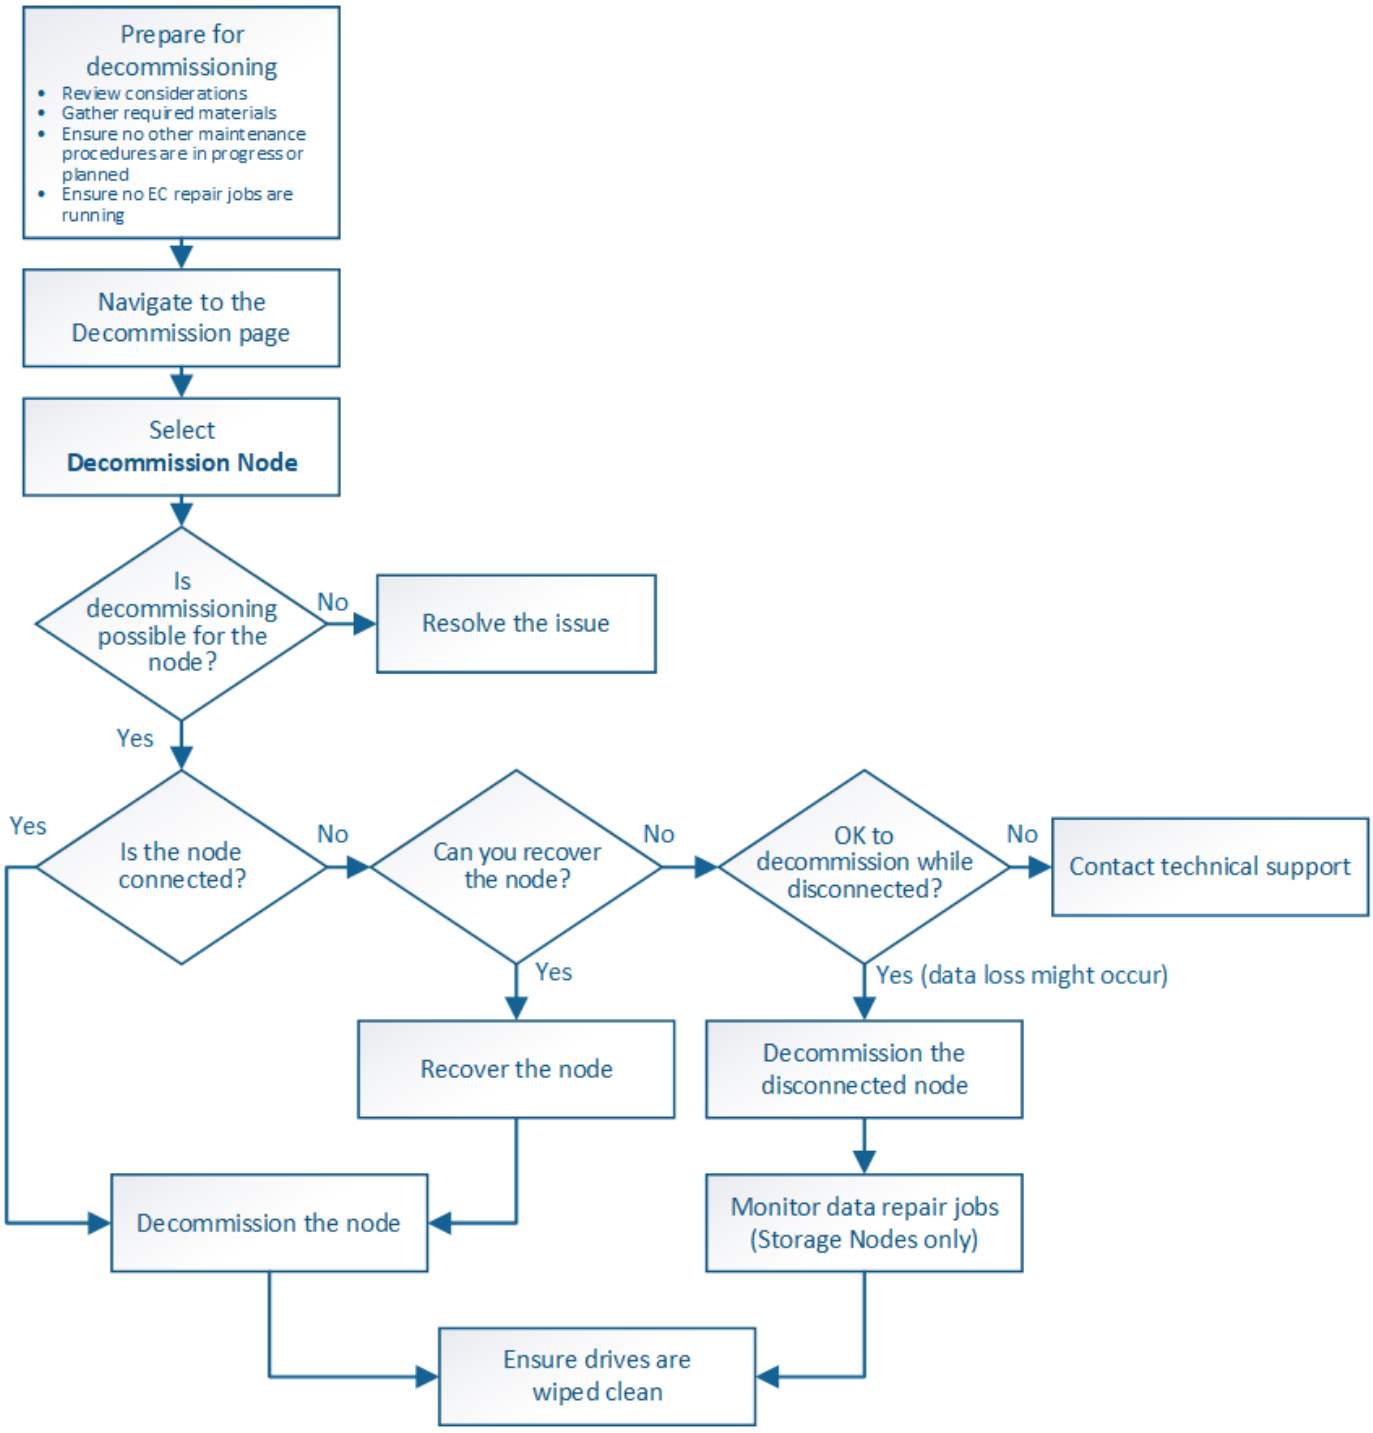

Considerations for grid node decommission

General considerations for grid node decommission

Before you start this procedure to decommission one or more nodes, you must understand the implications of removing each type of node. Upon the successful decommissioning of a node, its services will be disabled and the node will be automatically shut down.

You can't decommission a node if doing so will leave StorageGRID in an invalid state. The following rules are enforced:

- You can't decommission the primary Admin Node.
- You can't decommission Archive Nodes.
- You can't decommission an Admin Node or a Gateway Node if one of its network interfaces is part of a high availability (HA) group.
- You can't decommission a Storage Node if its removal would affect the ADC quorum.
- You can't decommission a Storage Node if it is required for the active ILM policy.
- You should not decommission more than 10 Storage Nodes in a single Decommission Node procedure.
- You can't decommission a connected node if your grid includes any disconnected nodes (nodes whose health is Unknown or Administratively Down). You must decommission or recover the disconnected nodes first.
- If your grid contains multiple disconnected nodes, the software requires you to decommission them all at the same time, which increases the potential for unexpected results.
- If a disconnected node can't be removed (for example, a Storage Node that is required for the ADC quorum), no other disconnected node can be removed.
- If you want to replace an older appliance with a newer appliance, consider cloning the appliance node instead of decommissioning the old node and adding the new node in an expansion.

> Don't remove a grid node's virtual machine or other resources until instructed to do so in decommission procedures.

Considerations for Admin Node or Gateway Node decommission

Review the following considerations before decommissioning an Admin Node or a Gateway Node.

- The decommission procedure requires exclusive access to some system resources, so you must confirm that no other maintenance procedures are running.
- You can't decommission the primary Admin Node.
- You can't decommission an Admin Node or a Gateway Node if one of its network interfaces is part of a high availability (HA) group. You must first remove the network interfaces from the HA group. See the instructions for managing HA groups.
- As required, you can safely change the ILM policy while decommissioning a Gateway Node or an Admin Node.
- If you decommission an Admin Node and single sign-on (SSO) is enabled for your StorageGRID system, you must remember to remove the node's relying party trust from Active Directory Federation Services (AD FS).
- If you use grid federation, ensure that the IP address of the node you are decommissioning was not specified for a grid federation connection.

Storage Nodes

Considerations for Storage Node decommission

If you plan to decommission a Storage Node, you must understand how StorageGRID manages the object data and metadata on that node.

The following considerations and restrictions apply when decommissioning Storage Nodes:

- The system must, at all times, include enough Storage Nodes to satisfy operational requirements, including the ADC quorum and the active ILM policy. To satisfy this restriction, you might need to add a new Storage Node in an expansion operation before you can decommission an existing Storage Node.
- If the Storage Node is disconnected when you decommission it, the system must reconstruct the data using data from the connected Storage Nodes, which can result in data loss.
- When you remove a Storage Node, large volumes of object data must be transferred over the network. Although these transfers should not affect normal system operations, they can have an impact on the total amount of network bandwidth consumed by the StorageGRID system.
- Tasks associated with Storage Node decommissioning are given a lower priority than tasks associated with normal system operations. This means that decommissioning does not interfere with normal StorageGRID system operations, and does not need to be scheduled for a period of system inactivity. Because decommissioning is performed in the background, it is difficult to estimate how long the process will take to complete. In general, decommissioning finishes more quickly when the system is quiet, or if only one Storage Node is being removed at a time.
- It might take days or weeks to decommission a Storage Node. Plan this procedure accordingly. While the decommission process is designed to not impact system operations, it can limit other procedures. In general, you should perform any planned system upgrades or expansions before you remove grid nodes.
- Decommission procedures that involve Storage Nodes can be paused during certain stages to allow other maintenance procedures to run if needed, and resumed once they are complete.
- You can’t run data repair operations on any grid nodes when a decommission task is running.
- You should not make any changes to the ILM policy while a Storage Node is being decommissioned.
- When you remove a Storage Node, data on the node is migrated to other grid nodes; however, this data is not completely removed from the decommissioned grid node. To permanently and securely remove data, you must wipe the decommissioned grid node’s drives after the decommission procedure is complete.
- When you decommission a Storage Node, the following alerts and alarms might be raised and you might receive related email and SNMP notifications:
  - **Unable to communicate with node** alert. This alert is triggered when you decommission a Storage Node that includes the ADC service. The alert is resolved when the decommission operation completes.
  - VSTU (Object Verification Status) alarm. This notice-level alarm indicates that the Storage Node is going into maintenance mode during the decommission process.
  - CASA (Data Store Status) alarm. This major-level alarm indicates that the Cassandra database is going down because services have stopped.

Related information

Restore object data to storage volume

Understand the ADC quorum

You might not be able to decommission certain Storage Nodes at a data center site if too few Administrative Domain Controller (ADC) services would remain after the decommissioning. This service, which is found on some Storage Nodes, maintains grid topology information and provides configuration services to the grid. The StorageGRID system requires a quorum of ADC services to be available at each site and at all times.

You can’t decommission a Storage Node if removing the node would cause the ADC quorum to no longer be

met. To satisfy the ADC quorum during a decommissioning, a minimum of three Storage Nodes at each data center site must have the ADC service. If a data center site has more than three Storage Nodes with the ADC service, a simple majority of those nodes must remain available after the decommissioning ((0.5 * *Storage Nodes with ADC*) + 1).

For example, suppose a data center site currently includes six Storage Nodes with ADC services and you want to decommission three Storage Nodes. Because of the ADC quorum requirement, you must complete two decommission procedures, as follows:

- In the first decommission procedure, you must ensure that four Storage Nodes with ADC services remain available ((0.5 * 6) +1) . This means that you can only decommission two Storage Nodes initially.
- In the second decommission procedure, you can remove the third Storage Node because the ADC quorum now only requires three ADC services to remain available ((0.5 * 4) + 1).

If you need to decommission a Storage Node but are unable to because of the ADC quorum requirement, you must add a new Storage Node in an expansion and specify that it should have an ADC service. Then, you can decommission the existing Storage Node.

Related information

Expand your grid

Review ILM policy and storage configuration

If you plan to decommission a Storage Node, you should review your StorageGRID system's ILM policy before starting the decommissioning process.

During decommissioning, all object data is migrated from the decommissioned Storage Node to other Storage Nodes.

> The ILM policy you have *during* the decommission will be the one used *after* the decommission. You must ensure this policy meets your data requirements both before you start the decommission and after the decommission is complete.

You should review the rules in the active ILM policy to ensure that the StorageGRID system will continue to have enough capacity of the correct type and in the correct locations to accommodate the decommissioning of a Storage Node.

Consider the following:

- Will it be possible for ILM evaluation services to copy object data such that ILM rules are satisfied?
- What happens if a site becomes temporarily unavailable while decommissioning is in progress? Can additional copies be made in an alternate location?
- How will the decommissioning process affect the final distribution of content? As described in Consolidate Storage Nodes, you should add new Storage Nodes before decommissioning old ones. If you add a larger replacement Storage Node after decommissioning a smaller Storage Node, the old Storage Nodes could be close to capacity and the new Storage Node could have almost no content. Most write operations for new object data would then be directed at the new Storage Node, reducing the overall efficiency of system operations.
- Will the system, at all times, include enough Storage Nodes to satisfy the active ILM policy?

> An ILM policy that can't be satisfied will lead to backlogs and alarms, and can halt operation of the StorageGRID system.

Verify that the proposed topology that will result from the decommissioning process satisfies the ILM policy by assessing the factors listed in the table.

Area to assess	Notes
Available capacity	Will there be enough storage capacity to accommodate all of the object data stored in the StorageGRID system, including the permanent copies of object data currently stored on the Storage Node to be decommissioned?Will there be enough capacity to handle the anticipated growth in stored object data for a reasonable interval of time after decommissioning is complete?
Location of storage	If enough capacity remains in the StorageGRID system as a whole, is the capacity in the right locations to satisfy the StorageGRID system's business rules?
Storage type	Will there be enough storage of the appropriate type after decommissioning is complete? For example, ILM rules might dictate that content be moved from one type of storage to another as content ages. If so, you must ensure that enough storage of the appropriate type is available in the final configuration of the StorageGRID system.

Related information

Manage objects with ILM

Expand your grid

Decommission disconnected Storage Nodes

You must understand what can happen if you decommission a Storage Node while it is disconnected (health is Unknown or Administratively Down).

When you decommission a Storage Node that is disconnected from the grid, StorageGRID uses data from other Storage Nodes to reconstruct the object data and metadata that was on the disconnected node. It does this by automatically starting data repair jobs at the end of the decommissioning process.

Before decommissioning a disconnected Storage Node, be aware of the following:

- You should never decommission a disconnected node unless you are sure it can't be brought online or recovered.

> Don't perform this procedure if you believe it might be possible to recover object data from the node. Instead, contact technical support to determine if node recovery is possible.

- If a disconnected Storage Node contains the only copy of an object, that object will be lost when you decommission the node. The data repair jobs can only reconstruct and recover objects if at least one replicated copy or enough erasure-coded fragments exist on Storage Nodes that are currently connected.

- When you decommission a disconnected Storage Node, the decommission procedure completes relatively quickly. However, the data repair jobs can take days or weeks to run and aren't monitored by the decommission procedure. You must manually monitor these jobs and restart them as needed. See Check data repair jobs.
- If you decommission more than one disconnected Storage Node at a time, data loss might occur. The system might not be able to reconstruct data if too few copies of object data, metadata, or erasure-coded fragments remain available.

> If you have more than one disconnected Storage Node that you can't recover, contact technical support to determine the best course of action.

Consolidate Storage Nodes

You can consolidate Storage Nodes to reduce the Storage Node count for a site or deployment while increasing storage capacity.

When you consolidate Storage Nodes, you expand the StorageGRID system to add new, larger capacity Storage Nodes and then decommission the old, smaller capacity Storage Nodes. During the decommission procedure, objects are migrated from the old Storage Nodes to the new Storage Nodes.

> If you are consolidating older and smaller appliances with new models or larger capacity appliances, you many use the node clone feature or the node clone procedure and the decommission procedure if you aren't doing a one-to-one replacement.

For example, you might add two new, larger capacity Storage Nodes to replace three older Storage Nodes. You would first use the expansion procedure to add the two new, larger Storage Nodes, and then use the decommission procedure to remove the three old, smaller capacity Storage Nodes.

By adding new capacity before removing existing Storage Nodes, you ensure a more balanced distribution of data across the StorageGRID system. You also reduce the possibility that an existing Storage Node might be pushed beyond the storage watermark level.

Related information

Expand your grid

Decommission multiple Storage Nodes

If you need to remove more than one Storage Node, you can decommission them either sequentially or in parallel.

- If you decommission Storage Nodes sequentially, you must wait for the first Storage Node to complete decommissioning before starting to decommission the next Storage Node.
- If you decommission Storage Nodes in parallel, the Storage Nodes simultaneously process decommission tasks for all Storage Nodes being decommissioned. This can result in a situation where all permanent copies of a file are marked as "read-only," temporarily disabling deletion in grids where this functionality is enabled.

Check data repair jobs

Before decommissioning a grid node, you must confirm that no data repair jobs are

active. If any repairs have failed, you must restart them and allow them to complete before performing the decommission procedure.

If you need to decommission a disconnected Storage Node, you will also complete these steps after the decommission procedure completes to ensure the data repair job has completed successfully. You must ensure that any erasure-coded fragments that were on the removed node have been restored successfully.

These steps only apply to systems that have erasure-coded objects.

1. Log in to the primary Admin Node:
   a. Enter the following command: `ssh admin@grid_node_IP`

      When you are logged in as root, the prompt changes from `$` to `#`.

   b. Enter the password listed in the `Passwords.txt` file.
   c. Enter the following command to switch to root: `su -`
   d. Enter the password listed in the `Passwords.txt` file.
2. Check for running repairs: `repair-data show-ec-repair-status`
   - If you have never run a data repair job, the output is `No job found`. You don't need to restart any repair jobs.
   - If the data repair job was run previously or is running currently, the output lists information for the repair. Each repair has a unique repair ID. Go to the next step.
   - You can also monitor the status of restoration jobs in process and view a history of restoration jobs completed in Grid Manager.

```
root@ADM1-0:~ # repair-data show-ec-repair-status
             Repair ID    Affected Nodes / Volumes                Start Time                  End Time       State  Estimated Bytes Affected  Bytes Repaired  Percentage
==========================================================================================================================================================
  4216507958013005550   DC1-S1-0-182 (Volumes: 2)   2022-08-17T21:37:30.051543   2022-08-17T21:37:37.320998   Completed                1015788876               0           0
 18214680851049518682   DC1-S1-0-182 (Volumes: 1)   2022-08-17T20:37:58.869362   2022-08-17T20:38:45.299688   Completed                         0               0         100
  7962734388032289010   DC1-S1-0-182 (Volumes: 0)   2022-08-17T20:42:29.578740                                Stopped                                                 Unknown
```

3. If the State for all repairs is `Completed`, you don't need to restart any repair jobs.
4. If the State for any repair is `Stopped`, you must restart that repair.
   a. Obtain the repair ID for the failed repair from the output.
   b. Run the `repair-data start-ec-node-repair` command.

      Use the `--repair-id` option to specify the Repair ID. For example, if you want to retry a repair with repair ID 949292, run this command: `repair-data start-ec-node-repair --repair-id 949292`

   c. Continue to track the status of EC data repairs until the State for all repairs is `Completed`.

Gather required materials

Before performing a grid node decommission, you must obtain the following information.

Item	Notes
Recovery Package `.zip` file	You must download the most recent Recovery Package `.zip` file (`sgws-recovery-package-id-revision.zip`). You can use the Recovery Package file to restore the system if a failure occurs.
`Passwords.txt` file	This file contains the passwords required to access grid nodes on the command line and is included in the Recovery Package.
Provisioning passphrase	The passphrase is created and documented when the StorageGRID system is first installed. The provisioning passphrase is not in the `Passwords.txt` file.
Description of StorageGRID system's topology before decommissioning	If available, obtain any documentation that describes the system's current topology.

Related information

Web browser requirements

Access Decommission Nodes page

When you access the Decommission Nodes page in the Grid Manager, you can see at a glance which nodes can be decommissioned.

Before you begin

- You are signed in to the Grid Manager using a supported web browser.
- You have the Maintenance or Root access permission.

Steps

1. Select **MAINTENANCE** > **Tasks** > **Decommission**.
2. Select **Decommission Nodes**.

   The Decommission Nodes page appears. From this page, you can:

   - Determine which grid nodes can be decommissioned currently.
   - See the health of all grid nodes
   - Sort the list in ascending or descending order by **Name**, **Site**, **Type**, or **Has ADC**.
   - Enter search terms to quickly find particular nodes. For example, this page shows grid nodes in two data centers. The Decommission Possible column indicates that you can decommission the Gateway Node, one of the five Storage Nodes, and the non-primary Admin Node.

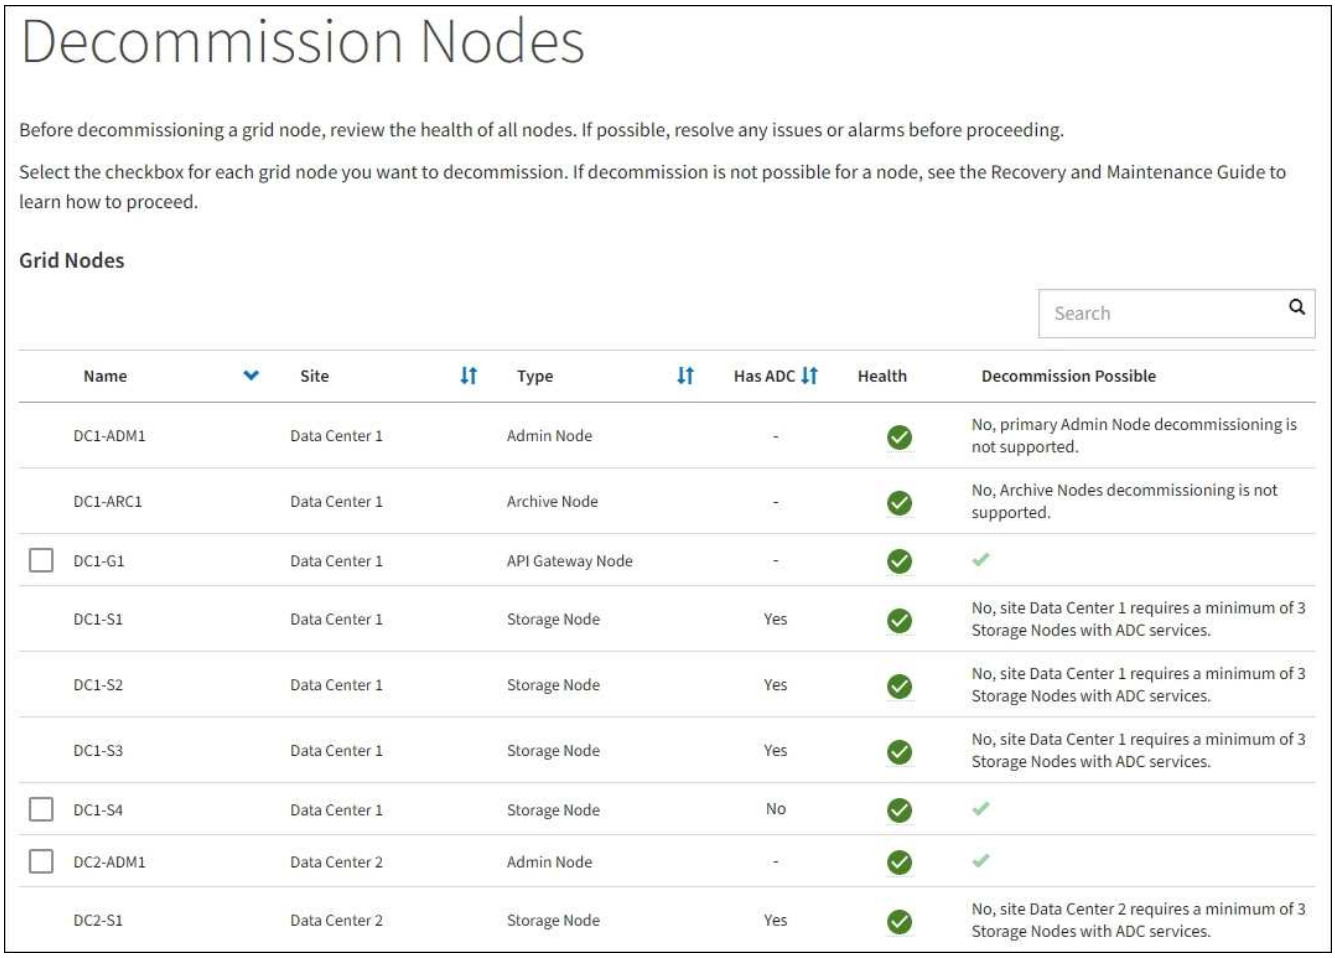

3. Review the **Decommission Possible** column for each node you want to decommission.

   If a grid node can be decommissioned, this column includes a green check mark, and the left column includes a checkbox. If a node can't be decommissioned, this column describes the issue. If there is more than one reason a node can't be decommissioned, the most critical reason is shown.

Decommission Possible reason	Description	Steps to resolve
No, node type decommissioning is not supported.	You can't decommission the primary Admin Node or an Archive Node.	None.
No, at least one grid node is disconnected. **Note:** This message is shown for connected grid nodes only.	You can't decommission a connected grid node if any grid node is disconnected. The **Health** column includes one of these icons for grid nodes that are disconnected: • (gray): Administratively Down • (blue): Unknown	Go to the step that lists the decommission procedure choices.

Decommission Possible reason	Description	Steps to resolve
No, one or more required nodes is currently disconnected and must be recovered. **Note:** This message is shown for disconnected grid nodes only.	You can't decommission a disconnected grid node if one or more required nodes is also disconnected (for example, a Storage Node that is required for the ADC quorum).	a. Review the Decommission Possible messages for all disconnected nodes. b. Determine which nodes can't be decommissioned because they are required. ◦ If the Health of a required node is Administratively Down, bring the node back online. ◦ If the health of a required node is Unknown, perform a node recovery procedure to recover the required node.
No, member of HA group(s): *x*. Before you can decommission this node, you must remove it from all HA groups.	You can't decommission an Admin Node or a Gateway Node if a node interface belongs to a high availability (HA) group.	Edit the HA group to remove the node's interface or remove the entire HA group. See the instructions for administering StorageGRID.
No, site *x* requires a minimum of *n* Storage Nodes with ADC services.	**Storage Nodes only.** You can't decommission a Storage Node if insufficient nodes would remain at the site to support ADC quorum requirements.	Perform an expansion. Add a new Storage Node to the site, and specify that it should have an ADC service. See information about the ADC quorum.

Decommission Possible reason	Description	Steps to resolve
No, one or more erasure coding profiles need at least *n* Storage Nodes. If the profile is not used in an ILM rule, you can deactivate it.	**Storage Nodes only.** You can't decommission a Storage Node unless enough nodes would remain for the existing erasure coding profiles. For example, if an erasure coding profile exists for 4+2 erasure coding, at least 6 Storage Nodes must remain.	For each affected erasure coding profile, perform one of the following steps, based on how the profile is being used: • **Used in the active ILM policy**: Perform an expansion. Add enough new Storage Nodes to allow erasure coding to continue. See the instructions for expanding your grid. • **Used in an ILM rule but not in the active ILM policy**: Edit or delete the rule and then deactivate the erasure coding profile. • **Not used in any ILM rule**: Deactivate the erasure coding profile. **Note:** An error message appears if you attempt to deactivate an erasure coding profile and object data is still associated with the profile. You might need to wait several weeks before trying the deactivation process again. Learn about deactivating an erasure coding profile in the instructions for managing objects with ILM.

4. If decommissioning is possible for the node, determine which procedure you need to perform:

If your grid includes…	Go to…
Any disconnected grid nodes	Decommission disconnected grid nodes
Only connected grid nodes	Decommission connected grid nodes

Decommission disconnected grid nodes

You might need to decommission a node that is not currently connected to the grid (one whose Health is Unknown or Administratively Down).

Before you begin

- You understand the requirements and considerations for decommissioning grid nodes.
- You have obtained all prerequisite items.
- You have ensured that no data repair jobs are active. See Check data repair jobs.
- You have confirmed that Storage Node recovery is not in progress anywhere in the grid. If it is, you must wait until any Cassandra rebuild performed as part of the recovery is complete. You can then proceed with decommissioning.
- You have ensured that other maintenance procedures will not be run while the node decommission procedure is running, unless the node decommission procedure is paused.
- The **Decommission Possible** column for the disconnected node or nodes you want to decommission includes a green check mark.
- You have the provisioning passphrase.

About this task

You can identify disconnected nodes by looking for Unknown (blue) or Administratively Down (gray) icons in the **Health** column. In the example, the Storage Node named DC1-S4 is disconnected; all of the other nodes are connected.

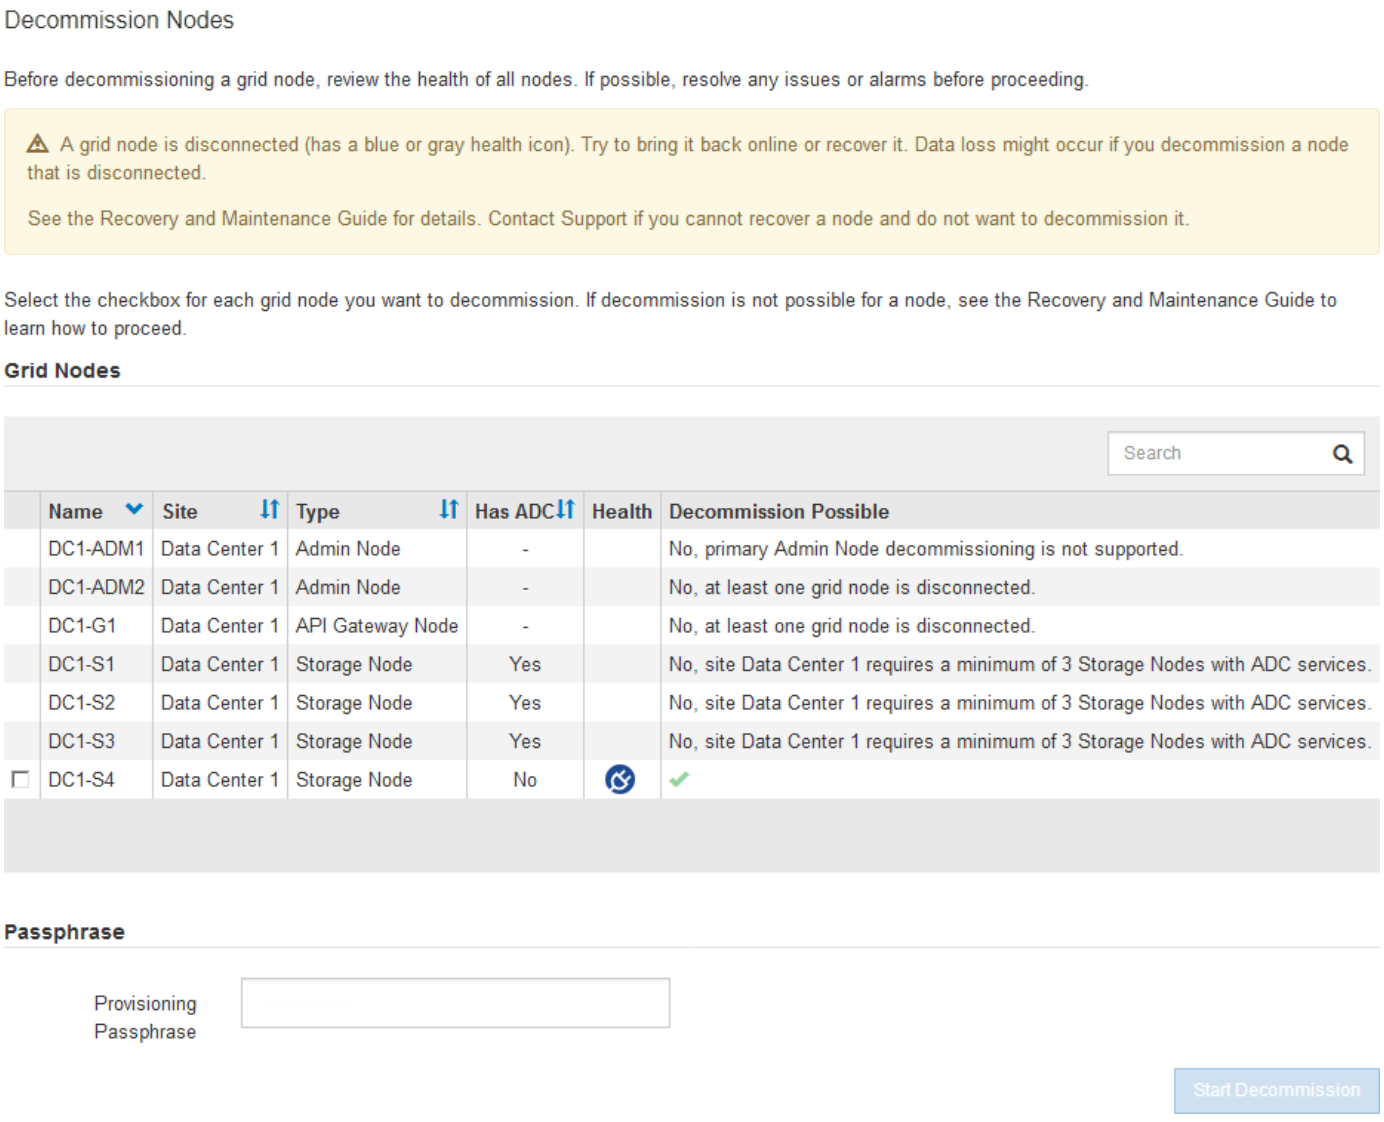

Before decommissioning any disconnected node, note the following:

- This procedure is primarily intended for removing a single disconnected node. If your grid contains multiple disconnected nodes, the software requires you to decommission them all at the same time, which increases the potential for unexpected results.

> Be careful when decommissioning more than one disconnected grid node at a time, especially if you are selecting multiple disconnected Storage Nodes.

- If a disconnected node can't be removed (for example, a Storage Node that is required for the ADC quorum), no other disconnected node can be removed.

Before decommissioning a disconnected **Storage Node**, note the following

- You should never decommission a disconnected Storage Node unless you are sure it can't be brought online or recovered.

> If you believe that object data can still be recovered from the node, don't perform this procedure. Instead, contact technical support to determine if node recovery is possible.

- If you decommission more than one disconnected Storage Node, data loss might occur. The system might not be able to reconstruct data if not enough object copies, erasure-coded fragments, or object metadata remain available.

> If you have more than one disconnected Storage Node that you can't recover, contact technical support to determine the best course of action.

- When you decommission a disconnected Storage Node, StorageGRID starts data repair jobs at the end of the decommissioning process. These jobs attempt to reconstruct the object data and metadata that was stored on the disconnected node.
- When you decommission a disconnected Storage Node, the decommission procedure completes relatively quickly. However, the data repair jobs can take days or weeks to run and aren't monitored by the decommission procedure. You must manually monitor these jobs and restart them as needed. See Check data repair jobs.
- If you decommission a disconnected Storage Node that contains the only copy of an object, the object will be lost. The data repair jobs can only reconstruct and recover objects if at least one replicated copy or enough erasure-coded fragments exist on Storage Nodes that are currently connected.

Before decommissioning a disconnected **Admin Node** or **Gateway Node**, note the following:

- When you decommission a disconnected Admin Node, you will lose the audit logs from that node; however, these logs should also exist on the primary Admin Node.
- You can safely decommission a Gateway Node while it is disconnected.

Steps

1. Attempt to bring any disconnected grid nodes back online or to recover them.

   See the recovery procedures for instructions.

2. If you are unable to recover a disconnected grid node and you want to decommission it while it is disconnected, select the checkbox for that node.

If your grid contains multiple disconnected nodes, the software requires you to decommission them all at the same time, which increases the potential for unexpected results.

Be careful when choosing to decommission more than one disconnected grid node at a time, especially if you are selecting multiple disconnected Storage Nodes. If you have more than one disconnected Storage Node that you can't recover, contact technical support to determine the best course of action.

3. Enter the provisioning passphrase.

   The **Start Decommission** button is enabled.

4. Click **Start Decommission**.

   A warning appears, indicating that you have selected a disconnected node and that object data will be lost if the node has the only copy of an object.

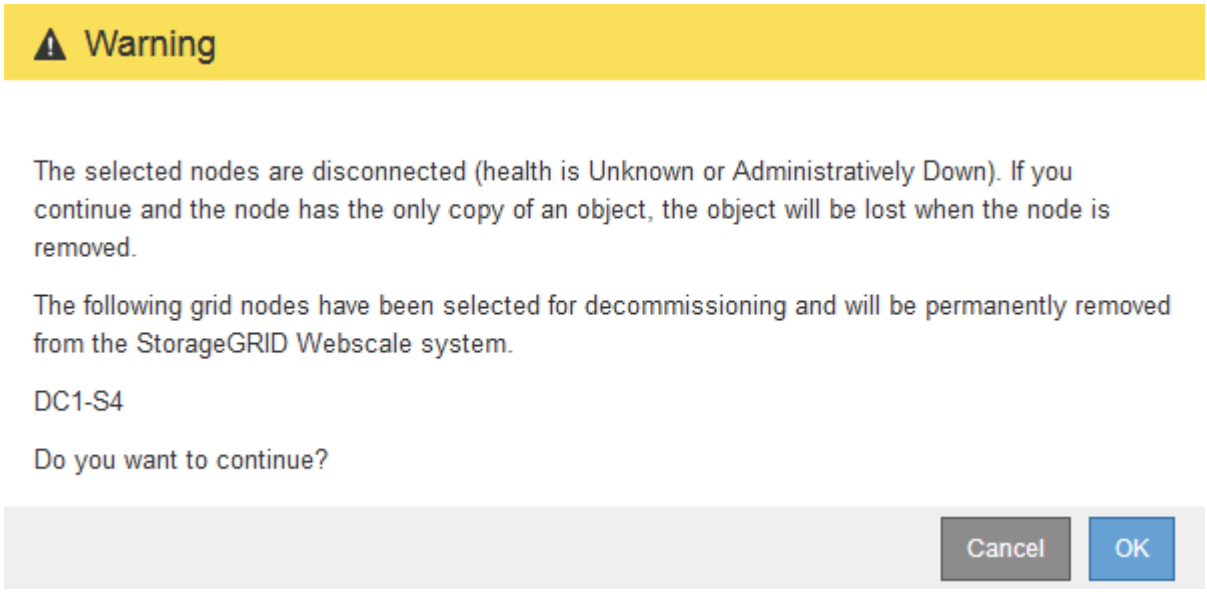

5. Review the list of nodes, and click **OK**.

   The decommission procedure starts, and the progress is displayed for each node. During the procedure, a new Recovery Package is generated containing the grid configuration change.

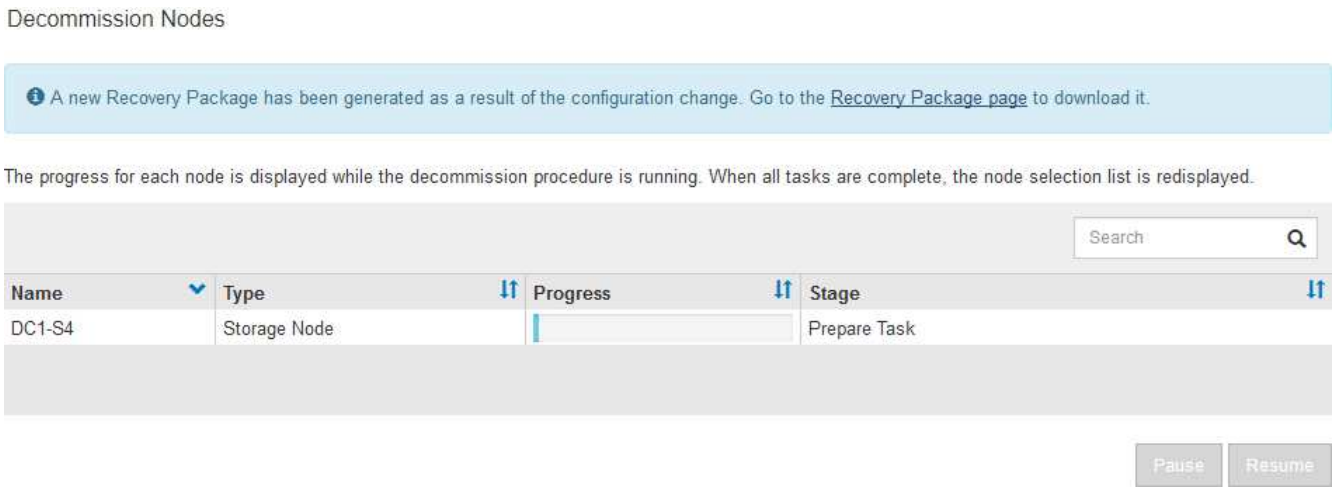

6. As soon as the new Recovery Package is available, click the link or select **MAINTENANCE** > **System** > **Recovery package** to access the Recovery Package page. Then, download the `.zip` file.

   See the instructions for downloading the Recovery Package.

   > Download the Recovery Package as soon as possible to ensure you can recover your grid if something goes wrong during the decommission procedure.

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

7. Periodically monitor the Decommission page to ensure that all selected nodes are decommissioned successfully.

   Storage Nodes can take days or weeks to decommission. When all tasks are complete, the node selection list is redisplayed with a success message. If you decommissioned a disconnected Storage Node, an information message indicates that the repair jobs have been started.

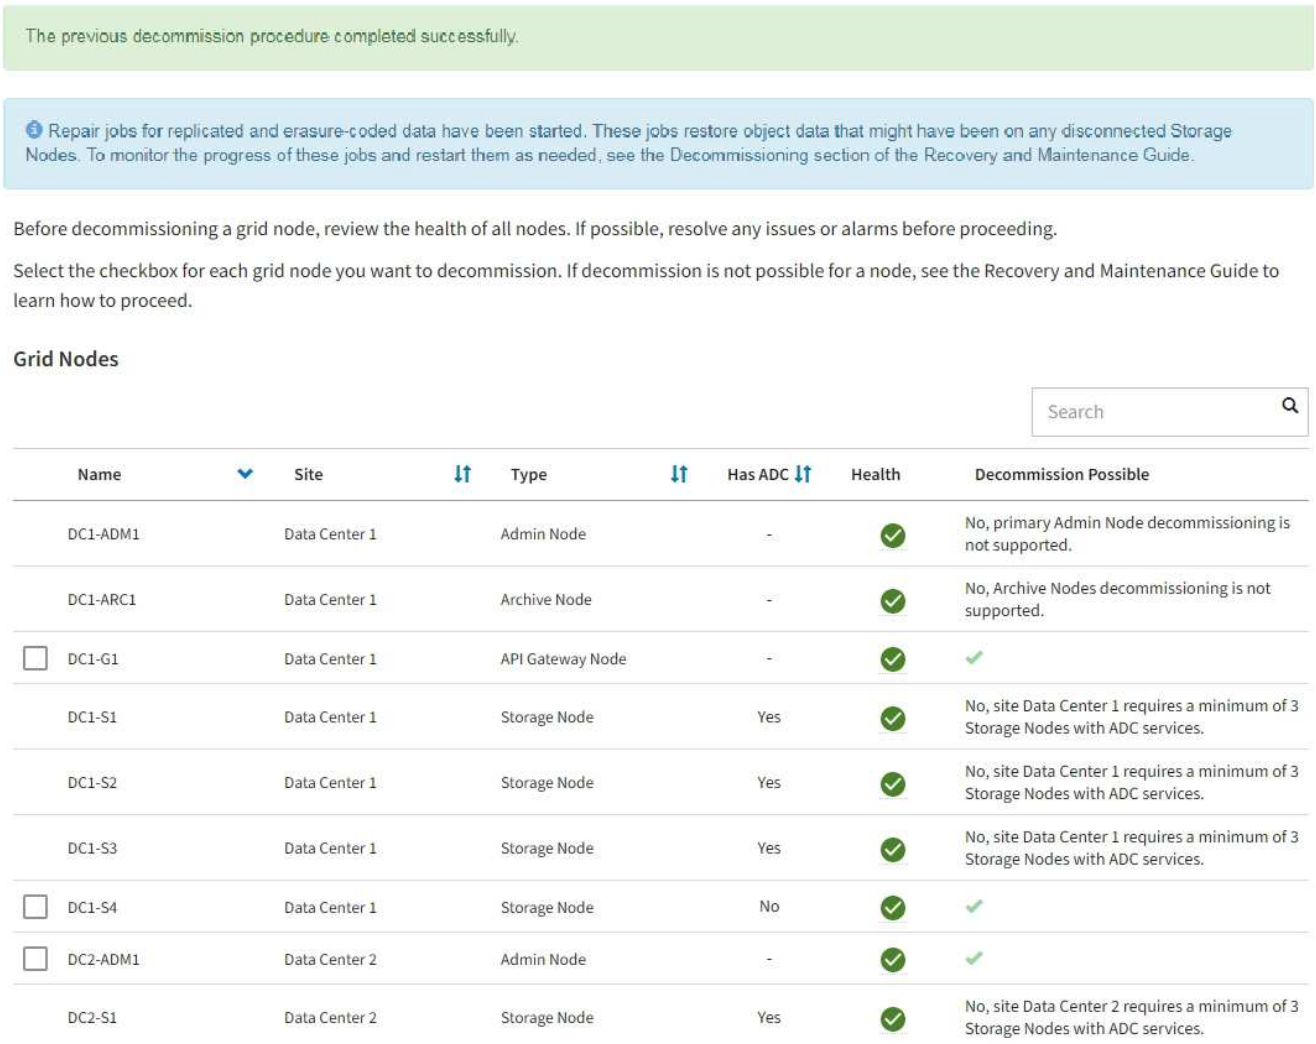

8. After the nodes have shut down automatically as part of the decommission procedure, remove any remaining virtual machines or other resources that are associated with the decommissioned node.

   Don’t perform this step until the nodes have shut down automatically.

9. If you are decommissioning a Storage Node, monitor the status of the **replicated data** and **erasure-coded (EC) data** repair jobs that are automatically started during the decommissioning process.

Replicated data

- To get an estimated percent completion for the replicated repair, add the `show-replicated-repair-status` option to the repair-data command.

  ```
  repair-data show-replicated-repair-status
  ```

- To determine if repairs are complete:
  1. Select **NODES** > ***Storage Node being repaired*** > **ILM**.
  2. Review the attributes in the Evaluation section. When repairs are complete, the **Awaiting - All** attribute indicates 0 objects.
- To monitor the repair in more detail:
  1. Select **SUPPORT** > **Tools** > **Grid topology**.
  2. Select ***grid*** > ***Storage Node being repaired*** > **LDR** > **Data Store**.
  3. Use a combination of the following attributes to determine, as well as possible, if replicated repairs are complete.

     Cassandra inconsistencies might be present, and failed repairs aren't tracked.

     - **Repairs Attempted (XRPA)**: Use this attribute to track the progress of replicated repairs. This attribute increases each time a Storage Node tries to repair a high-risk object. When this attribute does not increase for a period longer than the current scan period (provided by the **Scan Period — Estimated** attribute), it means that ILM scanning found no high-risk objects that need to be repaired on any nodes.

       High-risk objects are objects that are at risk of being completely lost. This does not include objects that don't satisfy their ILM configuration.

     - **Scan Period — Estimated (XSCM)**: Use this attribute to estimate when a policy change will be applied to previously ingested objects. If the **Repairs Attempted** attribute does not increase for a period longer than the current scan period, it is probable that replicated repairs are done. Note that the scan period can change. The **Scan Period — Estimated (XSCM)** attribute applies to the entire grid and is the maximum of all node scan periods. You can query the **Scan Period — Estimated** attribute history for the grid to determine an appropriate time frame.

Erasure coded (EC) data

To monitor the repair of erasure-coded data and retry any requests that might have failed:

1. Determine the status of erasure-coded data repairs:
   - Select **SUPPORT** > **Tools** > **Metrics** to view the estimated time to completion and the completion percentage for the current job. Then, select **EC Overview** in the Grafana section. Look at the **Grid EC Job Estimated Time to Completion** and **Grid EC Job Percentage Completed** dashboards.
   - Use this command to see the status of a specific `repair-data` operation:

     ```
     repair-data show-ec-repair-status --repair-id repair ID
     ```

   - Use this command to list all repairs:

```
repair-data show-ec-repair-status
```

The output lists information, including `repair ID`, for all previously and currently running repairs.

2. If the output shows that the repair operation failed, use the `--repair-id` option to retry the repair.

   This command retries a failed node repair, using the repair ID 6949309319275667690:

   ```
   repair-data start-ec-node-repair --repair-id 6949309319275667690
   ```

   This command retries a failed volume repair, using the repair ID 6949309319275667690:

   ```
   repair-data start-ec-volume-repair --repair-id 6949309319275667690
   ```

After you finish

As soon as the disconnected nodes have been decommissioned and all data repair jobs have been completed, you can decommission any connected grid nodes as required.

Then, complete these steps after you complete the decommission procedure:

- Ensure that the drives of the decommissioned grid node are wiped clean. Use a commercially available data wiping tool or service to permanently and securely remove data from the drives.
- If you decommissioned an appliance node and the data on the appliance was protected using node encryption, use the StorageGRID Appliance Installer to clear the key management server configuration (Clear KMS). You must clear the KMS configuration if you want to add the appliance to another grid. For instructions, see Monitor node encryption in maintenance mode.

Related information

Grid node recovery procedures

Decommission connected grid nodes

You can decommission and permanently remove nodes that are connected to the grid.

Before you begin

- You understand the requirements and considerations for decommissioning grid nodes.
- You have gathered all required materials.
- You have ensured that no data repair jobs are active.
- You have confirmed that Storage Node recovery is not in progress anywhere in the grid. If it is, wait until any Cassandra rebuild performed as part of the recovery is complete. You can then proceed with decommissioning.
- You have ensured that other maintenance procedures will not be run while the node decommission procedure is running, unless the node decommission procedure is paused.
- You have the provisioning passphrase.
- Grid nodes are connected.
- The **Decommission Possible** column for the node or nodes you want to decommission includes a green check mark.

The decommission will not start if one or more volumes are offline (unmounted) or if they are online (mounted) but in an error state.

If one or more volumes go offline while a decommission is in progress, the decommission process completes after these volumes have come back online.

- All grid nodes have Normal (green) health . If you see one of these icons in the **Health** column, you must try to resolve the issue:

Icon	Color	Severity
	Yellow	Notice
	Light orange	Minor
	Dark orange	Major
	Red	Critical

- If you previously decommissioned a disconnected Storage Node, the data repair jobs have all completed successfully. See Check data repair jobs.

Don’t remove a grid node’s virtual machine or other resources until instructed to do so in this procedure.

Steps

1. From the Decommission Nodes page, select the checkbox for each grid node you want to decommission.
2. Enter the provisioning passphrase.

   The **Start Decommission** button is enabled.

3. Click **Start Decommission**.
4. Review the list of nodes in the confirmation dialog, and click **OK**.

   The node decommission procedure starts, and the progress is displayed for each node. During the procedure, a new Recovery Package is generated to show the grid configuration change.

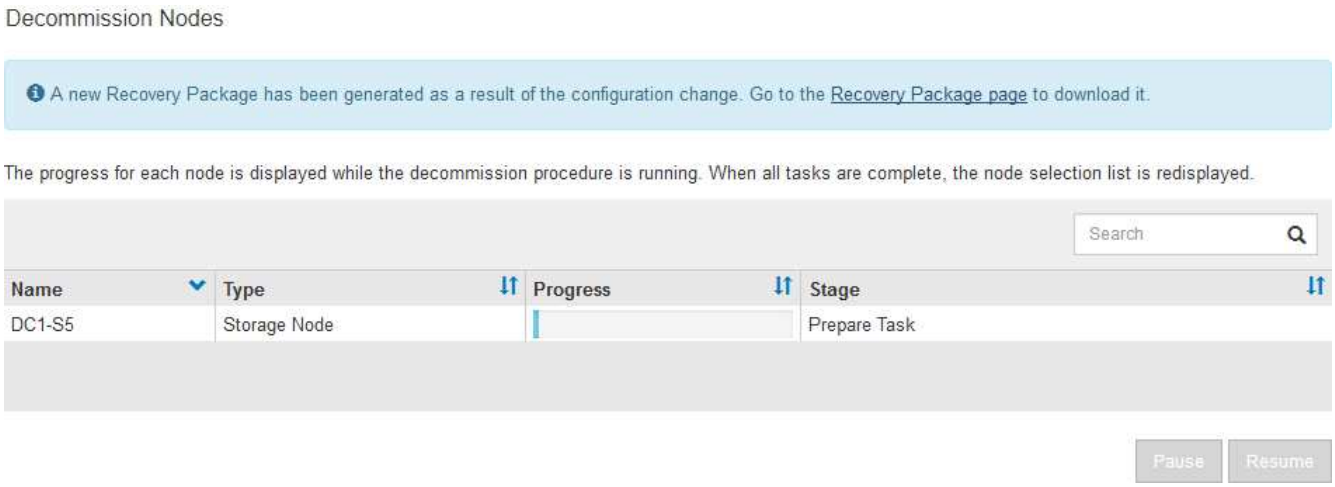

> Don't take a Storage Node offline after the decommission procedure has started. Changing the state might result in some content not being copied to other locations.

5. As soon as the new Recovery Package is available, click the link or select **MAINTENANCE** > **System** > **Recovery package** to access the Recovery Package page. Then, download the `.zip` file.

   See the instructions for downloading the Recovery Package.

   > Download the Recovery Package as soon as possible to ensure you can recover your grid if something goes wrong during the decommission procedure.

6. Periodically monitor the Decommission Nodes page to ensure that all selected nodes are decommissioned successfully.

   Storage Nodes can take days or weeks to decommission. When all tasks are complete, the node selection list is redisplayed with a success message.

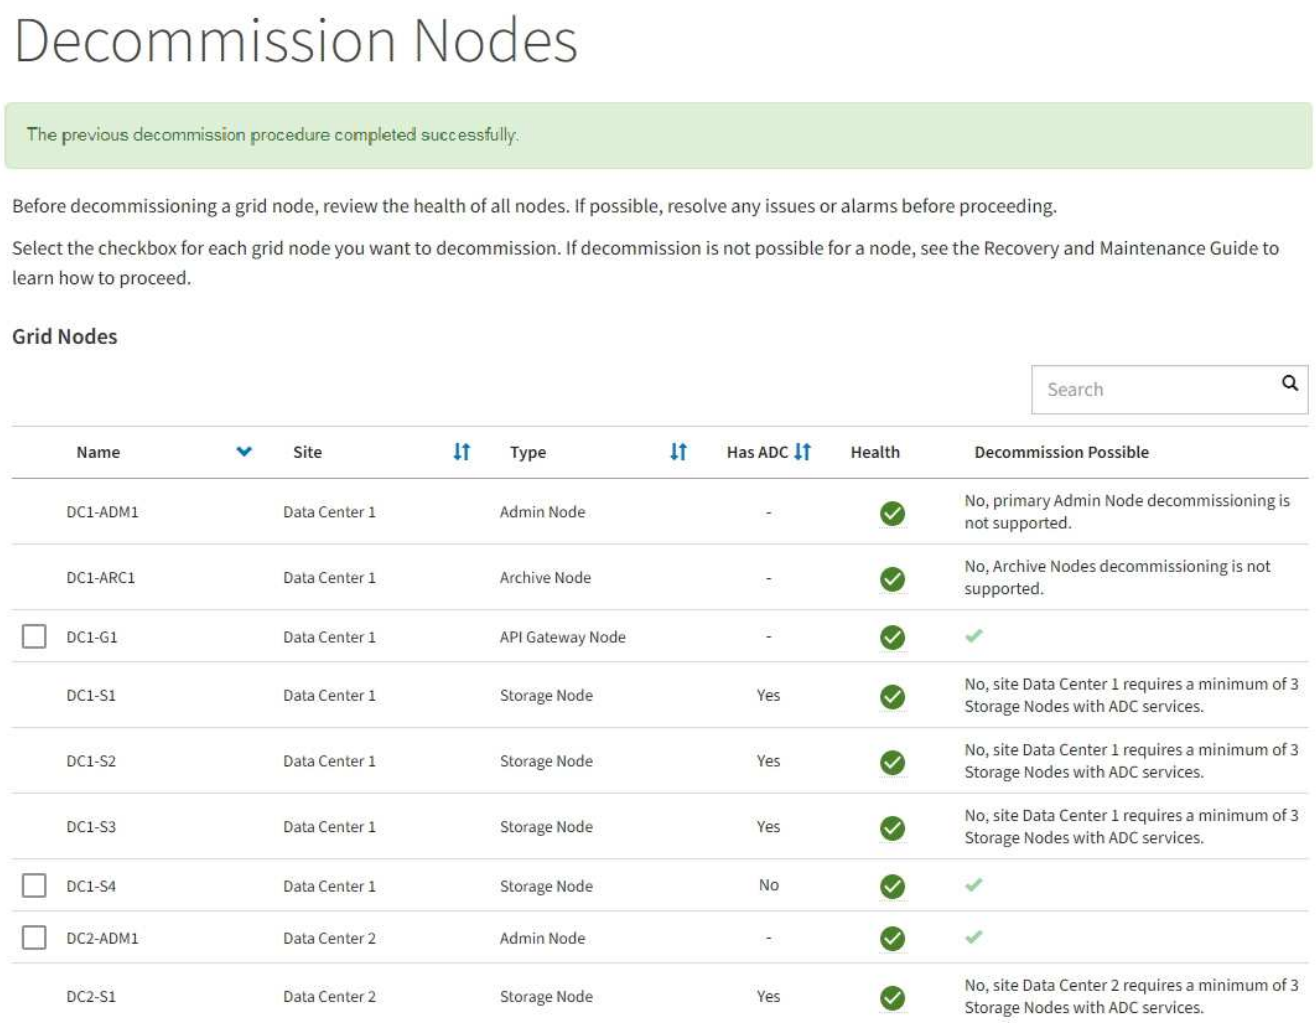

7. Follow the appropriate step for your platform. For example:
   - **Linux**: You might want to detach the volumes and delete the node configuration files you created during installation.
   - **VMware**: You might want to use the vCenter “Delete from Disk” option to delete the virtual machine. You might also need to delete any data disks that are independent of the virtual machine.
   - **StorageGRID appliance**: The appliance node automatically reverts to an undeployed state where you can access the StorageGRID Appliance Installer. You can power off the appliance or add it to another StorageGRID system.

Complete these steps after you complete the node decommission procedure:

- Ensure that the drives of the decommissioned grid node are wiped clean. Use a commercially available data wiping tool or service to permanently and securely remove data from the drives.
- If you decommissioned an appliance node and the data on the appliance was protected using node encryption, use the StorageGRID Appliance Installer to clear the key management server configuration (Clear KMS).You must clear the KMS configuration if you want to add the appliance to another grid. For instructions, see Monitor node encryption in maintenance mode.

Related information

Install Red Hat Enterprise Linux or CentOS

Pause and resume decommission process for Storage Nodes

If you need to perform a second maintenance procedure, you can pause the decommission procedure for a Storage Node during certain stages. After the other procedure is finished, you can resume decommissioning.

> The **Pause** button is enabled only when the ILM evaluation or erasure-coded data decommissioning stages are reached; however, ILM evaluation (data migration) will continue to run in the background.

Before you begin

- You are signed in to the Grid Manager using a supported web browser.
- You have the Maintenance or Root access permission.

Steps

1. Select **MAINTENANCE** > **Tasks** > **Decommission**.

   The Decommission page appears.

2. Select **Decommission Nodes**.

   The Decommission Nodes page appears. When the decommission procedure reaches either of the following stages, the **Pause** button is enabled.

   - Evaluating ILM
   - Decommissioning Erasure Coded data

3. Select **Pause** to suspend the procedure.

   The current stage is paused, and the **Resume** button is enabled.

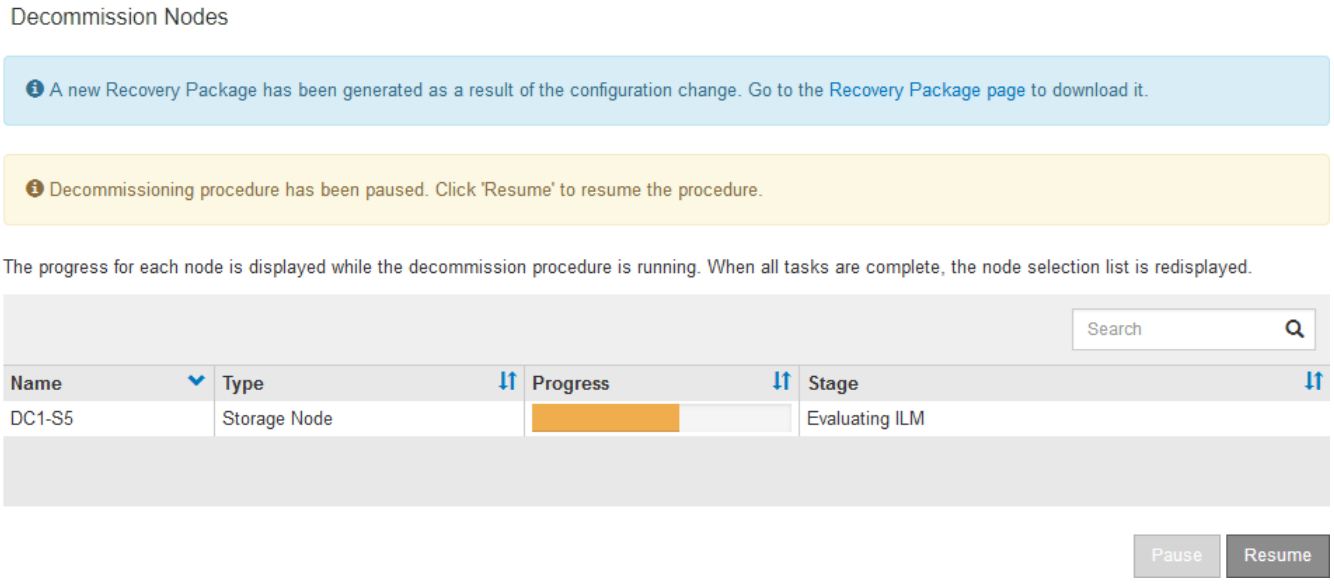

4. After the other maintenance procedure is finished, select **Resume** to proceed with the decommission.

Troubleshoot node decommissioning

If the node decommission procedure stops because of an error, you can take specific steps to troubleshoot the problem.

Before you begin

You are signed in to the Grid Manager using a supported web browser.

About this task

If you shut down the grid node being decommissioned, the task stops until the grid node is restarted. The grid node must be online.

Steps

1. Select **SUPPORT** > **Tools** > **Grid topology**.
2. In the Grid Topology tree, expand each Storage Node entry, and verify that the DDS and LDR services are both online.

   To perform Storage Node decommissioning, all nodes and all services need to be healthy at the start of an online node/site decommissioning.

3. To view the active grid tasks, select **primary Admin Node** > **CMN** > **Grid Tasks** > **Overview**.
4. Check the status of the decommissioning grid task.
   a. If the status of the decommissioning grid task indicates a problem with saving grid task bundles, select **primary Admin Node** > **CMN** > **Events** > **Overview**
   b. Check the number of Available Audit Relays.

      If the attribute Available Audit Relay is one or greater, the CMN service is connected to at least one ADC service. ADC services act as Audit Relays.

      The CMN service must be connected to at least one ADC service and a majority (50 percent plus one) of the StorageGRID system's ADC services must be available in order for a grid task to move from one stage of decommissioning to another and finish.

   c. If the CMN service is not connected to enough ADC services, ensure that Storage Nodes are online, and check network connectivity between the primary Admin Node and Storage Nodes.

Site decommission

Site decommission: Overview

You might need to remove a data center site from the StorageGRID system. To remove a site, you must decommission it.

The flowchart shows the high-level steps for decommissioning a site.

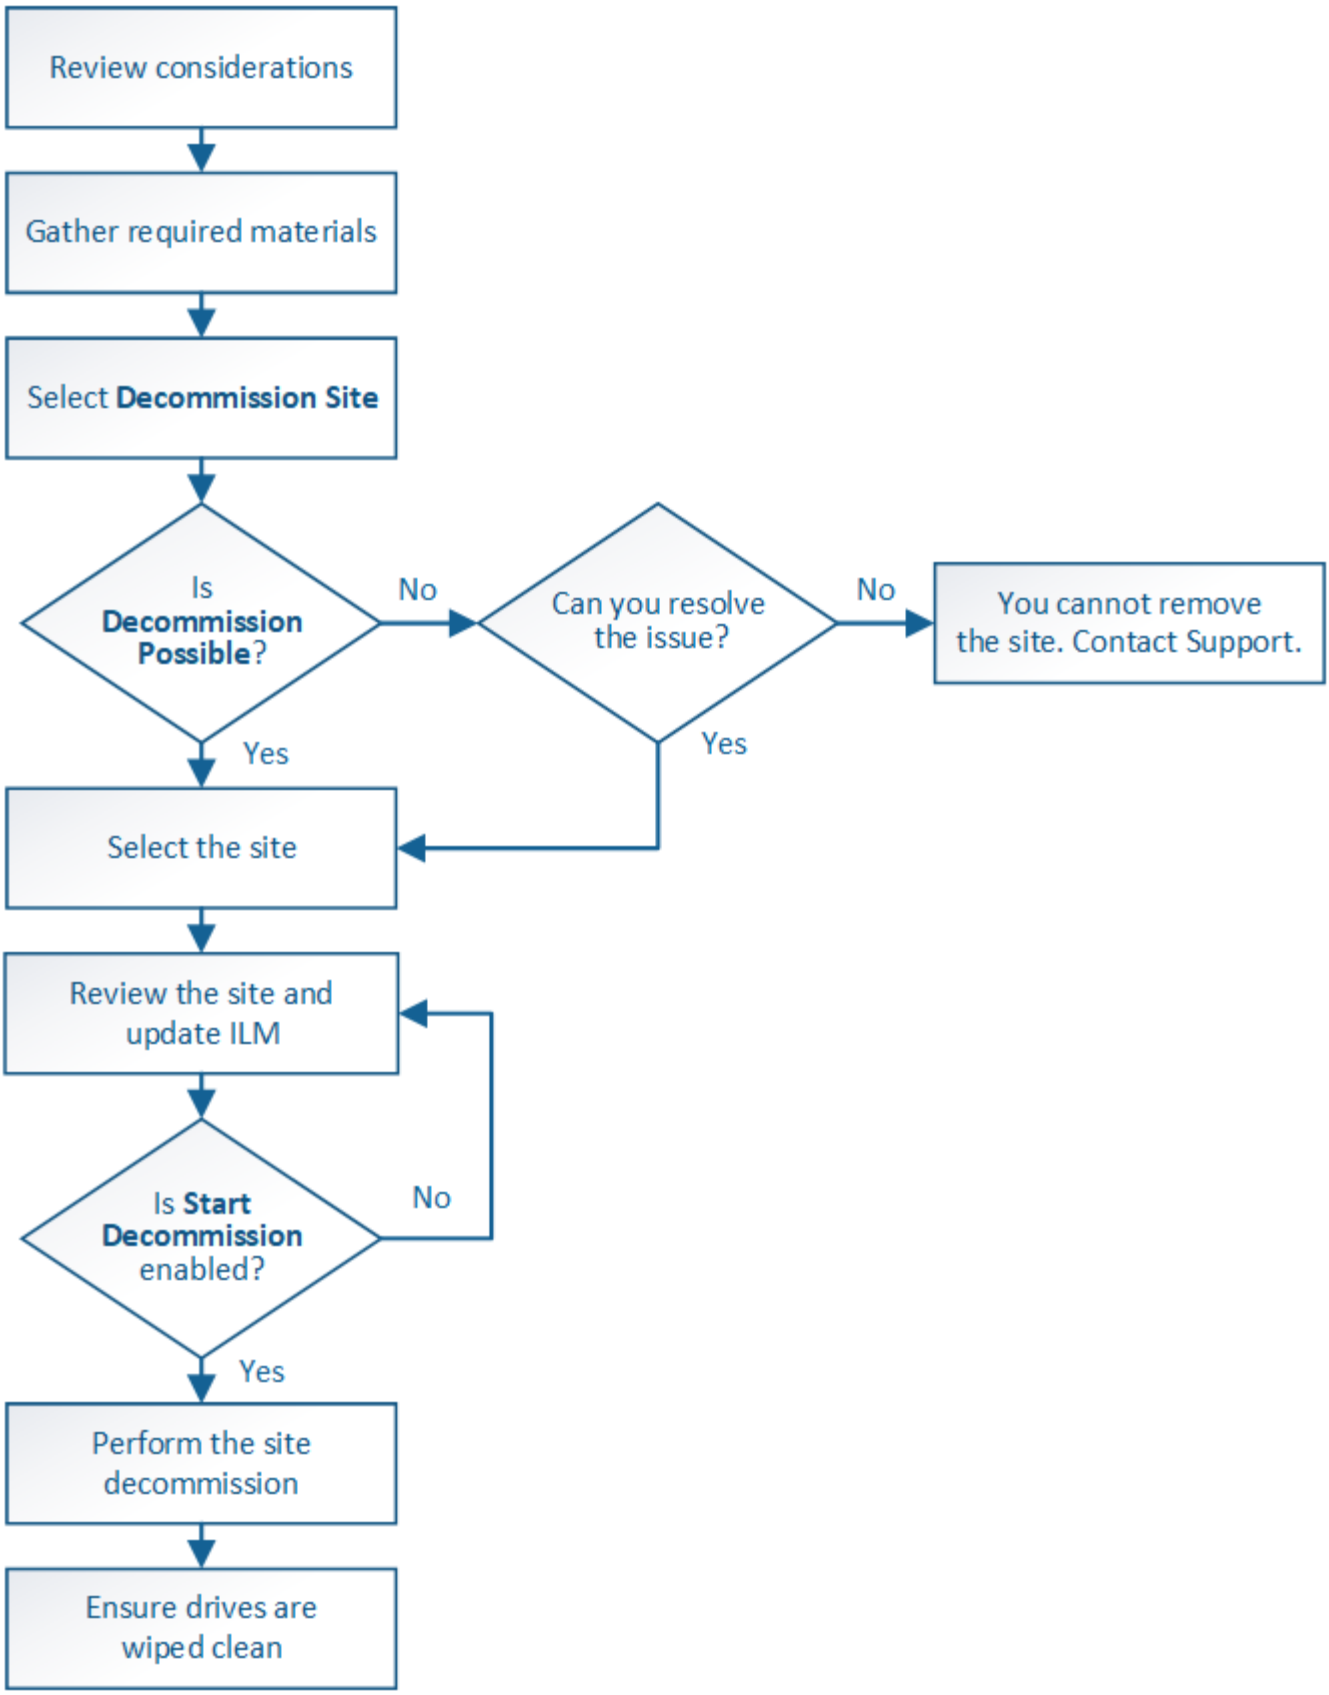

Considerations for removing a site

Before using the site decommission procedure to remove a site, you must review the considerations.

What happens when you decommission a site

When you decommission a site, StorageGRID permanently removes all nodes at the site and the site itself from the StorageGRID system.

When the site decommission procedure is complete:

- You can no longer use StorageGRID to view or access the site or any of the nodes at the site.
- You can no longer use any storage pools or erasure coding profiles that referred to the site. When StorageGRID decommissions a site, it automatically removes these storage pools and deactivates these erasure coding profiles.

Differences between connected site and disconnected site decommission procedures

You can use the site decommission procedure to remove a site in which all nodes are connected to StorageGRID (referred to as a connected site decommission) or to remove a site in which all nodes are disconnected from StorageGRID (referred to as a disconnected site decommission). Before you begin, you must understand the differences between these procedures.

> If a site contains a mixture of connected ( ) and disconnected nodes (or), you must bring all offline nodes back online.

- A connected site decommission allows you to remove an operational site from the StorageGRID system. For example, you can perform a connected site decommission to remove a site that is functional but no longer needed.
- When StorageGRID removes a connected site, it uses ILM to manage the object data at the site. Before you can start a connected site decommission, you must remove the site from all ILM rules and activate a new ILM policy. The ILM processes to migrate object data and the internal processes to remove a site can occur at the same time, but the best practice is to allow the ILM steps to complete before you start the actual decommission procedure.
- A disconnected site decommission allows you to remove a failed site from the StorageGRID system. For example, you can perform a disconnected site decommission to remove a site that has been destroyed by a fire or flood.

  When StorageGRID removes a disconnected site, it considers all nodes to be unrecoverable and makes no attempt to preserve data. However, before you can start a disconnected site decommission, you must remove the site from all ILM rules and activate a new ILM policy.

  > Before performing a disconnected site decommission procedure, you must contact your NetApp account representative. NetApp will review your requirements before enabling all steps in the Decommission Site wizard. You should not attempt a disconnected site decommission if you believe it might be possible to recover the site or to recover object data from the site.

General requirements for removing a connected or a disconnected site

Before removing a connected or disconnected site, you must be aware of the following requirements:

- You can’t decommission a site that includes the primary Admin Node.
- You can’t decommission a site that includes an Archive Node.

- You can't decommission a site if any of the nodes have an interface that belongs to a high availability (HA) group. You must either edit the HA group to remove the node's interface or remove the entire HA group.
- You can't decommission a site if it contains a mixture of connected ( ) and disconnected (or) nodes.
- You can't decommission a site if any node at any other site is disconnected (or).
- You can't start the site decommission procedure if an ec-node-repair operation is in progress. See Check data repair jobs to track repairs of erasure-coded data.
- While the site decommission procedure is running:
  - You can't create ILM rules that refer to the site being decommissioned. You also can't edit an existing ILM rule to refer to the site.
  - You can't perform other maintenance procedures, such as expansion or upgrade.

    If you need to perform another maintenance procedure during a connected site decommission, you can pause the procedure while the Storage Nodes are being removed. The **Pause** button is enabled only when the ILM evaluation or erasure-coded data decommissioning stages are reached; however, ILM evaluation (data migration) will continue to run in the background. After the second maintenance procedure is complete, you can resume decommissioning.

  - If you need to recover any node after starting the site decommission procedure, you must contact support.
- You can't decommission more than one site at a time.
- If the site includes one or more Admin Nodes and single sign-on (SSO) is enabled for your StorageGRID system, you must remove all relying party trusts for the site from Active Directory Federation Services (AD FS).

Requirements for information lifecycle management (ILM)

As part of removing a site, you must update your ILM configuration. The Decommission Site wizard guides you through a number of prerequisite steps to ensure the following:

- The site is not referred to by the active ILM policy. If it is, you must create and activate a new ILM policy with new ILM rules.
- No proposed ILM policy exists. If you have a proposed policy, you must delete it.
- No ILM rules refer to the site, even if those rules aren't used in the active or proposed policy. You must delete or edit all rules that refer to the site.

When StorageGRID decommissions the site, it will automatically deactivate any unused erasure coding profiles that refer to the site, and it will automatically delete any unused storage pools that refer to the site. If the All Storage Nodes storage pool exists (StorageGRID 11.6 and earlier), it is removed because it uses all sites.

Before you can remove a site, you might be required to create new ILM rules and activate a new ILM policy. These instructions assume that you have a good understanding of how ILM works and that you are familiar with creating storage pools, erasure coding profiles, ILM rules, and simulating and activating an ILM policy. See Manage objects with ILM.

Considerations for the object data at a connected site

If you are performing a connected site decommission, you must decide what to do with existing object data at the site when you create new ILM rules and a new ILM policy. You can do either or both of the following:

- Move object data from the selected site to one or more other sites in your grid.

  Example for moving data: Suppose you want to decommission a site in Raleigh because you added a new site in Sunnyvale. In this example, you want to move all object data from the old site to the new site. Before updating your ILM rules and ILM policy, you must review the capacity at both sites. You must ensure that the Sunnyvale site has enough capacity to accommodate the object data from the Raleigh site and that adequate capacity will remain in Sunnyvale for future growth.

  > To ensure that adequate capacity is available, you might need to expand your grid by adding storage volumes or Storage Nodes to an existing site or adding a new site before you perform this procedure.

- Delete object copies from the selected site.

  Example for deleting data: Suppose you currently use a 3-copy ILM rule to replicate object data across three sites. Before decommissioning a site, you can create an equivalent 2-copy ILM rule to store data at only two sites. When you activate a new ILM policy that uses the 2-copy rule, StorageGRID deletes the copies from the third site because they no longer satisfy ILM requirements. However, the object data will still be protected and the capacity of the two remaining sites will stay the same.

  > Never create a single-copy ILM rule to accommodate the removal of a site. An ILM rule that creates only one replicated copy for any time period puts data at risk of permanent loss. If only one replicated copy of an object exists, that object is lost if a Storage Node fails or has a significant error. You also temporarily lose access to the object during maintenance procedures such as upgrades.

Additional requirements for a connected site decommission

Before StorageGRID can remove a connected site, you must ensure the following:

- All nodes in your StorageGRID system must have a Connection State of **Connected** (); however, the nodes can have active alerts.

  > You can complete Steps 1-4 of the Decommission Site wizard if one or more nodes are disconnected. However, you can't complete Step 5 of the wizard, which starts the decommission process, unless all nodes are connected.

- If the site you plan to remove contains a Gateway Node or an Admin Node that is used for load balancing, you might need to expand your grid to add an equivalent new node at another site. Be sure clients can connect to the replacement node before starting the site decommission procedure.
- If the site you plan to remove contains any Gateway Node or Admin Nodes that are in an high availability (HA) group, you can complete Steps 1-4 of the Decommission Site wizard. However, you can't complete Step 5 of the wizard, which starts the decommission process, until you remove these nodes from all HA groups. If existing clients connect to an HA group that includes nodes from the site, you must ensure they can continue to connect to StorageGRID after the site is removed.
- If clients connect directly to Storage Nodes at the site you are planning to remove, you must ensure that

they can connect to Storage Nodes at other sites before starting the site decommission procedure.

- You must provide sufficient space on the remaining sites to accommodate any object data that will be moved because of changes to the active ILM policy. In some cases, you might need to expand your grid by adding Storage Nodes, storage volumes, or new sites before you can complete a connected site decommission.
- You must allow adequate time for the decommission procedure to complete. StorageGRID ILM processes might take days, weeks, or even months to move or delete object data from the site before the site can be decommissioned.

> Moving or deleting object data from a site might take days, weeks, or even months, depending on the amount of data at the site, the load on your system, network latencies, and the nature of the required ILM changes.

- Whenever possible, you should complete Steps 1-4 of the Decommission Site wizard as early as you can. The decommission procedure will complete more quickly and with fewer disruptions and performance impacts if you allow data to be moved from the site before starting the actual decommission procedure (by selecting **Start Decommission** in Step 5 of the wizard).

Additional requirements for a disconnected site decommission

Before StorageGRID can remove a disconnected site, you must ensure the following:

- You have contacted your NetApp account representative. NetApp will review your requirements before enabling all steps in the Decommission Site wizard.

> You should not attempt a disconnected site decommission if you believe it might be possible to recover the site or to recover any object data from the site. See How site recovery is performed by technical support.

- All nodes at the site must have a Connection State of one of the following:
  - **Unknown** ( ): For an unknown reason, a node is disconnected or services on the node are unexpectedly down. For example, a service on the node might be stopped, or the node might have lost its network connection because of a power failure or unexpected outage.
  - **Administratively Down** ( ): The node is not connected to the grid for an expected reason. For example, the node or services on the node have been gracefully shut down.
- All nodes at all other sites must have a Connection State of **Connected** ( ); however, these other nodes can have active alerts.
- You must understand that you will no longer be able to use StorageGRID to view or retrieve any object data that was stored at the site. When StorageGRID performs this procedure, it makes no attempt to preserve any data from the disconnected site.

> If your ILM rules and policy were designed to protect against the loss of a single site, copies of your objects still exist on the remaining sites.

- You must understand that if the site contained the only copy of an object, the object is lost and can't be retrieved.

Considerations for consistency controls when you remove a site

The consistency level for an S3 bucket or Swift container determines whether StorageGRID fully replicates object metadata to all nodes and sites before telling a client that object ingest was successful. Consistency controls provide a balance between the availability of the objects and the consistency of those objects across different Storage Nodes and sites.

When StorageGRID removes a site, it needs to ensure that no data is written to the site being removed. As a result, it temporarily overrides the consistency level for each bucket or container. After you start the site decommission process, StorageGRID temporarily uses strong-site consistency to prevent object metadata from being written to the site being removed.

As a result of this temporary override, be aware that any client write, update, and delete operations that occur during a site decommission can fail if multiple nodes become unavailable at the remaining sites.

Gather required materials

Before you decommission a site, you must obtain the following materials.

Item	Notes
Recovery Package `.zip` file	You must download the most recent Recovery Package `.zip` file (`sgws-recovery-package-id-revision.zip`). You can use the Recovery Package file to restore the system if a failure occurs. Download the Recovery Package
`Passwords.txt` file	This file contains the passwords required to access grid nodes on the command line and is included in the Recovery Package.
Provisioning passphrase	The passphrase is created and documented when the StorageGRID system is first installed. The provisioning passphrase is not in the `Passwords.txt` file.
Description of StorageGRID system's topology before decommissioning	If available, obtain any documentation that describes the system's current topology.

Related information

Web browser requirements

Step 1: Select Site

To determine if a site can be decommissioned, start by accessing the Decommission Site wizard.

Before you begin

- You have obtained all required materials.
- You have reviewed the considerations for removing a site.
- You are signed in to the Grid Manager using a supported web browser.

- You have the Root access permission, or the Maintenance and ILM permissions.

Steps

1. Select **MAINTENANCE** > **Tasks** > **Decommission**.
2. Select **Decommission Site**.

   Step 1 (Select Site) of the Decommission Site wizard appears. This step includes an alphabetic list of the sites in your StorageGRID system.

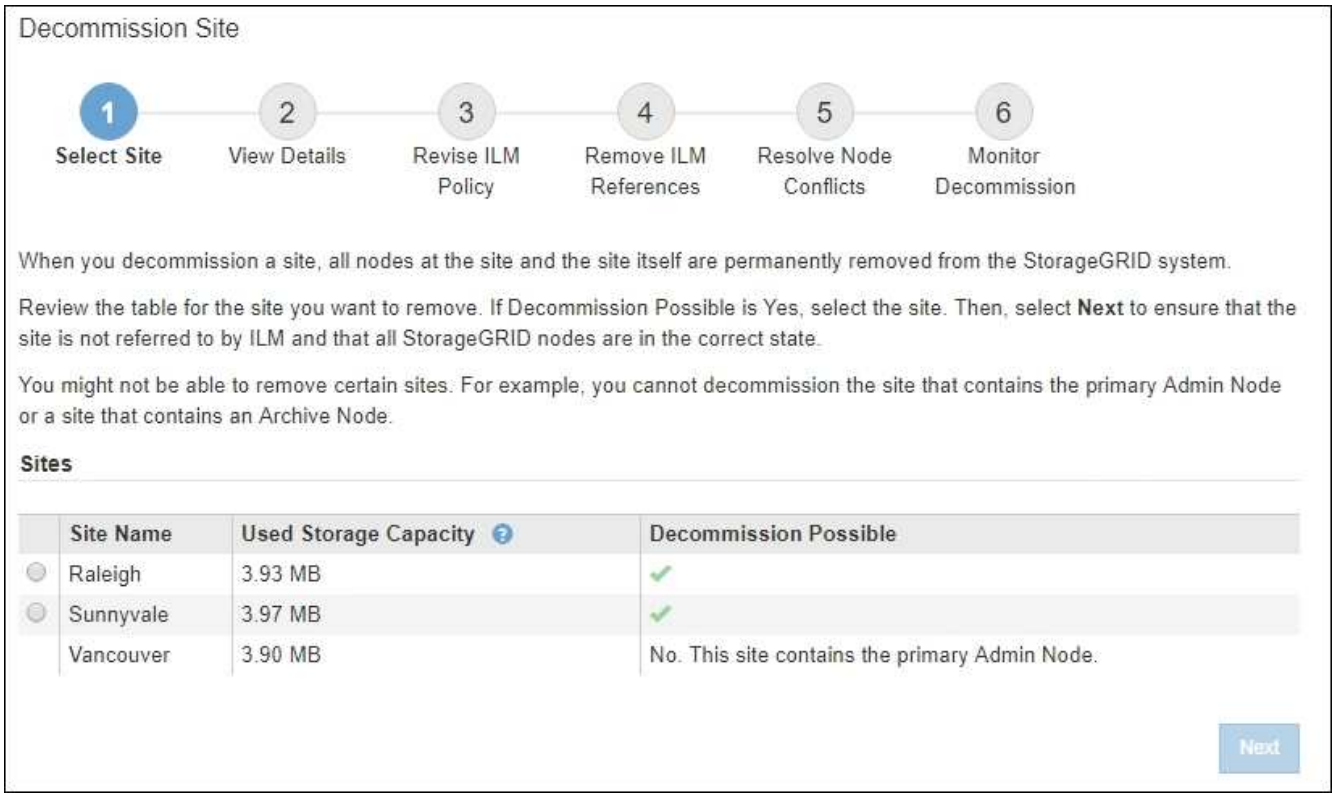

3. View the values in the **Used Storage Capacity** column to determine how much storage is currently being used for object data at each site.

   The Used Storage Capacity is an estimate. If nodes are offline, the Used Storage Capacity is the last known value for the site.

   - For a connected site decommission, this value represents how much object data will need to be moved to other sites or deleted by ILM before you can safely decommission this site.
   - For a disconnected site decommission, this value represents how much of your system's data storage will become inaccessible when you decommission this site.

   > If your ILM policy was designed to protect against the loss of a single site, copies of your object data should still exist on the remaining sites.

4. Review the reasons in the **Decommission Possible** column to determine which sites can be decommissioned currently.

   > If there is more than one reason a site can't be decommissioned, the most critical reason is shown.

Decommission Possible reason	Description	Next step
Green check mark ( )	You can decommission this site.	Go to the next step.
No. This site contains the primary Admin Node.	You can't decommission a site containing the primary Admin Node.	None. You can't perform this procedure.
No. This site contains one or more Archive Nodes.	You can't decommission a site containing an Archive Node.	None. You can't perform this procedure.
No. All nodes at this site are disconnected. Contact your NetApp account representative.	You can't perform a connected site decommission unless every node in the site is connected ( ).	If you want to perform a disconnected site decommission, you must contact your NetApp account representative, who will review your requirements and enable the rest of the Decommission Site wizard. **IMPORTANT**: Never take online nodes offline so that you can remove a site. You will lose data.

The example shows a StorageGRID system with three sites. The green check mark ( ) for the Raleigh and Sunnyvale sites indicates that you can decommission those sites. However, you can't decommission the Vancouver site because it contains the primary Admin Node.

5. If decommission is possible, select the radio button for the site.

   The **Next** button is enabled.

6. Select **Next**.

   Step 2 (View Details) appears.

Step 2: View Details

From Step 2 (View Details) of the Decommission Site wizard, you can review which nodes are included at the site, see how much space has been used on each Storage Node, and assess how much free space is available at the other sites in your grid.

Before you begin

Before decommissioning a site, you must review how much object data exists at the site.

- If you are performing a connected site decommission, you must understand how much object data currently exists at the site before updating ILM. Based on site capacities and your data protection needs, you can create new ILM rules to move data to other sites or to delete object data from the site.
- Perform any required Storage Node expansions before starting the decommission procedure if possible.

- If you are performing a disconnected site decommission, you must understand how much object data will become permanently inaccessible when you remove the site.

> If you are performing a disconnected site decommission, ILM can't move or delete object data. Any data that remains at the site will be lost. However, if your ILM policy was designed to protect against the loss of a single site, copies of your object data still exist on the remaining sites. See Enable site-loss protection.

Steps

1. From Step 2 (View Details), review any warnings related to the site you selected to remove.

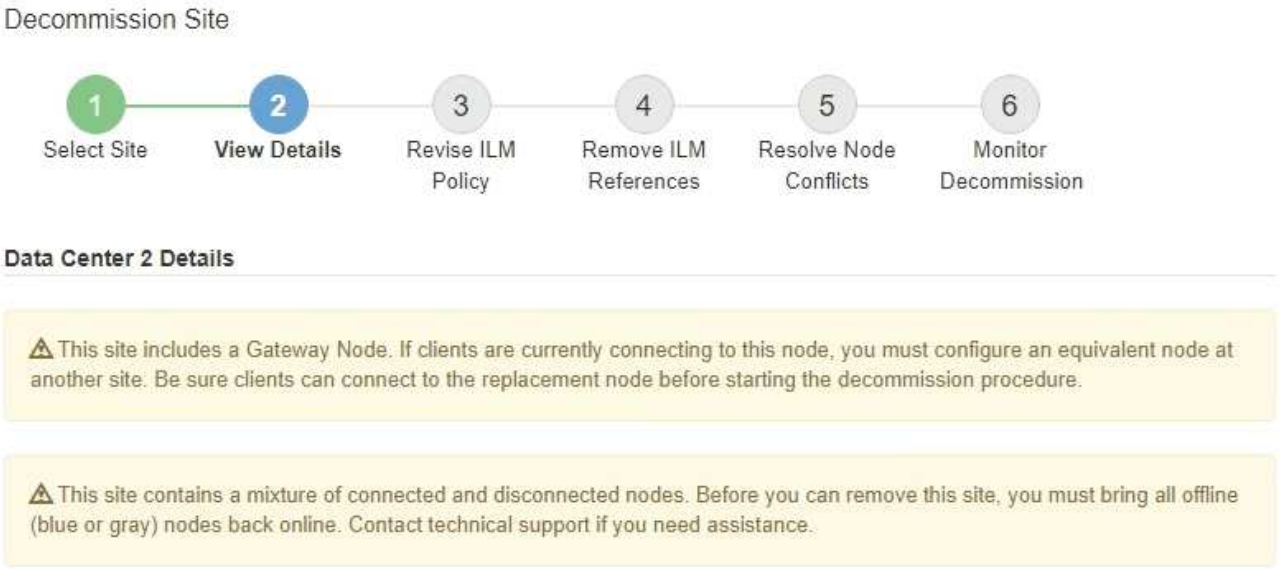

   A warning appears in these cases:

   - The site includes a Gateway Node. If S3 and Swift clients are currently connecting to this node, you must configure an equivalent node at another site. Be sure clients can connect to the replacement node before continuing with the decommission procedure.
   - The site contains a mixture of connected ( ) and disconnected nodes (or). Before you can remove this site, you must bring all offline nodes back online.

2. Review details about the site you selected to remove.

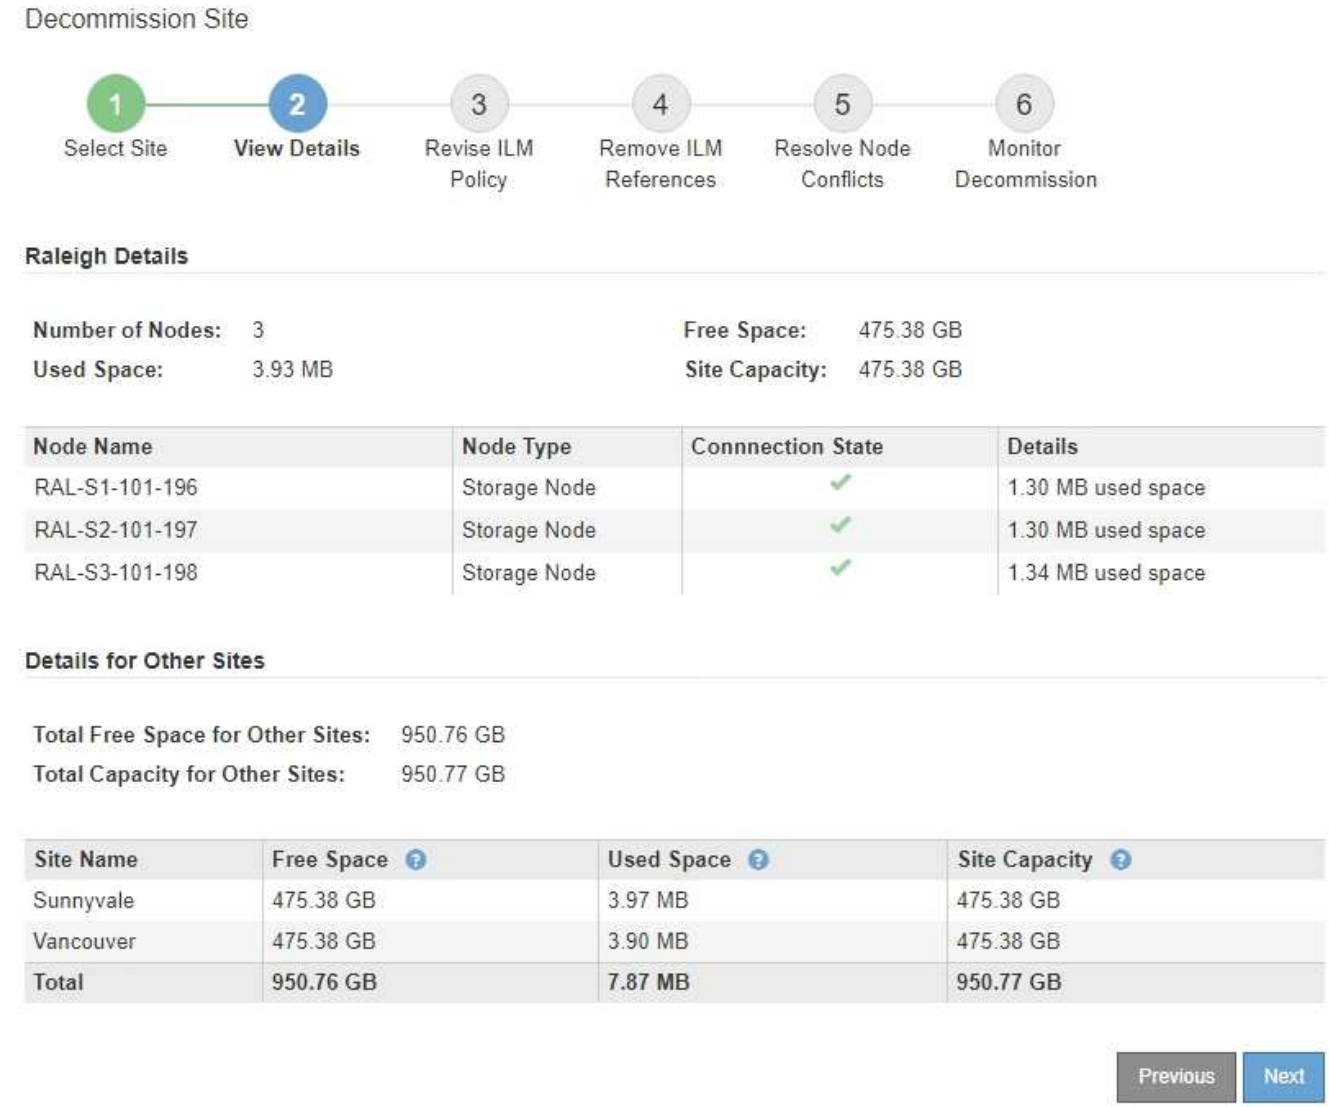

The following information is included for the selected site:

- Number of nodes
- The total used space, free space, and capacity of all Storage Nodes in the site.
  - For a connected site decommission, the **Used Space** value represents how much object data must be moved to other sites or deleted with ILM.
  - For a disconnected site decommission, the **Used Space** value indicates how much object data will become inaccessible when you remove the site.
- Node names, types, and connection states:
  - (Connected)
  - (Administratively Down)
  - (Unknown)
- Details about each node:
  - For each Storage Node, the amount of space that has been used for object data.

- For Admin Nodes and Gateway Nodes, whether the node is currently used in a high availability (HA) group. You can't decommission an Admin Node or a Gateway Node that is used in an HA group. Before you start the decommission, edit HA groups to remove all nodes at the site or remove the HA group if it only includes nodes from this site. For instructions, see Manage high availability (HA) groups.

3. In the Details for Other Sites section of the page, assess how much space is available at the other sites in your grid.

Details for Other Sites

Total Free Space for Other Sites: 950.76 GB
Total Capacity for Other Sites: 950.77 GB

Site Name	Free Space	Used Space	Site Capacity
Sunnyvale	475.38 GB	3.97 MB	475.38 GB
Vancouver	475.38 GB	3.90 MB	475.38 GB
Total	950.76 GB	7.87 MB	950.77 GB

If you are performing a connected site decommission and you plan to use ILM to move object data from the selected site (instead of just deleting it), you must ensure that the other sites have enough capacity to accommodate the moved data and that adequate capacity remains for future growth.

A warning appears if the **Used Space** for the site you want to remove is greater than the **Total Free Space for Other Sites**. To ensure that adequate storage capacity is available after the site is removed, you might need to perform an expansion before performing this procedure.

4. Select **Next**.

   Step 3 (Revise ILM Policy) appears.

Step 3: Revise ILM Policy

From Step 3 (Revise ILM Policy) of the Decommission Site wizard, you can determine if the site is referred to by the active ILM policy.

Before you begin

You have a good understanding of how ILM works and you are familiar with creating storage pools, erasure coding profiles, ILM rules, and simulating and activating an ILM policy. See Manage objects with ILM.

About this task

StorageGRID can't decommission a site if that site is referred to by any ILM rule in the active ILM policy.

If your current ILM policy refers to the site you want to remove, you must activate a new ILM policy that meets certain requirements. Specifically, the new ILM policy:

- Can't use a storage pool that refers to the site or uses the All Sites option.
- Can't use an erasure coding profile that refers to the site.
- Can't use the Make 2 Copies rule from StorageGRID 11.6 or earlier installations.

- Must be designed to fully protect all object data.

> Never create a single-copy ILM rule to accommodate the removal of a site. An ILM rule that creates only one replicated copy for any time period puts data at risk of permanent loss. If only one replicated copy of an object exists, that object is lost if a Storage Node fails or has a significant error. You also temporarily lose access to the object during maintenance procedures such as upgrades.

If you are performing a *connected site decommission*, you must consider how StorageGRID should manage the object data currently at the site you want to remove. Depending on your data protection requirements, the new rules can move existing object data to different sites or they can delete any extra object copies that are no longer needed.

Contact technical support if you need assistance designing the new policy.

Steps

1. From Step 3 (Revise ILM Policy), determine if any ILM rules in the active ILM policy refer to the site you selected to remove.
2. If no rules are listed, select **Next** to go to Step 4: Remove ILM References.
3. If one or more ILM rules are listed in the table, select the link next to **Active Policy Name**.

   The ILM policies page appears in a new browser tab. Use this tab to update ILM. The Decommission Site page will remain open on the other tab.

   a. If necessary, select **ILM** > **Storage pools** to create one or more storage pools that don't refer to the site.

   > For details, see the instructions for managing objects with information lifecycle management.

   b. If you plan to use erasure coding, select **ILM** > **Erasure coding** to create one or more erasure coding profiles.

   You must select storage pools that don't refer to the site.

   > Don't use the **All Storage Nodes** storage pool (StorageGRID 11.6 and earlier) in the erasure coding profiles.

4. Select **ILM** > **Rules** and clone each of the rules listed in the table for Step 3 (Revise ILM Policy).

   > For details, see the instructions for managing objects with information lifecycle management.

   a. Use names that will make it easy to select these rules in a new policy.

   b. Update the placement instructions.

   Remove any storage pools or erasure coding profiles that refer to the site and replace them with new storage pools or erasure coding profiles.

   > Don't use the **All Storage Nodes** storage pool in the new rules.

5. Select **ILM** > **Policies** and create a new policy that uses the new rules.

   For details, see the instructions for managing objects with information lifecycle management.

   a. Select the active policy, and select **Clone**.
   b. Specify a policy name and a reason for change.
   c. Select rules for the cloned policy.
      - Clear all rules listed for Step 3 (Revise ILM Policy) of the Decommission Site page.
      - Select a default rule that does not refer to the site.

        Don’t select the **Make 2 Copies** rule because that rule uses the **All Storage Nodes** storage pool, which is not allowed.

      - Select the other replacement rules you created. These rules should not refer to the site.

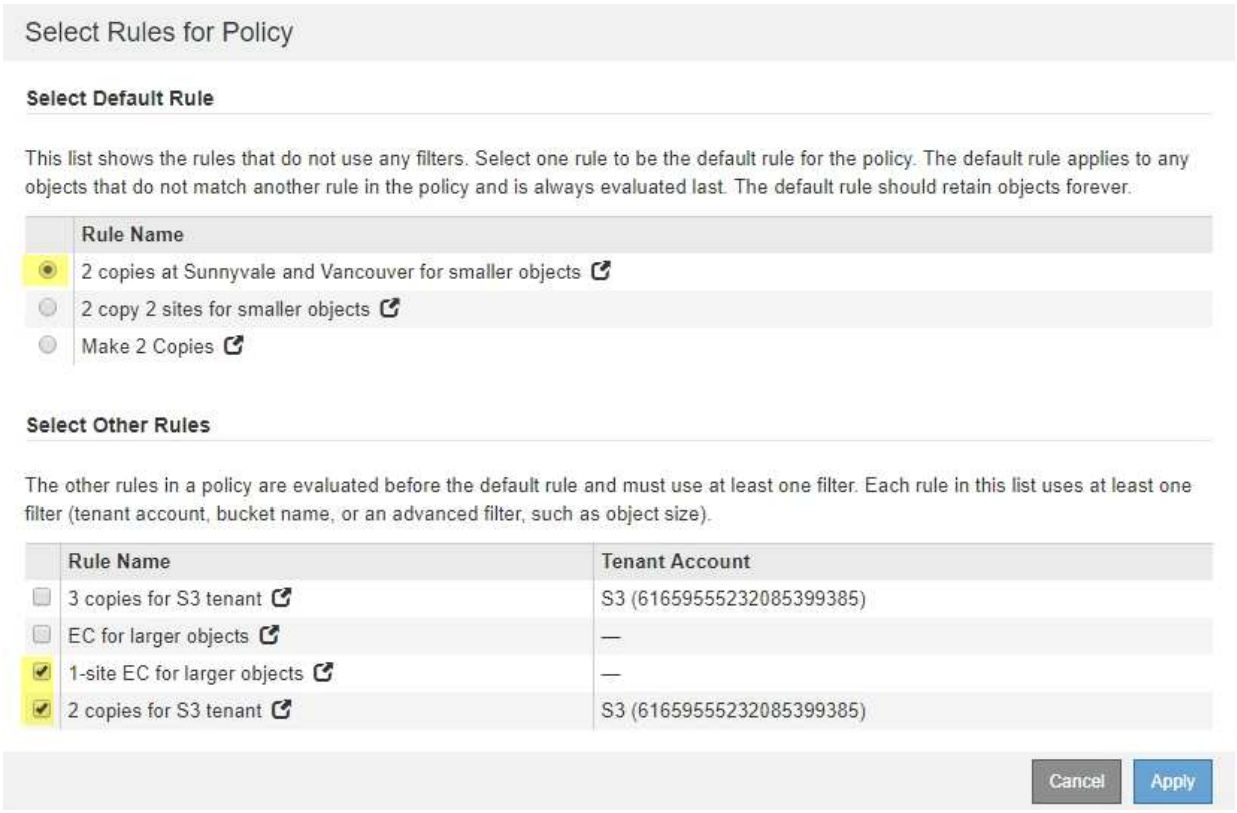

   d. Select **Apply**.
   e. Drag the rows to reorder the rules in the policy.

      You can’t move the default rule.

      You must confirm that the ILM rules are in the correct order. When the policy is activated, new and existing objects are evaluated by the rules in the order listed, starting at the top.

   f. Save the proposed policy.
6. Ingest test objects, and simulate the proposed policy to ensure that the correct rules are applied.

Errors in an ILM policy can cause unrecoverable data loss. Carefully review and simulate the policy before activating it to confirm that it will work as intended.

When you activate a new ILM policy, StorageGRID uses it to manage all objects, including existing objects and newly ingested objects. Before activating a new ILM policy, review any changes to the placement of existing replicated and erasure-coded objects. Changing an existing object's location might result in temporary resource issues when the new placements are evaluated and implemented.

7. Activate the new policy.

   If you are performing a connected site decommission, StorageGRID begins to remove object data from the selected site as soon as you activate the new ILM policy. Moving or deleting all object copies might take weeks. Although you can safely start a site decommission while object data still exists at the site, the decommission procedure will complete more quickly and with fewer disruptions and performance impacts if you allow data to be moved from the site before starting the actual decommission procedure (by selecting **Start Decommission** in Step 5 of the wizard).

8. Return to **Step 3 (Revise ILM Policy)** to ensure that no ILM rules in the new active policy refer to the site and the **Next** button is enabled.

   If any rules are listed, you must create and activate a new ILM policy before you can continue.

9. If no rules are listed, select **Next**.

   Step 4 (Remove ILM References) appears.

Step 4: Remove ILM References

From Step 4 (Remove ILM References) of the Decommission Site wizard, you can remove the proposed policy if one exists and delete or edit any unused ILM rules that still refer to the site.

About this task

You are prevented from starting the site decommission procedure in these cases:

- A proposed ILM policy exists. If you have a proposed policy, you must delete it.

- Any ILM rule refers to the site, even if that rule is not used in any ILM policy. You must delete or edit all rules that refer to the site.

Steps

1. If a proposed policy is listed, remove it.
   a. Select **Delete Proposed Policy**.
   b. Select **OK** in the confirmation dialog box.
2. Determine whether any unused ILM rules refer to the site.

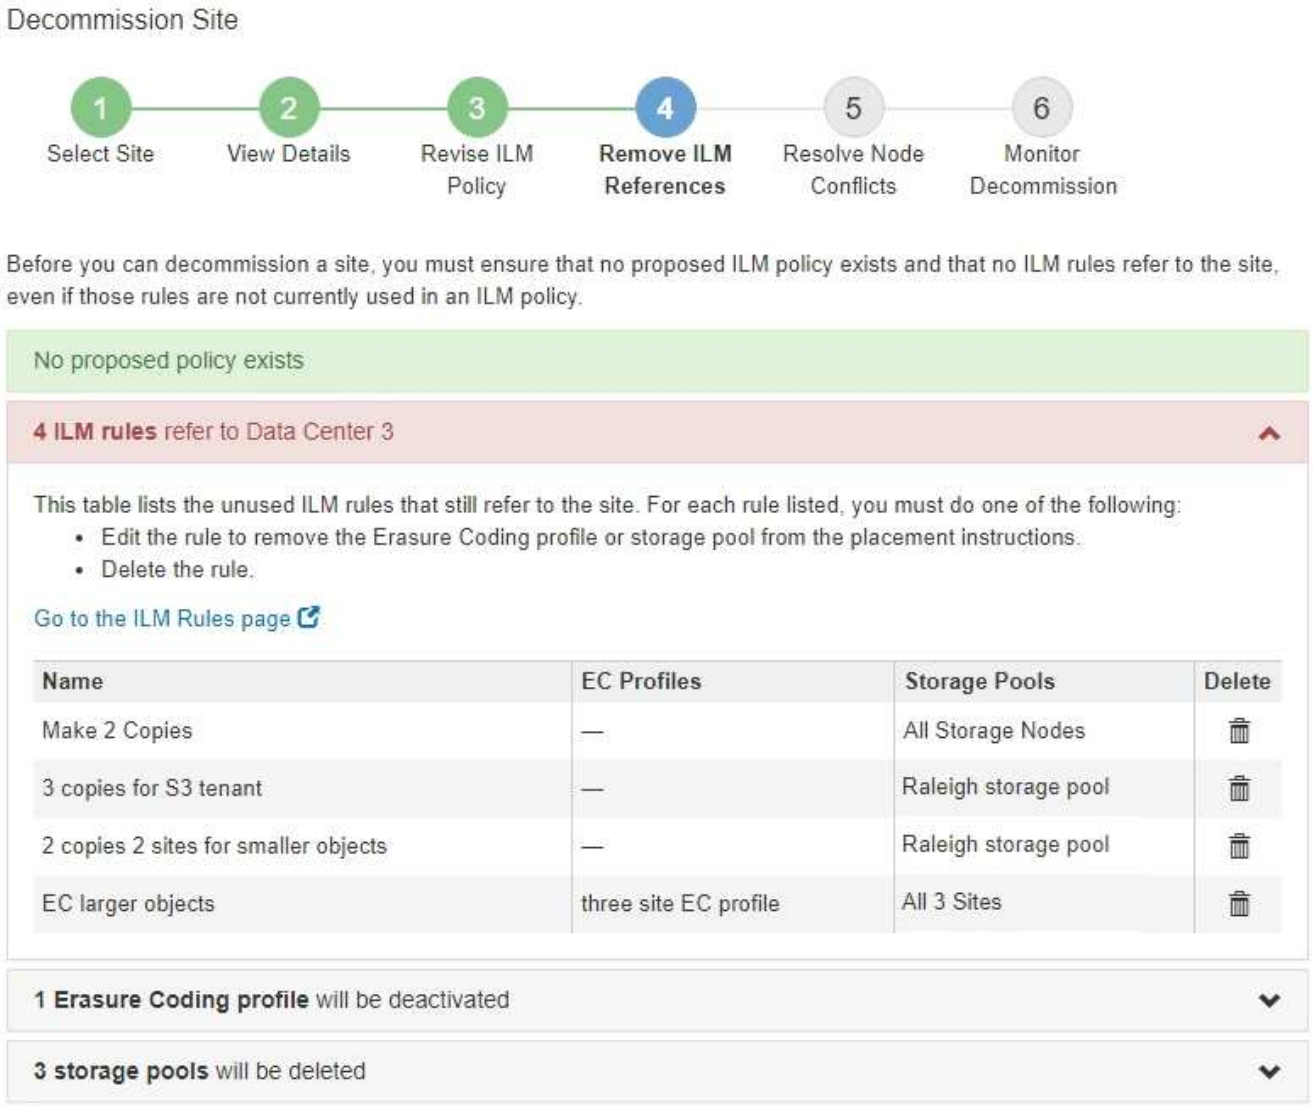

Any ILM rules that are listed still refer to the site but aren't used in any policy. In the example:

- The **Make 2 Copies** rule uses the **All Storage Nodes** storage pool (StorageGRID 11.6 and earlier), which uses the **All Sites** site.
- The unused **3 copies for S3 tenant** rule refers to the **Raleigh** storage pool.
- The unused **2 copy 2 sites for smaller objects** rule refers to the **Raleigh** storage pool.
- The unused **EC larger objects** rules uses the Raleigh site in the **All 3 Sites** erasure coding profile.
- If no ILM rules are listed, select **Next** to go to Step 5: Resolve Node Conflicts (and start decommission).

When StorageGRID decommissions the site, it will automatically deactivate any unused erasure coding profiles that refer to the site, and it will automatically delete any unused storage pools that refer to the site. The All Storage Nodes storage pool (StorageGRID 11.6 and earlier) is removed because it uses the All Sites site.

- If one or more ILM rules are listed, go to the next step.

3. Edit or delete each unused rule:
   - To edit a rule, go the ILM rules page and update all placements that use an erasure coding profile or storage pool that refers to the site. Then, return to **Step 4 (Remove ILM References)**.

     For details, see the instructions for managing objects with information lifecycle management.

   - To delete a rule, select the trash can icon and select **OK**.

     You must delete the **Make 2 Copies** rule before you can decommission a site.

4. Confirm that no proposed ILM policy exists, no unused ILM rules refer to the site, and the **Next** button is enabled.

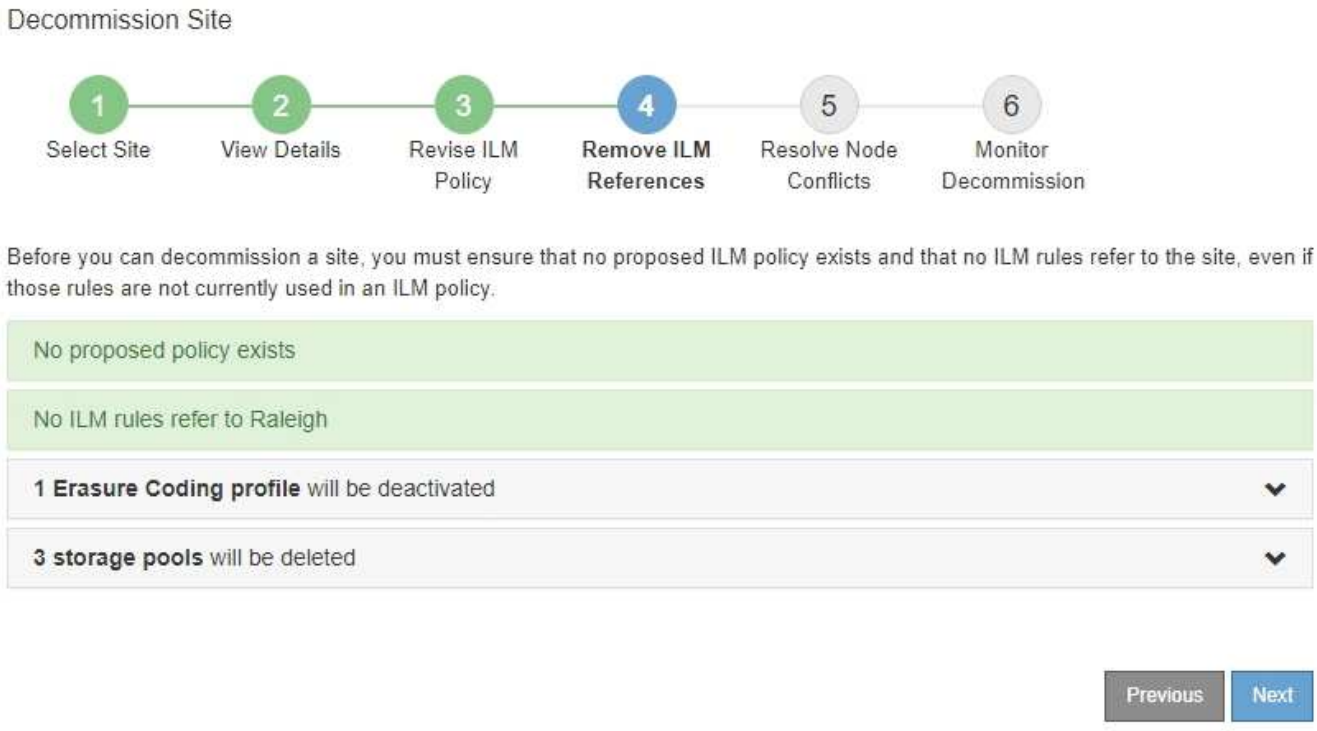

5. Select **Next**.

   Any remaining storage pools and erasure coding profiles that refer to the site will become invalid when the site is removed. When StorageGRID decommissions the site, it will automatically deactivate any unused erasure coding profiles that refer to the site, and it will automatically delete any unused storage pools that refer to the site. The All Storage Nodes storage pool (StorageGRID 11.6 and earlier) is removed because it uses the All Sites site.

   Step 5 (Resolve Node Conflicts) appears.

Step 5: Resolve Node Conflicts (and start decommission)

From Step 5 (Resolve Node Conflicts) of the Decommission Site wizard, you can determine if any nodes in your StorageGRID system are disconnected or if any nodes at the selected site belong to a high availability (HA) group. After any node conflicts are resolved, you start the decommission procedure from this page.

Before you begin

You must ensure that all nodes in your StorageGRID system are in the correct state, as follows:

- All nodes in your StorageGRID system must be connected ( ).

  If you are performing a disconnected site decommission, all nodes at the site you are removing must be disconnected, and all nodes at all other sites must be connected.

  The decommission will not start if one or more volumes are offline (unmounted) or if they are online (mounted) but in an error state.

  If one or more volumes go offline while a decommission is in progress, the decommission process completes after these volumes have come back online.

- No node at the site you are removing can have an interface that belongs to a high availability (HA) group.

About this task

If any node is listed for Step 5 (Resolve Node Conflicts), you must correct the issue before you can start the decommission.

Before starting the site decommission procedure from this page, review the following considerations:

- You must allow adequate time for the decommission procedure to complete.

  Moving or deleting object data from a site might take days, weeks, or even months, depending on the amount of data at the site, the load on your system, network latencies, and the nature of the required ILM changes.

- While the site decommission procedure is running:
  - You can't create ILM rules that refer to the site being decommissioned. You also can't edit an existing ILM rule to refer to the site.
  - You can't perform other maintenance procedures, such as expansion or upgrade.

    If you need to perform another maintenance procedure during a connected site decommission, you can pause the procedure while the Storage Nodes are being removed. The **Pause** button is enabled during the "Decommissioning Replicated and Erasure Coded Data" stage.

  - If you need to recover any node after starting the site decommission procedure, you must contact support.

Steps

1. Review the disconnected nodes section of Step 5 (Resolve Node Conflicts) to determine if any nodes in your StorageGRID system have a Connection State of Unknown ( ) or Administratively Down ( ).

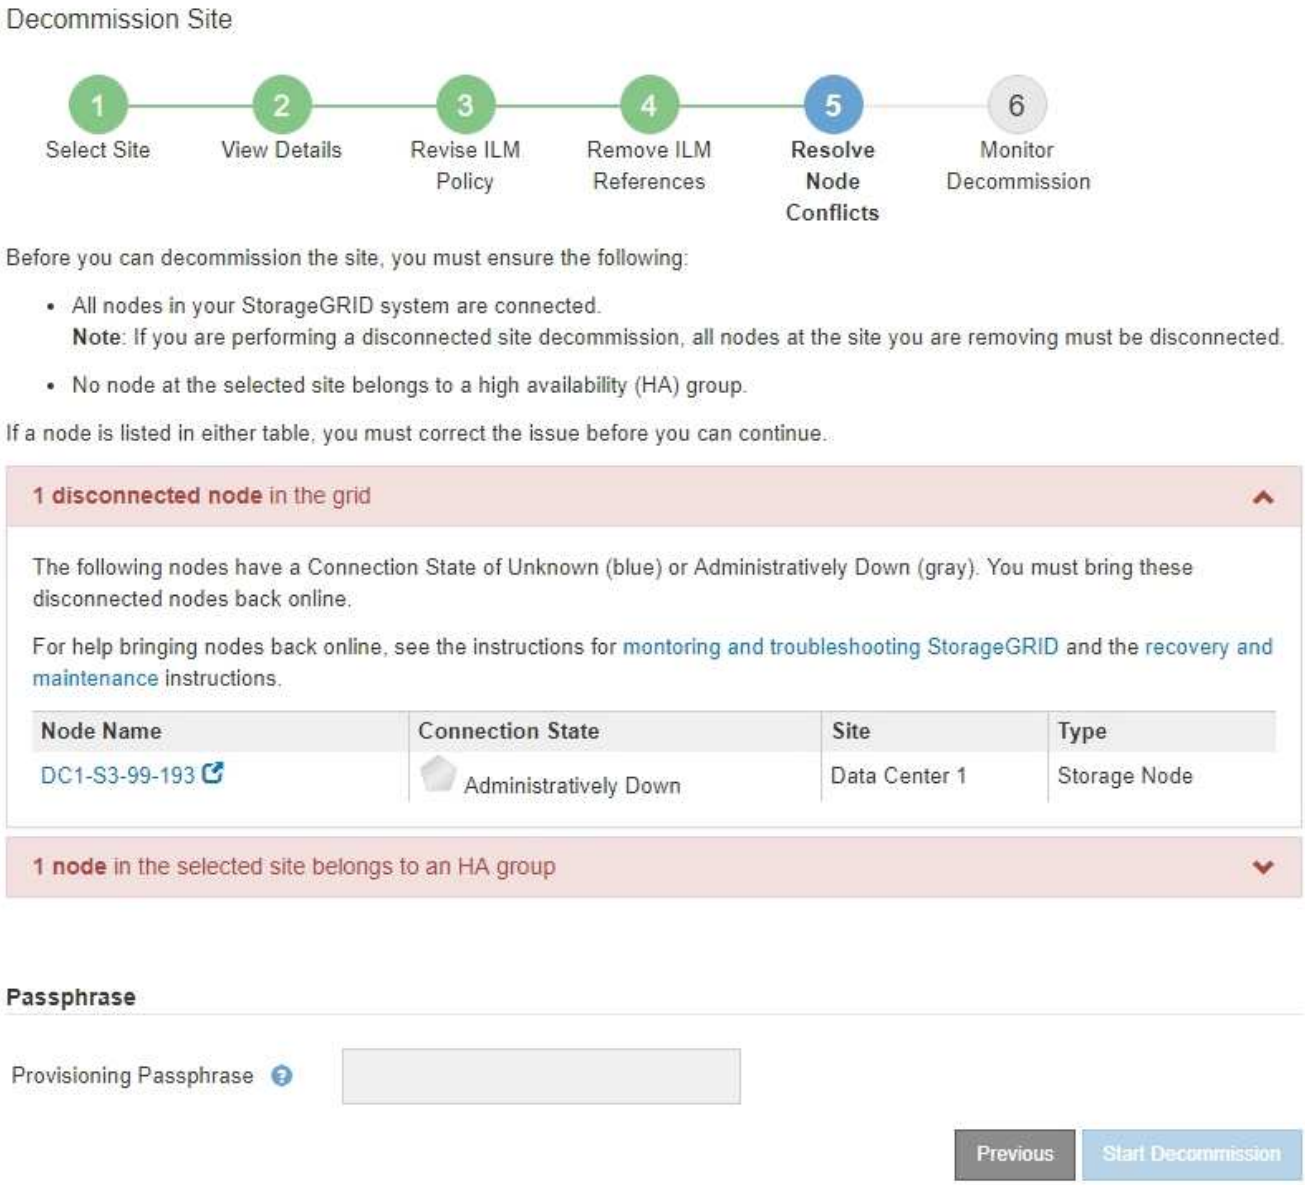

2. If any nodes are disconnected, bring them back online.

   See the Grid node procedures. Contact technical support if you need assistance.

3. When all disconnected nodes have been brought back online, review the HA groups section of Step 5 (Resolve Node Conflicts).

   This table lists any nodes at the selected site that belong to a high availability (HA) group.

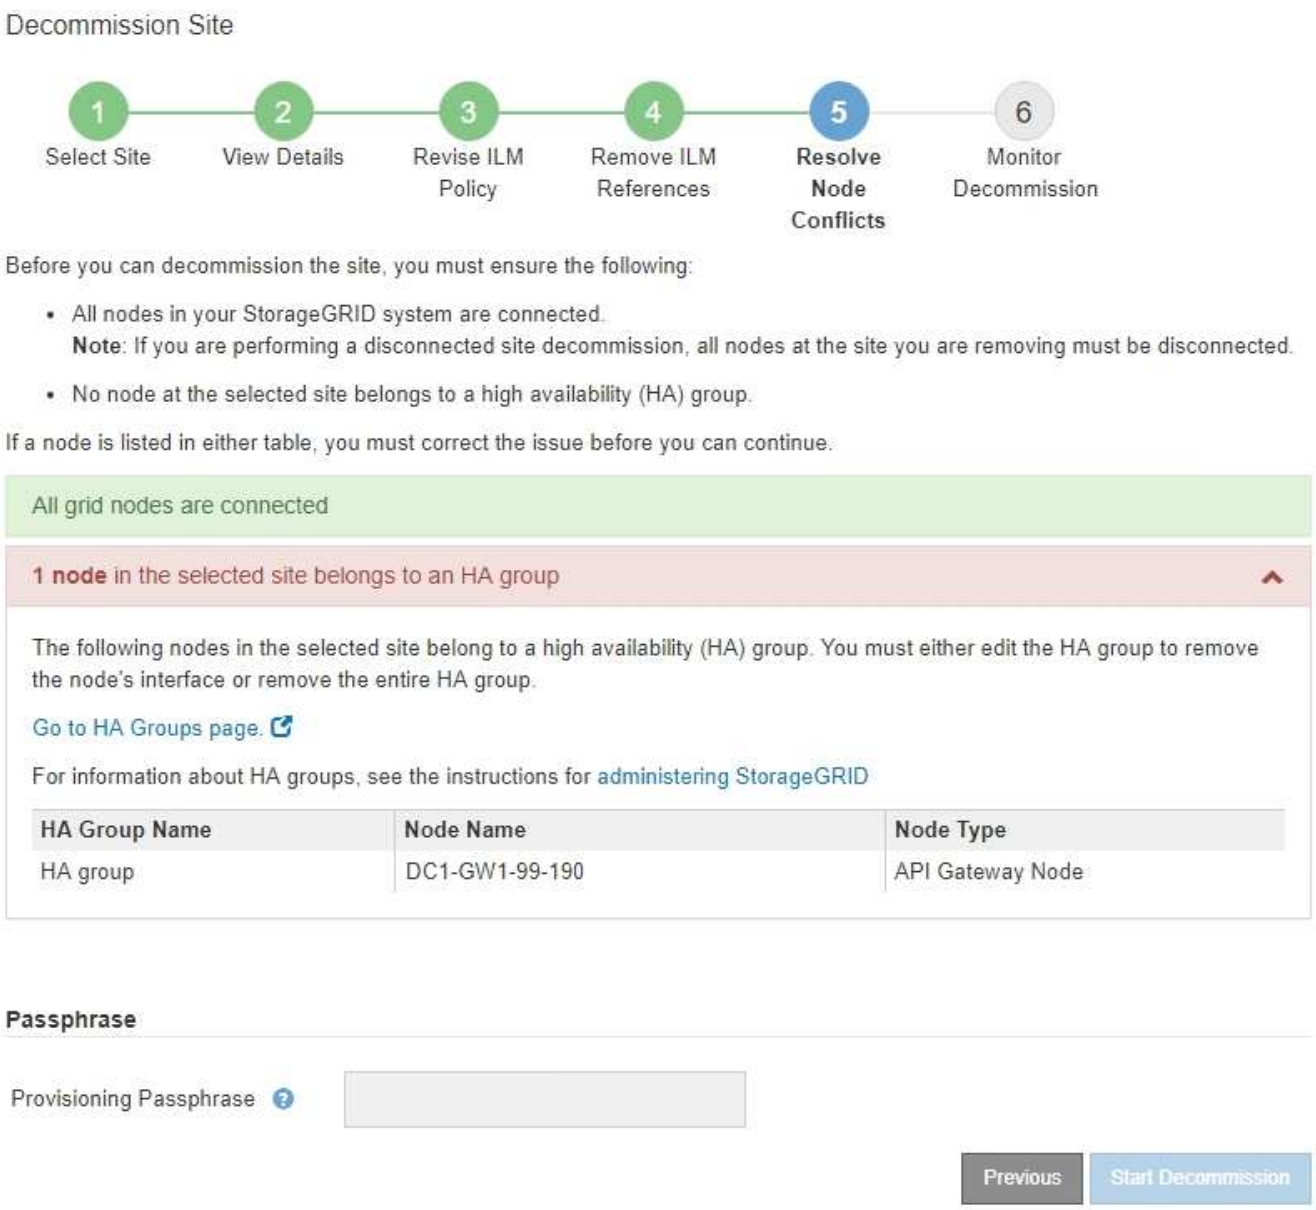

4. If any nodes are listed, do either of the following:
   - Edit each affected HA group to remove the node interface.
   - Remove an HA group that only includes nodes from this site. See the instructions for administering StorageGRID.

   If all nodes are connected and no nodes in the selected site are used in an HA group, the **Provisioning Passphrase** field is enabled.

5. Enter the provisioning passphrase.

   The **Start Decommission** button becomes enabled.

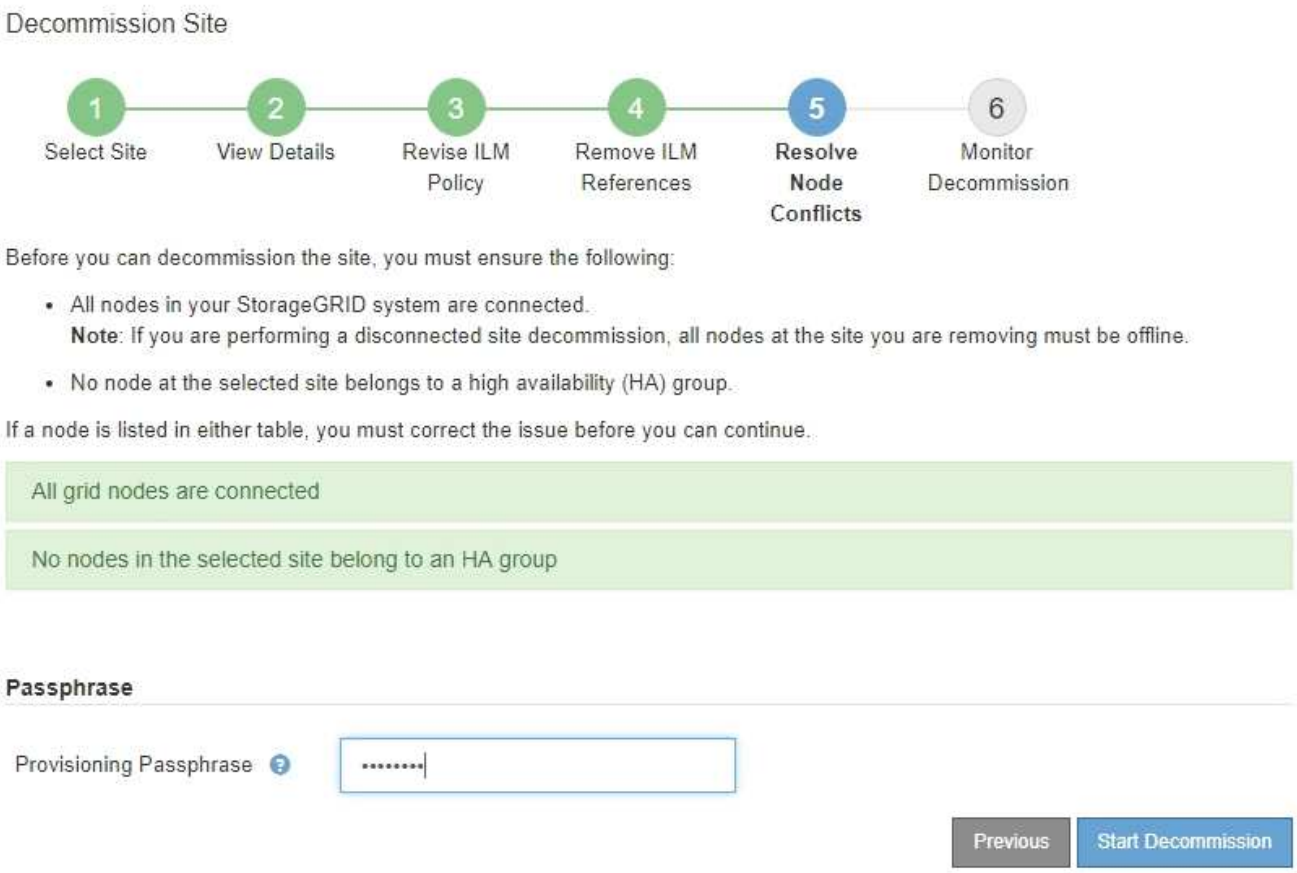

6. If you are ready to start the site decommission procedure, select **Start Decommission**.

   A warning lists the site and nodes that will be removed. You are reminded that it might take days, weeks, or even months to completely remove the site.

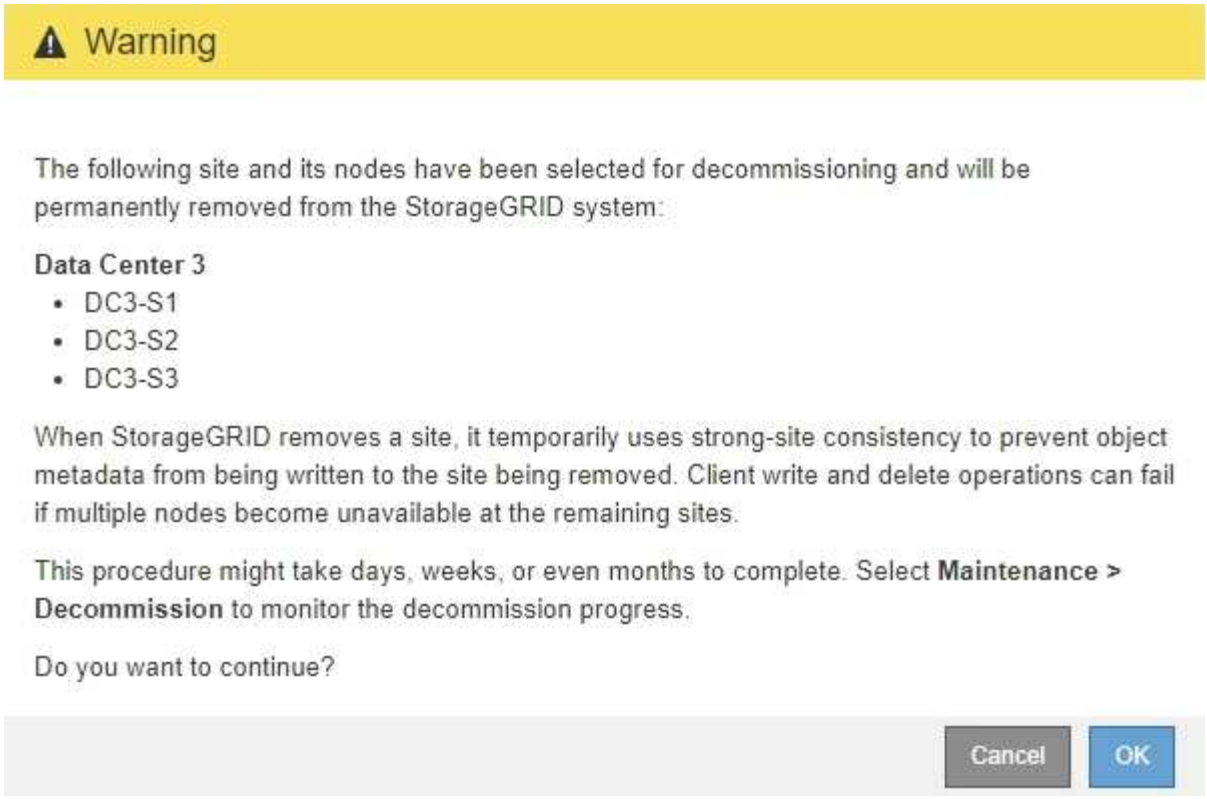

7. Review the warning. If you are ready to begin, select **OK**.

   A message appears as the new grid configuration is generated. This process might take some time, depending on the type and number of decommissioned grid nodes.

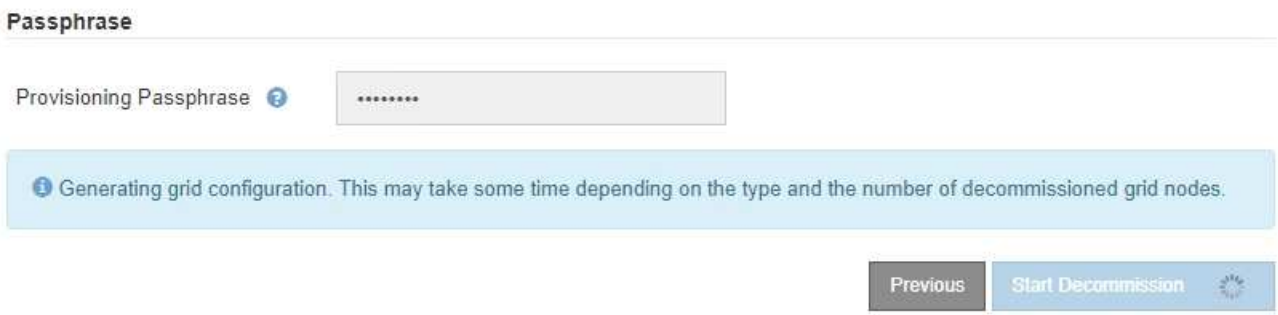

   When the new grid configuration has been generated, Step 6 (Monitor Decommission) appears.

   The **Previous** button remains disabled until the decommission is complete.

Step 6: Monitor Decommission

From Step 6 (Monitor Decommission) of the Decommission Site page wizard, you can monitor the progress as the site is removed.

About this task

When StorageGRID removes a connected site, it removes nodes in this order:

1. Gateway Nodes

2. Admin Nodes
3. Storage Nodes

When StorageGRID removes a disconnected site, it removes nodes in this order:

1. Gateway Nodes
2. Storage Nodes
3. Admin Nodes

Each Gateway Node or Admin Node might only require a few minutes or an hour to remove; however, Storage Nodes might take days or weeks.

Steps

1. As soon as a new Recovery Package has been generated, download the file.

> Download the Recovery Package as soon as possible to ensure you can recover your grid if something goes wrong during the decommission procedure.

   a. Select the link in the message, or select **MAINTENANCE** > **System** > **Recovery package**.
   b. Download the `.zip` file.

      See the instructions for downloading the Recovery Package.

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

2. Using the Data Movement chart, monitor the movement of object data from this site to other sites.

   Data movement started when you activated the new ILM policy in Step 3 (Revise ILM Policy). Data movement will occur throughout the decommission procedure.

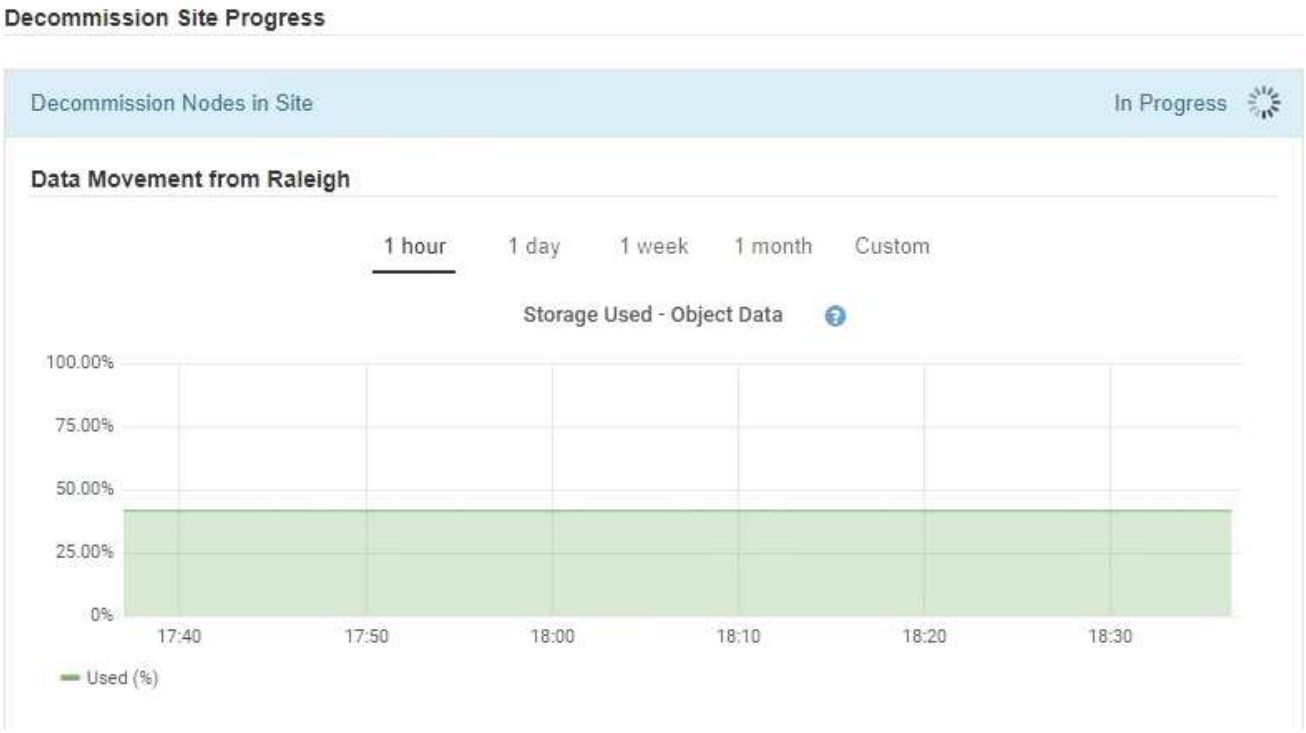

3. In the Node Progress section of the page, monitor the progress of the decommission procedure as nodes are removed.

   When a Storage Node is removed, each node goes through a series of stages. Although most of these stages occur quickly or even imperceptibly, you might need to wait days or even weeks for other stages to complete, based on how much data needs to moved. Additional time is required to manage erasure-coded data and re-evaluate ILM.

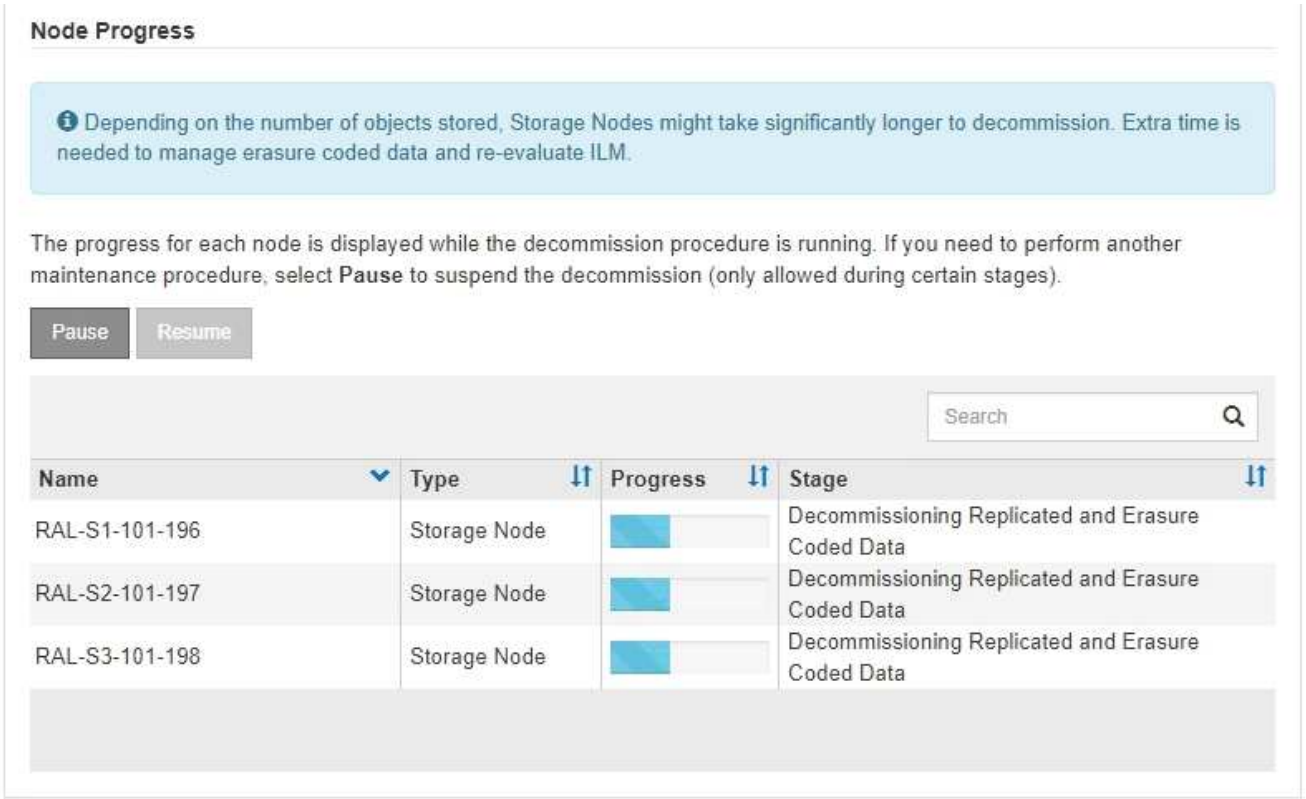

If you are monitoring the progress of a connected site decommission, refer to this table to understand the decommission stages for a Storage Node:

Stage	Estimated duration
Pending	Minute or less
Wait for Locks	Minutes
Prepare Task	Minute or less
Marking LDR Decommissioned	Minutes
Decommissioning Replicated and Erasure Coded Data	Hours, days, or weeks based on the amount of data **Note**: If you need to perform other maintenance activities, you can pause the site decommission during this stage.
LDR Set State	Minutes
Flush Audit Queues	Minutes to hours, based on the number of messages and network latency.
Complete	Minutes

If you are monitoring the progress of a disconnected site decommission, refer to this table to understand the decommission stages for a Storage Node:

Stage	Estimated duration
Pending	Minute or less
Wait for Locks	Minutes
Prepare Task	Minute or less
Disable External Services	Minutes
Certificate Revocation	Minutes
Node Unregister	Minutes
Storage Grade Unregister	Minutes
Storage Group Removal	Minutes
Entity Removal	Minutes

Stage	Estimated duration
Complete	Minutes

4. After all nodes have reached the Complete stage, wait for the remaining site decommission operations to complete.
   - During the **Repair Cassandra** step, StorageGRID makes any necessary repairs to the Cassandra clusters that remain in your grid. These repairs might take several days or more, depending on how many Storage Nodes remain in your grid.

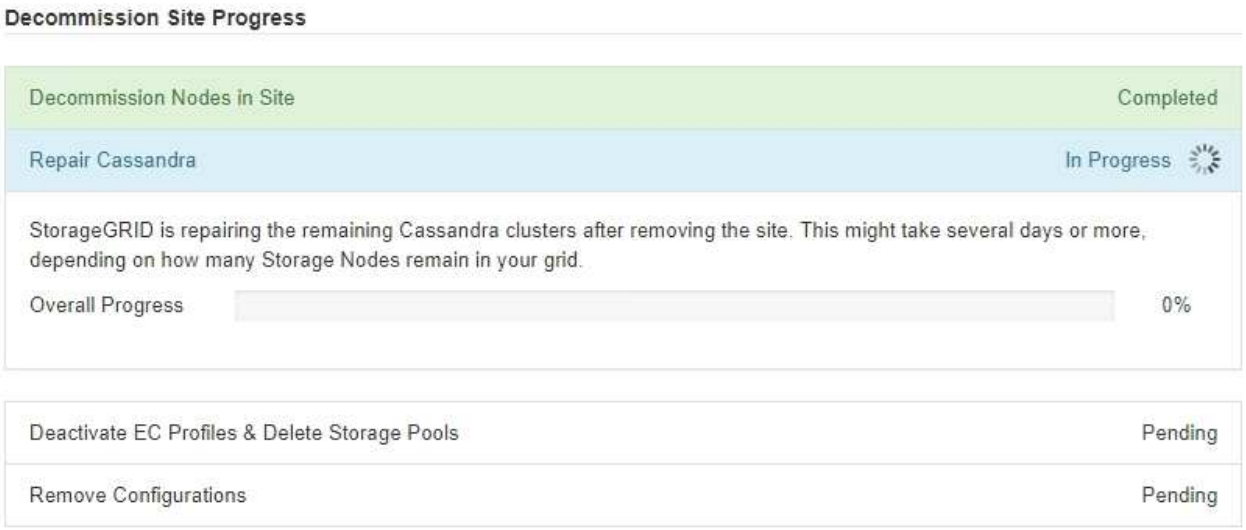

   - During the **Deactivate EC Profiles & Delete Storage Pools** step, the following ILM changes are made:
     - Any erasure coding profiles that referred to the site are deactivated.
     - Any Storage Pools that referred to the site are deleted.

> The All Storage Nodes storage pool (StorageGRID 11.6 and earlier) is also removed because it uses the All Sites site.

   - Finally, during the **Remove Configuration** step, any remaining references to the site and its nodes are removed from the rest of the grid.

5. When the decommission procedure has completed, the Decommission Site page shows a success message, and the removed site is no longer shown.

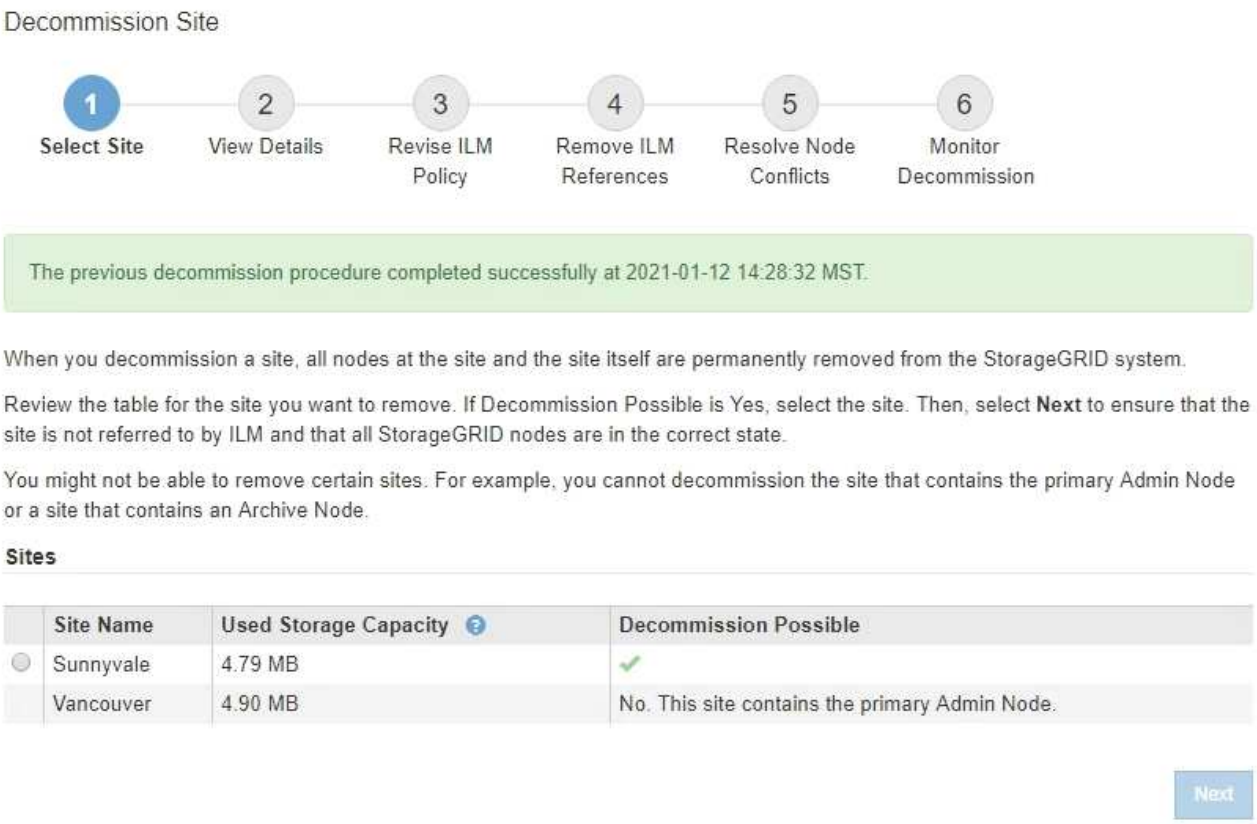

After you finish

Complete these tasks after you complete the site decommission procedure:

- Ensure that the drives of all Storage Nodes in the decommissioned site are wiped clean. Use a commercially available data wiping tool or service to permanently and securely remove data from the drives.
- If the site included one or more Admin Nodes and single sign-on (SSO) is enabled for your StorageGRID system, remove all relying party trusts for the site from Active Directory Federation Services (AD FS).
- After the nodes have been gracefully powered off automatically as part of the connected site decommission procedure, remove the associated virtual machines.

Rename grid, sites, and nodes

Rename grid, sites, and nodes: Overview

As required, you can change the display names that are shown throughout the Grid Manager for the entire grid, each site, and each node. You can update display names safely and whenever you need.

What is the rename procedure?

When you install StorageGRID initially, you specify a name for the grid, each site, and each node. These initial names are known as *system names*, and they are the names initially shown throughout StorageGRID.

System names are required for internal StorageGRID operations and can't be changed. However, you can use the rename procedure to define new *display names* for the grid, each site, and each node. These display names appear in various StorageGRID locations instead of (or in some cases, in addition to) the underlying system names.

Use the rename procedure to correct typos, to implement a different naming convention, or to indicate that a site and all of its nodes have been relocated. Unlike system names, display names can be updated whenever required and without impacting StorageGRID operations.

Where do system and display names appear?

The following table summarizes where system names and display names are shown in the StorageGRID user interface and in StorageGRID files.

Location	System name	Display name
Grid Manager pages	Shown unless the item is renamed	If an item is renamed, shown instead of the system name in these locations: • Dashboard • Nodes page • Configuration pages for high availability groups, load balancer endpoints, VLAN interfaces, key management servers, grid passwords, and firewall control • Alerts • Storage pool definitions • Object metadata lookup page • Pages related to maintenance procedures, including upgrade, hotfix, SANtricity OS upgrade, decommission, expansion, recovery, and object existence check • Support pages (logs and diagnostics) • Single sign-on page, next to the Admin Node hostname in the table for Admin Node details
NODES > **Overview** tab for a node	Always shown	Shown only if the item is renamed
Legacy pages in the Grid Manager (for example, **SUPPORT** > **Grid Topology**)	Shown	Not shown
node-health API	Always returned	Returned only if the item is renamed

Location	System name	Display name

Prompt when using SSH to access a node	Shown as the primary name unless the item has been renamed:
	
	`admin@SYSTEM-NAME: ~ $`
	
	Included in parentheses when the item is renamed:
	
	`admin@DISPLAY-NAME(SYSTEM-NAME):~ $`
		Shown as the primary name when the item is renamed:
		
		`admin@DISPLAY-NAME(SYSTEM-NAME):~ $`

`Passwords.txt` file in the Recovery Package	Shown as `Server Name`	Shown as `Display Name`
`/etc/hosts` file on all nodes For example: `10.96.99.128 SYSTEM-NAME 28989c59-a2c3-4d30-bb09-6879adf2437f DISPLAY-NAME localhost-grid # storagegrid-gen-host`	Always shown in the second column	When the item is renamed, shown in the fourth column
`topology-display-names.json`, included with AutoSupport data	Not included	Empty unless items have been renamed; otherwise, maps grid, site, and node IDs to their display names.

Display name requirements

Before using this procedure, review the requirements for display names.

Display names for nodes

Display names for nodes must follow these rules:

- Must be unique across your StorageGRID system.
- Can’t be the same as the system name for any other item in your StorageGRID system.
- Must contain at least 1 and no more than 32 characters.
- Can contain numbers, hyphens (-), and uppercase and lowercase letters.
- Can start or end with a letter or number, but can’t start or end with a hyphen.
- Can’t be all numbers.
- Are case-insensitive. For example, `DC1-ADM` and `dc1-adm` are considered to be duplicates.

You can rename a node with a display name that was previously used by a different node, as long as the

rename doesn’t result in a duplicate display name or system name.

Display names for grid and sites

Display names for the grid and sites follow the same rules with these exceptions:

- Can include spaces.
- Can include these special characters: `= - _ : , . @ !`
- Can start and end with the special characters, including hyphens.
- Can be all numbers or special characters.

Display name best practices

If you plan to rename multiple items, document your general naming scheme before using this procedure. Come up with a system that ensures that names are unique, consistent, and easy to understand at a glance.

You can use any naming convention that fits your organizational requirements. Consider these basic suggestions of what to include:

- **Site indicator**: If you have multiple sites, add a site code to each node name.
- **Node type**: Node names typically indicate the node’s type. You can use abbreviations like `s`, `adm`, `gw`, and `arc` (Storage Node, Admin Node, Gateway Node, and Archive Node).
- **Node number**: If a site contains more than one of a particular type of node, add a unique number to each node’s name.

Think twice before adding specific details to the names that are likely to change over time. For example, don’t include IP addresses in node names because these addresses can be changed. Similarly, rack locations or appliance model numbers can change if you move equipment or upgrade the hardware.

Example display names

Suppose your StorageGRID system has three data centers and has nodes of different types at each data center. Your display names might be as simple as these:

- **Grid**: `StorageGRID Deployment`
- **First site**: `Data Center 1`
  - `dc1-adm1`
  - `dc1-s1`
  - `dc1-s2`
  - `dc1-s3`
  - `dc1-gw1`
- **Second site**: `Data Center 2`
  - `dc2-adm2`
  - `dc2-s1`
  - `dc2-s2`

- dc2-s3

- **Third site**: Data Center 3
  - dc3-s1
  - dc3-s2
  - dc3-s3

Add or update display names

You can use this procedure to add or update the display names used for your grid, sites, and nodes. You can rename a single item, multiple items, or even all items at the same time. Defining or updating a display name does not affect StorageGRID operations in any way.

Before you begin

- From the **primary Admin Node**, you are signed in to the Grid Manager using a supported web browser.

> You can add or update display names from a non-primary Admin Node, but you must be signed in to the primary Admin Node to download a Recovery Package.

- You have the Maintenance or Root access permission.
- You have the provisioning passphrase.
- You understand the requirements and best practices for display names. See Rename grid, sites, and nodes: Overview.

How to rename grid, sites, or nodes

You can rename your StorageGRID system, one or more sites, or one or more nodes.

You can use a display name that was previously used by a different node, as long as the rename doesn't result in a duplicate display name or system name.

Select items to rename

To start, select the items you want to rename.

Steps

1. Select **MAINTENANCE** > **Tasks** > **Rename grid, sites, and nodes**.
2. For the **Select names** step, select the items you want to rename.

Item to change	Instruction
Names of everything (or almost everything) in your system	1. Select **Select all**. 2. Optionally clear any items you don't want to rename.
Name of the grid	Select the checkbox for the grid.

Item to change	Instruction
Name of a site and some or all of its nodes	1. Select the checkbox in the table header for the site. 2. Optionally, clear any nodes you don’t want to rename.
Name of a site	Select the checkbox for the site.
Name of a node	Select the checkbox for the node.

3. Select **Continue**.
4. Review the table, which includes the items you selected.
   - The **Display name** column shows the current name for each item. If the item has never been renamed, its display name is the same as its system name.
   - The **System name** column shows the name you entered for each item during installation. System names are used for internal StorageGRID operations and can’t be changed. For example, the system name for a node might be its hostname.
   - The **Type** column indicates the item’s type: Grid, Site, or the specific type of node.

Propose new names

For the **Propose new names** step, you can enter a display name for each item individually, or you can rename items in bulk.

Rename items individually

Follow these steps to enter a display name for each item you want to rename.

Steps

1. In the **Display name** field, enter a proposed display name for each item in the list.

   See Rename grid, sites, and nodes: Overview to learn the naming requirements.

2. To remove any items you don't want to rename, select ✕ in the **Remove from list** column.

   If you will not be proposing a new name for an item, you must remove it from the table.

3. When you have proposed new names for all items in the table, select **Rename**.

   A success message appears. The new display names are now used throughout Grid Manager.

Rename items in bulk

Use the bulk rename tool if item names share a common string that you want to replace with a different string.

Steps

1. For the **Propose new names** step, select **Use bulk rename tool**.

   The **Rename preview** includes all items that were shown for the **Propose new names** step. You can use the preview to see how display names will look after you replace a shared string.

2. In the **Existing string** field, enter the shared string you want to replace. For example, if the string you want to replace is `Data-Center-1`, enter **Data-Center-1**.

   As you type, your text is highlighted wherever it is found in the names on the left.

3. Select ✕ to remove any items that you don't want to rename with this tool.

   For example, suppose you want to rename all nodes that contain the string `Data-Center-1`, but you don't want to rename the `Data-Center-1` site itself. Select ✕ to remove the site from the rename preview.

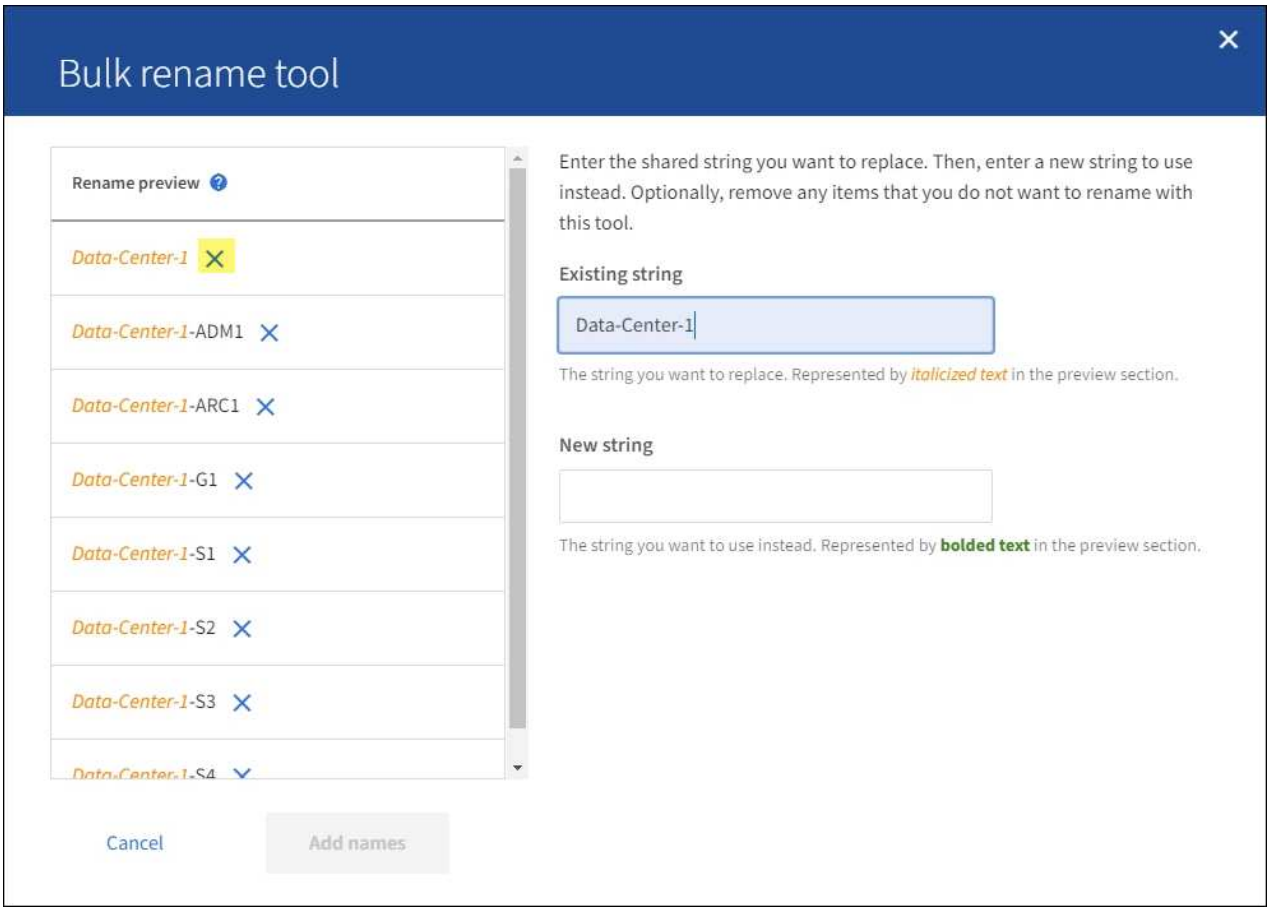

4. In the **New string** field, enter the replacement string you want to use instead. For example, enter **DC1**.

   See Rename grid, sites, and nodes: Overview to learn the naming requirements.

   As you enter the replacement string, the names on the left are updated, so you can verify that the new names will be correct.

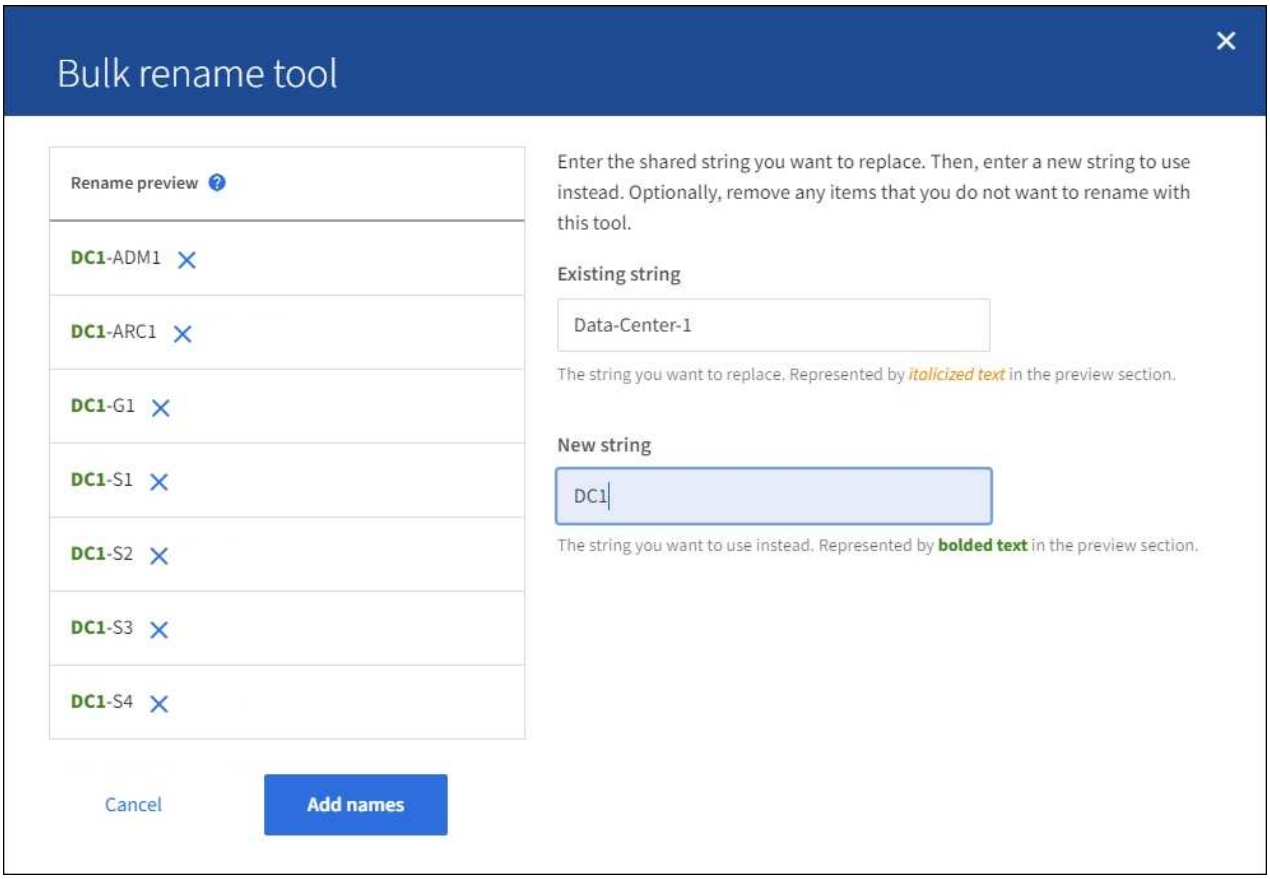

5. When you are satisfied with the names shown in the preview, select **Add names** to add the names to the table for the **Propose new names** step.
6. Make any additional changes required, or select ✕ to remove any items that you don't want to rename.
7. When you are ready to rename all items in the table, select **Rename**.

   A success message is shown. The new display names are now used throughout Grid Manager.

Download the recovery package

When you are done renaming items, download and save a new Recovery Package. The new display names for the items you renamed are included in the `Passwords.txt` file.

Steps

1. Enter the provisioning passphrase.
2. Select **Download Recovery Package**.

   The download starts immediately.

3. When the download completes, open the `Passwords.txt` file to see the server name for all nodes and the display names for any renamed nodes.
4. Copy the `sgws-recovery-package-id-revision.zip` file to two safe, secure, and separate locations.

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

5. Select **Finish** to return to the first step.

Revert display names back to system names

You can revert a renamed grid, site, or node back to its original system name. When you revert an item back to its system name, Grid Manager pages and other StorageGRID locations no longer show a **Display name** for that item. Only the item's system name is shown.

Steps

1. Select **MAINTENANCE** > **Tasks** > **Rename grid, sites, and nodes**.
2. For the **Select names** step, select any items you want to revert back to system names.
3. Select **Continue**.
4. For the **Propose new names** step, revert display names back to system names individually or in bulk.

Revert to system names individually

a. Copy each item's original system name and paste it into the **Display name** field, or select ✕ to remove any items you don't want to revert.

   To revert a display name, the system name must appear in the **Display name** field, but the name is case insensitive.

b. Select **Rename**.

   A success message appears. The display names for these items are no longer used.

Revert to system names in bulk

a. For the **Propose new names** step, select **Use bulk rename tool**.
b. In the **Existing string** field, enter the display name string you want to replace.
c. In the **New string** field, enter the system name string you want to use instead.
d. Select **Add names** to add the names to the table for the **Propose new names** step.
e. Confirm that each entry in the **Display name** field matches the name in the **System name** field. Make any changes or select ✕ to remove any items that you don't want to revert.

   To revert a display name, the system name must appear in the **Display name** field, but the name is case insensitive.

f. Select **Rename**.

   A success message is shown. The display names for these items are no longer used.

5. Download and save a new Recovery Package.

   Display names for the items you reverted are no longer included in the `Passwords.txt` file.

Network maintenance procedures

Update subnets for Grid Network

StorageGRID maintains a list of the network subnets used to communicate between grid nodes on the Grid Network (eth0). These entries include the subnets used for the Grid Network by each site in your StorageGRID system as well as any subnets used for NTP, DNS, LDAP, or other external servers accessed through the Grid Network gateway. When you add grid nodes or a new site in an expansion, you might need to update or add subnets to the Grid Network.

Before you begin

- You are signed in to the Grid Manager using a supported web browser.
- You have the Maintenance or Root access permission.
- You have the provisioning passphrase.
- You have the network addresses, in CIDR notation, of the subnets you want to configure.

About this task

If you are performing an expansion activity that includes adding a new subnet, you must add a new subnet to the Grid Network subnet list before you start the expansion procedure. Otherwise, you will have to cancel the expansion, add the new subnet, and start the expansion again.

Add a subnet

Steps

1. Select **MAINTENANCE** > **Network** > **Grid Network**.
2. Select **Add another subnet** to add a new subnet in CIDR notation.

   For example, enter `10.96.104.0/22`.

3. Enter the provisioning passphrase, and select **Save**.
4. Wait until the changes are applied, then download a new Recovery Package.
   a. Select **MAINTENANCE** > **System** > **Recovery package**.
   b. Enter the **Provisioning Passphrase**.

> The Recovery Package file must be secured because it contains encryption keys and passwords that can be used to obtain data from the StorageGRID system. It is also used to recover the primary Admin Node.

The subnets you have specified are configured automatically for your StorageGRID system.

Edit a subnet

Steps

1. Select **MAINTENANCE** > **Network** > **Grid Network**.
2. Select the subnet you want to edit and make the necessary changes.
3. Enter the Provisioning passphrase, and select **Save**.

4. Select **Yes** in the confirmation dialog box.
5. Wait until the changes are applied, then download a new Recovery Package.
    a. Select **MAINTENANCE** > **System** > **Recovery package**.
    b. Enter the **Provisioning Passphrase**.

Delete a subnet

Steps

1. Select **MAINTENANCE** > **Network** > **Grid Network**.
2. Select the delete icon ✕ next to the subnet.
3. Enter the Provisioning passphrase, and select **Save**.
4. Select **Yes** in the confirmation dialog box.
5. Wait until the changes are applied, then download a new Recovery Package.
    a. Select **MAINTENANCE** > **System** > **Recovery package**.
    b. Enter the **Provisioning Passphrase**.

Configure IP addresses

Configure IP addresses

You can perform network configuration by configuring IP addresses for grid nodes using the Change IP tool.

You must use the Change IP tool to make most changes to the networking configuration that was initially set during grid deployment. Manual changes using standard Linux networking commands and files might not propagate to all StorageGRID services, and might not persist across upgrades, reboots, or node recovery procedures.

The IP change procedure can be a disruptive procedure. Parts of the grid might be unavailable until the new configuration is applied.

If you are making changes to the Grid Network Subnet List only, use the Grid Manager to add or change the network configuration. Otherwise, use the Change IP tool if the Grid Manager is inaccessible due to a network configuration issue, or you are performing both a Grid Network routing change and other network changes at the same time.

If you want to change the Grid Network IP address for all nodes in the grid, use the special procedure for grid-wide changes.

Ethernet interfaces

The IP address assigned to eth0 is always the grid node's Grid Network IP address. The IP address assigned to eth1 is always the grid node's Admin Network IP address. The IP address assigned to eth2 is always the grid node's Client Network IP address.

Note that on some platforms, such as StorageGRID appliances, eth0, eth1, and eth2 might be aggregate interfaces composed of subordinate bridges or bonds of physical or VLAN interfaces. On these platforms, the **SSM** > **Resources** tab might show the Grid, Admin, and Client Network IP address assigned to other

interfaces in addition to eth0, eth1, or eth2.

DHCP

You can only set up DHCP during the deployment phase. You can't set up DHCP during configuration. You must use the IP address change procedures if you want to change IP addresses, subnet masks, and default gateways for a grid node. Using the Change IP tool will cause DHCP addresses to become static.

High availability (HA) groups

- If a Client Network interface is contained in an HA group, you can't change the Client Network IP address for that interface to an address that is outside of the subnet configured for the HA group.
- You can't change the Client Network IP address to the value of an existing virtual IP address assigned to an HA group configured on the Client Network interface.
- If a Grid network interface is contained in an HA group, you can't change the Grid network IP address for that interface to an address that is outside of the subnet configured for the HA group.
- You can't change the Grid Network IP address to the value of an existing virtual IP address assigned to an HA group configured on the Grid Network interface.

Change node network configuration

You can change the network configuration of one or more nodes using the Change IP tool. You can change the configuration of the Grid Network, or add, change, or remove the Admin or Client Networks.

Before you begin

You have the `Passwords.txt` file.

About this task

Linux: If you are adding a grid node to the Admin Network or Client Network for the first time, and you did not previously configure ADMIN_NETWORK_TARGET or CLIENT_NETWORK_TARGET in the node configuration file, you must do so now.

See the StorageGRID installation instructions for your Linux operating system:

- Install Red Hat Enterprise Linux or CentOS
- Install Ubuntu or Debian

Appliances: On StorageGRID appliances, if the Client or Admin Network was not configured in the StorageGRID Appliance Installer during the initial installation, the network can't be added by using only the Change IP tool. First, you must place the appliance in maintenance mode, configure the links, return the appliance to normal operating mode, and then use the Change IP tool to modify the network configuration. See the procedure for configuring network links.

You can change the IP address, subnet mask, gateway, or MTU value for one or more nodes on any network.

You can also add or remove a node from a Client Network or from an Admin Network:

- You can add a node to a Client Network or to an Admin Network by adding an IP address/subnet mask on that network to the node.
- You can remove a node from a Client Network or from an Admin Network by deleting the IP

address/subnet mask for the node on that network.

Nodes can't be removed from the Grid Network.

IP address swaps aren't allowed. If you must exchange IP addresses between grid nodes, you must use a temporary intermediate IP address.

If single sign-on (SSO) is enabled for your StorageGRID system and you are changing the IP address of an Admin Node, be aware that any relying party trust that was configured using the Admin Node's IP address (instead of its fully qualified domain name, as recommended) will become invalid. You will no longer be able to sign in to the node. Immediately after changing the IP address, you must update or reconfigure the node's relying party trust in Active Directory Federation Services (AD FS) with the new IP address. See the instructions for configuring SSO.

Any changes you make to the network using the Change IP tool are propagated to the installer firmware for the StorageGRID appliances. That way, if StorageGRID software is reinstalled on an appliance, or if an appliance is placed into maintenance mode, the networking configuration will be correct.

Steps

1. Log in to the primary Admin Node:
   a. Enter the following command: `ssh admin@primary_Admin_Node_IP`
   b. Enter the password listed in the `Passwords.txt` file.
   c. Enter the following command to switch to root: `su -`
   d. Enter the password listed in the `Passwords.txt` file.

      When you are logged in as root, the prompt changes from `$` to `#`.

2. Start the Change IP tool by entering the following command: `change-ip`
3. Enter the provisioning passphrase at the prompt.

   The main menu appears.

```
Welcome to the StorageGRID IP Change Tool.

Selected nodes: all

1:   SELECT NODES to edit
2:   EDIT IP/mask, gateway and MTU
3:   EDIT admin network subnet lists
4:   EDIT grid network subnet list
5:   SHOW changes
6:   SHOW full configuration, with changes highlighted
7:   VALIDATE changes
8:   SAVE changes, so you can resume later
9:   CLEAR all changes, to start fresh
10:  APPLY changes to the grid
0:   Exit

Selection:
```

4. Optionally select **1** to choose which nodes to update. Then select one of the following options:

- **1**: Single node — select by name
- **2**: Single node — select by site, then by name
- **3**: Single node — select by current IP
- **4**: All nodes at a site
- **5**: All nodes in the grid

  Note: If you want to update all nodes, allow "all" to remain selected.

After you make your selection, the main menu appears, with the **Selected nodes** field updated to reflect your choice. All subsequent actions are performed only on the nodes displayed.

5. On the main menu, select option **2** to edit IP/mask, gateway, and MTU information for the selected nodes.

   a. Select the network where you want to make changes:

      - **1**: Grid network
      - **2**: Admin network
      - **3**: Client network
      - **4**: All networks After you make your selection, the prompt shows the node name, network name (Grid, Admin, or Client), data type (IP/mask, Gateway, or MTU), and current value.

      Editing the IP address, prefix length, gateway, or MTU of a DHCP-configured interface will change the interface to static. When you select to change an interface configured by DHCP, a warning is displayed to inform you that the interface will change to static.

   Interfaces configured as `fixed` can’t be edited.

   a. To set a new value, enter it in the format shown for the current value.
   b. To leave the current value unchanged, press **Enter**.
   c. If the data type is `IP/mask`, you can delete the Admin or Client Network from the node by entering **d** or **0.0.0.0/0**.
   d. After editing all nodes you want to change, enter **q** to return to the main menu.

      Your changes are held until cleared or applied.

6. Review your changes by selecting one of the following options:
   - **5**: Shows edits in output that is isolated to show only the changed item. Changes are highlighted in green (additions) or red (deletions), as shown in the example output:

```
================================================================================
Site: RTP
================================================================================
username-x Grid    IP      [    172.16.0.239/21 ]: 172.16.0.240/21
username-x Grid    MTU     [              1400 ]: 9000
username-x Grid    MTU     [              1400 ]: 9000
username-x Grid    MTU     [              1400 ]: 9000
username-x Grid    MTU     [              1400 ]: 9000
username-x Grid    MTU     [              1400 ]: 9000
username-x Grid    MTU     [              1400 ]: 9000
username-x Grid    MTU     [              1400 ]: 9000
username-x Admin   IP      [    10.224.0.244/21 ]: 0.0.0.0/0
username-x Admin   IP      [    10.224.0.245/21 ]: 0.0.0.0/0
username-x Admin   IP      [    10.224.0.240/21 ]: 0.0.0.0/0
username-x Admin   IP      [    10.224.0.241/21 ]: 0.0.0.0/0
username-x Admin   IP      [    10.224.0.242/21 ]: 0.0.0.0/0
username-x Admin   IP      [    10.224.0.243/21 ]: 0.0.0.0/0
username-x Admin   Gateway [         10.224.0.1 ]: 0.0.0.0
username-x Admin   Gateway [         10.224.0.1 ]: 0.0.0.0
username-x Admin   Gateway [         10.224.0.1 ]: 0.0.0.0
username-x Admin   Gateway [         10.224.0.1 ]: 0.0.0.0
username-x Admin   Gateway [         10.224.0.1 ]: 0.0.0.0
username-x Admin   Gateway [         10.224.0.1 ]: 0.0.0.0
username-x Admin   MTU     [              1400 ]: 0
username-x Admin   MTU     [              1400 ]: 0
username-x Admin   MTU     [              1400 ]: 0
username-x Admin   MTU     [              1400 ]: 0
username-x Admin   MTU     [              1400 ]: 0
username-x Admin   MTU     [              1400 ]: 0
Press Enter to continue
```

- **6**: Shows edits in output that displays the full configuration. Changes are highlighted in green (additions) or red (deletions).

> Certain command line interfaces might show additions and deletions using strikethrough formatting. Proper display depends on your terminal client supporting the necessary VT100 escape sequences.

7. Select option **7** to validate all changes.

   This validation ensures that the rules for the Grid, Admin, and Client Networks, such as not using overlapping subnets, aren't violated.

   In this example, validation returned errors.

```
Validating new networking configuration... FAILED.

  DK-10-224-5-20-G1: The admin subnet 172.18.0.0/16 overlaps the 172.18.0.0/21 grid network.
  DK-10-224-5-22-S1: Duplicate Grid IP 172.16.5.18 (also in use by DK-10-224-5-21-ADM1)

You must correct these errors before you can apply any changes.
Checking for Grid Network IP address swaps... PASSED.

Press Enter to continue
```

   In this example, validation passed.

```
Validating new networking configuration... PASSED.
Checking for Grid Network IP address swaps... PASSED.

Press Enter to continue
```

8. Once validation passes, choose one of the following options:
   - **8**: Save unapplied changes.

     This option allows you to quit the Change IP tool and start it again later, without losing any unapplied changes.

   - **10**: Apply the new network configuration.
9. If you selected option **10**, choose one of the following options:
   - **apply**: Apply the changes immediately and automatically restart each node if necessary.

     If the new network configuration does not require any physical networking changes, you can select **apply** to apply the changes immediately. Nodes will be restarted automatically, if necessary. Nodes that need to be restarted will be displayed.

   - **stage**: Apply the changes the next time the nodes are restarted manually.

     If you need to make physical or virtual networking configuration changes for the new network configuration to function, you must use the **stage** option, shut down the affected nodes, make the necessary physical networking changes, and restart the affected nodes. If you select **apply** without first making these networking changes, the changes will usually fail.

     > If you use the **stage** option, you must restart the node as soon as possible after staging to minimize disruptions.

   - **cancel**: Don't make any network changes at this time.

     If you were unaware that the proposed changes require nodes to be restarted, you can defer the changes to minimize user impact. Selecting **cancel** returns you to the main menu and preserves your changes so you can apply them later.

     When you select **apply** or **stage**, a new network configuration file is generated, provisioning is performed, and nodes are updated with new working information.

     During provisioning, the output displays the status as updates are applied.

     ```
     Generating new grid networking description file...

     Running provisioning...

     Updating grid network configuration on Name
     ```

   After applying or staging changes, a new Recovery Package is generated as a result of the grid configuration change.

10. If you selected **stage**, follow these steps after provisioning is complete:
    a. Make the physical or virtual networking changes that are required.

       Physical networking changes: Make the necessary physical networking changes, safely shutting down the node if necessary.

Linux: If you are adding the node to an Admin Network or Client Network for the first time, ensure that you have added the interface as described in Linux: Add interfaces to existing node.

b. Restart the affected nodes.

11. Select **0** to exit the Change IP tool after your changes are complete.
12. Download a new Recovery Package from the Grid Manager.
    a. Select **MAINTENANCE** > **System** > **Recovery package**.
    b. Enter the provisioning passphrase.

Related information

SGF6112 storage appliances

SG6000 storage appliances

SG5700 storage appliances

SG100 and SG1000 services appliances

Add to or change subnet lists on Admin Network

You can add, delete, or change the subnets in the Admin Network Subnet List of one or more nodes.

Before you begin

- You have the `Passwords.txt` file.

You can add, delete, or change subnets to all nodes on the Admin Network Subnet List.

Steps

1. Log in to the primary Admin Node:
   a. Enter the following command: `ssh admin@`*`primary_Admin_Node_IP`*
   b. Enter the password listed in the `Passwords.txt` file.
   c. Enter the following command to switch to root: `su -`
   d. Enter the password listed in the `Passwords.txt` file.

   When you are logged in as root, the prompt changes from `$` to `#`.
2. Start the Change IP tool by entering the following command: `change-ip`
3. Enter the provisioning passphrase at the prompt.

   The main menu appears.

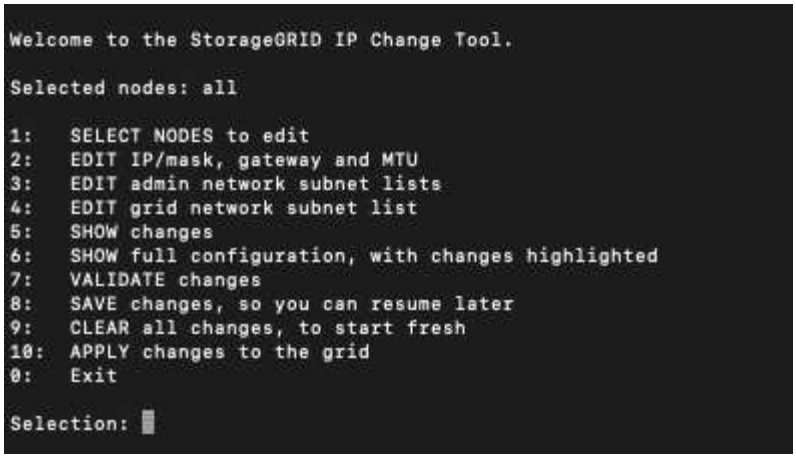

4. Optionally, limit the networks/nodes on which operations are performed. Choose one of the following:
   - Select the nodes to edit by choosing **1**, if you want to filter on specific nodes on which to perform the operation. Select one of the following options:
     - **1**: Single node (select by name)
     - **2**: Single node (select by site, then by name)
     - **3**: Single node (select by current IP)
     - **4**: All nodes at a site
     - **5**: All nodes in the grid
     - **0**: Go back
   - Allow “all” to remain selected. After the selection is made, the main menu screen appears. The Selected nodes field reflects your new selection, and now all operations selected will only be performed on this item.
5. On the main menu, select the option to edit subnets for the Admin Network (option **3**).
6. Choose one of the following:
   - Add a subnet by entering this command: `add CIDR`
   - Delete a subnet by entering this command: `del CIDR`
   - Set the list of subnets by entering this command: `set CIDR`

     For all commands, you can enter multiple addresses using this format: `add CIDR, CIDR`

     Example: `add 172.14.0.0/16, 172.15.0.0/16, 172.16.0.0/16`

     You can reduce the amount of typing required by using “up arrow” to recall previously typed values to the current input prompt, and then edit them if necessary.

     The example input below shows adding subnets to the Admin Network Subnet List:

```
Editing: Admin Network Subnet List for node DK-10-224-5-20-G1

Press <enter> to use the list as shown
Use up arrow to recall a previously typed value, which you can then edit
Use 'add <CIDR> [, <CIDR>]' to add subnets <CIDR> [, <CIDR>] to the list
Use 'del <CIDR> [, <CIDR>]' to delete subnets <CIDR> [, <CIDR>] from the list
Use 'set <CIDR> [, <CIDR>]' to set the list to the given list
Use q to complete the editing session early and return to the previous menu

DK-10-224-5-20-G1
  10.0.0.0/8
  172.19.0.0/16
  172.21.0.0/16
  172.20.0.0/16

[add/del/set/quit <CIDR>, ...]: add 172.14.0.0/16, 172.15.0.0/16
```

7. When ready, enter **q** to go back to the main menu screen. Your changes are held until cleared or applied.

> If you selected any of the "all" node selection modes in step 2, press **Enter** (without **q**) to get to the next node in the list.

8. Choose one of the following:

   - Select option **5** to show edits in output that is isolated to show only the changed item. Changes are highlighted in green (additions) or red (deletions), as shown in the example output below:

```
================================================================================
Site: Data Center 1
================================================================================
DC1-ADM1-105-154 Admin  Subnets                        add 172.17.0.0/16
                                                       del 172.16.0.0/16
                                [       172.14.0.0/16 ]
                                [       172.15.0.0/16 ]
                                [       172.17.0.0/16 ]
                                [       172.19.0.0/16 ]
                                [       172.20.0.0/16 ]
                                [       172.21.0.0/16 ]
Press Enter to continue
```

   - Select option **6** to show edits in output that displays the full configuration. Changes are highlighted in green (additions) or red (deletions). **Note:** Certain terminal emulators might show additions and deletions using strikethrough formatting.

     When you attempt to change the subnet list, the following message is displayed:

```
CAUTION: The Admin Network subnet list on the node might contain /32
subnets derived from automatically applied routes that aren't
persistent. Host routes (/32 subnets) are applied automatically if
the IP addresses provided for external services such as NTP or DNS
aren't reachable using default StorageGRID routing, but are reachable
using a different interface and gateway. Making and applying changes
to the subnet list will make all automatically applied subnets
persistent. If you don't want that to happen, delete the unwanted
subnets before applying changes. If you know that all /32 subnets in
the list were added intentionally, you can ignore this caution.
```

If you did not specifically assign the NTP and DNS server subnets to a network, StorageGRID creates a host route (/32) for the connection automatically. If, for example, you would rather have a /16 or /24 route for outbound connection to a DNS or NTP server, you should delete the automatically created /32 route and add the routes you want. If you don’t delete the automatically created host route, it will be persisted after you apply any changes to the subnet list.

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

9. Select option **7** to validate all staged changes.

   This validation ensures that the rules for the Grid, Admin, and Client Networks are followed, such as using overlapping subnets.

10. Optionally, select option **8** to save all staged changes and return later to continue making changes.

    This option allows you to quit the Change IP tool and start it again later, without losing any unapplied changes.

11. Do one of the following:
    - Select option **9** if you want to clear all changes without saving or applying the new network configuration.
    - Select option **10** if you are ready to apply changes and provision the new network configuration. During provisioning, the output displays the status as updates are applied as shown in the following sample output:

```
Generating new grid networking description file...

Running provisioning...

Updating grid network configuration on Name
```

12. Download a new Recovery Package from the Grid Manager.
    a. Select **MAINTENANCE** > **System** > **Recovery package**.
    b. Enter the provisioning passphrase.

Add to or change subnet lists on Grid Network

You can use the Change IP tool to add or change subnets on the Grid Network.

Before you begin

- You have the `Passwords.txt` file.

You can add, delete, or change subnets in the Grid Network Subnet List. Changes will affect routing on all nodes in the grid.

> If you are making changes to the Grid Network Subnet List only, use the Grid Manager to add or change the network configuration. Otherwise, use the Change IP tool if the Grid Manager is inaccessible due to a network configuration issue, or you are performing both a Grid Network routing change and other network changes at the same time.

Steps

1. Log in to the primary Admin Node:
   a. Enter the following command: `ssh admin@`*`primary_Admin_Node_IP`*
   b. Enter the password listed in the `Passwords.txt` file.
   c. Enter the following command to switch to root: `su -`
   d. Enter the password listed in the `Passwords.txt` file.

   When you are logged in as root, the prompt changes from `$` to `#`.
2. Start the Change IP tool by entering the following command: `change-ip`
3. Enter the provisioning passphrase at the prompt.

   The main menu appears.

```
Welcome to the StorageGRID IP Change Tool.

Selected nodes: all

1:   SELECT NODES to edit
2:   EDIT IP/mask, gateway and MTU
3:   EDIT admin network subnet lists
4:   EDIT grid network subnet list
5:   SHOW changes
6:   SHOW full configuration, with changes highlighted
7:   VALIDATE changes
8:   SAVE changes, so you can resume later
9:   CLEAR all changes, to start fresh
10:  APPLY changes to the grid
0:   Exit

Selection:
```

4. On the main menu, select the option to edit subnets for the Grid Network (option **4**).

   > Changes to the Grid Network Subnet List are grid-wide.

5. Choose one of the following:

   - Add a subnet by entering this command: `add CIDR`

- Delete a subnet by entering this command: `del CIDR`
- Set the list of subnets by entering this command: `set CIDR`

> For all commands, you can enter multiple addresses using this format: `add CIDR, CIDR`

Example: `add 172.14.0.0/16, 172.15.0.0/16, 172.16.0.0/16`

> You can reduce the amount of typing required by using “up arrow” to recall previously typed values to the current input prompt, and then edit them if necessary.

The example input below shows setting subnets for the Grid Network Subnet List:

```
Editing: Grid Network Subnet List

Press <enter> to use the list as shown
Use up arrow to recall a previously typed value, which you can then edit
Use 'add <CIDR> [, <CIDR>]' to add subnets <CIDR> [, <CIDR>] to the list
Use 'del <CIDR> [, <CIDR>]' to delete subnets <CIDR> [, <CIDR>] from the list
Use 'set <CIDR> [, <CIDR>]' to set the list to the given list
Use q to complete the editing session early and return to the previous menu

Grid Network Subnet List
  172.16.0.0/21
  172.17.0.0/21
  172.18.0.0/21
  192.168.0.0/21

[add/del/set/quit <CIDR>, ...]: set 172.30.0.0/21, 172.31.0.0/21, 192.168.0.0/21
```

6. When ready, enter **q** to go back to the main menu screen. Your changes are held until cleared or applied.
7. Choose one of the following:
   - Select option **5** to show edits in output that is isolated to show only the changed item. Changes are highlighted in green (additions) or red (deletions), as shown in the example output below:

```
================================================================================
Grid Network Subnet List (GNSL)
================================================================================
                                                    add 172.30.0.0/21
                                                    add 172.31.0.0/21
                                                    del 172.16.0.0/21
                                                    del 172.17.0.0/21
                                                    del 172.18.0.0/21
                              [      172.30.0.0/21 ]
                              [      172.31.0.0/21 ]
                              [    192.168.0.0/21 ]
Press Enter to continue
```

   - Select option **6** to show edits in output that displays the full configuration. Changes are highlighted in green (additions) or red (deletions).

> Certain command line interfaces might show additions and deletions using strikethrough formatting.

8. Select option **7** to validate all staged changes.

   This validation ensures that the rules for the Grid, Admin, and Client Networks are followed, such as using overlapping subnets.

9. Optionally, select option **8** to save all staged changes and return later to continue making changes.

   This option allows you to quit the Change IP tool and start it again later, without losing any unapplied changes.

10. Do one of the following:
    - Select option **9** if you want to clear all changes without saving or applying the new network configuration.
    - Select option **10** if you are ready to apply changes and provision the new network configuration. During provisioning, the output displays the status as updates are applied as shown in the following sample output:

      ```
      Generating new grid networking description file...

      Running provisioning...

      Updating grid network configuration on Name
      ```

11. If you selected option **10** when making Grid Network changes, select one of the following options:
    - **apply**: Apply the changes immediately and automatically restart each node if necessary.

      If the new network configuration will function simultaneously with the old network configuration without any external changes, you can use the **apply** option for a fully automated configuration change.

    - **stage**: Apply the changes the next time the nodes are restarted.

      If you need to make physical or virtual networking configuration changes for the new network configuration to function, you must use the **stage** option, shut down the affected nodes, make the necessary physical networking changes, and restart the affected nodes.

      > If you use the **stage** option, restart the node as soon as possible after staging to minimize disruptions.

    - **cancel**: Don't make any network changes at this time.

      If you were unaware that the proposed changes require nodes to be restarted, you can defer the changes to minimize user impact. Selecting **cancel** returns you to the main menu and preserves your changes so you can apply them later.

    After applying or staging changes, a new Recovery Package is generated as a result of the grid configuration change.

12. If configuration is stopped due to errors, the following options are available:
    - To terminate the IP change procedure and return to the main menu, enter **a**.

- To retry the operation that failed, enter **r**.
- To continue to the next operation, enter **c**.

  The failed operation can be retried later by selecting option **10** (Apply Changes) from the main menu. The IP change procedure will not be complete until all operations have completed successfully.

- If you had to manually intervene (to reboot a node, for example) and are confident that the action the tool thinks has failed was actually completed successfully, enter **f** to mark it as successful and move to the next operation.

13. Download a new Recovery Package from the Grid Manager.

    a. Select **MAINTENANCE** > **System** > **Recovery package**.

    b. Enter the provisioning passphrase.

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

Change IP addresses for all nodes in grid

If you need to change the Grid Network IP address for all nodes in the grid, you must follow this special procedure. You can't do a grid-wide Grid Network IP change using the procedure to change individual nodes.

Before you begin

- You have the `Passwords.txt` file.

To ensure that the grid starts up successfully, you must make all the changes at the same time.

> This procedure applies to the Grid Network only. You can't use this procedure to change IP addresses on the Admin or Client Networks.

If you want to change the IP addresses and MTU for the nodes at one site only, follow the Change node network configuration instructions.

Steps

1. Plan ahead for changes that you need to make outside of the Change IP tool, such as changes to DNS or NTP, and changes to the single sign-on (SSO) configuration, if used.

> If the existing NTP servers will not be accessible to the grid on the new IP addresses, add the new NTP servers before you perform the change-ip procedure.

> If the existing DNS servers will not be accessible to the grid on the new IP addresses, add the new DNS servers before you perform the change-ip procedure.

> If SSO is enabled for your StorageGRID system and any relying party trusts were configured using Admin Node IP addresses (instead of fully qualified domain names, as recommended), be prepared to update or reconfigure these relying party trusts in Active Directory Federation Services (AD FS) immediately after you change IP addresses. See Configure single sign-on.

If necessary, add the new subnet for the new IP addresses.

2. Log in to the primary Admin Node:
   a. Enter the following command: `ssh admin@primary_Admin_Node_IP`
   b. Enter the password listed in the `Passwords.txt` file.
   c. Enter the following command to switch to root: `su -`
   d. Enter the password listed in the `Passwords.txt` file.

      When you are logged in as root, the prompt changes from `$` to `#`.

3. Start the Change IP tool by entering the following command: `change-ip`
4. Enter the provisioning passphrase at the prompt.

   The main menu appears. By default, the `Selected nodes` field is set to `all`.

```
Welcome to the StorageGRID IP Change Tool.

Selected nodes: all

1:   SELECT NODES to edit
2:   EDIT IP/mask, gateway and MTU
3:   EDIT admin network subnet lists
4:   EDIT grid network subnet list
5:   SHOW changes
6:   SHOW full configuration, with changes highlighted
7:   VALIDATE changes
8:   SAVE changes, so you can resume later
9:   CLEAR all changes, to start fresh
10:  APPLY changes to the grid
0:   Exit

Selection:
```

5. On the main menu, select **2** to edit IP/subnet mask, gateway, and MTU information for all the nodes.
   a. Select **1** to make changes to the Grid Network.

      After you make your selection, the prompt shows the node names, Grid Network name, data type (IP/mask, Gateway, or MTU), and current values.

      Editing the IP address, prefix length, gateway, or MTU of a DHCP-configured interface will change the interface to static. A warning is displayed before each interface configured by DHCP.

      Interfaces configured as `fixed` can't be edited.

   b. To set a new value, enter it in the format shown for the current value.
   c. After editing all nodes you want to change, enter **q** to return to the main menu.

      Your changes are held until cleared or applied.

6. Review your changes by selecting one of the following options:
   - **5**: Shows edits in output that is isolated to show only the changed item. Changes are highlighted in green (additions) or red (deletions), as shown in the example output:

```
================================================================================
Site: RTP
================================================================================
username-x Grid    IP      [    172.16.0.239/21 ]: 172.16.0.240/21
username-x Grid    MTU     [              1400 ]: 9000
username-x Grid    MTU     [              1400 ]: 9000
username-x Grid    MTU     [              1400 ]: 9000
username-x Grid    MTU     [              1400 ]: 9000
username-x Grid    MTU     [              1400 ]: 9000
username-x Grid    MTU     [              1400 ]: 9000
username-x Grid    MTU     [              1400 ]: 9000
username-x Admin   IP      [    10.224.0.244/21 ]: 0.0.0.0/0
username-x Admin   IP      [    10.224.0.245/21 ]: 0.0.0.0/0
username-x Admin   IP      [    10.224.0.240/21 ]: 0.0.0.0/0
username-x Admin   IP      [    10.224.0.241/21 ]: 0.0.0.0/0
username-x Admin   IP      [    10.224.0.242/21 ]: 0.0.0.0/0
username-x Admin   IP      [    10.224.0.243/21 ]: 0.0.0.0/0
username-x Admin   Gateway [         10.224.0.1 ]: 0.0.0.0
username-x Admin   Gateway [         10.224.0.1 ]: 0.0.0.0
username-x Admin   Gateway [         10.224.0.1 ]: 0.0.0.0
username-x Admin   Gateway [         10.224.0.1 ]: 0.0.0.0
username-x Admin   Gateway [         10.224.0.1 ]: 0.0.0.0
username-x Admin   Gateway [         10.224.0.1 ]: 0.0.0.0
username-x Admin   MTU     [              1400 ]: 0
username-x Admin   MTU     [              1400 ]: 0
username-x Admin   MTU     [              1400 ]: 0
username-x Admin   MTU     [              1400 ]: 0
username-x Admin   MTU     [              1400 ]: 0
username-x Admin   MTU     [              1400 ]: 0
Press Enter to continue
```

- **6**: Shows edits in output that displays the full configuration. Changes are highlighted in green (additions) or red (deletions).

> Certain command line interfaces might show additions and deletions using strikethrough formatting. Proper display depends on your terminal client supporting the necessary VT100 escape sequences.

7. Select option **7** to validate all changes.

   This validation ensures that the rules for the Grid Network, such as not using overlapping subnets, aren't violated.

   In this example, validation returned errors.

```
Validating new networking configuration... FAILED.

  DK-10-224-5-20-G1: The admin subnet 172.18.0.0/16 overlaps the 172.18.0.0/21 grid network.
  DK-10-224-5-22-S1: Duplicate Grid IP 172.16.5.18 (also in use by DK-10-224-5-21-ADM1)

You must correct these errors before you can apply any changes.
Checking for Grid Network IP address swaps... PASSED.

Press Enter to continue
```

   In this example, validation passed.

```
Validating new networking configuration... PASSED.
Checking for Grid Network IP address swaps... PASSED.

Press Enter to continue
```

8. Once validation passes, select **10** to apply the new network configuration.
9. Select **stage** to apply the changes the next time the nodes are restarted.

> You must select **stage**. Don't perform a rolling restart, either manually or by selecting **apply** instead of **stage**; the grid will not start up successfully.

10. After your changes are complete, select **0** to exit the Change IP tool.
11. Shut down all nodes simultaneously.

> The entire grid must be shut down at once, so that all nodes are down at the same time.

12. Make the physical or virtual networking changes that are required.
13. Verify that all grid nodes are down.
14. Power on all nodes.
15. Once the grid starts up successfully:
    a. If you added new NTP servers, delete the old NTP server values.
    b. If you added new DNS servers, delete the old DNS server values.
16. Download the new Recovery Package from the Grid Manager.
    a. Select **MAINTENANCE** > **System** > **Recovery package**.
    b. Enter the provisioning passphrase.

Related information

- Add to or change subnet lists on Grid Network
- Shut down grid node

Add interfaces to existing node

Linux: Add Admin or Client interfaces to an existing node

Use these steps to add an interface on the Admin Network or the Client Network to a Linux node after it has been installed.

If you did not configure ADMIN_NETWORK_TARGET or CLIENT_NETWORK_TARGET in the node configuration file on the Linux host during installation, use this procedure to add the interface. For more information about the node configuration file, see the instructions for your Linux operating system:

- Install Red Hat Enterprise Linux or CentOS
- Install Ubuntu or Debian

You perform this procedure on the Linux server hosting the node that needs the new network assignment, not inside the node. This procedure only adds the interface to the node; a validation error occurs if you attempt to specify any other network parameters.

To provide addressing information, you must use the Change IP tool. See Change node network configuration.

Steps

1. Log in to the Linux server hosting the node.

2. Edit the node configuration file: `/etc/storagegrid/nodes/node-name.conf`.

   Don't specify any other network parameters, or a validation error will result.

   a. Add an entry for the new network target. For example:

      ```
      CLIENT_NETWORK_TARGET = bond0.3206
      ```

   b. Optional: Add an entry for the MAC address. For example:

      ```
      CLIENT_NETWORK_MAC = aa:57:61:07:ea:5c
      ```

3. Run the node validate command:

   ```
   sudo storagegrid node validate node-name
   ```

4. Resolve all validation errors.
5. Run the node reload command:

   ```
   sudo storagegrid node reload node-name
   ```

Linux: Add trunk or access interfaces to a node

You can add extra trunk or access interfaces to a Linux node after it has been installed. The interfaces you add are displayed on the VLAN interfaces page and the HA groups page.

Before you begin

- You have access to the instructions for installing StorageGRID on your Linux platform.
  - Install Red Hat Enterprise Linux or CentOS
  - Install Ubuntu or Debian
- You have the `Passwords.txt` file.
- You have specific access permissions.

Don't attempt to add interfaces to a node while a software upgrade, recovery procedure, or expansion procedure is active.

About this task

Use these steps to add one or more extra interfaces to a Linux node after the node has been installed. For example, you might want to add a trunk interface to an Admin or Gateway Node, so you can use VLAN interfaces to segregate the traffic belonging to different applications or tenants. Or, you might want to add an access interface to use in a high availability (HA) group.

If you add a trunk interface, you must configure a VLAN interface in StorageGRID. If you add an access interface, you can add the interface directly to an HA group; you don't need to configure a VLAN interface.

The node is unavailable for a brief time when you add interfaces. You should perform this procedure on one node at a time.

Steps

1. Log in to the Linux server hosting the node.
2. Using a text editor such as vim or pico, edit the node configuration file:

   `/etc/storagegrid/nodes/node-name.conf`

3. Add an entry to the file to specify the name and, optionally, the description of each extra interface you want to add to the node. Use this format.

   `INTERFACE_TARGET_nnnn=value`

   For *nnnn*, specify a unique number for each `INTERFACE_TARGET` entry you are adding.

   For *value*, specify the name of the physical interface on the bare-metal host. Then, optionally, add a comma and provide a description of the interface, which is displayed on the VLAN interfaces page and the HA groups page.

   For example:

   `INTERFACE_TARGET_0001=ens256, Trunk`

   > Don't specify any other network parameters, or a validation error will result.

4. Run the following command to validate your changes to the node configuration file:

   `sudo storagegrid node validate node-name`

   Address any errors or warnings before proceeding to the next step.

5. Run the following command to update the node's configuration:

   `sudo storagegrid node reload node-name`

After you finish

- If you added one or more trunk interfaces, go to configure VLAN interfaces to configure one or more VLAN interfaces for each new parent interface.
- If you added one or more access interfaces, go to configure high availability groups to add the new interfaces directly to HA groups.

VMware: Add trunk or access interfaces to a node

You can add a trunk or access interface to a VM node after the node has been installed. The interfaces you add are displayed on the VLAN interfaces page and the HA groups page.

Before you begin

- You have access to the instructions for installing StorageGRID on your VMware platform.
- You have Admin Node and Gateway Node VMware virtual machines.
- You have a network subnet that is not being used as Grid, Admin, or Client Network.
- You have the `Passwords.txt` file.

- You have specific access permissions.

> Don’t attempt to add interfaces to a node while a software upgrade, recovery procedure, or expansion procedure is active.

About this task

Use these steps to add one or more extra interfaces to a VMware node after the node has been installed. For example, you might want to add a trunk interface to an Admin or Gateway Node, so you can use VLAN interfaces to segregate the traffic belonging to different applications or tenants. Or you might want to add an access interface to use in a high availability (HA) group.

If you add a trunk interface, you must configure a VLAN interface in StorageGRID. If you add an access interface, you can add the interface directly to an HA group; you don’t need to configure a VLAN interface.

The node might be unavailable for a brief time when you add interfaces.

Steps

1. In vCenter, add a new network adapter (type VMXNET3) to an Admin Node and Gateway Node VM. Select **Connected** and **Connect At Power On** checkboxes.

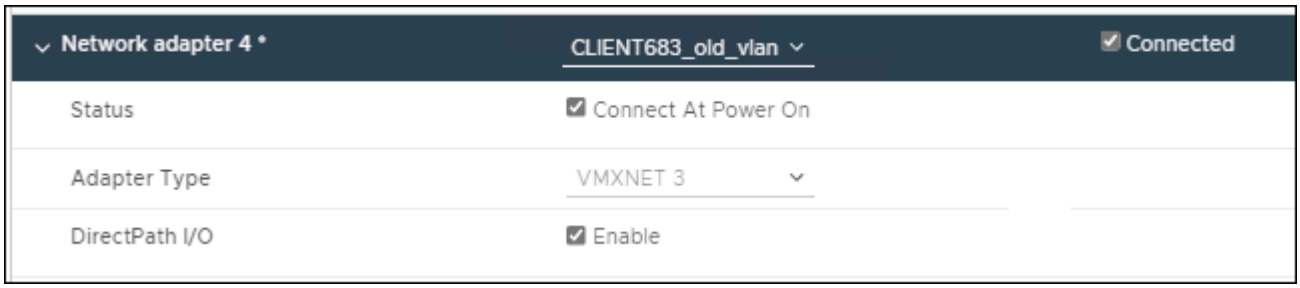

2. Use SSH to log in to the Admin Node or Gateway Node.
3. Use `ip link show` to confirm the new network interface ens256 is detected.

```
ip link show
1: lo: <LOOPBACK,UP,LOWER_UP> mtu 65536 qdisc noqueue state UNKNOWN mode
DEFAULT group default qlen 1000
    link/loopback 00:00:00:00:00:00 brd 00:00:00:00:00:00
2: eth0: <BROADCAST,MULTICAST,UP,LOWER_UP> mtu 1400 qdisc mq state UP mode
DEFAULT group default qlen 1000
    link/ether 00:50:56:a0:4e:5b brd ff:ff:ff:ff:ff:ff
3: eth1: <BROADCAST,MULTICAST> mtu 1500 qdisc noop state DOWN mode DEFAULT
group default qlen 1000
    link/ether 00:50:56:a0:fa:ce brd ff:ff:ff:ff:ff:ff
4: eth2: <BROADCAST,MULTICAST,UP,LOWER_UP> mtu 1400 qdisc mq state UP mode
DEFAULT group default qlen 1000
    link/ether 00:50:56:a0:d6:87 brd ff:ff:ff:ff:ff:ff
5: ens256: <BROADCAST,MULTICAST,UP,LOWER_UP> mtu 1500 qdisc mq master
ens256vrf state UP mode DEFAULT group default qlen 1000
    link/ether 00:50:56:a0:ea:88 brd ff:ff:ff:ff:ff:ff
```

After you finish

- If you added one or more trunk interfaces, go to configure VLAN interfaces to configure one or more VLAN interfaces for each new parent interface.
- If you added one or more access interfaces, go to configure high availability groups to add the new interfaces directly to HA groups.

Configure DNS servers

You can add, update, and remove DNS servers, so that you can use fully qualified domain name (FQDN) hostnames rather than IP addresses.

To use fully qualified domain names (FQDNs) instead of IP addresses when specifying hostnames for external destinations, specify the IP address of each DNS server you will use. These entries are used for AutoSupport, alert emails, SNMP notifications, platform services endpoints, Cloud Storage Pools, and more.

Before you begin

- You are signed in to the Grid Manager using a supported web browser.
- You have the Maintenance or Root access permission.
- You have the IP addresses of the DNS servers to configure.

About this task

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

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

Add a DNS server

Follow these steps to add a DNS server.

Steps

1. Select **MAINTENANCE** > **Network** > **DNS servers**.
2. Select **Add another server** to add a DNS server.
3. Select **Save**.

Modify a DNS server

Follow these steps to modify a DNS server.

Steps

1. Select **MAINTENANCE** > **Network** > **DNS servers**.
2. Select the IP address of the server name you want to edit and make the necessary changes.
3. Select **Save**.

Delete a DNS server

Follow these steps to delete an IP address of a DNS server.

Steps

1. Select **MAINTENANCE** > **Network** > **DNS servers**.
2. Select the delete icon ✕ next to the IP address.
3. Select **Save**.

Modify DNS configuration for single grid node

Rather than configure the DNS globally for the entire deployment, you can run a script to configure DNS differently for each grid node.

In general, you should use the **MAINTENANCE** > **Network** > **DNS servers** option on the Grid Manager to configure DNS servers. Only use the following script if you need to use different DNS servers for different grid nodes.

Steps

1. Log in to the primary Admin Node:
   a. Enter the following command: `ssh admin@primary_Admin_Node_IP`
   b. Enter the password listed in the `Passwords.txt` file.
   c. Enter the following command to switch to root: `su -`
   d. Enter the password listed in the `Passwords.txt` file.

      When you are logged in as root, the prompt changes from `$` to `#`.

   e. Add the SSH private key to the SSH agent. Enter: `ssh-add`
   f. Enter the SSH Access Password listed in the `Passwords.txt` file.
2. Log in to the node you want to update with a custom DNS configuration: `ssh node_IP_address`
3. Run the DNS setup script: `setup_resolv.rb.`

   The script responds with the list of supported commands.

```
Tool to modify external name servers

available commands:
  add search <domain>
               add a specified domain to search list
               e.g.> add search netapp.com
  remove search <domain>
               remove a specified domain from list
               e.g.> remove search netapp.com
  add nameserver <ip>
               add a specified IP address to the name server list
               e.g.> add nameserver 192.0.2.65
  remove nameserver <ip>
               remove a specified IP address from list
               e.g.> remove nameserver 192.0.2.65
  remove nameserver all
               remove all nameservers from list
  save         write configuration to disk and quit
  abort        quit without saving changes
  help         display this help message


Current list of name servers:
    192.0.2.64
Name servers inherited from global DNS configuration:
    192.0.2.126
    192.0.2.127
Current list of search entries:
    netapp.com

Enter command [`add search <domain>|remove search <domain>|add
nameserver <ip>`]
              [`remove nameserver <ip>|remove nameserver
all|save|abort|help`]
```

4. Add the IPv4 address of a server that provides domain name service for your network: `add <nameserver IP_address>`
5. Repeat the `add nameserver` command to add name servers.
6. Follow instructions as prompted for other commands.
7. Save your changes and exit the application: `save`
8. Close the command shell on the server: `exit`
9. For each grid node, repeat the steps from logging into the node through closing the command shell.
10. When you no longer require passwordless access to other servers, remove the private key from the SSH

agent. Enter: `ssh-add -D`

Manage NTP servers

You can add, update, or remove Network Time Protocol (NTP) servers to ensure that data is synchronized accurately between grid nodes in your StorageGRID system.

Before you begin

- You are signed in to the Grid Manager using a supported web browser.
- You have the Maintenance or Root access permission.
- You have the provisioning passphrase.
- You have the IPv4 addresses of the NTP servers to configure.

How StorageGRID uses NTP

The StorageGRID system uses the Network Time Protocol (NTP) to synchronize time between all grid nodes in the grid.

At each site, at least two nodes in the StorageGRID system are assigned the primary NTP role. They synchronize to a suggested minimum of four, and a maximum of six, external time sources and with each other. Every node in the StorageGRID system that is not a primary NTP node acts as an NTP client and synchronizes with these primary NTP nodes.

The external NTP servers connect to the nodes to which you previously assigned primary NTP roles. For this reason, specifying at least two nodes with primary NTP roles is recommended.

NTP server guidelines

Follow these guidelines to protect against timing issues:

- The external NTP servers connect to the nodes to which you previously assigned primary NTP roles. For this reason, specifying at least two nodes with primary NTP roles is recommended.
- Make sure at least two nodes at each site can access at least four external NTP sources. If only one node at a site can reach the NTP sources, timing issues will occur if that node goes down. In addition, designating two nodes per site as primary NTP sources ensures accurate timing if a site is isolated from the rest of the grid.
- The specified external NTP servers must use the NTP protocol. You must specify NTP server references of Stratum 3 or better to prevent issues with time drift.

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

Configure NTP servers

Follow these steps to add, update, or remove NTP servers.

Steps

1. Select **MAINTENANCE** > **Network** > **NTP servers**.
2. In the Servers section, add, update, or remove NTP server entries, as necessary.

   You should include at least four NTP servers, and you can specify up to six servers.

3. Enter the provisioning passphrase for your StorageGRID system, then select **Save**.

   The page is disabled until the configuration updates are complete.

> If all of your NTP servers fail the connection test after you save the new NTP servers, don't proceed. Contact technical support.

Resolve NTP server issues

If you encounter problems with the stability or availability of the NTP servers originally specified during installation, you can update the list of external NTP sources that the StorageGRID system uses by adding additional servers, or updating or removing existing servers.

Restore network connectivity for isolated nodes

Under certain circumstances, such as site- or grid-wide IP address changes, one or more groups of nodes might not be able to contact the rest of the grid.

About this task

In the Grid Manager (**SUPPORT** > **Tools** > **Grid topology**), if a node is gray, or if a node is blue with many of its services showing a status other than Running, you should check for node isolation.

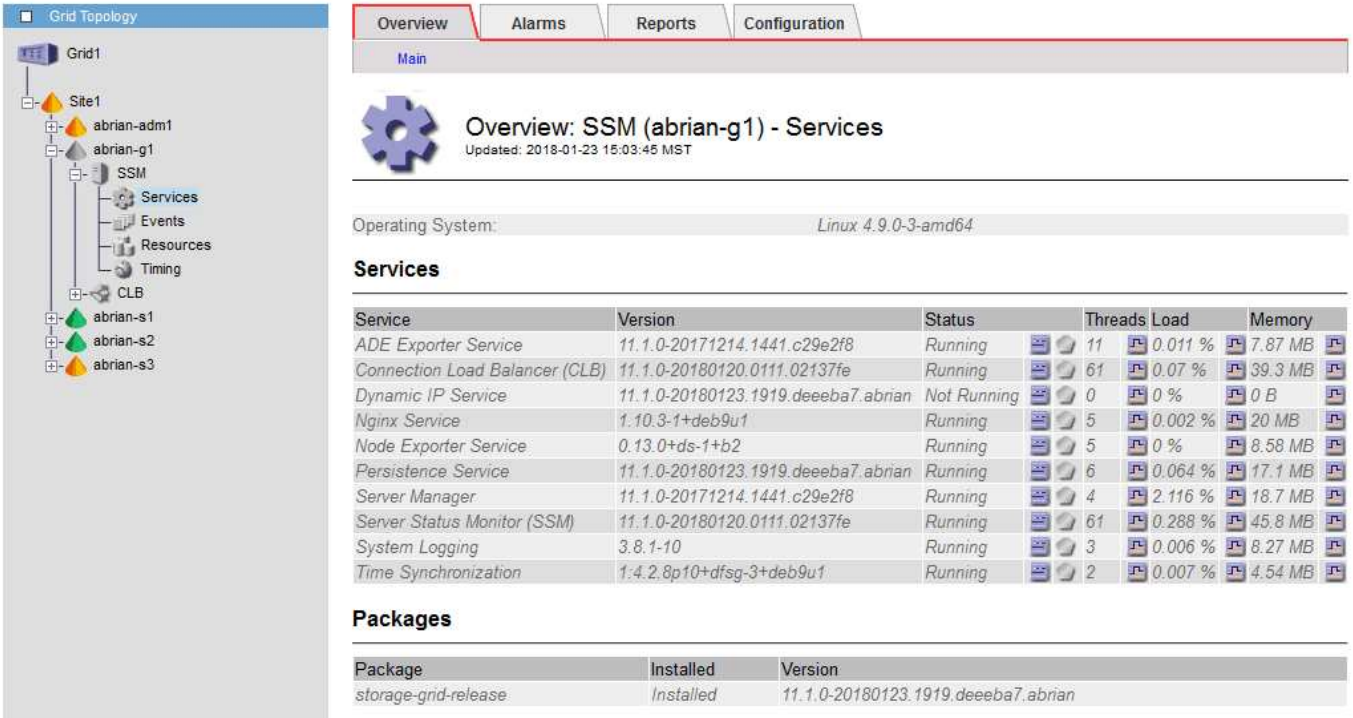

Some of the consequences of having isolated nodes include the following:

- If multiple nodes are isolated, you might not be able to sign in to or access the Grid Manager.

- If multiple nodes are isolated, the storage usage and quota values shown on the dashboard for the Tenant Manager might be out of date. The totals will be updated when network connectivity is restored.

To resolve the isolation issue, you run a command line utility on each isolated node or on one node in a group (all nodes in a subnet that does not contain the primary Admin Node) that is isolated from the grid. The utility provides the nodes with the IP address of a non-isolated node in the grid, which allows the isolated node or group of nodes to contact the entire grid again.

> If the multicast domain name system (mDNS) is disabled in the networks, the command line utility might have to be run on each isolated node.

Steps

1. Access the node and check `/var/local/log/dynip.log` for isolation messages.

   For example:

   ```
   [2018-01-09T19:11:00.545] UpdateQueue - WARNING -- Possible isolation,
   no contact with other nodes.
   If this warning persists, manual action might be required.
   ```

   If you are using the VMware console, it will contain a message that the node might be isolated.

   On Linux deployments, isolation messages would appear in `/var/log/storagegrid/node/<nodename>.log` files.

2. If the isolation messages are recurring and persistent, run the following command:

   `add_node_ip.py <address>`

   where `<address>` is the IP address of a remote node that is connected to the grid.

   ```
   # /usr/sbin/add_node_ip.py  10.224.4.210

   Retrieving local host information
   Validating remote node at address 10.224.4.210
   Sending node IP hint for 10.224.4.210 to local node
   Local node found on remote node. Update complete.
   ```

3. Verify the following for each node that was previously isolated:
   - The node's services have started.
   - The status of the Dynamic IP service is "Running" after you run the `storagegrid-status` command.
   - In the Grid Topology tree, the node no longer appears disconnected from the rest of the grid.

> If running the `add_node_ip.py` command does not solve the problem, there could be other networking issues that need to be resolved.

Host-level and middleware procedures

Linux: Migrate grid node to new host

You can migrate one or more StorageGRID nodes from one Linux host (the *source host*) to another Linux host (the *target host*) to perform host maintenance without impacting the functionality or availability of your grid.

For example, you might want to migrate a node to perform OS patching and reboot.

Before you begin

- You planned your StorageGRID deployment to include migration support.
  - Node container migration requirements for Red Hat Enterprise Linux or Centos
  - Node container migration requirements for Ubuntu or Debian
- The target host is already prepared for StorageGRID use.
- Shared storage is used for all per-node storage volumes
- Network interfaces have consistent names across hosts.

> In a production deployment, don't run more than one Storage Node on a single host. Using a dedicated host for each Storage Node provides an isolated failure domain.
>
> Other types of nodes, such as Admin Nodes or Gateway Nodes, can be deployed on the same host. However, if you have multiple nodes of the same type (two Gateway Nodes, for example), don't install all instances on the same host.

Export node from source host

As a first step, shut down the grid node and export it from the source Linux host.

Run the following commands on the *source host*.

Steps

1. Obtain the status of all nodes currently running on the source host.

   ```
   sudo storagegrid node status all
   ```

   Example output:

   ```
   Name Config-State Run-State
   DC1-ADM1 Configured Running
   DC1-ARC1 Configured Running
   DC1-GW1 Configured Running
   DC1-S1 Configured Running
   DC1-S2 Configured Running
   DC1-S3 Configured Running
   ```

2. Identify the name of the node you want to migrate, and stop it if its Run-State is Running.

```
sudo storagegrid node stop DC1-S3
```

Example output:

```
Stopping node DC1-S3
Waiting up to 630 seconds for node shutdown
```

3. Export the node from the source host.

```
sudo storagegrid node export DC1-S3
```

Example output:

```
Finished exporting node DC1-S3 to /dev/mapper/sgws-dc1-s3-var-local.
Use 'storagegrid node import /dev/mapper/sgws-dc1-s3-var-local' if you
want to import it again.
```

4. Make note of the `import` command suggested in the output.

   You will run this command on the target host in the next step.

Import node on target host

After exporting the node from the source host, you import and validate the node on the target host. Validation confirms that the node has access to the same block storage and network interface devices as it had on the source host.

Run the following commands on the *target host*.

Steps

1. Import the node on the target host.

```
sudo storagegrid node import /dev/mapper/sgws-dc1-s3-var-local
```

Example output:

```
Finished importing node DC1-S3 from /dev/mapper/sgws-dc1-s3-var-local.
You should run 'storagegrid node validate DC1-S3'
```

2. Validate the node configuration on the new host.

```
sudo storagegrid node validate DC1-S3
```

Example output:

```
Confirming existence of node DC1-S3... PASSED
Checking configuration file /etc/storagegrid/nodes/DC1-S3.conf for node
DC1-S3... PASSED
Checking for duplication of unique values... PASSED
```

3. If any validation errors occur, address them before starting the migrated node.

   For troubleshooting information, see the StorageGRID installation instructions for your Linux operating system.

   - Install Red Hat Enterprise Linux or CentOS
   - Install Ubuntu or Debian

Start migrated node

After you validate the migrated node, you start the node by running a command on the *target host*.

Steps

1. Start the node on the new host.

   ```
   sudo storagegrid node start DC1-S3
   ```

2. Sign in to the Grid Manager and verify that the status of the node is green with no alert.

   > Verifying that the status of the node is green ensures that the migrated node has fully restarted and rejoined the grid. If the status is not green, don't migrate any additional nodes so that you will not have more than one node out of service.

3. If you are unable to access the Grid Manager, wait for 10 minutes, then run the following command:

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

   Confirm that the migrated node has a Run-State of Running.

Archive Node maintenance for TSM middleware

Archive Nodes might be configured to target either tape through a TSM middleware server or the cloud through the S3 API. Once configured, an Archive Node's target can't be changed.

If the server hosting the Archive Node fails, replace the server and follow the appropriate recovery procedure.

Fault with archival storage devices

If you determine that there is a fault with the archival storage device that the Archive Node is accessing through Tivoli Storage Manager (TSM), take the Archive Node offline to limit the number of alarms displayed in the StorageGRID system. You can then use the administrative tools of the TSM server or the storage device, or both, to further diagnose and resolve the problem.

Take the Target component offline

Before undertaking any maintenance of the TSM middleware server that might result in it becoming unavailable to the Archive Node, take the Target component offline to limit the number of alarms that are triggered if the TSM middleware server becomes unavailable.

Before you begin

You are signed in to the Grid Manager using a supported web browser.

Steps

1. Select **SUPPORT** > **Tools** > **Grid topology**.
2. Select **Archive Node** > **ARC** > **Target** > **Configuration** > **Main**.
3. Change the value of Tivoli Storage Manager State to **Offline**, and click **Apply Changes**.
4. After maintenance is complete, change the value of Tivoli Storage Manager State to **Online**, and click **Apply Changes**.

Tivoli Storage Manager administrative tools

The dsmadmc tool is the administrative console for the TSM middleware server that is installed on the Archive Node. You can access the tool by typing `dsmadmc` at the command line of the server. Log in to the administrative console using the same administrative user name and password that is configured for the ARC service.

The `tsmquery.rb` script was created to generate status information from dsmadmc in a more readable form. You can run this script by entering the following command at the command line of the Archive Node: `/usr/local/arc/tsmquery.rb status`

For more information about the TSM administrative console dsmadmc, see the *Tivoli Storage Manager for Linux: Administrator's Reference*.

Object permanently unavailable

When the Archive Node requests an object from the Tivoli Storage Manager (TSM) server and the retrieval fails, the Archive Node retries the request after an interval of 10 seconds. If the object is permanently unavailable (for example, because the object is corrupted on tape), the TSM API has no way to indicate this to the Archive Node, so the Archive Node continues to retry the request.

When this situation occurs, an alarm is triggered, and the value continues to increase. To see the alarm, select **SUPPORT** > **Tools** > **Grid topology**. Then, select **Archive Node** > **ARC** > **Retrieve** > **Request Failures**.

If the object is permanently unavailable, you must identify the object and then manually cancel the Archive Node's request as described in the procedure, Determining if objects are permanently unavailable.

A retrieval can also fail if the object is temporarily unavailable. In this case, subsequent retrieval requests should eventually succeed.

If the StorageGRID system is configured to use an ILM rule that creates a single object copy and that copy can't be retrieved, the object is lost and can't be recovered. However, you must still follow the procedure to determine if the object is permanently unavailable to "clean up" the StorageGRID system, to cancel the Archive Node's request, and to purge metadata for the lost object.

Determining if objects are permanently unavailable

You can determine if objects are permanently unavailable by making a request using the TSM administrative console.

Before you begin

- You have specific access permissions.
- You have the `Passwords.txt` file.
- You have the IP address of an Admin Node.

About this task

This example is provided for your information. This procedure can't help you identify all failure conditions that might result in unavailable objects or tape volumes. For information about TSM administration, see TSM Server documentation.

Steps

1. Log in to an Admin Node:
   a. Enter the following command: `ssh admin@Admin_Node_IP`
   b. Enter the password listed in the `Passwords.txt` file.
2. Identify the object or objects that could not be retrieved by the Archive Node:
   a. Go to the directory containing the audit log files: `cd /var/local/audit/export`

      The active audit log file is named audit.log. Once a day, the active `audit.log` file is saved, and a new `audit.log` file is started. The name of the saved file indicates when it was saved, in the format `yyyy-mm-dd.txt`. After a day, the saved file is compressed and renamed, in the format `yyyy-mm-dd.txt.gz`, which preserves the original date.

   b. Search the relevant audit log file for messages indicating that an archived object could not be retrieved. For example, enter: `grep ARCE audit.log | less -n`

      When an object can't be retrieved from an Archive Node, the ARCE audit message (Archive Object Retrieve End) displays ARUN (archive middleware unavailable) or GERR (general error) in the result field. The following example line from the audit log shows that the ARCE message terminated with the result ARUN for CBID 498D8A1F681F05B3.

      ```
      [AUDT:[CBID(UI64):0x498D8A1F681F05B3][VLID(UI64):□20091127][RSLT(FC32
      ):ARUN][AVER(UI32):7]
      [ATIM(UI64):1350613602969243][ATYP(FC32):ARCE][ANID(UI32):13959984][A
      MID(FC32):ARCI]
      [ATID(UI64):4560349751312520631]]
      ```

      For more information see the instructions for understanding audit messages.

   c. Record the CBID of each object that had a request failure.

      You might also want to record the following additional information used by the TSM to identify objects saved by the Archive Node:

- **File Space Name**: Equivalent to the Archive Node ID. To find the Archive Node ID, select **SUPPORT** > **Tools** > **Grid topology**. Then, select **Archive Node** > **ARC** > **Target** > **Overview**.
- **High Level Name**: Equivalent to the volume ID assigned to the object by the Archive Node. The volume ID takes the form of a date (for example, `20091127`), and is recorded as the VLID of the object in archive audit messages.
- **Low Level Name**: Equivalent to the CBID assigned to an object by the StorageGRID system.

d. Log out of the command shell: `exit`

3. Check the TSM server to see if the objects identified in step 2 are permanently unavailable:

a. Log in to the administrative console of the TSM server: `dsmadmc`

Use the administrative user name and password that are configured for the ARC service. Enter the user name and password in the Grid Manager. (To see the user name, select **SUPPORT** > **Tools** > **Grid topology**. Then, select **Archive Node** > **ARC** > **Target** > **Configuration**.)

b. Determine if the object is permanently unavailable.

For example, you might search the TSM activity log for a data integrity error for that object. The following example shows a search of the activity log for the past day for an object with CBID `498D8A1F681F05B3`.

```
> query actlog begindate=-1 search=276C14E94082CC69
12/21/2008 05:39:15 ANR0548W Retrieve or restore
failed for session 9139359 for node DEV-ARC-20 (Bycast ARC)
processing file space /19130020 4 for file /20081002/
498D8A1F681F05B3 stored as Archive - data
integrity error detected. (SESSION: 9139359)
>
```

Depending on the nature of the error, the CBID might not be recorded in the TSM activity log. You might need to search the log for other TSM errors around the time of the request failure.

c. If an entire tape is permanently unavailable, identify the CBIDs for all objects stored on that volume: `query content TSM_Volume_Name`

where `TSM_Volume_Name` is the TSM name for the unavailable tape. The following is an example of the output for this command:

```
 > query content TSM-Volume-Name
Node Name     Type Filespace  FSID Client's Name for File Name
------------- ---- ---------- ---- --------------------------
DEV-ARC-20    Arch /19130020  216  /20081201/ C1D172940E6C7E12
DEV-ARC-20    Arch /19130020  216  /20081201/ F1D7FBC2B4B0779E
```

The `Client's Name for File Name` is the same as the Archive Node volume ID (or TSM "high level name") followed by the object's CBID (or TSM "low level name"). That is, the `Client's Name for File Name` takes the form `/Archive Node volume ID /CBID`. In the first line of the

example output, the `Client's Name for File Name` is `/20081201/ C1D172940E6C7E12`.

Recall also that the `Filespace` is the node ID of the Archive Node.

You will need the CBID of each object stored on the volume and the node ID of the Archive Node to cancel the retrieval request.

4. For each object that is permanently unavailable, cancel the retrieval request and issue a command to inform the StorageGRID system that the object copy was lost:

> Use the ADE Console with caution. If the console is used improperly, it is possible to interrupt system operations and corrupt data. Enter commands carefully, and only use the commands documented in this procedure.

   a. If you aren't already logged in to the Archive Node, log in as follows:
      i. Enter the following command: `ssh admin@grid_node_IP`
      ii. Enter the password listed in the `Passwords.txt` file.
      iii. Enter the following command to switch to root: `su -`
      iv. Enter the password listed in the `Passwords.txt` file.
   b. Access the ADE console of the ARC service: `telnet localhost 1409`
   c. Cancel the request for the object: `/proc/BRTR/cancel -c CBID`

      where `CBID` is the identifier of the object that can't be retrieved from the TSM.

      If the only copies of the object are on tape, the "bulk retrieval" request is canceled with a message, "1 requests canceled". If copies of the object exist elsewhere in the system, the object retrieval is processed by a different module so the response to the message is "0 requests canceled".

   d. Issue a command to notify the StorageGRID system that an object copy has been lost and that an additional copy must be made: `/proc/CMSI/Object_Lost CBID node_ID`

      where `CBID` is the identifier of the object that can't be retrieved from the TSM server, and `node_ID` is the node ID of the Archive Node where the retrieval failed.

      You must enter a separate command for each lost object copy: entering a range of CBIDs is not supported.

      In most cases, the StorageGRID system immediately begins to make additional copies of object data to ensure that the system's ILM policy is followed.

      However, if the ILM rule for the object specified that only one copy be made and that copy has now been lost, the object can't be recovered. In this case running the `Object_Lost` command purges the lost object's metadata from the StorageGRID system.

      When the `Object_Lost` command completes successfully, the following message is returned:

```
CLOC_LOST_ANS returned result 'SUCS'
```

> The `/proc/CMSI/Object_Lost` command is only valid for lost objects that are stored on Archive Nodes.

e. Exit the ADE Console: `exit`

f. Log out of the Archive Node: `exit`

5. Reset the value of Request Failures in the StorageGRID system:

   a. Go to **Archive Node** > **ARC** > **Retrieve** > **Configuration**, and select **Reset Request Failure Count**.

   b. Click **Apply Changes**.

Related information

Administer StorageGRID

Review audit logs

VMware: Configure virtual machine for automatic restart

If the virtual machine does not restart after VMware vSphere Hypervisor is restarted, you might need to configure the virtual machine for automatic restart.

You should perform this procedure if you notice that a virtual machine does not restart while you are recovering a grid node or performing another maintenance procedure.

Steps

1. In the VMware vSphere Client tree, select the virtual machine that is not started.
2. Right-click the virtual machine, and select **Power on**.
3. Configure VMware vSphere Hypervisor to restart the virtual machine automatically in future.

Grid node procedures

Grid node procedures: Overview

You might need to perform procedures on a specific grid node. While you can perform a few of these procedures from Grid Manager, most of the procedures require you to access Server Manager from the node's command line.

Server Manager runs on every grid node to supervise the starting and stopping of services and to ensure that services gracefully join and leave the StorageGRID system. Server Manager also monitors the services on every grid node and will automatically attempt to restart any services that report faults.

> You should access Server Manager only if technical support has directed you to do so.

> You must close the current command shell session and log out after you are finished with Server Manager. Enter: `exit`

View Server Manager status and version

For each grid node, you can view the current status and version of Server Manager

running on that grid node. You can also obtain the current status of all services running on that grid node.

Before you begin

You have the `Passwords.txt` file.

Steps

1. Log in to the grid node:
   a. Enter the following command: `ssh admin@grid_node_IP`
   b. Enter the password listed in the `Passwords.txt` file.
   c. Enter the following command to switch to root: `su -`
   d. Enter the password listed in the `Passwords.txt` file.

   When you are logged in as root, the prompt changes from `$` to `#`.

2. View the current status of Server Manager running on the grid node: **`service servermanager status`**

   The current status of Server Manager running on the grid node is reported (running or not). If Server Manager's status is `running`, the time it has been running since last it was started is listed. For example:

   ```
   servermanager running for 1d, 13h, 0m, 30s
   ```

3. View the current version of Server Manager running on a grid node: **`service servermanager version`**

   The current version is listed. For example:

   ```
   11.1.0-20180425.1905.39c9493
   ```

4. Log out of the command shell: **`exit`**

View current status of all services

You can view the current status of all services running on a grid node at any time.

Before you begin

You have the `Passwords.txt` file.

Steps

1. Log in to the grid node:
   a. Enter the following command: `ssh admin@grid_node_IP`
   b. Enter the password listed in the `Passwords.txt` file.
   c. Enter the following command to switch to root: `su -`

d. Enter the password listed in the `Passwords.txt` file.

When you are logged in as root, the prompt changes from `$` to `#`.

2. View the status of all services running on the grid node: `storagegrid-status`

   For example, the output for the primary Admin Node shows the current status of the AMS, CMN, and NMS services as Running. This output is updated immediately if the status of a service changes.

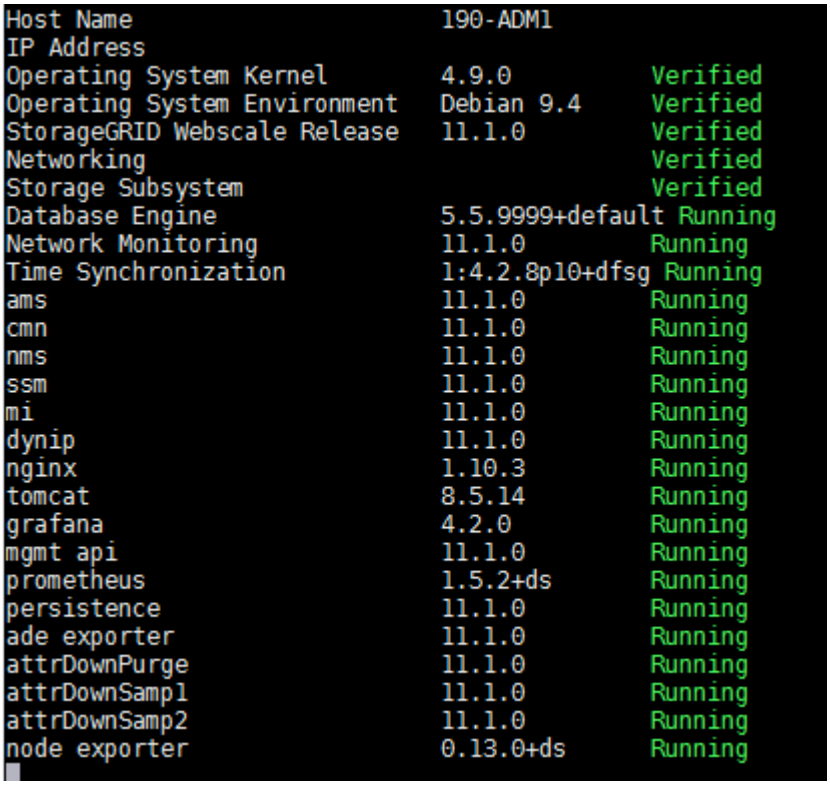

3. Return to the command line, press **Ctrl+C**.
4. Optionally, view a static report for all services running on the grid node: `/usr/local/servermanager/reader.rb`

   This report includes the same information as the continuously updated report, but it is not updated if the status of a service changes.

5. Log out of the command shell: `exit`

Start Server Manager and all services

You might need to start Server Manager, which also starts all services on the grid node.

Before you begin

You have the `Passwords.txt` file.

About this task

Starting Server Manager on a grid node where it is already running results in a restart of Server Manager and all services on the grid node.

Steps

1. Log in to the grid node:

   a. Enter the following command: `ssh admin@grid_node_IP`

   b. Enter the password listed in the `Passwords.txt` file.

   c. Enter the following command to switch to root: `su -`

   d. Enter the password listed in the `Passwords.txt` file.

   When you are logged in as root, the prompt changes from `$` to `#`.

2. Start Server Manager: `service servermanager start`

3. Log out of the command shell: `exit`

Restart Server Manager and all services

You might need to restart server manager and all services running on a grid node.

Before you begin

You have the `Passwords.txt` file.

Steps

1. Log in to the grid node:

   a. Enter the following command: `ssh admin@grid_node_IP`

   b. Enter the password listed in the `Passwords.txt` file.

   c. Enter the following command to switch to root: `su -`

   d. Enter the password listed in the `Passwords.txt` file.

   When you are logged in as root, the prompt changes from `$` to `#`.

2. Restart Server Manager and all services on the grid node: `service servermanager restart`

   Server Manager and all services on the grid node are stopped and then restarted.

   > Using the `restart` command is the same as using the `stop` command followed by the `start` command.

3. Log out of the command shell: `exit`

Stop Server Manager and all services

Server Manager is intended to run at all times, but you might need to stop Server Manager and all services running on a grid node.

Before you begin

You have the `Passwords.txt` file.

Steps

1. Log in to the grid node:

a. Enter the following command: `ssh admin@grid_node_IP`

b. Enter the password listed in the `Passwords.txt` file.

c. Enter the following command to switch to root: `su -`

d. Enter the password listed in the `Passwords.txt` file.

When you are logged in as root, the prompt changes from `$` to `#`.

2. Stop Server manager and all services running on the grid node: `service servermanager stop`

   Server Manager and all services running on the grid node are gracefully terminated. Services can take up to 15 minutes to shut down.

3. Log out of the command shell: `exit`

View current status of service

You can view the current status of a services running on a grid node at any time.

Before you begin

You have the `Passwords.txt` file.

Steps

1. Log in to the grid node:

   a. Enter the following command: `ssh admin@grid_node_IP`

   b. Enter the password listed in the `Passwords.txt` file.

   c. Enter the following command to switch to root: `su -`

   d. Enter the password listed in the `Passwords.txt` file.

   When you are logged in as root, the prompt changes from `$` to `#`.

2. View the current status of a service running on a grid node: `**service *servicename* status** The current status of the requested service running on the grid node is reported (running or not). For example:

   ```
   cmn running for 1d, 14h, 21m, 2s
   ```

3. Log out of the command shell: **`exit`**

Stop service

Some maintenance procedures require you to stop a single service while keeping other services on the grid node running. Only stop individual services when directed to do so by a maintenance procedure.

Before you begin

You have the `Passwords.txt` file.

About this task

When you use these steps to “administratively stop” a service, Server Manager will not automatically restart the service. You must either start the single service manually or restart Server Manager.

If you need to stop the LDR service on a Storage Node, be aware that it might take a while to stop the service if there are active connections.

Steps

1. Log in to the grid node:
   a. Enter the following command: `ssh admin@grid_node_IP`
   b. Enter the password listed in the `Passwords.txt` file.
   c. Enter the following command to switch to root: `su -`
   d. Enter the password listed in the `Passwords.txt` file.

   When you are logged in as root, the prompt changes from `$` to `#`.
2. Stop an individual service: `service servicename stop`

   For example:

   ```
   service ldr stop
   ```

   Services can take up to 11 minutes to stop.
3. Log out of the command shell: `exit`

Related information

Force service to terminate

Place appliance into maintenance mode

You must place the appliance into maintenance mode before performing specific maintenance procedures.

Before you begin

- You are signed in to the Grid Manager using a supported web browser.
- You have the Maintenance or Root access permission. For details, see the instructions for administering StorageGRID.

About this task

In rare instances, placing a StorageGRID appliance into maintenance mode might make the appliance unavailable for remote access.

The admin account password and SSH host keys for a StorageGRID appliance in maintenance mode remain the same as they were when the appliance was in service.

Steps

1. From the Grid Manager, select **NODES**.
2. From the tree view of the Nodes page, select the appliance Storage Node.
3. Select **Tasks**.

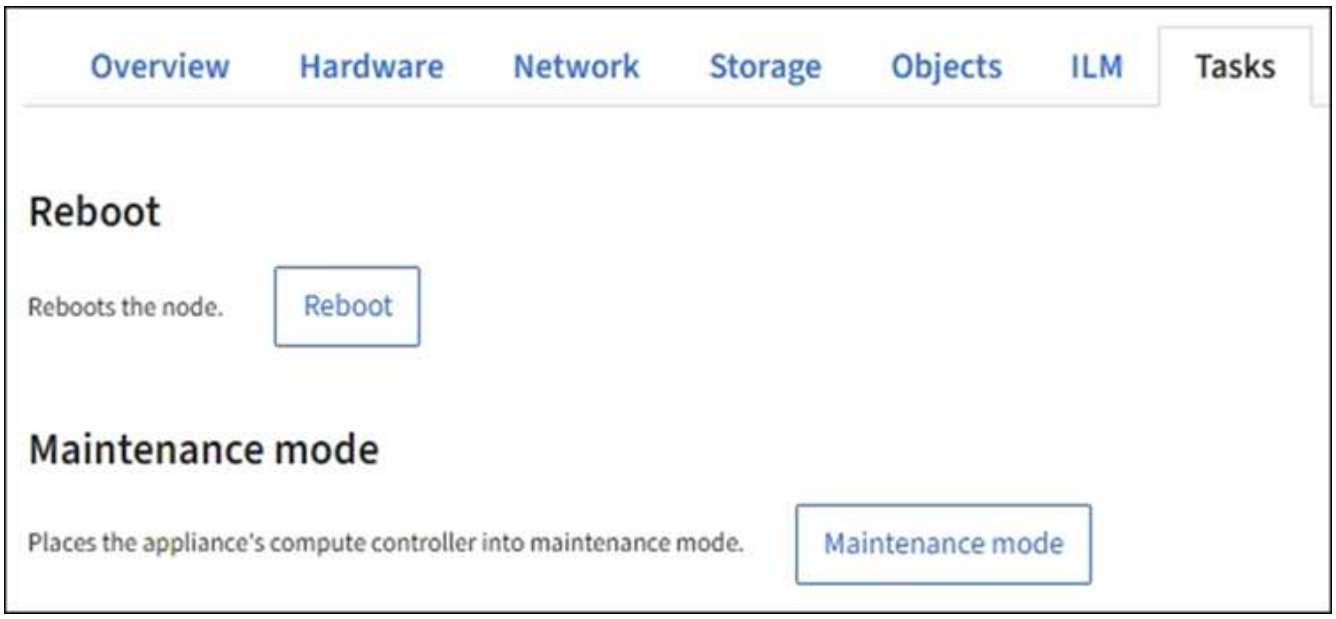

4. Select **Maintenance mode**. A confirmation dialog box appears.
5. Enter the provisioning passphrase, and select **OK**.

   A progress bar and a series of messages, including "Request Sent","Stopping StorageGRID", and "Rebooting", indicate that the appliance is completing the steps for entering maintenance mode.

   When the appliance is in maintenance mode, a confirmation message lists the URLs you can use to access the StorageGRID Appliance Installer.

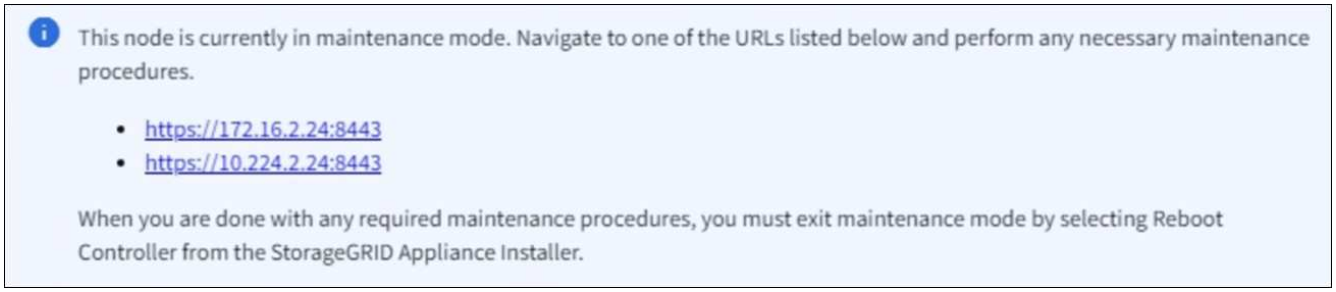

6. To access the StorageGRID Appliance Installer, browse to any of the URLs displayed.

   If possible, use the URL containing the IP address of the appliance's Admin Network port.

   > Accessing `https://169.254.0.1:8443` requires a direct connection to the local management port.

7. From the StorageGRID Appliance Installer, confirm that the appliance is in maintenance mode.

8. Perform any required maintenance tasks.
9. After completing maintenance tasks, exit maintenance mode and resume normal node operation. From the StorageGRID Appliance Installer, select **Advanced** > **Reboot Controller**, and then select **Reboot into StorageGRID**.

It can take up to 20 minutes for the appliance to reboot and rejoin the grid. To confirm that the reboot is complete and that the node has rejoined the grid, go back to the Grid Manager. The **Nodes** page should display a normal status (green check mark icon to the left of the node name) for the appliance node, indicating that no alerts are active and the node is connected to the grid.

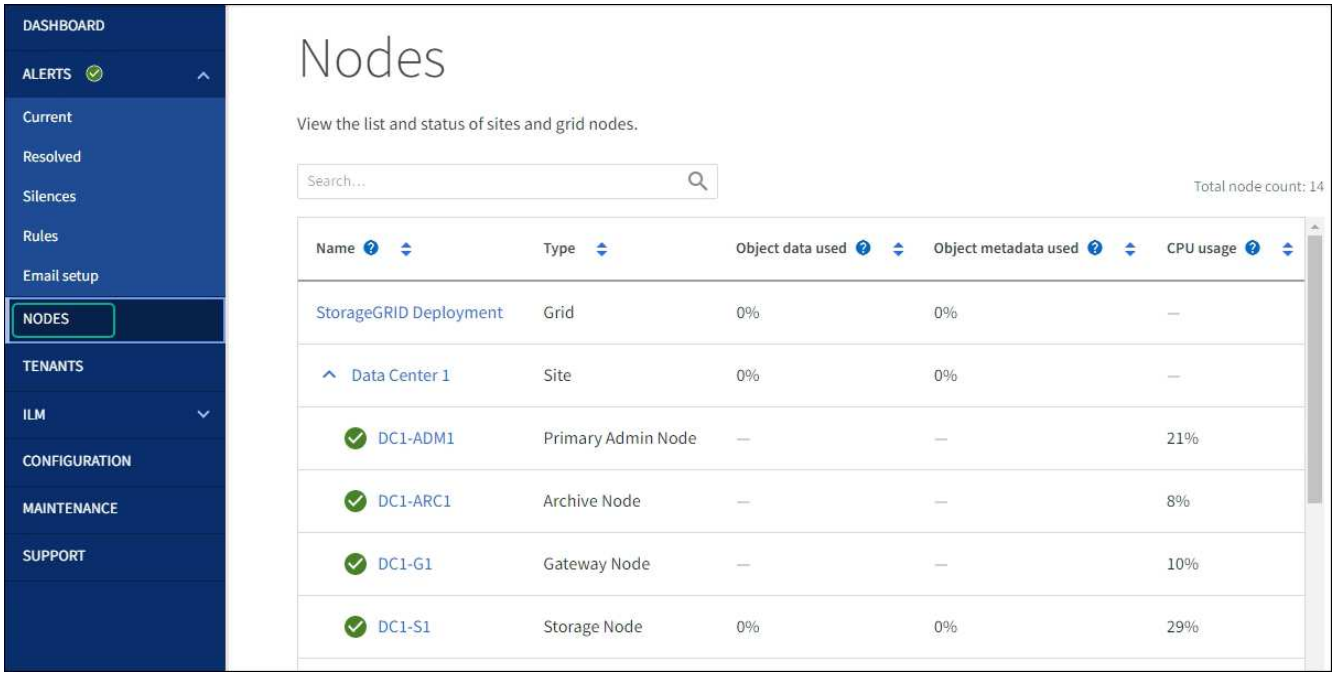

Force service to terminate

If you need to stop a service immediately, you can use the `force-stop` command.

Before you begin

You have the `Passwords.txt` file.

Steps

1. Log in to the grid node:

   a. Enter the following command: `ssh admin@grid_node_IP`

   b. Enter the password listed in the `Passwords.txt` file.

   c. Enter the following command to switch to root: `su -`

   d. Enter the password listed in the `Passwords.txt` file.

   When you are logged in as root, the prompt changes from `$` to `#`.

2. Manually force the service to terminate: `service servicename force-stop`

   For example:

   ```
   service ldr force-stop
   ```

   The system waits 30 seconds before terminating the service.

3. Log out of the command shell: `exit`

Start or restart service

You might need to start a service that has been stopped, or you might need to stop and restart a service.

Before you begin

You have the `Passwords.txt` file.

Steps

1. Log in to the grid node:

   a. Enter the following command: `ssh admin@grid_node_IP`

   b. Enter the password listed in the `Passwords.txt` file.

   c. Enter the following command to switch to root: `su -`

   d. Enter the password listed in the `Passwords.txt` file.

   When you are logged in as root, the prompt changes from `$` to `#`.

2. Decide which command to issue, based on whether the service is currently running or stopped.

   - If the service is currently stopped, use the `start` command to start the service manually: `service servicename start`

     For example:

     ```
     service ldr start
     ```

- If the service is currently running, use the `restart` command to stop the service and then restart it: `service servicename restart`

  For example:

  ```
  service ldr restart
  ```

  > Using the `restart` command is the same as using the `stop` command followed by the `start` command. You can issue `restart` even if the service is currently stopped.

3. Log out of the command shell: `exit`

Remove port remaps

If you want to configure an endpoint for the Load Balancer service, and you want to use a port that has already been configured as the Mapped-To Port of a port remap, you must first remove the existing port remap, or the endpoint will not be effective. You must run a script on each Admin Node and Gateway Node that has conflicting remapped ports to remove all of the node's port remaps.

About this task

This procedure removes all port remaps. If you need to keep some of the remaps, contact technical support.

For information about configuring load balancer endpoints, see Configuring load balancer endpoints.

> If the port remap provides client access, reconfigure the client to use a different port as an load balancer endpoint to avoid loss of service. Otherwise, removing the port mapping will result in loss of client access and should be scheduled appropriately.

> This procedure does not work for a StorageGRID system deployed as a container on bare metal hosts. See the instructions for removing port remaps on bare metal hosts.

Steps

1. Log in to the node.
   a. Enter the following command: `ssh -p 8022 admin@node_IP`

      Port 8022 is the SSH port of the base OS, while port 22 is the SSH port of the container engine running StorageGRID.

   b. Enter the password listed in the `Passwords.txt` file.
   c. Enter the following command to switch to root: `su -`
   d. Enter the password listed in the `Passwords.txt` file.

      When you are logged in as root, the prompt changes from `$` to `#`.

2. Run the following script: `remove-port-remap.sh`

3. Reboot the node.

   Follow the instructions for rebooting a grid node.

4. Repeat these steps on each Admin Node and Gateway Node that has conflicting remapped ports.

Remove port remaps on bare metal hosts

If you want to configure an endpoint for the Load Balancer service, and you want to use a port that has already been configured as the Mapped-To Port of a port remap, you must first remove the existing port remap, or the endpoint will not be effective.

About this task

If you are running StorageGRID on bare metal hosts, follow this procedure instead of the general procedure for removing port remaps. You must edit the node configuration file for each Admin Node and Gateway Node that has conflicting remapped ports to remove all of the node's port remaps and restart the node.

> This procedure removes all port remaps. If you need to keep some of the remaps, contact technical support.

For information about configuring load balancer endpoints, see the instructions for administering StorageGRID.

> This procedure can result in temporary loss of service as nodes are restarted.

Steps

1. Log in to the host supporting the node. Log in as root or with an account that has sudo permission.
2. Run the following command to temporarily disable the node: `sudo storagegrid node stop node-name`
3. Using a text editor such as vim or pico, edit the node configuration file for the node.

   The node configuration file can be found at `/etc/storagegrid/nodes/node-name.conf`.

4. Locate the section of the node configuration file that contains the port remaps.

   See the last two lines in the following example.

```
ADMIN_NETWORK_CONFIG = STATIC
ADMIN_NETWORK_ESL = 10.0.0.0/8, 172.19.0.0/16, 172.21.0.0/16
ADMIN_NETWORK_GATEWAY = 10.224.0.1
ADMIN_NETWORK_IP = 10.224.5.140
ADMIN_NETWORK_MASK = 255.255.248.0
ADMIN_NETWORK_MTU = 1400
ADMIN_NETWORK_TARGET = eth1
ADMIN_NETWORK_TARGET_TYPE = Interface
BLOCK_DEVICE_VAR_LOCAL = /dev/sda2
CLIENT_NETWORK_CONFIG = STATIC
CLIENT_NETWORK_GATEWAY = 47.47.0.1
CLIENT_NETWORK_IP = 47.47.5.140
CLIENT_NETWORK_MASK = 255.255.248.0
CLIENT_NETWORK_MTU = 1400
CLIENT_NETWORK_TARGET = eth2
CLIENT_NETWORK_TARGET_TYPE = Interface
GRID_NETWORK_CONFIG = STATIC
GRID_NETWORK_GATEWAY = 192.168.0.1
GRID_NETWORK_IP = 192.168.5.140
GRID_NETWORK_MASK = 255.255.248.0
GRID_NETWORK_MTU = 1400
GRID_NETWORK_TARGET = eth0
GRID_NETWORK_TARGET_TYPE = Interface
NODE_TYPE = VM_API_Gateway
PORT_REMAP = client/tcp/8082/443
PORT_REMAP_INBOUND = client/tcp/8082/443
```

5. Edit the PORT_REMAP and PORT_REMAP_INBOUND entries to remove port remaps.

```
PORT_REMAP =
PORT_REMAP_INBOUND =
```

6. Run the following command to validate your changes to the node configuration file for the node: `sudo storagegrid node validate node-name`

   Address any errors or warnings before proceeding to the next step.

7. Run the following command to restart the node without port remaps: `sudo storagegrid node start node-name`
8. Log in to the node as admin using the password listed in the `Passwords.txt` file.
9. Verify that the services start correctly.
   a. View a list of the statuses of all services on the server:`sudo storagegrid-status`

      The status is updated automatically.

b. Wait until all services have a status of either Running or Verified.

c. Exit the status screen: `Ctrl+C`

10. Repeat these steps on each Admin Node and Gateway Node that has conflicting remapped ports.

Reboot grid node

Reboot grid node: Overview

You can reboot a grid node from the Grid Manager or from the node's command shell.

When you reboot a grid node, the node shuts down and restarts. All services are restarted automatically.

If you plan to reboot Storage Nodes, note the following:

- If an ILM rule specifies an ingest behavior of Dual commit or Balanced and it is not possible to immediately create all required copies, StorageGRID immediately commits any newly ingested objects to two Storage Nodes on the same site and evaluates ILM later. If you want to reboot two or more Storage Nodes on a given site, you might not be able to access these objects for the duration of the reboot.
- To ensure you can access all objects while a Storage Node is rebooting, stop ingesting objects at a site for approximately one hour before rebooting the node.

Reboot grid node from Grid Manager

Rebooting a grid node from the Grid Manager issues the `reboot` command on the target node.

Before you begin

- You are signed in to the Grid Manager using a supported web browser.
- You have the Maintenance or Root access permission.
- You have the provisioning passphrase.

Steps

1. Select **NODES**.
2. Select the grid node you want to reboot.
3. Select the **Tasks** tab.

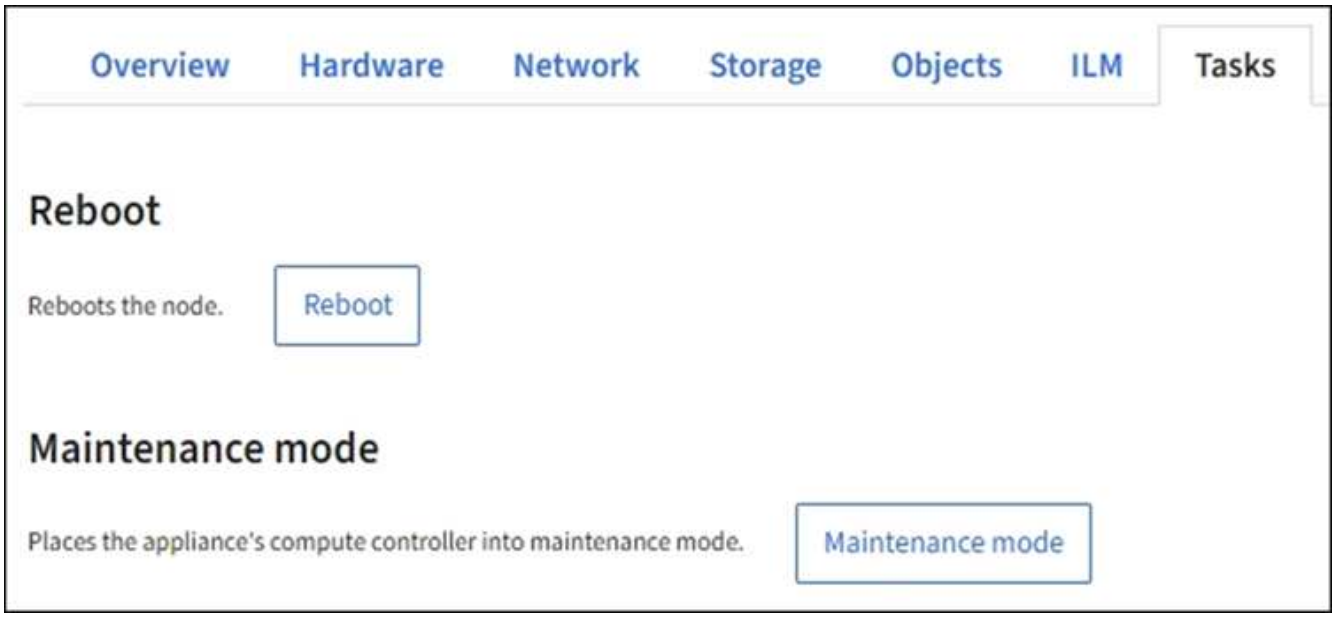

4. Select **Reboot**.

   A confirmation dialog box appears.

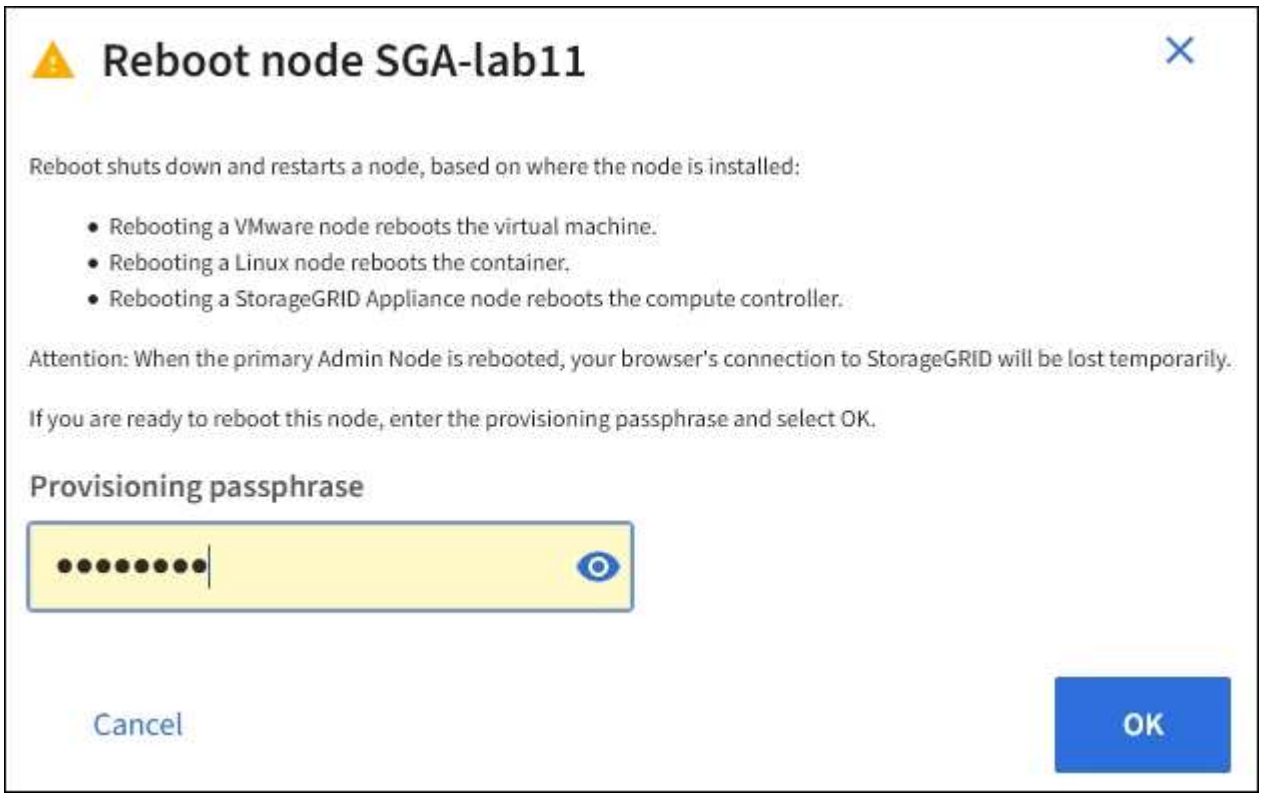

> If you are rebooting the primary Admin Node, the confirmation dialog box reminds you that your browser's connection to the Grid Manager will be lost temporarily when services are stopped.

5. Enter the provisioning passphrase, and select **OK**.
6. Wait for the node to reboot.

It might take some time for services to shut down.

When the node is rebooting, the gray icon (Administratively Down) appears on the left side of the **Nodes** page. When all services have started again and the node is successfully connected to the grid, the **Nodes** page should display a normal status (no icons to the left of the node name), indicating that no alerts are active and the node is connected to the grid.

Reboot grid node from command shell

If you need to monitor the reboot operation more closely or if you are unable to access the Grid Manager, you can log in to the grid node and run the Server Manager reboot command from the command shell.

Before you begin

You have the `Passwords.txt` file.

Steps

1. Log in to the grid node:
   a. Enter the following command: `ssh admin@grid_node_IP`
   b. Enter the password listed in the `Passwords.txt` file.
   c. Enter the following command to switch to root: `su -`
   d. Enter the password listed in the `Passwords.txt` file.

   When you are logged in as root, the prompt changes from `$` to `#`.
2. Optionally, stop services: `service servermanager stop`

   Stopping services is an optional, but recommended step. Services can take up to 15 minutes to shut down, and you might want to log in to the system remotely to monitor the shutdown process before you reboot the node in the next step.
3. Reboot the grid node: `reboot`
4. Log out of the command shell: `exit`

Shut down grid node

You can shut down a grid node from the node's command shell.

Before you begin

- You have the `Passwords.txt` file.

About this task

Before performing this procedure, review these considerations:

- In general, you should not shut down more than one node at a time to avoid disruptions.
- Don't shut down a node during a maintenance procedure unless explicitly instructed to do so by the documentation or by technical support.
- The shutdown process is based on where the node is installed, as follows:

- Shutting down a VMware node shuts down the virtual machine.
- Shutting down a Linux node shuts down the container.
- Shutting down a StorageGRID appliance node shuts down the compute controller.

- If you plan to shut down more than one Storage Node at a site, stop ingesting objects at that site for approximately one hour before shutting down the nodes.

  If any ILM rule uses the **Dual commit** ingest option (or if a rule uses the **Balanced** option and all required copies can't be created immediately), StorageGRID immediately commits any newly ingested objects to two Storage Nodes on the same site and evaluates ILM later. If more than one Storage Node at a site is shut down, you might not be able to access newly ingested objects for the duration of the shutdown. Write operations might also fail if too few Storage Nodes remain available at the site. See Manage objects with ILM.

Steps

1. Log in to the grid node:

   a. Enter the following command: `ssh admin@grid_node_IP`

   b. Enter the password listed in the `Passwords.txt` file.

   c. Enter the following command to switch to root: `su -`

   d. Enter the password listed in the `Passwords.txt` file.

   When you are logged in as root, the prompt changes from `$` to `#`.

2. Stop all services: `service servermanager stop`

   Services can take up to 15 minutes to shut down, and you might want to log in to the system remotely to monitor the shutdown process.

3. If the node is running on a VMware virtual machine or it is an appliance node, issue the shutdown command: `shutdown -h now`

   Perform this step regardless of the outcome of the `service servermanager stop` command.

   > After you issue the `shutdown -h now` command on an appliance node, you must power cycle the appliance to restart the node.

   For the appliance, this command shuts down the controller, but the appliance is still powered on. You must complete the next step.

4. If you are powering down an appliance node, follow the steps for your appliance.

SGF6112

a. Turn off the power to the appliance.

b. Wait for the blue power LED to turn off.

SG6000

a. Wait for the green Cache Active LED on the back of the storage controllers to turn off.

   This LED is on when cached data needs to be written to the drives. You must wait for this LED to turn off before you turn off power.

b. Turn off the power to the appliance, and wait for the blue power LED to turn off.

SG5700

a. Wait for the green Cache Active LED on the back of the storage controller to turn off.

   This LED is on when cached data needs to be written to the drives. You must wait for this LED to turn off before you turn off power.

b. Turn off the power to the appliance, and wait for all LED and seven-segment display activity to stop.

SG100 or SG1000

a. Turn off the power to the appliance.

b. Wait for the blue power LED to turn off.

Power down host

Before you power down a host, you must stop services on all grid nodes on that host.

Steps

1. Log in to the grid node:

   a. Enter the following command: `ssh admin@grid_node_IP`

   b. Enter the password listed in the `Passwords.txt` file.

   c. Enter the following command to switch to root: `su -`

   d. Enter the password listed in the `Passwords.txt` file.

   When you are logged in as root, the prompt changes from `$` to `#`.

2. Stop all services running on the node: `service servermanager stop`

   Services can take up to 15 minutes to shut down, and you might want to log in to the system remotely to monitor the shutdown process.

3. Repeat steps 1 and 2 for each node on the host.

4. If you have a Linux host:

   a. Log in to the host operating system.

b. Stop the node: `storagegrid node stop`

c. Shut down the host operating system.

5. If the node is running on a VMware virtual machine or it is an appliance node, issue the shutdown command: `shutdown -h now`

   Perform this step regardless of the outcome of the `service servermanager stop` command.

   > After you issue the `shutdown -h now` command on an appliance node, you must power cycle the appliance to restart the node.

   For the appliance, this command shuts down the controller, but the appliance is still powered on. You must complete the next step.

6. If you are powering down an appliance node, follow the steps for your appliance.

   SGF6112

   a. Turn off the power to the appliance.

   b. Wait for the blue power LED to turn off.

   SG6000

   a. Wait for the green Cache Active LED on the back of the storage controllers to turn off.

      This LED is on when cached data needs to be written to the drives. You must wait for this LED to turn off before you turn off power.

   b. Turn off the power to the appliance, and wait for the blue power LED to turn off.

   SG5700

   a. Wait for the green Cache Active LED on the back of the storage controller to turn off.

      This LED is on when cached data needs to be written to the drives. You must wait for this LED to turn off before you turn off power.

   b. Turn off the power to the appliance, and wait for all LED and seven-segment display activity to stop.

   SG100 or SG1000

   a. Turn off the power to the appliance.

   b. Wait for the blue power LED to turn off.

7. Log out of the command shell: `exit`

Related information

SGF6112 storage appliances

SG6000 storage appliances

SG5700 storage appliances

Power off and on all nodes in grid

You might need to shut down your entire StorageGRID system, for example, if you are moving a data center. These steps provide a high-level overview of the recommended sequence for performing a controlled shutdown and startup.

When you power off all nodes in a site or grid, you will not be able to access ingested objects while the Storage Nodes are offline.

Stop services and shut down grid nodes

Before you can power off a StorageGRID system, you must stop all services running on each grid node, and then shut down all VMware virtual machines, container engines, and StorageGRID appliances.

About this task

Stop services on Admin Nodes and Gateway Nodes first, and then stop services on Storage Nodes.

This approach allows you to use the primary Admin Node to monitor the status of the other grid nodes for as long as possible.

> If a single host includes more than one grid node, don't shut down the host until you have stopped all of the nodes on that host. If the host includes the primary Admin Node, shut down that host last.

> If required, you can migrate nodes from one Linux host to another to perform host maintenance without impacting the functionality or availability of your grid.

Steps

1. Stop all client applications from accessing the grid.
2. Log in to each Gateway Node:
   a. Enter the following command: `ssh admin@grid_node_IP`
   b. Enter the password listed in the `Passwords.txt` file.
   c. Enter the following command to switch to root: `su -`
   d. Enter the password listed in the `Passwords.txt` file.

   When you are logged in as root, the prompt changes from `$` to `#`.
3. Stop all services running on the node: `service servermanager stop`

   Services can take up to 15 minutes to shut down, and you might want to log in to the system remotely to monitor the shutdown process.
4. Repeat the previous two steps to stop the services on all Storage Nodes, Archive Nodes, and non-primary Admin Nodes.

   You can stop the services on these nodes in any order.

If you issue the `service servermanager stop` command to stop the services on an appliance Storage Node, you must power cycle the appliance to restart the node.

5. For the primary Admin Node, repeat the steps for logging into the node and stopping all services on the node.
6. For nodes that are running on Linux hosts:
   a. Log in to the host operating system.
   b. Stop the node: `storagegrid node stop`
   c. Shut down the host operating system.
7. For nodes that are running on VMware virtual machines and for appliance Storage Nodes, issue the shutdown command: `shutdown -h now`

   Perform this step regardless of the outcome of the `service servermanager stop` command.

   For the appliance, this command shuts down the compute controller, but the appliance is still powered on. You must complete the next step.

8. If you have appliance nodes, follow the steps for your appliance.

   SG100 or SG1000

   a. Turn off the power to the appliance.
   b. Wait for the blue power LED to turn off.

   SGF6112

   a. Turn off the power to the appliance.
   b. Wait for the blue power LED to turn off.

   SG6000

   a. Wait for the green Cache Active LED on the back of the storage controllers to turn off.

      This LED is on when cached data needs to be written to the drives. You must wait for this LED to turn off before you turn off power.

   b. Turn off the power to the appliance, and wait for the blue power LED to turn off.

   SG5700

   a. Wait for the green Cache Active LED on the back of the storage controller to turn off.

      This LED is on when cached data needs to be written to the drives. You must wait for this LED to turn off before you turn off power.

   b. Turn off the power to the appliance, and wait for all LED and seven-segment display activity to stop.

9. If required, log out of the command shell: `exit`

   The StorageGRID grid has now been shut down.

Start up grid nodes

If the entire grid has been shut down for more than 15 days, you must contact technical support before starting up any grid nodes. Don't attempt the recovery procedures that rebuild Cassandra data. Doing so might result in data loss.

If possible, power on the grid nodes in this order:

- Apply power to Admin Nodes first.
- Apply power to Gateway Nodes last.

If a host includes multiple grid nodes, the nodes will come back online automatically when you power on the host.

Steps

1. Power on the hosts for the primary Admin Node and any non-primary Admin Nodes.

   You will not be able to log in to the Admin Nodes until the Storage Nodes have been restarted.

2. Power on the hosts for all Archive Nodes and Storage Nodes.

   You can power on these nodes in any order.

3. Power on the hosts for all Gateway Nodes.
4. Sign in to the Grid Manager.
5. Select **NODES** and monitor the status of the grid nodes. Verify that there are no alert icons next to the node names.

Related information

- SG100 and SG1000 services appliances
- SG6000 storage appliances
- SG5700 storage appliances

Use a DoNotStart file

If you are performing various maintenance or configuration procedures under the direction of technical support, you might be asked to use a DoNotStart file to prevent services from starting when Server Manager is started or restarted.

You should add or remove a DoNotStart file only if technical support has directed you to do so.

To prevent a service from starting, place a DoNotStart file in the directory of the service you want to prevent from starting. At start-up, Server Manager looks for the DoNotStart file. If the file is present, the service (and any services dependent on it) is prevented from starting. When the DoNotStart file is removed, the previously stopped service will start on the next start or restart of Server Manager. Services aren't automatically started when the DoNotStart file is removed.

The most efficient way to prevent all services from restarting is to prevent the NTP service from starting. All services are dependent on the NTP service and can't run if the NTP service is not running.

Add DoNotStart file for service

You can prevent an individual service from starting by adding a DoNotStart file to that service's directory on a grid node.

Before you begin

You have the `Passwords.txt` file.

Steps

1. Log in to the grid node:
   a. Enter the following command: `ssh admin@grid_node_IP`
   b. Enter the password listed in the `Passwords.txt` file.
   c. Enter the following command to switch to root: `su -`
   d. Enter the password listed in the `Passwords.txt` file.

   When you are logged in as root, the prompt changes from `$` to `#`.

2. Add a DoNotStart file: `touch /etc/sv/service/DoNotStart`

   where `service` is the name of the service to be prevented from starting. For example,

   ```
   touch /etc/sv/ldr/DoNotStart
   ```

   A DoNotStart file is created. No file content is needed.

   When Server Manager or the grid node is restarted, Server Manager restarts, but the service does not.

3. Log out of the command shell: `exit`

Remove DoNotStart file for service

When you remove a DoNotStart file that is preventing a service from starting, you must start that service.

Before you begin

You have the `Passwords.txt` file.

Steps

1. Log in to the grid node:
   a. Enter the following command: `ssh admin@grid_node_IP`
   b. Enter the password listed in the `Passwords.txt` file.
   c. Enter the following command to switch to root: `su -`
   d. Enter the password listed in the `Passwords.txt` file.

   When you are logged in as root, the prompt changes from `$` to `#`.

2. Remove the DoNotStart file from the service directory: `rm /etc/sv/service/DoNotStart`

where `service` is the name of the service. For example,

```
rm /etc/sv/ldr/DoNotStart
```

3. Start the service: `service servicename start`
4. Log out of the command shell: `exit`

Troubleshoot Server Manager

If a problem arises when using Server Manager, check its log file.

Error messages related to Server Manager are captured in the Server Manager log file, which is located at: `/var/local/log/servermanager.log`

Check this file for error messages regarding failures. Escalate the issue to technical support if required. You might be asked to forward log files to technical support.

Service with an error state

If you detect that a service has entered an error state, attempt to restart the service.

Before you begin

You have the `Passwords.txt` file.

About this task

Server Manager monitors services and restarts any that have stopped unexpectedly. If a service fails, Server Manager attempts to restart it. If there are three failed attempts to start a service within five minutes, the service enters an error state. Server Manager does not attempt another restart.

Steps

1. Log in to the grid node:
   a. Enter the following command: `ssh admin@grid_node_IP`
   b. Enter the password listed in the `Passwords.txt` file.
   c. Enter the following command to switch to root: `su -`
   d. Enter the password listed in the `Passwords.txt` file.

   When you are logged in as root, the prompt changes from `$` to `#`.
2. Confirm the error state of the service: `service servicename status`

   For example:

   ```
   service ldr status
   ```

   If the service is in an error state, the following message is returned: `servicename in error state`. For example:

```
ldr in error state
```

If the service status is `disabled`, see the instructions for removing a DoNotStart file for a service.

3. Attempt to remove the error state by restarting the service: `service servicename restart`

   If the service fails to restart, contact technical support.

4. Log out of the command shell: `exit`

Maintain appliance configuration

Common procedures for node maintenance: Overview

Use these instructions to maintain your StorageGRID system.

About these instructions

These instructions describe procedures common to all nodes such as how to apply a software hotfix, recover grid nodes, recover a failed site, decommission grid nodes or an entire site, perform network maintenance, perform host-level and middleware maintenance procedures, and perform grid node procedures.

In these instructions, “Linux” refers to a Red Hat® Enterprise Linux®, Ubuntu®, CentOS, or Debian® deployment. Use the NetApp Interoperability Matrix Tool (IMT) to get a list of supported versions.

Before you begin

- You have a broad understanding of the StorageGRID system.
- You have reviewed your StorageGRID system’s topology and you understand the grid configuration.
- You understand that you must follow all instructions exactly and heed all warnings.
- You understand that maintenance procedures not described aren’t supported or require a services engagement.

Maintenance procedures for appliances

Specific maintenance procedures for each type of StorageGRID appliance are in the appliance maintenance sections:

- Maintain SGF6112 appliance
- Maintain SG6000 appliance
- Maintain SG5700 appliance
- Maintain SG100 and SG1000 appliances

Place appliance into maintenance mode

You must place the appliance into maintenance mode before performing specific maintenance procedures.

Before you begin

- You are signed in to the Grid Manager using a supported web browser.
- You have the Maintenance or Root access permission. For details, see the instructions for administering StorageGRID.

About this task

In rare instances, placing a StorageGRID appliance into maintenance mode might make the appliance unavailable for remote access.

> The admin account password and SSH host keys for a StorageGRID appliance in maintenance mode remain the same as they were when the appliance was in service.

Steps

1. From the Grid Manager, select **NODES**.
2. From the tree view of the Nodes page, select the appliance Storage Node.
3. Select **Tasks**.

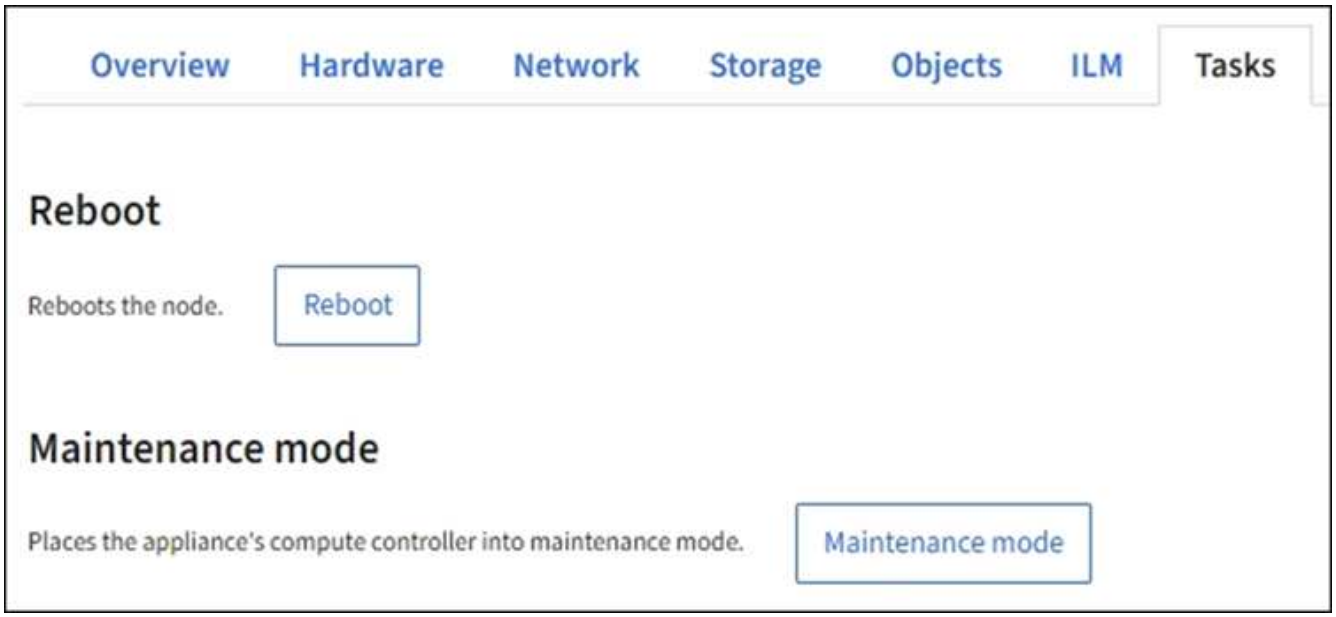

4. Select **Maintenance mode**. A confirmation dialog box appears.
5. Enter the provisioning passphrase, and select **OK**.

   A progress bar and a series of messages, including "Request Sent," "Stopping StorageGRID," and "Rebooting," indicate that the appliance is completing the steps for entering maintenance mode.

   When the appliance is in maintenance mode, a confirmation message lists the URLs you can use to access the StorageGRID Appliance Installer.

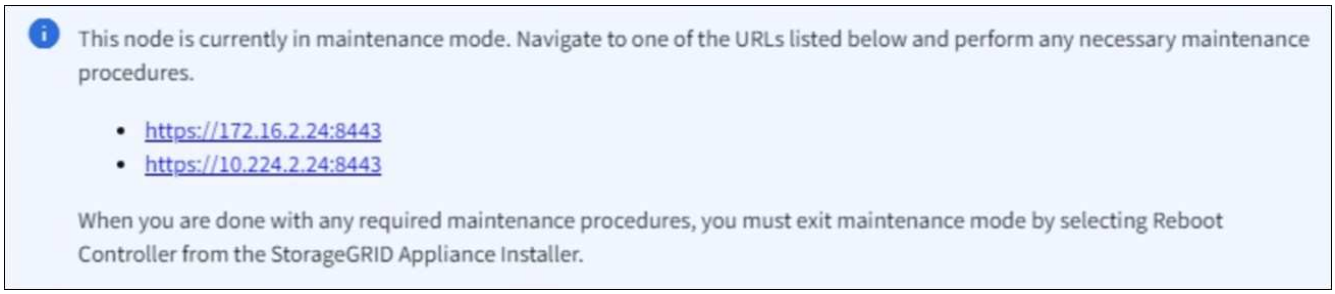

6. To access the StorageGRID Appliance Installer, browse to any of the URLs displayed.

   If possible, use the URL containing the IP address of the appliance's Admin Network port.

   > If you have a direct connection to the appliance's management port, use `https://169.254.0.1:8443` to access the StorageGRID Appliance Installer page.

7. From the StorageGRID Appliance Installer, confirm that the appliance is in maintenance mode.

8. Perform any required maintenance tasks.
9. After completing maintenance tasks, exit maintenance mode and resume normal node operation. From the StorageGRID Appliance Installer, select **Advanced** > **Reboot Controller**, and then select **Reboot into StorageGRID**.

   It can take up to 20 minutes for the appliance to reboot and rejoin the grid. To confirm that the reboot is complete and that the node has rejoined the grid, go back to the Grid Manager. The **NODES** page should display a normal status (green check mark icon to the left of the node name) for the appliance node, indicating that no alerts are active and the node is connected to the grid.

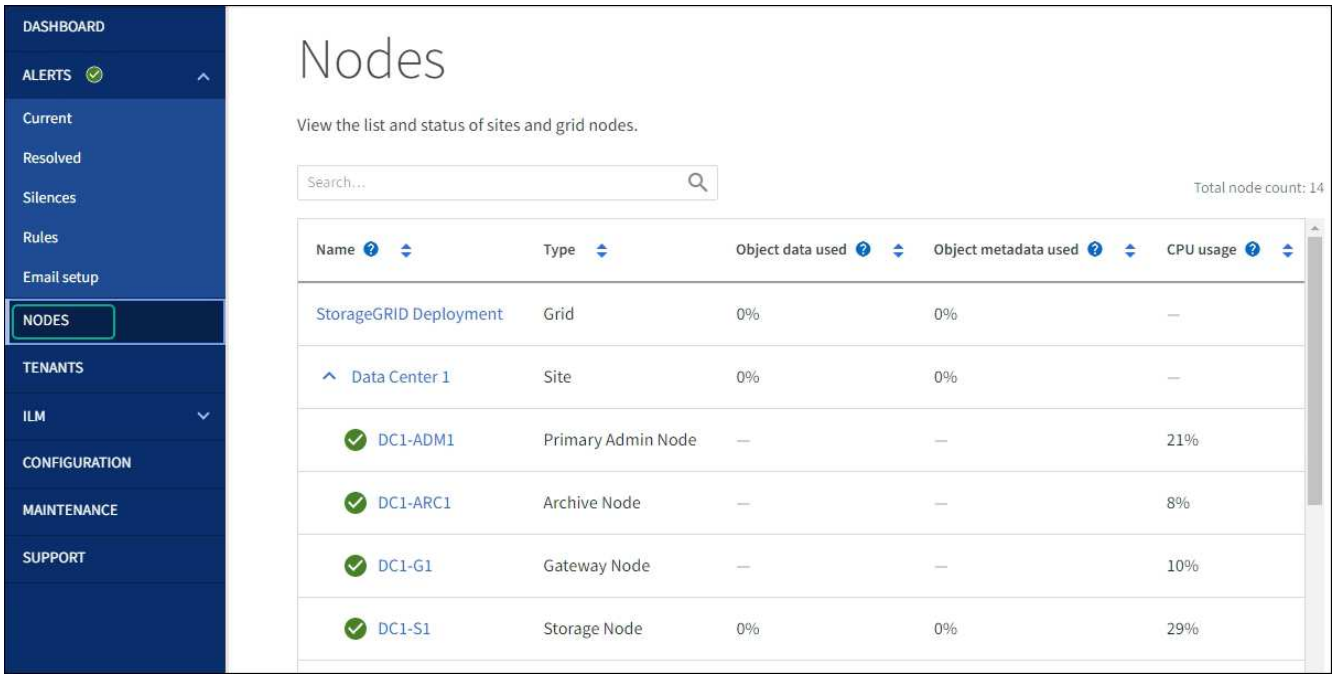

Change MTU setting

You can change the MTU setting that you assigned when you configured IP addresses for the appliance node.

About this task

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

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

To change the MTU setting without rebooting the appliance node, use the Change IP tool.

If the Client or Admin Network was not configured in the StorageGRID Appliance Installer during the initial installation, change the MTU setting using maintenance mode.

Change the MTU setting using the Change IP tool

Before you begin

You have the `Passwords.txt` file to use the Change IP tool.

Steps

Access the Change IP tool and update the MTU settings as described in Change node network configuration.

Change the MTU setting using maintenance mode

Change the MTU setting using maintenance mode if you are unable to access these settings using the Change IP tool.

Before you begin

The appliance has been placed maintenance mode.

Steps

1. From the StorageGRID Appliance Installer, select **Configure Networking** > **IP Configuration**.
2. Make the desired changes to the MTU settings for the Grid Network, Admin Network, and Client Network.
3. When you are satisfied with the settings, select **Save**.
4. If this procedure completed successfully and you have additional procedures to perform while the node is in maintenance mode, perform them now. When you are done, or if you experienced any failures and want to start over, select **Advanced** > **Reboot Controller**, and then select one of these options:
   - Select **Reboot into StorageGRID**
   - Select **Reboot into Maintenance Mode** to reboot the controller with the node remaining in maintenance mode. Select this option if you experienced any failures during the procedure and want to start over. After the node finishes rebooting into maintenance mode, restart from the appropriate step in the procedure that failed.

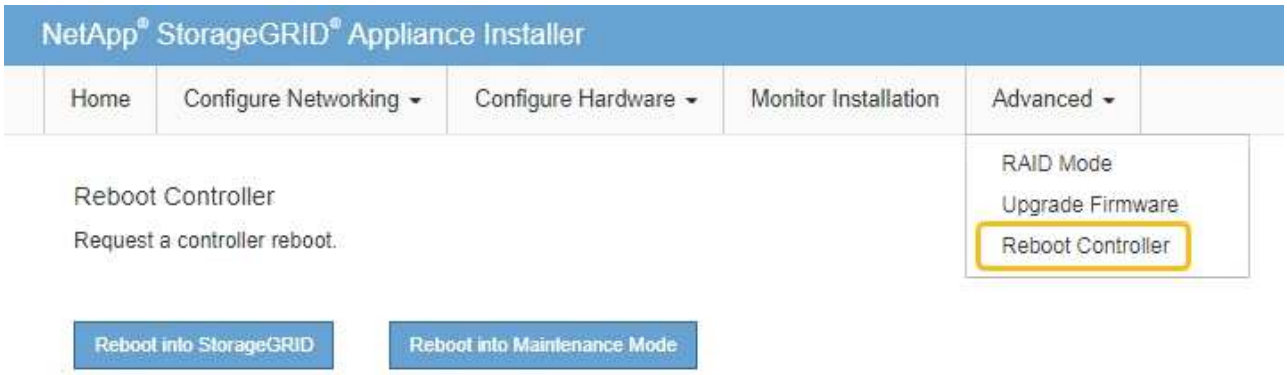

It can take up to 20 minutes for the appliance to reboot and rejoin the grid. To confirm that the reboot is complete and that the node has rejoined the grid, go back to the Grid Manager. The **NODES** page should display a normal status (green check mark icon to the left of the node name) for the appliance node, indicating that no alerts are active and the node is connected to the grid.

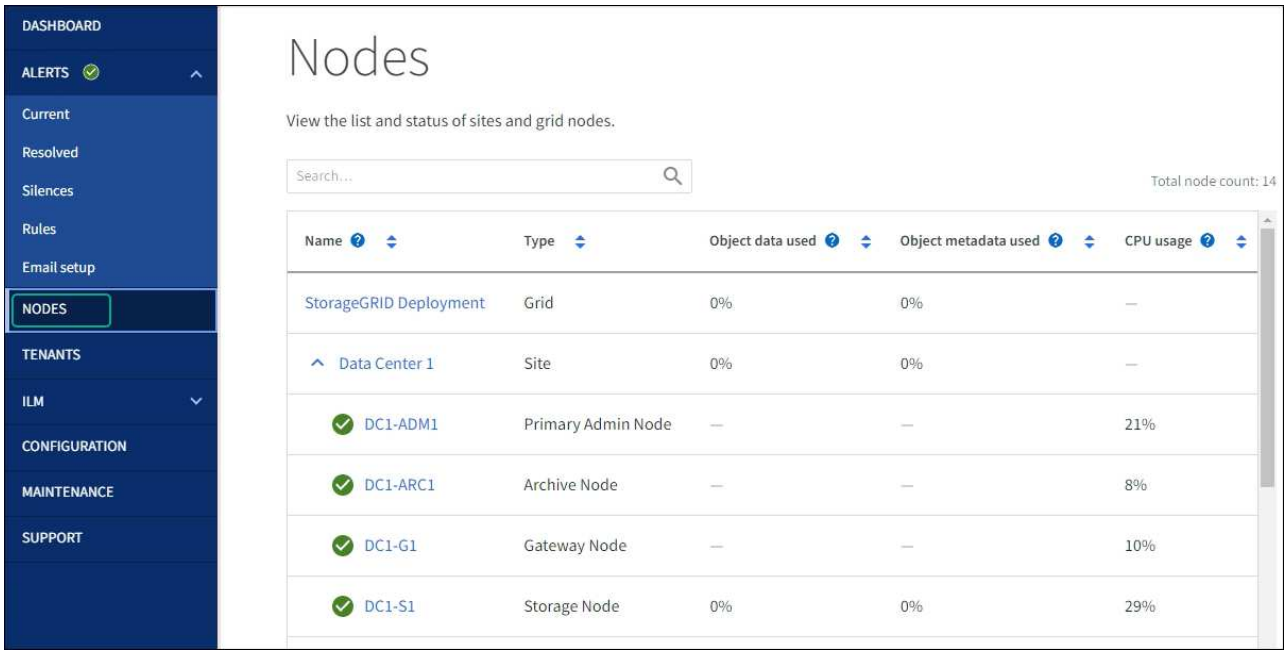

Check DNS server configuration

You can check and temporarily change the DNS servers that are currently in use by this appliance node.

Before you begin

The appliance has been placed maintenance mode.

About this task

You might need to change the DNS server settings if an encrypted appliance can't connect to the key management server (KMS) or KMS cluster because the hostname for the KMS was specified as a domain name instead of an IP address. Any changes that you make to the DNS settings for the appliance are temporary and are lost when you exit maintenance mode. To make these changes permanent, specify the DNS servers in Grid Manager (**MAINTENANCE** > **Network** > **DNS servers**).

- Temporary changes to the DNS configuration are necessary only for node-encrypted appliances where the KMS server is defined using a fully qualified domain name, instead of an IP address, for the hostname.
- When a node-encrypted appliance connects to a KMS using a domain name, it must connect to one of the DNS servers defined for the grid. One of these DNS servers then translates the domain name into an IP address.
- If the node can't reach a DNS server for the grid, or if you changed the grid-wide DNS settings when a node-encrypted appliance node was offline, the node is unable to connect to the KMS. Encrypted data on the appliance can't be decrypted until the DNS issue is resolved.

To resolve a DNS issue preventing KMS connection, specify the IP address of one or more DNS servers in the StorageGRID Appliance Installer. These temporary DNS settings allow the appliance to connect to the KMS and decrypt data on the node.

For example, if the DNS server for the grid changes while an encrypted node was offline, the node will not be able to reach the KMS when it comes back online, because it is still using the previous DNS values. Entering the new DNS server IP address in the StorageGRID Appliance Installer allows a temporary KMS connection to decrypt the node data.

Steps

1. From the StorageGRID Appliance Installer, select **Configure Networking** > **DNS Configuration**.
2. Verify that the DNS servers specified are correct.

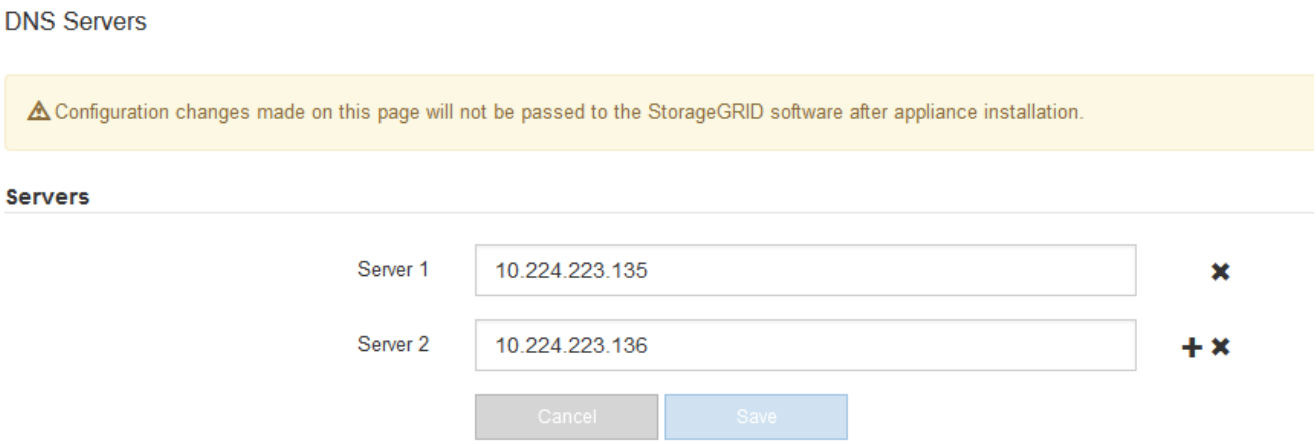

3. If required, change the DNS servers.

> Changes made to the DNS settings are temporary and are lost when you exit maintenance mode.

4. When you are satisfied with the temporary DNS settings, select **Save**.

   The node uses the DNS server settings specified on this page to reconnect to the KMS, allowing data on the node to be decrypted.

5. After node data is decrypted, reboot the node. From the StorageGRID Appliance Installer, select **Advanced** > **Reboot Controller**, and then select one of these options:
   - Select **Reboot into StorageGRID** to reboot the controller with the node rejoining the grid. Select this option if you are done working in maintenance mode and are ready to return the node to normal operation.
   - Select **Reboot into Maintenance Mode** to reboot the controller with the node remaining in maintenance mode. (This option is available only when the controller is in maintenance mode.) Select this option if there are additional maintenance operations you need to perform on the node before rejoining the grid.

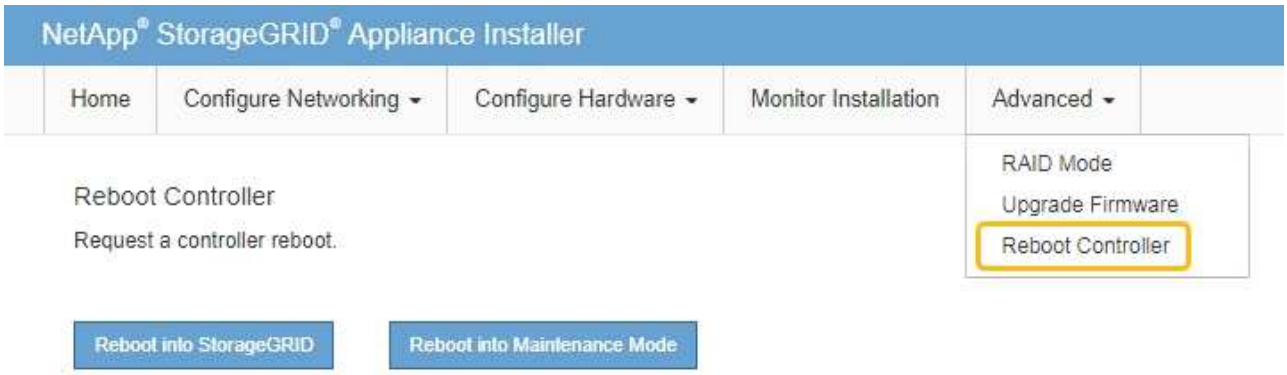

When the node reboots and rejoins the grid, it uses the system-wide DNS servers listed in the Grid Manager. After rejoining the grid, the appliance will no longer use the temporary DNS servers specified in the StorageGRID Appliance Installer while the appliance was in maintenance mode.

It can take up to 20 minutes for the appliance to reboot and rejoin the grid. To confirm that the reboot is complete and that the node has rejoined the grid, go back to the Grid Manager. The **NODES** page should display a normal status (green check mark icon to the left of the node name) for the appliance node, indicating that no alerts are active and the node is connected to the grid.

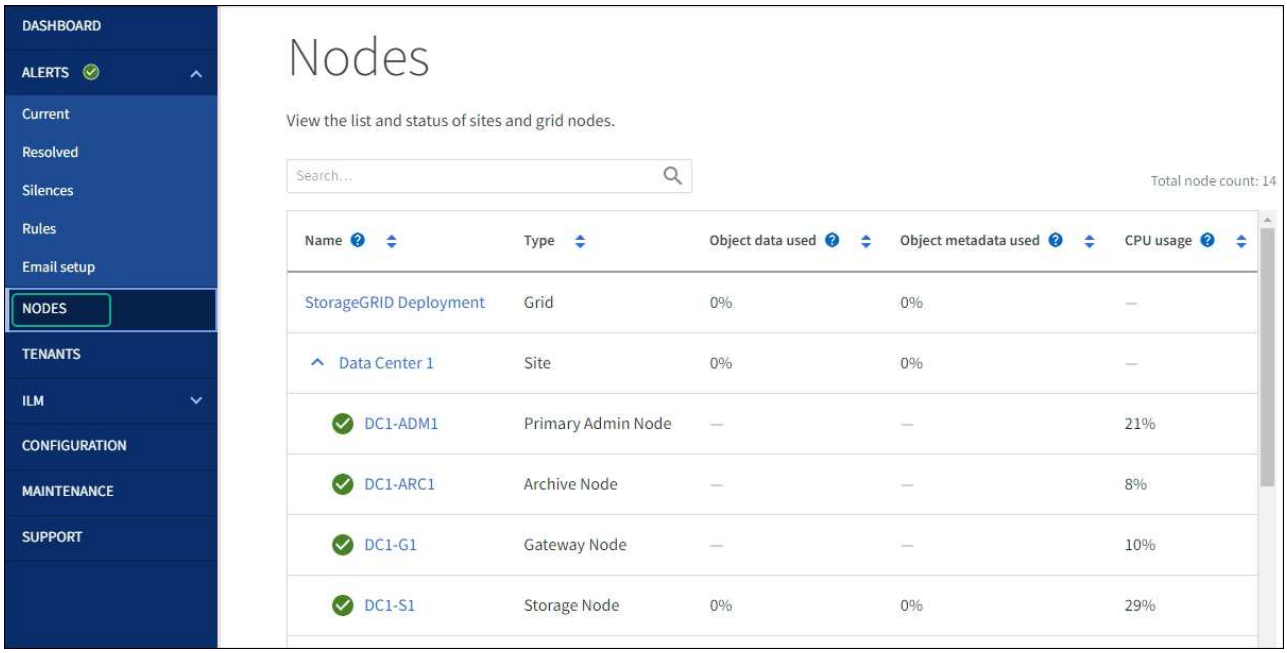

Update MAC address references

In some cases you might need to update MAC address references after the replacement of an appliance.

About this task

If any of the network interfaces on an appliance you are replacing are configured for DHCP, you might need to

update the permanent DHCP lease assignments on the DHCP servers to reference the MAC addresses of the replacement appliance. The update ensures the replacement appliance is assigned the expected IP addresses.

Steps

1. Locate the label on the front of the appliance. The label lists the MAC address for the BMC management port of the appliance.
2. To determine the MAC address for the Admin Network port, you must add **2** to the hexadecimal number on the label.

   For example, if the MAC address on the label ends in **09**, the MAC address for the Admin Port would end in **0B**. If the MAC address on the label ends in **(*y*)FF**, the MAC address for the Admin Port would end in **(*y*+1)01**.

   You can easily make this calculation by opening Calculator in Windows, setting it to Programmer mode, selecting Hex, typing the MAC address, then typing **+ 2 =**.

3. Ask your network administrator to associate the DNS/network and IP address for the appliance you removed with the MAC address for the replacement appliance.

   > ⓘ You must ensure that all IP addresses for the original appliance have been updated before you apply power to the replacement appliance. Otherwise, the appliance will obtain new DHCP IP addresses when it boots up and might not be able to reconnect to StorageGRID. This step applies to all StorageGRID networks that are attached to the appliance.

   > ⓘ If the original appliance used static IP address, the new appliance will automatically adopt the IP addresses of the appliance you removed.

Monitor node encryption in maintenance mode

If you enabled node encryption for the appliance during installation, you can monitor the node-encryption status of each appliance node, including the node-encryption state and key management server (KMS) details.

See Configure key management servers for information about implementing KMS for StorageGRID appliances.

Before you begin

- You enabled node encryption for the appliance during installation. You can't enable node encryption after the appliance is installed.
- You have placed the appliance into maintenance mode.

Steps

1. From the StorageGRID Appliance Installer, select **Configure Hardware** > **Node Encryption**.

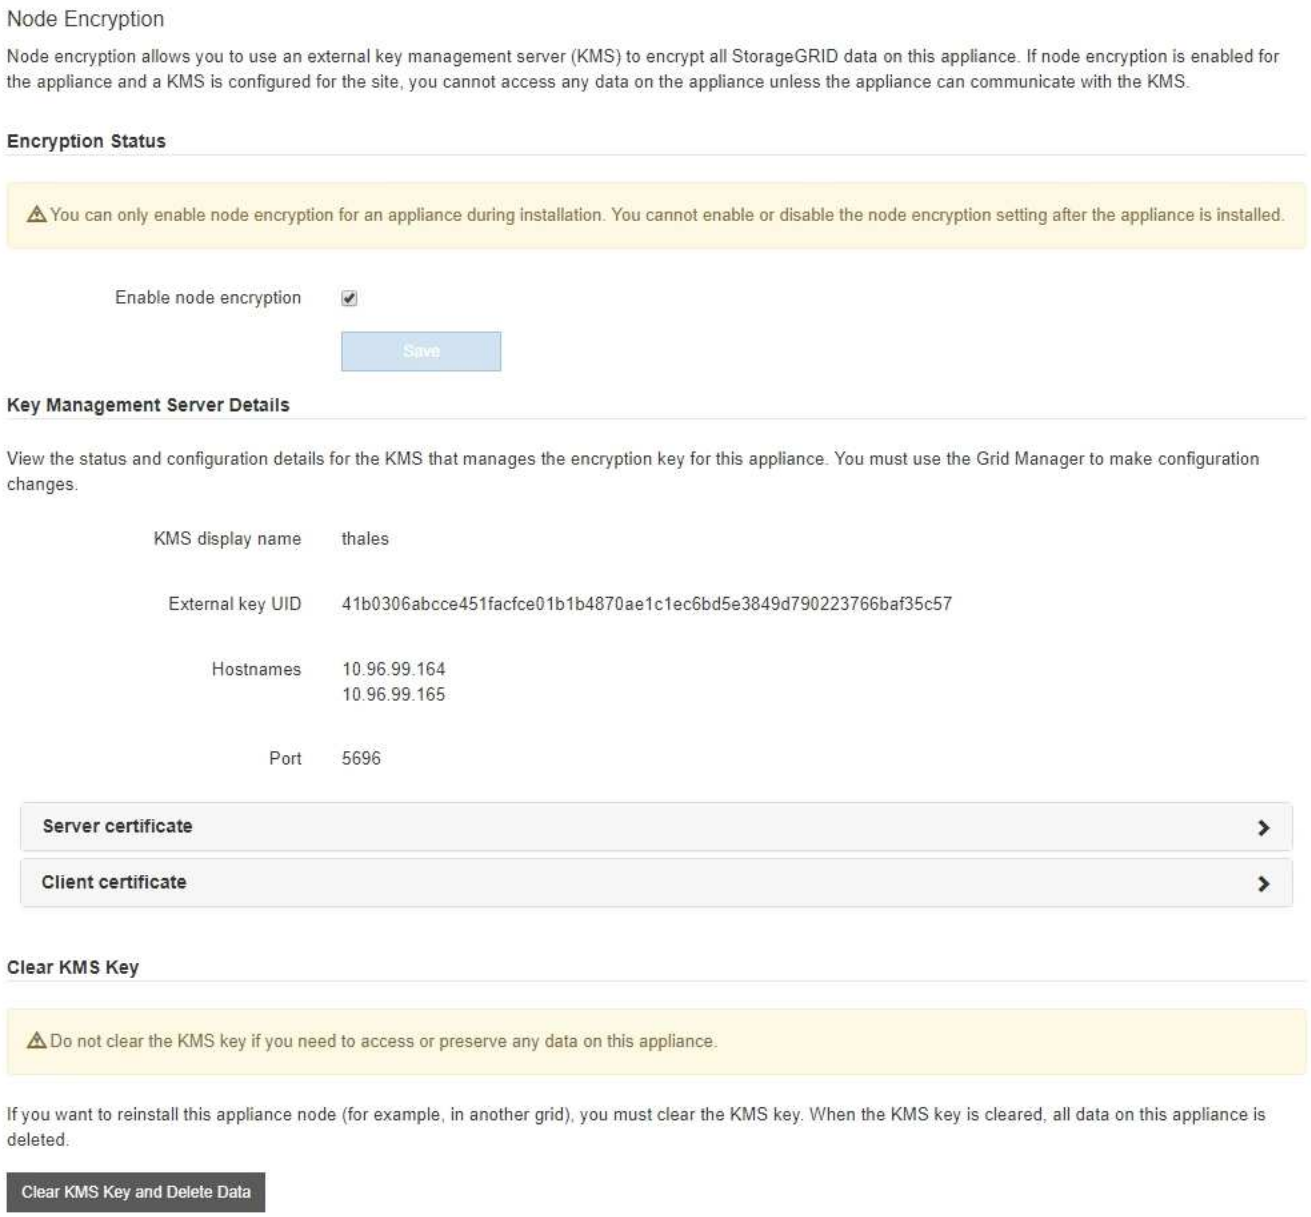

The Node Encryption page includes three sections:

- Encryption Status shows whether node encryption is enabled or disabled for the appliance.
- Key Management Server Details shows information about the KMS being used to encrypt the appliance. You can expand the server and client certificate sections to view certificate details and status.
  - To address issues with the certificates themselves, such as renewing expired certificates, see the instructions for configuring KMS.
  - If there are unexpected problems connecting to KMS hosts, verify that the DNS servers are correct and that appliance networking is correctly configured.
  - If you are unable to resolve your certificate issues, contact technical support.
- Clear KMS Key disables node encryption for the appliance, removes the association between the appliance and the key management server that was configured for the StorageGRID site, and deletes all data from the appliance. You must clear the KMS key before you can install the appliance into another StorageGRID system.

Clearing the KMS configuration deletes data from the appliance, rendering it permanently inaccessible. This data is not recoverable.

2. When you are done checking node-encryption status, reboot the node. From the StorageGRID Appliance Installer, select **Advanced** > **Reboot Controller**, and then select one of these options:
   - Select **Reboot into StorageGRID** to reboot the controller with the node rejoining the grid. Select this option if you are done working in maintenance mode and are ready to return the node to normal operation.
   - Select **Reboot into Maintenance Mode** to reboot the controller with the node remaining in maintenance mode. (This option is available only when the controller is in maintenance mode.) Select this option if there are additional maintenance operations you need to perform on the node before rejoining the grid.

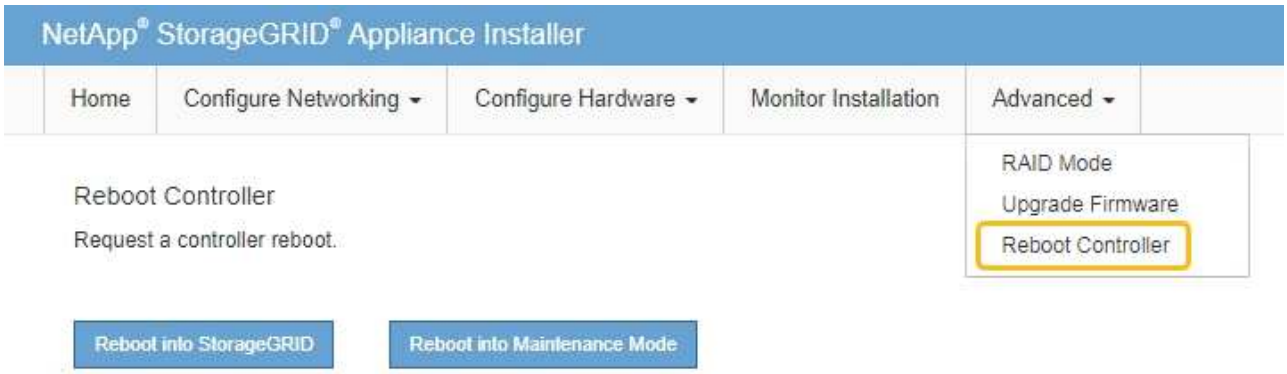

It can take up to 20 minutes for the appliance to reboot and rejoin the grid. To confirm that the reboot is complete and that the node has rejoined the grid, go back to the Grid Manager. The **NODES** page should display a normal status (green check mark icon to the left of the node name) for the appliance node, indicating that no alerts are active and the node is connected to the grid.

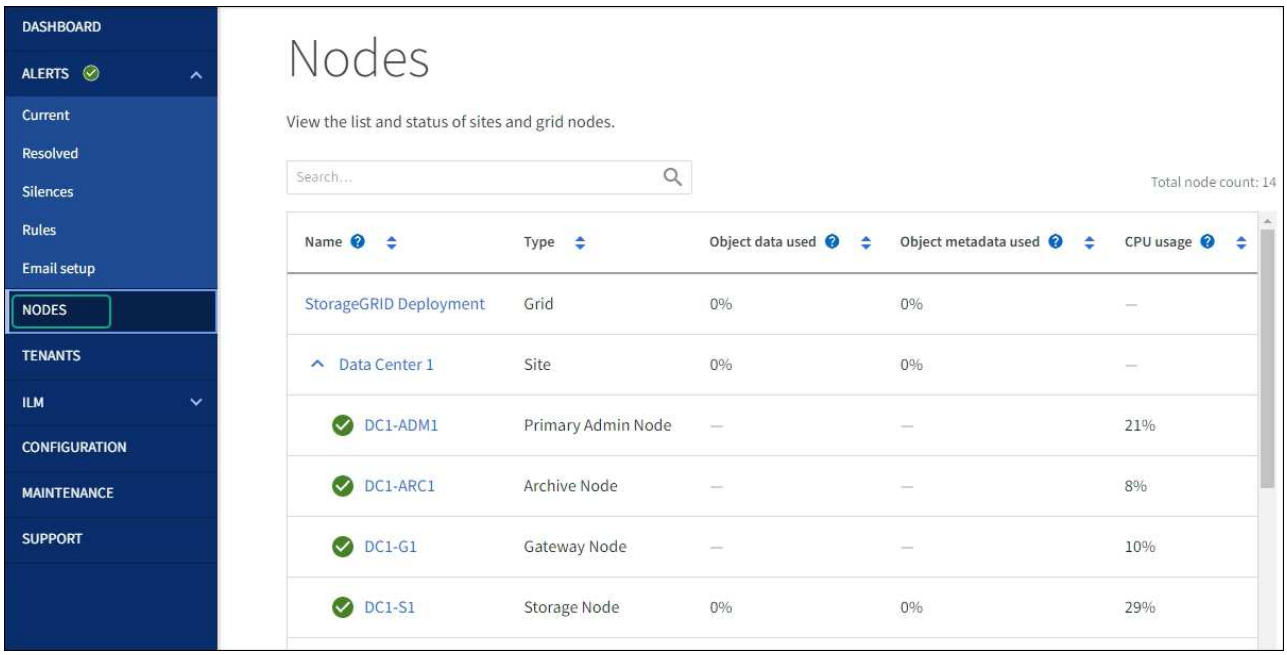

Clear key management server configuration

Clearing the key management server (KMS) configuration disables node encryption on your appliance. After clearing the KMS configuration, the data on your appliance is permanently deleted and is no longer accessible. This data is not recoverable.

Before you begin

If you need to preserve data on the appliance, you must either perform a node decommission procedure or clone the node before you clear the KMS configuration.

> When KMS is cleared, data on the appliance will be permanently deleted and no longer accessible. This data is not recoverable.

Decommission the node to move any data it contains to other nodes in StorageGRID.

About this task

Clearing the appliance KMS configuration disables node encryption, removing the association between the appliance node and the KMS configuration for the StorageGRID site. Data on the appliance is then deleted and the appliance is left in a pre-install state. This process can't be reversed.

You must clear the KMS configuration:

- Before you can install the appliance into another StorageGRID system, that does not use a KMS or that uses a different KMS.

  > Don't clear the KMS configuration if you plan to reinstall an appliance node in a StorageGRID system that uses the same KMS key.

- Before you can recover and reinstall a node where the KMS configuration was lost and the KMS key is not recoverable.
- Before returning any appliance that was previously in use at your site.
- After decommissioning a appliance that had node encryption enabled.

> Decommission the appliance before clearing KMS to move its data to other nodes in your StorageGRID system. Clearing KMS before decommissioning the appliance will result in data loss and might render the appliance inoperable.

Steps

1. Open a browser, and enter one of the IP addresses for the appliance's compute controller.

   `https://Controller_IP:8443`

   `Controller_IP` is the IP address of the compute controller (not the storage controller) on any of the three StorageGRID networks.

   The StorageGRID Appliance Installer Home page appears.

2. Select **Configure Hardware** > **Node Encryption**.

   > If the KMS configuration is cleared, data on the appliance will be permanently deleted. This data is not recoverable.

3. At the bottom of the window, select **Clear KMS Key and Delete Data**.
4. If you are sure that you want to clear the KMS configuration, type **`clear`** in the warning dialog box and select **Clear KMS Key and Delete Data**.

   The KMS encryption key and all data are deleted from the node, and the appliance reboots. This can take up to 20 minutes.

5. Open a browser, and enter one of the IP addresses for the appliance's compute controller.
   `https://Controller_IP:8443`

   `Controller_IP` is the IP address of the compute controller (not the storage controller) on any of the three StorageGRID networks.

   The StorageGRID Appliance Installer Home page appears.

6. Select **Configure Hardware** > **Node Encryption**.
7. Verify that node encryption is disabled and that the key and certificate information in **Key Management Server Details** and the **Clear KMS Key and Delete Data** control are removed from the window.

   Node encryption can't be reenabled on the appliance until it is reinstalled in a grid.

After you finish

After the appliance reboots and you have verified that KMS has been cleared and that the appliance is in a pre-install state, you can physically remove the appliance from your StorageGRID system. See the instructions for preparing the appliance for reinstallation.

Appliance node cloning

Appliance node cloning: Overview

You can clone an appliance node in StorageGRID to use an appliance of newer design or increased capabilities. Cloning transfers all information about the existing node to the new appliance, provides a hardware-upgrade process that is easy to perform, and provides an alternative to decommissioning and expansion for replacing appliances.

Appliance node cloning lets you easily replace an existing appliance node (source) in your grid with a compatible appliance (target) that is part of the same logical StorageGRID site. The process transfers all data to the new appliance, placing it in service to replace the old appliance node and leaving the old appliance in a pre-install state.

Why clone an appliance node?

You can clone an appliance node if you need to:

- Replace appliances that are reaching end-of-life.
- Upgrade existing nodes to take advantage of improved appliance technology.
- Increase grid storage capacity without changing the number of Storage Nodes in your StorageGRID system.
- Improve storage efficiency, such as by changing the RAID mode from DDP-8 to DDP-16, or to RAID-6.
- Efficiently implement node encryption to allow the use of external key management servers (KMS).

Which StorageGRID network is used?

Cloning transfers data from the source node directly to the target appliance over any of the three StorageGRID networks. The Grid Network is typically used, but you can also use the Admin Network or the Client Network if the source appliance is connected to these networks. Choose the network to use for cloning traffic that provides the best data-transfer performance without degrading StorageGRID network performance or data availability.

When you install the replacement appliance, you must specify temporary IP addresses for StorageGRID connection and data transfer. Because the replacement appliance will be part of the same networks as the appliance node it replaces, you must specify temporary IP addresses for each of these networks on the replacement appliance.

Target appliance compatibility

Replacement appliances must be the same type as the source node they are replacing and both must be part of the same logical StorageGRID site.

- A replacement services appliance can be different than the Admin Node or Gateway Node it is replacing.
  - You can clone an SG100 source node appliance to an SG1000 services target appliance to give the Admin Node or Gateway Node greater capability.
  - You can clone an SG1000 source node appliance to an SG100 services target appliance to redeploy the SG1000 for a more demanding application.

    For example, if an SG1000 source node appliance is being used as an Admin Node and you want to use it as a dedicated load-balancing node.

  - Replacing an SG1000 source node appliance with an SG100 services target appliance reduces the maximum speed of the network ports from 100-GbE to 25-GbE.
  - The SG100 and SG1000 appliances have different network connectors. Changing the appliance type might require replacing the cables or SFP modules.
- A replacement storage appliance must have greater capacity than the Storage Node it is replacing.
  - If the target storage appliance has the same number of drives as the source node, the drives in the target appliance must have greater capacity (in TB).
  - If you plan to use the same RAID mode on the target node as was used on the source node, or a less storage efficient RAID mode (for example, switching from RAID 6 to DDP), the drives in the target appliance must be larger (in TB) than the drives in the source appliance.
  - If the number of standard drives installed in a target storage appliance is less than the number of drives in the source node, due to installation of solid-state drives (SSDs), the overall storage capacity of the standard drives in the target appliance (in TB) must exceed the total functional drive capacity of all drives in the source Storage Node.

    For example, when cloning an SG5760 source Storage Node appliance with 60 drives to an SG6060 target appliance with 58 standard drives, larger drives should be installed in the SG6060 target appliance before cloning to maintain storage capacity. (The two drive slots containing SSDs in the target appliance aren’t included in the total appliance-storage capacity.)

    However, if a 60-drive SG5760 source node appliance is configured with SANtricity Dynamic Disk Pools DDP-8, configuring a 58-drive same-size-drive SG6060 target appliance with DDP-16 might make the SG6060 appliance a valid clone target due to its improved storage efficiency.

    You can view information about the current RAID mode of the source appliance node on the **NODES**

page in Grid Manager. Select the **Storage** tab for the appliance.

- The number of volumes in a target storage appliance must be greater than or equal to the number of volumes in the source node. You cannot clone a source node with 16 object store volumes (rangedb) to a target storage appliance with 12 object store volumes even if the target appliance has larger capacity than the source node. Most storage appliances have 16 object store volumes, except the SGF6112 storage appliance that has only 12 object store volumes.

What information is not cloned?

The following appliance configurations don't transfer to the replacement appliance during cloning. You must configure them during initial set up of the replacement appliance.

- BMC interface
- Network links
- Node encryption status
- SANtricity System Manager (for Storage Nodes)
- RAID mode (for Storage Nodes)

What issues prevent cloning?

If any of the following issues are encountered while cloning, the cloning process halts and an error message is generated:

- Wrong network configuration
- Lack of connectivity between the source and target appliances
- Source and target appliance incompatibility
- For Storage Nodes, a replacement appliance of insufficient capacity

You must resolve each issue for cloning to continue.

Considerations and requirements for appliance node cloning

Before cloning an appliance node, you must understand the considerations and requirements.

Hardware requirements for the replacement appliance

Ensure that the replacement appliance meets the following criteria:

- The source node (appliance being replaced) and the target (new) appliance must be the same type of appliance:
  - You can only clone an Admin Node appliance or a Gateway Node appliance to a new services appliance.
  - You can only clone a Storage Node appliance to a new storage appliance.
- For Admin Node or Gateway Node appliances, the source node appliance and the target appliance don't need to be the same type of appliance; however, changing the appliance type might require replacing the cables or SFP modules.

  For example, you can replace a SG1000 node appliance with a SG100 or replace a SG100 appliance with

a SG1000 appliance.

- For Storage Node appliances, the source node appliance and the target appliance don't need to be the same type of appliance; however:
  - The target appliance must have greater storage capacity than the source appliance.

    For example, you can replace a SG5700 node appliance with a SG6000 appliance.

  - The target appliance must have an equal or greater number of object storage volumes than the source appliance.

    For example, you cannot replace a SG6000 node appliance (16 object store volumes) with a SGF6112 appliance (12 object store volumes).

Contact your StorageGRID sales representative for help choosing compatible replacement appliances to clone specific appliance nodes in your StorageGRID installation.

Prepare to clone an appliance node

You must have the following information before you clone an appliance node:

- Obtain a temporary IP address for the Grid Network from your network administrator for use with the target appliance during initial installation. If the source node belongs to an Admin Network or Client Network, obtain temporary IP addresses for these networks.

  Temporary IP addresses are normally on the same subnet as the source node appliance being cloned and aren't needed after cloning completes. The source and target appliances must both connect to the primary Admin Node of your StorageGRID to establish a cloning connection.

- Determine which network to use for cloning data-transfer traffic that provides the best data-transfer performance without degrading StorageGRID network performance or data availability.

  > Using the 1-GbE Admin Network for clone data transfer results in slower cloning.

- Determine if node encryption using a key management server (KMS) will be used on the target appliance, so that you can enable node encryption during initial target appliance installation before cloning. You can check if node encryption is enabled on the source appliance node as described in the enabling node encryption.

  The source node and target appliance can have different node-encryption settings. Data decryption and encryption is performed automatically during data transfer and when the target node restarts and joins the grid.

- Determine if the RAID mode on the target appliance should be changed from its default setting, so you can specify this information during initial target appliance installation before cloning. You can view information about the current RAID mode of the source appliance node on the **NODES** page in Grid Manager. Select the **Storage** tab for the appliance.

  The source node and target appliance can have different RAID settings.

- Plan for sufficient time to complete the node cloning process. Several days might be required to transfer data from an operational Storage Node to a target appliance. Schedule cloning at a time that minimizes the impact to your business.

- You should only clone one appliance node at a time. Cloning can prevent you from performing other StorageGRID maintenance functions at the same time.
- After you have cloned an appliance node, you can use the source appliance that was returned to a pre-install state as the target to clone another compatible node appliance.

Clone appliance node

The cloning process might take several days to transfer data between the source node (appliance being replaced) and the target (new) appliance.

Before you begin

- You have installed the compatible target appliance into a cabinet or rack, connected all cables, and applied power.
- You have verified that the StorageGRID Appliance Installer version on the replacement appliance matches the software version of your StorageGRID system, upgrading the StorageGRID Appliance Installer firmware, if necessary.
- You have configured the target appliance, including configuring StorageGRID connections, SANtricity System Manager (storage appliances only), and the BMC interface.
  - When configuring StorageGRID connections, use the temporary IP addresses.
  - When configuring network links, use the final link configuration.

> Leave the StorageGRID Appliance Installer open after you complete initial target appliance configuration. You will return to the target appliance's installer page after you start the node cloning process.

- You have optionally enabled node encryption for the target appliance.
- You have optionally set the RAID mode for the target appliance (storage appliances only).
- You have reviewed the considerations and requirements for appliance node cloning.

You should clone only one appliance node at a time to maintain StorageGRID network performance and data availability.

Steps

1. Place the source node you are cloning into maintenance mode.
2. From the StorageGRID Appliance Installer on the source node, in the Installation section of the Home page, select **Enable Cloning**.

   The Primary Admin Node connection section is replaced with the Clone target node connection section.

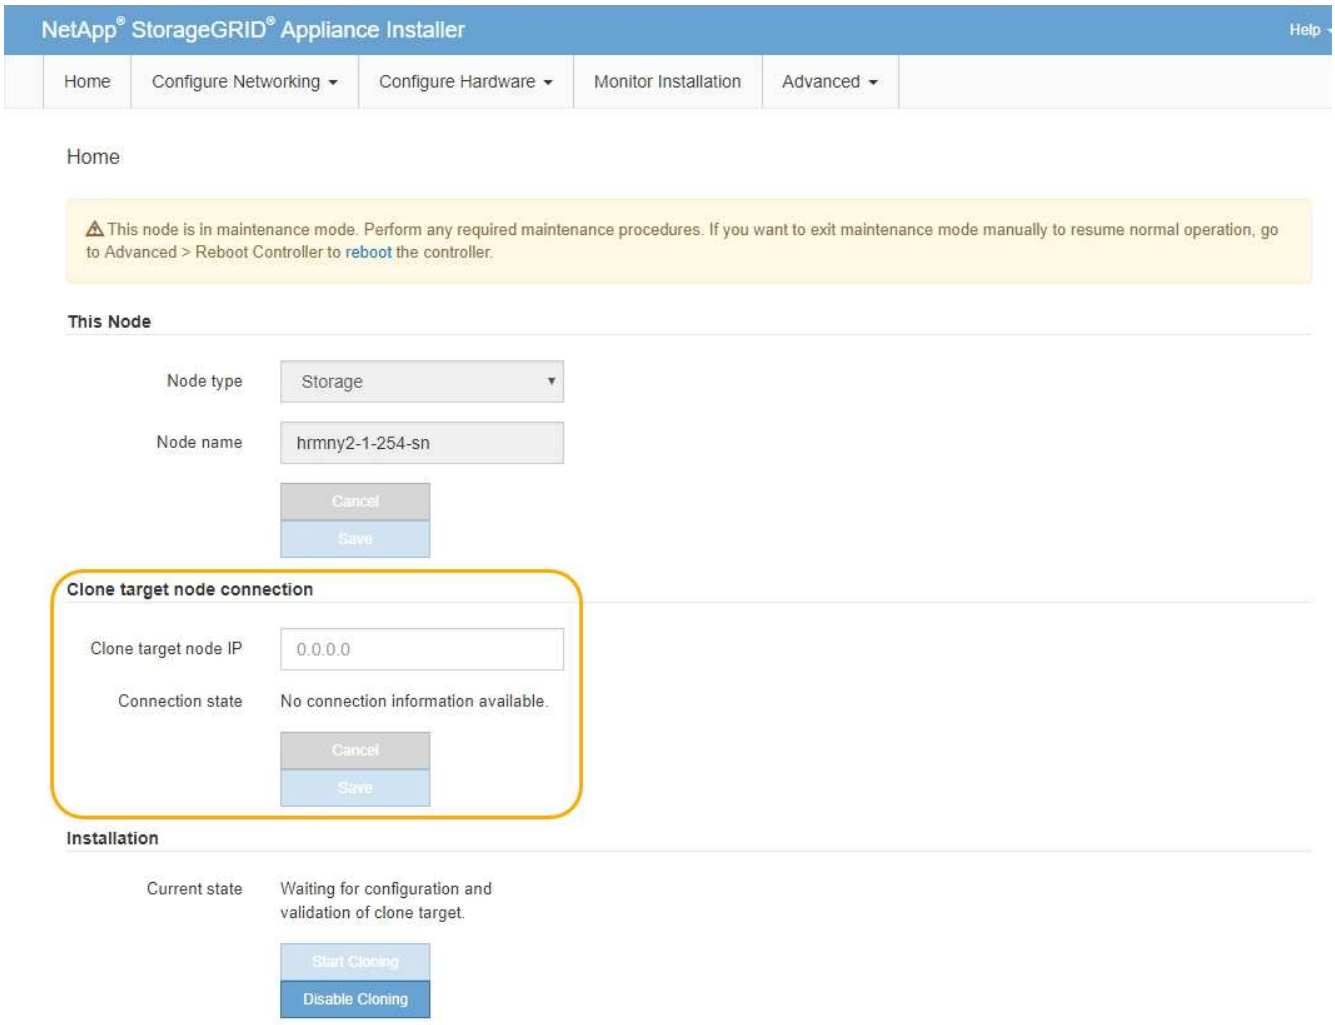

3. For **Clone target node IP**, enter the temporary IP address assigned to the target node for the network to use for clone data-transfer traffic, and then select **Save**.

   Typically, you enter the IP address for the Grid Network, but if you need to use a different network for clone data-transfer traffic, enter the IP address of the target node on that network.

   > Using the 1-GbE Admin Network for clone data transfer results in slower cloning.

   After the target appliance is configured and validated, in the Installation section, **Start Cloning** is enabled on the source node.

   If issues exist that prevent cloning, **Start Cloning** is not enabled and issues that you must resolve are listed as the **Connection state**. These issues are listed on the StorageGRID Appliance Installer Home page of both the source node and the target appliance. Only one issue displays at a time and the state automatically updates as conditions change. Resolve all cloning issues to enable **Start Cloning**.

   When **Start Cloning** is enabled, the **Current state** indicates the StorageGRID network that was selected for cloning traffic, along with information about using that network connection. See Considerations and requirements for appliance node cloning.

4. Select **Start Cloning** on the source node.
5. Monitor the cloning progress using the StorageGRID Appliance Installer on either the source or target

node.

The StorageGRID Appliance Installer on both the source and target nodes indicates the same status.

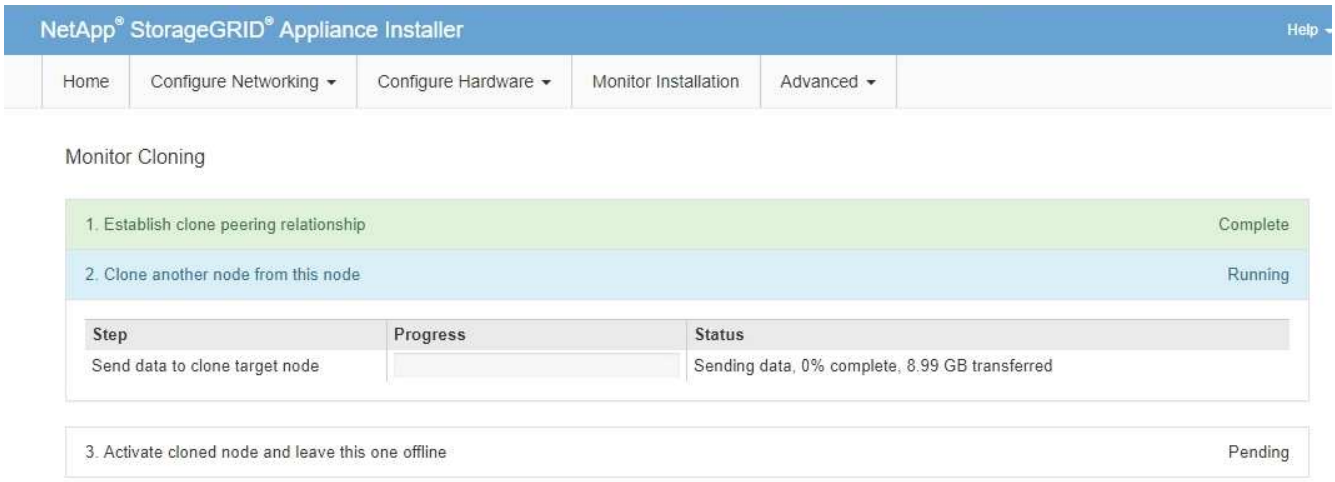

The Monitor Cloning page provides detailed progress for each stage of the cloning process:

- **Establish clone peering relationship** shows the progress of cloning set up and configuration.
- **Clone another node from this node** shows the progress of data transfer. (This part of the cloning process can take several days to complete.)
- **Activate cloned node and leave this one offline** shows the progress of transferring control to the target node and placing the source node in a pre-install state, after data transfer is complete.

6. If you need to terminate the cloning process and return the source node to service before cloning is complete, on the source node go to the StorageGRID Appliance Installer Home page and select **Advanced** > **Reboot Controller**, and then select **Reboot into StorageGRID**.

   If the cloning process is terminated:

   - The source node exits maintenance mode and rejoins StorageGRID.
   - The target node remains in the pre-install state. To restart cloning the source node, start the cloning process again from step 1.

When cloning successfully completes:

- The source and target nodes swap IP addresses:
  - The target node now uses the IP addresses originally assigned to the source node for the Grid, Admin, and Client Networks.
  - The source node now uses the temporary IP address initially assigned to the target node.
- The target node exits maintenance mode and joins StorageGRID, replacing the source node.
- The source appliance is in a pre-installed state, as if you had prepared it for reinstallation.

> If the appliance does not rejoin the grid, go to the StorageGRID Appliance Installer Home page for the source node, select **Advanced** > **Reboot Controller**, and then select **Reboot into Maintenance Mode**. After the source node reboots in maintenance mode, repeat the node cloning procedure.

- User data remains on the source appliance as a recovery option if an unexpected issue occurs with the target node. After the target node has successfully rejoined StorageGRID, user data on the source appliance is outdated and is no longer needed.

  Outdated user data is overwritten when you install or expand the source appliance as a new node in another grid.

  You can also reset the controller configuration on the source appliance to make this data inaccessible:

  1. Open the StorageGRID Appliance Installer for the source appliance using the temporary IP address initially assigned to the target node.
  2. Select **Help** > **Support and Debug Tools**.
  3. Select **Reset Storage Controller Configuration**.

  > If needed, contact technical support for assistance resetting the storage controller configuration.

  > Overwriting the data or resetting the controller configuration makes the outdated data difficult or impossible to retrieve; however, neither method securely removes the data from the source appliance. If a secure erase is required, use a data-wiping tool or service to permanently and securely remove data from the source appliance.

You can:

- Use the source appliance as a target for additional cloning operations: no additional configuration is required. This appliance already has the temporary IP address assigned that were originally specified for the first clone target.
- Install and set up the source appliance as a new appliance node.
- Discard the source appliance if it is no longer of use with StorageGRID.

Maintain SGF6112 hardware

Maintain SGF6112 appliance

You might need to perform maintenance procedures on your appliance. Procedures specific to maintaining your SGF6112 appliance are in this section.

The procedures in this section assume that the appliance has already been deployed as a Storage Node in a StorageGRID system.

Configuration maintenance procedures are performed using the Appliance Installer, Grid Manager, or BMC interface. These procedures include:

- Turn appliance identify LED on and off
- Locate appliance in data center
- Shut down the appliance
- Change link configuration of the appliance

Hardware maintenance procedures require the physical manipulation of specific SGF6112 components.

Drive firmware upgrade

The firmware on the drives in the SGF6112 is automatically checked every time the appliance is rebooted. When necessary, the firmware is automatically upgraded to the version expected by the current StorageGRID release. Usually, firmware upgrades occur during StorageGRID software upgrades. Any necessary drive firmware upgrades for existing StorageGRID versions will be included in hotfixes. Follow the instructions provided with each hotfix to ensure that the upgrade is applied to all drives that could benefit from it.

SANtricity System Manager is not needed to maintain the SGF6112 appliance.

General Maintenance procedures

See Common maintenance procedures for procedures that are the same for all appliances, such as applying a hotfix, recovering a node or site, and performing network maintenance.

See Set up appliance hardware for appliance maintenance procedures that are also performed during initial appliance installation and configuration.

Configuration maintenance procedures

Turn appliance identify LED on and off

The blue identify LED on the front and back of the appliance can be turned on to help locate the appliance in a data center.

Before you begin

You have the BMC IP address of the appliance you want to identify.

Steps

1. Access the appliance BMC interface.
2. Select **Server Identify**.

   The current status of the identify LED is selected.

3. Select **ON** or **OFF**, and then select **Perform Action**.

   When you select **ON**, the blue identify LEDs light on the front (shown) and rear of the appliance.

If a bezel is installed on the controller, it might be difficult to see the front identify LED.

4. Turn the LED on and off as needed.

Related information

Locate appliance in data center

Access BMC interface

Locate appliance in data center

Locate the appliance so that you can perform hardware maintenance or upgrades.

Before you begin

- You have determined which appliance requires maintenance.
- To help locate the appliance in your data center, turn on the blue identify LED.

Steps

1. Find the appliance in the data center.
   - Look for a lit blue identify LED on the front or rear of the appliance.

     The front identify LED is behind the front bezel and might be difficult to see if the bezel is installed.

     

   - Check the tags attached to the front of the appliance for a matching part number to confirm you have found the correct appliance.
2. Remove the front bezel, if one is installed, to access the front panel controls and indicators.
3. Turn off the blue identify LED if you used it to locate the appliance.
   - Press the identify LED switch on the appliance front panel.
   - Use the appliance BMC interface.

Shut down the SGF6112 appliance

Shut down the appliance to perform hardware maintenance.

Before you begin

- You have physically located the appliance.

About this task

To prevent service interruptions, shut down the appliance during a scheduled maintenance window when periods of service disruption are normally expected.

Steps

1. Shut down the appliance:

   > You must perform a controlled shut down of the appliance by entering the commands specified below. It is a best practice to perform a controlled shutdown when possible to avoid unnecessary alerts, ensure full logs are available, and avoid service disruptions.

   a. If you have not already logged into the grid node, log in using PuTTY or another ssh client:

      i. Enter the following command: `ssh admin@grid_node_IP`

      ii. Enter the password listed in the `Passwords.txt` file.

      iii. Enter the following command to switch to root: `su -`

      iv. Enter the password listed in the `Passwords.txt` file.

      When you are logged in as root, the prompt changes from `$` to `#`.

   b. Shut down the appliance:
   `shutdown -h now`

   This command might take up to 10 minutes to complete.

2. Use one of the following methods to verify that the appliance is powered off:
   - Look at the power LED on the front of the appliance and confirm that it is off.
   - Check the Power Control page of the BMC interface to confirm the appliance is off.

Power on SGF6112 and verify operation

Power on the controller after completing maintenance.

Before you begin

- You have installed the controller in a cabinet or rack and connected the data and power cables.

  Reinstall SGF6112 controller into cabinet or rack

- You have physically located the controller in the data center.

  Locate SGF6112 appliance in data center

Steps

1. Power on the appliance.

   You might have to remove the bezel to access the power switch; if so, remember to reinstall it afterwards.

2. Monitor the controller LEDs and boot-up codes using one of the following methods:
   - Press the power switch on the front of the controller.
   - Use the controller BMC interface:
     i. Access the controller BMC interface.

     Access BMC interface

ii. Select **Power Control**.

iii. Select **Power On** and then select **Perform Action**.

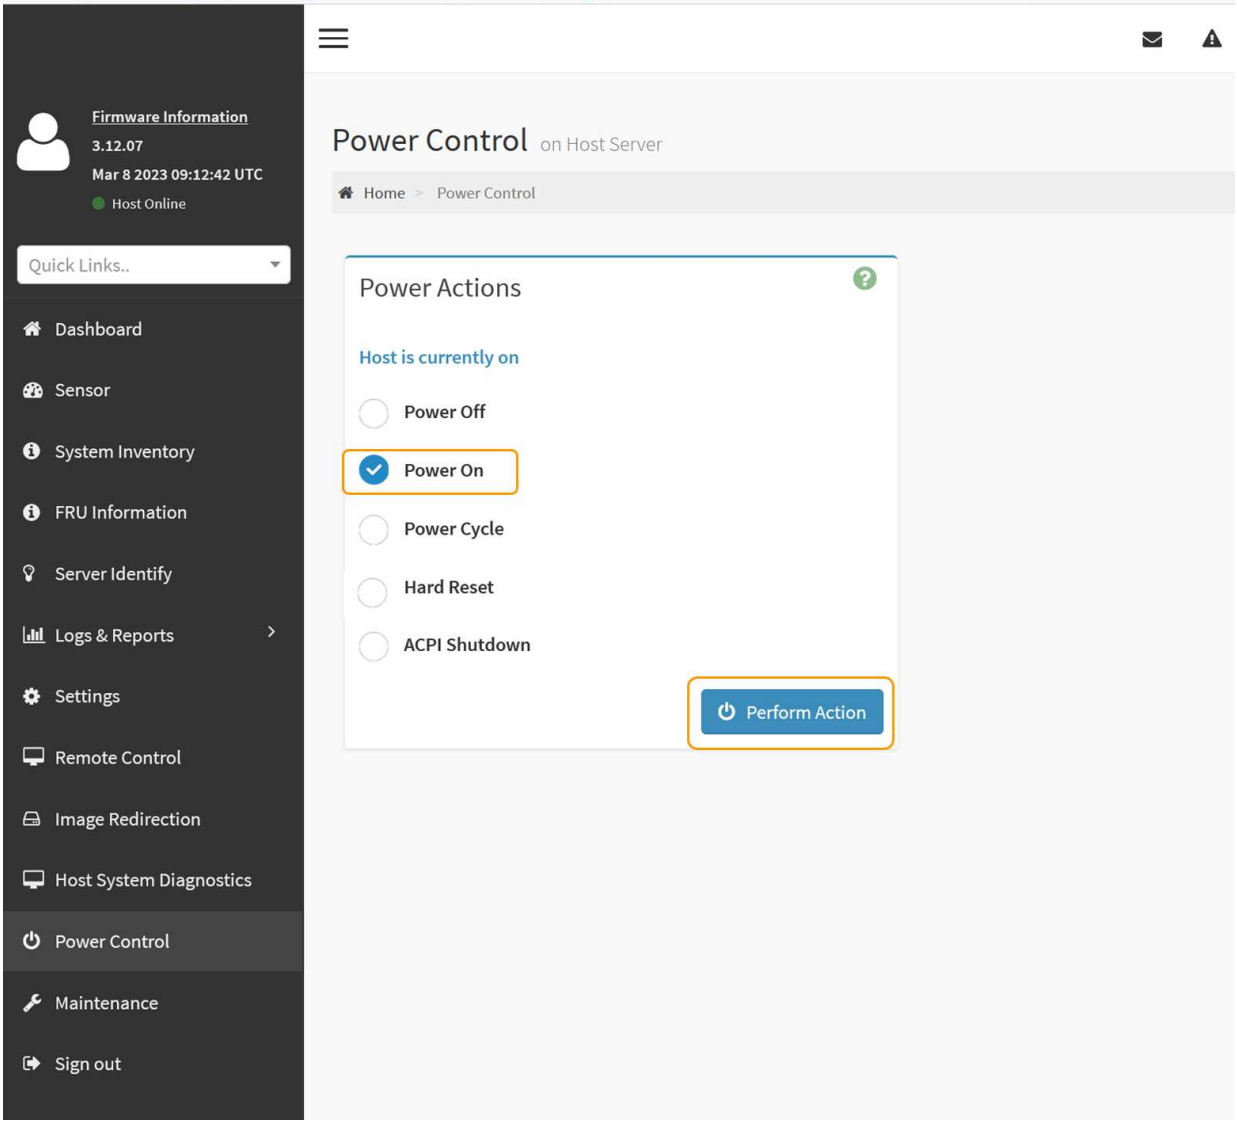

Use the BMC interface to monitor start-up status.

3. Confirm that the appliance controller displays in the Grid Manager and with no alerts.

   It might take up to 20 minutes for the controller to display in the Grid Manager.

   Don’t take another appliance node offline unless this appliance has a green icon.

4. Confirm that the new appliance is fully operational by logging in to the grid node using PuTTY or another ssh client:

   a. Enter the following command: `ssh Appliance_IP`

   b. Enter the password listed in the `Passwords.txt` file.

   c. Enter the following command to switch to root: `su -`

   d. Enter the password listed in the `Passwords.txt` file.

When you are logged in as root, the prompt changes from `$` to `#`.

Related information

View status indicators

Change link configuration of SGF6112 appliance

You can change the Ethernet link configuration of the appliance including the port bond mode, the network bond mode, and the link speed.

Before you begin

- You have placed the appliance into maintenance mode.

> In rare instances, placing a StorageGRID appliance into maintenance mode might make the appliance unavailable for remote access.

Steps

1. From the StorageGRID Appliance Installer, select **Configure Networking** > **Link Configuration**.
2. Make the desired changes to the link configuration.

   For more information about the options, see Configure network links.

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

   > You might lose your connection if you made changes to the network or link you are connected through. If you aren't reconnected within 1 minute, re-enter the URL for the StorageGRID Appliance Installer using one of the other IP addresses assigned to the appliance: **`https://appliance_IP:8443`**

4. Make any necessary changes to the IP addresses for the appliance.

   If you made changes to the VLAN settings, the subnet for the appliance might have changed. If you need to change the IP addresses for the appliance, see Configure StorageGRID IP addresses.

5. Select **Configure Networking** > **Ping Test** from the menu.
6. Use the Ping Test tool to check connectivity to IP addresses on any networks that might have been affected by the link configuration changes you made when configuring the appliance.

   In addition to any other tests you choose to perform, confirm that you can ping the Grid Network IP address of the primary Admin Node, and the Grid Network IP address of at least one other node. If necessary, return to the instructions for configuring network links, and correct any issues.

7. Once you are satisfied that your link configuration changes are working, reboot the node. From the StorageGRID Appliance Installer, select **Advanced** > **Reboot Controller**, and then select one of these options:
   - Select **Reboot into StorageGRID** to reboot the compute controller with the node rejoining the grid. Select this option if you are done working in maintenance mode and are ready to return the node to normal operation.
   - Select **Reboot into Maintenance Mode** to reboot the compute controller with the node remaining in maintenance mode. (This option is available only when the controller is in maintenance mode.) Select

this option if there are additional maintenance operations you need to perform on the node before it rejoins the grid.

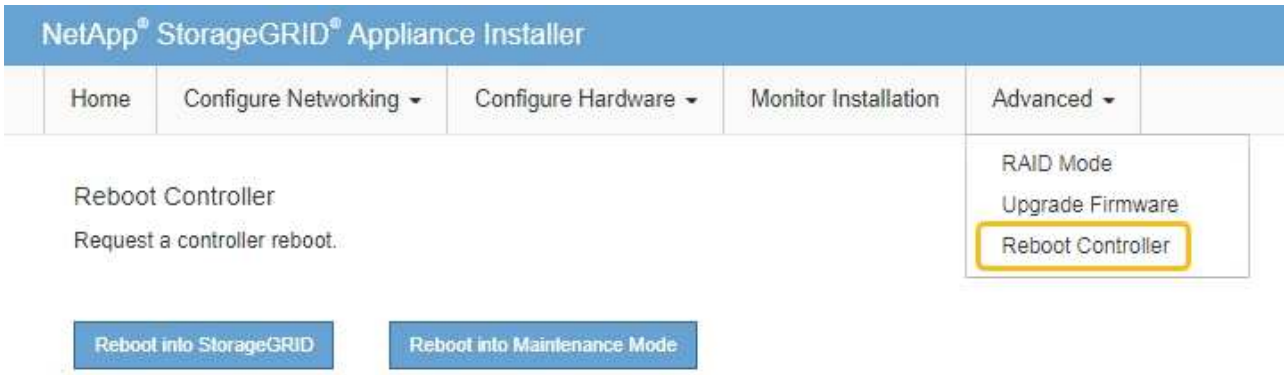

It can take up to 20 minutes for the appliance to reboot and rejoin the grid. To confirm that the reboot is complete and that the node has rejoined the grid, go back to the Grid Manager. The **NODES** page should display a normal status (green check mark icon to the left of the node name) for the appliance node, indicating that no alerts are active and the node is connected to the grid.

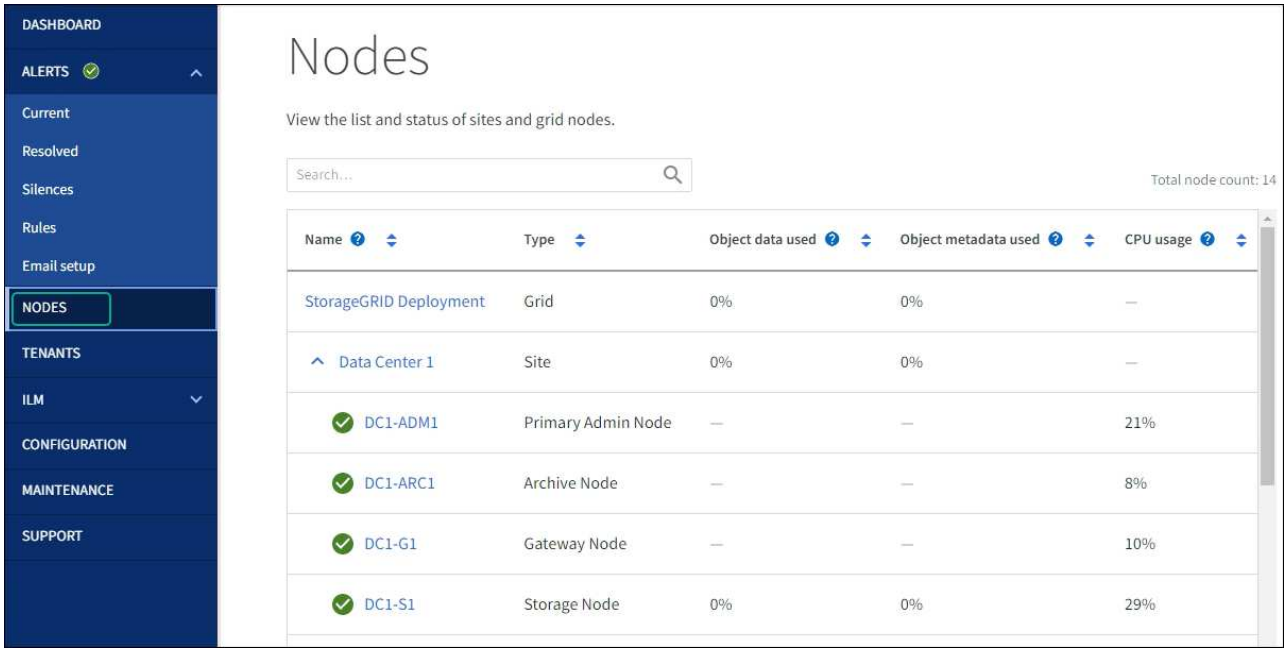

Hardware maintenance procedures

Verify component to replace in the SGF6112

If you are unsure which hardware component to replace in your appliance, complete this procedure to identify the component and the location of the appliance in the data center.

Before you begin

- You have the serial number of the storage appliance where the component needs to be replaced.

- You are signed in to the Grid Manager using a supported web browser.

About this task

Use this procedure to identify the appliance with failed hardware and which of the following replaceable hardware components is not operating properly:

- Power supplies
- Fans
- Solid state drives (SSDs)
- Network Interface Cards (NICs)
- CMOS battery

Steps

1. Identify the failed component and the name of the appliance that it is installed in.
   a. In Grid Manager, Select **ALERTS** > **Current**.

      The Alerts page appears.

   b. Select the alert to see the alert details.

      Select the alert, not the heading for a group of alerts.

   c. Record the node name and unique identifying label of the component that has failed.

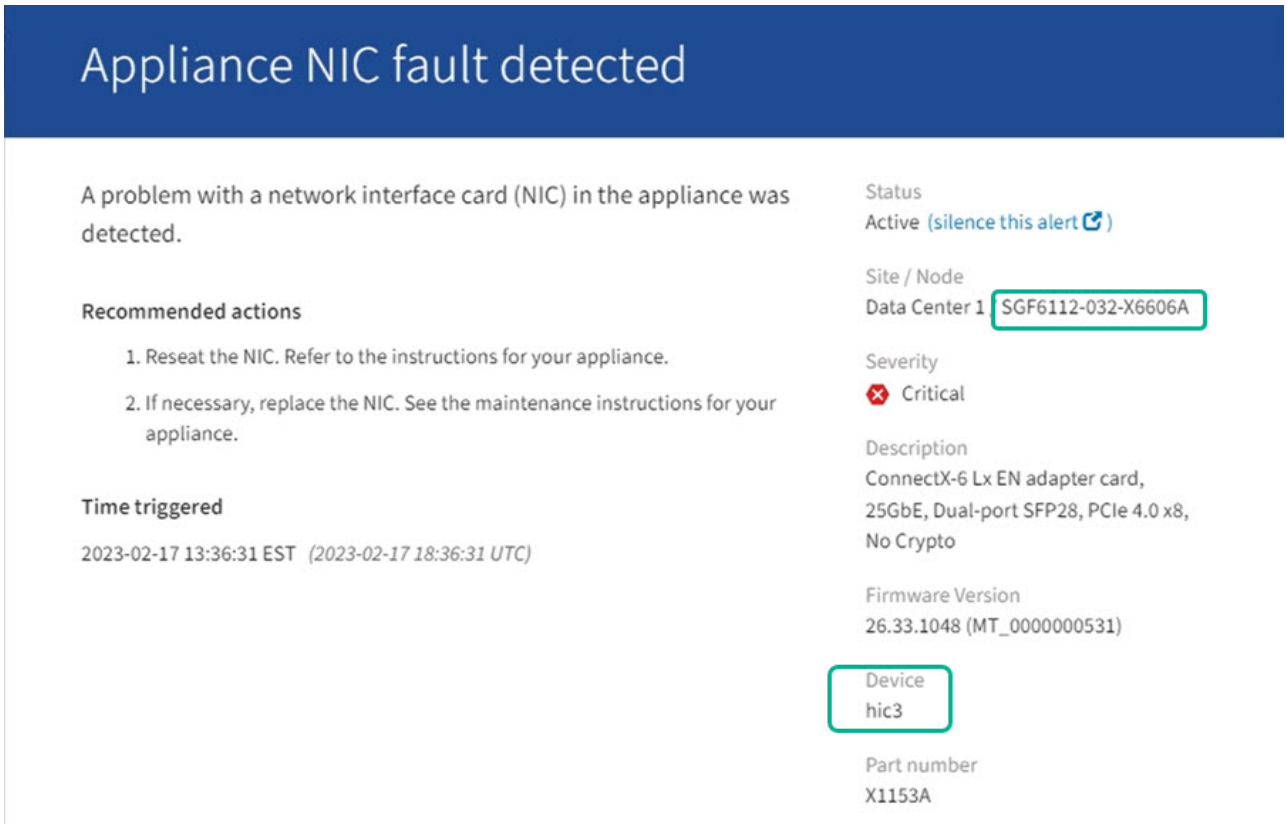

2. Identify the chassis with the component that needs to be replaced.
   a. From the Grid Manager, select **NODES**.

b. From the table on the Nodes page, select the appliance Storage Node name with the failed component.

c. Select the **Hardware** tab.

Check the **Compute controller serial number** in the StorageGRID Appliance section. Check if the serial number matches the serial number of the storage appliance where you are replacing the component. If the serial number matches, you have found the correct appliance.

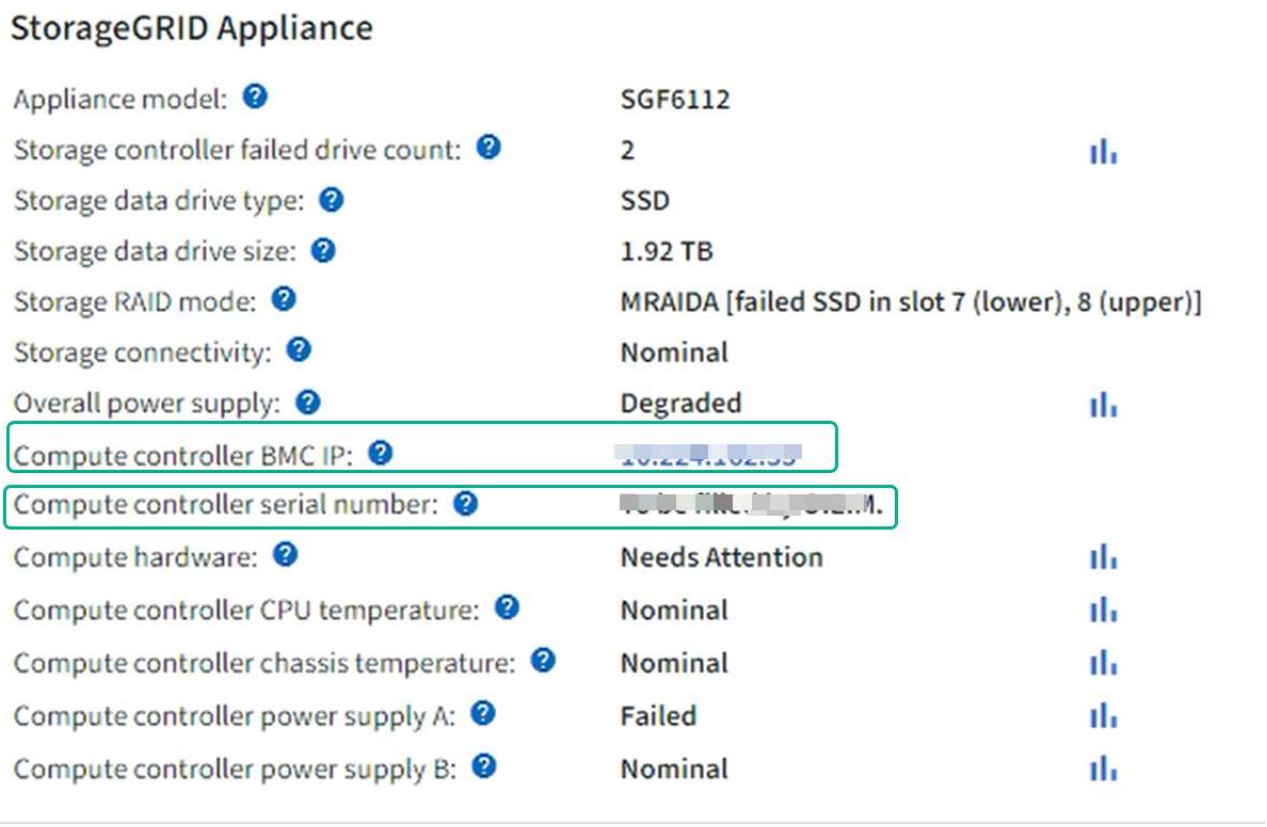

- If the StorageGRID Appliance section in Grid Manager does not display, the node selected is not a StorageGRID appliance. Select a different node from the tree view.
- If the serial numbers don't match, select a different node from the tree view.

3. After you locate the node where the component needs to be replaced, write down the appliance BMC IP address listed the StorageGRID Appliance section.

To help you locate the appliance in the data center, you can use the BMC IP address to turn on the appliance identify LED.

Turn the appliance identify LED on and off

Replace one or both power supplies in the SGF6112 appliance

The SGF6112 appliance has two power supplies for redundancy. If one of the power supplies fails, you must replace it as soon as possible to ensure that the appliance has redundant power. Both power supplies operating in the appliance must be of the same model and wattage.

Before you begin

- You have physically located the appliance with the power supply to be replaced.
- You have determined the location of the power supply to replace.
- If you are replacing only one power supply:
  - You have unpacked the replacement power supply unit and ensured that it is the same model and wattage as the power supply unit you are replacing.
  - You have confirmed that the other power supply is installed and running.
- If you are replacing both power supplies at the same time:
  - You have unpacked the replacement power supply units and ensured they are the same model and wattage.

About this task

The figure shows the two power supply units for the SGF6112. The power supplies are accessible from the back of the appliance.

Steps

1. If you are replacing only one power supply, you don't need to shut down the appliance. Go to the Unplug the power cord step. If you are replacing both power supplies at the same time, do the following before unplugging the power cords:
   a. Shut down the appliance.

   > If you have ever used an ILM rule that creates only one copy of an object and you are replacing both power supplies at the same time, you must replace the power supplies during a scheduled maintenance window. Otherwise, you might temporarily lose access to those objects during this procedure. See information about why you should not use single-copy replication.

2. Unplug the power cord from each power supply to be replaced.

   When viewed from the rear of the appliance, power supply A (PSU0) is on the right and power supply B (PSU1) is on the left.

3. Lift the handle on the first supply to be replaced.

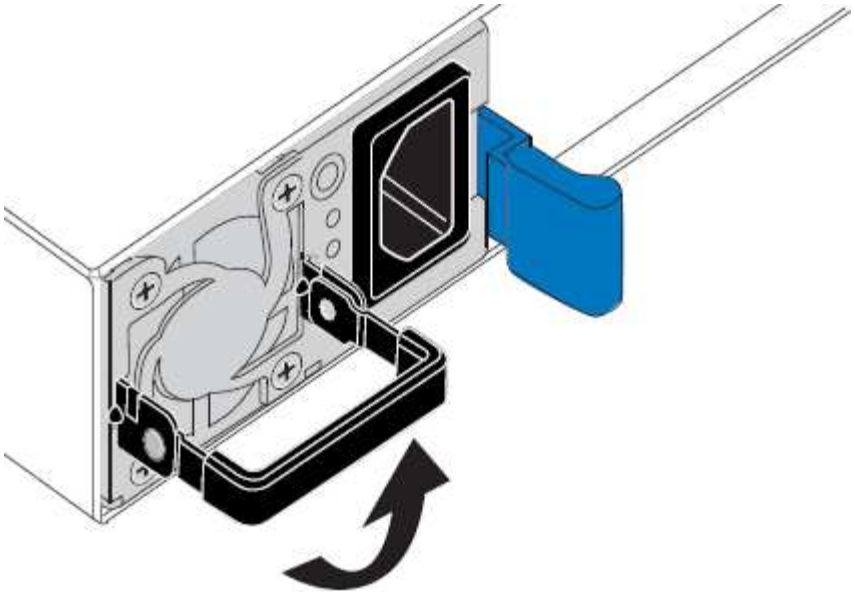

4. Press the blue latch and pull the power supply out.

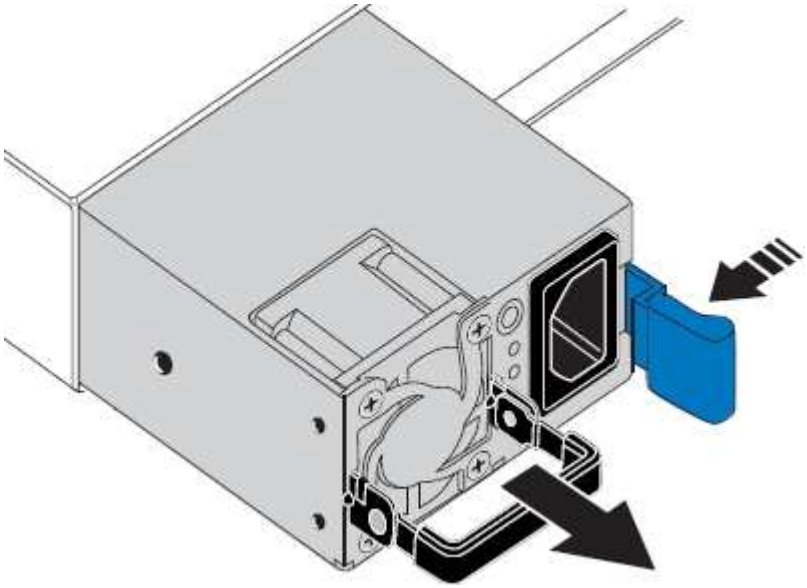

5. With the blue latch on the right, slide the replacement power supply into the chassis.

   Both installed power supplies must be the same model and wattage.

   Ensure that the blue latch is on the right side when you slide the replacement unit in.

   You will feel a click when the power supply is locked into place.

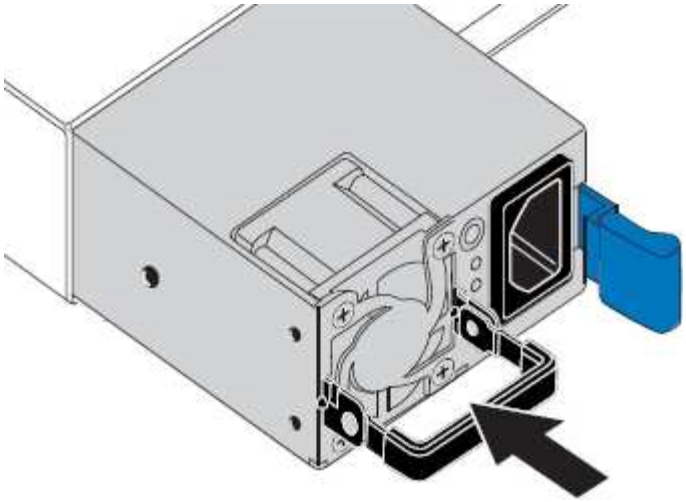

6. Push the handle back down against the body of the PSU.
7. If you are replacing both power supplies, repeat steps 2 though 6 to replace the second power supply.
8. Connect the power cords to the replaced units and apply power.

After replacing the part, return the failed part to NetApp, as described in the RMA instructions shipped with the kit. See the Part Return & Replacements page for further information.

Replace fan in an SGF6112 appliance

The SGF6112 appliance has eight cooling fans. If one of the fans fails, you must replace it as soon as possible to ensure that the appliance has proper cooling.

Before you begin

- You have the correct replacement fan.
- You have determined the location of the fan to replace.

  Verify component to replace

- You have physically located the SGF6112 appliance where you are replacing the fan in the data center.

  Locate appliance in data center

  > A controlled shutdown of the appliance is required before removing the appliance from the rack.

- You have disconnected all cables and removed the appliance cover.

  Remove SGF6112 cover

- You have confirmed that the other fans are installed and running.

About this task

To prevent service interruptions, confirm that all other Storage Nodes are connected to the grid before starting the fan replacement or replace the fan during a scheduled maintenance window when periods of service disruption are normally expected. See the information about monitoring node connection states.

If you have ever used an ILM rule that creates only one copy of an object, you must replace the fan during a scheduled maintenance window. Otherwise, you might temporarily lose access to those objects during this procedure. See information about why you should not use single-copy replication.

The appliance node will not be accessible while you replace the fan.

The photograph shows a fan for the appliance. The electrical connector is highlighted. The cooling fans are accessible after you take the top cover off of the appliance.

Each of the two power supply units also contain a fan. The power supply fans aren't included in this procedure.

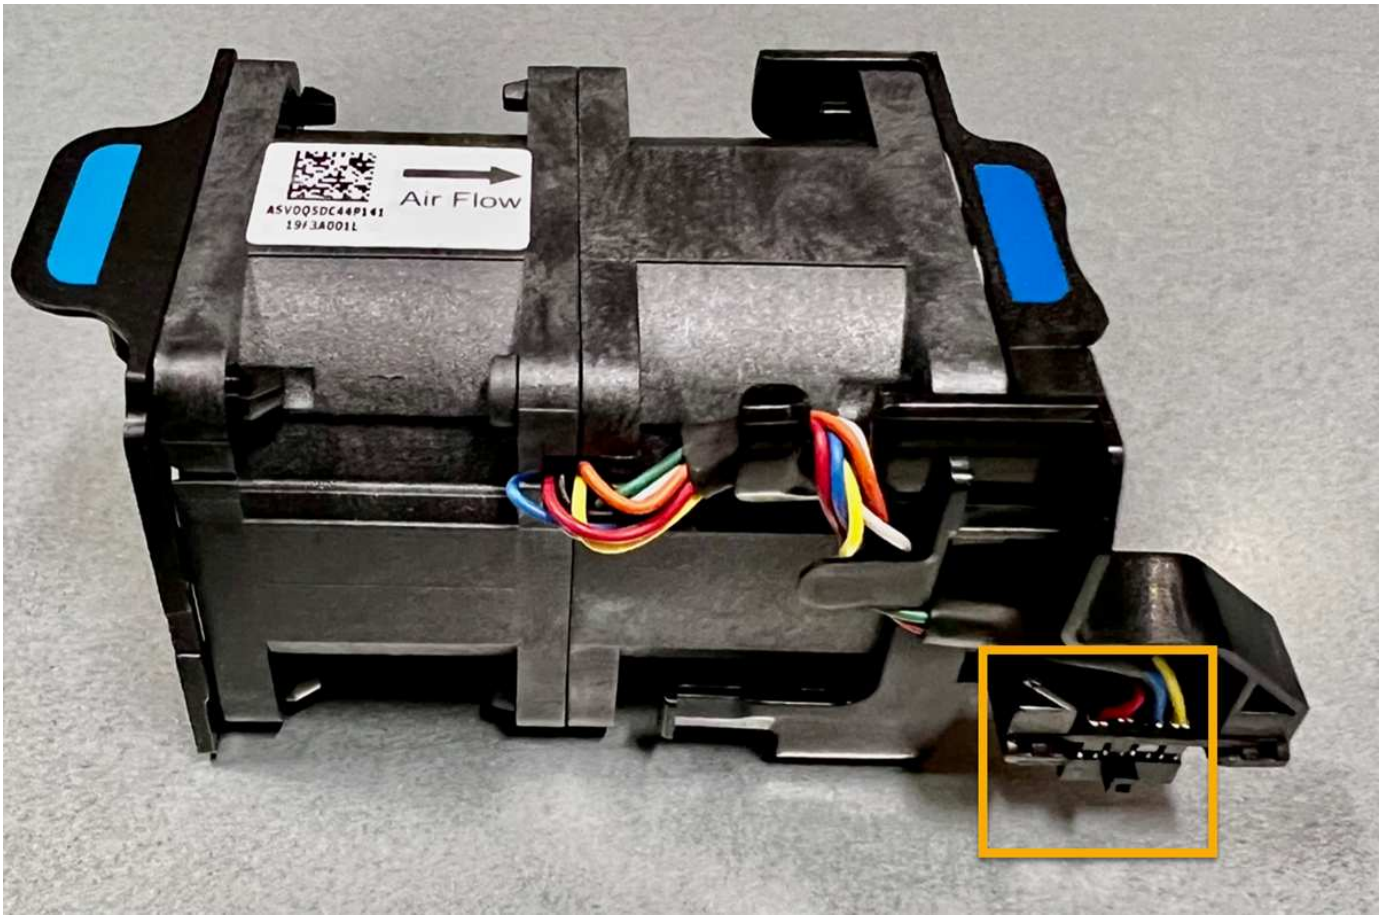

Steps

1. Wrap the strap end of the ESD wristband around your wrist, and secure the clip end to a metal ground to prevent static discharge.
2. Locate the fan that you need to replace.

   The eight fans are in the following positions in the chassis (front half of SGF6112 with top cover removed shown):

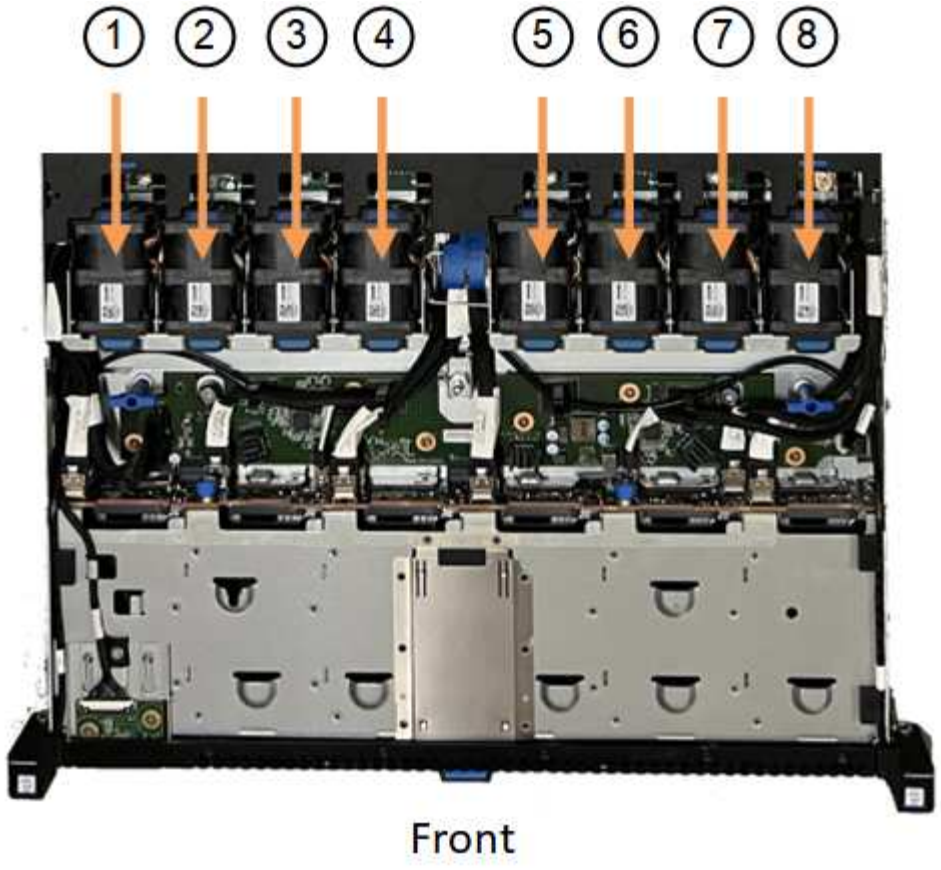

	Fan unit
1	Fan_SYS0
2	Fan_SYS1
3	Fan_SYS2
4	Fan_SYS3
5	Fan_SYS4
6	Fan_SYS5
7	Fan_SYS6
8	Fan_SYS7

3. Using the blue tabs on the fan, lift the failed fan out of the chassis.

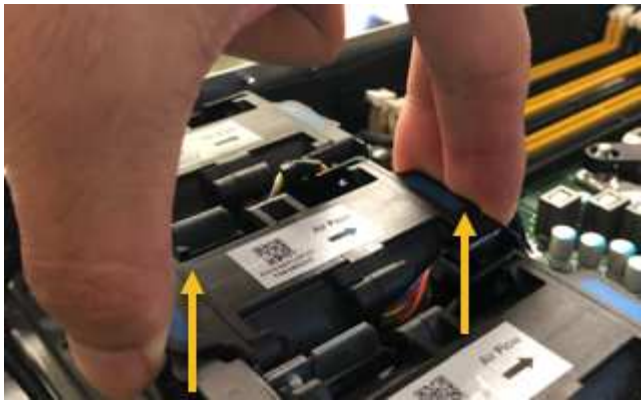

4. Slide the replacement fan into the open slot in the chassis.

   Be sure to align the connector on the fan with the socket in the circuit board.

5. Press the fan’s connector firmly into the circuit board (socket highlighted).

6. Put the top cover back on the appliance, and press the latch down to secure the cover in place.

7. Power on the appliance and monitor the appliance LEDs and boot-up codes.

   Use the BMC interface to monitor boot-up status.

8. Confirm that the appliance node appears in the Grid Manager and that no alerts appear.

After replacing the part, return the failed part to NetApp, as described in the RMA instructions shipped with the kit. See the Part Return & Replacements page for further information.

Replace drives in the SGF6112 appliance

The SGF6112 storage appliance contains 12 SSD drives. Data on the drives is protected by a RAID scheme that enables the appliance to recover from any single drive failure without having to copy data from another node.

The failure of a second drive before an initial drive failure has been corrected might require data be copied from other nodes to restore redundancy. This restoration of redundancy can take longer, and might be impossible, if single-copy ILM rules are in use or were used in the past, or if data redundancy has been impacted by failures on other nodes. Therefore, if one of the SGF6112 drives fails, you must replace it as soon as possible to ensure redundancy.

Before you begin

- You have physically located the appliance.
- You have verified which drive has failed by noting that the drive's left LED is solid amber or using the Grid Manager to view the alert caused by the failed drive.

  See the information about viewing status indicators to verify the failure.

- You have obtained the replacement drive.
- You have obtained proper ESD protection.

Steps

1. Verify that the drive's left fault LED is amber or use the drive slot ID from the alert to locate the drive.

   The twelve drives are in the following positions in the chassis (front of chassis with bezel removed shown):

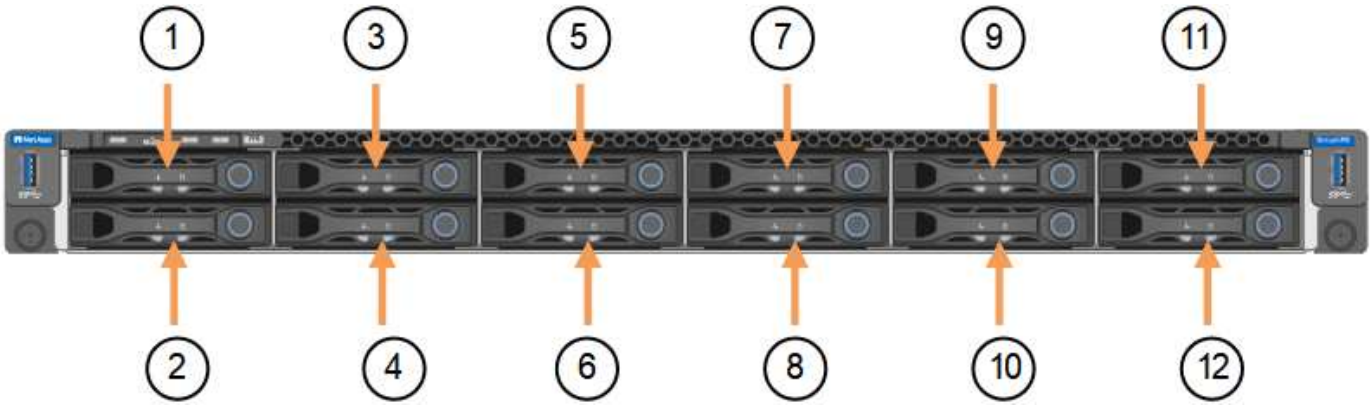

Position	Drive
1	HDD00
2	HDD01
3	HDD02
4	HDD03
5	HDD04
6	HDD05
7	HDD06
8	HDD07
9	HDD08
10	HDD09
11	HDD10
12	HDD11

You can also use the Grid Manager to monitor the status of the SSD drives. Select **NODES**. Then select `Storage Node` > **Hardware**. If a drive has failed, the Storage RAID Mode field contains a message about which drive has failed.

2. Wrap the strap end of the ESD wristband around your wrist, and secure the clip end to a metal ground to prevent static discharge.
3. Unpack the replacement drive, and set it on a static-free, level surface near the appliance.

   Save all packing materials.

4. Press the release button on the failed drive.

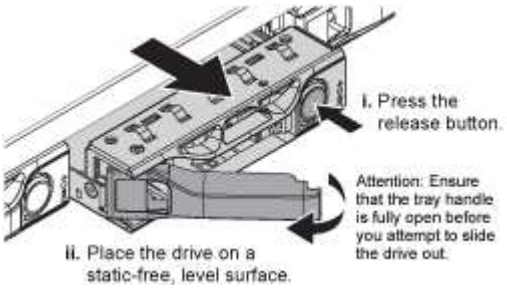

   The handle on the drive springs open partially, and the drive releases from the slot.

5. Open the handle, slide the drive out, and place it on a static-free, level surface.
6. Press the release button on the replacement drive before you insert it into the drive slot.

   The latch springs open.

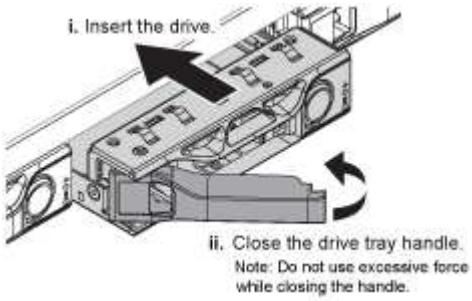

7. Insert the replacement drive in the slot, and then close the drive handle.

   Don't use excessive force while closing the handle.

   When the drive is fully inserted, you hear a click.

   The replaced drive is automatically rebuilt with mirrored data from the working drives. The drive LED should blink initially, but then stop blinking as soon as the system determines that the drive has enough capacity and is functional.

   You can check the status of the rebuild by using the Grid Manager.

8. If more than one drive failed and has been replaced, you might have alerts indicating that some volumes need to have data restored to them. If you receive an alert, before attempting volume recovery, select **NODES** > **`appliance Storage Node`** > **Hardware**. In the StorageGRID Appliance section of the page, verify that the Storage RAID mode is healthy or rebuilding. If the status lists one or more failed drives, correct this condition before attempting volume restoration.
9. In the Grid Manager, go to **NODES** > **`appliance Storage Node`** > **Hardware**. In the StorageGRID Appliance section of the page, verify that the Storage RAID mode is healthy.

After replacing the part, return the failed part to NetApp, as described in the RMA instructions shipped with the kit. See the Part Return & Replacements page for further information.

Replace NIC in the SGF6112

You might need to replace a Network Interface Card (NIC) in the SGF6112 if it is not functioning optimally or if it has failed.

Use these procedures to:

- Remove the NIC
- Reinstall the NIC

Remove the NIC

Before you begin

- You have the correct replacement NIC.

- You have determined the location of the NIC to replace.

  Verify location of component to replace

- You have physically located the SGF6112 appliance where you are replacing the NIC in the data center.

  Locate appliance in data center

  > A controlled shutdown of the appliance is required before removing the appliance from the rack.

- You have disconnected all cables and removed the appliance cover.

  Remove SGF6112 cover

About this task

To prevent service interruptions, confirm that all other Storage Nodes are connected to the grid before starting the Network Interface Card (NIC) replacement or replace the NIC during a scheduled maintenance window when periods of service disruption are normally expected. See the information about monitoring node connection states.

> If you have ever used an ILM rule that creates only one copy of an object, you must replace the NIC during a scheduled maintenance window. Otherwise, you might temporarily lose access to those objects during this procedure. See information about why you should not use single-copy replication.

Steps

1. Wrap the strap end of the ESD wristband around your wrist, and secure the clip end to a metal ground to prevent static discharge.
2. Locate the riser assembly that contains the NIC at the rear of the appliance.

   The three NICs in the SGF6112 are in two riser assemblies in the positions in the chassis shown in the photograph (Rear of SGF6112 with top cover removed shown):

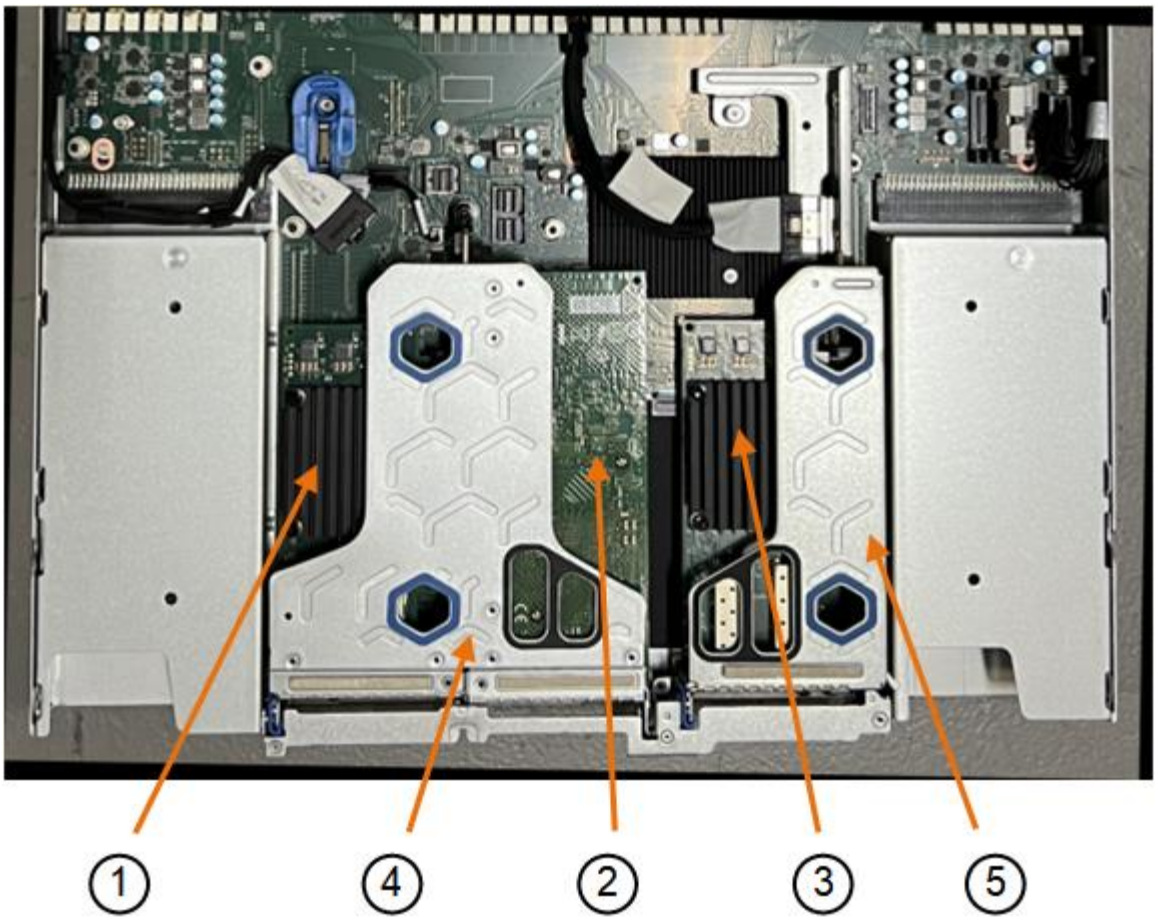

	Device or Part name	Description
1	hic1/hic2	10/25-GbE Ethernet network ports in the two-port riser assembly
2	mtc1/mtc2	1/10GBase-T management ports in the two-port riser assembly
3	hic3/hic4	10/25-GbE Ethernet network ports in the one-port riser assembly
4	Two-slot riser assembly	Support for one of the 10/25-GbE NICs and the 1/10GBase-T NIC
5	One-slot riser assembly	Support for one of the 10/25-GbE NICs

1. Grasp the riser assembly with the failed NIC through the blue-marked holes and carefully lift it upwards. Move the riser assembly toward the front of the chassis as you lift it to allow the external connectors in its installed NICs to clear the chassis.
2. Place the riser on a flat anti-static surface with the metal frame side down to access the NICs.
    - **Two-slot riser assembly with two NICs**

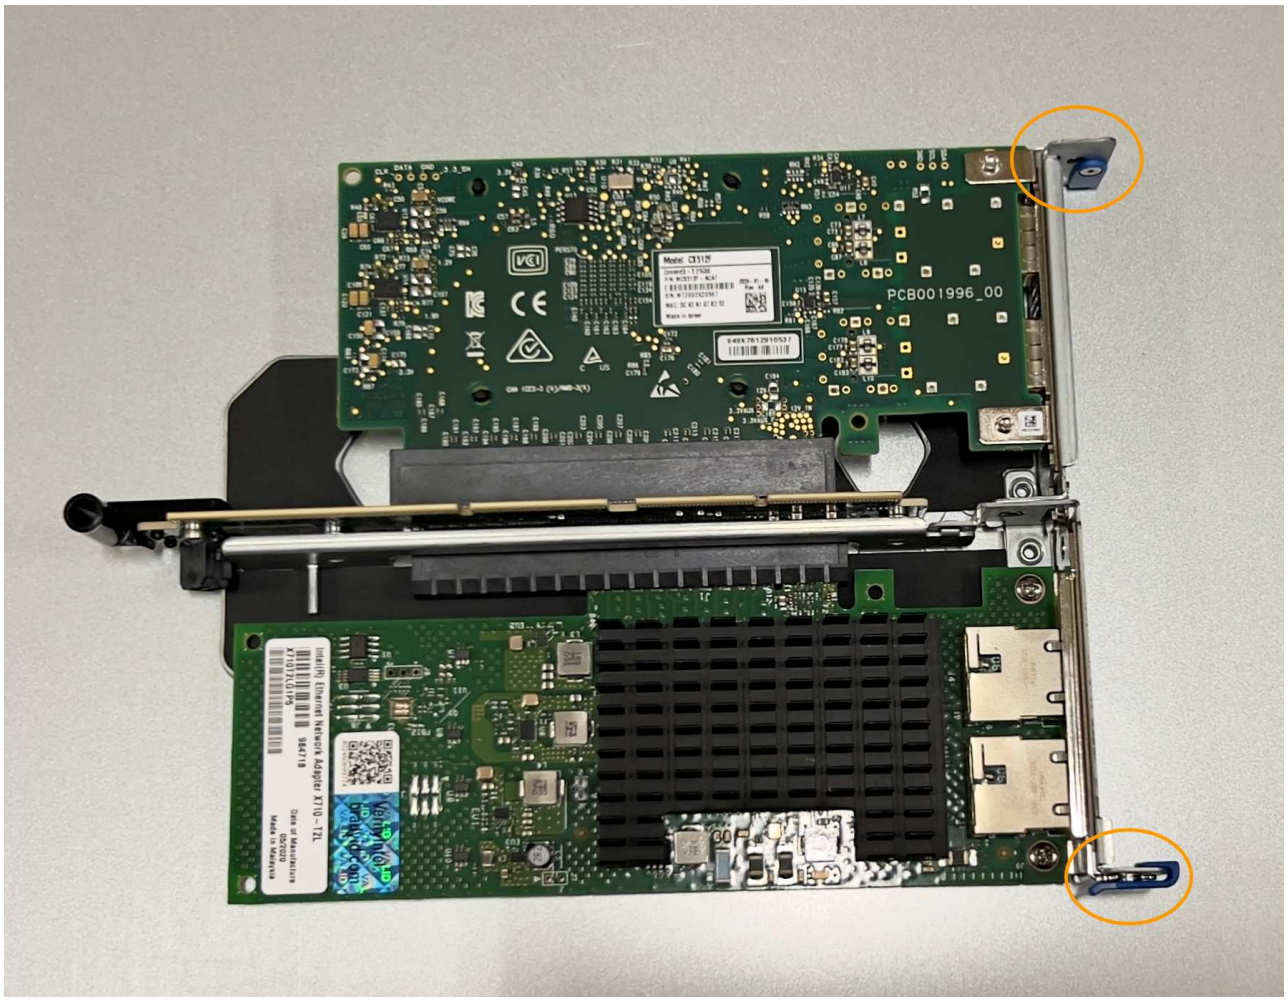

- **One-slot riser assembly with one NIC**

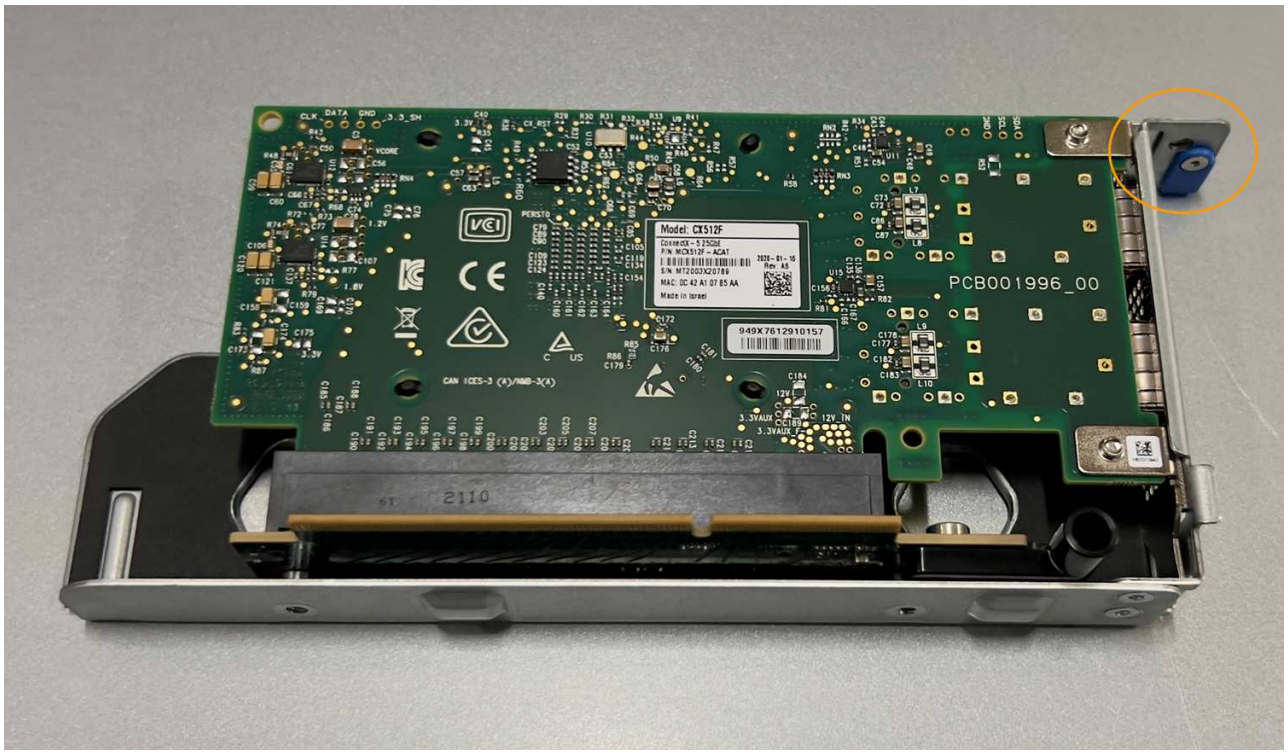

3. Open the blue latch (circled) on the NIC to be replaced and carefully remove the NIC from the riser assembly. Rock the NIC slightly to help remove the NIC from its connector. Don't use excessive force.
4. Place the NIC on a flat anti-static surface.

Reinstall the NIC

Install the replacement NIC into the same location as the one that was removed.

Before you begin

- You have the correct replacement NIC.
- You have removed the existing failed NIC.

Steps

1. Wrap the strap end of the ESD wristband around your wrist, and secure the clip end to a metal ground to prevent static discharge.
2. Remove the replacement NIC from its packaging.
3. If you are replacing one of the NICs in the two-slot riser assembly, do the following:
   a. Ensure the blue latch is in the open position.
   b. Align the NIC with its connector on the riser assembly. Carefully press the NIC into the connector until it is fully seated as shown in the photograph and close the blue latch.

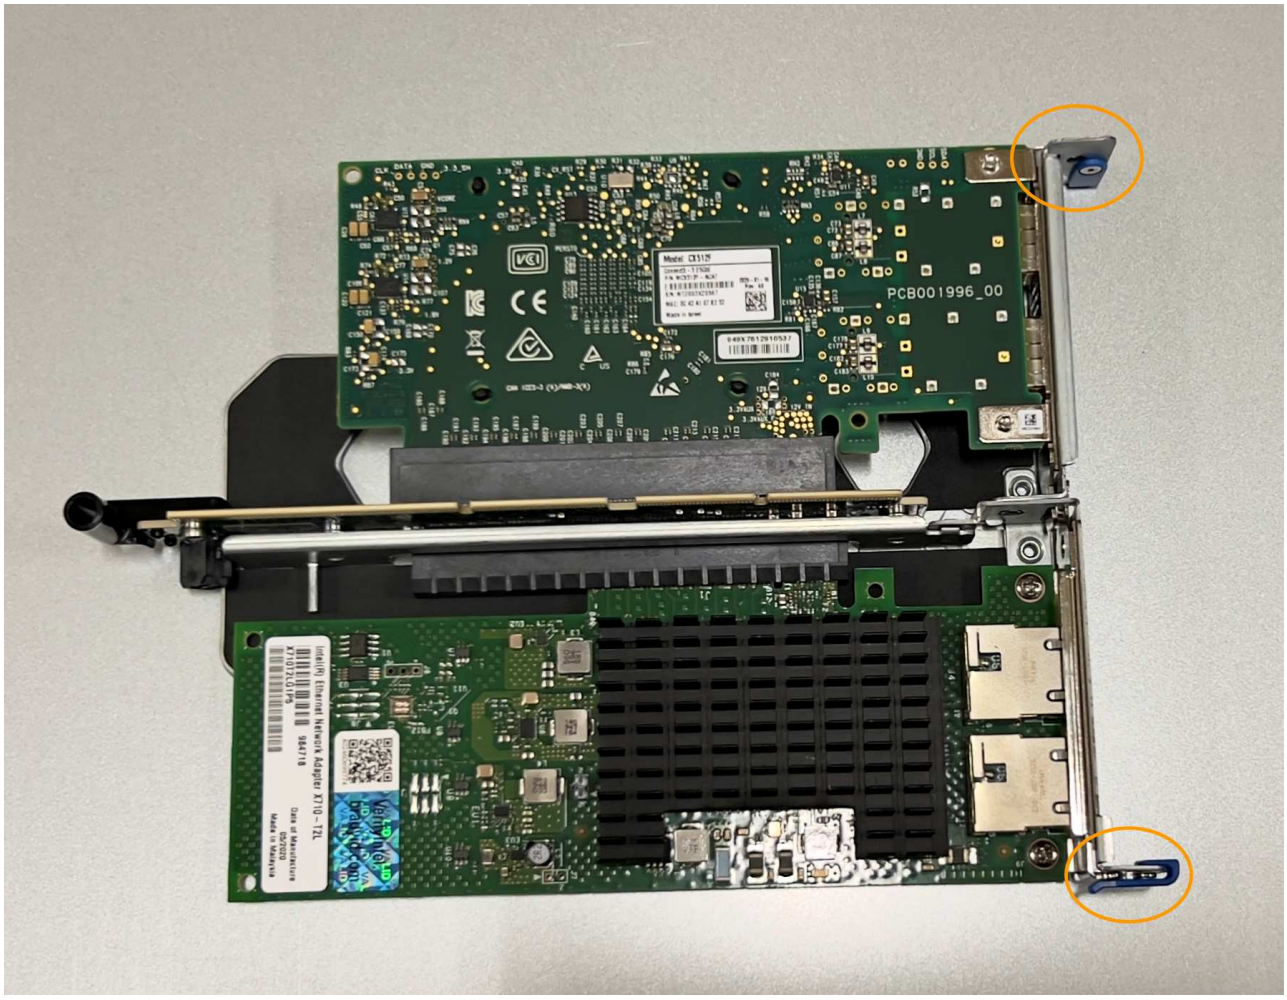

c. Locate the alignment hole on the two-slot riser assembly (circled) that aligns with a guide pin on the system board to ensure correct riser assembly positioning.

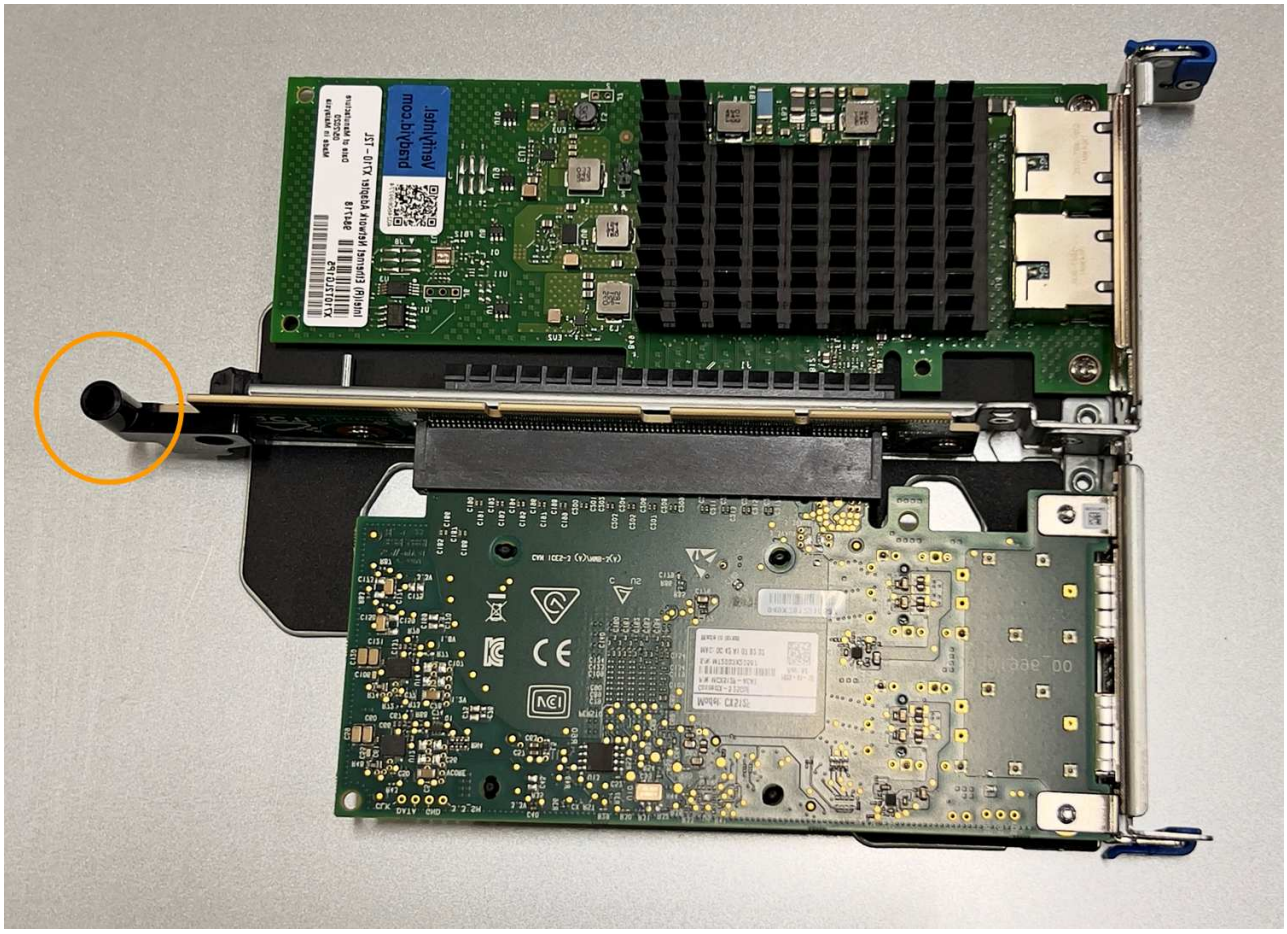

d. Locate the guide pin on the system board

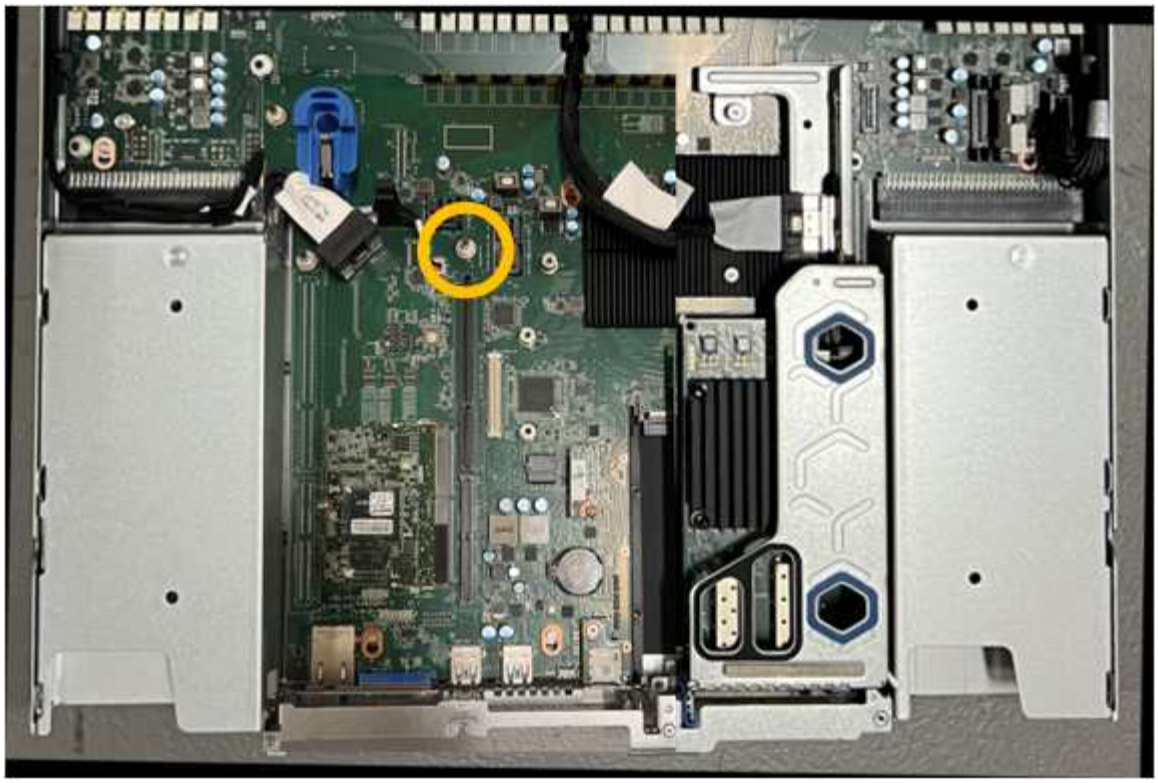

e. Position the riser assembly in the chassis, making sure that it aligns with the connector on the system board and guide pin.

f. Carefully press the two-slot riser assembly in place along its center line, next to the blue-marked holes, until it is fully seated.

4. If you are replacing the NIC in the one-slot riser assembly, do the following:

a. Ensure the blue latch is in the open position.

b. Align the NIC with its connector on the riser assembly. Carefully press the NIC into the connector until it is fully seated as shown in the photograph and close the blue latch.

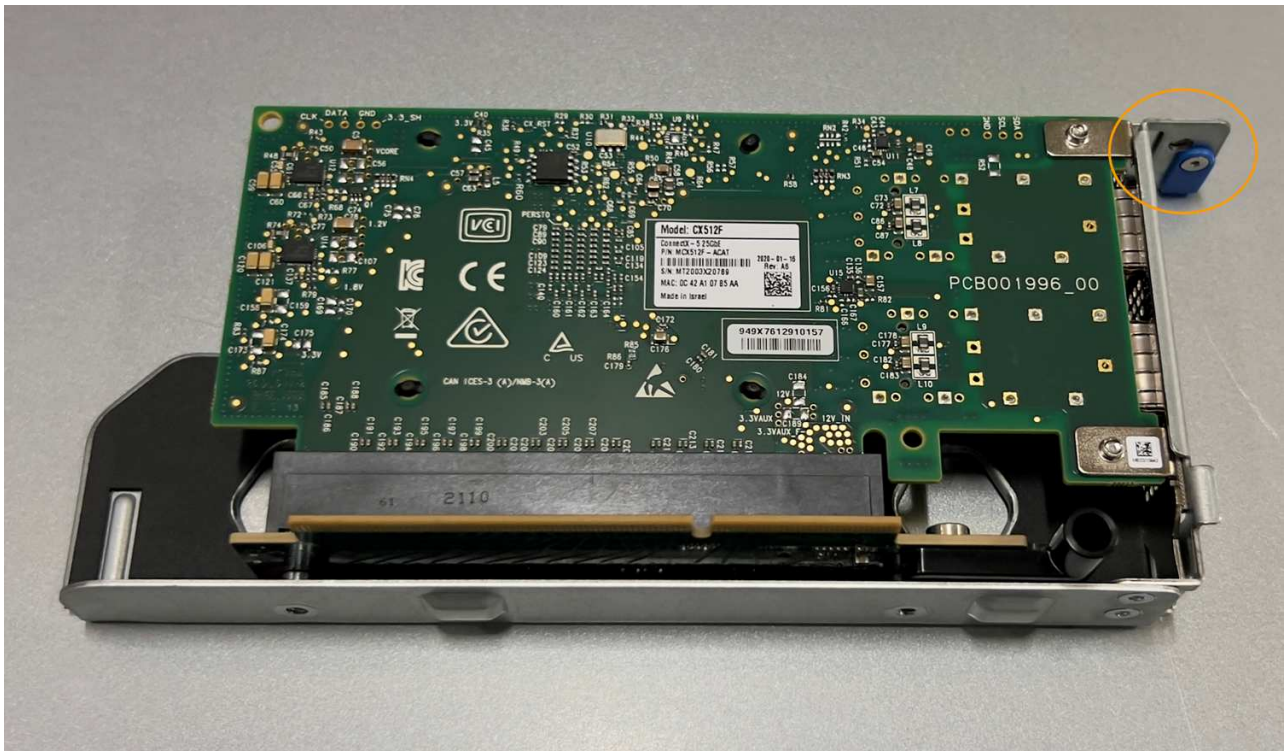

c. Locate the alignment hole on the one-slot riser assembly (circled) that aligns with a guide pin on the system board to ensure correct riser assembly positioning.

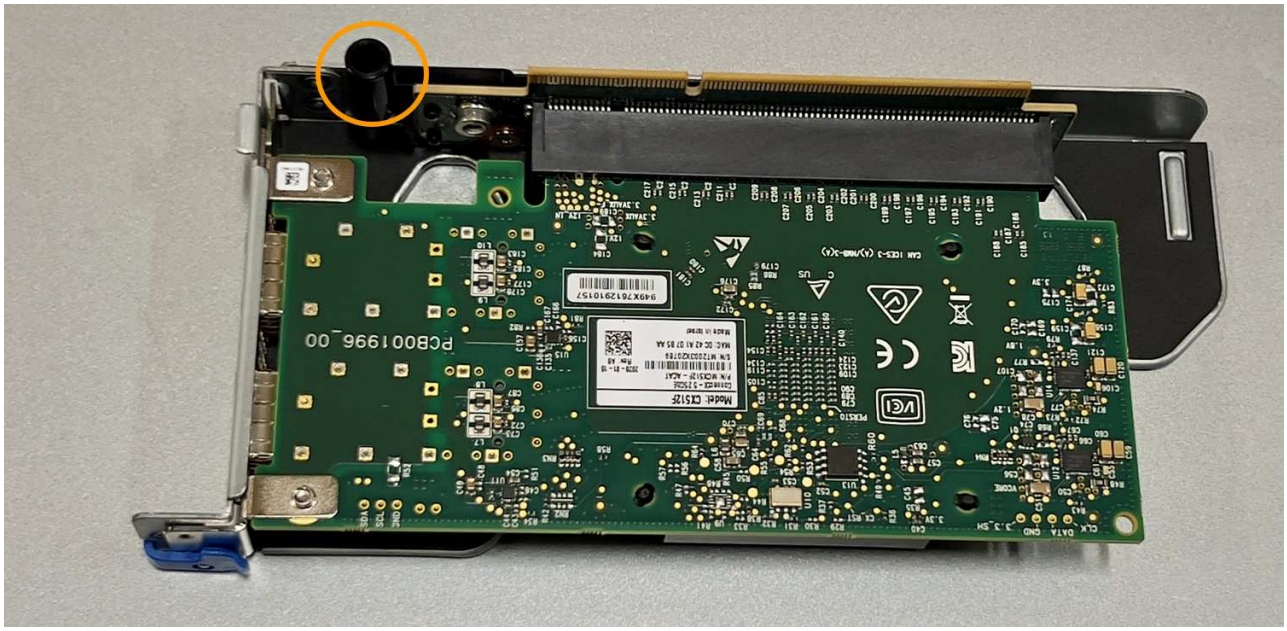

d. Locate the guide pin on the system board

e. Position the one-slot riser assembly in the chassis, making sure that it aligns with the connector on the system board and guide pin.

f. Carefully press the one-slot riser assembly in place along its center line, next to the blue-marked holes, until it is fully seated.

5. Remove the protective caps from the NIC ports where you will be reinstalling cables.

After you finish

If you have no other maintenance procedures to perform in the appliance, reinstall the appliance cover, return the appliance to the rack, attach cables, and apply power.

After replacing the part, return the failed part to NetApp, as described in the RMA instructions shipped with the kit. See the Part Return & Replacements page for further information.

Replace the CMOS battery in the SGF6112

Use this procedure to replace the CMOS coin cell battery on the system board.

Use these procedures to:

- Remove the CMOS battery
- Reinstall the CMOS battery

Remove the CMOS battery

Before you begin

- You have verified that the CMOS battery in the SGF6112 needs to be replaced.

  Verify component to replace

- You have physically located the SGF6112 appliance where you are replacing the CMOS battery in the data center.

Locate appliance in data center

> A controlled shutdown of the appliance is required before removing the appliance from the rack.

- You have disconnected all cables and removed the appliance cover.

Remove SGF6112 cover

About this task

To prevent service interruptions, confirm that all other Storage Nodes are connected to the grid before starting the CMOS battery replacement or replace the battery during a scheduled maintenance window when periods of service disruption are normally expected. See the information about monitoring node connection states.

> If you have ever used an ILM rule that creates only one copy of an object, you must replace the battery during a scheduled maintenance window. Otherwise, you might temporarily lose access to those objects during this procedure. See information about why you should not use single-copy replication.

Steps

1. Wrap the strap end of the ESD wristband around your wrist, and secure the clip end to a metal ground to prevent static discharge.
2. Locate the two-slot riser assembly at the rear of the appliance.

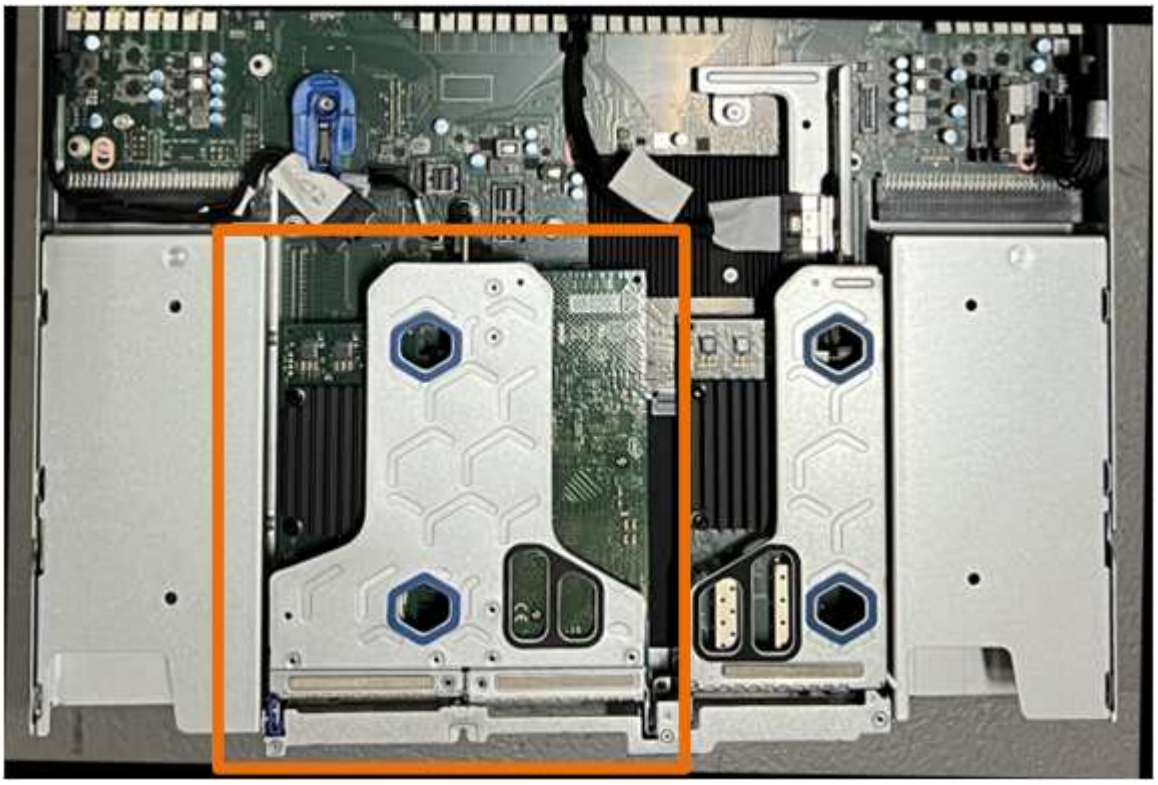

3. Grasp the riser assembly through the blue-marked holes and carefully lift it upwards. Move the riser assembly toward the front of the chassis as you lift it to allow the external connectors in its installed NICs to clear the chassis.
4. Place the riser on a flat anti-static surface with the metal frame side down.

5. Locate the CMOS battery on the system board in the position beneath the removed riser assembly.

6. Use your finger or a plastic pry tool to press the retaining clip (highlighted) away from the battery to spring it from the socket.

7. Remove the battery and dispose of it properly.

Reinstall the CMOS battery

Install the replacement CMOS battery into the socket on the system board.

Before you begin

- You have the correct replacement CMOS battery (CR2032).
- You have removed the failed CMOS battery.

Steps

1. Wrap the strap end of the ESD wristband around your wrist, and secure the clip end to a metal ground to prevent static discharge.
2. Remove the CMOS battery from its packaging.
3. Press the replacement battery into the empty socket on the system board with the positive (+) side up until the battery snaps in place.
4. Locate the alignment hole on the two-slot riser assembly (circled) that aligns with the guide pin on the system board to ensure correct riser assembly positioning.

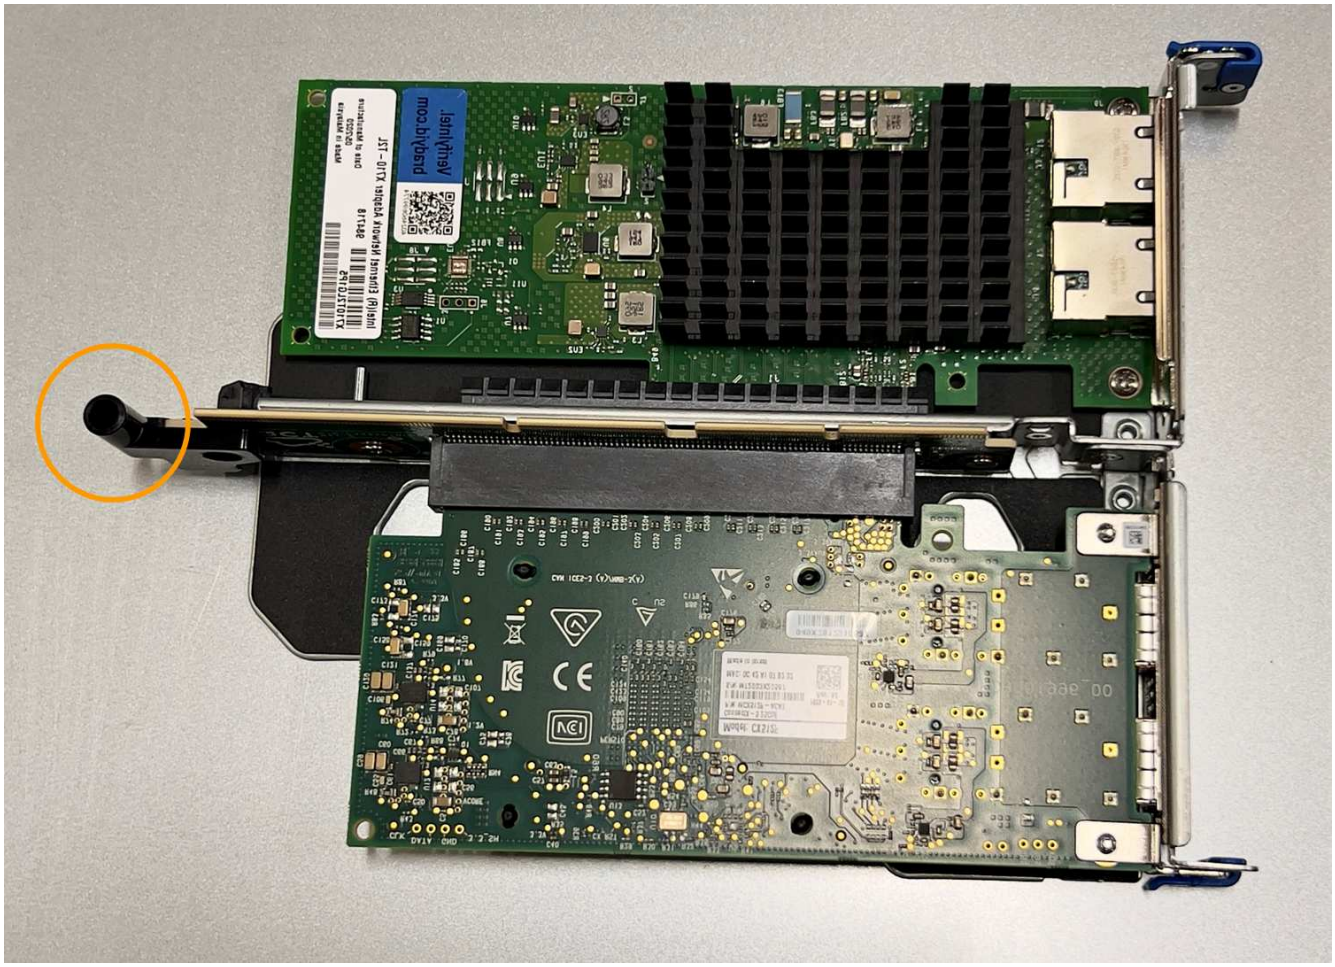

5. Locate the guide pin on the system board

6. Position the riser assembly in the chassis, making sure that it aligns with the connector on the system board and guide pin.
7. Carefully press the two-slot riser assembly in place along its center line, next to the blue-marked holes, until it is fully seated.
8. If you have no other maintenance procedures to perform in the appliance, reinstall the appliance cover, return the appliance to the rack, attach cables, and apply power.
9. If the appliance where you replaced the appliance used a key management server (KMS) to encrypt data, additional configuration might be required before the node can join the grid. If the node does not automatically join the grid, make sure that these configuration settings have transferred to the new appliance and manually configure any settings that don't have the expected configuration:
   - Configure StorageGRID connections
   - Configure node encryption for the appliance
10. Log in to the appliance:
    a. Enter the following command: `ssh admin@grid_node_IP`
    b. Enter the password listed in the `Passwords.txt` file.
    c. Enter the following command to switch to root: `su -`
    d. Enter the password listed in the `Passwords.txt` file.
11. Restore BMC network connectivity for the appliance. There are two options:
    - Use static IP, netmask, and gateway
    - Use DHCP to obtain an IP, netmask, and gateway
      a. To restore the BMC configuration to use a static IP, netmask, and gateway, enter the following commands:

```
run-host-command ipmitool lan set 1 ipsrc static
```

```
run-host-command ipmitool lan set 1 ipaddr Appliance_IP

run-host-command ipmitool lan set 1 netmask Netmask_IP

run-host-command ipmitool lan set 1 defgw ipaddr Default_gateway
```

b. To restore the BMC configuration to use DHCP to obtain an IP, netmask, and gateway, enter the following command:

```
run-host-command ipmitool lan set 1 ipsrc dhcp
```

12. After restoring BMC network connectivity, connect to the BMC interface to audit and restore any additional custom BMC configuration you might have applied. For example, you should confirm the settings for SNMP trap destinations and email notifications. See Configure BMC interface.
13. Confirm that the appliance node appears in the Grid Manager and that no alerts appear.

Replace SGF6112 appliance

You might need to replace the appliance if it is not functioning optimally or if it has failed.

Before you begin

- You have a replacement appliance with the same part number as the appliance you are replacing.
- You have labels to identify each cable that is connected to the appliance.
- You have physically located the appliance.

About this task

The StorageGRID node will not be accessible while you replace the appliance. If the appliance is functioning sufficiently, you can perform a controlled shutdown at the start of this procedure.

> If you are replacing the appliance before installing StorageGRID software, you might not be able to access the StorageGRID Appliance Installer immediately after completing this procedure. While you can access the StorageGRID Appliance Installer from other hosts on the same subnet as the appliance, you can't access it from hosts on other subnets. This condition should resolve itself within 15 minutes (when any ARP cache entries for the original appliance time out), or you can clear the condition immediately by purging any old ARP cache entries manually from the local router or gateway.

Steps

1. Display the current configurations of the appliance and record them.
   a. Log in to the appliance to be replaced:
      i. Enter the following command: `ssh admin@grid_node_IP`
      ii. Enter the password listed in the `Passwords.txt` file.
      iii. Enter the following command to switch to root: `su -`
      iv. Enter the password listed in the `Passwords.txt` file.

      When you are logged in as root, the prompt changes from `$` to `#`.

   b. Enter: **`run-host-command ipmitool lan print`** to display the current BMC configurations for the

appliance.

2. Shut down the appliance:

   Shut down SGF6112

3. If any of the network interfaces on this StorageGRID appliance are configured for DHCP, you may need to update the permanent DHCP lease assignments on the DHCP servers to reference the MAC addresses of the replacement appliance, to ensure the appliance is assigned the expected IP addresses. See Update MAC address references.

4. Remove and replace the appliance:

   a. Label the cables and then disconnect the cables and any network transceivers.

      To prevent degraded performance, don't twist, fold, pinch, or step on the cables.

   b. Remove the failed appliance from the cabinet or rack.

   c. Note the position of the replaceable components (two power supplies, eight cooling fans, three NICs, and twelve SSDs) in the failed appliance.

      The twelve drives are in the following positions in the chassis (front of chassis with bezel removed shown):

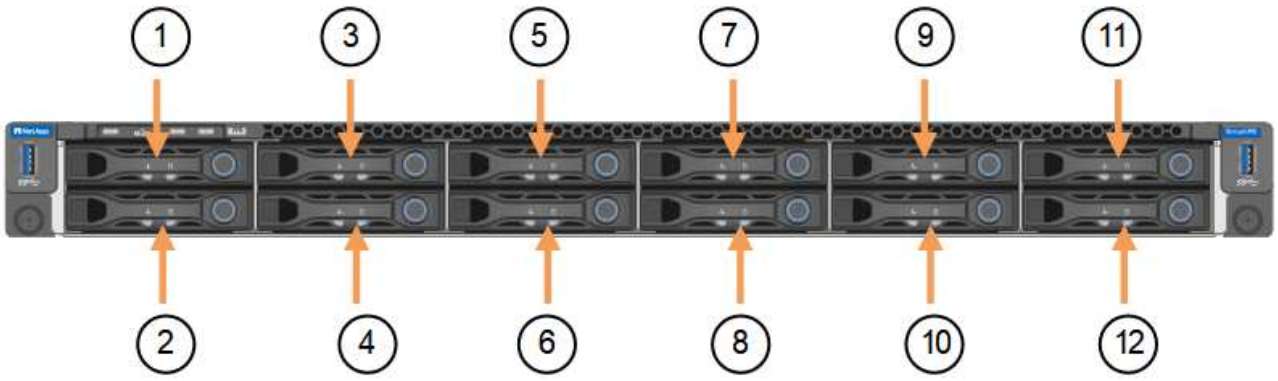

	Drive
1	HDD00
2	HDD01
3	HDD02
4	HDD03
5	HDD04
6	HDD05
7	HDD06

	Drive
8	HDD07
9	HDD08
10	HDD09
11	HDD10
12	HDD11

d. Transfer the replaceable components to the replacement appliance.

Follow the maintenance instructions provided for reinstalling the replaceable components.

> If you wish to retain the data on the drives, be sure to insert the SSD drives into the same drive slots they occupied in the failed appliance. If you don't, the Appliance Installer will display a warning message and you will have to put the drives into the correct slots and reboot the appliance before the appliance can rejoin the grid.

e. Install the replacement appliance into the cabinet or rack.

f. Replace the cables and any optical transceivers.

g. Power on the appliance and wait for it to rejoin the grid. If the appliance does not rejoin the grid, follow the guidance on the StorageGRID Appliance Installer home page to address any issues.

> To prevent data loss if the Appliance Installer indicates that physical hardware changes are required, such as moving disk drives to different slots, power down the appliance before making hardware changes.

5. If the appliance where you replaced the appliance used a key management server (KMS) to encrypt data, additional configuration might be required before the node can join the grid. If the node does not automatically join the grid, make sure that these configuration settings have transferred to the new appliance and manually configure any settings that don't have the expected configuration:
   - Configure StorageGRID connections
   - Configure node encryption for the appliance
6. Log in to the replaced appliance:
   a. Enter the following command: `ssh admin@grid_node_IP`
   b. Enter the password listed in the `Passwords.txt` file.
   c. Enter the following command to switch to root: `su -`
   d. Enter the password listed in the `Passwords.txt` file.
7. Restore BMC network connectivity for the replaced appliance. There are two options:
   - Use static IP, netmask, and gateway
   - Use DHCP to obtain an IP, netmask, and gateway

a. To restore the BMC configuration to use a static IP, netmask, and gateway, enter the following commands:

```
run-host-command ipmitool lan set 1 ipsrc static

run-host-command ipmitool lan set 1 ipaddr Appliance_IP

run-host-command ipmitool lan set 1 netmask Netmask_IP

run-host-command ipmitool lan set 1 defgw ipaddr Default_gateway
```

b. To restore the BMC configuration to use DHCP to obtain an IP, netmask, and gateway, enter the following command:

```
run-host-command ipmitool lan set 1 ipsrc dhcp
```

8. After restoring BMC network connectivity, connect to the BMC interface to audit and restore any additional custom BMC configuration you might have applied. For example, you should confirm the settings for SNMP trap destinations and email notifications. See Configure BMC interface.
9. Confirm that the appliance node appears in the Grid Manager and that no alerts appear.

Related information

Install appliance into a cabinet or rack (SGF6112)

View status indicators

View boot-up codes for appliance

After replacing the part, return the failed part to NetApp, as described in the RMA instructions shipped with the kit. See the Part Return & Replacements page for further information.

Relocate SGF6112 in cabinet or rack

Remove the SGF6112 from a cabinet or rack to access the top cover or to move the appliance to a different location, then reinstall the appliance into a cabinet or rack when hardware maintenance is complete.

Remove SGF6112 from cabinet or rack

Before you begin

- You have labels to identify each cable that is connected to the SGF6112.
- You have physically located the SGF6112 where you are performing maintenance in the data center.

  Locate SGF6112 in data center

- You have shut down the SGF6112.

  Shut down SGF6112

  Don’t shut down the appliance using the power switch.

Steps

1. Label and then disconnect the appliance power cables.
2. Wrap the strap end of the ESD wristband around your wrist, and secure the clip end to a metal ground to prevent static discharge.
3. Label and then disconnect the appliance data cables and any SFP+ or SFP28 transceivers.

   To prevent degraded performance, don't twist, fold, pinch, or step on the cables.

4. Loosen the two captive screws on the appliance front panel.

5. Slide the SGF6112 forward out of the rack until the mounting rails are fully extended and you hear the latches on both sides click.

   The appliance top cover is accessible.

6. Optional: If you are fully removing the appliance from the cabinet or rack, follow the instructions for the rail kit to remove the appliance from the rails.

After replacing the part, return the failed part to NetApp, as described in the RMA instructions shipped with the kit. See the Part Return & Replacements page for further information.

Reinstall SGF6112 into cabinet or rack

Before you begin

You have reinstalled the appliance cover.

Reinstall SGF6112 cover

Steps

1. Press the blue rail releases both rack rails at the same time and slide the SGF6112 into the rack until it is fully seated.

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

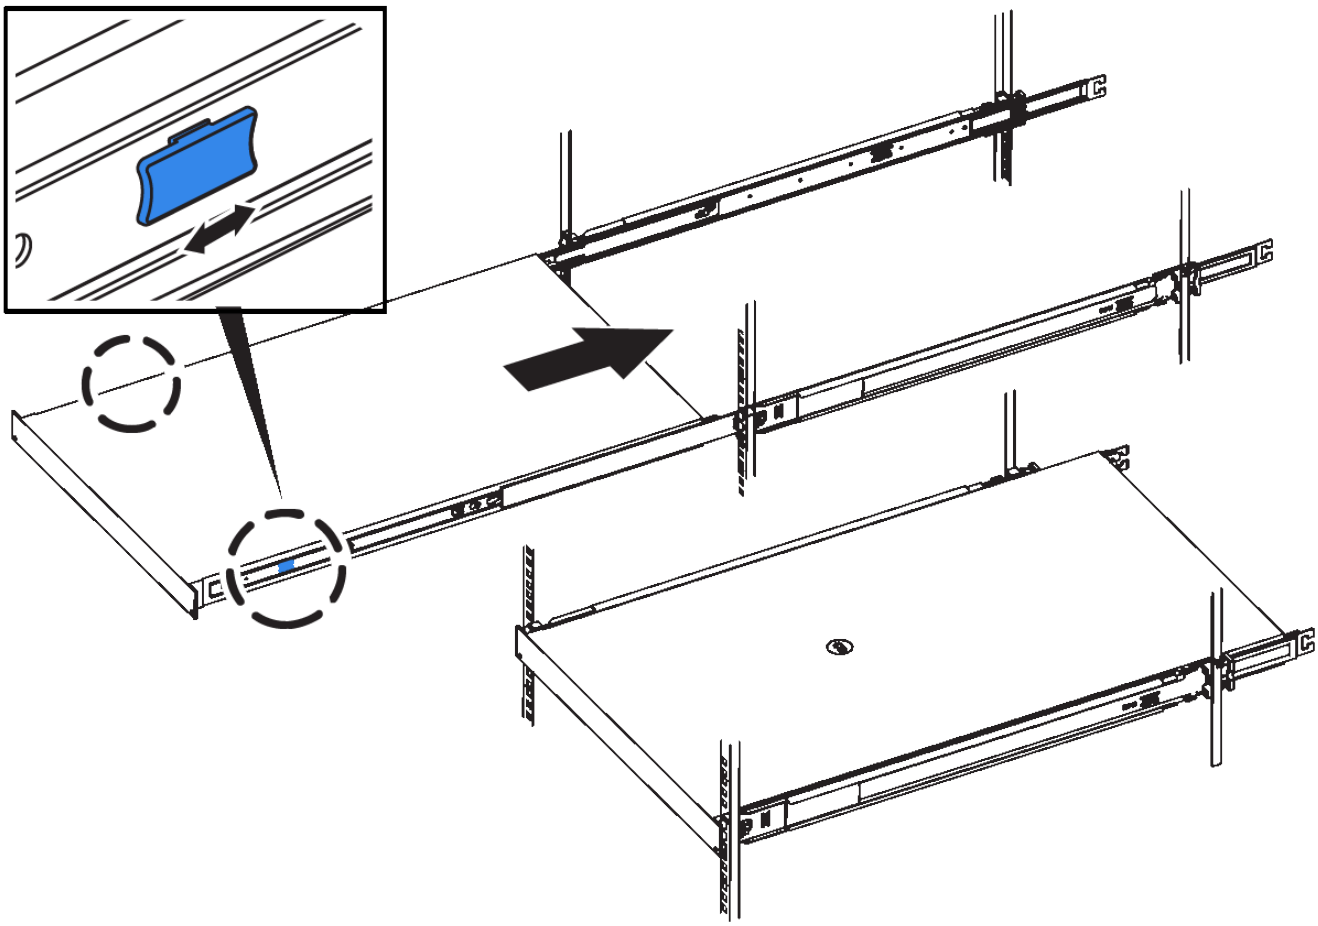

Don’t attach the front bezel until after you power on the controller.

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

3. Wrap the strap end of the ESD wristband around your wrist, and secure the clip end to a metal ground to prevent static discharge.
4. Reconnect the controller data cables and any SFP+ or SFP28 transceivers.

To prevent degraded performance, don’t twist, fold, pinch, or step on the cables.

Cable appliance (SGF6112)

5. Reconnect the controller power cables.

Connect power cords and apply power (SGF6112)

After you finish

The appliance can be restarted.

Power on SGF6112 and verify operation

Replace SGF6112 cover

Remove the appliance cover to access internal components for maintenance, and replace the cover when you are finished.

Remove cover

Before you begin

Remove the appliance from the cabinet or rack to access the top cover.

Remove SGF6112 from cabinet or rack

Steps

1. Make sure that the SGF6112 cover latch is not locked. If necessary, turn the blue plastic latch lock one-quarter turn in the unlock direction, as shown on the latch lock.
2. Rotate the latch up and back toward the rear of the SGF6112 chassis until it stops; then, carefully lift the cover from the chassis and set it aside.

Wrap the strap end of an ESD wristband around your wrist and secure the clip end to a metal ground to prevent static discharge when working inside the SGF6112.

Reinstall cover

Before you begin

You have completed all maintenance procedures inside the appliance.

Steps

1. With the cover latch open, hold the cover above the chassis and align the hole in the top cover latch with the pin in the chassis. When the cover is aligned, lower it onto the chassis.

2. Rotate the cover latch forward and down until it stops and the cover fully seats into the chassis. Verify that there are no gaps along the front edge of the cover.

   If the cover is not fully seated, you might not be able to slide the SGF6112 into the rack.

3. Optional: Turn the blue plastic latch lock one-quarter turn in the lock direction, as shown on the latch lock, to lock it.

After you finish

Reinstall the appliance in the cabinet or rack.

Reinstall SGF6112 into cabinet or rack

Maintain SG6000 hardware

Maintain SG6000 appliance

You might need to perform maintenance procedures on the SG6000 appliance.

Procedures specific to maintaining your SG6000 appliance are in this section and assume that the appliance has already been deployed as a Storage Node in a StorageGRID system.

See Common procedures for maintenance procedures that are used by all appliances.

See Set up hardware for maintenance procedures that are also performed during initial appliance installation and configuration.

To prevent service interruptions, confirm that all other Storage Nodes are connected to the grid before shutting down the appliance or shut down the appliance during a scheduled maintenance window when periods of service disruption are acceptable. See the information about monitoring node connection states.

> If you have ever used an ILM rule that creates only one copy of an object, you must shut down the appliance during a scheduled maintenance window. Otherwise, you might temporarily lose access to those objects during any maintenance procedure that takes a storage node out of service. See the information about managing objects with information lifecycle management.

Maintenance configuration procedures

Upgrade SANtricity OS on storage controllers

To ensure optimal functioning of the storage controller, you must upgrade to the latest maintenance release of the SANtricity OS that is qualified for your StorageGRID appliance.

Consult the NetApp Interoperability Matrix Tool (IMT) to determine which version you should be using.

Download the new SANtricity OS Software file from NetApp Downloads: StorageGRID Appliance.

Use one of the following procedures based on the version of SANtricity OS currently installed:

- If the storage controller is using SANtricity OS 08.42.20.00 (11.42) or newer, use the Grid Manager to perform the upgrade.

  Upgrade SANtricity OS on storage controllers using Grid Manager

- If the storage controller is using a SANtricity OS version older than 08.42.20.00 (11.42), use maintenance mode to perform the upgrade.

  Upgrade SANtricity OS on storage controllers using maintenance mode

> When upgrading the SANtricity OS for your storage appliance, you must follow the instructions in the StorageGRID documentation. If you use any other instructions, your appliance could become inoperable.

Upgrade SANtricity OS on storage controllers using Grid Manager

For storage controllers currently using SANtricity OS 08.42.20.00 (11.42) or newer, you must use the Grid Manager to apply an upgrade.

Before you begin

- You have consulted the NetApp Interoperability Matrix Tool (IMT) to confirm that the SANtricity OS version you are using for the upgrade is compatible with your appliance.
- You have the Maintenance or Root access permission.
- You are signed in to the Grid Manager using a supported web browser.
- You have the provisioning passphrase.
- You have access to NetApp Downloads: StorageGRID Appliance.

About this task

You can't perform other software updates (StorageGRID software upgrade or a hotfix) until you have completed the SANtricity OS upgrade process. If you attempt to start a hotfix or a StorageGRID software upgrade before the SANtricity OS upgrade process has finished, you are redirected to the SANtricity OS upgrade page.

The procedure will not be complete until the SANtricity OS upgrade has been successfully applied to all applicable nodes that have been selected for the upgrade. It might take more than 30 minutes to load the SANtricity OS on each node (sequentially) and up to 90 minutes to reboot each StorageGRID storage appliance. Any nodes in your grid that don't use SANtricity OS will not be affected by this procedure.

The following steps are only applicable when you are using the Grid Manager to perform the upgrade. The storage controllers in the appliance can't be upgraded using the Grid Manager when the controllers are using SANtricity OS older than 08.42.20.00 (11.42).

This procedure will automatically upgrade the NVSRAM to the most recent version associated with the SANtricity OS upgrade. You don't need to apply a separate NVSRAM upgrade file.

Be sure to apply the latest StorageGRID hotfix before you begin this procedure. See StorageGRID hotfix procedure for details.

Steps

1. Download the new SANtricity OS Software file from NetApp Downloads: StorageGRID Appliance.

   Be sure to choose the SANtricity OS version for your storage controllers.

2. Select **MAINTENANCE** > **System** > **Software update**.

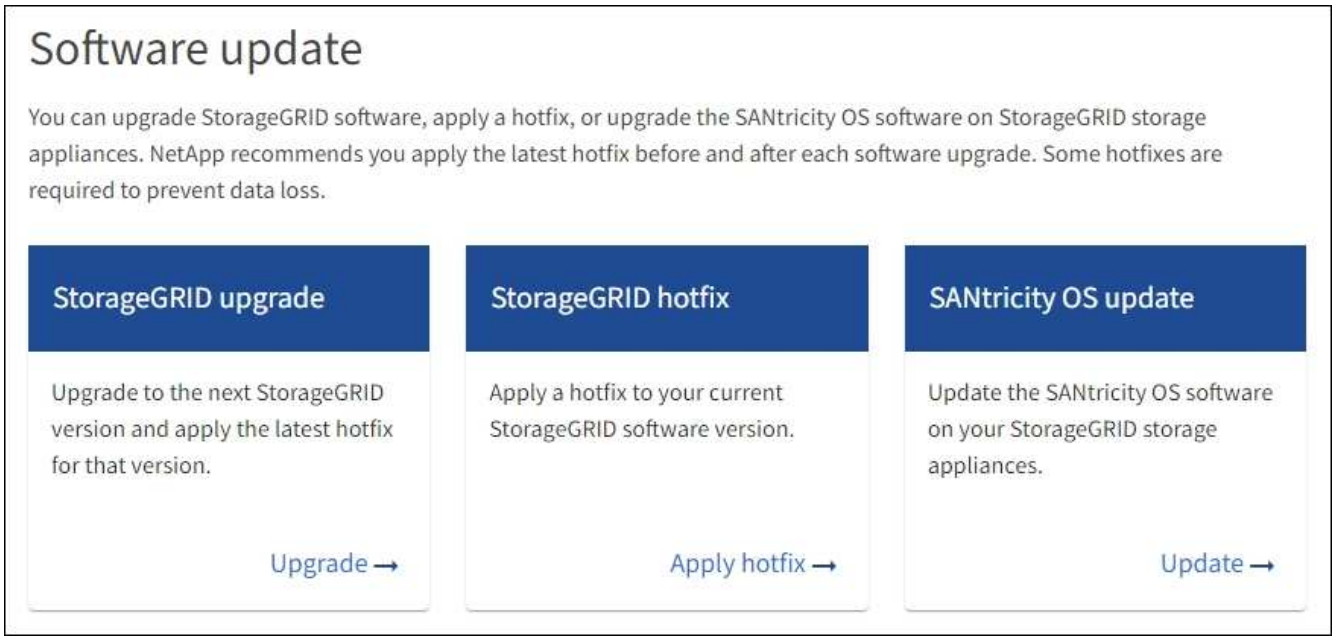

3. In the SANtricity OS update section, select **Update**.

   The SANtricity OS upgrade page appears and lists the details for each appliance node including:

   - Node name
   - Site
   - Appliance model
   - SANtricity OS version
   - Status
   - Last upgrade status

4. Review the information in the table for all of your upgradable appliances. Confirm that all storage controllers have **Nominal** status. If the status for any controller is **Unknown**, go to **Nodes** > ***appliance node*** > **Hardware** to investigate and resolve the issue.

5. Select the SANtricity OS upgrade file you downloaded from the NetApp Support Site.

   a. Select **Browse**.

   b. Locate and select the file.

   c. Select **Open**.

   The file is uploaded and validated. When the validation process is done, the file name is shown with a green check mark next to the **Browse** button. Don't change the file name because it is part of the verification process.

6. Enter the provisioning passphrase and select **Continue**.

   A warning box appears stating that your browser's connection might be lost temporarily as services on nodes that are upgraded are restarted.

7. Select **Yes** to stage the SANtricity OS upgrade file to the primary Admin Node.

   When the SANtricity OS upgrade starts:

   a. The health check is run. This process checks that no nodes have the status of Needs Attention.

      If any errors are reported, resolve them and select **Start** again.

   b. The SANtricity OS Upgrade Progress table appears. This table shows all Storage Nodes in your grid and the current stage of the upgrade for each node.

      The table shows all appliance Storage Nodes. Software-based Storage Nodes aren't displayed. Select **Approve** for all nodes that require the upgrade.

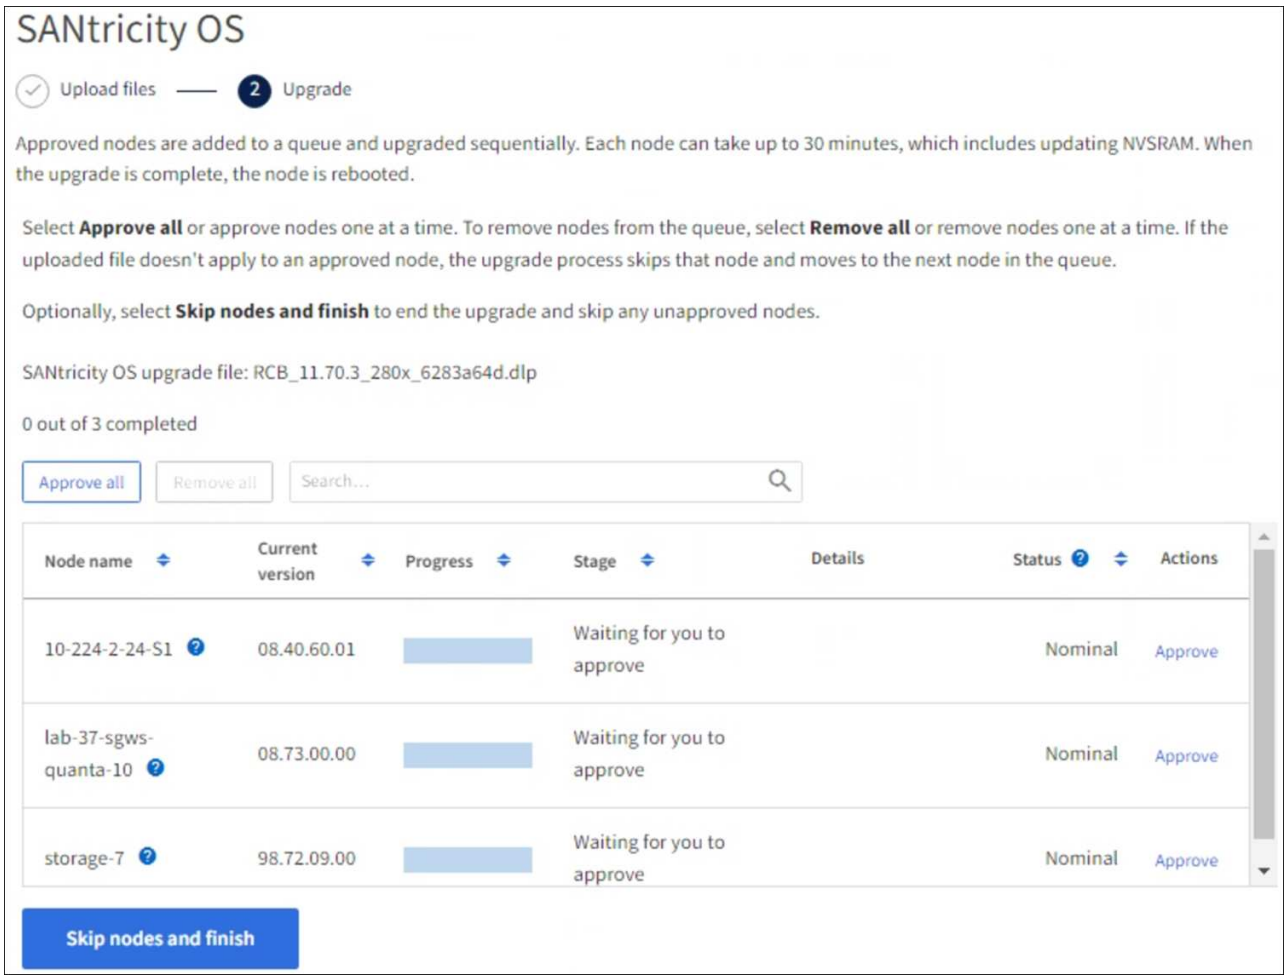

8. Optionally, sort the list of nodes in ascending or descending order by:
   - Node name
   - Current version
   - Progress
   - Stage
   - Status

   You can also enter a term in the Search box to search for specific nodes.

9. Approve the grid nodes you are ready to add to the upgrade queue. Approved nodes are upgraded one at a time.

   > Don’t approve the SANtricity OS upgrade for an appliance Storage Node unless you are sure the node is ready to be stopped and rebooted. When the SANtricity OS upgrade is approved on a node, the services on that node are stopped and the upgrade process begins. Later, when the node is finished upgrading, the appliance node is rebooted. These operations might cause service interruptions for clients that are communicating with the node.

   - Select the **Approve All** button to add all Storage Nodes to the SANtricity OS upgrade queue.

If the order in which nodes are upgraded is important, approve nodes or groups of nodes one at a time and wait until the upgrade is complete on each node before approving the next node.

- Select one or more **Approve** buttons to add one or more nodes to the SANtricity OS upgrade queue. The **Approve** button is disabled if the Status is not Nominal.

  After you select **Approve**, the upgrade process determines if the node can be upgraded. If a node can be upgraded, it is added to the upgrade queue.

  For some nodes, the selected upgrade file is intentionally not applied and you can complete the upgrade process without upgrading these specific nodes. Nodes intentionally not upgraded show a stage of Complete (upgrade attempted) and list the reason the node was not upgraded in the Details column.

10. If you need to remove a node or all nodes from the SANtricity OS upgrade queue, select **Remove** or **Remove All**.

    When the stage progresses beyond Queued, the **Remove** button is hidden and you can no longer remove the node from the SANtricity OS upgrade process.

11. Wait while the SANtricity OS upgrade is applied to each approved grid node.

    - If any node shows a stage of Error while the SANtricity OS upgrade is applied, the upgrade has failed for the node. With the assistance of technical support, you might need to place the appliance in maintenance mode to recover it.
    - If the firmware on the node is too old to be upgraded with the Grid Manager, the node shows a stage of Error with the details that you must use maintenance mode to upgrade SANtricity OS on the node. To resolve the error, do the following:
      a. Use maintenance mode to upgrade SANtricity OS on the node that shows a stage of Error.
      b. Use the Grid Manager to restart and complete the SANtricity OS upgrade.

    When the SANtricity OS upgrade is complete on all approved nodes, the SANtricity OS Upgrade Progress table closes and a green banner shows the number of nodes upgraded, and the date and time the upgrade completed.

12. If a node can’t be upgraded, note the reason shown in the Details column and take the appropriate action.

    The SANtricity OS upgrade process will not be complete until you approve the SANtricity OS upgrade on all the listed Storage Nodes.

Reason	Recommended action
Storage Node was already upgraded.	No further action required.
SANtricity OS upgrade is not applicable to this node.	The node does not have a storage controller that can be managed by the StorageGRID system. Complete the upgrade process without upgrading the node displaying this message.

Reason	Recommended action
SANtricity OS file is not compatible with this node.	The node requires a SANtricity OS file different than the one you selected. After completing the current upgrade, download the correct SANtricity OS file for the node and repeat the upgrade process.

13. If you want to end approving nodes and return to the SANtricity OS page to allow for an upload of a new SANtricity OS file, do the following:
    a. Select **Skip Nodes and Finish**.

       A warning appears asking if you are sure you want to finish the upgrade process without upgrading all applicable nodes.

    b. Select **OK** to return to the **SANtricity OS** page.
    c. When you are ready to continue approving nodes, download the SANtricity OS to restart the upgrade process.

       Nodes already approved and upgraded without errors remain upgraded.

14. Repeat this upgrade procedure for any nodes with a stage of Complete that require a different SANtricity OS upgrade file.

    For any nodes with a status of Needs Attention, use maintenance mode to perform the upgrade.

Related information

NetApp Interoperability Matrix Tool

Upgrade SANtricity OS on storage controllers using maintenance mode

Upgrade SANtricity OS on storage controllers using maintenance mode

For storage controllers currently using SANtricity OS older than 08.42.20.00 (11.42), you must use the maintenance mode procedure to apply an upgrade.

Before you begin

- You have consulted the NetApp Interoperability Matrix Tool (IMT) to confirm that the SANtricity OS version you are using for the upgrade is compatible with your appliance.
- If the StorageGRID appliance is running in a StorageGRID system, you have placed the SG6000-CN controller into maintenance mode.

  Maintenance mode interrupts the connection to the storage controller.

About this task

Don’t upgrade the SANtricity OS or NVSRAM in the E-Series controller on more than one StorageGRID appliance at a time.

Upgrading more than one StorageGRID appliance at a time might cause data unavailability, depending on your deployment model and ILM policies.

Steps

1. Confirm that the appliance is in maintenance mode.
2. From a service laptop, access SANtricity System Manager and sign in.
3. Download the new SANtricity OS Software file and NVSRAM file to the management client.

   The NVSRAM is specific to the StorageGRID appliance. Don't use the standard NVSRAM download.

4. Follow the instructions in the *Upgrading SANtricity OS* guide or the SANtricity System Manager online help to upgrade the firmware and NVSRAM.

   Activate the upgrade files immediately. Don't defer activation.

5. If this procedure completed successfully and you have additional procedures to perform while the node is in maintenance mode, perform them now. When you are done, or if you experienced any failures and want to start over, select **Advanced** > **Reboot Controller**, and then select one of these options:
   - Select **Reboot into StorageGRID**
   - Select **Reboot into Maintenance Mode** to reboot the controller with the node remaining in maintenance mode. Select this option if you experienced any failures during the procedure and want to start over. After the node finishes rebooting into maintenance mode, restart from the appropriate step in the procedure that failed.

   It can take up to 20 minutes for the appliance to reboot and rejoin the grid. To confirm that the reboot is complete and that the node has rejoined the grid, go back to the Grid Manager. The Nodes page should display a normal status (green check mark icon to the left of the node name) for the appliance node, indicating that no alerts are active and the node is connected to the grid.

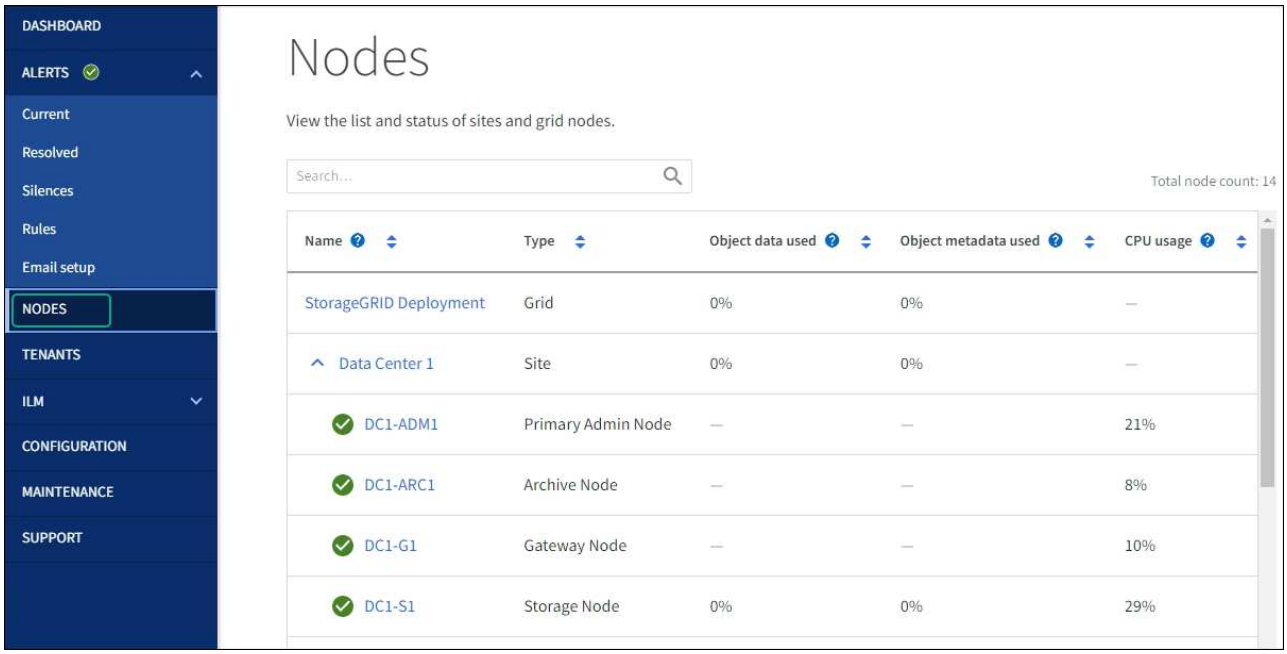

Related information

NetApp Interoperability Matrix Tool

Upgrade SANtricity OS on storage controllers using Grid Manager

Upgrade drive firmware using SANtricity System Manager

Upgrade SG6000 drive firmware using SANtricity System Manager using the online method

Use the SANtricity System Manager online method to upgrade the firmware on the drives in your appliance to make sure you have all the latest features and bug fixes.

Before you begin

- The storage appliance has an Optimal status.
- All drives have an Optimal status.

Don't upgrade the drive firmware on more than one StorageGRID appliance at a time. Doing so might cause data unavailability, depending on your deployment model and ILM policy.

About this task

The drives are upgraded one at a time while the appliance is performing I/O. This method does not require you to place the appliance in maintenance mode. However, system performance might be impacted and the upgrade might take several hours longer than the offline method.

Drives belonging to volumes that don't have redundancy must be updated using the offline method. The offline method should be used for any drive associated with flash read cache (for example, SSD drives in the SG6060), or any pool or volume group that is currently degraded. There are two types of drives: SSD and HDD. You must use the offline method to upgrade the firmware on the SSDs (for example, SSD drives in the SG6060). You can use either the online or offline method to upgrade the firmware on HDDs.

Steps

1. Access SANtricity System Manager using one of these methods:
   - Use the StorageGRID Appliance Installer and select **Advanced** > **SANtricity System Manager**
   - Use the Grid Manager and select **NODES** > ***Storage Node*** > **SANtricity System Manager**
   - Use SANtricity System Manager by browsing to the storage controller IP:

     `https://Storage_Controller_IP`

2. Enter the SANtricity System Manager administrator username and password, if required.
3. Verify the drive firmware version currently installed in the storage appliance:
   a. From SANtricity System Manager, select **SUPPORT** > **Upgrade Center**.
   b. Under Drive Firmware upgrade, select **Begin Upgrade**.

      The Upgrade Drive Firmware page displays the drive firmware files currently installed.

   c. Note the current drive firmware revisions and drive identifiers in the Current Drive Firmware column.

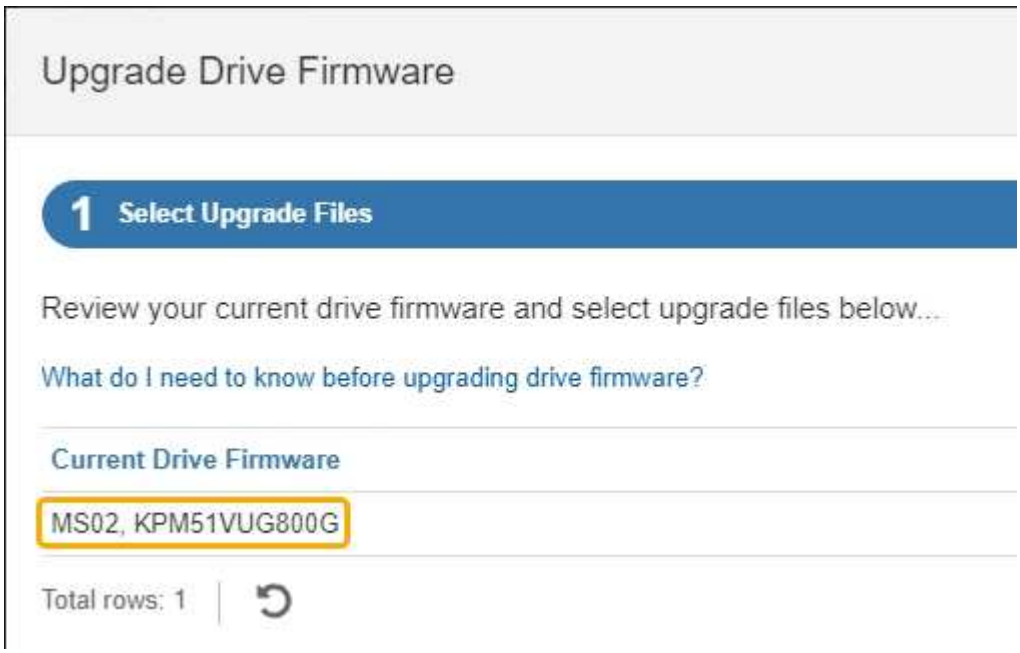

      In this example:

      - The drive firmware revision is **MS02**.
      - The drive identifier is **KPM51VUG800G**.

   d. Select **View drives** in the Associated Drives column to display where these drives are installed in your storage appliance.
   e. Close the Upgrade Drive Firmware window.
4. Download and prepare the available drive firmware upgrade:
   a. Under Drive Firmware upgrade, select **NetApp Support**.
   b. On the NetApp Support Site, select the **Downloads** tab, and then select **E-Series Disk Drive Firmware**.

      The E-Series Disk Firmware page displays.

c. Search for each **Drive Identifier** installed in your storage appliance and verify that each drive identifier has the latest firmware revision.

   - If the firmware revision is not a link, this drive identifier has the latest firmware revision.
   - If one or more drive part numbers are listed for a drive identifier, a firmware upgrade is available for these drives. You can select any link to download the firmware file.

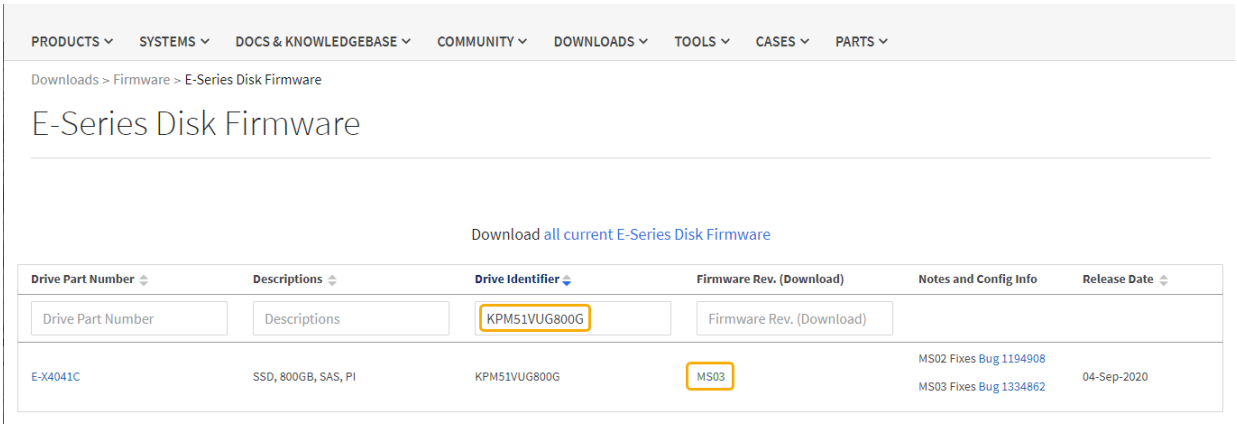

d. If a later firmware revision is listed, select the link in the Firmware Rev. (Download) column to download a `.zip` archive containing the firmware file.

e. Extract (unzip) the drive firmware archive files you downloaded from the Support site.

5. Install the drive firmware upgrade:

   a. From SANtricity System Manager, under Drive Firmware upgrade, select **Begin Upgrade**.

   b. Select **Browse**, and select the new drive firmware files that you downloaded from the Support site.

      Drive firmware files have a filename similar to `D_HUC101212CSS600_30602291_MS01_2800_0002.dlp`.

      You can select up to four drive firmware files, one at a time. If more than one drive firmware file is compatible with the same drive, you get a file conflict error. Decide which drive firmware file you want to use for the upgrade and remove the other one.

   c. Select **Next**.

      Select Drives lists the drives that you can upgrade with the selected firmware files.

      Only drives that are compatible appear.

      The selected firmware for the drive appears in the **Proposed Firmware** column. If you must change this firmware, select **Back**.

   d. Select **Upgrade all drives online** — Upgrades the drives that can support a firmware download while the storage array is processing I/O. You don't have to stop I/O to the associated volumes using these drives when you select this upgrade method.

      An online upgrade can take several hours longer than an offline upgrade.

      You must use the offline method to upgrade the firmware on SSDs.

   e. In the first column of the table, select the drive or drives you want to upgrade.

The best practice is to upgrade all drives of the same model to the same firmware revision.

f. Select **Start** and confirm that you want to perform the upgrade.

If you need to stop the upgrade, select **Stop**. Any firmware downloads currently in progress complete. Any firmware downloads that have not started are canceled.

> Stopping the drive firmware upgrade might result in data loss or unavailable drives.

g. (Optional) To see a list of what was upgraded, select **Save Log**.

The log file is saved in the downloads folder for your browser with the name `latest-upgrade-log-timestamp.txt`.

If required, troubleshoot driver firmware upgrade errors.

Upgrade SG6000 drive firmware using SANtricity System Manager using the offline method

Use the SANtricity System Manager offline method to upgrade the firmware on the drives in your appliance to make sure you have all the latest features and bug fixes.

Before you begin

- The storage appliance has an Optimal status.
- All drives have an Optimal status.
- You have placed the StorageGRID appliance into maintenance mode.

> While the appliance is in maintenance mode, I/O (input/output) activity to the storage controller is stopped to make disruptive storage operations safe.

> Don’t upgrade the drive firmware on more than one StorageGRID appliance at a time. Doing so might cause data unavailability, depending on your deployment model and ILM policy.

About this task

The drives are upgraded in parallel while the appliance is in maintenance mode. If the pool or volume group does not support redundancy or is degraded, you must use the offline method to upgrade the drive firmware. You should also use the offline method for any drive associated with flash read cache, or any pool or volume group that is currently degraded. The offline method upgrades firmware only while all I/O activity is stopped on the drives to be upgraded. To stop I/O activity, place the node into maintenance mode.

The offline method is faster than the online method and will be significantly faster when many drives in a single appliance need upgrades. However, it requires that nodes be taken out of service, which might require scheduling a maintenance window and monitoring progress. Choose the method that is the best fit for your operational procedures and the number of drives that need to be upgraded.

> There are two types of drives: SSD and HDD. You must use the offline method to upgrade the firmware on the SSDs (for example, SSD drives in the SG6060). You can use either the online or offline method to upgrade the firmware on HDDs.

Steps

1. Confirm that the appliance is in maintenance mode.

   > If you are upgrading the firmware in SSD drives that are part of a cache group, you must ensure that no I/O is sent to any cached volumes while the upgrade is in progress. When the appliance is in maintenance mode, no I/O is sent to any volumes while the upgrade is in progress.

2. Access SANtricity System Manager using one of these methods:
   - Use the StorageGRID Appliance Installer and select **Advanced** > **SANtricity System Manager**
   - Use the Grid Manager and select **NODES** > ***Storage Node*** > **SANtricity System Manager**
   - Use SANtricity System Manager by browsing to the storage controller IP:

     `https://Storage_Controller_IP`

3. Enter the SANtricity System Manager administrator username and password, if required.
4. Verify the drive firmware version currently installed in the storage appliance:
   a. From SANtricity System Manager, select **SUPPORT** > **Upgrade Center**.
   b. Under Drive Firmware upgrade, select **Begin Upgrade**.

      The Upgrade Drive Firmware page displays the drive firmware files currently installed.

   c. Note the current drive firmware revisions and drive identifiers in the Current Drive Firmware column.

      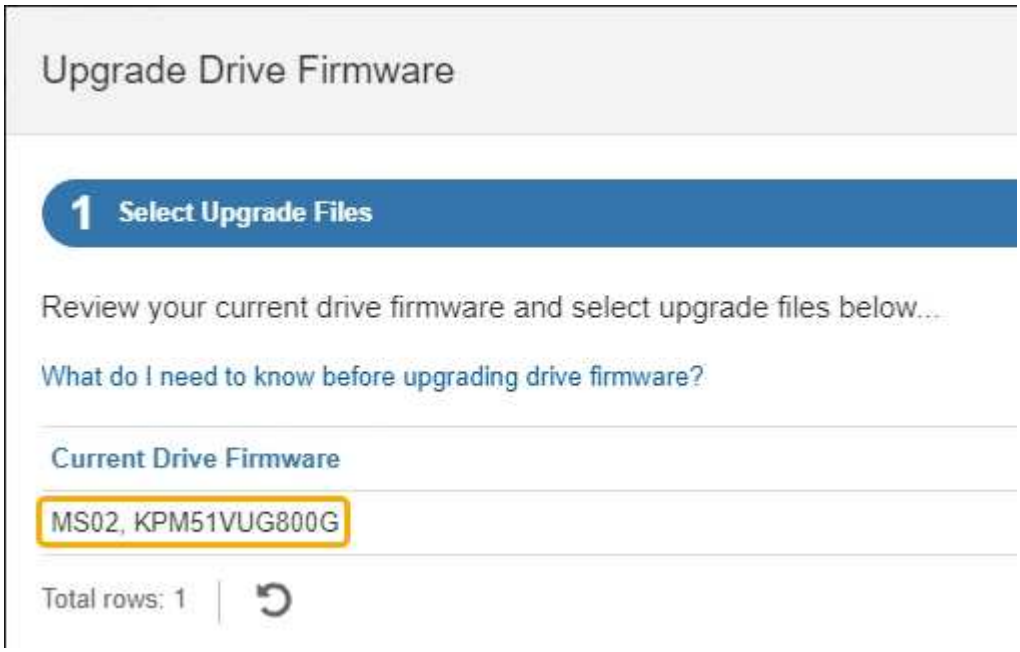

      In this example:

      - The drive firmware revision is **MS02**.
      - The drive identifier is **KPM51VUG800G**.

   d. Select **View drives** in the Associated Drives column to display where these drives are installed in your storage appliance.
   e. Close the Upgrade Drive Firmware window.

5. Download and prepare the available drive firmware upgrade:
   a. Under Drive Firmware upgrade, select **NetApp Support**.
   b. On the NetApp Support Site, select the **Downloads** tab, and then select **E-Series Disk Drive Firmware**.

      The E-Series Disk Firmware page displays.

   c. Search for each **Drive Identifier** installed in your storage appliance and verify that each drive identifier has the latest firmware revision.
      - If the firmware revision is not a link, this drive identifier has the latest firmware revision.
      - If one or more drive part numbers are listed for a drive identifier, a firmware upgrade is available for these drives. You can select any link to download the firmware file.

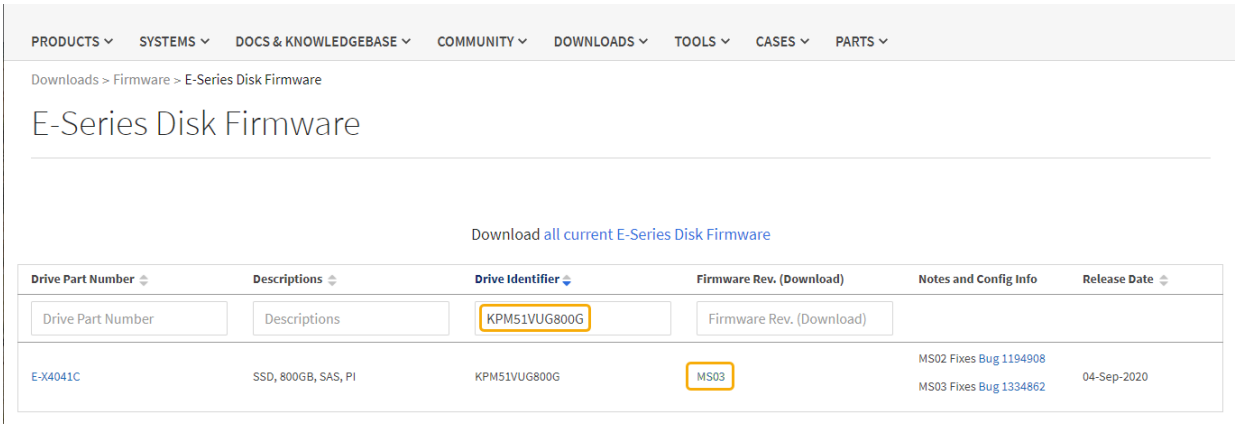

   d. If a later firmware revision is listed, select the link in the Firmware Rev. (Download) column to download a `.zip` archive containing the firmware file.
   e. Extract (unzip) the drive firmware archive files you downloaded from the Support site.
6. Install the drive firmware upgrade:
   a. From SANtricity System Manager, under Drive Firmware upgrade, select **Begin Upgrade**.
   b. Select **Browse**, and select the new drive firmware files that you downloaded from the Support site.

      Drive firmware files have a filename similar to `D_HUC101212CSS600_30602291_MS01_2800_0002.dlp`.

      You can select up to four drive firmware files, one at a time. If more than one drive firmware file is compatible with the same drive, you get a file conflict error. Decide which drive firmware file you want to use for the upgrade and remove the other one.

   c. Select **Next**.

      Select Drives lists the drives that you can upgrade with the selected firmware files.

      Only drives that are compatible appear.

      The selected firmware for the drive appears in the **Proposed Firmware** column. If you must change this firmware, select **Back**.

   d. Select **Upgrade all drives offline (parallel)** — Upgrades the drives that can support a firmware download only while all I/O activity is stopped on any volumes that use the drives.

> You must place the appliance into maintenance mode before using this method. You should use the **Offline** method to upgrade the drive firmware.

> If you want to use the Offline (parallel) upgrade, don't proceed unless you are certain that the appliance is in maintenance mode. Failure to place the appliance into maintenance mode before initiating an offline drive firmware update might cause data loss.

e. In the first column of the table, select the drive or drives you want to upgrade.

   The best practice is to upgrade all drives of the same model to the same firmware revision.

f. Select **Start** and confirm that you want to perform the upgrade.

   If you need to stop the upgrade, select **Stop**. Any firmware downloads currently in progress complete. Any firmware downloads that have not started are canceled.

   > Stopping the drive firmware upgrade might result in data loss or unavailable drives.

g. (Optional) To see a list of what was upgraded, select **Save Log**.

   The log file is saved in the downloads folder for your browser with the name `latest-upgrade-log-timestamp.txt`.

   If required, troubleshoot driver firmware upgrade errors.

7. After the procedure completes successfully, perform any additional maintenance procedures while the node is in maintenance mode. When you are done, or if you experienced any failures and want to start over, go to the StorageGRID Appliance Installer and select **Advanced** > **Reboot Controller**. Then select one of these options:
   - **Reboot into StorageGRID**.
   - **Reboot into Maintenance Mode**. Reboot the controller and keep the node in maintenance mode. Select this option if there were any failures during the procedure and you want to start over. After the node finishes rebooting into maintenance mode, restart from the appropriate step in the procedure that failed.

     It can take up to 20 minutes for the appliance to reboot and rejoin the grid. To confirm that the reboot is complete and that the node has rejoined the grid, go back to the Grid Manager. The Nodes page should display a normal status (green check mark icon to the left of the node name) for the appliance node, indicating that no alerts are active and the node is connected to the grid.

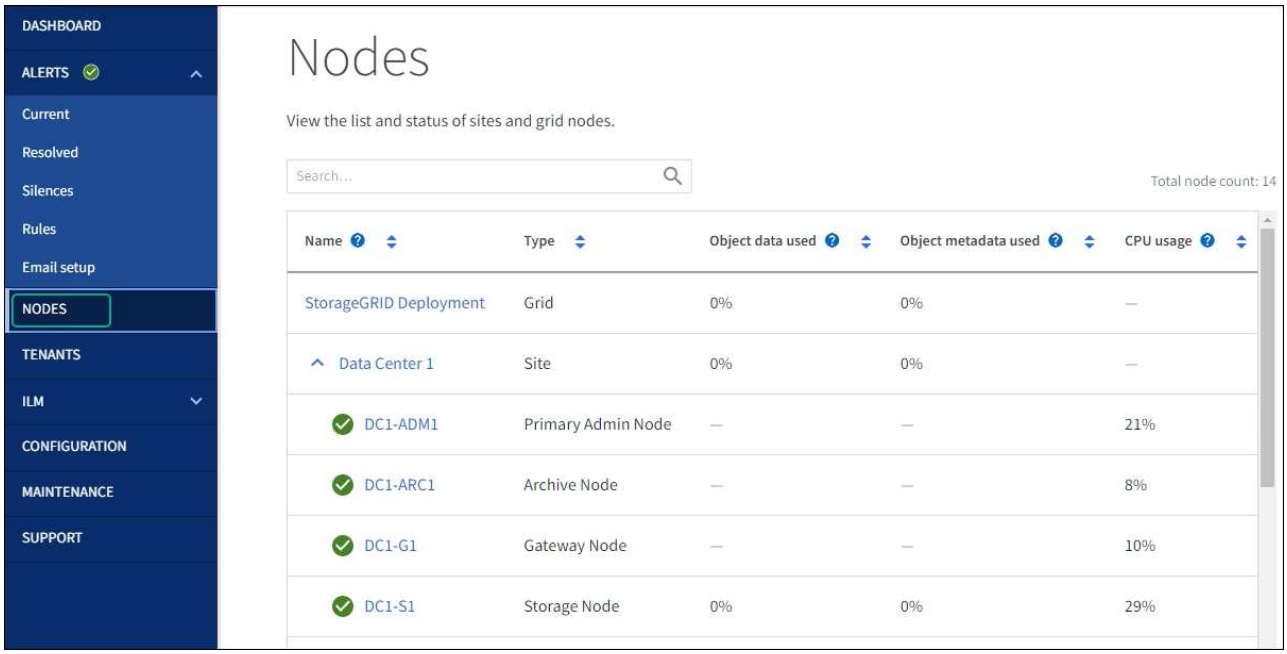

Troubleshoot drive firmware upgrade errors

Troubleshoot errors that can occur when using SANtricity System Manager to upgrade the firmware on the drives in your appliance.

- **Failed assigned drives**
  - One reason for the failure might be that the drive does not have the appropriate signature. Make sure that the affected drive is an authorized drive. Contact technical support for more information.
  - When replacing a drive, make sure that the replacement drive has a capacity equal to or greater than the failed drive you are replacing.
  - You can replace the failed drive while the storage array is receiving I/O.
- **Check storage array**
  - Make sure that an IP address has been assigned to each controller.
  - Make sure that all cables connected to the controller aren't damaged.
  - Make sure that all cables are tightly connected.
- **Integrated hot spare drives**

  This error condition must be corrected before you can upgrade the firmware.

- **Incomplete volume groups**

  If one or more volume groups or disk pools are incomplete, you must correct this error condition before you can upgrade the firmware.

- **Exclusive operations (other than background media/parity scan) currently running on any volume groups**

  If one or more exclusive operations are in progress, the operations must complete before the firmware can be upgraded. Use System Manager to monitor the progress of the operations.

- **Missing volumes**

  You must correct the missing volume condition before the firmware can be upgraded.

- **Either controller in a state other than Optimal**

  One of the storage array controllers needs attention. This condition must be corrected before the firmware can be upgraded.

- **Mismatched Storage Partition information between Controller Object Graphs**

  An error occurred while validating the data on the controllers. Contact technical support to resolve this issue.

- **SPM Verify Database Controller check fails**

  A storage partitions mapping database error occurred on a controller. Contact technical support to resolve this issue.

- **Configuration Database Validation (If supported by the storage array's controller version)**

  A configuration database error occurred on a controller. Contact technical support to resolve this issue.

- **MEL Related Checks**

  Contact technical support to resolve this issue.

- **More than 10 DDE Informational or Critical MEL events were reported in the last 7 days**

  Contact technical support to resolve this issue.

- **More than 2 Page 2C Critical MEL Events were reported in the last 7 days**

  Contact technical support to resolve this issue.

- **More than 2 Degraded Drive Channel Critical MEL events were reported in the last 7 days**

  Contact technical support to resolve this issue.

- **More than 4 critical MEL entries in the last 7 days**

  Contact technical support to resolve this issue.

Turn controller identify LED on and off

The blue identify LED on the front and back of the controller can be turned on to help locate the appliance in a data center.

Before you begin

You have the BMC IP address of the controller you want to identify.

Steps

1. Access the controller BMC interface.
2. Select **Server Identify**.

The current status of the identify LED is selected.

3. Select **ON** or **OFF**, and then select **Perform Action**.

   When you select **ON**, the blue identify LEDs light on the front (shown) and rear of the appliance.

   If a bezel is installed on the controller, it might be difficult to see the front identify LED.

4. Turn the LED on and off as needed.

Related information

Verify Fibre Channel HBA to replace

Locate controller in data center

Access BMC interface

Locate controller in data center

Locate the controller so that you can perform hardware maintenance or upgrades.

Before you begin

- You have determined which controller requires maintenance.

  (Optional) To help locate the controller in your data center, turn on the blue identify LED.

  Turn controller identify LED on and off

Steps

1. Find the controller requiring maintenance in the data center.
   - Look for a lit blue identify LED on the front or rear of the controller.

     The front identify LED is behind the controller front bezel and might be difficult to see if the bezel is installed.

- Check the tags attached to the front of each controller for a matching part number.

2. Remove the controller front bezel, if one is installed, to access the front panel controls and indicators.
3. Optional: Turn off the blue identify LED if you used it to locate the controller.
   - Press the identify LED switch on the controller front panel.
   - Use the controller BMC interface.

     Turn controller identify LED on and off

Related information

Remove Fibre Channel HBA

Remove SG6000-CN controller from cabinet or rack

Shut down SG6000-CN controller

Power controller on and off

Shut down SG6000-CN controller

Shut down the SG6000-CN controller to perform hardware maintenance.

Before you begin

- You have physically located the SG6000-CN controller requiring maintenance in the data center. See Locate controller in data center.

About this task

To prevent service interruptions, confirm that all other Storage Nodes are connected to the grid before shutting down the controller or shut down the controller during a scheduled maintenance window when periods of service disruption are acceptable. See the information about monitoring node connection states.

> If you have ever used an ILM rule that creates only one copy of an object, you must shut down the controller during a scheduled maintenance window. Otherwise, you might temporarily lose access to those objects during this procedure.
> See information about managing objects with information lifecycle management.

Steps

1. Shut down the SG6000-CN controller.

You must perform a controlled shut down of the appliance by entering the commands specified below. It is a best practice to perform a controlled shutdown when possible to avoid unnecessary alerts, ensure full logs are available, and avoid service disruptions.

a. If you have not already logged into the grid node, log in using PuTTY or another ssh client:
   i. Enter the following command: `ssh admin@grid_node_IP`
   ii. Enter the password listed in the `Passwords.txt` file.
   iii. Enter the following command to switch to root: `su -`
   iv. Enter the password listed in the `Passwords.txt` file.

   When you are logged in as root, the prompt changes from `$` to `#`.

b. Shut down the SG6000-CN controller:
   `shutdown -h now`

   This command might take up to 10 minutes to complete.

2. Use one of the following methods to verify that the SG6000-CN controller is powered off:
   - Look at the blue power LED on the front of the controller and confirm that it is off.

   - Look at the green LEDs on both power supplies in the rear of the controller and confirm that they blink at a regular rate (approximately one blink per second).

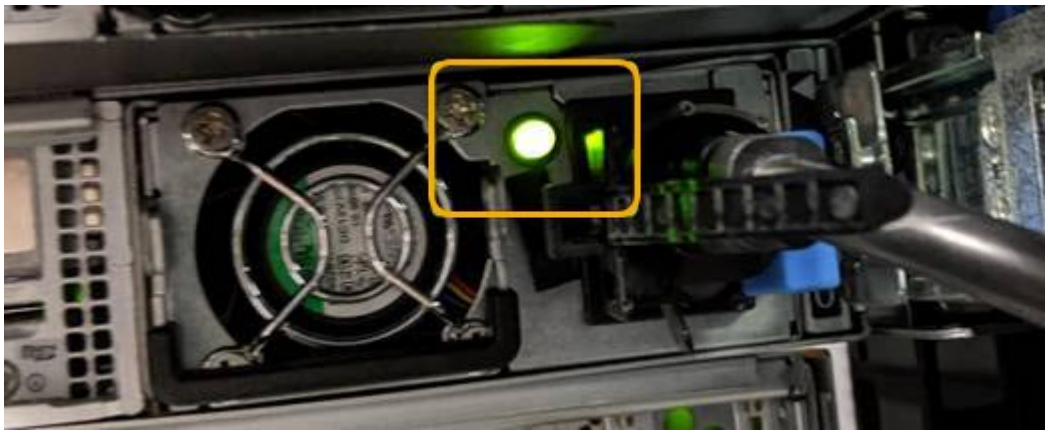

   - Use the controller BMC interface:

i. Access the controller BMC interface.

   Access BMC interface

ii. Select **Power Control**.

iii. Verify that the Power Actions indicates that the host is currently off.

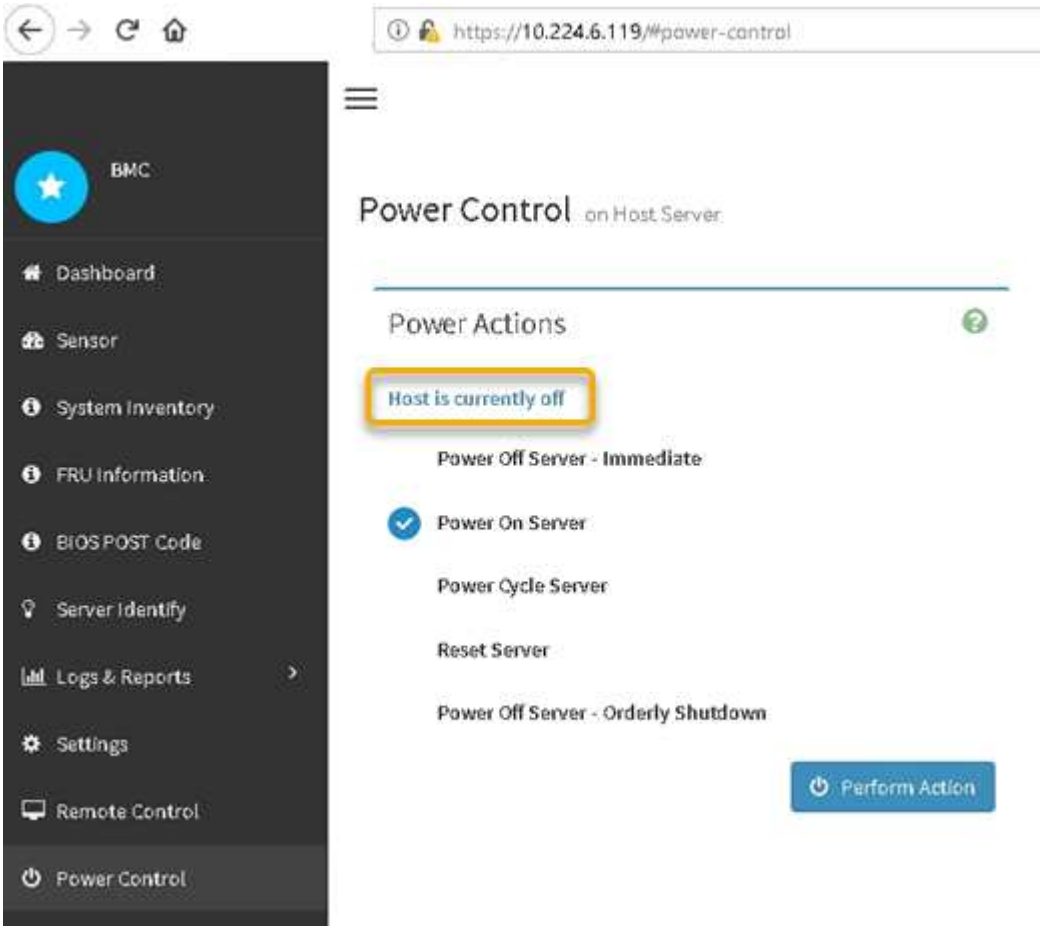

Related information

Remove SG6000-CN controller from cabinet or rack

Power on SG6000-CN controller and verify operation

Power on the controller after completing maintenance.

Before you begin

- You have installed the controller in a cabinet or rack and connected the data and power cables.

  Reinstall SG6000-CN controller into cabinet or rack

- You have physically located the controller in the data center.

  Locate controller in data center

Steps

1. Power on the SG6000-CN controller and monitor the controller LEDs and boot-up codes using one of the

following methods:

- Press the power switch on the front of the controller.

- Use the controller BMC interface:
  i. Access the controller BMC interface.

     Access BMC interface

  ii. Select **Power Control**.
  iii. Select **Power On Server** and then select **Perform Action**.

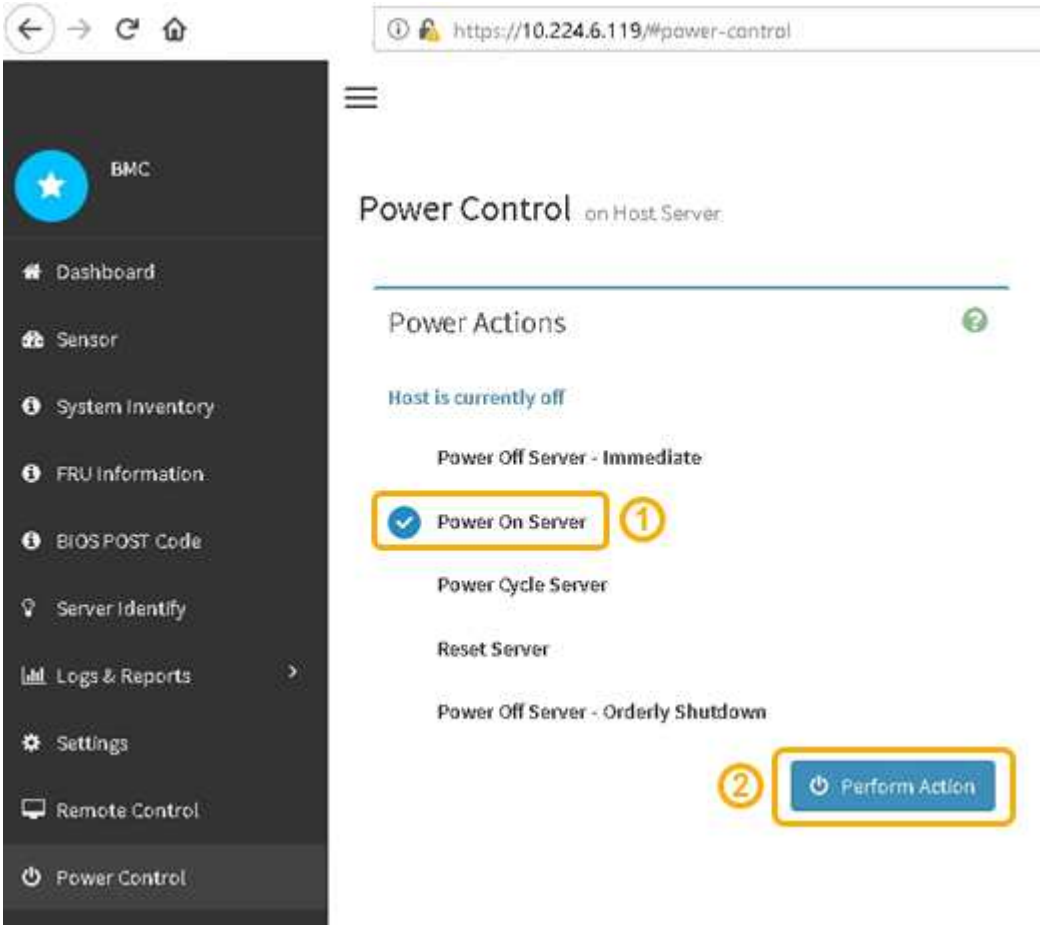

Use the BMC interface to monitor start-up status.

2. Confirm that the appliance controller displays in the Grid Manager and with no alerts.

   It might take up to 20 minutes for the controller to display in the Grid Manager.

3. Confirm that the new SG6000-CN controller is fully operational:
   a. Log in to the grid node using PuTTY or another ssh client:
      i. Enter the following command: `ssh admin@grid_node_IP`
      ii. Enter the password listed in the `Passwords.txt` file.
      iii. Enter the following command to switch to root: `su -`
      iv. Enter the password listed in the `Passwords.txt` file.

      When you are logged in as root, the prompt changes from `$` to `#`.

   b. Enter the following command and verify that it returns the expected output:
   `cat /sys/class/fc_host/*/port_state`

      Expected output:

      ```
      Online
      Online
      Online
      Online
      ```

      If the expected output is not returned, contact technical support.

   c. Enter the following command and verify that it returns the expected output:
   `cat /sys/class/fc_host/*/speed`

      Expected output:

      ```
      16 Gbit
      16 Gbit
      16 Gbit
      16 Gbit
      ```

      If the expected output is not returned, contact technical support.

   d. From the Nodes page in Grid Manager, make sure that the appliance node is connected to the grid and does not have any alerts.

      Don’t take another appliance node offline unless this appliance has a green icon.

4. Optional: Install the front bezel, if one was removed.

Related information

View status indicators

Change link configuration of SG6000-CN controller

You can change the Ethernet link configuration of the SG6000-CN controller. You can change the port bond mode, the network bond mode, and the link speed.

Before you begin

The appliance has been placed maintenance mode.

About this task

Options for changing the Ethernet link configuration of the SG6000-CN controller include:

- Changing **Port bond mode** from Fixed to Aggregate, or from Aggregate to Fixed
- Changing **Network bond mode** from Active-Backup to LACP, or from LACP to Active-Backup
- Enabling or disabling VLAN tagging, or changing the value of a VLAN tag
- Changing the link speed.

Steps

1. From the StorageGRID Appliance Installer, select **Configure Networking** > **Link Configuration**.

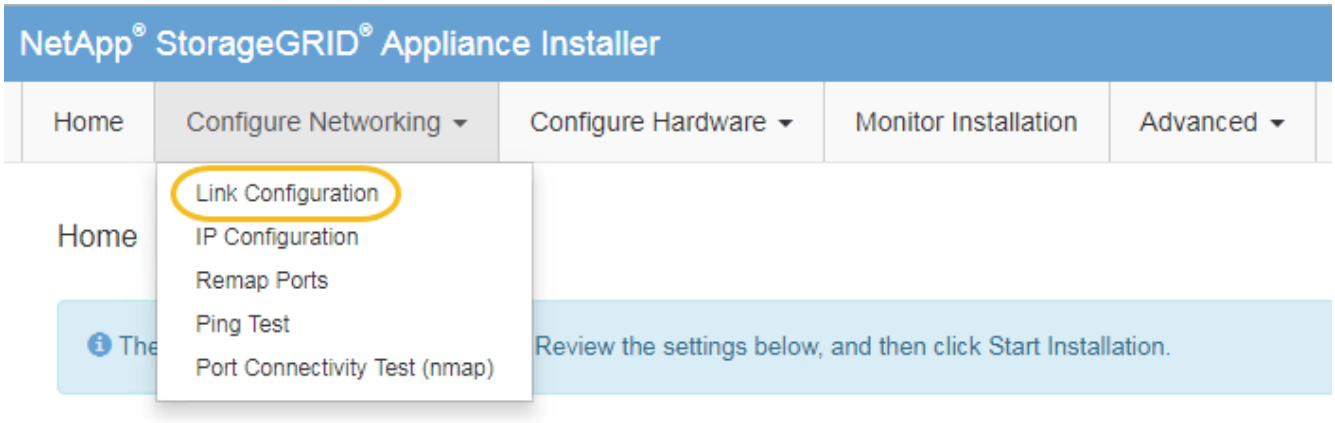

2. Make the desired changes to the link configuration.

   For more information about the options, see Configure network links.

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

   You might lose your connection if you made changes to the network or link you are connected through. If you aren't reconnected within 1 minute, re-enter the URL for the StorageGRID Appliance Installer using one of the other IP addresses assigned to the appliance:
   `https://Appliance_Controller_IP:8443`

   If you made changes to the VLAN settings, the subnet for the appliance might have changed. If you need to change the IP addresses for the appliance, follow the Configure IP addresses instructions.

   Configure StorageGRID IP addresses

4. Select **Configure Networking** > **Ping Test** from the menu.
5. Use the Ping Test tool to check connectivity to IP addresses on any networks that might have been

affected by the link configuration changes you made in the link configuration changes step.

In addition to any other tests you choose to perform, confirm that you can ping the Grid Network IP address of the primary Admin Node, and the Grid Network IP address of at least one other Storage Node. If necessary, return to the link configuration changes step and correct any link configuration issues.

6. When you are satisfied that your link configuration changes are working and you have additional procedures to perform while the node is in maintenance mode, perform them now. When you are done, or if you experienced any failures and want to start over, select **Advanced** > **Reboot Controller**, and then select one of these options:
   - Select **Reboot into StorageGRID**
   - Select **Reboot into Maintenance Mode** to reboot the controller with the node remaining in maintenance mode. Select this option if you experienced any failures during the procedure and want to start over. After the node finishes rebooting into maintenance mode, restart from the appropriate step in the procedure that failed.

It can take up to 20 minutes for the appliance to reboot and rejoin the grid. To confirm that the reboot is complete and that the node has rejoined the grid, go back to the Grid Manager. The **NODES** page should display a normal status (green check mark icon to the left of the node name) for the appliance node, indicating that no alerts are active and the node is connected to the grid.

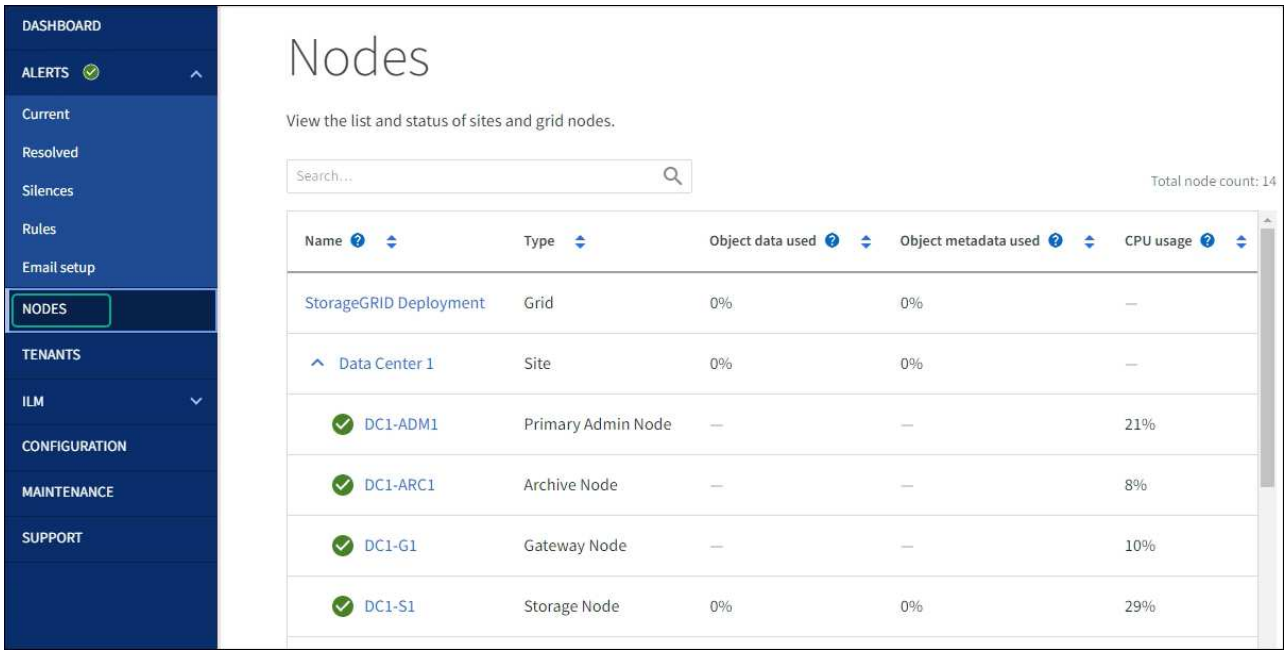

Hardware procedures

Add expansion shelf to deployed SG6060

To increase storage capacity, you can add one or two expansion shelves to an SG6060 that is already deployed in a StorageGRID system.

Before you begin

- You must have the provisioning passphrase.
- You must be running StorageGRID 11.4 or later.
- You have the expansion shelf and two SAS cables for each expansion shelf.
- You have physically located the storage appliance where you are adding the expansion shelf in the data center.

  Locate controller in data center

About this task

To add an expansion shelf, you perform these high-level steps:

- Install the hardware in the cabinet or rack.
- Place the SG6060 into maintenance mode.
- Connect the expansion shelf to the E2860 controller shelf or to another expansion shelf.
- Start the expansion using the StorageGRID Appliance Installer
- Wait until the new volumes are configured.

Completing the procedure for one or two expansion shelves should take one hour or less per appliance node. To minimize downtime, the following steps instruct you to install the new expansion shelves and drives before placing the SG6060 into maintenance mode. The remaining steps should take approximately 20 to 30 minutes per appliance node.

Steps

1. Follow the instructions for installing 60-drive shelves into a cabinet or rack.
2. Follow the instructions for installing the drives.
3. From the Grid Manager, place the SG6000-CN controller into maintenance mode.
4. Connect each expansion shelf to the E2860 controller shelf as shown in the diagram.

   This drawing shows two expansion shelves. If you have only one, connect IOM A to controller A and connect IOM B to controller B.

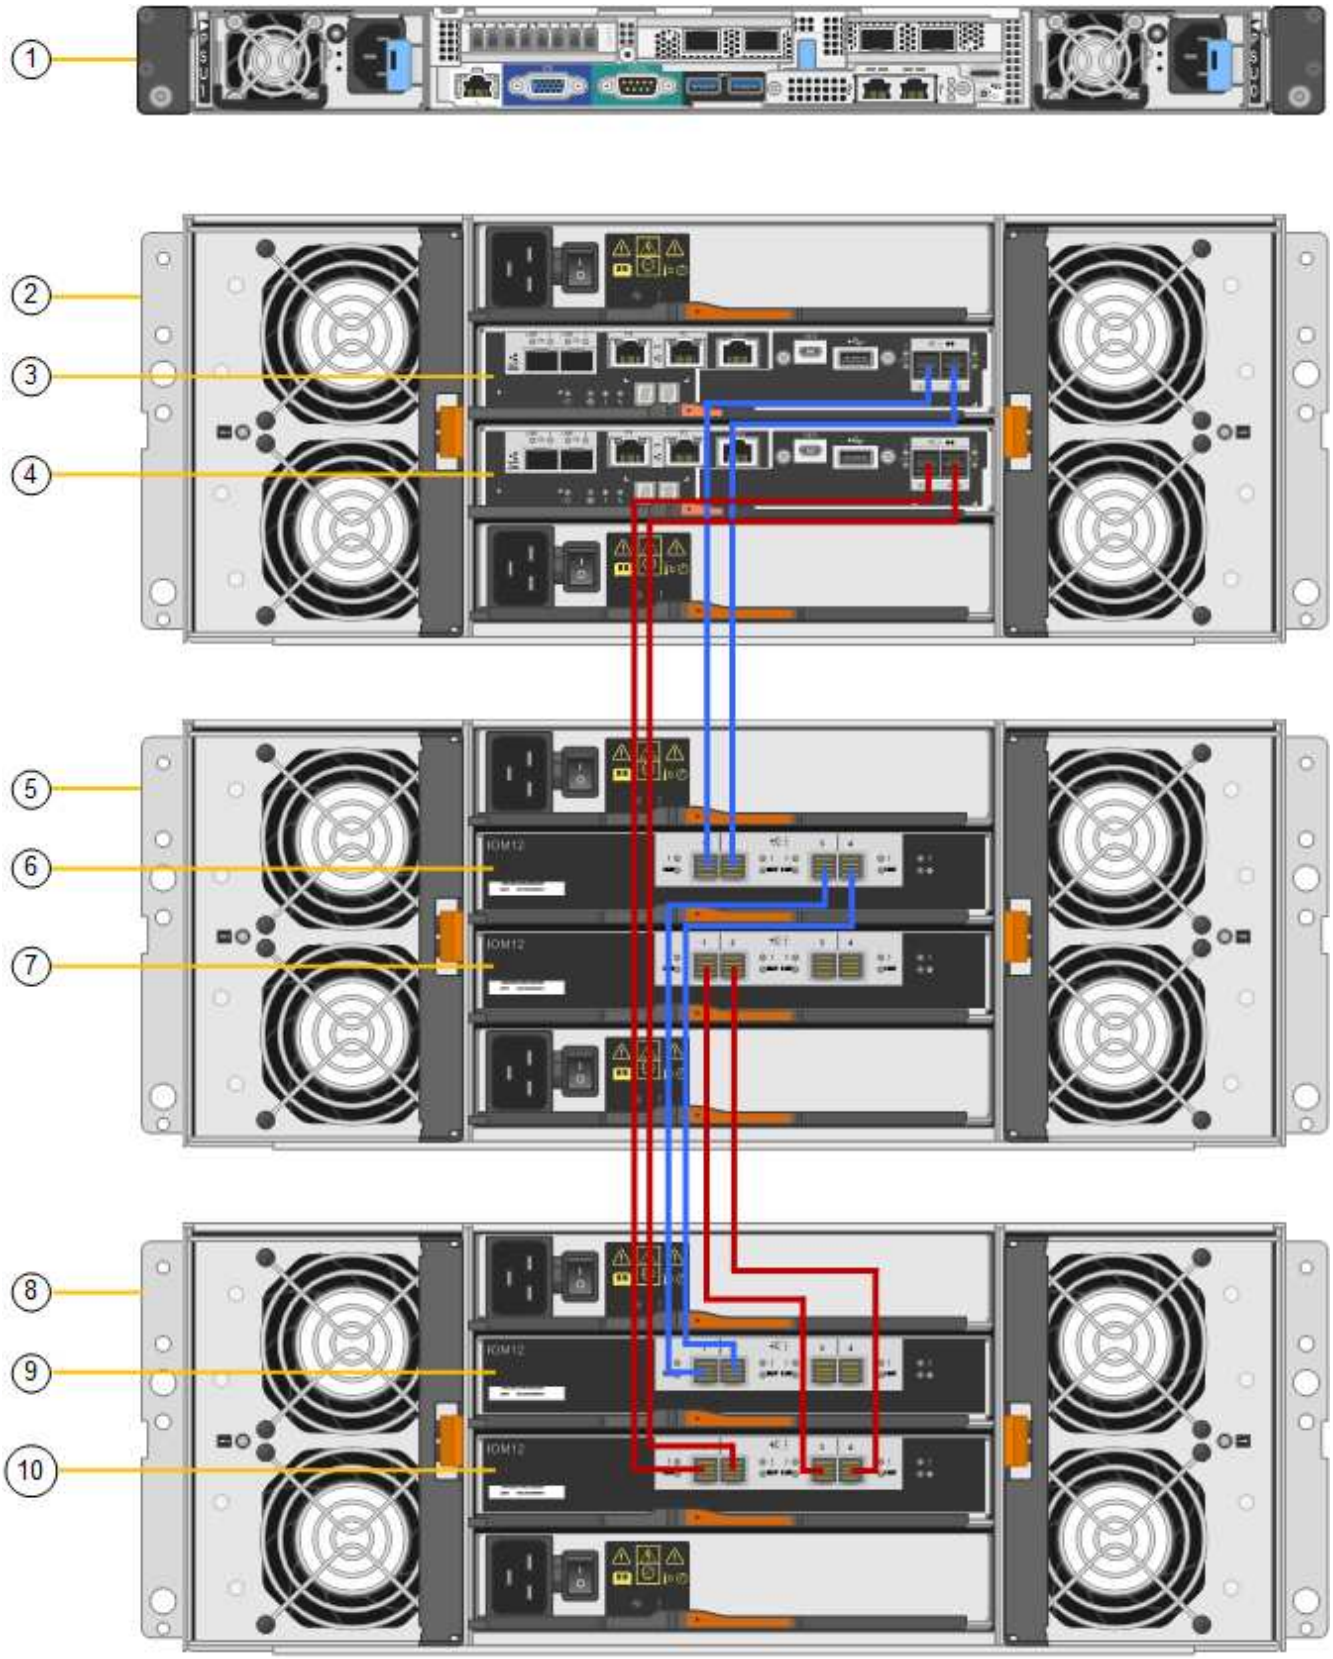

Callout	Description
1	SG6000-CN

Callout	Description
2	E2860 controller shelf
3	Controller A
4	Controller B
5	Expansion shelf 1
6	IOM A for expansion shelf 1
7	IOM B for expansion shelf 1
8	Expansion shelf 2
9	IOM A for expansion shelf 2
10	IOM B for expansion shelf 2

5. Connect the power cords and apply power to the expansion shelves.
   a. Connect a power cord to each of the two power supply units in each expansion shelf.
   b. Connect the two power cords in each expansion shelf to two different PDUs in the cabinet or rack.
   c. Turn on the two power switches for each expansion shelf.
      - Don’t turn off the power switches during the power-on process.
      - The fans in the expansion shelves might be very loud when they first start up. The loud noise during start-up is normal.
6. Monitor the Home page of the StorageGRID Appliance Installer.

   In approximately five minutes, the expansion shelves finish powering up and are detected by the system. The Home page shows the number of new expansion shelves detected, and the Start Expansion button is enabled.

   The screenshot shows examples of the messages that could appear on the Home page, depending on the number of existing or new expansion shelves, as follows:

   - The banner circled at the top of the page indicates the total number of expansion shelves detected.
     - The banner indicates the total number of expansion shelves, whether the shelves are configured and deployed or new and unconfigured.
     - If no expansion shelves are detected, the banner will not appear.
   - The message circled at the bottom of the page indicates an expansion is ready to be started.
     - The message indicates the number of new expansion shelves StorageGRID detects. “Attached” indicates that the shelf is detected. “Unconfigured” indicates that the shelf is new and not yet configured using the StorageGRID Appliance Installer.

Expansion shelves that are already deployed aren't included in this message. They are included in the count in the banner at the top of the page.

- The message will not appear if new expansion shelves aren't detected.

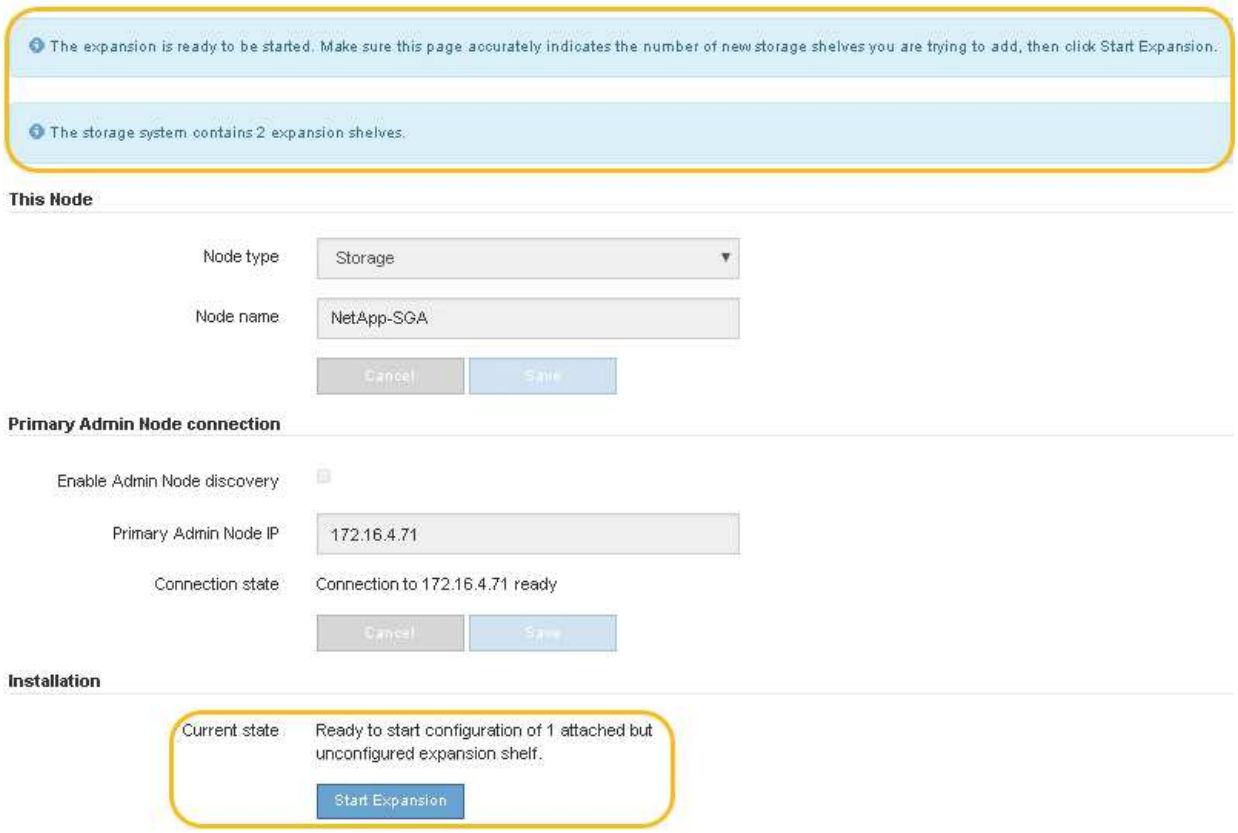

7. As necessary, resolve any issues described in the messages on the Home page.

   For example, use SANtricity System Manager to resolve any storage hardware issues.

8. Verify that the number of expansion shelves displayed on the Home page matches the number of expansion shelves you are adding.

   If the new expansion shelves have not been detected, verify that they are properly cabled and powered up.

9. Click **Start Expansion** to configure the expansion shelves and make them available for object storage.
10. Monitor the progress of the expansion shelf configuration.

    Progress bars appear on the web page, just as they do during initial installation.

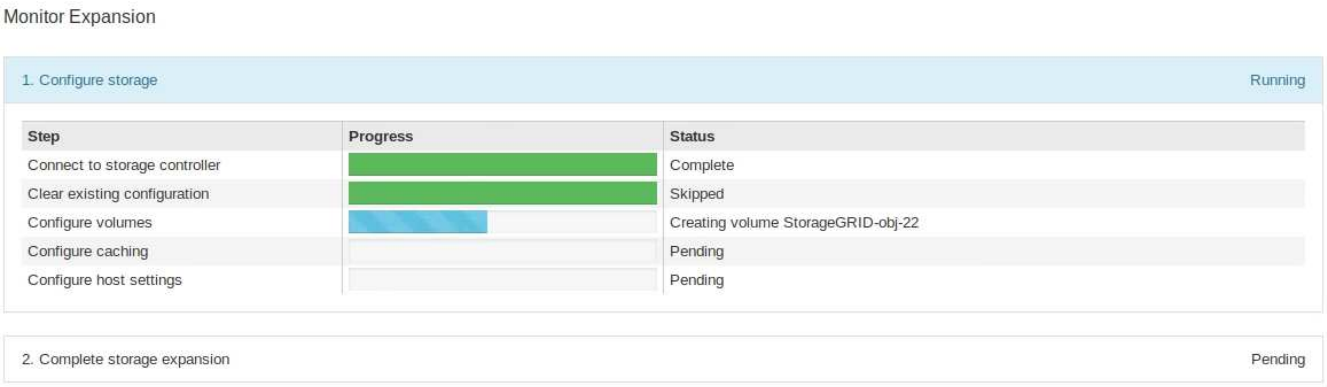

When configuration is complete, the appliance automatically reboots to exit maintenance mode and rejoin the grid. This process can take up to 20 minutes.

> To retry the expansion shelf configuration if it fails, go to the StorageGRID Appliance Installer, select **Advanced** > **Reboot Controller**, and then select **Reboot into Maintenance Mode**. After the node reboots, retry the expansion shelf configuration.

When the reboot is complete, the **Tasks** tab looks like the following screenshot:

11. Verify the status of the appliance Storage Node and the new expansion shelves.
    a. In the Grid Manager, select **NODES** and verify that the appliance Storage Node has a green check mark icon.

       The green check mark icon means that no alerts are active and the node is connected to the grid. For a description of node icons, see Monitor node connection states.

    b. Select the **Storage** tab and confirm that 16 new object stores are shown in the Object Storage table for each expansion shelf you added.
    c. Verify that each new expansion shelf has a shelf status of Nominal and a configuration status of Configured.

Replace storage controller in the SG6000

You might need to replace an E2800 series controller or an EF570 controller if it is not functioning optimally or if it has failed.

Before you begin

- You have a replacement controller with the same part number as the controller you are replacing.
- You have labels to identify each cable that is connected to the controller.
- You have an ESD wristband, or you have taken other antistatic precautions.
- You have a #1 Phillips screwdriver.
- You have physically located the storage appliance where you are replacing the controller in the data center.

  Locate controller in data center

> Don't rely on the E-Series instructions to replace a controller in the StorageGRID appliance, because the procedures aren't the same.

About this task

You can determine if you have a failed controller in two ways:

- The Recovery Guru in SANtricity System Manager directs you to replace the controller.
- The amber Attention LED on the controller is on, indicating that the controller has a fault.

> If both controllers in the shelf have their Attention LEDs on, contact technical support for assistance.

If your appliance contains two storage controllers, you can replace one of the controllers while your appliance is powered on and performing read/write operations, as long as the following conditions are true:

- The second controller in the shelf has Optimal status.
- The **OK to remove** field in the Details area of the Recovery Guru in SANtricity System Manager displays **Yes**, indicating that it is safe to remove this component.

> When possible, place the appliance into maintenance mode for this replacement procedure to minimize the potential impact of unforeseen errors or failures.

> If the second controller in the shelf does not have Optimal status or if the Recovery Guru indicates that it is not OK to remove the controller, contact technical support.

When you replace a controller, you must remove the battery from the original controller and install it in the replacement controller. In some cases, you might also need to remove the host interface card from the original controller and install it in the replacement controller.

> The storage controllers in most appliance models don't include host interface cards (HIC).

Step 1: Prepare the replacement controller

Prepare the replacement E2800A or E2800B controller.

Steps

1. Unpack the new controller, and set it on a flat, static-free surface.

   Save the packing materials to use when shipping the failed controller.

2. Locate the MAC address and FRU part number labels on the back of the replacement controller.

These figures show the E2800A controller and the E2800B controller. The procedure for replacing the E2800 series controllers and the EF570 controller is identical.

E2800A storage controller:

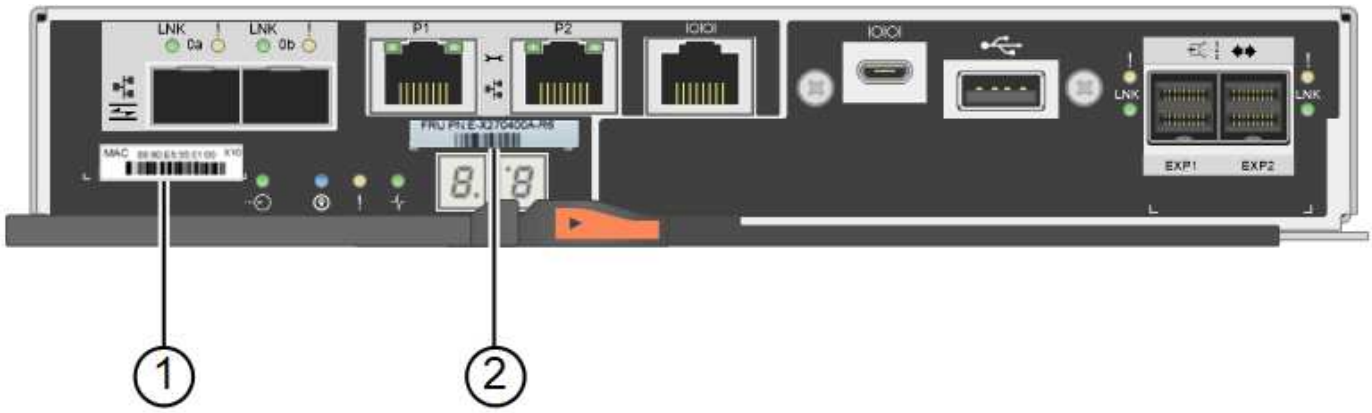

E2800B storage controller:

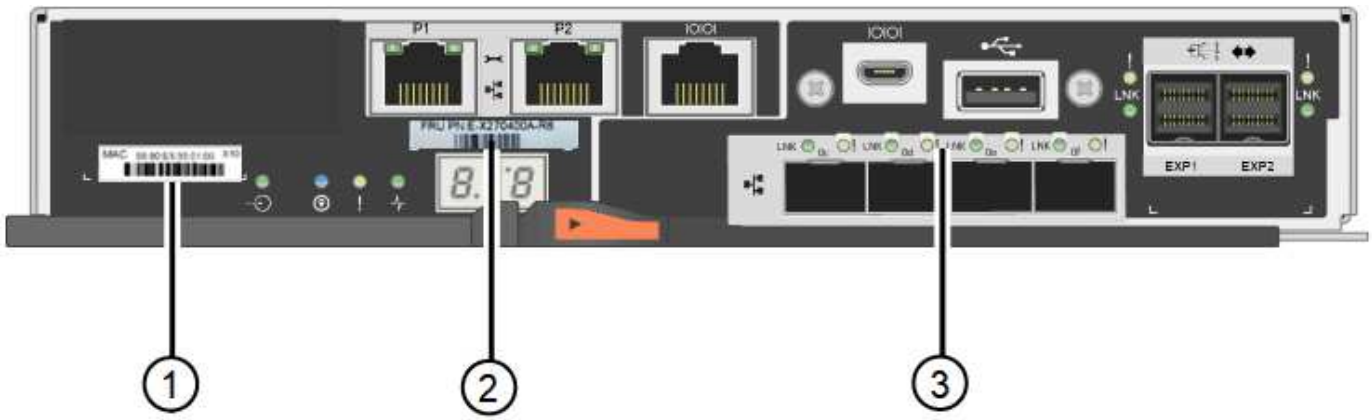

Label	component	Description
1	MAC address	The MAC address for management port 1 (“P1 on the E2800A and 0a on the E2800B”). If you used DHCP to obtain the original controller’s IP address, you will need this address to connect to the new controller.
2	FRU part number	The FRU part number. This number must match the replacement part number for the currently installed controller.
3	4-port HIC	The 4-port host interface card (HIC). This card must be moved to the new controller when you perform the replacement. **Note**: the E2800A controller does not have a HIC.

Ste 2: Take controller offline

Prepare to remove the failed controller and and take it offline.

Steps

1. Prepare to remove the controller. You use SANtricity System Manager to perform these steps.
   a. Confirm that the replacement part number for the failed controller is the same as the FRU part number for the replacement controller.

      When a controller has a fault and needs to be replaced, the replacement part number is displayed in the Details area of the Recovery Guru. If you need to find this number manually, you can look on the **Base** tab for the controller.

      Possible loss of data access — If the two part numbers aren't the same, don't attempt this procedure.

   b. Back up the configuration database.

      If a problem occurs when you remove a controller, you can use the saved file to restore your configuration.

   c. Collect support data for the appliance.

      Collecting support data before and after replacing a component ensures you can send a full set of logs to technical support if the replacement does not resolve the problem.

   d. Take the controller you plan to replace offline.
2. Power off the controller shelf.

Step 3: Remove the controller

Remove the controller from the appliance.

Steps

1. Put on an ESD wristband or take other antistatic precautions.
2. Label the cables and then disconnect the cables and SFPs.

   To prevent degraded performance, don't twist, fold, pinch, or step on the cables.
3. Release the controller from the appliance by squeezing the latch on the cam handle until it releases, and then open the cam handle to the right.
4. Using two hands and the cam handle, slide the controller out of the appliance.

   Always use two hands to support the weight of the controller.
5. Place the controller on a flat, static-free surface with the removable cover facing up.
6. Remove the cover by pressing down on the button and sliding the cover off.

Step 4: Move battery to the new controller

Remove the battery from the failed controller, and install it into the replacement controller.

Steps

1. Confirm that the green LED inside the controller (between the battery and the DIMMs) is off.

   If this green LED is on, the controller is still using battery power. You must wait for this LED to go off before removing any components.

Item	Description
1	Internal Cache Active LED
2	Battery

2. Locate the blue release latch for the battery.
3. Unlatch the battery by pushing the release latch down and away from the controller.

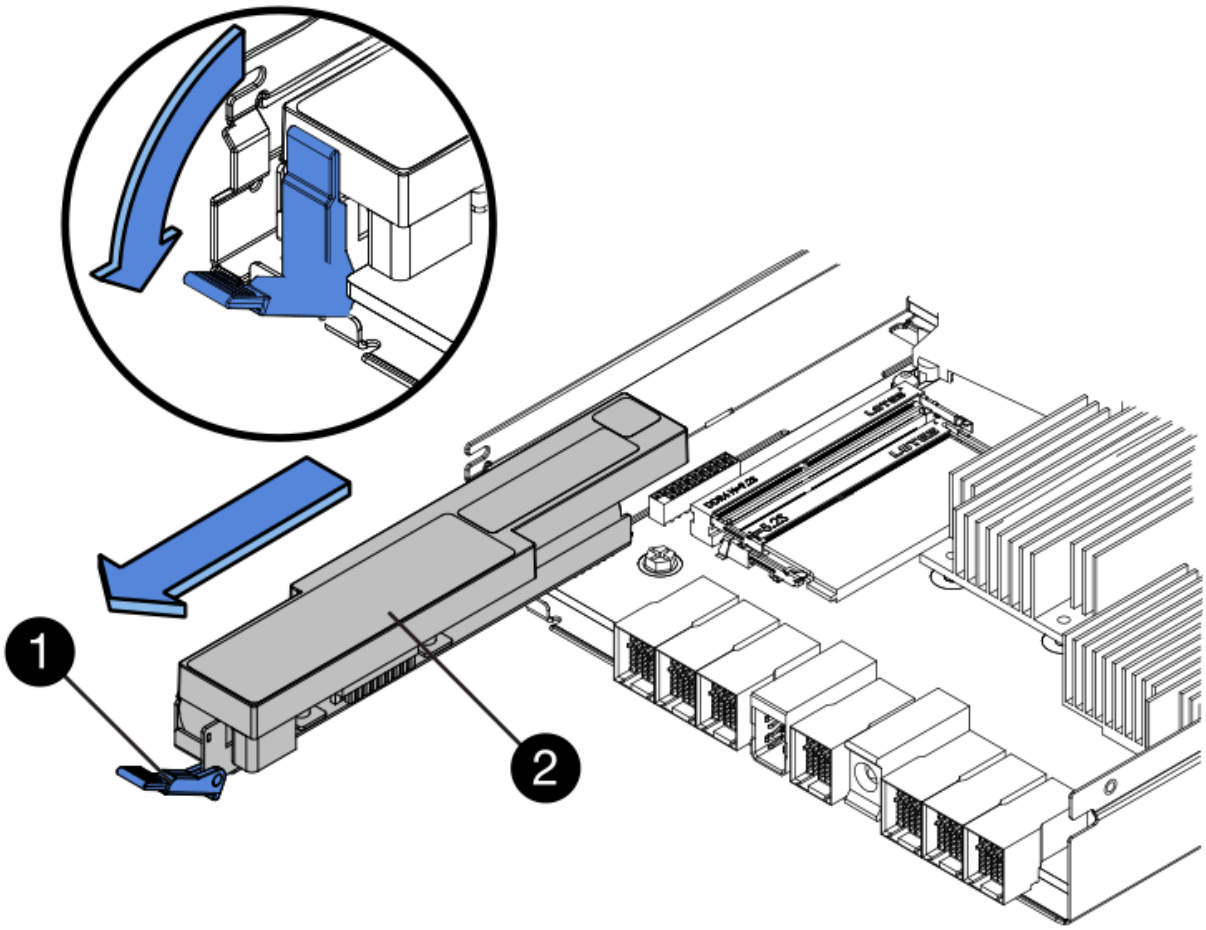

Item	Description
1	Battery release latch
2	Battery

4. Lift up on the battery, and slide it out of the controller.
5. Remove the cover from the replacement controller.
6. Orient the replacement controller so that the slot for the battery faces toward you.
7. Insert the battery into the controller at a slight downward angle.

   You must insert the metal flange at the front of the battery into the slot on the bottom of the controller, and slide the top of the battery beneath the small alignment pin on the left side of the controller.

8. Move the battery latch up to secure the battery.

   When the latch clicks into place, the bottom of the latch hooks into a metal slot on the chassis.

9. Turn the controller over to confirm that the battery is installed correctly.

Possible hardware damage — The metal flange at the front of the battery must be completely inserted into the slot on the controller (as shown in the first figure). If the battery is not installed correctly (as shown in the second figure), the metal flange might contact the controller board, causing damage.

- **Correct — The battery's metal flange is completely inserted into the slot on the controller:**

- **Incorrect — The battery's metal flange is not inserted into the slot on the controller:**

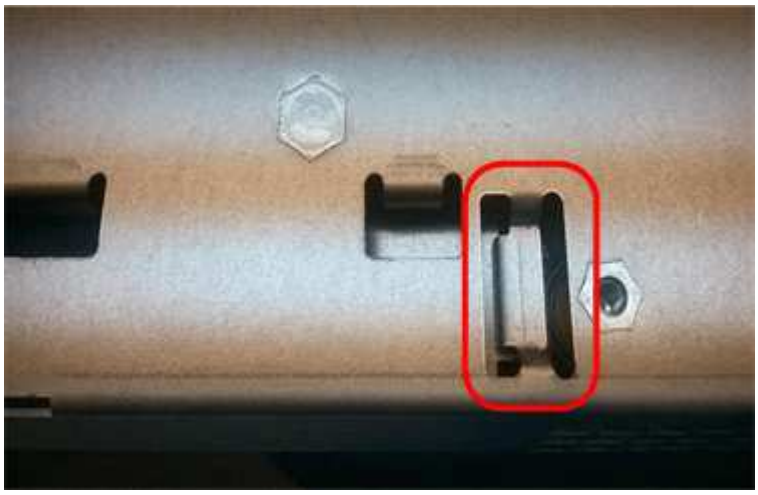

10. Replace the controller cover.

Step5: Move HIC to new controller, if needed

If the failed controller includes a host interface card (HIC), move the HIC from the failed controller to the replacement controller.

A separate HIC is used for the E2800B controller only. The HIC is mounted to the main controller board and includes two SPF connectors.

The illustrations in this procedure show a 2-port HIC. The HIC in your controller might have a different number of ports.

E2800A

An E2800A controller does not have a HIC.

Replace the E2800A controller cover, and go to Step 6: Replace controller

E2800B

Move the HIC from the failed E2800B controller to the replacement controller.

Steps

1. Remove any SFPs from the HIC.
2. Using a #1 Phillips screwdriver, remove the screws that attach the HIC faceplate to the controller.

   There are four screws: one on the top, one on the side, and two on the front.

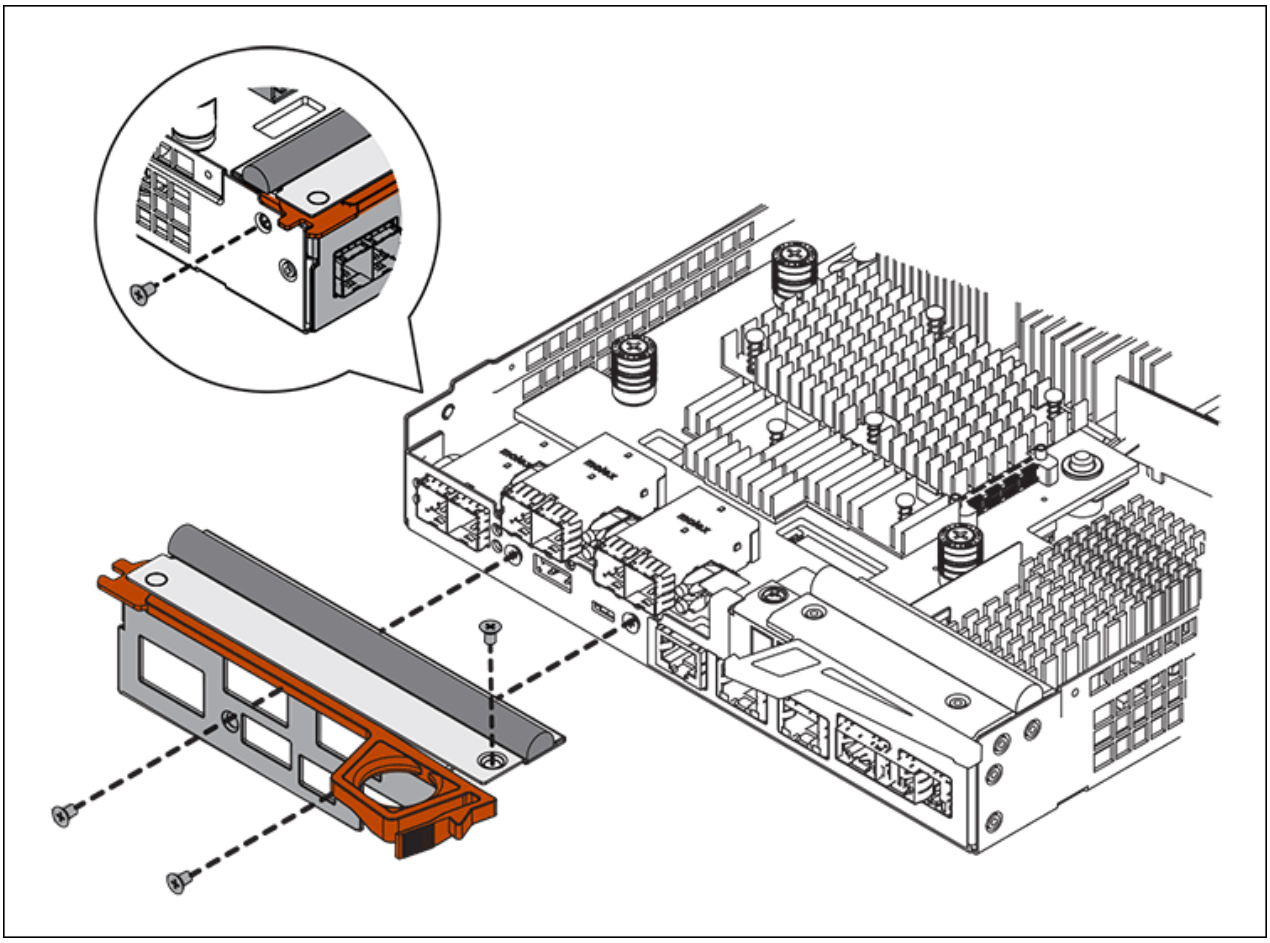

3. Remove the HIC faceplate.
4. Using your fingers or a Phillips screwdriver, loosen the three thumbscrews that secure the HIC to the controller card.
5. Carefully detach the HIC from the controller card by lifting the card up and sliding it back.

> Be careful not to scratch or bump the components on the bottom of the HIC or on the top of the controller card.

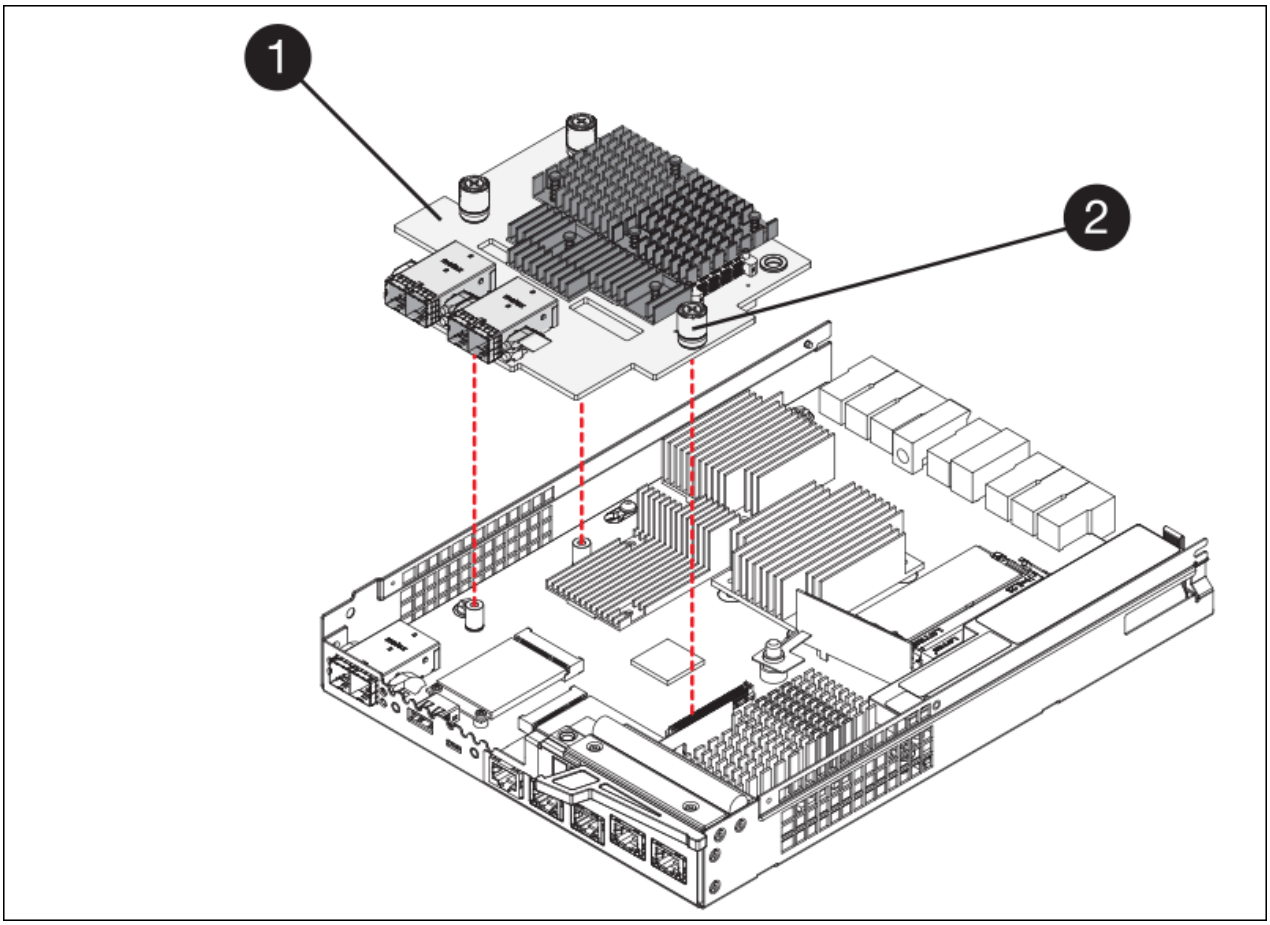

Label	Description
1	Host interface card
2	Thumbscrews

6. Place the HIC on a static-free surface.
7. Using a #1 Phillips screwdriver, remove the four screws that attach the blank faceplate to the replacement controller, and remove the faceplate.
8. Align the three thumbscrews on the HIC with the corresponding holes on the replacement controller, and align the connector on the bottom of the HIC with the HIC interface connector on the controller card.

   Be careful not to scratch or bump the components on the bottom of the HIC or on the top of the controller card.

9. Carefully lower the HIC into place, and seat the HIC connector by pressing gently on the HIC.

Possible equipment damage — Be careful not to pinch the gold ribbon connector for the controller LEDs between the HIC and the thumbscrews.

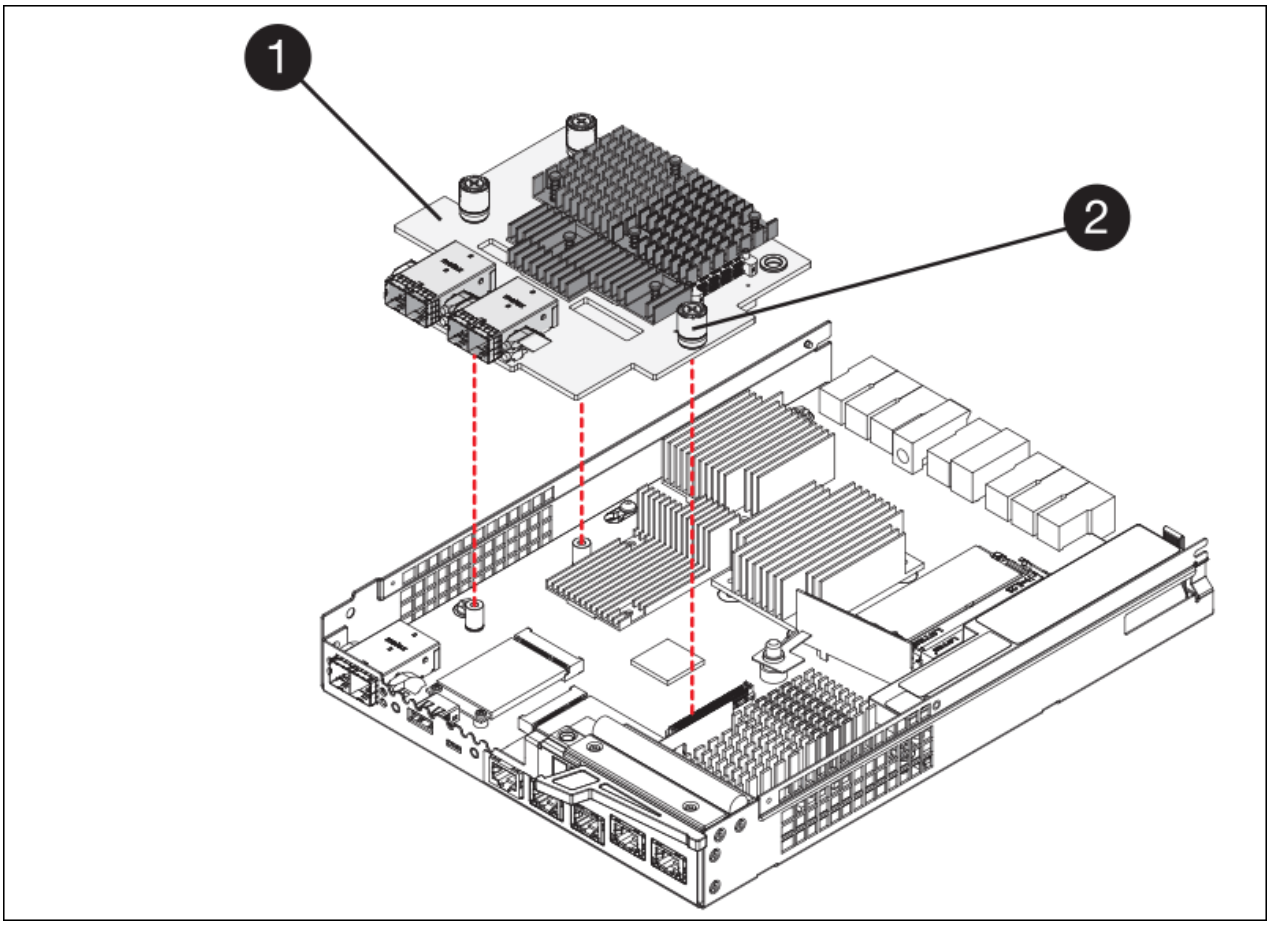

Label	Description
1	Host interface card
2	Thumbscrews

10. Hand-tighten the HIC thumbscrews.

    Don’t use a screwdriver, or you might over tighten the screws.

11. Using a #1 Phillips screwdriver, attach the HIC faceplate you removed from the original controller to the new controller with four screws.

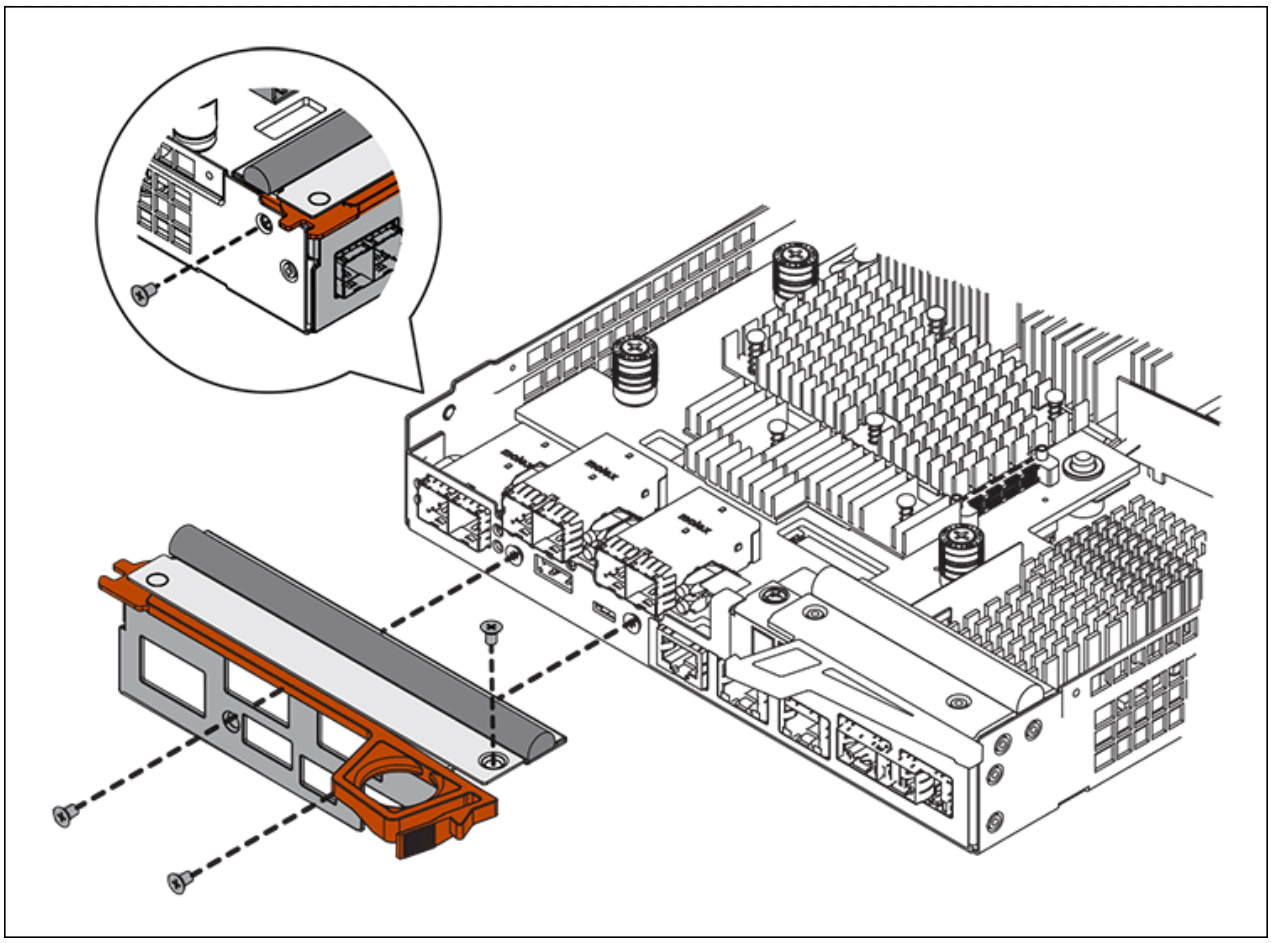

12. Reinstall any removed SFPs into the HIC.

Step 6: Replace controller

Install the replacement controller and verify that it has rejoined the grid.

Steps

1. Install the replacement controller into the appliance.
   a. Turn the controller over, so that the removable cover faces down.
   b. With the cam handle in the open position, slide the controller all the way into the appliance.
   c. Move the cam handle to the left to lock the controller in place.
   d. Replace the cables and SFPs.
   e. Power on the controller shelf.
   f. If the original controller used DHCP for the IP address, locate the MAC address on the label on the back of the replacement controller. Ask your network administrator to associate the DNS/network and IP address for the controller you removed with the MAC address for the replacement controller.

   > If the original controller did not use DHCP for the IP address, the new controller will adopt the IP address of the controller you removed.

2. Bring the controller online using SANtricity System Manager:

a. Select **Hardware**.
b. If the graphic shows the drives, select **Show back of shelf**.
c. Select the controller you want to place online.
d. Select **Place Online** from the context menu, and confirm that you want to perform the operation.
e. Verify that the seven-segment display shows a state of `99`.

3. Confirm that the new controller is Optimal, and collect support data.

After replacing the part, return the failed part to NetApp, as described in the RMA instructions shipped with the kit. See the Part Return & Replacements page for further information.

Related information

NetApp E-Series Systems Documentation Site

Replace hardware components in storage controller shelf

If a hardware problem occurs, you might need to replace a component in the storage controller shelf.

Before you begin

- You have the E-Series hardware replacement procedure.
- You have physically located the storage appliance where you are replacing storage shelf hardware components in the data center.

  Locate controller in data center

About this task

To replace the battery in the storage controller, see the steps in the instructions for replacing a storage controller. Those instructions describe how to remove a controller from the appliance, remove the battery from the controller, install the battery, and replace the controller.

For instructions for the other field replaceable units (FRUs) in the controller shelves, access the E-Series procedures for system maintenance.

FRU	See instructions
Battery	StorageGRID (these instructions): Replacing a storage controller
Drive	E-Series: • Replace drive (60-drive) • Replace drive (12-drive or 24-drive)
Power canister	E-Series • Replace power canister (60-drive) • Replace power supply (12-drive or 24-drive)

FRU	See instructions
Fan canister (60-drive shelves only)	E-Series: Replace fan canister (60-drive)
Drive drawer (60-drive shelves only)	E-Series: Replace drive drawer (60-drive)

After replacing the part, return the failed part to NetApp, as described in the RMA instructions shipped with the kit. See the Part Return & Replacements page for further information.

Related information

NetApp E-Series Systems Documentation Site

Replace storage controller

Replace hardware components in optional 60-drive expansion shelf

You might need to replace an input/output module, a power supply, or a fan in the expansion shelf.

Before you begin

- You have the E-Series hardware replacement procedure.
- You have physically located the storage appliance where you are replacing expansion shelf hardware components in the data center.

  Locate controller in data center

About this task

To replace an input/output module (IOM) in a 60-drive expansion shelf, see the steps in the instructions for replacing a storage controller.

To replace a power supply or a fan in a 60-drive expansion shelf, access the E-Series procedures for maintaining 60-drive hardware.

FRU	See E-Series instructions for
Input/output module (IOM)	Replacing an IOM
Power canister	Replace power canister (60-drive)
Fan canister	Replace fan canister (60-drive)

Replace SG6000-CN controller

You might need to replace the SG6000-CN controller if it is not functioning optimally or if it has failed.

Before you begin

- You have a replacement controller with the same part number as the controller you are replacing.
- You have labels to identify each cable that is connected to the controller.
- You have physically located the controller to replace in the data center.

  Locate controller in data center

About this task

The appliance Storage Node will not be accessible when you replace the SG6000-CN controller. If the SG6000-CN controller is functioning sufficiently, you can perform a controlled shutdown at the start of this procedure.

> If you are replacing the controller before installing StorageGRID software, you might not be able to access the StorageGRID Appliance Installer immediately after completing this procedure. While you can access the StorageGRID Appliance Installer from other hosts on the same subnet as the appliance, you can't access it from hosts on other subnets. This condition should resolve itself within 15 minutes (when any ARP cache entries for the original controller time out), or you can clear the condition immediately by purging any old ARP cache entries manually from the local router or gateway.

Steps

1. Display the current configurations of the appliance and record them.
   a. Log in to the appliance to be replaced:
      i. Enter the following command: `ssh admin@grid_node_IP`
      ii. Enter the password listed in the `Passwords.txt` file.
      iii. Enter the following command to switch to root: `su -`
      iv. Enter the password listed in the `Passwords.txt` file.

      When you are logged in as root, the prompt changes from `$` to `#`.

   b. Enter: **`run-host-command ipmitool lan print`** to display the current BMC configurations for the appliance.
2. If the SG6000-CN controller is functioning sufficiently to allow for a controlled shutdown, shut down the SG6000-CN controller.

   Shut down SG6000-CN controller

3. If any of the network interfaces on this StorageGRID appliance are configured for DHCP, you might need to update the permanent DHCP lease assignments on the DHCP servers to reference the MAC addresses of the replacement appliance. The update ensures the appliance is assigned the expected IP addresses. See Update MAC address references.
4. Remove and replace the SG6000-CN controller:
   a. Label the cables and then disconnect the cables and any SFP+ or SFP28 transceivers.

      > To prevent degraded performance, don't twist, fold, pinch, or step on the cables.

   b. Remove the failed controller from the cabinet or rack.
   c. Install the replacement controller into the cabinet or rack.

d. Replace the cables and any SFP+ or SFP28 transceivers.

e. Power on the controller and monitor the controller LEDs and boot-up codes.

5. If the appliance where you replaced the controller used a key management server (KMS) to encrypt data, additional configuration might be required before the node can join the grid. If the node does not automatically join the grid, make sure that these configuration settings have transferred to the new controller and manually configure any settings that don't have the expected configuration:
   - Configure network links
   - Configure StorageGRID IP addresses
   - Configure node encryption for the appliance
6. Log in to the appliance with the replaced controller:
   a. Enter the following command: `ssh admin@grid_node_IP`
   b. Enter the password listed in the `Passwords.txt` file.
   c. Enter the following command to switch to root: `su -`
   d. Enter the password listed in the `Passwords.txt` file.
7. Restore BMC network connectivity for the appliance. There are two options:
   - Use static IP, netmask, and gateway
   - Use DHCP to obtain an IP, netmask, and gateway
     a. To restore the BMC configuration to use a static IP, netmask, and gateway, enter the following commands:

```
run-host-command ipmitool lan set 1 ipsrc static

run-host-command ipmitool lan set 1 ipaddr Appliance_IP

run-host-command ipmitool lan set 1 netmask Netmask_IP

run-host-command ipmitool lan set 1 defgw ipaddr Default_gateway
```

     b. To restore the BMC configuration to use DHCP to obtain an IP, netmask, and gateway, enter the following command:

```
run-host-command ipmitool lan set 1 ipsrc dhcp
```

8. After restoring BMC network connectivity, connect to the BMC interface to audit and restore any additional custom BMC configuration you might have applied. For example, you should confirm the settings for SNMP trap destinations and email notifications. See Configure BMC interface.
9. Confirm that the appliance node appears in the Grid Manager and that no alerts appear.

After replacing the part, return the failed part to NetApp, as described in the RMA instructions shipped with the kit. See the Part Return & Replacements page for further information.

Related information

Install SG6000-CN into cabinet or rack

View status indicators

View boot-up codes for SG6000-CN controller

Replace one or both power supplies in the SG6000-CN controller

The SG6000-CN controller has two power supplies for redundancy. If one of the power supplies fails, you must replace it as soon as possible to ensure that the compute controller has redundant power. Both power supplies operating in the controller must be the same model and wattage.

Before you begin

- You have determined the physical location in the data center of the controller with the power supply to be replaced.

  Locating the controller in a data center

- If you are replacing only one power supply:
  - You have unpacked the replacement power supply unit and ensured that it is the same model and wattage as the power supply unit you are replacing.
  - You have confirmed that the other power supply is installed and running.
- If you are replacing both power supplies at the same time:
  - You have unpacked the replacement power supply units and ensured they are the same model and wattage.

About this task

The figure shows the two power supply units for the SG6000-CN controller, which are accessible from the back of the controller. Use this procedure to replace one or both of the power supplies. If you are replacing both power supplies, you must first perform a controlled shut down of the appliance.

Steps

1. If you are replacing only one power supply, you don't need to shut down the appliance. Go to the Unplug the power cord step. If you are replacing both power supplies at the same time, do the following before unplugging the power cords:
   a. Shut down the appliance.
2. Unplug the power cord from each power supply to be replaced.
3. Lift the cam handle on the first supply to be replaced.

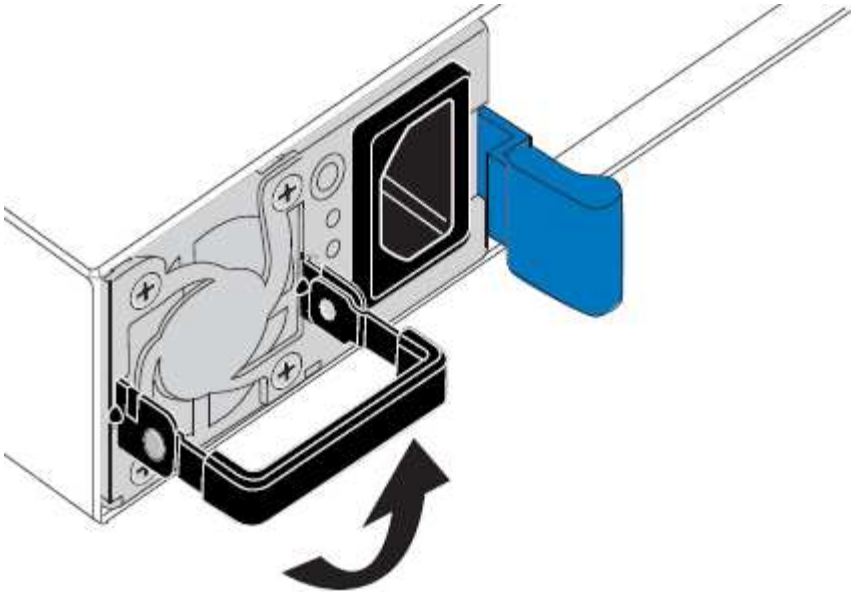

4. Press the blue latch and pull the power supply out.

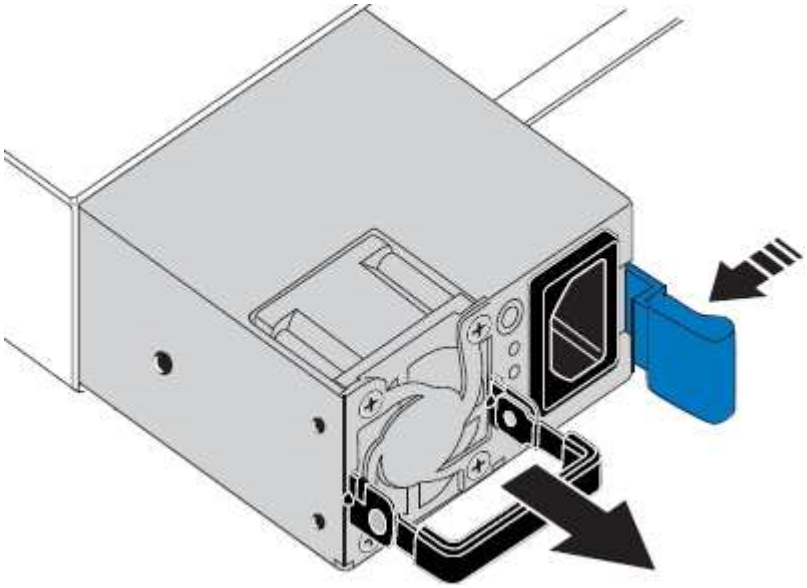

5. With the blue latch on the right, slide the replacement power supply into the chassis.

   Both power supplies must be the same model and wattage.

   Ensure that the blue latch is on the right side when you slide the replacement unit in.

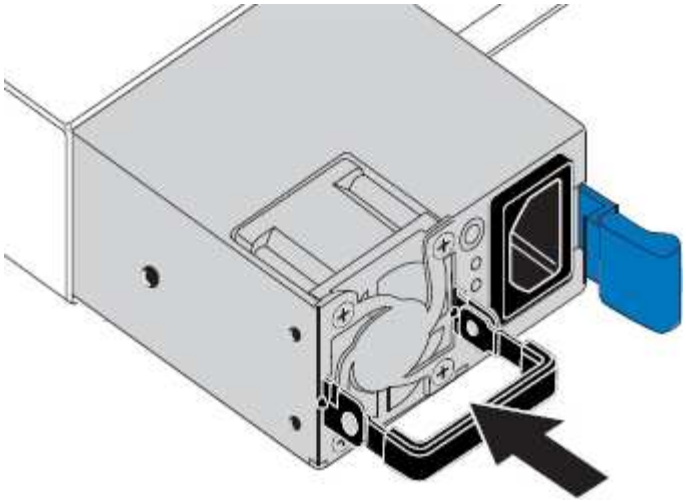

6. Push the cam handle down to secure the replacement power supply.
7. If you are replacing both power supplies, repeat steps 2 though 6 to replace the second power supply.
8. Connect the power cords to the replaced units and apply power.

After replacing the part, return the failed part to NetApp, as described in the RMA instructions shipped with the kit. See the Part Return & Replacements page for further information.

Replace controller in rack

Remove SG6000-CN controller from cabinet or rack

Remove the SG6000-CN controller from a cabinet or rack to access the top cover or to move the controller to a different location.

Before you begin

- You have labels to identify each cable that is connected to the SG6000-CN controller.
- You have physically located the SG6000-CN controller where you are performing maintenance in the data center.

  Locate controller in data center

- You have shut down the SG6000-CN controller.

  Shut down SG6000-CN controller

  Don’t shut down the controller using the power switch.

Steps

1. Label and then disconnect the controller power cables.
2. Wrap the strap end of the ESD wristband around your wrist, and secure the clip end to a metal ground to prevent static discharge.
3. Label and then disconnect the controller data cables and any SFP+ or SFP28 transceivers.

   To prevent degraded performance, don’t twist, fold, pinch, or step on the cables.

4. Loosen the two captive screws on the controller front panel.

5. Slide the SG6000-CN controller forward out of the rack until the mounting rails are fully extended and you hear the latches on both sides click.

   The controller top cover is accessible.

6. Optional: If you are fully removing the controller from the cabinet or rack, follow the instructions for the rail kit to remove the controller from the rails.

After replacing the part, return the failed part to NetApp, as described in the RMA instructions shipped with the kit. See the Part Return & Replacements page for further information.

Related information

Remove SG6000-CN controller cover

Reinstall SG6000-CN controller into cabinet or rack

Reinstall the controller into a cabinet or rack when hardware maintenance is complete.

Before you begin

You have reinstalled the controller cover.

Reinstall SG6000-CN controller cover

Steps

1. Press the blue rail releases both rack rails at the same time and slide the SG6000-CN controller into the rack until it is fully seated.

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

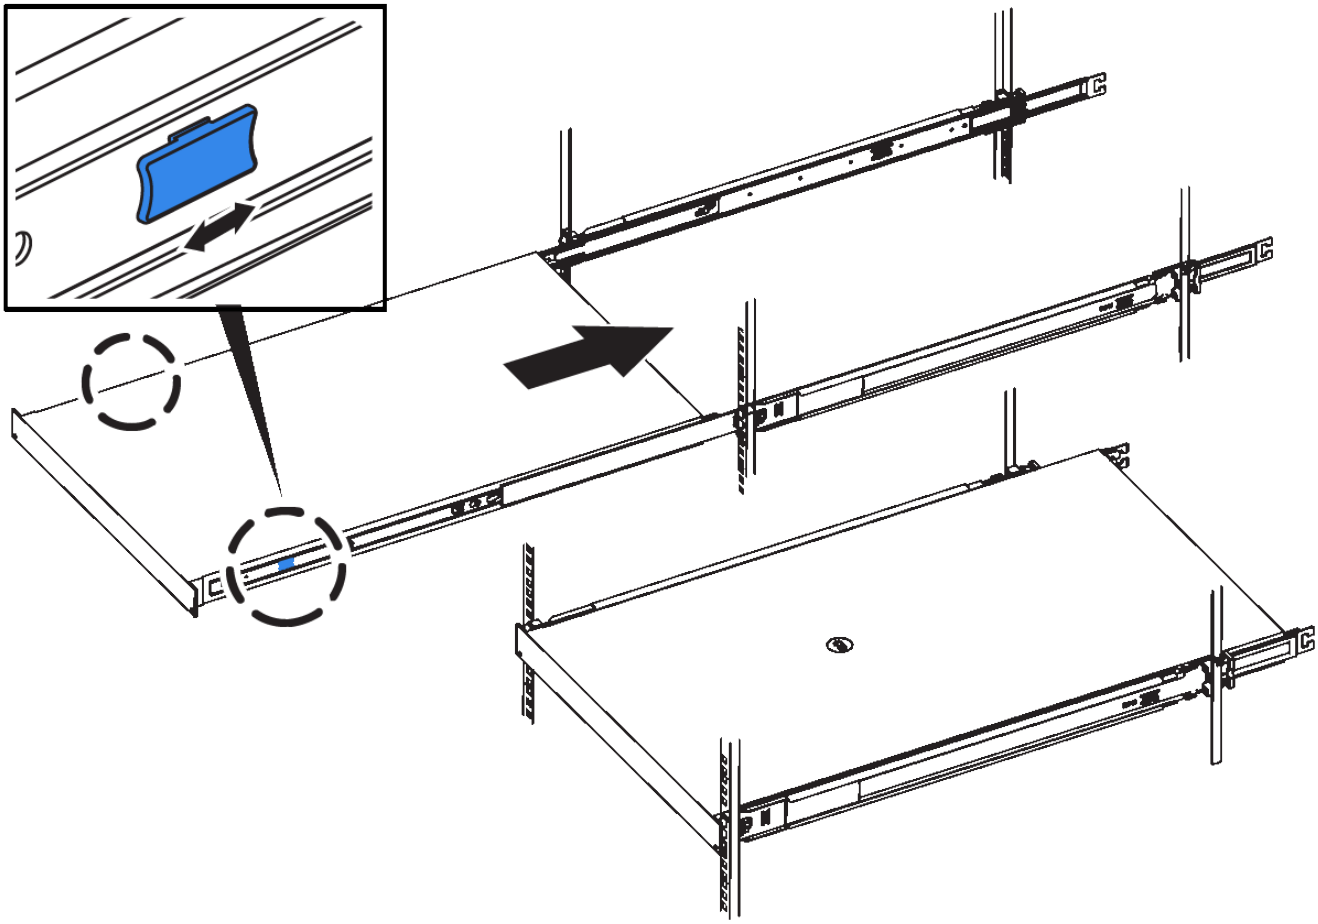

Don’t attach the front bezel until after you power on the controller.

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

3. Wrap the strap end of the ESD wristband around your wrist, and secure the clip end to a metal ground to prevent static discharge.
4. Reconnect the controller data cables and any SFP+ or SFP28 transceivers.

To prevent degraded performance, don’t twist, fold, pinch, or step on the cables.

Cable appliance (SG6000)

5. Reconnect the controller power cables.

Connect power cords and apply power (SG6000)

After you finish

The controller can be restarted.

Power on SG6000-CN controller and verify operation

Replace controller cover

Remove SG6000-CN controller cover

Remove the controller cover to access internal components for maintenance.

Before you begin

Remove the controller from the cabinet or rack to access the top cover.

Remove SG6000-CN controller from cabinet or rack

Steps

1. Make sure that the SG6000-CN controller cover latch is not locked. If necessary, turn the blue plastic latch lock one-quarter turn in the unlock direction, as shown on the latch lock.
2. Rotate the latch up and back toward the rear of the SG6000-CN controller chassis until it stops; then, carefully lift the cover from the chassis and set it aside.

Wrap the strap end of an ESD wristband around your wrist and secure the clip end to a metal ground to prevent static discharge when working inside the SG6000-CN controller.

After replacing the part, return the failed part to NetApp, as described in the RMA instructions shipped with the kit. See the Part Return & Replacements page for further information.

Related information

Remove Fibre Channel HBA

Reinstall SG6000-CN controller cover

Reinstall the controller cover when internal hardware maintenance is complete.

Before you begin

You have completed all maintenance procedures inside the controller.

Steps

1. With the cover latch open, hold the cover above the chassis and align the hole in the top cover latch with the pin in the chassis. When the cover is aligned, lower it onto the chassis.

2. Rotate the cover latch forward and down until it stops and the cover fully seats into the chassis. Verify that there are no gaps along the front edge of the cover.

   If the cover is not fully seated, you might not be able to slide the SG6000-CN controller into the rack.

3. Optional: Turn the blue plastic latch lock one-quarter turn in the lock direction, as shown on the latch lock, to lock it.

After you finish

Reinstall the controller in the cabinet or rack.

Reinstall SG6000-CN controller into cabinet or rack

Replace Fibre Channel HBA in SG6000-CN controller

Verify Fibre Channel HBA to replace

If you are unsure which Fibre Channel host bus adapter (HBA) to replace, complete this procedure to identify it.

Before you begin

- You have the serial number of the storage appliance or SG6000-CN controller where the Fibre Channel HBA needs to be replaced.

> If the serial number of the storage appliance containing the Fibre Channel HBA you are replacing starts with the letter Q, it will not be listed in the Grid Manager. You must check the tags attached to the front of each SG6000-CN controller in the data center until you find a match.

- You are signed in to the Grid Manager using a supported web browser.

Steps

1. From the Grid Manager, select **NODES**.
2. From the table on the Nodes page, select an appliance Storage Node.

3. Select the **Hardware** tab.

   Check the **Storage appliance chassis serial number** and the **Compute controller serial number** in the StorageGRID Appliance section. See if one of these serial numbers matches the serial number of the storage appliance where you are replacing the Fibre Channel HBA. If either serial number matches, you have found the correct appliance.

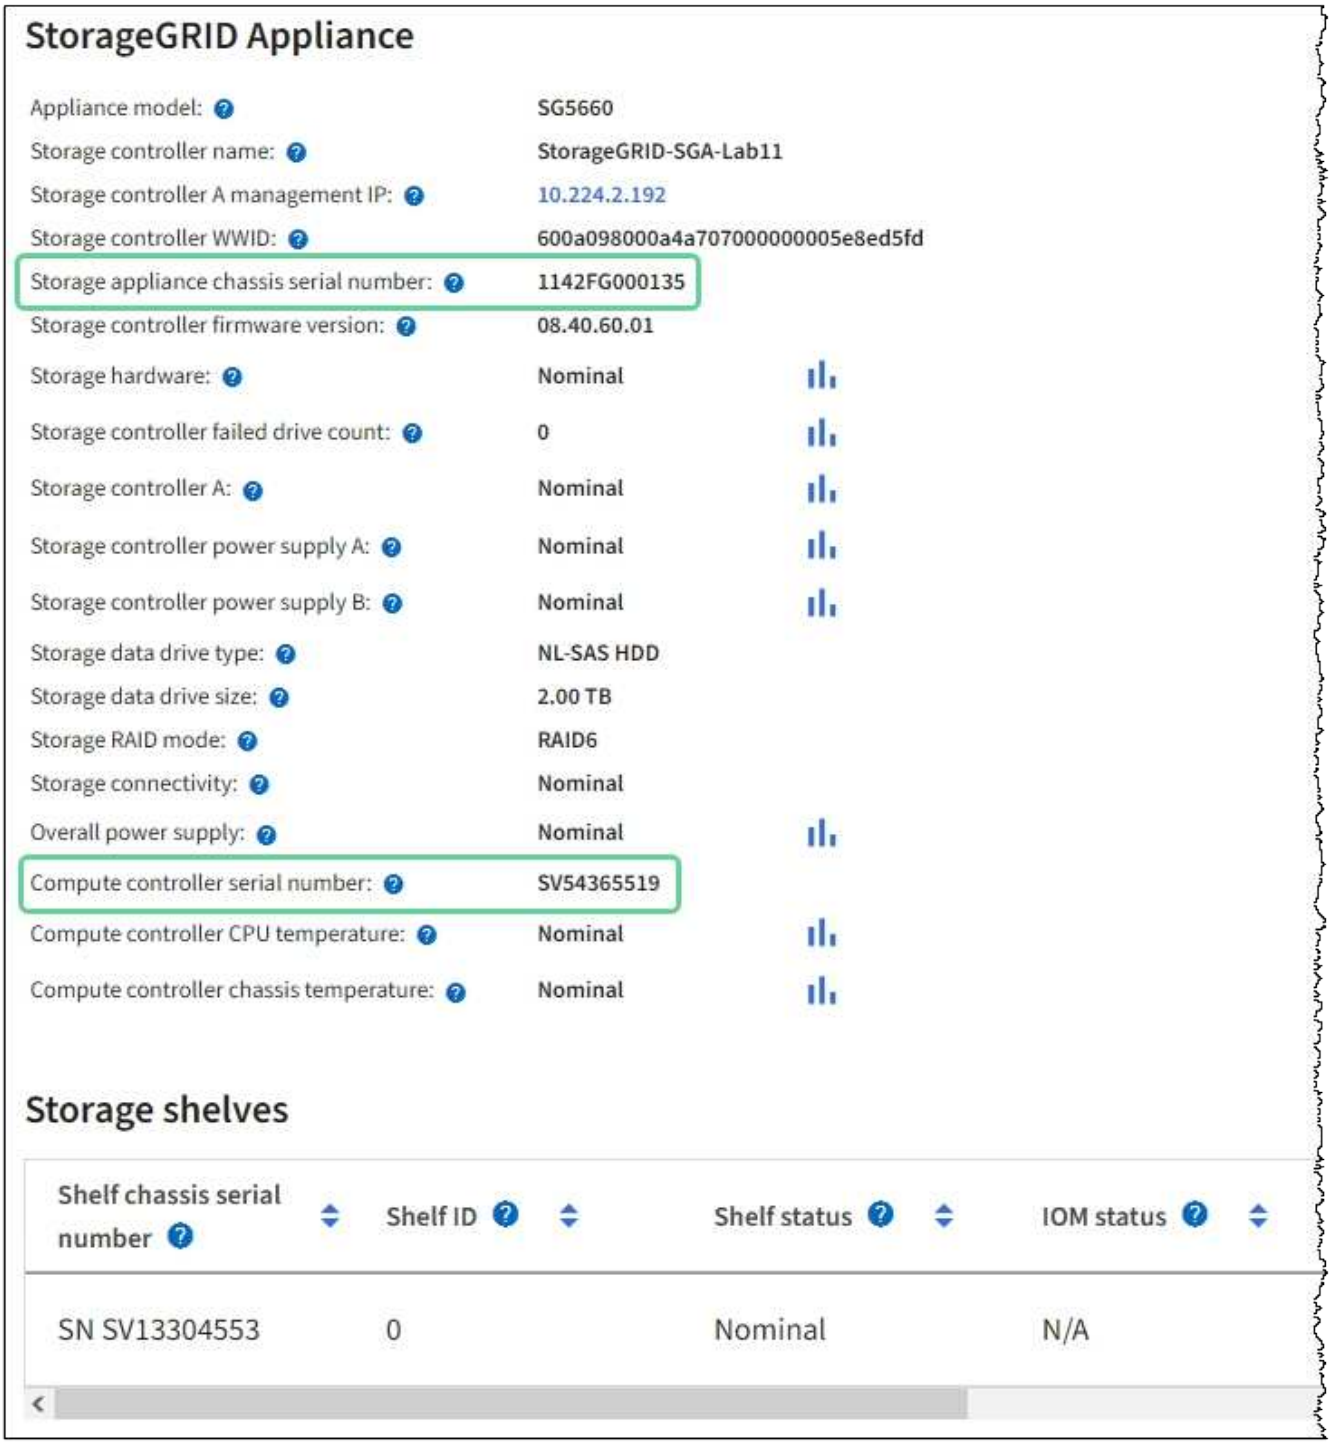

   - If the StorageGRID Appliance section does not display, the node selected is not a StorageGRID appliance. Select a different node from the tree view.
   - If the Appliance Model is not SG6060 or SG6060X, select a different node from the tree view.
   - If the serial numbers don't match, select a different node from the tree view.

4. After you locate the node where the Fibre Channel HBA needs to be replaced, write down the Compute controller BMC IP address listed the StorageGRID Appliance section.

   You can use this IP address to turn on the compute controller identify LED, to help you locate the appliance in the data center.

   Turn the controller identify LED on and off

Related information

Remove Fibre Channel HBA

Remove Fibre Channel HBA

You might need to replace the Fibre Channel host bus adapter (HBA) in the SG6000-CN controller if it is not functioning optimally or if it has failed.

Before you begin

- You have the correct replacement Fibre Channel HBA.
- You have determined which SG6000-CN controller contains the Fibre Channel HBA to replace.
- You have physically located the SG6000-CN controller in the data center.
- You have shut down the SG6000-CN controller.

  A controlled shutdown is required before you remove the controller from the rack.

- You have removed the controller from the cabinet or rack.
- You have removed the controller cover.

About this task

To prevent service interruptions, confirm that all other Storage Nodes are connected to the grid before starting the Fibre Channel HBA replacement or replace the adapter during a scheduled maintenance window when periods of service disruption are normally expected. See the information about monitoring node connection states.

If you have ever used an ILM rule that creates only one copy of an object, you must replace the Fibre Channel HBA during a scheduled maintenance window. Otherwise, you might temporarily lose access to those objects during this procedure.
See information about why you should not use single-copy replication.

Steps

1. Wrap the strap end of the ESD wristband around your wrist, and secure the clip end to a metal ground to prevent static discharge.
2. Locate the riser assembly at the rear of the controller that contains the Fibre Channel HBA.

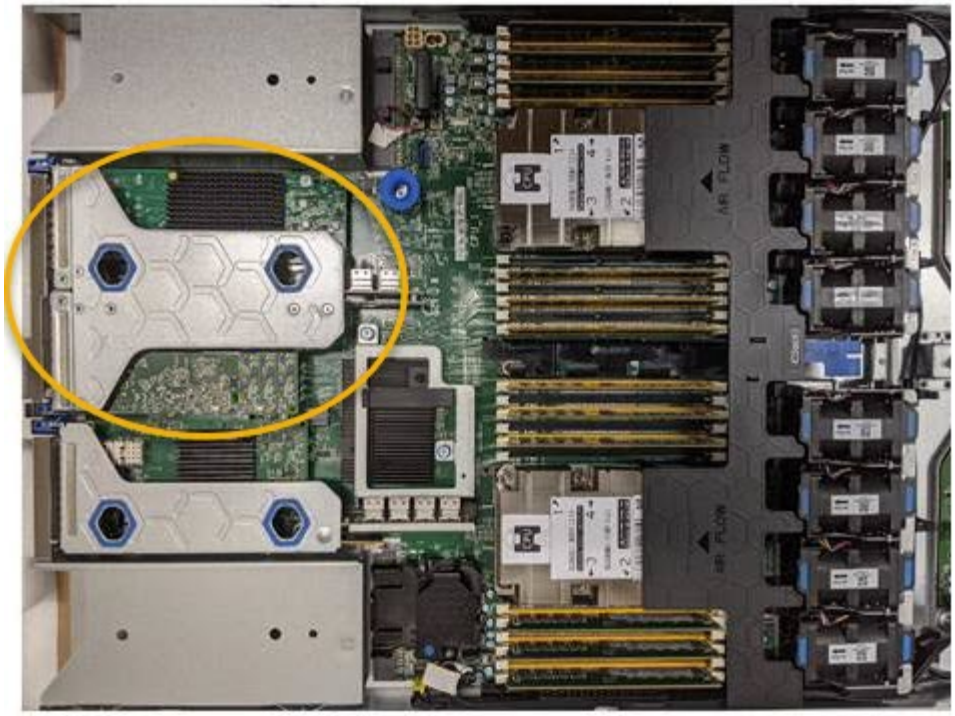

3. Grasp the riser assembly through the blue-marked holes and carefully lift it upwards. Move the riser assembly toward the front of the chassis as you lift it to allow the external connectors in its installed adapters to clear the chassis.
4. Place the riser card on a flat anti-static surface with the metal frame side down to access the adapters.

There are two adapters in the riser assembly: a Fibre Channel HBA and an Ethernet network adapter. The Fibre Channel HBA is indicated in the illustration.

5. Open the blue adapter latch (circled) and carefully remove the Fibre Channel HBA from the riser assembly. Rock the adapter slightly to help remove the adapter from its connector. Don’t use excessive force.
6. Place the adapter on a flat anti-static surface.

After you finish

Install the replacement Fibre Channel HBA.

After replacing the part, return the failed part to NetApp, as described in the RMA instructions shipped with the kit. See the Part Return & Replacements page for further information.

Reinstall Fibre Channel HBA

The replacement Fibre Channel HBA is installed into the same location as the one that was removed.

Before you begin

- You have the correct replacement Fibre Channel HBA.
- You have removed the existing Fibre Channel HBA.

  Remove Fibre Channel HBA

Steps

1. Wrap the strap end of the ESD wristband around your wrist, and secure the clip end to a metal ground to prevent static discharge.
2. Remove the replacement Fibre Channel HBA from its packaging.
3. With the blue adapter latch in the open position, align the Fibre Channel HBA with its connector on the riser assembly; then, carefully press the adapter into the connector until it is fully seated.

   

   There are two adapters in the riser assembly: a Fibre Channel HBA and an Ethernet network adapter. The Fibre Channel HBA is indicated in the illustration.

4. Locate the alignment hole on the riser assembly (circled) that aligns with a guide pin on the system board to ensure correct riser assembly positioning.

   

5. Position the riser assembly in the chassis, making sure that it aligns with the connector and guide pin on

the system board; then, insert the riser assembly.

6. Carefully press the riser assembly in place along its center line, next to the blue-marked holes, until it is fully seated.
7. Remove the protective caps from the Fibre Channel HBA ports where you will be reinstalling cables.

After you finish

If you have no other maintenance procedures to perform in the controller, reinstall the controller cover.

Reinstall SG6000-CN controller cover

Maintain SG5700 hardware

Maintain SG5700 appliance

You might need to upgrade the SANtricity OS Software on the E2800 controller, change the Ethernet link configuration of the E5700SG controller, replace the E2800 controller or the E5700SG controller, or replace specific components. The procedures in this section assume that the appliance has already been deployed as a Storage Node in a StorageGRID system.

Procedures specific to maintaining your SG5700 appliance are in this section.

See Common procedures for maintenance procedures that are used by all appliances.

See Set up hardware for maintenance procedures that are also performed during initial appliance installation and configuration.

Maintenance configuration procedures

Upgrade SANtricity OS on storage controller

To ensure optimal functioning of the storage controller, you must upgrade to the latest maintenance release of the SANtricity OS that is qualified for your StorageGRID appliance.

Consult the NetApp Interoperability Matrix Tool (IMT) to determine which version you should be using.

Download the new SANtricity OS Software file from NetApp Downloads: StorageGRID Appliance.

Use one of the following procedures based on the version of SANtricity OS currently installed:

- If the storage controller is using SANtricity OS 08.42.20.00 (11.42) or newer, use the Grid Manager to perform the upgrade.

  Upgrade SANtricity OS on storage controllers using Grid Manager

- If the storage controller is using a SANtricity OS version older than 08.42.20.00 (11.42), use maintenance mode to perform the upgrade.

  Upgrade SANtricity OS on E2800 controller using maintenance mode

Upgrade SANtricity OS on storage controllers using Grid Manager

For storage controllers currently using SANtricity OS 08.42.20.00 (11.42) or newer, you must use the Grid Manager to apply an upgrade.

Before you begin

- You have consulted the NetApp Interoperability Matrix Tool (IMT) to confirm that the SANtricity OS version you are using for the upgrade is compatible with your appliance.
- You have the Maintenance or Root access permission.
- You are signed in to the Grid Manager using a supported web browser.
- You have the provisioning passphrase.
- You have access to NetApp Downloads: StorageGRID Appliance.

About this task

You can't perform other software updates (StorageGRID software upgrade or a hotfix) until you have completed the SANtricity OS upgrade process. If you attempt to start a hotfix or a StorageGRID software upgrade before the SANtricity OS upgrade process has finished, you are redirected to the SANtricity OS upgrade page.

The procedure will not be complete until the SANtricity OS upgrade has been successfully applied to all applicable nodes that have been selected for the upgrade. It might take more than 30 minutes to load the SANtricity OS on each node (sequentially) and up to 90 minutes to reboot each StorageGRID storage appliance. Any nodes in your grid that don't use SANtricity OS will not be affected by this procedure.

> The following steps are only applicable when you are using the Grid Manager to perform the upgrade. The storage controllers in the appliance can't be upgraded using the Grid Manager when the controllers are using SANtricity OS older than 08.42.20.00 (11.42).

> This procedure will automatically upgrade the NVSRAM to the most recent version associated with the SANtricity OS upgrade. You don't need to apply a separate NVSRAM upgrade file.

> Be sure to apply the latest StorageGRID hotfix before you begin this procedure. See StorageGRID hotfix procedure for details.

Steps

1. Download the new SANtricity OS Software file from NetApp Downloads: StorageGRID Appliance.

   Be sure to choose the SANtricity OS version for your storage controllers.

2. Select **MAINTENANCE** > **System** > **Software update**.

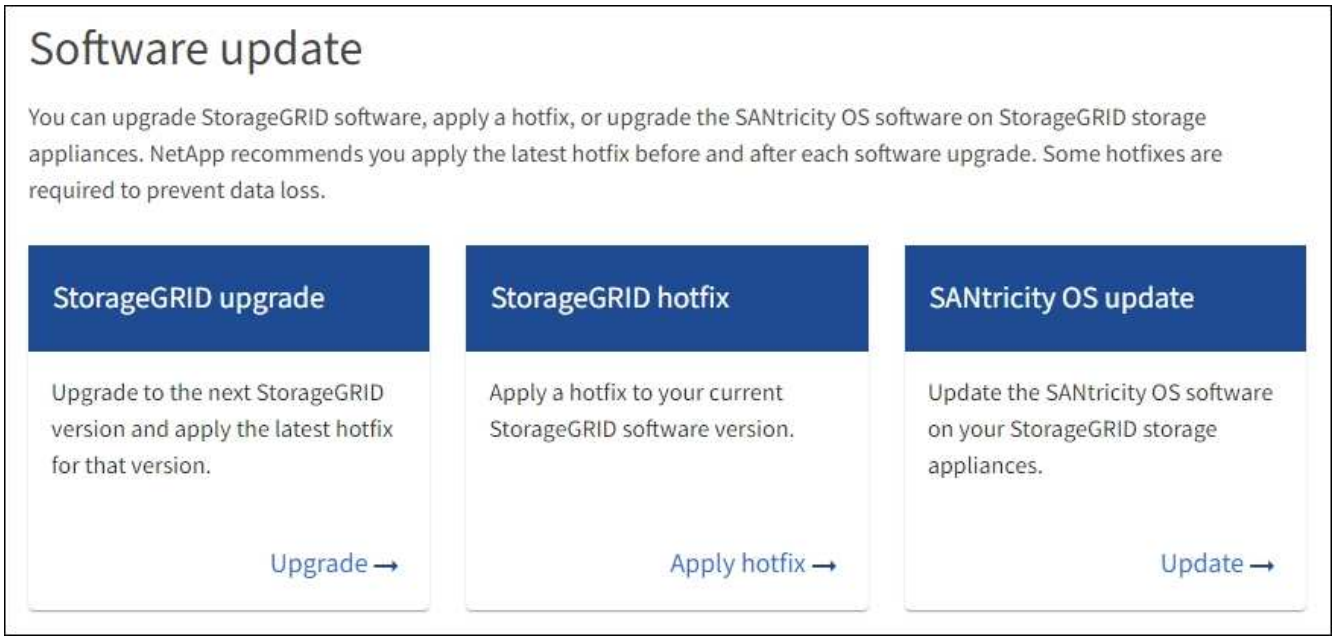

3. In the SANtricity OS update section, select **Update**.

   The SANtricity OS upgrade page appears and lists the details for each appliance node including:

   - Node name
   - Site
   - Appliance model
   - SANtricity OS version
   - Status
   - Last upgrade status

4. Review the information in the table for all of your upgradable appliances. Confirm that all storage controllers have **Nominal** status. If the status for any controller is **Unknown**, go to **Nodes** > ***appliance node*** > **Hardware** to investigate and resolve the issue.
5. Select the SANtricity OS upgrade file you downloaded from the NetApp Support Site.

   a. Select **Browse**.
   b. Locate and select the file.
   c. Select **Open**.

   The file is uploaded and validated. When the validation process is done, the file name is shown with a green check mark next to the **Browse** button. Don’t change the file name because it is part of the verification process.

6. Enter the provisioning passphrase and select **Continue**.

   A warning box appears stating that your browser’s connection might be lost temporarily as services on nodes that are upgraded are restarted.

7. Select **Yes** to stage the SANtricity OS upgrade file to the primary Admin Node.

   When the SANtricity OS upgrade starts:

a. The health check is run. This process checks that no nodes have the status of Needs Attention.

> If any errors are reported, resolve them and select **Start** again.

b. The SANtricity OS Upgrade Progress table appears. This table shows all Storage Nodes in your grid and the current stage of the upgrade for each node.

> The table shows all appliance Storage Nodes. Software-based Storage Nodes aren't displayed. Select **Approve** for all nodes that require the upgrade.

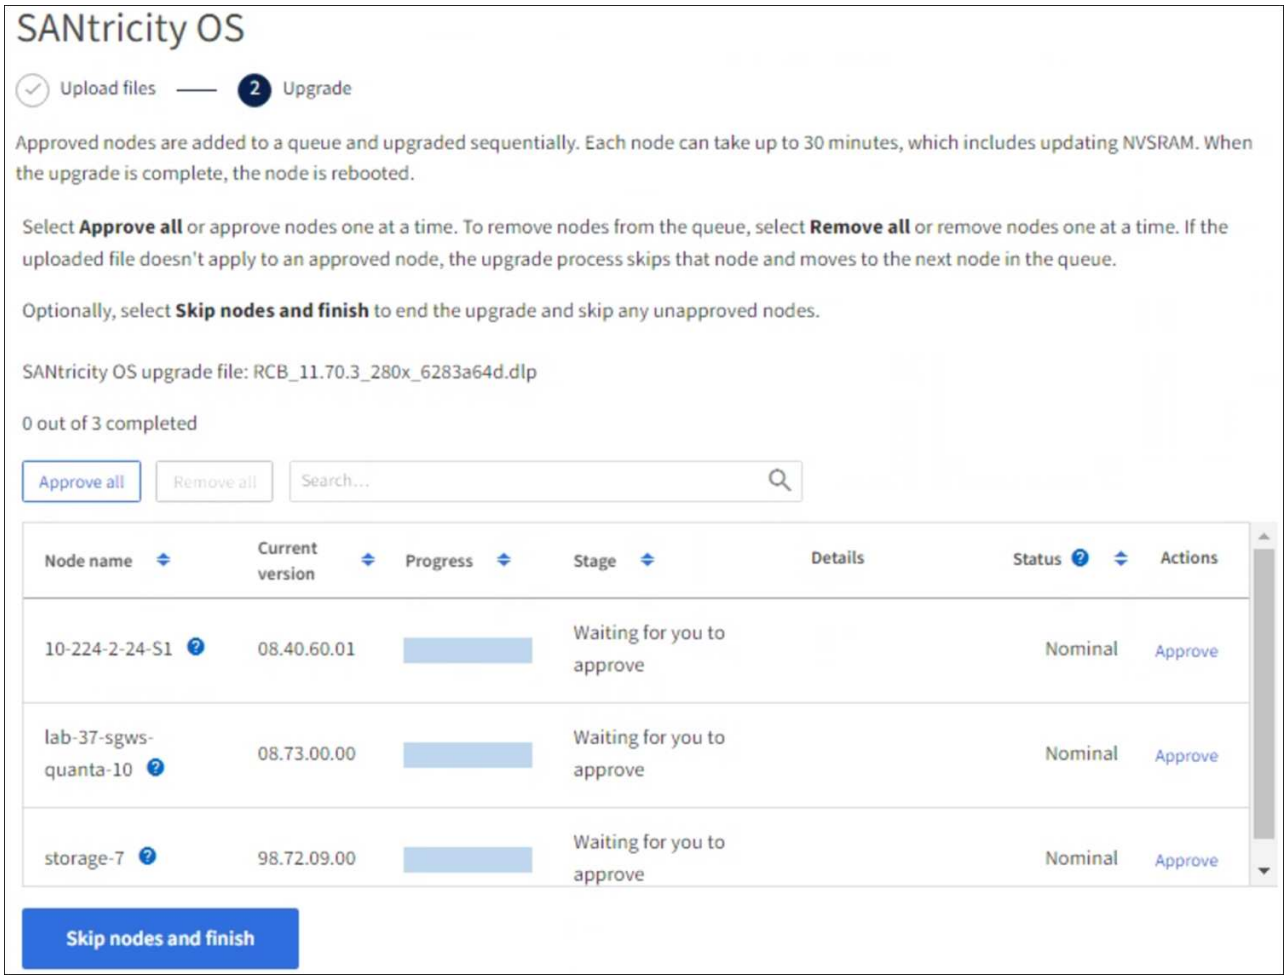

8. Optionally, sort the list of nodes in ascending or descending order by:
   - Node name
   - Current version
   - Progress
   - Stage
   - Status

   You can also enter a term in the Search box to search for specific nodes.

9. Approve the grid nodes you are ready to add to the upgrade queue. Approved nodes are upgraded one at a time.

> Don’t approve the SANtricity OS upgrade for an appliance Storage Node unless you are sure the node is ready to be stopped and rebooted. When the SANtricity OS upgrade is approved on a node, the services on that node are stopped and the upgrade process begins. Later, when the node is finished upgrading, the appliance node is rebooted. These operations might cause service interruptions for clients that are communicating with the node.

- Select the **Approve All** button to add all Storage Nodes to the SANtricity OS upgrade queue.

  > If the order in which nodes are upgraded is important, approve nodes or groups of nodes one at a time and wait until the upgrade is complete on each node before approving the next node.

- Select one or more **Approve** buttons to add one or more nodes to the SANtricity OS upgrade queue. The **Approve** button is disabled if the Status is not Nominal.

  After you select **Approve**, the upgrade process determines if the node can be upgraded. If a node can be upgraded, it is added to the upgrade queue.

  For some nodes, the selected upgrade file is intentionally not applied and you can complete the upgrade process without upgrading these specific nodes. Nodes intentionally not upgraded show a stage of Complete (upgrade attempted) and list the reason the node was not upgraded in the Details column.

10. If you need to remove a node or all nodes from the SANtricity OS upgrade queue, select **Remove** or **Remove All**.

    When the stage progresses beyond Queued, the **Remove** button is hidden and you can no longer remove the node from the SANtricity OS upgrade process.

11. Wait while the SANtricity OS upgrade is applied to each approved grid node.

    - If any node shows a stage of Error while the SANtricity OS upgrade is applied, the upgrade has failed for the node. With the assistance of technical support, you might need to place the appliance in maintenance mode to recover it.
    - If the firmware on the node is too old to be upgraded with the Grid Manager, the node shows a stage of Error with the details that you must use maintenance mode to upgrade SANtricity OS on the node. To resolve the error, do the following:

      a. Use maintenance mode to upgrade SANtricity OS on the node that shows a stage of Error.

      b. Use the Grid Manager to restart and complete the SANtricity OS upgrade.

    When the SANtricity OS upgrade is complete on all approved nodes, the SANtricity OS Upgrade Progress table closes and a green banner shows the number of nodes upgraded, and the date and time the upgrade completed.

12. If a node can’t be upgraded, note the reason shown in the Details column and take the appropriate action.

> The SANtricity OS upgrade process will not be complete until you approve the SANtricity OS upgrade on all the listed Storage Nodes.

Reason	Recommended action
Storage Node was already upgraded.	No further action required.
SANtricity OS upgrade is not applicable to this node.	The node does not have a storage controller that can be managed by the StorageGRID system. Complete the upgrade process without upgrading the node displaying this message.
SANtricity OS file is not compatible with this node.	The node requires a SANtricity OS file different than the one you selected. After completing the current upgrade, download the correct SANtricity OS file for the node and repeat the upgrade process.

13. If you want to end approving nodes and return to the SANtricity OS page to allow for an upload of a new SANtricity OS file, do the following:
    a. Select **Skip Nodes and Finish**.

       A warning appears asking if you are sure you want to finish the upgrade process without upgrading all applicable nodes.

    b. Select **OK** to return to the **SANtricity OS** page.
    c. When you are ready to continue approving nodes, download the SANtricity OS to restart the upgrade process.

       Nodes already approved and upgraded without errors remain upgraded.

14. Repeat this upgrade procedure for any nodes with a stage of Complete that require a different SANtricity OS upgrade file.

    For any nodes with a status of Needs Attention, use maintenance mode to perform the upgrade.

Related information

NetApp Interoperability Matrix Tool

Upgrade SANtricity OS on E2800 controller using maintenance mode

Upgrade SANtricity OS on E2800 controller using maintenance mode

For storage controllers currently using SANtricity OS older than 08.42.20.00 (11.42), you must use the maintenance mode procedure to apply an upgrade.

Before you begin

- You have consulted the NetApp Interoperability Matrix Tool (IMT) to confirm that the SANtricity OS version you are using for the upgrade is compatible with your appliance.
- You must place the E5700SG controller into maintenance mode, which interrupts the connection to the E2800 controller.

In rare instances, placing a StorageGRID appliance into maintenance mode might make the appliance unavailable for remote access.

About this task

Don't upgrade the SANtricity OS or NVSRAM in the E-Series controller on more than one StorageGRID appliance at a time.

Upgrading more than one StorageGRID appliance at a time might cause data unavailability, depending on your deployment model and ILM policies.

Steps

1. Confirm the appliance is in maintenance mode.
2. From a service laptop, access SANtricity System Manager and sign in.
3. Download the new SANtricity OS Software file and NVSRAM file to the management client.

   The NVSRAM is specific to the StorageGRID appliance. Don't use the standard NVSRAM download.

4. Follow the instructions in the *E2800 and E5700 SANtricity Software and Firmware Upgrade Guide* or the SANtricity System Manager online help to upgrade the E2800 controller's firmware and NVSRAM.

   Activate the upgrade files immediately. Don't defer activation.

5. If this procedure completed successfully and you have additional procedures to perform while the node is in maintenance mode, perform them now. When you are done, or if you experienced any failures and want to start over, select **Advanced** > **Reboot Controller**, and then select one of these options:
   - Select **Reboot into StorageGRID**
   - Select **Reboot into Maintenance Mode** to reboot the controller with the node remaining in maintenance mode. Select this option if you experienced any failures during the procedure and want to start over. After the node finishes rebooting into maintenance mode, restart from the appropriate step in the procedure that failed.

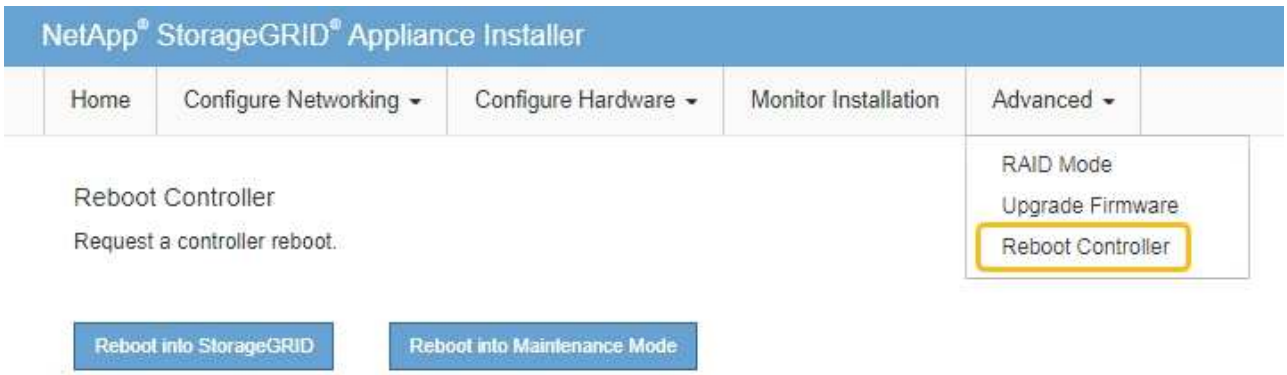

It can take up to 20 minutes for the appliance to reboot and rejoin the grid. To confirm that the reboot is complete and that the node has rejoined the grid, go back to the Grid Manager. The Nodes page

should display a normal status (green check mark icon to the left of the node name) for the appliance node, indicating that no alerts are active and the node is connected to the grid.

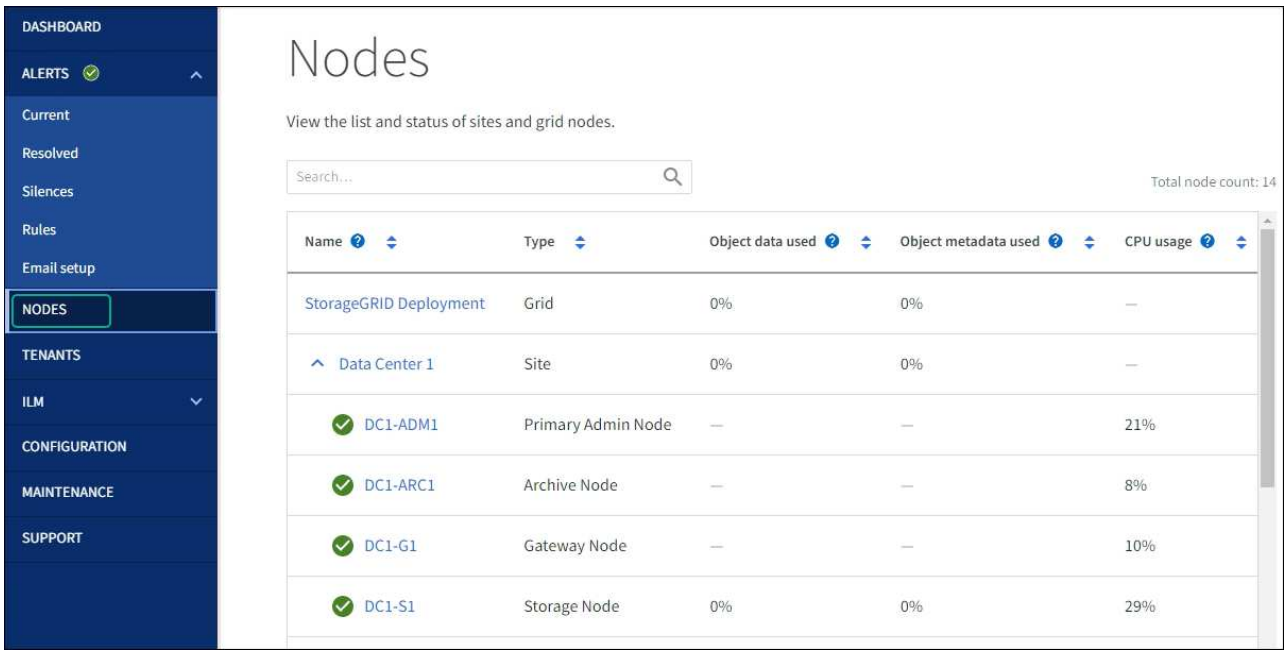

Related information

Upgrade SANtricity OS on storage controllers using Grid Manager

Upgrade drive firmware using SANtricity System Manager

Upgrade SG5700 drive firmware using SANtricity System Manager online method

Use the SANtricity System Manager online method to upgrade the firmware on the drives in your appliance to make sure you have all the latest features and bug fixes.

Before you begin

- The storage appliance has an Optimal status.
- All drives have an Optimal status.

> Don't upgrade the drive firmware on more than one StorageGRID appliance at a time. Doing so might cause data unavailability, depending on your deployment model and ILM policy.

About this task

The drives are upgraded one at a time while the appliance is performing I/O. This method does not require you to place the appliance in maintenance mode. However, system performance might be impacted and the upgrade might take several hours longer than the offline method.

> Drives belonging to volumes that don't have redundancy must be updated using the offline method. The offline method should be used for any drive associated with flash read cache, or any pool or volume group that is currently degraded. You must use the offline method when upgrading SSD drives.

Steps

1. Access SANtricity System Manager using one of these methods:
   - Use the StorageGRID Appliance Installer and select **Advanced** > **SANtricity System Manager**
   - Use the Grid Manager and select **NODES** > ***Storage Node*** > **SANtricity System Manager**
   - Use SANtricity System Manager by browsing to the storage controller IP:

     `https://Storage_Controller_IP`

2. Enter the SANtricity System Manager administrator username and password, if required.
3. Verify the drive firmware version currently installed in the storage appliance:
   a. From SANtricity System Manager, select **SUPPORT** > **Upgrade Center**.
   b. Under Drive Firmware upgrade, select **Begin Upgrade**.

      The Upgrade Drive Firmware page displays the drive firmware files currently installed.

   c. Note the current drive firmware revisions and drive identifiers in the Current Drive Firmware column.

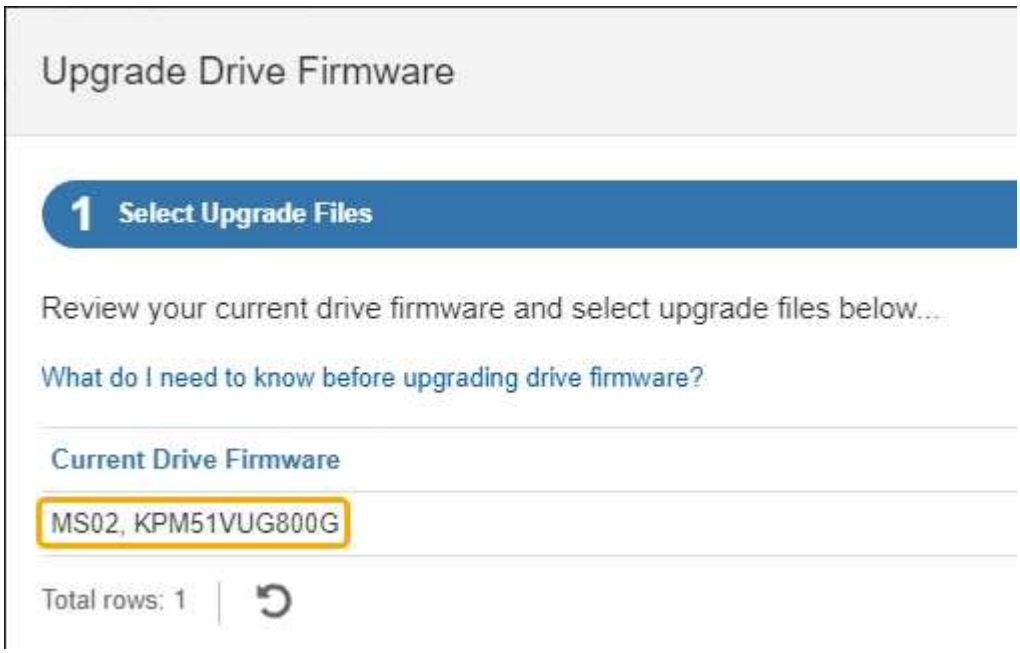

      In this example:

      - The drive firmware revision is **MS02**.
      - The drive identifier is **KPM51VUG800G**.

   d. Select **View drives** in the Associated Drives column to display where these drives are installed in your storage appliance.
   e. Close the Upgrade Drive Firmware window.
4. Download and prepare the available drive firmware upgrade:
   a. Under Drive Firmware upgrade, select **NetApp Support**.
   b. On the NetApp Support Site, select the **Downloads** tab, and then select **E-Series Disk Drive Firmware**.

      The E-Series Disk Firmware page displays.

c. Search for each **Drive Identifier** installed in your storage appliance and verify that each drive identifier has the latest firmware revision.

   - If the firmware revision is not a link, this drive identifier has the latest firmware revision.
   - If one or more drive part numbers are listed for a drive identifier, a firmware upgrade is available for these drives. You can select any link to download the firmware file.

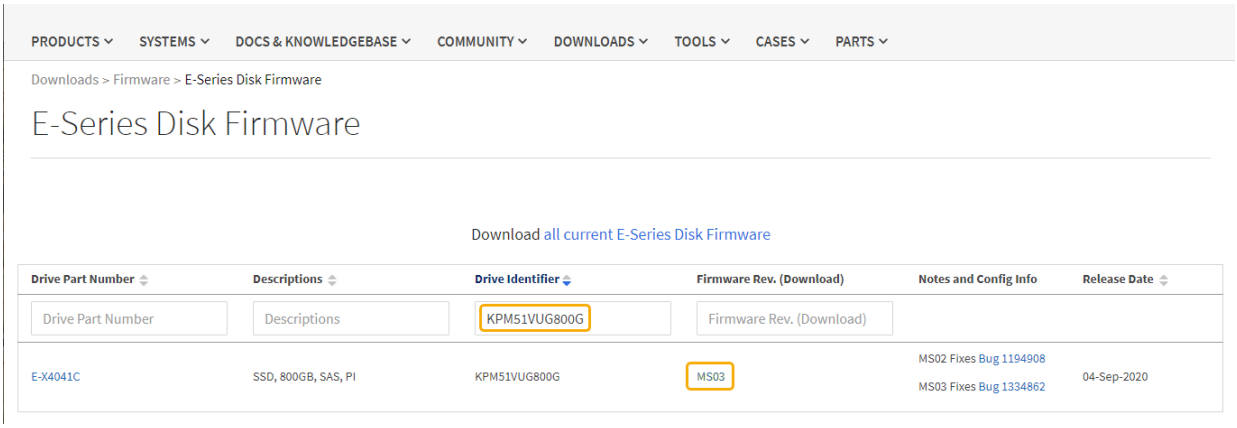

d. If a later firmware revision is listed, select the link in the Firmware Rev. (Download) column to download a `.zip` archive containing the firmware file.

e. Extract (unzip) the drive firmware archive files you downloaded from the Support site.

5. Install the drive firmware upgrade:

   a. From SANtricity System Manager, under Drive Firmware upgrade, select **Begin Upgrade**.

   b. Select **Browse**, and select the new drive firmware files that you downloaded from the Support site.

      Drive firmware files have a filename similar to `D_HUC101212CSS600_30602291_MS01_2800_0002.dlp`.

      You can select up to four drive firmware files, one at a time. If more than one drive firmware file is compatible with the same drive, you get a file conflict error. Decide which drive firmware file you want to use for the upgrade and remove the other one.

   c. Select **Next**.

      Select Drives lists the drives that you can upgrade with the selected firmware files.

      Only drives that are compatible appear.

      The selected firmware for the drive appears in the **Proposed Firmware** column. If you must change this firmware, select **Back**.

   d. Select **Upgrade all drives online** — Upgrades the drives that can support a firmware download while the storage array is processing I/O. You don't have to stop I/O to the associated volumes using these drives when you select this upgrade method.

      An online upgrade can take several hours longer than an offline upgrade.

   e. In the first column of the table, select the drive or drives you want to upgrade.

      The best practice is to upgrade all drives of the same model to the same firmware revision.

f. Select **Start** and confirm that you want to perform the upgrade.

   If you need to stop the upgrade, select **Stop**. Any firmware downloads currently in progress complete. Any firmware downloads that have not started are canceled.

   > Stopping the drive firmware upgrade might result in data loss or unavailable drives.

g. (Optional) To see a list of what was upgraded, select **Save Log**.

   The log file is saved in the downloads folder for your browser with the name `latest-upgrade-log-timestamp.txt`.

   If required, troubleshoot driver firmware upgrade errors.

Upgrade SG5700 drive firmware using SANtricity System Manager using offline method

Use the SANtricity System Manager online method to upgrade the firmware on the drives in your appliance to make sure you have all the latest features and bug fixes.

Before you begin

- The storage appliance has an Optimal status.
- All drives have an Optimal status.
- You have placed the StorageGRID appliance into maintenance mode.

> While the appliance is in maintenance mode, I/O (input/output) activity to the storage controller is stopped to make disruptive storage operations safe.

> Don’t upgrade the drive firmware on more than one StorageGRID appliance at a time. Doing so might cause data unavailability, depending on your deployment model and ILM policy.

About this task

The drives are upgraded in parallel while the appliance is in maintenance mode. If the pool or volume group does not support redundancy or is degraded, you must use the offline method to upgrade the drive firmware. You should also use the offline method for any drive associated with flash read cache, or any pool or volume group that is currently degraded. The offline method upgrades firmware only while all I/O activity is stopped on the drives to be upgraded. To stop I/O activity, place the node into maintenance mode.

The offline method is faster than the online method and will be significantly faster when many drives in a single appliance need upgrades. However, it requires that nodes be taken out of service, which might require scheduling a maintenance window and monitoring progress. Choose the method that is the best fit for your operational procedures and the number of drives that need to be upgraded.

Steps

1. Confirm that the appliance is in maintenance mode.
2. Access SANtricity System Manager using one of these methods:
   - Use the StorageGRID Appliance Installer and select **Advanced** > **SANtricity System Manager**
   - Use the Grid Manager and select **NODES** > ***Storage Node*** > **SANtricity System Manager**
   - Use SANtricity System Manager by browsing to the storage controller IP:

`https://Storage_Controller_IP`

3. Enter the SANtricity System Manager administrator username and password, if required.
4. Verify the drive firmware version currently installed in the storage appliance:
   a. From SANtricity System Manager, select **SUPPORT** > **Upgrade Center**.
   b. Under Drive Firmware upgrade, select **Begin Upgrade**.

      The Upgrade Drive Firmware page displays the drive firmware files currently installed.

   c. Note the current drive firmware revisions and drive identifiers in the Current Drive Firmware column.

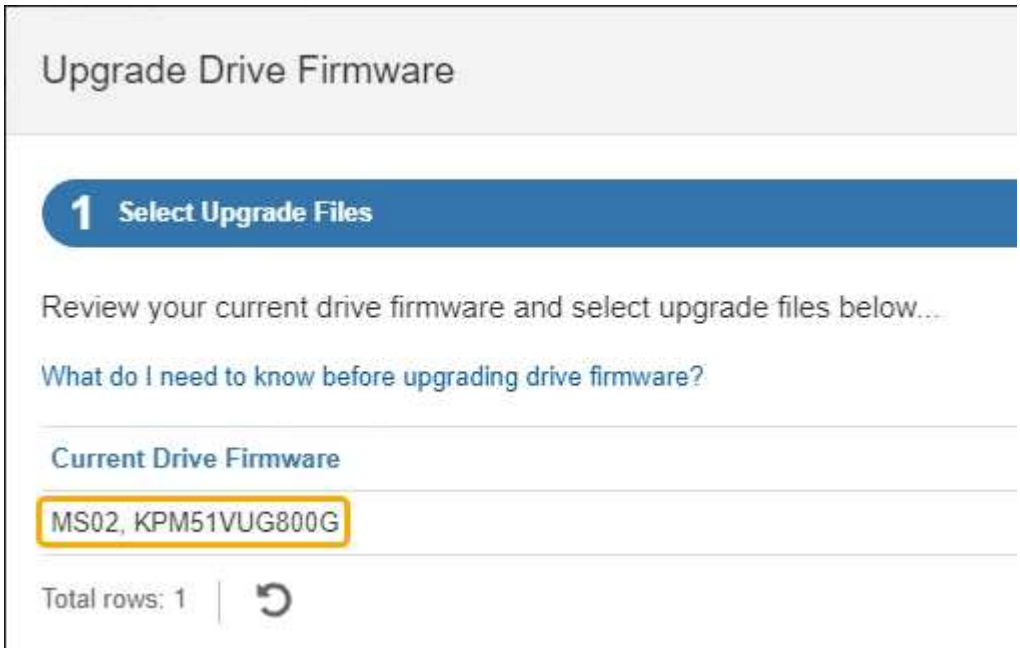

      In this example:

      - The drive firmware revision is **MS02**.
      - The drive identifier is **KPM51VUG800G**.

   d. Select **View drives** in the Associated Drives column to display where these drives are installed in your storage appliance.
   e. Close the Upgrade Drive Firmware window.
5. Download and prepare the available drive firmware upgrade:
   a. Under Drive Firmware upgrade, select **NetApp Support**.
   b. On the NetApp Support Site, select the **Downloads** tab, and then select **E-Series Disk Drive Firmware**.

      The E-Series Disk Firmware page displays.

   c. Search for each **Drive Identifier** installed in your storage appliance and verify that each drive identifier has the latest firmware revision.
      - If the firmware revision is not a link, this drive identifier has the latest firmware revision.
      - If one or more drive part numbers are listed for a drive identifier, a firmware upgrade is available for these drives. You can select any link to download the firmware file.

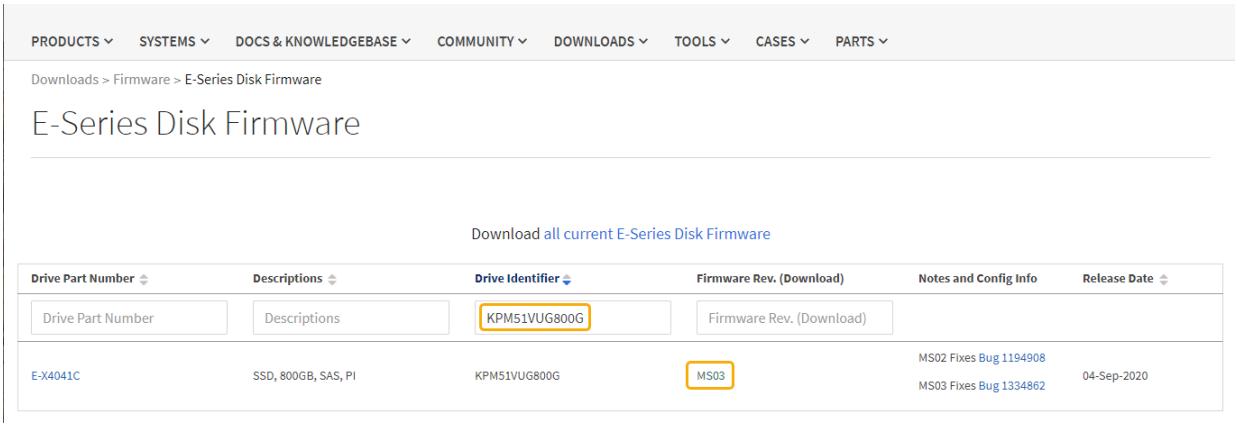

d. If a later firmware revision is listed, select the link in the Firmware Rev. (Download) column to download a `.zip` archive containing the firmware file.

e. Extract (unzip) the drive firmware archive files you downloaded from the Support site.

6. Install the drive firmware upgrade:

a. From SANtricity System Manager, under Drive Firmware upgrade, select **Begin Upgrade**.

b. Select **Browse**, and select the new drive firmware files that you downloaded from the Support site.

Drive firmware files have a filename similar to `D_HUC101212CSS600_30602291_MS01_2800_0002.dlp`.

You can select up to four drive firmware files, one at a time. If more than one drive firmware file is compatible with the same drive, you get a file conflict error. Decide which drive firmware file you want to use for the upgrade and remove the other one.

c. Select **Next**.

Select Drives lists the drives that you can upgrade with the selected firmware files.

Only drives that are compatible appear.

The selected firmware for the drive appears in the **Proposed Firmware** column. If you must change this firmware, select **Back**.

d. Select **Upgrade all drives offline (parallel)** — Upgrades the drives that can support a firmware download only while all I/O activity is stopped on any volumes that use the drives.

You must place the appliance into maintenance mode before using this method. You should use the **Offline** method to upgrade the drive firmware.

If you want to use the Offline (parallel) upgrade, don't proceed unless you are certain that the appliance is in maintenance mode. Failure to place the appliance into maintenance mode before initiating an offline drive firmware update might cause data loss.

e. In the first column of the table, select the drive or drives you want to upgrade.

The best practice is to upgrade all drives of the same model to the same firmware revision.

f. Select **Start** and confirm that you want to perform the upgrade.

If you need to stop the upgrade, select **Stop**. Any firmware downloads currently in progress complete. Any firmware downloads that have not started are canceled.

> Stopping the drive firmware upgrade might result in data loss or unavailable drives.

g. (Optional) To see a list of what was upgraded, select **Save Log**.

The log file is saved in the downloads folder for your browser with the name `latest-upgrade-log-timestamp.txt`.

If required, troubleshoot driver firmware upgrade errors.

7. After the procedure completes successfully, perform any additional maintenance procedures while the node is in maintenance mode. When you are done, or if you experienced any failures and want to start over, go to the StorageGRID Appliance Installer and select **Advanced** > **Reboot Controller**. Then select one of these options:
   - **Reboot into StorageGRID**.
   - **Reboot into Maintenance Mode**. Reboot the controller and keep the node in maintenance mode. Select this option if there were any failures during the procedure and you want to start over. After the node finishes rebooting into maintenance mode, restart from the appropriate step in the procedure that failed.

     It can take up to 20 minutes for the appliance to reboot and rejoin the grid. To confirm that the reboot is complete and that the node has rejoined the grid, go back to the Grid Manager. The Nodes page should display a normal status (green check mark icon to the left of the node name) for the appliance node, indicating that no alerts are active and the node is connected to the grid.

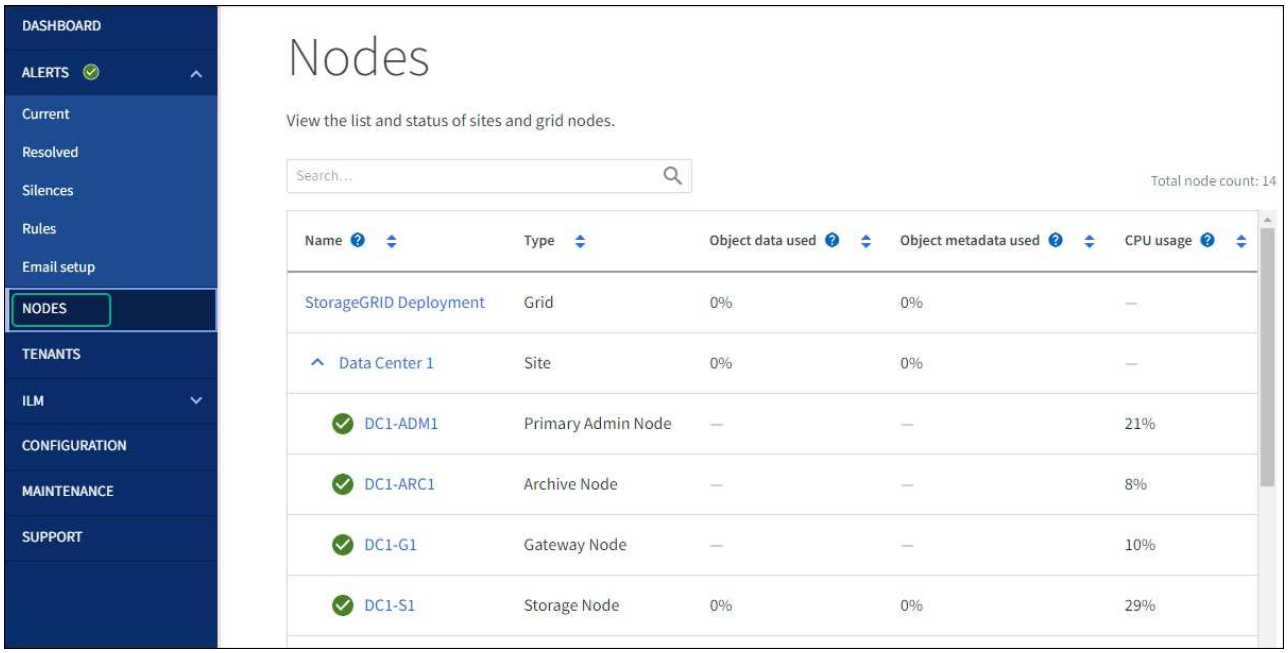

Troubleshoot drive firmware upgrade errors

Troubleshoot errors that can occur when using SANtricity System Manager to upgrade the firmware on the drives in your appliance.

- **Failed assigned drives**
  - One reason for the failure might be that the drive does not have the appropriate signature. Make sure that the affected drive is an authorized drive. Contact technical support for more information.
  - When replacing a drive, make sure that the replacement drive has a capacity equal to or greater than the failed drive you are replacing.
  - You can replace the failed drive while the storage array is receiving I/O.
- **Check storage array**
  - Make sure that an IP address has been assigned to each controller.
  - Make sure that all cables connected to the controller aren't damaged.
  - Make sure that all cables are tightly connected.
- **Integrated hot spare drives**

  This error condition must be corrected before you can upgrade the firmware.

- **Incomplete volume groups**

  If one or more volume groups or disk pools are incomplete, you must correct this error condition before you can upgrade the firmware.

- **Exclusive operations (other than background media/parity scan) currently running on any volume groups**

  If one or more exclusive operations are in progress, the operations must complete before the firmware can be upgraded. Use System Manager to monitor the progress of the operations.

- **Missing volumes**

  You must correct the missing volume condition before the firmware can be upgraded.

- **Either controller in a state other than Optimal**

  One of the storage array controllers needs attention. This condition must be corrected before the firmware can be upgraded.

- **Mismatched Storage Partition information between Controller Object Graphs**

  An error occurred while validating the data on the controllers. Contact technical support to resolve this issue.

- **SPM Verify Database Controller check fails**

  A storage partitions mapping database error occurred on a controller. Contact technical support to resolve this issue.

- **Configuration Database Validation (If supported by the storage array's controller version)**

  A configuration database error occurred on a controller. Contact technical support to resolve this issue.

- **MEL Related Checks**

  Contact technical support to resolve this issue.

- **More than 10 DDE Informational or Critical MEL events were reported in the last 7 days**

  Contact technical support to resolve this issue.

- **More than 2 Page 2C Critical MEL Events were reported in the last 7 days**

  Contact technical support to resolve this issue.

- **More than 2 Degraded Drive Channel Critical MEL events were reported in the last 7 days**

  Contact technical support to resolve this issue.

- **More than 4 critical MEL entries in the last 7 days**

  Contact technical support to resolve this issue.

Change link configuration of E5700SG controller

You can change the Ethernet link configuration of the E5700SG controller. You can change the port bond mode, the network bond mode, and the link speed.

Before you begin

Place E5700SG controller into maintenance mode.

> In rare instances, placing a StorageGRID appliance into maintenance mode might make the appliance unavailable for remote access.

About this task

Options for changing the Ethernet link configuration of the E5700SG controller include:

- Changing **Port bond mode** from Fixed to Aggregate, or from Aggregate to Fixed
- Changing **Network bond mode** from Active-Backup to LACP, or from LACP to Active-Backup
- Enabling or disabling VLAN tagging, or changing the value of a VLAN tag
- Changing the link speed from 10-GbE to 25-GbE, or from 25-GbE to 10-GbE

Steps

1. Select **Configure Networking** > **Link Configuration** from the menu.

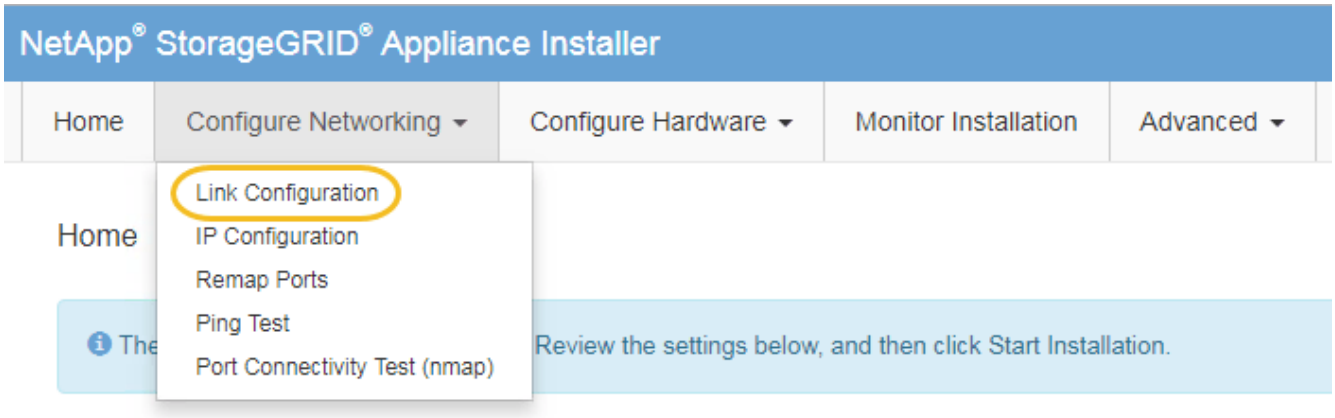

2. Make the desired changes to the link configuration.

   For more information about the options, see Configure network links.

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

   > You might lose your connection if you made changes to the network or link you are connected through. If you aren't reconnected within 1 minute, re-enter the URL for the StorageGRID Appliance Installer using one of the other IP addresses assigned to the appliance:
   > `https://E5700SG_Controller_IP:8443`

   If you made changes to the VLAN settings, the subnet for the appliance might have changed. If you need to change the IP addresses for the appliance, follow the configure StorageGRID IP addresses instructions.

4. From the StorageGRID Appliance Installer, select **Configure Networking** > **Ping Test**.
5. Use the Ping Test tool to check connectivity to IP addresses on any networks that might have been affected by the link configuration changes you made in the Change link configuration step.

   In addition to any other tests you choose to perform, confirm that you can ping the grid IP address of the primary Admin Node, and the grid IP address of at least one other Storage Node. If necessary, correct any link configuration issues.

6. Once you are satisfied that your link configuration changes are working, reboot the node. From the StorageGRID Appliance Installer, select **Advanced** > **Reboot Controller**, and then select one of these options:
   - Select **Reboot into StorageGRID** to reboot the controller with the node rejoining the grid. Select this option if you are done working in maintenance mode and are ready to return the node to normal operation.
   - Select **Reboot into Maintenance Mode** to reboot the controller with the node remaining in maintenance mode. (This option is available only when the controller is in maintenance mode.) Select this option if there are additional maintenance operations you need to perform on the node before rejoining the grid.

   It can take up to 20 minutes for the appliance to reboot and rejoin the grid. To confirm that the reboot is complete and that the node has rejoined the grid, go back to the Grid Manager. The **NODES** page should display a normal status (green check mark icon to the left of the node name) for the

appliance node, indicating that no alerts are active and the node is connected to the grid.

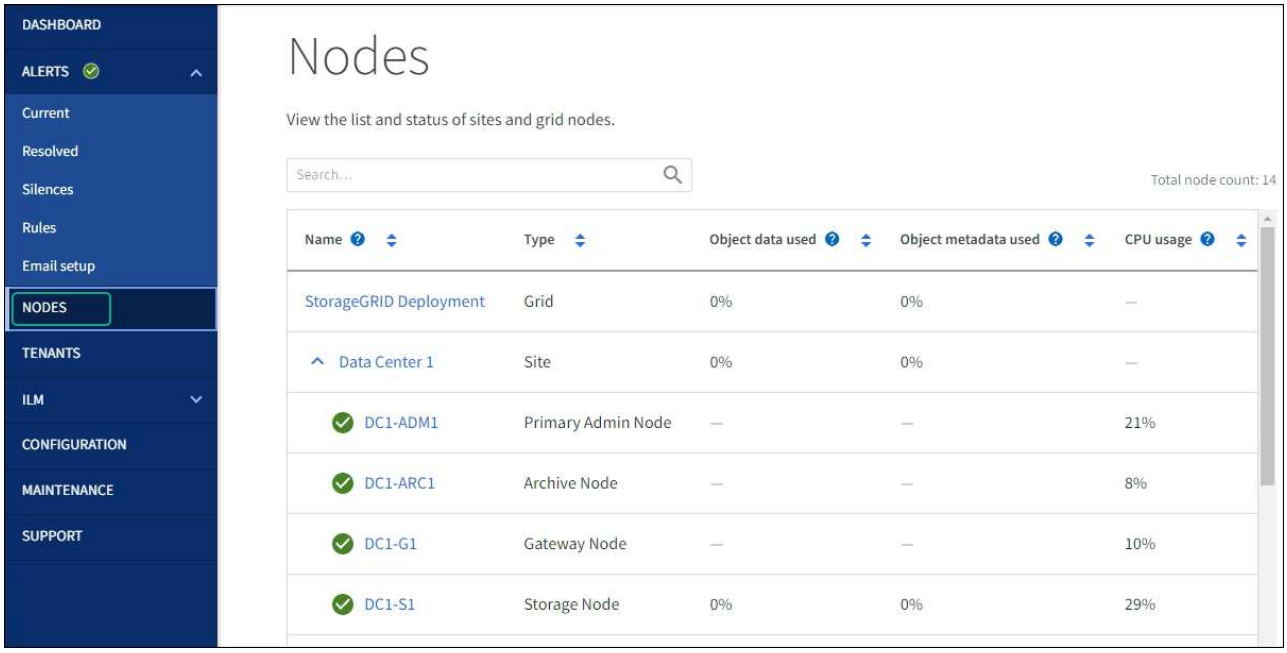

Hardware procedures

Replace E2800 series storage controller in the SG5700

You might need to replace the E2800 series controller if it is not functioning optimally or if it has failed.

About this task

- You have a replacement controller with the same part number as the controller you are replacing.

> Don’t rely on the E-Series instructions to replace a controller in the StorageGRID appliance, because the procedures aren’t the same.

- You have labels to identify each cable that is connected to the controller.
- If all drives are secured, you have reviewed the steps in the simplex E2800 series controller replacement procedure, which include downloading and installing E-Series SANtricity System Manager from the NetApp Support Site and then using the Enterprise Management Window (EMW) to unlock the secured drives after you have replaced the controller.

> You will not be able to use the appliance until you unlock the drives with the saved key.

- You must have specific access permissions.
- You must be signed in to the Grid Manager using a supported web browser.

About this task

You can determine if you have a failed controller canister in two ways:

- The Recovery Guru in SANtricity System Manager directs you to replace the controller.
- The amber Attention LED on the controller is on, indicating that the controller has a fault.

The appliance Storage Node will not be accessible when you replace the controller. If the E2800 series controller is functioning sufficiently, you can place the E5700SG controller into maintenance mode.

When you replace a controller, you must remove the battery from the original controller and install it in the replacement controller. In some cases, you might also need to remove the host interface card from the original controller and install it in the replacement controller.

The storage controllers in most appliance models don't include host interface cards (HIC).

Step 1: Prepare to remove the controller

These figures show the E2800A controller and the E2800B controller. The procedure for replacing the E2800 series controllers and the EF570 controller is identical.

E2800A storage controller:

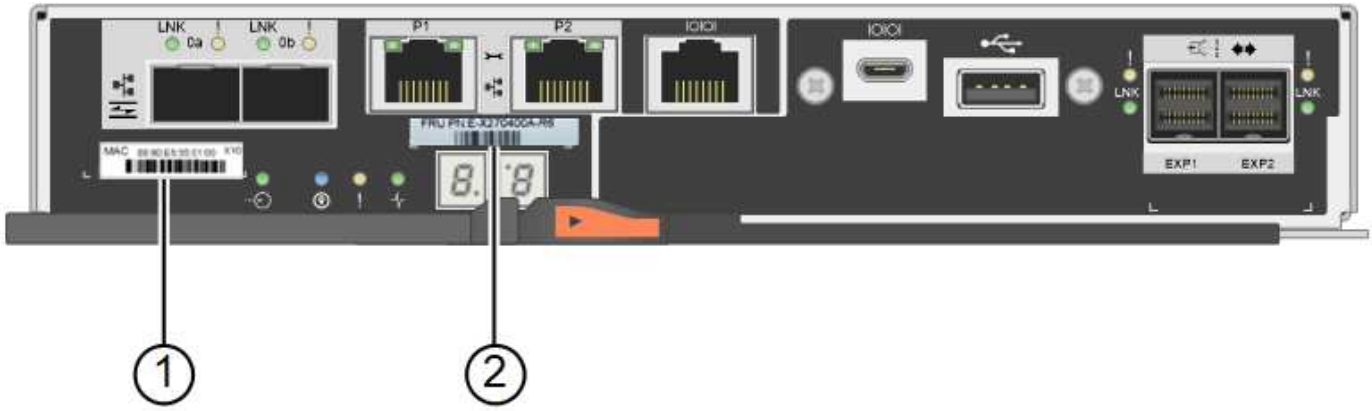

E2800B storage controller:

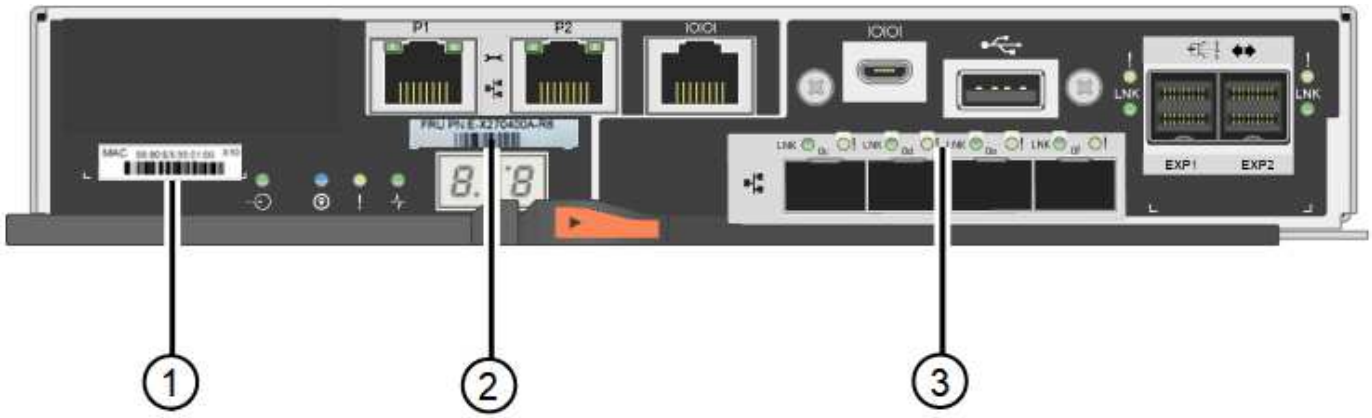

Label	component	Description
1	MAC address	The MAC address for management port 1 ("P1 on the E2800A and 0a on the E2800B"). If you used DHCP to obtain the original controller's IP address, you will need this address to connect to the new controller.

Label	component	Description
2	FRU part number	The FRU part number. This number must match the replacement part number for the currently installed controller.
3	4-port HIC	The 4-port host interface card (HIC). This card must be moved to the new controller when you perform the replacement. **Note**: the E2800A controller does not have a HIC.

Follow the instructions in the E2800 controller replacement procedure to prepare to remove the controller.

You use SANtricity System Manager to perform these steps.

Steps

1. Make a note of which version of SANtricity OS software is currently installed on the controller.
2. Make a note of which version of NVSRAM is currently installed.
3. If the Drive Security feature is enabled, be sure a saved key exists and that you know the pass phrase required to install it.

> **Possible loss of data access —** If all drives in the appliance are security enabled, the new controller will not be able to access the appliance until you unlock the secured drives using the Enterprise Management Window in SANtricity System Manager.

4. Back up the configuration database.

   If a problem occurs when you remove a controller, you can use the saved file to restore your configuration.

5. Collect support data for the appliance.

> Collecting support data before and after replacing a component ensures you can send a full set of logs to technical support if the replacement does not resolve the problem.

Step 2: Take the controller offline

Take the controller offline and confirm that all operations are complete.

Steps

1. If the StorageGRID appliance is running in a StorageGRID system, place the E5700SG controller into maintenance mode.
2. If the E2800 controller is functioning sufficiently to allow for a controlled shutdown, confirm that all operations have completed.
3. From the home page of SANtricity System Manager, select **View Operations in Progress**.
4. Confirm that all operations have completed.
5. Power off the controller shelf.

Step 3: Remove the controller

Remove the controller from the appliance.

Steps

1. Put on an ESD wristband or take other antistatic precautions.
2. Label the cables and then disconnect the cables and SFPs.

   To prevent degraded performance, don’t twist, fold, pinch, or step on the cables.

3. Release the controller from the appliance by squeezing the latch on the cam handle until it releases, and then open the cam handle to the right.
4. Using two hands and the cam handle, slide the controller out of the appliance.

   Always use two hands to support the weight of the controller.

5. Place the controller on a flat, static-free surface with the removable cover facing up.
6. Remove the cover by pressing down on the button and sliding the cover off.

Step 4: Move battery to the new controller

Remove the battery from the failed controller, and install it into the replacement controller.

Steps

1. Confirm that the green LED inside the controller (between the battery and the DIMMs) is off.

   If this green LED is on, the controller is still using battery power. You must wait for this LED to go off before removing any components.

Item	Description
1	Internal Cache Active LED
2	Battery

2. Locate the blue release latch for the battery.
3. Unlatch the battery by pushing the release latch down and away from the controller.

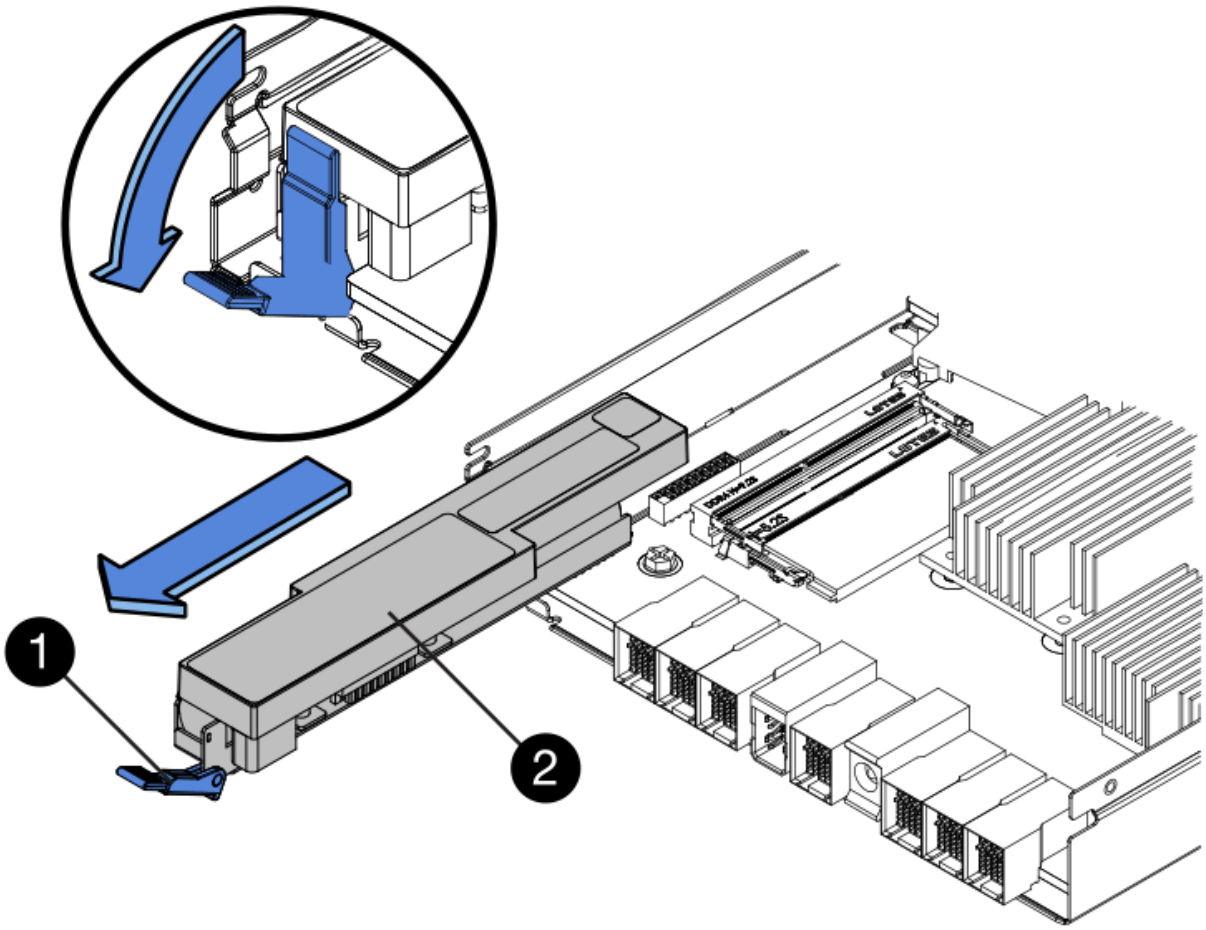

Item	Description
1	Battery release latch
2	Battery

4. Lift up on the battery, and slide it out of the controller.
5. Remove the cover from the replacement controller.
6. Orient the replacement controller so that the slot for the battery faces toward you.
7. Insert the battery into the controller at a slight downward angle.

   You must insert the metal flange at the front of the battery into the slot on the bottom of the controller, and slide the top of the battery beneath the small alignment pin on the left side of the controller.

8. Move the battery latch up to secure the battery.

   When the latch clicks into place, the bottom of the latch hooks into a metal slot on the chassis.

9. Turn the controller over to confirm that the battery is installed correctly.

Possible hardware damage — The metal flange at the front of the battery must be completely inserted into the slot on the controller (as shown in the first figure). If the battery is not installed correctly (as shown in the second figure), the metal flange might contact the controller board, causing damage.

- **Correct — The battery's metal flange is completely inserted into the slot on the controller:**

- **Incorrect — The battery's metal flange is not inserted into the slot on the controller:**

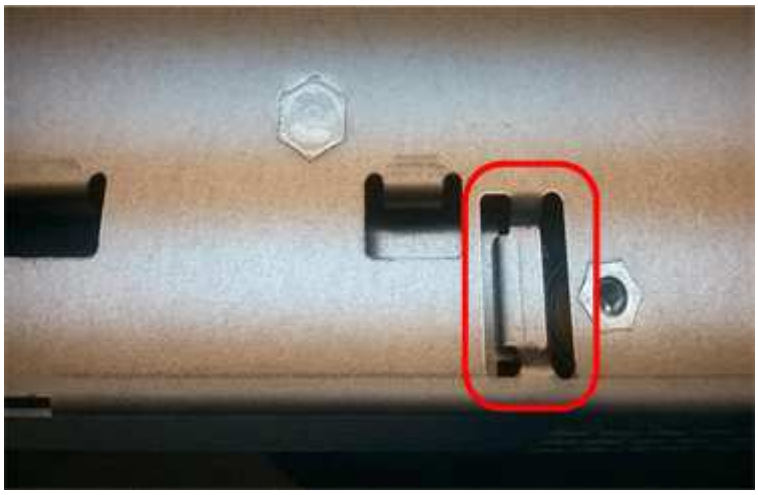

10. Replace the controller cover.

Step5: Move HIC to new controller, if needed

If the failed controller includes a host interface card (HIC), move the HIC from the failed controller to the replacement controller.

A separate HIC is used for the E2800B controller only. The HIC is mounted to the main controller board and includes two SPF connectors.

The illustrations in this procedure show a 2-port HIC. The HIC in your controller might have a different number of ports.

E2800A

An E2800A controller does not have a HIC.

Replace the E2800A controller cover, and go to Step 6: Replace controller

E2800B

Move the HIC from the failed E2800B controller to the replacement controller.

Steps

1. Remove any SFPs from the HIC.
2. Using a #1 Phillips screwdriver, remove the screws that attach the HIC faceplate to the controller.

   There are four screws: one on the top, one on the side, and two on the front.

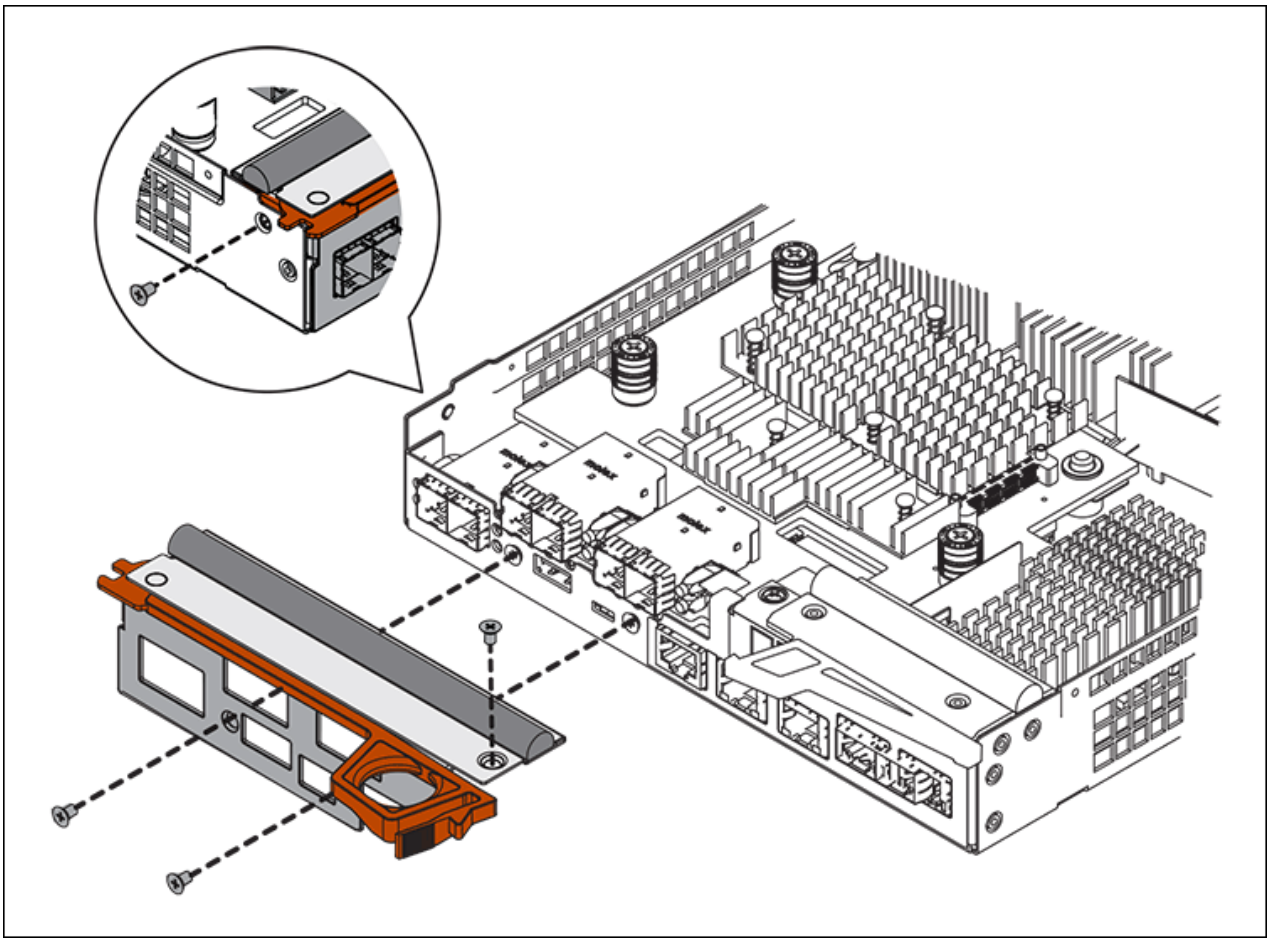

3. Remove the HIC faceplate.
4. Using your fingers or a Phillips screwdriver, loosen the three thumbscrews that secure the HIC to the controller card.
5. Carefully detach the HIC from the controller card by lifting the card up and sliding it back.

> Be careful not to scratch or bump the components on the bottom of the HIC or on the top of the controller card.

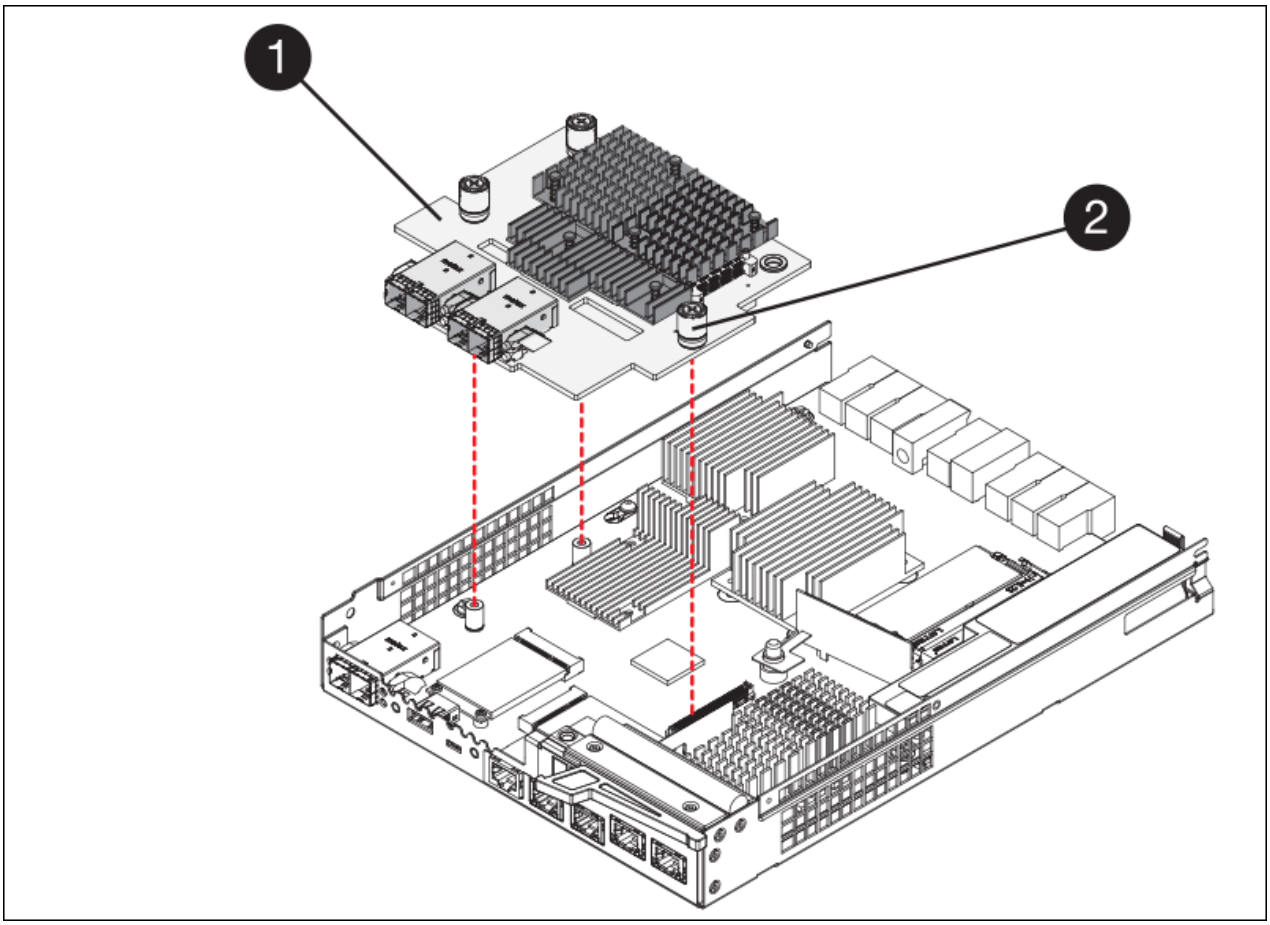

Label	Description
1	Host interface card
2	Thumbscrews

6. Place the HIC on a static-free surface.
7. Using a #1 Phillips screwdriver, remove the four screws that attach the blank faceplate to the replacement controller, and remove the faceplate.
8. Align the three thumbscrews on the HIC with the corresponding holes on the replacement controller, and align the connector on the bottom of the HIC with the HIC interface connector on the controller card.

   Be careful not to scratch or bump the components on the bottom of the HIC or on the top of the controller card.

9. Carefully lower the HIC into place, and seat the HIC connector by pressing gently on the HIC.

Possible equipment damage — Be careful not to pinch the gold ribbon connector for the controller LEDs between the HIC and the thumbscrews.

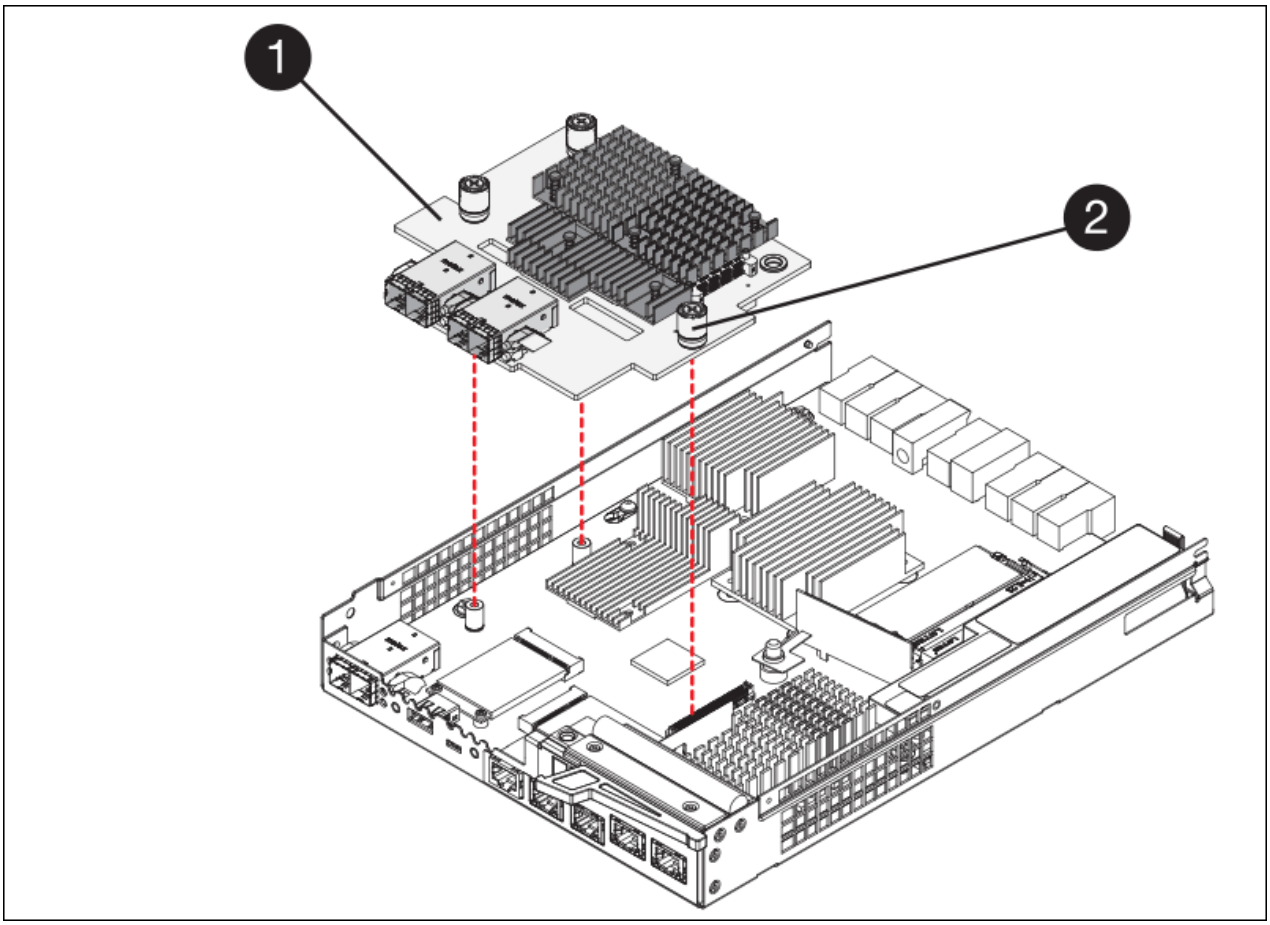

Label	Description
1	Host interface card
2	Thumbscrews

10. Hand-tighten the HIC thumbscrews.

    Don’t use a screwdriver, or you might over tighten the screws.

11. Using a #1 Phillips screwdriver, attach the HIC faceplate you removed from the original controller to the new controller with four screws.

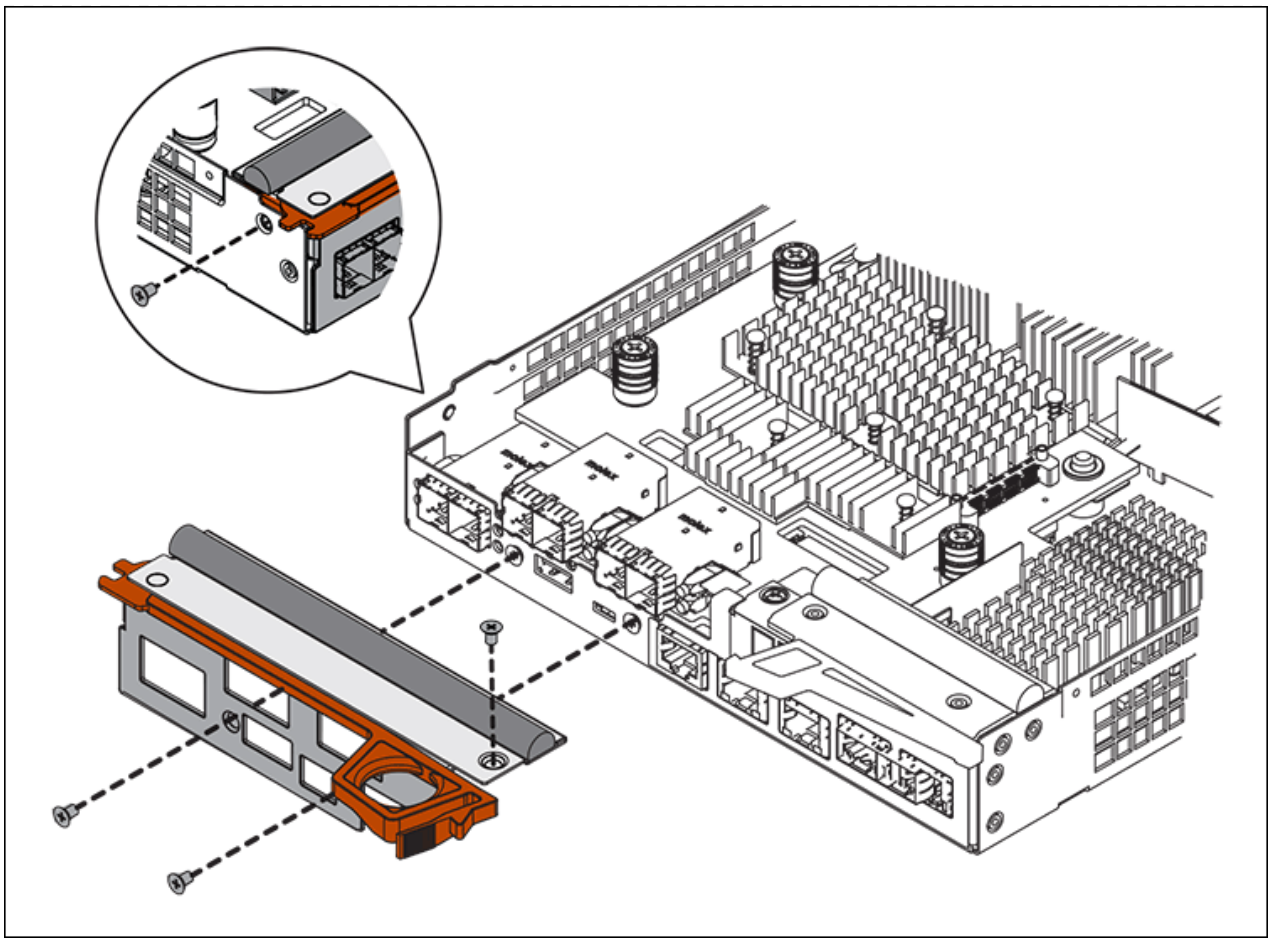

12. Reinstall any removed SFPs into the HIC.

Step 6: Replace controller

Install the replacement controller and verify that it has rejoined the grid.

Steps

1. Install the replacement controller into the appliance.
   a. Turn the controller over, so that the removable cover faces down.
   b. With the cam handle in the open position, slide the controller all the way into the appliance.
   c. Move the cam handle to the left to lock the controller in place.
   d. Replace the cables and SFPs.
   e. Power on the controller shelf.
   f. Wait for the E2800 controller to reboot. Verify that the seven-segment display shows a state of `99`.
   g. Determine how you will assign an IP address to the replacement controller.

> The steps for assigning an IP address to the replacement controller depend on whether you connected management port 1 to a network with a DHCP server and on whether all drives are secured.

If management port 1 is connected to a network with a DHCP server, the new controller will obtain its

IP address from the DHCP server. This value might be different than the original controller's IP address.

2. If the appliance uses secured drives, follow the instructions in the E2800 controller replacement procedure to import the drive security key.
3. Return the appliance to normal operating mode. From the StorageGRID Appliance Installer, select **Advanced** > **Reboot Controller**, and then select **Reboot into StorageGRID**.

4. During the reboot, monitor the node's status to determine when it has rejoined the grid.

   The appliance reboots and rejoins the grid. This process can take up to 20 minutes.

5. Confirm that the reboot is complete and that the node has rejoined the grid. In the Grid Manager, verify that the Nodes page displays a normal status (green check mark icon to the left of the node name) for the appliance node, indicating that no alerts are active and the node is connected to the grid.

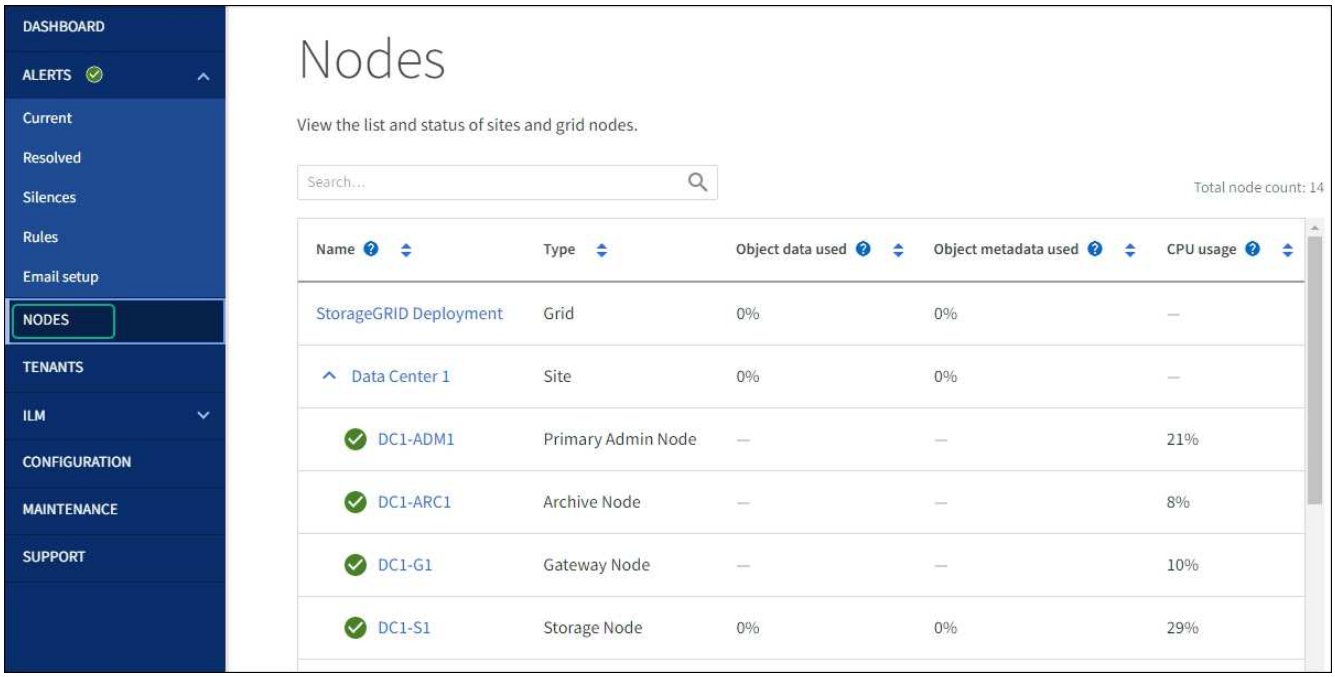

6. From SANtricity System Manager, confirm that the new controller is Optimal, and collect support data.

After replacing the part, return the failed part to NetApp, as described in the RMA instructions shipped with the kit. See the Part Return & Replacements page for further information.

Related information

NetApp E-Series Systems Documentation Site

Replace E5700SG compute controller

You might need to replace the E5700SG controller if it is not functioning optimally or if it has failed.

Before you begin

- You have a replacement controller with the same part number as the controller you are replacing.
- You have downloaded the E-Series instructions for replacing a failed E5700 controller.

> Use the E-Series instructions for reference only if you need more details to perform a specific step. Don't rely on the E-Series instructions to replace a controller in the StorageGRID appliance, because the procedures aren't the same. For example, the E-Series instructions for the E5700 controller describe how to remove the battery and the host interface card (HIC) from a failed controller and install them in a replacement controller. These steps don't apply to the E5700SG controller.

- You have labels to identify each cable that is connected to the controller.

About this task

The appliance Storage Node will not be accessible when you replace the controller. If the E5700SG controller is functioning sufficiently, you can perform a controlled shutdown at the start of this procedure.

> If you are replacing the controller before installing StorageGRID software, you might not be able to access the StorageGRID Appliance Installer immediately after completing this procedure. While you can access the StorageGRID Appliance Installer from other hosts on the same subnet as the appliance, you can't access it from hosts on other subnets. This condition should resolve itself within 15 minutes (when any ARP cache entries for the original controller time out), or you can clear the condition immediately by purging any old ARP cache entries manually from the local router or gateway.

Steps

1. Shut down the E5700SG controller.
   a. Log in to the grid node:
      i. Enter the following command: `ssh admin@grid_node_IP`
      ii. Enter the password listed in the `Passwords.txt` file.
      iii. Enter the following command to switch to root: `su -`
      iv. Enter the password listed in the `Passwords.txt` file.

      When you are logged in as root, the prompt changes from `$` to `#`.

   b. Shut down the E5700SG controller:
      `shutdown -h now`

c. Wait for any data in cache memory to be written to the drives.

   The green Cache Active LED on the back of the E2800 controller is on when cached data needs to be written to the drives. You must wait for this LED to turn off.

2. Turn off the power.
   a. From the home page of SANtricity System Manager, select **View Operations in Progress**.
   b. Confirm that all operations have completed.
   c. Turn off both power switches on the appliance.
   d. Wait for all LEDs to turn off.
3. If the StorageGRID networks attached to the controller use DHCP servers:
   a. Note the MAC addresses for the ports on the replacement controller (located on labels on the controller).
   b. Ask your network administrator to update the IP address settings for the original controller to reflect the MAC addresses for the replacement controller.

   > You must ensure that the IP addresses for the original controller have been updated before you apply power to the replacement controller. Otherwise, the controller will obtain new DHCP IP addresses when it boots up and might not be able to reconnect to StorageGRID. This step applies to all StorageGRID networks that are attached to the controller.

4. Remove the controller from the appliance:
   a. Put on an ESD wristband or take other antistatic precautions.
   b. Label the cables and then disconnect the cables and SFPs.

      > To prevent degraded performance, don’t twist, fold, pinch, or step on the cables.

   c. Release the controller from the appliance by squeezing the latch on the cam handle until it releases, and then open the cam handle to the right.
   d. Using two hands and the cam handle, slide the controller out of the appliance.

      > Always use two hands to support the weight of the controller.

5. Install the replacement controller into the appliance.
   a. Turn the controller over, so that the removable cover faces down.
   b. With the cam handle in the open position, slide the controller all the way into the appliance.
   c. Move the cam handle to the left to lock the controller in place.
   d. Replace the cables and SFPs.
6. Power on the appliance, and monitor the controller LEDs and seven-segment displays.

   After the controllers have successfully booted up, the seven-segment displays should show the following:

   - E2800 series controller:

     The final state is `99`.

- E5700SG controller:

  The final state is `HA`.

7. Confirm that the appliance Storage Node appears in the Grid Manager and that no alarms appear.

After replacing the part, return the failed part to NetApp, as described in the RMA instructions shipped with the kit. See the Part Return & Replacements page for further information.

Related information

NetApp E-Series Systems Documentation Site

Replace other hardware components

You might need to replace a controller battery, drive, fan, or power supply, in the StorageGRID appliance.

Before you begin

- You have the E-Series hardware replacement procedure.
- The appliance has been placed into maintenance mode if the component replacement procedure requires that you shut down the appliance.

About this task

To replace the battery in the E2800 controller, see the instructions in these instructions for replacing the E2800 controller. Those instructions describe how to remove the controller from the appliance, remove the battery from the controller, install the battery, and replace the controller.

To replace a drive, power-fan canister, fan canister, power canister, or drive drawer in the appliance, access the E-Series procedures for maintaining E2800 hardware.

SG5712 component replacement instructions

FRU	See E-Series instructions for
Drive	Replacing a drive in E2800 12-drive or 24-drive shelves
Power-fan canister	Replacing a power-fan canister in E2800 shelves

SG5760 component replacement instructions

FRU	See E-Series instructions for
Drive	Replacing a drive in E2860 shelves
Power canister	Replacing a power canister in E2860 shelves
Fan canister	Replacing a fan canister in E2860 shelves
Drive drawer	Replacing a drive drawer in E2860 shelves

Related information

Replace E2800 controller

NetApp E-Series Systems Documentation Site

Maintain SG5600 appliance

The SG5600 appliance has reached End Of Support. Contact your NetApp Sales Representative for hardware refresh options.

If you need to perform maintenance procedures on SG5600 hardware, use the StorageGRID 11.6 instructions.

Maintain SG100 and SG1000 hardware

Maintain SG100 and SG1000 appliances

You might need to perform maintenance procedures on the appliance. The procedures in this section assume that the appliance has already been deployed as a Gateway Node or an Admin Node in a StorageGRID system.

Procedures specific to maintaining your SG100 or SG1000 appliance are in this section.

See Common procedures for maintenance procedures that are used by all appliances.

See Set up hardware for maintenance procedures that are also performed during initial appliance installation and configuration.

Maintenance configuration procedures

Turn controller identify LED on and off

The blue identify LED on the front and back of the controller can be turned on to help locate the appliance in a data center.

Before you begin

You have the BMC IP address of the controller you want to identify.

Steps

1. Access the appliance BMC interface.
2. Select **Server Identify**.

   The current status of the identify LED is selected.

3. Select **ON** or **OFF**, and then select **Perform Action**.

   When you select **ON**, the blue identify LEDs light on the front (shown) and rear of the appliance.

If a bezel is installed on the controller, it might be difficult to see the front identify LED.

4. Turn the LED on and off as needed.

Related information

Locate controller in data center

Access BMC interface

Locate controller in data center

Locate the controller so that you can perform hardware maintenance or upgrades.

Before you begin

- You have determined which controller requires maintenance.
- (Optional) To help locate the controller in your data center, turn on the blue identify LED.

Steps

1. Find the controller requiring maintenance in the data center.
   - Look for a lit blue identify LED on the front or rear of the controller.

     The front identify LED is behind the controller front bezel and might be difficult to see if the bezel is installed.

     

   - Check the tags attached to the front of each controller for a matching part number.

2. Remove the controller front bezel, if one is installed, to access the front panel controls and indicators.
3. Optional: Turn off the blue identify LED if you used it to locate the controller.
   - Press the identify LED switch on the controller front panel.
   - Use the controller BMC interface.

Shut down the services appliance

Shut down the services appliance to perform hardware maintenance.

Before you begin

- You have physically located the services appliance requiring maintenance in the data center.

  Locating the controller in a data center.

About this task

To prevent service interruptions, shut down the services appliance during a scheduled maintenance window when periods of service disruption are normally expected.

Steps

1. Shut down the appliance:

   > You must perform a controlled shut down of the appliance by entering the commands specified below. It is a best practice to perform a controlled shutdown when possible to avoid unnecessary alerts, ensure full logs are available, and avoid service disruptions.

   a. If you have not already logged into the grid node, log in using PuTTY or another ssh client:
      i. Enter the following command: `ssh admin@grid_node_IP`
      ii. Enter the password listed in the `Passwords.txt` file.
      iii. Enter the following command to switch to root: `su -`
      iv. Enter the password listed in the `Passwords.txt` file.

      When you are logged in as root, the prompt changes from `$` to `#`.

   b. Shut down the services appliance:
      `shutdown -h now`

      This command might take up to 10 minutes to complete.

2. Use one of the following methods to verify that the appliance is powered off:
   - Look at the power LED on the front of the appliance and confirm that it is off.
   - Check the Power Control page of the BMC interface to confirm the appliance is off.

Change link configuration of services appliance

You can change the Ethernet link configuration of the services appliance. You can change the port bond mode, the network bond mode, and the link speed.

Before you begin

- You have placed the appliance into maintenance mode.

> In rare instances, placing a StorageGRID appliance into maintenance mode might make the appliance unavailable for remote access.

About this task

Options for changing the Ethernet link configuration of the services appliance include:

- Changing **Port bond mode** from Fixed to Aggregate, or from Aggregate to Fixed
- Changing **Network bond mode** from Active-Backup to LACP, or from LACP to Active-Backup
- Enabling or disabling VLAN tagging, or changing the value of a VLAN tag
- Changing the link speed

Steps

1. From the StorageGRID Appliance Installer, select **Configure Networking** > **Link Configuration**.

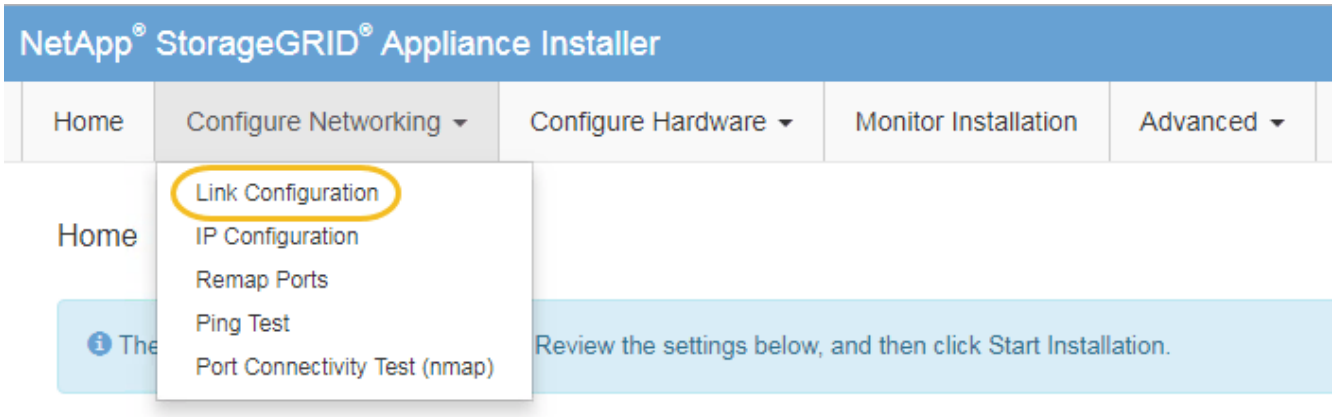

2. Make the desired changes to the link configuration.

   For more information about the options, see Configure network links.

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

> You might lose your connection if you made changes to the network or link you are connected through. If you aren't reconnected within 1 minute, re-enter the URL for the StorageGRID Appliance Installer using one of the other IP addresses assigned to the appliance:
> `https://services_appliance_IP:8443`

4. Make any necessary changes to the IP addresses for the appliance.

   If you made changes to the VLAN settings, the subnet for the appliance might have changed. If you need to change the IP addresses for the appliance, see Configure StorageGRID IP addresses.

5. Select **Configure Networking** > **Ping Test** from the menu.
6. Use the Ping Test tool to check connectivity to IP addresses on any networks that might have been affected by the link configuration changes you made when configuring the appliance.

In addition to any other tests you choose to perform, confirm that you can ping the Grid Network IP address of the primary Admin Node, and the Grid Network IP address of at least one other node. If necessary, return to the instructions for configuring network links, and correct any issues.

7. Once you are satisfied that your link configuration changes are working, reboot the node. From the StorageGRID Appliance Installer, select **Advanced** > **Reboot Controller**, and then select one of these options:
   - Select **Reboot into StorageGRID** to reboot the controller with the node rejoining the grid. Select this option if you are done working in maintenance mode and are ready to return the node to normal operation.
   - Select **Reboot into Maintenance Mode** to reboot the controller with the node remaining in maintenance mode. (This option is available only when the controller is in maintenance mode.) Select this option if there are additional maintenance operations you need to perform on the node before rejoining the grid.

It can take up to 20 minutes for the appliance to reboot and rejoin the grid. To confirm that the reboot is complete and that the node has rejoined the grid, go back to the Grid Manager. The **NODES** page should display a normal status (green check mark icon to the left of the node name) for the appliance node, indicating that no alerts are active and the node is connected to the grid.

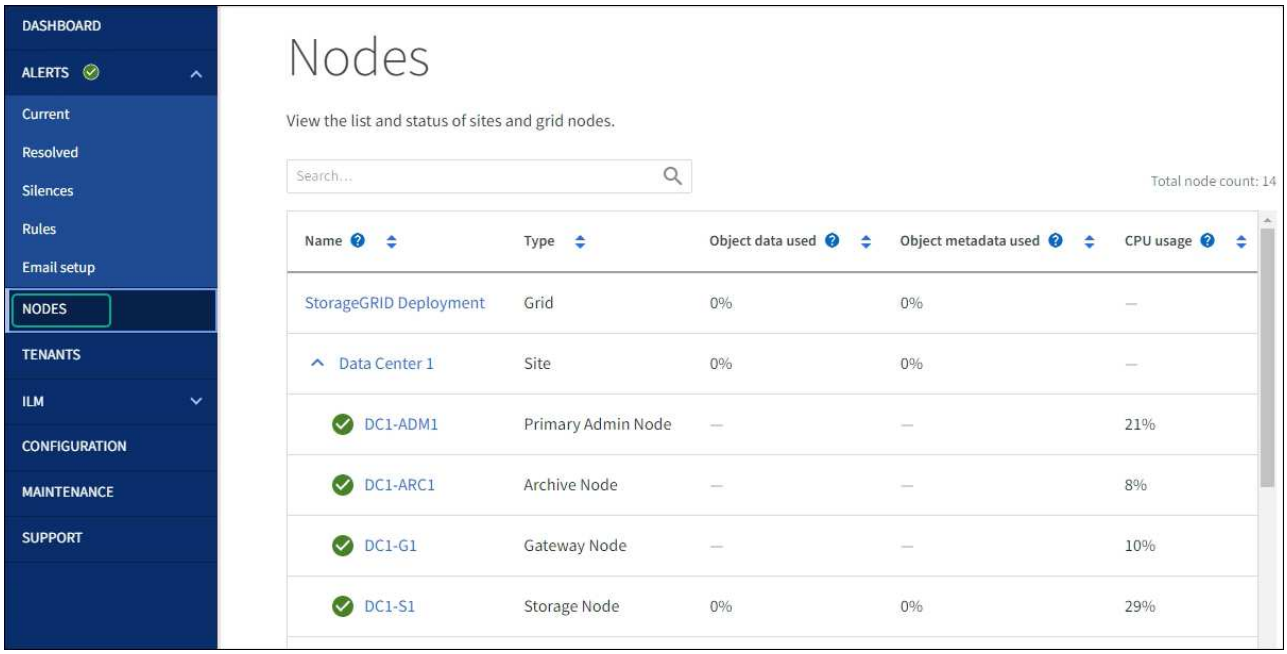

Hardware procedures

Replace one or both power supplies in the services appliance

The services appliance has two power supplies for redundancy. If one of the power supplies fails, you must replace it as soon as possible to ensure that the compute controller has redundant power. Both power supplies operating in the controller must be the same model and wattage.

Before you begin

- You have physically located the controller with the power supply to be replaced.
- If you are replacing only one power supply:
  - You have unpacked the replacement power supply unit and ensured that it is the same model and wattage as the power supply unit you are replacing.
  - You have confirmed that the other power supply is installed and running.
- If you are replacing both power supplies at the same time:
  - You have unpacked the replacement power supply units and ensured they are the same model and wattage.

About this task

The figure shows the two power supply units for the SG100, which are accessible from the back of the appliance.

The power supplies for the SG1000 are identical.

Steps

1. If you are replacing only one power supply, you don't need to shut down the appliance. Go to the Unplug the power cord step. If you are replacing both power supplies at the same time, do the following before unplugging the power cords:
   a. Shut down the appliance.
2. Unplug the power cord from each power supply to be replaced.
3. Lift the cam handle on the first supply to be replaced.

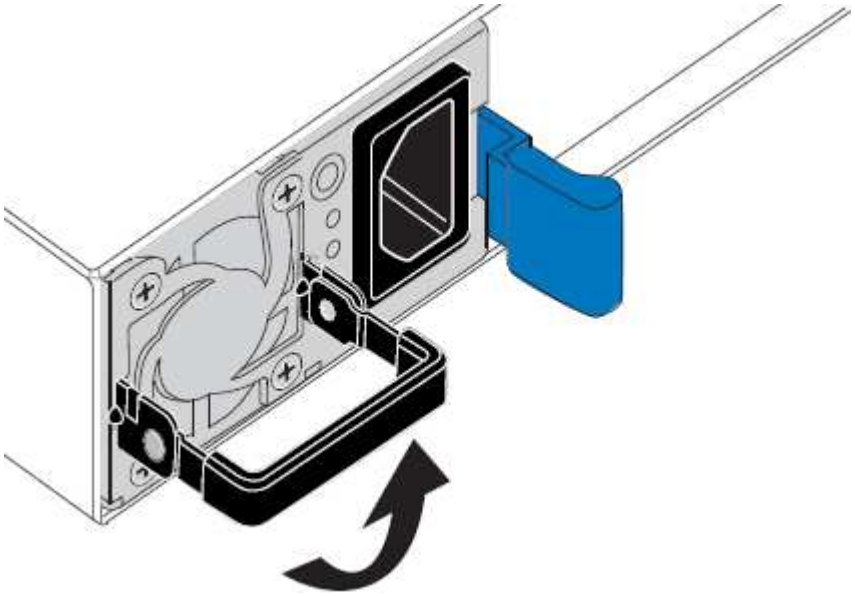

4. Press the blue latch and pull the power supply out.

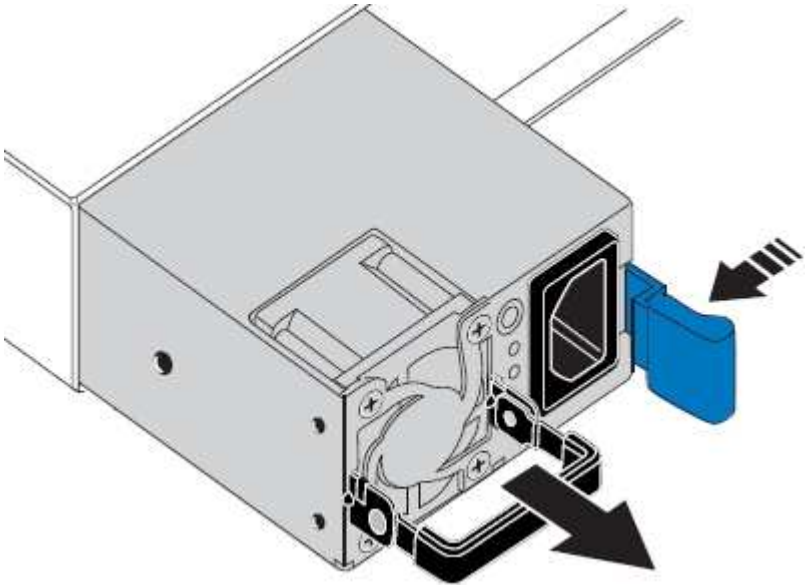

5. With the blue latch on the right, slide the replacement power supply into the chassis.

   Both power supplies must be the same model and wattage.

   Ensure that the blue latch is on the right side when you slide the replacement unit in.

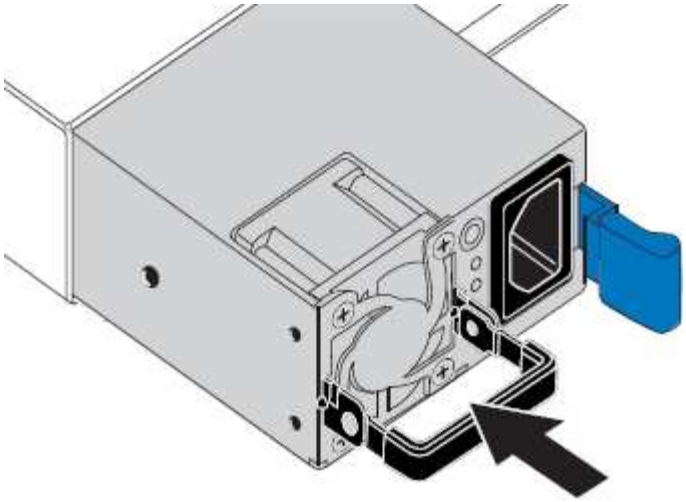

6. Push the cam handle down to secure the replacement power supply.
7. If you are replacing both power supplies, repeat steps 2 though 6 to replace the second power supply.
8. Connect the power cords to the replaced units and apply power.

After replacing the part, return the failed part to NetApp, as described in the RMA instructions shipped with the kit. See the Part Return & Replacements page for further information.

Replace fan in services appliance

The services appliance has eight cooling fans. If one of the fans fails, you must replace it as soon as possible to ensure that the appliance has proper cooling.

Before you begin

- You have unpacked the replacement fan.
- You have physically located the appliance.
- You have confirmed that the other fans are installed and running.

About this task

The appliance node will not be accessible while you replace the fan.

The photograph shows a fan for the services appliance. The cooling fans are accessible after you take the top cover off the appliance.

Each of the two power supply units also contain a fan. Those fans aren't included in this procedure.

Steps

1. Shut down the appliance.
   a. Log in to the grid node:
      i. Enter the following command: `ssh admin@grid_node_IP`
      ii. Enter the password listed in the `Passwords.txt` file.
      iii. Enter the following command to switch to root: `su -`
      iv. Enter the password listed in the `Passwords.txt` file.

      When you are logged in as root, the prompt changes from `$` to `#`.

   b. Shut down the services appliance:
      `shutdown -h now`
2. Use one of two methods to verify that the power for the services appliance is off:
   - The power indicator LED on the front of the appliance is off.
   - The Power Control page of the BMC interface indicates that the appliance is off.
3. Pull the appliance out of the rack.
4. Lift the latch on the top cover and remove the cover from the appliance.
5. Locate the fan that failed.

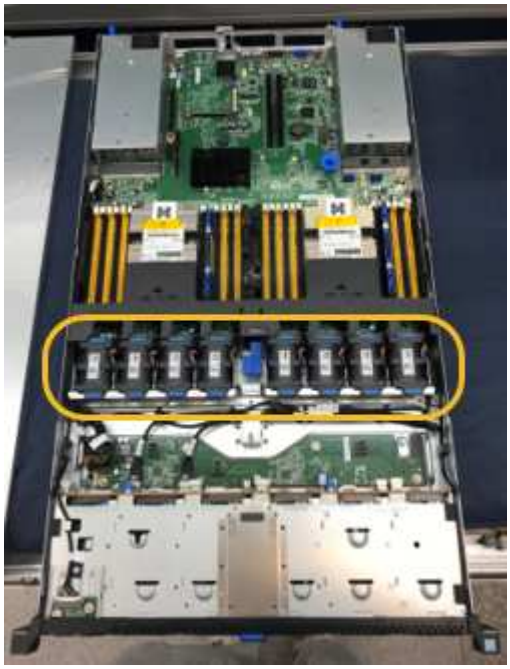

6. Lift the failed fan out of the chassis.

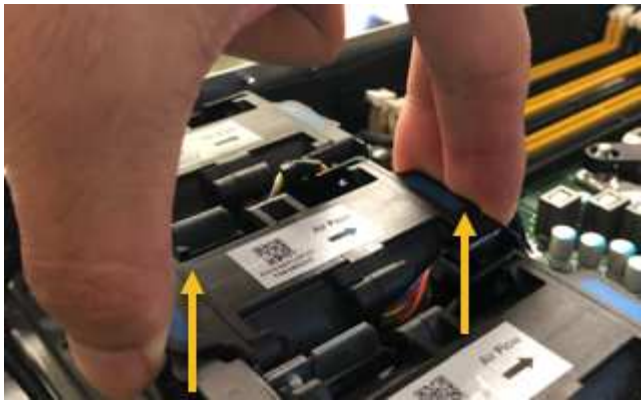

7. Slide the replacement fan into the open slot in the chassis.

   Line up the edge of the fan with the guide pin. The pin is circled in the photograph.

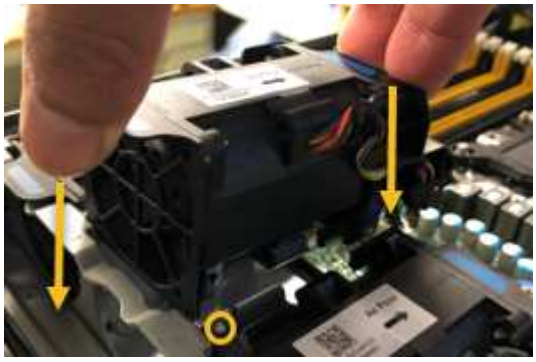

8. Press the fan's connector firmly into the circuit board.

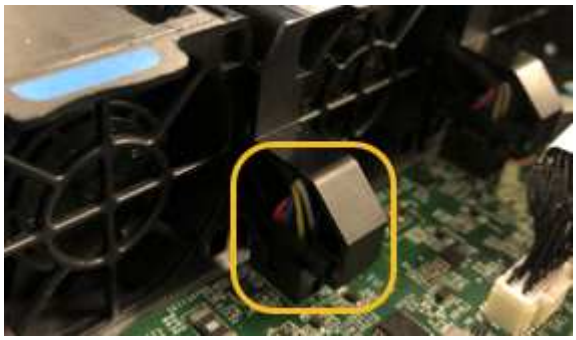

9. Put the top cover back on the appliance, and press the latch down to secure the cover in place.
10. Power on the appliance and monitor the controller LEDs and boot-up codes.

    Use the BMC interface to monitor boot-up status.

11. Confirm that the appliance node appears in the Grid Manager and that no alerts appear.

After replacing the part, return the failed part to NetApp, as described in the RMA instructions shipped with the kit. See the Part Return & Replacements page for further information.

Replace drive in services appliance

The SSDs in the services appliance contain the StorageGRID operating system. Additionally, when the appliance is configured as an Admin Node, the SSDs also contain

audit logs, metrics, and database tables. The drives are mirrored using RAID1 for redundancy. If one of the drives fails, you must replace it as soon as possible to ensure redundancy.

Before you begin

- You have physically located the appliance.
- You have verified which drive has failed by noting that its left LED is blinking amber.

  The two SSDs are placed in the slots as shown in the following diagram:

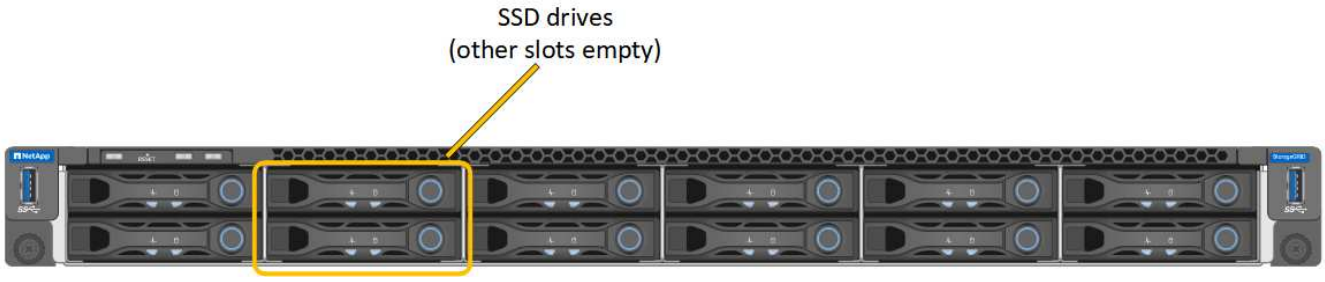

> If you remove the working drive, you will bring down the appliance node. See the information about viewing status indicators to verify the failure.

- You have obtained the replacement drive.
- You have obtained proper ESD protection.

Steps

1. Verify that the drive to be replaced has the left LED blinking amber. If a drive issue was reported in the Grid Manager or BMC UIs, HDD02 or HDD2 refer to the drive in the upper slot, and HDD03 or HDD3 refer to the drive in the lower slot.

   You can also use the Grid Manager to monitor the status of the SSDs. Select **NODES**. Then select `Appliance Node` > **Hardware**. If a drive has failed, the Storage RAID Mode field contains a message about which drive has failed.

2. Wrap the strap end of the ESD wristband around your wrist, and secure the clip end to a metal ground to prevent static discharge.
3. Unpack the replacement drive, and set it on a static-free, level surface near the appliance.

   Save all packing materials.

4. Press the release button on the failed drive.

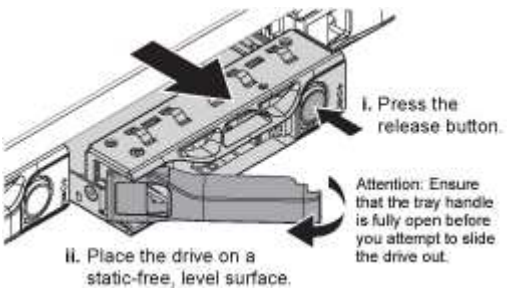

   The handle on the drive springs open partially, and the drive releases from the slot.

5. Open the handle, slide the drive out, and place it on a static-free, level surface.
6. Press the release button on the replacement drive before you insert it into the drive slot.

   The latch springs open.

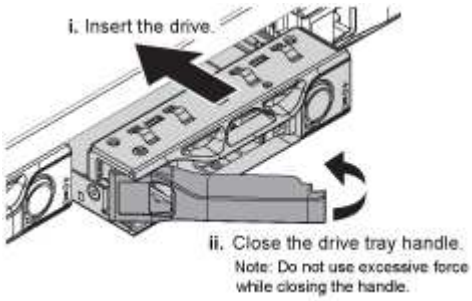

7. Insert the replacement drive in the slot, and then close the drive handle.

   Don’t use excessive force while closing the handle.

   When the drive is fully inserted, you hear a click.

   The drive is automatically rebuilt with mirrored data from the working drive. You can check the status of the rebuild by using the Grid Manager. Select **NODES**. Then select **`Appliance Node`** > **Hardware**. The Storage RAID Mode field contains a “rebuilding” message until the drive is completely rebuilt.

After replacing the part, return the failed part to NetApp, as described in the RMA instructions shipped with the kit. See the Part Return & Replacements page for further information.

Replace services appliance

You might need to replace the appliance if it is not functioning optimally or if it has failed.

Before you begin

- You have a replacement appliance with the same part number as the appliance you are replacing.
- You have labels to identify each cable that is connected to the appliance.
- You have physically located the appliance.

About this task

The StorageGRID node will not be accessible while you replace the appliance. If the appliance is functioning sufficiently, you can perform a controlled shutdown at the start of this procedure.

If you are replacing the appliance before installing StorageGRID software, you might not be able to access the StorageGRID Appliance Installer immediately after completing this procedure. While you can access the StorageGRID Appliance Installer from other hosts on the same subnet as the appliance, you can’t access it from hosts on other subnets. This condition should resolve itself within 15 minutes (when any ARP cache entries for the original appliance time out), or you can clear the condition immediately by purging any old ARP cache entries manually from the local router or gateway.

Steps

1. Display the current configurations of the appliance and record them.

a. Log in to the appliance to be replaced:

   i. Enter the following command: `ssh admin@grid_node_IP`

   ii. Enter the password listed in the `Passwords.txt` file.

   iii. Enter the following command to switch to root: `su -`

   iv. Enter the password listed in the `Passwords.txt` file.

   When you are logged in as root, the prompt changes from `$` to `#`.

b. Enter: **`run-host-command ipmitool lan print`** to display the current BMC configurations for the appliance.

2. Shut down the appliance: `shutdown -h now`

3. If any of the network interfaces on this StorageGRID appliance are configured for DHCP, you might need to update the permanent DHCP lease assignments on the DHCP servers to reference the MAC addresses of the replacement appliance. The update ensures the appliance is assigned the expected IP addresses. See Update MAC address references.

4. Remove and replace the appliance:

   a. Label the cables and then disconnect the cables and any network transceivers.

   To prevent degraded performance, don't twist, fold, pinch, or step on the cables.

   b. Remove the failed appliance from the cabinet or rack.

   c. Transfer the two power supplies, eight cooling fans, and two SSDs from the failed appliance to the replacement appliance.

   The two SSDs are placed in the slots as shown in the following diagram:

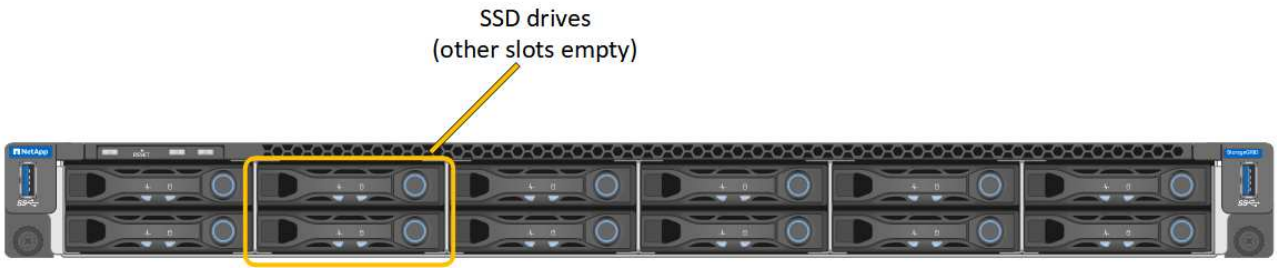

   HDD02 or HDD2 refer to the drive in the upper slot, and HDD03 or HDD3 refer to the drive in the lower slot.

   Follow the instructions provided for replacing these components.

   d. Install the replacement appliance into the cabinet or rack.

   e. Replace the cables and any optical transceivers.

   f. Power on the appliance and wait for it to rejoin the grid.

   g. Confirm that the appliance node appears in the Grid Manager and that no alerts appear.

5. Log in to the replaced appliance:

   a. Enter the following command: `ssh admin@grid_node_IP`

b. Enter the password listed in the `Passwords.txt` file.

c. Enter the following command to switch to root: `su -`

d. Enter the password listed in the `Passwords.txt` file.

6. Restore BMC network connectivity for the replaced appliance. There are two options:
   - Use static IP, netmask, and gateway
   - Use DHCP to obtain an IP, netmask, and gateway

   a. To restore the BMC configuration to use a static IP, netmask, and gateway, enter the following commands:

   ```
   run-host-command ipmitool lan set 1 ipsrc static

   run-host-command ipmitool lan set 1 ipaddr Appliance_IP

   run-host-command ipmitool lan set 1 netmask Netmask_IP

   run-host-command ipmitool lan set 1 defgw ipaddr Default_gateway
   ```

   b. To restore the BMC configuration to use DHCP to obtain an IP, netmask, and gateway, enter the following command:

   ```
   run-host-command ipmitool lan set 1 ipsrc dhcp
   ```

7. After restoring BMC network connectivity, connect to the BMC interface to audit and restore any additional custom BMC configuration you might have applied. For example, you should confirm the settings for SNMP trap destinations and email notifications. See Configure BMC interface.
8. Confirm that the appliance node appears in the Grid Manager and that no alerts appear.

After replacing the part, return the failed part to NetApp, as described in the RMA instructions shipped with the kit. See the Part Return & Replacements page for further information.

Related information

View status indicators

View boot-up codes for appliance