

StorageGRID® 11.2

Recovery and Maintenance Guide

August 2019 | 215-13582_2019-08_en-us
doccomments@netapp.com

NetApp®

Contents

Introduction to StorageGRID recovery and maintenance 6
Web browser requirements 6
Hotfix procedure 7
What happens when you apply a hotfix 7
Hotfix planning and preparation 7
How your system is affected when you apply a hotfix 8
Obtaining the required materials for a hotfix 8
Downloading the hotfix file 9
Checking the system's condition before applying a hotfix 9
Applying the hotfix 10
Recovery procedures 15
Reviewing warnings and preconditions for node recovery 15
Gathering required materials for grid node recovery 16
Downloading and extracting the StorageGRID installation files 17
Selecting a recovery procedure 20
Recovering from Storage Node failures 20
Recovering a Storage Node that has been down more than 15 days 21
Recovering a StorageGRID appliance Storage Node 22
Recovering from storage volume failure where the system drive is intact 38
Recovering from system drive failure and possible storage volume failure 47
Recovering from Admin Node failures 59
Recovering from primary Admin Node failures 60
Recovering from non-primary Admin Node failures 68
Recovering from API Gateway Node failures 75
Replacing the API Gateway Node 76
Selecting Start Recovery to configure the API Gateway Node 77
Recovering from Archive Node failures 78
Replacing the Archive Node 79
Selecting Start Recovery to configure the Archive Node 79
Resetting Archive Node connection to the cloud 81
All grid node types: Replacing a VMware node 81
Removing the failed grid node in VMware vSphere 81
Deploying the recovery grid node in VMware vSphere 82
All grid node types: Replacing a Linux node 85
Deploying new Linux hosts 86
Restoring grid nodes to the host 86
Performing additional recovery steps, if required 91
Decommission procedure 93
Preparing for decommissioning 94
Reviewing the considerations for decommissioning 95

Gathering required materials for decommissioning 99
Checking data repair jobs 100
Performing the decommission 101
Navigating to the Decommission page 102
Decommissioning disconnected grid nodes 104
Decommissioning connected grid nodes 108
Pausing and resuming the decommission process for Storage Nodes 111
Troubleshooting decommissioning 112
Network maintenance procedures 113
Updating subnets for the Grid Network 113
Configuring IP addresses 114
Modifying a node's network configuration 115
Adding to or changing subnet lists on the Admin Network 119
Adding to or changing subnet lists on the Grid Network 122
Linux: Adding interfaces to an existing node 126
Changing IP addresses for all nodes on the Grid Network 127
Configuring DNS servers 130
Modifying the DNS configuration for a single grid node 131
Configuring NTP servers 132
Restoring network connectivity for isolated nodes 133
Host-level and middleware procedures 136
Linux: Migrating a grid node to a new host 136
Linux: Exporting the node from the source host 136
Linux: Importing the node on the target host 137
Linux: Starting the migrated node 138
Archive Node maintenance for TSM middleware 138
Fault with archival storage devices 138
Tivoli Storage Manager administrative tools 139
Object permanently unavailable 139
VMware: Configuring a virtual machine for automatic restart 142
Grid node procedures 144
Viewing Server Manager status and version 144
Viewing current status of all services 145
Starting Server Manager and all services 146
Restarting Server Manager and all services 147
Stopping Server Manager and all services 147
Viewing current status of a service 148
Stopping a service 149
Forcing a service to terminate 149
Starting or restarting a service 150
Rebooting a grid node 151
Rebooting a grid node from the Grid Manager 151
Rebooting a grid node from the command shell 152
Shutting down a grid node 153
Powering down a host 154

Powering off and on all nodes in the grid 155
Stopping services and shutting down grid nodes 155
Starting up the grid nodes 156
Using a DoNotStart file 158
Adding a DoNotStart file for a service 158
Removing a DoNotStart file for a service 159
Troubleshooting Server Manager 159
Accessing the Server Manager log file 159
Service with an error state 159
Copyright 161
Trademark 162
How to send comments about documentation and receive update notifications 163

Introduction to StorageGRID recovery and maintenance

The recovery and maintenance procedures for StorageGRID include applying a software hotfix, recovering grid nodes, decommissioning grid nodes, performing network maintenance, performing host-level and middleware maintenance procedures, and performing Server Manager procedures.

All recovery and maintenance activities require a broad understanding of the StorageGRID system. You should review your StorageGRID system's topology to ensure that you understand the grid configuration.

You must follow all instructions exactly and heed all warnings.

Maintenance procedures not described are not supported or require a services engagement.

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

Note: "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.

Related information

Grid primer
Administering StorageGRID
SG6000 appliance installation and maintenance
SG5700 appliance installation and maintenance
SG5600 appliance installation and maintenance
NetApp Interoperability Matrix Tool

Web browser requirements

You must use a supported web browser.

Web browser	Minimum supported version
Google Chrome	70
Microsoft Internet Explorer	11 (Native Mode)
Mozilla Firefox	63

You should set the browser window to a recommended width.

Browser width	Pixels
Minimum	1024
Optimum	1280

Hotfix procedure

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 not included in a feature or service pack release. The same changes will be included in a future StorageGRID release.

What happens when you apply a hotfix

Applying a hotfix is similar to upgrading the software. You must download a file and install it on all grid nodes. Similar to a software upgrade, the hotfix is applied first to the primary Admin Node. Then, the hotfix is applied to all other grid nodes in your StorageGRID system. Unlike a software upgrade, you do not need to start the process on individual grid nodes.

While all 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.

If a failure occurs when the hotfix is applied to a specific node, the system continues to apply the hotfix to other nodes. If a failure occurs on a specific node, you must resolve the issue and retry the process. You can safely retry the hotfix process as many times as required until all nodes have been updated.

Hotfixes take significantly less time to install than full software upgrades. In addition, when StorageGRID applies a hotfix, it does not update grid provisioning, so you do not need to download a new Recovery Package.

How hotfixes are applied for recovery and expansion

Once a hotfix has been applied to your grid, the primary Admin Node automatically installs the same hotfix version to any nodes affected by recovery operations or added in an expansion. The exception to this is the primary Admin Node itself. If you need to recover the primary Admin Node, you must ensure that the recovered node is running the same software version as all other grid nodes. If a hotfix was applied previously, you must reapply the same version of the hotfix to the recovered primary Admin Node.

Related tasks

Configuring the replacement primary Admin Node on page 62

Hotfix planning and preparation

You must plan before applying a hotfix to ensure minimal disruption to your StorageGRID system.

Steps

1. How your system is affected when you apply a hotfix on page 8
2. Obtaining the required materials for a hotfix on page 8
3. Downloading the hotfix file on page 9
4. Checking the system's condition before applying a hotfix on page 9

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 API 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.

Alarms might be triggered

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

Many emails are generated

When you apply a hotfix to grid nodes, email notifications are generated when nodes are stopped and restarted. To avoid excessive emails, you can disable email notifications before the hotfix is applied to the first node and re-enable notifications after the hotfix has been applied to all nodes.

Configuration changes are restricted

When applying a hotfix to StorageGRID:

- Do not make any grid configuration changes until the hotfix has been applied to all nodes.
- Do not update the ILM configuration until the hotfix has been applied to all nodes.

Obtaining the required materials for a hotfix

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

Item	Notes
StorageGRID hotfix file	You must download the StorageGRID hotfix file to a service laptop.
Service laptop	The service laptop must have: • Network port • Supported web browser • SSH client (for example, PuTTY)
`Passwords.txt` file	Optional and used only if you are applying a hotfix manually using the SSH client. The `Passwords.txt` file is included in the SAID package, which 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.

Item	Notes
Related documentation	• Readme 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. • Instructions for administering StorageGRID

Related tasks

Downloading the hotfix file on page 9

Related information

Administering StorageGRID

Downloading the hotfix file

You must download the hotfix file to a service laptop before you can apply the hotfix.

Steps

1. Go to the Software Download page on the NetApp Support Site.

   NetApp Downloads: Software

2. Sign in using the username and password for your NetApp account.
3. Scroll to **StorageGRID (formerly StorageGRID Webscale)**, select **Hotfix**, and click **Go**.
4. Click **View & Download**.
5. Locate the hotfix version you need to install.

   Note: 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.

6. Review the readme file for the hotfix, and make sure you understand any special requirements.
7. Click **Download**, and accept the End User License Agreement.
8. Right click `hotfix-install-`*`version`*, and select **Save Link As**. (Click **Save Target As** if you are using Windows Explorer.)

   Note: Do not change the name of this file.

   Note: If you are using a Mac, the file might be automatically saved as a `.txt` file. If it is, you must rename the file without the `.txt` extension.

9. Select a location on the service laptop, and click **Save**.

Related tasks

Configuring the replacement primary Admin Node on page 62

Checking the system's condition before applying a hotfix

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

Steps

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

2. If possible, ensure that the system is running normally and that all grid nodes are operational.
3. Check for and resolve any active alarms, if possible.

   For information on specific alarms, see the instructions for troubleshooting StorageGRID.
4. Ensure there are no active maintenance procedures in progress, such as upgrade, recovery, expansion, or decommission.

   If another procedure is in progress, you must wait for it to complete. You cannot start the hotfix procedure if another maintenance procedure is in progress.

Related information

Troubleshooting StorageGRID

Applying the hotfix

The hotfix is applied first to the primary Admin Node, and then it is automatically applied to other grid nodes until all nodes are running the same software version.

Before you begin

- You have reviewed all of the considerations and completed all of the steps in "Hotfix planning and preparation."
- You must have the provisioning passphrase.
- You must have specific access permissions. For details, see the instructions for administering StorageGRID.

Steps

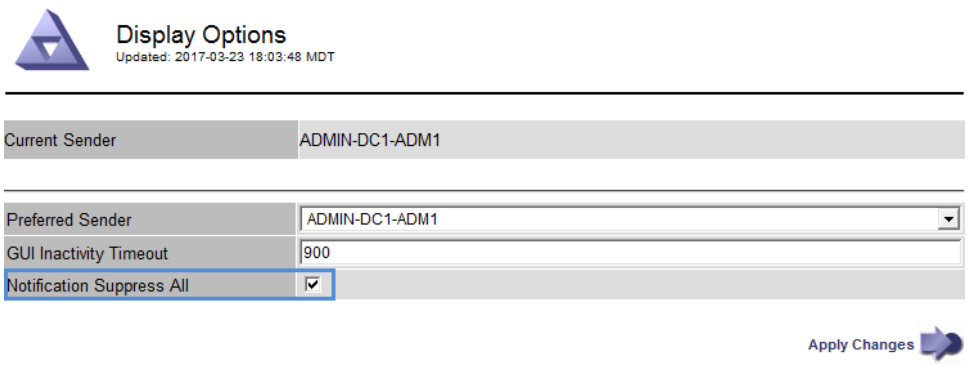
1. Sign in to the Grid Manager using a supported browser.
2. Optionally, disable email notifications.

   You can disable email notifications before applying the hotfix to avoid receiving excessive email notifications about node outages and hotfix processes.

   a. Select **Configuration** > **Display Options**.

   b. Select the **Notification Suppress All** check box

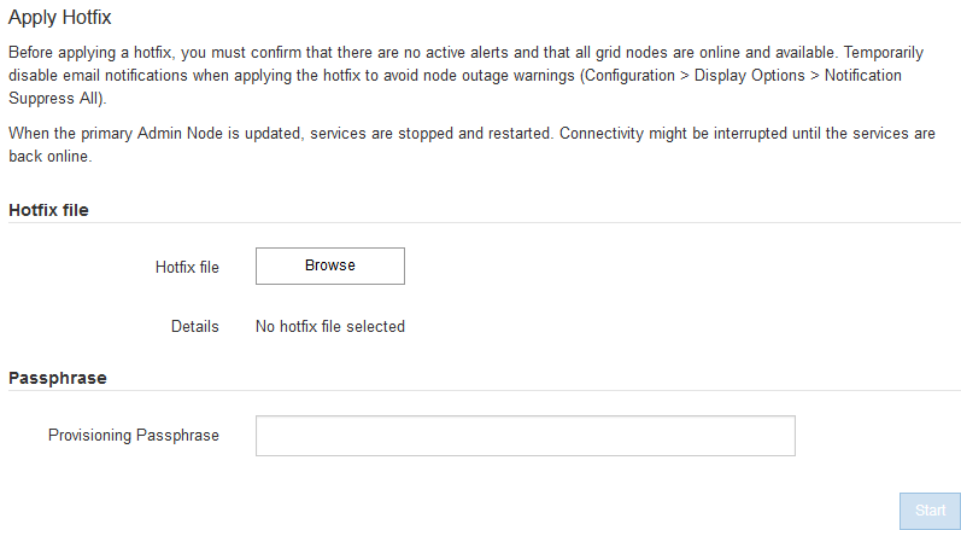
      When this check box is selected, all email notifications are suppressed, including those unrelated to the hotfix.



c. Click **Apply Changes**.

3. Select **Maintenance > Apply Hotfix**.

   The Apply Hotfix page appears.



4. Select the hotfix file you downloaded.

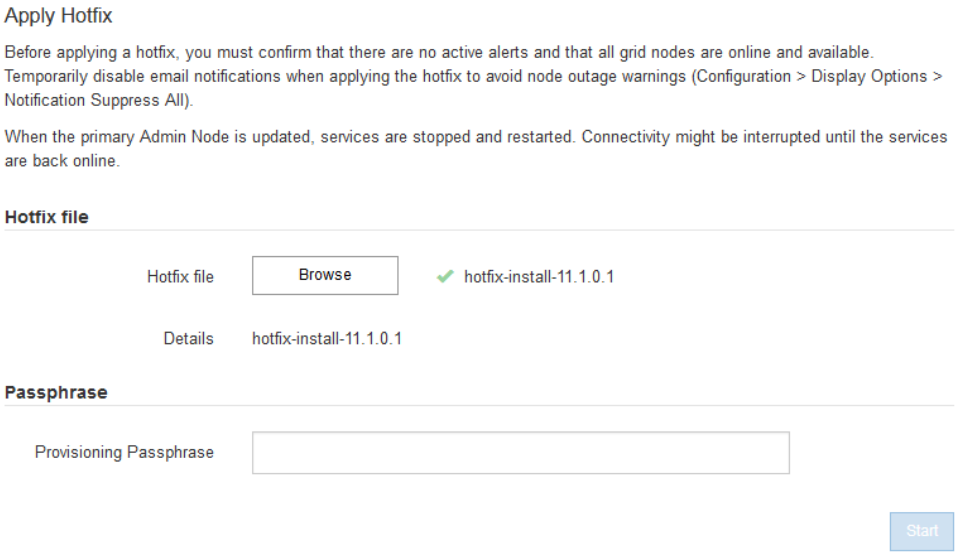
   a. Click **Browse**.

   b. Locate and select the file.

      `hotfix-install-`*`version`*

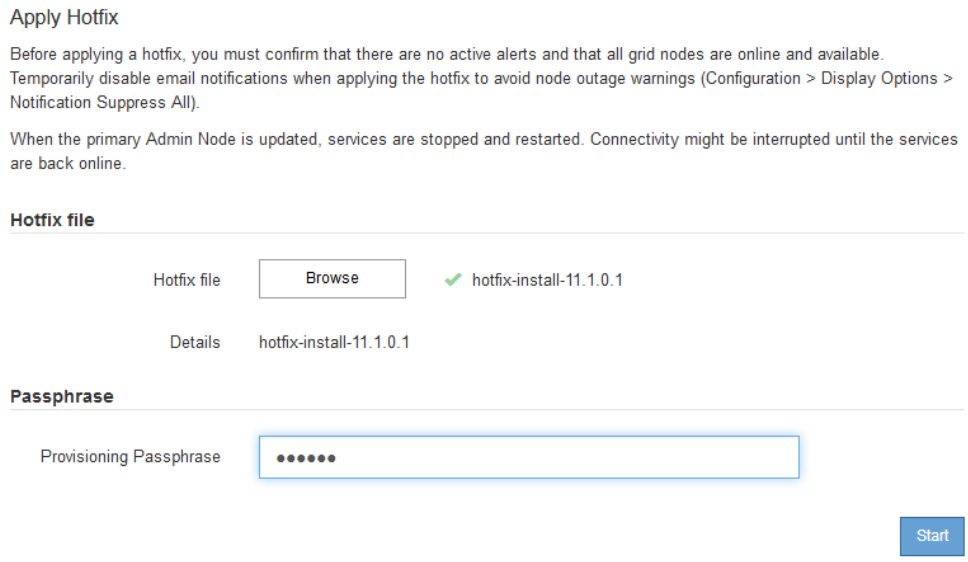
   c. Click **Open**.

      The file is uploaded and validated. When the validation process is done, the file name is shown in the Details field.



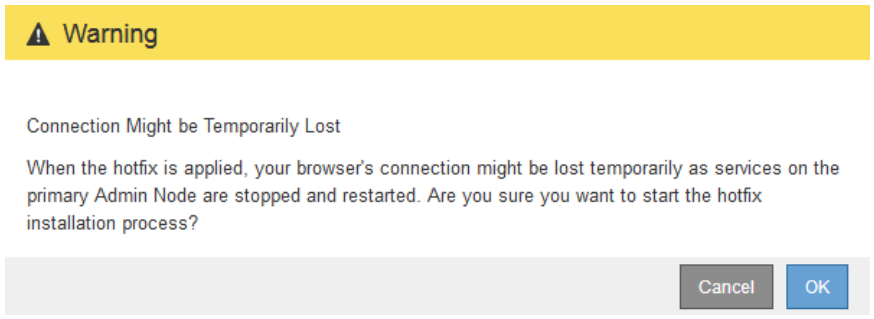
5. Enter the provisioning passphrase in the text box.

The **Start** button becomes enabled.



6. Click **Start**.

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



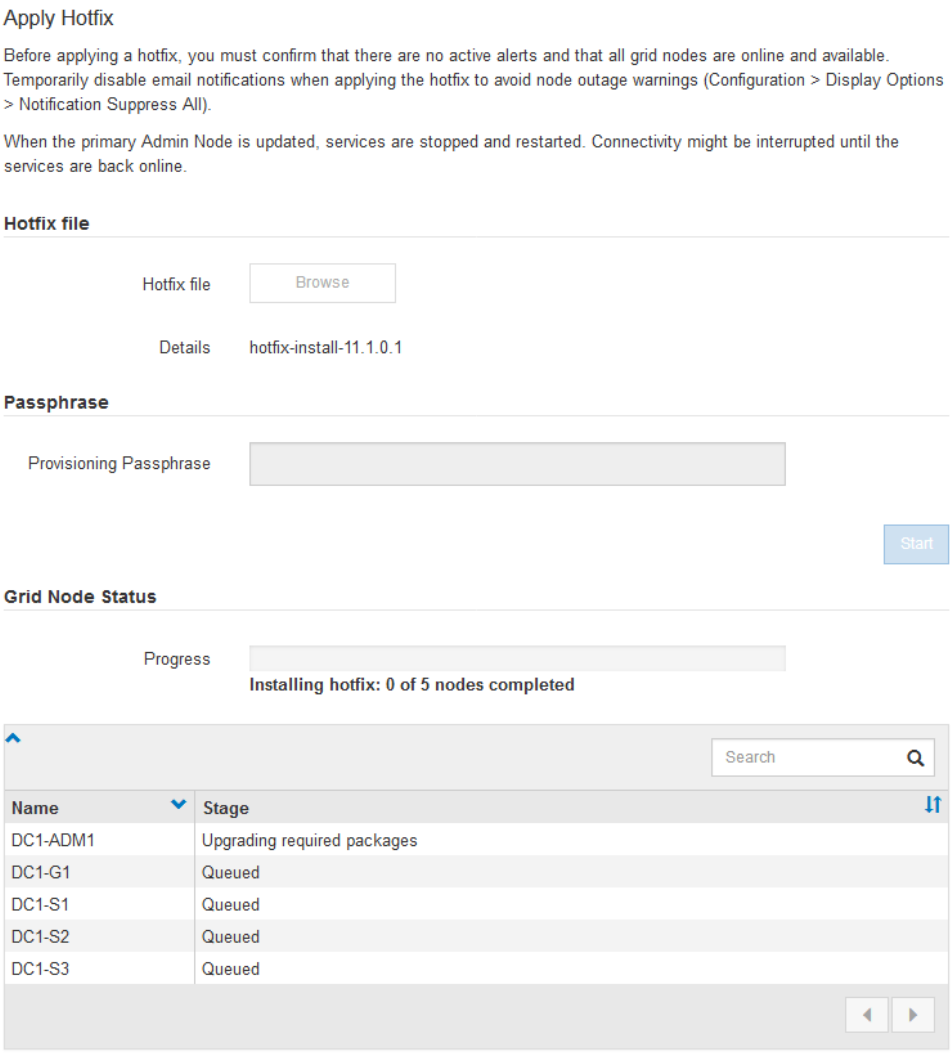
7. Click **OK** to start applying the hotfix.

   When the hotfix starts:

   a. The pre-hotfix validations are run.

      Note: If any errors are reported, resolve them and click **Start** again.

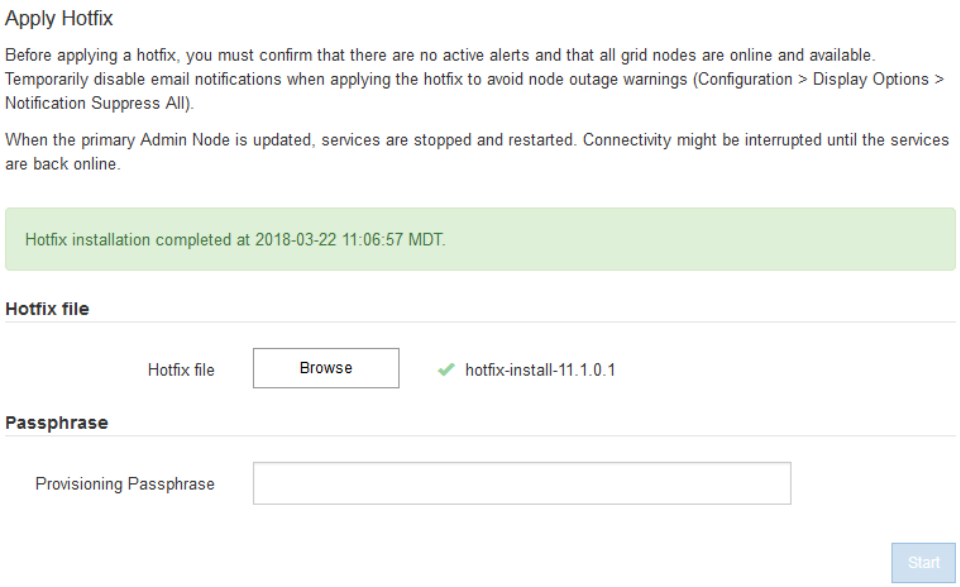
   b. The Grid Node Status table appears. This table shows all nodes in your grid and the current stage of the hotfix installation for each node.



8. Optionally, sort the list of nodes in ascending or descending order by **Name** or **Stage**. Or, enter a term in the **Search** box to search for specific nodes.
9. Wait while the hotfix is automatically applied to each grid node.

   The stages include Queued, Copying the hotfix file, Stopping node services, Upgrading required packages, Starting node services, and Complete.

   When the hotfix has been installed on all nodes, the Grid Nodes Status table closes and a green banner shows the date and time the hotfix was completed.



10. If a failure occurs while the hotfix is being applied to a specific node, resolve the issue, and repeat these steps.

    The procedure will not be complete until the hotfix has been successfully applied to all required nodes. You can safely retry the hotfix process as many time as required until it is complete.

11. Re-enable email notifications if you suppressed them.

    a. Select **Configuration > Display Options**.

    b. Unselect the **Notification Suppress All** check box.

    c. Click **Apply Changes**.

Related concepts

Hotfix planning and preparation on page 7

Related information

Administering StorageGRID

Recovery procedures

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

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 host, and then restore the StorageGRID node to the host. After grid node software is running, you restore the node's data and connections, so the node can resume operation.

Note: 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 should 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.

Reviewing warnings and preconditions for node recovery

Always recover a failed grid node as soon as possible. Review all warnings and preconditions for node recovery before beginning.

About this task

You must recover a failed grid node as soon as possible. A failed grid node might reduce the redundancy of data in the StorageGRID system, leaving you vulnerable to the risk of permanent data loss if another node fails. Operating with failed grid nodes can impact the efficiency of day-to-day operations, increase recovery time (when queues develop that need to be cleared before recovery is complete), and reduce your ability to monitor system operations.

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

- The failed physical or virtual hardware has been replaced and configured.
- 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.

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, API Gateway Nodes, Archive Nodes, appliance Storage Nodes, and virtual Storage Nodes.

Attention: Do not 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 original IP address or an unused IP address.

Gathering 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 the StorageGRID installation archive for your platform. See "Downloading and extracting the StorageGRID installation files" for instructions. **Note:** You do not 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 • Supported browser • SSH client (for example, PuTTY)
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. **Note:** If the primary Admin Node is operating normally, you can download the Recovery Package from the Grid Manager. Select **Maintenance** > **Recovery Package**.
`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	For the current supported versions of your platform, see the Interoperability Matrix Tool. *NetApp Interoperability Matrix Tool* Go to the platform vendor's website for documentation.

Related tasks

Downloading and extracting the StorageGRID installation files on page 17

Related references

Web browser requirements on page 6

Downloading and extracting the StorageGRID installation files

Before you can recover StorageGRID grid nodes, you must download the software and extract the files to your service laptop.

About this task

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 Grid Manager, go to **Help > About**.
2. Go to the Software Download page on the NetApp Support Site.

   NetApp Downloads: Software
3. Sign in using the username and password for your NetApp account.
4. Scroll to **StorageGRID (formerly StorageGRID Webscale)**, select **All Platforms**, and click **Go**.
5. Select the StorageGRID release, and click **View & Download**.
6. From the Software Download section of the page, click **CONTINUE**, and accept the End User License Agreement.
7. 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
VMware	• `StorageGRID-Webscale-`*`version`*`-VMware-`*`uniqueID`*`.tgz` • `StorageGRID-Webscale-`*`version`*`-VMware-`*`uniqueID`*`.zip`
Red Hat Enterprise Linux or CentOS	• `StorageGRID-Webscale-`*`version`*`-RPM-`*`uniqueID`*`.tgz` • `StorageGRID-Webscale-`*`version`*`-RPM-`*`uniqueID`*`.zip`
Ubuntu or Debian	• `StorageGRID-Webscale-`*`version`*`-DEB-`*`uniqueID`*`.tgz` • `StorageGRID-Webscale-`*`version`*`-DEB-`*`uniqueID`*`.zip`
OpenStack or other hypervisor	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.

8. Download the archive file.
9. Extract the archive file.
10. Choose the files you need, based on your platform and which grid nodes you need to recover.

    The paths listed in the tables are relative to the top-level directory installed by the archive file.

Table 1: VMware files

Filename	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 API Gateway Nodes.
`/vsphere/vsphere-storage.ovf` `/vsphere/vsphere-storage.mf`	The template file (`.ovf`) and manifest file (`.mf`) for deploying non-appliance Storage Nodes.
Deployment scripting tools	
`/vsphere/deploy-vsphere-ovftool.sh`	A Bash shell script used to automate the deployment of virtual grid nodes.
`/vsphere/deploy-vsphere-ovftool-sample.ini`	A sample configuration file for use with the `deploy-vsphere-ovftool.sh` script.
`/vsphere/configure-sga.py`	A Python script used to automate the configuration of StorageGRID appliances.
`/vsphere/configure-storagegrid.py`	A Python script used to automate the configuration of a StorageGRID system.
`/vsphere/configure-storagegrid.sample.json`	A sample 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.

Table 2: Red Hat Enterprise Linux or CentOS 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 hosts.
`/rpms/StorageGRID-Webscale-Service-`*`version-SHA`*`.rpm`	RPM package for installing the StorageGRID host service on your hosts.
Deployment scripting tools	
`/rpms/configure-storagegrid.py`	A Python script used to automate the configuration of a StorageGRID system.
`/rpms/configure-storagegrid.sample.json`	A sample configuration file for use with the `configure-storagegrid.py` script.
`/rpms/configure-storagegrid.blank.json`	A blank configuration file for use with the `configure-storagegrid.py` script.
`/rpms/configure-sga.py`	A Python script used to automate the configuration of StorageGRID appliances.
`/rpms/extras/ansible`	An Ansible role and example Ansible playbook for configuring the hosts for StorageGRID container deployment. You can customize the playbook or role as necessary.

Table 3: Ubuntu or Debian 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 free license that does not provide any support entitlement for the product..
`/debs/storagegrid-webscale-images-`*`version-SHA`*`.deb`	DEB package for installing the StorageGRID node images on your hosts.
`/debs/storagegrid-webscale-service-`*`version-SHA`*`.deb`	DEB package for installing the StorageGRID host service on your hosts.
Deployment scripting tools	
`/debs/configure-storagegrid.py`	A Python script used to automate the configuration of a StorageGRID system.
`/debs/configure-storagegrid.sample.json`	A sample 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/configure-sga.py`	A Python script used to automate the configuration of StorageGRID appliances.
`/debs/extras/ansible`	An Ansible role and example Ansible playbook for configuring the hosts for StorageGRID container deployment. You can customize the playbook or role as necessary.

Related information

VMware installation
Red Hat Enterprise Linux or CentOS installation
Ubuntu or Debian installation

Selecting a recovery procedure

You must select the correct recovery procedure for the type of node that has failed. For Storage Nodes, the procedures vary based on the type and duration of the failure. For Admin Nodes, the procedures vary based on whether you need to recover the primary Admin Node or a non-primary Admin Node.

About this task

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 should 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.

Grid node	Recovery procedure
Storage Node	Go to *Recovering from Storage Node failures* on page 20. Then, select the specific recovery procedure for your Storage Node failure.
Admin Node	Go to *Recovering from Admin Node failures* on page 59. Then, select the recovery procedure for a primary or non-primary Admin Node.
API Gateway Node	Go to *Recovering from API Gateway Node failures* on page 75.
Archive Node	Go to *Recovering from Archive Node failures* on page 78.

Recovering from Storage Node failures

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.	Contact technical support.	Recovering more than one Storage Node might affect the integrity of the Cassandra database. Technical support can determine when it is safe to begin recovery of a second Storage Node.
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.	
A Storage Node has been offline for more than 15 days.	*Recovering a Storage Node that has been down more than 15 days* on page 21	This procedure is required to ensure Cassandra database integrity.

Issue	Action	Notes
An appliance Storage Node has failed.	*Recovering a StorageGRID appliance Storage Node* on page 22	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	*Recovering from storage volume failure where the system drive is intact* on page 38	This procedure is used for virtual Storage Nodes.
The system drive has failed. Storage volumes might have failed.	*Recovering from system drive failure and possible storage volume failure* on page 47	The node replacement procedure varies depending on the deployment platform.

Recovering a Storage Node that has been 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 decommissioning. (In the Grid Manager, go to **Maintenance > Decommission**.)
- You have checked that an expansion is not in progress. (In the Grid Manager go to **Maintenance > 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 cannot 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.

Caution: Do not perform this procedure if more than one Storage Node is offline, or if more than one Storage Node has a failure of any kind. Doing so might result in data loss. Contact technical support.

Caution: Do not rebuild Cassandra on more than one Storage Node within a 15-day period. Rebuilding Cassandra on two or more Storage Nodes within 15 days of each other might result in data loss. Contact technical support.

Steps

1. If necessary, power on the Storage Node that needs to be recovered.
2. From the service laptop, 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 `#`.

Note: 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.

Recovering from system drive failure and possible storage volume failure on page 47

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** > **Grid Topology**. Then, select ***`site`*** > ***`Storage Node`*** > **SSM** > **Services**. Verify that the Cassandra service displays `Not Running`.

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

   d. 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 storage services are running, you will be prompted to stop them. Enter: **y**
   - Review the warnings in the script. If none of them apply, confirm that you want to rebuild Cassandra. Enter: **y**

5. After the rebuild completes, perform the following checks:

   a. In the Grid Manager, select **Support** > **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.”

Related tasks

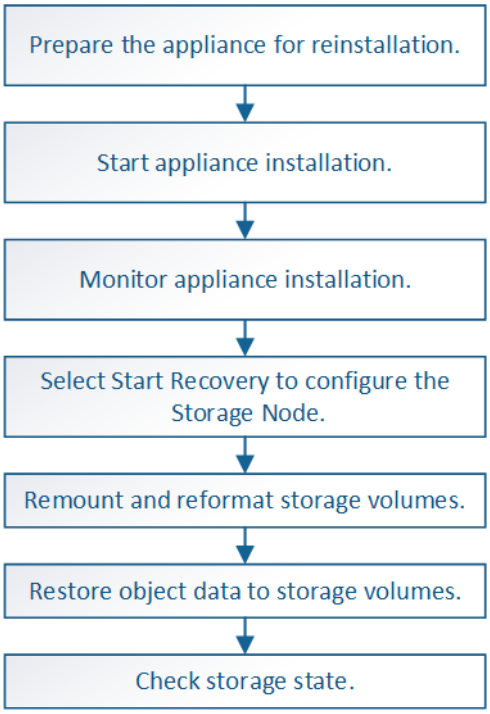
Recovering from system drive failure and possible storage volume failure on page 47

Recovering a StorageGRID appliance Storage Node

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.

About this task

You must prepare the appliance and reinstall software, configure the node to rejoin the grid, reformat storage, and restore object data.



Caution: Do not perform this procedure if more than one Storage Node is offline, or if more than one Storage Node has a failure of any kind. Doing so might result in data loss. Contact technical support.

Caution: Do not rebuild Cassandra on more than one Storage Node within a 15-day period. Rebuilding Cassandra on two or more Storage Nodes within 15 days of each other might result in data loss. Contact technical support.

Attention: If ILM rules are configured to store only one replicated copy and the copy exists on a Storage Node that has failed, you will not be able to recover the object. However, you must still perform the procedure to restore object data to a storage volume to purge lost object information from the database.

Note: If you encounter a DDS alarm during recovery, see troubleshooting instructions to recover from the alarm by rebuilding Cassandra. After Cassandra is rebuilt, alarms should clear. If alarms do not clear, contact technical support.

Note: For hardware maintenance procedures, such as instructions for replacing a controller, see the installation and maintenance instructions for your appliance.

Steps

1. Preparing an appliance Storage Node for reinstallation on page 24
2. Starting StorageGRID appliance installation on page 24
3. Monitoring StorageGRID appliance installation on page 26
4. Selecting Start Recovery to configure the appliance Storage Node on page 28
5. Remounting and reformatting appliance storage volumes ("Manual Steps") on page 30
6. Restoring object data to a storage volume for an appliance on page 34
7. Checking the storage state after recovering an appliance Storage Node on page 37

Related information

SG6000 appliance installation and maintenance
SG5700 appliance installation and maintenance
SG5600 appliance installation and maintenance

Preparing an appliance Storage Node for reinstallation

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

Steps

1. From the service laptop, 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.

   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.

Starting 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.

For instructions for completing these prerequisite tasks, see the installation and maintenance instructions for your StorageGRID appliance.

- You have a service laptop with 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 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:

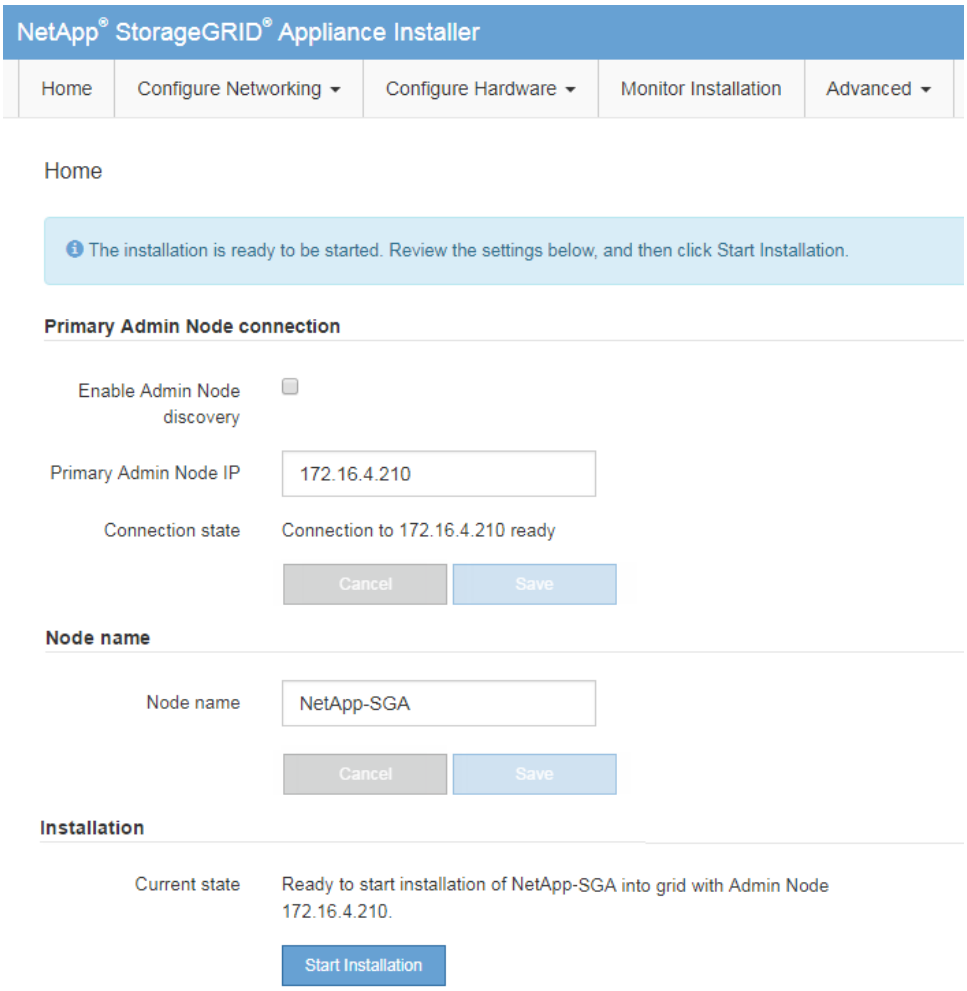
Option	Description
Manual IP entry	**a.** Unselect the **Enable Admin Node discovery** check box. **b.** Enter the IP address manually. **c.** Click **Save**. **d.** Wait while the connection state for the new IP address becomes “ready.”
Automatic discovery of all connected primary Admin Nodes	**a.** Select the **Enable Admin Node discovery** check box. **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 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 installation and 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.

Note: If you need to access the Monitor Installation page manually, click **Monitor Installation** from the menu bar.

Related information

SG6000 appliance installation and maintenance
SG5700 appliance installation and maintenance
SG5600 appliance installation and maintenance

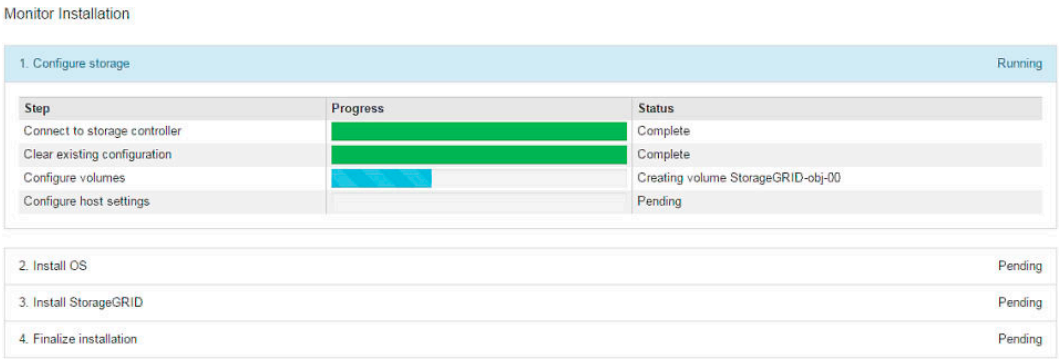
Monitoring StorageGRID appliance installation

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.

> **Note:** The installer ensures that tasks completed in a previous install are not re-run. If you are re-running an installation, any tasks that do not 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.

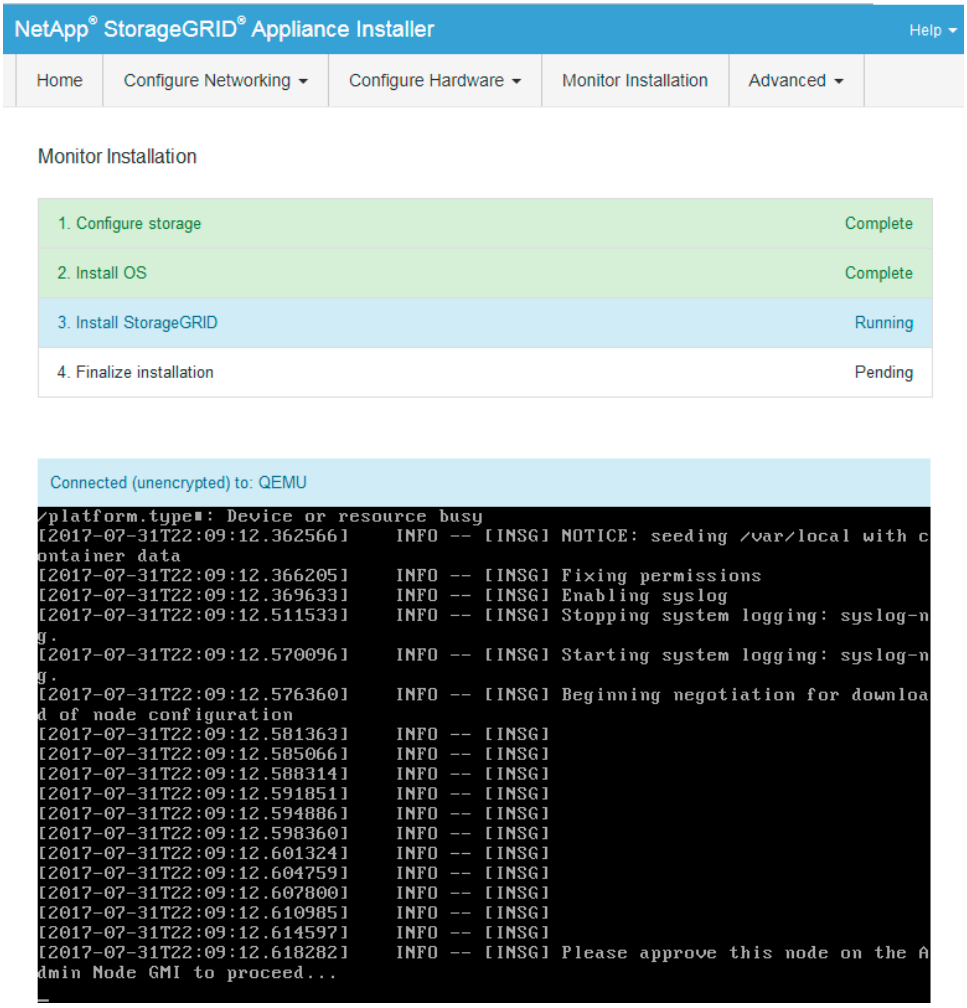
   1. Configure storage

   During this stage, the installer connects to the storage controller, clears any existing configuration, communicates with SANtricity software to configure volumes, 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.

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 the procedure to configure the appliance Storage Node.

Selecting Start Recovery to configure the appliance Storage Node

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

Before you begin

- You must be signed in to the Grid Manager using a supported browser.
- You must have a service laptop.
- You must have Maintenance or Root Access permissions.
- You must have the provisioning passphrase.
- You must have deployed a recovery appliance Storage Node.
- You must know the start date of any repair jobs for erasure-coded data.
- You must have verified that the Storage Node has not been rebuilt within the last 15 days.

Steps

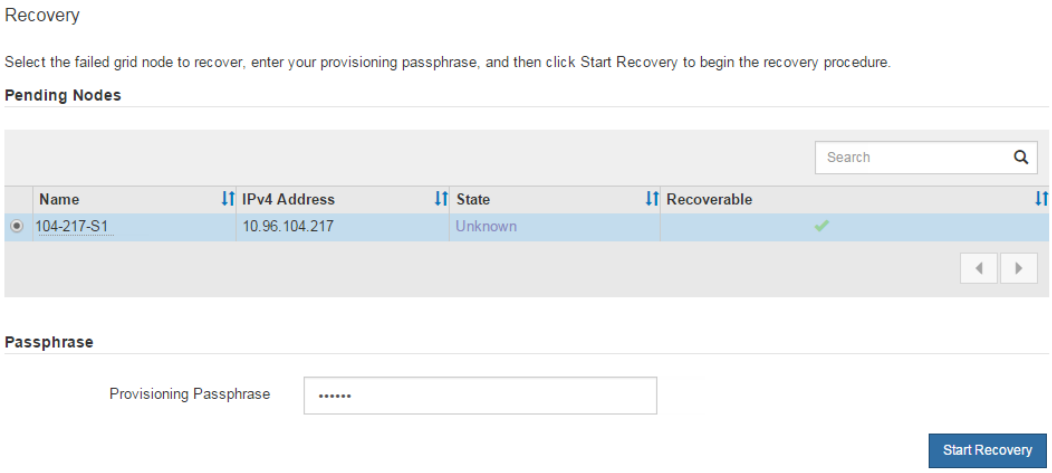
1. From the Grid Manager, select **Maintenance > Recovery**.

2. Select the grid node you want to recover in the **Pending Nodes** list.

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

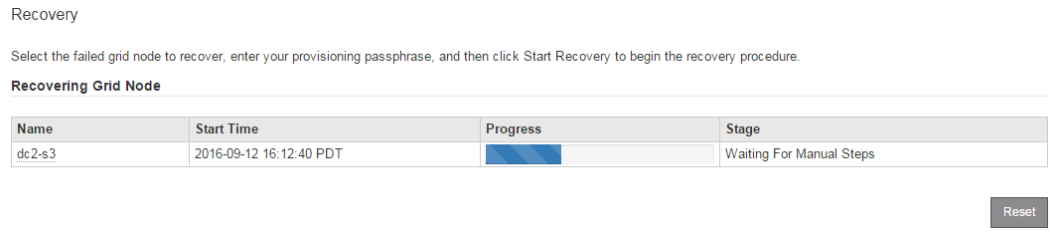
3. Enter the **Provisioning Passphrase**.

4. Click **Start Recovery**.

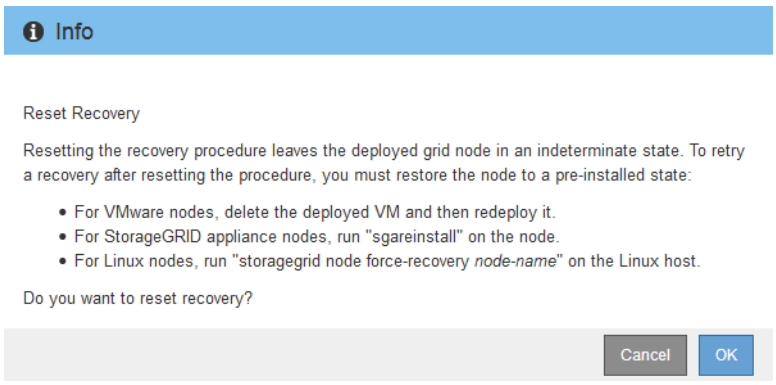


5. Monitor the progress of the recovery in the **Recovering Grid Node** table.

   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.



Note: At any point during the recovery, you can click **Reset** to start a new recovery. An Info 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.

Remounting and reformatting appliance storage volumes ("Manual Steps")

You must manually run two scripts to remount preserved storage volumes and reformat failed storage volumes. The first script remounts volumes that are properly formatted as StorageGRID storage volumes. The second script reformats unmounted volumes, and rebuilds the Cassandra database on the Storage Node as required.

Before you begin

- You have already replaced the hardware for any failed storage volumes that you know require replacement.
  This procedure might help you identify additional failed storage volumes.
- You have checked that a Storage Node decommissioning is not in progress, or you have paused decommissioning. (In the Grid Manager, go to **Maintenance > Decommission**.)
- You have checked that an expansion is not in progress. (In the Grid Manager go to **Maintenance > Expansion**.)

Caution: 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. Do not 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 first script, `sn-remount-volumes`, 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.
- Review the script output and resolve any issues. Then, rerun the script until you are satisfied that all valid storage volumes have been remounted.
- Run the second script, `sn-recovery-postinstall.sh`, to reformat any storage volumes that the first script could not mount or that were found to be improperly formatted; rebuild Cassandra, if needed; and start services.

  Note: All data on these storage volumes is lost, and must be recovered from other locations in the grid if it is possible.

Steps

1. From the service laptop, 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.

Note: 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`

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

Note: 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.

Example

```
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.

Do not continue to the next step if you believe that the data remaining on
this volume cannot 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.

Do not continue to the next step if you believe that the data remaining on
this volume cannot 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 uninitialized or the disk's superblock was corrupted. When the script cannot 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 do not 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.

   Attention: 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` 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 are not 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.

      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, replace or remove the disk and run the `sn-remount-volumes` script again to ensure the issue is resolved.

      As needed, you can find the node ID for the Storage Node you are recovering at the top of the script (the “configured LDR noid”). You can look up node IDs for other Storage Nodes in the Grid Manager. Select **Support** > **Grid Topology** > ***`Site`*** > ***`Storage Node`*** > **LDR** > **Overview**.

Caution: If you are unable to resolve the issue, contact technical support. If you run the `sn-recovery-postinstall.sh` script, the storage volume will be reformatted, which might cause data loss.

c. Review the messages for devices that could not be mounted, and make a note of the device name for each failed storage volume.

Note: 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 in the next procedure (restoring object data).

d. Run the `sn-remount-volumes` script again to ensure that all valid storage volumes have been remounted.

Attention: If a storage volume could not 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).

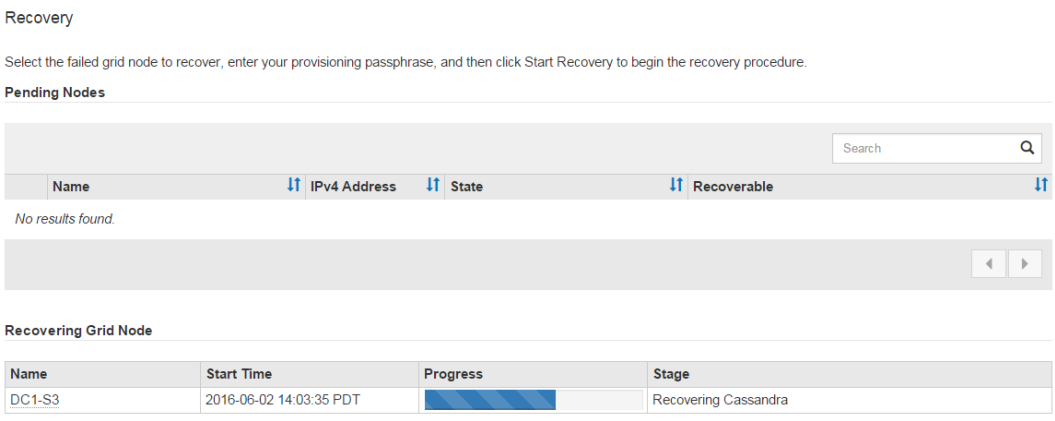
Caution: Do not run the `sn-recovery-postinstall` script if you believe that the data remaining on this volume cannot be rebuilt from elsewhere in the grid (for example, if your ILM policy uses a rule that makes only one copy or if volumes have failed on multiple nodes). Instead, contact technical support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

4. Run the second script to reformat all unmounted (failed) storage volumes, rebuild Cassandra if required, and start services on the Storage Node:

`sn-recovery-postinstall.sh`

5. As the script runs, monitor the Recovery page in the Grid Manager, and review the command line output, which provides more detailed status information.

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



6. Return to the Monitor Install page of the StorageGRID Appliance Installer by entering `http://Controller_IP:8080`, using the IP address of the compute controller.

The Monitor Install page shows the installation progress while the script is running.

After you finish

Restore object data to any storage volumes that were formatted by `sn-recovery-postinstall.sh`, as described in the next procedure. You can restore object data at the same time Cassandra is being rebuilt.

Related tasks

Reviewing warnings for Storage Node system drive recovery on page 48
Restoring object data to a storage volume for an appliance on page 34

Restoring object data to a storage volume for an appliance

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

Before you begin

- You must have confirmed that the recovered Storage Node displays in green in the Grid Topology tree and shows all services as Online.

About this task

Object data is restored from other Storage Nodes and Archive Nodes, assuming that the grid's ILM rules were configured such that object copies are available.

Attention: If ILM rules are configured to store only one replicated copy and that copy exists on a Storage Node that has failed, you will not be able to recover the object. However, you must still perform the procedure to restore object data to a storage volume to purge lost object information from the database.

Note: If the only remaining copy of object data is located on an Archive Node, object data is retrieved from the Archive Node. Due to the latency associated with retrievals from external archival storage systems, restoring object data to a Storage Node from an Archive Node takes longer than restoring copies from other Storage Nodes.

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. You use different options with the `repair-data` script, based on whether you are restoring replicated data or erasure coded data, as follows:

- **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
```

- **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
```

  Repairs of erasure coded data can begin while some Storage Nodes are offline. Repair will complete after all nodes are available. You can track repairs of erasure coded data with this command:

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

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

Steps

1. From the service laptop, 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 host name 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`
3. If all storage volumes have failed, repair the entire node. (If only some volumes have failed, go to the next step.)

   Attention: You cannot run `repair-data` operations for more than one node at the same time. To recover multiple nodes, contact technical support.

   - If your grid includes replicated data, use the `repair-data start-replicated-node-repair` command with the `--nodes` option to repair the entire Storage Node.
     This command repairs the replicated data on Storage Node named SG-DC-SN3:

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

     Note: As object data is restored, the LOST (Lost Objects) alarm is triggered if the StorageGRID system cannot locate replicated object data. Alarms might be triggered on Storage Nodes throughout the system. You should determine the cause of the loss and if recovery is possible. See the instructions for troubleshooting StorageGRID.
   - If your grid contains erasure coded data, use the `repair-data start-ec-node-repair` command with the `--nodes` option 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.

     Note: Repairs of erasure coded data can begin while some Storage Nodes are offline. Repair will complete after all nodes are available.
   - If your grid has both replicated and erasure coded data, run both commands.
4. If only some of the volumes have failed, repair the affected volumes.

   Enter the volume IDs in hexadecimal, where 0000 is the first volume and 000F is the sixteenth volume. You can specify one volume or a range of volumes.

   - If your grid contains replicated data, use the `start-replicated-volume-repair` command with the `--nodes` and `--volume-range` options.
     This command restores replicated data to all volumes in the range 0003 to 000B on a Storage Node named SG-DC-SN3:

     ```
     repair-data start-replicated-volume-repair --nodes SG-DC-SN3 --
     volume-range 0003,000B
     ```

For replicated data, you can run more than one `repair-data` operation at the same time for the same node. You might want to do this if you need to restore two volumes that are not in a range, such as 0000 and 000A.

Note: As object data is restored, the LOST (Lost Objects) alarm is triggered if the StorageGRID system cannot locate replicated object data. Alarms might be triggered on Storage Nodes throughout the system. You should determine the cause of the loss and if recovery is possible. See the instructions for troubleshooting StorageGRID.

- If your grid contains erasure coded data, use the `start-ec-volume-repair` command with the `--nodes` and `--volume-range` options.

  This command restores erasure coded data to a single volume 000A on a Storage Node named SG-DC-SN3:

  ```
  repair-data start-ec-volume-repair --nodes SG-DC-SN3 --volume-
  range 000A
  ```

  For erasure coded data, you must wait until one `repair-data start ec-volume-repair` operation completes before starting a second `repair-data` operation for the same node.

  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.

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

- If your grid has both replicated and erasure coded data, run both commands.

5. Monitor the repair of replicated data.

   a. Select **Support** > **Grid Topology**.

   b. Select ***`deployment`*** > **Overview** > **ILM Activity**.

   c. Use a combination of the following attributes to determine, as well as possible, if replicated repairs are complete.

      Note: Cassandra inconsistencies might be present, and failed repairs are not 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.

        Note: High-risk objects are objects that are at risk of being completely lost. This does not include objects that do not 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 is at the deployment level and is the maximum of all node scan periods. You can query the **Scan Period – Estimated** attribute history at the deployment level to determine an appropriate time frame for your grid.

6. Monitor the repair of erasure coded data, and retry any requests that might have failed.

   a. Determine the status of erasure coded data repairs:

- 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.

  ```
  root@DC1-ADM1:~ # repair-data show-ec-repair-status

   Repair ID   Scope                  Start Time  End Time  State  Est Bytes Affected Bytes Repaired  Retry Repair
  ================================================================================================================
   949283   DC1-S-99-10(Volumes: 1,2) 2016-11-30T15:27:06.9  Success   17359           17359           No
   949292   DC1-S-99-10(Volumes: 1,2) 2016-11-30T15:37:06.9  Failure   17359           0               Yes
   949294   DC1-S-99-10(Volumes: 1,2) 2016-11-30T15:47:06.9  Failure   17359           0               Yes
   949299   DC1-S-99-10(Volumes: 1,2) 2016-11-30T15:57:06.9  Failure   17359           0               Yes
  ```

b. 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 83930030303133434:

   ```
   repair-data start-ec-node-repair --repair-id 83930030303133434
   ```

   This command retries a failed volume repair, using the repair ID 83930030303133434:

   ```
   repair-data start-ec-volume-repair --repair-id 83930030303133434
   ```

Checking the storage state after recovering an 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 must be signed in to the Grid Manager using a supported browser.
- The Storage Node has been recovered, and data recovery is complete.

Steps

1. Select **Support** > **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**.

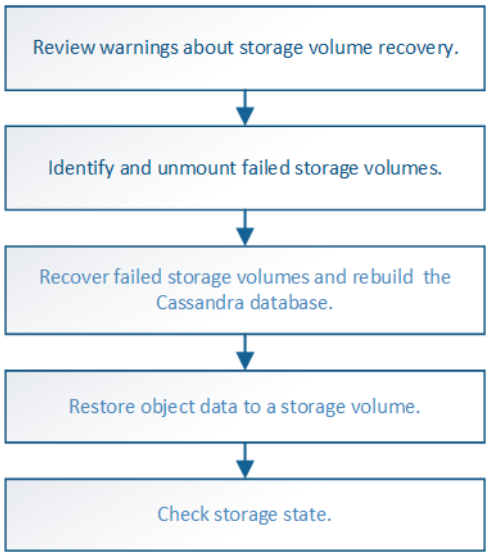
   d. Click the **Overview** tab and confirm that the values of **Storage State – Desired** and **Storage State – Current** are updated to Online.

Recovering from storage volume failure where the system drive is intact

You must complete a series of tasks to recover a virtual 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.

About this task

This recovery procedure applies to virtual Storage Nodes only. If storage volumes have failed on an appliance Storage Node, use the procedure for "Recovering a StorageGRID appliance Storage Node."



Steps

1. Reviewing warnings about storage volume recovery on page 38
2. Identifying and unmounting failed storage volumes on page 39
3. Recovering failed storage volumes and rebuilding the Cassandra database on page 41
4. Restoring object data to a storage volume where the system drive is intact on page 43
5. Checking the storage state after recovering storage volumes on page 47

Related tasks

Recovering a StorageGRID appliance Storage Node on page 22

Reviewing warnings about storage volume recovery

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

About this task

The storage volumes (or rangedbs) in a Storage Node are identified by a hexadecimal number from 0000 to 000F, which is known as the volume ID. The first object store (volume 0) on each Storage Node uses 3 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.

Caution: Do not perform this procedure if more than one Storage Node is offline, or if more than one Storage Node has a failure of any kind. Doing so might result in data loss. Contact technical support.

Caution: When you recover a failed storage volume, the script prompts you to rebuild the Cassandra database if rebuilding Cassandra is necessary. Do not rebuild Cassandra on more than one Storage Node within a 15-day period. Rebuilding Cassandra on two or more nodes within 15 days of each other may result in data loss. Contact technical support.

Attention: If ILM rules are configured to store only one replicated copy and the copy exists on a Storage Node that has failed, you will not be able to recover the object. However, you must still perform the procedure to restore object data to a storage volume to purge lost object information from the database.

Note: If you encounter a DDS alarm during recovery, see troubleshooting instructions to recover from the alarm by rebuilding Cassandra. After Cassandra is rebuilt, alarms should clear. If alarms do not clear, contact technical support.

Related tasks

Reviewing warnings and preconditions for node recovery on page 15

Identifying and unmounting 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 must be signed in to the Grid Manager using a supported 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. While you should run file system checks for consistency on a normal basis, only perform this procedure for detecting failed volumes on a large file system when necessary, such as in cases of power loss.

Note: 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.

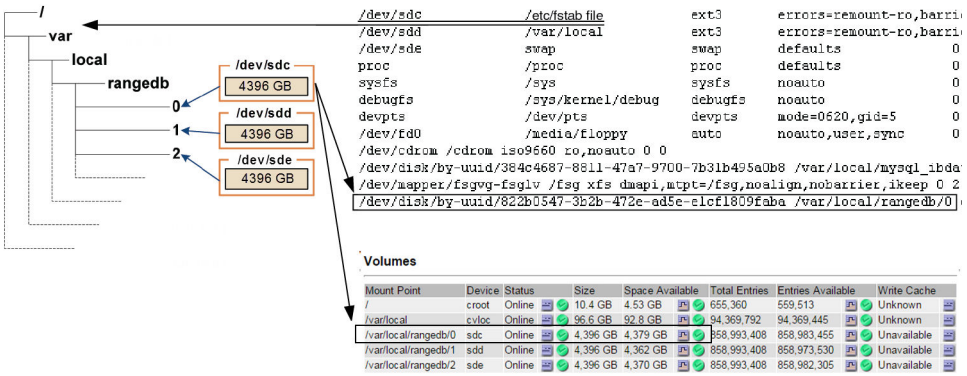
Note: Repair the hardware and properly attach the disks before running the `reboot` command.

Caution: 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 cannot 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:

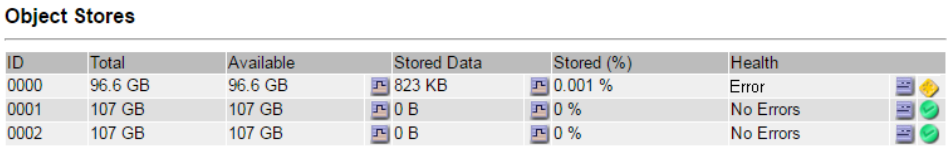


Steps

1. Complete the following steps to record the failed storage volumes and their device names:

   a. Select **Support** > **Grid Topology**.

   b. Select ***site*** > ***failed Storage Node*** > **LDR** > **Storage** > **Overview** > **Main**, and look for object stores with alarms.



   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, from 0000 to 000F. 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

If you cannot determine the volume number and device name of failed storage volumes, log in to an equivalent Storage Node and determine the mapping of volumes to device names on that server.

Storage Nodes are usually added in pairs, with identical hardware and storage configurations. Examine the `/etc/fstab` file on the equivalent Storage Node to identify the device names that correspond to each storage volume. Identify and record the device name for each failed storage volume.

2. From the service laptop, 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. Stop the LDR service and unmount the failed storage volumes:

   a. Stop LDR services:

   ```
   service ldr stop
   ```

   b. If rangedb/0 needs recovery, stop Cassandra before unmounting rangedb/0:

   ```
   service cassandra stop
   ```

   c. Unmount the failed storage volume:

   ```
   umount /var/local/rangedb/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.

Recovering failed storage volumes and rebuilding the 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 must have the `Passwords.txt` file.
- The system drives on the server must be intact.
- The cause of the failure must have been identified and, if necessary, replacement storage hardware must already have been acquired.
- The total size of the replacement storage must be the same as the original.
- You have checked that a Storage Node decommissioning is not in progress, or you have paused decommissioning. (In the Grid Manager, go to **Maintenance** > **Decommission**.)

- You have checked that an expansion is not in progress. (In the Grid Manager go to **Maintenance > Expansion**.)
- You have reviewed the warnings about storage volume recovery.
  Reviewing warnings about storage volume recovery on page 38

Steps

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

   After you replace the storage, make sure you rescan or reboot to make sure that it is recognized by the operating system, but do not remount the volumes. The storage is remounted and added to `/etc/fstab` in a later step.

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

   Note: 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.

4. Reformat any failed storage volumes and rebuild the Cassandra database if it is necessary. Enter:

   `reformat_storage_block_devices.rb`

   - If storage services are running, you will be prompted to stop them. Enter: **y**
   - 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 to `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 do not want to reformat it.

       Note: 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.

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

```
root@DC1-S1:~ # reformat_storage_block_devices.rb
Storage services must be stopped before running this script.
Stop storage services [y/N]? y
Shutting down storage services.
Storage services stopped.
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 c817f87f-f989-4a21-8f03-b6f42180063f
Skipping in use device /dev/sdg
All devices processed
Running: /usr/local/ldr/setup_rangedb.sh 12075630
Cassandra does not need rebuilding.
Starting services.

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

Related tasks

Reviewing warnings about storage volume recovery on page 38

Restoring object data to a storage volume where the system drive is intact

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

Before you begin

- You must have confirmed that the recovered Storage Node displays in green in the Grid Topology tree and shows all services as Online.

About this task

Object data is restored from other Storage Nodes and Archive Nodes, assuming that the grid's ILM rules were configured such that object copies are available.

Attention: If ILM rules are configured to store only one replicated copy and that copy exists on a Storage Node that has failed, you will not be able to recover the object. However, you must still perform the procedure to restore object data to a storage volume to purge lost object information from the database.

Note: If the only remaining copy of object data is located on an Archive Node, object data is retrieved from the Archive Node. Due to the latency associated with retrievals from external archival storage systems, restoring object data to a Storage Node from an Archive Node takes longer than restoring copies from other Storage Nodes.

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. You use different options with the `repair-data` script, based on whether you are restoring replicated data or erasure coded data, as follows:

- **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
```

- **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
```

  Repairs of erasure coded data can begin while some Storage Nodes are offline. Repair will complete after all nodes are available. You can track repairs of erasure coded data with this command:

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

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

Steps

1. From the service laptop, 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 host name 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`
3. If all storage volumes have failed, repair the entire node. (If only some volumes have failed, go to the next step.)

   Attention: You cannot run `repair-data` operations for more than one node at the same time. To recover multiple nodes, contact technical support.

   - If your grid includes replicated data, use the `repair-data start-replicated-node-repair` command with the `--nodes` option to repair the entire Storage Node.
     This command repairs the replicated data on Storage Node named SG-DC-SN3:

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

     Note: As object data is restored, the LOST (Lost Objects) alarm is triggered if the StorageGRID system cannot locate replicated object data. Alarms might be triggered on Storage Nodes throughout the system. You should determine the cause of the loss and if recovery is possible. See the instructions for troubleshooting StorageGRID.

   - If your grid contains erasure coded data, use the `repair-data start-ec-node-repair` command with the `--nodes` option 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.

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

- If your grid has both replicated and erasure coded data, run both commands.

4. If only some of the volumes have failed, repair the affected volumes.

   Enter the volume IDs in hexadecimal, where 0000 is the first volume and 000F is the sixteenth volume. You can specify one volume or a range of volumes.

   - If your grid contains replicated data, use the `start-replicated-volume-repair` command with the `--nodes` and `--volume-range` options.

     This command restores replicated data to all volumes in the range 0003 to 000B on a Storage Node named SG-DC-SN3:

     ```
     repair-data start-replicated-volume-repair --nodes SG-DC-SN3 --
     volume-range 0003,000B
     ```

     For replicated data, you can run more than one `repair-data` operation at the same time for the same node. You might want to do this if you need to restore two volumes that are not in a range, such as 0000 and 000A.

     Note: As object data is restored, the LOST (Lost Objects) alarm is triggered if the StorageGRID system cannot locate replicated object data. Alarms might be triggered on Storage Nodes throughout the system. You should determine the cause of the loss and if recovery is possible. See the instructions for troubleshooting StorageGRID.

   - If your grid contains erasure coded data, use the `start-ec-volume-repair` command with the `--nodes` and `--volume-range` options.

     This command restores erasure coded data to a single volume 000A on a Storage Node named SG-DC-SN3:

     ```
     repair-data start-ec-volume-repair --nodes SG-DC-SN3 --volume-
     range 000A
     ```

     For erasure coded data, you must wait until one `repair-data start ec-volume-repair` operation completes before starting a second `repair-data` operation for the same node.

     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.

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

   - If your grid has both replicated and erasure coded data, run both commands.

5. Monitor the repair of replicated data.

   a. Select **Support** > **Grid Topology**.

   b. Select ***`deployment`*** > **Overview** > **ILM Activity**.

   c. Use a combination of the following attributes to determine, as well as possible, if replicated repairs are complete.

      Note: Cassandra inconsistencies might be present, and failed repairs are not 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.

  Note: High-risk objects are objects that are at risk of being completely lost. This does not include objects that do not 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 is at the deployment level and is the maximum of all node scan periods. You can query the **Scan Period – Estimated** attribute history at the deployment level to determine an appropriate time frame for your grid.

6. Monitor the repair of erasure coded data, and retry any requests that might have failed.

   a. Determine the status of erasure coded data repairs:

   - 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.

     ```
     root@DC1-ADM1:~ # repair-data show-ec-repair-status

      Repair ID   Scope                  Start Time  End Time  State  Est Bytes Affected Bytes Repaired  Retry Repair
     ================================================================================================================
      949283   DC1-S-99-10(Volumes: 1,2) 2016-11-30T15:27:06.9  Success   17359            17359           No
      949292   DC1-S-99-10(Volumes: 1,2) 2016-11-30T15:37:06.9  Failure   17359            0               Yes
      949294   DC1-S-99-10(Volumes: 1,2) 2016-11-30T15:47:06.9  Failure   17359            0               Yes
      949299   DC1-S-99-10(Volumes: 1,2) 2016-11-30T15:57:06.9  Failure   17359            0               Yes
     ```

   b. 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 83930030303133434:

      ```
      repair-data start-ec-node-repair --repair-id 83930030303133434
      ```

      This command retries a failed volume repair, using the repair ID 83930030303133434:

      ```
      repair-data start-ec-volume-repair --repair-id 83930030303133434
      ```

Related information

Administering StorageGRID
Troubleshooting StorageGRID

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

Steps

1. Select **Support** > **Grid Topology**.
2. Check the values of ***`Recovered Storage Node`*** > **LDR** > **Storage** > **Storage State – Desired** and **Storage State – Current**.

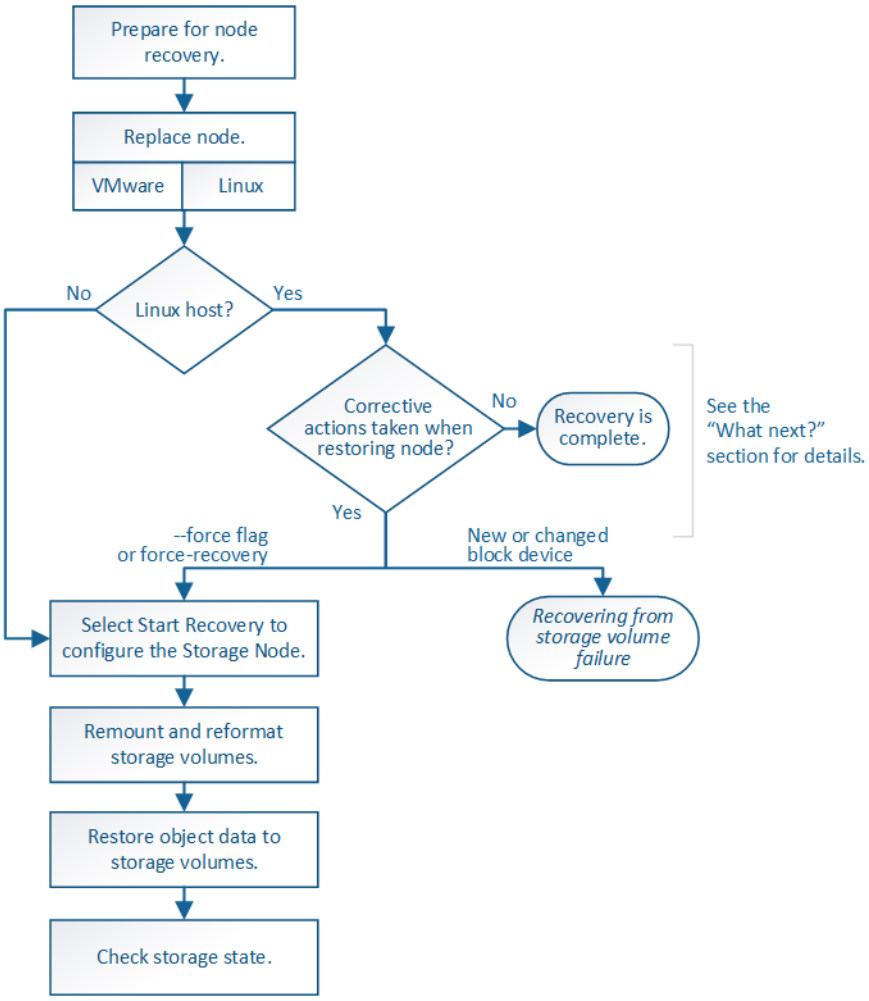
   The value of both attributes should be Online.
3. If the Storage State – Desired is set to Read-only, complete the following steps:
   a. Click the **Configuration** tab.
   b. From the **Storage State – Desired** drop-down list, select **Online**.
   c. Click **Apply Changes**.
   d. Click the **Overview**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

Recovering from system drive failure and possible storage volume failure

If the system drive on a virtual 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.

About this task

This recovery procedure applies to virtual Storage Nodes only.



Steps

1. Reviewing warnings for Storage Node system drive recovery on page 48
2. Replacing the Storage Node on page 49
3. Selecting Start Recovery to configure the Storage Node on page 50
4. Remounting and reformatting storage volumes ("Manual Steps") on page 51
5. Restoring object data to a storage volume on page 55
6. Checking the storage state after recovering a Storage Node system drive on page 59

Reviewing warnings for Storage Node system drive recovery

Before recovering a failed system drive of a Storage Node, you must review the following warnings.

About this task

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.

Caution: Do not perform this procedure if more than one Storage Node is offline, or if more than one Storage Node has a failure of any kind. Doing so might result in data loss. Contact technical support.

Caution: Do not rebuild Cassandra on more than one Storage Node within a 15-day period. Rebuilding Cassandra on two or more Storage Nodes within 15 days of each other might result in data loss. Contact technical support.

Caution: 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 “Recovering from loss of storage volumes where the system drive is intact” for instructions.

Attention: If ILM rules are configured to store only one replicated copy and the copy exists on a Storage Node that has failed, you will not be able to recover the object. However, you must still perform the procedure to restore object data to a storage volume to purge lost object information from the database.

Note: If you encounter a DDS alarm during recovery, see troubleshooting instructions to recover from the alarm by rebuilding Cassandra. After Cassandra is rebuilt, alarms should clear. If alarms do not clear, contact technical support.

Related tasks

Reviewing warnings and preconditions for node recovery on page 15
Recovering from storage volume failure where the system drive is intact on page 38

Replacing the Storage Node

If the system drive has failed, you must first replace the Storage Node.

About this task

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.

Linux: If you are not sure if your system drive has failed, follow the instructions to replace the node to determine which recovery steps are required.

Linux: The node replacement procedure will help you determine if grid node recovery is complete after you complete that procedure, and if not, which steps you need to take next.

Platform	Procedure
VMware	*Replacing a VMware node* on page 81
Linux	*Replacing a Linux node* on page 85
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.

Selecting Start Recovery to configure the Storage Node

After replacing the Storage Node, you must select Start Recovery in the Grid Manager to configure the new node as a replacement for the failed node.

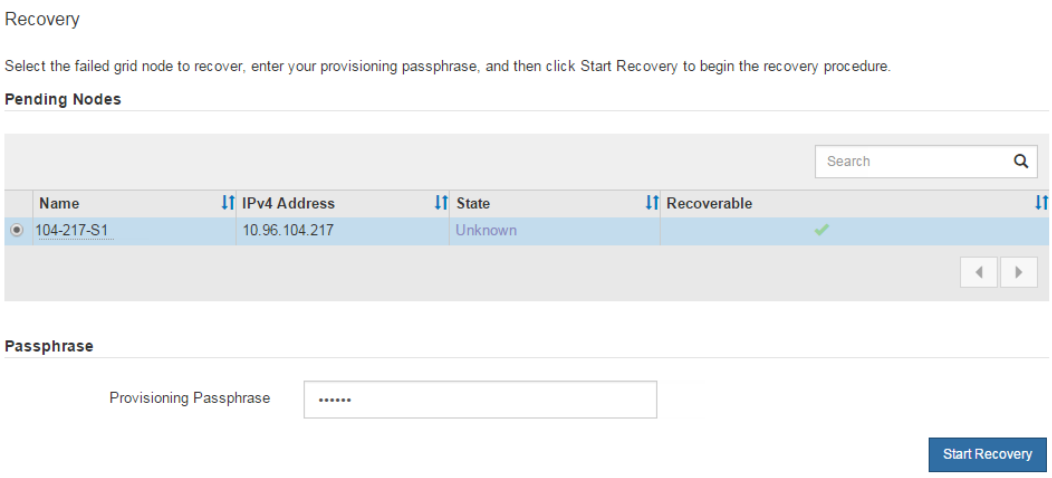
Before you begin

- You must be signed in to the Grid Manager using a supported browser.
- You must have a service laptop.
- You must have Maintenance or Root Access permissions.
- You must have the provisioning passphrase.
- You must have deployed and configured the replacement node.
- You must know the start date of any repair jobs for erasure-coded data.
- You must have verified that the Storage Node has not been rebuilt within the last 15 days.

Steps

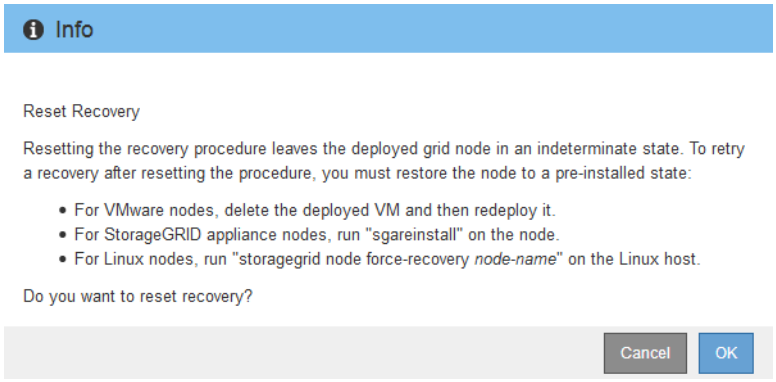
1. From the Grid Manager, select **Maintenance > Recovery**.
2. Select the grid node you want to recover in the **Pending Nodes** list.

   Nodes appear in the list after they fail, but you cannot 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.

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

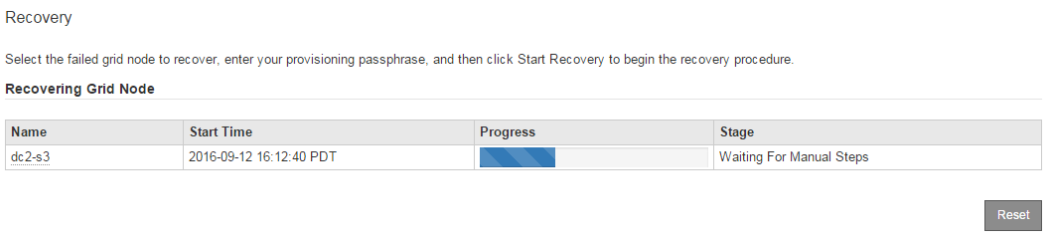


If you want to retry the recovery after resetting the procedure, you must restore the node to a pre-installed state, as follows:

- **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 stage “Waiting for Manual Steps” stage, go to the next task in the recovery procedure, to remount and reformat storage volumes.



Remounting and reformatting storage volumes (“Manual Steps”)

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

Before you begin

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

  This procedure might help you identify additional failed storage volumes.
- You have checked that a Storage Node decommissioning is not in progress, or you have paused decommissioning. (In the Grid Manager, go to **Maintenance > Decommission**.)
- You have checked that an expansion is not in progress. (In the Grid Manager go to **Maintenance > Expansion**.)
- You have reviewed the warnings for Storage Node system drive recovery.

  Reviewing warnings for Storage Node system drive recovery on page 48

Caution: 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. Do not 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 first script, `sn-remount-volumes`, 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.
- Review the script output and resolve any issues. Then, rerun the script until you are satisfied that all valid storage volumes have been remounted.
- Run the second script, `sn-recovery-postinstall.sh`, to reformat any storage volumes that the first script could not mount or that were found to be improperly formatted; rebuild Cassandra, if needed; and start services.

  Note: All data on these storage volumes is lost, and must be recovered from other locations in the grid if it is possible.

Steps

1. From the service laptop, 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.

   Note: 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`

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

   Note: 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.

Example

```
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.

Do not continue to the next step if you believe that the data remaining on
this volume cannot 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.

Do not continue to the next step if you believe that the data remaining on
this volume cannot 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 uninitialized or the disk's superblock was corrupted. When the script cannot 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 do not 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.

   Attention: 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` 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 are not 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.

      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, replace or remove the disk and run the `sn-remount-volumes` script again to ensure the issue is resolved.

      As needed, you can find the node ID for the Storage Node you are recovering at the top of the script (the "configured LDR noid"). You can look up node IDs for other Storage Nodes in the Grid Manager. Select **Support** > **Grid Topology** > ***`Site`*** > ***`Storage Node`*** > **LDR** > **Overview**.

      Caution: If you are unable to resolve the issue, contact technical support. If you run the `sn-recovery-postinstall.sh` script, the storage volume will be reformatted, which might cause data loss.

   c. Review the messages for devices that could not be mounted, and make a note of the device name for each failed storage volume.

      Note: 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 in the next procedure (restoring object data).

   d. Run the `sn-remount-volumes` script again to ensure that all valid storage volumes have been remounted.

   Attention: If a storage volume could not 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).

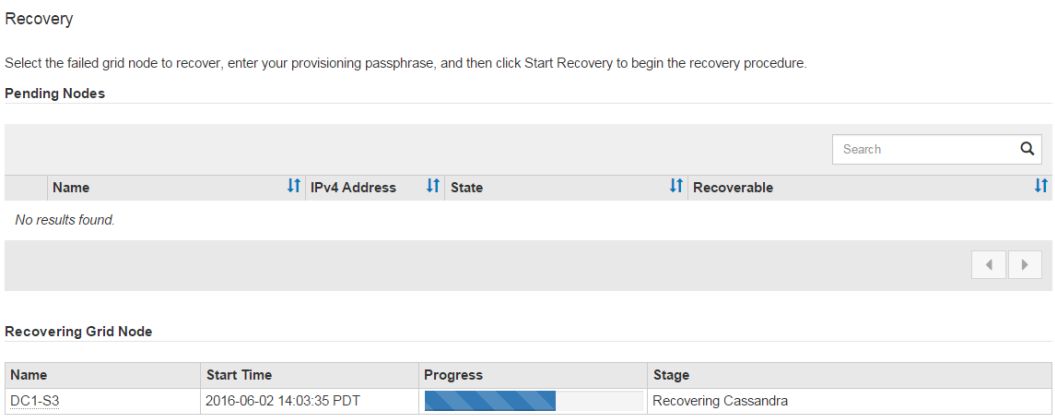
   Caution: Do not run the `sn-recovery-postinstall` script if you believe that the data remaining on this volume cannot be rebuilt from elsewhere in the grid (for example, if your ILM policy uses a rule that makes only one copy or if volumes have failed on multiple nodes). Instead, contact technical support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

4. Run the second script to reformat all unmounted (failed) storage volumes, rebuild Cassandra if required, and start services on the Storage Node:

   `sn-recovery-postinstall.sh`

5. As the script runs, monitor the Recovery page in the Grid Manager, and review the command line output, which provides more detailed status information.

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



After you finish

Restore object data to any storage volumes that were formatted by `sn-recovery-postinstall.sh`, as described in the next procedure. You can restore object data at the same time Cassandra is being rebuilt.

Related tasks

Reviewing warnings for Storage Node system drive recovery on page 48
Restoring object data to a storage volume on page 55

Restoring object data to a storage volume

After recovering a storage volume on a Storage Node where the system drive also failed and was recovered, you can restore object data to the recovered storage volume from other Storage Nodes and Archive Nodes.

Before you begin

- You must have confirmed that the recovered Storage Node displays in green in the Grid Topology tree and shows all services as Online.

About this task

Object data is restored from other Storage Nodes and Archive Nodes, assuming that the grid's ILM rules were configured such that object copies are available.

Attention: If ILM rules are configured to store only one replicated copy and that copy exists on a Storage Node that has failed, you will not be able to recover the object. However, you must still perform the procedure to restore object data to a storage volume to purge lost object information from the database.

Note: If the only remaining copy of object data is located on an Archive Node, object data is retrieved from the Archive Node. Due to the latency associated with retrievals from external

archival storage systems, restoring object data to a Storage Node from an Archive Node takes longer than restoring copies from other Storage Nodes.

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. You use different options with the `repair-data` script, based on whether you are restoring replicated data or erasure coded data, as follows:

- **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
```

- **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
```

  Repairs of erasure coded data can begin while some Storage Nodes are offline. Repair will complete after all nodes are available. You can track repairs of erasure coded data with this command:

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

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

Steps

1. From the service laptop, 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 host name 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`
3. If all storage volumes have failed, repair the entire node. (If only some volumes have failed, go to the next step.)

   Attention: You cannot run `repair-data` operations for more than one node at the same time. To recover multiple nodes, contact technical support.

   - If your grid includes replicated data, use the `repair-data start-replicated-node-repair` command with the `--nodes` option to repair the entire Storage Node.

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

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

Note: As object data is restored, the LOST (Lost Objects) alarm is triggered if the StorageGRID system cannot locate replicated object data. Alarms might be triggered on Storage Nodes throughout the system. You should determine the cause of the loss and if recovery is possible. See the instructions for troubleshooting StorageGRID.

- If your grid contains erasure coded data, use the `repair-data start-ec-node-repair` command with the `--nodes` option 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.

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

- If your grid has both replicated and erasure coded data, run both commands.

4. If only some of the volumes have failed, repair the affected volumes.

   Enter the volume IDs in hexadecimal, where 0000 is the first volume and 000F is the sixteenth volume. You can specify one volume or a range of volumes.

   - If your grid contains replicated data, use the `start-replicated-volume-repair` command with the `--nodes` and `--volume-range` options.
     This command restores replicated data to all volumes in the range 0003 to 000B on a Storage Node named SG-DC-SN3:

     ```
     repair-data start-replicated-volume-repair --nodes SG-DC-SN3 --
     volume-range 0003,000B
     ```

     For replicated data, you can run more than one `repair-data` operation at the same time for the same node. You might want to do this if you need to restore two volumes that are not in a range, such as 0000 and 000A.

     Note: As object data is restored, the LOST (Lost Objects) alarm is triggered if the StorageGRID system cannot locate replicated object data. Alarms might be triggered on Storage Nodes throughout the system. You should determine the cause of the loss and if recovery is possible. See the instructions for troubleshooting StorageGRID.

   - If your grid contains erasure coded data, use the `start-ec-volume-repair` command with the `--nodes` and `--volume-range` options.
     This command restores erasure coded data to a single volume 000A on a Storage Node named SG-DC-SN3:

     ```
     repair-data start-ec-volume-repair --nodes SG-DC-SN3 --volume-
     range 000A
     ```

     For erasure coded data, you must wait until one `repair-data start ec-volume-repair` operation completes before starting a second `repair-data` operation for the same node.

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.

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

- If your grid has both replicated and erasure coded data, run both commands.

5. Monitor the repair of replicated data.
   a. Select **Support > Grid Topology**.
   b. Select ***`deployment`*** **> Overview > ILM Activity**.
   c. Use a combination of the following attributes to determine, as well as possible, if replicated repairs are complete.

      Note: Cassandra inconsistencies might be present, and failed repairs are not 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.

        Note: High-risk objects are objects that are at risk of being completely lost. This does not include objects that do not 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 is at the deployment level and is the maximum of all node scan periods. You can query the **Scan Period – Estimated** attribute history at the deployment level to determine an appropriate time frame for your grid.

6. Monitor the repair of erasure coded data, and retry any requests that might have failed.
   a. Determine the status of erasure coded data repairs:
      - 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.

        ```
        root@DC1-ADM1:~ # repair-data show-ec-repair-status

         Repair ID   Scope                  Start Time  End Time  State  Est Bytes Affected Bytes Repaired  Retry Repair
        ================================================================================================================
         949283   DC1-S-99-10(Volumes: 1,2) 2016-11-30T15:27:06.9  Success   17359           17359           No
         949292   DC1-S-99-10(Volumes: 1,2) 2016-11-30T15:37:06.9  Failure   17359           0               Yes
         949294   DC1-S-99-10(Volumes: 1,2) 2016-11-30T15:47:06.9  Failure   17359           0               Yes
         949299   DC1-S-99-10(Volumes: 1,2) 2016-11-30T15:57:06.9  Failure   17359           0               Yes
        ```

   b. 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 83930030303133434:

```
repair-data start-ec-node-repair --repair-id 83930030303133434
```

This command retries a failed volume repair, using the repair ID 83930030303133434:

```
repair-data start-ec-volume-repair --repair-id 83930030303133434
```

Related information

Administering StorageGRID
Troubleshooting StorageGRID

Checking the storage state after recovering a 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 must be signed in to the Grid Manager using a supported browser.
- The Storage Node has been recovered, and data recovery is complete.

Steps

1. Select **Support** > **Grid Topology**.
2. Check the values of ***`Recovered Storage Node`*** > **LDR** > **Storage** > **Storage State – Desired** and **Storage State – Current**.

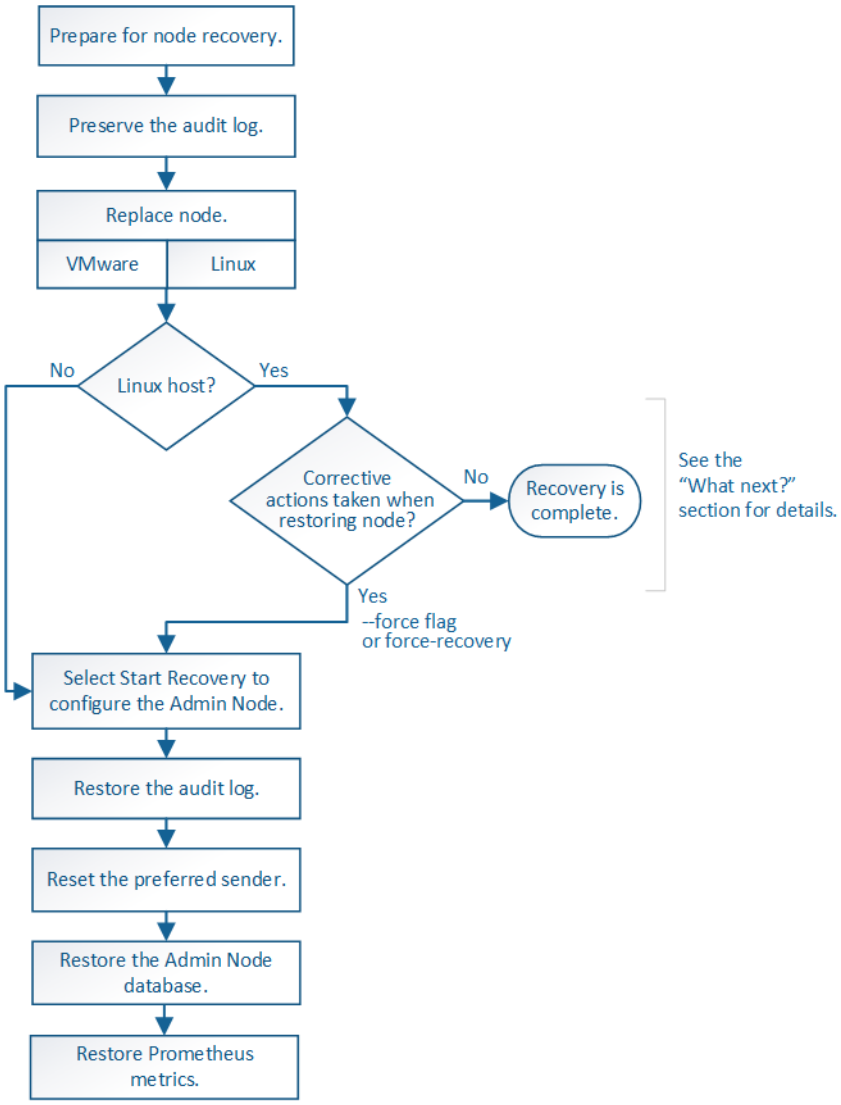
   The value of both attributes should be Online.
3. If the Storage State – Desired is set to Read-only, complete the following steps:
   a. Click the **Configuration** tab.
   b. From the **Storage State – Desired** drop-down list, select **Online**.
   c. Click **Apply Changes**.
   d. Click the **Overview**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

Recovering from Admin Node failures

The recovery process for an Admin Node depends on whether it is the primary Admin Node or a non-primary Admin Node.

About this task

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

- Recovering from primary Admin Node failures on page 60
- Recovering from non-primary Admin Node failures on page 68

Recovering from primary Admin Node failures

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.

About this task

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 cannot 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.

Attention: 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.

Steps

1. Copying audit logs from the failed primary Admin Node on page 61
2. Replacing the primary Admin Node on page 62
3. Configuring the replacement primary Admin Node on page 62
4. Restoring the audit log on the recovered primary Admin Node on page 63
5. Resetting the preferred sender on the recovered primary Admin Node on page 64
6. Restoring the Admin Node database when recovering a primary Admin Node on page 65
7. Restoring Prometheus metrics when recovering a primary Admin Node on page 66

Copying audit logs from the 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 are not 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.

Note: If the audit logs are not accessible on the failed Admin Node now, you might be able to access them later, for example, after host recovery.

Steps

1. From the service laptop, 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 file on the recovered Admin Node when you copy it to the recovered Admin Node.

   Rename `audit.log` to a unique numbered file name such as *`yyyy-mm-dd.txt.1`*. For example, you can rename the `audit.log` file to `2015-10-25.txt.1`

   ```
   cd /var/local/audit/export
   ```

```
ls -l
mv audit.log 2015-10-25.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
   ```

Replacing the primary Admin Node

To recover a primary Admin Node, you must first replace the node.

About this task

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.

Linux: The node replacement procedure will help you determine if grid node recovery is complete after you complete that procedure, and if not, which steps you need to take next.

Platform	Procedure
VMware	*Replacing a VMware node* on page 81
Linux	*Replacing a Linux node* on page 85
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.

Configuring the replacement primary Admin Node

The replacement node must be configured as the primary Admin Node for your StorageGRID system.

Before you begin

- The virtual machine for the primary Admin Node is deployed, powered on, and initialized.
- 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 Server 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.

   Note: To configure a relying party trust, see the instructions for administering StorageGRID. 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 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 do not 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.

Related concepts

Hotfix procedure on page 7

Related information

Administering StorageGRID

Restoring the audit log on the 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 must be installed and running.
- You must 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 cannot 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.

Steps

1. From the service laptop, 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`**

After you finish

You must also restore any pre-existing client access to the audit share. For more information, see the instructions for administering StorageGRID.

Related information

Administering StorageGRID

Resetting the preferred sender on the recovered primary Admin Node

If the primary Admin Node you are recovering is currently set as the preferred sender of notifications and AutoSupport messages, you must reconfigure this setting in the StorageGRID system.

Before you begin

- You must be signed in to the Grid Manager using a supported browser.

- You must have specific access permissions. For details, see the instructions for administering StorageGRID.
- The recovered Admin Node must be installed and running.

Steps

1. Select **Configuration** > **Display Options**.
2. Select the recovered Admin Node from the **Preferred Sender** drop-down list.
3. Click **Apply Changes**.

Related information

Administering StorageGRID

Restoring the Admin Node database when recovering a primary Admin Node

If you want to retain the historical information about attribute values and alarms on a primary Admin Node that has failed, you can restore the Admin Node database. This database can only be restored if your StorageGRID system includes another Admin Node.

Before you begin

- The recovered Admin Node must be installed and running.
- The StorageGRID system must include at least two Admin Nodes.
- You must have the `Passwords.txt` file.
- You must have the provisioning passphrase.

About this task

If an Admin Node fails, the historical information about attribute values and alarms that are stored in its Admin Node database are lost. When you recover the Admin Node, the software installation process creates a new database for the NMS service. After the recovered Admin Node is started, it records attribute and audit information for all services 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 this 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 cannot restore the Admin Node database.

Note: Copying the Admin Node database may 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. 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. 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 database from the source Admin Node to the recovered Admin Node:

      ```
      /usr/local/mi/bin/mi-clone-db.sh Source_Admin_Node_IP
      ```

   f. 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. The following status appears:

      ```
      Database cloned, starting services
      starting mi ... done
      ```

   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 MI service on the source Admin Node.

   ```
   service mi start
   ```

Restoring Prometheus metrics when recovering a 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 must be installed and running.
- The StorageGRID system must include at least two Admin Nodes.
- You must have the `Passwords.txt` file.
- You must 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 cannot restore the Prometheus database.

Note: 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
   ```

Recovering from non-primary Admin Node failures

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.

Steps

1. Copying audit logs from the failed non-primary Admin Node on page 68
2. Replacing a non-primary Admin Node on page 69
3. Selecting Start Recovery to configure a non-primary Admin Node on page 69
4. Restoring the audit log on the recovered non-primary Admin Node on page 71
5. Resetting the preferred sender on the recovered non-primary Admin Node on page 72
6. Restoring the Admin Node database when recovering a non-primary Admin Node on page 72
7. Restoring Prometheus metrics when recovering a non-primary Admin Node on page 74

Copying audit logs from the 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.

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 are not 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.

Note: If the audit logs are not accessible on the failed Admin Node now, you might be able to access them later, for example, after host recovery.

Steps

1. From the service laptop, 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 file on the recovered Admin Node when you copy it to the recovered Admin Node.

   Rename `audit.log` to a unique numbered file name such as *`yyyy-mm-dd.txt.1`*. For example, you can rename the `audit.log` file to `2015-10-25.txt.1`

   ```
   cd /var/local/audit/export
   ls -l
   mv audit.log 2015-10-25.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
   ```

Replacing a non-primary Admin Node

To recover a non-primary Admin Node, you first must replace the node.

About this task

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.

Linux: The node replacement procedure will help you determine if grid node recovery is complete after you complete that procedure, and if not, which steps you need to take next.

Platform	Procedure
VMware	*Replacing a VMware node* on page 81
Linux	*Replacing a Linux node* on page 85
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.

Selecting Start Recovery to configure a 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

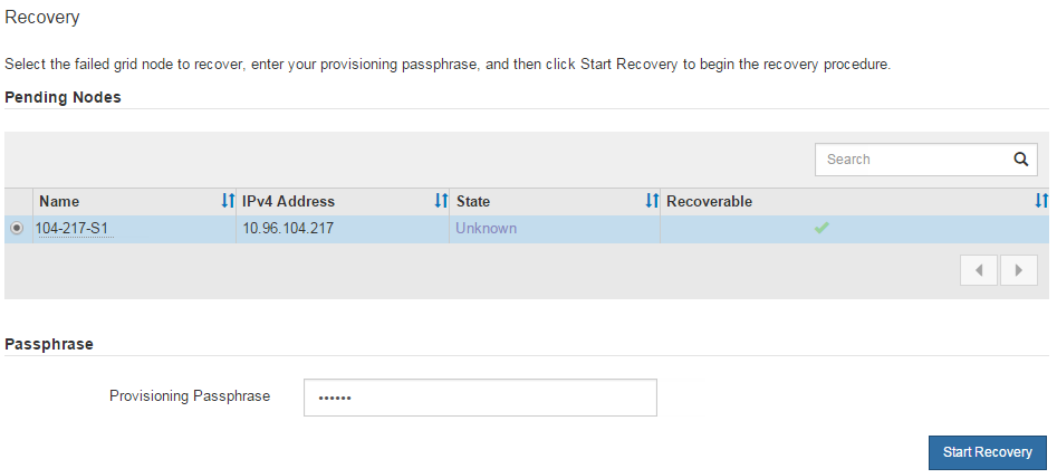
- You must be signed in to the Grid Manager using a supported browser.
- You must have a service laptop.

- You must have Maintenance or Root Access permissions.
- You must have the provisioning passphrase.
- You must have deployed and configured the replacement node.

Steps

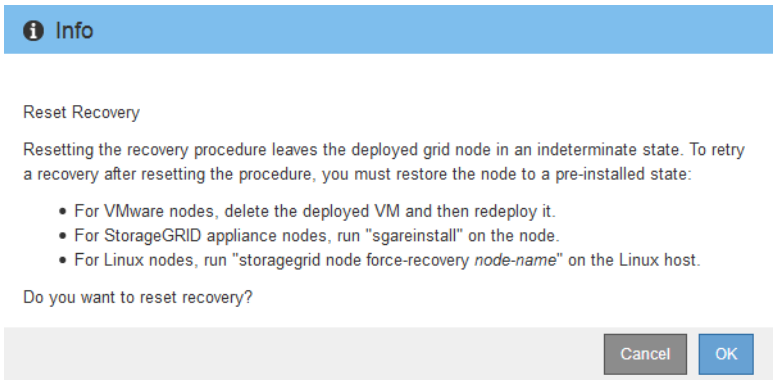
1. From the Grid Manager, select **Maintenance > Recovery**.
2. Select the grid node you want to recover in the **Pending Nodes** list.

   Nodes appear in the list after they fail, but you cannot 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.

   Note: At any point during the recovery, you can click **Reset** to start a new recovery. An Info 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. 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 Server 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.

   Note: To configure a relying party trust, see the instructions for administering StorageGRID. 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.

Related information

Administering StorageGRID

Restoring the audit log on the 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 must be installed and running.
- You must 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 cannot 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.

Steps

1. From the service laptop, 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`

After you finish

You must also restore any pre-existing client access to the audit share. For more information, see the instructions for administering StorageGRID.

Related information

Administering StorageGRID

Resetting the preferred sender on the recovered non-primary Admin Node

If the non-primary Admin Node you are recovering is currently set as the preferred sender of notifications and AutoSupport messages, you must reconfigure this setting in the StorageGRID system.

Before you begin

- You must be signed in to the Grid Manager using a supported browser.
- You must have specific access permissions. For details, see the instructions for administering StorageGRID.
- The recovered Admin Node must be installed and running.

Steps

1. Select **Configuration** > **Display Options**.
2. Select the recovered Admin Node from the **Preferred Sender** drop-down list.
3. Click **Apply Changes**.

Related information

Administering StorageGRID

Restoring the Admin Node database when recovering a non-primary Admin Node

If you want to retain the historical information about attribute values and alarms 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 must be installed and running.
- The StorageGRID system must include at least two Admin Nodes.

- You must have the `Passwords.txt` file.
- You must have the provisioning passphrase.

About this task

If an Admin Node fails, the historical information about attribute values and alarms that are stored in its Admin Node database are lost. When you recover the Admin Node, the software installation process creates a new database for the NMS service. After the recovered Admin Node is started, it records attribute and audit information for all services as if you had performed a new installation of the StorageGRID system.

If you restored a non-primary Admin Node, you can restore this historical information by copying the Admin Node database from the primary Admin Node (the *source Admin Node*) to the recovered node.

Note: Copying the Admin Node database may 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 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. 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. 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 database from the source Admin Node to the recovered Admin Node:

      ```
      /usr/local/mi/bin/mi-clone-db.sh Source_Admin_Node_IP
      ```

f. 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. The following status appears:

```
Database cloned, starting services
starting mi ... done
```

g. 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 MI service on the source Admin Node.

```
service mi start
```

Restoring Prometheus metrics when recovering a 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 must be installed and running.
- The StorageGRID system must include at least two Admin Nodes.
- You must have the `Passwords.txt` file.
- You must 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.

Note: 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:

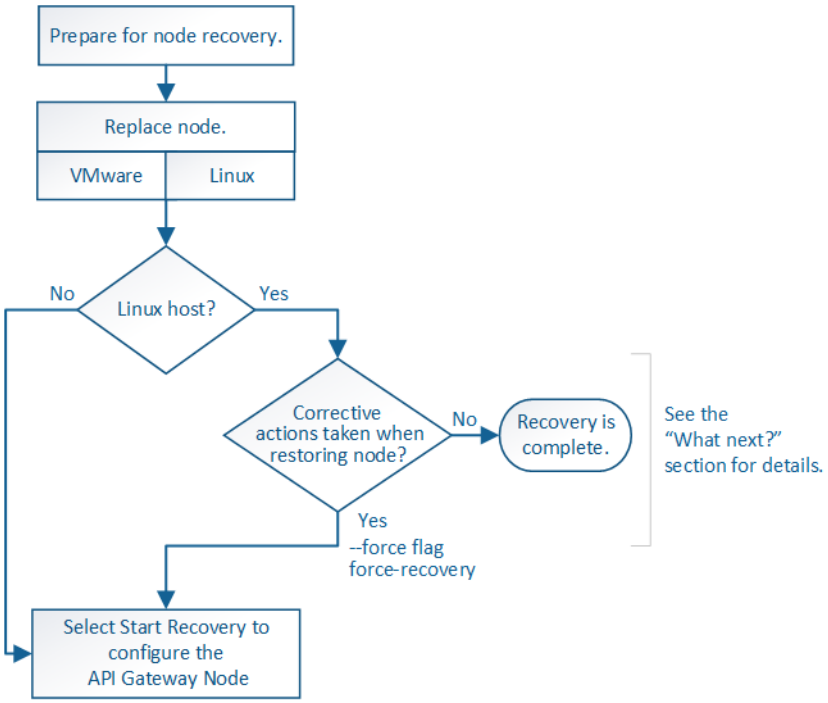
   `ssh-add -D`

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

   `service prometheus start`

Recovering from API Gateway Node failures

You must complete a sequence of tasks in exact order to recover from an API Gateway Node failure.



Steps

1. Replacing the API Gateway Node on page 76
2. Selecting Start Recovery to configure the API Gateway Node on page 77

Replacing the API Gateway Node

To recover an API Gateway Node, you must first replace the node.

About this task

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.

Linux: The node replacement procedure will help you determine if grid node recovery is complete after you complete that procedure, and if not, which steps you need to take next.

Platform	Procedure
VMware	*Replacing a VMware node* on page 81
Linux	*Replacing a Linux node* on page 85
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.

Selecting Start Recovery to configure the API Gateway Node

After replacing the API Gateway Node, you must select Start Recovery in the Grid Manager to configure the new node as a replacement for the failed node.

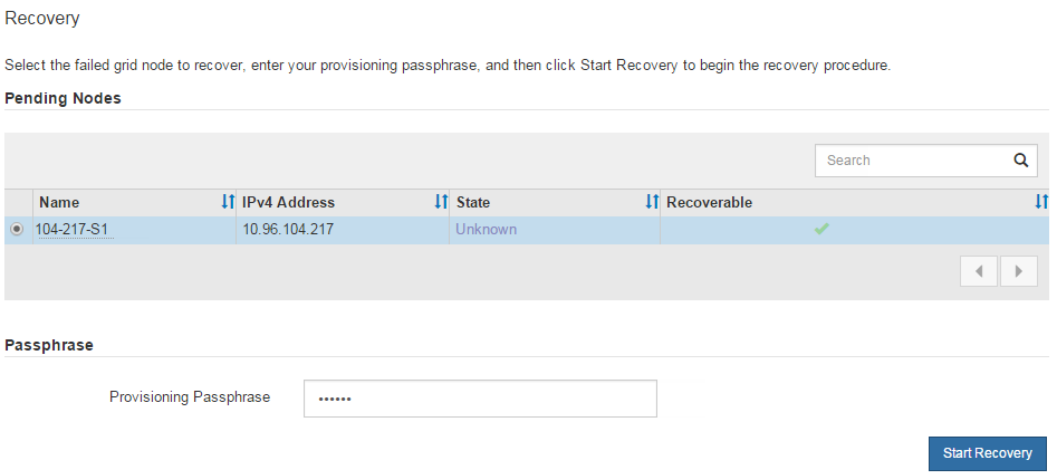
Before you begin

- You must be signed in to the Grid Manager using a supported browser.
- You must have a service laptop.
- You must have Maintenance or Root Access permissions.
- You must have the provisioning passphrase.
- You must have deployed and configured the replacement node.

Steps

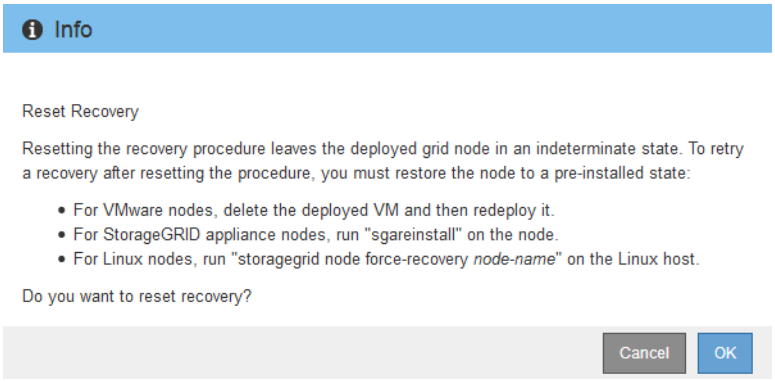
1. From the Grid Manager, select **Maintenance** > **Recovery**.
2. Select the grid node you want to recover in the **Pending Nodes** list.

   Nodes appear in the list after they fail, but you cannot 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.

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



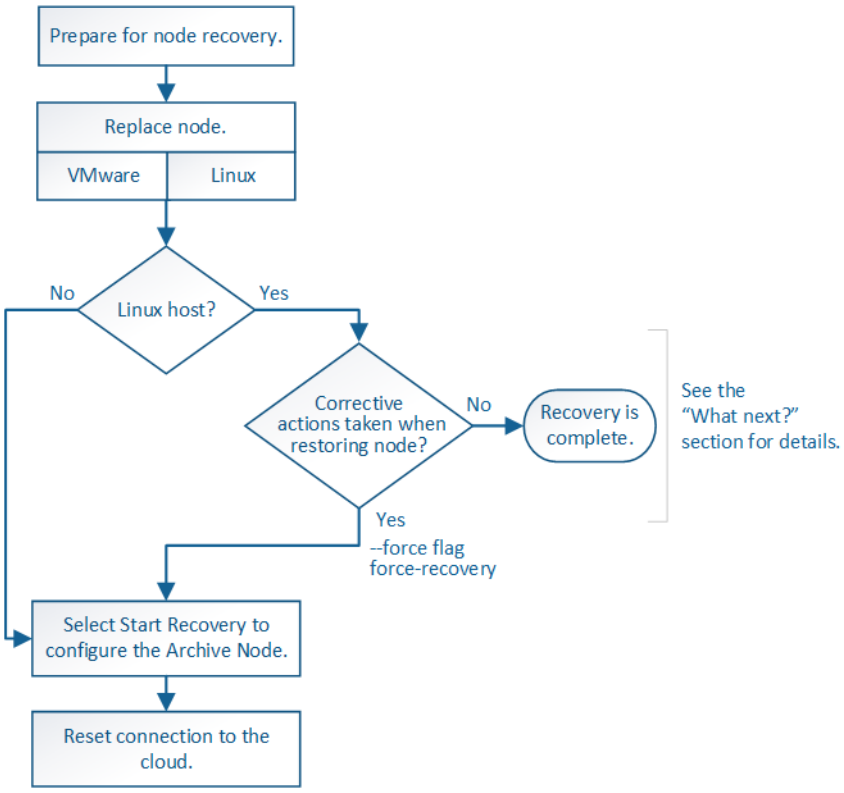
If you want to retry the recovery after resetting the procedure, you must restore the node to a pre-installed state, as follows:

- **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`*

Recovering from Archive Node failures

You must complete a sequence of tasks in exact order to recover from an Archive Node failure.



About this task

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.

Steps

1. Replacing the Archive Node on page 79
2. Selecting Start Recovery to configure the Archive Node on page 79
3. Resetting Archive Node connection to the cloud on page 81

Replacing the Archive Node

To recover an Archive Node, you must first replace the node.

About this task

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.

Linux: The node replacement procedure will help you determine if grid node recovery is complete after you complete that procedure, and if not, which steps you need to take next.

Platform	Procedure
VMware	*Replacing a VMware node* on page 81
Linux	*Replacing a Linux node* on page 85
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.

Selecting Start Recovery to configure the Archive Node

After replacing the 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

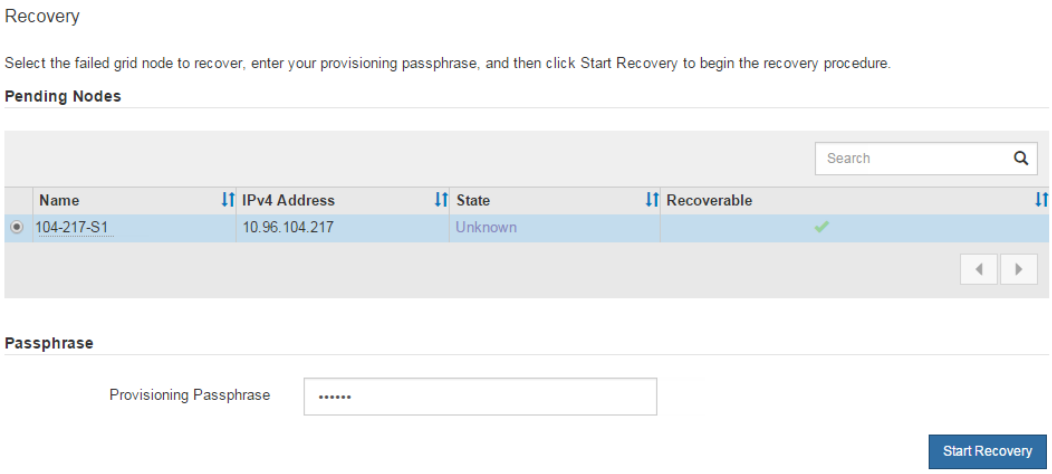
- You must be signed in to the Grid Manager using a supported browser.
- You must have a service laptop.
- You must have Maintenance or Root Access permissions.
- You must have the provisioning passphrase.

- You must have deployed and configured the replacement node.

Steps

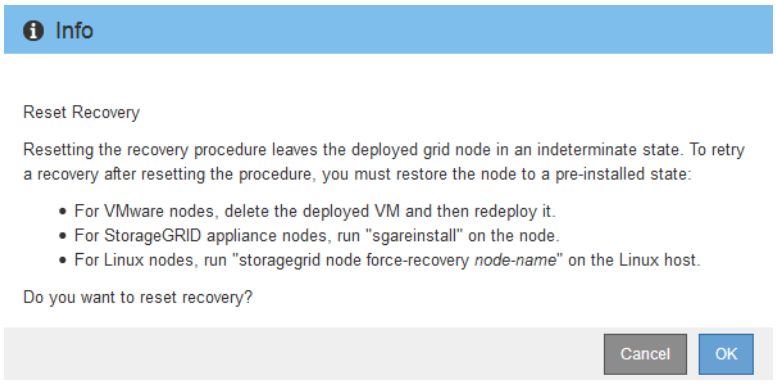
1. From the Grid Manager, select **Maintenance > Recovery**.
2. Select the grid node you want to recover in the **Pending Nodes** list.

   Nodes appear in the list after they fail, but you cannot 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.

   Note: At any point during the recovery, you can click **Reset** to start a new recovery. An Info 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
```

Resetting 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.

Before you begin

Note: If your Archive Node connects to external storage through TSM middleware, then the node resets itself automatically and you do not need to reconfigure.

You must be signed in to the Grid Manager using a supported browser.

Steps

1. Select **Support** > **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: Replacing a VMware node

When you recover a failed grid node of any type that was hosted on VMware, you must use VMware vSphere to first remove the failed node and then deploy a recovery node. This procedure is one step of the grid node recovery process.

Before you begin

You must have determined that the virtual machine cannot be restored, and must be replaced.

About this task

This procedure is only performed as one step in the process of recovering non-appliance Storage Nodes, primary or non-primary Admin Nodes, API Gateway Nodes, or Archive Nodes. The steps are identical regardless of the type of grid node you are recovering.

Steps

1. Removing the failed grid node in VMware vSphere on page 81
2. Deploying the recovery grid node in VMware vSphere on page 82

Removing the failed grid node in VMware vSphere

You must remove the virtual machine associated with the failed grid node using the VMware vSphere Web Client before you can recover it.

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. Right-click the virtual machine and select **Edit Settings** tab, and note the settings in use. Click 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.

Deploying the recovery grid node in VMware vSphere

You start the recovery of a grid node by deploying a new virtual machine using the VMware vSphere Web Client.

Before you begin

- You have the `.ovf` and `.mf` files for the grid nodes you are deploying:

Filename	Description
`vsphere-non-primary-admin.ovf` `vsphere-non-primary-admin.mf`	The template file and manifest file for deploying non-primary Admin Nodes.
`vsphere-archive.ovf` `vsphere-archive.mf`	The template file and manifest file for deploying Archive Nodes.
`vsphere-gateway.ovf` `vsphere-gateway.mf`	The template file and manifest file for deploying API Gateway Nodes.
`vsphere-storage.ovf` `vsphere-storage.mf`	The template file and manifest file for deploying Storage Nodes.

- You have placed all of these files in the same directory.
- You have the StorageGRID Virtual Machine Disk (`.vmdk`) file, and it is in the same folder as the `.ovf` and `.mf` files:
  `NetApp-SG-`*`version-SHA`*`.vmdk`

  Note: The same `.vmdk` file is used for all types of nodes.

Caution: You must deploy the new VM using the same StorageGRID version as is currently running on the grid.

Steps

1. Open VMware vSphere Web Client, and sign in.
2. Navigate to the vApp or resource pool where you want to deploy the StorageGRID grid, and select **Actions > All vCenter Actions > Deploy OVF Template**.
3. Select the `vsphere-`*`node`*`.ovf` and `NetApp-SG-`*`version-SHA`*`.vmdk` files.
4. Specify the name of the virtual machine.

   The best practice is to use the same name for the virtual machine as you used for the grid node.
5. In the **Network Mapping** page, select the networks to use by associating a network port to each network. The Grid Network is required. The Admin and Client Networks are optional. Select the Grid Network to use, and then choose the following as applicable:

- If you are planning to use the Admin Network, assign the Admin Network adapter to a network in the vSphere environment.
- If you are planning to use the Client Network, assign the Client Network adapter to a network in the vSphere environment.
- If you do not plan to use an Admin Network or Client Network, assign their network adapters to the same network as the Grid Network.

6. Provide the required StorageGRID information in the **Properties** page, and click **Finish**.

a. Enter the **Node Name**.

Attention: If you are performing this task because you are recovering a grid node, you must use the same name for the replacement node that was used for the node you are recovering.

b. Enter the **Primary Admin IP**.

If you omit the primary Admin Node IP address, the IP address will be automatically discovered if the primary Admin Node, or at least one other grid node with ADMIN_IP configured, is present on the same subnet. However, it is recommended to set the primary Admin Node IP address here.

c. In the **Grid Network (eth0)** section, under **Grid Network IP configuration**, select STATIC or DHCP.

- If you select STATIC, enter the **Grid Network IP**, **Grid Network mask**, and **Grid Network gateway**.
- If you select DHCP, the **Grid Network IP**, **Grid Network mask**, and **Grid Network gateway** are automatically assigned.

d. In the **Admin Network (eth1)** section, under **Admin Network IP configuration**, select STATIC, DHCP, or DISABLED.

- If you select STATIC, enter the **Admin Network IP**, **Admin Network mask**, and **Admin Network gateway**.
- If you select STATIC, enter the **Admin network external subnet list**. You must also configure a gateway.
- If you select DHCP, the **Admin Network IP**, **Admin Network mask**, and **Admin Network gateway** are automatically assigned.
- If you do not want to use the Admin Network (eth1), select DISABLED and enter **0.0.0.0** for the Admin Network IP. You can leave the other fields blank.

e. In the **Client Network (eth2)** section, under **Client Network IP configuration**, select STATIC, DHCP, or DISABLED.

- If you select STATIC, enter the **Client Network IP**, **Client Network mask**, and **Client Network gateway**.
- If you select DHCP, the **Client Network IP**, **Client Network mask**, and **Client Network gateway** are automatically assigned.
- If you do not want to use the Client Network (eth2), select DISABLED and enter **0.0.0.0** for the Client Network IP. You can leave the other fields blank.

7. Click **Next** and then **Finish** to start the upload of the virtual machine.

8. If this is not a full node recovery, perform these steps after deployment is complete:

a. Right-click the virtual machine, and select the **Edit Settings** tab.

b. Select each default virtual hard disk that has been designated for storage, and click the **Remove** button located at the top of the tab.

c. Depending on your data recovery circumstances, add new virtual disks according to your storage requirements, or reattach any virtual hard disks preserved from the previously removed failed grid node, or both.

Note the following important guidelines:

- In general, if you are adding new disks you should use the same type of storage device that was in use before node recovery.
- The Storage Node OVF provided defines several VMDKs for storage. You should remove these and assign appropriate VMDKs or RDMs for storage before powering up the node. VMDKs are more commonly used in VMware environments and are easier to manage, while RDMs may provide better performance for workloads that use larger object sizes (for example, greater than 100 MB).
- If the host for the StorageGRID node will use storage assigned from a NetApp All Flash FAS (AFF) system, you must confirm that the FlexVol does not have a tiering policy enabled.

  Attention: Never assign storage for a StorageGRID node from a FlexVol with an active tiering policy. Node outages might occur if storage used by the StorageGRID node is tiered to a capacity tier.

9. If you want to remap ports used by a node:

a. If you specified DISABLED for the Client network IP configuration, you must enter **0.0.0.0.** for the Client Network IP under the under the **Client Network (eth2)** section. Completing this field is required.

b. Right-click on the VM, and select **Edit Settings**.

c. Select **vApp Options**.

d. In the **Authoring** section, expand **Properties** and scroll down until you see PORT_REMAP_INBOUND and PORT_REMAP.

You might need to remap a port if your enterprise networking policies restrict access to one or more ports that are used by StorageGRID. See “Internal grid node communications” or “External client communications” for the list of ports used by StorageGRID.

e. To symmetrically map both inbound and outbound communications for a port, select **PORT_REMAP** and click **Edit**.

Enter the port mapping as *<network type>*/*<protocol>*/*<default port used by grid node>*/*<new port>*, where network type is grid, admin, or client, and protocol is tcp or udp.

Example

To remap ssh traffic from port 22 to port 3022, enter the following:

```
client/tcp/22/3022
```

Click **OK**.

Note: If only PORT_REMAP is set, the mapping that you specify applies to both inbound and outbound communications. If PORT_REMAP_INBOUND is also specified, PORT_REMAP applies only to outbound communications.

f. To specify the port used for inbound communications to the node, select **PORT_REMAP_INBOUND** and click **Edit**.

Enter the port mapping as *<network type>*/*<protocol>*/*<remapped inbound port>*/*<default inbound port used by grid node>*, where network type is grid, admin, or client, and protocol is tcp or udp.

Example

To remap inbound SSH traffic that is sent to port 3022 so that it is received at port 22 by the grid node, enter the following:

```
client/tcp/3022/22
```

Click **OK**.

Note: If you specify PORT_REMAP_INBOUND and do not specify a value for PORT_REMAP, outbound communications for the port are unchanged.

10. Power on the virtual machine.

After you finish

To complete the recovery, return to the procedure for the failure you are addressing.

Type of recovery	Reference
Primary Admin Node	*Configuring the replacement primary Admin Node* on page 62
Non-primary Admin Node	*Selecting Start Recovery to configure a non-primary Admin Node* on page 69
API Gateway Node	*Selecting Start Recovery to configure the API Gateway Node* on page 77
Archive Node	*Selecting Start Recovery to configure the Archive Node* on page 79
Storage Node (virtual)	*Selecting Start Recovery to configure the Storage Node* on page 50

All grid node types: Replacing a Linux node

If a failure requires that you deploy new physical or virtual hosts or reinstall Linux on an existing host, you must deploy and configure replacement hosts before you can recover individual grid nodes. This procedure is one step of the grid node recovery process for all types of grid node.

About this task

"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.

This procedure is only performed as one step in the process of recovering virtual Storage Nodes, primary or non-primary Admin Nodes, API 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.

Steps

1. *Deploying new Linux hosts* on page 86
2. *Restoring grid nodes to the host* on page 86
3. *Performing additional recovery steps, if required* on page 91

Related information

NetApp Interoperability Matrix Tool

Deploying 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 "Preparing the hosts" in the StorageGRID installation instructions for your Linux operating system.

This procedure includes steps to accomplish the following tasks:

1. Install Linux.
2. Configure the host network.
3. Configure host storage.
4. Install Docker.
5. Install the StorageGRID host service.

   Attention: Stop after you complete the "Install StorageGRID host service" task in the installation instructions. Do not 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 disk 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 host for the StorageGRID node will use storage assigned from a NetApp All Flash FAS (AFF) system, you must confirm that the FlexVol does not have a tiering policy enabled.

  Attention: Never assign storage for a StorageGRID node from a FlexVol with an active tiering policy. Node outages might occur if storage used by the StorageGRID node is tiered to a capacity tier.

Related information

Red Hat Enterprise Linux or CentOS installation
Ubuntu or Debian installation

Restoring grid nodes to the host

To restore a failed grid node to a new Linux host, you restore the node configuration file using the appropriate commands.

When doing a fresh 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.

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.

Steps

1. Restoring and validating grid nodes on page 87
2. Starting the StorageGRID host service on page 90
3. Recovering nodes that fail to start normally on page 90

Restoring and validating 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 metadata 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 recreate 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 on the location of the `/var/local` volume for a node.

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 are not 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`

Note: Any nodes imported with the `--force` flag will require additional recovery steps before they can rejoin the grid, as described in "Performing additional recovery steps, if required."

3. For grid nodes that do not have a `/var/local` volume, recreate the node's configuration file to restore it to the host.

Follow the guidelines in "Creating node configuration files" in the installation instructions.

Attention: 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).

Attention: 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 recreating the configuration file for a node, be sure to follow all of the guidelines in "Fixing 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.

Related concepts

Fixing missing network interface errors on page 89
Fixing missing block device errors on page 89
Performing additional recovery steps, if required on page 91

Related information

Red Hat Enterprise Linux or CentOS installation
Ubuntu or Debian installation

Fixing 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>: <SOME>_NETWORK_TARGET = <host-interface-name>
       <node-name>: Interface '<host-interface-name>' does not exist
```

This error means that the `/etc/storagegrid/nodes/<node-name>.conf` file maps the StorageGRID `<SOME>` 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 "Deploying 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 `<SOME>_NETWORK_TARGET` to match an existing host interface. Make sure the interface provides access to the appropriate physical network port or VLAN.

Related concepts

Deploying new Linux hosts on page 86

Fixing 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 "Deploying 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 recreate 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 from the tables in the "Storage requirements" section of the installation instructions for your Linux operating system. Review the recommendations in "Configuring host storage" before proceeding with the block device replacement.

Attention: 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.

Caution: Run the following command only for new block storage devices. Do not 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
```

Related concepts

Deploying new Linux hosts on page 86

Related information

Red Hat Enterprise Linux or CentOS installation
Ubuntu or Debian installation

Starting the StorageGRID host service

To start your StorageGRID nodes, and ensure they restart after a host reboot, you must enable and start the StorageGRID host service.

Steps

1. Run the following commands on each host:

```
sudo systemctl enable storagegrid
sudo systemctl start storagegrid
```

2. Run the following command to ensure the deployment is proceeding:

```
sudo storagegrid node status node-name
```

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

```
sudo storagegrid node start node-name
```

3. If you have previously enabled and started the StorageGRID host service (or if you are unsure if the service has been enabled and started), also run the following command:

```
sudo systemctl reload-or-restart storagegrid
```

Recovering nodes that fail to start normally

If a StorageGRID node does not rejoin the grid normally and does not show up as recoverable, it may be corrupted. You can force the node into recovery mode.

To force the node into recovery mode:

```
sudo storagegrid node force-recovery node-name
```

Tip: Before issuing this command, confirm that the node’s network configuration is correct; it may have failed to rejoin the grid due to incorrect network interface mappings or an incorrect Grid Network IP address or gateway.

Attention: After issuing the `storagegrid node force-recovery` *`node-name`* command, you must perform additional recovery steps for *`node-name`*.

Related concepts

Performing additional recovery steps, if required on page 91

Performing 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 may 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	*Configuring the replacement primary Admin Node* on page 62
Non-primary Admin Node	*Selecting Start Recovery to configure a non-primary Admin Node* on page 69
API Gateway Node	*Selecting Start Recovery to configure the API Gateway Node* on page 77
Archive Node	*Selecting Start Recovery to configure the Archive Node* on page 79
Storage Node (virtual): • 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`	*Selecting Start Recovery to configure the Storage Node* on page 50

Type of recovery	Next step
Storage Node (virtual): • You added a new block device. • 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.	*Recovering from storage volume failure where the system drive is intact* on page 38

If you need to recover more than one grid node on a host, recover any primary Admin Node first to prevent the recovery of other nodes from stalling as they try to contact the primary Admin Node.

Decommission procedure

You might need to permanently remove grid nodes from the StorageGRID system. To remove a grid node, you must decommission it.

You can decommission the following types of grid nodes:

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

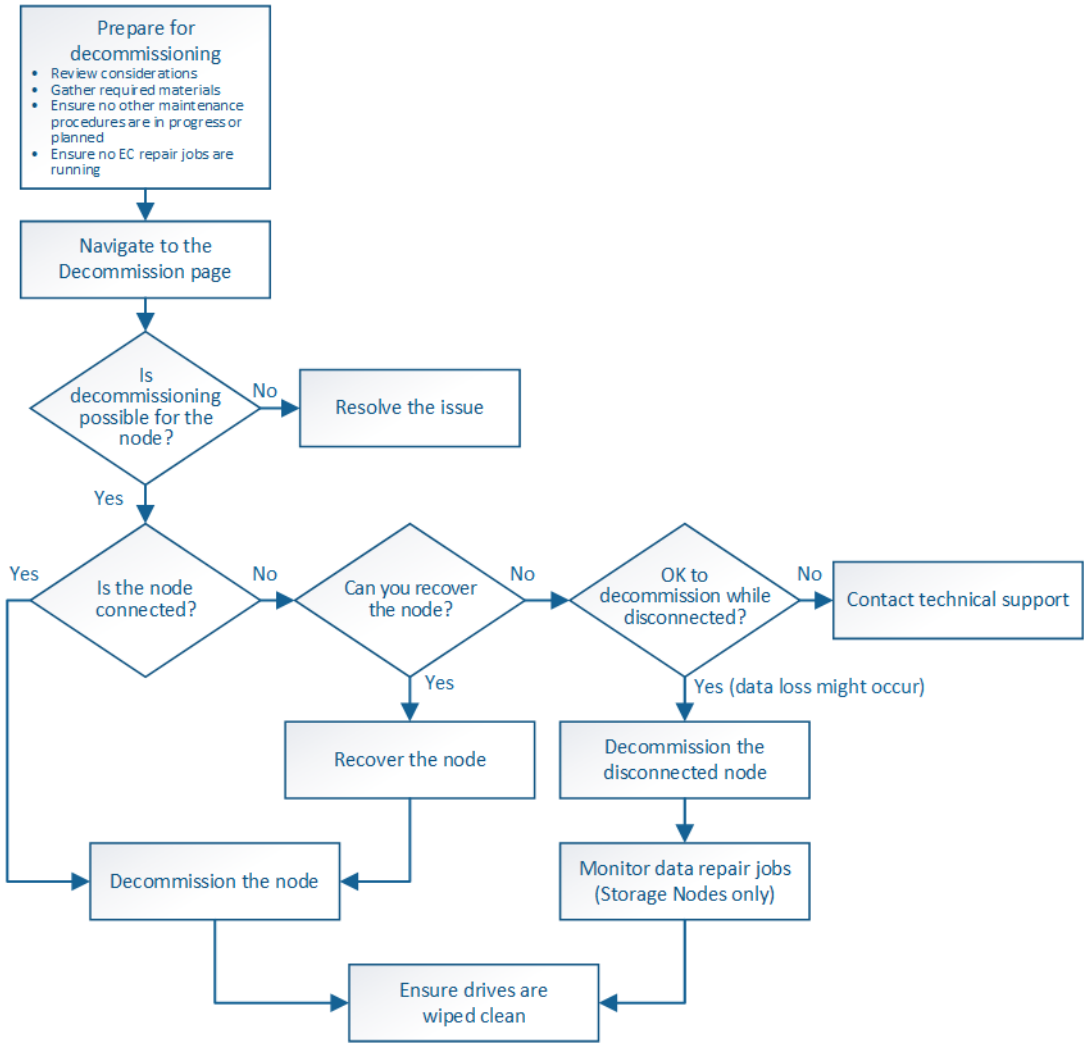
You cannot decommission the primary Admin Node or an Archive Node. You also cannot decommission all grid nodes at a site or a site itself.

In general, you should decommission grid nodes only while they are connected to the StorageGRID grid and all nodes are in normal health (have green icons on the **Nodes** pages and on the **Maintenance > Decommission** page). However, if required, you can decommission a grid node that is no longer connected to the grid (health is Unknown or Administratively Down). Before removing a disconnected node, make sure you understand the implications and restrictions of that process, as described in “Reviewing the considerations for decommissioning.”

Use this procedure when:

- 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 an API Gateway Node.
- You no longer require a non-primary Admin Node.
- Your grid includes a disconnected node that you cannot recover or bring back online.

The flowchart shows the high-level steps for decommissioning grid nodes.



Related concepts

Reviewing the considerations for decommissioning on page 95

Preparing for decommissioning

You must review the considerations for removing grid nodes, obtain required materials, and confirm no repair jobs are active for erasure-coded data.

Steps

1. Reviewing the considerations for decommissioning on page 95
2. Gathering required materials for decommissioning on page 99
3. Checking data repair jobs on page 100

Reviewing the considerations for decommissioning

You must understand the implications of removing grid nodes before you start this procedure. Removing a grid node impacts data security, the ILM policy, other maintenance procedures, and system operations.

Impact of decommissioning on data security

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.

Impact of decommissioning on the ILM policy

As required, you can safely change the ILM policy while decommissioning API Gateway Nodes and Admin Nodes. However, you should not make any changes to the ILM policy while a Storage Node is being decommissioned.

Impact of decommissioning on other maintenance procedures

The decommission procedure requires exclusive access to some system resources, so you must confirm that no other maintenance procedures are running. 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.

Impact of decommissioning on system operations

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. Therefore, 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.

Upon the successful decommissioning of a node, its services will be disabled and the node will be automatically shut down.

Impact of removing an Admin Node if SSO is enabled

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).

Considerations for decommissioning Storage Nodes

If you plan to decommission a Storage Node, you must understand how StorageGRID manages the object data and metadata on that node.

The following restrictions apply when decommissioning Storage Nodes:

- 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.

- You cannot run data repair operations on grid nodes when a decommission task is running.
- 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.

Related concepts

Understanding the ADC quorum on page 96
Reviewing the ILM policy and storage configuration on page 97
Decommissioning disconnected Storage Nodes on page 98
Consolidating Storage Nodes on page 98
Decommissioning multiple Storage Nodes on page 98

Related tasks

Restoring object data to a storage volume on page 55

Understanding 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 cannot 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

Expanding a StorageGRID system

Reviewing the 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. Then, ILM processes continue in a manner that satisfies the active ILM policy.

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 "Consolidating Storage Nodes," you should add new Storage Nodes before decommissioning old ones. If you add a larger replacement Storage Node after a 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?

  Note: An ILM policy that cannot 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 concepts

Consolidating Storage Nodes on page 98

Related information

Expanding a StorageGRID system

Decommissioning 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 on 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 cannot be brought online or recovered.
- 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 are not monitored by the decommission procedure. You must manually monitor these jobs and restart them as needed.
- If you attempt to decommission more than one disconnected Storage Node at a time, you increase the risk of unexpected results and data loss. The system might not be able to reconstruct data if too few copies of object data, metadata, or erasure-coded fragments remain available.

Consolidating 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 older, smaller capacity Storage Nodes. During the decommission procedure, objects are migrated from the old Storage Nodes to the new Storage Nodes.

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 older, smaller capacity Storage Nodes.

If a Storage Node is being decommissioned as part of a consolidation operation, you must add the new Storage Nodes to the StorageGRID system before decommissioning the existing Storage Node.

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

Expanding a StorageGRID system

Decommissioning 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.

Gathering required materials for decommissioning

Before you begin the decommission procedure, you must obtain the following information.

Item	Notes
Service laptop	The service laptop must have the following: • Network port • Supported browser • SSH client (for example, PuTTY)
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.
Description of StorageGRID topology after decommissioning	After decommissioning grid nodes, create an updated description of the system's new topology.

Related tasks

Downloading the Recovery Package on page 99

Related references

Web browser requirements on page 6

Downloading the Recovery Package

You must download an updated copy of the Recovery Package file before and after making grid topology changes to the StorageGRID system and before and after upgrading the software. The Recovery Package file allows you to restore the system if a failure occurs.

Before you begin

- You must be signed in to the Grid Manager using a supported browser.
- You must have the provisioning passphrase.
- You must have specific access permissions. For details, see information about controlling system access with administration user accounts and groups.

Steps

1. Select **Maintenance** > **Recovery Package**.
2. Enter the provisioning passphrase, and click **Start Download**.

   The download starts immediately.
3. When the download completes:
   a. Open the `.zip` file.
   b. Confirm it includes a `gpt-backup` directory and an inner `.zip` file.
   c. Extract the inner `.zip` file.
   d. Confirm you can open the `Passwords.txt` file.
4. Copy the downloaded Recovery Package file (`.zip`) to two safe, secure, and separate locations.

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

Related information

Administering StorageGRID

Checking data repair jobs

Before removing a grid node, you must confirm that no data repair jobs are active. If any repairs have failed, you must restart them 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. You must ensure that any erasure-coded fragments that were on the removed node have been restored successfully.

About this task

These steps only apply to systems that have erasure-coded objects.

Steps

1. From the service laptop, 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 do not 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.

```
root@DC1-ADM1:~ # repair-data show-ec-repair-status

 Repair ID   Scope                  Start Time  End Time  State  Est Bytes Affected Bytes Repaired  Retry Repair
=================================================================================================================
 949283   DC1-S-99-10(Volumes: 1,2) 2016-11-30T15:27:06.9  Success   17359            17359           No
 949292   DC1-S-99-10(Volumes: 1,2) 2016-11-30T15:37:06.9  Failure   17359            0               Yes
 949294   DC1-S-99-10(Volumes: 1,2) 2016-11-30T15:47:06.9  Failure   17359            0               Yes
 949299   DC1-S-99-10(Volumes: 1,2) 2016-11-30T15:57:06.9  Failure   17359            0               Yes
```

3. If the State for all repairs is `Success`, you do not need to restart any repair jobs.
4. If the State for any repair is `Failure`, 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 `Success`.

Performing the decommission

You can decommission and permanently remove one or more Storage Nodes, API Gateway Nodes, or non-primary Admin Nodes from the StorageGRID system.

About this task

You cannot decommission a grid node if doing so will leave the StorageGRID in an invalid state. The following rules are enforced:

- You cannot decommission the primary Admin Node.
- You cannot decommission Archive Nodes.
- You cannot decommission a Storage Node if its removal would affect the ADC quorum.
- You cannot decommission a Storage Node if it is required for the active ILM policy.
- You cannot decommission a connected grid node if your grid includes any disconnected grid nodes (nodes whose health is Unknown or Administratively Down). If your grid includes any disconnected nodes, you must bring them back online, recover them, or decommission them while they are disconnected.

Steps

1. Navigating to the Decommission page on page 102
2. Decommissioning disconnected grid nodes on page 104
3. Decommissioning connected grid nodes on page 108
4. Pausing and resuming the decommission process for Storage Nodes on page 111
5. Troubleshooting decommissioning on page 112

Related concepts

Understanding the ADC quorum on page 96
Reviewing the ILM policy and storage configuration on page 97

Navigating to the Decommission page

When you access the Decommission page in the Grid Manager, you can see at a glance which nodes can be decommissioned.

Before you begin

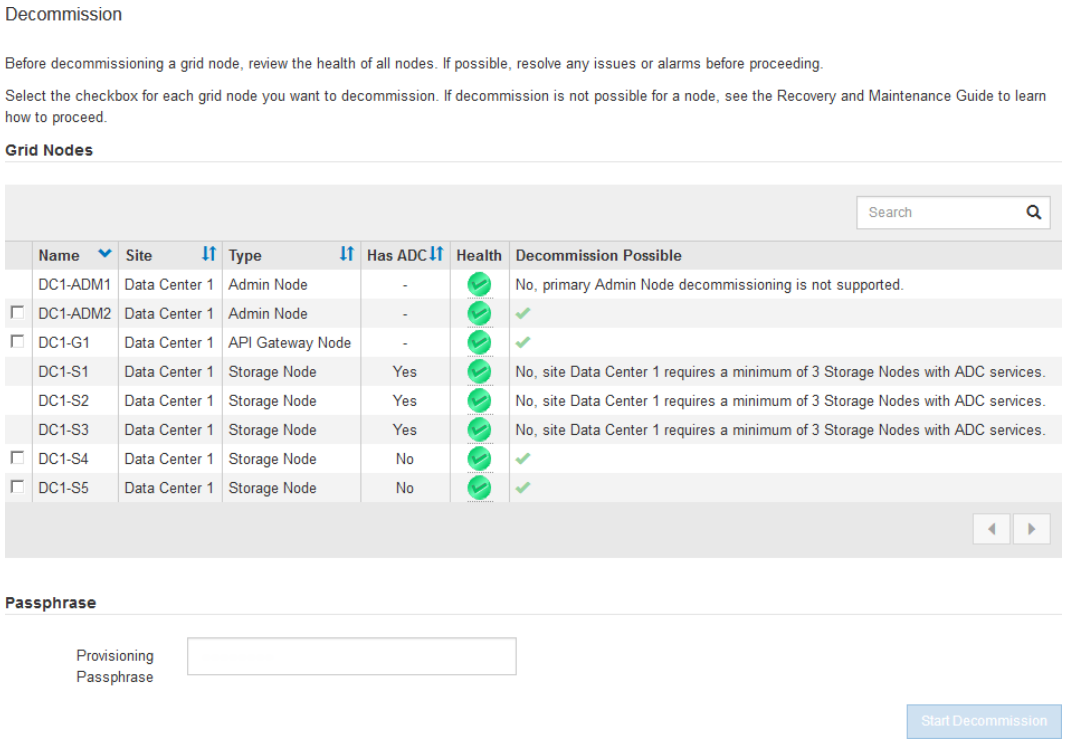
- You have reviewed and understand the requirements and considerations for decommissioning grid nodes.
- You have obtained all prerequisite items.
- No data report jobs are active. See “Checking data repair jobs.”
- You have confirmed that Storage Node recovery is not in progress anywhere in the grid. 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 decommission procedure is running, unless the decommission procedure is paused.
- You must be signed in to the Grid Manager using a supported browser.
- You must have Maintenance or Root Access permissions.
- You must have the provisioning passphrase.

Steps

1. Select **Maintenance > Decommission**.

   The Decommission page appears, listing all grid nodes in a table. From this page, you can:

   - Determine which grid nodes can be decommissioned currently.
   - See the health of all grid nodes (assuming there are no disconnected nodes)
   - Sort the list in ascending or descending order by **Name**, **Site**, **Type**, or **Has ADC**.
   - Enter search terms to quickly find a particular node or group of nodes.
   - Use the page controls to move forward and backward through the list.



2. 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-most column includes a check box. If a grid cannot be decommissioned, this column describes the issue. If there is more than one reason a node cannot be decommissioned, the most critical reason is shown.

Decommission Possible message	**Description**	**Steps to resolve**
`No, node type decommissioning is not supported.`	StorageGRID does not allow you to decommission the primary Admin Node or an Archive Node.	None.
`No, at least one grid node is disconnected.`	You cannot 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 step *3*.

Decommission Possible message	Description	Steps to resolve
`No, site x requires a minimum of n Storage Nodes with ADC services.`	**Storage Nodes only.** You cannot decommission a Storage Node if not enough 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 "Understanding the ADC quorum".
`No, the ILM policy requires a minimum of n Storage Nodes for erasure coding.`	**Storage Nodes only.** If the active ILM policy uses rules that specify objects be erasure coded, you cannot decommission a Storage Node if not enough nodes would remain to support the erasure coding scheme selected in the Erasure Coded profile.	• Create a new active ILM policy. Only use rules that make replicated copies or select an Erasure Coding profile that uses a different erasure coding scheme. See the ILM section of *Administering StorageGRID*. • Perform an expansion. Add new Storage Nodes so that the current Erasure Coding profile can continue to be used. See "Understanding the ADC quorum" and *Expanding StorageGRID*.

3. If decommissioning is possible for the node, determine which procedure you need to perform:

If your grid includes...	Go to...
Any disconnected grid nodes	*Decommissioning disconnected grid nodes* on page 104
Only connected grid nodes	*Decommissioning connected grid nodes* on page 108

Related concepts

Understanding the ADC quorum on page 96

Related tasks

Checking data repair jobs on page 100

Related information

Administering StorageGRID
Expanding a StorageGRID system

Decommissioning disconnected grid nodes

You might need to decommission a grid node that is not currently connected to the grid (one whose Health is Unknown or Administratively Down).

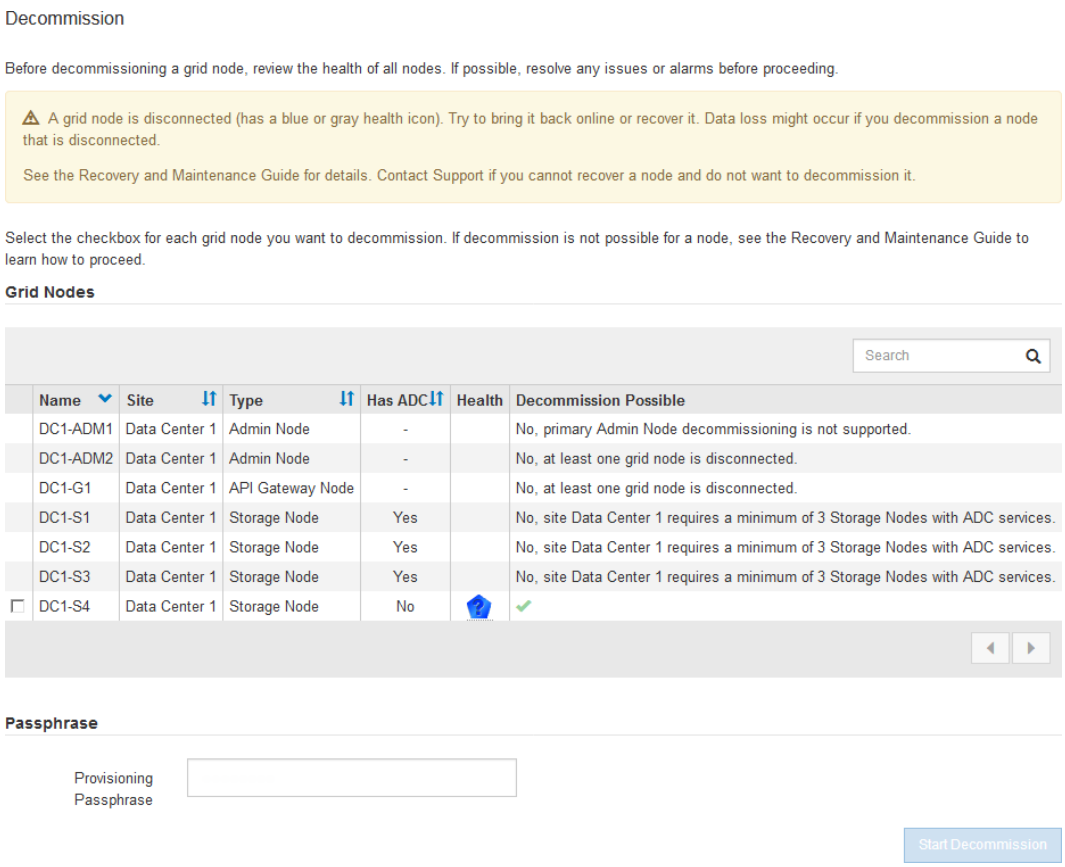
Before you begin

Before decommissioning a disconnected node, confirm the following:

- The **Decommission Possible** column for the disconnected node or nodes you want to decommission includes a green check mark.

About this task

You can identify disconnected nodes by looking for Unknown (blue) or Administratively Down (gray) icons in the **Health** column. No health icons are shown for nodes that are connected. In addition, the **Decommission Possible** shows `No, at least one grid node is disconnected` for the connected nodes. You cannot decommission a connected node if any nodes are disconnected.



Before decommissioning a disconnected node, note the following:

- You should never decommission a disconnected node unless you are sure it cannot be brought online or recovered.
- You can safely decommission an API Gateway Node while it is disconnected.
- If you decommission an Admin Node that is disconnected, you will lose the audit logs from that node; however, these logs should also exist on the primary Admin Node.
- When you decommission a Storage Node that is disconnected, 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 are not monitored by the decommission procedure. You must manually monitor these jobs and restart them as needed.
- If you decommission a Storage Node that is disconnected and that node 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.

- If you attempt to decommission more than one disconnected Storage Node at a time, you increase the risk of unexpected results and data loss. The system might not be able to reconstruct data if too few copies of object data, metadata, or EC fragments remain available.

Attention: Do not remove a grid node's virtual machine or other resources until instructed to do so in this procedure.

Steps

1. Attempt to bring any disconnected grid nodes back online or to recover them.

   See “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 check box for that node.

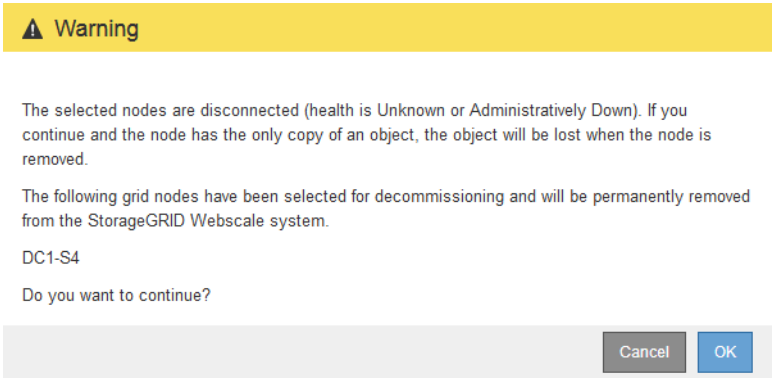
   Attention: Be very careful when selecting to decommission more than one disconnected grid node at a time, especially if you are selecting multiple disconnected Storage Nodes. If you have more than one disconnected Storage Node that you cannot recover, contact technical support to determine the best course of action.

3. Enter the provisioning passphrase.

   The **Start Decommission** button is enabled.

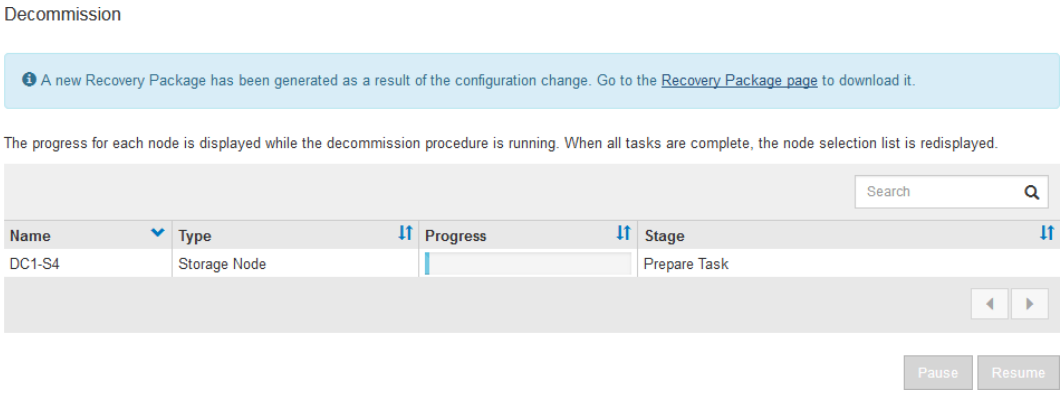
4. Click **Start Decommission**.

   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 to show the grid configuration change.



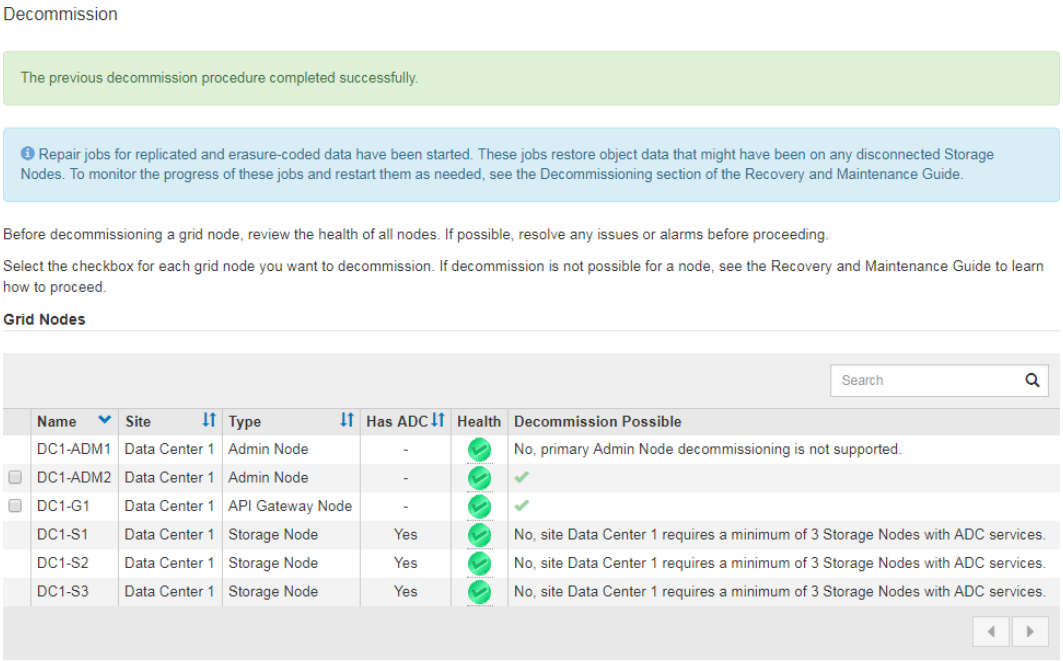
6. As soon as the new Recovery Package is available, click the link or select **Maintenance** > **Recovery Package** to access the Recovery Package page. Then, download the `.zip` file.

   See “Downloading the Recovery Package” for instructions.

   Note: Download the Recovery Package as soon as possible to ensure you can recover your grid if something goes wrong during the decommission procedure.

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. Remove any remaining virtual machines or other resources that are associated with the decommissioned node.

9. If you are decommissioning a Storage Node, monitor the status of the data repair jobs that are automatically started during the decommissioning process.

   a. Select **Support** > **Grid Topology**.

b. Select **StorageGRID deployment** at the top of the Grid Topology tree.

c. On the Overview tab, locate the ILM Activity section.

d. Use a combination of the following attributes to monitor repairs and to determine as well as possible if replicated repairs are complete:

- Use the **Repairs Attempted (XRPA)** attribute to track the progress of replicated repairs. This attribute increases each time the LDR service 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.

  Note: High-risk objects are objects that are at risk of being completely lost. This does not include objects that do not satisfy their ILM configuration.

- Use the **Scan Period - Estimated (XSCM)** 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 is at the Summary level and is the maximum of all node scan periods. You can query the **Scan Period - Estimated** attribute history at the Summary level to determine an appropriate timeframe for your grid.

e. Use the `repair-data show-ec-repair-status` command to track repairs of erasure coded data. Use the `repair-data start-ec-node-repair` command with the `--repair-id` option to restart a failed repair.

See “Checking data repair jobs” for instructions.

10. Continue to track the status of EC data repairs until all repair jobs have completed successfully.

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.

After you finish

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.

Related concepts

Recovery procedures on page 15

Related tasks

Downloading the Recovery Package on page 99
Checking data repair jobs on page 100

Decommissioning connected grid nodes

You can decommission and permanently remove grid nodes that are connected to the grid.

Before you begin

Before decommissioning a connected node, confirm the following:

- The grid only includes connected grid nodes.
- The **Decommission Possible** column for the node or nodes you want to decommission includes a green check mark.

- 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 or alarm:
  - (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 “Checking data repair jobs.”

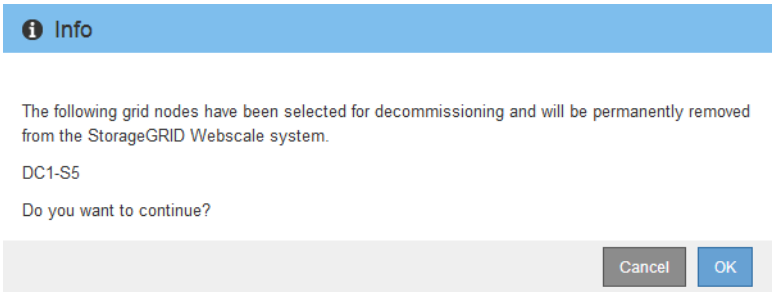
Attention: Do not remove a grid node's virtual machine or other resources until instructed to do so in this procedure.

Steps

1. From the **Decommission** page, select the check box for each grid node you want to decommission.
2. Enter the provisioning passphrase.

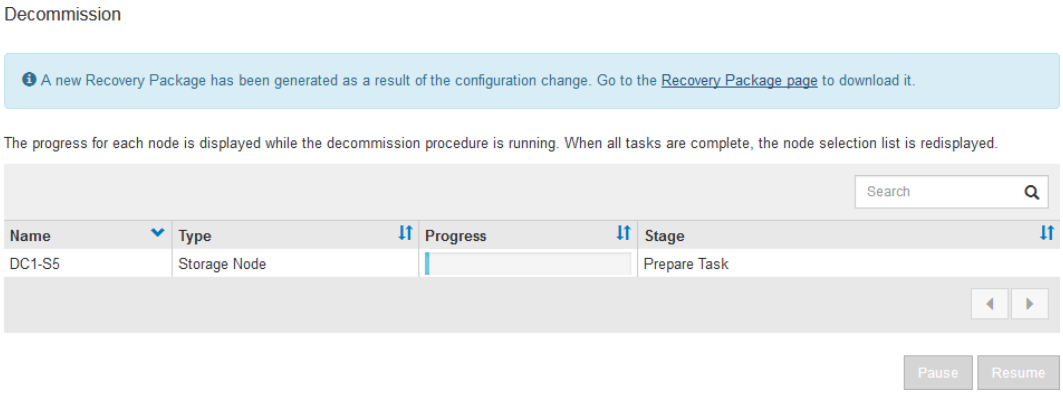
   The **Start Decommission** button is enabled.
3. Click **Start Decommission**.

   A confirmation dialog box appears.



4. Review the list of selected 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 to show the grid configuration change.



Note: Do not take a Storage Node offline after the decommission procedure has started. (That is, do not change the setting of **Storage State-Desired** from **LDR > Storage > Configuration**.) Changing the state might result in some content not being copied to other locations.

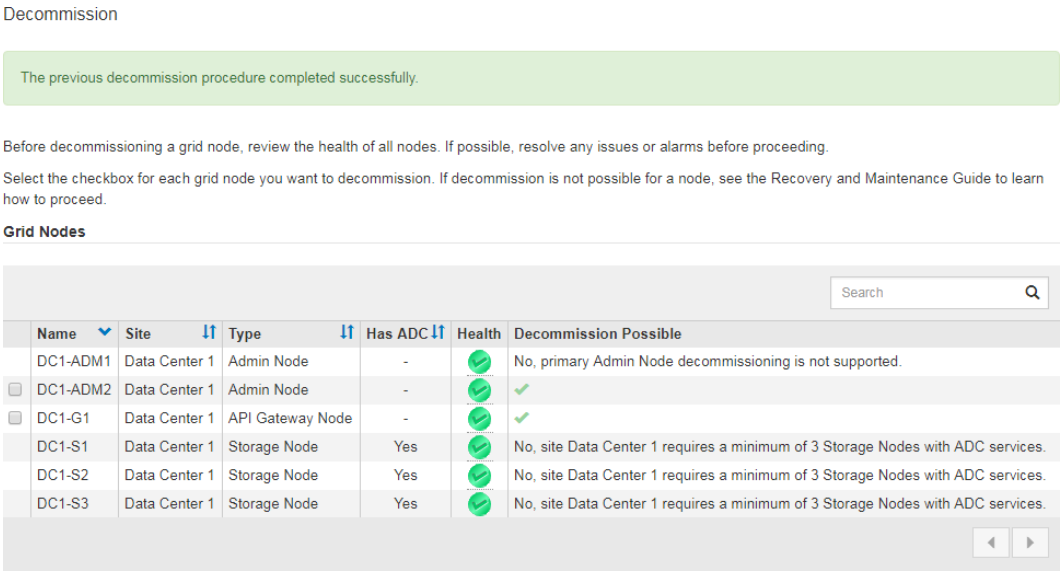
5. As soon as the new Recovery Package is available, click the link or select **Maintenance > Recovery Package** to access the Recovery Package page. Then, download the `.zip` file.

   See "Downloading the Recovery Package" for instructions.

   Note: 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 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:

   - **Linux**: Detach the volumes and delete the node configuration files you created during installation.
   - **VMware**: Use the vCenter "Delete from Disk" option to delete the virtual Storage Node.

- **StorageGRID appliance**: The appliance Storage 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 grid.

After you finish

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.

Related tasks

Checking data repair jobs on page 100
Downloading the Recovery Package on page 99

Related information

Red Hat Enterprise Linux or CentOS installation

Pausing and resuming the decommission process for Storage Nodes

If necessary, you can pause the decommission procedure for a Storage Node during certain stages. You must pause decommissioning on a Storage Node before you can start a second maintenance procedure. After the other procedure is finished, you can resume decommissioning.

Before you begin

- You must be signed in to the Grid Manager using a supported browser.
- You must have Maintenance or Root Access permissions.

Steps

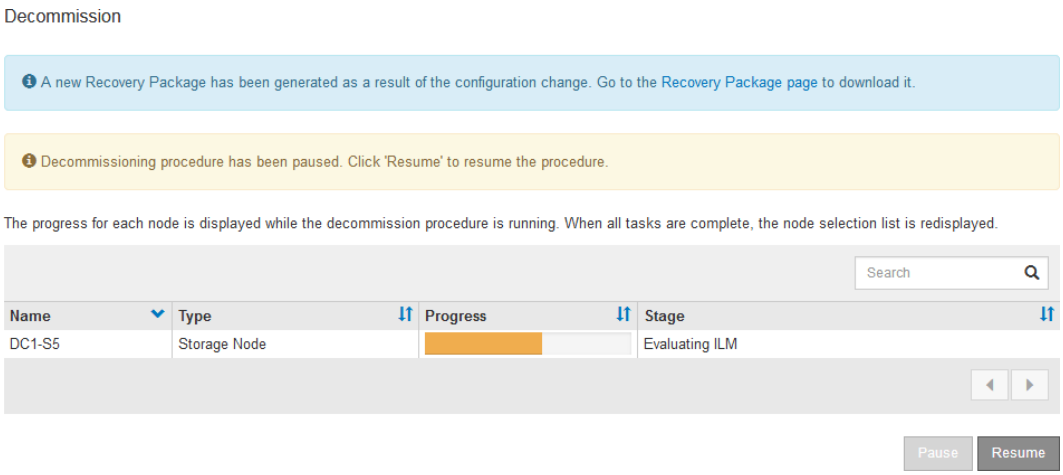
1. Select **Maintenance** > **Decommission**.

   The Decommission page appears. When the decommission procedure reaches either of the following stages, the **Pause** button is enabled.

   - Evaluating ILM
   - Decommissioning Erasure Coded data

2. Click **Pause** to halt the procedure.

   The current stage is paused, and the **Resume** button is enabled.



3. After the other procedure is finished, click **Resume** to proceed with the decommission.

Troubleshooting decommissioning

If the decommissioning process stops because of an error, you can take specific steps to troubleshoot the problem.

Before you begin

You must be signed in to the Grid Manager using a supported 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 > 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, the StorageGRID system's DDS services (hosted by Storage Nodes) must be online. This is a requirement of the ILM re-evaluation.
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.

Network maintenance procedures

You can configure the list of subnets on the Grid Network or update IP addresses, DNS servers, or NTP servers for your StorageGRID system.

Choices

- Updating subnets for the Grid Network on page 113
- Configuring IP addresses on page 114
- Configuring DNS servers on page 130
- Configuring NTP servers on page 132
- Restoring network connectivity for isolated nodes on page 133

Updating subnets for 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. 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

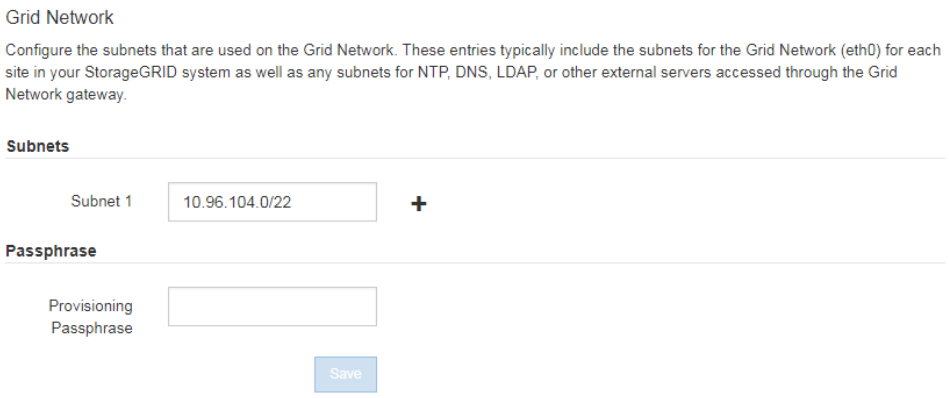
- You must be signed in to the Grid Manager using a supported browser.
- You must have a service laptop.
- You must have Maintenance or Root Access permissions.
- You must have the provisioning passphrase.
- You must have the network addresses, in CIDR notation, of the subnets you want to configure.

About this task

If you are performing an expansion activity that includes adding a new site, you must add the new Grid subnet before you start the expansion procedure.

Steps

1. Select **Maintenance** > **Grid Network**.



2. In the Subnets list, click the plus sign to add a new subnet in CIDR notation.

   For example, enter 10.96.104.0/22.

3. Enter the provisioning passphrase, and click **Save**.

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

Configuring IP addresses

You can perform network configuration by configuring IP addresses for grid nodes using the Change IP tool.

About this task

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.

Note: 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.

Attention: The IP change procedure can be a disruptive procedure. Parts of the grid might be unavailable until the new configuration is applied.

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 cannot 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.

Choices

- Modifying a node's network configuration on page 115
- Adding to or changing subnet lists on the Admin Network on page 119
- Adding to or changing subnet lists on the Grid Network on page 122
- Linux: Adding interfaces to an existing node on page 126
- Changing IP addresses for all nodes on the Grid Network on page 127

Modifying a node's network configuration

You can modify the network configuration of one or more nodes using the Change IP tool. You can modify the configuration of the Grid Network, or add, modify, or remove the Admin or Client Networks.

Before you begin

- You must have a service laptop.
- You must have the `Passwords.txt` file.

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.

Appliances: On StorageGRID appliances, if the Client or Admin Network was not configured in the StorageGRID Appliance Installer during the initial installation, the network cannot be added by using only the Change IP tool. First, you must place the appliance's compute controller 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 in the installation and maintenance instructions for your appliance model.

About this task

You can change the IP address, mask, or gateway 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/mask on that network to the node.

  If you are changing a node IP to an address for which no reservation has been obtained, the address used is a static IP.

- You can remove a node from a Client Network or from an Admin Network by deleting the IP/mask for the node on that network.

  Nodes cannot be removed from the Grid Network.

Attention: IP address swaps are not allowed. If you must exchange IP addresses between grid nodes, you must use a temporary intermediate IP address.

Attention: 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 recreate the node's relying party trust in Active Directory Federation Services (AD FS) with the new IP address. See the instructions for administering StorageGRID.

Note: 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.

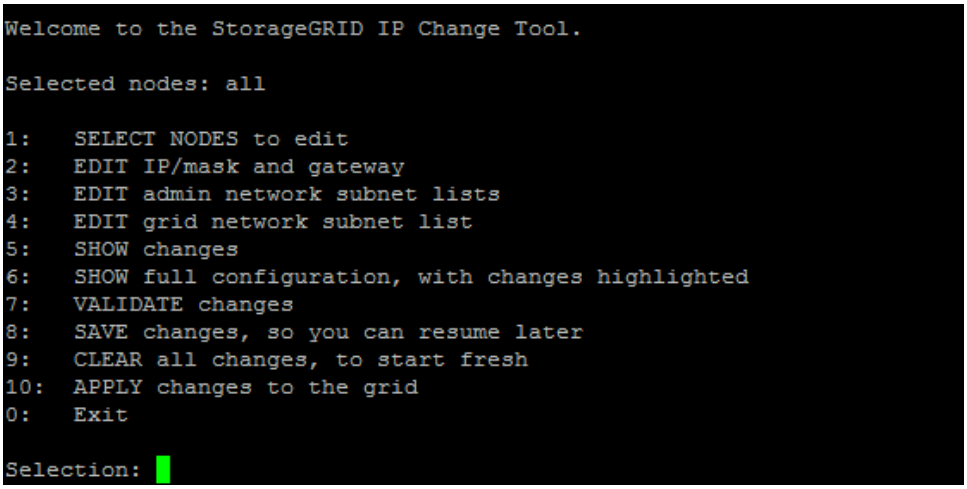
Steps

1. From the service laptop, 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 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 screen appears, with the **Selected nodes** field updated to reflect your choice. All subsequent actions are performed only on those nodes.

5. On the main menu screen, select option **2** to edit IP/mask and gateway 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 or Gateway), and current value.

      Editing the IP address, prefix length, or gateway 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' cannot be edited.

   b. To set a new value, enter it in the format shown for the current value.

   c. To leave the current value unchanged, press **Enter**.

   d. 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**.

   e. 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 Lab 1
================================================================================
DK-10-224-5-22-S1   Grid    IP      [     172.16.5.22/21 ]: 172.16.5.16/21
DK-10-224-5-20-G1   Admin   IP      [     10.224.5.20/21 ]: 0.0.0.0/0
DK-10-224-5-21-ADM1 Admin   IP      [     10.224.5.21/21 ]: 0.0.0.0/0
DK-10-224-5-22-S1   Admin   IP      [     10.224.5.22/21 ]: 0.0.0.0/0
DK-10-224-5-23-S2   Admin   IP      [     10.224.5.23/21 ]: 0.0.0.0/0
DK-10-224-5-24-S3   Admin   IP      [     10.224.5.24/21 ]: 0.0.0.0/0
DK-10-224-5-20-G1   Admin   Gateway [         10.224.0.1 ]: 0.0.0.0
DK-10-224-5-21-ADM1 Admin   Gateway [         10.224.0.1 ]: 0.0.0.0
DK-10-224-5-22-S1   Admin   Gateway [         10.224.0.1 ]: 0.0.0.0
DK-10-224-5-23-S2   Admin   Gateway [         10.224.0.1 ]: 0.0.0.0
DK-10-224-5-24-S3   Admin   Gateway [         10.224.0.1 ]: 0.0.0.0
DK-10-224-5-20-G1   Admin   Subnets                         del 10.0.0.0/8
                                                            del 172.19.0.0/16
                                                            del 172.21.0.0/16
                                                            del 172.20.0.0/16
DK-10-224-5-21-ADM1 Admin   Subnets                         del 10.0.0.0/8
                                                            del 172.19.0.0/16
                                                            del 172.21.0.0/16
DK-10-224-5-22-S1   Admin   Subnets                         del 10.0.0.0/8
                                                            del 172.19.0.0/16
                                                            del 172.21.0.0/16
DK-10-224-5-23-S2   Admin   Subnets                         del 10.0.0.0/8
                                                            del 172.19.0.0/16
                                                            del 172.21.0.0/16
DK-10-224-5-24-S3   Admin   Subnets                         del 10.0.0.0/8
                                                            del 172.19.0.0/16
                                                            del 172.21.0.0/16
                                    [       (None)       ]
Press Enter to continue
```

- **6**: Shows edits in output that displays the full configuration. Changes are highlighted in green (additions) or red (deletions).

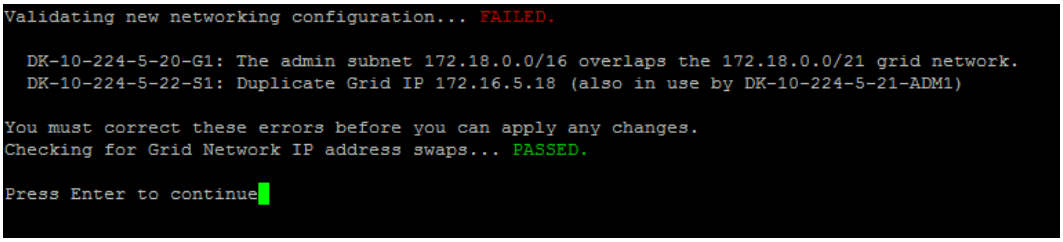
  Note: Certain command line interfaces may show additions and deletions using strikethrough formatting. Proper display depends on your terminal client supporting the necessary VT100 escape sequences.

7. Select option **7** to validate all changes.

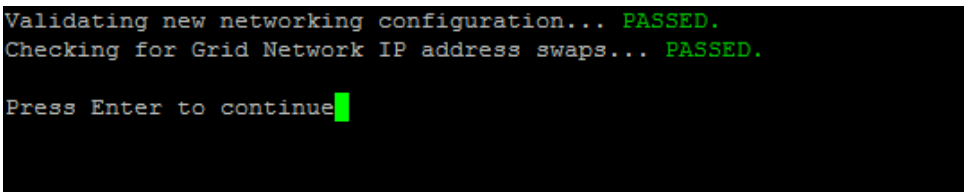
   This validation ensures that the rules for the Grid, Admin, and Client Networks, such as not using overlapping subnets, are not violated.

   Example

   In this example, validation returned errors.



   In this example, validation passed.



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.

     Attention: If you use the **stage** option, you must restart the node as soon as possible after staging to minimize disruptions.

- **cancel**: Do not 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 "Adding interfaces to an existing node."

    b. Restart the affected nodes.

11. Select **0** to exit the Change IP tool after your changes are complete.

12. Download the new Recovery Package from the Grid Manager. Select **Maintenance > Recovery Package** and enter the provisioning passphrase.

Related tasks

Linux: Adding interfaces to an existing node on page 126
Configuring IP addresses on page 114

Related information

Red Hat Enterprise Linux or CentOS installation
Ubuntu or Debian installation
SG6000 appliance installation and maintenance
SG5700 appliance installation and maintenance
Administering StorageGRID

Adding to or changing subnet lists on the 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 must have a service laptop.
- You must have the `Passwords.txt` file.

About this task

You can add, delete, or change subnets to all nodes on the Admin Network Subnet List.

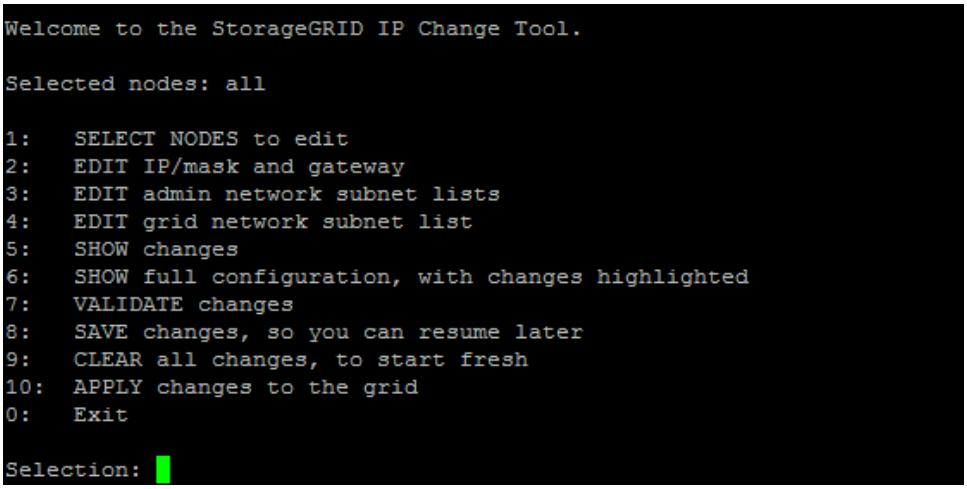
Steps

1. From the service laptop, 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 screen 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`**

     Note: 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`

     Tip: 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 sample 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

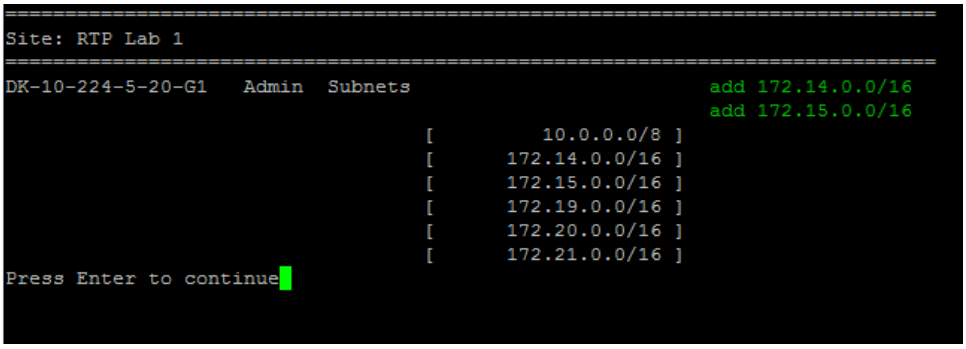
DK-10-224-5-20-G1
  10.0.0.0/8
  172.19.0.0/16
  172.21.0.0/16
  172.20.0.0/16

[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.

   Note: If you selected any of the "all" node selection modes in step 2, you must 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 sample output below:



- 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.

9. Select option **7** to validate all staged changes.
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. After successfully applying the changes, download the new provisioning repository by accessing the new package on the **Maintenance > Recovery Package** page.

Related tasks

Configuring IP addresses on page 114

Adding to or changing subnet lists on the Grid Network

You can use the Change IP tool to add or change subnets on the Grid Network.

Before you begin

- You must have a service laptop.
- You must have the `Passwords.txt` file.

About this task

You can add, delete, or change subnets in the Grid Network Subnet List. Changes will affect routing on all nodes in the grid.

Note: 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.

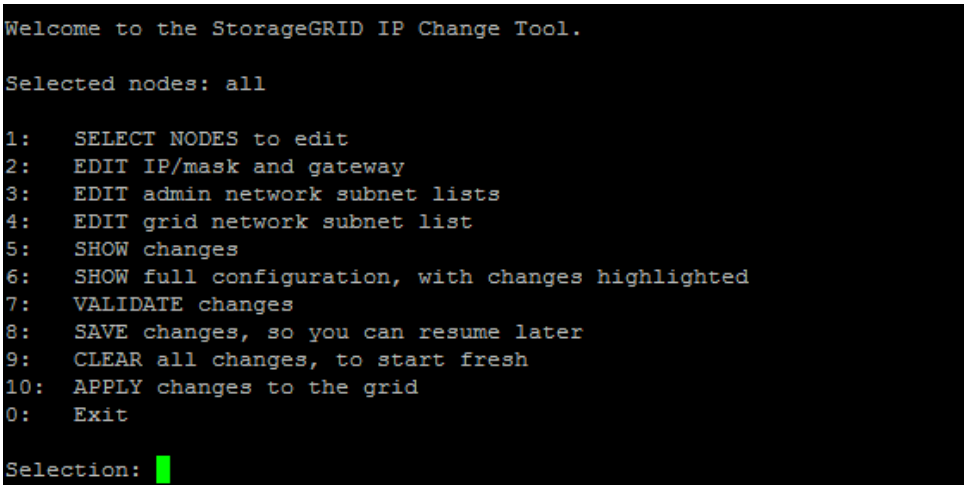
Steps

1. From the service laptop, 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 screen appears.



4. On the main menu, select the option to edit subnets for the Grid Network (option **4**).
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`**

     Note: 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
```

Tip: 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 sample 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 sample 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).

     Note: 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. (Optional) 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.

      Attention: If you use the **stage** option, you must restart the node as soon as possible after staging to minimize disruptions.

    - **cancel**: Do not 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 abort 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. After successfully applying the changes, download the new provisioning repository by accessing the new package on the **Maintenance > Recovery Package** page.

Related tasks

Configuring IP addresses on page 114

Linux: Adding interfaces to an existing node

You must add an interface to an existing node only if you did not configure ADMIN_NETWORK_TARGET or CLIENT_NETWORK_TARGET in the node configuration file on the Linux host during installation.

About this task

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 the StorageGRID installation instructions for your Linux operating system for more information on the node configuration file.

Steps

1. Log in to the Linux server hosting the node that needs the new network assignment.
2. Edit the node configuration file at `/etc/storagegrid/nodes/<node-name>.conf`.

   Attention: Do not specify any other network parameters, or a validation error will result.

   a. Add the new network target.

   Example

   ```
   CLIENT_NETWORK_TARGET = bond0.3206
   ```

   b. Optional: Add a MAC address.

   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>`

Related tasks

Modifying a node's network configuration on page 115

Related information

Red Hat Enterprise Linux or CentOS installation
Ubuntu or Debian installation

Changing IP addresses for all nodes on the Grid Network

You can use the Change IP tool to change the IP addresses of all nodes on the Grid Network, across multiple sites.

Before you begin

- You must have a service laptop.
- You must have the `Passwords.txt` file.

About this task

To ensure that the grid starts up successfully, you must make all the changes at once.

If you want to change the IP addresses for the nodes at one site only, follow the steps for modifying a node's network configuration.

Note: This procedure applies to the Grid Network only. You cannot use this procedure to change IP addresses on the Admin or Client Networks.

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.

   #### Example

   - 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 recreate these relying party trusts in Active Directory Federation Services (AD FS) immediately after you change IP addresses. See the instructions for administering StorageGRID.
   - If necessary, add the new subnet for the new IP addresses.

2. From the service laptop, 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.

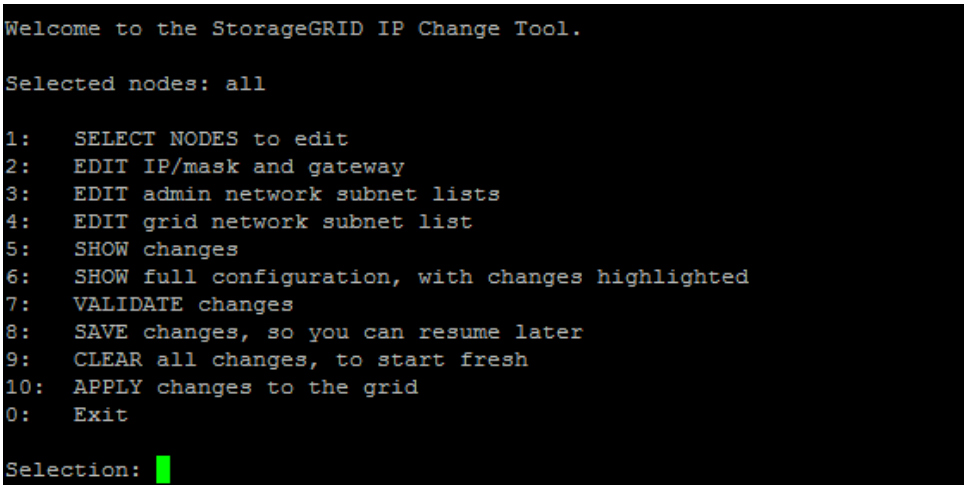
   When you are logged in as root, the prompt changes from `$` to `#`.

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`.



5. On the main menu, select **2** to edit IP/mask and gateway 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 or Gateway), and current values.

      Editing the IP address, prefix length, or gateway 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' cannot 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 Lab 1
================================================================================
DK-10-224-5-22-S1   Grid   IP      [     172.16.5.22/21 ]: 172.16.5.16/21
DK-10-224-5-20-G1   Admin  IP      [     10.224.5.20/21 ]: 0.0.0.0/0
DK-10-224-5-21-ADM1 Admin  IP      [     10.224.5.21/21 ]: 0.0.0.0/0
DK-10-224-5-22-S1   Admin  IP      [     10.224.5.22/21 ]: 0.0.0.0/0
DK-10-224-5-23-S2   Admin  IP      [     10.224.5.23/21 ]: 0.0.0.0/0
DK-10-224-5-24-S3   Admin  IP      [     10.224.5.24/21 ]: 0.0.0.0/0
DK-10-224-5-20-G1   Admin  Gateway [         10.224.0.1 ]: 0.0.0.0
DK-10-224-5-21-ADM1 Admin  Gateway [         10.224.0.1 ]: 0.0.0.0
DK-10-224-5-22-S1   Admin  Gateway [         10.224.0.1 ]: 0.0.0.0
DK-10-224-5-23-S2   Admin  Gateway [         10.224.0.1 ]: 0.0.0.0
DK-10-224-5-24-S3   Admin  Gateway [         10.224.0.1 ]: 0.0.0.0
DK-10-224-5-20-G1   Admin  Subnets                       del 10.0.0.0/8
                                                         del 172.19.0.0/16
                                                         del 172.21.0.0/16
                                                         del 172.20.0.0/16
DK-10-224-5-21-ADM1 Admin  Subnets                       del 10.0.0.0/8
                                                         del 172.19.0.0/16
                                                         del 172.21.0.0/16
DK-10-224-5-22-S1   Admin  Subnets                       del 10.0.0.0/8
                                                         del 172.19.0.0/16
                                                         del 172.21.0.0/16
DK-10-224-5-23-S2   Admin  Subnets                       del 10.0.0.0/8
                                                         del 172.19.0.0/16
                                                         del 172.21.0.0/16
DK-10-224-5-24-S3   Admin  Subnets                       del 10.0.0.0/8
                                                         del 172.19.0.0/16
                                                         del 172.21.0.0/16
                                  [       (None)       ]
Press Enter to continue
```

- **6**: Shows edits in output that displays the full configuration. Changes are highlighted in green (additions) or red (deletions).

  Note: Certain command line interfaces may 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, are not violated.

   Example

   In this example, validation returned errors.

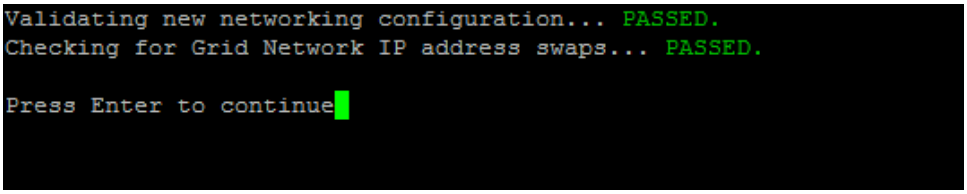
```
Validating new networking configuration... FAILED.

  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.



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.

   Attention: You must select **stage**. Do not 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.

    Attention: 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. Select **Maintenance > Recovery Package** and enter the provisioning passphrase.

Related tasks

Modifying a node's network configuration on page 115
Adding to or changing subnet lists on the Grid Network on page 122
Shutting down a grid node on page 153

Related information

Administering StorageGRID

Configuring DNS servers

You can add, remove, and update domain name system (DNS) servers, so that you can use fully qualified domain name (FQDN) hostnames rather than IP addresses.

Before you begin

- You must be signed in to the Grid Manager using a supported browser.
- You must have a service laptop.
- You must have Maintenance or Root Access permissions.
- You must have the IP addresses of the DNS servers to configure.

About this task

Specifying DNS server information allows you to use fully qualified domain name (FQDN) hostnames rather than IP addresses for email notifications and AutoSupport. Specifying at least two DNS servers is recommended.

Attention: Provide between two to six IP addresses for DNS servers. It is recommended to select DNS servers that each site can access locally in the event of network islanding. This is to ensure an islanded site continues to have access to the DNS service. After configuring the grid-wide DNS server list, you may further customize the DNS server list for each node. For more information, see "Modifying the DNS configuration for a single node.".

If the DNS server information is omitted or incorrectly configured, a DNST alarm is triggered on each grid node's SSM service. The alarm clears when DNS is configured correctly and the new server information has reached all grid nodes.

Steps

1. Select **Maintenance** > **DNS Servers**.
2. In the **Servers** section, add update, or remove DNS server entries, as necessary.

   The best practice is to specify at least two DNS servers per site. You can specify up to six DNS servers.
3. Click **Save**.

Modifying the DNS configuration for a single grid node

Rather than configure the Domain Name System (DNS) globally for the entire deployment, you can run a script to configure DNS differently for each grid node.

About this task

In general, you should use the **Maintenance** > **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. From the service laptop, log in to the primary Admin Node or the Archive 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. 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]
```

3. Add the IP address of a server that provides domain name service for your network: **add nameserver *IP_address***
4. Repeat the command to add name servers.
5. Follow instructions as prompted for other commands.
6. Save your changes and exit the application: **save**
7. Close the command shell on the server: **exit**
8. Repeat these steps for each grid node.

Configuring 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 must be signed in to the Grid Manager using a supported browser.
- You must have a service laptop.
- You must have Maintenance or Root Access permissions.
- You must have the provisioning passphrase.
- You must have the IP addresses of the NTP servers to configure.

About this task

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.

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

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.

> **Note:** When specifying the external NTP source for a production-level StorageGRID installation, do not use the Windows Time (W32Time) service on a version of Windows earlier than Windows Server 2016. The time service on earlier versions of Windows is not sufficiently accurate and is not supported by Microsoft for use in high-accuracy environments, such as StorageGRID.
>
> *Support boundary to configure the Windows Time service for high-accuracy environments*

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.

Steps

1. Select **Maintenance** > **NTP Servers**.
2. In the **Servers** section, add update, or remove NTP server entries, as necessary.

   You should include at least 4 NTP servers, and you can specify up to 6 servers.
3. In the **Provisioning Passphrase** text box, enter the provisioning passphrase for your StorageGRID system and click **Save**.

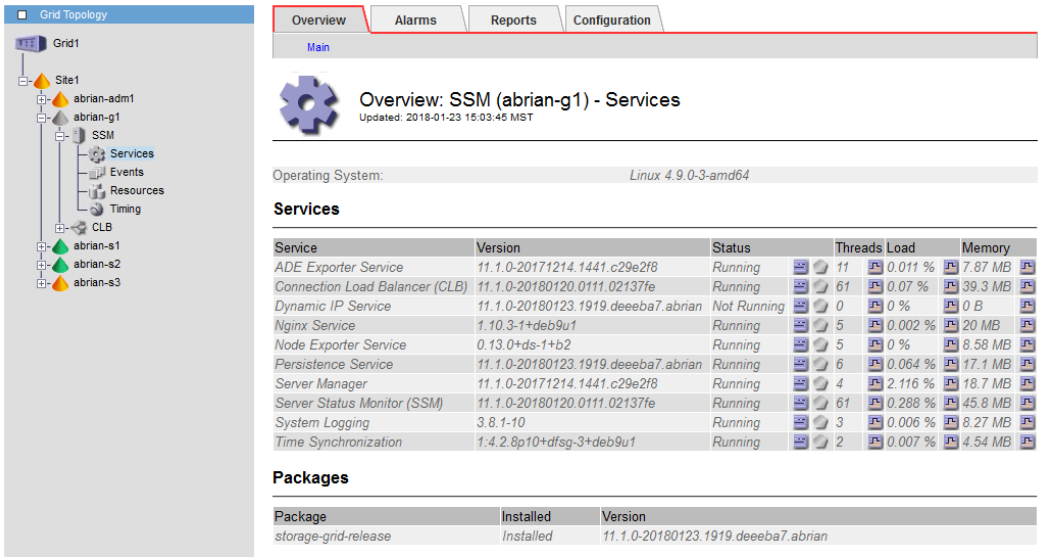
   The status of the procedure is displayed at the top of the page. The page is disabled until the configuration updates are complete.

Restoring 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, 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 node or on one node in a group 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.

Steps

1. Access the node and check `/var/local/log/dynip.log` for isolation messages.

   Example

   ```
    [2018-01-09T19:11:00.545] UpdateQueue - WARNING -- Possible
   isolation, no contact with other nodes.
   If this warning persists, manual action may 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.

   Example

   ```
   # /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.

   Attention: 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

Some maintenance procedures are specific to Linux or VMware deployments of StorageGRID, or are specific to other components of the StorageGRID solution.

Linux: Migrating a grid node to a new host

You can migrate StorageGRID nodes from one Linux host to another to perform host maintenance (such as OS patching and reboot) without impacting the functionality or availability of your grid.

You migrate one or more nodes from one Linux host (the “source host”) to another Linux host (the “target host”). The target host must have previously been prepared for StorageGRID use.

Attention: You can use this procedure only if you planned your StorageGRID deployment to include migration support.

To migrate a grid node to a new host, both of the following conditions must be true:

- Shared storage is used for all per-node storage volumes
- Network interfaces have consistent names across hosts

For more information, see “Node migration requirements” in the StorageGRID installation instructions for your Linux operating system.

Related concepts

Deploying new Linux hosts on page 86

Related information

Red Hat Enterprise Linux or CentOS installation
Ubuntu or Debian installation

Linux: Exporting the node from the source host

Shut down the grid node and export it from the source Linux host.

About this task

Run the following command on the source Linux host.

Steps

1. Obtain the status of all nodes currently running on the source host.

```
sudo storagegrid node status all
```

```
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
```

```
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
```

```
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. Take note of the import command suggested in the output of the export command.

   You will run this command on the target host in the next step.

Linux: Importing the node on the target host

After exporting the node from the source host, you import and validate the node on the target Linux host. Validation confirms that the node has access to the same block storage and network interface devices as it had on the source host.

About this task

Run the following command on the target Linux host.

Steps

1. Import the node on the target host.

```
sudo storagegrid node import /dev/mapper/sgws-dc1-s3-var-local
```

```
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
```

```
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.

Related information

Red Hat Enterprise Linux or CentOS installation
Ubuntu or Debian installation

Linux: Starting the migrated node

After you validate the migrated node, you start the node by running a command on the target Linux host.

Steps

1. Start the node on the new host.

```
sudo storagegrid node start DC1-S3
```

```
Starting node DC1-S3
```

2. In the Grid Manager, verify that the status of the node is green with no alarms raised against it.

   Attention: 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, do not migrate any additional nodes so that you will not have more than one node out of service.

   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 cannot 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.

Taking 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 must be signed in to the Grid Manager using a supported browser.

Steps

1. Select **Support** > **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, go to **Support** > **Grid Topology** > ***Archive Node*** > **ARC** > **Retrieve** > **Request Failures**.

A retrieval can fail if the object is temporarily unavailable. In this case, subsequent retrieval requests should eventually succeed.

If the object is permanently unavailable, you must perform the “determining if objects are unavailable” procedure to identify the object, and then manually cancel the Archive Node’s request.

If the StorageGRID system is configured to use an ILM rule with only one active content placement instruction, copies of an object are not made.

If an object is lost, it cannot be recovered; however, you must still determine if object is permanently unavailable to “clean up” the StorageGRID system, to cancel the Archive Node’s request, and purge metadata for the lost object.

Related tasks

Determining if objects are permanently unavailable on page 139

Related information

Administering StorageGRID

Determining if objects are permanently unavailable

You can determine if objects are permanently unavailable by making a request using the TSM administrative console.

About this task

This example is provided for your information only; this procedure cannot help you identify all failure conditions that may result in unavailable objects or tape volumes. For information about TSM administration, see TSM Server documentation.

Steps

1. Identify the object or objects that could not be retrieved by the Archive Node:

   a. From the service laptop, 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. 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.

   c. Search the relevant audit log file for messages indicating that a retrieval failure occurred. For example, enter:

      `grep ARCE audit.log | less -n`

      When a retrieval fails, 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][AMID(FC32):ARCI][ATID(UI64):
      4560349751312520631]]
      ```

      See the instructions for understanding audit messages.

   d. Record the CBID of each object with 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
        Select **Support > Grid Topology**. Then, select ***`Archive Node`*** **> ARC > Target > Overview**.
        The file space name is the Archive Node's node ID.

      - 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 (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.

   e. Log out of the command shell:

      `exit`

2. Check the TSM server to see if the objects identified in step 1 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. (Select **Support** > **Grid Topology**. Then, select ***`Archive Node`*** > **ARC** > **Target** > **Configuration**.)

b. An example of a way that you might discover that an object is permanently unavailable is to 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)
>
```

Note that depending on the nature of the error, the CBID might not be recorded in the TSM activity log. It might be necessary 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” is the Archive Node volume ID (TSM “high level name”) followed by the object’s CBID (TSM “low level name). That is: /*`Archive Node volume ID`* /*`CBID`* or, in the first line of this example: /20081201/ C1D172940E6C7E12

Recall also that the “Filespace Name” 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 in step 3.

3. For each object that is permanently unavailable, cancel the retrieval request and inform the StorageGRID system that the object copy was lost:

Attention: 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 are not 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 cannot 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. Notify the StorageGRID system that an object copy has been lost and an additional copy must be made of the indicated object:

`/proc/CMSI/Object_Lost CBID node_ID`

where *CBID* is the identifier of the object that cannot be retrieved from the TSM server.

For Archive Nodes, you cannot use a range of CBIDs.

node_ID is the node ID of the Archive Node where the retrieval failed.

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. In a StorageGRID system configured to use an ILM rule with only one active content placement instruction, copies of an object are not made. If an object is lost, it cannot 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, it returns the message CLOC_LOST_ANS returned result 'SUCS'.

e. Exit the ADE Console:

`exit`

f. Log out of the Archive Node:

`exit`

4. 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

Administering StorageGRID

Understanding audit messages

VMware: Configuring a 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.

About this task

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

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.

Attention: You should access Server Manager only if technical support has directed you to do so.

Note: You must close the current command shell session and log out after you are finished with Server Manager. Enter: **`exit`**

Choices

- Viewing Server Manager status and version on page 144
- Viewing current status of all services on page 145
- Starting Server Manager and all services on page 146
- Restarting Server Manager and all services on page 147
- Stopping Server Manager and all services on page 147
- Viewing current status of a service on page 148
- Stopping a service on page 149
- Forcing a service to terminate on page 149
- Starting or restarting a service on page 150
- Rebooting a grid node on page 151
- Shutting down a grid node on page 153
- Powering down a host on page 154
- Powering off and on all nodes in the grid on page 155
- Using a DoNotStart file on page 158
- Troubleshooting Server Manager on page 159

Viewing 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 must have the `Passwords.txt` file.

Steps

1. From the service laptop, 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
   ```

   This status is the equivalent of the status shown in the header of the local console display.

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`**

Viewing 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 must have the `Passwords.txt` file.

Steps

1. From the service laptop, 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.

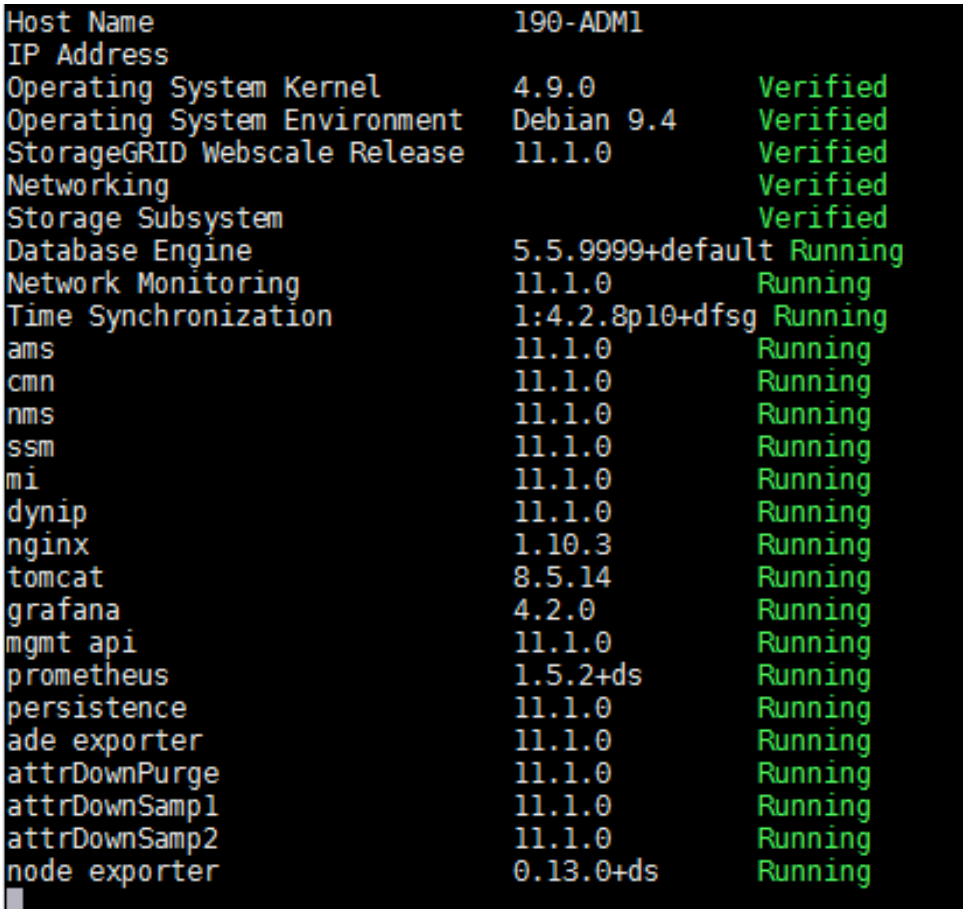
   c. Enter the following command to switch to root: `su -`

   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`**

Starting 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 must 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. From the service laptop, 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`**

Restarting Server Manager and all services

You might need to restart server manager and all services running on a grid node.

Before you begin

You must have the `Passwords.txt` file.

Steps

1. From the service laptop, 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.

   Note: 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`**

Stopping 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 must have the `Passwords.txt` file.

About this task

The only scenario that requires you to stop Server Manager while keeping the operating system running is when you need to integrate Server Manager to other services. If there is a requirement to stop the Server Manager for servicing of the hardware or reconfiguration of the server, the entire server should be halted.

Steps

1. From the service laptop, 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.
3. Log out of the command shell: **`exit`**

Viewing current status of a service

You can view the current status of a services running on a grid node at any time.

Before you begin

You must have the `Passwords.txt` file.

Steps

1. From the service laptop, 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`**

Stopping a 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 must 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. From the service laptop, 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
   ```

   Note: Services can take up to 11 minutes to stop.
3. Log out of the command shell: **`exit`**

Related tasks

Forcing a service to terminate on page 149

Forcing a service to terminate

If you need to stop a service immediately, you can use the force-stop command.

Before you begin

You must have the `Passwords.txt` file.

Steps

1. From the service laptop, 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`**

Starting or restarting a 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 must have the `Passwords.txt` file.

Steps

1. From the service laptop, 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
     ```

Note: 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`**

Rebooting a grid node

You can reboot a grid node from the Grid Manager or from the node's command shell.

About this task

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:

- StorageGRID immediately commits any newly ingested objects to two Storage Nodes on the same site. 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.

Choices

- Rebooting a grid node from the Grid Manager on page 151
- Rebooting a grid node from the command shell on page 152

Rebooting a grid node from the Grid Manager

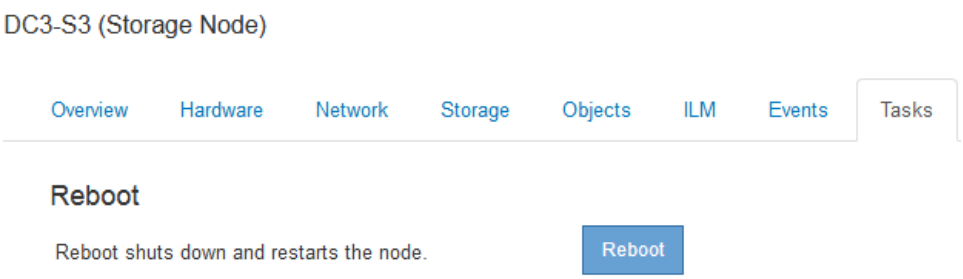
Rebooting a grid node from the Grid Manager issues the `reboot` command on the target node.

Before you begin

- You must be signed in to the Grid Manager using a supported browser.
- You must have Maintenance or Root Access permissions.
- You must have the provisioning passphrase.

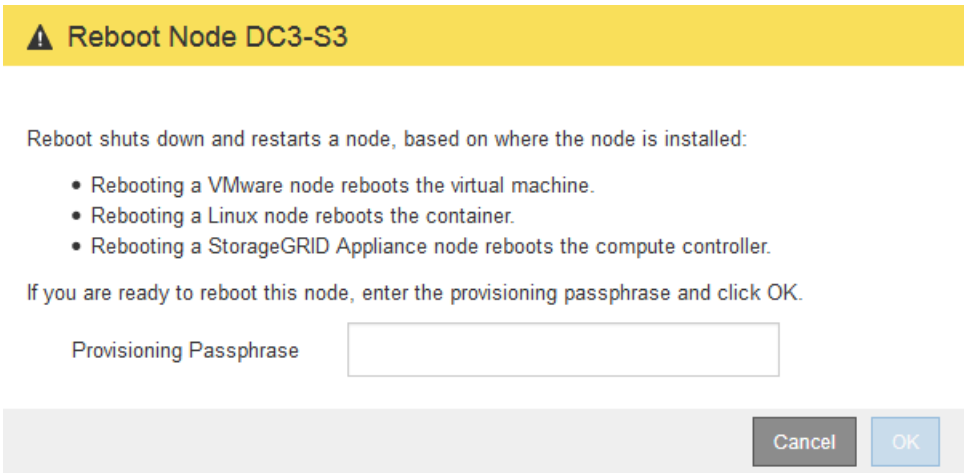
Steps

1. Select **Nodes**.
2. Select the grid node you want to reboot.
3. Select the **Tasks** tab.



4. Click **Reboot**.

   A confirmation dialog appears.



   Note: If you are rebooting the primary Admin Node, the confirmation dialog 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 click **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, the icon changes back to its original color.

Rebooting a grid node from the 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 into the grid node and run the Server Manager `reboot` command from the command shell.

Before you begin

You must have the `Passwords.txt` file.

Steps

1. From the service laptop, 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 some time 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`**

Shutting down a grid node

You can shut down a grid node from the node's command shell.

Before you begin

You must 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.
- Do not 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 Storage Nodes, note the following:
  - StorageGRID immediately commits any newly ingested objects to two Storage Nodes on the same site. If you want to shut down two or more Storage Nodes on a given site, you might not be able to access these objects for the duration of the shutdown.
  - To ensure you can access all objects when a Storage Node is shut down, stop ingesting objects at a site for approximately one hour before shutting down the node.

Steps

1. From the service laptop, 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 some time to shut down, and you might want to log in to the system remotely to monitor the shutdown process.
3. Power off the grid node by running the following command:

```
shutdown -h now
```

4. Log out of the command shell: **`exit`**

Powering down a host

Before you power down a host, you must stop services on all grid nodes on that host.

Steps

1. From the service laptop, 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 grid node: **`service servermanager stop`**

   Wait a minimum of 30 minutes for the command to complete.

   Note: 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.
3. Repeat steps 1 and 2 for each grid 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 Storage Node, 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.
6. If you are powering down an appliance Storage Node:
   - For the SG6000 appliance
     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 the blue power LED to turn off.
   - For the SG5700 appliance
     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.

7. Log out of the command shell: **`exit`**

Powering off and on all nodes in the grid

You might need to shut down your entire StorageGRID grid, 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.

About this task

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.

Stopping services and shutting 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, Docker containers, and StorageGRID appliances.

About this task

If possible, you should stop services on the grid nodes in this order:

- Stop services on API Gateway Nodes first.
- Stop services on the primary Admin Node last.

This approach allows you to use the primary Admin Node to monitor the status of the other grid nodes for as long as possible.

Note: If a single host includes more than one grid node, do not shut down the host until you have stopped all of the services on that host. If the host includes the primary Admin Node, shut down that host last.

Note: 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.

Linux: Migrating a grid node to a new host on page 136

Steps

1. Stop all client applications from accessing the grid.

2. Log in to each API 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 grid node: **`service servermanager stop`**

   Wait a minimum of 30 minutes for the command to complete.

4. Repeat steps *2* and *3* 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.

   Note: 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. Repeat steps *2* and *3* to stop the services on the primary Admin 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 Storage Nodes:

   - For the SG6000 appliance

     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 the blue power LED to turn off.

   - For the SG5700 appliance

     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.

Starting up the grid nodes

Follow this sequence to start up the grid nodes after a complete shutdown.

Before you begin

Caution: If the entire grid has been shut down for more than 15 days, you must contact technical support before starting up any grid nodes. Do not attempt the recovery procedures that rebuild Cassandra data. Doing so might result in data loss.

About this task

If possible, you should power on the grid nodes in this order:

- Apply power to Admin Nodes first.
- Apply power to API Gateway Nodes last.

Note: If a host includes multiple grid nodes, the nodes will come back online automatically when you power on the host.

Steps

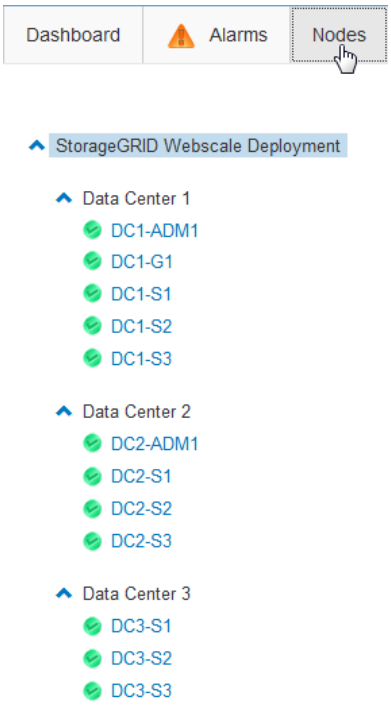
1. Power on the hosts for the primary Admin Node and any non-primary Admin Nodes.

   Note: 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 API Gateway Nodes.
4. Sign into the Grid Manager.
5. Click **Nodes**, and monitor the status of the grid nodes. Verify that all nodes return to “green” status.



Using 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.

Attention: 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 are not 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 cannot run if the NTP service is not running.

Adding a DoNotStart file for a 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 must have the `Passwords.txt` file.

Steps

1. From the service laptop, 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`**

Removing a DoNotStart file for a service

When you remove a DoNotStart file that is preventing a service from starting, you must start that service.

Before you begin

You must have the `Passwords.txt` file.

Steps

1. From the service laptop, 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`**

Troubleshooting Server Manager

Technical support might direct you to troubleshooting tasks to determine the source of Server Manager-related problems.

Accessing the Server Manager log file

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 must 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 goes down, fails to start, and enters an error state. Server Manager does not attempt another restart.

Steps

1. From the service laptop, 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
   ```

   Note: 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`**

Related tasks

Removing a DoNotStart file for a service on page 159

Copyright

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

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

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

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

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

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

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

Trademark

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

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

How to send comments about documentation and receive update notifications

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

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

doccomments@netapp.com

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

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

You can also contact us in the following ways:

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