



Grid node recovery procedures

StorageGRID

NetApp

October 03, 2025

This PDF was generated from <https://docs.netapp.com/us-en/storagegrid-115/maintain/warnings-and-considerations-for-grid-node-recovery.html> on October 03, 2025. Always check docs.netapp.com for the latest.

Table of Contents

Grid node recovery procedures	1
Warnings and considerations for grid node recovery	1
Preconditions for recovering grid nodes	1
Order of node recovery if a server hosting more than one grid node fails	2
IP addresses for recovered nodes	2
Gathering required materials for grid node recovery	2
Downloading and extracting the StorageGRID installation files	3
Selecting a node recovery procedure	8
Recovering from Storage Node failures	8
Recovering a Storage Node that has been down more than 15 days	10
Recovering a StorageGRID appliance Storage Node	11
Recovering from storage volume failure where the system drive is intact	31
Recovering from system drive failure	44
Recovering from Admin Node failures	61
Recovering from primary Admin Node failures	63
Recovering from non-primary Admin Node failures	71
Recovering from Gateway Node failures	78
Replacing a Gateway Node	79
Selecting Start Recovery to configure a Gateway Node	80
Recovering from Archive Node failures	82
Replacing an Archive Node	83
Selecting Start Recovery to configure an Archive Node	83
Resetting Archive Node connection to the cloud	85
All grid node types: Replacing a VMware node	85
All grid node types: Replacing a Linux node	86
Deploying new Linux hosts	87
Restoring grid nodes to the host	87
What's next: Performing additional recovery steps, if required	92
Replacing a failed node with a services appliance	93
Installing a services appliance (platform change only)	94
Preparing an appliance for reinstallation (platform replacement only)	94
Starting software installation on a services appliance	95
Monitoring services appliance installation	99

Grid node recovery procedures

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

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

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

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



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

Related information

[How site recovery is performed by technical support](#)

Warnings and considerations for grid node recovery

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



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

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

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

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

Preconditions for recovering grid nodes

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

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

- The StorageGRID Appliance Installer version on the replacement appliance matches the software version of your StorageGRID system, as described in hardware installation and maintenance for verifying and upgrading the StorageGRID Appliance Installer version.
 - [SG100 & SG1000 services appliances](#)
 - [SG5600 storage appliances](#)
 - [SG5700 storage appliances](#)
 - [SG6000 storage appliances](#)
- If you are recovering a grid node other than the primary Admin Node, there is connectivity between the grid node being recovered and the primary Admin Node.

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

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

IP addresses for recovered nodes

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 current 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	<p>If you need to recover a grid node, you need the StorageGRID installation archive for your platform.</p> <p>Note: You do not need to download files if you are recovering failed storage volumes on a Storage Node.</p>

Item	Notes
Recovery Package .zip file	<p>Obtain a copy of the most recent Recovery Package .zip file: sgws-recovery-package-<i>id-revision</i>.zip</p> <p>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.</p> <p>If the primary Admin Node is operating normally, you can download the Recovery Package from the Grid Manager. Select Maintenance > System > Recovery Package.</p> <p>If you cannot access the Grid Manager, you can find encrypted copies of the Recovery Package on some Storage Nodes that contain the ADC service. On each Storage Node, examine this location for the Recovery Package: /var/local/install/sgws-recovery-package-grid-<i>id-revision</i>.zip.gpg Use the Recovery Package with the highest revision number.</p>
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	<p>For the current supported versions of your platform, see the Interoperability Matrix Tool.</p> <p>NetApp Interoperability Matrix Tool</p> <p>Go to the platform vendor's website for documentation.</p>

Related information

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

[Web browser requirements](#)

Downloading and extracting the StorageGRID installation files

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

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 NetApp Downloads page for StorageGRID.

NetApp Downloads: StorageGRID

3. Select the version of StorageGRID that is currently running on the grid.

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

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

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

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

Platform	Installation archive
VMware	StorageGRID-Webscale-version-VMware-uniqueID.zip StorageGRID-Webscale-version-VMware-uniqueID.tgz
Red Hat Enterprise Linux or CentOS	StorageGRID-Webscale-version-RPM-uniqueID.zip StorageGRID-Webscale-version-RPM-uniqueID.tgz
Ubuntu or Debian or Appliances	StorageGRID-Webscale-version-DEB-uniqueID.zip StorageGRID-Webscale-version-DEB-uniqueID.tgz
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.

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

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

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

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

Path and file name	Description
./vsphere/NetApp-SG-version-SHA.vmdk	The virtual machine disk file that is used as a template for creating grid node virtual machines.
./vsphere/vsphere-primary-admin.ovf ./vsphere/vsphere-primary-admin.mf	The Open Virtualization Format template file (.ovf) and manifest file (.mf) for deploying the primary Admin Node.
./vsphere/vsphere-non-primary-admin.ovf ./vsphere/vsphere-non-primary-admin.mf	The template file (.ovf) and manifest file (.mf) for deploying non-primary Admin Nodes.
/vsphere/vsphere-archive.ovf ./vsphere/vsphere-archive.mf	The template file (.ovf) and manifest file (.mf) for deploying Archive Nodes.
./vsphere/vsphere-gateway.ovf ./vsphere/vsphere-gateway.mf	The template file (.ovf) and manifest file (.mf) for deploying Gateway Nodes.
./vsphere/vsphere-storage.ovf ./vsphere/vsphere-storage.mf	The template file (.ovf) and manifest file (.mf) for deploying virtual machine-based Storage Nodes.
Deployment scripting tool	Description
./vsphere/deploy-vsphere-ovftool.sh	A Bash shell script used to automate the deployment of virtual grid nodes.
./vsphere/deploy-vsphere-ovftool-sample.ini	A sample configuration file for use with the deploy-vsphere-ovftool.sh script.
./vsphere/configure-storagegrid.py	A Python script used to automate the configuration of a StorageGRID system.
./vsphere/configure-sga.py	A Python script used to automate the configuration of StorageGRID appliances.
./vsphere/storagegrid-ssoauth.py	An example Python script that you can use to sign in to the Grid Management API when single sign-on is enabled.
./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.

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

Path and file name	Description
./rpms/README	A text file that describes all of the files contained in the StorageGRID download file.
./rpms/NLF000000.txt	A free license that does not provide any support entitlement for the product.
./rpms/StorageGRID-Webscale-Images-version-SHA.rpm	RPM package for installing the StorageGRID node images on your RHEL or CentOS hosts.
./rpms/StorageGRID-Webscale-Service-version-SHA.rpm	RPM package for installing the StorageGRID host service on your RHEL or CentOS hosts.
Deployment scripting tool	Description
./rpms/configure-storagegrid.py	A Python script used to automate the configuration of a StorageGRID system.
./rpms/configure-sga.py	A Python script used to automate the configuration of StorageGRID appliances.
./rpms/configure-storagegrid.sample.json	A sample configuration file for use with the configure-storagegrid.py script.
./rpms/storagegrid-ssoauth.py	An example Python script that you can use to sign in to the Grid Management API when single sign-on is enabled.
./rpms/configure-storagegrid.blank.json	A blank configuration file for use with the configure-storagegrid.py script.
./rpms/extras/ansible	Example Ansible role and playbook for configuring RHEL or CentOS hosts for StorageGRID container deployment. You can customize the role or playbook as necessary.

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

Path and file name	Description
./debs/README	A text file that describes all of the files contained in the StorageGRID download file.
./debs/NLF000000.txt	A non-production NetApp License File that you can use for testing and proof of concept deployments.

Path and file name	Description
./debs/storagegrid-webscale-images-version-SHA.deb	DEB package for installing the StorageGRID node images on Ubuntu or Debian hosts.
./debs/storagegrid-webscale-images-version-SHA.deb.md5	MD5 checksum for the file /debs/storagegrid-webscale-images-version-SHA.deb
./debs/storagegrid-webscale-service-version-SHA.deb	DEB package for installing the StorageGRID host service on Ubuntu or Debian hosts.
Deployment scripting tool	Description
./debs/configure-storagegrid.py	A Python script used to automate the configuration of a StorageGRID system.
./debs/configure-sga.py	A Python script used to automate the configuration of StorageGRID appliances.
./debs/storagegrid-ssoauth.py	An example Python script that you can use to sign in to the Grid Management API when single sign-on is enabled.
./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/extras/ansible	Example Ansible role and playbook for configuring Ubuntu or Debian hosts for StorageGRID container deployment. You can customize the role or playbook as necessary.

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

Path and file name	Description
./debs/storagegrid-webscale-images-version-SHA.deb	DEB package for installing the StorageGRID node images on your appliances.
./debs/storagegrid-webscale-images-version-SHA.deb.md5	Checksum of the DEB installation package used by the StorageGRID Appliance Installer to validate that the package is intact after upload.

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

Related information

[Install VMware](#)

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

[Install Ubuntu or Debian](#)

Selecting a node recovery procedure

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

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



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

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
<ul style="list-style-type: none"> More than one Storage Node has failed. A second Storage Node has failed less than 15 days after a Storage Node failure or recovery. <p>This includes the case where a Storage Node fails while recovery of another Storage Node is still in progress.</p>	<p>You must contact technical support.</p>	<p>If all failed Storage Nodes are at the same site, it might be necessary to perform a site recovery procedure.</p> <p>Technical support will assess your situation and develop a recovery plan.</p> <p>How site recovery is performed by technical support</p> <p>Recovering more than one Storage Node (or more than one Storage Node within 15 days) might affect the integrity of the Cassandra database, which can cause data loss.</p> <p>Technical support can determine when it is safe to begin recovery of a second Storage Node.</p> <p>Note: If more than one Storage Node that contains the ADC service fails at a site, you lose any pending platform service requests for that site.</p>
A Storage Node has been offline for more than 15 days.	Recovering a Storage Node that has been down more than 15 days	This procedure is required to ensure Cassandra database integrity.
An appliance Storage Node has failed.	Recovering a StorageGRID appliance Storage Node	The recovery procedure for appliance Storage Nodes is the same for all failures.
One or more storage volumes have failed, but the system drive is intact	Recovering from storage volume failure where the system drive is intact	This procedure is used for software-based Storage Nodes.
The system drive has failed.	Recovering from system drive failure	The node replacement procedure depends on the deployment platform and on whether any storage volumes have also failed.



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

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.

What you'll need

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

About this task

Storage Nodes have a Cassandra database that includes object metadata. If a Storage Node has not been able to communicate with other Storage Nodes for more than 15 days, StorageGRID assumes that node's Cassandra database is stale. The Storage Node 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.



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



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



If more than one Storage Node at a site has failed, a site recovery procedure might be required. Contact technical support.

How site recovery is performed by technical support

Steps

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

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



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

3. Perform the following checks on the Storage Node:

- Issue this command: `nodetool status`

The output should be `Connection refused`

- In the Grid Manager, select **Support > Tools > Grid Topology**.
- Select **site > Storage Node > SSM > Services**. Verify that the Cassandra service displays **Not Running**.
- Select **Storage Node > SSM > Resources**. Verify that there is no error status in the Volumes section.
- 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**



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

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

- In the Grid Manager, select **Support > Tools > Grid Topology**.
- Select **site > recovered Storage Node > SSM > Services**.
- Confirm that all services are running.
- Select **DDS > Data Store**.
- Confirm that the **Data Store Status** is “Up” and the **Data Store State** is “Normal.”

Related information

[Recovering from system drive failure](#)

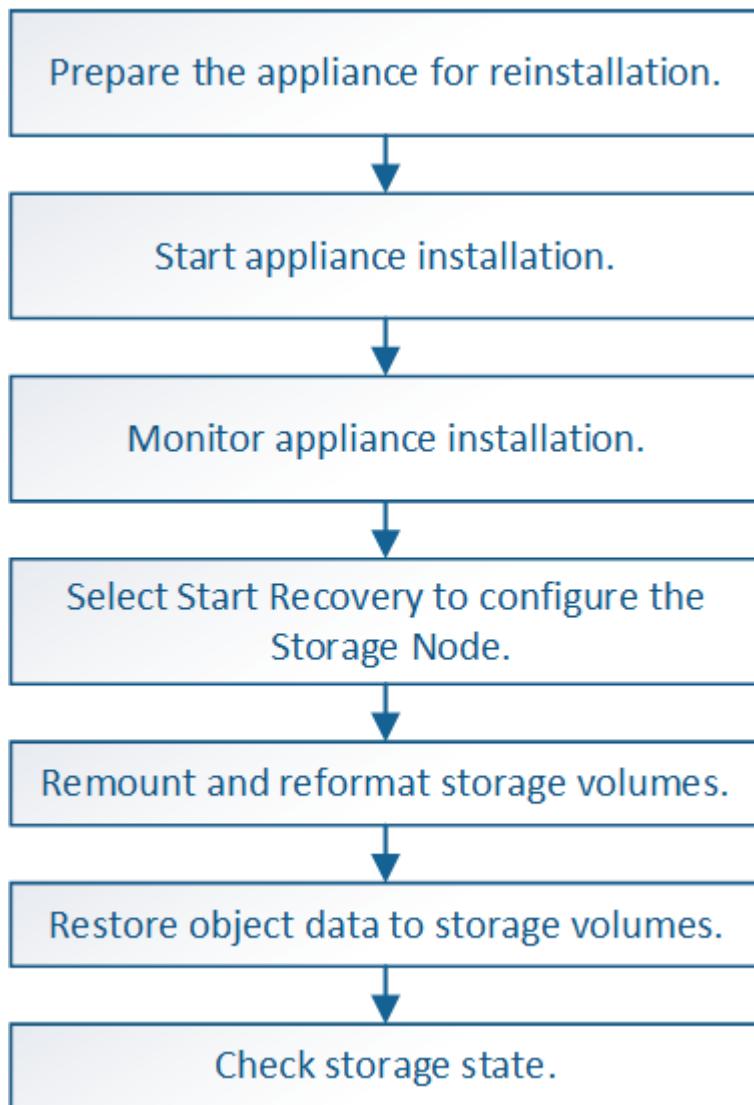
Recovering a StorageGRID appliance Storage Node

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.



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



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



If more than one Storage Node at a site has failed, a site recovery procedure might be required. Contact technical support.

How site recovery is performed by technical support



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



If you encounter a Services: Status - Cassandra (SVST) alarm during recovery, see the monitoring and 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.



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

Related information

[Monitor & troubleshoot](#)

[SG6000 storage appliances](#)

[SG5700 storage appliances](#)

[SG5600 storage appliances](#)

Steps

- [Preparing an appliance Storage Node for reinstallation](#)
- [Starting StorageGRID appliance installation](#)
- [Monitoring StorageGRID appliance installation](#)
- [Selecting Start Recovery to configure an appliance Storage Node](#)
- [Remounting and reformatting appliance storage volumes \(“Manual Steps”\)](#)
- [Restoring object data to a storage volume for an appliance](#)
- [Checking the storage state after recovering an appliance Storage Node](#)

Preparing an appliance Storage Node for reinstallation

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

1. Log in to the failed Storage Node:

- Enter the following command: `ssh admin@grid_node_IP`
- Enter the password listed in the `Passwords.txt` file.
- Enter the following command to switch to root: `su -`
- 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.

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

Starting StorageGRID appliance installation

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

What you'll need

- The appliance has been installed in a rack, connected to your networks, and powered on.
- Network links and IP addresses have been configured for the appliance using the StorageGRID Appliance Installer.
- You know the IP address of the primary Admin Node for the StorageGRID grid.
- All Grid Network subnets listed on the IP Configuration page of the StorageGRID Appliance Installer have been defined in the Grid Network Subnet List on the primary Admin Node.
- You have completed these prerequisite tasks by following the installation and maintenance instructions for your storage appliance:
 - [SG5600 storage appliances](#)
 - [SG5700 storage appliances](#)
 - [SG6000 storage appliances](#)
- You are using a supported web browser.
- You know one of the IP addresses assigned to the compute controller in the appliance. You can use the IP address for the Admin Network (management port 1 on the controller), the Grid Network, or the Client Network.

About this task

To install StorageGRID on an appliance Storage Node:

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

Steps

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

`https://Controller_IP:8443`

The StorageGRID Appliance Installer Home page appears.

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

The StorageGRID Appliance Installer can discover this IP address automatically, assuming the primary

Admin Node, or at least one other grid node with ADMIN_IP configured, is present on the same subnet.

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

Option	Steps
Manual IP entry	<ol style="list-style-type: none">Unselect the Enable Admin Node discovery check box.Enter the IP address manually.Click Save.Wait while the connection state for the new IP address becomes “ready.”
Automatic discovery of all connected primary Admin Nodes	<ol style="list-style-type: none">Select the Enable Admin Node discovery check box.From the list of discovered IP addresses, select the primary Admin Node for the grid where this appliance Storage Node will be deployed.Click Save.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**.

NetApp® StorageGRID® Appliance Installer

[Home](#)[Configure Networking](#)[Configure Hardware](#)[Monitor Installation](#)[Advanced](#)

Home

 The installation is ready to be started. Review the settings below, and then click Start Installation.

Primary Admin Node connection

Enable Admin Node discovery

Primary Admin Node IP

Connection state Connection to 172.16.4.210 ready

[Cancel](#)[Save](#)

Node name

Node name

[Cancel](#)[Save](#)

Installation

Current state Ready to start installation of NetApp-SGA into grid with Admin Node 172.16.4.210.

[Start Installation](#)

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



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

Related information

[SG100 & SG1000 services appliances](#)

[SG6000 storage appliances](#)

[SG5700 storage appliances](#)

[SG5600 storage appliances](#)

Monitoring StorageGRID appliance installation

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

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

The Monitor Installation page shows the installation progress.

Monitor Installation

1. Configure storage			Running
Step	Progress	Status	
Connect to storage controller	<div style="width: 100%; background-color: green;"></div>	Complete	
Clear existing configuration	<div style="width: 100%; background-color: green;"></div>	Complete	
Configure volumes	<div style="width: 25%; background-color: blue;"></div>	Creating volume StorageGRID-obj-00	
Configure host settings	<div style="width: 0%; background-color: grey;"></div>	Pending	

2. Install OS			Pending
3. Install StorageGRID			Pending
4. Finalize installation			Pending

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



The installer ensures that tasks completed in a previous install 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 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.

Home	Configure Networking ▾	Configure Hardware ▾	Monitor Installation	Advanced ▾	
------	------------------------	----------------------	----------------------	------------	--

Monitor Installation

1. Configure storage	Complete
2. Install OS	Complete
3. Install StorageGRID	Running
4. Finalize installation	Pending

Connected (unencrypted) to: QEMU

```
/platform.type=: Device or resource busy
[2017-07-31T22:09:12.362566]    INFO -- [INSG] NOTICE: seeding /var/local with c
ontainer data
[2017-07-31T22:09:12.366205]    INFO -- [INSG] Fixing permissions
[2017-07-31T22:09:12.369633]    INFO -- [INSG] Enabling syslog
[2017-07-31T22:09:12.511533]    INFO -- [INSG] Stopping system logging: syslog-n
g.
[2017-07-31T22:09:12.570096]    INFO -- [INSG] Starting system logging: syslog-n
g.
[2017-07-31T22:09:12.576360]    INFO -- [INSG] Beginning negotiation for downloa
d of node configuration
[2017-07-31T22:09:12.581363]    INFO -- [INSG]
[2017-07-31T22:09:12.585066]    INFO -- [INSG]
[2017-07-31T22:09:12.588314]    INFO -- [INSG]
[2017-07-31T22:09:12.591851]    INFO -- [INSG]
[2017-07-31T22:09:12.594886]    INFO -- [INSG]
[2017-07-31T22:09:12.598360]    INFO -- [INSG]
[2017-07-31T22:09:12.601324]    INFO -- [INSG]
[2017-07-31T22:09:12.604759]    INFO -- [INSG]
[2017-07-31T22:09:12.607800]    INFO -- [INSG]
[2017-07-31T22:09:12.610985]    INFO -- [INSG]
[2017-07-31T22:09:12.614597]    INFO -- [INSG]
[2017-07-31T22:09:12.618282]    INFO -- [INSG] Please approve this node on the A
dmin Node GMI to proceed...

```

4. Go to the procedure to configure the appliance Storage Node.

Selecting Start Recovery to configure an appliance Storage Node

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

What you'll need

- You must be signed in to the Grid Manager using a supported browser.
- You must have the Maintenance or Root Access permission.
- 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 > Maintenance Tasks > Recovery**.
2. Select the grid node you want to recover in the Pending Nodes list.

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

3. Enter the **Provisioning Passphrase**.

4. Click **Start Recovery**.

Recovery

Select the failed grid node to recover, enter your provisioning passphrase, and then click Start Recovery to begin the recovery procedure.

Pending Nodes

	Name	IPv4 Address	State	Recoverable	
<input checked="" type="radio"/>	104-217-S1	10.96.104.217	Unknown		

Passphrase

Provisioning Passphrase

.....

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.

Recovery

Select the failed grid node to recover, enter your provisioning passphrase, and then click Start Recovery to begin the recovery procedure.

Recovering Grid Node

Name	Start Time	Progress	Stage
dc2-s3	2016-09-12 16:12:40 PDT	<div style="width: 25%;"><div style="width: 100%;"></div></div>	Waiting For Manual Steps

Reset



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.

Reset Recovery

Resetting the recovery procedure leaves the deployed grid node in an indeterminate state. To retry a recovery after resetting the procedure, you must restore the node to a pre-installed state:

- For VMware nodes, delete the deployed VM and then redeploy it.
- For StorageGRID appliance nodes, run "sgareinstall" on the node.
- For Linux nodes, run "storagegrid node force-recovery *node-name*" on the Linux host.

Do you want to reset recovery?

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 any failed storage volumes. The first script remounts volumes that are properly formatted as StorageGRID storage volumes. The second script reformats any unmounted volumes, rebuilds the Cassandra database, if needed, and starts services.

What you'll need

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

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

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

 Contact technical support if more than one Storage Node is offline or if a Storage Node in this grid has been rebuilt in the last 15 days. 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 `sn-remount-volumes` script to remount properly formatted storage volumes. When this script runs, it does the following:

- Mounts and unmounts each storage volume to replay the XFS journal.
- Performs an XFS file consistency check.
- If the file system is consistent, determines if the storage volume is a properly formatted StorageGRID storage volume.
- If the storage volume is properly formatted, remounts the storage volume. Any existing data on the volume remains intact.
- Review the script output and resolve any issues.
- Run the `sn-recovery-postinstall.sh` script. When this script runs, it does the following.



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

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



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

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

Steps

1. Log in to the recovered Storage Node:

- Enter the following command: `ssh admin@grid_node_IP`
- Enter the password listed in the `Passwords.txt` file.
- Enter the following command to switch to root: `su -`
- Enter the password listed in the `Passwords.txt` file.

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

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



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

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

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

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



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

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

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

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

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

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.



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

- a. Check to make sure that the results include an entry for all of the volumes you expected. If any volumes 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.

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

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



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

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



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

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

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



If a storage volume cannot 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).



Do not run the `sn-recovery-postinstall.sh` script if you believe that the data remaining on a failed storage 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 `sn-recovery-postinstall.sh` script: `sn-recovery-postinstall.sh`

This script reformats any storage volumes that could not be mounted or that were found to be improperly formatted; rebuilds the Cassandra database on the node, if needed; and starts the services on the Storage Node.

Be aware of the following:

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



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



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

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

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

Recovery

Select the failed grid node to recover, enter your provisioning passphrase, and then click Start Recovery to begin the recovery procedure.

Pending Nodes

Name	IPv4 Address	State	Recoverable	Actions
No results found.				

Recovering Grid Node

Name	Start Time	Progress	Stage
DC1-S3	2016-06-02 14:03:35 PDT	<div style="width: 50%;"><div style="width: 100%;"></div></div>	Recovering Cassandra

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

Related information

[Reviewing warnings for Storage Node system drive recovery](#)

[Restoring object data to a storage volume for an appliance](#)

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.

What you'll need

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

About this task

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



If an ILM rule was configured to store only one replicated copy and that copy existed on a storage volume that failed, you will not be able to recover the object.



If the only remaining copy of an object is in a Cloud Storage Pool, StorageGRID must issue multiple requests to the Cloud Storage Pool endpoint to restore object data. Before performing this procedure, contact technical support for help in estimating the recovery time frame and the associated costs.



If the only remaining copy of an object is on an Archive Node, object data is retrieved from the Archive Node. 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
```

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

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

Steps

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

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

2. Use the `/etc/hosts` file to find the hostname of the Storage Node for the restored storage volumes. To see a list of all nodes in the grid, enter the following: `cat /etc/hosts`
3. If all storage volumes have failed, repair the entire node. (If only some volumes have failed, go to the next step.)



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 a Storage Node named SG-DC-SN3:

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



As object data is restored, the **Objects Lost** alert is triggered if the StorageGRID system cannot locate replicated object data. Alerts might be triggered on Storage Nodes throughout the system. You should determine the cause of the loss and if recovery is possible. See the instructions for monitoring and 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.



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. For example, 0000 is the first volume and 000F is the sixteenth volume. You can specify one volume, a range of volumes, or multiple volumes that are not in a sequence.

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

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

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

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

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

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

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

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



As object data is restored, the **Objects Lost** alert is triggered if the StorageGRID system cannot locate replicated object data. Alerts might be triggered on Storage Nodes throughout the system. You should determine the cause of the loss and if recovery is possible. See the instructions for monitoring and troubleshooting StorageGRID.

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

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

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

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

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

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

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

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



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

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

5. Monitor the repair of replicated data.
 - a. Select **Nodes > Storage Node being repaired > ILM**.
 - b. Use the attributes in the Evaluation section to determine if repairs are complete.

When repairs are complete, the Awaiting - All attribute indicates 0 objects.

- c. To monitor the repair in more detail, select **Support > Tools > Grid Topology**.
- d. Select **grid > Storage Node being repaired > LDR > Data Store**.
- e. Use a combination of the following attributes to determine, as well as possible, if replicated repairs are complete.



Cassandra inconsistencies might be present, and failed repairs 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.



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 applies to the entire grid and is the maximum of all node scan periods. You can query the **Scan Period — Estimated** attribute history for the grid to determine an appropriate time frame.

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/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

[Monitor & troubleshoot](#)

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.

What you'll need

- 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

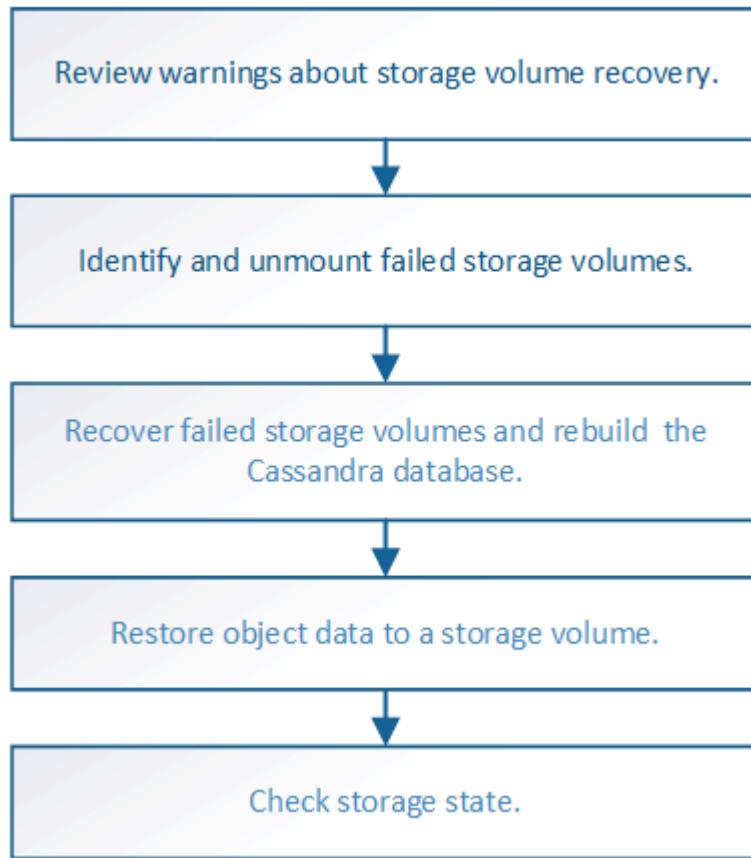
1. Select **Support > Tools > Grid Topology**.
2. Check the values of **Recovered Storage Node > LDR > Storage > Storage State — Desired** and **Storage State — Current**.
The value of both attributes should be Online.
3. If the Storage State — Desired is set to Read-only, complete the following steps:
 - a. Click the **Configuration** tab.
 - b. From the **Storage State — Desired** drop-down list, select **Online**.
 - c. Click **Apply Changes**.
 - d. Click the **Overview** tab and confirm that the values of **Storage State — Desired** and **Storage State — Current** are updated to Online.

Recovering from storage volume failure where the system drive is intact

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

About this task

This recovery procedure applies to software-based Storage Nodes only. If storage volumes have failed on an appliance Storage Node, use the procedure for “Recovering a StorageGRID appliance Storage Node.”



Related information

[Recovering a StorageGRID appliance Storage Node](#)

Steps

- Reviewing warnings about storage volume recovery
- Identifying and unmounting failed storage volumes
- Recovering failed storage volumes and rebuilding the Cassandra database
- Restoring object data to a storage volume where the system drive is intact
- Checking the storage state after recovering storage volumes

Reviewing warnings about storage volume recovery

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

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

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

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

When Cassandra is rebuilt, the system uses information from other Storage Nodes. If too many Storage Nodes are offline, some Cassandra data might not be available. If Cassandra has been rebuilt recently, Cassandra data might not yet be consistent across the grid. Data loss can occur if Cassandra is rebuilt when too many Storage Nodes are offline or if two or more Storage Nodes are rebuilt within 15 days of each other.



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



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



If more than one Storage Node at a site has failed, a site recovery procedure might be required. Contact technical support.

[How site recovery is performed by technical support](#)



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



If you encounter a Services: Status - Cassandra (SVST) alarm during recovery, see the monitoring and 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 information

[Monitor & troubleshoot](#)

[Warnings and considerations for grid node recovery](#)

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.

What you'll need

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.



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



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

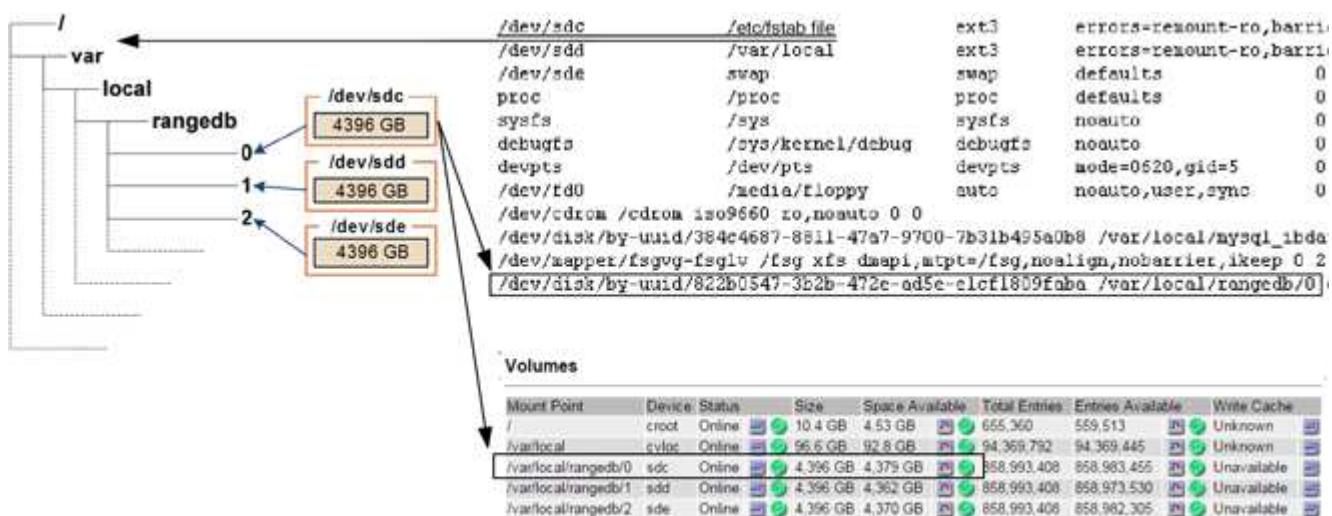


Identify failed storage volumes carefully. You will use this information to verify which volumes must be reformatted. Once a volume has been reformatted, data on the volume cannot be recovered.

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

Object Stores

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

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

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

Volumes

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

2. Log in to the failed Storage Node:

- Enter the following command: `ssh admin@grid_node_IP`
- Enter the password listed in the `Passwords.txt` file.
- Enter the following command to switch to root: `su -`
- Enter the password listed in the `Passwords.txt` file.

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

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

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

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

4. If prompted, press **y** to stop the storage services on the Storage Node.



If the storage services are already stopped, you are not prompted. The Cassandra service is stopped only for volume 0.

```
root@Storage-180:~ # sn-unmount-volume 0
Storage services (ldr, chunk, dds, cassandra) are not down.
Storage services must be stopped before running this script.
Stop storage services [y/N]? y
Shutting down storage services.
Storage services stopped.
Unmounting /var/local/rangedb/0
/var/local/rangedb/0 is unmounted.
```

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

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.

- 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 the node decommission procedure. (In the Grid Manager, select **Maintenance > Maintenance Tasks > Decommission**.)
- You have checked that an expansion is not in progress. (In the Grid Manager, select **Maintenance > Maintenance Tasks > Expansion**.)
- You have reviewed the warnings about storage volume recovery.

Reviewing warnings about storage volume recovery

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. 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 `#`.

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



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

2. 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: `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.



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



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

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

```
root@DC1-S1:~ # reformat_storage_block_devices.rb
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 information

[Reviewing warnings about storage volume recovery](#)

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.

What you'll need

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

About this task

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



If an ILM rule was configured to store only one replicated copy and that copy existed on a storage volume that failed, you will not be able to recover the object.



If the only remaining copy of an object is in a Cloud Storage Pool, StorageGRID must issue multiple requests to the Cloud Storage Pool endpoint to restore object data. Before performing this procedure, contact technical support for help in estimating the recovery time frame and the associated costs.



If the only remaining copy of an object is on an Archive Node, object data is retrieved from the Archive Node. 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
```



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

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

Steps

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

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

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

2. Use the `/etc/hosts` file to find the hostname of the Storage Node for the restored storage volumes. To see a list of all nodes in the grid, enter the following: `cat /etc/hosts`
3. If all storage volumes have failed, repair the entire node. (If only some volumes have failed, go to the next step.)



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 a Storage Node named SG-DC-SN3:

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



As object data is restored, the **Objects Lost** alert is triggered if the StorageGRID system cannot locate replicated object data. Alerts might be triggered on Storage Nodes throughout the system. You should determine the cause of the loss and if recovery is possible. See the instructions for monitoring and 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.



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. For example, `0000` is the first volume and `000F` is the sixteenth volume. You can specify one volume, a range of volumes, or multiple volumes that are not in a sequence.

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

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

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

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

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

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

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

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



As object data is restored, the **Objects Lost** alert is triggered if the StorageGRID system cannot locate replicated object data. Alerts might be triggered on Storage Nodes throughout the system. You should determine the cause of the loss and if recovery is possible. See the instructions for monitoring and troubleshooting StorageGRID.

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

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

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

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

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

Multiple volumes not in a sequence: This command restores erasure coded data to volumes 000A,

000C, and 000E on a Storage Node named SG-DC-SN3:

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

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



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

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

5. Monitor the repair of replicated data.

- Select **Nodes > Storage Node being repaired > ILM**.
- Use the attributes in the Evaluation section to determine if repairs are complete.

When repairs are complete, the Awaiting - All attribute indicates 0 objects.

- To monitor the repair in more detail, select **Support > Tools > Grid Topology**.
- Select **grid > Storage Node being repaired > LDR > Data Store**.
- Use a combination of the following attributes to determine, as well as possible, if replicated repairs are complete.



Cassandra inconsistencies might be present, and failed repairs 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.



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 applies to the entire grid and is the maximum of all node scan periods. You can query the **Scan Period — Estimated** attribute history for the grid to determine an appropriate time frame.

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

- 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/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
```

- 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

[Administer StorageGRID](#)

[Monitor & troubleshoot](#)

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.

What you'll need

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

Recovering from system drive failure

If the system drive on a software-based Storage Node has failed, the Storage Node is not available to the StorageGRID system. You must complete a specific set of tasks to recover from a system drive failure.

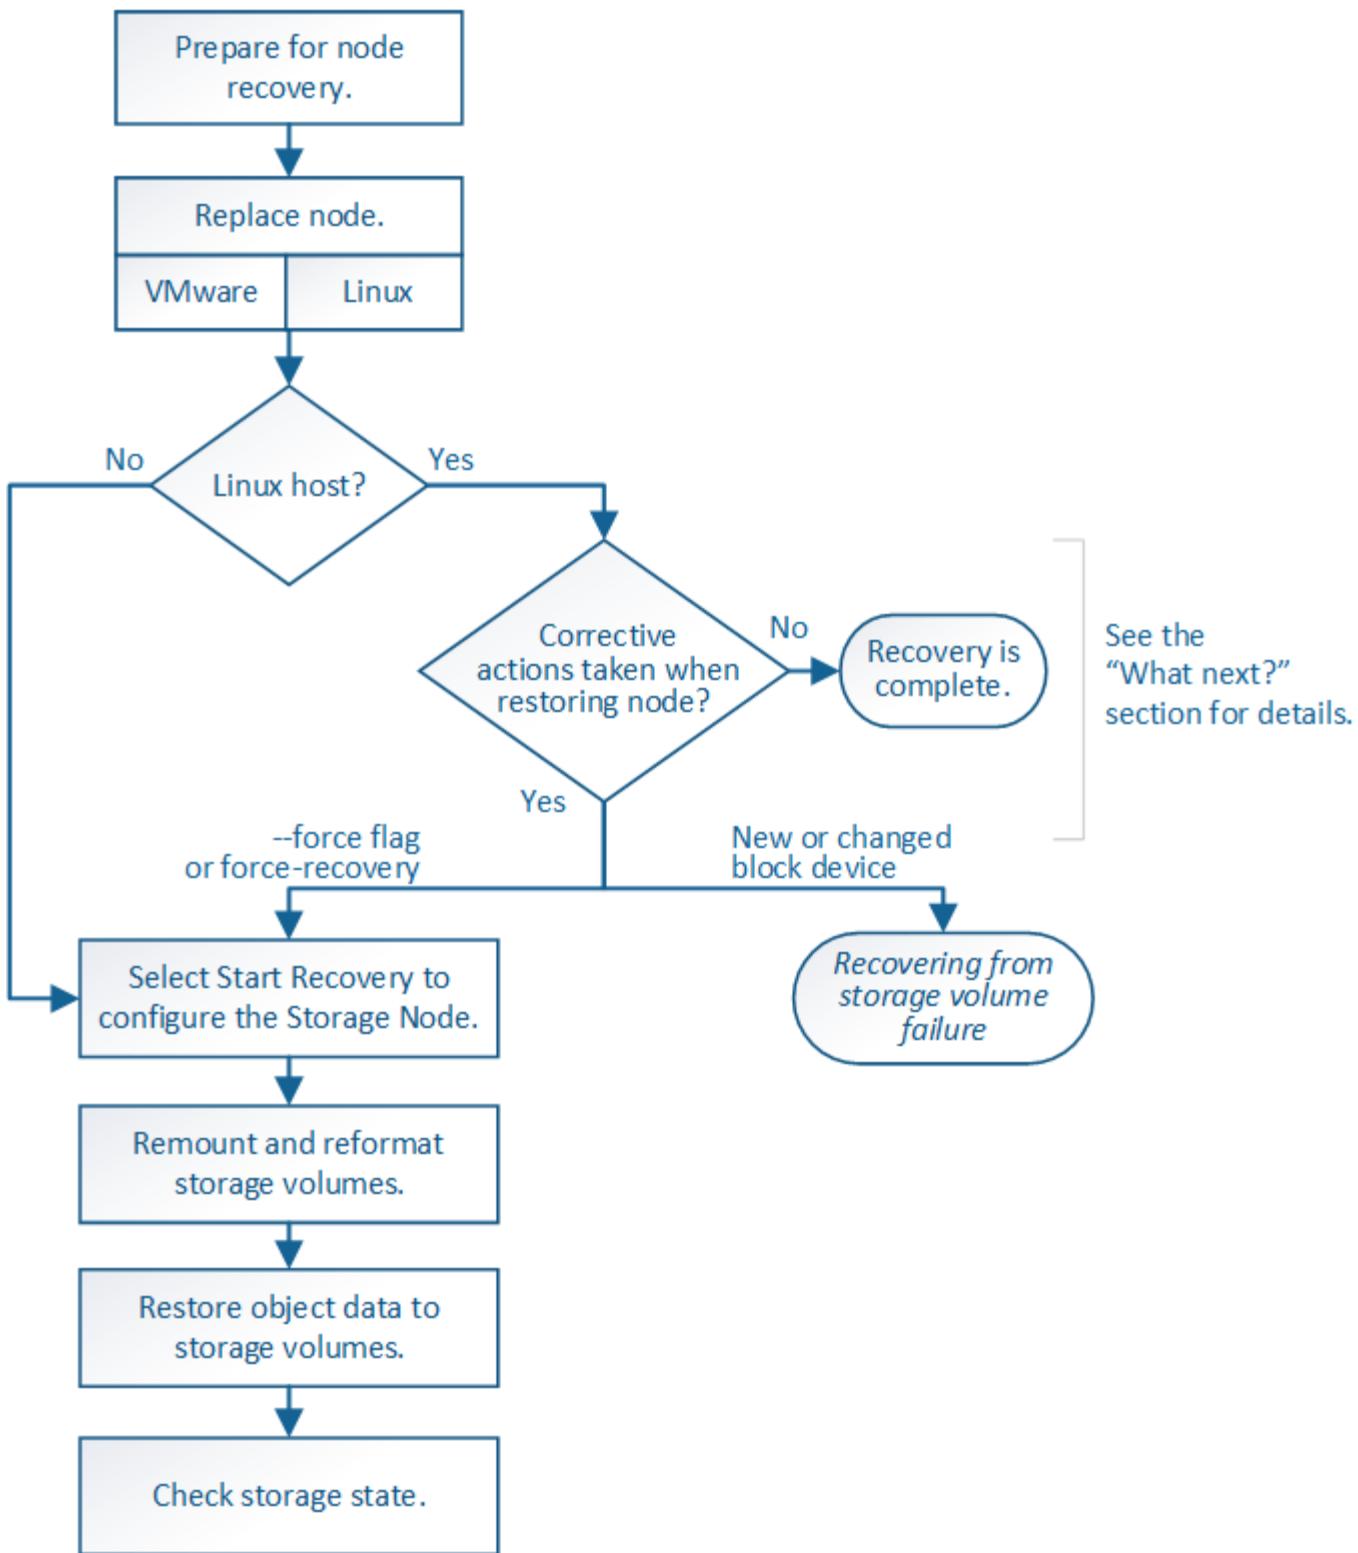
About this task

Use this procedure to recover from a system drive failure on a software-based Storage Node. This procedure includes the steps to follow if any storage volumes also failed or cannot be remounted.



This procedure applies to software-based Storage Nodes only. You must follow a different procedure to recover an appliance Storage Node.

[Recovering a StorageGRID appliance Storage Node](#)



Steps

- Reviewing warnings for Storage Node system drive recovery
- Replacing the Storage Node
- Selecting Start Recovery to configure a Storage Node
- Remounting and reformatting storage volumes (“Manual Steps”)
- Restoring object data to a storage volume, if required

- [Checking the storage state after recovering a Storage Node system drive](#)

Reviewing warnings for Storage Node system drive recovery

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

Storage Nodes have a Cassandra database that includes object metadata. The Cassandra database might be rebuilt in the following circumstances:

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

When Cassandra is rebuilt, the system uses information from other Storage Nodes. If too many Storage Nodes are offline, some Cassandra data might not be available. If Cassandra has been rebuilt recently, Cassandra data might not yet be consistent across the grid. Data loss can occur if Cassandra is rebuilt when too many Storage Nodes are offline or if two or more Storage Nodes are rebuilt within 15 days of each other.



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



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



If more than one Storage Node at a site has failed, a site recovery procedure might be required. Contact technical support.

How site recovery is performed by technical support



If this Storage Node is in read-only maintenance mode to allow for the retrieval of objects by another Storage Node with failed storage volumes, recover volumes on the Storage Node with failed storage volumes before recovering this failed Storage Node. See the instructions for recovering from loss of storage volumes where the system drive is intact.



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



If you encounter a Services: Status - Cassandra (SVST) alarm during recovery, see the monitoring and 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 information

[Monitor & troubleshoot](#)

[Warnings and considerations for grid node recovery](#)

[Recovering from storage volume failure where the system drive is intact](#)

Replacing the Storage Node

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

You must select the node replacement procedure for your platform. The steps to replace a node are the same for all types of grid nodes.



This procedure applies to software-based Storage Nodes only. You must follow a different procedure to recover an appliance Storage Node.

Recovering a StorageGRID appliance Storage Node

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.

Platform	Procedure
VMware	Replacing a VMware node
Linux	Replacing a Linux node
OpenStack	NetApp-provided virtual machine disk files and scripts for OpenStack are no longer supported for recovery operations. If you need to recover a node running in an OpenStack deployment, download the files for your Linux operating system. Then, follow the procedure for replacing a Linux node.

Selecting Start Recovery to configure a Storage Node

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

What you'll need

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

About this task

If the Storage Node is installed as a container on a Linux host, you must perform this step only if one of these is true:

- You had to use the `--force` flag to import the node, or you issued `storagegrid node force-recovery node-name`
- You had to do a full node reinstall, or you needed to restore `/var/local`.

Steps

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

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

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

3. Enter the **Provisioning Passphrase**.

4. Click **Start Recovery**.

Recovery

Select the failed grid node to recover, enter your provisioning passphrase, and then click Start Recovery to begin the recovery procedure.

Pending Nodes

	Name	IPv4 Address	State	Recoverable	
<input checked="" type="radio"/>	104-217-S1	10.96.104.217	Unknown		

Passphrase

Provisioning Passphrase

.....

Start Recovery

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



While the recovery procedure is running, you can click **Reset** to start a new recovery. An Info dialog box appears, indicating that the node will be left in an indeterminate state if you reset the procedure.



Reset Recovery

Resetting the recovery procedure leaves the deployed grid node in an indeterminate state. To retry a recovery after resetting the procedure, you must restore the node to a pre-installed state:

- For VMware nodes, delete the deployed VM and then redeploy it.
- For StorageGRID appliance nodes, run "sgareinstall" on the node.
- For Linux nodes, run "storagegrid node force-recovery *node-name*" on the Linux host.

Do you want to reset recovery?

Cancel

OK

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.

Recovery

Select the failed grid node to recover, enter your provisioning passphrase, and then click Start Recovery to begin the recovery procedure.

Recovering Grid Node

Name	Start Time	Progress	Stage
dc2-s3	2016-09-12 16:12:40 PDT	<div style="width: 20%;"><div style="width: 100%;"></div></div>	Waiting For Manual Steps

[Reset](#)

Related information

[Preparing an appliance for reinstallation \(platform replacement only\)](#)

Remounting and reformatting storage volumes (“Manual Steps”)

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

What you’ll need

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

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

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

[Reviewing warnings for Storage Node system drive recovery](#)



Contact technical support if more than one Storage Node is offline or if a Storage Node in this grid has been rebuilt in the last 15 days. 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 `sn-remount-volumes` script to remount properly formatted storage volumes. When this script runs, it does the following:
 - Mounts and unmounts each storage volume to replay the XFS journal.
 - Performs an XFS file consistency check.
 - If the file system is consistent, determines if the storage volume is a properly formatted StorageGRID storage volume.
 - If the storage volume is properly formatted, remounts the storage volume. Any existing data on the volume remains intact.
- Review the script output and resolve any issues.
- Run the `sn-recovery-postinstall.sh` script. When this script runs, it does the following.



Do not reboot a Storage Node during recovery before running `sn-recovery-postinstall.sh` (see the step for [post-install script](#)) to reformat the failed storage volumes and restore object metadata. Rebooting the Storage Node before `sn-recovery-postinstall.sh` completes causes errors for services that attempt to start and causes StorageGRID appliance nodes to exit maintenance mode.

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



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

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

Steps

1. Log in to the recovered Storage Node:

- Enter the following command: `ssh admin@grid_node_IP`
- Enter the password listed in the `Passwords.txt` file.
- Enter the following command to switch to root: `su -`
- Enter the password listed in the `Passwords.txt` file.

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

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



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

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

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

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



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

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

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

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

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

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.



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

- a. Check to make sure that the results include an entry for all of the volumes you expected. If any volumes 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.

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

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



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

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



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

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

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



If a storage volume cannot 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).



Do not run the `sn-recovery-postinstall.sh` script if you believe that the data remaining on a failed storage 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 `sn-recovery-postinstall.sh` script: `sn-recovery-postinstall.sh`

This script reformats any storage volumes that could not be mounted or that were found to be improperly formatted; rebuilds the Cassandra database on the node, if needed; and starts the services on the Storage Node.

Be aware of the following:

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



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



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

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

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

Recovery

Select the failed grid node to recover, enter your provisioning passphrase, and then click Start Recovery to begin the recovery procedure.

Pending Nodes

				Search 
Name	IPv4 Address	State	Recoverable	Actions
<i>No results found.</i>				

Recovering Grid Node

Name	Start Time	Progress	Stage
DC1-S3	2016-06-02 14:03:35 PDT		Recovering Cassandra

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

Related information

[Reviewing warnings for Storage Node system drive recovery](#)

[Restoring object data to a storage volume, if required](#)

Restoring object data to a storage volume, if required

If the `sn-recovery-postinstall.sh` script is needed to reformat one or more failed storage volumes, you must restore object data to the reformatted storage volume from other Storage Nodes and Archive Nodes. These steps are not required unless one or more storage volumes were reformatted.

What you'll need

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

About this task

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



If an ILM rule was configured to store only one replicated copy and that copy existed on a storage volume that failed, you will not be able to recover the object.



If the only remaining copy of an object is in a Cloud Storage Pool, StorageGRID must issue multiple requests to the Cloud Storage Pool endpoint to restore object data. Before performing this procedure, contact technical support for help in estimating the recovery time frame and the associated costs.



If the only remaining copy of an object is on an Archive Node, object data is retrieved from the Archive Node. 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
```



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

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

Steps

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

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

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

2. Use the `/etc/hosts` file to find the hostname of the Storage Node for the restored storage volumes. To see a list of all nodes in the grid, enter the following: `cat /etc/hosts`
3. If all storage volumes have failed, repair the entire node. (If only some volumes have failed, go to the next step.)



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 a Storage Node named SG-DC-SN3:

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



As object data is restored, the **Objects Lost** alert is triggered if the StorageGRID system cannot locate replicated object data. Alerts might be triggered on Storage Nodes throughout the system. You should determine the cause of the loss and if recovery is possible. See the instructions for monitoring and 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.



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. For example, `0000` is the first volume and `000F` is the sixteenth volume. You can specify one volume, a range of volumes, or multiple volumes that are not in a sequence.

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

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

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

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

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

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

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

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



As object data is restored, the **Objects Lost** alert is triggered if the StorageGRID system cannot locate replicated object data. Alerts might be triggered on Storage Nodes throughout the system. You should determine the cause of the loss and if recovery is possible. See the instructions for monitoring and troubleshooting StorageGRID.

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

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

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

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

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

Multiple volumes not in a sequence: This command restores erasure coded data to volumes 000A,

000C, and 000E on a Storage Node named SG-DC-SN3:

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

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



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

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

5. Monitor the repair of replicated data.

- Select **Nodes > Storage Node being repaired > ILM**.
- Use the attributes in the Evaluation section to determine if repairs are complete.

When repairs are complete, the Awaiting - All attribute indicates 0 objects.

- To monitor the repair in more detail, select **Support > Tools > Grid Topology**.
- Select **grid > Storage Node being repaired > LDR > Data Store**.
- Use a combination of the following attributes to determine, as well as possible, if replicated repairs are complete.



Cassandra inconsistencies might be present, and failed repairs 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.



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 applies to the entire grid and is the maximum of all node scan periods. You can query the **Scan Period — Estimated** attribute history for the grid to determine an appropriate time frame.

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

- 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/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
```

- 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

[Administer StorageGRID](#)

[Monitor & troubleshoot](#)

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.

What you'll need

- 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

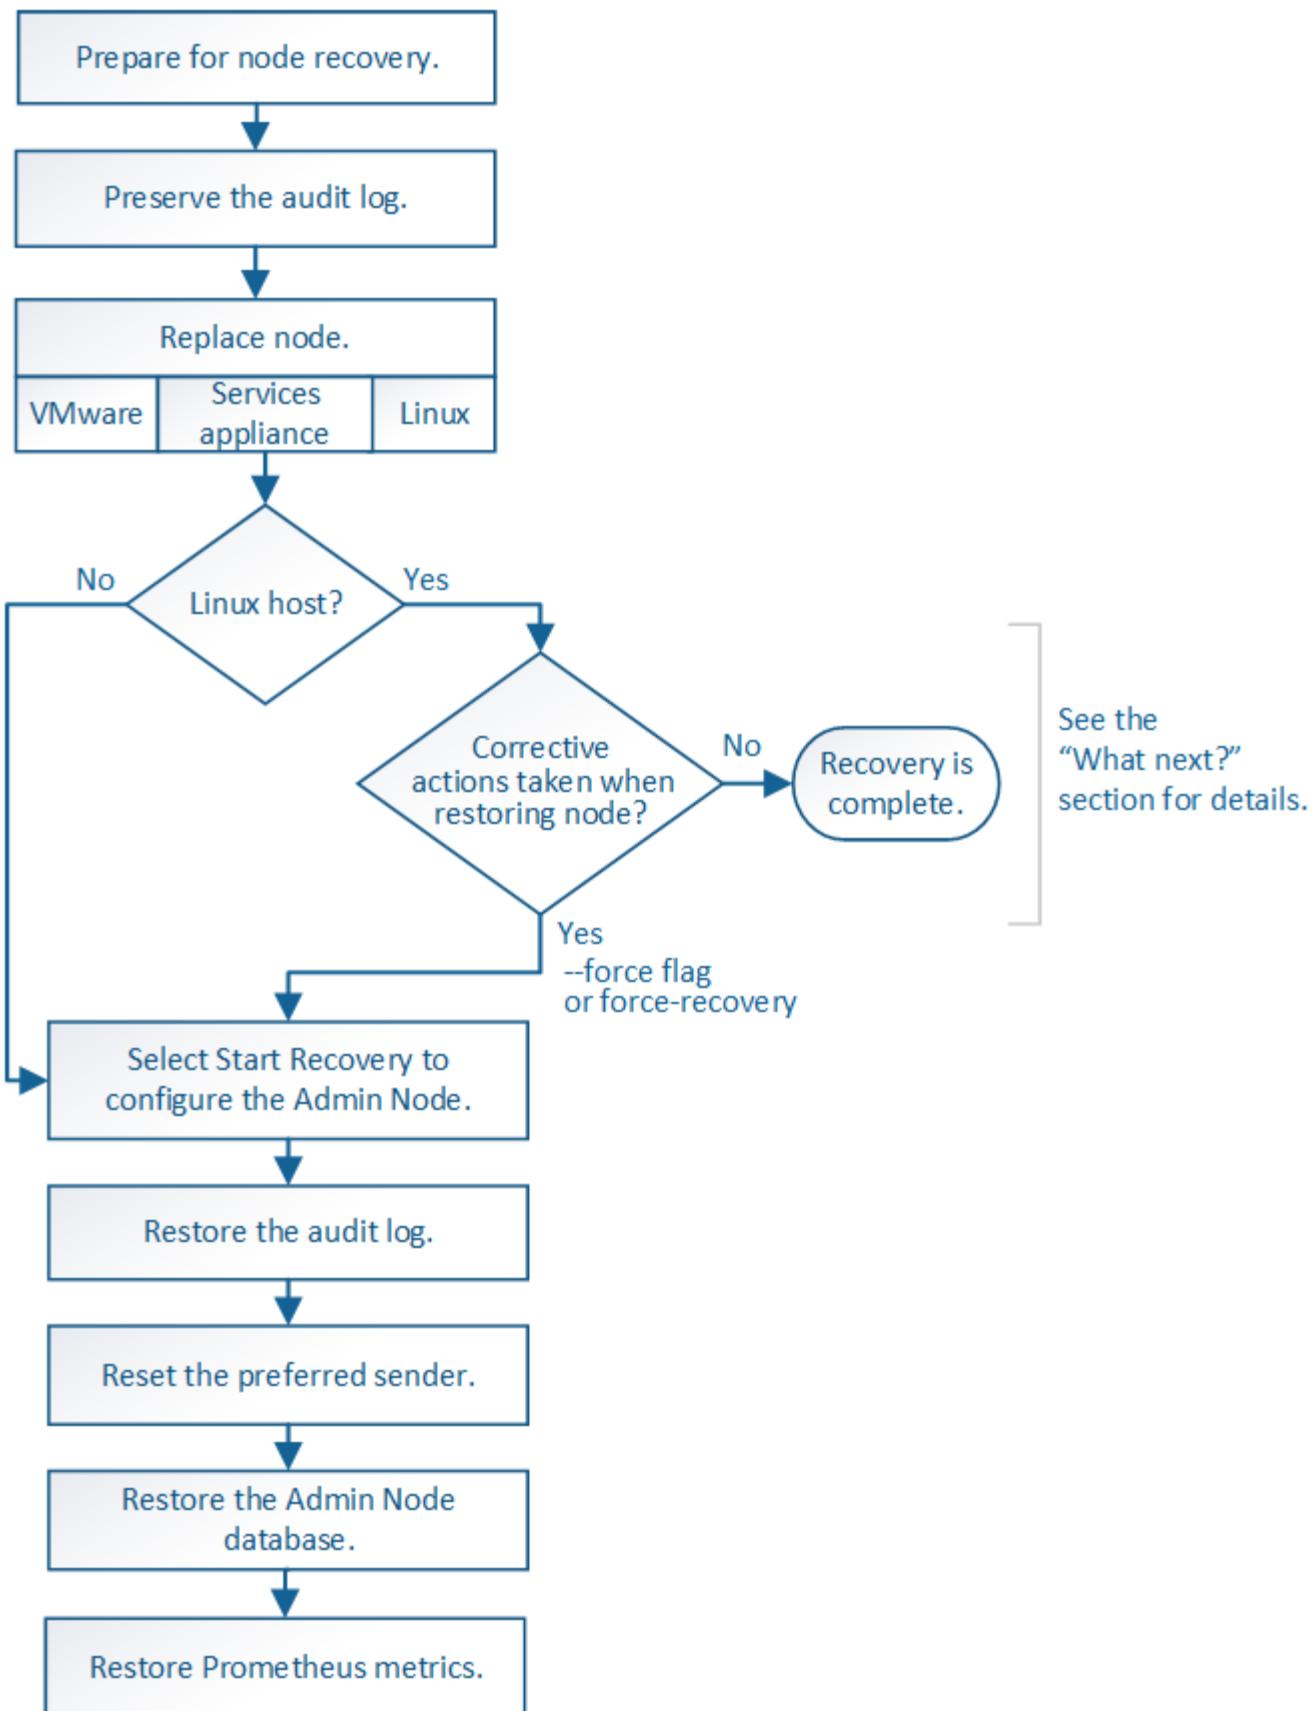
1. Select **Support > Tools > Grid Topology**.
2. Check the values of **Recovered Storage Node > LDR > Storage > Storage State — Desired** and **Storage State — Current**.
The value of both attributes should be Online.
3. If the Storage State — Desired is set to Read-only, complete the following steps:
 - a. Click the **Configuration** tab.
 - b. From the **Storage State — Desired** drop-down list, select **Online**.
 - c. Click **Apply Changes**.
 - d. Click the **Overview** tab and confirm that the values of **Storage State — Desired** and **Storage State — Current** are updated to Online.

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.

Related information

Choices

- Recovering from primary Admin Node failures
- Recovering from non-primary Admin Node failures

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.

 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

- [Copying audit logs from the failed primary Admin Node](#)
- [Replacing the primary Admin Node](#)
- [Configuring the replacement primary Admin Node](#)
- [Restoring the audit log on the recovered primary Admin Node](#)
- [Resetting the preferred sender on the recovered primary Admin Node](#)
- [Restoring the Admin Node database when recovering a primary Admin Node](#)
- [Restoring Prometheus metrics when recovering a primary Admin Node](#)

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.

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.



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.

1. Log in to the failed Admin Node if possible. Otherwise, log in to the primary Admin Node or another Admin Node, if available.
 - a. Enter the following command: `ssh admin@grid_node_IP`
 - b. Enter the password listed in the `Passwords.txt` file.
 - c. Enter the following command to switch to root: `su -`
 - d. Enter the password listed in the `Passwords.txt` file.

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

2. Stop the AMS service to prevent it from creating a new log file:`service ams stop`
3. Rename the audit.log file so that it does not overwrite the existing file when you copy it to the recovered Admin Node.

Rename audit.log to a unique numbered file name 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/exportls -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 physical or virtual hardware.

You can replace a failed primary Admin Node with a primary Admin Node running on the same platform, or you can replace a primary Admin Node running on VMware or a Linux host with a primary Admin Node hosted on a services appliance.

Use the procedure that matches the replacement platform you select for the node. After you complete the node replacement procedure (which is suitable for all node types), that procedure will direct you to the next step for primary Admin Node recovery.

Replacement platform	Procedure
VMware	Replacing a VMware node

Replacement platform	Procedure
Linux	Replacing a Linux node
SG100 and SG1000 services appliances	Replacing a services appliance
OpenStack	NetApp-provided virtual machine disk files and scripts for OpenStack are no longer supported for recovery operations. If you need to recover a node running in an OpenStack deployment, download the files for your Linux operating system. Then, follow the procedure for replacing a Linux node.

Configuring the replacement primary Admin Node

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

What you'll need

- For primary Admin Nodes hosted on virtual machines, the virtual machine must be deployed, powered on, and initialized.
- For primary Admin Nodes hosted on a services appliance, you have replaced the appliance and have installed software. See the installation guide for your appliance.

[SG100 & SG1000 services appliances](#)

- You must have the latest backup of the Recovery Package file (`sgws-recovery-package-id-revision.zip`).
- You must have the provisioning passphrase.

Steps

1. Open your web browser and navigate to `https://primary_admin_node_ip`.

NetApp® StorageGRID®

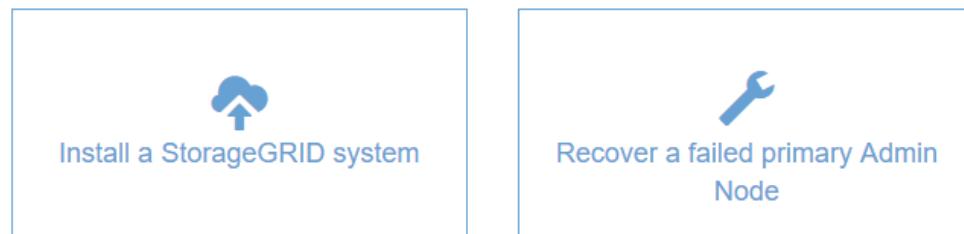
Help ▾

Install

Welcome

Use this page to install a new StorageGRID system, or recover a failed primary Admin Node for an existing system.

Info: Note: You must have access to a StorageGRID license, network configuration and grid topology information, and NTP settings to complete the installation. You must have the latest version of the Recovery Package file to complete a primary Admin Node recovery.



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.



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 information

[Administer StorageGRID](#)

[StorageGRID hotfix procedure](#)

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.

- 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.

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.

1. Log in to the recovered Admin Node:
 - a. Enter the following command: `ssh admin@recovery_Admin_Node_IP`
 - b. Enter the password listed in the `Passwords.txt` file.
 - c. Enter the following command to switch to root: `su -`
 - d. Enter the password listed in the `Passwords.txt` file.
- After you are logged in as root, the prompt changes from `$` to `#`.
2. Check which audit files have been preserved: `cd /var/local/audit/export`
3. Copy the preserved audit log files to the recovered Admin Node: `scp admin@grid_node_IP:/var/local/tmp/saved-audit-logs/YYYY* .`
- When prompted, enter the password for admin.
4. For security, delete the audit logs from the failed grid node after verifying that they have been copied successfully to the recovered Admin Node.
5. Update the user and group settings of the audit log files on the recovered Admin Node: `chown ams-user:bycast *`

6. Log out as root: `exit`

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

Related information

[Administer 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 alert notifications, alarm notifications, and AutoSupport messages, you must reconfigure this setting.

What you'll need

- You must be signed in to the Grid Manager using a supported browser.
- You must have specific access permissions.
- The recovered Admin Node must be installed and running.

Steps

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

Related information

[Administer StorageGRID](#)

Restoring the Admin Node database when recovering a primary Admin Node

If you want to retain the historical information about attributes, alarms, and alerts on a primary Admin Node that has failed, you can restore the Admin Node database. You can only restore this database if your StorageGRID system includes another Admin Node.

- 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.

If an Admin Node fails, the historical information stored in its Admin Node database is lost. This database includes the following information:

- Alert history
- Alarm history
- Historical attribute data, which is used in the charts and text reports available from the **Support > Tools > Grid Topology** page.

When you recover an Admin Node, the software installation process creates an empty Admin Node database on the recovered node. However, the new database only includes information for servers and services that are

currently part of the system or added later.

If you restored a primary Admin Node and your StorageGRID system has another Admin Node, you can restore the historical information by copying the Admin Node database from a non-primary Admin Node (the *source Admin Node*) to the recovered primary Admin Node. If your system has only a primary Admin Node, you cannot restore the Admin Node database.



Copying the Admin Node database might take several hours. Some Grid Manager features will be unavailable while services are stopped on the source Admin Node.

1. Log in to the source Admin Node:

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

2. From the source Admin Node, stop the MI service: `service mi stop`

3. From the source Admin Node, stop the Management Application Program Interface (mgmt-api) service: `service mgmt-api stop`

4. Complete the following steps on the recovered Admin Node:

- a. Log in to the recovered Admin Node:
 - i. Enter the following command: `ssh admin@grid_node_IP`
 - ii. Enter the password listed in the `Passwords.txt` file.
 - iii. Enter the following command to switch to root: `su -`
 - iv. Enter the password listed in the `Passwords.txt` file.
- b. Stop the MI service: `service mi stop`
- c. Stop the mgmt-api service: `service mgmt-api stop`
- d. Add the SSH private key to the SSH agent. Enter: `ssh-add`
- e. Enter the SSH Access Password listed in the `Passwords.txt` file.
- f. Copy the database from the source Admin Node to the recovered Admin Node:
`/usr/local/mi/bin/mi-clone-db.sh Source_Admin_Node_IP`
- g. When prompted, confirm that you want to overwrite the MI database on the recovered Admin Node.

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

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

5. Restart the services on the source Admin Node: `service servermanager start`

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.

- 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.

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.



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.

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.

Related information

[SG100 & SG1000 services appliances](#)

Steps

- [Copying audit logs from the failed non-primary Admin Node](#)
- [Replacing a non-primary Admin Node](#)
- [Selecting Start Recovery to configure a non-primary Admin Node](#)
- [Restoring the audit log on the recovered non-primary Admin Node](#)
- [Resetting the preferred sender on the recovered non-primary Admin Node](#)
- [Restoring the Admin Node database when recovering a non-primary Admin Node](#)
- [Restoring Prometheus metrics when recovering a non-primary Admin Node](#)

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.

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.



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.

1. Log in to the failed Admin Node if possible. Otherwise, log in to the primary Admin Node or another Admin Node, if available.

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

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

2. Stop the AMS service to prevent it from creating a new log file: `service ams stop`
3. Rename the `audit.log` file so that it does not overwrite the existing file when you copy it to the recovered Admin Node.

Rename `audit.log` to a unique numbered file name 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/exportls -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 physical or virtual hardware.

You can replace a failed non-primary Admin Node with a non-primary Admin Node running on the same platform, or you can replace a non-primary Admin Node running on VMware or a Linux host with a non-primary Admin Node hosted on a services appliance.

Use the procedure that matches the replacement platform you select for the node. After you complete the node replacement procedure (which is suitable for all node types), that procedure will direct you to the next step for non-primary Admin Node recovery.

Replacement platform	Procedure
VMware	Replacing a VMware node
Linux	Replacing a Linux node
SG100 and SG1000 services appliances	Replacing a services appliance

Replacement platform	Procedure
OpenStack	NetApp-provided virtual machine disk files and scripts for OpenStack are no longer supported for recovery operations. If you need to recover a node running in an OpenStack deployment, download the files for your Linux operating system. Then, follow the procedure for replacing a Linux node.

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.

What you'll need

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

Steps

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

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

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

Recovery

Select the failed grid node to recover, enter your provisioning passphrase, and then click Start Recovery to begin the recovery procedure.

Pending Nodes

	Name	IPv4 Address	State	Recoverable	
<input checked="" type="radio"/>	104-217-S1	10.96.104.217	Unknown		

Passphrase

Provisioning Passphrase

.....

Start Recovery

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



While the recovery procedure is running, you can click **Reset** to start a new recovery. An Info dialog box appears, indicating that the node will be left in an indeterminate state if you reset the procedure.



Reset Recovery

Resetting the recovery procedure leaves the deployed grid node in an indeterminate state. To retry a recovery after resetting the procedure, you must restore the node to a pre-installed state:

- For VMware nodes, delete the deployed VM and then redeploy it.
- For StorageGRID appliance nodes, run "sgareinstall" on the node.
- For Linux nodes, run "storagegrid node force-recovery *node-name*" on the Linux host.

Do you want to reset recovery?

Cancel

OK

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

- **VMware:** Delete the deployed virtual grid node. Then, when you are ready to restart the recovery, redeploy the node.
- **Linux:** Restart the node by running this command on the Linux host: `storagegrid node force-recovery node-name`
- **Appliance:** If you want to retry the recovery after resetting the procedure, you must restore the appliance node to a pre-installed state by running `sgareinstall` on the node.

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.



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

[Administer StorageGRID](#)

[Preparing an appliance for reinstallation \(platform replacement only\)](#)

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.

- 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.

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.

1. Log in to the recovered Admin Node:

- Enter the following command: `+ ssh admin@recovery_Admin_Node_IP`
- Enter the password listed in the `Passwords.txt` file.
- Enter the following command to switch to root: `su -`
- 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:broadcast *
```

6. Log out as root: `exit`

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

Related information

[Administer 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 alert notifications, alarm notifications, and AutoSupport messages, you must

reconfigure this setting in the StorageGRID system.

What you'll need

- You must be signed in to the Grid Manager using a supported browser.
- You must have specific access permissions.
- The recovered Admin Node must be installed and running.

Steps

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

Related information

[Administer StorageGRID](#)

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

If you want to retain the historical information about attributes, alarms, and alerts on a non-primary Admin Node that has failed, you can restore the Admin Node database from the primary Admin Node.

- 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.

If an Admin Node fails, the historical information stored in its Admin Node database is lost. This database includes the following information:

- Alert history
- Alarm history
- Historical attribute data, which is used in the charts and text reports available from the **Support > Tools > Grid Topology** page.

When you recover an Admin Node, the software installation process creates an empty Admin Node database on the recovered node. However, the new database only includes information for servers and services that are currently part of the system or added later.

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



Copying the Admin Node database might take several hours. Some Grid Manager features will be unavailable while services are stopped on the source node.

1. Log in to the source Admin Node:
 - a. Enter the following command: `ssh admin@grid_node_IP`
 - b. Enter the password listed in the `Passwords.txt` file.

- c. Enter the following command to switch to root: `su -`
- d. Enter the password listed in the `Passwords.txt` file.

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

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

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

6. Restart the services on the source Admin Node: `service servermanager start`

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.

- 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.

If an Admin Node fails, the metrics maintained in the Prometheus database on the Admin Node are lost. When you recover the Admin Node, the software installation process creates a new Prometheus database. After the recovered Admin Node is started, it records metrics as if you had performed a new installation of the StorageGRID system.

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



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

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:

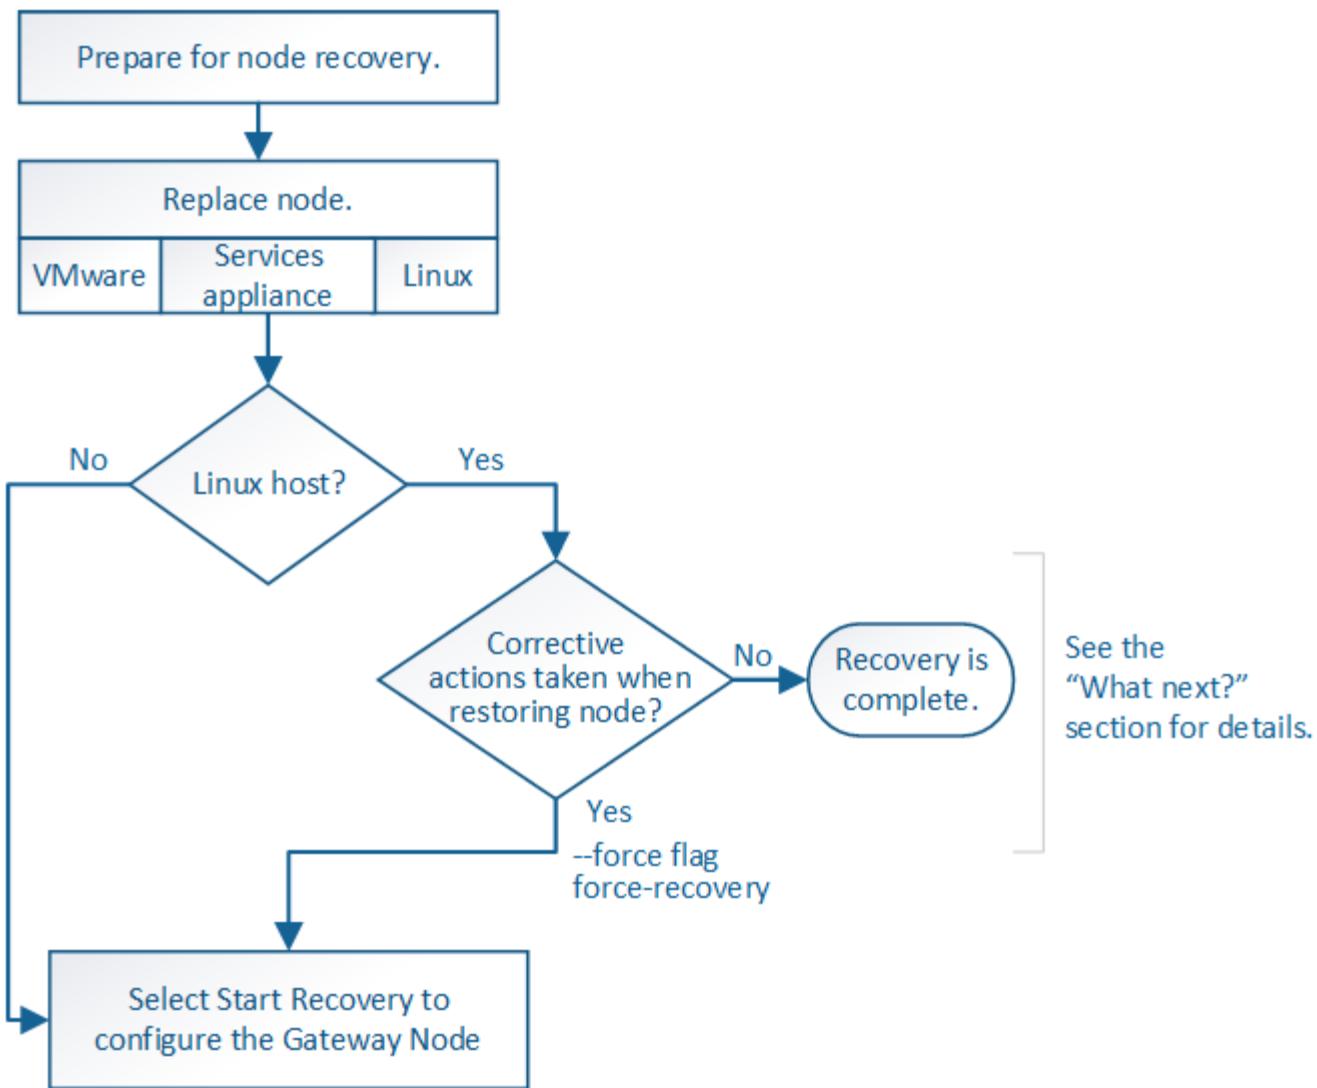
Database cloned, starting services

- 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 Gateway Node failures

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



Related information

[SG100 & SG1000 services appliances](#)

Steps

- [Replacing a Gateway Node](#)
- [Selecting Start Recovery to configure a Gateway Node](#)

Replacing a Gateway Node

You can replace a failed Gateway Node with a Gateway Node running on the same physical or virtual hardware, or you can replace a Gateway Node running on VMware or a Linux host with a Gateway Node hosted on a services appliance.

The node replacement procedure you must follow depends on which platform will be used by the replacement node. After you complete the node replacement procedure (which is suitable for all node types), that procedure will direct you to the next step for Gateway Node recovery.

Replacement platform	Procedure
VMware	Replacing a VMware node
Linux	Replacing a Linux node
SG100 and SG1000 services appliances	Replacing a services appliance
OpenStack	NetApp-provided virtual machine disk files and scripts for OpenStack are no longer supported for recovery operations. If you need to recover a node running in an OpenStack deployment, download the files for your Linux operating system. Then, follow the procedure for replacing a Linux node.

Selecting Start Recovery to configure a Gateway Node

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

What you'll need

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

Steps

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

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

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

Recovery

Select the failed grid node to recover, enter your provisioning passphrase, and then click Start Recovery to begin the recovery procedure.

Pending Nodes



	Name	IPv4 Address	State	Recoverable	
<input checked="" type="radio"/>	104-217-S1	10.96.104.217	Unknown	<input checked="" type="checkbox"/>	

Passphrase

Provisioning Passphrase

.....

Start Recovery

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



While the recovery procedure is running, you can click **Reset** to start a new recovery. An Info dialog box appears, indicating that the node will be left in an indeterminate state if you reset the procedure.

Info

Reset Recovery

Resetting the recovery procedure leaves the deployed grid node in an indeterminate state. To retry a recovery after resetting the procedure, you must restore the node to a pre-installed state:

- For VMware nodes, delete the deployed VM and then redeploy it.
- For StorageGRID appliance nodes, run "sgareinstall" on the node.
- For Linux nodes, run "storagegrid node force-recovery *node-name*" on the Linux host.

Do you want to reset recovery?

Cancel

OK

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

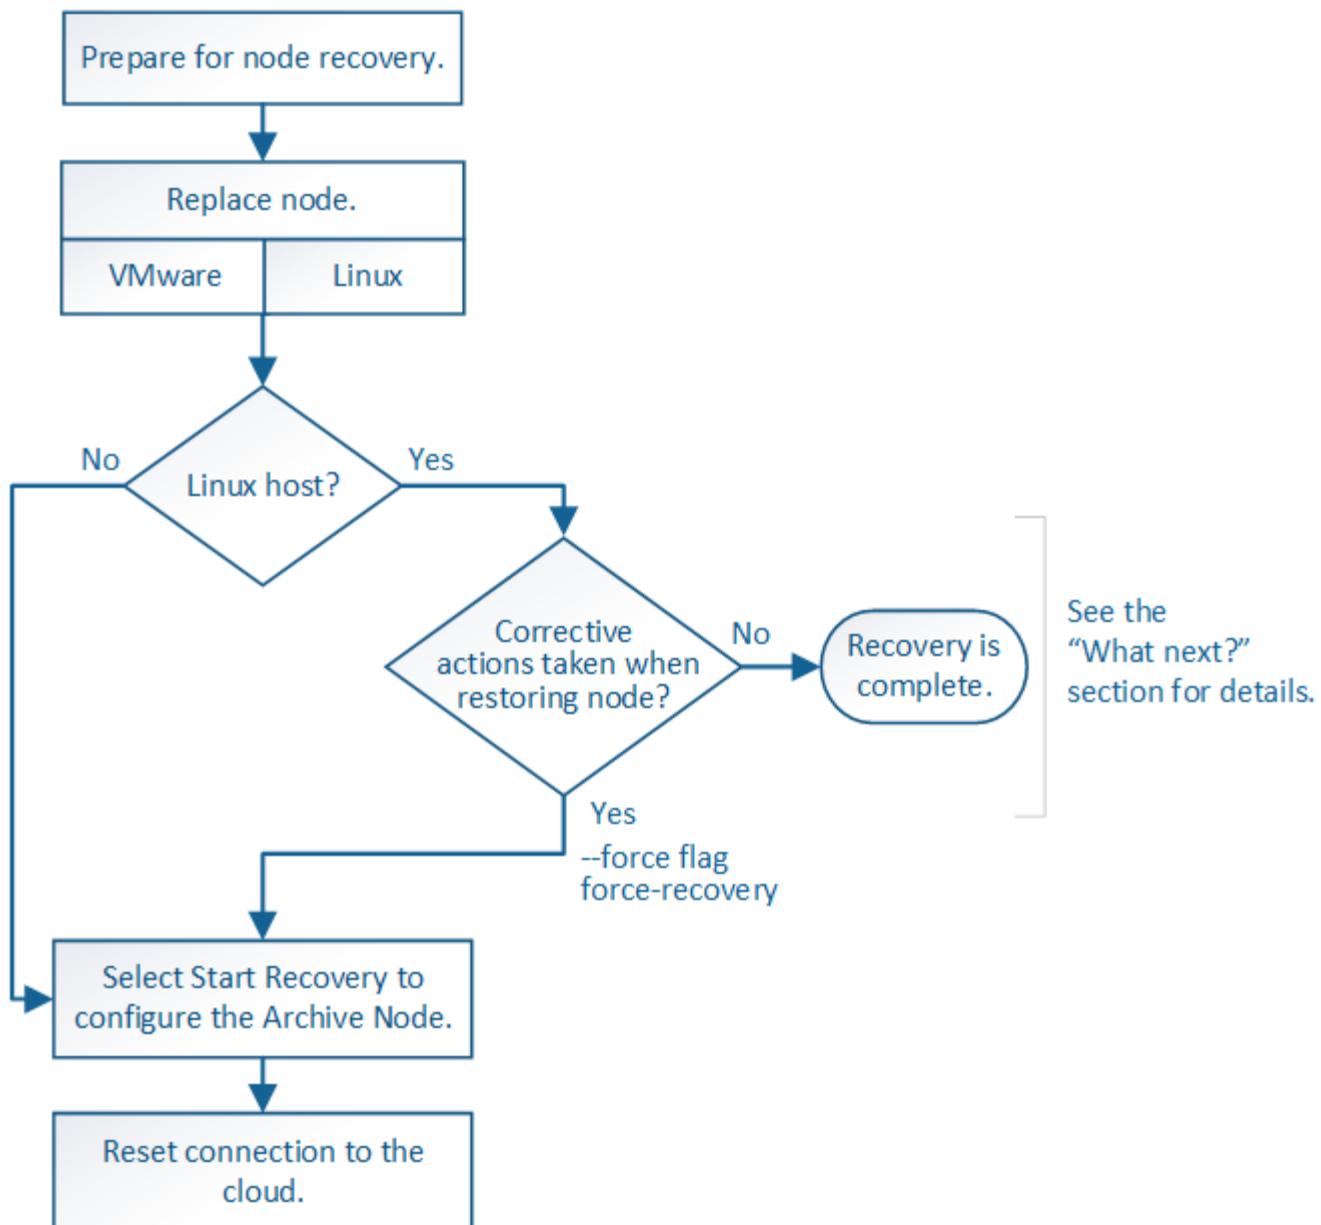
- **VMware:** Delete the deployed virtual grid node. Then, when you are ready to restart the recovery, redeploy the node.
- **Linux:** Restart the node by running this command on the Linux host: `storagegrid node force-recovery node-name`
- **Appliance:** If you want to retry the recovery after resetting the procedure, you must restore the appliance node to a pre-installed state by running `sgareinstall` on the node.

Related information

[Preparing an appliance for reinstallation \(platform replacement only\)](#)

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

- [Replacing an Archive Node](#)
- [Selecting Start Recovery to configure an Archive Node](#)
- [Resetting Archive Node connection to the cloud](#)

Replacing an Archive Node

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

You must select the node replacement procedure for your platform. The steps to replace a node are the same for all types of grid nodes.

Platform	Procedure
VMware	Replacing a VMware node
Linux	Replacing a Linux node
OpenStack	NetApp-provided virtual machine disk files and scripts for OpenStack are no longer supported for recovery operations. If you need to recover a node running in an OpenStack deployment, download the files for your Linux operating system. Then, follow the procedure for replacing a Linux node.

Selecting Start Recovery to configure an Archive Node

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

What you'll need

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

Steps

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

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

3. Enter the **Provisioning Passphrase**.

4. Click **Start Recovery**.

Recovery

Select the failed grid node to recover, enter your provisioning passphrase, and then click Start Recovery to begin the recovery procedure.

Pending Nodes

	Name	IPv4 Address	State	Recoverable	
<input checked="" type="radio"/>	104-217-S1...	10.96.104.217	Unknown		

Passphrase

Provisioning Passphrase

.....

Start Recovery

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



While the recovery procedure is running, you can click **Reset** to start a new recovery. An Info dialog box appears, indicating that the node will be left in an indeterminate state if you reset the procedure.

i Info

Reset Recovery

Resetting the recovery procedure leaves the deployed grid node in an indeterminate state. To retry a recovery after resetting the procedure, you must restore the node to a pre-installed state:

- For VMware nodes, delete the deployed VM and then redeploy it.
- For StorageGRID appliance nodes, run "sgareinstall" on the node.
- For Linux nodes, run "storagegrid node force-recovery *node-name*" on the Linux host.

Do you want to reset recovery?

Cancel

OK

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.



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

What you'll need

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

Steps

1. Select **Support > Tools > Grid Topology**.
2. Select **Archive Node > ARC > Target**.
3. Edit the **Access Key** field by entering an incorrect value and click **Apply Changes**.
4. Edit the **Access Key** field by entering the correct value and click **Apply Changes**.

All grid node types: Replacing a VMware node

When you recover a failed StorageGRID node that was hosted on VMware, you must remove the failed node and deploy a recovery node.

What you'll need

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

About this task

You use the VMware vSphere Web Client to first remove the virtual machine associated with the failed grid node. Then, you can deploy a new virtual machine.

This procedure is only one step in the grid node recovery process. The node removal and deployment procedure is the same for all VMware nodes, including Admin Nodes, Storage Nodes, Gateway Nodes, and Archive Nodes.

Steps

1. Log in to VMware vSphere Web Client.
2. Navigate to the failed grid node virtual machine.
3. Make a note of all of the information required to deploy the recovery node.
 - a. Right-click the virtual machine, select the **Edit Settings** tab, and note the settings in use.
 - b. Select the **vApp Options** tab to view and record the grid node network settings.
4. If the failed grid node is a Storage Node, determine if any of the virtual hard disks used for data storage are undamaged and preserve them for reattachment to the recovered grid node.
5. Power off the virtual machine.
6. Select **Actions > All vCenter Actions > Delete from Disk** to delete the virtual machine.
7. Deploy a new virtual machine to be the replacement node, and connect it to one or more StorageGRID networks.

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



After deploying the new node, you can add new virtual disks according to your storage requirements, reattach any virtual hard disks preserved from the previously removed failed grid node, or both.

For instructions:

[Install VMware](#) > Deploying a StorageGRID node as a virtual machine

8. Complete the node recovery procedure, based on the type of node you are recovering.

Type of node	Go to
Primary Admin Node	Configuring the replacement primary Admin Node
Non-primary Admin Node	Selecting Start Recovery to configure a non-primary Admin Node
Gateway Node	Selecting Start Recovery to configure a Gateway Node
Storage Node	Selecting Start Recovery to configure a Storage Node
Archive Node	Selecting Start Recovery to configure an Archive Node

All grid node types: Replacing a Linux node

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

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

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

If more than one grid node is hosted on a physical or virtual Linux host, you can recover the grid nodes in any order. However, recovering a primary Admin Node first, if present, prevents the recovery of other grid nodes from stalling as they try to contact the primary Admin Node to register for recovery.

1. [Deploying new Linux hosts](#)
2. [Restoring grid nodes to the host](#)
3. [What's next: Performing additional recovery steps, if required](#)

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 for 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.



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



Never use FabricPool to tier any data related to StorageGRID back to StorageGRID itself. Tiering StorageGRID data back to StorageGRID increases troubleshooting and operational complexity.

Related information

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

[Install Ubuntu or Debian](#)

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

- [Restoring and validating grid nodes](#)
- [Starting the StorageGRID host service](#)
- [Recovering nodes that fail to start normally](#)

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 data volumes, as described in the StorageGRID installation instructions for your Linux operating system. Importing the node restores its node configuration file to the host.

If it is not possible to import missing nodes, you must 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
<hr/>	
dc1-adm1	<code>/dev/mapper/sgws-adm1-var-local</code>
dc1-gw1	<code>/dev/mapper/sgws-gw1-var-local</code>
dc1-sn1	<code>/dev/mapper/sgws-sn1-var-local</code>
dc1-arc1	<code>/dev/mapper/sgws-arc1-var-local</code>

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`



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.



When you re-create the configuration file for a node, you must use the same name for the replacement node that was used for the node you are recovering. For Linux deployments, ensure that the configuration file name contains the node name. You should use the same network interfaces, block device mappings, and IP addresses when possible. This practice minimizes the amount of data that needs to be copied to the node during recovery, which could make the recovery significantly faster (in some cases, minutes rather than weeks).



If you use any new block devices (devices that the StorageGRID node did not use previously) as values for any of the configuration variables that start with `BLOCK_DEVICE_` when you are 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 information

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

[Install Ubuntu or Debian](#)

[Fixing missing network interface errors](#)

[Fixing missing block device errors](#)

What's next: Performing additional recovery steps, if required

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: GRID_NETWORK_TARGET = host-interface-name `node-name: Interface 'host-interface-name' does not exist'
```

The error could be reported for the Grid Network, the Admin Network, or the Client Network. This error means that the `/etc/storagegrid/nodes/node-name.conf` file maps the indicated StorageGRID network to the host interface named `host-interface-name`, but there is no interface with that name on the current host.

If you receive this error, verify that you completed the steps in “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 the `GRID_NETWORK_TARGET`, the `ADMIN_NETWORK_TARGET`, or the `CLIENT_NETWORK_TARGET` to match an existing host interface.

Make sure the host interface provides access to the appropriate physical network port or VLAN, and that the interface does not directly reference a bond or bridge device. You must either configure a VLAN (or other virtual interface) on top of the bond device on the host, or use a bridge and virtual Ethernet (veth) pair.

Related information

[Deploying new Linux hosts](#)

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.

 If you must provide a new block storage device for any of the configuration file variables starting with `BLOCK_DEVICE_` because the original block device was lost with the failed host, ensure the new block device is unformatted before attempting further recovery procedures. The new block device will be unformatted if you are using shared storage and have created a new volume. If you are unsure, run the following command against any new block storage device special files.

 Run the following command only for new block storage devices. 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 information

[Deploying new Linux hosts](#)

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

[Install Ubuntu or Debian](#)

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.

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
```



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.



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

Related information

[What's next: Performing additional recovery steps, if required](#)

What's next: 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

Type of recovery	Next step
Non-primary Admin Node	Selecting Start Recovery to configure a non-primary Admin Node
Gateway Node	Selecting Start Recovery to configure a Gateway Node
Archive Node	Selecting Start Recovery to configure an Archive Node
Storage Node (software-based): <ul style="list-style-type: none"> • If you had to use the <code>--force</code> flag to import the node, or you issued <code>storagegrid node force-recovery node-name</code> • If you had to do a full node reinstall, or you needed to restore <code>/var/local</code> 	Selecting Start Recovery to configure a Storage Node
Storage Node (software-based): <ul style="list-style-type: none"> • If you added a new block device. • If, for any <PURPOSE>, the value of the <code>BLOCK_DEVICE_<PURPOSE></code> 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

Replacing a failed node with a services appliance

You can use an SG100 or SG1000 services appliance to recover a failed Gateway Node, a failed non-primary Admin Node, or a failed primary Admin Node that was hosted on VMware, a Linux host, or a services appliance. This procedure is one step of the grid node recovery procedure.

What you'll need

- You must have determined that one of the following situations is true:
 - The virtual machine hosting the node cannot be restored.
 - The physical or virtual Linux host for the grid node has failed, and must be replaced.
 - The services appliance hosting the grid node must be replaced.
- You must make sure that the StorageGRID Appliance Installer version on the services appliance matches the software version of your StorageGRID system, as described in hardware installation and maintenance for verifying and upgrading the StorageGRID Appliance Installer version.

[SG100 & SG1000 services appliances](#)



Do not deploy both an SG100 and an SG1000 service appliance in the same site. Unpredictable performance might result.

About this task

You can use an SG100 or SG1000 services appliance to recover a failed grid node in the following cases:

- The failed node was hosted on VMware or Linux (platform change)
- The failed node was hosted on a services appliance (platform replacement)

Steps

- [Installing a services appliance \(platform change only\)](#)
- [Preparing an appliance for reinstallation \(platform replacement only\)](#)
- [Starting software installation on a services appliance](#)
- [Monitoring services appliance installation](#)

Installing a services appliance (platform change only)

When you are recovering a failed grid node that was hosted on VMware or a Linux host and you are using an SG100 or SG1000 services appliance for the replacement node, you must first install the new appliance hardware using the same node name as the failed node.

You must have the following information about the failed node:

- **Node name:** You must install the services appliance using the same node name as the failed node.
- **IP addresses:** You can assign the services appliance the same IP addresses as the failed node, which is the preferred option, or you can select a new unused IP address on each network.

Perform this procedure only if you are recovering a failed node that was hosted on VMware or Linux and are replacing it with a node hosted on a services appliance.

1. Follow the instructions for installing a new SG100 or SG1000 services appliance.
2. When prompted for a node name, use the node name of the failed node.

Related information

[SG100 & SG1000 services appliances](#)

Preparing an appliance for reinstallation (platform replacement only)

When recovering a grid node that was hosted on a services appliance, you must first prepare the appliance for reinstallation of StorageGRID software.

Perform this procedure only if you are replacing a failed node that was hosted on a services appliance. Do not follow these steps if the failed node was originally hosted on VMware or a Linux host.

1. Log in to the failed grid node:
 - a. Enter the following command: `ssh admin@grid_node_IP`

- b. Enter the password listed in the `Passwords.txt` file.
- c. Enter the following command to switch to root: `su -`
- d. Enter the password listed in the `Passwords.txt` file.

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

2. Prepare the appliance for the installation of StorageGRID software. Enter: `sgareinstall`

3. When prompted to continue, enter: `y`

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

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

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

Starting software installation on a services appliance

To install a Gateway Node or Admin Node on an SG100 or SG1000 services appliance, you use the StorageGRID Appliance Installer, which is included on the appliance.

What you'll need

- The appliance must be installed in a rack, connected to your networks, and powered on.
- Network links and IP addresses must be configured for the appliance using the StorageGRID Appliance Installer.
- If you are installing a Gateway Node or non-primary Admin Node, you know the IP address of the primary Admin Node for the StorageGRID grid.
- All Grid Network subnets listed on the IP Configuration page of the StorageGRID Appliance Installer must be 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 an SG100 or SG1000 services appliance.

- You must be using a supported web browser.
- You must know one of the IP addresses assigned to the appliance. You can use the IP address for the Admin Network, the Grid Network, or the Client Network.
- If you are installing a primary Admin Node, you have the Ubuntu or Debian install files for this version of StorageGRID available.



A recent version of StorageGRID software is preloaded onto the services appliance during manufacturing. If the preloaded version of software matches the version being used in your StorageGRID deployment, you do not need the installation files.

About this task

To install StorageGRID software on an SG100 or SG1000 services appliance:

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

Steps

1. Open a browser and enter one of the IP addresses for the SG100 or SG1000 services appliance.

`https://Controller_IP:8443`

The StorageGRID Appliance Installer Home page appears.

NetApp® StorageGRID® Appliance Installer

Help ▾

Home Configure Networking ▾ Configure Hardware ▾ Monitor Installation Advanced ▾

Home

This Node

Node type	Gateway
Node name	NetApp-SGA

Primary Admin Node connection

Enable Admin Node discovery	<input checked="" type="checkbox"/> Uncheck to manually enter the Primary Admin Node IP
Connection state	Admin Node discovery is in progress

Installation

Current state	Unable to start installation. The Admin Node connection is not ready.
---------------	---

2. To install a Primary Admin Node:

- In the This Node section, for **Node Type**, select **Primary Admin**.
- In the **Node Name** field, enter the same name that was used for the node you are recovering, and click **Save**.
- In the Installation section, check the software version listed under Current state

If the version of software that is ready to install is correct, skip ahead to the [Installation step](#).

- If you need to upload a different version of software, under the **Advanced** menu, select **Upload StorageGRID Software**.

The [Upload StorageGRID Software](#) page appears.

Upload StorageGRID Software

If this node is the primary Admin Node of a new deployment, you must use this page to upload the StorageGRID software installation package, unless the version of the software you want to install has already been uploaded. If you are adding this node to an existing deployment, you can avoid network traffic by uploading the installation package that matches the software version running on the existing grid. If you do not upload the correct package, the node obtains the software from the grid's primary Admin Node during installation.

Current StorageGRID Installation Software

Version None

Package Name None

Upload StorageGRID Installation Software

Software Package

[Browse](#)

Checksum File

[Browse](#)

- e. Click **Browse** to upload the **Software Package** and **Checksum File** for StorageGRID software.

The files are automatically uploaded after you select them.

- f. Click **Home** to return to the StorageGRID Appliance Installer Home page.

3. To install a Gateway Node or non-Primary Admin Node:

- a. In the This Node section, for **Node Type**, select **Gateway** or **Non-Primary Admin**, depending on the type of node you are restoring.
- b. In the **Node Name** field, enter the same name that was used for the node you are recovering, and click **Save**.
- c. In the Primary Admin Node connection section, determine whether you need to specify the IP address for the primary Admin Node.

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

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

Option	Description
Manual IP entry	<ol style="list-style-type: none"> 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."

Option	Description
Automatic discovery of all connected primary Admin Nodes	<ol style="list-style-type: none"> <li data-bbox="638 164 1454 196">a. Select the Enable Admin Node discovery check box. <li data-bbox="638 213 1454 308">b. From the list of discovered IP addresses, select the primary Admin Node for the grid where this services appliance will be deployed. <li data-bbox="638 325 817 356">c. Click Save. <li data-bbox="638 373 1470 447">d. Wait while the connection state for the new IP address becomes “ready.”

4. In the Installation section, confirm that the current state is Ready to start installation of node name and that the **Start Installation** button is enabled.

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

5. From the StorageGRID Appliance Installer home page, click **Start Installation**.

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



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

Related information

[SG100 & SG1000 services appliances](#)

Monitoring services appliance installation

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

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

The Monitor Installation page shows the installation progress.

Monitor Installation

1. Configure storage	Complete	
2. Install OS	Running	
Step	Progress	Status
Obtain installer binaries	<div style="width: 100%; background-color: green;"></div>	Complete
Configure installer	<div style="width: 100%; background-color: green;"></div>	Complete
Install OS	<div style="width: 50%; background-color: #0072BD;"></div>	Installer VM running
3. Install StorageGRID	Pending	
4. Finalize installation	Pending	

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



The installer ensures that tasks completed in a previous install 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 clears any existing configuration from the drives, and configures host settings.

- **2. Install OS**

During this stage, the installer copies the base operating system image for StorageGRID from the primary Admin Node to the appliance or installs the base operating system from the installation package for the primary Admin Node.

3. Continue monitoring the installation progress until one of the following occurs:

- For appliance Gateway Nodes or non-primary appliance Admin Nodes, the **Install StorageGRID** stage pauses and a message appears on the embedded console, prompting you to approve this node on the Admin Node using the Grid Manager.

Monitor Installation

1. Configure storage	Complete
2. Install OS	Complete
3. Install StorageGRID	Running
4. Finalize installation	Pending

Connected (unencrypted) to: QEMU

```
/platform.type=: Device or resource busy
[2017-07-31T22:09:12.362566]     INFO -- [INSG] NOTICE: seeding /var/local with c
ontainer data
[2017-07-31T22:09:12.366205]     INFO -- [INSG] Fixing permissions
[2017-07-31T22:09:12.369633]     INFO -- [INSG] Enabling syslog
[2017-07-31T22:09:12.511533]     INFO -- [INSG] Stopping system logging: syslog-n
g.
[2017-07-31T22:09:12.570096]     INFO -- [INSG] Starting system logging: syslog-n
g.
[2017-07-31T22:09:12.576360]     INFO -- [INSG] Beginning negotiation for downloa
d of node configuration
[2017-07-31T22:09:12.581363]     INFO -- [INSG]
[2017-07-31T22:09:12.585066]     INFO -- [INSG]
[2017-07-31T22:09:12.588314]     INFO -- [INSG]
[2017-07-31T22:09:12.591851]     INFO -- [INSG]
[2017-07-31T22:09:12.594886]     INFO -- [INSG]
[2017-07-31T22:09:12.598360]     INFO -- [INSG]
[2017-07-31T22:09:12.601324]     INFO -- [INSG]
[2017-07-31T22:09:12.604759]     INFO -- [INSG]
[2017-07-31T22:09:12.607800]     INFO -- [INSG]
[2017-07-31T22:09:12.610985]     INFO -- [INSG]
[2017-07-31T22:09:12.614597]     INFO -- [INSG]
[2017-07-31T22:09:12.618282]     INFO -- [INSG] Please approve this node on the A
dmin Node GMI to proceed...
```

- For appliance primary Admin Nodes, a fifth phase (Load StorageGRID Installer) appears. If the fifth phase is in progress for more than 10 minutes, refresh the page manually.

NetApp® StorageGRID® Appliance Installer

Help ▾

Home Configure Networking ▾ Configure Hardware ▾ Monitor Installation Advanced ▾

Monitor Installation

1. Configure storage	Complete	
2. Install OS	Complete	
3. Install StorageGRID	Complete	
4. Finalize installation	Complete	
5. Load StorageGRID Installer	Running	
Step	Progress	Status
Starting StorageGRID Installer		Do not refresh. You will be redirected when the installer is ready

4. Go on to the next step of the recovery process for the type of appliance grid node that you are recovering.

Type of recovery	Reference
Gateway Node	Selecting Start Recovery to configure a Gateway Node
Non-primary Admin Node	Selecting Start Recovery to configure a non-primary Admin Node
Primary Admin Node	Configuring the replacement primary Admin Node

Copyright information

Copyright © 2025 NetApp, Inc. All Rights Reserved. Printed in the U.S. No part of this document covered by copyright may be reproduced in any form or by any means—graphic, electronic, or mechanical, including photocopying, recording, taping, or storage in an electronic retrieval system—without prior written permission of the copyright owner.

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

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

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

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

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

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

Trademark information

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