

Expand a grid

StorageGRID 11.8

NetApp May 10, 2024

This PDF was generated from https://docs.netapp.com/us-en/storagegrid-118/expand/index.html on May 10, 2024. Always check docs.netapp.com for the latest.

Table of Contents

Expand a grid	1
Expand a grid: Overview	1
Plan StorageGRID expansion	2
Gather required materials	12
Add storage volumes	18
Add grid nodes or site	26
Configure expanded system	40
Troubleshoot expansion	49

Expand a grid

Expand a grid: Overview

You can expand the capacity or capabilities of your StorageGRID system without interrupting system operations.

A StorageGRID expansion allows you to add:

- Storage volumes to Storage Nodes
- New grid nodes to an existing site
- An entire new site

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

Follow the steps for the type of expansion you're performing:

Add storage volumes

Follow the steps for adding storage volumes to Storage Nodes.

Add grid nodes

- 1. Follow the steps for adding grid nodes to an existing site.
- 2. Update the subnets.
- 3. Deploy grid nodes:
 - Appliances
 - VMware
 - Linux



"Linux" refers to a Red Hat Enterprise Linux, Ubuntu, or Debian deployment. For a list of supported versions, see the NetApp Interoperability Matrix Tool (IMT).

- 4. Perform the expansion.
- 5. Configure the expanded system.

Add new site

- 1. Follow the steps for Adding a new site.
- 2. Update the subnets.
- 3. Deploy grid nodes:
 - Appliances
 - VMware
 - Linux



"Linux" refers to a Red Hat Enterprise Linux, Ubuntu, or Debian deployment. For a list of supported versions, see the NetApp Interoperability Matrix Tool (IMT).

- 4. Perform the expansion.
- 5. Configure the expanded system.

Plan StorageGRID expansion

Add storage capacity

Guidelines for adding object capacity

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

Guidelines for adding storage volumes

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

- You must examine your current ILM rules to determine where and when to add storage volumes to increase the storage available for replicated objects or erasure-coded objects.
- You can't increase the metadata capacity of your system by adding storage volumes because object metadata is stored only on volume 0.
- Each software-based Storage Node can support a maximum of 16 storage volumes. If you need to add capacity beyond that, you must add new Storage Nodes.
- You can add one or two expansion shelves to each SG6060 appliance. Each expansion shelf adds 16 storage volumes. With both expansion shelves installed, the SG6060 can support a total of 48 storage volumes.
- You can't add storage volumes to any other storage appliance.
- You can't increase the size of an existing storage volume.
- You can't add storage volumes to a Storage Node at the same time you are performing a system upgrade, recovery operation, or another expansion.

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

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

Guidelines for adding Storage Nodes

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

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

Guidelines for ADC service on Storage Nodes

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

- The StorageGRID system requires a quorum of ADC services to be available at each site and at all times.
- At least three Storage Nodes at each site must include the ADC service.

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

Add storage capacity for replicated objects

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

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

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

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

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

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

Add storage capacity for erasure-coded objects

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

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

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

General recommendation for adding storage capacity for erasure-coded objects

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

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

To better understand the factors that led to this recommendation or to develop a more precise plan for your site, see Considerations for rebalancing erasure-coded data. For a custom recommendation optimized for your situation, contact your NetApp Professional Services consultant.

Considerations for rebalancing erasure-coded data

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

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

What is EC rebalancing?

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

The EC rebalance procedure:

- Only moves erasure-coded object data. It does not move replicated object data.
- Redistributes the data within a site. It does not move data between sites.
- Redistributes data among all Storage Nodes at a site. It does not redistribute data within storage volumes.
- Does not consider the replicated data usage on each Storage Node when determining where to move erasure-coded data.
- Redistributes erasure-coded data evenly between Storage Nodes without considering the relative capacities of each node.
- Will not distribute erasure-coded data to Storage Nodes that are more than 80% full.
- Might decrease the performance of ILM operations and S3 and Swift client operations when it runs—additional resources are required to redistribute the erasure-coding fragments.

When the EC rebalance procedure is complete:

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

end up with approximately the same amount of erasure-coded data.

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



When to rebalance erasure-coded data

Consider the following scenario:

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



Rebalance is not required if you add enough nodes

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



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

Rebalance is required if you can't add enough nodes

1

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



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



Recommendations for EC rebalancing

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

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

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

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

How EC rebalance procedure interacts with other maintenance tasks

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

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

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

How EC rebalance procedure interacts with ILM

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

Add metadata capacity

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

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

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

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

For more information, see the following:

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

Guidelines for increasing metadata capacity

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

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

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

See the description of what Metadata Reserved Space is.

How metadata is redistributed when you add Storage Nodes

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

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

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



a _ .

Reserved metadata space

Add grid nodes to add capabilities to your system

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

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

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

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

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

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

during an expansion.

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

Add a new site

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

Guidelines for adding a site

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

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

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

For more information, see the following:

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

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

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

Gather required materials

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

Item	Notes
StorageGRID installation archive	If you are adding new grid nodes or a new site, you must download and extract the StorageGRID installation archive. You must use the same version that is currently running on the grid.
	For details, see the instructions for downloading and extracting the StorageGRID installation files.
	Note: You don't need to download files if you are adding new storage volumes to existing Storage Nodes or installing a new StorageGRID appliance.
Service laptop	The service laptop has the following:
	Network port
	 SSH client (for example, PuTTY)
	Supported web browser
Passwords.txt file	Contains the passwords required to access grid nodes on the command line. Included in the Recovery Package.
Provisioning passphrase	The passphrase is created and documented when the StorageGRID system is first installed. The provisioning passphrase is not in the Passwords.txt file.
StorageGRID documentation	Administer StorageGRID
	Release notes
	 Installation instructions for your platform
	 Install StorageGRID on Red Hat Enterprise Linux
	 Install StorageGRID on Ubuntu or Debian
	 Install StorageGRID on VMware
Current documentation for your platform	For supported versions, see the Interoperability Matrix Tool (IMT).

Download and extract the StorageGRID installation files

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

About this task

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

Steps

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

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

Use the .zip file if you are running ${\tt V}$	Nindows on the service laptop.
--	--------------------------------

Platform	Installation archive
Red Hat Enterprise Linux	StorageGRID-Webscale-version-RPM-uniqueID.zip
	StorageGRID-Webscale-version-RPM-uniqueID.tgz
Ubuntu or Debian or Appliances	StorageGRID-Webscale-version-DEB-uniqueID.zip
	StorageGRID-Webscale-version-DEB-uniqueID.tgz
VMware	StorageGRID-Webscale-version-VMware-uniqueID.zip
	StorageGRID-Webscale-version-VMware-uniqueID.tgz
OpenStack/other Hypervisor	To expand an existing deployment on OpenStack, you must deploy a virtual machine running one of the supported Linux distributions listed above and follow the appropriate instructions for Linux.

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

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

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

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

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

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

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

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

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

Path and file name	Description
./vsphere/README	A text file that describes all of the files contained in the StorageGRID download file.
./vsphere/NLF000000.txt	A free license that does not provide any support entitlement for the product.
./vsphere/NetApp-SG-version-SHA.vmdk	The virtual machine disk file that is used as a template for creating grid node virtual machines.
<pre>./vsphere/vsphere-primary-admin.ovf ./vsphere/vsphere-primary-admin.mf</pre>	The Open Virtualization Format template file (.ovf) and manifest file (.mf) for deploying the primary Admin Node.
<pre>./vsphere/vsphere-non-primary-admin.ovf ./vsphere/vsphere-non-primary-admin.mf</pre>	The template file $(.ovf)$ and manifest file $(.mf)$ for deploying non-primary Admin Nodes.
<pre>./vsphere/vsphere-archive.ovf ./vsphere/vsphere-archive.mf</pre>	The template file $(.ovf)$ and manifest file $(.mf)$ for deploying Archive Nodes.
<pre>./vsphere/vsphere-gateway.ovf ./vsphere/vsphere-gateway.mf</pre>	The template file $(.ovf)$ and manifest file $(.mf)$ for deploying Gateway Nodes.
<pre>./vsphere/vsphere-storage.ovf ./vsphere/vsphere-storage.mf</pre>	The template file (.ovf) and manifest file (.mf) for deploying virtual machine-based Storage Nodes.
Deployment scripting tool	Description
./vsphere/deploy-vsphere-ovftool.sh	A Bash shell script used to automate the deployment of virtual grid nodes.

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

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

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

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



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

Verify hardware and networking

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

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

See Internal grid node communications.

- The primary Admin Node can communicate with all expansion servers that are intended to host the StorageGRID system.
- If any of the new nodes has a Grid Network IP address on a subnet not previously used, you have already added the new subnet to the Grid Network subnet list. Otherwise, you will have to cancel the expansion, add the new subnet, and start the procedure again.
- You aren't using network address translation (NAT) on the Grid Network between grid nodes or between StorageGRID sites. When you use private IPv4 addresses for the Grid Network, those addresses must be directly routable from every grid node at every site. Using NAT to bridge the Grid Network across a public network segment is supported only if you use a tunneling application that is transparent to all nodes in the grid, meaning the grid nodes require no knowledge of public IP addresses.

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

Add storage volumes

Add storage volumes to Storage Nodes

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

Before you begin

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



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

About this task

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



Object metadata is always stored in object store 0.

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

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

Steps

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

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

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

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

Volumes

Mount point 💡 💠	Device 🔞 🗘	Status 😮 ≑	Size 😧 💠	Available 🕝 ≑	Write cache status 🥝 💠
1	croot	Online	21.00 GB	14.73 GB 👖	Unknown
/var/local	cvloc	Online	85.86 GB	80.94 GB 📊	Unknown
/var/local/rangedb/0	sdc	Online	107.32 GB	107.17 GB 1	Enabled
/var/local/rangedb/1	sdd	Online	107.32 GB	107.18 GB 1	Enabled
/var/local/rangedb/2	sde	Online	107.32 GB	107.18 GB 1	Enabled
/var/local/rangedb/3	sdf	Online	107.32 GB	107.18 GB 1	Enabled
/var/local/rangedb/4	sdg	Online	107.32 GB	107.18 GB	Enabled

Object stores

ID 😧 韋	Size 😢 💠	Available 💡 💠	Replicated data 💡 💠	EC data 💡 🜲	Object data (%) 💡 💠	Health 😮 💠
0000	107.32 GB	96.44 GB 1	1.55 MB 1	0 bytes 📊	0.00%	No Errors
0001	107.32 GB	107.18 GB	0 bytes 👖	0 bytes 👖	0.00%	No Errors
0002	107.32 GB	107.18 GB 11	0 bytes 👖	0 bytes 👖	0.00%	No Errors
0003	107.32 GB	107.18 GB	0 bytes 👖	0 bytes II	0.00%	No Errors
0004	107.32 GB	107.18 GB 1	0 bytes 👖	0 bytes 👖	0.00%	No Errors

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

VMware: Add storage volumes to Storage Node

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

Before you begin

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



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

About this task

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

Steps

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

The hard disks are typically configured as Virtual Machine Disks (VMDKs). VMDKs are more commonly used and are easier to manage, while RDMs might provide better performance for workloads that use larger object sizes (for example, greater than 100 MB). For more information about adding hard disks to virtual machines, see the VMware vSphere documentation.

3. Restart the virtual machine by using the **Restart Guest OS** option in the VMware vSphere Client, or by entering the following command in an ssh session to the virtual machine:sudo reboot



Don't use **Power Off** or **Reset** to restart the virtual machine.

- 4. Configure the new storage for use by the Storage Node:
 - a. Log in to the grid node:
 - i. Enter the following command: ssh admin@grid node IP
 - ii. Enter the password listed in the Passwords.txt file.
 - iii. Enter the following command to switch to root: su -
 - iv. Enter the password listed in the Passwords.txt file. When you are logged in as root, the prompt changes from \$ to #.

b. Configure the new storage volumes:

sudo add_rangedbs.rb

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

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

The setup_rangedbs.sh script runs automatically.

- 5. Check that the services start correctly:
 - a. View a list of the status of all services on the server:

sudo storagegrid-status

The status is updated automatically.

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

Ctrl+C

- 6. Verify that the Storage Node is online:
 - a. Sign in to the Grid Manager using a supported web browser.
 - b. Select SUPPORT > Tools > Grid topology.
 - c. Select *site > Storage Node > LDR > Storage*.
 - d. Select the **Configuration** tab and then the **Main** tab.
 - e. If the Storage State Desired drop-down list is set to Read-only or Offline, select Online.
 - f. Select Apply Changes.
- 7. To see the new object stores:
 - a. Select NODES > site > Storage Node > Storage.
 - b. View the details in the **Object Stores** table.

Result

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

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

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

Before you begin

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



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

About this task

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

Steps

1. Install the new storage hardware.

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

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

You might have to rescan for devices.

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

sudo storagegrid node stop <node-name>

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

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

```
NODE_TYPE = VM_Storage_Node
ADMIN_IP = 10.1.0.2
BLOCK_DEVICE_VAR_LOCAL = /dev/mapper/sgws-sn1-var-local
BLOCK_DEVICE_RANGEDB_00 = /dev/mapper/sgws-sn1-rangedb-0
BLOCK_DEVICE_RANGEDB_01 = /dev/mapper/sgws-sn1-rangedb-1
BLOCK_DEVICE_RANGEDB_02 = /dev/mapper/sgws-sn1-rangedb-2
BLOCK_DEVICE_RANGEDB_03 = /dev/mapper/sgws-sn1-rangedb-3
GRID_NETWORK_TARGET = bond0.1001
ADMIN_NETWORK_TARGET = bond0.1002
CLIENT_NETWORK_TARGET = bond0.1003
GRID_NETWORK_IP = 10.1.0.3
GRID_NETWORK_MASK = 255.255.0
GRID_NETWORK_GATEWAY = 10.1.0.1
```

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

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

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

```
NODE TYPE = VM Storage Node
ADMIN IP = 10.1.0.2
BLOCK DEVICE VAR LOCAL = /dev/mapper/sgws-sn1-var-local
BLOCK DEVICE RANGEDB 00 = /dev/mapper/sgws-sn1-rangedb-0
BLOCK DEVICE RANGEDB 01 = /dev/mapper/sgws-sn1-rangedb-1
BLOCK DEVICE RANGEDB 02 = /dev/mapper/sqws-snl-rangedb-2
BLOCK DEVICE RANGEDB 03 = /dev/mapper/sqws-snl-rangedb-3
BLOCK DEVICE RANGEDB 04 = /dev/mapper/sgws-sn1-rangedb-4
BLOCK DEVICE RANGEDB 05 = /dev/mapper/sgws-sn1-rangedb-5
BLOCK DEVICE RANGEDB 06 = /dev/mapper/sgws-sn1-rangedb-6
BLOCK_DEVICE_RANGEDB_07 = /dev/mapper/sgws-sn1-rangedb-7
GRID NETWORK TARGET = bond0.1001
ADMIN NETWORK TARGET = bond0.1002
CLIENT NETWORK TARGET = bond0.1003
GRID NETWORK IP = 10.1.0.3
GRID NETWORK MASK = 255.255.255.0
GRID NETWORK GATEWAY = 10.1.0.1
```

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

sudo storagegrid node validate <node-name>

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

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

i

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

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

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

sudo storagegrid node start <node-name>

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

sudo storagegrid-status

The status is updated automatically.

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

Ctrl+C

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

sudo add_rangedbs.rb

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

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

The setup_rangedbs.sh script runs automatically.

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

- d. Select the Configuration tab and then the Main tab.
- e. If the Storage State Desired drop-down list is set to Read-only or Offline, select Online.
- f. Click Apply Changes.
- 15. To see the new object stores:
 - a. Select NODES > site > Storage Node > Storage.
 - b. View the details in the **Object Stores** table.

Result

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

Add grid nodes or site

Add grid nodes to existing site or add new site

Follow this procedure to add grid nodes to existing sites or to add a new site. You can only perform one type of expansion at a time.

Before you begin

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



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

Steps

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

Update subnets for Grid Network

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

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

Before you begin

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

• You have the network addresses, in CIDR notation, of the subnets you want to configure.

About this task

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

Steps

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

For example, enter 10.96.104.0/22.

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



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

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

Deploy new grid nodes

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

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

VMware: Deploy grid nodes

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

Steps

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

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

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

Linux: Deploy grid nodes

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

expansion nodes in the same way you deployed grid nodes during installation.

Before you begin

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



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

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

Steps

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

After you finish

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

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

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

Before you begin

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

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

• All Grid Network subnets listed on the IP Configuration page of the StorageGRID Appliance Installer have been defined in the Grid Network Subnet List on the primary Admin Node.

- The StorageGRID Appliance Installer firmware on the replacement appliance is compatible with the StorageGRID software version currently running on your grid. If the versions aren't compatible, you must upgrade the StorageGRID Appliance Installer firmware.
- You have a service laptop with a supported web browser.
- You know one of the IP addresses assigned to the appliance's compute controller. You can use the IP address for any attached StorageGRID network.

About this task

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

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

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



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

Steps

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

https://Controller_IP:8443

The StorageGRID Appliance Installer Home page appears.

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

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

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

Option	Description
Manual IP entry	a. Clear the Enable Admin Node discovery checkbox.
	b. Enter the IP address manually.
	c. Click Save.
	d. Wait for the connection state for the new IP address to become ready.

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

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

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

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

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

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

NetApp®	StorageGRID [®] Applian	ce Installer			
Home	Configure Networking -	Configure Hardware 🗸	Monitor Installation	Advanced 🗸	

Home

Interinstallation 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	172.16.4.210	
Connection state	Connection to 172.16.4.210 ready	
	Cancel Save	
Node name		
Node name	NetApp-SGA	
	Cancel Save	
Installation		
Current state	Ready to start installation of NetApp-SGA into grid wi 172.16.4.210.	th Admin Node
	Start Installation	

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

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



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

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

The Monitor Installation page shows the installation progress.

Monitor Installation

1. Configure storage			Running
Step	Progress	Status	
Connect to storage controller		Complete	
Clear existing configuration		Complete	
Configure volumes	and the second second	Creating volume StorageGRID-obj-00	
Configure host settings		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 aren't re-run. If you are rerunning an installation, any tasks that don't need to be re-run are shown with a green status bar and a status of "Skipped."

9. Review the progress of first two installation stages.

1. Configure appliance

During this stage, one of the following processes occurs:

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

2. Install OS

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

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



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

NetApp [®]	StorageGRID [®] Applian	ce Installer			Help 🔻
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 resou	irce 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		0 0 0
[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		

Perform expansion

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

Before you begin

- You are signed in to the Grid Manager using a supported web browser.
- You have the provisioning passphrase.
- You have deployed all of the grid nodes that are being added in this expansion.
- You have the Maintenance or Root access permission.

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

About this task

Performing the expansion includes these main user tasks:

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



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

Steps

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

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

+ Approve X Remove Grid Network MAC Address	11	Name rleo-010-096	L† 106-151	Type Storage Node	11	Platform VMware VM	11	Grid 10.96.	Search Network IPv4 Address 106.151/22	Q •
+ Approve X Remove Grid Network MAC Address	11	Name	11	Туре	11	Platform	ļ1	Grid	Search Network IPv4 Address	م •
+ Approve X Remove									Search	۹
irid nodes are listed as pending un	til they	are assigned t	to a site, co	onfigured, and	d approv	ed.				
ending Nodes										
Configure Expansion										
pprove and configure grid nodes,	so that	they are adde	d correctly	to your Stora	ageGRID	system.				
an units and an afferture wild up day	e e the et i	- have a set of a	بالاممينية ال	te con Char						
(*)										

2. Select Configure Expansion.

The Site Selection dialog box appears.

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

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

Pending Nodes				
Grid nodes are listed as	rleo-010-096 Storage Node	-106-151		
+ Approve 🗱 Remo	Network Grid Network	10 06 106 151/22	10.96.10/1	
Grid Network MA	Admin Network Client Network	Namé (1	Type	
00:50:56:a7:7a:c0	Hardware VMware VM			
00:50:56:a7:0f:2e	4 CPUs 8 GB RAM			
Approved Nodes	Disks 55 GB 55 GB 55 GB			



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

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

The grid node configuration form appears.

c. As required, modify the general settings:

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

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

- d. As required, modify the settings for the Grid Network, Admin Network, and Client Network.
 - **IPv4 Address (CIDR)**: The CIDR network address for the network interface. For example: 172.16.10.100/24



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

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

The approved grid node moves to the Approved Nodes list.

- To modify the properties of an approved grid node, select its radio button, and select Edit.
- To move an approved grid node back to the Pending Nodes list, select its radio button, and select **Reset**.
- To permanently remove an approved grid node, power the node off. Then, select its radio button, and select **Remove**.
- f. Repeat these steps for each pending grid node you want to approve.



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

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

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



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

s the status of gr	id configu	ration tasks r	equire	d to change the grid topology. The	ese grid cor	nfiguration tas	ks are ru	un automatica	ally by the StorageG	RID system.
Installing grid node	es									In Progre
rid Node Stat	us ion and co	nfiguration st	atus o	f each grid node included in the e	xpansion.					
									Search	a
Name	11	Site	ţţ	Grid Network IPv4 Address	~	Progress	11	Stage		41
leo-010-096-106-1	51	Data Center 1	e L	10.96.106.151/22		N. 19		Waiting for Dy	namic IP Service peers	
leo-010-096-106-1	56	Data Center 1	ų.	10.96.106.156/22				Waiting for NT	P to synchronize	
Initial configuratio	n									Pendi
Distributing the ne	w grid node	's certificates to	the Sto	rageGRID system.						Pendi
Assigning Storage	Nodes to sto	orage grade								Pendi
Starting services o	n the new gr	rid nodes								Pend

 $\overline{\mathbf{i}}$

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

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

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

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

- a. Select the download link.
- b. Enter the provisioning passphrase, and select Start Download.
- c. When the download completes, open the .zip file and confirm that you can access the contents, including the Passwords.txt file.
- d. Copy the downloaded Recovery Package file (.zip) to two safe, secure, and separate locations.



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

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



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

Adding Storage Nodes

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

, starting ser	vices on th	e new gri	d nodes					In Progres
irid Node	Status							
ists the ins	tallation	and con	figura	tion status of each grid node incl	uded in the exp	ansion.		
n Do not r	eboot any	Storage	Nodes	during Step 4. The "Starting Cassandr	a and streaming o	data" stage might t	take nours, especially if exi	isting Storage
Nodes con	tain a larg	e amount	t of obje	ect metadata.			• months	0
Nodes con	tain a larg	e amount	t of obje	ect metadata.			Search	۹
Nodes con	tain a larg	e amount Site	t of obje	ect metadata. Grid Network IPv4 Address 🗸	Progress ‡†	Stage	Search	D II
Nodes con Name rleo-010-096	tain a larg 11 -106-151	e amount Site Data Ce	t of obje 11 enter 1	Grid Network IPv4 Address v 10.96.106.151/22	Progress 1†	Stage Starting Cassan	Search dra and streaming data (20.	Q It 4% streamed)

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

Adding site

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

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

After you finish

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

Configure expanded system

Configuration steps after expansion

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

About this task

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

Steps

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



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

- 2. If you added one or more Storage Nodes to an existing site:
 - View storage pool details to confirm that each node you added is included in the expected storage pools and used in the expected ILM rules.
 - Confirm that the ILM policy meets the new requirements. If rule changes are required, create new rules and update the ILM policy. If the rules are already correct, activate a new policy with no rule changes to ensure StorageGRID uses the new nodes.
 - Verify that the Storage Node is active and able to ingest objects.
 - If you were unable to add the recommended number of Storage Nodes, rebalance erasure-coded data.
 See

Rebalance erasure-coded data after adding Storage Nodes.

- 3. If you added a Gateway Node:
 - If high availability (HA) groups are used for client connections, optionally add the Gateway Node to an HA group. Select CONFIGURATION > Network > High availability groups to review the list of existing HA groups and to add the new node. See Configure high availability groups.
- 4. If you added an Admin Node:
 - a. If single sign-on (SSO) is enabled for your StorageGRID system, create a relying party trust for the new Admin Node. You can't sign in to the node until you create this relying party trust. See Configure single sign-on.
 - b. If you plan to use the Load Balancer service on Admin Nodes, optionally add the new Admin Node to an HA group. Select CONFIGURATION > Network > High availability groups to review the list of existing HA groups and to add the new node. See Configure high availability groups.
 - c. Optionally, copy the Admin Node database from the primary Admin Node to the expansion Admin Node if you want to keep the attribute and audit information consistent on each Admin Node. See Copy the Admin Node database.
 - d. Optionally, copy the Prometheus database from the primary Admin Node to the expansion Admin Node if you want to keep the historical metrics consistent on each Admin Node. See Copy Prometheus metrics.
 - e. Optionally, copy the existing audit logs from the primary Admin Node to the expansion Admin Node if you want to keep the historical log information consistent on each Admin Node. See Copy audit logs.
- 5. To check if expansion nodes were added with an untrusted Client Network or to change whether a node's Client Network is untrusted or trusted, go to **CONFIGURATION** > **Security** > **Firewall control**.

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

6. Configure the DNS.

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

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

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

Verify that Storage Node is active

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

Steps

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

Copy Admin Node database

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

Before you begin

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

About this task

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

- Alert history
- · Alarm history

Historical attribute data, which is used in the charts and text reports available from the SUPPORT > Tools
 > Grid topology page

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



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

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

Steps

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

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

- h. When you no longer require passwordless access to other servers, remove the private key from the SSH agent. Enter:ssh-add -D
- 3. Restart the services on the primary Admin Node: service servermanager start

Copy Prometheus metrics

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

Before you begin

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

About this task

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



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

Steps

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

f. When prompted, press **Enter** to confirm that you want to destroy the new Prometheus database on the new Admin Node.

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

Database cloned, starting services

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

ssh-add -D

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

service prometheus start

Copy audit logs

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

Before you begin

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

About this task

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

By default, audit information is sent to the audit log on Admin Nodes. You can skip these steps if either of the following applies:



- You configured an external syslog server and audit logs are now being sent to the syslog server instead of to Admin Nodes.
- You explicitly specified that audit messages should be saved only on the local nodes that generated them.

See Configure audit messages and log destinations for details.

Steps

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

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

- 2. Stop the AMS service to prevent it from creating a new file: service ams stop
- 3. Navigate to the audit export directory:

cd /var/local/log

4. Rename the source audit.log file to ensure that it does not overwrite the file on the expansion Admin Node you are copying it to:

```
ls -l
mv audit.log _new_name_.txt
```

5. Copy all audit log files to the destination location on the expansion Admin Node:

```
scp -p * IP address:/var/local/log
```

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

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

8. Start the AMS service:

service ams start

9. Log out from the server:

exit

- 10. Log in to the expansion Admin Node:
 - a. Enter the following command: ssh admin@expansion Admin Node IP
 - b. Enter the password listed in the Passwords.txt file.
 - c. Enter the following command to switch to root: su -
 - d. Enter the password listed in the <code>Passwords.txt</code> file.

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

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

cd /var/local/log
chown ams-user:bycast *

12. Log out from the server:

exit

Rebalance erasure-coded data after adding Storage Nodes

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

Before you begin

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

What happens when this procedure runs

Before starting the procedure, note the following:

- The EC rebalance procedure will not start if one or more volumes are offline (unmounted) or if they are online (mounted) but in an error state.
- The EC rebalance procedure temporarily reserves a large amount of storage. Storage alerts might be triggered, but will resolve when the rebalance is complete. If there is not enough storage for the reservation, the EC rebalance procedure will fail. Storage reservations are released when the EC rebalance procedure s, whether the procedure failed or succeeded.
- If a volume goes offline while the EC rebalance procedure is in process, the rebalance procedure will terminate. Any data fragments that were already moved will remain in their new locations, and no data will be lost.

You can rerun the procedure after all volumes are back online.

• When the EC rebalance procedure is running, the performance of ILM operations and S3 and Swift client operations might be impacted.



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

- During this procedure all nodes have a storage capacity limit of 80%. Nodes that exceed this limit, but still store below the target data partition, are excluded from:
 - The site imbalance value
 - Any job completion conditions



The target data partition is calculated by dividing the total data for a site by the number of nodes.

- Job completion conditions. The EC rebalance procedure is considered complete when any of the following is true:
 - It can't move any more erasure-coded data.
 - $\circ\,$ The data in all nodes is within a 5% deviation of the target data partition.

• The procedure has been running for 30 days.

Steps

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

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

3. Start the procedure:

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

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

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

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

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

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

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

rebalance-data status



To get help on the rebalance-data command:

rebalance-data --help

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

Use the Site Imbalance value to assess how unbalanced erasure-code data usage is across the Storage Nodes at the site. This value can range from 1.0 to 0, with 0 indicating that erasure-coding data usage is completely balanced across all Storage Nodes at the site.

The EC rebalance job is considered complete and will stop when the data in all nodes is within a 5% deviation of the target data partition.

• If State is Success, optionally review object storage to see the updated details for the site.

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

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

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

- d. View the status of the new procedure. If State is still Failure, contact technical support.
- 7. If the EC rebalance procedure is generating too much load (for example, ingest operations are affected), pause the procedure.

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

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

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



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

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

Troubleshoot expansion

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

Before you contact technical support, collect the required log files to assist in troubleshooting.

Steps

- 1. Connect to the expansion node that has experienced failures:
 - a. Enter the following command:ssh -p 8022 admin@grid_node_IP



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

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

- c. Enter the following command to switch to root: su -
- d. Enter the password listed in the ${\tt Passwords.txt}$ file.

After you log in as root, the prompt changes from \$ to #.

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

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

Copyright information

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

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

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

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

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

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

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

Trademark information

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