StorageGRID® 11.3

Expansion Guide

November 2019 | 215-14211_2019-11_en-us
doccomments@netapp.com

NetApp®

Contents

Expansion overview 4
Planning and preparation 5
Reviewing the options and limitations of each type of expansion 5
Gathering required materials 7
Web browser requirements 8
Downloading and extracting the StorageGRID installation files 8
Verifying hardware and networking 11
Adding storage volumes to Storage Nodes 13
VMware: Adding storage volumes to a Storage Node 14
Linux: Adding direct-attached or SAN volumes to a Storage Node 16
Adding grid nodes to an existing site or adding a new site 20
Updating subnets for the Grid Network 20
Deploying new grid nodes 21
VMware: Deploying grid nodes 21
Linux: Deploying grid nodes 25
Appliances: Deploying Storage, Gateway, or non-primary Admin Nodes ... 26
Performing the expansion 30
Configuring your expanded StorageGRID system 36
Verifying that the Storage Node is active 40
Copying the Admin Node database 40
Copying Prometheus metrics 42
Copying audit logs 43
Contacting technical support 45
Copyright 46
Trademark information 47
How to send comments about documentation and receive update notifications 48

Expansion overview

You can expand a StorageGRID system by adding storage to Storage Nodes, adding new grid nodes to an existing site, or adding a new data center site. You can perform expansions without interrupting the operation of your current system.

As shown in the workflow diagram, the basic steps for performing each of the three types of expansions vary slightly, depending on whether your StorageGRID deployment uses hosts running VMware or Linux.

Note: NetApp-provided virtual machine disk files and scripts for new installations or expansions of StorageGRID on OpenStack are no longer supported. To expand an existing deployment on OpenStack, refer to the steps for your Linux distribution.

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

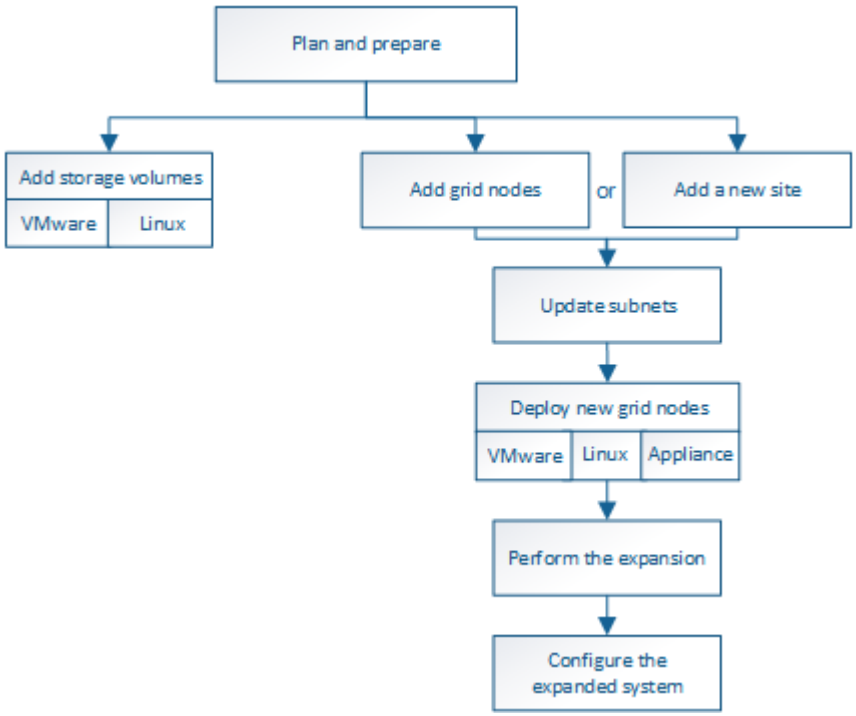

Related tasks

Planning and preparation on page 5
Adding storage volumes to Storage Nodes on page 13
Adding grid nodes to an existing site or adding a new site on page 20

Related information

NetApp Interoperability Matrix Tool

Planning and preparation

You must plan and prepare for the StorageGRID expansion, which includes reviewing the limitations of each type of expansion, obtaining required materials, and installing and configuring any new hardware and networks.

Reviewing the options and limitations of each type of expansion

Before performing an expansion, you must understand the options and limitations of each type of expansion.

Adding storage volumes to Storage Nodes

You can increase the overall storage capacity of your grid by adding additional storage volumes (LUNs) to existing non-appliance Storage Nodes. You cannot expand the storage capacity of an appliance Storage Node. The maximum storage available on each StorageGRID appliance Storage Node is fixed, based on the appliance model and original configuration.

When preparing to add storage volumes, be aware of the following limitations:

- Each node can support a maximum of 16 storage volumes. If you need to add capacity beyond that, you must add new Storage Nodes.
- You cannot add capacity to an existing storage volume.
- You cannot add storage volumes to a Storage Node at the same time you are performing a system upgrade, recovery operation, or another expansion.
- Before adding storage volumes, you must review the active information lifecycle management (ILM) policy and rules. In most cases, you must add storage volumes to more than one Storage Node to satisfy ILM requirements for replicated or erasure-coded copies.

See “Adding storage volumes to Storage Nodes” for instructions.

Adding grid nodes to existing sites

You can expand storage capacity and add redundancy to a StorageGRID system by adding new grid nodes to existing data center sites. In a single expansion, you can add one or more of the following types of nodes to one or more existing sites:

- Non-primary Admin Nodes
- Storage Nodes
- Archive 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 cannot add a primary Admin Node during an expansion.
- Before adding Storage Nodes, you must confirm that all data-repair operations performed as part of a recovery are complete. See the steps for checking data repair jobs in the recovery and maintenance instructions.

- Before adding Storage Nodes, you must review the active ILM policy and rules and the StorageGRID system's current topology.
  - If an ILM rule creates replicated copies, you must consider the number of copies created and how they are distributed to storage pools. For example, if an ILM rule creates one replicated copy at data center site 1 (DC1) and one copy at data center site 2 (DC2), you must add an equivalent number of Storage Nodes at each site to increase the overall capacity of the StorageGRID system.
  - If you are using or want to use erasure coding, you must understand the Storage Node and site requirements for the erasure coding scheme you want to use. If you are adding Storage Nodes because the system's existing Storage Nodes are nearing capacity, you must add enough Storage Nodes to support the erasure coding scheme in use. See the ILM information to learn the requirements for each supported erasure coding scheme.
- If you are adding Archive Nodes, note that each Archive Node supports only one interface, either tape through Tivoli Storage Manager (TSM) middleware or the cloud through the S3 API. To support both tape and S3, you must deploy two Archive Nodes.
- If the **New Node Client Network Default** option is set to to **Untrusted** on the Untrusted Client Networks page (at **Configuration > Untrusted Client Network**), client applications that connect to expansion nodes using the Client Network must connect using a load balancer endpoint port. See the instructions for administering StorageGRID to change the setting for the new node and to configure load balancer endpoints.

Adding a data center site

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

When preparing to add a data center site, be aware of the following limitations:

- You can only add one data center site per expansion operation.
- You cannot add grid nodes to an existing data center site as part of the same expansion.
- The new data center site must include at least three Storage Nodes.
- 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 are not stored to the new data center site until the expansion procedure is complete. See the information about ILM rules.
- When you add a new site, a copy of the Cassandra database (used for object metadata) is saved to the Storage Nodes at the new site. For this reason, adding a new site does not increase the number of objects you can store in the grid. To increase the maximum number of objects you can store in your StorageGRID system, you must increase the number of Storage Nodes per site, which increases the amount of space available for object metadata.

Related tasks

Adding storage volumes to Storage Nodes on page 13
Adding grid nodes to an existing site or adding a new site on page 20

Related information

Recovery and maintenance

Gathering required materials

Before you perform an expansion operation, you must gather the materials listed in the following table.

Item	Notes
StorageGRID installation archive	If you are adding new grid nodes or a new data center 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 do not need to download files if you are adding new storage volumes to existing Storage Nodes or installing new StorageGRID storage appliances or services appliances.
Service laptop	The service laptop must meet the following requirements: • Network port • SSH client (for example, PuTTY) • Supported browser
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	• *Administering StorageGRID* • *StorageGRID Release Notes* • Installation instructions for your platform
Current documentation for your platform	For supported versions, see the Interoperability Matrix.

Related tasks

Downloading and extracting the StorageGRID installation files on page 8

Related references

Web browser requirements on page 8

Related information

Administering StorageGRID
StorageGRID release notes
VMware installation
Red Hat Enterprise Linux or CentOS installation
Ubuntu or Debian installation
NetApp Interoperability Matrix Tool

Web browser requirements

You must use a supported web browser.

Web browser	Minimum supported version
Google Chrome	74
Microsoft Internet Explorer	11 (Native Mode)
Mozilla Firefox	67

You should set the browser window to a recommended width.

Browser width	Pixels
Minimum	1024
Optimum	1280

Downloading and extracting the StorageGRID installation files

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

About this task

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

Steps

1. Go to the NetApp Downloads page for StorageGRID.

   NetApp Downloads: StorageGRID

2. Select the version of StorageGRID that is currently running on the grid.
3. Sign in using the username and password for your NetApp account.
4. Read and accept the End User License Agreement.

   The downloads page for the version you selected appears. The page contains columns for new installation files, upgrade files, and NAS Bridge.

5. In the **New install files** column, select the `.tgz` or `.zip` file for your platform.

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

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

Platform	Installation archive
VMware	• `StorageGRID-Webscale-`*`version`*`-VMware-`*`uniqueID`*`.zip` • `StorageGRID-Webscale-`*`version`*`-VMware-`*`uniqueID`*`.tgz`

Platform	Installation archive
Red Hat Enterprise Linux or CentOS	• `StorageGRID-Webscale-`*`version`*`-RPM-`*`uniqueID`*`.zip` • `StorageGRID-Webscale-`*`version`*`-RPM-`*`uniqueID`*`.tgz`
Ubuntu or Debian Appliance	• `StorageGRID-Webscale-`*`version`*`-DEB-`*`uniqueID`*`.zip` • `StorageGRID-Webscale-`*`version`*`-DEB-`*`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. 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 tables are relative to the top-level directory installed by the archive file.

Table 1: VMware files

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

Filename	Description
`/vsphere/configure-sga.py`	A Python script used to automate the configuration of StorageGRID appliances.
`/vsphere/configure-storagegrid.py`	A Python script used to automate the configuration of a StorageGRID system.
`/vsphere/configure-storagegrid.sample.json`	A sample configuration file for use with the `configure-storagegrid.py` script.
`/vsphere/configure-storagegrid.blank.json`	A blank configuration file for use with the `configure-storagegrid.py` script.

Table 2: Red Hat Enterprise Linux or CentOS files

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

Table 3: Ubuntu or Debian files

Path and file name	Description
`/debs/README`	A text file that describes all of the files contained in the StorageGRID download file.
`/debs/NLF000000.txt`	A free license that does not provide any support entitlement for the product.
`/debs/storagegrid-webscale-images-`*`version-SHA`*`.deb`	DEB package for installing the StorageGRID node images on your hosts.

Path and file name	Description
`/debs/storagegrid-webscale-service-`*`version-SHA`*`.deb`	DEB package for installing the StorageGRID host service on your hosts.
Deployment scripting tools	
`/debs/configure-storagegrid.py`	A Python script used to automate the configuration of a StorageGRID system.
`/debs/configure-storagegrid.sample.json`	A sample configuration file for use with the `configure-storagegrid.py` script.
`/debs/configure-storagegrid.blank.json`	A blank configuration file for use with the `configure-storagegrid.py` script.
`/debs/configure-sga.py`	A Python script used to automate the configuration of StorageGRID appliances.
`/debs/extras/ansible`	An Ansible role and example playbook for configuring the hosts for StorageGRID container deployment. You can customize the playbook or role as necessary.

Table 4: Appliance files

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

Note: For appliance installation, these files are usually not required. The appliance can download the required files from the Primary Admin Node.

Verifying hardware and networking

Before you begin the expansion of your StorageGRID system, you must ensure that you have installed and configured the necessary hardware to support the new grid nodes or new site.

For information about supported versions, see the Interoperability Matrix.

You must also verify network connectivity between servers at the data center site, and confirm that the primary Admin Node can communicate with all expansion servers that are intended to host the StorageGRID system.

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

You can use network address translation (NAT) between external clients and grid nodes, such as to provide a public IP address for a Gateway Node. However, you must not use 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 a public network segment is supported only when you employ a tunneling application that is transparent to all nodes in the grid, meaning the grid nodes require no knowledge of public IP addresses.

Related tasks

Updating subnets for the Grid Network on page 20

Related information

NetApp Interoperability Matrix Tool

Adding storage volumes to Storage Nodes

You can expand the storage capacity of non-appliance Storage Nodes that have fewer than 16 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.

About this task

The underlying storage of a Storage Node is divided into a fixed number of 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.

Note: Object metadata is always stored to object store 0.

Each object store corresponds to a mount point on the same Storage Node. That is, 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.

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

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

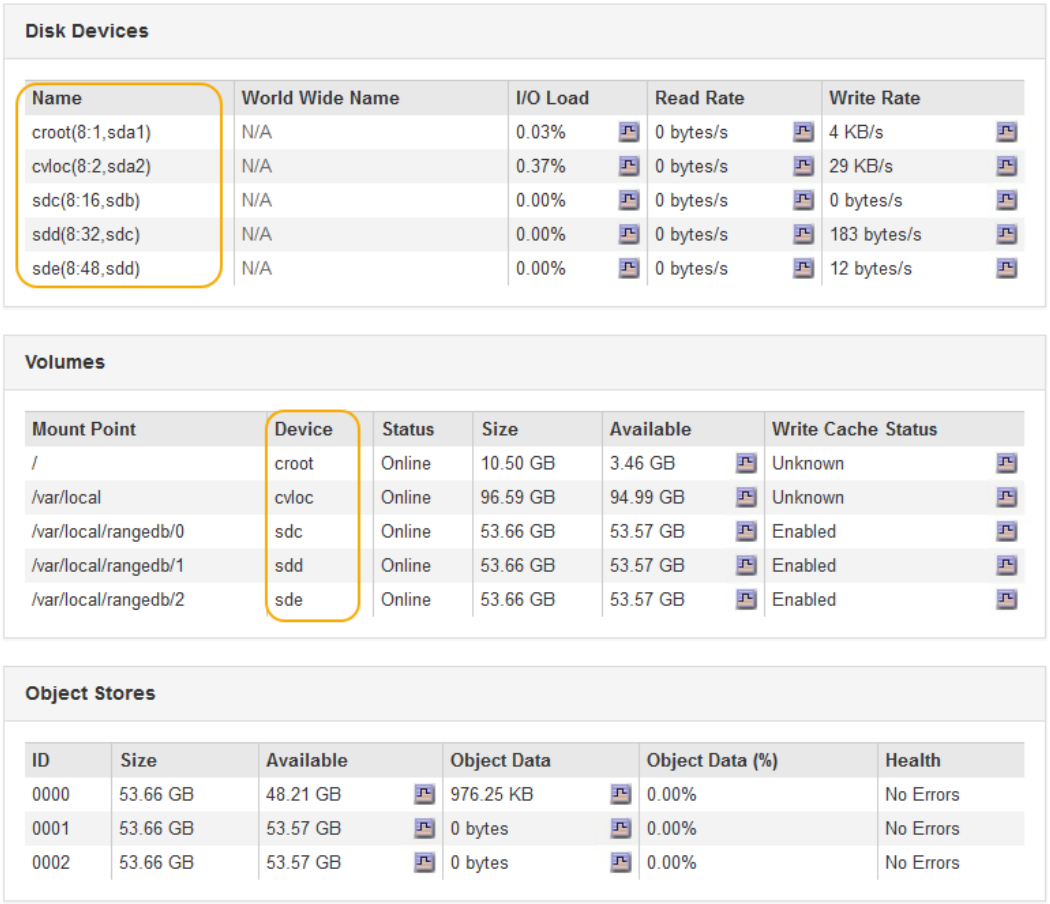

Disk Devices

Name	World Wide Name	I/O Load	Read Rate	Write Rate
croot(8:1,sda1)	N/A	0.03%	0 bytes/s	4 KB/s
cvloc(8:2,sda2)	N/A	0.37%	0 bytes/s	29 KB/s
sdc(8:16,sdb)	N/A	0.00%	0 bytes/s	0 bytes/s
sdd(8:32,sdc)	N/A	0.00%	0 bytes/s	183 bytes/s
sde(8:48,sdd)	N/A	0.00%	0 bytes/s	12 bytes/s

Volumes

Mount Point	Device	Status	Size	Available	Write Cache Status
/	croot	Online	10.50 GB	3.46 GB	Unknown
/var/local	cvloc	Online	96.59 GB	94.99 GB	Unknown
/var/local/rangedb/0	sdc	Online	53.66 GB	53.57 GB	Enabled
/var/local/rangedb/1	sdd	Online	53.66 GB	53.57 GB	Enabled
/var/local/rangedb/2	sde	Online	53.66 GB	53.57 GB	Enabled

Object Stores

ID	Size	Available	Object Data	Object Data (%)	Health
0000	53.66 GB	48.21 GB	976.25 KB	0.00%	No Errors
0001	53.66 GB	53.57 GB	0 bytes	0.00%	No Errors
0002	53.66 GB	53.57 GB	0 bytes	0.00%	No Errors

3. Follow the instructions for your platform to add new storage volumes to the Storage Node.

Choices

- VMware: Adding storage volumes to a Storage Node on page 14
- Linux: Adding direct-attached or SAN volumes to a Storage Node on page 16

VMware: Adding storage volumes to a 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 must have access to the instructions for installing StorageGRID for VMware deployments.
- You must have the `Passwords.txt` file.
- You must have specific access permissions. For details, see information about controlling system access with administration user accounts and groups.

Attention: Do not 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 may 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** option in the VMware vSphere Client, or by entering the following command in an ssh session to the virtual machine:

   ```
   sudo reboot
   ```

   Note: Do not 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 storage volumes are formatted.
   a. When asked, enter **y** to stop storage services.

      The storage services are stopped, and the `setup_rangedb.sh` script is run. After the volumes are ready for use as rangedbs, the services start again.
5. Check that the services start correctly:
   a. View a listing 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 browser.
   b. Select **Support** > **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**.
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 now use the expanded capacity of the Storage Nodes to save object data.

Related information

VMware installation

Linux: Adding direct-attached or SAN volumes to a 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 must have access to the instructions for installing StorageGRID for Red Hat Enterprise Linux or CentOS deployments.
- You must have the `Passwords.txt` file.
- You must have specific access permissions. For details, see information about controlling system access with administration user accounts and groups.

Attention: Do not 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 disk 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 Red Hat Enterprise Linux or CentOS deployments.

3. Log into 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, this section is in **bold**.

   Example

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

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

   Make sure to start at the next `BLOCK_DEVICE_RANGEDB_NN`. Do not leave a gap.

   - For the example above, start at `BLOCK_DEVICE_RANGEDB_04`.
   - In the example below, four new block storage volumes have been added to the node. The new mappings are in **bold**.

   Example

   ```
   NODE_TYPE = VM_Storage_Node
   ADMIN_IP = 10.1.0.2
   BLOCK_DEVICE_VAR_LOCAL = /dev/mapper/sgws-sn1-var-local
   ```

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

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

   `sudo storagegrid node validate <node-name>`

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

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

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

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

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

    `sudo storagegrid node start <node-name>`

11. Log in to the Storage Node as admin using the password listed in the `Passwords.txt` file.

12. Check that the services start correctly:

    a. View a listing 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 storage volumes are formatted.

d. When asked, enter **y** to stop storage services.

The storage services are stopped, and the `setup_rangedbs.sh` script is run. After the volumes are ready for use as rangedbs, the services start again.

14. Check that the services start correctly:

a. View a listing 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
```

15. Verify that the Storage Node is online:

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

You must have Maintenance or Root Access permissions.

b. Select **Support** > **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**.

16. 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.

Related information

Red Hat Enterprise Linux or CentOS installation
Ubuntu or Debian installation

Adding grid nodes to an existing site or adding a new site

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

Before you begin

- You must have root or maintenance permissions. For details, see information about controlling system access with administration user accounts and groups.
- All existing nodes in the grid must be up and running across all sites.
- Any previous expansion, upgrade, decommissioning, or recovery procedures must be complete.

  Note: You are prevented from starting an expansion while another expansion, upgrade, recovery, or active decommission procedure is in progress.

Steps

1. Updating subnets for the Grid Network on page 20
2. Deploying new grid nodes on page 21
3. Performing the expansion on page 30
4. Configuring your expanded StorageGRID system on page 36

Related information

Administering StorageGRID

Updating subnets for the Grid Network

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

Before you begin

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

About this task

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

Steps

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

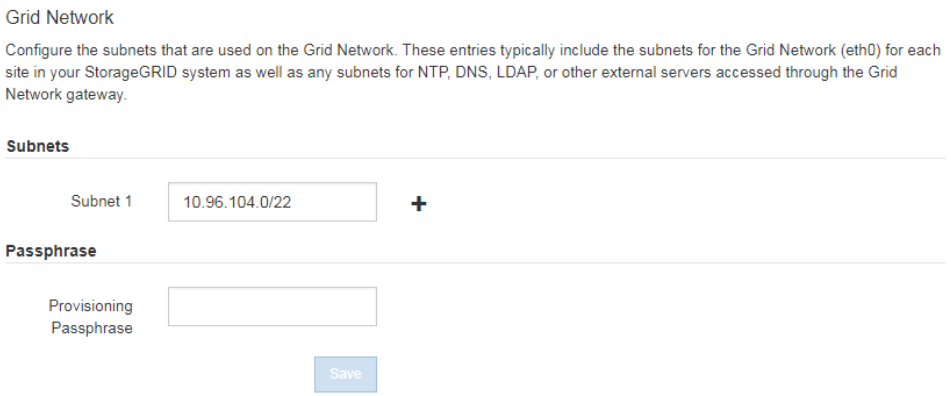

2. In the Subnets list, click the plus sign to add a new subnet in CIDR notation.

   For example, enter 10.96.104.0/22.

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

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

Deploying new grid nodes

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

When you expand the grid, the nodes you add do not have to match the existing node types. You can add VMware nodes, Linux containerized nodes, or appliance nodes.

Choices

- VMware: Deploying grid nodes on page 21
- Linux: Deploying grid nodes on page 25
- Appliances: Deploying Storage, Gateway, or non-primary Admin Nodes on page 26

VMware: Deploying grid nodes

You must deploy a virtual machine in VMware vSphere for each grid node.

Before you begin

- You have the instructions for installing StorageGRID for VMware, and you have reviewed the hardware, software, virtual machine, and storage and performance requirements.
- You have a plan to minimize failure domains. For example, you should not deploy all Gateway Nodes on a single virtual machine server.

  Attention: In a production deployment, do not run more than one Storage Node on a single virtual machine server. Using a dedicated virtual machine host for each Storage Node provides an isolated failure domain.

- If a StorageGRID node is deployed in a virtual machine with storage assigned from a NetApp AFF system, confirm that the volume does not have a FabricPool tiering policy enabled.

Disabling FabricPool tiering for volumes used with StorageGRID nodes simplifies troubleshooting and storage operations.

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

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

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

- You have placed all of these files in the same directory.

Steps

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

   The best practice is to use the same name for the virtual machine as you used for the grid node.
5. In the **Network Mapping** page, select the networks to use by associating a network port to each network. The Grid Network is required. The Admin and Client Networks are optional. Select the Grid Network to use, and then choose the following as applicable:
   - If you are planning to use the Admin Network, assign the Admin Network adapter to a network in the vSphere environment.
   - If you are planning to use the Client Network, assign the Client Network adapter to a network in the vSphere environment.
   - If you do not plan to use an Admin Network or Client Network, assign their network adapters to the same network as the Grid Network.
6. Provide the required StorageGRID information in the **Properties** page, and click **Finish**.
   a. Enter the **Node Name**.

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

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

      If you omit the primary Admin Node IP address, the IP address will be automatically discovered if the primary Admin Node, or at least one other grid node with ADMIN_IP

configured, is present on the same subnet. However, it is recommended to set the primary Admin Node IP address here.

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

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

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

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

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

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

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

8. Assign storage to the virtual machine.

It is recommended that you use 3 or more storage volumes for each Storage Node, with each storage volume being 4 TB or larger. You must assign at least 4 TB to volume 0. See “Storage and performance requirements” for more information.

Attention: The Storage Node OVF provided defines several VMDKs for storage. Unless these VMDKs meet your storage requirements, you should remove them and assign appropriate VMDKs or RDMs for storage before powering up the node. VMDKs are more commonly used in VMware environments and are easier to manage, while RDMs may provide better performance for workloads that use larger object sizes (for example, greater than 100 MB).

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

Attention: If you remap any ports, you cannot use the same ports to configure load balancer endpoints. If you want to configure load balancer endpoints and have already remapped ports, follow the steps in the recovery and maintenance instructions for removing port remaps.

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

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

c. Select **vApp Options**.

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

You might need to remap a port if your enterprise networking policies restrict access to one or more ports that are used by StorageGRID. See the information about internal grid node communications or external communications for the list of ports used by StorageGRID.

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

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

Example

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

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

Click **OK**.

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

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

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

Example

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

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

Click **OK**.

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

10. Power on the virtual machine.

Related tasks

Performing the expansion on page 30

Related information

VMware installation

Linux: Deploying 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.
- 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.

  Attention: In a production deployment, do not 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 host for the StorageGRID node will use storage assigned from a NetApp AFF system, confirm that the volume does not have a FabricPool tiering policy enabled. Disabling FabricPool tiering for volumes used with StorageGRID nodes simplifies troubleshooting and storage operations.

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

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.

Related tasks

Performing the expansion on page 30

Related information

Red Hat Enterprise Linux or CentOS installation

Ubuntu or Debian installation

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.
- Network links, IP addresses, and port remapping (if necessary) have been configured for the appliance using the StorageGRID Appliance Installer.

  Attention: If you have remapped any ports, you cannot use the same ports to configure load balancer endpoints. You can create endpoints using remapped ports, but those endpoints will be remapped to the original CLB ports and service, not the Load Balancer service. Follow the steps in the recovery and maintenance instructions for removing port remaps.

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

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

- 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.

  Note: 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. This script applies only to Storage Nodes.

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.** Unselect the **Enable Admin Node discovery** check box. **b.** Enter the IP address manually. **c.** Click **Save**. **d.** Wait for the connection state for the new IP address to become ready.
Automatic discovery of all connected primary Admin Nodes	**a.** Select the **Enable Admin Node discovery** check box. **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 click **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 installation and maintenance instructions for your appliance.

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

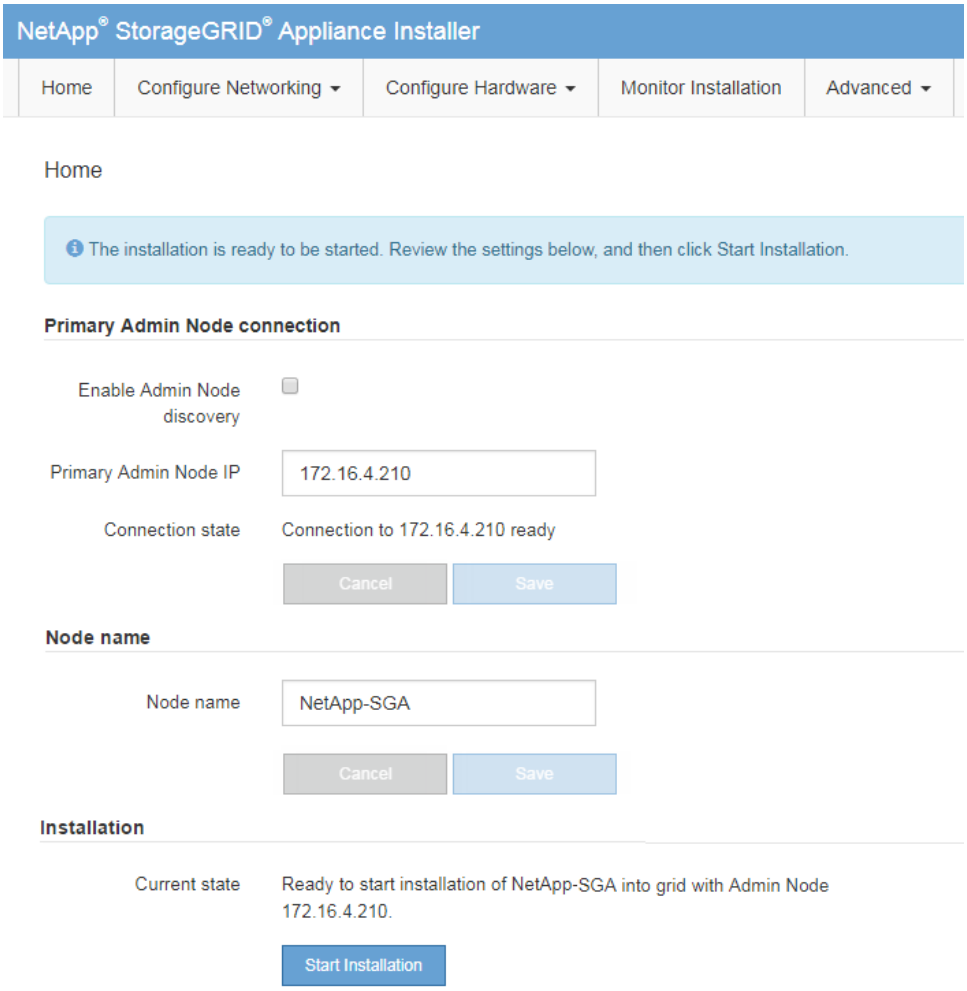

The Current state changes to “Installation is in progress,” and the Monitor Installation page is displayed.

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

   Note: 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. This script applies only to Storage Nodes.

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

   The Monitor Installation page shows the installation progress.

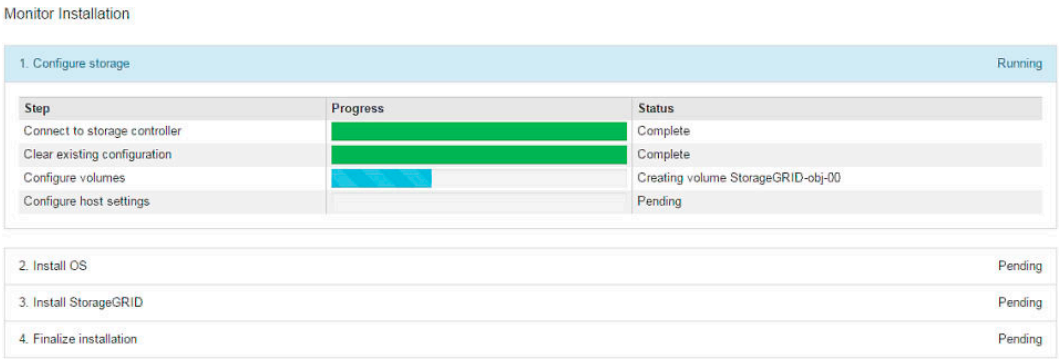

The blue status bar indicates which task is currently in progress. Green status bars indicate tasks that have completed successfully.

Note: The installer ensures that tasks completed in a previous install are not re-run. If you are re-running an installation, any tasks that do not need to be re-run are shown with a green status bar and a status of "Skipped."

9. Review the progress of first two installation stages.

1. Configure appliance

During this stage, one of the following processes occurs, depending on the appliance type:

- For an appliance Storage Node, the installer connects to the storage controller, clears any existing configuration, communicates with SANtricity software to configure volumes, and configures host settings.
- For an appliance Gateway Node or Admin Node, 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 you are prompted to use the Grid Manager to approve the node.

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

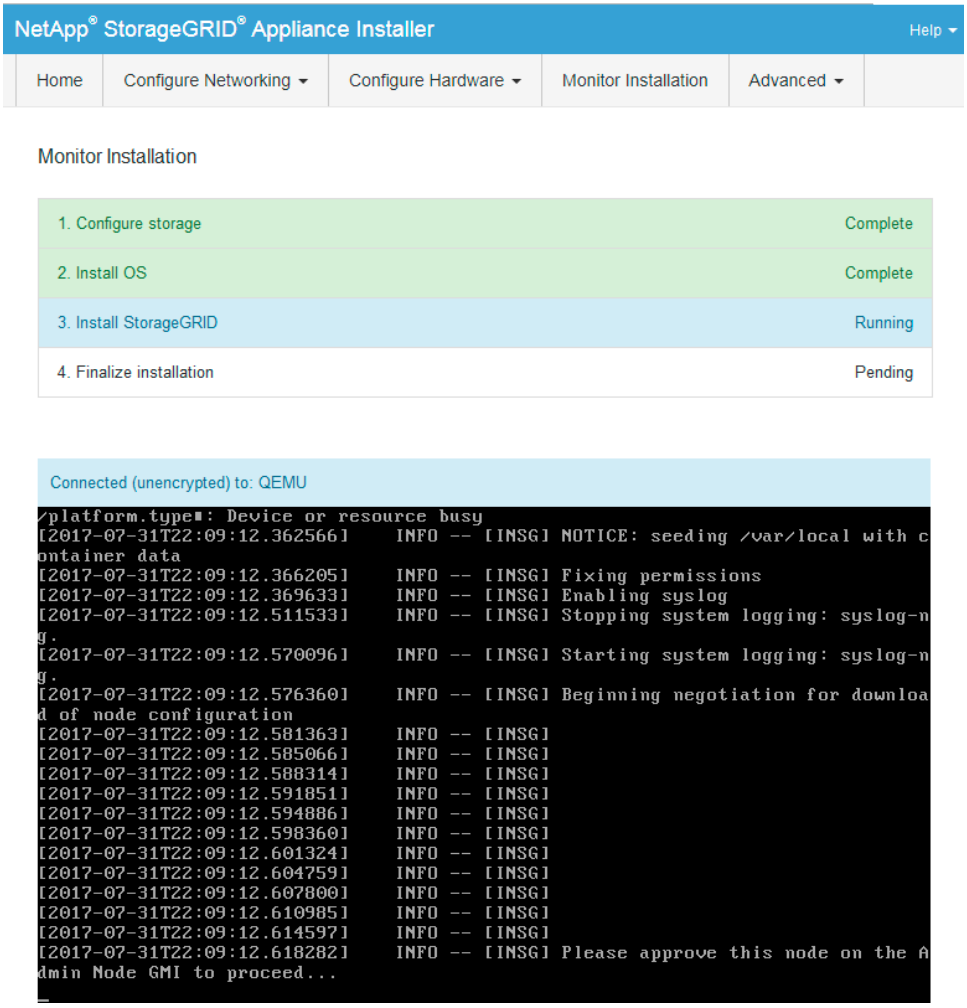

Related information

SG5700 appliance installation and maintenance
SG5600 appliance installation and maintenance
SG6000 appliance installation and maintenance
SG1000 appliance installation and maintenance

Performing the expansion

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

Before you begin

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

- If you are adding Storage Nodes, you must have confirmed that all data-repair operations performed as part of a recovery are complete. See the steps for checking data repair jobs in the recovery and maintenance instructions.

About this task

Performing the expansion includes these phases:

1. You configure the expansion by specifying whether you are adding new grid nodes or a new site and approving the grid nodes you want to add.
2. You start the expansion.
3. While the expansion process is running, you download a new Recovery Package file.
4. You monitor the status of the grid configuration tasks, which run automatically. The set of tasks depends on what types of grid nodes are being added and on whether a new site is being added.

   Attention: Some tasks 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 relatively empty. However, if the Cassandra database includes a large amount of object metadata, this stage might take several hours or longer. You can look at the "streamed" percentage shown during the "Starting Cassandra and streaming data" stage to determine how complete the Cassandra streaming operation is.

Steps

1. Select **Maintenance > Expansion**.

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

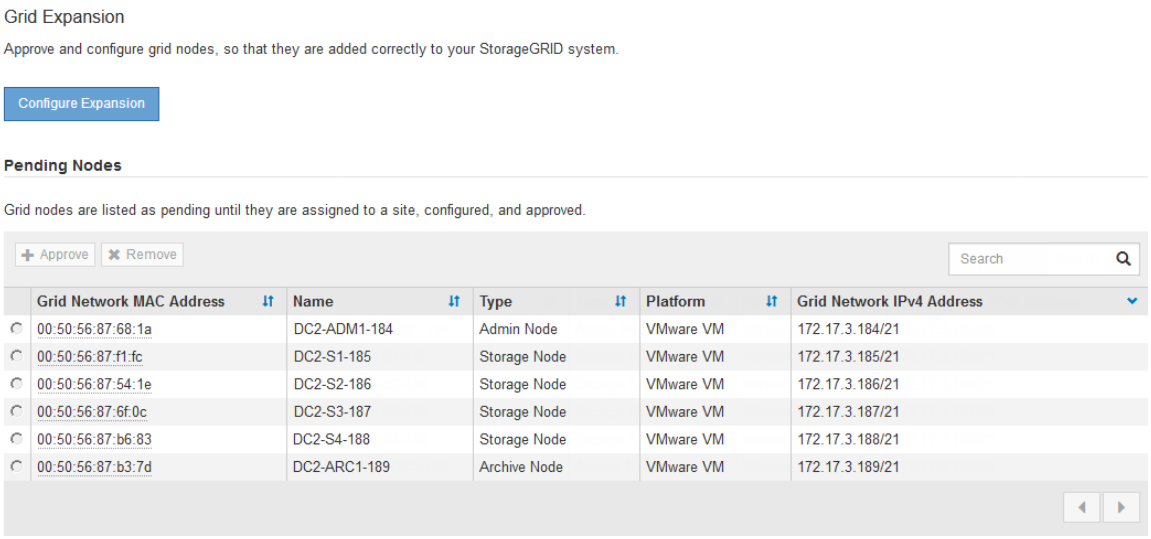

2. Click **Configure Expansion**.

   The Site Selection dialog box appears.

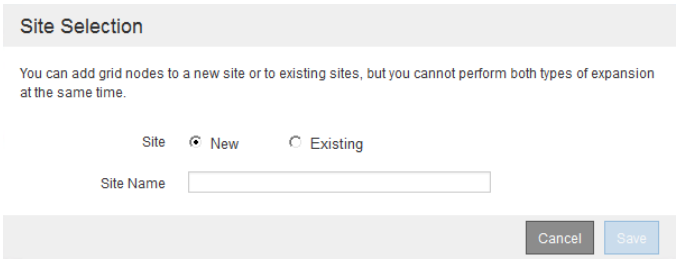

3. Select the type of expansion you are starting:
   - If you are adding a new site, select **New**, and enter the name of the new site.
   - If you are adding grid nodes to an existing site, select **Existing**.
4. Click **Save**.
5. Review the **Pending Nodes** list, and confirm that it shows all of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

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

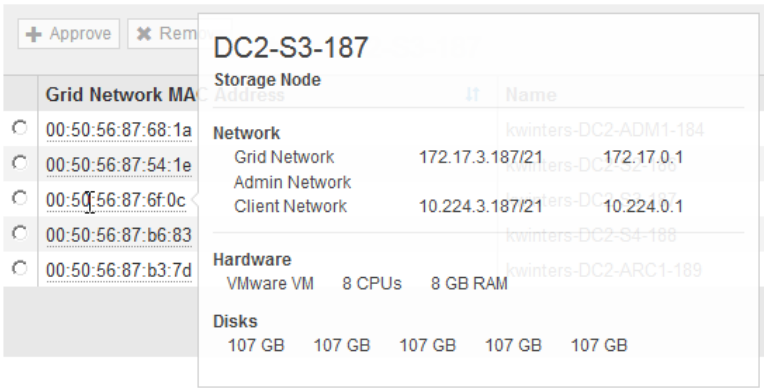

Note: If a grid node is missing, confirm that it was deployed successfully.

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

      The grid node configuration form appears.

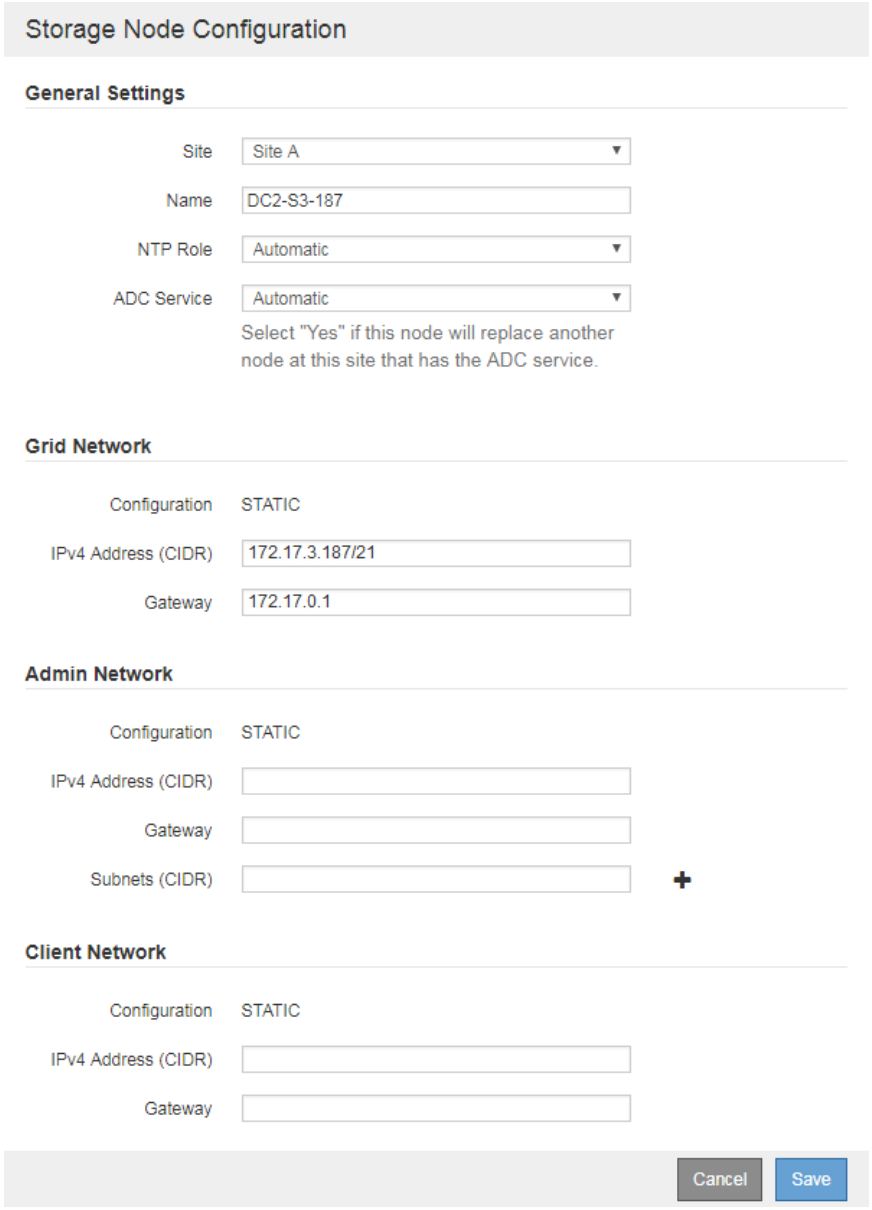

c. As required, modify the general settings:

- **Site**: The name of the site the grid node will be associated with. If you are adding a new site, all nodes are added to the new site.
- **Name**: The host name that will be assigned to the node, and the name that will be displayed in the Grid Manager.
- **NTP Role**: The Network Time Protocol (NTP) role of the grid node. The options are **Automatic**, **Primary**, and **Client**. Selecting **Automatic** assigns the Primary role to Admin Nodes, Storage Nodes with ADC services, Gateway Nodes, and any grid nodes that have non-static IP addresses. All other grid nodes are assigned the Client role.

  Note: Assign the Primary NTP role to at least two nodes at each site. This provides redundant system access to external timing sources.

- **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 cannot add the ADC service to a node after it is deployed.

- If you are adding this node to replace a Storage Node, select **Yes** if the node you are replacing includes the ADC service. Because you cannot 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.
- Otherwise, select **Automatic** to let the system determine whether this node requires the ADC service.

Learn about the ADC quorum in the recovery and maintenance instructions.

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
- **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. Click **Save**.

The approved grid node moves to the Approved Nodes list.

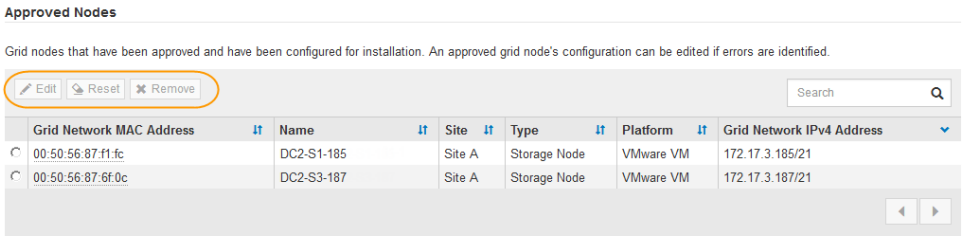

- To modify the properties of an approved grid node, select its radio button, and click **Edit**.
- To move an approved grid node back to the Pending Nodes list, select its radio button, and click **Reset**.
- To permanently remove an approved grid node, power the node off. Then, select its radio button,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

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

Note: 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 click **Expand**.

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

Note: During this process, for appliances the StorageGRID Appliance Installer shows installation moving from Stage 3 to Stage 4, Finalize Installation. When Stage 4 completes, the controller is rebooted.

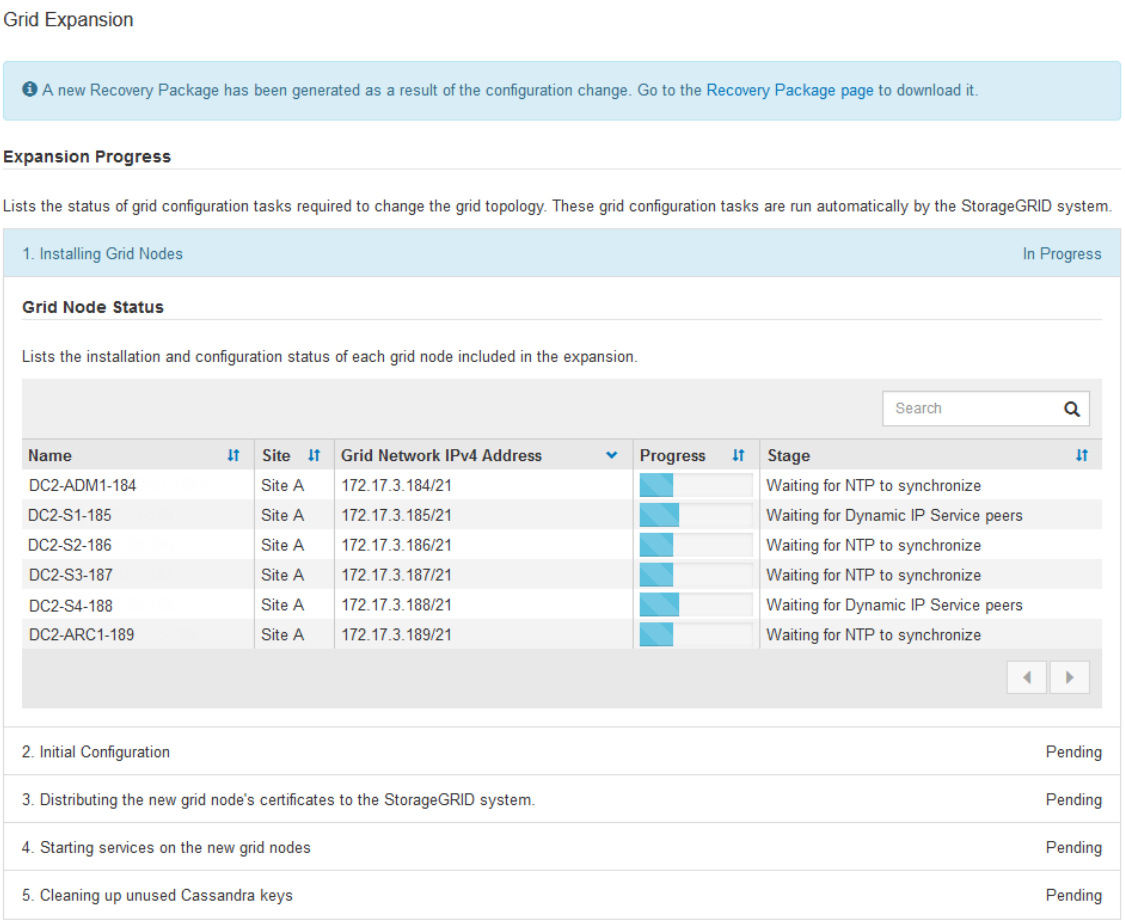

Note: A site expansion includes an additional task to configure Cassandra for the new site.

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. Click the download link.

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

   c. When the download completes, open the `.zip` file and confirm it includes a `gpt-backup` directory and a `*_SAID.zip` file. Then, extract the `*_SAID.zip` file, go to the `/GID*_REV*` directory, and confirm you can open the `passwords.txt` file.

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

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

9. If you are adding one or more Storage Nodes, monitor the progress of the “Starting Cassandra and streaming data” stage by reviewing the percentage shown in the status message.

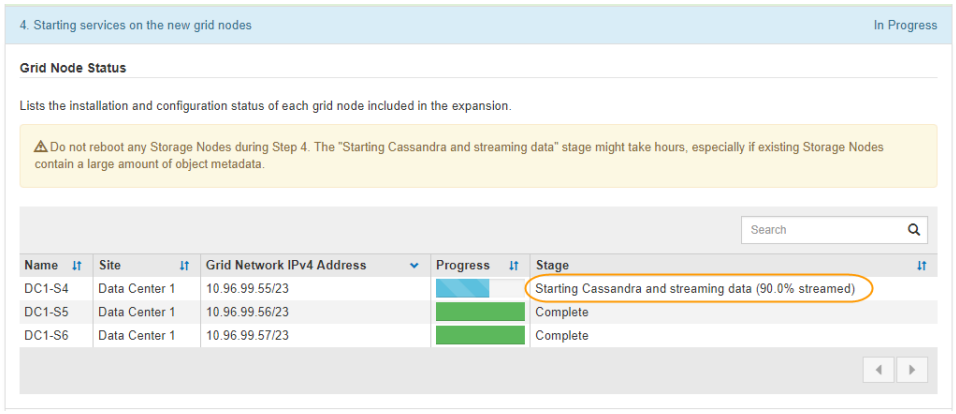

This percentage estimates how complete the Cassandra streaming operation is, based on the total amount of Cassandra data available and the amount that has already been written to the new node.

Attention: Do not reboot any Storage Nodes during Step 4 (Starting services on the new grid nodes). The “Starting Cassandra and streaming data” stage might take hours to complete for each new Storage Node, especially if existing Storage Nodes contain a large amount of object metadata.

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, you must perform additional integration and configuration steps.

Related tasks

Configuring your expanded StorageGRID system on page 36

Related information

Recovery and maintenance

Configuring your expanded StorageGRID system

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 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 grid nodes added during the expansion.

Steps

1. Complete the following configuration tasks for each type of grid node you added during the expansion:

Type of grid node	Configuration tasks
Storage Nodes	**a.** Verify the information lifecycle management (ILM) storage pool configuration. You must verify that the expansion Storage Nodes are included in a storage pool used by a rule in the active ILM policy. Otherwise, the new storage will not be used by the StorageGRID system. See the instructions for administering StorageGRID. **b.** Verify that the Storage Node is ingesting objects. See the instructions for verifying that the Storage Node is active.
Gateway Nodes	**a.** If High Availability Groups are used for client connections, add the Gateway Nodes to an HA group. Go to **Configuration > High Availability Groups** to review the list of existing HA groups and to add the new nodes. For more information, see the instructions for administering StorageGRID.

Type of grid node	Configuration tasks
Admin Nodes	**a.** If single sign-on is enabled for your StorageGRID system, you must create a relying party trust in Active Directory Federation Services (AD FS) for the new Admin Node. You cannot sign in to the node until you create this relying party trust. To learn how to create relying party trusts for Admin Nodes, see "Configuring single sign-on" in the instructions for administering StorageGRID. **b.** If you plan to use the Load Balancer service on Admin Nodes, you might need to add the Admin Nodes to High Availability Groups. Go to **Configuration > High Availability Groups** to review the list of existing HA groups and to add the new nodes. For more information, see the instructions for administering StorageGRID. **c.** Copy the Admin Node database. 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. For more information, see "Copying the Admin Node database." **d.** Copy the Prometheus metrics. 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. For more information, see "Copying Prometheus metrics." **e.** Copy the audit logs. 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 "Copying the audit logs." **f.** Configure access to audit shares. Optionally, you can configure access to the system for auditing purposes through an NFS or a CIFS file share. See the instructions for administering StorageGRID. **Note:** Audit export through CIFS/Samba has been deprecated and will be removed in a future StorageGRID release. **g.** Change the preferred sender for email notifications. Optionally, you can update your configuration to make the expansion Admin Node the preferred sender. Otherwise, an existing Admin Node configured as the preferred sender continues to send notifications and AutoSupport messages. See the instructions for administering StorageGRID.

Type of grid node	Configuration tasks
Archive Nodes	**a.** Configure the Archive Node's connection to the targeted external archival storage system. When you complete the expansion, Archive Nodes are in an alarm state until you configure connection information through the **ARC** > **Target** component. **b.** Update the ILM policy. You must update your ILM policy in order to archive object data through the new Archive Node. **c.** Configure custom alarms. You should establish custom alarms for the attributes that are used to monitor the speed and efficiency of object data retrieval from Archive Nodes. For more information, see the instructions for administering StorageGRID.

2. 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 > Untrusted Client Network**.

   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 the instructions for administering StorageGRID for more information.

3. Configure the Domain Name System (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 information about modifying the DNS configuration for a single grid node in the recovery and maintenance instructions.

   The best practice is for the grid-wide DNS server list to contain some DNS servers that are accessible locally from each site. If you just added a new site, add new DNS servers for the site to the grid-wide DNS configuration.

   Attention: Provide two to six IP addresses for DNS servers. You should select DNS servers that each site can access locally in the event of network islanding. This is to ensure an islanded site continues to have access to the DNS service. After configuring the grid-wide DNS server list, you can further customize the DNS server list for each node. For details, see information about modifying the DNS configuration in the recovery and maintenance instructions.

4. If required, update the Network Time Protocol (NTP) external sources list.

   For some expansions, you must update the NTP configuration to ensure time synchronization is available. For example:

   - If you added a new site, you might need to add new NTP sources that are accessible to that site.
   - If you chose to configure the NTP settings separately for each grid node, you must manually update the NTP settings.

   If you do need to update the NTP sources list, make it the last step of your expansion.

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

   For more information, see the recovery and maintenance instructions.

Related tasks

Verifying that the Storage Node is active on page 40
Copying the Admin Node database on page 40
Copying Prometheus metrics on page 42
Copying audit logs on page 43

Related information

Administering StorageGRID
Upgrading StorageGRID
Recovery and maintenance

Verifying that the 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 browser.
2. Select **Nodes** > ***Expansion Storage Node*** > **Storage**.
3. Hover over the **Storage Used - Object Data** graph to view the value for **Total**, which is total storage used.
4. Verify that the value of **Total** is increasing as you move your cursor to the right on the graph.

Copying the 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 retain historical information about attributes.

Before you begin

- You must have completed the required expansion steps to add an Admin Node.
- You must have the `Passwords.txt` file.
- You must 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 attribute information for servers and services that are currently part of the system or added later. Historical NMS data is not available unless you manually copy the existing database from the primary Admin Node to the expansion Admin Node.

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

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

Steps

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

      The database and its historical data are copied to the expansion Admin Node. When the copy operation is done, the script starts the expansion Admin Node.
   h. When you no longer require passwordless access to other servers, remove the private key from the SSH agent. Enter:
      `ssh-add -D`

3. Restart the services on the primary Admin Node:

   ```
   service servermanager start
   ```

Copying 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 must be installed and running.
- You must have the `Passwords.txt` file.
- You must 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.

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

Steps

1. Log in to the source Admin Node:
   a. Enter the following command: `ssh admin@grid_node_IP`
   b. Enter the password listed in the `Passwords.txt` file.
   c. Enter the following command to switch to root: `su -`
   d. Enter the password listed in the `Passwords.txt` file.
2. From the source Admin Node, stop the Prometheus service:

   ```
   service prometheus stop
   ```

3. Complete the following steps on the 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`

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

About this task

To make the historical audit messages from other Admin Nodes available on the expansion Admin Node, you must copy the audit log files manually from the primary Admin Node, or another existing Admin Node, to the expansion Admin Node.

Steps

1. From the service laptop, log in to the primary Admin Node:

a. Enter the following command: `ssh admin@`*`primary_Admin_Node_IP`*

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

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

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

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

2. Stop the AMS service to prevent it from creating a new file:

`service ams stop`

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

`cd /var/local/audit/export`

`ls -l`

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

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

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

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

6. Restore the original `audit.log` file:

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

7. Start the AMS service:

   ```
   service ams start
   ```

8. Log out from the server:

   ```
   exit
   ```

9. From the service laptop, log in to the expansion Admin Node:

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

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

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

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

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

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

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

11. Log out from the server:

    ```
    exit
    ```

Contacting technical support

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

About this task

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

Steps

1. Connect to the expansion node that has experienced failures:
   a. Enter the following command:
      `ssh -p 8022 admin@grid_node_IP`
      Note: Port 8022 is the SSH port of the base OS, while port 22 is the SSH port of the Docker container running StorageGRID.
   b. Enter the password listed in the `Passwords.txt` file.
   c. Enter the following command to switch to root: `su -`
   d. Enter the password listed in the `Passwords.txt` file.

   Once you are logged in as root, the prompt changes from `$` to `#`.
2. Depending on the stage the installation reached, retrieve any of the following logs that are available on the grid node:

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

Copyright

Trademark information

How to send comments about documentation and receive update notifications

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

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

doccomments@netapp.com

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

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

You can also contact us in the following ways:

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