



Add grid nodes or site

StorageGRID

NetApp
June 10, 2022

Table of Contents

- Add grid nodes or site 1
 - Add grid nodes to existing site or add new site 1
 - Update subnets for Grid Network 1
 - Deploy new grid nodes 2
 - Perform expansion 8

Add grid nodes or site

Add grid nodes to existing site or add new site

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

What you'll need

- 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.

What you'll need

- 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**.

Grid Network

Configure the subnets that are used on the Grid Network. These entries typically include the subnets for the Grid Network (eth0) for each site in your StorageGRID system as well as any subnets for NTP, DNS, LDAP, or other external servers accessed through the Grid Network gateway.

Subnets

Subnet 1 +

Passphrase

Provisioning
Passphrase

Save

2. In the Subnets list, select 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 select **Save**.

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

Deploy new grid nodes

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

When you expand the grid, the nodes you add do not 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.

What you'll need

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



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

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:

```
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.

What you'll need

- The appliance has been installed in a rack or cabinet, connected to your networks, and powered on.
- You have used the StorageGRID Appliance Installer to complete all of the “configure hardware” steps in the appliance installation and maintenance instructions.
 - [SG100 and SG1000 services appliances](#)
 - [Configure hardware \(SG5600\)](#)
 - [Configure hardware \(SG5700\)](#)
 - [Configure hardware \(SG6000\)](#)

Configuring 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. For example, version 3.6 of the StorageGRID Appliance Installer is compatible with version 11.6 of StorageGRID. (If the versions are not 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 | <ol style="list-style-type: none">a. Unselect the Enable Admin Node discovery check box.b. Enter the IP address manually.c. Click Save.d. Wait for the connection state for the new IP address to become ready. |

| Option | Description |
|--|---|
| Automatic discovery of all connected primary Admin Nodes | <ol style="list-style-type: none"> 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 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 installation and maintenance instructions for your appliance.

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

Home

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

Primary Admin Node connection

Enable Admin Node discovery

Primary Admin Node IP

Connection state

Connection to 172.16.4.210 ready

Node name

Node name

Installation

Current state

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

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




- 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.

- 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.

| 1. Configure storage | | | Running |
|-------------------------------|---|------------------------------------|---------|
| Step | Progress | Status | |
| Connect to storage controller |  | Complete | |
| Clear existing configuration |  | Complete | |
| Configure volumes |  | 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 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:

- For a storage appliance, the installer connects to the storage controller, clears any existing configuration, communicates with SANtricity software 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.

Monitor Installation

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

Connected (unencrypted) to: QEMU

```

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

```

Perform expansion

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

What you'll need

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

- 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 a new site, you must review and update ILM rules before starting the expansion procedure to ensure that object copies are not 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 the ILM rule 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. See the instructions for [managing objects with ILM](#).

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 stages, which run automatically. The set of stages depends on what types of grid nodes are being added and on whether a new site is being added.



Some 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. Do not 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 all nodes that are ready to be added.

Grid Expansion

Approve and configure grid nodes, so that they are added correctly to your StorageGRID system.

[Configure Expansion](#)

Pending Nodes

Grid nodes are listed as pending until they are assigned to a site, configured, and approved.

| | Grid Network MAC Address | Name | Type | Platform | Grid Network IPv4 Address |
|--|--------------------------|--------------|--------------|-----------|---------------------------|
| | 00:50:56:87:68:1a | DC2-ADM1-184 | Admin Node | VMware VM | 172.17.3.184/21 |
| | 00:50:56:87:f1:fc | DC2-S1-185 | Storage Node | VMware VM | 172.17.3.185/21 |
| | 00:50:56:87:54:1e | DC2-S2-186 | Storage Node | VMware VM | 172.17.3.186/21 |
| | 00:50:56:87:6f:0c | DC2-S3-187 | Storage Node | VMware VM | 172.17.3.187/21 |
| | 00:50:56:87:b6:83 | DC2-S4-188 | Storage Node | VMware VM | 172.17.3.188/21 |
| | 00:50:56:87:b3:7d | DC2-ARC1-189 | Archive Node | VMware VM | 172.17.3.189/21 |

2. Select **Configure Expansion**.

The Site Selection dialog box appears.

Site Selection

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.

Site New Existing

Site Name

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. Select **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.

| Grid Network MA |
|---|
| <input type="radio"/> 00:50:56:87:68:1a |
| <input type="radio"/> 00:50:56:87:54:1e |
| <input type="radio"/> 00:50:56:87:6f:0c |
| <input type="radio"/> 00:50:56:87:b6:83 |
| <input type="radio"/> 00:50:56:87:b3:7d |

DC2-S3-187

Storage Node

| Address | Name |
|---------|------|
|---------|------|

Network

| | | |
|----------------|-----------------|------------|
| Grid Network | 172.17.3.187/21 | 172.17.0.1 |
| Admin Network | | |
| Client Network | 10.224.3.187/21 | 10.224.0.1 |

Hardware

VMware VM 8 CPUs 8 GB RAM

Disks

| | | | | |
|--------|--------|--------|--------|--------|
| 107 GB | 107 GB | 107 GB | 107 GB | 107 GB |
|--------|--------|--------|--------|--------|



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. Select **Approve**.

The grid node configuration form appears.

Storage Node Configuration

General Settings

| | |
|-------------|---|
| Site | <input type="text" value="Site A"/> |
| Name | <input type="text" value="DC2-S3-187"/> |
| NTP Role | <input type="text" value="Automatic"/> |
| ADC Service | <input type="text" value="Automatic"/> |

Select "Yes" if this node will replace another node at this site that has the ADC service.

Grid Network

| | |
|---------------------|--|
| Configuration | STATIC |
| IPv4 Address (CIDR) | <input type="text" value="172.17.3.187/21"/> |
| Gateway | <input type="text" value="172.17.0.1"/> |

Admin Network

| | |
|---------------------|------------------------|
| Configuration | STATIC |
| IPv4 Address (CIDR) | <input type="text"/> |
| Gateway | <input type="text"/> |
| Subnets (CIDR) | <input type="text"/> + |

Client Network

| | |
|---------------------|----------------------|
| Configuration | STATIC |
| IPv4 Address (CIDR) | <input type="text"/> |
| Gateway | <input type="text"/> |

Cancel

Save

c. As required, modify the general settings:

- **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 hostname 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.



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 [here](#).

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. Select **Save**.

The approved grid node moves to the Approved Nodes list.

Approved Nodes

Grid nodes that have been approved and have been configured for installation. An approved grid node's configuration can be edited if errors are identified.

| Grid Network MAC Address | Name | Site | Type | Platform | Grid Network IPv4 Address |
|---|------------|--------|--------------|-----------|---------------------------|
| <input type="radio"/> 00:50:56:87:f1:fc | DC2-S1-185 | Site A | Storage Node | VMware VM | 172.17.3.185/21 |
| <input type="radio"/> 00:50:56:87:6f:0c | DC2-S3-187 | Site A | Storage Node | VMware VM | 172.17.3.187/21 |

Passphrase

Enter the provisioning passphrase to change the grid topology of your StorageGRID system.

Provisioning Passphrase

- 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 node are in progress, the Grid Node Status section lists the current status for each grid node.



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.

Grid Expansion

i A new Recovery Package has been generated as a result of the configuration change. Go to the [Recovery Package](#) page to download it.

Expansion Progress

Lists the status of grid configuration tasks required to change the grid topology. These grid configuration tasks are run automatically by the StorageGRID system.

1. Installing Grid Nodes In Progress

Grid Node Status

Lists the installation and configuration status of each grid node included in the expansion.

Q

| Name | Site | Grid Network IPv4 Address | Progress | Stage |
|--------------|--------|---------------------------|---|--------------------------------------|
| DC2-ADM1-184 | Site A | 172.17.3.184/21 | <div style="width: 100%; height: 10px; background-color: #00bcd4;"></div> | Waiting for NTP to synchronize |
| DC2-S1-185 | Site A | 172.17.3.185/21 | <div style="width: 100%; height: 10px; background-color: #00bcd4;"></div> | Waiting for Dynamic IP Service peers |
| DC2-S2-186 | Site A | 172.17.3.186/21 | <div style="width: 100%; height: 10px; background-color: #00bcd4;"></div> | Waiting for NTP to synchronize |
| DC2-S3-187 | Site A | 172.17.3.187/21 | <div style="width: 100%; height: 10px; background-color: #00bcd4;"></div> | Waiting for NTP to synchronize |
| DC2-S4-188 | Site A | 172.17.3.188/21 | <div style="width: 100%; height: 10px; background-color: #00bcd4;"></div> | Waiting for Dynamic IP Service peers |
| DC2-ARC1-189 | Site A | 172.17.3.189/21 | <div style="width: 100%; height: 10px; background-color: #00bcd4;"></div> | Waiting for NTP to synchronize |

| | |
|---|---------|
| 2. Initial Configuration | Pending |
| 3. Distributing the new grid node's certificates to the StorageGRID system. | Pending |
| 4. Starting services on the new grid nodes | Pending |
| 5. Cleaning up unused Cassandra keys | Pending |



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 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.



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. Follow the instructions for adding a Storage Node to an existing site or adding a new site.

Add Storage Node to existing site

If you are adding one or more Storage Nodes to an existing site, monitor the progress of the "Starting Cassandra and streaming data" stage by reviewing the percentage shown in the status message.

4. Starting services on the new grid nodes In Progress

Grid Node Status

Lists the installation and configuration status of each grid node included in the expansion.

⚠ Do not reboot any Storage Nodes during Step 4. The "Starting Cassandra and streaming data" stage might take hours, especially if existing Storage Nodes contain a large amount of object metadata.

Q

| Name ↑↓ | Site ↑↓ | Grid Network IPv4 Address ▼ | Progress ↑↓ | Stage ↑↓ |
|----------------------|----------------------|--|--|--|
| DC1-S4 | Data Center 1 | 10.96.99.55/23 | <div style="width: 90%; height: 10px; background: linear-gradient(to right, #00a0e3, #00a0e3);"></div> | Starting Cassandra and streaming data (90.0% streamed) |
| DC1-S5 | Data Center 1 | 10.96.99.56/23 | <div style="width: 100%; height: 10px; background: #008000;"></div> | Complete |
| DC1-S6 | Data Center 1 | 10.96.99.57/23 | <div style="width: 100%; height: 10px; background: #008000;"></div> | Complete |

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.



Do not 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.

Add new 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.



Do not 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.

- 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. See [Configuration steps after expansion](#).

Copyright Information

Copyright © 2022 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.

RESTRICTED RIGHTS LEGEND: Use, duplication, or disclosure by the government is subject to restrictions as set forth in subparagraph (c)(1)(ii) of the Rights in Technical Data and Computer Software clause at DFARS 252.277-7103 (October 1988) and FAR 52-227-19 (June 1987).

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.