



**StorageGRID® Webscale 10.4**

# **Expansion Guide for Red Hat® Enterprise Linux® Deployments**

April 2017 | 215-11763\_A0  
[doccomments@netapp.com](mailto:doccomments@netapp.com)



# Contents

<b>Planning and preparation .....</b>	<b>4</b>
Supported expansions .....	4
Limitations of the expansion procedure .....	5
Required materials .....	6
Downloading and extracting the StorageGRID Webscale installation files .....	7
Verifying hardware and networking .....	8
Noting the StorageGRID Webscale software version .....	8
Web browser requirements .....	9
<b>Adding storage volumes to Storage Nodes .....</b>	<b>10</b>
What storage volumes are .....	10
Adding direct-attached or SAN volumes to a Storage Node .....	10
<b>Adding grid nodes to an existing site or adding a new site .....</b>	<b>14</b>
Adding a new network or updating the Grid network configuration .....	14
Preparing new RHEL hosts .....	15
Deploying new nodes on RHEL hosts .....	16
Deploying StorageGRID Webscale appliance Storage Nodes .....	16
Connecting to the StorageGRID Webscale Appliance Installer .....	17
Configuring network connections .....	18
Setting the primary Admin Node IP address .....	25
Installing and monitoring the software installation on the appliance .....	25
Performing the expansion .....	28
<b>Configuring your expanded StorageGRID Webscale system .....</b>	<b>32</b>
Verifying that the Storage Node is active .....	33
Copying audit logs .....	34
Copying the Admin Node database .....	35
<b>Contacting technical support .....</b>	<b>37</b>
<b>Glossary .....</b>	<b>38</b>
<b>Copyright information .....</b>	<b>45</b>
<b>Trademark information .....</b>	<b>46</b>
<b>How to send comments about documentation and receive update</b>	
<b>notifications .....</b>	<b>47</b>
<b>Index .....</b>	<b>48</b>

## Planning and preparation

---

You must understand all of the steps required to successfully expand a StorageGRID Webscale system without interrupting system operations.

Read this guide in full before you begin an expansion. It is assumed that you are familiar with the StorageGRID Webscale system and that the expansion has been planned in advance. It is also assumed that all required hardware is installed, connected, and configured to specifications.

This guide is intended for technical personnel trained to install and support the StorageGRID Webscale system. An advanced level of computer literacy is required, including knowledge of Linux/UNIX command shells, networking, virtualization, and server hardware setup and configuration.

**Attention:** Before starting an expansion, you must ensure that any grid tasks from previous provisioning revisions are complete. You cannot start an expansion while a previous expansion, upgrade, or decommissioning procedure is in progress.

### Related information

[\*StorageGRID Webscale 10.4 Grid Primer\*](#)

[\*StorageGRID Webscale 10.4 Administrator Guide\*](#)

## Supported expansions

You can expand your StorageGRID Webscale system to meet your storage requirements by adding storage to virtual Storage Nodes, adding grid nodes to an existing site, or adding a new data center site.

### Adding storage volumes to virtual Storage Nodes

You can add storage volumes (LUNs) to a virtual Storage Node to increase its storage capacity. In some cases, you must add equivalent storage capacity to more than one Storage Node at a time.

You cannot add storage volumes to a Storage Node during upgrade.

**Note:** You cannot expand the available storage on a StorageGRID Webscale Appliance by adding a shelf with additional drives. The maximum storage available on each StorageGRID Webscale appliance Storage Node is fixed, based on the appliance model and original configuration.

### Adding grid nodes to data center sites

You can expand your StorageGRID Webscale system by adding the following types of grid nodes to existing data center site:

- **Admin Nodes:** While each grid can have only one primary Admin Node, you can add multiple non-primary Admin Nodes to each site.
- **Storage Nodes:** In many cases, you must add more than one Storage Node at a time.
- **Archive Nodes:** Each Archive Node supports one interface, either tape through Tivoli Storage Manager (TSM) middleware or the cloud through the S3 API.
- **API Gateway Nodes:** An API Gateway Node provides a load balancing interface to StorageGRID Webscale through which applications can connect to the system.

## Adding data center sites

You can expand your StorageGRID Webscale system by adding a new data center site. When adding a new site, be aware of the following:

- 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 Webscale system remains operational during the expansion, you must verify your information lifecycle management (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. For more information about ILM rules, see the *Administrator Guide*.

### Related tasks

[Adding storage volumes to Storage Nodes](#) on page 10

[Adding grid nodes to an existing site or adding a new site](#) on page 14

### Related information

[StorageGRID Webscale 10.4 Administrator Guide](#)

## Limitations of the expansion procedure

Before you expand your StorageGRID Webscale system, you should be aware of the limitations when adding Admin Nodes, Storage Nodes, and data center sites.

These limitations are enforced by the Grid Expansion maintenance procedure in the StorageGRID Webscale system.

### Admin Nodes

You can add a non-primary Admin Node to a data center site. The primary Admin Node is always deployed during the initial installation; you cannot add a primary Admin Node using the expansion procedure.

### Storage Nodes

Storage expansion requirements depend on the StorageGRID Webscale system's information lifecycle management (ILM) rules, the type of copies made (replicated or erasure-coded), and the StorageGRID Webscale system's topology.

- If an ILM rule creates replicated copies, you must consider the number of copies created and their distribution through storage pools. For example, if an ILM rule requires one replicated copy at data center site 1 (DC1) and one copy at data center site 2 (DC2), you must add an equivalent amount of storage at each site to increase the overall capacity of the StorageGRID Webscale system.
- If an ILM rule creates erasure-coded copies, you must consider the Erasure Coding profile in use and the reason you are adding Storage Nodes. If you are adding Storage Nodes because the system's existing Storage Nodes are nearing capacity, it is not enough to just add new Storage Nodes to the system. You must add enough Storage Nodes to support the erasure-coding scheme in use. For example, to support an Erasure Coding profile using a 6+3 scheme, you must add at least nine Storage Nodes to the system.

## Required materials

Gather the materials listed in the following table to prepare for expansion.

**Important:** You must perform expansion operations using the version of StorageGRID Webscale that is currently running on the grid. Make sure the version shown in the file name for the installation package matches the version of StorageGRID Webscale that is currently installed. To determine the installed version, sign in to the Grid Management Interface, and click **Help > About**.

Item	Notes
StorageGRID Webscale installation package	<p>You must download one of the following StorageGRID Webscale installation packages and extract the files to your service laptop.</p> <ul style="list-style-type: none"> <li>StorageGRID-Webscale-version-rpms-timestamp.SHA.tgz</li> <li>StorageGRID-Webscale-version-rpms-timestamp.SHA.zip</li> </ul>
Service laptop	<p>The StorageGRID Webscale system is installed through a service laptop.</p> <p>The service laptop must have:</p> <ul style="list-style-type: none"> <li>Network port</li> <li>SSH client (for example, PuTTY)</li> <li>Supported browser</li> </ul>
Provisioning passphrase	<p>The passphrase is created and documented when the StorageGRID Webscale system is first installed. The provisioning passphrase is not in the <code>Passwords.txt</code> file.</p>
StorageGRID Webscale documentation	<ul style="list-style-type: none"> <li><i>Administrator Guide</i></li> <li><i>Release Notes</i></li> <li><i>Software Installation Guide</i></li> </ul>

### Related tasks

[Downloading and extracting the StorageGRID Webscale installation files](#) on page 7

### Related references

[Web browser requirements](#) on page 9

### Related information

[StorageGRID Webscale 10.4 Administrator Guide](#)

[StorageGRID Webscale 10.4 Release Notes](#)

[StorageGRID Webscale 10.4 Software Installation Guide for Red Hat Enterprise Linux Deployments](#)

## Downloading and extracting the StorageGRID Webscale installation files

Before you can use StorageGRID Webscale features, you must download the software.

### About this task

You must download one of the following StorageGRID Webscale installation packages and extract the files to your service laptop.

- `StorageGRID-Webscale-version-rpms-timestamp.SHA.tgz`
- `StorageGRID-Webscale-version-rpms-timestamp.SHA.zip`

The compressed files contain the RPM files and scripts for RHEL.

**Important:** You must expand using the same version of StorageGRID Webscale that is currently running on the grid.

### Steps

1. Go to the Software Download page on the NetApp Support Site.  
*[NetApp Downloads: Software](#)*
2. If required, sign in using the username and password for your NetApp account.
3. Scroll to **StorageGRID Webscale**, select **All Platforms**, and click **Go**.  
**Note:** Be sure to select **StorageGRID Webscale**, not **StorageGRID**.
4. Select the StorageGRID Webscale release, and click **View & Download**.
5. From the Software Download section of the page, click **CONTINUE**, and accept the End User License Agreement.
6. Download the appropriate file.
7. Extract the compressed files.
8. Choose the files you need from the following list.

The set of files you need depends on your planned grid topology and how you will deploy your StorageGRID Webscale grid.

**Note:** The paths listed in the table are relative to the top-level directory installed by the install package

Path and file name	Description
<code>/rpms/README</code>	A text file that describes all of the files contained in the StorageGRID Webscale download file.
<code>/rpms/NLF000000.txt</code>	A non-production NetApp License File that you can use for testing and proof of concept deployments.
<code>/rpms/StorageGRID-Webscale-Images-version-SHA.rpm</code>	RPM package for installing the StorageGRID Webscale node images on your RHEL hosts.

Path and file name	Description
<code>/rpms/StorageGRID-Webscale-Service-version-SHA.rpm</code>	RPM package for installing the StorageGRID Webscale host service on your RHEL hosts.
<b>Deployment scripting tools</b>	
<code>/rpms/configure-storagegrid.py</code>	A Python script used to automate the configuration of a StorageGRID Webscale system.
<code>/rpms/configure-storagegrid.sample.json</code>	A sample configuration file for use with the <code>configure-storagegrid.py</code> script.
<code>/rpms/configure-storagegrid.blank.json</code>	A blank configuration file for use with the <code>configure-storagegrid.py</code> script.
<code>/rpms/configure-sga.py</code>	A Python script used to automate the configuration of StorageGRID Webscale appliances.
<code>/rpms/extras/ansible</code>	An Ansible role and example Ansible playbook for configuring RHEL hosts for StorageGRID Webscale container deployment. You can customize the playbook or role as necessary.

## Verifying hardware and networking

Before you begin the expansion of your StorageGRID Webscale system, you must ensure that you have installed and configured the necessary hardware to support the StorageGRID Webscale system.

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 Webscale system.

If your expansion includes adding a new site, you must first add any new subnets to the Grid Network configuration in the Grid Management Interface.

For information about supported versions, see the Interoperability Matrix.

### Related information

[NetApp Interoperability Matrix Tool](#)

## Noting the StorageGRID Webscale software version

For each type of grid node you plan to add, you must know which version of StorageGRID Webscale is currently installed. At the end of the expansion procedure, you must ensure that all grid nodes of the same type are using the same version.

### Steps

1. Sign in to the Grid Management Interface using a supported browser.
2. Select **Grid**.
3. Select **site > node > SSM > Services**.
4. Click the **Main** tab.



5. Scroll to the **Packages** section at the bottom of the page.
6. Note the version number listed under the Version attribute.

When you complete the expansion, you will compare this version to the version installed on the grid nodes you added. You must ensure the version numbers are the same and that any required hotfixes have been installed on the new grid nodes.

7. Repeat these steps for each type of grid node you plan to add.

## Web browser requirements

You must use a supported web browser.

Web browser	Minimum supported version
Google Chrome	54
Microsoft Internet Explorer	11 (Native Mode)
Mozilla Firefox	50

You should set the browser window to a recommended width.

Browser width	Pixels
Minimum	1024
Optimum	1280

## Adding storage volumes to Storage Nodes

You can expand the storage capacity of virtual Storage Nodes with fewer than 16 storage volumes by adding additional storage volumes. Storage Nodes support a maximum of 16 storage volumes.

The maximum storage available on each StorageGRID Webscale appliance Storage Node is fixed, based on the appliance model and original configuration. You cannot expand the available storage on a StorageGRID Webscale appliance by adding a shelf with additional drives.

**Important:** Do not attempt to add storage volumes to a Storage Node while a maintenance procedure, such as a software upgrade, is active.

### What storage volumes are

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 Webscale system and mounted to store objects. Storage volumes are also referred to as *object stores*.

You can view the list of storage volumes on a Storage Node using the Grid Management Interface. Each object store corresponds to a mount point for the same Storage Node. That is, the object store with an ID of 0000 corresponds to the `/var/local/rangedb/0` mount point.

To see the object stores	To see the mount points
Go to <b>Grid</b> > <i>site</i> > <i>Storage Node</i> > <b>LDR</b> > <b>Storage</b> > <b>Overview</b> > <b>Object Stores</b> .	Go to <b>Grid</b> > <i>site</i> > <i>Storage Node</i> > <b>SSM</b> > <b>Resources</b> > <b>Overview</b> > <b>Volumes</b> .

Object metadata is always stored to object store 0.

### Adding direct-attached or SAN volumes to a Storage Node

If a Storage Node includes fewer than 16 storage volumes, you can add volumes to that Storage Node to increase its storage capacity. Each Storage Node can support up to 16 storage volumes.

#### Before you begin

- You must have the `Passwords.txt` file.
- To perform this task, you need specific access permissions. For details, see information about controlling system access with administration user accounts and groups.

#### About this task

To increase the total usable storage capacity of the StorageGRID Webscale system, generally you must add storage capacity to more than one Storage Node.

**Important:** To add storage volumes to a Storage Node, you must first shut down the node. In most cases, 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 RHEL 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 Webscale node migration feature, make the new volumes visible to other RHEL hosts that are migration targets for this Storage Node.

For more information, see the *Software Installation Guide for Red Hat Enterprise Linux Deployments*.

3. Log into the RHEL 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 RHEL host.  
You might have to rescan for devices.

5. Run the following command:

```
sudo storagegrid node stop <node-name>
```

**Note:** This command temporarily disables Storage Node *<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 in step 2.

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 Storage Node `<node-name>`:

```
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 RHEL 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.

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:

```
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. Accept the formatting: `y`

If any of the drives have previously been formatted, you are asked if you want to reformat them.

- c. Choose one of the following:

- To reformat: y
- To skip reformatting: n

14. Again, check that the services start correctly:

- a. View a listing of the status of all services on the server:

**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 Management Interface using a supported browser.

**Note:** Sign in using an account with root or Maintenance permissions.

- b. Select **Grid**.
- c. Select **site > Storage Node > LDR > Storage**.
- d. Click 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. View information about the new storage volume:

- a. Select **Grid**.
- b. Select **site > Storage Node > LDR > Storage**.
- c. Click the **Overview** tab and then the **Main** tab.
- d. View the details in the **Object Stores** table.

#### Related information

[\*StorageGRID Webscale 10.4 Administrator Guide\*](#)

[\*StorageGRID Webscale 10.4 Software Installation Guide for Red Hat Enterprise Linux Deployments\*](#)

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

---

You can add grid nodes to a deployed site, or you can add a new site in a single expansion procedure.

### Before you begin

- Any previously started expansion, upgrade, or decommissioning operation must be completed before you start a new expansion.
- To perform this task, you must have root or maintenance permissions. For details, see information about controlling system access with administration user accounts and groups.

### Steps

1. [Adding a new network or updating the Grid network configuration](#) on page 14
2. [Preparing new RHEL hosts](#) on page 15
3. [Deploying new nodes on RHEL hosts](#) on page 16
4. [Deploying StorageGRID Webscale appliance Storage Nodes](#) on page 16
5. [Performing the expansion](#) on page 28

### Related information

[StorageGRID Webscale 10.4 Administrator Guide](#)

## Adding a new network or updating the Grid network configuration

StorageGRID Webscale maintains a list of network subnets used to communicate between grid nodes on the Grid network (eth0). You can add or configure the subnets that are used on the Grid network. These entries typically include the subnets for the Grid network for each site in your StorageGRID Webscale system.

### Before you begin

- You must be signed in to the Grid Management Interface using a supported browser.
- You must have a service laptop.
- You must have Maintenance or Root Access permissions.
- You must have the provisioning passphrase.
- You must have the IP addresses, in CIDR notation, of the Grid networks to configure.

### About this task

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

### Steps

1. Select **Maintenance > Grid Networks**.  
Currently configured Grid networks are displayed.

2. In the Networks list, add or remove Grid networks to match your grid topology. Click the plus sign to add another Grid network in CIDR notation.
3. Enter the provisioning passphrase, and click **Save**.

The Grid networks you have specified are configured automatically for your StorageGRID Webscale system.

## Preparing new RHEL hosts

If you need additional RHEL hosts to support the CPU, RAM, and storage requirements of the StorageGRID Webscale nodes you want to add to your grid, you prepare them in the same way you prepared the hosts when you first installed them.

Follow the procedures in “Preparing the RHEL hosts” in the *StorageGRID Webscale Software Installation Guide for RHEL Deployments*. Follow only “Preparing the RHEL hosts”; stop when you get to “Performing the installation.”

**Note:** As an alternative, you can deploy additional StorageGRID Webscale nodes on your existing RHEL hosts. If you choose this approach, ensure you have enough CPU, RAM, and storage capacity for the additional nodes.

### Related information

[StorageGRID Webscale 10.4 Software Installation Guide for Red Hat Enterprise Linux Deployments](#)

## Deploying new nodes on RHEL hosts

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

### About this task

To deploy expansion nodes, follow the procedures in “Deploying StorageGRID Webscale nodes on RHEL hosts” in the *StorageGRID Webscale Software Installation Guide for RHEL Deployments*, which include the following steps.

### Steps

1. Create node configuration files.
2. Validate the StorageGRID Webscale node configuration.
3. If you are adding nodes to a new RHEL host, start the StorageGRID Webscale host service.
4. If you are adding nodes to a RHEL host that is already hosting other grid nodes, start the new nodes using the storagegrid host service CLI.

### Example

```
sudo storagegrid node start [<node name> [ <node name>...]]
```

### After you finish

The system pauses deployment while awaiting the node-specific configuration. The system does not automatically continue the deployment. You must go to [Performing the expansion](#) on page 28 to finish adding the new nodes to your grid.

### Related information

[StorageGRID Webscale 10.4 Software Installation Guide for Red Hat Enterprise Linux Deployments](#)

## Deploying StorageGRID Webscale appliance Storage Nodes

If your grid includes StorageGRID Webscale appliances, you deploy them after deploying virtual grid nodes.

### Before you begin

You must have already deployed the primary Admin Node.

### About this task

Deploying a Storage Node appliance involves several major tasks.

**Note:** Optionally, you can use the `configure-sga.py` script to configure and install StorageGRID Webscale appliance Storage Nodes.

### Steps

1. [Connecting to the StorageGRID Webscale Appliance Installer](#) on page 17
2. [Configuring network connections](#) on page 18
3. [Setting the primary Admin Node IP address](#) on page 25



#### 4. [Installing and monitoring the software installation on the appliance](#) on page 25

### Connecting to the StorageGRID Webscale Appliance Installer

To start the appliance software installation, connect to the StorageGRID Webscale Appliance Installer. The installer enables you to configure the connections between the appliance and the Admin network, Grid network, and optional Client network. You can also enter the primary Admin Node IP address if necessary, and monitor the installation progress.

#### Steps

1. Open a browser, and enter the IP address that was provisioned during the physical installation for Management port 1 on the E5600SG controller.

**http://Management\_Port\_IP:8080**

This IP address might be displayed on the appliance's seven-segment display if the StorageGRID Webscale Appliance Installer supports the display of IP addresses.

The StorageGRID Webscale Appliance Installer Home page appears.

NetApp® StorageGRID® Webscale Appliance Installer 1.7.0   Home | Configure Network Connections | Monitor Install | Upgrade Installer | Logs ▼

## StorageGRID® Webscale Appliance Installer Home

**Install procedure:**

1. [Configure Admin network](#) - **Complete**
2. [Configure Grid network](#) - **Complete**
3. [Set StorageGRID® Installer IP](#) - **Connection OK**

StorageGRID® Installer IP:	172.16.4.81	Update
StorageGRID® Node Name:	NetApp-SGA	Update
4. [Begin StorageGRID® node NetApp-SGA install](#)
5. [Monitor install](#)

**Note:** This example shows version 1.7 of the StorageGRID Webscale Appliance Installer. You can locate the version number in the header bar for all versions later than 1.3.

During the initial setup, the status for procedure steps might be “Not Complete.”

2. Verify the status of these items:
  - a. **Configure Admin network:** The Admin network for this StorageGRID Webscale Appliance. To enter or update this information, see “Configuring network connections.”
  - b. **Configure Grid network:** The Grid network for this StorageGRID Webscale Appliance. To enter or update this information, see “Configuring network connections.”
  - c. **Set StorageGRID Installer IP:** The primary Admin Node IP address for this StorageGRID Webscale Appliance. To enter or update this information, see “Setting the primary Admin Node IP address.”

#### Related tasks

[Setting the primary Admin Node IP address](#) on page 25

## Configuring network connections

You use the StorageGRID Webscale Appliance Installer to enter networking information for the E5600SG controller in the StorageGRID Webscale appliance.


### About this task

When you configure network connections for the StorageGRID Webscale appliance, you specify which networks are connected to the E5600SG controller, the IP addresses for the controller's ports, and the routes for each network. If the E5600SG controller includes version 1.7 or later of the StorageGRID Webscale Appliance Installer, you can also specify what type of bonding should be used for the four 10-GbE optical ports on the controller.

**Note:** To see which version of the Appliance Installer is running, look at the title bar of the Appliance Installer web page. If you want to use a later version, contact technical support.

### Steps

1. From the title bar of the StorageGRID Webscale Appliance Installer, click **Configure Network Connections**.



The Configure Network Connections page appears. The version is shown in the header bar.

## Configure Network Connections

The Client network is optional, and is currently enabled. [Disable](#)

### Configure 10GbE Ports

Select how to configure the four 10GbE ports on the StorageGRID® appliance controller:

- **Fixed Port Configuration:** Ports 2 and 4 are used for the Grid network. Ports 1 and 3 are used for the Client network, if it is enabled. You must select whether the ports for each network are bonded using active-backup mode (default) or LACP (802.3ad) mode. You can optionally specify a VLAN ID for either or both networks.
- **Aggregate Port Configuration:** All four 10GbE ports are added to a single LACP (802.3ad) bond. You must specify a VLAN ID for the Grid network and the Client network if it is enabled.

<b>Port Configuration:</b>	<input checked="" type="radio"/> Fixed	<input type="radio"/> Aggregate	
<b>Grid Network:</b>	<input checked="" type="radio"/> Active-Backup	<input type="radio"/> LACP	<input type="checkbox"/> VLAN ID
<b>Client Network:</b>	<input checked="" type="radio"/> Active-Backup	<input type="radio"/> LACP	<input type="checkbox"/> VLAN ID

[Submit](#) [Cancel](#)

### Admin Network

The Admin network uses Management Port 1 on the StorageGRID® appliance controller (the left 1GbE RJ45 port).

#### Loaded DHCP IP configuration

IP	Netmask
10.224.2.24	255.255.248.0

[Use static IP/netmask](#)

Routes	
default via 10.224.0.1	<a href="#">Delete route</a>
10.224.0.0/21 via 10.224.2.24	

[Add route](#)

Ping IP on Admin Network:  [Test](#)

### Grid Network

The Grid network uses a network bond of iSCSI/FC/10GbE Ports on the StorageGRID® appliance controller.

The Grid network is currently using Fixed Port Configuration. 10GbE Ports 2 and 4 are bonded using active-backup mode, with no VLAN ID.

#### Loaded static IP configuration

IP	Netmask
172.16.1.190	255.255.248.0

[Use DHCP for IP/netmask](#) [Change Static IP/netmask](#)

Routes	
47.0.0.0/8 via 172.16.0.1	<a href="#">Delete route</a>
172.17.0.0/21 via 172.16.0.1	<a href="#">Delete route</a>
172.18.0.0/21 via 172.16.0.1	<a href="#">Delete route</a>
172.20.0.0/24 via 172.16.0.2	<a href="#">Delete route</a>
192.168.0.0/21 via 172.16.0.1	<a href="#">Delete route</a>
172.16.0.0/21 via 172.16.1.190	

[Add route](#)

Ping IP on Grid Network:  [Test](#)

### Client Network

The Client network uses a network bond of iSCSI/FC/10GbE Ports on the StorageGRID® appliance controller.

The Client network is currently using Fixed Port Configuration. 10GbE Ports 1 and 3 are bonded using active-backup mode, with no VLAN ID.

#### Loaded DHCP IP configuration

IP	Netmask
47.47.7.213	255.255.248.0

[Use static IP/netmask](#)

Routes	
47.47.0.0/21 via 47.47.7.213	

[Add route](#)

Ping IP on Client Network:  [Test](#)

2. If you are not using a Client network to provide client protocol access to the grid, click the **Disable** button in the Configure Networks section.

## Configure Network Connections

The Client network is optional, and is currently enabled. [Disable](#)

The Client Network section of the page and the Client network settings for the 10-GbE ports are no longer shown. You can click **Enable** to re-enable these settings.

### Related information

[NetApp Support](#)

### Configuring the 10-GbE ports on the controller

You can use the StorageGRID Webscale Appliance Installer to configure the four 10-GbE optical ports on the E5600SG controller.

#### Before you begin

If you plan to use Link Aggregation Control Protocol (LACP IEEE 802.3ad), a virtual LAN (VLAN), or both:

- You must be using version 1.7 or later of the StorageGRID Webscale Appliance Installer.
- The 10-GbE optical ports on the E5600SG controller must be connected to switches that can support these features.
- You must understand how to configure the switches to use these features.

#### About this task

The E5600SG controller in the StorageGRID Webscale appliance includes four 10-GbE optical ports, which connect to the Grid network and, optionally, to the Client network. The ports are bonded together in groups to provide redundant paths to the controller, as follows:

- By default, and for all versions of the StorageGRID Webscale Appliance Installer, controller ports 2 and 4 use active-backup bonding for the Grid network, and ports 1 and 3 use active-backup bonding for the Client network (if used).  
In *active-backup mode*, pairs of ports are bonded together. If the active port fails, its backup port automatically provides a failover connection to the controller and to the appliance Storage Node. In the default configuration (Fixed Port), port 4 provides a backup path for port 2 (Grid network), and port 3 provides a backup path for port 1 (Client network).
- Starting with version 1.7 of the StorageGRID Webscale Appliance Installer, you can configure the four ports to use LACP (802.3ad) instead of active-backup mode.  
In *LACP mode*, the ports are grouped together as a single logical channel between the controller and the network, allowing for higher throughput. In addition, if a port fails, the other ports in the LACP group provide a failover connection to the controller and the appliance Storage Node.

If you have version 1.7 or later of the StorageGRID Webscale Appliance Installer, you can specify how the ports should be grouped, as follows:

- **Fixed Port Configuration** (default) – Ports 2 and 4 control the Grid network, and ports 1 and 3 control the Client network, if enabled. The ports can be bonded as pairs in active-backup mode (default) or in LACP mode.

If you do not need redundant connections, you can use only one port for each network.

**Note:** If you choose to use one port, be aware that an alarm will be raised in the Grid Management Interface after StorageGRID Webscale is installed, indicating that a cable is unplugged. You can safely acknowledge this alarm to clear it.

- **Aggregate Port Configuration** – All four ports are grouped in a single LACP bond, allowing all ports to be used for Grid network and Client network traffic. Selecting Aggregate Port Configuration significantly increases the throughput for each network and provides additional failover paths. When using Aggregate Port Configuration, you must specify a unique VLAN ID for each network. This VLAN ID will be added to each network packet to ensure that network traffic is routed to the correct network.

If you do not need the full bandwidth of four 10-GbE ports, you can use one, two, or three ports, as follows:

- If you choose to use one port, use Port 1.
- If you choose to use two ports, use Ports 1 and 3.
- If you choose to use three ports, use Ports 1, 2, and 3.

This strategy maximizes the chance that some network connectivity will remain available if a single 10-GbE interface chip fails on the E5600SG controller.

**Note:** If you choose to use fewer than four ports, be aware that one or more alarms will be raised in the Grid Management Interface after StorageGRID Webscale is installed, indicating that cables are unplugged. You can safely acknowledge the alarms to clear them.

The table summarizes the options for configuring the 10-GbE ports. These options are available only if the E5600SG controller includes version 1.7 or later of the StorageGRID Webscale Appliance Installer. If your controller includes an earlier version, you cannot disable the Client network, configure the ports for LACP mode, or define VLAN IDs. Go to the next section to configure the Admin, Grid, and Client networks.

Port configuration	Type of bond	Client network enabled	Client network disabled
Fixed (default)	Active-Backup (default)	<ul style="list-style-type: none"> <li>• Ports 2 and 4 use an active-backup bond for the Grid network.</li> <li>• Ports 1 and 3 use an active-backup bond for the Client network.</li> <li>• VLAN IDs can be specified for both networks for the convenience of the network administrator.</li> </ul>	<ul style="list-style-type: none"> <li>• Ports 2 and 4 use an active-backup bond for the Grid network.</li> <li>• Ports 1 and 3 are not used.</li> <li>• A VLAN ID is optional.</li> </ul>
	LACP (802.3ad)	<ul style="list-style-type: none"> <li>• Ports 2 and 4 use an LACP bond for the Grid network.</li> <li>• Ports 1 and 3 use an LACP bond for the Client network.</li> <li>• VLAN IDs can be specified for both networks for the convenience of the network administrator.</li> </ul>	<ul style="list-style-type: none"> <li>• Ports 2 and 4 use an LACP bond for the Grid network.</li> <li>• Ports 1 and 3 are not used.</li> <li>• A VLAN ID is optional.</li> </ul>

Port configuration	Type of bond	Client network enabled	Client network disabled
Aggregate	LACP (802.3ad) only	<ul style="list-style-type: none"> <li>All four ports use a single LACP bond. Packets for the Grid network and the Client network use the same ports.</li> <li>The two VLAN IDs allow Grid network packets to be segregated from Client network packets.</li> </ul>	<ul style="list-style-type: none"> <li>All four port use a single LACP bond, and all are used for the Grid network.</li> <li>The VLAN ID identifies Grid network packets.</li> </ul>

### Steps

1. If the E5600SG controller includes version 1.7 or later of the StorageGRID Webscale Appliance Installer, refer to the table and specify how you want to configure the four 10-GbE ports on the E5600SG controller.

This screen shot shows the default setting (**Fixed** and **Active-Backup** selected). You can optionally specify a VLAN ID for either network.

### Configure 10GbE Ports

Select how to configure the four 10GbE ports on the StorageGRID® appliance controller:

- **Fixed Port Configuration:** Ports 2 and 4 are used for the Grid network. Ports 1 and 3 are used for the Client network, if it is enabled. You must select whether the ports for each network are bonded using active-backup mode (default) or LACP (802.3ad) mode. You can optionally specify a VLAN ID for either or both networks.
- **Aggregate Port Configuration:** All four 10GbE ports are added to a single LACP (802.3ad) bond. You must specify a VLAN ID for the Grid network and the Client network if it is enabled.

Port Configuration:	<input checked="" type="radio"/> Fixed	<input type="radio"/> Aggregate	
Grid Network:	<input checked="" type="radio"/> Active-Backup	<input type="radio"/> LACP	<input type="checkbox"/> VLAN ID
Client Network:	<input checked="" type="radio"/> Active-Backup	<input type="radio"/> LACP	<input type="checkbox"/> VLAN ID
<input type="button" value="Submit"/> <input type="button" value="Cancel"/>			

This screen shot shows **Aggregate** and **LACP** selected for both networks. You must specify a unique VLAN ID for each network.

### Configure 10GbE Ports

Select how to configure the four 10GbE ports on the StorageGRID® appliance controller:

- **Fixed Port Configuration:** Ports 2 and 4 are used for the Grid network. Ports 1 and 3 are used for the Client network, if it is enabled. You must select whether the ports for each network are bonded using active-backup mode (default) or LACP (802.3ad) mode. You can optionally specify a VLAN ID for either or both networks.
- **Aggregate Port Configuration:** All four 10GbE ports are added to a single LACP (802.3ad) bond. You must specify a VLAN ID for the Grid network and the Client network if it is enabled.

Port Configuration:	<input type="radio"/> Fixed	<input checked="" type="radio"/> Aggregate	
Grid Network:	<input type="radio"/> Active-Backup	<input checked="" type="radio"/> LACP	<input checked="" type="checkbox"/> VLAN ID <input type="text" value="328"/>
Client Network:	<input type="radio"/> Active-Backup	<input checked="" type="radio"/> LACP	<input checked="" type="checkbox"/> VLAN ID <input type="text" value="332"/>
<input type="button" value="Submit"/> <input type="button" value="Cancel"/>			

VLAN IDs can be integers between 1 and 4094.

2. When you are done configuring the ports and entering VLAN IDs, click **Submit**.

The Grid Network and Client Network sections of the page are updated. The current configuration for the 10-GbE ports is shown, along with the network's VLAN ID, if entered. These examples show that the Grid and Client networks are using Aggregate Port Configuration with VLAN IDs of 328 and 332.

## Grid Network

The Grid network uses a network bond of iSCSI/FC/10GbE Ports on the StorageGRID® appliance controller.

The Grid network is currently using Aggregate Port Configuration, with a VLAN ID of 328.

## Client Network

The Client network uses a network bond of iSCSI/FC/10GbE Ports on the StorageGRID® appliance controller.

The Client network is currently using Aggregate Port Configuration, with a VLAN ID of 332.

### Configuring the Admin, Grid, and Client networks

You use the StorageGRID Webscale Appliance Installer to enter networking information for the three type of networks that can be connected to the appliance.

#### About this task

Based on how the controller was connected during installation, you might need to configure three networks.

- **Admin (or management) network:** The Admin network allows you to access the StorageGRID Webscale Appliance Installer, which is included on the E5600SG controller. When you configure the Admin network, you specify the IP address and routes for Management Port 1 on the E5600SG controller (the leftmost 1-GbE RJ45 port).

## Admin Network

This network uses Management Port 1 on the StorageGRID® appliance controller (the left 1GbE RJ45 port).

### Loaded static IP configuration

IP	Netmask
10.224.2.24	255.255.248.0

Routes	
default via 10.224.0.1	<input type="button" value="Delete route"/>
10.224.0.0/21 via 10.224.2.24	

Ping IP on Admin Network:

- **Grid network:** During installation, the Grid network allows the appliance to communicate with the installation API on the primary Admin Node. When installation is complete, the Grid network allows the appliance to communicate with the grid. You must specify the IP address and routes for the appliance Storage Node on this network.
- **Client network** (optional): The Client network provides client protocol access to the grid. If the Client network is enabled, you must specify the IP address and routes for the appliance Storage Node on this network.

## Client Network

The Client network uses a network bond of iSCSI/FC/10GbE Ports on the StorageGRID® appliance controller.

The Client network is currently using Fixed Port Configuration. 10GbE Ports 1 and 3 are bonded using active-backup mode, with no VLAN ID.

### Loaded DHCP IP configuration

IP	Netmask
47.47.7.213	255.255.248.0

Use static IP/netmask

Routes
47.47.0.0/21 via 47.47.7.213

Add route

Ping IP on Client Network:

## Steps

- If you are using static IP address allocation, complete the following steps:
  - Click **Change Static IP/netmask**.  
The button name changes to **Save Changes** and a pop-up appears.
  - Enter the IP address for the network, and click **Save Changes**.  
Route information based on the specified IP appears.
  - In the pop-up, click **OK**.
  - If needed, edit the route, and click **Save route**.  
The default route must be configured on the network connection interface.
- If you are using DHCP for IP address allocation, click **Use DHCP for IP/Netmask**.
- Click **Test** next to the **Ping IP on network** text box to verify that the IP address is reachable.

The following message appears for a successful ping:

```
PING OF "10.223.0.1" SUCCESSFUL.

PING 10.224.1.196 (10.224.1.196) from 10.224.1.92 br1: 56(84) bytes
of data.
64 bytes from 10.224.1.196: icmp_req=1 ttl=64 time=0.433 ms

--- 10.224.1.196 ping statistics ---
1 packets transmitted, 1 received, 0% packet loss, time 0ms
rrt min/avg/max/mdev = 0.433/0.433/0.433/0.000 ms
```

The following message appears for a failed ping:

```
PING OF "10.223.0.1" FAILED.

PING 10.223.0.1 (10.223.0.1) from 10.224.1.92 br1: 56(84) bytes of
data.

--- 10.223.0.1 ping statistics ---
1 packets transmitted, 0 received, 100% packet loss, time 0ms
```

- Click **OK** to acknowledge the ping message.



## Setting the primary Admin Node IP address

You can use the StorageGRID Webscale Appliance Installer to set or correct the IP address of the primary Admin Node if necessary.

### Steps

1. Click **Home** to navigate to the main StorageGRID Webscale Appliance Installer.
2. Verify that the **StorageGRID Installer IP** text box shows the IP address of the primary Admin Node.

If this IP address is not shown or should be changed, enter the correct address and click **Update**.



## StorageGRID® Webscale Appliance Installer Home

### Install procedure:

1. [Configure Admin network](#) - **Complete**
2. [Configure Grid network](#) - **Complete**
3. Set StorageGRID® Installer IP - **Connection OK**

StorageGRID® Installer IP:	172.16.4.81	Update
StorageGRID® Node Name:	NetApp-SGA	Update
4. [Begin StorageGRID® node NetApp-SGA install](#)
5. [Monitor install](#)

The Installer 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. If the IP address in the field was automatically populated, (**discovered**) appears next to the field as shown in this example:

3. Set StorageGRID® Installer IP - **Connection OK**

StorageGRID® Installer IP (discovered):	10.96.104.212	Update
StorageGRID® Node Name:	NetApp-SGA	Update
4. [Begin StorageGRID® node NetApp-SGA install](#)

3. Click **OK** in the confirmation dialog box.
4. Verify that the correct **Node Name** is shown in the **StorageGRID Node Name** text box.

To change the name of this Storage Node, enter a new name and click **Update**.

The Node Name is shown on the Grid Nodes page in the Grid Management Interface, and the name is assigned to the Storage Node in the StorageGRID Webscale system.

## Installing and monitoring the software installation on the appliance

The StorageGRID Webscale Appliance Installer provides status until the StorageGRID Webscale software installation is complete.

### Before you begin

You must have already entered the network configuration information and entered the primary Admin Node IP address.

### About this task

When the software installation is complete, the "Convert installation to bare-metal" step starts. After that step is complete, the appliance reboots itself.

**Important:** If your grid includes multiple StorageGRID Webscale appliance Storage Nodes, install the software on no more than five nodes at one time. To avoid bandwidth issues, do not click **Begin StorageGRID node install** *node\_name* for additional appliances until the first appliances appear in the Grid Management Interface.

### Steps

1. From the StorageGRID Webscale Appliance Installer, next to the Set StorageGRID Webscale Installer IP option, click **Begin StorageGRID node install** *node\_name*:

NetApp® StorageGRID® Webscale Appliance Installer 1.7.0    Home | Configure Network Connections | Monitor Install | Upgrade Installer | Logs ▼

## StorageGRID® Webscale Appliance Installer Home

### Install procedure:

1. Configure Admin network - **Complete**
2. Configure Grid network - **Complete**
3. Set StorageGRID® Installer IP - **Connection OK**

StorageGRID® Installer IP:	172.16.4.81	Update
StorageGRID® Node Name:	NetApp-SGA	Update
4. **Begin StorageGRID® node NetApp-SGA install**
5. Monitor install

A list of install operations appears. You can review the installation progress as each operation status changes from “Not started” to “Completed.” The status refreshes every five seconds.

NetApp® StorageGRID® Webscale Appliance Installer 1.7.0

```
Install storage node MM-85-S4-SGA-lab5 - Started
  Initializing installation system - Completed
  Establish connection to array - Completed
    Array status - OK
  Reset array configuration - Not started
  Check operational status of all drives - Not started
  Create primary disk pool - Not started
  Compute volume sizes - Not started
  Create volumes - Not started
  Configure host - Not started
  Create LUN mappings - Not started
  Rename array - Completed
  Create multipath configuration file - Completed
  Save current array configuration - Completed
  Rescan SCSI ports and reload devices - Completed
    Request rescan of SCSI hosts - Completed
    Request restart of multipath service - Completed
  Obtain installer binaries - Completed
    Retrieve installer kernel - Completed
    Retrieve installer initrd - Completed
  Generate StorageGRID installation script config files - Completed
  Install OS image - Completed
    Monitoring - Completed
  Install StorageGRID software - Completed
    Monitoring - Completed
  Convert installation to bare-metal - Completed
  Boot into StorageGRID - Reboot initiated
```

**Note:** The Installer ensures that certain tasks completed in a previous install are not re-run. On the progress page, any tasks that do not need to be re-run remain in a “Not started” state.

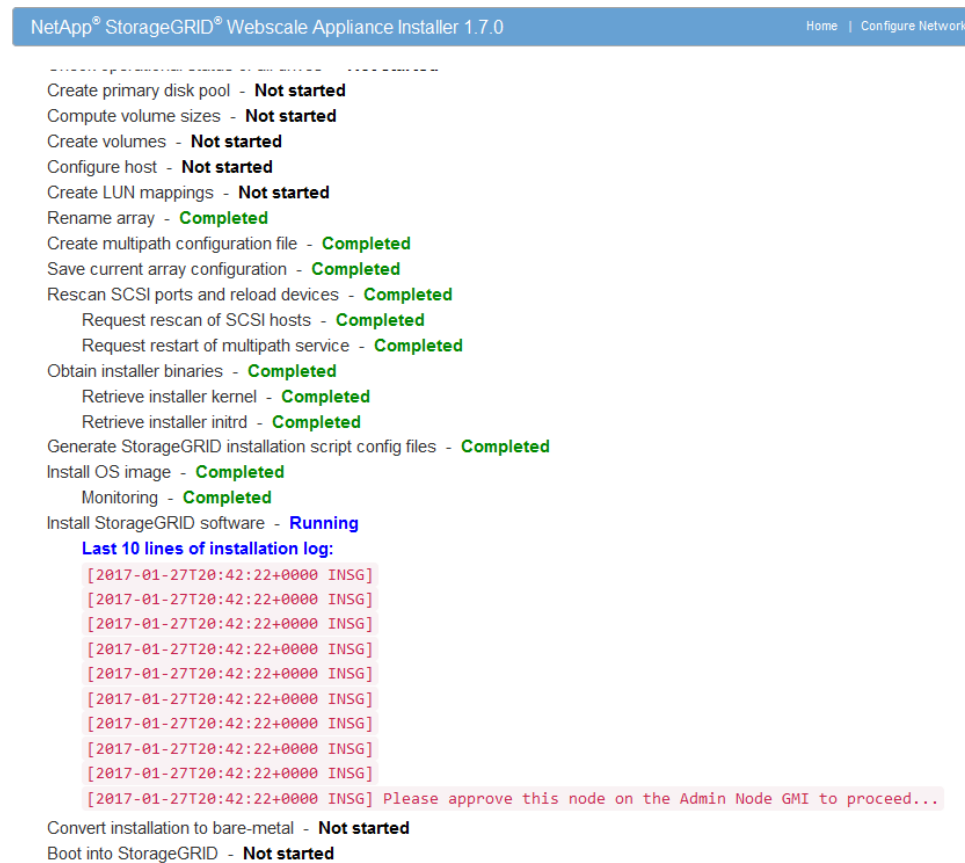
2. To monitor progress, return to the StorageGRID Webscale Appliance Installer.

The Deploy Grid Nodes section shows the installation progress for the Storage Node appliance.

3. Review the installer.

The following occurs:

- When the operating system is being installed, a thumbnail image of the installation screen appears next to the list of operations.
- The StorageGRID Webscale Installer Deploy Grid Nodes status bar changes to blue, indicating a job in progress, and then to green, indicating successful completion.
- The installer shows the last 10 lines of the installation log, which updates every five seconds.
- For new installations, after the last operation, “Boot into StorageGRID” is displayed on the last line, indicating a status of initiating the reboot.



```

NetApp® StorageGRID® Webscale Appliance Installer 1.7.0
Home | Configure Network

-----
Create primary disk pool - Not started
Compute volume sizes - Not started
Create volumes - Not started
Configure host - Not started
Create LUN mappings - Not started
Rename array - Completed
Create multipath configuration file - Completed
Save current array configuration - Completed
Rescan SCSI ports and reload devices - Completed
    Request rescan of SCSI hosts - Completed
    Request restart of multipath service - Completed
Obtain installer binaries - Completed
    Retrieve installer kernel - Completed
    Retrieve installer initrd - Completed
Generate StorageGRID installation script config files - Completed
Install OS image - Completed
    Monitoring - Completed
Install StorageGRID software - Running
Last 10 lines of installation log:
[2017-01-27T20:42:22+0000 INSG]
[2017-01-27T20:42:22+0000 INSG]
[2017-01-27T20:42:22+0000 INSG]
[2017-01-27T20:42:22+0000 INSG]
[2017-01-27T20:42:22+0000 INSG]
[2017-01-27T20:42:22+0000 INSG]
[2017-01-27T20:42:22+0000 INSG]
[2017-01-27T20:42:22+0000 INSG]
[2017-01-27T20:42:22+0000 INSG]
[2017-01-27T20:42:22+0000 INSG] Please approve this node on the Admin Node GMI to proceed...
Convert installation to bare-metal - Not started
Boot into StorageGRID - Not started

```

4. Repeat this process for any additional appliances.

### After you finish

The deployment stops while awaiting the node-specific configuration. The system does not automatically continue the deployment. You must proceed to the next task to trigger it.

## Performing the expansion

When you perform an expansion, the new grid nodes or the new site is added to your existing StorageGRID Webscale deployment. You must download an updated copy of the Recovery Package file while the expansion process is completing.

### Before you begin

- You must be signed in to the Grid Management Interface using a supported browser.
- You must have a service laptop.
- You must have Maintenance or Root Access permissions.
- You must have the provisioning passphrase.
- You must have deployed the virtual or appliance grid nodes that are being added in this expansion.

### About this task

The expansion process includes three phases:

1. To start the expansion, you specify whether you are adding new grid nodes or a new site, approve the grid nodes you want to add, enter the provisioning passphrase, and click **Expand**.
2. While the expansion process is running, you download a new Recovery Package file and monitor the expansion steps on each grid node.
3. When the expansion is complete, you verify that all new grid nodes have the correct version of StorageGRID Webscale. Then, you complete additional integration and configuration steps.

During the expansion, you can look at the Expansion Progress section to see the grid configuration tasks required to change the grid topology, along with the status of each task. The set of grid configuration tasks depends on the grid nodes being added, and whether a new site is being added. Some grid configuration tasks might take a significant amount of time to run on a large grid. These grid configuration tasks are run automatically in the background.

## Grid Expansion

## 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 Webscale system.

1. Installing Grid Nodes	Completed
2. Initial Configuration	Completed
3. Configuring Cassandra for the new site	Completed
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.

Name	Grid Network IPv4 Address	Progress	Stage
kgregory-adm1-98-245	10.96.98.245/23	<div></div>	Complete
kgregory-s1-98-246	10.96.98.246/23	<div></div>	Complete
kgregory-s2-98-247	10.96.98.247/23	<div></div>	Complete
kgregory-s3-98-248	10.96.98.248/23	<div></div>	Expanding the Cassandra cluster

5. Distributing the new grid node's certificates to the StorageGRID Webscale system.	Pending
--	---------

The example includes the following tasks:

1. Installing grid nodes
2. Initial configuration
3. Configuring Cassandra for the new site
 

**Note:** If the Cassandra database is relatively empty, this step does not take more than a few minutes or hours. However, if the Cassandra database has a large quantity of objects, this step might take several hours or longer.
4. Starting services on the new grid nodes
5. Distributing the new grid node's certificates to the StorageGRID Webscale system

When each task is in progress, the Grid Node Status section lists the current status for each grid node.

## Steps

1. Select **Maintenance > Expansion**.
2. Click **Configure Expansion**.
3. Select the type of expansion you are starting:
  - If you are adding grid nodes to an existing site, select the existing site to add the grid nodes to from the **Site** drop-down list.
  - If you are adding a new site, select **--New Site--** from the **Site** drop-down list, and enter the name of the new site.
4. Click **Save**.
5. Approve the pending grid nodes:
  - a. Review the **Pending Nodes** list, and confirm that it shows all of the grid nodes you deployed.
 

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

- b. Select the radio button next to the pending grid node you want to approve.
- c. Click **Approve**.

The Grid Node Configuration page appears, with the appropriate node name in the title.

- d. In **General Settings**, modify settings for the following properties, as necessary:
  - **Site:** The name of the site the grid node will be associated with.
  - **Name:** The name that will be displayed in the Grid Management Interface.
  - **NTP Role:** The grid node's Network Time Protocol (NTP) role. The options are **Automatic**, **Primary**, and **Client**. Selecting **Automatic** assigns the **Primary** option to Admin Nodes, the **Primary** option to Storage Nodes that have an ADC service, and the **Client** option to all other grid node types.
 

**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:** This field is only available if the Node Configuration dialog box was opened for Storage Nodes. For Storage Nodes, this option determines whether the selected node will run the Administrative Domain Controller service. Select **Automatic** to have this option applied automatically by the system as required, or select **Yes** or **No** to explicitly set this option for the grid node. For example, because you cannot add an ADC service to a node after it has been deployed, you might need to select **Yes** if this expansion will be followed by a decommissioning of Storage Nodes containing the ADC service at this site.
- e. In **Grid Network**, modify settings for the following properties, as necessary:
  - **IPv4 Address (CIDR):** The CIDR network address for the eth0 Grid Network interface. For example: 172.16.10.100/24
  - **Gateway:** The default gateway of the grid node. For example: 172.16.10.1
- f. If you want to configure the Admin Network for the grid node, add or update settings in the **Admin Network** section, as necessary.
 

Enter the destination subnetworks of the routes out of this interface in the **Subnets (CIDR)** text box.
- g. If you want to configure the Client Network for the grid node, add or update settings in the **Client Network** section, as necessary.
- h. Click **Save**.

The approved grid node entry moves to the Approved Nodes list.

**Note:** If you need to modify the properties of an approved grid node, select its radio button and click **Edit**.

- i. Repeat these steps for each pending grid node you want to approve.
6. When you have approved the new grid nodes, enter the **Provisioning Passphrase**.
7. Click **Expand**.
 

The page updates to display the status of the expansion procedure.
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 Webscale 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**.  
The download starts immediately.
- c. When the download completes, open the .zip file and confirm it includes a `gpt-backup` directory and a second .zip file. Then, extract this inner .zip file and confirm you can open the `passwords.txt` file.
- d. Copy the downloaded Recovery Package file (.zip) to two safe, secure, and separate locations.  
**Important:** The Recovery Package file must be secured because it contains encryption keys and passwords that can be used to obtain data from the StorageGRID Webscale system.

9. When all expansion tasks are complete, verify that all new nodes have the correct version of StorageGRID Webscale and that any applicable hotfixes were automatically applied.
  - a. Select **Grid**.
  - b. Select **site > Node > SSM > Services**.
  - c. Click the **Main** tab.
  - d. Scroll to the **Packages** section at the bottom of the page.
  - e. Confirm that the version number listed under the Version attribute is the same as the version number for existing grid nodes of the same type.
  - f. If the versions are different, apply any applicable hotfixes manually.  
See the hotfix release notes for installation instructions.

#### 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 Webscale system](#)* on page 32

## Configuring your expanded StorageGRID Webscale system

---

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

### About this task

The following list describes the post-expansion configuration tasks for each grid node type. You only need to complete the tasks that apply to 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. If the expansion added Storage Nodes, complete the following configuration:
  - 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 the active ILM policy. Otherwise, the new storage will not be used by the StorageGRID Webscale system. For information on configuring ILM policies, see the *Administrator Guide*.
  - b. Verify that the Storage Node is active.  
 You must verify that the Storage Nodes you add are ingesting objects. For more information, see [Verifying that the Storage Node is active](#) on page 33.
2. If the expansion added Admin Nodes, complete the following configuration:
  - a. 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](#) on page 35.
  - b. 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. For more information, see [Copying the audit logs](#) on page 34.
  - c. Optionally, configure access to audit shares.  
 You can configure access to the system for auditing purposes, either through a Common Internet File System (CIFS) or Network File System (NFS ) file share. For more information, see the *Administrator Guide*.
3. If the expansion added Archive Nodes, complete the following configuration:
  - 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. For more information, see the *Administrator Guide*.
  - b. Update the ILM policy.  
 You must update your ILM policy in order to archive object data through the new Archive Node. For more information, see the *Administrator Guide*.
  - 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 *Administrator Guide*.

#### 4. Configure the Domain Name System (DNS).

If you chose to configure the DNS settings separately for each grid node, you must specify the DNS information for the expansion Admin Node. See information about modifying the DNS configuration for a single grid node in the *Maintenance Guide*.

**Important:** 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 more information, see “Modifying the DNS configuration for a single grid node” in the *Maintenance Guide*.

#### 5. Change the preferred sender for email notifications.

If you want to change the preferred sender, you can update your configuration to make the expansion Admin Node the preferred sender. Otherwise, the existing Admin Node configured as the preferred sender continues to send notifications and AutoSupport messages. For information on configuring the preferred sender, see the *Administrator Guide*.

#### 6. 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.

**Important:** 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 network islanded from the rest of the grid.

For more information, see the *Maintenance Guide*.

#### Related information

[StorageGRID Webscale 10.4 Administrator Guide](#)


[StorageGRID Webscale 10.4 Maintenance Guide for Red Hat Enterprise Linux Deployments](#)

## Verifying that the Storage Node is active

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

#### Steps

1. Sign in to the Grid Management Interface using a supported browser.
2. Select **Grid**.
3. Select **Site > Expansion Storage Node > LDR > Storage**.

4. Select the **Overview** tab and then the **Main** tab.
5. For the **Total Data** attribute, click the chart icon .

The StorageGRID Webscale system displays a graph that shows the total data on the Storage Node for the last day.

6. Verify that the value of the **Total Data** attribute is increasing.

## 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 Webscale 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.

Once logged in as root, the prompt changes from `$` to `#`.
2. Stop the AMS service to prevent it from creating a new file:
 

```
/etc/init.d/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:

```
/etc/init.d/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.

Once 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
```

## 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 Webscale 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 Management Interface is inaccessible.

Use these steps to stop the MI 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:
 

```
/etc/init.d/mi 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:
 

```
/etc/init.d/mi stop
```
  - c. Add the SSH private key to the SSH agent. Enter:
 

```
ssh-add
```
  - d. Enter the SSH Access Password listed in the `Passwords.txt` file.
  - e. Copy the database from the source Admin Node to the expansion Admin Node:
 

```
/usr/local/mi/bin/mi-clone-db.sh Source_Admin_Node_IP
```
  - f. When prompted, confirm that you want to overwrite the MI database on the 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. The following status appears:

```
Database cloned, starting services
starting mi ... done
```
  - g. When you no longer require passwordless access to other servers, remove the private key from the SSH agent. Enter:
 

```
ssh-add -D
```
3. Restart the MI service on the source Admin Node.
 

```
/etc/init.d/mi start
```

## 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.

### Step

1. Gather the following files from each RHEL host supporting an expansion node that has experienced failures:
  - `/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)

## Glossary

---

**ACL**

Access control list. Specifies which users or groups of users are allowed to access an object and what operations are permitted, for example, read, write, and execute.

**active-backup mode**

A method for bonding two physical ports together for redundancy.

**ADC service**

Administrative Domain Controller. The ADC service maintains topology information, provides authentication services, and responds to queries from the LDR, CMN, and CLB services. The ADC service is present on each of the first three Storage Nodes installed at a site.

**ADE**

Asynchronous Distributed Environment. Proprietary development environment used as a framework for services within the StorageGRID Webscale system.

**Admin Node**

The Admin Node provides services for the web interface, system configuration, and audit logs. See also, *primary Admin Node*.

**Amazon S3**

Proprietary web service from Amazon for the storage and retrieval of data.

**AMS service**

Audit Management System. The AMS service monitors and logs all audited system events and transactions to a text log file. The AMS service is present on the Admin Node.

**API Gateway Node**

An API Gateway Node provides load balancing functionality to the StorageGRID Webscale system and is used to distribute the workload when multiple client applications are performing ingest and retrieval operations. API Gateway Nodes include a Connection Load Balancer (CLB) service.

**ARC service**

Archive. The ARC service provides the management interface with which you configure connections to external archival storage such as the cloud through an S3 interface or tape through TSM middleware. The ARC service is present on the Archive Node.

**Archive Node**

The Archive Node manages the archiving of object data to an external archival storage system.

**atom**

Atoms are the lowest level component of the container data structure, and generally encode a single piece of information.

**audit message**

Information about an event occurring in the StorageGRID Webscale system that is captured and logged to a file.

**Base64**

A standardized data encoding algorithm that enables 8-bit data to be converted into a format that uses a smaller character set, enabling it to safely pass through legacy systems

that can process only basic (low order) ASCII text excluding control characters. See RFC 2045 for more details.

### **bundle**

A structured collection of configuration information used internally by various components of the StorageGRID Webscale system. Bundles are structured in container format.

### **Cassandra**

An open-source database that is scalable and distributed, provides high availability, and handles large amounts of data across multiple servers.

### **CBID**

Content Block Identifier. A unique internal identifier of a piece of content within the StorageGRID Webscale system.

### **CDMI**

Cloud Data Management Interface. An industry-standard defined by SNIA that includes a RESTful interface for object storage. For more information, see [www.snia.org/cdmf](http://www.snia.org/cdmf).

### **CIDR**

Classless Inter-Domain Routing. A notation used to compactly describe a subnet mask used to define a range of IP addresses. In CIDR notation, the subnet mask is expressed as an IP address in dotted decimal notation, followed by a slash and the number of bits in the subnet. For example, 192.0.2.0/24.

### **CLB service**

Connection Load Balancer. The CLB service provides a gateway into the StorageGRID Webscale system for client applications connecting through HTTP. The CLB service is part of the API Gateway Node.

### **Cloud Data Management Interface**

See *CDMI*.

### **CMN service**

Configuration Management Node. The CMN service manages system-wide configurations and grid tasks. The CMN service is present on the primary Admin Node.

### **CMS service**

Content Management System. The CMS service carries out the operations of the active ILM policy's ILM rules, determining how object data is protected over time. The CMS service is present on the Storage Node.

### **command**

In HTTP, an instruction in the request header such as GET, HEAD, DELETE, OPTIONS, POST, or PUT. Also known as an HTTP method.

### **container**

Created when an object is split into segments. A container object lists the header information for all segments of the split object and is used by the LDR service to assemble the segmented object when it is retrieved by a client application.

### **content block ID**

See *CBID*.

### **content handle**

See *UUID*.

### **CSTR**

Null-terminated, variable-length string.

**DC**

Data Center site.

**DDS service**

Distributed Data Store. The DDS service interfaces with the distributed key-value store and manages object metadata. It distributes metadata copies to multiple instances of the distributed key-value store so that metadata is always protected against loss.

**distributed key value store**

Data storage and retrieval that unlike a traditional relational database manages data across grid nodes.

**DNS**

Domain Name System.

**enablement layer**

Used during installation to customize the Linux operating system installed on each grid node. Only the packages needed to support the services hosted on the grid node are retained, which minimizes the overall footprint occupied by the operating system and maximizes the security of each grid node.

**Fibre Channel**

A networking technology primarily used for storage.

**Grid ID signed text block**

A Base64 encoded block of cryptographically signed data that contains the grid ID. See also, *provisioning*.

**grid node**

The basic software building block for the StorageGRID Webscale system, for example, Admin Node or Storage Node. Each grid node type consists of a set of services that perform a specialized set of tasks.

**grid task**

System-wide scripts used to trigger various actions that implement specific changes to the StorageGRID Webscale system. For example, most maintenance and expansion procedures involve running grid tasks. Grid tasks are typically long-term operations that span many entities within the StorageGRID Webscale system. See also, *Task Signed Text Block*.

**ILM**

Information Lifecycle Management. A process of managing content storage location and duration based on content value, cost of storage, performance access, regulatory compliance, and other factors. See also, *Admin Node* and *storage pool*.

**LACP**

Link Aggregation Control Protocol. A method for bundling two or more physical ports together to form a single logical channel.

**LAN**

Local Area Network. A network of interconnected computers that is restricted to a small area, such as a building or campus. A LAN can be considered a node to the Internet or other wide area network.

**latency**

Time duration for processing a transaction or transmitting a unit of data from end to end. When evaluating system performance, both throughput and latency need to be considered. See also, *throughput*.



**LDR service**

Local Distribution Router. The LDR service manages the storage and transfer of content within the StorageGRID Webscale system. The LDR service is present on the Storage Node.

**LUN**

See *object store*.

**mDNS**

Multicast Domain Name System. A system for resolving IP addresses in a small network where no DNS server has been installed.

**metadata**

Information related to or describing an object stored in the StorageGRID Webscale system; for example, ingest time.

**MLAG**

Multi-Chassis Link Aggregation Group. A type of link aggregation group that uses two (and sometimes more) switches to provide redundancy in case one of the switches fails.

**MTU**

Maximum transmission unit. The largest size packet or frame that can be sent in any transmission.

**namespace**

A set whose elements are unique names. There is no guarantee that a name in one namespace is not repeated in a different namespace.

**nearline**

A term describing data storage that is neither “online” (implying that it is instantly available, like spinning disk) nor “offline” (which can include offsite storage media). An example of a nearline data storage location is a tape that is loaded in a tape library, but is not mounted.

**NFS**

Network File System. A protocol (developed by SUN Microsystems) that enables access to network files as if they were on local disks.

**NMS service**

Network Management System. The NMS service provides a web-based interface for managing and monitoring the StorageGRID Webscale system. The NMS service is present on the Admin Node. See also, *Admin Node*.

**node ID**

An identification number assigned to a service within the StorageGRID Webscale system. Each service (such as an NMS service or ADC service) must have a unique node ID. The number is set during system configuration and tied to authentication certificates.

**NTP**

Network Time Protocol. A protocol used to synchronize distributed clocks over a variable latency network, such as the Internet.

**object**

An artificial construct used to describe a system that divides content into data and metadata.

**object segmentation**

A StorageGRID Webscale process that splits a large object into a collection of small objects (segments) and creates a segment container to track the collection. The segment container contains the UUID for the collection of small objects as well as the header

information for each small object in the collection. All of the small objects in the collection are the same size. See also, *segment container*.

### **object storage**

An approach to storing data where the data is accessed by unique identifiers and not by a user-defined hierarchy of directories and files. Each object has both data (for example, a picture) and metadata (for example, the date the picture was taken). Object storage operations act on entire objects as opposed to reading and writing bytes as is commonly done with files, and provided via APIs or HTTP instead of NAS (CIFS/NFS) or block protocols (iSCSI/ FC/FCOE).

### **object store**

A configured file system on a disk volume. The configuration includes a specific directory structure and resources initialized at system installation.

### **OID**

Object Identifier. The unique identifier of an object.

### **primary Admin Node**

Admin Node that hosts the CMN service. Each StorageGRID Webscale system has only one primary Admin Node. See also, *Admin Node*.

### **provisioning**

The process of generating a new or updated Recovery Package and GPT repository. See also, *SAID*.

### **quorum**

A simple majority:  $50\% + 1$ . Some system functionality requires a quorum of the total number of a particular service type.

### **Recovery Package**

A .zip file containing deployment-specific files and software needed to install, expand, upgrade, and maintain a StorageGRID Webscale system. The package also contains system-specific configuration and integration information, including server hostnames and IP addresses, and highly confidential passwords needed during system maintenance, upgrade, and expansion. See also, *SAID*.

### **SAID**

Software Activation and Integration Data. The component in the Recovery Package that includes the `Passwords.txt` file.

### **SATA**

Serial Advanced Technology Attachment. A connection technology used to connect server and storage devices.

### **SCSI**

Small Computer System Interface. A connection technology used to connect servers and peripheral devices, such as storage systems.

### **segment container**

An object created by the StorageGRID Webscale system during the segmentation process. Object segmentation splits a large object into a collection of small objects (segments) and creates a segment container to track the collection. A segment container contains the UUID for the collection of segmented objects as well as the header information for each segment in the collection. When assembled, the collection of segments creates the original object. See also, *object segmentation*.

### **server**

Used when specifically referring to hardware. Might also refer to a virtual machine.

**service**

A unit of the StorageGRID Webscale system, such as the ADC service, NMS service, or SSM service. Each service performs unique tasks critical to the normal operations of a StorageGRID Webscale system.

**SQL**

Structured Query Language. An industry-standard interface language for managing relational databases. An SQL database is one that supports the SQL interface.

**ssh**

Secure Shell. A UNIX shell program and supporting protocols used to log in to a remote computer and run commands over an authenticated and encrypted channel.

**SSL**

Secure Socket Layer. The original cryptographic protocol used to enable secure communications over the Internet. See also, *TLS*.

**SSM service**

Server Status Monitor. A component of the StorageGRID Webscale software that monitors hardware conditions and reports to the NMS service. Every grid node runs an instance of the SSM service.

**Storage Node**

The Storage Node provides storage capacity and services to store, move, verify, and retrieve objects stored on disks.

**storage pool**

The element of an ILM rule that determines the location where an object is stored.

**storage volume**

See *object store*

**StorageGRID**

A registered trademark of NetApp, Inc., used for an object storage grid architecture and software system.

**Task Signed Text Block**

A Base64 encoded block of cryptographically signed data that provides the set of instructions that define a grid task.

**TCP/IP**

Transmission Control Protocol/Internet Protocol. A process of encapsulating and transmitting packet data over a network. It includes positive acknowledgment of transmissions.

**throughput**

The amount of data that can be transmitted or the number of transactions that can be processed by a system or subsystem in a given period of time. See also, *latency*.

**Tivoli Storage Manager**

IBM storage middleware product that manages storage and retrieval of data from removable storage resources.

**TLS**

Transport Layer Security. A cryptographic protocol used to enable secure communications over the Internet. See RFC 2246 for more details.

**transfer syntax**

The parameters, such as the byte order and compression method, needed to exchange data between systems.

**URI**

Universal Resource Identifier. A generic set of all names or addresses used to refer to resources that can be served from a computer system. These addresses are represented as short text strings.

**UTC**

A language-independent international abbreviation, UTC is neither English nor French. It means both “Coordinated Universal Time” and “Temps Universel Coordonné.” UTC refers to the standard time common to every place in the world.

**UUID**

Universally Unique Identifier. Unique identifier for each piece of content in the StorageGRID Webscale system. UUIDs provide client applications with a content handle that permits them to access content in a way that does not interfere with the StorageGRID Webscale system’s management of that same content. A 128-bit number that is guaranteed to be unique. See RFC 4122 for more details.

**virtual machine (VM)**

A software platform that enables the installation of an operating system and software, substituting for a physical server and permitting the sharing of physical server resources among several virtual servers.

**VLAN**

Virtual local area network (or virtual LAN). A group of devices that are located on different LAN segments but are configured to communicate as if they were attached to the same network switch.

**WAN**

Wide area network. A network of interconnected computers that covers a large geographic area, such as a country.

**XFS**

A scalable, high-performance journaled file system originally developed by Silicon Graphics.

**XML**

Extensible Markup Language. A text format for the extensible representation of structured information; classified by type and managed like a database. XML has the advantages of being verifiable, human readable, and easily interchangeable between different systems.

## Copyright information

---

Copyright © 1994–2017 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

---

Active IQ, AltaVault, Arch Design, ASUP, AutoSupport, Campaign Express, Clustered Data ONTAP, Customer Fitness, Data ONTAP, DataMotion, Element, Fitness, Flash Accel, Flash Cache, Flash Pool, FlexArray, FlexCache, FlexClone, FlexPod, FlexScale, FlexShare, FlexVol, FPolicy, Fueled by SolidFire, GetSuccessful, Helix Design, LockVault, Manage ONTAP, MetroCluster, MultiStore, NetApp, NetApp Insight, OnCommand, ONTAP, ONTAPI, RAID DP, RAID-TEC, SANscreen, SANshare, SANtricity, SecureShare, Simplicity, Simulate ONTAP, Snap Creator, SnapCenter, SnapCopy, SnapDrive, SnapIntegrator, SnapLock, SnapManager, SnapMirror, SnapMover, SnapProtect, SnapRestore, Snapshot, SnapValidator, SnapVault, SolidFire, SolidFire Helix, StorageGRID, SyncMirror, Tech OnTap, Unbound Cloud, and WAFL and other names are trademarks or registered trademarks of NetApp, Inc., in the United States, and/or other countries. All other brands or products are trademarks or registered trademarks of their respective holders and should be treated as such. A current list of NetApp trademarks is available on the web.

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

## How to send comments about documentation and receive update notifications

---

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

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

[doccomments@netapp.com](mailto: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., 495 East Java Drive, Sunnyvale, CA 94089 U.S.
- Telephone: +1 (408) 822-6000
- Fax: +1 (408) 822-4501
- Support telephone: +1 (888) 463-8277

# Index

10-GbE ports  
     configuring on controller [20](#)

## A

active-backup mode  
     10-GbE controller ports [20](#)  
 Admin Node  
     configuring after expansion [32](#)  
     copying audit logs [34](#)  
     copying database [35](#)  
 appliance  
     configuring 10-GbE ports [20](#)  
     connecting to the StorageGRID Webscale Appliance Installer [17](#)  
     deploying [16](#)  
     installing [25](#)  
 Archive Node  
     configuring after expansion [32](#)  
 audit logs  
     copying [34](#)

## B

browsers  
     supported [9](#)

## C

comments  
     how to send feedback about documentation [47](#)  
 connectivity  
     configuring networks [18, 23](#)  
 controller ports  
     configuring [20](#)

## D

database  
     copying from primary Admin Node [35](#)  
 deploying  
     appliance Storage Nodes [16](#)  
 DNS  
     configuring after expansion [32](#)  
 documentation  
     how to receive automatic notification of changes to [47](#)  
     how to send feedback about [47](#)  
 downloading installation files  
     steps for [7](#)

## E

email notifications  
     configuring after expansion [32](#)  
 expansion  
     adding data center sites [4](#)

adding grid nodes [4](#)  
 adding storage volumes [4](#)  
 configuring system after [32](#)  
 limitations of [5](#)  
 performing [28](#)  
 planning for [4](#)  
 required materials for [6](#)  
 steps to complete after [32](#)  
 StorageGRID Webscale version [6](#)  
 types of [4](#)  
 expansion nodes  
     deploying in RHEL [16](#)

## F

feedback  
     how to send comments about documentation [47](#)

## G

Grid network  
     adding subnets [14](#)  
 grid nodes  
     deploying in RHEL [16](#)  
     site expansion and [14](#)

## H

hardware  
     verifying [8](#)

## I

information  
     how to send feedback about improving documentation [47](#)  
 installation  
     setting IP address [25](#)  
 installation files  
     downloading [7](#)  
 installer  
     configuring 10-GbE ports [20](#)

## L

LACP mode  
     10-GbE ports [20](#)  
 limitations  
     to expansion [5](#)  
 logs, audit  
     copying [34](#)  
 LUN [10](#)

## N

network  
     adding or updating configuration [14](#)



- network connections
  - configuring [18, 23](#)
- networking
  - verifying [8](#)
- nodes
  - deploying with StorageGRID Webscale Appliance Installer [16](#)
- NTP
  - configuring after expansion [32](#)

## P

- preparing
  - RHEL hosts for expansion [15](#)

## R

- RHEL
  - deploying new grid nodes [16](#)
- RHEL hosts
  - preparing [15](#)

## S

- site expansion
  - adding grid nodes [14](#)
- software
  - checking version of [8](#)
  - network connections [18, 23](#)
  - setting installer IP address [25](#)
- storage
  - adding [10](#)
- Storage Node
  - adding storage to [10](#)

- configuring after expansion [32](#)
  - deploying [16](#)
  - verifying active status of [33](#)
- storage volumes
  - described [10](#)
- StorageGRID Webscale
  - version for expansion [6](#)
- StorageGRID Webscale Appliance Installer
  - connecting to [17](#)
- subnets
  - adding for Grid network [14](#)
- suggestions
  - how to send feedback about documentation [47](#)

## T

- technical support
  - contacting [37](#)
- Twitter
  - how to receive automatic notification of documentation changes [47](#)

## V

- version
  - checking installed [8](#)
- VLAN support
  - 10-GbE controller ports [20](#)

## W

- web browsers
  - supported [9](#)