



Install and upgrade software

StorageGRID

NetApp

October 03, 2025

This PDF was generated from <https://docs.netapp.com/us-en/storagegrid-116/upgrade/index.html> on October 03, 2025. Always check docs.netapp.com for the latest.

Table of Contents

Install and upgrade software	1
Upgrade StorageGRID software	1
Upgrade StorageGRID software: Overview	1
What's new in StorageGRID 11.6	1
Removed or deprecated features	8
Changes to the Grid Management API	8
Changes to the Tenant Management API	10
Plan and prepare for upgrade	10
Upgrade StorageGRID software	19
Troubleshoot upgrade issues	30
Increase Metadata Reserved Space setting	32
Install Red Hat Enterprise Linux or CentOS	34
Install Red Hat Enterprise Linux or CentOS: Overview	34
Plan and prepare for Red Hat or CentOS installation	35
Deploy virtual grid nodes (Red Hat or CentOS)	55
Configure the grid and complete installation (Red Hat or CentOS)	75
Automate the installation (Red Hat Enterprise Linux or CentOS)	89
Overview of the installation REST API	91
Where to go next	92
Troubleshoot installation issues	94
Example /etc/sysconfig/network-scripts	94
Install Ubuntu or Debian	97
Install Ubuntu or Debian: Overview	97
Plan and prepare for Ubuntu or Debian installation	98
Deploy virtual grid nodes (Ubuntu or Debian)	118
Configure grid and complete installation (Ubuntu or Debian)	138
Automate the installation (Ubuntu or Debian)	152
Overview of the installation REST API	154
Where to go next	155
Troubleshoot installation issues	157
Example /etc/network/interfaces	157
Install VMware	159
Install VMware: Overview	159
Plan and prepare for VMware installation	160
Deploy virtual machine grid nodes (VMware)	167
Configure the grid and complete installation (VMware)	176
Automate the installation (VMware)	190
Overview of the installation REST API	203
Where to go next	204
Troubleshoot installation issues	205

Install and upgrade software

Upgrade StorageGRID software

Upgrade StorageGRID software: Overview

Use these instructions to upgrade a StorageGRID system to a new release.

About these instructions

These instructions describe what's new in StorageGRID 11.6 and provide step-by-step instructions for upgrading all nodes in your StorageGRID system to the new release.

Before you begin

Review these topics to learn about the new features and enhancements in StorageGRID 11.6, determine whether any features have been deprecated or removed, and find out about changes to StorageGRID APIs.

- [What's new in StorageGRID 11.6](#)
- [Removed or deprecated features](#)
- [Changes to the Grid Management API](#)
- [Changes to the Tenant Management API](#)

What's new in StorageGRID 11.6

This release of StorageGRID introduces the following features.

Usability enhancements

The Grid Manager user interface was substantially redesigned to improve the user experience.

- A new sidebar replaces the pull-down menus in the old user interface.
- Several menus were reorganized to keep related options together. For example, the **CONFIGURATION** menu includes a new **Security** section for the Certificates, Key management server, Proxy settings, and Untrusted Client Networks options.
- A **Search** field in the header bar allows you to quickly navigate to Grid Manager pages.
- The summary table on the **Nodes** page provides high-level information for all sites and nodes, such as object data used and object metadata used, and includes a new search field. Alert icons are displayed next to any nodes with active alerts.
- New wizards guide you through more complex configurations, such as the workflows for admin groups, admin users, tenants, load balancer endpoints, and high availability (HA) groups.
- All UI pages were restyled with updated fonts, button styles, and table formats.



Unless there was a functional change, the screenshots in the StorageGRID 11.6 Doc Site were not updated to reflect the new Grid Manager page styling.

See the following:

- [Administer StorageGRID](#)
- [Monitor and troubleshoot](#)

Multiple VLAN interfaces

You can now create virtual LAN (VLAN) interfaces for Admin Nodes and Gateway Nodes. You can use VLAN interfaces in HA groups and load balancer endpoints to isolate and partition client traffic for security, flexibility, and performance.

- The new **Create a VLAN interface** wizard guides you through the process of entering a VLAN ID and choosing a parent interface on one or more nodes. A parent interface can be the Grid Network, the Client Network, or an additional trunk interface for the VM or bare-metal host. See [Configure VLAN interfaces](#).
- You can now add extra trunk or access interfaces to a node. If you add a trunk interface, you must configure a VLAN interface. If you add an access interface, you can add the interface directly to an HA group; you do not need to configure a VLAN interface. See the following:
 - **Linux (before installing the node):** [Installation enhancements](#)
 - **Linux (after installing the node):** [Linux: Add trunk or access interfaces to a node](#)
 - **VMware (after installing the node):** [Collect information about your deployment environment](#)

Can use Azure AD for identity federation

You can now select Azure Active Directory (Azure AD) as the identity source when configuring identity federation for the Grid Manager or the Tenant Manager. See [Use identity federation](#).

Can use Azure AD and PingFederate for SSO

You can now select Azure AD or PingFederate as the SSO type when configuring single sign-on (SSO) for your grid. You can then use sandbox mode to configure and test the Azure AD enterprise applications or PingFederate service provider (SP) connections to each StorageGRID Admin Node. See [Configure single sign-on](#).

Centralized certificate management

- The new Certificates page (**CONFIGURATION > Security > Certificates**) consolidates information about all StorageGRID security certificates to a single location. You can manage StorageGRID global, grid CA, and client certificates from the new page or view information about other certificates, such as those used for load balancer endpoints, tenants, and identity federation. See [About security certificates](#).
- As part of this change, the following global certificates were renamed:
 - The **Management Interface Server Certificate** is now the **Management interface certificate**.
 - The **Object Storage API Service Endpoints Server Certificate** (also called the Storage API Server Certificate) is now the **S3 and Swift API certificate**.
 - The **Internal CA Certificate**, **System CA Certificate**, **CA Certificate**, and **Default CA certificate** are now consistently referred to as the **Grid CA certificate**.

Other Grid Manager enhancements

- **Updates to high availability (HA) groups.** A wizard now guides you through the process of creating an HA group. See [Configure high availability groups](#).
 - In addition to selecting interfaces on the Grid Network (eth0) or Client Network (eth2), you can now select VLAN interfaces or any access interfaces you have added to the node.

- You can now specify a priority order for the interfaces. You can choose the primary interface and rank each backup interface in order.
- If any S3, Swift, administrative, or tenant clients will access the VIP addresses for the HA group from a different subnet, you can now provide the IP address for the gateway.
- **Updates to load balancer endpoints.** A new wizard guides you through the process of creating a load balancer endpoint. See [Configure load balancer endpoints](#).
- You now select the client type (S3 or Swift) when you first create the endpoint, instead of adding this detail after the endpoint is created.
- You can now use the global **StorageGRID S3 and Swift certificate** for a load balancer endpoint instead of uploading or generating a separate certificate.



This global certificate was previously used for connections to the deprecated CLB service and to Storage Nodes. If you want to use the global certificate for a load balancer endpoint, you must upload a custom certificate on the S3 and Swift API certificate page.

New Tenant Manager features

- **New Experimental S3 Console.** Available as a link from the Buckets page in Tenant Manager, the new experimental S3 Console lets S3 tenant users view and manage the objects in their buckets. See [Use Experimental S3 Console](#).



The experimental S3 Console has not been fully tested and is not intended for bulk management of objects or for use in a production environment. Tenants should only use S3 Console when performing functions for a small number of objects or when using proof-of-concept or non-production grids.

- **Can delete multiple S3 buckets.** Tenant users can now delete more than one S3 bucket at a time. Each bucket that you want to delete must be empty. See [Delete S3 bucket](#).
- **Updates to Tenant accounts permission.** Admin users who belong to a group with the Tenant accounts permission can now view existing traffic classification policies. Previously, users were required to have Root access permission to view these metrics.

New upgrade and hotfix process

- The **StorageGRID Upgrade** page was redesigned (**MAINTENANCE > System > Software update > StorageGRID upgrade**).
- After the upgrade to StorageGRID 11.6 completes, you can use the Grid Manager to upgrade to a future release and apply a hotfix for that release at the same time. The StorageGRID upgrade page will show the recommended upgrade path and will link directly to the correct download pages.
- A new **Check for software updates** check box on the AutoSupport page (**SUPPORT > Tools > AutoSupport**) lets you control this functionality. You can disable the check for available software updates if your system does not have WAN access. See [Configure AutoSupport > Disable checks for software updates](#).



For the upgrade to StorageGRID 11.6, you can optionally use a script to upgrade and apply a hotfix at the same time. See [NetApp Knowledge Base: How to run combined major upgrade and hotfix script for StorageGRID](#).

- You can now pause a SANtricity OS upgrade and skip upgrading some nodes if you need to finish the upgrade later. See the instructions for your storage appliance:

- [Upgrade SANtricity OS on storage controllers using the Grid Manager \(SG5600\)](#)
- [Upgrade SANtricity OS on storage controllers using the Grid Manager \(SG5700\)](#)
- [Upgrade SANtricity OS on storage controllers using the Grid Manager \(SG6000\)](#)

External syslog server support

- You can now configure an external syslog server if you want to save and manage audit messages and a subset of StorageGRID logs remotely (**CONFIGURATION > Monitoring > Audit and syslog server**). After an external syslog server is configured, you can save audit messages and certain log files locally, remotely, or both. By configuring the destinations of your audit information, you can reduce network traffic on your Admin Nodes. See [Configure audit messages and log destinations](#).
- Related to this functionality, new check boxes on the Logs page (**SUPPORT > Tools > Logs**) allow you to specify which types of logs you want to collect, such as certain application logs, audit logs, logs used for network debugging, and Prometheus database logs. See [Collect log files and system data](#).

S3 Select

You can now optionally allow S3 tenants to issue `SelectObjectContent` requests to individual objects. S3 Select provides an efficient way to search through large amounts of data without having to deploy a database and associated resources to enable searches. It also reduces the cost and latency of retrieving data. See [Manage S3 Select for tenant accounts](#) and [Use S3 Select](#).

Grafana charts for S3 Select operations were also added. See [Review support metrics](#).

S3 Object Lock default bucket retention period

When using S3 Object Lock, you can now specify a default retention period for the bucket. The default retention period applies to any objects added to the bucket that do not have their own retention settings. See [Use S3 Object Lock](#).

Google Cloud Platform support

You can now use the Google Cloud Platform (GCP) as an endpoint for Cloud Storage Pools and the CloudMirror platform service. See [Specify the URN for a platform services endpoint](#) and [Create a Cloud Storage Pool](#).

AWS C2S support

You can now use AWS Commercial Cloud Services (C2S) endpoints for CloudMirror replication. See [Create platform services endpoint](#).

Other S3 changes

- **GET Object and HEAD Object support for multipart objects.** Previously, StorageGRID did not support the `partNumber` request parameter in GET Object or HEAD Object requests. You can now issue GET and HEAD requests to retrieve a specific part of a multipart object. GET and HEAD Object also support the `x-amz-mp-parts-count` response element to indicate how many parts an object has.
- **Changes to "Available" consistency control.** The "Available" consistency control now behaves the same as the "read-after-new-write" consistency level, but provides eventual consistency for HEAD and GET operations. The "Available" consistency control offers higher availability for HEAD and GET operations than

“read-after-new-write” if Storage Nodes are unavailable. Differs from Amazon S3 consistency guarantees for HEAD and GET operations.

Use S3

Performance enhancements

- **Storage Nodes can support 2 billion objects.** The underlying directory structure on Storage Nodes was optimized for better scalability and performance. Storage Nodes now use additional subdirectories to store up to two billion replicated objects and to maximize performance. Node subdirectories are modified when you upgrade to StorageGRID 11.6, but existing objects are not redistributed to the new directories.
- **ILM-driven delete performance increased for high-performance appliances.** The resources and throughput used to perform ILM delete operations now adapt to the size and capability of each StorageGRID appliance node. For SG5600 appliances, the throughput is the same as for StorageGRID 11.5. For SG5700 appliances, small improvements were made to ILM delete performance. For SG6000 appliances, which have more RAM and more CPUs, ILM deletes are now processed much more quickly. The improvements are especially noticeable on all-flash SGF6024 appliances.
- **Storage volume watermarks optimized.** In previous releases, the settings of the three Storage Volume Watermarks applied to every storage volume on every Storage Node. StorageGRID can now optimize these watermarks for each storage volume, based on the size of the Storage Node and the relative capacity of the volume. See [What are storage volume watermarks](#).

Optimized watermarks are automatically applied to all new and most upgraded StorageGRID 11.6 systems. The optimized watermarks will be larger than the previous default settings.

If you use custom watermarks, the **Low read-only watermark override** alert might be triggered after you upgrade. This alert lets you know if your custom watermark settings are too small. See [Troubleshoot Low read-only watermark override alerts](#).

As part of this change, two Prometheus metrics were added:

- storagegrid_storage_volume_minimum_optimized_soft_READONLY_watermark
- storagegrid_storage_volume_maximum_optimized_soft_READONLY_watermark

- **Maximum allowed metadata space increased.** The maximum allowed metadata space for Storage Nodes was increased to 3.96 TB (from 2.64 TB) for higher-capacity nodes, which are nodes with an actual reserved space for metadata of more than 4 TB. This new value allows more object metadata to be stored on certain Storage Nodes and can increase StorageGRID metadata capacity by up to 50%.



If you have not already done so, and if your Storage Nodes have enough RAM and sufficient space on volume 0, you can [manually increase the Metadata Reserved Space setting up to 8 TB after you install or upgrade](#).

- [Manage object metadata storage > Allowed metadata space](#)
- [Increase Metadata Reserved Space setting](#)

Enhancements to maintenance procedures and support tools

- **Can change node console passwords.** You now can use the Grid Manager to change node console passwords (**CONFIGURATION > Access control > Grid passwords**). These passwords are used to log in to a node as “admin” using SSH, or to the root user on a VM/physical console connection. See [Change node console passwords](#).

- **New Object existence check wizard.** You can now verify object integrity with an easy-to-use Object existence check wizard (**MAINTENANCE > Tasks > Object existence check**), which replaces the foreground verification procedure. The new procedure takes one third of the time or less to complete and can verify multiple nodes at the same time. See [Verify object integrity](#).
- **"Estimated time to completion" chart for EC rebalance and EC repair jobs.** You can now view the estimated time to completion and the completion percentage for a current EC rebalance or EC repair job.
- **Estimated percent complete for replicated data repairs.** You can now add the `show-replicated-repair-status` option to the `repair-data` command to see an estimated percent completion for a replicated repair.



The `show-replicated-repair-status` option is available for technical preview in StorageGRID 11.6. This feature is under development, and the value returned might be incorrect or delayed. To determine if a repair is complete, continue to use **Awaiting - All**, **Repairs Attempted (XRPA)**, and **Scan Period — Estimated (XSCM)** as described in the recovery procedures.

- The results on the Diagnostics page (**SUPPORT > Tools > Diagnostics**) are now sorted by severity and then alphabetically.
- Prometheus and Grafana were updated to newer versions with modified interfaces and charts. As part of this change, the labels in some metrics were changed.
 - If you have custom queries that used the labels from `node_network_up`, you should now use the labels from `node_network_info` instead.
 - If you also used the label name `interface` from any of the `node_network` metrics, you should now use the `device` label instead.
- Previously, Prometheus metrics were stored on Admin Nodes for 31 days. Now, metrics are stored until the space reserved for Prometheus data is full, which can significantly increase how long historical metrics are available.

When the `/var/local/mysql_ibdata/` volume reaches capacity, the oldest metrics are deleted first.

Installation enhancements

- You now have the option to use Podman as a container during the installation of Red Hat Enterprise Linux. Previously, StorageGRID only supported a Docker container.
- The API schemas for StorageGRID are now included in the installation archives for the Red Hat Enterprise Linux/CentOS, Ubuntu/Debian, and VMware platforms. After extracting the archive, you can find the schemas in the `/extras/api-schemas` folder.
- The `BLOCK_DEVICE_RANGEDB` key in the node configuration file for bare-metal deployments should now contain three digits instead of two. That is, instead of `BLOCK_DEVICE_RANGEDB_nn`, you should specify `BLOCK_DEVICE_RANGEDB_nnn`.

For compatibility with existing deployments, two-digit keys are still supported for upgraded nodes.

- You can optionally add one or more instances of the new `INTERFACES_TARGET_nnnn` key to the node configuration file for bare-metal deployments. Each key provides the name and description of a physical interface on the bare-metal host, which will be displayed on the VLAN interfaces page and the HA groups page.
 - [Create node configuration files for Red Hat Enterprise Linux or CentOS deployments](#)

- [Create node configuration files for Ubuntu or Debian deployments](#)

New alerts

The following new alerts were added for StorageGRID 11.6:

- Audit logs are being added to the in-memory queue
- Cassandra table corruption
- EC rebalance failure
- EC repair failure
- EC repair stalled
- Expiration of global server certificate for S3 and Swift API
- External syslog CA certificate expiration
- External syslog client certificate expiration
- External syslog server certificate expiration
- External syslog server forwarding error
- Identity federation synchronization failure for a tenant
- Legacy CLB load balancer activity detected
- Logs are being added to the on-disk queue
- Low read-only watermark override
- Low tmp directory free space
- Object existence check failed
- Object existence check stalled
- S3 PUT Object size too large

See the [Alerts reference](#).

Changes to audit messages

- A new **BUID** field was added to the ORLM: Object Rules Met audit message. The **BUID** field shows the bucket ID, which is used for internal operations. The new field appears only if the message status is PRGD.
- A new **SGRP** field was added to the following audit messages. The **SGRP** field is present only if an object was deleted at a different site than where it was ingested.
 - IDEL: ILM Initiated Delete
 - OVWR: Object Overwrite
 - SDEL: S3 DELETE
 - WDEL: Swift DELETE

See [Review audit logs](#).

StorageGRID documentation changes

The look and feel of the StorageGRID 11.6 documentation site has been modified and now uses GitHub as the underlying platform.

NetApp appreciates feedback on content and encourages users to take advantage of the new “Request doc changes” function available on every page of the product documentation. The documentation platform also offers an embedded content contribution function for GitHub users.

Take a look and contribute to this documentation. You can edit, request a change, or simply send feedback.

Removed or deprecated features

Some features were removed or deprecated in this release. Review these items to understand whether you need to update client applications or modify your configuration before you upgrade.

Alarm system and alarm-based APIs deprecated

As of the StorageGRID 11.6 release, the legacy alarm system is deprecated. The user interface and APIs for the legacy alarm system will be removed in a future release.



If you are still using legacy alarms, plan to fully transition to the alert system after upgrading to StorageGRID 11.6. See [Manage alerts and alarms: Overview](#) to learn more about alerts.

The 11.6 release deprecates all alarm-based APIs. The following APIs are affected by this change:

- GET /grid/alarms: Fully deprecated
- GET /grid/health/topology: Fully deprecated
- GET /grid/health: The `alarm-counts` section of the response is deprecated

Future releases will not support 5 TiB maximum object size for PUT Object

In future StorageGRID releases, the maximum size for a single PUT Object operation will be 5 GiB, instead of 5 TiB. You can use multipart upload for objects that are larger than 5 GiB, up to a maximum of 5 TiB (5,497,558,138,880 bytes).

To help you transition clients to using smaller object sizes in PUT Object, the **S3 PUT Object size too large** alert will be triggered in StorageGRID 11.6 if an S3 client attempts to upload an object that exceeds 5 GiB.

NAS Bridge feature deprecated

The NAS Bridge feature previously entered limited access with the StorageGRID 11.4 release. The NAS Bridge feature remains at limited access and is deprecated as of StorageGRID 11.6.

NAS Bridge 11.4 remains the final release and will continue to be compatible with StorageGRID 11.6. Review the [NetApp Interoperability Matrix Tool](#) for continued compatibility between NAS Bridge 11.4 and StorageGRID versions.

Review the NetApp Support site for the [support schedule for NAS Bridge](#).

Changes to the Grid Management API

StorageGRID 11.6 uses version 3 of the Grid Management API. Version 3 deprecates version 2; however, version 1 and version 2 are still supported.



You can continue to use version 1 and version 2 of the management API with StorageGRID 11.6; however, support for these versions of the API will be removed in a future release of StorageGRID. After upgrading to StorageGRID 11.6, the deprecated v1 and v2 APIs can be deactivated using the `PUT /grid/config/management` API.

To learn more, go to [Use the Grid Management API](#).

Can access Swagger docs for private API operations

You can now access the Swagger docs for the private API from the Grid Manager. To see the available operations, select the Grid Manager help icon and select **API Documentation**. Then, select **Go to private API documentation** from the StorageGRID Management API page.

StorageGRID private APIs are subject to change without notice. StorageGRID private endpoints also ignore the API version of the request.

Alarm-based APIs deprecated

The 11.6 release deprecates all alarm-based APIs. The following APIs are affected by this change:

- `GET /grid/alarms`: Fully deprecated
- `GET /grid/health/topology`: Fully deprecated
- `GET /grid/health`: The `alarm-counts` section of the response is deprecated

Can import S3 access keys

You can now use the Grid Management API to import S3 access keys for tenant users. For example, you can migrate access keys from another S3 provider to StorageGRID or use this feature to keep user credentials the same between grids.



When this feature is enabled, any Grid Manager user with the `Change tenant root password` permission has full access to tenant data. Disable this feature immediately after use to protect tenant data.

New account operations

Three new `grid/account` API operations have been added:

- `POST /grid/account-enable-s3-key-import`: This request enables the Import S3 Credentials feature. You must have the `Root` access permission to enable this feature.
- `POST /grid/accounts/{id}/users/{user_id}/s3-access-keys`: This request imports S3 credentials for a given user in a tenant account. You must have the `Root` access or `Change tenant root password` permission, and you must know the user ID and the tenant account ID.
- `POST /grid/account-disable-s3-key-import`: This request disables the Import S3 Credentials feature. You must have the `Root` access permission to disable this feature.

PATCH method deprecated

The PATCH method has now been deprecated and will be removed in a future release. As required, perform a PUT operation to replace a resource instead of using a PATCH operation to modify the resource.

Additions to `grid/logs/collect` endpoint

Four boolean values have been added to the `grid/logs/collect` endpoint:

- `applicationLogs`: Application-specific logs that technical support uses most frequently for troubleshooting. The logs collected are a subset of the available application logs. The default is `true`.
- `auditLogs`: Logs containing the audit messages generated during normal system operation. The default is `true`.
- `networkTrace`: Logs used for network debugging. The default is `false`.
- `prometheusDatabase`: Time series metrics from the services on all nodes. The default is `false`.

New `node-details /grid/service-ids` endpoint

The new `/grid/service-ids` endpoint provides a mapping of node UUIDs to the associated node names, service IDs, and service types.

Can retrieve grid node console passwords

You can now use `POST /grid/node-console-passwords` to retrieve the list of grid nodes and their associated console passwords.

Changes to the Tenant Management API

StorageGRID 11.6 uses version 3 of the Tenant Management API. Version 3 deprecates version 2; however, version 1 and version 2 are still supported.

 You can continue to use version 1 and version 2 of the management API with StorageGRID 11.6; however, support for these versions of the API will be removed in a future release of StorageGRID. After upgrading to StorageGRID 11.6, the deprecated v1 and v2 APIs can be deactivated using the `PUT /grid/config/management` API.

To learn more, go to [Understand the Tenant Management API](#).

PATCH method deprecated

The PATCH method has now been deprecated and will be removed in a future release. As required, perform a PUT operation to replace a resource instead of using a PATCH operation to modify the resource.

Plan and prepare for upgrade

Estimate the time to complete an upgrade

When planning an upgrade to StorageGRID 11.6, you must consider when to upgrade, based on how long the upgrade might take. You must also be aware of which operations you can and cannot perform during each stage of the upgrade.

About this task

The time required to complete a StorageGRID upgrade depends on a variety of factors such as client load and hardware performance.

The table summarizes the main upgrade tasks and lists the approximate time required for each task. The steps after the table provide instructions you can use to estimate the upgrade time for your system.

Upgrade task	Description	Approximate time required	During this task
Start Upgrade Service	Upgrade prechecks are run, the software file is distributed, and the upgrade service is started.	3 minutes per grid node, unless validation errors are reported	As required, you can run the upgrade prechecks manually before the scheduled upgrade maintenance window.
Upgrade Grid Nodes (primary Admin Node)	The primary Admin Node is stopped, upgraded, and restarted.	30 minutes to 1 hour, with SG100 and SG1000 appliance nodes requiring the most time.	You cannot access the primary Admin Node. Connection errors are reported, which you can ignore.
Upgrade Grid Nodes (all other nodes)	The software on all other grid nodes is upgraded, in the order in which you approve the nodes. Every node in your system will be brought down one at a time for several minutes each.	<p>15 minutes to 1 hour per node, with appliance nodes requiring the most time</p> <p>Note: For appliance nodes, the StorageGRID Appliance Installer is automatically updated to the latest release.</p>	<ul style="list-style-type: none"> • Do not change the grid configuration. • Do not change the audit level configuration. • Do not update the ILM configuration. • You are prevented from performing other maintenance procedures, such as hotfix, decommission, or expansion. <p>Note: If you need to perform a recovery, contact technical support.</p>
Enable Features	The new features for the new version are enabled.	Less than 5 minutes	<ul style="list-style-type: none"> • Do not change the grid configuration. • Do not change the audit level configuration. • Do not update the ILM configuration. • You cannot perform another maintenance procedure.

Upgrade task	Description	Approximate time required	During this task
Upgrade Database	The upgrade process checks each node to verify that the Cassandra database does not need to be updated.	10 seconds per node or a few minutes for the entire grid	<p>The upgrade from StorageGRID 11.5 to 11.6 does not require a Cassandra database upgrade; however, the Cassandra service will be stopped and restarted on each Storage Node.</p> <p>For future StorageGRID feature releases, the Cassandra database update step might take several days to complete.</p>
Final Upgrade Steps	Temporary files are removed and the upgrade to the new release completes.	5 minutes	When the Final Upgrade Steps task completes, you can perform all maintenance procedures.

Steps

1. Estimate the time required to upgrade all grid nodes.
 - a. Multiply the number of nodes in your StorageGRID system by 1 hour/node.

As a general rule, appliance nodes take longer to upgrade than software-based nodes.
 - b. Add 1 hour to this time to account for the time required to download the .upgrade file, run precheck validations, and complete the final upgrade steps.
2. If you have Linux nodes, add 15 minutes for each node to account for the time required to download and install the RPM or DEB package.
3. Calculate the total estimated time for the upgrade by adding the results of steps 1 and 2.

Example: Estimated time to upgrade to StorageGRID 11.6

Suppose your system has 14 grid nodes, of which 8 are Linux nodes.

1. Multiply 14 by 1 hour/node.
2. Add 1 hour to account for the download, precheck, and final steps.

The estimated time to upgrade all nodes is 15 hours.

3. Multiply 8 by 15 minutes/node to account for the time to install the RPM or DEB package on the Linux nodes.

The estimated time for this step is 2 hours.

4. Add the values together.

You should allow up to 17 hours to complete the upgrade of your system to StorageGRID 11.6.0.

How your system is affected during the upgrade

You must understand how your StorageGRID system will be affected during the upgrade.

StorageGRID upgrades are non-disruptive

The StorageGRID system can ingest and retrieve data from client applications throughout the upgrade process. Grid nodes are brought down one at a time during the upgrade, so there is not a time when all grid nodes are unavailable.

To allow for continued availability, you must ensure that objects are stored redundantly using the appropriate ILM policies. You must also ensure that all external S3 or Swift clients are configured to send requests to one of the following:

- A StorageGRID endpoint configured as a high availability (HA) group
- A high availability third-party load balancer
- Multiple Gateway Nodes for each client
- Multiple Storage Nodes for each client

Appliance firmware is upgraded

During the StorageGRID 11.6 upgrade:

- All StorageGRID appliance nodes are automatically upgraded to StorageGRID Appliance Installer firmware version 3.6.
- SG6060 and SGF6024 appliances are automatically upgraded to BIOS firmware version 3B07.EX and BMC firmware version BMC 3.93.07.
- SG100 and SG1000 appliances are automatically upgraded to BIOS firmware version 3B12.EC and BMC firmware version 4.67.07.

Alerts might be triggered

Alerts might be triggered when services start and stop and when the StorageGRID system is operating as a mixed-version environment (some grid nodes running an earlier version, while others have been upgraded to a later version). Other alerts might be triggered after the upgrade completes.

For example, you might see the **Unable to communicate with node** alert when services are stopped, or you might see the **Cassandra communication error** alert when some nodes have been upgraded to StorageGRID 11.6 but other nodes are still running StorageGRID 11.5. In general, these alerts will clear when the upgrade completes.

The **ILM placement unachievable** alert might be triggered when Storage Nodes are stopped during the upgrade to StorageGRID 11.6. This alert might persist for 1 day after the upgrade completes.

If you use custom values for the storage volume watermarks, the **Low read-only watermark override** alert might be triggered after the upgrade is complete. See [Troubleshoot Low read-only watermark override alerts](#) for details.

After the upgrade completes, you can review any upgrade-related alerts by selecting **Recently resolved alerts** or **Current alerts** from the Grid Manager Dashboard.

Many SNMP notifications are generated

Be aware that a large number of SNMP notifications might be generated when grid nodes are stopped and restarted during the upgrade. To avoid excessive notifications, unselect the **Enable SNMP Agent Notifications** check box (**CONFIGURATION > Monitoring > SNMP agent**) to disable SNMP notifications before you start the upgrade. Then, re-enable notifications after the upgrade is complete.

Configuration changes are restricted



The list of restricted configuration changes can change from release to release. When upgrading to another StorageGRID release, refer to the list in the appropriate upgrade instructions.

Until the **Enable New Feature** task completes:

- Do not make any grid configuration changes.
- Do not change the audit level configuration or configure an external syslog server.
- Do not enable or disable any new features.
- Do not update the ILM configuration. Otherwise, you might experience inconsistent and unexpected ILM behavior.
- Do not apply a hotfix or recover a grid node.
- You cannot manage HA groups, VLAN interfaces, or load balancer endpoints while you are upgrading to StorageGRID 11.6.

Until the **Final Upgrade Steps** task completes:

- Do not perform an expansion procedure.
- Do not perform a decommission procedure.

You cannot view bucket details or manage buckets from the Tenant Manager

During the upgrade to StorageGRID 11.6 (that is, while the system is operating as a mixed-version environment), you cannot view bucket details or manage buckets using the Tenant Manager. One of the following errors appears on the Buckets page in Tenant Manager:

- “You cannot use this API while you are upgrading to 11.6.”
- “You cannot view bucket versioning details in the Tenant Manager while you are upgrading to 11.6.”

This error will resolve after the upgrade to 11.6 is complete.

Workaround

While the 11.6 upgrade is in progress, use the following tools to view bucket details or manage buckets, instead of using the Tenant Manager:

- To perform standard S3 operations on a bucket, use either the S3 REST API or the Tenant Management API.
- To perform StorageGRID custom operations on a bucket (for example, viewing and modifying the bucket consistency level, enabling or disabling last access time updates, or configuring search integration), use the Tenant Management API.

See [Understand Tenant Management API](#) and [Use S3](#) for instructions.

Impact of an upgrade on groups and user accounts

You must understand the impact of the StorageGRID upgrade, so that you can update groups and user accounts appropriately after the upgrade is complete.

Changes to group permissions and options

After upgrading to StorageGRID 11.6, optionally select the following updated or new permissions and options (**CONFIGURATION > Access control > Admin groups**).

Permission or option	Description
Tenant accounts	In addition to allowing users to create, edit, and remove tenant accounts, this permission now allows admin users to view existing traffic classification policies (CONFIGURATION > Network > Traffic classification).

See [Manage admin groups](#).

Verify the installed version of StorageGRID

Before starting the upgrade, you must verify that the previous version of StorageGRID is currently installed with the latest available hotfix applied.

About this task

Before you upgrade to StorageGRID 11.6, your grid must have StorageGRID 11.5 installed. If you are currently using a previous version of StorageGRID, you must install all previous upgrade files along with their latest hotfixes (strongly recommended) until your grid's current version is StorageGRID 11.5.x.y.

One possible upgrade path is shown in the [example](#).



NetApp strongly recommends that you apply the latest hotfix for each StorageGRID version before upgrading to the next version and that you also apply the latest hotfix for each new version you install. In some cases, you must apply a hotfix to avoid the risk of data loss. See [NetApp Downloads: StorageGRID](#) and the release notes for each hotfix to learn more.

Note that you can run a script to update from 11.3.0.13+ to 11.4.0.y in one step and from 11.4.0.7+ to 11.5.0.y in one step. See [NetApp Knowledge Base: How to run combined major upgrade and hotfix script for StorageGRID](#).

Steps

1. Sign in to the Grid Manager using a [supported web browser](#).
2. From the top of the Grid Manager, select **Help > About**.
3. Verify that **Version** is 11.5.x.y.

In the StorageGRID 11.5.x.y version number:

- The **major release** has an x value of 0 (11.5.0).
- A **hotfix**, if one has been applied, has a y value (for example, 11.5.0.1).

4. If **Version** is not 11.5.x.y, go to [NetApp Downloads: StorageGRID](#) to download the files for each previous

release, including the latest hotfix for each release.

5. Obtain the upgrade instructions for each release you downloaded. Then, perform the software upgrade procedure for that release, and apply the latest hotfix for that release (strongly recommended).

See the [StorageGRID hotfix procedure](#).

Example: Upgrade to StorageGRID 11.5 from version 11.3.0.8

The following example shows the steps to upgrade from StorageGRID version 11.3.0.8 to version 11.5 in preparation for a StorageGRID 11.6 upgrade.



Optionally, you can run a script to combine steps 2 and 3 (update from 11.3.0.13+ to 11.4.0.y) and to combine steps 4 and 5 (update from 11.4.0.7+ to 11.5.0.y). See [NetApp Knowledge Base: How to run combined major upgrade and hotfix script for StorageGRID](#).

Download and install software in the following sequence to prepare your system for upgrade:

1. Apply the latest StorageGRID 11.3.0.y hotfix.
2. Upgrade to the StorageGRID 11.4.0 major release.
3. Apply the latest StorageGRID 11.4.0.y hotfix.
4. Upgrade to the StorageGRID 11.5.0 major release.
5. Apply the latest StorageGRID 11.5.0.y hotfix.

Obtain the required materials for a software upgrade

Before you begin the software upgrade, you must obtain all required materials so you can complete the upgrade successfully.

Item	Notes
StorageGRID upgrade files	Download the StorageGRID upgrade files to your service laptop.
Service laptop	The service laptop must have: <ul style="list-style-type: none">• Network port• SSH client (for example, PuTTY)
Supported web browser	Browser support typically changes for each StorageGRID release. Make sure your browser is compatible with the new StorageGRID version.
Recovery Package (.zip) file	Download the Recovery Package before upgrading, and save the file in a safe location. The Recovery Package file allows you to restore the system if a failure occurs.
Passwords.txt file	This file is included in the SAID package, which is part of the Recovery Package .zip file. You must obtain the latest version of the Recovery Package.

Item	Notes
Provisioning passphrase	The passphrase is created and documented when the StorageGRID system is first installed. The provisioning passphrase is not listed in the <code>Passwords.txt</code> file.
Related documentation	<ul style="list-style-type: none"> • Release notes for StorageGRID 11.6. Be sure to read these carefully before starting the upgrade. • Instructions for administering StorageGRID. • If you are upgrading a Linux deployment, the StorageGRID installation instructions for your Linux platform: <ul style="list-style-type: none"> ◦ Install Red Hat Enterprise Linux or CentOS ◦ Install Ubuntu or Debian • Other StorageGRID documentation, as required.

Download the StorageGRID upgrade files

You must download one or more files, depending on where your nodes are installed.

- **All platforms:** `.upgrade` file

If any nodes are deployed on Linux hosts, you must also download an RPM or DEB archive, which you will install before you start the upgrade:

- **Red Hat Enterprise Linux or CentOS:** An additional RPM file (`.zip` or `.tgz`)
- **Ubuntu or Debian:** An additional DEB file (`.zip` or `.tgz`)

Steps

1. Go to [NetApp Downloads: StorageGRID](#).
2. Select the button for downloading the latest release, or select another version from the drop-down menu and select **Go**.

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

3. Sign in with the username and password for your NetApp account.
4. If a Caution/MustRead statement appears, read it and select the check box.

This statement appears if there is a required hotfix for the release.

5. Read the End User License Agreement, select the check box, and then select **Accept & Continue**.

The downloads page for the version you selected appears. The page contains three columns:

- Install StorageGRID
- Upgrade StorageGRID
- Support files for StorageGRID Appliances

6. In the **Upgrade StorageGRID** column, select and download the `.upgrade` archive.

Every platform requires the `.upgrade` archive.

7. If any nodes are deployed on Linux hosts, also download the RPM or DEB archive in either `.tgz` or `.zip` format. Select the `.zip` file if you are running Windows on the service laptop.

- Red Hat Enterprise Linux or CentOS

`StorageGRID-Webscale-version-RPM-uniqueID.zip`

`StorageGRID-Webscale-version-RPM-uniqueID.tgz`

- Ubuntu or Debian

`StorageGRID-Webscale-version-DEB-uniqueID.zip`

`StorageGRID-Webscale-version-DEB-uniqueID.tgz`



No additional files are required for the SG100 or SG1000.

Download the Recovery Package

The Recovery Package file allows you to restore the StorageGRID system if a failure occurs. Download the current Recovery Package file before making grid topology changes to the StorageGRID system or before upgrading software. Then, download a new copy of the Recovery Package after making grid topology changes or after upgrading software.

What you'll need

- You must be signed in to the Grid Manager using a [supported web browser](#).
- You must have the provisioning passphrase.
- You must have specific access permissions.

Steps

1. Select **Maintenance > System > Recovery Package**.
2. Enter the provisioning passphrase, and select **Start Download**.

The download starts immediately.

3. When the download completes:
 - a. Open the `.zip` file.
 - b. Confirm it includes a `gpt-backup` directory and an inner `.zip` file.
 - c. Extract the inner `.zip` file.
 - d. Confirm you can open the `Passwords.txt` file.
4. Copy the downloaded Recovery Package file (`.zip`) to two safe, secure, and separate locations.



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

Check the system's condition

Before upgrading a StorageGRID system, you must verify the system is ready to accommodate the upgrade. You must ensure that the system is running normally and that

all grid nodes are operational.

Steps

1. Sign in to the Grid Manager using a [supported web browser](#).
2. Check for and resolve any active alerts.

For information about specific alerts, see the [Alerts reference](#).

3. Confirm that no conflicting grid tasks are active or pending.
 - a. Select **SUPPORT > Tools > Grid topology**.
 - b. Select **site > primary Admin Node > CMN > Grid Tasks > Configuration**.

Information lifecycle management evaluation (ILME) tasks are the only grid tasks that can run concurrently with the software upgrade.

- c. If any other grid tasks are active or pending, wait for them to finish or release their lock.



Contact technical support if a task does not finish or release its lock.

4. Refer to [Internal grid node communications](#) and [External communications](#) to ensure that all required ports for StorageGRID 11.6 are opened before you upgrade.

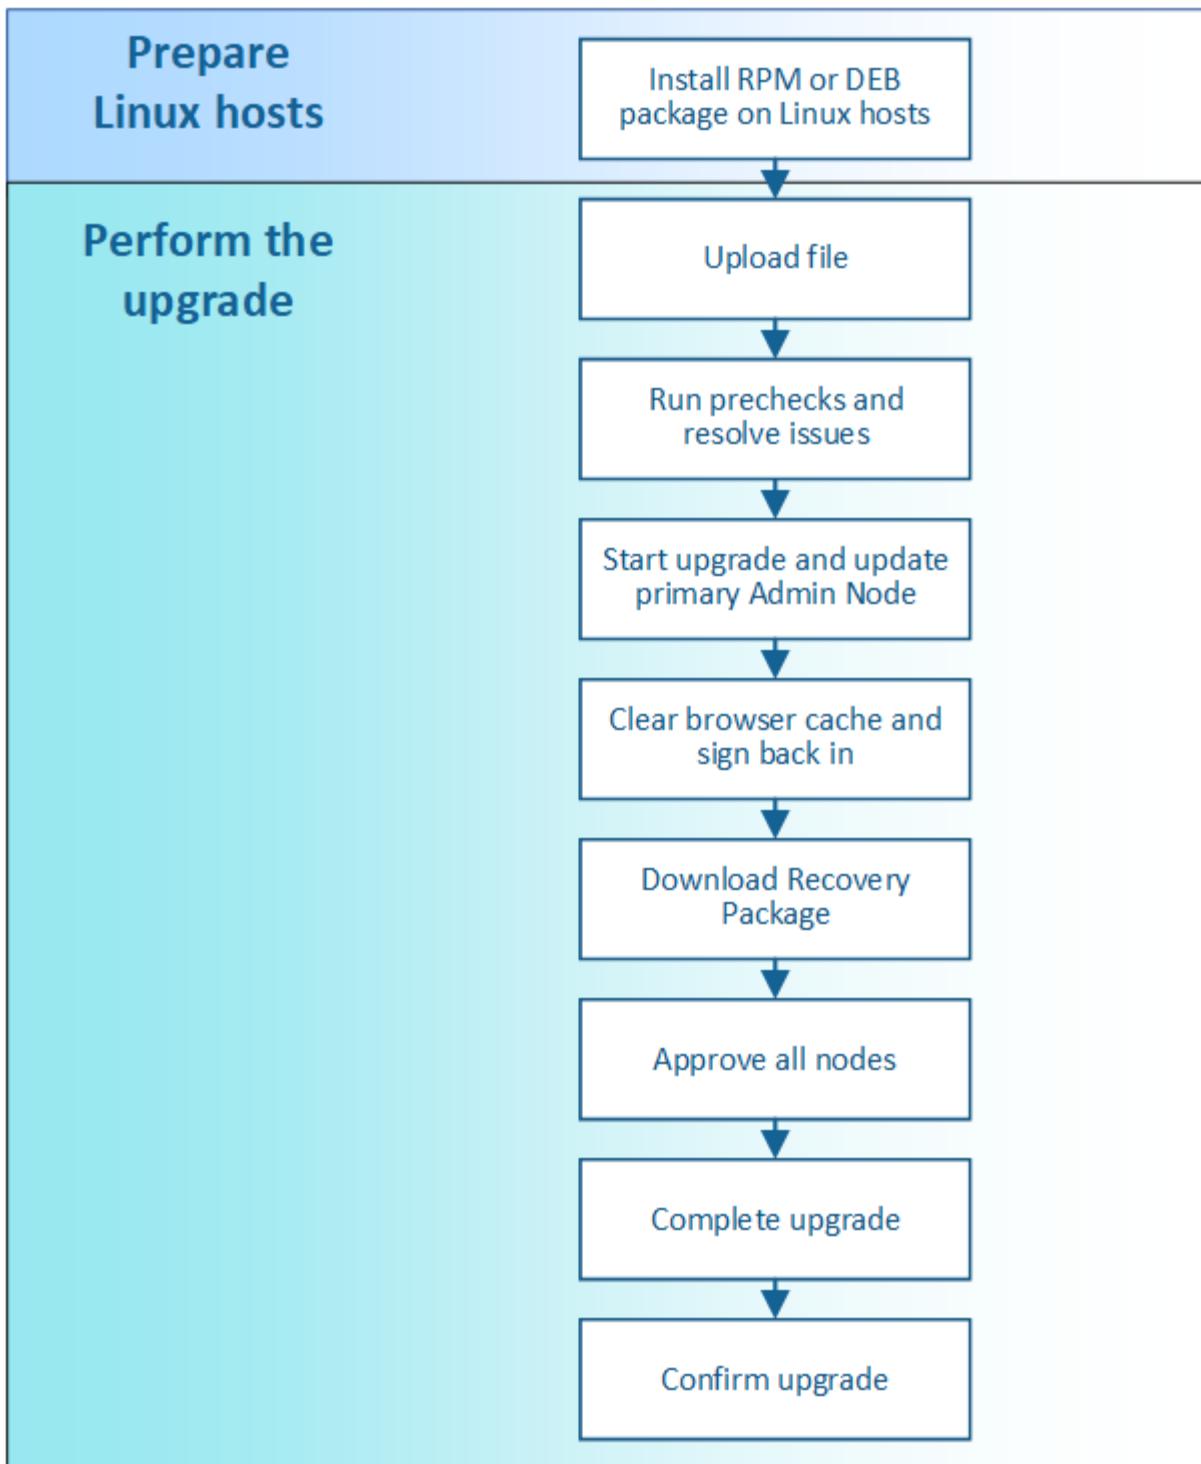


If you have opened any custom firewall ports, you are notified during the upgrade precheck. You must contact technical support before proceeding with the upgrade.

Upgrade StorageGRID software

Upgrade workflow

Before starting the upgrade, review the general workflow. The StorageGRID Upgrade page guides you through each upgrade step.



1. If any StorageGRID nodes are deployed on Linux hosts, [install the RPM or DEB package on each host](#) before you start the upgrade.
2. From the primary Admin Node, access the StorageGRID Upgrade page and upload the upgrade file.
3. Optionally run upgrade prechecks to detect and resolve any issues before you start the actual upgrade.
4. Start the upgrade, which runs prechecks and upgrades the primary Admin Node automatically. You cannot access the Grid Manager while the primary Admin Node is being upgraded. Audit logs will also be unavailable. This upgrade can take up to 30 minutes.
5. After the primary Admin Node has been upgraded, clear your web browser's cache, sign back in, and return to the StorageGRID Upgrade page.

6. Download a new Recovery Package.
7. Approve the grid nodes. You can approve individual grid nodes, groups of grid nodes, or all grid nodes.



Do not approve the upgrade for a grid node unless you are sure that node is ready to be stopped and rebooted.

8. Resume operations. When all grid nodes have been upgraded, new features are enabled and you can resume operations. You must wait to perform a decommission or expansion procedure until the background **Upgrade Database** task and the **Final Upgrade Steps** task have completed.
9. When the upgrade is complete, confirm the software version and apply any hotfixes.

Related information

[Estimate the time to complete an upgrade](#)

Linux: Install the RPM or DEB package on all hosts

If any StorageGRID nodes are deployed on Linux hosts, you must install an additional RPM or DEB package on each of these hosts before you start the upgrade.

What you'll need

You must have downloaded one of the following .tgz or .zip files from the NetApp Downloads page for StorageGRID.



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

Linux platform	Additional file (choose one)
Red Hat Enterprise Linux or CentOS	<ul style="list-style-type: none">• StorageGRID-Webscale-version-RPM-uniqueID.zip• StorageGRID-Webscale-version-RPM-uniqueID.tgz
Ubuntu or Debian	<ul style="list-style-type: none">• StorageGRID-Webscale-version-DEB-uniqueID.zip• StorageGRID-Webscale-version-DEB-uniqueID.tgz

Steps

1. Extract the RPM or DEB packages from the installation file.
2. Install the RPM or DEB packages on all Linux hosts.

See the steps for installing StorageGRID host services in the installation instructions for your Linux platform.

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

The new packages are installed as additional packages. Do not remove the existing packages.

Perform the upgrade

When you are ready to perform the upgrade, you select the .upgrade archive and enter the provisioning passphrase. As an option, you can run the upgrade prechecks before performing the actual upgrade.

What you'll need

You have reviewed all of the considerations and completed all of the planning and preparation steps.

Upload the upgrade file

1. Sign in to the Grid Manager using a [supported web browser](#).
2. Select **Maintenance > System > Software Update**.

The Software Update page appears.

3. Select **StorageGRID Upgrade**.
4. On the StorageGRID Upgrade page, select the .upgrade archive.
 - a. Select **Browse**.
 - b. Locate and select the file: NetApp_StorageGRID_11.6.0_Software_uniqueID.upgrade
 - c. Select **Open**.

The file is uploaded and validated. When the validation process is done, a green checkmark appears next to the upgrade file name.

5. Enter the provisioning passphrase in the text box.

The **Run Prechecks** and **Start Upgrade** buttons become enabled.

The screenshot shows the StorageGRID Upgrade page. At the top, there is a note about confirming active alerts and grid nodes. Below that, the 'Upgrade file' section shows a 'Browse' button with the selected file 'NetApp_StorageGRID_11.6.0_Software_20211206.1924.c35b8bf.upgrade'. The 'Upgrade Version' is listed as 'StorageGRID® 11.6.0'. In the 'Passphrase' section, a 'Provisioning Passphrase' field contains several asterisks. At the bottom, there are two buttons: 'Run Prechecks' and 'Start Upgrade'.

Run prechecks

Optionally, you can validate the condition of your system before you start the actual upgrade. Selecting **Run Prechecks** allows you to detect and resolve issues before starting the upgrade. The same prechecks are

performed when you start the upgrade. Precheck failures will stop the upgrade process and some might require technical support involvement to resolve.

1. Select **Run Prechecks**.
2. Wait for the prechecks to complete.
3. Follow the instructions to resolve any precheck errors that are reported.



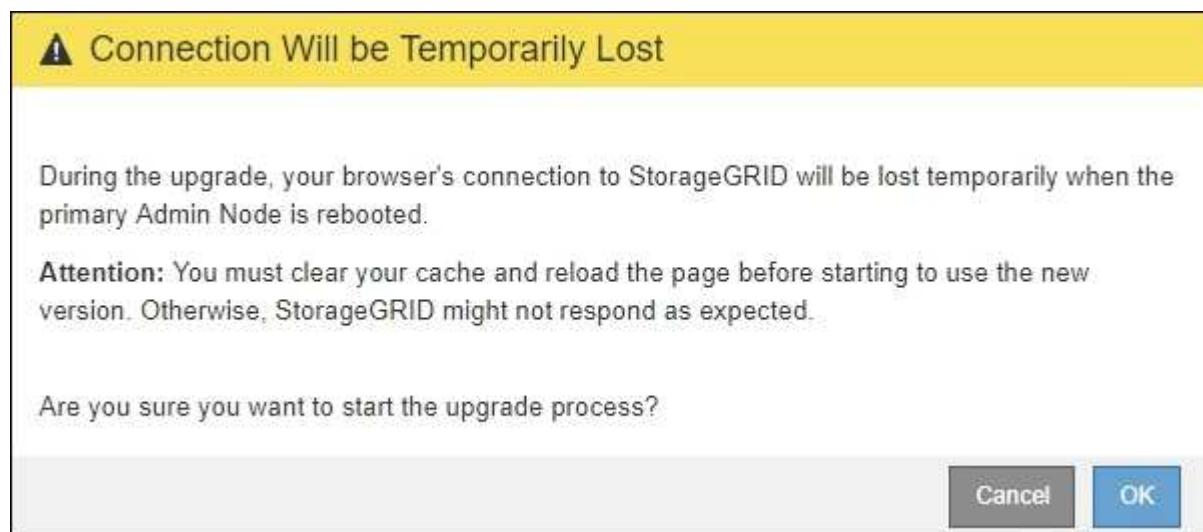
If you have opened any custom firewall ports, you are notified during the precheck validation. You must contact technical support before proceeding with the upgrade.

Start upgrade and update primary Admin Node

When the upgrade starts, upgrade prechecks are performed, and the primary Admin Node is upgraded, which includes stopping services, upgrading the software, and restarting services. You cannot access the Grid Manager while the primary Admin Node is being upgraded. Audit logs will also be unavailable. This upgrade can take up to 30 minutes.

1. When you are ready to perform the upgrade, select **Start Upgrade**.

A warning appears to remind you that your browser's connection will be lost when the primary Admin Node is rebooted.



2. Select **OK** to acknowledge the warning and start the upgrade process.
3. Wait for the upgrade prechecks to be performed and for the primary Admin Node to be upgraded.



If any precheck errors are reported, resolve them and select **Start Upgrade** again.

While the primary Admin Node is being upgraded, multiple **503: Service Unavailable** and **Problem connecting to the server** messages appear, which you can ignore.

! Error

503: Service Unavailable

Service Unavailable

The StorageGRID API service is not responding. Please try again later. If the problem persists, contact Technical Support.

4 additional copies of this message are not shown.

OK

! Error

Problem connecting to the server

Unable to communicate with the server. Please reload the page and try again. Contact technical support if the problem persists.

2 additional copies of this message are not shown.

OK

4. When you see the **400: Bad Request** message, go to the next step. The Admin Node upgrade is complete.

! Error

400: Bad Request

Clear your web browser's cache and reload the page to continue the upgrade.

2 additional copies of this message are not shown.

OK

Clear browser cache and sign back in

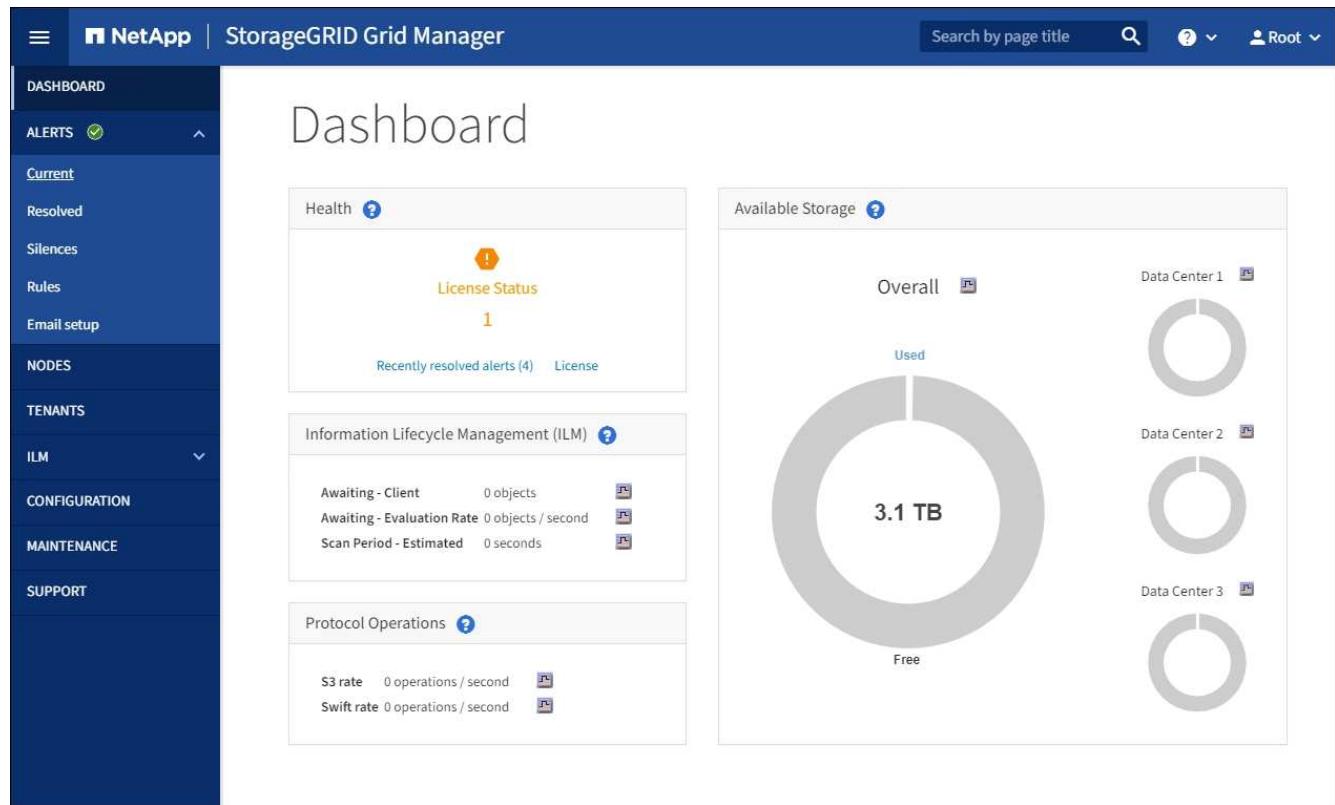
1. After the primary Admin Node has been upgraded, clear your web browser's cache and sign back in.

For instructions, see the documentation for your web browser.



You must clear the web browser's cache to remove outdated resources used by the previous version of the software.

The redesigned Grid Manager interface appears, which indicates that the primary Admin Node has been upgraded.



2. From the sidebar, select **MAINTENANCE** to open the Maintenance menu.
3. In the **System** section, select **Software update**.
4. In the **StorageGRID Upgrade** section, select **Upgrade**.
5. Review the Upgrade Progress section on the StorageGRID Upgrade page, which provides information about each major upgrade task.
 - a. **Start Upgrade Service** is the first upgrade task. During this task, the software file is distributed to the grid nodes, and the upgrade service is started.
 - b. When the **Start Upgrade Service** task is complete, the **Upgrade Grid Nodes** task starts.
 - c. While the **Upgrade Grid Nodes** task is in progress, the Grid Node Status table appears and shows the upgrade stage for each grid node in your system.

Download Recovery Package and upgrade all grid nodes

1. After the grid nodes appear in the Grid Node Status table, but before approving any grid nodes, [download a new copy of the Recovery Package](#).



You must download a new copy of the Recovery Package file after you upgrade the software version on the primary Admin Node. The Recovery Package file allows you to restore the system if a failure occurs.

2. Review the information in the Grid Node Status table. Grid nodes are arranged in sections by type: Admin Nodes, API Gateway Nodes, Storage Nodes, and Archive Nodes.

Upgrade Progress

Start Upgrade Service	Completed
Upgrade Grid Nodes	In Progress 

Grid Node Status

You must approve all grid nodes to complete an upgrade, but you can update grid nodes in any order.

During the upgrade of a node, the services on that node are stopped. Later, the node is rebooted. Do not click Approve for a node unless you are sure the node is ready to be stopped and rebooted.

When you are ready to add grid nodes to the upgrade queue, click one or more Approve buttons to add individual nodes to the queue, click the Approve All button at the top of the nodes table to add all nodes of the same type, or click the top-level Approve All button to add all nodes in the grid.

If necessary, you can remove nodes from the upgrade queue before node services are stopped by clicking Remove or Remove All.

Approve All

Remove All

▼ Admin Nodes

▼ API Gateway Nodes

Approve All **Remove All**

▲ Storage Nodes

Approve All **Remove All**

Search 

Site	Name	Progress	Stage	Error	Action
ALT-ADM1-177	ALT-S1-175	<div style="width: 20%;"></div>	Waiting for you to approve	 	Approve
ALT-ADM1-177	ALT-S2-174	<div style="width: 20%;"></div>	Waiting for you to approve	 	Approve
ALT-ADM1-177	ALT-S3-173	<div style="width: 20%;"></div>	Waiting for you to approve	 	Approve

▼ Archive Nodes

A grid node can be in one of these stages when this page first appears:

- Done (primary Admin Node only)
- Preparing upgrade
- Software download queued
- Downloading
- Waiting for you to approve

3. Approve the grid nodes you are ready to add to the upgrade queue.



When the upgrade starts on a grid node, the services on that node are stopped. Later, the grid node is rebooted. To avoid service interruptions for client applications that are communicating with the node, do not approve the upgrade for a node unless you are sure that node is ready to be stopped and rebooted. As required, schedule a maintenance window or notify customers.

You must upgrade all grid nodes in your StorageGRID system, but you can customize the upgrade sequence. You can approve individual grid nodes, groups of grid nodes, or all grid nodes.

If the order in which nodes are upgraded is important, approve nodes or groups of nodes one at a time and wait until the upgrade is complete on each node before approving the next node or group of nodes.

- Select one or more **Approve** buttons to add one or more individual nodes to the upgrade queue. If you approve more than one node of the same type, the nodes will be upgraded one at a time.
- Select the **Approve All** button within each section to add all nodes of the same type to the upgrade queue.
- Select the top-level **Approve All** button to add all nodes in the grid to the upgrade queue.
- Select **Remove** or **Remove All** to remove a node or all nodes from the upgrade queue. You cannot remove a node when its Stage reaches **Stopping services**. The **Remove** button is hidden.

Site	Name	Progress	Stage	Error	Action
ALT-ADM1-177	ALT-S1-175	<div style="width: 20%; background-color: #0072bc; height: 10px; display: inline-block;"></div>	Queued		Remove
ALT-ADM1-177	ALT-S2-174	<div style="width: 40%; background-color: #0072bc; height: 10px; display: inline-block;"></div>	Stopping services		
ALT-ADM1-177	ALT-S3-173	<div style="width: 20%; background-color: #0072bc; height: 10px; display: inline-block;"></div>	Queued		Remove

4. Wait for each node to proceed through the upgrade stages, which include Queued, Stopping services, Stopping container, Cleaning up Docker images, Upgrading base OS packages, Rebooting, Performing steps after reboot, Starting services, and Done.



When an appliance node reaches the Upgrading base OS packages stage, the StorageGRID Appliance Installer software on the appliance is updated. This automated process ensures that the StorageGRID Appliance Installer version remains in sync with the StorageGRID software version.

Complete upgrade

When all grid nodes have completed the upgrade stages, the **Upgrade Grid Nodes** task is shown as Completed. The remaining upgrade tasks are performed automatically and in the background.

1. As soon as the **Enable Features** task is complete (which occurs quickly), optionally start using the new features in the upgraded StorageGRID version.
2. During the **Upgrade Database** task, the upgrade process checks each node to verify that the Cassandra database does not need to be updated.



The upgrade from StorageGRID 11.5 to 11.6 does not require a Cassandra database upgrade; however, the Cassandra service will be stopped and restarted on each Storage Node. For future StorageGRID feature releases, the Cassandra database update step might take several days to complete.

3. When the **Upgrade Database** task has completed, wait a few minutes for the **Final Upgrade Steps** task to complete.

When the Final Upgrade Steps task has completed, the upgrade is done.

Confirm upgrade

1. Confirm that the upgrade completed successfully.
 - a. From the top of the Grid Manager, select the help icon and select **About**.
 - b. Confirm that the displayed version is what you would expect.
 - c. Select **MAINTENANCE > System > Software update**.
 - d. In the **StorageGRID upgrade** section, select **Upgrade**.
 - e. Confirm that the green banner shows that the software upgrade was completed on the date and time you expect.

StorageGRID Upgrade

1 Select files — 2 Run prechecks — 3 Upgrade primary Admin Node — 4 Upgrade other nodes

Before updating software, confirm that your StorageGRID system has no active alerts and that all nodes are connected to the grid.

 StorageGRID upgrade completed at 2021-12-06 16:29:20 MST

Current version : 11.6.0

Update Path

No upgrade or hotfix available

Upload files

Upload the upgrade file for the new version. Then, if a hotfix is available for the new version, upload the hotfix. The hotfix will be automatically applied as part of the upgrade.

Upgrade file :

Hotfix for new version (if available) :

2. From the StorageGRID Upgrade page, determine if any hotfixes are available for the current StorageGRID version.



If no Update Path is shown, your browser might not be able to reach the NetApp Support Site. Or, the **Check for software updates** check box on the AutoSupport page (**SUPPORT > Tools > AutoSupport**) might be disabled.

3. If a hotfix is available, download the file. Then, use the [StorageGRID hotfix procedure](#) to apply the hotfix.

4. Verify that grid operations have returned to normal:

- Check that the services are operating normally and that there are no unexpected alerts.
- Confirm that client connections to the StorageGRID system are operating as expected.

Troubleshoot upgrade issues

Upgrade does not complete

If the upgrade does not complete successfully, you might be able to resolve the issue yourself. If you cannot resolve an issue, you should gather the required information before contacting technical support.

The following sections describe how to recover from situations where the upgrade has partially failed. Contact technical support if you cannot resolve an upgrade issue.

Upgrade precheck errors

To detect and resolve issues, you can manually run the upgrade prechecks before starting the actual upgrade. Most precheck errors provide information about how to resolve the issue. If you need help, contact technical support.

Provisioning failures

If the automatic provisioning process fails, contact technical support.

Grid node crashes or fails to start

If a grid node crashes during the upgrade process or fails to start successfully after the upgrade finishes, contact technical support to investigate and to correct any underlying issues.

Ingest or data retrieval is interrupted

If data ingest or retrieval is unexpectedly interrupted when you are not upgrading a grid node, contact technical support.

Database upgrade errors

If the database upgrade fails with an error, retry the upgrade. If it fails again, contact technical support.

Related information

[Checking the system's condition before upgrading software](#)

Troubleshoot user interface issues

You might see issues with the Grid Manager or the Tenant Manager after upgrading to a new version of StorageGRID software.

Web interface does not respond as expected

The Grid Manager or the Tenant Manager might not respond as expected after StorageGRID software is upgraded.

If you experience issues with the web interface:

- Make sure you are using a [supported web browser](#).



Browser support typically changes for each StorageGRID release.

- Clear your web browser cache.

Clearing the cache removes outdated resources used by the previous version of StorageGRID software, and permits the user interface to operate correctly again. For instructions, see the documentation for your web browser.

“Docker image availability check” error messages

When attempting to start the upgrade process, you might receive an error message that states “The following issues were identified by the Docker image availability check validation suite.” All issues must be resolved before you can complete the upgrade.

Contact technical support if you are unsure of the changes required to resolve the identified issues.

Message	Cause	Solution
Unable to determine upgrade version. Upgrade version info file {file_path} did not match the expected format.	The upgrade package is corrupt.	Re-upload the upgrade package, and try again. If the problem persists, contact technical support.
Upgrade version info file {file_path} was not found. Unable to determine upgrade version.	The upgrade package is corrupt.	Re-upload the upgrade package, and try again. If the problem persists, contact technical support.
Unable to determine currently installed release version on {node_name}.	A critical file on the node is corrupt.	Contact technical support.
Connection error while attempting to list versions on {node_name}	The node is offline or the connection was interrupted.	Check to make sure that all nodes are online and reachable from the primary Admin Node, and try again.
The host for node {node_name} does not have StorageGRID {upgrade_version} image loaded. Images and services must be installed on the host before the upgrade can proceed.	<p>The RPM or DEB packages for the upgrade have not been installed on the host where the node is running, or the images are still in the process of being imported.</p> <p>Note: This error only applies to nodes that are running as containers on Linux.</p>	<p>Check to make sure that the RPM or DEB packages have been installed on all Linux hosts where nodes are running. Make sure the version is correct for both the service and the images file. Wait a few minutes, and try again.</p> <p>See Linux: Install RPM or DEB package on all hosts.</p>
Error while checking node {node_name}	An unexpected error occurred.	Wait a few minutes, and try again.
Uncaught error while running prechecks. {error_string}	An unexpected error occurred.	Wait a few minutes, and try again.

Increase Metadata Reserved Space setting

After you upgrade to StorageGRID 11.6, you might be able to increase the Metadata Reserved Space system setting if your Storage Nodes meet specific requirements for RAM and available space.

What you'll need

- You must be signed in to the Grid Manager using a [supported web browser](#).
- You must have the Root Access permission or the Grid Topology Page Configuration and Other Grid Configuration permissions.
- You have completed the StorageGRID 11.6 upgrade.

About this task

You might be able to manually increase the system-wide Metadata Reserved Space setting up to 8 TB after upgrading to StorageGRID 11.6. Reserving additional metadata space after the 11.6 upgrade will simplify future hardware and software upgrades.

You can only increase the value of the system-wide Metadata Reserved Space setting if both of these statements are true:

- The Storage Nodes at any site in your system each have 128 GB or more RAM.
- The Storage Nodes at any site in your system each have sufficient available space on storage volume 0.

Be aware that if you increase this setting, you will simultaneously reduce the space available for object storage on storage volume 0 of all Storage Nodes. For this reason, you might prefer to set the Metadata Reserved Space to a value smaller than 8 TB, based on your expected object metadata requirements.



In general, it is better to use a higher value instead of a lower value. If the Metadata Reserved Space setting is too large, you can decrease it later. In contrast, if you increase the value later, the system might need to move object data to free up space.

For a detailed explanation of how the Metadata Reserved Space setting affects the allowed space for object metadata storage on a particular Storage Node, go to [Manage object metadata storage](#).

Steps

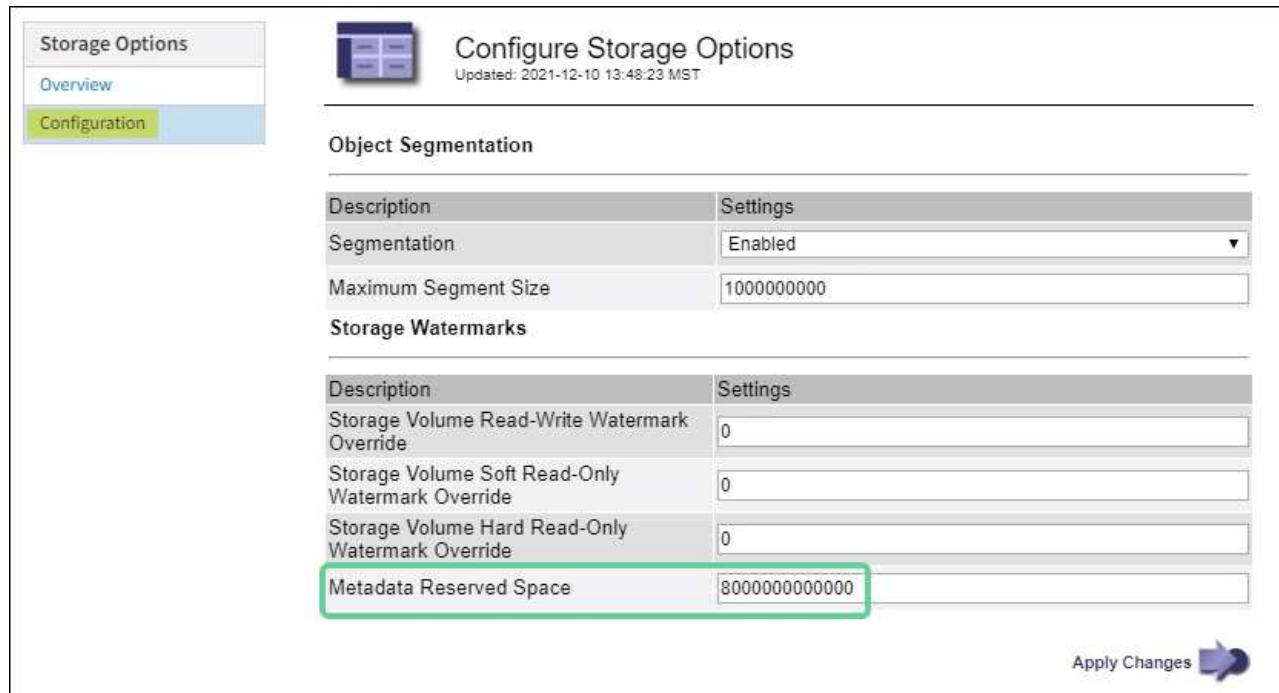
1. Sign in to the Grid Manager using a [supported web browser](#).
2. Determine the current Metadata Reserved Space setting.
 - a. Select **CONFIGURATION > System > Storage options**.
 - b. In the Storage Watermarks section, note the value of **Metadata Reserved Space**.
3. Ensure you have enough available space on storage volume 0 of each Storage Node to increase this value.
 - a. Select **NODES**.
 - b. Select the first Storage Node in the grid.
 - c. Select the Storage tab.
 - d. In the Volumes section, locate the **/var/local/rangedb/0** entry.
 - e. Confirm that the Available value is equal to or greater than difference between the new value you want to use and the current Metadata Reserved Space value.
- For example, if the Metadata Reserved Space setting is currently 4 TB and you want to increase it to 6 TB, the Available value must be 2 TB or greater.
- f. Repeat these steps for all Storage Nodes.
 - If one or more Storage Nodes do not have enough available space, the Metadata Reserved Space value cannot be increased. Do not continue with this procedure.
 - If each Storage Node has enough available space on volume 0, go to the next step.
4. Ensure you have at least 128 GB of RAM on each Storage Node.
 - a. Select **NODES**.
 - b. Select the first Storage Node in the grid.

- c. Select the **Hardware** tab.
- d. Hover your cursor over the Memory Usage chart. Ensure that **Total Memory** is at least 128 GB.
- e. Repeat these steps for all Storage Nodes.
 - If one or more Storage Nodes do not have enough available total memory, the **Metadata Reserved Space** value cannot be increased. Do not continue with this procedure.
 - If each Storage Node has at least 128 GB of total memory, go to the next step.

5. Update the Metadata Reserved Space setting.

- a. Select **CONFIGURATION > System > Storage options**.
- b. Select the Configuration tab.
- c. In the Storage Watermarks section, select **Metadata Reserved Space**.
- d. Enter the new value.

For example, to enter 8 TB, which is the maximum supported value, enter **80000000000000** (8, followed by 12 zeros)



The screenshot shows the 'Configure Storage Options' page. The 'Storage Options' sidebar on the left has 'Configuration' selected. The main area is titled 'Configure Storage Options' with a timestamp 'Updated: 2021-12-10 13:48:23 MST'. The 'Object Segmentation' section is above the 'Storage Watermarks' section. In the 'Storage Watermarks' section, the 'Metadata Reserved Space' row is highlighted with a green border. The 'Settings' column for 'Metadata Reserved Space' contains the value '80000000000000'.

Description	Settings
Segmentation	Enabled
Maximum Segment Size	1000000000
Storage Watermarks	
Description	Settings
Storage Volume Read-Write Watermark Override	0
Storage Volume Soft Read-Only Watermark Override	0
Storage Volume Hard Read-Only Watermark Override	0
Metadata Reserved Space	80000000000000

- e. Select **Apply Changes**.

Install Red Hat Enterprise Linux or CentOS

Install Red Hat Enterprise Linux or CentOS: Overview

Installing a StorageGRID system in a Red Hat Enterprise Linux (RHEL) or CentOS Linux environment includes three primary steps.

1. **Preparation:** During planning and preparation, you perform the following tasks:
 - Learn about the hardware and storage requirements for StorageGRID.
 - Learn about the specifics of [StorageGRID networking](#) so you can configure your network appropriately.

- Identify and prepare the physical or virtual servers you plan to use to host your StorageGRID grid nodes.
- On the servers you have prepared:
 - Install Linux
 - Configure the host network
 - Configure host storage
 - Install container engine
 - Install the StorageGRID host services

2. **Deployment:** Deploy grid nodes using the appropriate user interface. When you deploy grid nodes, they are created as part of the StorageGRID system and connected to one or more networks.

- a. Use the Linux command line and node configuration files to deploy software-based grid nodes on the hosts you prepared in step 1.
- b. Use the StorageGRID Appliance Installer to deploy StorageGRID appliance nodes.



Hardware-specific installation and integration instructions are not included in the StorageGRID installation procedure. To learn how to install StorageGRID appliances, see the installation and maintenance instructions for your appliance.

3. **Configuration:** When all nodes have been deployed, use the Grid Manager to configure the grid and complete the installation.

These instructions recommend a standard approach for deploying and configuring a StorageGRID system. See also the information about the following alternative approaches:

- Use a standard orchestration framework such as Ansible, Puppet, or Chef to install RHEL or CentOS, configure networking and storage, install the container engine and the StorageGRID host service, and deploy virtual grid nodes.
- Automate the deployment and configuration of the StorageGRID system using a Python configuration script (provided in the installation archive).
- Automate the deployment and configuration of appliance grid nodes with a Python configuration script (available from the installation archive or from the StorageGRID Appliance Installer).
- If you are an advanced developer of StorageGRID deployments, use the installation REST APIs to automate the installation of StorageGRID grid nodes.

Plan and prepare for Red Hat or CentOS installation

Before you install (Red Hat or CentOS)

Before deploying grid nodes and configuring the StorageGRID grid, you must be familiar with the steps and requirements for completing the procedure.

The StorageGRID deployment and configuration procedures assume that you are familiar with the architecture and operation of the StorageGRID system.

You can deploy a single site or multiple sites at one time; however, all sites must meet the minimum requirement of having at least three Storage Nodes.

Before starting a StorageGRID installation, you must:

- Understand StorageGRID's compute requirements, including the minimum CPU and RAM requirements for each node.
- Understand how StorageGRID supports multiple networks for traffic separation, security, and administrative convenience, and have a plan for which networks you intend to attach to each StorageGRID node.

See the StorageGRID networking guidelines.

- Understand the storage and performance requirements of each type of grid node.
- Identify a set of servers (physical, virtual, or both) that, in aggregate, provide sufficient resources to support the number and type of StorageGRID nodes you plan to deploy.
- Understand the requirements for node migration, if you want to perform scheduled maintenance on physical hosts without any service interruption.
- Gather all networking information in advance. Unless you are using DHCP, gather the IP addresses to assign to each grid node, and the IP addresses of the domain name system (DNS) and network time protocol (NTP) servers that will be used.
- Install, connect, and configure all required hardware, including any StorageGRID appliances, to specifications.



Hardware-specific installation and integration instructions are not included in the StorageGRID installation procedure. To learn how to install StorageGRID appliances, see the installation and maintenance instructions for your appliance.

- Decide which of the available deployment and configuration tools you want to use.

Related information

[Networking guidelines](#)

[SG100 and SG1000 services appliances](#)

[SG6000 storage appliances](#)

[SG5700 storage appliances](#)

[SG5600 storage appliances](#)

Required materials

Before you install StorageGRID, you must gather and prepare required materials.

Item	Notes
NetApp StorageGRID license	You must have a valid, digitally signed NetApp license. Note: A non-production license, which can be used for testing and proof of concept grids, is included in the StorageGRID installation archive.
StorageGRID installation archive	You must download the StorageGRID installation archive and extract the files .

Item	Notes
Service laptop	<p>The StorageGRID system is installed through a service laptop.</p> <p>The service laptop must have:</p> <ul style="list-style-type: none"> • Network port • SSH client (for example, PuTTY) • Supported web browser
StorageGRID documentation	<ul style="list-style-type: none"> • Release notes • Instructions for administering StorageGRID

Related information

[NetApp Interoperability Matrix Tool](#)

Download and extract the StorageGRID installation files

You must download the StorageGRID installation archive and extract the required files.

Steps

1. Go to the [NetApp Downloads page for StorageGRID](#).
2. Select the button for downloading the latest release, or select another version from the drop-down menu and select **Go**.
3. Sign in with the username and password for your NetApp account.
4. If a Caution/MustRead statement appears, read it and select the check box.



You must apply any required hotfixes after you install the StorageGRID release. For more information, see the [hotfix procedure in the recovery and maintenance instructions](#).

5. Read the End User License Agreement, select the check box, and then select **Accept & Continue**.
6. In the **Install StorageGRID** column, select the .tgz or .zip file for Red Hat Enterprise Linux or CentOS.



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

7. Save and extract the archive file.
8. Choose the files you need from the following list.

The files you need depend on your planned grid topology and how you will deploy your StorageGRID system.



The paths listed in the table are relative to the top-level directory installed by the extracted installation archive

Path and file name	Description
./rpms/README	A text file that describes all of the files contained in the StorageGRID download file.
./rpms/NLF000000.txt	A free license that does not provide any support entitlement for the product.
./rpms/StorageGRID-Webscale-Images-version-SHA.rpm	RPM package for installing the StorageGRID node images on your RHEL or CentOS hosts.
./rpms/StorageGRID-Webscale-Service-version-SHA.rpm	RPM package for installing the StorageGRID host service on your RHEL or CentOS hosts.
Deployment scripting tool	Description
./rpms/configure-storagegrid.py	A Python script used to automate the configuration of a StorageGRID system.
./rpms/configure-sga.py	A Python script used to automate the configuration of StorageGRID appliances.
./rpms/configure-storagegrid.sample.json	An example configuration file for use with the configure-storagegrid.py script.
./rpms/storagegrid-ssoauth.py	An example Python script that you can use to sign in to the Grid Management API when single sign-on is enabled.
./rpms/configure-storagegrid.blank.json	A blank configuration file for use with the configure-storagegrid.py script.
./rpms/extras/ansible	Example Ansible role and playbook for configuring RHEL or CentOS hosts for StorageGRID container deployment. You can customize the role or playbook as necessary.
./rpms/extras/api-schemas	<p>API schemas for StorageGRID.</p> <p>Note: Before you perform an upgrade, you can use these schemas to confirm that any code you have written to use StorageGRID management APIs will be compatible with the new StorageGRID release if you do not have a non-production StorageGRID environment for upgrade compatibility testing.</p>

CPU and RAM requirements

Before installing StorageGRID software, verify and configure the hardware so that it is

ready to support the StorageGRID system.

For information about supported servers, see the Interoperability Matrix.

Each StorageGRID node requires the following minimum resources:

- CPU cores: 8 per node
- RAM: At least 24 GB per node, and 2 to 16 GB less than the total system RAM, depending on the total RAM available and the amount of non-StorageGRID software running on the system

Ensure that the number of StorageGRID nodes you plan to run on each physical or virtual host does not exceed the number of CPU cores or the physical RAM available. If the hosts are not dedicated to running StorageGRID (not recommended), be sure to consider the resource requirements of the other applications.

 Monitor your CPU and memory usage regularly to ensure that these resources continue to accommodate your workload. For example, doubling the RAM and CPU allocation for virtual Storage Nodes would provide similar resources to those provided for StorageGRID appliance nodes. Additionally, if the amount of metadata per node exceeds 500 GB, consider increasing the RAM per node to 48 GB or more. For information about managing object metadata storage, increasing the Metadata Reserved Space setting, and monitoring CPU and memory usage, see the instructions for administering, monitoring, and upgrading StorageGRID.

If hyperthreading is enabled on the underlying physical hosts, you can provide 8 virtual cores (4 physical cores) per node. If hyperthreading is not enabled on the underlying physical hosts, you must provide 8 physical cores per node.

If you are using virtual machines as hosts and have control over the size and number of VMs, you should use a single VM for each StorageGRID node and size the VM accordingly.

For production deployments, you should not run multiple Storage Nodes on the same physical storage hardware or virtual host. Each Storage Node in a single StorageGRID deployment should be in its own isolated failure domain. You can maximize the durability and availability of object data if you ensure that a single hardware failure can only impact a single Storage Node.

See also the information about storage requirements.

Related information

[NetApp Interoperability Matrix Tool](#)

[Storage and performance requirements](#)

[Administer StorageGRID](#)

[Monitor and troubleshoot](#)

[Upgrade software](#)

Storage and performance requirements

You must understand the storage requirements for StorageGRID nodes, so you can provide enough space to support the initial configuration and future storage expansion.

StorageGRID nodes require three logical categories of storage:

- **Container pool** — Performance-tier (10K SAS or SSD) storage for the node containers, which will be assigned to the container engine storage driver when you install and configure the container engine on the hosts that will support your StorageGRID nodes.
- **System data** — Performance-tier (10K SAS or SSD) storage for per-node persistent storage of system data and transaction logs, which the StorageGRID host services will consume and map into individual nodes.
- **Object data** — Performance-tier (10K SAS or SSD) storage and capacity-tier (NL-SAS/SATA) bulk storage for the persistent storage of object data and object metadata.

You must use RAID-backed block devices for all storage categories. Non-redundant disks, SSDs, or JBODs are not supported. You can use shared or local RAID storage for any of the storage categories; however, if you want to use StorageGRID's node migration capability, you must store both system data and object data on shared storage.

Performance requirements

The performance of the volumes used for the container pool, system data, and object metadata significantly impacts the overall performance of the system. You should use performance-tier (10K SAS or SSD) storage for these volumes to ensure adequate disk performance in terms of latency, input/output operations per second (IOPS), and throughput. You can use capacity-tier (NL-SAS/SATA) storage for the persistent storage of object data.

The volumes used for the container pool, system data, and object data must have write-back caching enabled. The cache must be on a protected or persistent media.

Requirements for hosts that use NetApp ONTAP storage

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

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

Number of hosts required

Each StorageGRID site requires a minimum of three Storage Nodes.

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

Other types of nodes, such as Admin Nodes or Gateway Nodes, can be deployed on the same hosts, or they can be deployed on their own dedicated hosts as required.

Number of storage volumes for each host

The following table shows the number of storage volumes (LUNs) required for each host and the minimum size required for each LUN, based on which nodes will be deployed on that host.

The maximum tested LUN size is 39 TB.



These numbers are for each host, not for the entire grid.

LUN purpose	Storage category	Number of LUNs	Minimum size/LUN
Container engine storage pool	Container pool	1	Total number of nodes × 100 GB
/var/local volume	System data	1 for each node on this host	90 GB
Storage Node	Object data	3 for each Storage Node on this host Note: A software-based Storage Node can have 1 to 16 storage volumes; at least 3 storage volumes are recommended.	12 TB (4 TB/LUN) See Storage requirements for Storage Nodes for more information.
Admin Node audit logs	System data	1 for each Admin Node on this host	200 GB
Admin Node tables	System data	1 for each Admin Node on this host	200 GB



Depending on the audit level configured, the size of user inputs such as S3 object key name, and how much audit log data you need to preserve, you might need to increase the size of the audit log LUN on each Admin Node. As a general rule, a grid generates approximately 1 KB of audit data per S3 operation, which would mean that a 200 GB LUN would support 70 million operations per day or 800 operations per second for two to three days.

Minimum storage space for a host

The following table shows the minimum storage space required for each type of node. You can use this table to determine the minimum amount of storage you must provide to the host in each storage category, based on which nodes will be deployed on that host.



Disk snapshots cannot be used to restore grid nodes. Instead, refer to the recovery and maintenance procedures for each type of node.

Type of node	Container pool	System data	Object data
Storage Node	100 GB	90 GB	4,000 GB
Admin Node	100 GB	490 GB (3 LUNs)	<i>not applicable</i>
Gateway Node	100 GB	90 GB	<i>not applicable</i>

Type of node	Container pool	System data	Object data
Archive Node	100 GB	90 GB	<i>not applicable</i>

Example: Calculating the storage requirements for a host

Suppose you plan to deploy three nodes on the same host: one Storage Node, one Admin Node, and one Gateway Node. You should provide a minimum of nine storage volumes to the host. You will need a minimum of 300 GB of performance-tier storage for the node containers, 670 GB of performance-tier storage for system data and transaction logs, and 12 TB of capacity-tier storage for object data.

Type of node	LUN purpose	Number of LUNs	LUN size
Storage Node	Container engine storage pool	1	300 GB (100 GB/node)
Storage Node	/var/local volume	1	90 GB
Storage Node	Object data	3	12 TB (4 TB/LUN)
Admin Node	/var/local volume	1	90 GB
Admin Node	Admin Node audit logs	1	200 GB
Admin Node	Admin Node tables	1	200 GB
Gateway Node	/var/local volume	1	90 GB
Total		9	Container pool: 300 GB System data: 670 GB Object data: 12,000 GB

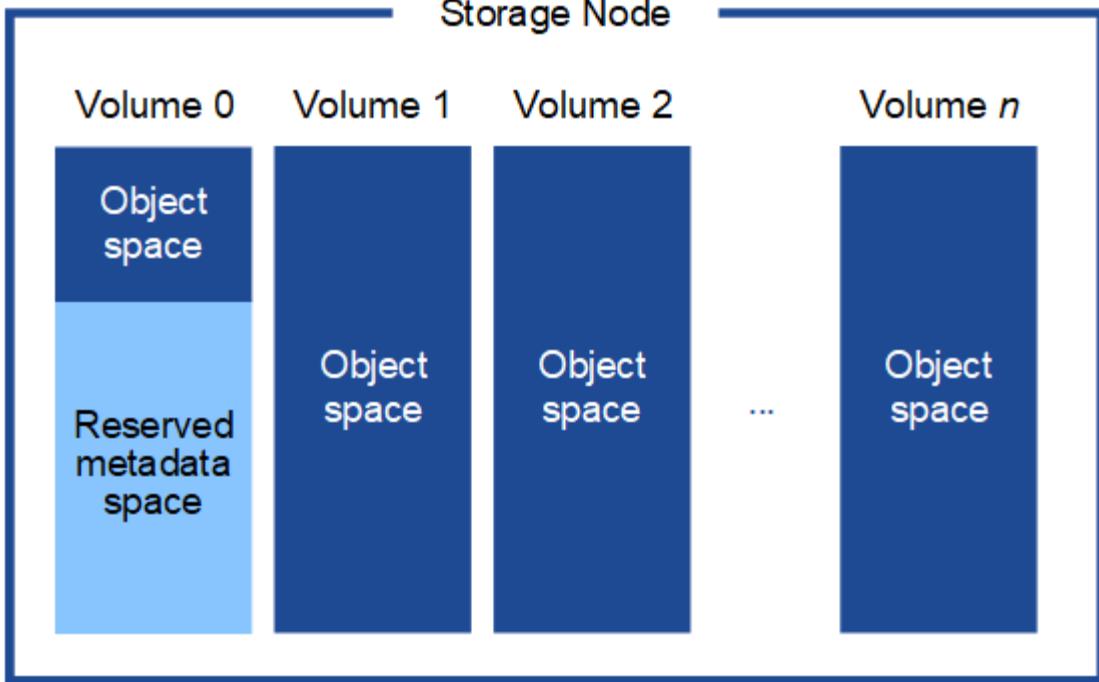
Storage requirements for Storage Nodes

A software-based Storage Node can have 1 to 16 storage volumes—3 or more storage volumes are recommended. Each storage volume should be 4 TB or larger.



An appliance Storage Node can have up to 48 storage volumes.

As shown in the figure, StorageGRID reserves space for object metadata on storage volume 0 of each Storage Node. Any remaining space on storage volume 0 and any other storage volumes in the Storage Node are used exclusively for object data.



To provide redundancy and to protect object metadata from loss, StorageGRID stores three copies of the metadata for all objects in the system at each site. The three copies of object metadata are evenly distributed across all Storage Nodes at each site.

When you assign space to volume 0 of a new Storage Node, you must ensure there is adequate space for that node's portion of all object metadata.

- At a minimum, you must assign at least 4 TB to volume 0.



If you use only one storage volume for a Storage Node and you assign 4 TB or less to the volume, the Storage Node might enter the Storage Read-Only state on startup and store object metadata only.

- If you are installing a new StorageGRID 11.6 system and each Storage Node has 128 GB or more of RAM, you should assign 8 TB or more to volume 0. Using a larger value for volume 0 can increase the space allowed for metadata on each Storage Node.
- When configuring different Storage Nodes for a site, use the same setting for volume 0 if possible. If a site contains Storage Nodes of different sizes, the Storage Node with the smallest volume 0 will determine the metadata capacity of that site.

For details, go to [Manage object metadata storage](#).

Related information

[Node container migration requirements](#)

[Recover and maintain](#)

Node container migration requirements

The node migration feature allows you to manually move a node from one host to another. Typically, both hosts are in the same physical data center.

Node migration allows you to perform physical host maintenance without disrupting grid operations. You simply move all StorageGRID nodes, one at a time, to another host before taking the physical host offline. Migrating nodes requires only a short downtime for each node and should not affect operation or availability of grid services.

If you want to use the StorageGRID node migration feature, your deployment must meet additional requirements:

- Consistent network interface names across hosts in a single physical data center
- Shared storage for StorageGRID metadata and object repository volumes that is accessible by all hosts in a single physical data center. For example, you might use NetApp E-Series storage arrays.

If you are using virtual hosts and the underlying hypervisor layer supports VM migration, you might want to use this capability instead of StorageGRID's node migration feature. In this case, you can ignore these additional requirements.

Before performing migration or hypervisor maintenance, shut down the nodes gracefully. See the instructions for [shutting down a grid node](#).

VMware Live Migration not supported

OpenStack Live Migration and VMware live vMotion cause the virtual machine clock time to jump and are not supported for grid nodes of any type. Though rare, incorrect clock times can result in loss of data or configuration updates.

Cold migration is supported. In cold migration, you shut down the StorageGRID nodes before migrating them between hosts. See the instructions for [shutting down a grid node](#).

Consistent network interface names

In order to move a node from one host to another, the StorageGRID host service needs to have some confidence that the external network connectivity the node has at its current location can be duplicated at the new location. It gets this confidence through the use of consistent network interface names in the hosts.

Suppose, for example, that StorageGRID NodeA running on Host1 has been configured with the following interface mappings:

eth0 → bond0.1001

eth1 → bond0.1002

eth2 → bond0.1003

The lefthand side of the arrows corresponds to the traditional interfaces as viewed from within a StorageGRID container (that is, the Grid, Admin, and Client Network interfaces, respectively). The righthand side of the arrows corresponds to the actual host interfaces providing these networks, which are three VLAN interfaces subordinate to the same physical interface bond.

Now, suppose you want to migrate NodeA to Host2. If Host2 also has interfaces named bond0.1001, bond0.1002, and bond0.1003, the system will allow the move, assuming that the like-named interfaces will provide the same connectivity on Host2 as they do on Host1. If Host2 does not have interfaces with the same names, the move will not be allowed.

There are many ways to achieve consistent network interface naming across multiple hosts; see [Configuring the host network](#) for some examples.

Shared storage

In order to achieve rapid, low-overhead node migrations, the StorageGRID node migration feature does not physically move node data. Instead, node migration is performed as a pair of export and import operations, as follows:

1. During the “node export” operation, a small amount of persistent state data is extracted from the node container running on HostA and cached on that node’s system data volume. Then, the node container on HostA is deinstantiated.
2. During the “node import” operation, the node container on HostB that uses the same network interface and block storage mappings that were in effect on HostA is instantiated. Then, the cached persistent state data is inserted into the new instance.

Given this mode of operation, all of the node’s system data and object storage volumes must be accessible from both HostA and HostB for the migration to be allowed, and to work. In addition, they must have been mapped into the node using names that are guaranteed to refer to the same LUNs on HostA and HostB.

The following example shows one solution for block device mapping for a StorageGRID Storage Node, where DM multipathing is in use on the hosts, and the alias field has been used in `/etc/multipath.conf` to provide consistent, friendly block device names available on all hosts.

`/var/local` → `/dev/mapper/sgws-sn1-var-local`

`rangedb0` → `/dev/mapper/sgws-sn1-rangedb0`

`rangedb1` → `/dev/mapper/sgws-sn1-rangedb1`

`rangedb2` → `/dev/mapper/sgws-sn1-rangedb2`

`rangedb3` → `/dev/mapper/sgws-sn1-rangedb3`

Deployment tools

You might benefit from automating all or part of the StorageGRID installation.

Automating the deployment might be useful in any of the following cases:

- You already use a standard orchestration framework, such as Ansible, Puppet, or Chef, to deploy and configure physical or virtual hosts.
- You intend to deploy multiple StorageGRID instances.
- You are deploying a large, complex StorageGRID instance.

The StorageGRID host service is installed by a package and driven by configuration files that can be created interactively during a manual installation, or prepared ahead of time (or programmatically) to enable automated installation using standard orchestration frameworks. StorageGRID provides optional Python scripts for

automating the configuration of StorageGRID appliances, and the whole StorageGRID system (the “grid”). You can use these scripts directly, or you can inspect them to learn how to use the [StorageGRID installation REST API](#) in grid deployment and configuration tools you develop yourself.

If you are interested in automating all or part of your StorageGRID deployment, review [Automate the installation](#) before beginning the installation process.

Prepare the hosts (Red Hat or CentOS)

Install Linux

You must install Linux on all grid hosts. Use the [NetApp Interoperability Matrix Tool](#) to get a list of supported versions.

Steps

1. Install Linux on all physical or virtual grid hosts according to the distributor’s instructions or your standard procedure.



If you are using the standard Linux installer, NetApp recommends selecting the “compute node” software configuration, if available, or “minimal install” base environment. Do not install any graphical desktop environments.

2. Ensure that all hosts have access to package repositories, including the Extras channel.

You might need these additional packages later in this installation procedure.

3. If swap is enabled:
 - a. Run the following command: `$ sudo swapoff --all`
 - b. Remove all swap entries from `/etc/fstab` to persist the settings.



Failing to disable swap entirely can severely lower performance.

Configure the host network (Red Hat Enterprise Linux or CentOS)

After completing the Linux installation on your hosts, you might need to perform some additional configuration to prepare a set of network interfaces on each host that are suitable for mapping into the StorageGRID nodes you will deploy later.

What you’ll need

- You have reviewed the [StorageGRID networking guidelines](#).
- You have reviewed the information about [node container migration requirements](#).
- If you are using virtual hosts, you have read the [considerations and recommendations for MAC address cloning](#) before configuring the host network.



If you are using VMs as hosts, you should select VMXNET 3 as the virtual network adapter. The VMware E1000 network adapter has caused connectivity issues with StorageGRID containers deployed on certain distributions of Linux.

About this task

Grid nodes must be able to access the Grid Network and, optionally, the Admin and Client Networks. You provide this access by creating mappings that associate the host's physical interface to the virtual interfaces for each grid node. When creating host interfaces, use friendly names to facilitate deployment across all hosts, and to enable migration.

The same interface can be shared between the host and one or more nodes. For example, you might use the same interface for host access and node Admin Network access, to facilitate host and node maintenance. Although the same interface can be shared between the host and individual nodes, all must have different IP addresses. IP addresses cannot be shared between nodes or between the host and any node.

You can use the same host network interface to provide the Grid Network interface for all StorageGRID nodes on the host; you can use a different host network interface for each node; or you can do something in between. However, you would not typically provide the same host network interface as both the Grid and Admin Network interfaces for a single node, or as the Grid Network interface for one node and the Client Network interface for another.

You can complete this task in many ways. For example, if your hosts are virtual machines and you are deploying one or two StorageGRID nodes for each host, you can simply create the correct number of network interfaces in the hypervisor, and use a 1-to-1 mapping. If you are deploying multiple nodes on bare metal hosts for production use, you can leverage the Linux networking stack's support for VLAN and LACP for fault tolerance and bandwidth sharing. The following sections provide detailed approaches for both of these examples. You do not need to use either of these examples; you can use any approach that meets your needs.

 Do not use bond or bridge devices directly as the container network interface. Doing so could prevent node start-up caused by a kernel issue with the use of MACVLAN with bond and bridge devices in the container namespace. Instead, use a non-bond device, such as a VLAN or virtual Ethernet (veth) pair. Specify this device as the network interface in the node configuration file.

Related information

[Creating node configuration files](#)

Considerations and recommendations for MAC address cloning

MAC address cloning causes the container to use the MAC address of the host, and the host to use the MAC address of either an address you specify or a randomly generated one. You should use MAC address cloning to avoid the use of promiscuous mode network configurations.

Enabling MAC cloning

In certain environments, security can be enhanced through MAC address cloning because it enables you to use a dedicated virtual NIC for the Admin Network, Grid Network, and Client Network. Having the container use the MAC address of the dedicated NIC on the host allows you to avoid using promiscuous mode network configurations.

 MAC address cloning is intended to be used with virtual server installations and might not function properly with all physical appliance configurations.

 If a node fails to start due to a MAC cloning targeted interface being busy, you might need to set the link to "down" before starting node. Additionally, it is possible that the virtual environment might prevent MAC cloning on a network interface while the link is up. If a node fails to set the MAC address and start due to an interface being busy, setting the link to "down" before starting the node might fix the issue.

MAC address cloning is disabled by default and must be set by node configuration keys. You should enable it when you install StorageGRID.

There is one key for each network:

- ADMIN_NETWORK_TARGET_TYPE_INTERFACE_CLONE_MAC
- GRID_NETWORK_TARGET_TYPE_INTERFACE_CLONE_MAC
- CLIENT_NETWORK_TARGET_TYPE_INTERFACE_CLONE_MAC

Setting the key to "true" causes the container to use the MAC address of the host's NIC. Additionally, the host will then use the MAC address of the specified container network. By default, the container address is a randomly generated address, but if you have set one using the _NETWORK_MAC node configuration key, that address is used instead. The host and container will always have different MAC addresses.



Enabling MAC cloning on a virtual host without also enabling promiscuous mode on the hypervisor might cause Linux host networking using the host's interface to stop working.

MAC cloning use cases

There are two use cases to consider with MAC cloning:

- MAC cloning not enabled: When the _CLONE_MAC key in the node configuration file is not set, or set to "false," the host will use the host NIC MAC and the container will have a StorageGRID-generated MAC unless a MAC is specified in the _NETWORK_MAC key. If an address is set in the _NETWORK_MAC key, the container will have the address specified in the _NETWORK_MAC key. This configuration of keys requires the use of promiscuous mode.
- MAC cloning enabled: When the _CLONE_MAC key in the node configuration file is set to "true," the container uses the host NIC MAC, and the host uses a StorageGRID-generated MAC unless a MAC is specified in the _NETWORK_MAC key. If an address is set in the _NETWORK_MAC key, the host uses the specified address instead of a generated one. In this configuration of keys, you should not use promiscuous mode.



If you do not want to use MAC address cloning and would rather allow all interfaces to receive and transmit data for MAC addresses other than the ones assigned by the hypervisor, ensure that the security properties at the virtual switch and port group levels are set to **Accept** for Promiscuous Mode, MAC Address Changes, and Forged Transmits. The values set on the virtual switch can be overridden by the values at the port group level, so ensure that settings are the same in both places.

To enable MAC cloning, see the [instructions for creating node configuration files](#).

MAC cloning example

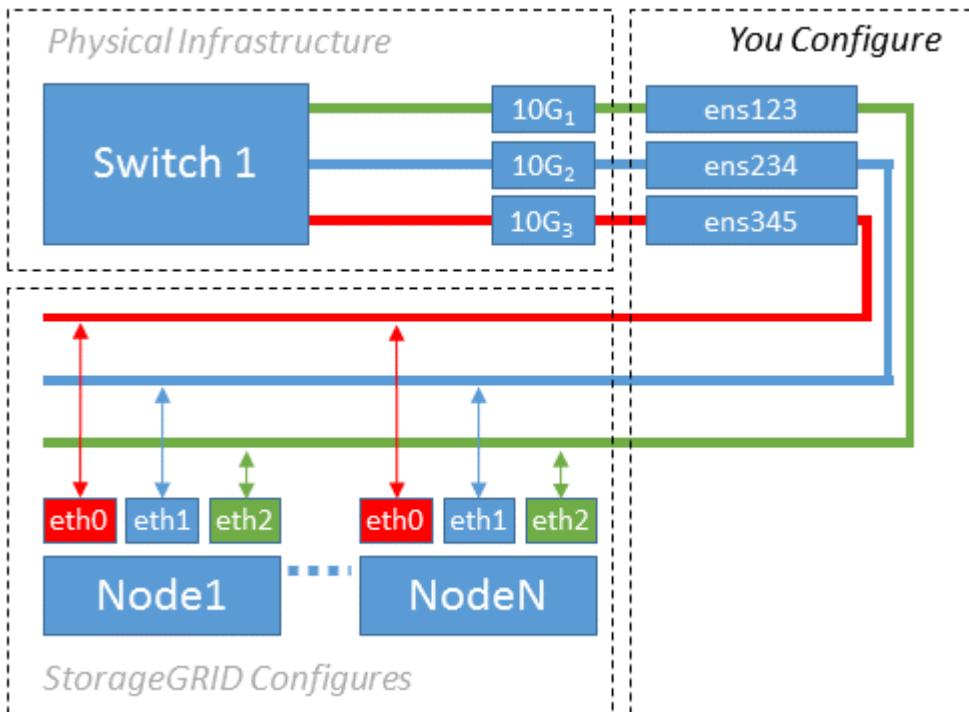
Example of MAC cloning enabled with a host having MAC address of 11:22:33:44:55:66 for the interface ens256 and the following keys in the node configuration file:

- ADMIN_NETWORK_TARGET = ens256
- ADMIN_NETWORK_MAC = b2:9c:02:c2:27:10
- ADMIN_NETWORK_TARGET_TYPE_INTERFACE_CLONE_MAC = true

Result: the host MAC for ens256 is b2:9c:02:c2:27:10 and the Admin Network MAC is 11:22:33:44:55:66

Example 1: 1-to-1 mapping to physical or virtual NICs

Example 1 describes a simple physical interface mapping that requires little or no host-side configuration.



The Linux operating system creates the `ensXYZ` interfaces automatically during installation or boot, or when the interfaces are hot-added. No configuration is required other than ensuring that the interfaces are set to come up automatically after boot. You do have to determine which `ensXYZ` corresponds to which StorageGRID network (Grid, Admin, or Client) so you can provide the correct mappings later in the configuration process.

Note that the figure shows multiple StorageGRID nodes; however, you would normally use this configuration for single-node VMs.

If Switch 1 is a physical switch, you should configure the ports connected to interfaces 10G1 through 10G3 for access mode, and place them on the appropriate VLANs.

Example 2: LACP bond carrying VLANs

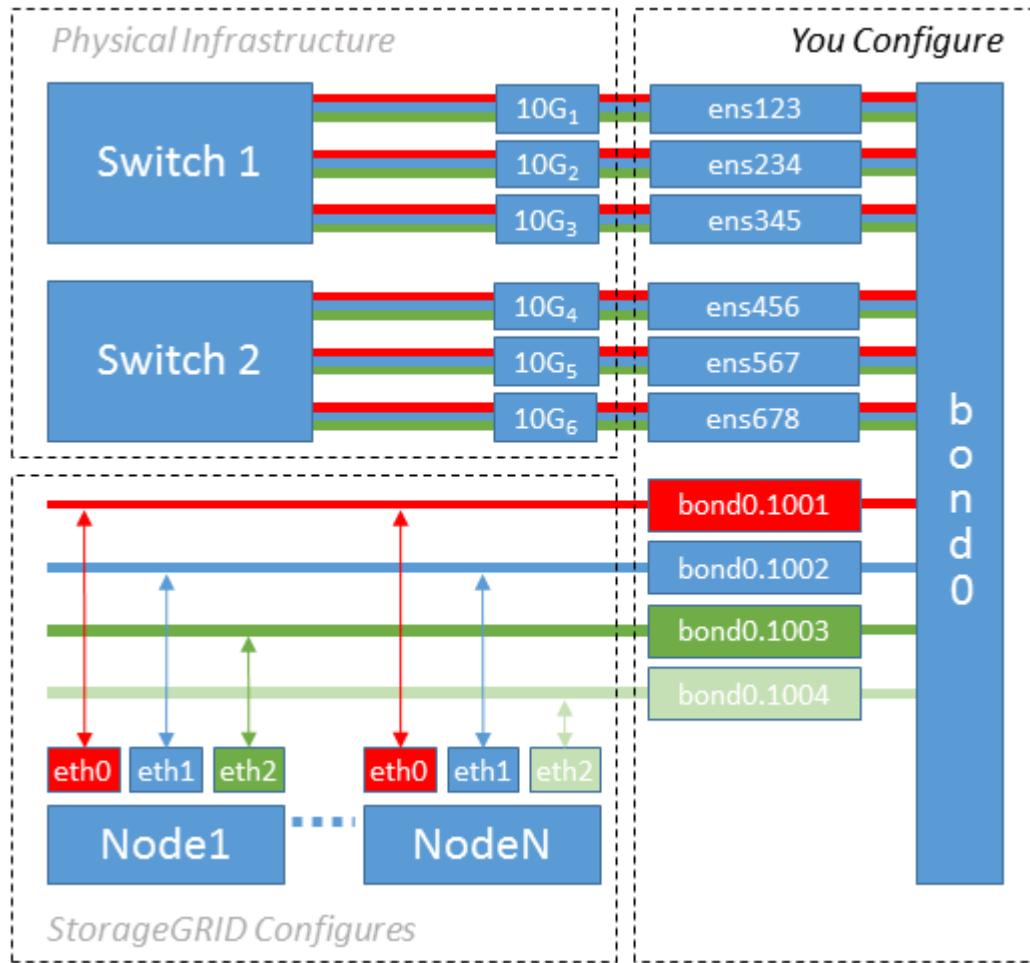
About this task

Example 2 assumes you are familiar with bonding network interfaces and with creating VLAN interfaces on the Linux distribution you are using.

Example 2 describes a generic, flexible, VLAN-based scheme that facilitates the sharing of all available network bandwidth across all nodes on a single host. This example is particularly applicable to bare metal hosts.

To understand this example, suppose you have three separate subnets for the Grid, Admin, and Client Networks at each data center. The subnets are on separate VLANs (1001, 1002, and 1003) and are presented to the host on a LACP-bonded trunk port (`bond0`). You would configure three VLAN interfaces on the bond: `bond0.1001`, `bond0.1002`, and `bond0.1003`.

If you require separate VLANs and subnets for node networks on the same host, you can add VLAN interfaces on the bond and map them into the host (shown as bond0.1004 in the illustration).



Steps

1. Aggregate all physical network interfaces that will be used for StorageGRID network connectivity into a single LACP bond.

Use the same name for the bond on every host. For example, `bond0`.

2. Create VLAN interfaces that use this bond as their associated “physical device,” using the standard VLAN interface naming convention `physdev-name.VLAN_ID`.

Note that steps 1 and 2 require appropriate configuration on the edge switches terminating the other ends of the network links. The edge switch ports must also be aggregated into a LACP port channel, configured as a trunk, and allowed to pass all required VLANs.

Sample interface configuration files for this per-host networking configuration scheme are provided.

Related information

[Example /etc/sysconfig/network-scripts](#)

Configure host storage

You must allocate block storage volumes to each host.

What you'll need

You have reviewed the following topics, which provide information you need to accomplish this task:

[Storage and performance requirements](#)

[Node container migration requirements](#)

About this task

When allocating block storage volumes (LUNs) to hosts, use the tables in “Storage requirements” to determine the following:

- Number of volumes required for each host (based on the number and types of nodes that will be deployed on that host)
- Storage category for each volume (that is, System Data or Object Data)
- Size of each volume

You will use this information as well as the persistent name assigned by Linux to each physical volume when you deploy StorageGRID nodes on the host.



You do not need to partition, format, or mount any of these volumes; you just need to ensure they are visible to the hosts.

Avoid using “raw” special device files (`/dev/sdb`, for example) as you compose your list of volume names. These files can change across reboots of the host, which will impact proper operation of the system. If you are using iSCSI LUNs and device mapper multipathing, consider using multipath aliases in the `/dev/mapper` directory, especially if your SAN topology includes redundant network paths to the shared storage. Alternatively, you can use the system-created softlinks under `/dev/disk/by-path/` for your persistent device names.

For example:

```
ls -l
$ ls -l /dev/disk/by-path/
total 0
lrwxrwxrwx 1 root root 9 Sep 19 18:53 pci-0000:00:07.1-ata-2 -> ../../sr0
lrwxrwxrwx 1 root root 9 Sep 19 18:53 pci-0000:03:00.0-scsi-0:0:0:0 ->
../../sda
lrwxrwxrwx 1 root root 10 Sep 19 18:53 pci-0000:03:00.0-scsi-0:0:0:0-part1
-> ../../sda1
lrwxrwxrwx 1 root root 10 Sep 19 18:53 pci-0000:03:00.0-scsi-0:0:0:0-part2
-> ../../sda2
lrwxrwxrwx 1 root root 9 Sep 19 18:53 pci-0000:03:00.0-scsi-0:0:1:0 ->
../../sdb
lrwxrwxrwx 1 root root 9 Sep 19 18:53 pci-0000:03:00.0-scsi-0:0:2:0 ->
../../sdc
lrwxrwxrwx 1 root root 9 Sep 19 18:53 pci-0000:03:00.0-scsi-0:0:3:0 ->
../../sdd
```

Results will differ for each installation.

Assign friendly names to each of these block storage volumes to simplify the initial StorageGRID installation and future maintenance procedures. If you are using the device mapper multipath driver for redundant access to shared storage volumes, you can use the alias field in your /etc/multipath.conf file.

For example:

```
multipaths {
    multipath {
        wwid 3600a09800059d6df00005df2573c2c30
        alias docker-storage-volume-hostA
    }
    multipath {
        wwid 3600a09800059d6df00005df3573c2c30
        alias sgws-adm1-var-local
    }
    multipath {
        wwid 3600a09800059d6df00005df4573c2c30
        alias sgws-adm1-audit-logs
    }
    multipath {
        wwid 3600a09800059d6df00005df5573c2c30
        alias sgws-adm1-tables
    }
    multipath {
        wwid 3600a09800059d6df00005df6573c2c30
        alias sgws-gw1-var-local
    }
    multipath {
        wwid 3600a09800059d6df00005df7573c2c30
        alias sgws-sn1-var-local
    }
    multipath {
        wwid 3600a09800059d6df00005df7573c2c30
        alias sgws-sn1-rangedb-0
    }
    ...
}
```

This will cause the aliases to appear as block devices in the /dev/mapper directory on the host, allowing you to specify a friendly, easily-validated name whenever a configuration or maintenance operation requires specifying a block storage volume.



If you are setting up shared storage to support StorageGRID node migration and using device mapper multipathing, you can create and install a common `/etc/multipath.conf` on all co-located hosts. Just make sure to use a different container engine storage volume on each host. Using aliases and including the target hostname in the alias for each container engine storage volume LUN will make this easy to remember and is recommended.

Related information

[Configure container engine storage volume](#)

Configure container engine storage volume

Before installing the container engine (Docker or Podman), you might need to format the storage volume and mount it.

About this task

You can skip these steps if you plan to use local storage for the Docker or Podman storage volume and have sufficient space available on the host partition containing `/var/lib/docker` for Docker and `/var/lib/containers` for Podman.



Podman is supported only on Red Hat Enterprise Linux (RHEL).

Steps

1. Create a file system on the container engine storage volume:

```
sudo mkfs.ext4 container-engine-storage-volume-device
```

2. Mount the container engine storage volume:

- For Docker:

```
sudo mkdir -p /var/lib/docker
sudo mount container-storage-volume-device /var/lib/docker
```

- For Podman:

```
sudo mkdir -p /var/lib/containers
sudo mount container-storage-volume-device /var/lib/containers
```

3. Add an entry for `container-storage-volume-device` to `/etc/fstab`.

This step ensures that the storage volume will remount automatically after host reboots.

Install Docker

The StorageGRID system runs on Red Hat Enterprise Linux or CentOS as a collection of containers. If you have chosen to use the Docker container engine, follow these steps to install Docker. Otherwise, [Install Podman](#).

Steps

1. Install Docker by following the instructions for your Linux distribution.



If Docker is not included with your Linux distribution, you can download it from the Docker website.

2. Ensure Docker has been enabled and started by running the following two commands:

```
sudo systemctl enable docker
```

```
sudo systemctl start docker
```

3. Confirm you have installed the expected version of Docker by entering the following:

```
sudo docker version
```

The Client and Server versions must be 1.11.0 or later.

Install Podman

The StorageGRID system runs on Red Hat Enterprise Linux as a collection of containers. If you have chosen to use the Podman container engine, follow these steps to install Podman. Otherwise, [install Docker](#).



Podman is supported only on Red Hat Enterprise Linux (RHEL).

Steps

1. Install Podman and Podman-Docker by following the instructions for your Linux distribution.



You must also install the Podman-Docker package when you install Podman.

2. Confirm you have installed the expected version of Podman and Podman-Docker by entering the following:

```
sudo docker version
```



The Podman-Docker package allows you to use Docker commands.

The Client and Server versions must be 3.2.3 or later.

```
Version: 3.2.3
API Version: 3.2.3
Go Version: go1.15.7
Built: Tue Jul 27 03:29:39 2021
OS/Arch: linux/amd64
```

Install StorageGRID host services

You use the StorageGRID RPM package to install the StorageGRID host services.

About this task

These instructions describe how to install the host services from the RPM packages. As an alternative, you can use the Yum repository metadata included in the installation archive to install the RPM packages remotely. See the Yum repository instructions for your Linux operating system.

Steps

1. Copy the StorageGRID RPM packages to each of your hosts, or make them available on shared storage.

For example, place them in the `/tmp` directory, so you can use the example command in the next step.

2. Log in to each host as root or using an account with sudo permission, and run the following commands in the order specified:

```
sudo yum --nogpgcheck localinstall /tmp/StorageGRID-Webscale-Images-
version-SHA.rpm
```

```
sudo yum --nogpgcheck localinstall /tmp/StorageGRID-Webscale-Service-
version-SHA.rpm
```



You must install the Images package first, and the Service package second.



If you placed the packages in a directory other than `/tmp`, modify the command to reflect the path you used.

Deploy virtual grid nodes (Red Hat or CentOS)

Create node configuration files for Red Hat Enterprise Linux or CentOS deployments

Node configuration files are small text files that provide the information the StorageGRID host service needs to start a node and connect it to the appropriate network and block storage resources. Node configuration files are used for virtual nodes and are not used for appliance nodes.

Where do I put the node configuration files?

You must place the configuration file for each StorageGRID node in the `/etc/storagegrid/nodes` directory on the host where the node will run. For example, if you plan to run one Admin Node, one Gateway Node, and one Storage Node on HostA, you must place three node configuration files in `/etc/storagegrid/nodes` on HostA. You can create the configuration files directly on each host using a text editor, such as vim or nano, or you can create them elsewhere and move them to each host.

What do I name the node configuration files?

The names of the configuration files are significant. The format is `node-name.conf`, where `node-name` is a name you assign to the node. This name appears in the StorageGRID Installer and is used for node maintenance operations, such as node migration.

Node names must follow these rules:

- Must be unique
- Must start with a letter
- Can contain the characters A through Z and a through z
- Can contain the numbers 0 through 9
- Can contain one or more hyphens (-)
- Must be no more than 32 characters, not including the `.conf` extension

Any files in `/etc/storagegrid/nodes` that do not follow these naming conventions will not be parsed by the host service.

If you have a multi-site topology planned for your grid, a typical node naming scheme might be:

```
site-nodetype-nodenumber.conf
```

For example, you might use `dc1-adm1.conf` for the first Admin Node in Data Center 1, and `dc2-sn3.conf` for the third Storage Node in Data Center 2. However, you can use any scheme you like, as long as all node names follow the naming rules.

What is in a node configuration file?

The configuration files contain key/value pairs, with one key and one value per line. For each key/value pair, you must follow these rules:

- The key and the value must be separated by an equal sign (=) and optional whitespace.
- The keys can contain no spaces.
- The values can contain embedded spaces.
- Any leading or trailing whitespace is ignored.

Some keys are required for every node, while others are optional or only required for certain node types.

The table defines the acceptable values for all supported keys. In the middle column:

R: required

BP: best practice

O: optional

Key	R, BP, or O?	Value
ADMIN_IP	BP	<p>Grid Network IPv4 address of the primary Admin Node for the grid to which this node belongs. Use the same value you specified for GRID_NETWORK_IP for the grid node with NODE_TYPE = VM_Admin_Node and ADMIN_ROLE = Primary. If you omit this parameter, the node attempts to discover a primary Admin Node using mDNS.</p> <p>How grid nodes discover the primary Admin Node</p> <p>Note: This value is ignored, and might be prohibited, on the primary Admin Node.</p>
ADMIN_NETWORK_CONFIG	O	DHCP, STATIC, or DISABLED
ADMIN_NETWORK_ESL	O	<p>Comma-separated list of subnets in CIDR notation to which this node should communicate via the Admin Network gateway.</p> <p>Example: 172.16.0.0/21,172.17.0.0/21</p>
ADMIN_NETWORK_GATEWAY	O (R)	<p>IPv4 address of the local Admin Network gateway for this node. Must be on the subnet defined by ADMIN_NETWORK_IP and ADMIN_NETWORK_MASK. This value is ignored for DHCP-configured networks.</p> <p>Note: This parameter is required if ADMIN_NETWORK_ESL is specified.</p> <p>Examples:</p> <p>1.1.1.1</p> <p>10.224.4.81</p>
ADMIN_NETWORK_IP	O	<p>IPv4 address of this node on the Admin Network. This key is only required when ADMIN_NETWORK_CONFIG = STATIC; do not specify it for other values.</p> <p>Examples:</p> <p>1.1.1.1</p> <p>10.224.4.81</p>

Key	R, BP, or O?	Value
ADMIN_NETWORK_MAC	O	<p>The MAC address for the Admin Network interface in the container.</p> <p>This field is optional. If omitted, a MAC address will be generated automatically.</p> <p>Must be 6 pairs of hexadecimal digits separated by colons.</p> <p>Example: b2:9c:02:c2:27:10</p>
ADMIN_NETWORK_MASK	O	<p>IPv4 netmask for this node, on the Admin Network. This key is only required when ADMIN_NETWORK_CONFIG = STATIC; do not specify it for other values.</p> <p>Examples:</p> <p>255.255.255.0</p> <p>255.255.248.0</p>
ADMIN_NETWORK_MTU	O	<p>The maximum transmission unit (MTU) for this node on the Admin Network. Do not specify if ADMIN_NETWORK_CONFIG = DHCP. If specified, the value must be between 1280 and 9216. If omitted, 1500 is used.</p> <p>If you want to use jumbo frames, set the MTU to a value suitable for jumbo frames, such as 9000. Otherwise, keep the default value.</p> <p>IMPORTANT: The MTU value of the network must match the value configured on the switch port the node is connected to. Otherwise, network performance issues or packet loss might occur.</p> <p>Examples:</p> <p>1500</p> <p>8192</p>

Key	R, BP, or O?	Value
ADMIN_NETWORK_TARGET	BP	<p>Name of the host device that you will use for Admin Network access by the StorageGRID node. Only network interface names are supported. Typically, you use a different interface name than what was specified for GRID_NETWORK_TARGET or CLIENT_NETWORK_TARGET.</p> <p>Note: Do not use bond or bridge devices as the network target. Either configure a VLAN (or other virtual interface) on top of the bond device, or use a bridge and virtual Ethernet (veth) pair.</p> <p>Best practice: Specify a value even if this node will not initially have an Admin Network IP address. Then you can add an Admin Network IP address later, without having to reconfigure the node on the host.</p> <p>Examples:</p> <p>bond0.1002</p> <p>ens256</p>
ADMIN_NETWORK_TARGET_TYPE	O	<p>Interface</p> <p>(This is the only supported value.)</p>
ADMIN_NETWORK_TARGET_TYPE_INTERFACE_CLONE_MAC	BP	<p>True or False</p> <p>Set the key to "true" to cause the StorageGRID container use the MAC address of the host target interface on the Admin Network.</p> <p>Best practice: In networks where promiscuous mode would be required, use the ADMIN_NETWORK_TARGET_TYPE_INTERFACE_CLONE_MAC key instead.</p> <p>For more details on MAC cloning:</p> <p>Considerations and recommendations for MAC address cloning (Red Hat Enterprise Linux or CentOS)</p> <p>Considerations and recommendations for MAC address cloning (Ubuntu or Debian)</p>
ADMIN_ROLE	R	<p>Primary or Non-Primary</p> <p>This key is only required when NODE_TYPE = VM_Admin_Node; do not specify it for other node types.</p>

Key	R, BP, or O?	Value
BLOCK_DEVICE_AUDIT_LOGS	R	<p>Path and name of the block device special file this node will use for persistent storage of audit logs. This key is only required for nodes with NODE_TYPE = VM_Admin_Node; do not specify it for other node types.</p> <p>Examples:</p> <pre>/dev/disk/by-path/pci-0000:03:00.0-scsi-0:0:0:0 /dev/disk/by-id/wwn-0x600a09800059d6df000060d757b475fd /dev/mapper/sgws-adm1-audit-logs</pre>

Key	R, BP, or O?	Value
BLOCK_DEVICE_RANGEDB_000	R	Path and name of the block device special file this node will use for persistent object storage. This key is only required for nodes with NODE_TYPE = VM_Storage_Node; do not specify it for other node types.
BLOCK_DEVICE_RANGEDB_001		Only BLOCK_DEVICE_RANGEDB_000 is required; the rest are optional. The block device specified for BLOCK_DEVICE_RANGEDB_000 must be at least 4 TB; the others can be smaller.
BLOCK_DEVICE_RANGEDB_002		
BLOCK_DEVICE_RANGEDB_003		
BLOCK_DEVICE_RANGEDB_004		
BLOCK_DEVICE_RANGEDB_005		Note: For compatibility with existing deployments, two-digit keys are supported for upgraded nodes.
BLOCK_DEVICE_RANGEDB_006		Examples:
BLOCK_DEVICE_RANGEDB_007		/dev/disk/by-path/pci-0000:03:00.0-scsi-0:0:0:0
BLOCK_DEVICE_RANGEDB_008		/dev/disk/by-id/wwn-0x600a09800059d6df000060d757b475fd
BLOCK_DEVICE_RANGEDB_009		/dev/mapper/sgws-sn1-rangedb-000
BLOCK_DEVICE_RANGEDB_010		
BLOCK_DEVICE_RANGEDB_011		
BLOCK_DEVICE_RANGEDB_012		
BLOCK_DEVICE_RANGEDB_013		
BLOCK_DEVICE_RANGEDB_014		
BLOCK_DEVICE_RANGEDB_015		

Key	R, BP, or O?	Value
BLOCK_DEVICE_TABLES	R	<p>Path and name of the block device special file this node will use for persistent storage of database tables. This key is only required for nodes with NODE_TYPE = VM_Admin_Node; do not specify it for other node types.</p> <p>Examples:</p> <pre>/dev/disk/by-path/pci-0000:03:00.0-scsi-0:0:0:0 /dev/disk/by-id/wwn-0x600a09800059d6df000060d757b475fd /dev/mapper/sgws-adm1-tables</pre>
BLOCK_DEVICE_VAR_LOCAL	R	<p>Path and name of the block device special file this node will use for its /var/local persistent storage.</p> <p>Examples:</p> <pre>/dev/disk/by-path/pci-0000:03:00.0-scsi-0:0:0:0 /dev/disk/by-id/wwn-0x600a09800059d6df000060d757b475fd /dev/mapper/sgws-sn1-var-local</pre>
CLIENT_NETWORK_CONFIG	O	DHCP, STATIC, or DISABLED
CLIENT_NETWORK_GATEWAY	O	<p>IPv4 address of the local Client Network gateway for this node, which must be on the subnet defined by CLIENT_NETWORK_IP and CLIENT_NETWORK_MASK. This value is ignored for DHCP-configured networks.</p> <p>Examples:</p> <pre>1.1.1.1 10.224.4.81</pre>

Key	R, BP, or O?	Value
CLIENT_NETWORK_IP	O	<p>IPv4 address of this node on the Client Network. This key is only required when CLIENT_NETWORK_CONFIG = STATIC; do not specify it for other values.</p> <p>Examples:</p> <p>1.1.1.1</p> <p>10.224.4.81</p>
CLIENT_NETWORK_MAC	O	<p>The MAC address for the Client Network interface in the container.</p> <p>This field is optional. If omitted, a MAC address will be generated automatically.</p> <p>Must be 6 pairs of hexadecimal digits separated by colons.</p> <p>Example: b2:9c:02:c2:27:20</p>
CLIENT_NETWORK_MASK	O	<p>IPv4 netmask for this node on the Client Network. This key is only required when CLIENT_NETWORK_CONFIG = STATIC; do not specify it for other values.</p> <p>Examples:</p> <p>255.255.255.0</p> <p>255.255.248.0</p>
CLIENT_NETWORK_MTU	O	<p>The maximum transmission unit (MTU) for this node on the Client Network. Do not specify if CLIENT_NETWORK_CONFIG = DHCP. If specified, the value must be between 1280 and 9216. If omitted, 1500 is used.</p> <p>If you want to use jumbo frames, set the MTU to a value suitable for jumbo frames, such as 9000. Otherwise, keep the default value.</p> <p>IMPORTANT: The MTU value of the network must match the value configured on the switch port the node is connected to. Otherwise, network performance issues or packet loss might occur.</p> <p>Examples:</p> <p>1500</p> <p>8192</p>

Key	R, BP, or O?	Value
CLIENT_NETWORK_TARGET	BP	<p>Name of the host device that you will use for Client Network access by the StorageGRID node. Only network interface names are supported. Typically, you use a different interface name than what was specified for GRID_NETWORK_TARGET or ADMIN_NETWORK_TARGET.</p> <p>Note: Do not use bond or bridge devices as the network target. Either configure a VLAN (or other virtual interface) on top of the bond device, or use a bridge and virtual Ethernet (veth) pair.</p> <p>Best practice: Specify a value even if this node will not initially have a Client Network IP address. Then you can add a Client Network IP address later, without having to reconfigure the node on the host.</p> <p>Examples:</p> <p>bond0.1003</p> <p>ens423</p>
CLIENT_NETWORK_TARGET_TYPE	O	<p>Interface</p> <p>(This is only supported value.)</p>
CLIENT_NETWORK_TARGET_TYPE_INTERFACE_CLONE_MAC	BP	<p>True or False</p> <p>Set the key to "true" to cause the StorageGRID container to use the MAC address of the host target interface on the Client Network.</p> <p>Best practice: In networks where promiscuous mode would be required, use the CLIENT_NETWORK_TARGET_TYPE_INTERFACE_CLONE_MAC key instead.</p> <p>For more details on MAC cloning:</p> <p>Considerations and recommendations for MAC address cloning (Red Hat Enterprise Linux or CentOS)</p> <p>Considerations and recommendations for MAC address cloning (Ubuntu or Debian)</p>
GRID_NETWORK_CONFIG	BP	<p>STATIC or DHCP</p> <p>(Defaults to STATIC if not specified.)</p>

Key	R, BP, or O?	Value
GRID_NETWORK_GATEWAY	R	<p>IPv4 address of the local Grid Network gateway for this node, which must be on the subnet defined by GRID_NETWORK_IP and GRID_NETWORK_MASK. This value is ignored for DHCP-configured networks.</p> <p>If the Grid Network is a single subnet with no gateway, use either the standard gateway address for the subnet (X.Y.Z.1) or this node's GRID_NETWORK_IP value; either value will simplify potential future Grid Network expansions.</p>
GRID_NETWORK_IP	R	<p>IPv4 address of this node on the Grid Network. This key is only required when GRID_NETWORK_CONFIG = STATIC; do not specify it for other values.</p> <p>Examples:</p> <p>1.1.1.1</p> <p>10.224.4.81</p>
GRID_NETWORK_MAC	O	<p>The MAC address for the Grid Network interface in the container.</p> <p>This field is optional. If omitted, a MAC address will be generated automatically.</p> <p>Must be 6 pairs of hexadecimal digits separated by colons.</p> <p>Example: b2:9c:02:c2:27:30</p>
GRID_NETWORK_MASK	O	<p>IPv4 netmask for this node on the Grid Network. This key is only required when GRID_NETWORK_CONFIG = STATIC; do not specify it for other values.</p> <p>Examples:</p> <p>255.255.255.0</p> <p>255.255.248.0</p>

Key	R, BP, or O?	Value
GRID_NETWORK_MTU	O	<p>The maximum transmission unit (MTU) for this node on the Grid Network. Do not specify if GRID_NETWORK_CONFIG = DHCP. If specified, the value must be between 1280 and 9216. If omitted, 1500 is used.</p> <p>If you want to use jumbo frames, set the MTU to a value suitable for jumbo frames, such as 9000. Otherwise, keep the default value.</p> <p>IMPORTANT: The MTU value of the network must match the value configured on the switch port the node is connected to. Otherwise, network performance issues or packet loss might occur.</p> <p>IMPORTANT: For the best network performance, all nodes should be configured with similar MTU values on their Grid Network interfaces. The Grid Network MTU mismatch alert is triggered if there is a significant difference in MTU settings for the Grid Network on individual nodes. The MTU values do not have to be the same for all network types.</p> <p>Examples:</p> <p>1500 8192</p>
GRID_NETWORK_TARGET	R	<p>Name of the host device that you will use for Grid Network access by the StorageGRID node. Only network interface names are supported. Typically, you use a different interface name than what was specified for ADMIN_NETWORK_TARGET or CLIENT_NETWORK_TARGET.</p> <p>Note: Do not use bond or bridge devices as the network target. Either configure a VLAN (or other virtual interface) on top of the bond device, or use a bridge and virtual Ethernet (veth) pair.</p> <p>Examples:</p> <p>bond0.1001</p> <p>ens192</p>
GRID_NETWORK_TARGET_TYPE	O	<p>Interface</p> <p>(This is the only supported value.)</p>

Key	R, BP, or O?	Value
GRID_NETWORK_TARGET_TYPE_INTERFACE_CLONE_MAC	BP	<p>True or False</p> <p>Set the value of the key to "true" to cause the StorageGRID container to use the MAC address of the host target interface on the Grid Network.</p> <p>Best practice: In networks where promiscuous mode would be required, use the GRID_NETWORK_TARGET_TYPE_INTERFACE_CLONE_MAC key instead.</p> <p>For more details on MAC cloning:</p> <p>Considerations and recommendations for MAC address cloning (Red Hat Enterprise Linux or CentOS)</p> <p>Considerations and recommendations for MAC address cloning (Ubuntu or Debian)</p>
INTERFACES_TARGET_nnn	O	<p>Name and optional description for an extra interface you want to add to this node. You can add multiple extra interfaces to each node.</p> <p>For <i>nnnn</i>, specify a unique number for each INTERFACES_TARGET entry you are adding.</p> <p>For the value, specify the name of the physical interface on the bare-metal host. Then, optionally, add a comma and provide a description of the interface, which is displayed on the VLAN interfaces page and the HA groups page.</p> <p>For example: INTERFACES_TARGET_01=ens256, Trunk</p> <p>If you add a trunk interface, you must configure a VLAN interface in StorageGRID. If you add an access interface, you can add the interface directly to an HA group; you do not need to configure a VLAN interface.</p>

Key	R, BP, or O?	Value
MAXIMUM_RAM	O	<p>The maximum amount of RAM that this node is allowed to consume. If this key is omitted, the node has no memory restrictions. When setting this field for a production-level node, specify a value that is at least 24 GB and 16 to 32 GB less than the total system RAM.</p> <p>Note: The RAM value affects a node's actual metadata reserved space. See the instructions for administering StorageGRID for a description of what Metadata Reserved Space is.</p> <p>The format for this field is <number><unit>, where <unit> can be b, k, m, or g.</p> <p>Examples:</p> <p>24g</p> <p>38654705664b</p> <p>Note: If you want to use this option, you must enable kernel support for memory cgroups.</p>
NODE_TYPE	R	<p>Type of node:</p> <p>VM_Admin_Node VM_Storage_Node VM_Archive_Node VM_API_Gateway</p>

Key	R, BP, or O?	Value
PORT_REMAP	O	<p>Remaps any port used by a node for internal grid node communications or external communications. Remapping ports is necessary if enterprise networking policies restrict one or more ports used by StorageGRID, as described in “Internal grid node communications” or “External communications.”</p> <p>IMPORTANT: Do not remap the ports you are planning to use to configure load balancer endpoints.</p> <p>Note: If only PORT_REMAP is set, the mapping that you specify is used for both inbound and outbound communications. If PORT_REMAP_INBOUND is also specified, PORT_REMAP applies only to outbound communications.</p> <p>The format used is: <network type>/<protocol>/<default port used by grid node>/<new port>, where <network type> is grid, admin, or client, and protocol is tcp or udp.</p> <p>For example:</p> <pre>PORT_REMAP = client/tcp/18082/443</pre>
PORT_REMAP_INBOUND	O	<p>Remaps inbound communications to the specified port. If you specify PORT_REMAP_INBOUND but do not specify a value for PORT_REMAP, outbound communications for the port are unchanged.</p> <p>IMPORTANT: Do not remap the ports you are planning to use to configure load balancer endpoints.</p> <p>The format used is: <network type>/<protocol:>/<remapped port >/<default port used by grid node>, where <network type> is grid, admin, or client, and protocol is tcp or udp.</p> <p>For example:</p> <pre>PORT_REMAP_INBOUND = grid/tcp/3022/22</pre>

Related information

[Networking guidelines](#)

How grid nodes discover the primary Admin Node

Grid nodes communicate with the primary Admin Node for configuration and management. Each grid node must know the IP address of the primary Admin Node on the Grid Network.

To ensure that a grid node can access the primary Admin Node, you can do either of the following when deploying the node:

- You can use the ADMIN_IP parameter to enter the primary Admin Node's IP address manually.
- You can omit the ADMIN_IP parameter to have the grid node discover the value automatically. Automatic discovery is especially useful when the Grid Network uses DHCP to assign the IP address to the primary Admin Node.

Automatic discovery of the primary Admin Node is accomplished using a multicast Domain Name System (mDNS). When the primary Admin Node first starts up, it publishes its IP address using mDNS. Other nodes on the same subnet can then query for the IP address and acquire it automatically. However, because multicast IP traffic is not normally routable across subnets, nodes on other subnets cannot acquire the primary Admin Node's IP address directly.

If you use automatic discovery:

 • You must include the ADMIN_IP setting for at least one grid node on any subnets that the primary Admin Node is not directly attached to. This grid node will then publish the primary Admin Node's IP address for other nodes on the subnet to discover with mDNS.

• Ensure that your network infrastructure supports passing multi-cast IP traffic within a subnet.

Example node configuration files

You can use the example node configuration files to help set up the node configuration files for your StorageGRID system. The examples show node configuration files for all types of grid nodes.

For most nodes, you can add Admin and Client Network addressing information (IP, mask, gateway, and so on) when you configure the grid using the Grid Manager or the Installation API. The exception is the primary Admin Node. If you want to browse to the Admin Network IP of the primary Admin Node to complete grid configuration (because the Grid Network is not routed, for example), you must configure the Admin Network connection for the primary Admin Node in its node configuration file. This is shown in the example.

 In the examples, the Client Network target has been configured as a best practice, even though the Client Network is disabled by default.

Example for primary Admin Node

Example file name: /etc/storagegrid/nodes/dc1-adm1.conf

Example file contents:

```

NODE_TYPE = VM_Admin_Node
ADMIN_ROLE = Primary
BLOCK_DEVICE_VAR_LOCAL = /dev/mapper/dc1-adm1-var-local
BLOCK_DEVICE_AUDIT_LOGS = /dev/mapper/dc1-adm1-audit-logs
BLOCK_DEVICE_TABLES = /dev/mapper/dc1-adm1-tables
GRID_NETWORK_TARGET = bond0.1001
ADMIN_NETWORK_TARGET = bond0.1002
CLIENT_NETWORK_TARGET = bond0.1003

GRID_NETWORK_IP = 10.1.0.2
GRID_NETWORK_MASK = 255.255.255.0
GRID_NETWORK_GATEWAY = 10.1.0.1

ADMIN_NETWORK_CONFIG = STATIC
ADMIN_NETWORK_IP = 192.168.100.2
ADMIN_NETWORK_MASK = 255.255.248.0
ADMIN_NETWORK_GATEWAY = 192.168.100.1
ADMIN_NETWORK_ESL = 192.168.100.0/21,172.16.0.0/21,172.17.0.0/21

```

Example for Storage Node

Example file name: /etc/storagegrid/nodes/dc1-sn1.conf

Example file contents:

```

NODE_TYPE = VM_Storage_Node
ADMIN_IP = 10.1.0.2
BLOCK_DEVICE_VAR_LOCAL = /dev/mapper/dc1-sn1-var-local
BLOCK_DEVICE_RANGEDB_00 = /dev/mapper/dc1-sn1-rangedb-0
BLOCK_DEVICE_RANGEDB_01 = /dev/mapper/dc1-sn1-rangedb-1
BLOCK_DEVICE_RANGEDB_02 = /dev/mapper/dc1-sn1-rangedb-2
BLOCK_DEVICE_RANGEDB_03 = /dev/mapper/dc1-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

```

Example for Archive Node

Example file name: /etc/storagegrid/nodes/dc1-arcl.conf

Example file contents:

```
NODE_TYPE = VM_Archive_Node
ADMIN_IP = 10.1.0.2
BLOCK_DEVICE_VAR_LOCAL = /dev/mapper/dc1-arc1-var-local
GRID_NETWORK_TARGET = bond0.1001
ADMIN_NETWORK_TARGET = bond0.1002
CLIENT_NETWORK_TARGET = bond0.1003

GRID_NETWORK_IP = 10.1.0.4
GRID_NETWORK_MASK = 255.255.255.0
GRID_NETWORK_GATEWAY = 10.1.0.1
```

Example for Gateway Node

Example file name: /etc/storagegrid/nodes/dc1-gw1.conf

Example file contents:

```
NODE_TYPE = VM_API_Gateway
ADMIN_IP = 10.1.0.2
BLOCK_DEVICE_VAR_LOCAL = /dev/mapper/dc1-gw1-var-local
GRID_NETWORK_TARGET = bond0.1001
ADMIN_NETWORK_TARGET = bond0.1002
CLIENT_NETWORK_TARGET = bond0.1003
GRID_NETWORK_IP = 10.1.0.5
GRID_NETWORK_MASK = 255.255.255.0
GRID_NETWORK_GATEWAY = 10.1.0.1
```

Example for a non-primary Admin Node

Example file name: /etc/storagegrid/nodes/dc1-adm2.conf

Example file contents:

```
NODE_TYPE = VM_Admin_Node
ADMIN_ROLE = Non-Primary
ADMIN_IP = 10.1.0.2
BLOCK_DEVICE_VAR_LOCAL = /dev/mapper/dc1-adm2-var-local
BLOCK_DEVICE_AUDIT_LOGS = /dev/mapper/dc1-adm2-audit-logs
BLOCK_DEVICE_TABLES = /dev/mapper/dc1-adm2-tables
GRID_NETWORK_TARGET = bond0.1001
ADMIN_NETWORK_TARGET = bond0.1002
CLIENT_NETWORK_TARGET = bond0.1003

GRID_NETWORK_IP = 10.1.0.6
GRID_NETWORK_MASK = 255.255.255.0
GRID_NETWORK_GATEWAY = 10.1.0.1
```

Validate the StorageGRID configuration

After creating configuration files in `/etc/storagegrid/nodes` for each of your StorageGRID nodes, you must validate the contents of those files.

To validate the contents of the configuration files, run the following command on each host:

```
sudo storagegrid node validate all
```

If the files are correct, the output shows **PASSED** for each configuration file, as shown in the example.

```
Checking for misnamed node configuration files... PASSED
Checking configuration file for node dc1-adm1... PASSED
Checking configuration file for node dc1-gw1... PASSED
Checking configuration file for node dc1-sn1... PASSED
Checking configuration file for node dc1-sn2... PASSED
Checking configuration file for node dc1-sn3... PASSED
Checking for duplication of unique values between nodes... PASSED
```



For an automated installation, you can suppress this output by using the `-q` or `--quiet` options in the `storagegrid` command (for example, `storagegrid --quiet...`). If you suppress the output, the command will have a non-zero exit value if any configuration warnings or errors were detected.

If the configuration files are incorrect, the issues are shown as **WARNING** and **ERROR**, as shown in the example. If any configuration errors are found, you must correct them before you continue with the installation.

```
Checking for misnamed node configuration files...
  WARNING: ignoring /etc/storagegrid/nodes/dc1-adm1
  WARNING: ignoring /etc/storagegrid/nodes/dc1-sn2.conf.keep
  WARNING: ignoring /etc/storagegrid/nodes/my-file.txt
Checking configuration file for node dc1-adm1...
  ERROR: NODE_TYPE = VM_Foo_Node
    VM_Foo_Node is not a valid node type.  See *.conf.sample
  ERROR: ADMIN_ROLE = Foo
    Foo is not a valid admin role.  See *.conf.sample
  ERROR: BLOCK_DEVICE_VAR_LOCAL = /dev/mapper/sgws-gw1-var-local
    /dev/mapper/sgws-gw1-var-local is not a valid block device
Checking configuration file for node dc1-gw1...
  ERROR: GRID_NETWORK_TARGET = bond0.1001
    bond0.1001 is not a valid interface.  See `ip link show`
  ERROR: GRID_NETWORK_IP = 10.1.3
    10.1.3 is not a valid IPv4 address
  ERROR: GRID_NETWORK_MASK = 255.248.255.0
    255.248.255.0 is not a valid IPv4 subnet mask
Checking configuration file for node dc1-sn1...
  ERROR: GRID_NETWORK_GATEWAY = 10.2.0.1
    10.2.0.1 is not on the local subnet
  ERROR: ADMIN_NETWORK_ESL = 192.168.100.0/21,172.16.0foo
    Could not parse subnet list
Checking configuration file for node dc1-sn2... PASSED
Checking configuration file for node dc1-sn3... PASSED
Checking for duplication of unique values between nodes...
  ERROR: GRID_NETWORK_IP = 10.1.0.4
    dc1-sn2 and dc1-sn3 have the same GRID_NETWORK_IP
  ERROR: BLOCK_DEVICE_VAR_LOCAL = /dev/mapper/sgws-sn2-var-local
    dc1-sn2 and dc1-sn3 have the same BLOCK_DEVICE_VAR_LOCAL
  ERROR: BLOCK_DEVICE_RANGEDB_00 = /dev/mapper/sgws-sn2-rangedb-0
    dc1-sn2 and dc1-sn3 have the same BLOCK_DEVICE_RANGEDB_00
```

Start the StorageGRID host service

To start your StorageGRID nodes, and ensure they restart after a host reboot, you must enable and start the StorageGRID host service.

Steps

1. Run the following commands on each host:

```
sudo systemctl enable storagegrid
sudo systemctl start storagegrid
```

2. Run the following command to ensure the deployment is proceeding:

```
sudo storagegrid node status node-name
```

For any node that returns a status of “Not-Running” or “Stopped”, run the following command:

```
sudo storagegrid node start node-name
```

3. If you have previously enabled and started the StorageGRID host service (or if you are unsure if the service has been enabled and started), also run the following command:

```
sudo systemctl reload-or-restart storagegrid
```

Configure the grid and complete installation (Red Hat or CentOS)

Navigate to the Grid Manager

You use the Grid Manager to define all of the information required to configure your StorageGRID system.

What you'll need

The primary Admin Node must be deployed and have completed the initial startup sequence.

Steps

1. Open your web browser and navigate to one of the following addresses:

```
https://primary_admin_node_ip
```

```
client_network_ip
```

Alternatively, you can access the Grid Manager on port 8443:

```
https://primary_admin_node_ip:8443
```



You can use the IP address for the primary Admin Node IP on the Grid Network or on the Admin Network, as appropriate for your network configuration.

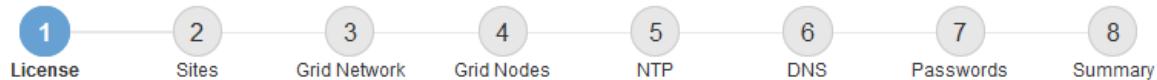
1. Click **Install a StorageGRID system**.

The page used to configure a StorageGRID system appears.

NetApp® StorageGRID®

Help ▾

Install	
---------	--



License

Enter a grid name and upload the license file provided by NetApp for your StorageGRID system.

Grid Name

License File

Specify the StorageGRID license information

You must specify the name for your StorageGRID system and upload the license file provided by NetApp.

Steps

1. On the License page, enter a meaningful name for your StorageGRID system in **Grid Name**.

After installation, the name is displayed at the top of the Nodes menu.

2. Click **Browse**, locate the NetApp License File (NLUnique_id.txt), and click **Open**.

The license file is validated, and the serial number and licensed storage capacity are displayed.



The StorageGRID installation archive includes a free license that does not provide any support entitlement for the product. You can update to a license that offers support after installation.

NetApp® StorageGRID®

Help ▾

Install	
---------	--



License

Enter a grid name and upload the license file provided by NetApp for your StorageGRID system.

Grid Name

New License File

License Serial Number

Storage Capacity (TB)

3. Click **Next**.

Add sites

You must create at least one site when you are installing StorageGRID. You can create additional sites to increase the reliability and storage capacity of your StorageGRID system.

Steps

1. On the Sites page, enter the **Site Name**.
2. To add additional sites, click the plus sign next to the last site entry and enter the name in the new **Site Name** text box.

Add as many additional sites as required for your grid topology. You can add up to 16 sites.

NetApp® StorageGRID®

Help ▾

Install

1 License 2 Sites 3 Grid Network 4 Grid Nodes 5 NTP 6 DNS 7 Passwords 8 Summary

Sites

In a single-site deployment, infrastructure and operations are centralized in one site.

In a multi-site deployment, infrastructure can be distributed asymmetrically across sites, and proportional to the needs of each site. Typically, sites are located in geographically different locations. Having multiple sites also allows the use of distributed replication and erasure coding for increased availability and resiliency.

Site Name 1	Raleigh	x
Site Name 2	Atlanta	+ x

3. Click **Next**.

Specify Grid Network subnets

You must specify the subnets that are used on the Grid Network.

About this task

The subnet entries include the subnets for the Grid Network for each site in your StorageGRID system, along with any subnets that need to be reachable via the Grid Network.

If you have multiple grid subnets, the Grid Network gateway is required. All grid subnets specified must be reachable through this gateway.

Steps

1. Specify the CIDR network address for at least one Grid Network in the **Subnet 1** text box.
2. Click the plus sign next to the last entry to add an additional network entry.

If you have already deployed at least one node, click **Discover Grid Networks Subnets** to automatically

populate the Grid Network Subnet List with the subnets reported by grid nodes that have registered with the Grid Manager.

NetApp® StorageGRID®

Help ▾

Install

1 License 2 Sites 3 Grid Network 4 Grid Nodes 5 NTP 6 DNS 7 Passwords 8 Summary

Grid Network

You must specify 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 system. Select Discover Grid Networks to automatically add subnets based on the network configuration of all registered nodes.

Note: You must manually add any subnets for NTP, DNS, LDAP, or other external servers accessed through the Grid Network gateway.

Subnet 1 172.16.0.0/21 +

Discover Grid Network subnets

3. Click **Next**.

Approve pending grid nodes

You must approve each grid node before it can join the StorageGRID system.

What you'll need

You have deployed all virtual and StorageGRID appliance grid nodes.



It is more efficient to perform one single installation of all the nodes, rather than installing some nodes now and some nodes later.

Steps

1. Review the Pending Nodes list, and confirm that it shows all of the grid nodes you deployed.



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

2. Select the radio button next to a pending node you want to approve.



Grid Nodes

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

Pending Nodes

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

						Search
	Grid Network MAC Address	Name	Type	Platform	Grid Network IPv4 Address	
<input checked="" type="radio"/>	50:6b:4b:42:d7:00	NetApp-SGA	Storage Node	StorageGRID Appliance	172.16.5.20/21	

Approved Nodes

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

						Search
	Grid Network MAC Address	Name	Site	Type	Platform	Grid Network IPv4 Address
<input checked="" type="radio"/>	00:50:56:87:42:ff	dc1-adm1	Raleigh	Admin Node	VMware VM	172.16.4.210/21
<input checked="" type="radio"/>	00:50:56:87:c0:16	dc1-s1	Raleigh	Storage Node	VMware VM	172.16.4.211/21
<input checked="" type="radio"/>	00:50:56:87:79:ee	dc1-s2	Raleigh	Storage Node	VMware VM	172.16.4.212/21
<input checked="" type="radio"/>	00:50:56:87:db:9c	dc1-s3	Raleigh	Storage Node	VMware VM	172.16.4.213/21
<input checked="" type="radio"/>	00:50:56:87:62:38	dc1-g1	Raleigh	API Gateway Node	VMware VM	172.16.4.214/21

3. Click **Approve**.

4. In General Settings, modify settings for the following properties, as necessary:

Storage Node Configuration

General Settings

Site	<input type="text" value="Raleigh"/>
Name	<input type="text" value="NetApp-SGA"/>
NTP Role	<input type="text" value="Automatic"/>
ADC Service	<input type="text" value="Automatic"/>

Grid Network

Configuration	STATIC
IPv4 Address (CIDR)	<input type="text" value="172.16.5.20/21"/>
Gateway	<input type="text" value="172.16.5.20"/>

Admin Network

Configuration	STATIC
IPv4 Address (CIDR)	<input type="text" value="10.224.5.20/21"/>
Gateway	<input type="text" value="10.224.0.1"/>
Subnets (CIDR)	<input type="text" value="10.0.0.0/8"/> ×
	<input type="text" value="172.19.0.0/16"/> ×
	<input type="text" value="172.21.0.0/16"/> + ×

Client Network

Configuration	STATIC
IPv4 Address (CIDR)	<input type="text" value="47.47.5.20/21"/>
Gateway	<input type="text" value="47.47.0.1"/>

- **Site:** The name of the site with which this grid node will be associated.
- **Name:** The name that will be assigned to the node, and the name that will be displayed in the Grid Manager. The name defaults to the name you specified when you configured the node. During this step of the installation process, you can change the name as required.



After you complete the installation, you cannot change the name of the node.



For a VMware node, you can change the name here, but this action will not change the name of the virtual machine in vSphere.

- **NTP Role:** The Network Time Protocol (NTP) role of the grid node. The options are **Automatic**, **Primary**, and **Client**. Selecting **Automatic** assigns the Primary role to Admin Nodes, Storage Nodes with ADC services, Gateway Nodes, and any grid nodes that have non-static IP addresses. All other grid nodes are assigned the Client role.



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

- **ADC service** (Storage Nodes only): Select **Automatic** to let the system determine whether the node requires the Administrative Domain Controller (ADC) service. The ADC service keeps track of the location and availability of grid services. At least three Storage Nodes at each site must include the ADC service. You cannot add the ADC service to a node after it is deployed.

5. In Grid Network, modify settings for the following properties as necessary:

- **IPv4 Address (CIDR):** The CIDR network address for the Grid Network interface (eth0 inside the container). For example: 192.168.1.234/21
- **Gateway:** The Grid Network gateway. For example: 192.168.0.1

The gateway is required if there are multiple grid subnets.



If you selected DHCP for the Grid Network configuration and you change the value here, the new value will be configured as a static address on the node. You must make sure the resulting IP address is not within a DHCP address pool.

6. If you want to configure the Admin Network for the grid node, add or update the settings in the Admin Network section as necessary.

Enter the destination subnets of the routes out of this interface in the **Subnets (CIDR)** text box. If there are multiple Admin subnets, the Admin gateway is required.



If you selected DHCP for the Admin Network configuration and you change the value here, the new value will be configured as a static address on the node. You must make sure the resulting IP address is not within a DHCP address pool.

Appliances: For a StorageGRID appliance, if the Admin Network was not configured during the initial installation using the StorageGRID Appliance Installer, it cannot be configured in this Grid Manager dialog box. Instead, you must follow these steps:

- a. Reboot the appliance: In the Appliance Installer, select **Advanced > Reboot**.
Rebooting can take several minutes.
- b. Select **Configure Networking > Link Configuration** and enable the appropriate networks.
- c. Select **Configure Networking > IP Configuration** and configure the enabled networks.
- d. Return to the Home page and click **Start Installation**.
- e. In the Grid Manager: If the node is listed in the Approved Nodes table, reset the node.
- f. Remove the node from the Pending Nodes table.
- g. Wait for the node to reappear in the Pending Nodes list.

- h. Confirm that you can configure the appropriate networks. They should already be populated with the information you provided on the IP Configuration page.

For additional information, see the installation and maintenance instructions for your appliance model.

7. If you want to configure the Client Network for the grid node, add or update the settings in the Client Network section as necessary. If the Client Network is configured, the gateway is required, and it becomes the default gateway for the node after installation.



If you selected DHCP for the Client Network configuration and you change the value here, the new value will be configured as a static address on the node. You must make sure the resulting IP address is not within a DHCP address pool.

Appliances: For a StorageGRID appliance, if the Client Network was not configured during the initial installation using the StorageGRID Appliance Installer, it cannot be configured in this Grid Manager dialog box. Instead, you must follow these steps:

- a. Reboot the appliance: In the Appliance Installer, select **Advanced > Reboot**.

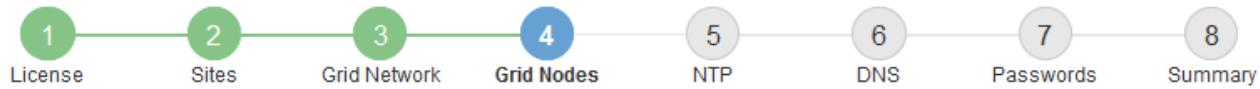
Rebooting can take several minutes.

- b. Select **Configure Networking > Link Configuration** and enable the appropriate networks.
- c. Select **Configure Networking > IP Configuration** and configure the enabled networks.
- d. Return to the Home page and click **Start Installation**.
- e. In the Grid Manager: If the node is listed in the Approved Nodes table, reset the node.
- f. Remove the node from the Pending Nodes table.
- g. Wait for the node to reappear in the Pending Nodes list.
- h. Confirm that you can configure the appropriate networks. They should already be populated with the information you provided on the IP Configuration page.

For additional information, see the installation and maintenance instructions for your appliance.

8. Click **Save**.

The grid node entry moves to the Approved Nodes list.



Grid Nodes

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

Pending Nodes

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

		<input type="button" value="Approve"/>	<input type="button" value="Remove"/>			Search	<input type="button" value=""/>				
	Grid Network MAC Address		Name		Type		Platform		Grid Network IPv4 Address		▼
<i>No results found.</i>											

Approved Nodes

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

		<input type="button" value="Edit"/>	<input type="button" value="Reset"/>	<input type="button" value="Remove"/>			Search	<input type="button" value=""/>				
	Grid Network MAC Address		Name		Site		Type		Platform		Grid Network IPv4 Address	▼
<input type="radio"/>	00:50:56:87:42:ff		dc1-adm1		Raleigh		Admin Node		VMware VM		172.16.4.210/21	
<input type="radio"/>	00:50:56:87:c0:16		dc1-s1		Raleigh		Storage Node		VMware VM		172.16.4.211/21	
<input type="radio"/>	00:50:56:87:79:ee		dc1-s2		Raleigh		Storage Node		VMware VM		172.16.4.212/21	
<input type="radio"/>	00:50:56:87:db:9c		dc1-s3		Raleigh		Storage Node		VMware VM		172.16.4.213/21	
<input type="radio"/>	00:50:56:87:62:38		dc1-g1		Raleigh		API Gateway Node		VMware VM		172.16.4.214/21	
<input type="radio"/>	50:6b:4b:42:d7:00		NetApp-SGA		Raleigh		Storage Node		StorageGRID Appliance		172.16.5.20/21	

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

You must approve all nodes that you want in the grid. However, you can return to this page at any time before you click **Install** on the Summary page. You can modify the properties of an approved grid node by selecting its radio button and clicking **Edit**.

10. When you are done approving grid nodes, click **Next**.

Specify Network Time Protocol server information

You must specify the Network Time Protocol (NTP) configuration information for the StorageGRID system, so that operations performed on separate servers can be kept synchronized.

About this task

You must specify IPv4 addresses for the NTP servers.

You must specify external NTP servers. The specified NTP servers must use the NTP protocol.

You must specify four NTP server references of Stratum 3 or better to prevent issues with time drift.



When specifying the external NTP source for a production-level StorageGRID installation, do not use the Windows Time (W32Time) service on a version of Windows earlier than Windows Server 2016. The time service on earlier versions of Windows is not sufficiently accurate and is not supported by Microsoft for use in high-accuracy environments, such as StorageGRID.

[Support boundary to configure the Windows Time service for high-accuracy environments](#)

The external NTP servers are used by the nodes to which you previously assigned Primary NTP roles.



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

Steps

1. Specify the IPv4 addresses for at least four NTP servers in the **Server 1** to **Server 4** text boxes.
2. If necessary, select the plus sign next to the last entry to add additional server entries.

NetApp® StorageGRID®

Help ▾

Install

1 License 2 Sites 3 Grid Network 4 Grid Nodes 5 NTP 6 DNS 7 Passwords 8 Summary

Network Time Protocol

Enter the IP addresses for at least four Network Time Protocol (NTP) servers, so that operations performed on separate servers are kept in sync.

Server 1	10.60.248.183
Server 2	10.227.204.142
Server 3	10.235.48.111
Server 4	0.0.0.0

+

3. Select **Next**.

Specify Domain Name System server information

You must specify Domain Name System (DNS) information for your StorageGRID system, so that you can access external servers using hostnames instead of IP addresses.

About this task

Specifying DNS server information allows you to use Fully Qualified Domain Name (FQDN) hostnames rather than IP addresses for email notifications and AutoSupport. Specifying at least two DNS servers is

recommended.



Provide two to six IPv4 addresses for DNS servers. You should select DNS servers that each site can access locally in the event of network islanding. This is to ensure an islanded site continues to have access to the DNS service. After configuring the grid-wide DNS server list, you can further customize the DNS server list for each node. For details, see the information about modifying the DNS configuration in the recovery and maintenance instructions.

If the DNS server information is omitted or incorrectly configured, a DNST alarm is triggered on each grid node's SSM service. The alarm clears when DNS is configured correctly and the new server information has reached all grid nodes.

Steps

1. Specify the IPv4 address for at least one DNS server in the **Server 1** text box.
2. If necessary, select the plus sign next to the last entry to add additional server entries.

NetApp® StorageGRID®

Install

Help ▾

1 License 2 Sites 3 Grid Network 4 Grid Nodes 5 NTP 6 DNS 7 Passwords 8 Summary

Domain Name Service

Enter the IP address for at least one Domain Name System (DNS) server, so that server hostnames can be used instead of IP addresses. Specifying at least two DNS servers is recommended. Configuring DNS enables server connectivity, email notifications, and NetApp AutoSupport.

Server 1: 10.224.223.130

Server 2: 10.224.223.136

The best practice is to specify at least two DNS servers. You can specify up to six DNS servers.

3. Select **Next**.

Specify the StorageGRID system passwords

As part of installing your StorageGRID system, you need to enter the passwords to use to secure your system and perform maintenance tasks.

About this task

Use the Install passwords page to specify the provisioning passphrase and the grid management root user password.

- The provisioning passphrase is used as an encryption key and is not stored by the StorageGRID system.
- You must have the provisioning passphrase for installation, expansion, and maintenance procedures, including downloading the Recovery Package. Therefore, it is important that you store the provisioning passphrase in a secure location.
- You can change the provisioning passphrase from the Grid Manager if you have the current one.

- The grid management root user password may be changed using the Grid Manager.
- Randomly generated command line console and SSH passwords are stored in the Passwords.txt file in the Recovery Package.

Steps

1. In **Provisioning Passphrase**, enter the provisioning passphrase that will be required to make changes to the grid topology of your StorageGRID system.

Store the provisioning passphrase in a secure place.



If after the installation completes and you want to change the provisioning passphrase later, you can use the Grid Manager. Select **CONFIGURATION > Access control > Grid passwords**.

2. In **Confirm Provisioning Passphrase**, reenter the provisioning passphrase to confirm it.
3. In **Grid Management Root User Password**, enter the password to use to access the Grid Manager as the “root” user.

Store the password in a secure place.

4. In **Confirm Root User Password**, reenter the Grid Manager password to confirm it.

NetApp® StorageGRID®

Help ▾

Install

1 License 2 Sites 3 Grid Network 4 Grid Nodes 5 NTP 6 DNS 7 Passwords 8 Summary

Passwords

Enter secure passwords that meet your organization's security policies. A text file containing the command line passwords must be downloaded during the final installation step.

Provisioning Passphrase
Confirm Provisioning Passphrase
Grid Management Root User Password
Confirm Root User Password

Create random command line passwords.

5. If you are installing a grid for proof of concept or demo purposes, optionally deselect the **Create random command line passwords** check box.

For production deployments, random passwords should always be used for security reasons. Deselect **Create random command line passwords** only for demo grids if you want to use default passwords to access grid nodes from the command line using the “root” or “admin” account.



You are prompted to download the Recovery Package file (`sgws-recovery-package-id-revision.zip`) after you click **Install** on the Summary page. You must [download this file](#) to complete the installation. The passwords required to access the system are stored in the `Passwords.txt` file, contained in the Recovery Package file.

6. Click **Next**.

Review your configuration and complete installation

You must carefully review the configuration information you have entered to ensure that the installation completes successfully.

Steps

1. View the **Summary** page.

NetApp® StorageGRID®

Help ▾

Install

1 License 2 Sites 3 Grid Network 4 Grid Nodes 5 NTP 6 DNS 7 Passwords 8 Summary

Summary

Verify that all of the grid configuration information is correct, and then click **Install**. You can view the status of each grid node as it installs. Click the **Modify** links to go back and change the associated information.

General Settings

Grid Name	Grid1	Modify License
Passwords	Auto-generated random command line passwords	Modify Passwords

Networking

NTP	10.60.248.183 10.227.204.142 10.235.48.111	Modify NTP
DNS	10.224.223.130 10.224.223.136	Modify DNS
Grid Network	172.16.0.0/21	Modify Grid Network

Topology

Topology	Atlanta	Modify Sites	Modify Grid Nodes					
	Raleigh	dc1-adm1	dc1-g1	dc1-s1	dc1-s2	dc1-s3	NetApp-SGA	

2. Verify that all of the grid configuration information is correct. Use the **Modify** links on the Summary page to go back and correct any errors.
3. Click **Install**.



If a node is configured to use the Client Network, the default gateway for that node switches from the Grid Network to the Client Network when you click **Install**. If you lose connectivity, you must ensure that you are accessing the primary Admin Node through an accessible subnet. See [Networking guidelines](#) for details.

4. Click **Download Recovery Package**.

When the installation progresses to the point where the grid topology is defined, you are prompted to download the Recovery Package file (.zip), and confirm that you can successfully access the contents of this file. You must download the Recovery Package file so that you can recover the StorageGRID system if one or more grid nodes fail. The installation continues in the background, but you cannot complete the installation and access the StorageGRID system until you download and verify this file.

5. Verify that you can extract the contents of the .zip file, and then save it in two safe, secure, and separate locations.



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

6. Select the **I have successfully downloaded and verified the Recovery Package file** check box, and click **Next**.

Download Recovery Package

Before proceeding, you must download the Recovery Package file. This file is necessary to recover the StorageGRID system if a failure occurs.

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.

After you have verified the contents, copy the Recovery Package file to two safe, secure, and separate locations. The Recovery Package file must be secured because it contains encryption keys and passwords that can be used to obtain data from the StorageGRID system.

The Recovery Package is required for recovery procedures and must be stored in a secure location.

Download Recovery Package

I have successfully downloaded and verified the Recovery Package file.

If the installation is still in progress, the status page appears. This page indicates the progress of the installation for each grid node.

Installation Status

If necessary, you may [Download the Recovery Package file](#) again.

Name	Site	Grid Network IPv4 Address	Progress	Stage
dc1-adm1	Site1	172.16.4.215/21	<div style="width: 100%;"><div style="width: 100%;"></div></div>	Starting services
dc1-g1	Site1	172.16.4.216/21	<div style="width: 100%;"><div style="width: 100%;"></div></div>	Complete
dc1-s1	Site1	172.16.4.217/21	<div style="width: 100%;"><div style="width: 100%;"></div></div>	Waiting for Dynamic IP Service peers
dc1-s2	Site1	172.16.4.218/21	<div style="width: 100%;"><div style="width: 100%;"></div></div>	Downloading hotfix from primary Admin if needed
dc1-s3	Site1	172.16.4.219/21	<div style="width: 100%;"><div style="width: 100%;"></div></div>	Downloading hotfix from primary Admin if needed

When the Complete stage is reached for all grid nodes, the sign-in page for the Grid Manager appears.

7. Sign in to the Grid Manager using the "root" user and the password you specified during the installation.

Post-installation guidelines

After completing grid node deployment and configuration, follow these guidelines for DHCP addressing and network configuration changes.

- If DHCP was used to assign IP addresses, configure a DHCP reservation for each IP address on the networks being used.

You can only set up DHCP during the deployment phase. You cannot set up DHCP during configuration.



Nodes reboot when their IP addresses change, which can cause outages if a DHCP address change affects multiple nodes at the same time.

- You must use the Change IP procedures if you want to change IP addresses, subnet masks, and default gateways for a grid node. See [Configure IP addresses](#).
- If you make networking configuration changes, including routing and gateway changes, client connectivity to the primary Admin Node and other grid nodes might be lost. Depending on the networking changes applied, you might need to re-establish these connections.

Automate the installation (Red Hat Enterprise Linux or CentOS)

You can automate the installation of the StorageGRID host service and the configuration of grid nodes.

Automating the deployment might be useful in any of the following cases:

- You already use a standard orchestration framework, such as Ansible, Puppet, or Chef, to deploy and configure physical or virtual hosts.
- You intend to deploy multiple StorageGRID instances.
- You are deploying a large, complex StorageGRID instance.

The StorageGRID host service is installed by a package and driven by configuration files. You can create the configuration files using one of these methods:

- [Create the configuration files](#) interactively during a manual installation.
- Prepare the configuration files ahead of time (or programmatically) to enable automated installation using standard orchestration frameworks, as described in this article.

StorageGRID provides optional Python scripts for automating the configuration of StorageGRID appliances and the entire StorageGRID system (the “grid”). You can use these scripts directly, or you can inspect them to learn how to use the StorageGRID Installation REST API in grid deployment and configuration tools you develop yourself.

Automate the installation and configuration of the StorageGRID host service

You can automate the installation of the StorageGRID host service using standard orchestration frameworks such as Ansible, Puppet, Chef, Fabric, or SaltStack.

The StorageGRID host service is packaged in an RPM and is driven by configuration files that you can prepare ahead of time (or programmatically) to enable automated installation. If you already use a standard orchestration framework to install and configure RHEL or CentOS, adding StorageGRID to your playbooks or

recipes should be straightforward.

See the example Ansible role and playbook in the `/extras` folder supplied with the installation archive. The Ansible playbook shows how the `storagegrid` role prepares the host and installs StorageGRID onto the target servers. You can customize the role or playbook as necessary.



The example playbook does not include the steps required to create network devices before starting the StorageGRID host service. Add these steps before finalizing and using the playbook.

You can automate all of the steps for preparing the hosts and deploying virtual grid nodes.

Automate the configuration of StorageGRID

After deploying the grid nodes, you can automate the configuration of the StorageGRID system.

What you'll need

- You know the location of the following files from the installation archive.

Filename	Description
<code>configure-storagegrid.py</code>	Python script used to automate the configuration
<code>configure-storagegrid.sample.json</code>	Sample configuration file for use with the script
<code>configure-storagegrid.blank.json</code>	Blank configuration file for use with the script

- You have created a `configure-storagegrid.json` configuration file. To create this file, you can modify the sample configuration file (`configure-storagegrid.sample.json`) or the blank configuration file (`configure-storagegrid.blank.json`).

About this task

You can use the `configure-storagegrid.py` Python script and the `configure-storagegrid.json` configuration file to automate the configuration of your StorageGRID system.



You can also configure the system using the Grid Manager or the Installation API.

Steps

1. Log in to the Linux machine you are using to run the Python script.
2. Change to the directory where you extracted the installation archive.

For example:

```
cd StorageGRID-Webscale-version/platform
```

where `platform` is `debs`, `rpms`, or `vsphere`.

3. Run the Python script and use the configuration file you created.

For example:

```
./configure-storagegrid.py ./configure-storagegrid.json --start-install
```

Result

A Recovery Package `.zip` file is generated during the configuration process, and it is downloaded to the directory where you are running the installation and configuration process. You must back up the Recovery Package file so that you can recover the StorageGRID system if one or more grid nodes fails. For example, copy it to a secure, backed up network location and to a secure cloud storage location.



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

If you specified that random passwords should be generated, you need to extract the `Passwords.txt` file and look for the passwords required to access your StorageGRID system.

```
#####
##### The StorageGRID "recovery package" has been downloaded as: #####
#####           ./sgws-recovery-package-994078-rev1.zip           #####
#####   Safeguard this file as it will be needed in case of a   #####
#####           StorageGRID node recovery.                      #####
#################################################################
```

Your StorageGRID system is installed and configured when a confirmation message is displayed.

```
StorageGRID has been configured and installed.
```

Related information

[Overview of the installation REST API](#)

Overview of the installation REST API

StorageGRID provides the StorageGRID Installation API for performing installation tasks.

The API uses the Swagger open source API platform to provide the API documentation. Swagger allows both developers and non-developers to interact with the API in a user interface that illustrates how the API responds to parameters and options. This documentation assumes that you are familiar with standard web technologies and the JSON (JavaScript Object Notation) data format.



Any API operations you perform using the API Docs webpage are live operations. Be careful not to create, update, or delete configuration data or other data by mistake.

Each REST API command includes the API's URL, an HTTP action, any required or optional URL parameters, and an expected API response.

StorageGRID Installation API

The StorageGRID Installation API is only available when you are initially configuring your StorageGRID system, and in the event that you need to perform a primary Admin Node recovery. The Installation API can be accessed over HTTPS from the Grid Manager.

To access the API documentation, go to the installation web page on the primary Admin Node and select **Help > API Documentation** from the menu bar.

The StorageGRID Installation API includes the following sections:

- **config** — Operations related to the product release and versions of the API. You can list the product release version and the major versions of the API supported by that release.
- **grid** — Grid-level configuration operations. You can get and update grid settings, including grid details, Grid Network subnets, grid passwords, and NTP and DNS server IP addresses.
- **nodes** — Node-level configuration operations. You can retrieve a list of grid nodes, delete a grid node, configure a grid node, view a grid node, and reset a grid node's configuration.
- **provision** — Provisioning operations. You can start the provisioning operation and view the status of the provisioning operation.
- **recovery** — Primary Admin Node recovery operations. You can reset information, upload the Recover Package, start the recovery, and view the status of the recovery operation.
- **recovery-package** — Operations to download the Recovery Package.
- **schemas** — API schemas for advanced deployments
- **sites** — Site-level configuration operations. You can create, view, delete, and modify a site.

Where to go next

After completing an installation, you must perform a series of integration and configuration steps. Some steps are required; others are optional.

Required tasks

- Create a tenant account for each client protocol (Swift or S3) that will be used to store objects on your StorageGRID system.
- Control system access by configuring groups and user accounts. Optionally, you can configure a federated identity source (such as Active Directory or OpenLDAP), so you can import administration groups and users. Or, you can create local groups and users.
- Integrate and test the S3 or Swift API client applications you will use to upload objects to your StorageGRID system.
- When you are ready, configure the information lifecycle management (ILM) rules and ILM policy you want to use to protect object data.



When you install StorageGRID, the default ILM policy, Baseline 2 Copies Policy, is active. This policy includes the stock ILM rule (Make 2 Copies), and it applies if no other policy has been activated.

- If your installation includes appliance Storage Nodes, use SANtricity software to complete the following tasks:

- Connect to each StorageGRID appliance.
- Verify receipt of AutoSupport data.
- If your StorageGRID system includes any Archive Nodes, configure the Archive Node's connection to the target external archival storage system.



If any Archive Nodes will use Tivoli Storage Manager as the external archival storage system, you must also configure Tivoli Storage Manager.

- Review and follow the StorageGRID system hardening guidelines to eliminate security risks.
- Configure email notifications for system alerts.

Optional tasks

- If you want to receive notifications from the (legacy) alarm system, configure mailing lists and email notifications for alarms.
- Update grid node IP addresses if they have changed since you planned your deployment and generated the Recovery Package. See information about changing IP addresses in the recovery and maintenance instructions.
- Configure storage encryption, if required.
- Configure storage compression to reduce the size of stored objects, if required.
- Configure audit client access. You can configure access to the system for auditing purposes through an NFS or a CIFS file share. See the instructions for administering StorageGRID.



Audit export through CIFS/Samba has been deprecated and will be removed in a future StorageGRID release.

Related information

[Administer StorageGRID](#)

[Use S3](#)

[Use Swift](#)

[Manage objects with ILM](#)

[Monitor and troubleshoot](#)

[Recover and maintain](#)

[SG100 and SG1000 services appliances](#)

[SG5600 storage appliances](#)

[SG5700 storage appliances](#)

[SG6000 storage appliances](#)

[Release notes](#)

[System hardening](#)

[Review audit logs](#)

[Upgrade software](#)

Troubleshoot installation issues

If any problems occur while installing your StorageGRID system, you can access the installation log files. Technical support might also need to use the installation log files to resolve issues.

The following installation log files are available from the container that is running each node:

- `/var/local/log/install.log` (found on all grid nodes)
- `/var/local/log/gdu-server.log` (found on the primary Admin Node)

The following installation log files are available from the host:

- `/var/log/storagegrid/daemon.log`
- `/var/log/storagegrid/nodes/node-name.log`

To learn how to access the log files, see the instructions for monitoring and troubleshooting StorageGRID. For help troubleshooting appliance installation issues, see the installation and maintenance instructions for your appliances. If you need additional help, contact technical support.

Related information

[Monitor and troubleshoot](#)

[SG100 and SG1000 services appliances](#)

[SG6000 storage appliances](#)

[SG5700 storage appliances](#)

[SG5600 storage appliances](#)

[NetApp Support](#)

Example `/etc/sysconfig/network-scripts`

You can use the example files to aggregate four Linux physical interfaces into a single LACP bond and then establish three VLAN interfaces subtending the bond for use as StorageGRID Grid, Admin, and Client Network interfaces.

Physical interfaces

Note that the switches at the other ends of the links must also treat the four ports as a single LACP trunk or port channel, and must pass at least the three referenced VLANs with tags.

`/etc/sysconfig/network-scripts/ifcfg-ens160`

```
TYPE=Ethernet
NAME=ens160
UUID=011b17dd-642a-4bb9-acae-d71f7e6c8720
DEVICE=ens160
ONBOOT=yes
MASTER=bond0
SLAVE=yes
```

/etc/sysconfig/network-scripts/ifcfg-ens192

```
TYPE=Ethernet
NAME=ens192
UUID=e28eb15f-76de-4e5f-9a01-c9200b58d19c
DEVICE=ens192
ONBOOT=yes
MASTER=bond0
SLAVE=yes
```

/etc/sysconfig/network-scripts/ifcfg-ens224

```
TYPE=Ethernet
NAME=ens224
UUID=b0e3d3ef-7472-4cde-902c-ef4f3248044b
DEVICE=ens224
ONBOOT=yes
MASTER=bond0
SLAVE=yes
```

/etc/sysconfig/network-scripts/ifcfg-ens256

```
TYPE=Ethernet
NAME=ens256
UUID=7cf7aabc-3e4b-43d0-809a-1e2378faa4cd
DEVICE=ens256
ONBOOT=yes
MASTER=bond0
SLAVE=yes
```

Bond interface

/etc/sysconfig/network-scripts/ifcfg-bond0

```
DEVICE=bond0
TYPE=Bond
BONDING_MASTER=yes
NAME=bond0
ONBOOT=yes
BONDING_OPTS=mode=802.3ad
```

VLAN interfaces

/etc/sysconfig/network-scripts/ifcfg-bond0.1001

```
VLAN=yes
TYPE=Vlan
DEVICE=bond0.1001
PHYSDEV=bond0
VLAN_ID=1001
REORDER_HDR=0
BOOTPROTO=none
UUID=296435de-8282-413b-8d33-c4dd40fca24a
ONBOOT=yes
```

/etc/sysconfig/network-scripts/ifcfg-bond0.1002

```
VLAN=yes
TYPE=Vlan
DEVICE=bond0.1002
PHYSDEV=bond0
VLAN_ID=1002
REORDER_HDR=0
BOOTPROTO=none
UUID=dbaaec72-0690-491c-973a-57b7dd00c581
ONBOOT=yes
```

/etc/sysconfig/network-scripts/ifcfg-bond0.1003

```
VLAN=yes
TYPE=Vlan
DEVICE=bond0.1003
PHYSDEV=bond0
VLAN_ID=1003
REORDER_HDR=0
BOOTPROTO=none
UUID=d1af4b30-32f5-40b4-8bb9-71a2fbf809a1
ONBOOT=yes
```

Install Ubuntu or Debian

Install Ubuntu or Debian: Overview

Installing a StorageGRID system in an Ubuntu or Debian environment includes three primary steps.

- 1. Preparation:** During planning and preparation, you perform the following tasks:
 - Learn about the hardware and storage requirements for StorageGRID.
 - Learn about the specifics of [StorageGRID networking](#) so you can configure your network appropriately.
 - Identify and prepare the physical or virtual servers you plan to use to host your StorageGRID grid nodes.
 - On the servers you have prepared:
 - Install Linux
 - Configure the host network
 - Configure host storage
 - Install Docker
 - Install the StorageGRID host services
- 2. Deployment:** Deploy grid nodes using the appropriate user interface. When you deploy grid nodes, they are created as part of the StorageGRID system and connected to one or more networks.
 - a. Use the Linux command line and node configuration files to deploy virtual grid nodes on the hosts you prepared in step 1.
 - b. Use the StorageGRID Appliance Installer to deploy StorageGRID appliance nodes.



Hardware-specific installation and integration instructions are not included in the StorageGRID installation procedure. To learn how to install StorageGRID appliances, see the installation and maintenance instructions for your appliance.

- 3. Configuration:** When all nodes have been deployed, use the Grid Manager to configure the grid and complete the installation.

These instructions recommend a standard approach for deploying and configuring a StorageGRID system in an Ubuntu or Debian environment. See also the information about the following alternative approaches:

- Use a standard orchestration framework such as Ansible, Puppet, or Chef to install Ubuntu or Debian, configure networking and storage, install Docker and the StorageGRID host service, and deploy virtual grid nodes.
- Automate the deployment and configuration of the StorageGRID system using a Python configuration script (provided in the installation archive).
- Automate the deployment and configuration of appliance grid nodes with a Python configuration script (available from the installation archive or from the StorageGRID Appliance Installer).
- If you are an advanced developer of StorageGRID deployments, use the installation REST APIs to automate the installation of StorageGRID grid nodes.

Plan and prepare for Ubuntu or Debian installation

Before you install (Ubuntu or Debian)

Before deploying grid nodes and configuring the StorageGRID grid, you must be familiar with the steps and requirements for completing the procedure.

The StorageGRID deployment and configuration procedures assume that you are familiar with the architecture and operation of the StorageGRID system.

You can deploy a single site or multiple sites at one time; however, all sites must meet the minimum requirement of having at least three Storage Nodes.

Before starting a StorageGRID installation, you must:

- Understand StorageGRID's compute requirements, including the minimum CPU and RAM requirements for each node.
- Understand how StorageGRID supports multiple networks for traffic separation, security, and administrative convenience, and have a plan for which networks you intend to attach to each StorageGRID node.

See the StorageGRID networking guidelines.

- Understand the storage and performance requirements of each type of grid node.
- Identify a set of servers (physical, virtual, or both) that, in aggregate, provide sufficient resources to support the number and type of StorageGRID nodes you plan to deploy.
- Understand the requirements for node migration, if you want to perform scheduled maintenance on physical hosts without any service interruption.
- Gather all networking information in advance. Unless you are using DHCP, gather the IP addresses to assign to each grid node, and the IP addresses of the domain name system (DNS) and network time protocol (NTP) servers that will be used.
- Install, connect, and configure all required hardware, including any StorageGRID appliances, to specifications.



Hardware-specific installation and integration instructions are not included in the StorageGRID installation procedure. To learn how to install StorageGRID appliances, see the installation and maintenance instructions for your appliance.

- Decide which of the available deployment and configuration tools you want to use.

Related information

Networking guidelines

[SG100 and SG1000 services appliances](#)

[SG6000 storage appliances](#)

[SG5700 storage appliances](#)

[SG5600 storage appliances](#)

[Node container migration requirements](#)

Required materials

Before you install StorageGRID, you must gather and prepare required materials.

Item	Notes
NetApp StorageGRID license	You must have a valid, digitally signed NetApp license. Note: A non-production license, which can be used for testing and proof of concept grids, is included in the StorageGRID installation archive.
StorageGRID installation archive	You must download the StorageGRID installation archive and extract the files.
Service laptop	The StorageGRID system is installed through a service laptop. The service laptop must have: <ul style="list-style-type: none">• Network port• SSH client (for example, PuTTY)• Supported web browser
StorageGRID documentation	<ul style="list-style-type: none">• Release notes• Instructions for administering StorageGRID

Related information

[NetApp Interoperability Matrix Tool](#)

Download and extract the StorageGRID installation files

You must download the StorageGRID installation archive and extract the required files.

Steps

1. Go to the [NetApp Downloads page for StorageGRID](#).
2. Select the button for downloading the latest release, or select another version from the drop-down menu

and select **Go**.

3. Sign in with the username and password for your NetApp account.
4. If a Caution/MustRead statement appears, read it and select the check box.



You must apply any required hotfixes after you install the StorageGRID release. For more information, see the [hotfix procedure in the recovery and maintenance instructions](#)

5. Read the End User License Agreement, select the check box, and then select **Accept & Continue**.

The downloads page for the version you selected appears. The page contains three columns:

6. In the **Install StorageGRID** column, select the .tgz or .zip file for Ubuntu or Debian.



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

7. Save and extract the archive file.
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 grid.



The paths listed in the table are relative to the top-level directory installed by the extracted installation archive.

Path and file name	Description
./debs/README	A text file that describes all of the files contained in the StorageGRID download file.
./debs/NLF000000.txt	A non-production NetApp License File that you can use for testing and proof of concept deployments.
./debs/storagegrid-webscale-images-version-SHA.deb	DEB package for installing the StorageGRID node images on Ubuntu or Debian hosts.
./debs/storagegrid-webscale-images-version-SHA.deb.md5	MD5 checksum for the file /debs/storagegrid-webscale-images-version-SHA.deb.
./debs/storagegrid-webscale-service-version-SHA.deb	DEB package for installing the StorageGRID host service on Ubuntu or Debian hosts.
Deployment scripting tool	Description
./debs/configure-storagegrid.py	A Python script used to automate the configuration of a StorageGRID system.
./debs/configure-sga.py	A Python script used to automate the configuration of StorageGRID appliances.

Path and file name	Description
./debs/storagegrid-ssoauth.py	An example Python script that you can use to sign in to the Grid Management API when single sign-on is enabled.
./debs/configure-storagegrid.sample.json	An example configuration file for use with the configure-storagegrid.py script.
./debs/configure-storagegrid.blank.json	A blank configuration file for use with the configure-storagegrid.py script.
./debs/extras/ansible	Example Ansible role and playbook for configuring Ubuntu or Debian hosts for StorageGRID container deployment. You can customize the role or playbook as necessary.
./debs/extras/api-schemas	<p>API schemas for StorageGRID.</p> <p>Note: Before you perform an upgrade, you can use these schemas to confirm that any code you have written to use StorageGRID management APIs will be compatible with the new StorageGRID release if you do not have a non-production StorageGRID environment for upgrade compatibility testing.</p>

Related information

[Recover and maintain](#)

CPU and RAM requirements

Before installing StorageGRID software, verify and configure the hardware so that it is ready to support the StorageGRID system.

For information about supported servers, see the Interoperability Matrix.

Each StorageGRID node requires the following minimum resources:

- CPU cores: 8 per node
- RAM: At least 24 GB per node, and 2 to 16 GB less than the total system RAM, depending on the total RAM available and the amount of non-StorageGRID software running on the system

Ensure that the number of StorageGRID nodes you plan to run on each physical or virtual host does not exceed the number of CPU cores or the physical RAM available. If the hosts are not dedicated to running StorageGRID (not recommended), be sure to consider the resource requirements of the other applications.



Monitor your CPU and memory usage regularly to ensure that these resources continue to accommodate your workload. For example, doubling the RAM and CPU allocation for virtual Storage Nodes would provide similar resources to those provided for StorageGRID appliance nodes. Additionally, if the amount of metadata per node exceeds 500 GB, consider increasing the RAM per node to 48 GB or more. For information about managing object metadata storage, increasing the Metadata Reserved Space setting, and monitoring CPU and memory usage, see the instructions for administering, monitoring, and upgrading StorageGRID.

If hyperthreading is enabled on the underlying physical hosts, you can provide 8 virtual cores (4 physical cores) per node. If hyperthreading is not enabled on the underlying physical hosts, you must provide 8 physical cores per node.

If you are using virtual machines as hosts and have control over the size and number of VMs, you should use a single VM for each StorageGRID node and size the VM accordingly.

For production deployments, you should not run multiple Storage Nodes on the same physical storage hardware or virtual host. Each Storage Node in a single StorageGRID deployment should be in its own isolated failure domain. You can maximize the durability and availability of object data if you ensure that a single hardware failure can only impact a single Storage Node.

See also the information about storage requirements.

Related information

[NetApp Interoperability Matrix Tool](#)

[Storage and performance requirements](#)

[Administer StorageGRID](#)

[Monitor and troubleshoot](#)

[Upgrade software](#)

Storage and performance requirements

You must understand the storage requirements for StorageGRID nodes, so you can provide enough space to support the initial configuration and future storage expansion.

StorageGRID nodes require three logical categories of storage:

- **Container pool** — Performance-tier (10K SAS or SSD) storage for the node containers, which will be assigned to the Docker storage driver when you install and configure Docker on the hosts that will support your StorageGRID nodes.
- **System data** — Performance-tier (10K SAS or SSD) storage for per-node persistent storage of system data and transaction logs, which the StorageGRID host services will consume and map into individual nodes.
- **Object data** — Performance-tier (10K SAS or SSD) storage and capacity-tier (NL-SAS/SATA) bulk storage for the persistent storage of object data and object metadata.

You must use RAID-backed block devices for all storage categories. Non-redundant disks, SSDs, or JBODs are not supported. You can use shared or local RAID storage for any of the storage categories; however, if you want to use StorageGRID's node migration capability, you must store both system data and object data on shared storage.

Performance requirements

The performance of the volumes used for the container pool, system data, and object metadata significantly impacts the overall performance of the system. You should use performance-tier (10K SAS or SSD) storage for these volumes to ensure adequate disk performance in terms of latency, input/output operations per second (IOPS), and throughput. You can use capacity-tier (NL-SAS/SATA) storage for the persistent storage of object data.

The volumes used for the container pool, system data, and object data must have write-back caching enabled. The cache must be on a protected or persistent media.

Requirements for hosts that use NetApp ONTAP storage

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

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

Number of hosts required

Each StorageGRID site requires a minimum of three Storage Nodes.

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

Other types of nodes, such as Admin Nodes or Gateway Nodes, can be deployed on the same hosts, or they can be deployed on their own dedicated hosts as required.

Number of storage volumes for each host

The following table shows the number of storage volumes (LUNs) required for each host and the minimum size required for each LUN, based on which nodes will be deployed on that host.

The maximum tested LUN size is 39 TB.

 These numbers are for each host, not for the entire grid.

LUN purpose	Storage category	Number of LUNs	Minimum size/LUN
Container engine storage pool	Container pool	1	Total number of nodes × 100 GB
/var/local volume	System data	1 for each node on this host	90 GB

LUN purpose	Storage category	Number of LUNs	Minimum size/LUN
Storage Node	Object data	3 for each Storage Node on this host Note: A software-based Storage Node can have 1 to 16 storage volumes; at least 3 storage volumes are recommended.	12 TB (4 TB/LUN) See storage requirements for Storage Nodes for more information.
Admin Node audit logs	System data	1 for each Admin Node on this host	200 GB
Admin Node tables	System data	1 for each Admin Node on this host	200 GB

 Depending on the audit level configured, the size of user inputs such as S3 object key name, and how much audit log data you need to preserve, you might need to increase the size of the audit log LUN on each Admin Node. As a general rule, a grid generates approximately 1 KB of audit data per S3 operation, which would mean that a 200 GB LUN would support 70 million operations per day or 800 operations per second for two to three days.

Minimum storage space for a host

The following table shows the minimum storage space required for each type of node. You can use this table to determine the minimum amount of storage you must provide to the host in each storage category, based on which nodes will be deployed on that host.

 Disk snapshots cannot be used to restore grid nodes. Instead, refer to the recovery and maintenance procedures for each type of node.

Type of node	Container pool	System data	Object data
Storage Node	100 GB	90 GB	4,000 GB
Admin Node	100 GB	490 GB (3 LUNs)	<i>not applicable</i>
Gateway Node	100 GB	90 GB	<i>not applicable</i>
Archive Node	100 GB	90 GB	<i>not applicable</i>

Example: Calculating the storage requirements for a host

Suppose you plan to deploy three nodes on the same host: one Storage Node, one Admin Node, and one Gateway Node. You should provide a minimum of nine storage volumes to the host. You will need a minimum of 300 GB of performance-tier storage for the node containers, 670 GB of performance-tier storage for system data and transaction logs, and 12 TB of capacity-tier storage for object data.

Type of node	LUN purpose	Number of LUNs	LUN size
Storage Node	Docker storage pool	1	300 GB (100 GB/node)
Storage Node	/var/local volume	1	90 GB
Storage Node	Object data	3	12 TB (4 TB/LUN)
Admin Node	/var/local volume	1	90 GB
Admin Node	Admin Node audit logs	1	200 GB
Admin Node	Admin Node tables	1	200 GB
Gateway Node	/var/local volume	1	90 GB
Total		9	Container pool: 300 GB System data: 670 GB Object data: 12,000 GB

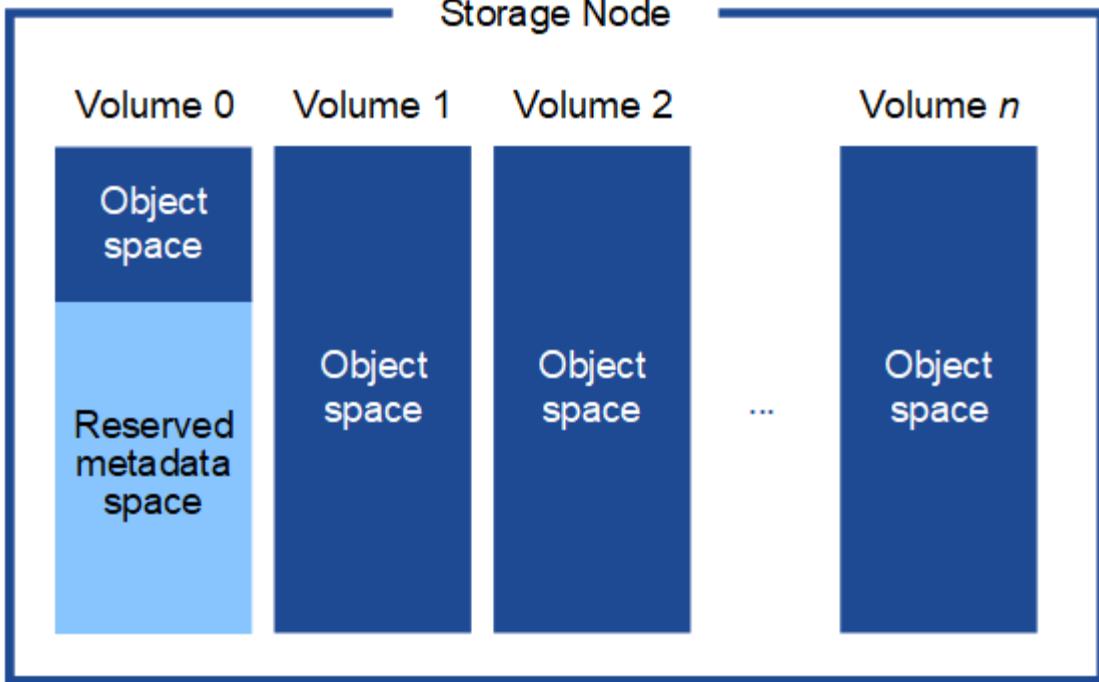
Storage requirements for Storage Nodes

A software-based Storage Node can have 1 to 16 storage volumes—3 or more storage volumes are recommended. Each storage volume should be 4 TB or larger.



An appliance Storage Node can have up to 48 storage volumes.

As shown in the figure, StorageGRID reserves space for object metadata on storage volume 0 of each Storage Node. Any remaining space on storage volume 0 and any other storage volumes in the Storage Node are used exclusively for object data.



To provide redundancy and to protect object metadata from loss, StorageGRID stores three copies of the metadata for all objects in the system at each site. The three copies of object metadata are evenly distributed across all Storage Nodes at each site.

When you assign space to volume 0 of a new Storage Node, you must ensure there is adequate space for that node's portion of all object metadata.

- At a minimum, you must assign at least 4 TB to volume 0.



If you use only one storage volume for a Storage Node and you assign 4 TB or less to the volume, the Storage Node might enter the Storage Read-Only state on startup and store object metadata only.

- If you are installing a new StorageGRID 11.6 system and each Storage Node has 128 GB or more of RAM, you should assign 8 TB or more to volume 0. Using a larger value for volume 0 can increase the space allowed for metadata on each Storage Node.
- When configuring different Storage Nodes for a site, use the same setting for volume 0 if possible. If a site contains Storage Nodes of different sizes, the Storage Node with the smallest volume 0 will determine the metadata capacity of that site.

For details, go to [Manage object metadata storage](#).

Related information

[Node container migration requirements](#)

[Recover and maintain](#)

Node container migration requirements

The node migration feature allows you to manually move a node from one host to another. Typically, both hosts are in the same physical data center.

Node migration allows you to perform physical host maintenance without disrupting grid operations. You simply move all StorageGRID nodes, one at a time, to another host before taking the physical host offline. Migrating nodes requires only a short downtime for each node and should not affect operation or availability of grid services.

If you want to use the StorageGRID node migration feature, your deployment must meet additional requirements:

- Consistent network interface names across hosts in a single physical data center
- Shared storage for StorageGRID metadata and object repository volumes that is accessible by all hosts in a single physical data center. For example, you might use NetApp E-Series storage arrays.

If you are using virtual hosts and the underlying hypervisor layer supports VM migration, you might want to use this capability instead of StorageGRID's node migration feature. In this case, you can ignore these additional requirements.

Before performing migration or hypervisor maintenance, shut down the nodes gracefully. See the instructions for [shutting down a grid node](#).

VMware Live Migration not supported

OpenStack Live Migration and VMware live vMotion cause the virtual machine clock time to jump and are not supported for grid nodes of any type. Though rare, incorrect clock times can result in loss of data or configuration updates.

Cold migration is supported. In cold migration, you shut down the StorageGRID nodes before migrating them between hosts. See the instructions for [shutting down a grid node](#).

Consistent network interface names

In order to move a node from one host to another, the StorageGRID host service needs to have some confidence that the external network connectivity the node has at its current location can be duplicated at the new location. It gets this confidence through the use of consistent network interface names in the hosts.

Suppose, for example, that StorageGRID NodeA running on Host1 has been configured with the following interface mappings:

eth0 → bond0.1001

eth1 → bond0.1002

eth2 → bond0.1003

The lefthand side of the arrows corresponds to the traditional interfaces as viewed from within a StorageGRID container (that is, the Grid, Admin, and Client Network interfaces, respectively). The righthand side of the arrows corresponds to the actual host interfaces providing these networks, which are three VLAN interfaces subordinate to the same physical interface bond.

Now, suppose you want to migrate NodeA to Host2. If Host2 also has interfaces named bond0.1001, bond0.1002, and bond0.1003, the system will allow the move, assuming that the like-named interfaces will provide the same connectivity on Host2 as they do on Host1. If Host2 does not have interfaces with the same names, the move will not be allowed.

There are many ways to achieve consistent network interface naming across multiple hosts; see [Configure the host network](#) for some examples.

Shared storage

In order to achieve rapid, low-overhead node migrations, the StorageGRID node migration feature does not physically move node data. Instead, node migration is performed as a pair of export and import operations, as follows:

Steps

1. During the “node export” operation, a small amount of persistent state data is extracted from the node container running on HostA and cached on that node’s system data volume. Then, the node container on HostA is deinstantiated.
2. During the “node import” operation, the node container on HostB that uses the same network interface and block storage mappings that were in effect on HostA is instantiated. Then, the cached persistent state data is inserted into the new instance.

Given this mode of operation, all of the node’s system data and object storage volumes must be accessible from both HostA and HostB for the migration to be allowed, and to work. In addition, they must have been mapped into the node using names that are guaranteed to refer to the same LUNs on HostA and HostB.

The following example shows one solution for block device mapping for a StorageGRID Storage Node, where DM multipathing is in use on the hosts, and the alias field has been used in `/etc/multipath.conf` to provide consistent, friendly block device names available on all hosts.

<code>/var/local</code>	 <code>/dev/mapper/sgws-sn1-var-local</code>
<code>rangedb0</code>	 <code>/dev/mapper/sgws-sn1-rangedb0</code>
<code>rangedb1</code>	 <code>/dev/mapper/sgws-sn1-rangedb1</code>
<code>rangedb2</code>	 <code>/dev/mapper/sgws-sn1-rangedb2</code>
<code>rangedb3</code>	 <code>/dev/mapper/sgws-sn1-rangedb3</code>

Deployment tools

You might benefit from automating all or part of the StorageGRID installation.

Automating the deployment might be useful in any of the following cases:

- You already use a standard orchestration framework, such as Ansible, Puppet, or Chef, to deploy and configure physical or virtual hosts.
- You intend to deploy multiple StorageGRID instances.
- You are deploying a large, complex StorageGRID instance.

The StorageGRID host service is installed by a package and driven by configuration files that can be created interactively during a manual installation, or prepared ahead of time (or programmatically) to enable automated

installation using standard orchestration frameworks. StorageGRID provides optional Python scripts for automating the configuration of StorageGRID appliances, and the whole StorageGRID system (the “grid”). You can use these scripts directly, or you can inspect them to learn how to use the StorageGRID Installation REST API in grid deployment and configuration tools you develop yourself.

If you are interested in automating all or part of your StorageGRID deployment, review [Automate the installation](#) before beginning the installation process.

Prepare the hosts (Ubuntu or Debian)

Install Linux

You must install Linux on all grid hosts. Use the [NetApp Interoperability Matrix Tool](#) to get a list of supported versions.

Steps

1. Install Linux on all physical or virtual grid hosts according to the distributor’s instructions or your standard procedure.



Do not install any graphical desktop environments. When installing Ubuntu, you must select **standard system utilities**. Selecting **OpenSSH server** is recommended to enable ssh access to your Ubuntu hosts. All other options can remain unselected.

2. Ensure that all hosts have access to Ubuntu or Debian package repositories.

3. If swap is enabled:

- a. Run the following command: `$ sudo swapoff --all`

- b. Remove all swap entries from `/etc/fstab` to persist the settings.



Failing to disable swap entirely can severely lower performance.

Understand AppArmor profile installation

If you are operating in a self-deployed Ubuntu environment and using the AppArmor mandatory access control system, the AppArmor profiles associated with packages you install on the base system might be blocked by the corresponding packages installed with StorageGRID.

By default, AppArmor profiles are installed for packages that you install on the base operating system. When you run these packages from the StorageGRID system container, the AppArmor profiles are blocked. The DHCP, MySQL, NTP, and tcdump base packages conflict with AppArmor, and other base packages might also conflict.

You have two choices for handling AppArmor profiles:

- Disable individual profiles for the packages installed on the base system that overlap with the packages in the StorageGRID system container. When you disable individual profiles, an entry appears in the StorageGRID log files indicating that AppArmor is enabled.

Use the following commands:

```
sudo ln -s /etc/apparmor.d/<profile.name> /etc/apparmor.d/disable/  
sudo apparmor_parser -R /etc/apparmor.d/<profile.name>
```

Example:

```
sudo ln -s /etc/apparmor.d/bin.ping /etc/apparmor.d/disable/  
sudo apparmor_parser -R /etc/apparmor.d/bin.ping
```

- Disable AppArmor altogether. For Ubuntu 9.10 or later, follow the instructions in the Ubuntu online community: [Disable AppArmor](#).

Once you disable AppArmor, no entries indicating that AppArmor is enabled will appear in the StorageGRID log files.

Configure the host network (Ubuntu or Debian)

After completing the Linux installation on your hosts, you might need to perform some additional configuration to prepare a set of network interfaces on each host that are suitable for mapping into the StorageGRID nodes you will deploy later.

What you'll need

- You have reviewed the [StorageGRID networking guidelines](#).
- You have reviewed the information about [node container migration requirements](#).
- If you are using virtual hosts, you have read the [considerations and recommendations for MAC address cloning](#) before configuring the host network.



If you are using VMs as hosts, you should select VMXNET 3 as the virtual network adapter. The VMware E1000 network adapter has caused connectivity issues with StorageGRID containers deployed on certain distributions of Linux.

About this task

Grid nodes must be able to access the Grid Network and, optionally, the Admin and Client Networks. You provide this access by creating mappings that associate the host's physical interface to the virtual interfaces for each grid node. When creating host interfaces, use friendly names to facilitate deployment across all hosts, and to enable migration.

The same interface can be shared between the host and one or more nodes. For example, you might use the same interface for host access and node Admin Network access, to facilitate host and node maintenance. Although the same interface can be shared between the host and individual nodes, all must have different IP addresses. IP addresses cannot be shared between nodes or between the host and any node.

You can use the same host network interface to provide the Grid Network interface for all StorageGRID nodes on the host; you can use a different host network interface for each node; or you can do something in between. However, you would not typically provide the same host network interface as both the Grid and Admin Network interfaces for a single node, or as the Grid Network interface for one node and the Client Network interface for another.

You can complete this task in many ways. For example, if your hosts are virtual machines and you are

deploying one or two StorageGRID nodes for each host, you can simply create the correct number of network interfaces in the hypervisor, and use a 1-to-1 mapping. If you are deploying multiple nodes on bare metal hosts for production use, you can leverage the Linux networking stack's support for VLAN and LACP for fault tolerance and bandwidth sharing. The following sections provide detailed approaches for both of these examples. You do not need to use either of these examples; you can use any approach that meets your needs.

 Do not use bond or bridge devices directly as the container network interface. Doing so could prevent node start-up caused by a kernel issue with the use of MACVLAN with bond and bridge devices in the container namespace. Instead, use a non-bond device, such as a VLAN or virtual Ethernet (veth) pair. Specify this device as the network interface in the node configuration file.

Considerations and recommendations for MAC address cloning

MAC address cloning causes the container to use the MAC address of the host, and the host to use the MAC address of either an address you specify or a randomly generated one. You should use MAC address cloning to avoid the use of promiscuous mode network configurations.

Enabling MAC cloning

In certain environments, security can be enhanced through MAC address cloning because it enables you to use a dedicated virtual NIC for the Admin Network, Grid Network, and Client Network. Having the container use the MAC address of the dedicated NIC on the host allows you to avoid using promiscuous mode network configurations.

 MAC address cloning is intended to be used with virtual server installations and might not function properly with all physical appliance configurations.

 If a node fails to start due to a MAC cloning targeted interface being busy, you might need to set the link to "down" before starting node. Additionally, it is possible that the virtual environment might prevent MAC cloning on a network interface while the link is up. If a node fails to set the MAC address and start due to an interface being busy, setting the link to "down" before starting the node might fix the issue.

MAC address cloning is disabled by default and must be set by node configuration keys. You should enable it when you install StorageGRID.

There is one key for each network:

- `ADMIN_NETWORK_TARGET_TYPE_INTERFACE_CLONE_MAC`
- `GRID_NETWORK_TARGET_TYPE_INTERFACE_CLONE_MAC`
- `CLIENT_NETWORK_TARGET_TYPE_INTERFACE_CLONE_MAC`

Setting the key to "true" causes the container to use the MAC address of the host's NIC. Additionally, the host will then use the MAC address of the specified container network. By default, the container address is a randomly generated address, but if you have set one using the `_NETWORK_MAC` node configuration key, that address is used instead. The host and container will always have different MAC addresses.

 Enabling MAC cloning on a virtual host without also enabling promiscuous mode on the hypervisor might cause Linux host networking using the host's interface to stop working.

MAC cloning use cases

There are two use cases to consider with MAC cloning:

- MAC cloning not enabled: When the `_CLONE_MAC` key in the node configuration file is not set, or set to "false," the host will use the host NIC MAC and the container will have a StorageGRID-generated MAC unless a MAC is specified in the `_NETWORK_MAC` key. If an address is set in the `_NETWORK_MAC` key, the container will have the address specified in the `_NETWORK_MAC` key. This configuration of keys requires the use of promiscuous mode.
- MAC cloning enabled: When the `_CLONE_MAC` key in the node configuration file is set to "true," the container uses the host NIC MAC, and the host uses a StorageGRID-generated MAC unless a MAC is specified in the `_NETWORK_MAC` key. If an address is set in the `_NETWORK_MAC` key, the host uses the specified address instead of a generated one. In this configuration of keys, you should not use promiscuous mode.

 If you do not want to use MAC address cloning and would rather allow all interfaces to receive and transmit data for MAC addresses other than the ones assigned by the hypervisor, ensure that the security properties at the virtual switch and port group levels are set to **Accept** for Promiscuous Mode, MAC Address Changes, and Forged Transmits. The values set on the virtual switch can be overridden by the values at the port group level, so ensure that settings are the same in both places.

To enable MAC cloning, see the [instructions for creating node configuration files](#).

MAC cloning example

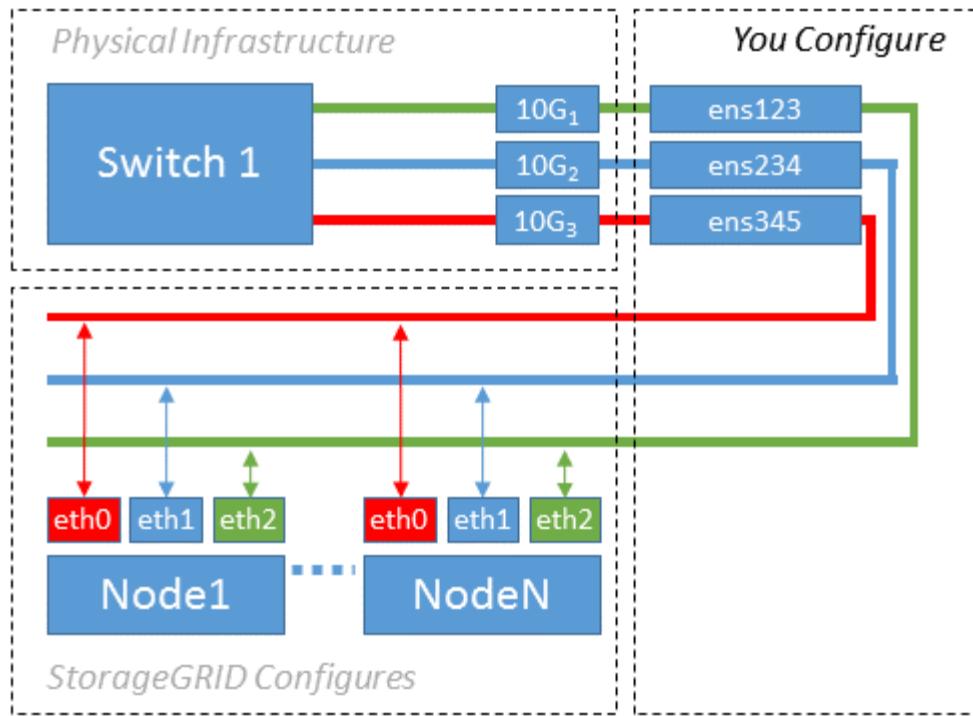
Example of MAC cloning enabled with a host having MAC address of 11:22:33:44:55:66 for the interface ens256 and the following keys in the node configuration file:

- `ADMIN_NETWORK_TARGET = ens256`
- `ADMIN_NETWORK_MAC = b2:9c:02:c2:27:10`
- `ADMIN_NETWORK_TARGET_TYPE_INTERFACE_CLONE_MAC = true`

Result: the host MAC for ens256 is b2:9c:02:c2:27:10 and the Admin Network MAC is 11:22:33:44:55:66

Example 1: 1-to-1 mapping to physical or virtual NICs

Example 1 describes a simple physical interface mapping that requires little or no host-side configuration.



The Linux operating system creates the ensXYZ interfaces automatically during installation or boot, or when the interfaces are hot-added. No configuration is required other than ensuring that the interfaces are set to come up automatically after boot. You do have to determine which ensXYZ corresponds to which StorageGRID network (Grid, Admin, or Client) so you can provide the correct mappings later in the configuration process.

Note that the figure shows multiple StorageGRID nodes; however, you would normally use this configuration for single-node VMs.

If Switch 1 is a physical switch, you should configure the ports connected to interfaces 10G₁ through 10G₃ for access mode, and place them on the appropriate VLANs.

Example 2: LACP bond carrying VLANs

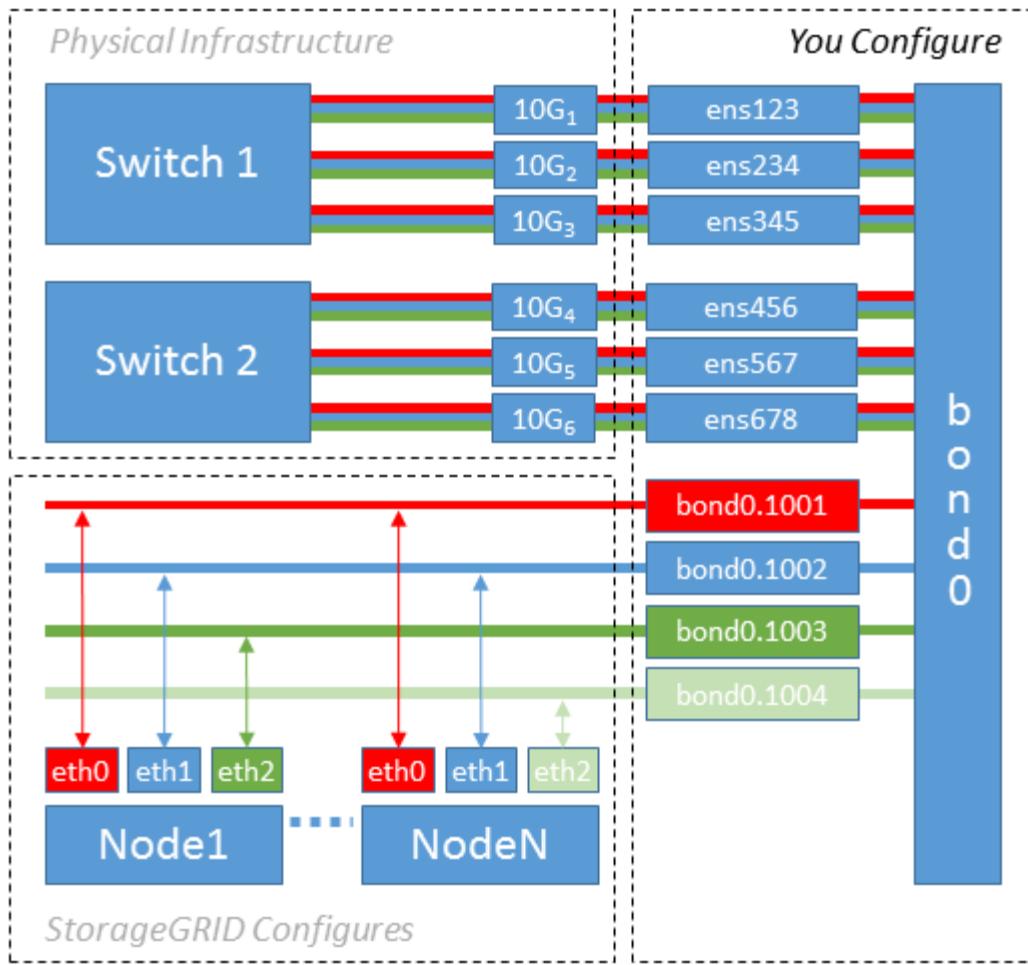
Example 2 assumes you are familiar with bonding network interfaces and with creating VLAN interfaces on the Linux distribution you are using.

About this task

Example 2 describes a generic, flexible, VLAN-based scheme that facilitates the sharing of all available network bandwidth across all nodes on a single host. This example is particularly applicable to bare metal hosts.

To understand this example, suppose you have three separate subnets for the Grid, Admin, and Client Networks at each data center. The subnets are on separate VLANs (1001, 1002, and 1003) and are presented to the host on a LACP-bonded trunk port (bond0). You would configure three VLAN interfaces on the bond: bond0.1001, bond0.1002, and bond0.1003.

If you require separate VLANs and subnets for node networks on the same host, you can add VLAN interfaces on the bond and map them into the host (shown as bond0.1004 in the illustration).



Steps

1. Aggregate all physical network interfaces that will be used for StorageGRID network connectivity into a single LACP bond.

Use the same name for the bond on every host, for example, bond0.

2. Create VLAN interfaces that use this bond as their associated “physical device,” using the standard VLAN interface naming convention `physdev-name.VLAN_ID`.

Note that steps 1 and 2 require appropriate configuration on the edge switches terminating the other ends of the network links. The edge switch ports must also be aggregated into a LACP port channel, configured as a trunk, and allowed to pass all required VLANs.

Sample interface configuration files for this per-host networking configuration scheme are provided.

Related information

[Example /etc/network/interfaces](#)

Configure host storage

You must allocate block storage volumes to each host.

What you'll need

You have reviewed the following topics, which provide information you need to accomplish this task:

Storage and performance requirements

Node container migration requirements

About this task

When allocating block storage volumes (LUNs) to hosts, use the tables in “Storage requirements” to determine the following:

- Number of volumes required for each host (based on the number and types of nodes that will be deployed on that host)
- Storage category for each volume (that is, System Data or Object Data)
- Size of each volume

You will use this information as well as the persistent name assigned by Linux to each physical volume when you deploy StorageGRID nodes on the host.



You do not need to partition, format, or mount any of these volumes; you just need to ensure they are visible to the hosts.

Avoid using “raw” special device files (`/dev/sdb`, for example) as you compose your list of volume names. These files can change across reboots of the host, which will impact proper operation of the system. If you are using iSCSI LUNs and device mapper multipathing, consider using multipath aliases in the `/dev/mapper` directory, especially if your SAN topology includes redundant network paths to the shared storage.

Alternatively, you can use the system-created softlinks under `/dev/disk/by-path/` for your persistent device names.

For example:

```
ls -l
$ ls -l /dev/disk/by-path/
total 0
lrwxrwxrwx 1 root root 9 Sep 19 18:53 pci-0000:00:07.1-ata-2 -> ../../sr0
lrwxrwxrwx 1 root root 9 Sep 19 18:53 pci-0000:03:00.0-scsi-0:0:0:0 ->
 ../../sda
lrwxrwxrwx 1 root root 10 Sep 19 18:53 pci-0000:03:00.0-scsi-0:0:0:0-part1
-> ../../sda1
lrwxrwxrwx 1 root root 10 Sep 19 18:53 pci-0000:03:00.0-scsi-0:0:0:0-part2
-> ../../sda2
lrwxrwxrwx 1 root root 9 Sep 19 18:53 pci-0000:03:00.0-scsi-0:0:1:0 ->
 ../../sdb
lrwxrwxrwx 1 root root 9 Sep 19 18:53 pci-0000:03:00.0-scsi-0:0:2:0 ->
 ../../sdc
lrwxrwxrwx 1 root root 9 Sep 19 18:53 pci-0000:03:00.0-scsi-0:0:3:0 ->
 ../../sdd
```

Results will differ for each installation.

Assign friendly names to each of these block storage volumes to simplify the initial StorageGRID installation and future maintenance procedures. If you are using the device mapper multipath driver for redundant access

to shared storage volumes, you can use the `alias` field in your `/etc/multipath.conf` file.

For example:

```
multipaths {
    multipath {
        wwid 3600a09800059d6df00005df2573c2c30
        alias docker-storage-volume-hostA
    }
    multipath {
        wwid 3600a09800059d6df00005df3573c2c30
        alias sgws-adm1-var-local
    }
    multipath {
        wwid 3600a09800059d6df00005df4573c2c30
        alias sgws-adm1-audit-logs
    }
    multipath {
        wwid 3600a09800059d6df00005df5573c2c30
        alias sgws-adm1-tables
    }
    multipath {
        wwid 3600a09800059d6df00005df6573c2c30
        alias sgws-gw1-var-local
    }
    multipath {
        wwid 3600a09800059d6df00005df7573c2c30
        alias sgws-sn1-var-local
    }
    multipath {
        wwid 3600a09800059d6df00005df7573c2c30
        alias sgws-sn1-rangedb-0
    }
    ...
}
```

This will cause the aliases to appear as block devices in the `/dev/mapper` directory on the host, allowing you to specify a friendly, easily-validated name whenever a configuration or maintenance operation requires specifying a block storage volume.

 If you are setting up shared storage to support StorageGRID node migration and using device mapper multipathing, you can create and install a common `/etc/multipath.conf` on all co-located hosts. Just make sure to use a different Docker storage volume on each host. Using aliases and including the target hostname in the alias for each Docker storage volume LUN will make this easy to remember and is recommended.

Related information

[Storage and performance requirements](#)

Node container migration requirements

Configure the Docker storage volume

Before installing Docker, you might need to format the Docker storage volume and mount it on `/var/lib/docker`.

About this task

You can skip these steps if you plan to use local storage for the Docker storage volume and have sufficient space available on the host partition containing `/var/lib`.

Steps

1. Create a file system on the Docker storage volume:

```
sudo mkfs.ext4 docker-storage-volume-device
```

2. Mount the Docker storage volume:

```
sudo mkdir -p /var/lib/docker
sudo mount docker-storage-volume-device /var/lib/docker
```

3. Add an entry for `docker-storage-volume-device` to `/etc/fstab`.

This step ensures that the storage volume will remount automatically after host reboots.

Install Docker

The StorageGRID system runs on Linux as a collection of Docker containers. Before you can install StorageGRID, you must install Docker.

Steps

1. Install Docker by following the instructions for your Linux distribution.



If Docker is not included with your Linux distribution, you can download it from the Docker website.

2. Ensure Docker has been enabled and started by running the following two commands:

```
sudo systemctl enable docker
```

```
sudo systemctl start docker
```

3. Confirm you have installed the expected version of Docker by entering the following:

```
sudo docker version
```

The Client and Server versions must be 1.11.0 or later.

Related information

[Configure host storage](#)

Install StorageGRID host services

You use the StorageGRID DEB package to install the StorageGRID host services.

About this task

These instructions describe how to install the host services from the DEB packages. As an alternative, you can use the APT repository metadata included in the installation archive to install the DEB packages remotely. See the APT repository instructions for your Linux operating system.

Steps

1. Copy the StorageGRID DEB packages to each of your hosts, or make them available on shared storage.

For example, place them in the `/tmp` directory, so you can use the example command in the next step.

2. Log in to each host as root or using an account with sudo permission, and run the following commands.

You must install the `images` package first, and the `service` package second. If you placed the packages in a directory other than `/tmp`, modify the command to reflect the path you used.

```
sudo dpkg --install /tmp/storagegrid-webscale-images-version-SHA.deb
```

```
sudo dpkg --install /tmp/storagegrid-webscale-service-version-SHA.deb
```



Python 2.7 must already be installed before the StorageGRID packages can be installed. The `sudo dpkg --install /tmp/storagegrid-webscale-images-version-SHA.deb` command will fail until you have done so.

Deploy virtual grid nodes (Ubuntu or Debian)

Create node configuration files for Ubuntu or Debian deployments

Node configuration files are small text files that provide the information the StorageGRID host service needs to start a node and connect it to the appropriate network and block storage resources. Node configuration files are used for virtual nodes and are not used for appliance nodes.

Where do I put the node configuration files?

You must place the configuration file for each StorageGRID node in the `/etc/storagegrid/nodes` directory on the host where the node will run. For example, if you plan to run one Admin Node, one Gateway Node, and one Storage Node on HostA, you must place three node configuration files in `/etc/storagegrid/nodes` on HostA. You can create the configuration files directly on each host using a text editor, such as vim or nano, or you can create them elsewhere and move them to each host.

What do I name the node configuration files?

The names of the configuration files are significant. The format is `node-name.conf`, where `node-name` is a name you assign to the node. This name appears in the StorageGRID Installer and is used for node maintenance operations, such as node migration.

Node names must follow these rules:

- Must be unique
- Must start with a letter
- Can contain the characters A through Z and a through z
- Can contain the numbers 0 through 9
- Can contain one or more hyphens (-)
- Must be no more than 32 characters, not including the `.conf` extension

Any files in `/etc/storagegrid/nodes` that do not follow these naming conventions will not be parsed by the host service.

If you have a multi-site topology planned for your grid, a typical node naming scheme might be:

```
site-nodetype-nodenumber.conf
```

For example, you might use `dc1-adm1.conf` for the first Admin Node in Data Center 1, and `dc2-sn3.conf` for the third Storage Node in Data Center 2. However, you can use any scheme you like, as long as all node names follow the naming rules.

What is in a node configuration file?

The configuration files contain key/value pairs, with one key and one value per line. For each key/value pair, you must follow these rules:

- The key and the value must be separated by an equal sign (=) and optional whitespace.
- The keys can contain no spaces.
- The values can contain embedded spaces.
- Any leading or trailing whitespace is ignored.

Some keys are required for every node, while others are optional or only required for certain node types.

The table defines the acceptable values for all supported keys. In the middle column:

R: required

BP: best practice

O: optional

Key	R, BP, or O?	Value
ADMIN_IP	BP	<p>Grid Network IPv4 address of the primary Admin Node for the grid to which this node belongs. Use the same value you specified for GRID_NETWORK_IP for the grid node with NODE_TYPE = VM_Admin_Node and ADMIN_ROLE = Primary. If you omit this parameter, the node attempts to discover a primary Admin Node using mDNS.</p> <p>How grid nodes discover the primary Admin Node</p> <p>Note: This value is ignored, and might be prohibited, on the primary Admin Node.</p>
ADMIN_NETWORK_CONFIG	O	DHCP, STATIC, or DISABLED
ADMIN_NETWORK_ESL	O	<p>Comma-separated list of subnets in CIDR notation to which this node should communicate via the Admin Network gateway.</p> <p>Example: 172.16.0.0/21,172.17.0.0/21</p>
ADMIN_NETWORK_GATEWAY	O (R)	<p>IPv4 address of the local Admin Network gateway for this node. Must be on the subnet defined by ADMIN_NETWORK_IP and ADMIN_NETWORK_MASK. This value is ignored for DHCP-configured networks.</p> <p>Note: This parameter is required if ADMIN_NETWORK_ESL is specified.</p> <p>Examples:</p> <p>1.1.1.1</p> <p>10.224.4.81</p>
ADMIN_NETWORK_IP	O	<p>IPv4 address of this node on the Admin Network. This key is only required when ADMIN_NETWORK_CONFIG = STATIC; do not specify it for other values.</p> <p>Examples:</p> <p>1.1.1.1</p> <p>10.224.4.81</p>

Key	R, BP, or O?	Value
ADMIN_NETWORK_MAC	O	<p>The MAC address for the Admin Network interface in the container.</p> <p>This field is optional. If omitted, a MAC address will be generated automatically.</p> <p>Must be 6 pairs of hexadecimal digits separated by colons.</p> <p>Example: b2:9c:02:c2:27:10</p>
ADMIN_NETWORK_MASK	O	<p>IPv4 netmask for this node, on the Admin Network. This key is only required when ADMIN_NETWORK_CONFIG = STATIC; do not specify it for other values.</p> <p>Examples:</p> <p>255.255.255.0</p> <p>255.255.248.0</p>
ADMIN_NETWORK_MTU	O	<p>The maximum transmission unit (MTU) for this node on the Admin Network. Do not specify if ADMIN_NETWORK_CONFIG = DHCP. If specified, the value must be between 1280 and 9216. If omitted, 1500 is used.</p> <p>If you want to use jumbo frames, set the MTU to a value suitable for jumbo frames, such as 9000. Otherwise, keep the default value.</p> <p>IMPORTANT: The MTU value of the network must match the value configured on the switch port the node is connected to. Otherwise, network performance issues or packet loss might occur.</p> <p>Examples:</p> <p>1500</p> <p>8192</p>

Key	R, BP, or O?	Value
ADMIN_NETWORK_TARGET	BP	<p>Name of the host device that you will use for Admin Network access by the StorageGRID node. Only network interface names are supported. Typically, you use a different interface name than what was specified for GRID_NETWORK_TARGET or CLIENT_NETWORK_TARGET.</p> <p>Note: Do not use bond or bridge devices as the network target. Either configure a VLAN (or other virtual interface) on top of the bond device, or use a bridge and virtual Ethernet (veth) pair.</p> <p>Best practice: Specify a value even if this node will not initially have an Admin Network IP address. Then you can add an Admin Network IP address later, without having to reconfigure the node on the host.</p> <p>Examples:</p> <p>bond0.1002</p> <p>ens256</p>
ADMIN_NETWORK_TARGET_TYPE	O	<p>Interface</p> <p>(This is the only supported value.)</p>
ADMIN_NETWORK_TARGET_TYPE_INTERFACE_CLONE_MAC	BP	<p>True or False</p> <p>Set the key to "true" to cause the StorageGRID container use the MAC address of the host target interface on the Admin Network.</p> <p>Best practice: In networks where promiscuous mode would be required, use the ADMIN_NETWORK_TARGET_TYPE_INTERFACE_CLONE_MAC key instead.</p> <p>For more details on MAC cloning:</p> <p>Considerations and recommendations for MAC address cloning (Red Hat Enterprise Linux or CentOS)</p> <p>Considerations and recommendations for MAC address cloning (Ubuntu or Debian)</p>
ADMIN_ROLE	R	<p>Primary or Non-Primary</p> <p>This key is only required when NODE_TYPE = VM_Admin_Node; do not specify it for other node types.</p>

Key	R, BP, or O?	Value
BLOCK_DEVICE_AUDIT_LOGS	R	<p>Path and name of the block device special file this node will use for persistent storage of audit logs. This key is only required for nodes with NODE_TYPE = VM_Admin_Node; do not specify it for other node types.</p> <p>Examples:</p> <pre>/dev/disk/by-path/pci-0000:03:00.0-scsi-0:0:0:0 /dev/disk/by-id/wwn-0x600a09800059d6df000060d757b475fd /dev/mapper/sgws-adm1-audit-logs</pre>

Key	R, BP, or O?	Value
BLOCK_DEVICE_RANGEDB_000	R	Path and name of the block device special file this node will use for persistent object storage. This key is only required for nodes with NODE_TYPE = VM_Storage_Node; do not specify it for other node types.
BLOCK_DEVICE_RANGEDB_001		Only BLOCK_DEVICE_RANGEDB_000 is required; the rest are optional. The block device specified for BLOCK_DEVICE_RANGEDB_000 must be at least 4 TB; the others can be smaller.
BLOCK_DEVICE_RANGEDB_002		
BLOCK_DEVICE_RANGEDB_003		
BLOCK_DEVICE_RANGEDB_004		
BLOCK_DEVICE_RANGEDB_005		Note: For compatibility with existing deployments, two-digit keys are supported for upgraded nodes.
BLOCK_DEVICE_RANGEDB_006		Examples:
BLOCK_DEVICE_RANGEDB_007		/dev/disk/by-path/pci-0000:03:00.0-scsi-0:0:0:0
BLOCK_DEVICE_RANGEDB_008		/dev/disk/by-id/wwn-0x600a09800059d6df000060d757b475fd
BLOCK_DEVICE_RANGEDB_009		/dev/mapper/sgws-sn1-rangedb-000
BLOCK_DEVICE_RANGEDB_010		
BLOCK_DEVICE_RANGEDB_011		
BLOCK_DEVICE_RANGEDB_012		
BLOCK_DEVICE_RANGEDB_013		
BLOCK_DEVICE_RANGEDB_014		
BLOCK_DEVICE_RANGEDB_015		

Key	R, BP, or O?	Value
BLOCK_DEVICE_TABLES	R	<p>Path and name of the block device special file this node will use for persistent storage of database tables. This key is only required for nodes with NODE_TYPE = VM_Admin_Node; do not specify it for other node types.</p> <p>Examples:</p> <pre>/dev/disk/by-path/pci-0000:03:00.0-scsi-0:0:0:0 /dev/disk/by-id/wwn-0x600a09800059d6df000060d757b475fd /dev/mapper/sgws-adm1-tables</pre>
BLOCK_DEVICE_VAR_LOCAL	R	<p>Path and name of the block device special file this node will use for its /var/local persistent storage.</p> <p>Examples:</p> <pre>/dev/disk/by-path/pci-0000:03:00.0-scsi-0:0:0:0 /dev/disk/by-id/wwn-0x600a09800059d6df000060d757b475fd /dev/mapper/sgws-sn1-var-local</pre>
CLIENT_NETWORK_CONFIG	O	DHCP, STATIC, or DISABLED
CLIENT_NETWORK_GATEWAY	O	<p>IPv4 address of the local Client Network gateway for this node, which must be on the subnet defined by CLIENT_NETWORK_IP and CLIENT_NETWORK_MASK. This value is ignored for DHCP-configured networks.</p> <p>Examples:</p> <pre>1.1.1.1 10.224.4.81</pre>

Key	R, BP, or O?	Value
CLIENT_NETWORK_IP	O	<p>IPv4 address of this node on the Client Network. This key is only required when CLIENT_NETWORK_CONFIG = STATIC; do not specify it for other values.</p> <p>Examples:</p> <p>1.1.1.1</p> <p>10.224.4.81</p>
CLIENT_NETWORK_MAC	O	<p>The MAC address for the Client Network interface in the container.</p> <p>This field is optional. If omitted, a MAC address will be generated automatically.</p> <p>Must be 6 pairs of hexadecimal digits separated by colons.</p> <p>Example: b2:9c:02:c2:27:20</p>
CLIENT_NETWORK_MASK	O	<p>IPv4 netmask for this node on the Client Network. This key is only required when CLIENT_NETWORK_CONFIG = STATIC; do not specify it for other values.</p> <p>Examples:</p> <p>255.255.255.0</p> <p>255.255.248.0</p>
CLIENT_NETWORK_MTU	O	<p>The maximum transmission unit (MTU) for this node on the Client Network. Do not specify if CLIENT_NETWORK_CONFIG = DHCP. If specified, the value must be between 1280 and 9216. If omitted, 1500 is used.</p> <p>If you want to use jumbo frames, set the MTU to a value suitable for jumbo frames, such as 9000. Otherwise, keep the default value.</p> <p>IMPORTANT: The MTU value of the network must match the value configured on the switch port the node is connected to. Otherwise, network performance issues or packet loss might occur.</p> <p>Examples:</p> <p>1500</p> <p>8192</p>

Key	R, BP, or O?	Value
CLIENT_NETWORK_TARGET	BP	<p>Name of the host device that you will use for Client Network access by the StorageGRID node. Only network interface names are supported. Typically, you use a different interface name than what was specified for GRID_NETWORK_TARGET or ADMIN_NETWORK_TARGET.</p> <p>Note: Do not use bond or bridge devices as the network target. Either configure a VLAN (or other virtual interface) on top of the bond device, or use a bridge and virtual Ethernet (veth) pair.</p> <p>Best practice: Specify a value even if this node will not initially have a Client Network IP address. Then you can add a Client Network IP address later, without having to reconfigure the node on the host.</p> <p>Examples:</p> <p>bond0.1003</p> <p>ens423</p>
CLIENT_NETWORK_TARGET_TYPE	O	<p>Interface</p> <p>(This is only supported value.)</p>
CLIENT_NETWORK_TARGET_TYPE_INTERFACE_CLONE_MAC	BP	<p>True or False</p> <p>Set the key to "true" to cause the StorageGRID container to use the MAC address of the host target interface on the Client Network.</p> <p>Best practice: In networks where promiscuous mode would be required, use the CLIENT_NETWORK_TARGET_TYPE_INTERFACE_CLONE_MAC key instead.</p> <p>For more details on MAC cloning:</p> <p>Considerations and recommendations for MAC address cloning (Red Hat Enterprise Linux or CentOS)</p> <p>Considerations and recommendations for MAC address cloning (Ubuntu or Debian)</p>
GRID_NETWORK_CONFIG	BP	<p>STATIC or DHCP</p> <p>(Defaults to STATIC if not specified.)</p>

Key	R, BP, or O?	Value
GRID_NETWORK_GATEWAY	R	<p>IPv4 address of the local Grid Network gateway for this node, which must be on the subnet defined by GRID_NETWORK_IP and GRID_NETWORK_MASK. This value is ignored for DHCP-configured networks.</p> <p>If the Grid Network is a single subnet with no gateway, use either the standard gateway address for the subnet (X.Y.Z.1) or this node's GRID_NETWORK_IP value; either value will simplify potential future Grid Network expansions.</p>
GRID_NETWORK_IP	R	<p>IPv4 address of this node on the Grid Network. This key is only required when GRID_NETWORK_CONFIG = STATIC; do not specify it for other values.</p> <p>Examples:</p> <p>1.1.1.1</p> <p>10.224.4.81</p>
GRID_NETWORK_MAC	O	<p>The MAC address for the Grid Network interface in the container.</p> <p>This field is optional. If omitted, a MAC address will be generated automatically.</p> <p>Must be 6 pairs of hexadecimal digits separated by colons.</p> <p>Example: b2:9c:02:c2:27:30</p>
GRID_NETWORK_MASK	O	<p>IPv4 netmask for this node on the Grid Network. This key is only required when GRID_NETWORK_CONFIG = STATIC; do not specify it for other values.</p> <p>Examples:</p> <p>255.255.255.0</p> <p>255.255.248.0</p>

Key	R, BP, or O?	Value
GRID_NETWORK_MTU	O	<p>The maximum transmission unit (MTU) for this node on the Grid Network. Do not specify if GRID_NETWORK_CONFIG = DHCP. If specified, the value must be between 1280 and 9216. If omitted, 1500 is used.</p> <p>If you want to use jumbo frames, set the MTU to a value suitable for jumbo frames, such as 9000. Otherwise, keep the default value.</p> <p>IMPORTANT: The MTU value of the network must match the value configured on the switch port the node is connected to. Otherwise, network performance issues or packet loss might occur.</p> <p>IMPORTANT: For the best network performance, all nodes should be configured with similar MTU values on their Grid Network interfaces. The Grid Network MTU mismatch alert is triggered if there is a significant difference in MTU settings for the Grid Network on individual nodes. The MTU values do not have to be the same for all network types.</p> <p>Examples:</p> <p>1500 8192</p>
GRID_NETWORK_TARGET	R	<p>Name of the host device that you will use for Grid Network access by the StorageGRID node. Only network interface names are supported. Typically, you use a different interface name than what was specified for ADMIN_NETWORK_TARGET or CLIENT_NETWORK_TARGET.</p> <p>Note: Do not use bond or bridge devices as the network target. Either configure a VLAN (or other virtual interface) on top of the bond device, or use a bridge and virtual Ethernet (veth) pair.</p> <p>Examples:</p> <p>bond0.1001</p> <p>ens192</p>
GRID_NETWORK_TARGET_TYPE	O	<p>Interface</p> <p>(This is the only supported value.)</p>

Key	R, BP, or O?	Value
GRID_NETWORK_TARGET_TYPE_INTERFACE_CLONE_MAC	BP	<p>True or False</p> <p>Set the value of the key to "true" to cause the StorageGRID container to use the MAC address of the host target interface on the Grid Network.</p> <p>Best practice: In networks where promiscuous mode would be required, use the GRID_NETWORK_TARGET_TYPE_INTERFACE_CLONE_MAC key instead.</p> <p>For more details on MAC cloning:</p> <p>Considerations and recommendations for MAC address cloning (Red Hat Enterprise Linux or CentOS)</p> <p>Considerations and recommendations for MAC address cloning (Ubuntu or Debian)</p>
INTERFACES_TARGET_nnn	O	<p>Name and optional description for an extra interface you want to add to this node. You can add multiple extra interfaces to each node.</p> <p>For <i>nnnn</i>, specify a unique number for each INTERFACES_TARGET entry you are adding.</p> <p>For the value, specify the name of the physical interface on the bare-metal host. Then, optionally, add a comma and provide a description of the interface, which is displayed on the VLAN interfaces page and the HA groups page.</p> <p>For example: INTERFACES_TARGET_01=ens256, Trunk</p> <p>If you add a trunk interface, you must configure a VLAN interface in StorageGRID. If you add an access interface, you can add the interface directly to an HA group; you do not need to configure a VLAN interface.</p>

Key	R, BP, or O?	Value
MAXIMUM_RAM	O	<p>The maximum amount of RAM that this node is allowed to consume. If this key is omitted, the node has no memory restrictions. When setting this field for a production-level node, specify a value that is at least 24 GB and 16 to 32 GB less than the total system RAM.</p> <p>Note: The RAM value affects a node's actual metadata reserved space. See the instructions for administering StorageGRID for a description of what Metadata Reserved Space is.</p> <p>The format for this field is <number><unit>, where <unit> can be b, k, m, or g.</p> <p>Examples:</p> <p>24g</p> <p>38654705664b</p> <p>Note: If you want to use this option, you must enable kernel support for memory cgroups.</p>
NODE_TYPE	R	<p>Type of node:</p> <p>VM_Admin_Node VM_Storage_Node VM_Archive_Node VM_API_Gateway</p>

Key	R, BP, or O?	Value
PORT_REMAP	O	<p>Remaps any port used by a node for internal grid node communications or external communications. Remapping ports is necessary if enterprise networking policies restrict one or more ports used by StorageGRID, as described in “Internal grid node communications” or “External communications.”</p> <p>IMPORTANT: Do not remap the ports you are planning to use to configure load balancer endpoints.</p> <p>Note: If only PORT_REMAP is set, the mapping that you specify is used for both inbound and outbound communications. If PORT_REMAP_INBOUND is also specified, PORT_REMAP applies only to outbound communications.</p> <p>The format used is: <network type>/<protocol>/<default port used by grid node>/<new port>, where <network type> is grid, admin, or client, and protocol is tcp or udp.</p> <p>For example:</p> <pre>PORT_REMAP = client/tcp/18082/443</pre>
PORT_REMAP_INBOUND	O	<p>Remaps inbound communications to the specified port. If you specify PORT_REMAP_INBOUND but do not specify a value for PORT_REMAP, outbound communications for the port are unchanged.</p> <p>IMPORTANT: Do not remap the ports you are planning to use to configure load balancer endpoints.</p> <p>The format used is: <network type>/<protocol:>/<remapped port >/<default port used by grid node>, where <network type> is grid, admin, or client, and protocol is tcp or udp.</p> <p>For example:</p> <pre>PORT_REMAP_INBOUND = grid/tcp/3022/22</pre>

Related information

[Networking guidelines](#)

How grid nodes discover the primary Admin Node

Grid nodes communicate with the primary Admin Node for configuration and management. Each grid node must know the IP address of the primary Admin Node on the Grid Network.

To ensure that a grid node can access the primary Admin Node, you can do either of the following when deploying the node:

- You can use the ADMIN_IP parameter to enter the primary Admin Node's IP address manually.
- You can omit the ADMIN_IP parameter to have the grid node discover the value automatically. Automatic discovery is especially useful when the Grid Network uses DHCP to assign the IP address to the primary Admin Node.

Automatic discovery of the primary Admin Node is accomplished using a multicast Domain Name System (mDNS). When the primary Admin Node first starts up, it publishes its IP address using mDNS. Other nodes on the same subnet can then query for the IP address and acquire it automatically. However, because multicast IP traffic is not normally routable across subnets, nodes on other subnets cannot acquire the primary Admin Node's IP address directly.

If you use automatic discovery:

 • You must include the ADMIN_IP setting for at least one grid node on any subnets that the primary Admin Node is not directly attached to. This grid node will then publish the primary Admin Node's IP address for other nodes on the subnet to discover with mDNS.

• Ensure that your network infrastructure supports passing multi-cast IP traffic within a subnet.

Example node configuration files

You can use the example node configuration files to help set up the node configuration files for your StorageGRID system. The examples show node configuration files for all types of grid nodes.

For most nodes, you can add Admin and Client Network addressing information (IP, mask, gateway, and so on) when you configure the grid using the Grid Manager or the Installation API. The exception is the primary Admin Node. If you want to browse to the Admin Network IP of the primary Admin Node to complete grid configuration (because the Grid Network is not routed, for example), you must configure the Admin Network connection for the primary Admin Node in its node configuration file. This is shown in the example.

 In the examples, the Client Network target has been configured as a best practice, even though the Client Network is disabled by default.

Example for primary Admin Node

Example file name: /etc/storagegrid/nodes/dc1-adm1.conf

Example file contents:

```

NODE_TYPE = VM_Admin_Node
ADMIN_ROLE = Primary
BLOCK_DEVICE_VAR_LOCAL = /dev/mapper/dc1-adm1-var-local
BLOCK_DEVICE_AUDIT_LOGS = /dev/mapper/dc1-adm1-audit-logs
BLOCK_DEVICE_TABLES = /dev/mapper/dc1-adm1-tables
GRID_NETWORK_TARGET = bond0.1001
ADMIN_NETWORK_TARGET = bond0.1002
CLIENT_NETWORK_TARGET = bond0.1003

GRID_NETWORK_IP = 10.1.0.2
GRID_NETWORK_MASK = 255.255.255.0
GRID_NETWORK_GATEWAY = 10.1.0.1

ADMIN_NETWORK_CONFIG = STATIC
ADMIN_NETWORK_IP = 192.168.100.2
ADMIN_NETWORK_MASK = 255.255.248.0
ADMIN_NETWORK_GATEWAY = 192.168.100.1
ADMIN_NETWORK_ESL = 192.168.100.0/21,172.16.0.0/21,172.17.0.0/21

```

Example for Storage Node

Example file name: /etc/storagegrid/nodes/dc1-sn1.conf

Example file contents:

```

NODE_TYPE = VM_Storage_Node
ADMIN_IP = 10.1.0.2
BLOCK_DEVICE_VAR_LOCAL = /dev/mapper/dc1-sn1-var-local
BLOCK_DEVICE_RANGEDB_00 = /dev/mapper/dc1-sn1-rangedb-0
BLOCK_DEVICE_RANGEDB_01 = /dev/mapper/dc1-sn1-rangedb-1
BLOCK_DEVICE_RANGEDB_02 = /dev/mapper/dc1-sn1-rangedb-2
BLOCK_DEVICE_RANGEDB_03 = /dev/mapper/dc1-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

```

Example for Archive Node

Example file name: /etc/storagegrid/nodes/dc1-arcl.conf

Example file contents:

```
NODE_TYPE = VM_Archive_Node
ADMIN_IP = 10.1.0.2
BLOCK_DEVICE_VAR_LOCAL = /dev/mapper/dc1-arc1-var-local
GRID_NETWORK_TARGET = bond0.1001
ADMIN_NETWORK_TARGET = bond0.1002
CLIENT_NETWORK_TARGET = bond0.1003

GRID_NETWORK_IP = 10.1.0.4
GRID_NETWORK_MASK = 255.255.255.0
GRID_NETWORK_GATEWAY = 10.1.0.1
```

Example for Gateway Node

Example file name: /etc/storagegrid/nodes/dc1-gw1.conf

Example file contents:

```
NODE_TYPE = VM_API_Gateway
ADMIN_IP = 10.1.0.2
BLOCK_DEVICE_VAR_LOCAL = /dev/mapper/dc1-gw1-var-local
GRID_NETWORK_TARGET = bond0.1001
ADMIN_NETWORK_TARGET = bond0.1002
CLIENT_NETWORK_TARGET = bond0.1003
GRID_NETWORK_IP = 10.1.0.5
GRID_NETWORK_MASK = 255.255.255.0
GRID_NETWORK_GATEWAY = 10.1.0.1
```

Example for a non-primary Admin Node

Example file name: /etc/storagegrid/nodes/dc1-adm2.conf

Example file contents:

```
NODE_TYPE = VM_Admin_Node
ADMIN_ROLE = Non-Primary
ADMIN_IP = 10.1.0.2
BLOCK_DEVICE_VAR_LOCAL = /dev/mapper/dc1-adm2-var-local
BLOCK_DEVICE_AUDIT_LOGS = /dev/mapper/dc1-adm2-audit-logs
BLOCK_DEVICE_TABLES = /dev/mapper/dc1-adm2-tables
GRID_NETWORK_TARGET = bond0.1001
ADMIN_NETWORK_TARGET = bond0.1002
CLIENT_NETWORK_TARGET = bond0.1003

GRID_NETWORK_IP = 10.1.0.6
GRID_NETWORK_MASK = 255.255.255.0
GRID_NETWORK_GATEWAY = 10.1.0.1
```

Validate the StorageGRID configuration

After creating configuration files in `/etc/storagegrid/nodes` for each of your StorageGRID nodes, you must validate the contents of those files.

To validate the contents of the configuration files, run the following command on each host:

```
sudo storagegrid node validate all
```

If the files are correct, the output shows **PASSED** for each configuration file, as shown in the example.

```
Checking for misnamed node configuration files... PASSED
Checking configuration file for node dc1-adm1... PASSED
Checking configuration file for node dc1-gw1... PASSED
Checking configuration file for node dc1-sn1... PASSED
Checking configuration file for node dc1-sn2... PASSED
Checking configuration file for node dc1-sn3... PASSED
Checking for duplication of unique values between nodes... PASSED
```



For an automated installation, you can suppress this output by using the `-q` or `--quiet` options in the `storagegrid` command (for example, `storagegrid --quiet...`). If you suppress the output, the command will have a non-zero exit value if any configuration warnings or errors were detected.

If the configuration files are incorrect, the issues are shown as **WARNING** and **ERROR**, as shown in the example. If any configuration errors are found, you must correct them before you continue with the installation.

```
Checking for misnamed node configuration files...
  WARNING: ignoring /etc/storagegrid/nodes/dc1-adm1
  WARNING: ignoring /etc/storagegrid/nodes/dc1-sn2.conf.keep
  WARNING: ignoring /etc/storagegrid/nodes/my-file.txt
Checking configuration file for node dc1-adm1...
  ERROR: NODE_TYPE = VM_Foo_Node
    VM_Foo_Node is not a valid node type.  See *.conf.sample
  ERROR: ADMIN_ROLE = Foo
    Foo is not a valid admin role.  See *.conf.sample
  ERROR: BLOCK_DEVICE_VAR_LOCAL = /dev/mapper/sgws-gw1-var-local
    /dev/mapper/sgws-gw1-var-local is not a valid block device
Checking configuration file for node dc1-gw1...
  ERROR: GRID_NETWORK_TARGET = bond0.1001
    bond0.1001 is not a valid interface.  See `ip link show`
  ERROR: GRID_NETWORK_IP = 10.1.3
    10.1.3 is not a valid IPv4 address
  ERROR: GRID_NETWORK_MASK = 255.248.255.0
    255.248.255.0 is not a valid IPv4 subnet mask
Checking configuration file for node dc1-sn1...
  ERROR: GRID_NETWORK_GATEWAY = 10.2.0.1
    10.2.0.1 is not on the local subnet
  ERROR: ADMIN_NETWORK_ESL = 192.168.100.0/21,172.16.0foo
    Could not parse subnet list
Checking configuration file for node dc1-sn2... PASSED
Checking configuration file for node dc1-sn3... PASSED
Checking for duplication of unique values between nodes...
  ERROR: GRID_NETWORK_IP = 10.1.0.4
    dc1-sn2 and dc1-sn3 have the same GRID_NETWORK_IP
  ERROR: BLOCK_DEVICE_VAR_LOCAL = /dev/mapper/sgws-sn2-var-local
    dc1-sn2 and dc1-sn3 have the same BLOCK_DEVICE_VAR_LOCAL
  ERROR: BLOCK_DEVICE_RANGEDB_00 = /dev/mapper/sgws-sn2-rangedb-0
    dc1-sn2 and dc1-sn3 have the same BLOCK_DEVICE_RANGEDB_00
```

Start the StorageGRID host service

To start your StorageGRID nodes, and ensure they restart after a host reboot, you must enable and start the StorageGRID host service.

Steps

1. Run the following commands on each host:

```
sudo systemctl enable storagegrid
sudo systemctl start storagegrid
```

2. Run the following command to ensure the deployment is proceeding:

```
sudo storagegrid node status node-name
```

For any node that returns a status of “Not Running” or “Stopped”, run the following command:

```
sudo storagegrid node start node-name
```

3. If you have previously enabled and started the StorageGRID host service (or if you are unsure if the service has been enabled and started), also run the following command:

```
sudo systemctl reload-or-restart storagegrid
```

Configure grid and complete installation (Ubuntu or Debian)

Navigate to the Grid Manager

You use the Grid Manager to define all of the information required to configure your StorageGRID system.

What you'll need

The primary Admin Node must be deployed and have completed the initial startup sequence.

Steps

1. Open your web browser and navigate to one of the following addresses:

```
https://primary_admin_node_ip
```

```
client_network_ip
```

Alternatively, you can access the Grid Manager on port 8443:

```
https://primary_admin_node_ip:8443
```



You can use the IP address for the primary Admin Node IP on the Grid Network or on the Admin Network, as appropriate for your network configuration.

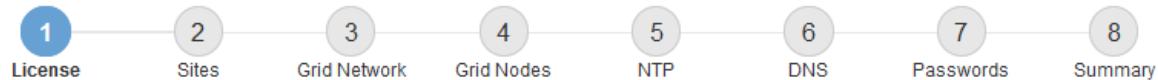
1. Click **Install a StorageGRID system**.

The page used to configure a StorageGRID grid appears.

NetApp® StorageGRID®

Help ▾

Install	
---------	--



License

Enter a grid name and upload the license file provided by NetApp for your StorageGRID system.

Grid Name

License File

Specify the StorageGRID license information

You must specify the name for your StorageGRID system and upload the license file provided by NetApp.

Steps

1. On the License page, enter a meaningful name for your StorageGRID system in **Grid Name**.

After installation, the name is displayed at the top of the Nodes menu.

2. Click **Browse**, locate the NetApp License File (NLUnique_id.txt), and click **Open**.

The license file is validated, and the serial number and licensed storage capacity are displayed.



The StorageGRID installation archive includes a free license that does not provide any support entitlement for the product. You can update to a license that offers support after installation.

NetApp® StorageGRID®

Help ▾

Install	
---------	--



License

Enter a grid name and upload the license file provided by NetApp for your StorageGRID system.

Grid Name

New License File

License Serial Number

Storage Capacity (TB)

3. Click **Next**.

Add sites

You must create at least one site when you are installing StorageGRID. You can create additional sites to increase the reliability and storage capacity of your StorageGRID system.

1. On the Sites page, enter the **Site Name**.
2. To add additional sites, click the plus sign next to the last site entry and enter the name in the new **Site Name** text box.

Add as many additional sites as required for your grid topology. You can add up to 16 sites.

NetApp® StorageGRID®

Help ▾

Install

1 License 2 Sites 3 Grid Network 4 Grid Nodes 5 NTP 6 DNS 7 Passwords 8 Summary

Sites

In a single-site deployment, infrastructure and operations are centralized in one site.

In a multi-site deployment, infrastructure can be distributed asymmetrically across sites, and proportional to the needs of each site. Typically, sites are located in geographically different locations. Having multiple sites also allows the use of distributed replication and erasure coding for increased availability and resiliency.

Site Name 1	Raleigh	X
Site Name 2	Atlanta	+ X

3. Click **Next**.

Specify Grid Network subnets

You must specify the subnets that are used on the Grid Network.

About this task

The subnet entries include the subnets for the Grid Network for each site in your StorageGRID system, along with any subnets that need to be reachable via the Grid Network.

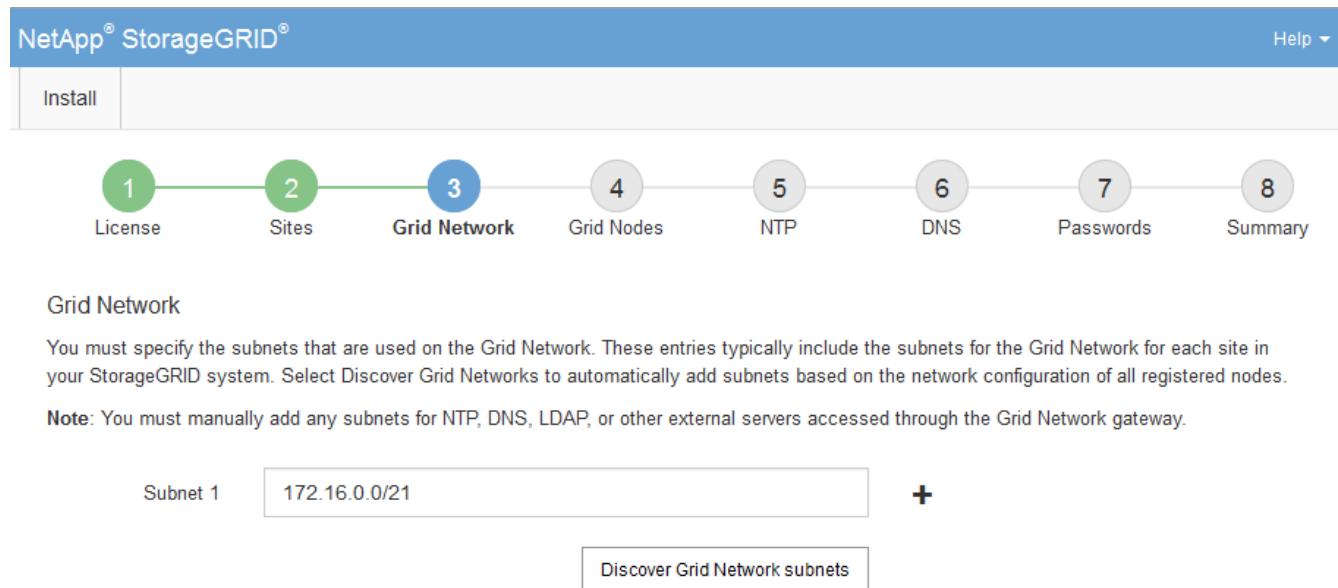
If you have multiple grid subnets, the Grid Network gateway is required. All grid subnets specified must be reachable through this gateway.

Steps

1. Specify the CIDR network address for at least one Grid Network in the **Subnet 1** text box.
2. Click the plus sign next to the last entry to add an additional network entry.

If you have already deployed at least one node, click **Discover Grid Networks Subnets** to automatically populate the Grid Network Subnet List with the subnets reported by grid nodes that have registered with

the Grid Manager.



The screenshot shows the 'Grid Network' step (Step 3) of the NetApp StorageGRID installation process. The top navigation bar includes 'Help ▾' and a 'Install' button. Below the navigation is a horizontal progress bar with numbered circles (1-8) and labels: License, Sites, Grid Network (highlighted in blue), Grid Nodes, NTP, DNS, Passwords, and Summary. The 'Grid Network' section contains a table with one row, 'Subnet 1', with the value '172.16.0.0/21'. A '+' button is to the right of the table, and a 'Discover Grid Network subnets' button is below it.

3. Click **Next**.

Approve pending grid nodes

You must approve each grid node before it can join the StorageGRID system.

What you'll need

You have deployed all virtual and StorageGRID appliance grid nodes.



It is more efficient to perform one single installation of all the nodes, rather than installing some nodes now and some nodes later.

Steps

1. Review the Pending Nodes list, and confirm that it shows all of the grid nodes you deployed.



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

2. Select the radio button next to a pending node you want to approve.



Grid Nodes

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

Pending Nodes

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

						Search
	Grid Network MAC Address	Name	Type	Platform	Grid Network IPv4 Address	
<input checked="" type="radio"/>	50:6b:4b:42:d7:00	NetApp-SGA	Storage Node	StorageGRID Appliance	172.16.5.20/21	

Approved Nodes

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

						Search
	Grid Network MAC Address	Name	Site	Type	Platform	Grid Network IPv4 Address
<input checked="" type="radio"/>	00:50:56:87:42:ff	dc1-adm1	Raleigh	Admin Node	VMware VM	172.16.4.210/21
<input checked="" type="radio"/>	00:50:56:87:c0:16	dc1-s1	Raleigh	Storage Node	VMware VM	172.16.4.211/21
<input checked="" type="radio"/>	00:50:56:87:79:ee	dc1-s2	Raleigh	Storage Node	VMware VM	172.16.4.212/21
<input checked="" type="radio"/>	00:50:56:87:db:9c	dc1-s3	Raleigh	Storage Node	VMware VM	172.16.4.213/21
<input checked="" type="radio"/>	00:50:56:87:62:38	dc1-g1	Raleigh	API Gateway Node	VMware VM	172.16.4.214/21

3. Click **Approve**.

4. In General Settings, modify settings for the following properties, as necessary:

Storage Node Configuration

General Settings

Site	<input type="text" value="Raleigh"/>
Name	<input type="text" value="NetApp-SGA"/>
NTP Role	<input type="text" value="Automatic"/>
ADC Service	<input type="text" value="Automatic"/>

Grid Network

Configuration	STATIC
IPv4 Address (CIDR)	<input type="text" value="172.16.5.20/21"/>
Gateway	<input type="text" value="172.16.5.20"/>

Admin Network

Configuration	STATIC
IPv4 Address (CIDR)	<input type="text" value="10.224.5.20/21"/>
Gateway	<input type="text" value="10.224.0.1"/>
Subnets (CIDR)	<input type="text" value="10.0.0.0/8"/> ×
	<input type="text" value="172.19.0.0/16"/> ×
	<input type="text" value="172.21.0.0/16"/> + ×

Client Network

Configuration	STATIC
IPv4 Address (CIDR)	<input type="text" value="47.47.5.20/21"/>
Gateway	<input type="text" value="47.47.0.1"/>

- **Site:** The name of the site with which this grid node will be associated.
- **Name:** The name that will be assigned to the node, and the name that will be displayed in the Grid Manager. The name defaults to the name you specified when you configured the node. During this step of the installation process, you can change the name as required.



After you complete the installation, you cannot change the name of the node.



For a VMware node, you can change the name here, but this action will not change the name of the virtual machine in vSphere.

- **NTP Role:** The Network Time Protocol (NTP) role of the grid node. The options are **Automatic**, **Primary**, and **Client**. Selecting **Automatic** assigns the Primary role to Admin Nodes, Storage Nodes with ADC services, Gateway Nodes, and any grid nodes that have non-static IP addresses. All other grid nodes are assigned the Client role.



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

- **ADC service** (Storage Nodes only): Select **Automatic** to let the system determine whether the node requires the Administrative Domain Controller (ADC) service. The ADC service keeps track of the location and availability of grid services. At least three Storage Nodes at each site must include the ADC service. You cannot add the ADC service to a node after it is deployed.

5. In Grid Network, modify settings for the following properties as necessary:

- **IPv4 Address (CIDR):** The CIDR network address for the Grid Network interface (eth0 inside the container). For example: 192.168.1.234/21
- **Gateway:** The Grid Network gateway. For example: 192.168.0.1

The gateway is required if there are multiple grid subnets.



If you selected DHCP for the Grid Network configuration and you change the value here, the new value will be configured as a static address on the node. You must make sure the resulting IP address is not within a DHCP address pool.

6. If you want to configure the Admin Network for the grid node, add or update the settings in the Admin Network section as necessary.

Enter the destination subnets of the routes out of this interface in the **Subnets (CIDR)** text box. If there are multiple Admin subnets, the Admin gateway is required.



If you selected DHCP for the Admin Network configuration and you change the value here, the new value will be configured as a static address on the node. You must make sure the resulting IP address is not within a DHCP address pool.

Appliances: For a StorageGRID appliance, if the Admin Network was not configured during the initial installation using the StorageGRID Appliance Installer, it cannot be configured in this Grid Manager dialog box. Instead, you must follow these steps:

- a. Reboot the appliance: In the Appliance Installer, select **Advanced > Reboot**.
Rebooting can take several minutes.
- b. Select **Configure Networking > Link Configuration** and enable the appropriate networks.
- c. Select **Configure Networking > IP Configuration** and configure the enabled networks.
- d. Return to the Home page and click **Start Installation**.
- e. In the Grid Manager: If the node is listed in the Approved Nodes table, reset the node.
- f. Remove the node from the Pending Nodes table.
- g. Wait for the node to reappear in the Pending Nodes list.

- h. Confirm that you can configure the appropriate networks. They should already be populated with the information you provided on the IP Configuration page.

For additional information, see the installation and maintenance instructions for your appliance model.

- 7. If you want to configure the Client Network for the grid node, add or update the settings in the Client Network section as necessary. If the Client Network is configured, the gateway is required, and it becomes the default gateway for the node after installation.



If you selected DHCP for the Client Network configuration and you change the value here, the new value will be configured as a static address on the node. You must make sure the resulting IP address is not within a DHCP address pool.

Appliances: For a StorageGRID appliance, if the Client Network was not configured during the initial installation using the StorageGRID Appliance Installer, it cannot be configured in this Grid Manager dialog box. Instead, you must follow these steps:

- a. Reboot the appliance: In the Appliance Installer, select **Advanced > Reboot**.

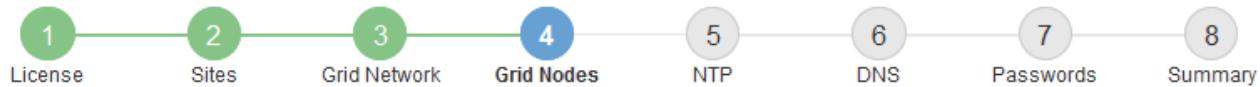
Rebooting can take several minutes.

- b. Select **Configure Networking > Link Configuration** and enable the appropriate networks.
- c. Select **Configure Networking > IP Configuration** and configure the enabled networks.
- d. Return to the Home page and click **Start Installation**.
- e. In the Grid Manager: If the node is listed in the Approved Nodes table, reset the node.
- f. Remove the node from the Pending Nodes table.
- g. Wait for the node to reappear in the Pending Nodes list.
- h. Confirm that you can configure the appropriate networks. They should already be populated with the information you provided on the IP Configuration page.

For additional information, see the installation and maintenance instructions for your appliance.

- 8. Click **Save**.

The grid node entry moves to the Approved Nodes list.



Grid Nodes

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

Pending Nodes

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

							Approve	Remove	Search	Search
	Grid Network MAC Address	Name	Type	Platform	Grid Network IPv4 Address					
<i>No results found.</i>										

Approved Nodes

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

							Edit	Reset	Remove	Search	Search
	Grid Network MAC Address	Name	Site	Type	Platform	Grid Network IPv4 Address					
<input type="radio"/>	00:50:56:87:42:ff	dc1-adm1	Raleigh	Admin Node	VMware VM	172.16.4.210/21					
<input type="radio"/>	00:50:56:87:c0:16	dc1-s1	Raleigh	Storage Node	VMware VM	172.16.4.211/21					
<input type="radio"/>	00:50:56:87:79:ee	dc1-s2	Raleigh	Storage Node	VMware VM	172.16.4.212/21					
<input type="radio"/>	00:50:56:87:db:9c	dc1-s3	Raleigh	Storage Node	VMware VM	172.16.4.213/21					
<input type="radio"/>	00:50:56:87:62:38	dc1-g1	Raleigh	API Gateway Node	VMware VM	172.16.4.214/21					
<input type="radio"/>	50:6b:4b:42:d7:00	NetApp-SGA	Raleigh	Storage Node	StorageGRID Appliance	172.16.5.20/21					

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

You must approve all nodes that you want in the grid. However, you can return to this page at any time before you click **Install** on the Summary page. You can modify the properties of an approved grid node by selecting its radio button and clicking **Edit**.

10. When you are done approving grid nodes, click **Next**.

Specify Network Time Protocol server information

You must specify the Network Time Protocol (NTP) configuration information for the StorageGRID system, so that operations performed on separate servers can be kept synchronized.

About this task

You must specify IPv4 addresses for the NTP servers.

You must specify external NTP servers. The specified NTP servers must use the NTP protocol.

You must specify four NTP server references of Stratum 3 or better to prevent issues with time drift.



When specifying the external NTP source for a production-level StorageGRID installation, do not use the Windows Time (W32Time) service on a version of Windows earlier than Windows Server 2016. The time service on earlier versions of Windows is not sufficiently accurate and is not supported by Microsoft for use in high-accuracy environments, such as StorageGRID.

[Support boundary to configure the Windows Time service for high-accuracy environments](#)

The external NTP servers are used by the nodes to which you previously assigned Primary NTP roles.



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

Steps

1. Specify the IPv4 addresses for at least four NTP servers in the **Server 1** to **Server 4** text boxes.
2. If necessary, select the plus sign next to the last entry to add additional server entries.

NetApp® StorageGRID®

Help ▾

Install

1 License 2 Sites 3 Grid Network 4 Grid Nodes 5 NTP 6 DNS 7 Passwords 8 Summary

Network Time Protocol

Enter the IP addresses for at least four Network Time Protocol (NTP) servers, so that operations performed on separate servers are kept in sync.

Server 1	10.60.248.183
Server 2	10.227.204.142
Server 3	10.235.48.111
Server 4	0.0.0.0

+

3. Select **Next**.

Related information

[Networking guidelines](#)

Specify Domain Name System server information

You must specify Domain Name System (DNS) information for your StorageGRID system, so that you can access external servers using hostnames instead of IP addresses.

About this task

Specifying DNS server information allows you to use Fully Qualified Domain Name (FQDN) hostnames rather than IP addresses for email notifications and AutoSupport. Specifying at least two DNS servers is recommended.



Provide two to six IPv4 addresses for DNS servers. You should select DNS servers that each site can access locally in the event of network islanding. This is to ensure an islanded site continues to have access to the DNS service. After configuring the grid-wide DNS server list, you can further customize the DNS server list for each node. For details, see the information about modifying the DNS configuration in the recovery and maintenance instructions.

If the DNS server information is omitted or incorrectly configured, a DNST alarm is triggered on each grid node's SSM service. The alarm clears when DNS is configured correctly and the new server information has reached all grid nodes.

Steps

1. Specify the IPv4 address for at least one DNS server in the **Server 1** text box.
2. If necessary, select the plus sign next to the last entry to add additional server entries.

NetApp® StorageGRID®

Install

Help ▾

1 License 2 Sites 3 Grid Network 4 Grid Nodes 5 NTP 6 DNS 7 Passwords 8 Summary

Domain Name Service

Enter the IP address for at least one Domain Name System (DNS) server, so that server hostnames can be used instead of IP addresses. Specifying at least two DNS servers is recommended. Configuring DNS enables server connectivity, email notifications, and NetApp AutoSupport.

Server 1	10.224.223.130	✖
Server 2	10.224.223.136	✚ ✖

The best practice is to specify at least two DNS servers. You can specify up to six DNS servers.

3. Select **Next**.

Specify the StorageGRID system passwords

As part of installing your StorageGRID system, you need to enter the passwords to use to secure your system and perform maintenance tasks.

About this task

Use the Install passwords page to specify the provisioning passphrase and the grid management root user password.

- The provisioning passphrase is used as an encryption key and is not stored by the StorageGRID system.
- You must have the provisioning passphrase for installation, expansion, and maintenance procedures, including downloading the Recovery Package. Therefore, it is important that you store the provisioning

passphrase in a secure location.

- You can change the provisioning passphrase from the Grid Manager if you have the current one.
- The grid management root user password may be changed using the Grid Manager.
- Randomly generated command line console and SSH passwords are stored in the Passwords.txt file in the Recovery Package.

Steps

1. In **Provisioning Passphrase**, enter the provisioning passphrase that will be required to make changes to the grid topology of your StorageGRID system.

Store the provisioning passphrase in a secure place.



If after the installation completes and you want to change the provisioning passphrase later, you can use the Grid Manager. Select **CONFIGURATION > Access control > Grid passwords**.

2. In **Confirm Provisioning Passphrase**, reenter the provisioning passphrase to confirm it.
3. In **Grid Management Root User Password**, enter the password to use to access the Grid Manager as the “root” user.

Store the password in a secure place.

4. In **Confirm Root User Password**, reenter the Grid Manager password to confirm it.

NetApp® StorageGRID®

Install

1 License 2 Sites 3 Grid Network 4 Grid Nodes 5 NTP 6 DNS 7 Passwords 8 Summary

Passwords

Enter secure passwords that meet your organization's security policies. A text file containing the command line passwords must be downloaded during the final installation step.

Provisioning Passphrase:

Confirm Provisioning Passphrase:

Grid Management Root User Password:

Confirm Root User Password:

Create random command line passwords.

5. If you are installing a grid for proof of concept or demo purposes, optionally deselect the **Create random command line passwords** check box.

For production deployments, random passwords should always be used for security reasons. Deselect

Create random command line passwords only for demo grids if you want to use default passwords to access grid nodes from the command line using the “root” or “admin” account.



You are prompted to download the Recovery Package file (`sgws-recovery-package-id-revision.zip`) after you click **Install** on the Summary page. You must [download this file](#) to complete the installation. The passwords required to access the system are stored in the `Passwords.txt` file, contained in the Recovery Package file.

6. Click **Next**.

Review your configuration and complete installation

You must carefully review the configuration information you have entered to ensure that the installation completes successfully.

Steps

1. View the **Summary** page.

NetApp® StorageGRID®

Help ▾

Install

1 License 2 Sites 3 Grid Network 4 Grid Nodes 5 NTP 6 DNS 7 Passwords 8 Summary

Summary

Verify that all of the grid configuration information is correct, and then click **Install**. You can view the status of each grid node as it installs. Click the [Modify](#) links to go back and change the associated information.

General Settings

Grid Name	Grid1	Modify License
Passwords	Auto-generated random command line passwords	Modify Passwords

Networking

NTP	10.60.248.183 10.227.204.142 10.235.48.111	Modify NTP
DNS	10.224.223.130 10.224.223.136	Modify DNS
Grid Network	172.16.0.0/21	Modify Grid Network

Topology

Topology	Atlanta	Modify Sites Modify Grid Nodes
	Raleigh	
	dc1-adm1 dc1-g1 dc1-s1 dc1-s2 dc1-s3 NetApp-SGA	

2. Verify that all of the grid configuration information is correct. Use the [Modify](#) links on the Summary page to go back and correct any errors.
3. Click **Install**.



If a node is configured to use the Client Network, the default gateway for that node switches from the Grid Network to the Client Network when you click **Install**. If you lose connectivity, you must ensure that you are accessing the primary Admin Node through an accessible subnet. See [Networking guidelines](#) for details.

4. Click **Download Recovery Package**.

When the installation progresses to the point where the grid topology is defined, you are prompted to download the Recovery Package file (.zip), and confirm that you can successfully access the contents of this file. You must download the Recovery Package file so that you can recover the StorageGRID system if one or more grid nodes fail. The installation continues in the background, but you cannot complete the installation and access the StorageGRID system until you download and verify this file.

5. Verify that you can extract the contents of the .zip file, and then save it in two safe, secure, and separate locations.



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

6. Select the **I have successfully downloaded and verified the Recovery Package file** check box, and click **Next**.

Download Recovery Package

Before proceeding, you must download the Recovery Package file. This file is necessary to recover the StorageGRID system if a failure occurs.

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.

After you have verified the contents, copy the Recovery Package file to two safe, secure, and separate locations. The Recovery Package file must be secured because it contains encryption keys and passwords that can be used to obtain data from the StorageGRID system.

The Recovery Package is required for recovery procedures and must be stored in a secure location.

Download Recovery Package

I have successfully downloaded and verified the Recovery Package file.

If the installation is still in progress, the status page appears. This page indicates the progress of the installation for each grid node.

Installation Status

If necessary, you may [Download the Recovery Package file again](#).

Name	Site	Grid Network IPv4 Address	Progress	Stage
dc1-adm1	Site1	172.16.4.215/21	<div style="width: 80%;"><div style="width: 100%;"></div></div>	Starting services
dc1-g1	Site1	172.16.4.216/21	<div style="width: 100%;"><div style="width: 100%;"></div></div>	Complete
dc1-s1	Site1	172.16.4.217/21	<div style="width: 50%;"><div style="width: 100%;"></div></div>	Waiting for Dynamic IP Service peers
dc1-s2	Site1	172.16.4.218/21	<div style="width: 10%;"><div style="width: 100%;"></div></div>	Downloading hotfix from primary Admin if needed
dc1-s3	Site1	172.16.4.219/21	<div style="width: 10%;"><div style="width: 100%;"></div></div>	Downloading hotfix from primary Admin if needed

When the Complete stage is reached for all grid nodes, the sign-in page for the Grid Manager appears.

7. Sign in to the Grid Manager using the "root" user and the password you specified during the installation.

Post-installation guidelines

After completing grid node deployment and configuration, follow these guidelines for DHCP addressing and network configuration changes.

- If DHCP was used to assign IP addresses, configure a DHCP reservation for each IP address on the networks being used.

You can only set up DHCP during the deployment phase. You cannot set up DHCP during configuration.



Nodes reboot when their IP addresses change, which can cause outages if a DHCP address change affects multiple nodes at the same time.

- You must use the Change IP procedures if you want to change IP addresses, subnet masks, and default gateways for a grid node. See [Configure IP addresses](#).
- If you make networking configuration changes, including routing and gateway changes, client connectivity to the primary Admin Node and other grid nodes might be lost. Depending on the networking changes applied, you might need to re-establish these connections.

Automate the installation (Ubuntu or Debian)

You can automate the installation of the StorageGRID host service and the configuration of grid nodes.

About this task

Automating the deployment might be useful in any of the following cases:

- You already use a standard orchestration framework, such as Ansible, Puppet, or Chef, to deploy and configure physical or virtual hosts.
- You intend to deploy multiple StorageGRID instances.
- You are deploying a large, complex StorageGRID instance.

The StorageGRID host service is installed by a package and driven by configuration files that can be created interactively during a manual installation, or prepared ahead of time (or programmatically) to enable automated installation using standard orchestration frameworks. StorageGRID provides optional Python scripts for automating the configuration of StorageGRID appliances, and the whole StorageGRID system (the "grid"). You can use these scripts directly, or you can inspect them to learn how to use the StorageGRID Installation REST API in grid deployment and configuration tools you develop yourself.

Automate the installation and configuration of the StorageGRID host service

You can automate the installation of the StorageGRID host service using standard orchestration frameworks such as Ansible, Puppet, Chef, Fabric, or SaltStack.

The StorageGRID host service is packaged in a DEB and is driven by configuration files that can be prepared ahead of time (or programmatically) to enable automated installation. If you already use a standard orchestration framework to install and configure Ubuntu or Debian, adding StorageGRID to your playbooks or recipes should be straightforward.

You can automate these tasks:

1. Installing Linux
2. Configuring Linux
3. Configuring host network interfaces to meet StorageGRID requirements
4. Configuring host storage to meet StorageGRID requirements
5. Installing Docker
6. Installing the StorageGRID host service
7. Creating StorageGRID node configuration files in `/etc/storagegrid/nodes`
8. Validating StorageGRID node configuration files
9. Starting the StorageGRID host service

Example Ansible role and playbook

Example Ansible role and playbook are supplied with the installation archive in the `/extras` folder. The Ansible playbook shows how the `storagegrid` role prepares the hosts and installs StorageGRID onto the target servers. You can customize the role or playbook as necessary.

Automate the configuration of StorageGRID

After deploying the grid nodes, you can automate the configuration of the StorageGRID system.

What you'll need

- You know the location of the following files from the installation archive.

Filename	Description
<code>configure-storagegrid.py</code>	Python script used to automate the configuration
<code>configure-storagegrid.sample.json</code>	Sample configuration file for use with the script
<code>configure-storagegrid.blank.json</code>	Blank configuration file for use with the script

- You have created a `configure-storagegrid.json` configuration file. To create this file, you can modify the sample configuration file (`configure-storagegrid.sample.json`) or the blank configuration file (`configure-storagegrid.blank.json`).

About this task

You can use the `configure-storagegrid.py` Python script and the `configure-storagegrid.json` configuration file to automate the configuration of your StorageGRID system.



You can also configure the system using the Grid Manager or the Installation API.

Steps

1. Log in to the Linux machine you are using to run the Python script.
2. Change to the directory where you extracted the installation archive.

For example:

```
cd StorageGRID-Webscale-version/platform
```

where `platform` is `debs`, `rpms`, or `vsphere`.

3. Run the Python script and use the configuration file you created.

For example:

```
./configure-storagegrid.py ./configure-storagegrid.json --start-install
```

Result

A Recovery Package `.zip` file is generated during the configuration process, and it is downloaded to the directory where you are running the installation and configuration process. You must back up the Recovery Package file so that you can recover the StorageGRID system if one or more grid nodes fails. For example, copy it to a secure, backed up network location and to a secure cloud storage location.



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

If you specified that random passwords should be generated, you need to extract the `Passwords.txt` file and look for the passwords required to access your StorageGRID system.

```
#####
##### The StorageGRID "recovery package" has been downloaded as: #####
#####           ./sgws-recovery-package-994078-rev1.zip           #####
#####   Safeguard this file as it will be needed in case of a   #####
#####           StorageGRID node recovery.                      #####
#################################################################
```

Your StorageGRID system is installed and configured when a confirmation message is displayed.

```
StorageGRID has been configured and installed.
```

Related information

[Overview of the installation REST API](#)

Overview of the installation REST API

StorageGRID provides the StorageGRID Installation API for performing installation tasks.

The API uses the Swagger open source API platform to provide the API documentation. Swagger allows both developers and non-developers to interact with the API in a user interface that illustrates how the API responds to parameters and options. This documentation assumes that you are familiar with standard web technologies

and the JSON (JavaScript Object Notation) data format.



Any API operations you perform using the API Docs webpage are live operations. Be careful not to create, update, or delete configuration data or other data by mistake.

Each REST API command includes the API's URL, an HTTP action, any required or optional URL parameters, and an expected API response.

StorageGRID Installation API

The StorageGRID Installation API is only available when you are initially configuring your StorageGRID system, and in the event that you need to perform a primary Admin Node recovery. The Installation API can be accessed over HTTPS from the Grid Manager.

To access the API documentation, go to the installation web page on the primary Admin Node and select **Help > API Documentation** from the menu bar.

The StorageGRID Installation API includes the following sections:

- **config** — Operations related to the product release and versions of the API. You can list the product release version and the major versions of the API supported by that release.
- **grid** — Grid-level configuration operations. You can get and update grid settings, including grid details, Grid Network subnets, grid passwords, and NTP and DNS server IP addresses.
- **nodes** — Node-level configuration operations. You can retrieve a list of grid nodes, delete a grid node, configure a grid node, view a grid node, and reset a grid node's configuration.
- **provision** — Provisioning operations. You can start the provisioning operation and view the status of the provisioning operation.
- **recovery** — Primary Admin Node recovery operations. You can reset information, upload the Recover Package, start the recovery, and view the status of the recovery operation.
- **recovery-package** — Operations to download the Recovery Package.
- **schemas** — API schemas for advanced deployments
- **sites** — Site-level configuration operations. You can create, view, delete, and modify a site.

Related information

[Automating the installation](#)

Where to go next

After completing an installation, you must perform a series of integration and configuration steps. Some steps are required; others are optional.

Required tasks

- Create a tenant account for each client protocol (Swift or S3) that will be used to store objects on your StorageGRID system.
- Control system access by configuring groups and user accounts. Optionally, you can configure a federated identity source (such as Active Directory or OpenLDAP), so you can import administration groups and users. Or, you can create local groups and users.
- Integrate and test the S3 or Swift API client applications you will use to upload objects to your

StorageGRID system.

- When you are ready, configure the information lifecycle management (ILM) rules and ILM policy you want to use to protect object data.



When you install StorageGRID, the default ILM policy, Baseline 2 Copies Policy, is active. This policy includes the stock ILM rule (Make 2 Copies), and it applies if no other policy has been activated.

- If your installation includes appliance Storage Nodes, use SANtricity software to complete the following tasks:
 - Connect to each StorageGRID appliance.
 - Verify receipt of AutoSupport data.
- If your StorageGRID system includes any Archive Nodes, configure the Archive Node's connection to the target external archival storage system.



If any Archive Nodes will use Tivoli Storage Manager as the external archival storage system, you must also configure Tivoli Storage Manager.

- Review and follow the StorageGRID system hardening guidelines to eliminate security risks.
- Configure email notifications for system alerts.

Optional tasks

- If you want to receive notifications from the (legacy) alarm system, configure mailing lists and email notifications for alarms.
- Update grid node IP addresses if they have changed since you planned your deployment and generated the Recovery Package. See information about changing IP addresses in the recovery and maintenance instructions.
- Configure storage encryption, if required.
- Configure storage compression to reduce the size of stored objects, if required.
- Configure audit client access. You can configure access to the system for auditing purposes through an NFS or a CIFS file share. See the instructions for administering StorageGRID.



Audit export through CIFS/Samba has been deprecated and will be removed in a future StorageGRID release.

Related information

[Administer StorageGRID](#)

[Use S3](#)

[Use Swift](#)

[Manage objects with ILM](#)

[Monitor and troubleshoot](#)

[Recover and maintain](#)

[SG100 and SG1000 services appliances](#)

[SG5600 storage appliances](#)

[SG5700 storage appliances](#)

[SG6000 storage appliances](#)

[Release notes](#)

[System hardening](#)

[Review audit logs](#)

[Upgrade software](#)

Troubleshoot installation issues

If any problems occur while installing your StorageGRID system, you can access the installation log files. Technical support might also need to use the installation log files to resolve issues.

The following installation log files are available from the container that is running each node:

- `/var/local/log/install.log` (found on all grid nodes)
- `/var/local/log/gdu-server.log` (found on the primary Admin Node)

The following installation log files are available from the host:

- `/var/log/storagegrid/daemon.log`
- `/var/log/storagegrid/nodes/<node-name>.log`

To learn how to access the log files, see the instructions for monitoring and troubleshooting StorageGRID. For help troubleshooting appliance installation issues, see the installation and maintenance instructions for your appliances. If you need additional help, contact technical support.

Related information

[Monitor and troubleshoot](#)

[SG100 and SG1000 services appliances](#)

[SG6000 storage appliances](#)

[SG5700 storage appliances](#)

[SG5600 storage appliances](#)

[NetApp Support](#)

Example `/etc/network/interfaces`

The `/etc/network/interfaces` file includes three sections, which define the physical

interfaces, bond interface, and VLAN interfaces. You can combine the three example sections into a single file, which will aggregate four Linux physical interfaces into a single LACP bond and then establish three VLAN interfaces subtending the bond for use as StorageGRID Grid, Admin, and Client Network interfaces.

Physical interfaces

Note that the switches at the other ends of the links must also treat the four ports as a single LACP trunk or port channel, and must pass at least the three referenced VLANs with tags.

```
# loopback interface
auto lo
iface lo inet loopback

# ens160 interface
auto ens160
iface ens160 inet manual
    bond-master bond0
    bond-primary en160

# ens192 interface
auto ens192
iface ens192 inet manual
    bond-master bond0

# ens224 interface
auto ens224
iface ens224 inet manual
    bond-master bond0

# ens256 interface
auto ens256
iface ens256 inet manual
    bond-master bond0
```

Bond interface

```
# bond0 interface
auto bond0
iface bond0 inet manual
    bond-mode 4
    bond-miimon 100
    bond-slaves ens160 ens192 ens224 ens256
```

VLAN interfaces

```
# 1001 vlan
auto bond0.1001
iface bond0.1001 inet manual
vlan-raw-device bond0

# 1002 vlan
auto bond0.1002
iface bond0.1002 inet manual
vlan-raw-device bond0

# 1003 vlan
auto bond0.1003
iface bond0.1003 inet manual
vlan-raw-device bond0
```

Install VMware

Install VMware: Overview

Installing a StorageGRID system in a VMware environment includes three primary steps.

1. Preparation: During planning and preparation, you perform the following tasks:

- Learn about the hardware, software, virtual machine, storage, and performance requirements for StorageGRID.
- Learn about the specifics of [StorageGRID networking](#) so you can configure your network appropriately.
- Identify and prepare the physical servers you plan to use to host your StorageGRID grid nodes.
- On the servers you have prepared:
 - Install VMware vSphere Hypervisor
 - Configure the ESX hosts
 - Install and configure VMware vSphere and vCenter

2. Deployment: Deploy grid nodes using the VMware vSphere Web Client. When you deploy grid nodes, they are created as part of the StorageGRID system and connected to one or more networks.

- a. Use the VMware vSphere Web Client, a .vmdk file, and a set of .ovf file templates to deploy the software-based nodes as virtual machines (VMs) on the servers you prepared in step 1.
- b. Use the StorageGRID Appliance Installer to deploy StorageGRID appliance nodes.



Hardware-specific installation and integration instructions are not included in the StorageGRID installation procedure. To learn how to install StorageGRID appliances, see the installation and maintenance instructions for your appliance.

3. Configuration: When all nodes have been deployed, use the Grid Manager to configure the grid and complete the installation.

These instructions recommend a standard approach for deploying and configuring a StorageGRID system in a VMware environment. See also the information about the following alternative approaches:

- Use the `deploy-vsphere-ovftool.sh` Bash script (available from the installation archive) to deploy grid nodes in VMware vSphere.
- Automate the deployment and configuration of the StorageGRID system using a Python configuration script (provided in the installation archive).
- Automate the deployment and configuration of appliance grid nodes with a Python configuration script (available from the installation archive or from the StorageGRID Appliance Installer).
- If you are an advanced developer of StorageGRID deployments, use the installation REST APIs to automate the installation of StorageGRID grid nodes.

Plan and prepare for VMware installation

Before you install (VMware)

Before deploying grid nodes and configuring the StorageGRID grid, you must be familiar with the steps and requirements for completing the procedure.

The StorageGRID deployment and configuration procedures assume that you are familiar with the architecture and operational functionality of the StorageGRID system.

You can deploy a single site or multiple sites at one time; however, all sites must meet the minimum requirement of having at least three Storage Nodes.



StorageGRID does not support the use of virtual storage area networks (vSANs), because the underlying disk protection is not hardware RAID.

Before starting the node deployment and grid configuration procedure, you must:

- Plan the StorageGRID deployment.
- Install, connect, and configure all required hardware, including any StorageGRID appliances, to specifications.



Hardware-specific installation and integration instructions are not included in the StorageGRID installation procedure. To learn how to install StorageGRID appliances, see the installation and maintenance instructions for your appliance.

- Understand the [available network options](#) and how each network option should be implemented on grid nodes.
- Gather all networking information in advance. Unless you are using DHCP, gather the IP addresses to assign to each grid node, and the IP addresses of the domain name system (DNS) and network time protocol (NTP) servers that will be used.
- Decide which of the available deployment and configuration tools you want to use.

Related information

[SG100 and SG1000 services appliances](#)

[SG6000 storage appliances](#)

SG5700 storage appliances

SG5600 storage appliances

Required materials

Before you install StorageGRID, you must gather and prepare required materials.

Item	Notes
NetApp StorageGRID license	You must have a valid, digitally signed NetApp license. Note: The StorageGRID installation archive includes a free license that does not provide any support entitlement for the product.
StorageGRID installation archive	You must download the StorageGRID installation archive and extract the files .
VMware software and documentation	During installation, you use VMware vSphere Web Client to deploy virtual grid nodes on virtual machines. For supported versions, see the Interoperability Matrix .
Service laptop	The StorageGRID system is installed through a service laptop. The service laptop must have: <ul style="list-style-type: none">• Network port• SSH client (for example, PuTTY)• Supported web browser
StorageGRID documentation	<ul style="list-style-type: none">• Release notes• Instructions for administering StorageGRID

Related information

[NetApp Interoperability Matrix Tool](#)

Download and extract the StorageGRID installation files

You must download the StorageGRID installation archives and extract the files..

Steps

1. Go to the [NetApp Downloads page for StorageGRID](#).
2. Select the button for downloading the latest release, or select another version from the drop-down menu and select **Go**.
3. Sign in with the username and password for your NetApp account.
4. If a Caution/MustRead statement appears, read it and select the check box.



You must apply any required hotfixes after you install the StorageGRID release. For more information, see the [hotfix procedure in the recovery and maintenance instructions](#)

5. Read the End User License Agreement, select the check box, and then select **Accept & Continue**.
6. In the **Install StorageGRID** column, select the .tgz or .zip file for VMware.



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

7. Save and extract the archive file.
8. Choose the files you need from the following list.

The files you need depend on your planned grid topology and how you will deploy your StorageGRID system.



The paths listed in the table are relative to the top-level directory installed by the extracted installation archive.

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

Path and file name	Description
./vsphere/deploy-vsphere-ovftool-sample.ini	An example configuration file for use with the <code>deploy-vsphere-ovftool.sh</code> script.
./vsphere/configure-storagegrid.py	A Python script used to automate the configuration of a StorageGRID system.
./vsphere/configure-sga.py	A Python script used to automate the configuration of StorageGRID appliances.
./vsphere/storagegrid-ssoauth.py	An example Python script that you can use to sign in to the Grid Management API when single sign-on is enabled.
./vsphere/configure-storagegrid.sample.json	An example configuration file for use with the <code>configure-storagegrid.py</code> script.
./vsphere/configure-storagegrid.blank.json	A blank configuration file for use with the <code>configure-storagegrid.py</code> script.
./vsphere/extras/api-schemas	<p>API schemas for StorageGRID.</p> <p>Note: Before you perform an upgrade, you can use these schemas to confirm that any code you have written to use StorageGRID management APIs will be compatible with the new StorageGRID release if you do not have a non-production StorageGRID environment for upgrade compatibility testing.</p>

Related information

[Recover and maintain](#)

Software requirements

You can use a virtual machine to host any type of StorageGRID grid node. One virtual machine is required for each grid node installed on the VMware server.

VMware vSphere Hypervisor

You must install VMware vSphere Hypervisor on a prepared physical server. The hardware must be configured correctly (including firmware versions and BIOS settings) before you install VMware software.

- Configure networking in the hypervisor as required to support networking for the StorageGRID system you are installing.

Networking guidelines

- Ensure that the datastore is large enough for the virtual machines and virtual disks that are required to host the grid nodes.

- If you create more than one datastore, name each datastore so that you can easily identify which datastore to use for each grid node when you create virtual machines.

ESX host configuration requirements



You must properly configure the network time protocol (NTP) on each ESX host. If the host time is incorrect, negative effects, including data loss, could occur.

VMware configuration requirements

You must install and configure VMware vSphere and vCenter before deploying StorageGRID grid nodes.

For supported versions of VMware vSphere Hypervisor and VMware vCenter Server software, see the Interoperability Matrix.

For the steps required to install these VMware products, see the VMware documentation.

Related information

[NetApp Interoperability Matrix Tool](#)

CPU and RAM requirements

Before installing StorageGRID software, verify and configure the hardware so that it is ready to support the StorageGRID system.

For information about supported servers, see the Interoperability Matrix.

Each StorageGRID node requires the following minimum resources:

- CPU cores: 8 per node
- RAM: At least 24 GB per node, and 2 to 16 GB less than the total system RAM, depending on the total RAM available and the amount of non-StorageGRID software running on the system

Ensure that the number of StorageGRID nodes you plan to run on each physical or virtual host does not exceed the number of CPU cores or the physical RAM available. If the hosts are not dedicated to running StorageGRID (not recommended), be sure to consider the resource requirements of the other applications.



Monitor your CPU and memory usage regularly to ensure that these resources continue to accommodate your workload. For example, doubling the RAM and CPU allocation for virtual Storage Nodes would provide similar resources to those provided for StorageGRID appliance nodes. Additionally, if the amount of metadata per node exceeds 500 GB, consider increasing the RAM per node to 48 GB or more. For information about managing object metadata storage, increasing the Metadata Reserved Space setting, and monitoring CPU and memory usage, see the instructions for administering, monitoring, and upgrading StorageGRID.

If hyperthreading is enabled on the underlying physical hosts, you can provide 8 virtual cores (4 physical cores) per node. If hyperthreading is not enabled on the underlying physical hosts, you must provide 8 physical cores per node.

If you are using virtual machines as hosts and have control over the size and number of VMs, you should use a single VM for each StorageGRID node and size the VM accordingly.

For production deployments, you should not run multiple Storage Nodes on the same physical storage

hardware or virtual host. Each Storage Node in a single StorageGRID deployment should be in its own isolated failure domain. You can maximize the durability and availability of object data if you ensure that a single hardware failure can only impact a single Storage Node.

See also the information about storage requirements.

Related information

[NetApp Interoperability Matrix Tool](#)

[Storage and performance requirements](#)

[Administer StorageGRID](#)

[Monitor and troubleshoot](#)

[Upgrade software](#)

Storage and performance requirements

You must understand the storage and performance requirements for StorageGRID nodes hosted by virtual machines, so you can provide enough space to support the initial configuration and future storage expansion.

Performance requirements

The performance of the OS volume and of the first storage volume significantly impacts the overall performance of the system. Ensure that these provide adequate disk performance in terms of latency, input/output operations per second (IOPS), and throughput.

All StorageGRID nodes require that the OS drive and all storage volumes have write-back caching enabled. The cache must be on a protected or persistent media.

Requirements for virtual machines that use NetApp ONTAP storage

If you are deploying a StorageGRID node as a virtual machine with storage assigned from a NetApp ONTAP system, you have confirmed that the volume does not have a FabricPool tiering policy enabled. For example, if a StorageGRID node is running as a virtual machine on a VMWare host, ensure the volume backing the datastore for the node does not have a FabricPool tiering policy enabled. Disabling FabricPool tiering for volumes used with StorageGRID nodes simplifies troubleshooting and storage operations.



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

Number of virtual machines required

Each StorageGRID site requires a minimum of three Storage Nodes.



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

Other types of nodes, such as Admin Nodes or Gateway Nodes, can be deployed on the same virtual machine

host, or they can be deployed on their own dedicated virtual machine hosts as required. However, if you have multiple nodes of the same type (two Gateway Nodes, for example), do not install all instances on the same virtual machine host.

Storage requirements by node type

In a production environment, the virtual machines for StorageGRID grid nodes must meet different requirements, depending on the types of nodes.



Disk snapshots cannot be used to restore grid nodes. Instead, refer to the recovery and maintenance procedures for each type of node.

Node Type	Storage
Admin Node	100 GB LUN for OS 200 GB LUN for Admin Node tables 200 GB LUN for Admin Node audit log
Storage Node	100 GB LUN for OS 3 LUNs for each Storage Node on this host Note: A Storage Node can have 1 to 16 storage LUNs; at least 3 storage LUNs are recommended. Minimum size per LUN: 4 TB Maximum tested LUN size: 39 TB.
Gateway Node	100 GB LUN for OS
Archive Node	100 GB LUN for OS



Depending on the audit level configured, the size of user inputs such as S3 object key name, and how much audit log data you need to preserve, you might need to increase the size of the audit log LUN on each Admin Node. As a general rule, a grid generates approximately 1 KB of audit data per S3 operation, which would mean that a 200 GB LUN would support 70 million operations per day or 800 operations per second for two to three days.

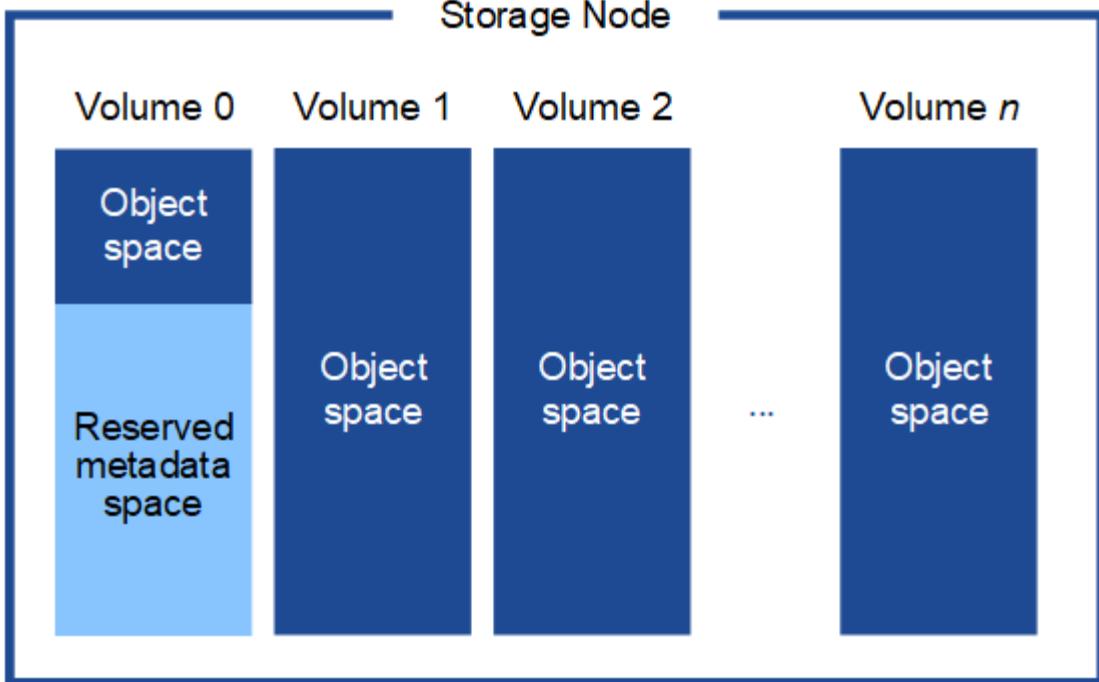
Storage requirements for Storage Nodes

A software-based Storage Node can have 1 to 16 storage volumes—3 or more storage volumes are recommended. Each storage volume should be 4 TB or larger.



An appliance Storage Node can have up to 48 storage volumes.

As shown in the figure, StorageGRID reserves space for object metadata on storage volume 0 of each Storage Node. Any remaining space on storage volume 0 and any other storage volumes in the Storage Node are used exclusively for object data.



To provide redundancy and to protect object metadata from loss, StorageGRID stores three copies of the metadata for all objects in the system at each site. The three copies of object metadata are evenly distributed across all Storage Nodes at each site.

When you assign space to volume 0 of a new Storage Node, you must ensure there is adequate space for that node's portion of all object metadata.

- At a minimum, you must assign at least 4 TB to volume 0.



If you use only one storage volume for a Storage Node and you assign 4 TB or less to the volume, the Storage Node might enter the Storage Read-Only state on startup and store object metadata only.

- If you are installing a new StorageGRID 11.6 system and each Storage Node has 128 GB or more of RAM, you should assign 8 TB or more to volume 0. Using a larger value for volume 0 can increase the space allowed for metadata on each Storage Node.
- When configuring different Storage Nodes for a site, use the same setting for volume 0 if possible. If a site contains Storage Nodes of different sizes, the Storage Node with the smallest volume 0 will determine the metadata capacity of that site.

For details, go to [Manage object metadata storage](#).

Related information

[Recover and maintain](#)

Deploy virtual machine grid nodes (VMware)

Collect information about your deployment environment

Before deploying grid nodes, you must collect information about your network configuration and VMware environment.



It is more efficient to perform one single installation of all the nodes, rather than installing some nodes now and some nodes later.

VMware information

You must access the deployment environment and collect information about the VMware environment; the networks that were created for the Grid, Admin, and Client Networks; and the storage volume types you plan to use for Storage Nodes.

You must collect information about your VMware environment, including the following:

- The username and password for a VMware vSphere account that has appropriate permissions to complete the deployment.
- Host, datastore, and network configuration information for each StorageGRID grid node virtual machine.



VMware live vMotion causes the virtual machine clock time to jump and is not supported for grid nodes of any type. Though rare, incorrect clock times can result in loss of data or configuration updates.

Grid Network information

You must collect information about the VMware network created for the StorageGRID Grid Network (required), including:

- The network name.
- If you are not using DHCP, the required networking details for each grid node (IP address, gateway, and network mask).
- If you are not using DHCP, the IP address of the primary Admin Node on the Grid Network. See “How grid nodes discover the primary Admin Node” for more information.

Admin Network information

For nodes that will be connected to the optional StorageGRID Admin Network, you must collect information about the VMware network created for this network, including:

- The network name.
- The method used to assign IP addresses, either static or DHCP.
- If you are using static IP addresses, the required networking details for each grid node (IP address, gateway, network mask).
- The external subnet list (ESL) for the Admin Network.

Client Network information

For nodes that will be connected to the optional StorageGRID Client Network, you must collect information about the VMware network created for this network, including:

- The network name.
- The method used to assign IP addresses, either static or DHCP.
- If you are using static IP addresses, the required networking details for each grid node (IP address, gateway, network mask).

Information about additional interfaces

You can optionally add trunk or access interfaces to the VM in vCenter after you install the node. For example, you might want to add a trunk interface to an Admin or Gateway Node, so you can use VLAN interfaces to segregate the traffic belonging to different applications or tenants. Or, you might want to add an access interface to use in a high availability (HA) group.

The interfaces you add are displayed on the VLAN interfaces page and on the HA groups page in the Grid Manager.

- If you add a trunk interface, configure one or more VLAN interfaces for each new parent interface. See [configure VLAN interfaces](#).
- If you add an access interface, you must add it directly to HA groups. See [configure high availability groups](#).

Storage volumes for virtual Storage Nodes

You must collect the following information for virtual machine-based Storage Nodes:

- The number and size of storage volumes (storage LUNs) you plan to add. See “Storage and performance requirements.”

Grid configuration information

You must collect information to configure your grid:

- Grid license
- Network Time Protocol (NTP) server IP addresses
- Domain Name System (DNS) server IP addresses

Related information

[How grid nodes discover the primary Admin Node](#)

[Storage and performance requirements](#)

How grid nodes discover the primary Admin Node

Grid nodes communicate with the primary Admin Node for configuration and management. Each grid node must know the IP address of the primary Admin Node on the Grid Network.

To ensure that a grid node can access the primary Admin Node, you can do either of the following when deploying the node:

- You can use the ADMIN_IP parameter to enter the primary Admin Node’s IP address manually.
- You can omit the ADMIN_IP parameter to have the grid node discover the value automatically. Automatic discovery is especially useful when the Grid Network uses DHCP to assign the IP address to the primary Admin Node.

Automatic discovery of the primary Admin Node is accomplished using a multicast Domain Name System (mDNS). When the primary Admin Node first starts up, it publishes its IP address using mDNS. Other nodes on the same subnet can then query for the IP address and acquire it automatically. However, because multicast IP traffic is not normally routable across subnets, nodes on other subnets cannot acquire the primary Admin

Node's IP address directly.

If you use automatic discovery:

- You must include the ADMIN_IP setting for at least one grid node on any subnets that the primary Admin Node is not directly attached to. This grid node will then publish the primary Admin Node's IP address for other nodes on the subnet to discover with mDNS.
- Ensure that your network infrastructure supports passing multi-cast IP traffic within a subnet.

Deploy a StorageGRID node as a virtual machine

You use VMware vSphere Web Client to deploy each grid node as a virtual machine. During deployment, each grid node is created and connected to one or more StorageGRID networks.

If you need to deploy any StorageGRID appliance Storage Nodes, see the installation and maintenance instructions for the appliance.

Optionally, you can remap node ports or increase CPU or memory settings for the node before powering it on.

What you'll need

- You have reviewed how to [plan and prepare for installation](#), and you understand the requirements for software, CPU and RAM, and storage and performance.
- You are familiar with VMware vSphere Hypervisor and have experience deploying virtual machines in this environment.



The open-vm-tools package, an open-source implementation similar to VMware Tools, is included with the StorageGRID virtual machine. You do not need to install VMware Tools manually.

- You have downloaded and extracted the correct version of the StorageGRID installation archive for VMware.



If you are deploying the new node as part of an expansion or recovery operation, you must use the version of StorageGRID that is currently running on the grid.

- You have the StorageGRID Virtual Machine Disk (.vmdk) file:

NetApp-SG-version-SHA.vmdk

- You have the .ovf and .mf files for each type of grid node you are deploying:

Filename	Description
vsphere-primary-admin.ovf	The template file and manifest file for the primary Admin Node.
vsphere-primary-admin.mf	

Filename	Description
vsphere-non-primary-admin.ovf	The template file and manifest file for a non-primary Admin Node.
vsphere-non-primary-admin.mf	
vsphere-archive.ovf	The template file and manifest file for an Archive Node.
vsphere-archive.mf	
vsphere-gateway.ovf	The template file and manifest file for a Gateway Node.
vsphere-gateway.mf	
vsphere-storage.ovf	The template file and manifest file for a Storage Node.
vsphere-storage.mf	

- The .vdmk, .ovf, and .mf files are all in the same directory.
- You have a plan to minimize failure domains. For example, you should not deploy all Gateway Nodes on a single virtual machine server.



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

- If you are deploying a node as part of an expansion or recovery operation, you have the [instructions for expanding a StorageGRID system](#) or the [recovery and maintenance instructions](#).
- If you are deploying a StorageGRID node as a virtual machine with storage assigned from a NetApp ONTAP system, you have confirmed that the volume does not have a FabricPool tiering policy enabled. For example, if a StorageGRID node is running as a virtual machine on a VMWare host, ensure the volume backing the datastore for the node does not have a FabricPool tiering policy enabled. Disabling FabricPool tiering for volumes used with StorageGRID nodes simplifies troubleshooting and storage operations.



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

About this task

Follow these instructions to initially deploy VMware nodes, add a new VMware node in an expansion, or replace a VMware node as part of a recovery operation. Except as noted in the steps, the node deployment procedure is the same for all node types, including Admin Nodes, Storage Nodes, Gateway Nodes, and Archive Nodes.

If you are installing a new StorageGRID system:

- You must deploy the primary Admin Node before you deploy any other grid node.
- You must ensure that each virtual machine can connect to the primary Admin Node over the Grid Network.
- You must deploy all grid nodes before configuring the grid.

If you are performing an expansion or recovery operation:

- You must ensure that the new virtual machine can connect to the primary Admin Node over the Grid Network.

If you need to remap any of the node's ports, do not power on the new node until the port remap configuration is complete.

Steps

1. Using VCenter, deploy an OVF template.

If you specify a URL, point to a folder containing the following files. Otherwise, select each of these files from a local directory.

```
NetApp-SG-version-SHA.vmdk  
vsphere-node.ovf  
vsphere-node.mf
```

For example, if this is the first node you are deploying, use these files to deploy the primary Admin Node for your StorageGRID system:

```
NetApp-SG-version-SHA.vmdk  
sphere-primary-admin.ovf  
sphere-primary-admin.mf
```

2. Provide a name for the virtual machine.

The standard practice is to use the same name for both the virtual machine and the grid node.

3. Place the virtual machine in the appropriate vApp or resource pool.
4. If you are deploying the primary Admin Node, read and accept the End User License Agreement.

Depending on your version of vCenter, the order of the steps will vary for accepting the End User License Agreement, specifying the name of the virtual machine, and selecting a datastore.

5. Select storage for the virtual machine.

If you are deploying a node as part of recovery operation, perform the instructions in the [storage recovery step](#) to add new virtual disks, reattach virtual hard disks from the failed grid node, or both.

When deploying a Storage Node, use 3 or more storage volumes, with each storage volume being 4 TB or larger. You must assign at least 4 TB to volume 0.



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



Some StorageGRID installations might use larger, more active storage volumes than typical virtualized workloads. You might need to tune some hypervisor parameters, such as MaxAddressableSpaceTB, to achieve optimal performance. If you encounter poor performance, contact your virtualization support resource to determine whether your environment could benefit from workload-specific configuration tuning.

6. Select networks.

Determine which StorageGRID networks the node will use by selecting a destination network for each source network.

- The Grid Network is required. You must select a destination network in the vSphere environment.
- If you use the Admin Network, select a different destination network in the vSphere environment. If you do not use the Admin Network, select the same destination you selected for the Grid Network.
- If you use the Client Network, select a different destination network in the vSphere environment. If you do not use the Client Network, select the same destination you selected for the Grid Network.

7. Under **Customize Template**, configure the required StorageGRID node properties.

a. Enter the **Node name**.



If you are recovering a grid node, you must enter the name of the node you are recovering.

b. In the **Grid Network (eth0)** section, select STATIC or DHCP for the **Grid network IP configuration**.

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

c. In the **Primary Admin IP** field, enter the IP address of the primary Admin Node for the Grid Network.



This step does not apply if the node you are deploying is the primary Admin Node.

If you omit the primary Admin Node IP address, the IP address will be automatically discovered if the primary Admin Node, or at least one other grid node with ADMIN_IP configured, is present on the same subnet. However, it is recommended to set the primary Admin Node IP address here.

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

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

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

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

8. Review the virtual machine configuration and make any changes necessary.

9. When you are ready to complete, select **Finish** to start the upload of the virtual machine.

10. If you deployed this node as part of recovery operation and this is not a full-node recovery, perform these steps after deployment is complete:

- a. Right-click the virtual machine, and select **Edit Settings**.
- b. Select each default virtual hard disk that has been designated for storage, and select **Remove**.
- c. Depending on your data recovery circumstances, add new virtual disks according to your storage requirements, reattach any virtual hard disks preserved from the previously removed failed grid node, or both.

Note the following important guidelines:

- If you are adding new disks you should use the same type of storage device that was in use before node recovery.
- The Storage Node .ovf file defines several VMDKs for storage. Unless these VMDKs meet your storage requirements, you should remove them and assign appropriate VMDKs or RDMs for storage before powering up the node. VMDKs are more commonly used in VMware environments and are easier to manage, while RDMs may provide better performance for workloads that use larger object sizes (for example, greater than 100 MB).

11. If you need to remap the ports used by this node, follow these steps.

You might need to remap a port if your enterprise networking policies restrict access to one or more ports that are used by StorageGRID. See the [networking guidelines](#) for the ports used by StorageGRID.



Do not remap the ports used in load balancer endpoints.

- a. Select the new VM.
- b. From the Configure tab, select **Settings > vApp Options**. The location of **vApp Options** depends on the version of vCenter.
- c. In the **Properties** table, locate **PORT_REMAP_INBOUND** and **PORT_REMAP**.
- d. To symmetrically map both inbound and outbound communications for a port, select **PORT_REMAP**.



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

- i. Scroll back to the top of the table, and select **Edit**.
- ii. On the Type tab, select **User configurable**, and select **Save**.
- iii. Select **Set Value**.
- iv. Enter the port mapping:

```
<network type>/<protocol>/<default port used by grid node>/<new port>
```

<network type> is grid, admin, or client, and <protocol> is tcp or udp.

For example, to remap ssh traffic from port 22 to port 3022, enter:

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

v. Select **OK**.

e. To specify the port used for inbound communications to the node, select **PORT_REMAP_INBOUND**.



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

i. Scroll back to the top of the table, and select **Edit**.

ii. On the Type tab, select **User configurable**, and select **Save**.

iii. Select **Set Value**.

iv. Enter the port mapping:

```
<network type>/<protocol>/<remapped inbound port>/<default inbound port used by grid node>
```

<network type> is grid, admin, or client, and <protocol> is tcp or udp.

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

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

v. Select **OK**.

12. If you want to increase the CPU or memory for the node from the default settings:

a. Right-click the virtual machine, and select **Edit Settings**.

b. Change the number of CPUs or the amount of memory as required.

Set the **Memory Reservation** to the same size as the **Memory** allocated to the virtual machine.

c. Select **OK**.

13. Power on the virtual machine.

After you finish

If you deployed this node as part of an expansion or recovery procedure, return to those instructions to complete the procedure.

Configure the grid and complete installation (VMware)

Navigate to the Grid Manager

You use the Grid Manager to define all of the information required to configure your StorageGRID system.

What you'll need

The primary Admin Node must be deployed and have completed the initial startup sequence.

Steps

1. Open your web browser and navigate to one of the following addresses:

`https://primary_admin_node_ip`

`client_network_ip`

Alternatively, you can access the Grid Manager on port 8443:

`https://primary_admin_node_ip:8443`



You can use the IP address for the primary Admin Node IP on the Grid Network or on the Admin Network, as appropriate for your network configuration.

2. Click **Install a StorageGRID system**.

The page used to configure a StorageGRID grid appears.

NetApp® StorageGRID®

Help ▾

Install

1 License 2 Sites 3 Grid Network 4 Grid Nodes 5 NTP 6 DNS 7 Passwords 8 Summary

License

Enter a grid name and upload the license file provided by NetApp for your StorageGRID system.

Grid Name

License File

Specify the StorageGRID license information

You must specify the name for your StorageGRID system and upload the license file provided by NetApp.

Steps

1. On the License page, enter a meaningful name for your StorageGRID system in **Grid Name**.

After installation, the name is displayed at the top of the Nodes menu.

2. Click **Browse**, locate the NetApp License File (NLUnique_id.txt) and click **Open**.

The license file is validated, and the serial number and licensed storage capacity are displayed.



The StorageGRID installation archive includes a free license that does not provide any support entitlement for the product. You can update to a license that offers support after installation.

NetApp® StorageGRID®

Help ▾

Install

1 License 2 Sites 3 Grid Network 4 Grid Nodes 5 NTP 6 DNS 7 Passwords 8 Summary

License

Enter a grid name and upload the license file provided by NetApp for your StorageGRID system.

Grid Name	Grid1
New License File	<input type="button" value="Browse"/>
License Serial Number	950719
Storage Capacity (TB)	240

3. Click **Next**.

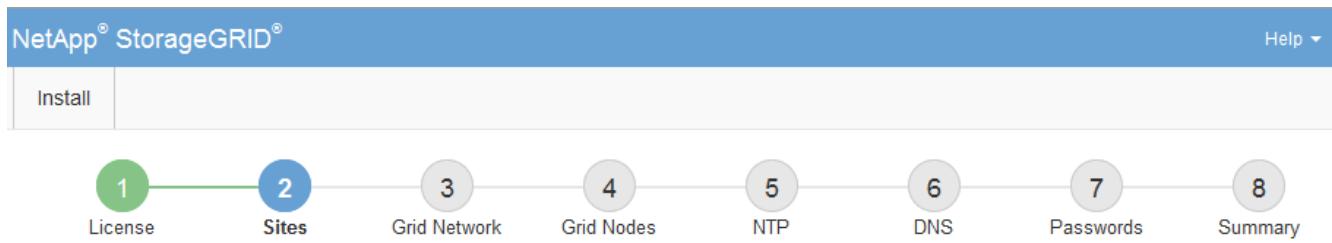
Add sites

You must create at least one site when you are installing StorageGRID. You can create additional sites to increase the reliability and storage capacity of your StorageGRID system.

Steps

1. On the Sites page, enter the **Site Name**.
2. To add additional sites, click the plus sign next to the last site entry and enter the name in the new **Site Name** text box.

Add as many additional sites as required for your grid topology. You can add up to 16 sites.



3. Click **Next**.

Specify Grid Network subnets

You must specify the subnets that are used on the Grid Network.

About this task

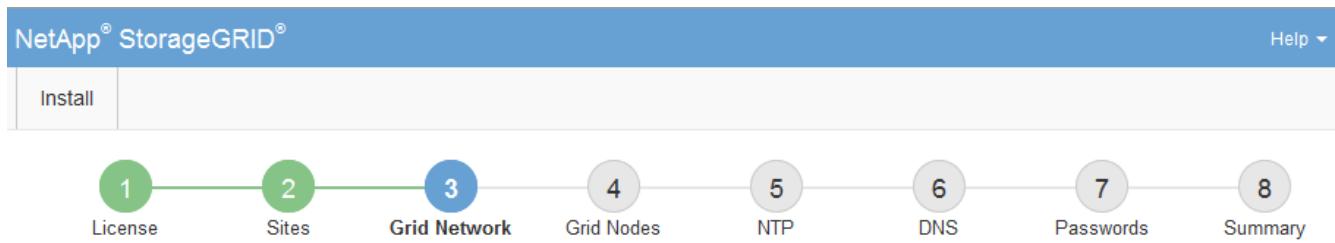
The subnet entries include the subnets for the Grid Network for each site in your StorageGRID system, along with any subnets that need to be reachable via the Grid Network.

If you have multiple grid subnets, the Grid Network gateway is required. All grid subnets specified must be reachable through this gateway.

Steps

1. Specify the CIDR network address for at least one Grid Network in the **Subnet 1** text box.
2. Click the plus sign next to the last entry to add an additional network entry.

If you have already deployed at least one node, click **Discover Grid Networks Subnets** to automatically populate the Grid Network Subnet List with the subnets reported by grid nodes that have registered with the Grid Manager.



Grid Network

You must specify 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 system. Select Discover Grid Networks to automatically add subnets based on the network configuration of all registered nodes.

Note: You must manually add any subnets for NTP, DNS, LDAP, or other external servers accessed through the Grid Network gateway.

Subnet 1

+

Discover Grid Network subnets

3. Click **Next**.

Approve pending grid nodes

You must approve each grid node before it can join the StorageGRID system.

What you'll need

You have deployed all virtual and StorageGRID appliance grid nodes.



It is more efficient to perform one single installation of all the nodes, rather than installing some nodes now and some nodes later.

Steps

1. Review the Pending Nodes list, and confirm that it shows all of the grid nodes you deployed.



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

2. Select the radio button next to a pending node you want to approve.



Grid Nodes

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

Pending Nodes

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

					<input type="button" value="Approve"/>	<input type="button" value="Remove"/>	Search <input type="text"/>
	Grid Network MAC Address	Name	Type	Platform	Grid Network IPv4 Address		
<input checked="" type="radio"/>	50:6b:4b:42:d7:00	NetApp-SGA	Storage Node	StorageGRID Appliance	172.16.5.20/21	<input type="button" value="<"/>	<input type="button" value=">"/>

Approved Nodes

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

					<input type="button" value="Edit"/>	<input type="button" value="Reset"/>	<input type="button" value="Remove"/>	Search <input type="text"/>
	Grid Network MAC Address	Name	Site	Type	Platform	Grid Network IPv4 Address		
<input checked="" type="radio"/>	00:50:56:87:42:ff	dc1-adm1	Raleigh	Admin Node	VMware VM	172.16.4.210/21	<input type="button" value="<"/>	
<input checked="" type="radio"/>	00:50:56:87:c0:16	dc1-s1	Raleigh	Storage Node	VMware VM	172.16.4.211/21	<input type="button" value=">"/>	
<input checked="" type="radio"/>	00:50:56:87:79:ee	dc1-s2	Raleigh	Storage Node	VMware VM	172.16.4.212/21		
<input checked="" type="radio"/>	00:50:56:87:db:9c	dc1-s3	Raleigh	Storage Node	VMware VM	172.16.4.213/21		
<input checked="" type="radio"/>	00:50:56:87:62:38	dc1-g1	Raleigh	API Gateway Node	VMware VM	172.16.4.214/21		

3. Click **Approve**.

4. In General Settings, modify settings for the following properties, as necessary:

Storage Node Configuration

General Settings

Site	<input type="text" value="Raleigh"/>
Name	<input type="text" value="NetApp-SGA"/>
NTP Role	<input type="text" value="Automatic"/>
ADC Service	<input type="text" value="Automatic"/>

Grid Network

Configuration	STATIC
IPv4 Address (CIDR)	<input type="text" value="172.16.5.20/21"/>
Gateway	<input type="text" value="172.16.5.20"/>

Admin Network

Configuration	STATIC
IPv4 Address (CIDR)	<input type="text" value="10.224.5.20/21"/>
Gateway	<input type="text" value="10.224.0.1"/>
Subnets (CIDR)	<input type="text" value="10.0.0.0/8"/> ×
	<input type="text" value="172.19.0.0/16"/> ×
	<input type="text" value="172.21.0.0/16"/> + ×

Client Network

Configuration	STATIC
IPv4 Address (CIDR)	<input type="text" value="47.47.5.20/21"/>
Gateway	<input type="text" value="47.47.0.1"/>

- **Site:** The name of the site with which this grid node will be associated.
- **Name:** The name that will be assigned to the node, and the name that will be displayed in the Grid Manager. The name defaults to the name you specified when you configured the node. During this step of the installation process, you can change the name as required.



After you complete the installation, you cannot change the name of the node.



For a VMware node, you can change the name here, but this action will not change the name of the virtual machine in vSphere.

- **NTP Role:** The Network Time Protocol (NTP) role of the grid node. The options are **Automatic**, **Primary**, and **Client**. Selecting **Automatic** assigns the Primary role to Admin Nodes, Storage Nodes with ADC services, Gateway Nodes, and any grid nodes that have non-static IP addresses. All other grid nodes are assigned the Client role.



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

- **ADC service** (Storage Nodes only): Select **Automatic** to let the system determine whether the node requires the Administrative Domain Controller (ADC) service. The ADC service keeps track of the location and availability of grid services. At least three Storage Nodes at each site must include the ADC service. You cannot add the ADC service to a node after it is deployed.

5. In Grid Network, modify settings for the following properties as necessary:

- **IPv4 Address (CIDR):** The CIDR network address for the Grid Network interface (eth0 inside the container). For example: 192.168.1.1234/21
- **Gateway:** The Grid Network gateway. For example: 192.168.0.1



The gateway is required if there are multiple grid subnets.



If you selected DHCP for the Grid Network configuration and you change the value here, the new value will be configured as a static address on the node. You must make sure the resulting IP address is not within a DHCP address pool.

6. If you want to configure the Admin Network for the grid node, add or update the settings in the Admin Network section as necessary.

Enter the destination subnets of the routes out of this interface in the **Subnets (CIDR)** text box. If there are multiple Admin subnets, the Admin gateway is required.



If you selected DHCP for the Admin Network configuration and you change the value here, the new value will be configured as a static address on the node. You must make sure the resulting IP address is not within a DHCP address pool.

Appliances: For a StorageGRID appliance, if the Admin Network was not configured during the initial installation using the StorageGRID Appliance Installer, it cannot be configured in this Grid Manager dialog box. Instead, you must follow these steps:

- a. Reboot the appliance: In the Appliance Installer, select **Advanced > Reboot**.
Rebooting can take several minutes.
- b. Select **Configure Networking > Link Configuration** and enable the appropriate networks.
- c. Select **Configure Networking > IP Configuration** and configure the enabled networks.
- d. Return to the Home page and click **Start Installation**.
- e. In the Grid Manager: If the node is listed in the Approved Nodes table, reset the node.
- f. Remove the node from the Pending Nodes table.
- g. Wait for the node to reappear in the Pending Nodes list.

- h. Confirm that you can configure the appropriate networks. They should already be populated with the information you provided on the IP Configuration page.

For additional information, see the installation and maintenance instructions for your appliance model.

- 7. If you want to configure the Client Network for the grid node, add or update the settings in the Client Network section as necessary. If the Client Network is configured, the gateway is required, and it becomes the default gateway for the node after installation.



If you selected DHCP for the Client Network configuration and you change the value here, the new value will be configured as a static address on the node. You must make sure the resulting IP address is not within a DHCP address pool.

Appliances: For a StorageGRID appliance, if the Client Network was not configured during the initial installation using the StorageGRID Appliance Installer, it cannot be configured in this Grid Manager dialog box. Instead, you must follow these steps:

- a. Reboot the appliance: In the Appliance Installer, select **Advanced > Reboot**.

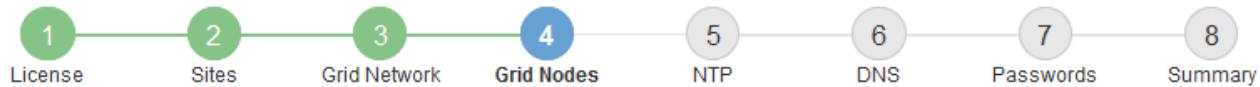
Rebooting can take several minutes.

- b. Select **Configure Networking > Link Configuration** and enable the appropriate networks.
- c. Select **Configure Networking > IP Configuration** and configure the enabled networks.
- d. Return to the Home page and click **Start Installation**.
- e. In the Grid Manager: If the node is listed in the Approved Nodes table, reset the node.
- f. Remove the node from the Pending Nodes table.
- g. Wait for the node to reappear in the Pending Nodes list.
- h. Confirm that you can configure the appropriate networks. They should already be populated with the information you provided on the IP Configuration page.

For additional information, see the installation and maintenance instructions for your appliance.

- 8. Click **Save**.

The grid node entry moves to the Approved Nodes list.



Grid Nodes

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

Pending Nodes

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

		<input type="button" value="Approve"/>	<input type="button" value="Remove"/>			Search	<input type="button" value=""/>
	Grid Network MAC Address	Name	Type	Platform	Grid Network IPv4 Address		
<i>No results found.</i>							
						<input type="button" value=""/>	<input type="button" value=""/>

Approved Nodes

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

		<input type="button" value="Edit"/>	<input type="button" value="Reset"/>	<input type="button" value="Remove"/>			Search	<input type="button" value=""/>
	Grid Network MAC Address	Name	Site	Type	Platform	Grid Network IPv4 Address		
<input type="radio"/>	00:50:56:87:42:ff	dc1-adm1	Raleigh	Admin Node	VMware VM	172.16.4.210/21		
<input type="radio"/>	00:50:56:87:c0:16	dc1-s1	Raleigh	Storage Node	VMware VM	172.16.4.211/21		
<input type="radio"/>	00:50:56:87:79:ee	dc1-s2	Raleigh	Storage Node	VMware VM	172.16.4.212/21		
<input type="radio"/>	00:50:56:87:db:9c	dc1-s3	Raleigh	Storage Node	VMware VM	172.16.4.213/21		
<input type="radio"/>	00:50:56:87:62:38	dc1-g1	Raleigh	API Gateway Node	VMware VM	172.16.4.214/21		
<input type="radio"/>	50:6b:4b:42:d7:00	NetApp-SGA	Raleigh	Storage Node	StorageGRID Appliance	172.16.5.20/21		
							<input type="button" value=""/>	<input type="button" value=""/>

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

You must approve all nodes that you want in the grid. However, you can return to this page at any time before you click **Install** on the Summary page. You can modify the properties of an approved grid node by selecting its radio button and clicking **Edit**.

10. When you are done approving grid nodes, click **Next**.

Specify Network Time Protocol server information

You must specify the Network Time Protocol (NTP) configuration information for the StorageGRID system, so that operations performed on separate servers can be kept synchronized.

About this task

You must specify IPv4 addresses for the NTP servers.

You must specify external NTP servers. The specified NTP servers must use the NTP protocol.

You must specify four NTP server references of Stratum 3 or better to prevent issues with time drift.



When specifying the external NTP source for a production-level StorageGRID installation, do not use the Windows Time (W32Time) service on a version of Windows earlier than Windows Server 2016. The time service on earlier versions of Windows is not sufficiently accurate and is not supported by Microsoft for use in high-accuracy environments, such as StorageGRID.

[Support boundary to configure the Windows Time service for high-accuracy environments](#)

The external NTP servers are used by the nodes to which you previously assigned Primary NTP roles.



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

Perform additional checks for VMware, such as ensuring that the hypervisor uses the same NTP source as the virtual machine, and using VMTools to disable the time sync between the hypervisor and StorageGRID virtual machines.

Steps

1. Specify the IPv4 addresses for at least four NTP servers in the **Server 1 to Server 4** text boxes.
2. If necessary, select the plus sign next to the last entry to add additional server entries.

NetApp® StorageGRID®

Help ▾

Install

1 License 2 Sites 3 Grid Network 4 Grid Nodes 5 NTP 6 DNS 7 Passwords 8 Summary

Network Time Protocol

Enter the IP addresses for at least four Network Time Protocol (NTP) servers, so that operations performed on separate servers are kept in sync.

Server 1	10.60.248.183
Server 2	10.227.204.142
Server 3	10.235.48.111
Server 4	0.0.0.0

+

3. Select **Next**.

Specify Domain Name System server information

You must specify Domain Name System (DNS) information for your StorageGRID system, so that you can access external servers using hostnames instead of IP addresses.

About this task

Specifying DNS server information allows you to use Fully Qualified Domain Name (FQDN) hostnames rather than IP addresses for email notifications and AutoSupport. Specifying at least two DNS servers is recommended.

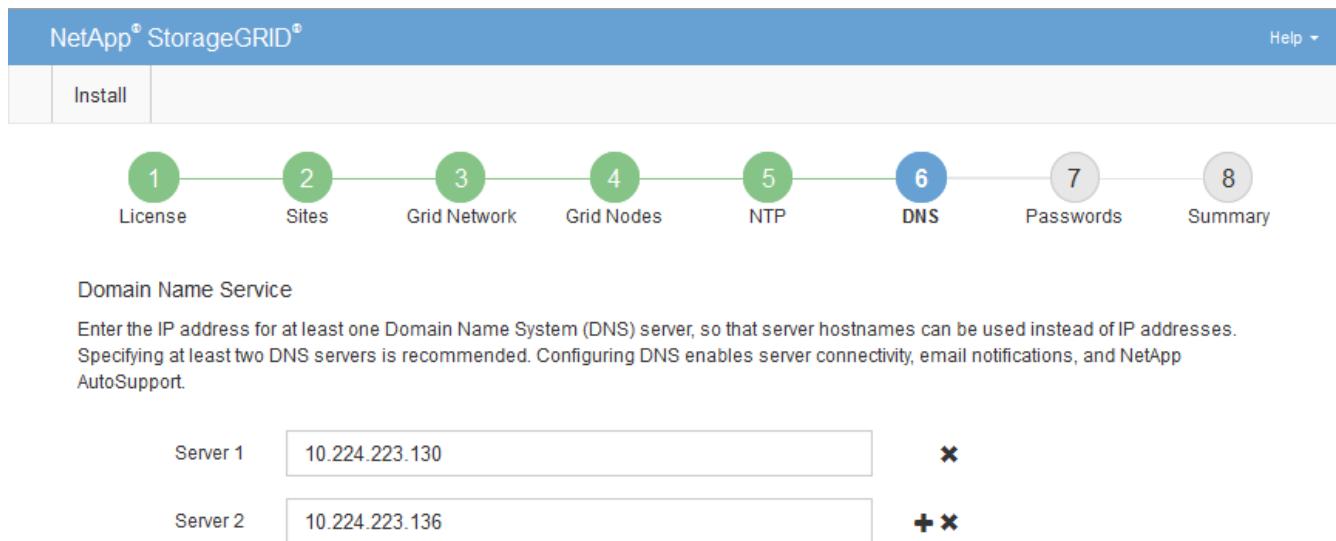


Provide two to six IPv4 addresses for DNS servers. You should select DNS servers that each site can access locally in the event of network islanding. This is to ensure an islanded site continues to have access to the DNS service. After configuring the grid-wide DNS server list, you can further customize the DNS server list for each node. For details, see the information about modifying the DNS configuration in the recovery and maintenance instructions.

If the DNS server information is omitted or incorrectly configured, a DNST alarm is triggered on each grid node's SSM service. The alarm clears when DNS is configured correctly and the new server information has reached all grid nodes.

Steps

1. Specify the IPv4 address for at least one DNS server in the **Server 1** text box.
2. If necessary, select the plus sign next to the last entry to add additional server entries.



The best practice is to specify at least two DNS servers. You can specify up to six DNS servers.

3. Select **Next**.

Related information

[Recover and maintain](#)

Specify the StorageGRID system passwords

As part of installing your StorageGRID system, you need to enter the passwords to use to secure your system and perform maintenance tasks.

About this task

Use the Install passwords page to specify the provisioning passphrase and the grid management root user password.

- The provisioning passphrase is used as an encryption key and is not stored by the StorageGRID system.
- You must have the provisioning passphrase for installation, expansion, and maintenance procedures, including downloading the Recovery Package. Therefore, it is important that you store the provisioning passphrase in a secure location.
- You can change the provisioning passphrase from the Grid Manager if you have the current one.
- The grid management root user password may be changed using the Grid Manager.
- Randomly generated command line console and SSH passwords are stored in the `Passwords.txt` file in the Recovery Package.

Steps

1. In **Provisioning Passphrase**, enter the provisioning passphrase that will be required to make changes to the grid topology of your StorageGRID system.

Store the provisioning passphrase in a secure place.



If after the installation completes and you want to change the provisioning passphrase later, you can use the Grid Manager. Select **CONFIGURATION > Access control > Grid passwords**.

2. In **Confirm Provisioning Passphrase**, reenter the provisioning passphrase to confirm it.
3. In **Grid Management Root User Password**, enter the password to use to access the Grid Manager as the “root” user.

Store the password in a secure place.

4. In **Confirm Root User Password**, reenter the Grid Manager password to confirm it.

NetApp® StorageGRID®

Help ▾

Install

1 License 2 Sites 3 Grid Network 4 Grid Nodes 5 NTP 6 DNS 7 Passwords 8 Summary

Passwords

Enter secure passwords that meet your organization's security policies. A text file containing the command line passwords must be downloaded during the final installation step.

Provisioning Passphrase
Confirm Provisioning Passphrase
Grid Management Root User Password
Confirm Root User Password

Create random command line passwords.

5. If you are installing a grid for proof of concept or demo purposes, optionally deselect the **Create random command line passwords** check box.

For production deployments, random passwords should always be used for security reasons. Deselect **Create random command line passwords** only for demo grids if you want to use default passwords to access grid nodes from the command line using the “root” or “admin” account.



You are prompted to download the Recovery Package file (`sgws-recovery-package-id-revision.zip`) after you click **Install** on the Summary page. You must [download this file](#) to complete the installation. The passwords required to access the system are stored in the `Passwords.txt` file, contained in the Recovery Package file.

6. Click **Next**.

Review your configuration and complete installation

You must carefully review the configuration information you have entered to ensure that the installation completes successfully.

Steps

1. View the **Summary** page.

NetApp® StorageGRID®

Help ▾

Install

1 License 2 Sites 3 Grid Network 4 Grid Nodes 5 NTP 6 DNS 7 Passwords 8 Summary

Summary

Verify that all of the grid configuration information is correct, and then click **Install**. You can view the status of each grid node as it installs. Click the **Modify** links to go back and change the associated information.

General Settings

Grid Name	Grid1	Modify License
Passwords	Auto-generated random command line passwords	Modify Passwords

Networking

NTP	10.60.248.183 10.227.204.142 10.235.48.111	Modify NTP
DNS	10.224.223.130 10.224.223.136	Modify DNS
Grid Network	172.16.0.0/21	Modify Grid Network

Topology

Topology	Atlanta	Modify Sites	Modify Grid Nodes				
	Raleigh	dc1-adm1	dc1-g1	dc1-s1	dc1-s2	dc1-s3	NetApp-SGA

2. Verify that all of the grid configuration information is correct. Use the **Modify** links on the Summary page to go back and correct any errors.

3. Click **Install**.



If a node is configured to use the Client Network, the default gateway for that node switches from the Grid Network to the Client Network when you click **Install**. If you lose connectivity, you must ensure that you are accessing the primary Admin Node through an accessible subnet. See [Networking guidelines](#) for details.

4. Click **Download Recovery Package**.

When the installation progresses to the point where the grid topology is defined, you are prompted to download the Recovery Package file (.zip), and confirm that you can successfully access the contents of this file. You must download the Recovery Package file so that you can recover the StorageGRID system if one or more grid nodes fail. The installation continues in the background, but you cannot complete the installation and access the StorageGRID system until you download and verify this file.

5. Verify that you can extract the contents of the .zip file, and then save it in two safe, secure, and separate locations.



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

6. Select the **I have successfully downloaded and verified the Recovery Package file** check box, and click **Next**.

Download Recovery Package

Before proceeding, you must download the Recovery Package file. This file is necessary to recover the StorageGRID system if a failure occurs.

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.

After you have verified the contents, copy the Recovery Package file to two safe, secure, and separate locations. The Recovery Package file must be secured because it contains encryption keys and passwords that can be used to obtain data from the StorageGRID system.

The Recovery Package is required for recovery procedures and must be stored in a secure location.

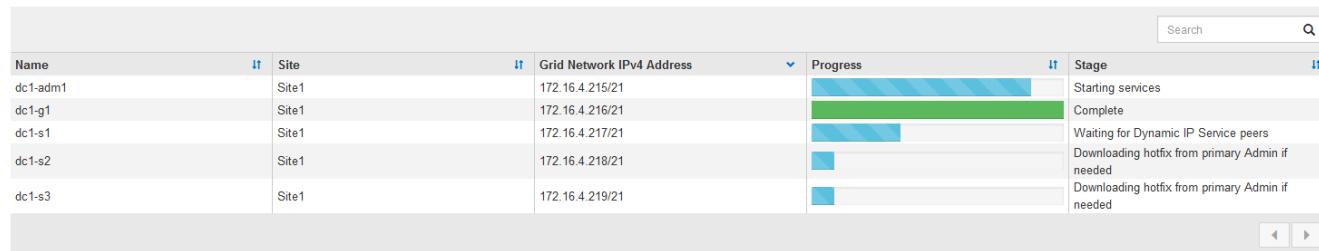
Download Recovery Package

I have successfully downloaded and verified the Recovery Package file.

If the installation is still in progress, the status page appears. This page indicates the progress of the installation for each grid node.

Installation Status

If necessary, you may [Download the Recovery Package file](#) again.



Name	Site	Grid Network IPv4 Address	Progress	Stage
dc1-adm1	Site1	172.16.4.215/21	<div style="width: 80%;"></div>	Starting services
dc1-g1	Site1	172.16.4.216/21	<div style="width: 100%; background-color: green;"></div>	Complete
dc1-s1	Site1	172.16.4.217/21	<div style="width: 70%;"></div>	Waiting for Dynamic IP Service peers
dc1-s2	Site1	172.16.4.218/21	<div style="width: 10%;"></div>	Downloading hotfix from primary Admin if needed
dc1-s3	Site1	172.16.4.219/21	<div style="width: 10%;"></div>	Downloading hotfix from primary Admin if needed

When the Complete stage is reached for all grid nodes, the sign-in page for the Grid Manager appears.

7. Sign in to the Grid Manager using the “root” user and the password you specified during the installation.

Post-installation guidelines

After completing grid node deployment and configuration, follow these guidelines for DHCP addressing and network configuration changes.

- If DHCP was used to assign IP addresses, configure a DHCP reservation for each IP address on the networks being used.

You can only set up DHCP during the deployment phase. You cannot set up DHCP during configuration.



Nodes reboot when their IP addresses change, which can cause outages if a DHCP address change affects multiple nodes at the same time.

- You must use the Change IP procedures if you want to change IP addresses, subnet masks, and default gateways for a grid node. See [Configure IP addresses](#).
- If you make networking configuration changes, including routing and gateway changes, client connectivity to the primary Admin Node and other grid nodes might be lost. Depending on the networking changes applied, you might need to re-establish these connections.

Automate the installation (VMware)

You can use VMware vSphere to automate the deployment of grid nodes. You can also automate the configuration of StorageGRID.

Automate grid node deployment

Use VMware vSphere to automate the deployment of grid nodes.

What you'll need

- You have access to a Linux/Unix system with Bash 3.2 or later.
- You have VMware OVF Tool 4.1 installed and correctly configured.
- You know the username and password required to access VMware vSphere using the OVF Tool.
- You know the virtual infrastructure (VI) URL for the location in vSphere where you want to deploy the StorageGRID virtual machines. This URL will typically be a vApp, or Resource Pool. For example: `vi://vcenter.example.com/vi/sgws`



You can use the VMware ovftool utility to determine this value (see the ovftool documentation for details).



If you are deploying to a vApp, the virtual machines will not start automatically the first time, and you must power them on manually.

- You have collected all the required information for the configuration file. See [Collect information about your deployment environment](#) for information.
- You have access to the following files from the VMware installation archive for StorageGRID:

Filename	Description
NetApp-SG-version-SHA.vmdk	<p>The virtual machine disk file that is used as a template for creating grid node virtual machines.</p> <p>Note: This file must be in the same folder as the .ovf and .mf files.</p>
vsphere-primary-admin.ovf vsphere-primary-admin.mf	<p>The Open Virtualization Format template file (.ovf) and manifest file (.mf) for deploying the primary Admin Node.</p>
vsphere-non-primary-admin.ovf vsphere-non-primary-admin.mf	<p>The template file (.ovf) and manifest file (.mf) for deploying non-primary Admin Nodes.</p>
vsphere-archive.ovf vsphere-archive.mf	<p>The template file (.ovf) and manifest file (.mf) for deploying Archive Nodes.</p>
vsphere-gateway.ovf vsphere-gateway.mf	<p>The template file (.ovf) and manifest file (.mf) for deploying Gateway Nodes.</p>
vsphere-storage.ovf vsphere-storage.mf	<p>The template file (.ovf) and manifest file (.mf) for deploying virtual machine-based Storage Nodes.</p>
deploy-vsphere-ovftool.sh	<p>The Bash shell script used to automate the deployment of virtual grid nodes.</p>
deploy-vsphere-ovftool-sample.ini	<p>The sample configuration file for use with the deploy-vsphere-ovftool.sh script.</p>

Define the configuration file for your deployment

You specify the information needed to deploy virtual grid nodes for StorageGRID in a configuration file, which is used by the `deploy-vsphere-ovftool.sh` Bash script. You can modify a sample configuration file, so that you do not have to create the file from scratch.

Steps

1. Make a copy of the sample configuration file (`deploy-vsphere-ovftool.sample.ini`). Save the new

file as `deploy-vsphere-ovftool.ini` in the same directory as `deploy-vsphere-ovftool.sh`.

2. Open `deploy-vsphere-ovftool.ini`.
3. Enter all of the information required to deploy VMware virtual grid nodes.

See [Configuration file settings](#) for information.

4. When you have entered and verified all of the necessary information, save and close the file.

Configuration file settings

The `deploy-vsphere-ovftool.ini` configuration file contains the settings that are required to deploy virtual grid nodes.

The configuration file first lists global parameters, and then lists node-specific parameters in sections defined by node name. When the file is used:

- *Global parameters* are applied to all grid nodes.
- *Node-specific parameters* override global parameters.

Global parameters

Global parameters are applied to all grid nodes, unless they are overridden by settings in individual sections. Place the parameters that apply to multiple nodes in the global parameter section, and then override these settings as necessary in the sections for individual nodes.

- **OVFTOOL_ARGUMENTS:** You can specify OVFTOOL_ARGUMENTS as global settings, or you can apply arguments individually to specific nodes. For example:

```
OVFTOOL_ARGUMENTS = --powerOn --noSSLVerify --diskMode=eagerZeroedThick  
--datastore='datastore_name'
```

You can use the `--powerOffTarget` and `--overwrite` options to shut down and replace existing virtual machines.



You should deploy nodes to different datastores and specify OVFTOOL_ARGUMENTS for each node, instead of globally.

- **SOURCE:** The path to the StorageGRID virtual machine template (.vmdk) file and the .ovf and .mf files for individual grid nodes. This defaults to the current directory.

```
SOURCE = /downloads/StorageGRID-Webscale-version/vsphere
```

- **TARGET:** The VMware vSphere virtual infrastructure (vi) URL for the location where StorageGRID will be deployed. For example:

```
TARGET = vi://vcenter.example.com/vm/sgws
```

- **GRID_NETWORK_CONFIG**: The method used to acquire IP addresses, either STATIC or DHCP. The default is STATIC. If all or most of the nodes use the same method for acquiring IP addresses, you can specify that method here. You can then override the global setting by specifying different settings for one or more individual nodes. For example:

```
GRID_NETWORK_CONFIG = DHCP
```

- **GRID_NETWORK_TARGET**: The name of an existing VMware network to use for the Grid Network. If all or most of the nodes use the same network name, you can specify it here. You can then override the global setting by specifying different settings for one or more individual nodes. For example:

```
GRID_NETWORK_TARGET = SG-Admin-Network
```

- **GRID_NETWORK_MASK**: The network mask for the Grid Network. If all or most of the nodes use the same network mask, you can specify it here. You can then override the global setting by specifying different settings for one or more individual nodes. For example:

```
GRID_NETWORK_MASK = 255.255.255.0
```

- **GRID_NETWORK_GATEWAY**: The network gateway for the Grid Network. If all or most of the nodes use the same network gateway, you can specify it here. You can then override the global setting by specifying different settings for one or more individual nodes. For example:

```
GRID_NETWORK_GATEWAY = 10.1.0.1
```

- **GRID_NETWORK_MTU**: Optional. The maximum transmission unit (MTU) on the Grid Network. If specified, the value must be between 1280 and 9216. For example:

```
GRID_NETWORK_MTU = 8192
```

If omitted, 1400 is used.

If you want to use jumbo frames, set the MTU to a value suitable for jumbo frames, such as 9000. Otherwise, keep the default value.



The MTU value of the network must match the value configured on the switch port the node is connected to. Otherwise, network performance issues or packet loss might occur.



For the best network performance, all nodes should be configured with similar MTU values on their Grid Network interfaces. The **Grid Network MTU mismatch** alert is triggered if there is a significant difference in MTU settings for the Grid Network on individual nodes. The MTU values do not have to be the same for all network types.

- **ADMIN_NETWORK_CONFIG**: The method used to acquire IP addresses, either DISABLED, STATIC, or DHCP. The default is DISABLED. If all or most of the nodes use the same method for acquiring IP

addresses, you can specify that method here. You can then override the global setting by specifying different settings for one or more individual nodes. For example:

```
ADMIN_NETWORK_CONFIG = STATIC
```

- **ADMIN_NETWORK_TARGET**: The name of an existing VMware network to use for the Admin Network. This setting is required unless the Admin Network is disabled. If all or most of the nodes use the same network name, you can specify it here. You can then override the global setting by specifying different settings for one or more individual nodes. For example:

```
ADMIN_NETWORK_TARGET = SG-Admin-Network
```

- **ADMIN_NETWORK_MASK**: The network mask for the Admin Network. This setting is required if you are using static IP addressing. If all or most of the nodes use the same network mask, you can specify it here. You can then override the global setting by specifying different settings for one or more individual nodes. For example:

```
ADMIN_NETWORK_MASK = 255.255.255.0
```

- **ADMIN_NETWORK_GATEWAY**: The network gateway for the Admin Network. This setting is required if you are using static IP addressing and you specify external subnets in the **ADMIN_NETWORK_ESL** setting. (That is, it is not required if **ADMIN_NETWORK_ESL** is empty.) If all or most of the nodes use the same network gateway, you can specify it here. You can then override the global setting by specifying different settings for one or more individual nodes. For example:

```
ADMIN_NETWORK_GATEWAY = 10.3.0.1
```

- **ADMIN_NETWORK_ESL**: The external subnet list (routes) for the Admin Network, specified as a comma-separated list of CIDR route destinations. If all or most of the nodes use the same external subnet list, you can specify it here. You can then override the global setting by specifying different settings for one or more individual nodes. For example:

```
ADMIN_NETWORK_ESL = 172.16.0.0/21,172.17.0.0/21
```

- **ADMIN_NETWORK_MTU**: Optional. The maximum transmission unit (MTU) on the Admin Network. Do not specify if **ADMIN_NETWORK_CONFIG** = **DHCP**. If specified, the value must be between 1280 and 9216. If omitted, 1400 is used. If you want to use jumbo frames, set the MTU to a value suitable for jumbo frames, such as 9000. Otherwise, keep the default value. If all or most of the nodes use the same MTU for the Admin Network, you can specify it here. You can then override the global setting by specifying different settings for one or more individual nodes. For example:

```
ADMIN_NETWORK_MTU = 8192
```

- **CLIENT_NETWORK_CONFIG**: The method used to acquire IP addresses, either **DISABLED**, **STATIC**, or

DHCP. The default is DISABLED. If all or most of the nodes use the same method for acquiring IP addresses, you can specify that method here. You can then override the global setting by specifying different settings for one or more individual nodes. For example:

```
CLIENT_NETWORK_CONFIG = STATIC
```

- **CLIENT_NETWORK_TARGET**: The name of an existing VMware network to use for the Client Network. This setting is required unless the Client Network is disabled. If all or most of the nodes use the same network name, you can specify it here. You can then override the global setting by specifying different settings for one or more individual nodes. For example:

```
CLIENT_NETWORK_TARGET = SG-Client-Network
```

- **CLIENT_NETWORK_MASK**: The network mask for the Client Network. This setting is required if you are using static IP addressing. If all or most of the nodes use the same network mask, you can specify it here. You can then override the global setting by specifying different settings for one or more individual nodes. For example:

```
CLIENT_NETWORK_MASK = 255.255.255.0
```

- **CLIENT_NETWORK_GATEWAY**: The network gateway for the Client Network. This setting is required if you are using static IP addressing. If all or most of the nodes use the same network gateway, you can specify it here. You can then override the global setting by specifying different settings for one or more individual nodes. For example:

```
CLIENT_NETWORK_GATEWAY = 10.4.0.1
```

- **CLIENT_NETWORK_MTU**: Optional. The maximum transmission unit (MTU) on the Client Network. Do not specify if CLIENT_NETWORK_CONFIG = DHCP. If specified, the value must be between 1280 and 9216. If omitted, 1400 is used. If you want to use jumbo frames, set the MTU to a value suitable for jumbo frames, such as 9000. Otherwise, keep the default value. If all or most of the nodes use the same MTU for the Client Network, you can specify it here. You can then override the global setting by specifying different settings for one or more individual nodes. For example:

```
CLIENT_NETWORK_MTU = 8192
```

- **PORT_REMAP**: Remaps any port used by a node for internal grid node communications or external communications. Remapping ports is necessary if enterprise networking policies restrict one or more ports used by StorageGRID. For the list of ports used by StorageGRID, see internal grid node communications and external communications in [Networking guidelines](#).



Do not remap the ports you are planning to use to configure load balancer endpoints.



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

The format used is: *network type/protocol/default port used by grid node/new port*, where network type is grid, admin, or client, and protocol is tcp or udp.

For example:

```
PORT_REMAP = client/tcp/18082/443
```

If used alone, this example setting symmetrically maps both inbound and outbound communications for the grid node from port 18082 to port 443. If used in conjunction with PORT_REMAP_INBOUND, this example setting maps outbound communications from port 18082 to port 443.

- **PORT_REMAP_INBOUND:** Remaps inbound communications for the specified port. If you specify PORT_REMAP_INBOUND but do not specify a value for PORT_REMAP, outbound communications for the port are unchanged.



Do not remap the ports you are planning to use to configure load balancer endpoints.

The format used is: *network type/protocol/_default port used by grid node/new port*, where network type is grid, admin, or client, and protocol is tcp or udp.

For example:

```
PORT_REMAP_INBOUND = client/tcp/443/18082
```

This example takes traffic that is sent to port 443 to pass an internal firewall and directs it to port 18082, where the grid node is listening for S3 requests.

Node-specific parameters

Each node is in its own section of the configuration file. Each node requires the following settings:

- The section head defines the node name that will be displayed in the Grid Manager. You can override that value by specifying the optional NODE_NAME parameter for the node.
- **NODE_TYPE:** VM_Admin_Node, VM_Storage_Node, VM_Archive_Node, or VM_API_Gateway_Node
- **GRID_NETWORK_IP:** The IP address for the node on the Grid Network.
- **ADMIN_NETWORK_IP:** The IP address for the node on the Admin Network. Required only if the node is attached to the Admin Network and ADMIN_NETWORK_CONFIG is set to STATIC.
- **CLIENT_NETWORK_IP:** The IP address for the node on the Client Network. Required only if the node is attached to the Client Network and CLIENT_NETWORK_CONFIG for this node is set to STATIC.
- **ADMIN_IP:** The IP address for the primary Admin node on the Grid Network. Use the value that you specify as the GRID_NETWORK_IP for the primary Admin Node. If you omit this parameter, the node attempts to discover the primary Admin Node IP using mDNS. For more information, see [How grid nodes discover the primary Admin Node](#).



The ADMIN_IP parameter is ignored for the primary Admin Node.

- Any parameters that were not set globally. For example, if a node is attached to the Admin Network and you did not specify ADMIN_NETWORK parameters globally, you must specify them for the node.

Primary Admin Node

The following additional settings are required for the primary Admin Node:

- **NODE_TYPE**: VM_Admin_Node
- **ADMIN_ROLE**: Primary

This example entry is for a primary Admin Node that is on all three networks:

```
[DC1-ADM1]
ADMIN_ROLE = Primary
NODE_TYPE = VM_Admin_Node

GRID_NETWORK_IP = 10.1.0.2
ADMIN_NETWORK_IP = 10.3.0.2
CLIENT_NETWORK_IP = 10.4.0.2
```

The following additional setting is optional for the primary Admin Node:

- **DISK**: By default, Admin Nodes are assigned two additional 200 GB hard disks for audit and database use. You can increase these settings using the DISK parameter. For example:

```
DISK = INSTANCES=2, CAPACITY=300
```



For Admin nodes, INSTANCES must always equal 2.

Storage Node

The following additional setting is required for Storage Nodes:

- **NODE_TYPE**: VM_Storage_Node

This example entry is for a Storage Node that is on the Grid and Admin Networks, but not on the Client Network. This node uses the ADMIN_IP setting to specify the primary Admin Node's IP address on the Grid Network.

```
[DC1-S1]
NODE_TYPE = VM_Storage_Node

GRID_NETWORK_IP = 10.1.0.3
ADMIN_NETWORK_IP = 10.3.0.3

ADMIN_IP = 10.1.0.2
```

This second example entry is for a Storage Node on a Client Network where the customer's enterprise networking policy states that an S3 client application is only permitted to access the Storage Node using either port 80 or 443. The example configuration file uses PORT_REMAP to enable the Storage Node to send and receive S3 messages on port 443.

```
[DC2-S1]
NODE_TYPE = VM_Storage_Node

GRID_NETWORK_IP = 10.1.1.3
CLIENT_NETWORK_IP = 10.4.1.3
PORT_REMAP = client/tcp/18082/443

ADMIN_IP = 10.1.0.2
```

The last example creates a symmetric remapping for ssh traffic from port 22 to port 3022, but explicitly sets the values for both inbound and outbound traffic.

```
[DC1-S3]
NODE_TYPE = VM_Storage_Node

GRID_NETWORK_IP = 10.1.1.3

PORT_REMAP = grid/tcp/22/3022
PORT_REMAP_INBOUND = grid/tcp/3022/22

ADMIN_IP = 10.1.0.2
```

The following additional setting is optional for Storage Nodes:

- **DISK:** By default, Storage Nodes are assigned three 4 TB disks for RangeDB use. You can increase these settings with the DISK parameter. For example:

```
DISK = INSTANCES=16, CAPACITY=4096
```

Archive Node

The following additional setting is required for Archive Nodes:

- **NODE_TYPE**: VM_Archive_Node

This example entry is for an Archive Node that is on the Grid and Admin Networks, but not on the Client Network.

```
[DC1-ARC1]
NODE_TYPE = VM_Archive_Node

GRID_NETWORK_IP = 10.1.0.4
ADMIN_NETWORK_IP = 10.3.0.4

ADMIN_IP = 10.1.0.2
```

Gateway Node

The following additional setting is required for Gateway Nodes:

- **NODE_TYPE**: VM_API_Gateway

This example entry is for an example Gateway Node on all three networks. In this example, no Client Network parameters were specified in the global section of the configuration file, so they must be specified for the node:

```
[DC1-G1]
NODE_TYPE = VM_API_Gateway

GRID_NETWORK_IP = 10.1.0.5
ADMIN_NETWORK_IP = 10.3.0.5

CLIENT_NETWORK_CONFIG = STATIC
CLIENT_NETWORK_TARGET = SG-Client-Network
CLIENT_NETWORK_MASK = 255.255.255.0
CLIENT_NETWORK_GATEWAY = 10.4.0.1
CLIENT_NETWORK_IP = 10.4.0.5

ADMIN_IP = 10.1.0.2
```

Non-primary Admin Node

The following additional settings are required for non-primary Admin Nodes:

- **NODE_TYPE**: VM_Admin_Node
- **ADMIN_ROLE**: Non-Primary

This example entry is for a non-primary Admin Node that is not on the Client Network:

```
[DC2-ADM1]
ADMIN_ROLE = Non-Primary
NODE_TYPE = VM_Admin_Node

GRID_NETWORK_TARGET = SG-Grid-Network
GRID_NETWORK_IP = 10.1.0.6
ADMIN_NETWORK_IP = 10.3.0.6

ADMIN_IP = 10.1.0.2
```

The following additional setting is optional for non-primary Admin Nodes:

- **DISK**: By default, Admin Nodes are assigned two additional 200 GB hard disks for audit and database use. You can increase these settings using the DISK parameter. For example:

```
DISK = INSTANCES=2, CAPACITY=300
```



For Admin nodes, INSTANCES must always equal 2.

Run the Bash script

You can use the `deploy-vsphere-ovftool.sh` Bash script and the `deploy-vsphere-ovftool.ini` configuration file you modified to automate the deployment of StorageGRID grid nodes in VMware vSphere.

What you'll need

- You have created a `deploy-vsphere-ovftool.ini` configuration file for your environment.

You can use the help available with the Bash script by entering the help commands (`-h`/`--help`). For example:

```
./deploy-vsphere-ovftool.sh -h
```

or

```
./deploy-vsphere-ovftool.sh --help
```

Steps

1. Log in to the Linux machine you are using to run the Bash script.
2. Change to the directory where you extracted the installation archive.

For example:

```
cd StorageGRID-Webscale-version/vsphere
```

3. To deploy all grid nodes, run the Bash script with the appropriate options for your environment.

For example:

```
./deploy-vsphere-ovftool.sh --username=user --password=pwd ./deploy-vsphere-ovftool.ini
```

4. If a grid node failed to deploy because of an error, resolve the error and rerun the Bash script for only that node.

For example:

```
./deploy-vsphere-ovftool.sh --username=user --password=pwd --single -node="DC1-S3" ./deploy-vsphere-ovftool.ini
```

The deployment is complete when the status for each node is “Passed.”

Deployment Summary

node	attempts	status
DC1-ADM1	1	Passed
DC1-G1	1	Passed
DC1-S1	1	Passed
DC1-S2	1	Passed
DC1-S3	1	Passed

Automate the configuration of StorageGRID

After deploying the grid nodes, you can automate the configuration of the StorageGRID system.

What you'll need

- You know the location of the following files from the installation archive.

Filename	Description
configure-storagegrid.py	Python script used to automate the configuration
configure-storagegrid.sample.json	Sample configuration file for use with the script

Filename	Description
configure-storagegrid.blank.json	Blank configuration file for use with the script

- You have created a `configure-storagegrid.json` configuration file. To create this file, you can modify the sample configuration file (`configure-storagegrid.sample.json`) or the blank configuration file (`configure-storagegrid.blank.json`).

You can use the `configure-storagegrid.py` Python script and the `configure-storagegrid.json` configuration file to automate the configuration of your StorageGRID system.



You can also configure the system using the Grid Manager or the Installation API.

Steps

1. Log in to the Linux machine you are using to run the Python script.
2. Change to the directory where you extracted the installation archive.

For example:

```
cd StorageGRID-Webscale-version/platform
```

where `platform` is `debs`, `rpm`s, or `vsphere`.

3. Run the Python script and use the configuration file you created.

For example:

```
./configure-storagegrid.py ./configure-storagegrid.json --start-install
```

Result

A Recovery Package .zip file is generated during the configuration process, and it is downloaded to the directory where you are running the installation and configuration process. You must back up the Recovery Package file so that you can recover the StorageGRID system if one or more grid nodes fails. For example, copy it to a secure, backed up network location and to a secure cloud storage location.



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

If you specified that random passwords should be generated, you need to extract the `Passwords.txt` file and look for the passwords required to access your StorageGRID system.

```
#####
##### The StorageGRID "recovery package" has been downloaded as: #####
#####           ./sgws-recovery-package-994078-rev1.zip           #####
##### Safeguard this file as it will be needed in case of a      #####
#####           StorageGRID node recovery.                      #####
#####
```

Your StorageGRID system is installed and configured when a confirmation message is displayed.

StorageGRID has been configured and installed.

Related information

[Navigate to the Grid Manager](#)

[Overview of the installation REST API](#)

Overview of the installation REST API

StorageGRID provides the StorageGRID Installation API for performing installation tasks.

The API uses the Swagger open source API platform to provide the API documentation. Swagger allows both developers and non-developers to interact with the API in a user interface that illustrates how the API responds to parameters and options. This documentation assumes that you are familiar with standard web technologies and the JSON (JavaScript Object Notation) data format.



Any API operations you perform using the API Docs webpage are live operations. Be careful not to create, update, or delete configuration data or other data by mistake.

Each REST API command includes the API's URL, an HTTP action, any required or optional URL parameters, and an expected API response.

StorageGRID Installation API

The StorageGRID Installation API is only available when you are initially configuring your StorageGRID system, and in the event that you need to perform a primary Admin Node recovery. The Installation API can be accessed over HTTPS from the Grid Manager.

To access the API documentation, go to the installation web page on the primary Admin Node and select **Help > API Documentation** from the menu bar.

The StorageGRID Installation API includes the following sections:

- **config** — Operations related to the product release and versions of the API. You can list the product release version and the major versions of the API supported by that release.
- **grid** — Grid-level configuration operations. You can get and update grid settings, including grid details, Grid Network subnets, grid passwords, and NTP and DNS server IP addresses.
- **nodes** — Node-level configuration operations. You can retrieve a list of grid nodes, delete a grid node, configure a grid node, view a grid node, and reset a grid node's configuration.

- **provision** — Provisioning operations. You can start the provisioning operation and view the status of the provisioning operation.
- **recovery** — Primary Admin Node recovery operations. You can reset information, upload the Recover Package, start the recovery, and view the status of the recovery operation.
- **recovery-package** — Operations to download the Recovery Package.
- **schemas** — API schemas for advanced deployments
- **sites** — Site-level configuration operations. You can create, view, delete, and modify a site.

Where to go next

After completing an installation, you must perform a series of integration and configuration steps. Some steps are required; others are optional.

Required tasks

- Configure VMware vSphere Hypervisor for automatic restart.

You must configure the hypervisor to restart the virtual machines when the server restarts. Without an automatic restart, the virtual machines and grid nodes remain shut down after the server restarts. For details, see the VMware vSphere Hypervisor documentation.

- Create a tenant account for each client protocol (Swift or S3) that will be used to store objects on your StorageGRID system.
- Control system access by configuring groups and user accounts. Optionally, you can configure a federated identity source (such as Active Directory or OpenLDAP), so you can import administration groups and users. Or, you can create local groups and users.
- Integrate and test the S3 or Swift API client applications you will use to upload objects to your StorageGRID system.
- When you are ready, configure the information lifecycle management (ILM) rules and ILM policy you want to use to protect object data.



When you install StorageGRID, the default ILM policy, Baseline 2 Copies Policy, is active. This policy includes the stock ILM rule (Make 2 Copies), and it applies if no other policy has been activated.

- If your installation includes appliance Storage Nodes, use SANtricity software to complete the following tasks:
 - Connect to each StorageGRID appliance.
 - Verify receipt of AutoSupport data.
- If your StorageGRID system includes any Archive Nodes, configure the Archive Node's connection to the target external archival storage system.



If any Archive Nodes will use Tivoli Storage Manager as the external archival storage system, you must also configure Tivoli Storage Manager.

- Review and follow the StorageGRID system hardening guidelines to eliminate security risks.
- Configure email notifications for system alerts.

Optional tasks

- If you want to receive notifications from the (legacy) alarm system, configure mailing lists and email notifications for alarms.
- Update grid node IP addresses if they have changed since you planned your deployment and generated the Recovery Package. See information about changing IP addresses in the recovery and maintenance instructions.
- Configure storage encryption, if required.
- Configure storage compression to reduce the size of stored objects, if required.
- Configure audit client access. You can configure access to the system for auditing purposes through an NFS or a CIFS file share. See the instructions for administering StorageGRID.



Audit export through CIFS/Samba has been deprecated and will be removed in a future StorageGRID release.

Related information

[Administer StorageGRID](#)

[Use S3](#)

[Use Swift](#)

[Manage objects with ILM](#)

[Monitor and troubleshoot](#)

[Recover and maintain](#)

[SG100 and SG1000 services appliances](#)

[SG5600 storage appliances](#)

[SG5700 storage appliances](#)

[SG6000 storage appliances](#)

[Release notes](#)

[System hardening](#)

[Review audit logs](#)

[Upgrade software](#)

Troubleshoot installation issues

If any problems occur while installing your StorageGRID system, you can access the installation log files.

The following are the main installation log files, which technical support might need to resolve issues.

- `/var/local/log/install.log` (found on all grid nodes)

- `/var/local/log/gdu-server.log` (found on the primary Admin Node)

To learn how to access the log files, see the [instructions for monitoring and troubleshooting StorageGRID](#). For help troubleshooting appliance installation issues, see the installation and maintenance instructions for your appliances. If you need additional help, contact technical support.

[SG100 and SG1000 services appliances](#)

[SG6000 storage appliances](#)

[SG5700 storage appliances](#)

[SG5600 storage appliances](#)

[NetApp Support](#)

Virtual machine resource reservation requires adjustment

OVF files include a resource reservation designed to ensure that each grid node has sufficient RAM and CPU to operate efficiently. If you create virtual machines by deploying these OVF files on VMware and the predefined number of resources are not available, the virtual machines will not start.

About this task

If you are certain that the VM host has sufficient resources for each grid node, manually adjust the resources allocated for each virtual machine, and then try starting the virtual machines.

Steps

1. In the VMware vSphere Hypervisor client tree, select the virtual machine that is not started.
2. Right-click the virtual machine, and select **Edit Settings**.
3. From the Virtual Machines Properties window, select the **Resources** tab.
4. Adjust the resources allocated to the virtual machine:
 - a. Select **CPU**, and then use the Reservation slider to adjust the MHz reserved for this virtual machine.
 - b. Select **Memory**, and then use the Reservation slider to adjust the MB reserved for this virtual machine.
5. Click **OK**.
6. Repeat as required for other virtual machines hosted on the same VM host.

Copyright information

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

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

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

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

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

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

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

Trademark information

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