



Upgrade StorageGRID software

StorageGRID 11.7

NetApp
April 10, 2024

Table of Contents

- Upgrade StorageGRID software 1
 - Upgrade StorageGRID software: Overview 1
 - What's new in StorageGRID 11.7 1
 - Removed or deprecated features 5
 - Changes to the Grid Management API 7
 - Changes to the Tenant Management API 8
 - Plan and prepare for upgrade 8
 - Upgrade software 16
 - Troubleshoot upgrade issues 24

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 11.7 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 11.7, determine whether any features have been deprecated or removed, and find out about changes to StorageGRID APIs.

- [What's new in StorageGRID 11.7](#)
- [Removed or deprecated features](#)
- [Changes to the Grid Management API](#)
- [Changes to the Tenant Management API](#)

What's new in StorageGRID 11.7

This release of StorageGRID introduces the following features and functional changes.

New features

Grid federation

You can configure a grid federation connection between two StorageGRID systems to clone tenant account information and replicate bucket objects between the grids for disaster recovery. See [What is grid federation?](#), [What is account clone](#), and [What is cross-grid replication](#).

Improved read availability

The read-after-new-write (default) consistency control was improved to be more available. GET/HEAD requests for non-existent objects will succeed with up to one Storage Node offline at each site. Buckets are no longer required to be set to the Available consistency control for this scenario. For example, applications that check existence of an object before creation will properly function with read-after-new-write even during software upgrade when one Storage Node is offline.

Rename grid, sites, and nodes

A new maintenance procedure lets you change the display names that are shown throughout the Grid Manager. You can update display names safely and whenever you need. See [Rename grid, sites, and nodes](#).

FabricPool and S3 setup wizard

The FabricPool and S3 setup wizard guides you through each step of configuring StorageGRID for use with ONTAP FabricPool or other S3 client application and produces a file you can use when entering required

values in the other application. See [Use FabricPool setup wizard](#) and [Use S3 setup wizard](#).

Related to this change, a banner is now displayed on the dashboard to remind new users to configure [S3 endpoint domain names](#) for S3 virtual-hosted-style requests and set up [email notifications for alerts](#).

Firewall controls

The Firewall control page enables you to manage the external access of ports on nodes in your grid, and to define host addresses and IP subnets that are allowed access to closed ports. The new page also includes the untrusted Client Network settings, which now allow you to select additional ports you want open when untrusted Client Network is configured. See [Configure internal firewall](#).

Enhanced security policies

You can now determine which protocols and ciphers are used to establish secure TLS connections with client applications and secure SSH connections to internal StorageGRID services. See [Manage the TLS and SSH policy](#).

Load balancer endpoint changes

When [configuring load balancer endpoints](#), you can now:

- Allow all tenants to access the endpoint (default), or specify a list of allowed or blocked tenants to provide better security isolation between tenants and their endpoints.
- Use the **Node Type** binding mode to require clients to use the IP address (or corresponding FQDN) of any Admin Node or the IP address of any Gateway Node, based on the type of node you select.

SGF6112 all-flash appliance

The new StorageGRID SGF6112 storage appliance features a compact design with compute controller and storage controller integrated into a 1U chassis. The appliance supports 12 SSD NVMe drives with a storage capacity of up to 15.3 TB per drive. The SSD drives are in a RAID that provides resilient object storage. See [SGF6112 appliance: Overview](#).

Other Grid Manager enhancements

ILM enhancements

The improved ILM wizard makes it easier to specify filters, enter time periods and placements, and view retention diagrams. Erasure-coding profiles are created automatically when you select a storage pool and an EC scheme for a placement. For new StorageGRID 11.7 installations (not upgrades), a storage pool is automatically created for each site and the new **1 Copy Per Site** default rule ensures that new multi-site installations will have site-loss protection by default. See [Manage objects with ILM](#).

Customizable dashboard

You can now configure custom dashboards for the Grid Manager. See [View and manage the dashboard](#).

Volume restoration UI

Storage volume restoration lets you restore object data if a storage volume fails. For StorageGRID 11.7, you can start volume restoration from Grid Manager in addition to the existing method of entering commands manually. Using Grid Manager is now the preferred method for restoring object data. See [Restore object data using Grid Manager](#).

Upgrade and hotfix UI

When you upgrade to StorageGRID 11.7, you can apply the latest 11.7 hotfix at the same time. The StorageGRID upgrade page shows the recommended upgrade path and links directly to the correct download pages. See [Perform upgrade](#).

Units for storage values

You can now select base 10 or base 2 units for the storage values displayed in the Grid Manager and Tenant Manager. Select the user drop-down in the upper right of the Grid Manager or Tenant Manager, then select **User preferences**.

Access MIB from Grid Manager

You can now access SNMP-compliant MIB files from the Grid Manager using the SNMP agent page. See [Access MIB files](#).

Custom storage grades for new nodes

When you perform an expansion to add a new site or new Storage Nodes, you can now assign a custom storage grade to each new node. See [Perform expansion](#).

Tenant Manager updates

Cross-grid replication

Tenant accounts that have permission to use a [grid federation connection](#) can clone tenant groups, users, and S3 keys from one grid to another and use cross-grid replication to replicate bucket objects between two grids. See [Clone tenant groups and users](#) and [Manage cross-grid replication](#).

Delete all objects from bucket

Tenant Manager users can now delete all objects in a bucket, so the bucket can be deleted. See [Delete objects in bucket](#).

S3 Object Lock default retention

Tenant Manager users can now enable and configure default retention when creating S3 Object Lock buckets. See [Create an S3 bucket](#).

S3 updates

S3 Object Lock governance mode

When specifying the S3 Object Lock settings for an object or the default retention settings for a bucket, you can now use governance mode. This retention mode allows users with special permission to bypass certain retention settings. See [Use S3 Object Lock to retain objects](#) and [Use S3 REST API to configure S3 Object Lock](#).

S3 group policy for ransomware mitigation

When added as the group policy for an S3 tenant account, the sample policy helps mitigate ransomware attacks. It prevents older object versions from being permanently deleted. See [Create groups for an S3 tenant](#).

NewerNoncurrentVersions threshold for S3 buckets

The `NewerNoncurrentVersions` action in the bucket lifecycle configuration specifies the number of noncurrent versions retained in a versioned S3 bucket. This threshold overrides lifecycle rules provided by ILM. See [How objects are deleted](#).

S3 Select updates

S3 `SelectObjectContent` now offers support for Parquet objects. In addition, you can now use S3 Select with Admin and Gateway load balancer endpoints that are bare metal nodes running a kernel with cgroup v2 enabled. See [S3 SelectObjectContent](#).

Other enhancements

Certificate subject optional

The certificate subject field is now optional. If this field is left blank, the generated certificate uses the first domain name or IP address as the subject common name (CN). See [Manage security certificates](#).

ILM audit message category and new messages

An audit message category was added for ILM operations and includes the IDEL, LKCU, and ORLM messages. This new category is set to **Normal**. See [ILM operations audit messages](#).

In addition, new audit messages were added to support new 11.7 functionality:

- [BROR: Bucket Read Only Request](#)
- [CGRR: Cross-Grid Replication Request](#)
- [EBDL: Empty Bucket Delete](#)
- [EBKR: Empty Bucket Request](#)
- [S3SL: S3 Select Request](#)

New alerts

The following new alerts were added for StorageGRID 11.7:

- Appliance DAS drive fault detected
- Appliance DAS drive rebuilding
- Appliance fan fault detected
- Appliance NIC fault detected
- Appliance SSD critical warning
- AutoSupport message failed to send
- Cassandra oversize write error
- Cross-grid replication permanent request failure
- Cross-grid replication resources unavailable
- Debug performance impact
- Expiration of grid federation certificate

- FabricPool bucket has unsupported bucket consistency setting
- Firewall configuration failure
- Grid federation connection failure
- Storage appliance fan fault detected
- Storage Node not in desired storage state
- Storage volume needs attention
- Storage volume needs to be restored
- Storage volume offline
- Trace configuration enabled
- Volume Restoration failed to start replicated data repair

Documentation changes

- A new quick reference summarizes how StorageGRID supports Amazon Simple Storage Service (S3) APIs. See [Quick reference: Supported S3 API requests](#).
- The new [StorageGRID quick start](#) lists the high-level steps for configuring and using a StorageGRID system and provides links to the relevant instructions.
- The appliance hardware installation instructions were combined and consolidated for ease of use. A quick start was added as a high-level guide to hardware installation. See [Quick start for hardware installation](#).
- The maintenance instructions common to all appliance models were combined, consolidated, and moved to the maintenance section of the doc site. See [Common node maintenance: Overview](#).
- The maintenance instructions specific to each appliance model were also moved to the maintenance section:

[Maintain SGF6112 hardware](#)

[Maintain SG6000 hardware](#)

[Maintain SG5700 hardware](#)

[Maintain SG100 and SG1000 hardware](#)

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.

Connection Load Balancer (CLB) service removed

The Connection Load Balancer (CLB) service on Gateway Nodes was deprecated in StorageGRID 11.4 and has now been completely removed from the software. To distribute incoming network connections from client applications to Storage Nodes, you can configure load balancer endpoints for the Load Balancer service, which is included on all Admin Nodes and Gateway Nodes, or you can integrate a third-party load balancer. See [Considerations for load balancing](#).

If custom certificates were set up for the S3 or Swift API in the existing StorageGRID version, the CLB ports

8082, 8083, 8084, and 8085 will be automatically converted to load balancer endpoints during the upgrade to StorageGRID 11.7.

SG5600 appliance is End of Support

The SG5600 appliance has reached End Of Support. Contact your NetApp Sales Representative for hardware refresh options.

If you need to perform maintenance procedures on SG5600 hardware, use the [StorageGRID 11.6 instructions](#).

Swift support deprecated

As of the StorageGRID 11.7 release, support for Swift client applications has been deprecated. The user interface and APIs that support Swift client applications will be removed in a future release.

Archive Node support deprecated

Support for Archive Nodes (for both archiving to the cloud using the S3 API and archiving to tape using TSM middleware) is deprecated and will be removed in a future release. Moving objects from an Archive Node to an external archival storage system has been replaced by ILM Cloud Storage Pools, which offer more functionality.

See:

- [Migrate objects to a Cloud Storage Pool](#)
- [Use Cloud Storage Pools](#)

In addition, you should remove Archive Nodes from the active ILM policy in StorageGRID 11.7 or earlier. Removing object data stored on Archive Nodes will simplify future upgrades. See [Working with ILM rules and ILM policies](#).

Audit export through CIFS/Samba removed

Audit export through CIFS/Samba was deprecated in StorageGRID Webscale 11.1 and has now been removed. As required, you can [use an external syslog server](#) or [configure audit client access for NFS](#).

Option to specify a storage pool as a temporary location removed

Previously, when you created an ILM rule with an object placement that includes a single storage pool, you were prompted to specify a second storage pool to use as a temporary location. Starting with StorageGRID 11.7, this option has been removed.

Grid Manager options moved or removed

Several Grid Manager options were moved or removed.

- The [Compress stored objects](#) option was moved to **CONFIGURATION > System > Object compression**.
- The **Network Transfer Encryption** internal connection setting was removed and replaced by the [TLS and SSH policies](#) tab on the new **CONFIGURATION > Security > Security settings** page.



The AES256-SHA option was the default in StorageGRID 11.6 and is the only setting available in StorageGRID 11.7. The AES128-SHA value is ignored in the Grid Management API. During the StorageGRID 11.7 upgrade, the network transfer encryption algorithm is set to AES256-SHA.

- The **Stored object encryption**, **Prevent client modification**, and **Enable HTTP for Storage Node connections** options were moved to the [Network and objects tab](#) on the new **CONFIGURATION > Security > Security settings** page.
- The [Browser inactivity timeout](#) option was moved to the new **CONFIGURATION > Security > Security settings** page.
- The [Link cost](#) option was moved to **SUPPORT > Other > Link cost**.
- The list of NMS entities was moved to **SUPPORT > Other > NMS entities**.
- The **Stored Object Hashing** option was removed. The **SHA-1** and **SHA-256** settings are no longer used for internal background verification because they require additional CPU resources over MD5 and packet CRC32 check.
- The **Preferred sender** option was removed. If your StorageGRID deployment includes multiple Admin Nodes, the primary Admin Node is the preferred sender for alert notifications, AutoSupport messages, SNMP traps and informs, and legacy alarm notifications. If the primary Admin Node becomes unavailable, notifications are temporarily sent by other Admin Nodes. See [What is an Admin Node?](#).
- The [Untrusted Client Network settings](#) were moved to **CONFIGURATION > Firewall control**.

S3 endpoint domain name format restrictions

The use of IP addresses as endpoint domain names is unsupported. Future releases will prevent the configuration. If you need to use IP addresses for endpoint domain names, contact technical support. See [S3 endpoint domain names](#).

User initiated Volume Lost command removed

The `proc/CMSI/Volume_Lost` has been removed. Use the `repair-data start-replicated-volume-repair` command to restore replicated data for a volume.

Changes to the Grid Management API

StorageGRID 11.7 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 11.7; however, support for these versions of the API will be removed in a future release of StorageGRID. After upgrading to StorageGRID 11.7, 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](#).

Display names now included in responses to node-health requests

Related to the new [Rename grid, sites, and nodes procedure](#), after you rename a site or node, the item's name (its system name) and its display name are both returned by the **node-health** API.

Can create bucket and access keys for new S3 tenant

New `s3Bucket` and `s3AccessKey` options were added to the **accounts** API. When you create an S3 tenant account using the Grid Management API, you can optionally create a bucket for that tenant as well as the access key ID and secret key for the tenant's root user.

Can change storage state for Storage Node

You can use the new **node-storage-state** API endpoints to determine and change the state of the storage in a Storage Node (online, offline, read-only).

Changes to the Tenant Management API

StorageGRID 11.7 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 11.7; however, support for these versions of the API will be removed in a future release of StorageGRID. After upgrading to StorageGRID 11.7, the deprecated v1 and v2 APIs can be deactivated using the `PUT /grid/config/management` API.

New endpoints for grid federation

You can use the **grid-federation-connections** API endpoints to list grid federation connections for the current tenant and to clear the last cross-grid replication error for the current tenant and selected grid federation connection.

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

Plan and prepare for upgrade

Estimate the time to complete an upgrade

When planning an upgrade to StorageGRID 11.7, 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 can't 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
Run prechecks and upgrade primary Admin Node	The upgrade prechecks are run, and the primary Admin Node is stopped, upgraded, and restarted.	30 minutes to 1 hour, with SG100 and SG1000 appliance nodes requiring the most time. Unresolved precheck errors will increase this time.	You can't access the primary Admin Node. Connection errors might be reported, which you can ignore. Running the upgrade prechecks before starting the upgrade lets you resolve any errors before the scheduled upgrade maintenance window.
Start upgrade service	The software file is distributed, and the upgrade service is started.	3 minutes per grid node	
Upgrade other grid 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.	15 minutes to 1 hour per node, with appliance nodes requiring the most time Note: For appliance nodes, the StorageGRID Appliance Installer is automatically updated to the latest release.	<ul style="list-style-type: none"> • Don't change the grid configuration. • Don't change the audit level configuration. • Don't update the ILM configuration. • You are prevented from performing other maintenance procedures, such as hotfix, decommission, or expansion. Note: If you need to perform a recovery, contact technical support.
Enable features	The new features for the new version are enabled.	Less than 5 minutes	<ul style="list-style-type: none"> • Don't change the grid configuration. • Don't change the audit level configuration. • Don't update the ILM configuration. • You can't perform another maintenance procedure.
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 11.6 to 11.7 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 11.7

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 11.7.0.



As required, you can split the maintenance window into smaller windows by approving subsets of grid nodes to upgrade in multiple sessions. For example, you might prefer to upgrade the nodes at site A in one session and then upgrade the nodes at site B in a later session. If you choose to perform the upgrade in more than one session, be aware that you can't start using the new features until all nodes have been upgraded.

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. If you approve all nodes of the same type to upgrade (for example, Storage Nodes), the nodes are brought down one at a time, so there is no time when all grid nodes or all grid nodes of a certain type are unavailable.

To allow for continued availability, ensure that your ILM policy contains rules that specify storing multiple copies of each object. You must also ensure that all external S3 or Swift clients are configured to send requests to one of the following:

- A high availability (HA) group virtual IP address
- 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 11.7 upgrade:

- All StorageGRID appliance nodes are automatically upgraded to StorageGRID Appliance Installer firmware version 3.7.
- SG6060 and SGF6024 appliances are automatically upgraded to BIOS firmware version 3B07.EX and BMC firmware version 3.97.07.
- SG100 and SG1000 appliances are automatically upgraded to BIOS firmware version 3B12.EC and BMC firmware version 4.71.07.

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 11.7 but other nodes are still running StorageGRID 11.6. 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 11.7. This alert might persist for 1 day after the upgrade completes.

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, clear the **Enable SNMP Agent Notifications** checkbox (**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



This list applies specifically to upgrades from StorageGRID 11.6 to StorageGRID 11.7. If you're upgrading to another StorageGRID release, refer to the list of restricted changes in the upgrade instructions for that release.

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

- Don't make any grid configuration changes.
- Don't enable or disable any new features. For example, don't attempt to create a grid federation connection until both StorageGRID systems have been updated to StorageGRID 11.7.

- Don't update the ILM configuration. Otherwise, you might experience inconsistent and unexpected ILM behavior.
- Don't apply a hotfix or recover a grid node.



Contact technical support if you need to recover a node during upgrade.

- You should not manage HA groups, VLAN interfaces, or load balancer endpoints while you're upgrading to StorageGRID 11.7.
- Don't delete any HA groups until the upgrade to StorageGRID 11.7 is complete. Virtual IP addresses in other HA groups might become inaccessible.

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

- Don't perform an expansion procedure.
- Don't perform a decommission procedure.

You can't view bucket details or manage buckets from the Tenant Manager

During the upgrade to StorageGRID 11.7 (that is, while the system is operating as a mixed-version environment), you can't view bucket details or manage buckets using the Tenant Manager. One of the following errors appears on the Buckets page in Tenant Manager:

- You can't use this API while you're upgrading to 11.7.
- You can't view bucket versioning details in the Tenant Manager while you're upgrading to 11.7.

This error will resolve after the upgrade to 11.7 is complete.

Workaround

While the 11.7 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.

TLS ciphers or SSH configurations might change

If TLS ciphers or SSH configurations have been manually changed or are inconsistent across nodes, all nodes will be overwritten to be either Legacy Compatibility or Modern Compatibility after upgrade. If you used `fips-ciphers.sh` in StorageGRID 11.6, the Common Criteria policy is applied to all nodes. Otherwise, the Legacy Compatibility policy is applied. If you require Common Criteria validated configurations, you must use the Common Criteria policy or the FIPS strict policy. If you did not use `fips-ciphers.sh`, you should use the new Modern Compatibility setting after upgrade completes. To configure ciphers, go to **CONFIGURATION > Security > Security settings** and select **TLS and SSH policies**.

CLB ports might be converted to load balancer endpoints

The legacy Connection Load Balancer (CLB) service has been removed in StorageGRID 11.7. If CLB configuration is detected during upgrade prechecks, the **Legacy CLB load balancer activity detected** alert will be triggered. If custom certificates were set up for the S3 or Swift API in the existing StorageGRID version,

the CLB ports 8082, 8083, 8084, and 8085 will be converted to load balancer endpoints during upgrade to StorageGRID 11.7.

See also [Considerations for load balancing](#).

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 11.7, optionally assign the following new permission to tenant user groups.

Permission	Description
Tenant Manager > Manage objects with S3 Console	<p>When combined with the Manage all buckets permission, this permission allows users to access the Experimental S3 Console from the Buckets page.</p> <p>Users who have this permission but who don't have the Manage all buckets permission can still navigate directly to the Experimental S3 Console.</p>

See [Tenant management permissions](#).

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 11.7, your grid must have StorageGRID 11.6 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 11.6.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 11.6.x.y.

In the StorageGRID 11.6.x.y version number:

- The **major release** has an x value of 0 (11.6.0).
 - A **hotfix**, if one has been applied, has a y value (for example, 11.6.0.1).
4. If **Version** is not 11.6.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 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 11.6 from version 11.3.0.8

The following example shows the steps to upgrade from StorageGRID version 11.3.0.8 to version 11.6 in preparation for a StorageGRID 11.7 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.
6. Upgrade to the StorageGRID 11.6.0 major release.
7. Apply the latest StorageGRID 11.6.0.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
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.

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.
Linux RPM or DEB archive	<p>If any nodes are deployed on Linux hosts, you must download and install the RPM or DEB package on all hosts before you start the upgrade.</p> <p>Important: Ensure that your operating system is upgraded to Linux kernel 4.15 or higher.</p>
StorageGRID documentation	<ul style="list-style-type: none"> • Release notes for StorageGRID 11.7 (sign in required). Be sure to read these carefully before starting the upgrade. • StorageGRID software upgrade resolution guide for the major version you are upgrading to (sign in required) • Other StorageGRID 11.7 documentation, as required.

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.
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 11.7 are opened before you upgrade.

The following [internal ports](#) must be open before you upgrade to StorageGRID 11.7:

Port	Description
1055	Used for firewall controls knocking protocol.
8011 10342	Before upgrading, confirm this port is open between all grid nodes on the Grid Network.
18086	TCP port used for new LDR service. Before upgrading, confirm this port is open from all grid nodes to all Storage Nodes.




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 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 and the hotfix file, if required.
3. Download the current Recovery Package.
4. Run upgrade prechecks to detect and resolve any issues before you start the actual upgrade.
5. Start the upgrade, which runs prechecks and upgrades the primary Admin Node automatically. You can't 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.
6. After the primary Admin Node has been upgraded, download a new Recovery Package.

7. Approve the grid nodes. You can approve individual grid nodes, groups of grid nodes, or all grid nodes.



Don't 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.

Related information

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

Linux: Download and install the RPM or DEB package on all hosts

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

Download upgrade, Linux, and hotfix files

When you perform a StorageGRID upgrade from the Grid Manager, you are prompted to download the upgrade archive and any required hotfix as the first step. However, if you need to download files to upgrade Linux hosts, you can save time by downloading all required files in advance.

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 notice appears, make note of the hotfix number, and select the checkbox.
5. Read the End User License Agreement, select the checkbox, and then select **Accept & Continue**.

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

6. From the second column (**Upgrade StorageGRID**), download two files:
 - The upgrade archive for the latest release (this is the file in the section labeled **VMware, SG1000, or SG100 Primary Admin Node**). While this file is not needed until you perform the upgrade, downloading it now will save time.
 - An 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`

7. If you needed to agree to a Caution/MustRead notice because of a required hotfix, download the hotfix:
 - a. Go back to [NetApp Downloads: StorageGRID](#).
 - b. Select the hotfix number from the drop-down.
 - c. Agree to the Caution notice and EULA again.
 - d. Download and save the hotfix and its README.

You will be prompted to upload the hotfix file on the StorageGRID Upgrade page when you start the upgrade.

Install archive on all Linux hosts

Perform these steps before upgrading StorageGRID software.

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:

- [Red Hat Enterprise Linux or CentOS: Install StorageGRID host services](#)
- [Ubuntu or Debian: Install StorageGRID host services](#)

The new packages are installed as additional packages. Don't remove the existing packages.

Perform the upgrade

You can upgrade to StorageGRID 11.7 and apply the latest hotfix for that release at the same time. The StorageGRID upgrade page provides the recommended upgrade path and links directly to the correct download pages.

Before you begin

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

Access StorageGRID Upgrade page

As a first step, access the StorageGRID Upgrade page in the Grid Manager.

Steps

1. Sign in to the Grid Manager using a [supported web browser](#).
2. Select **MAINTENANCE** > **System** > **Software update**.
3. From the StorageGRID upgrade tile, select **Upgrade**.

Select files

The update path on the StorageGRID Upgrade page indicates which major versions (for example, 11.7.0) and hotfixes (for example, 11.7.0.1) you must install to get to the latest StorageGRID release. You should install the recommended versions and hotfixes in the order shown.



If no update path is shown, your browser might not be able to access the NetApp Support Site, or the **Check for software updates** checkbox on the AutoSupport page (**SUPPORT > Tools > AutoSupport**) might be disabled.

Steps

1. For the **Select files** step, review the update path.
2. From the Download files section, select each **Download** link to download the required files from the NetApp Support Site.

If no update path is shown, go to the [NetApp Downloads: StorageGRID](#) to determine if a new version or hotfix is available and to download the files you need.



If you needed to download and install an RPM or DEB package on all Linux hosts, you might already have the StorageGRID upgrade and hotfix files listed in the update path.

3. Select **Browse** to upload the version upgrade file to StorageGRID:
`NetApp_StorageGRID_11.7.0_Software_uniqueID.upgrade`

When the upload and validation process is done, a green check mark appears next to the file name.

4. If you downloaded a hotfix file, select **Browse** to upload that file. The hotfix will be automatically applied as part of the version upgrade.
5. Select **Continue**.

Run prechecks

Running prechecks allows you to detect and resolve any upgrade issues before you start upgrading your grid.

Steps

1. For the **Run prechecks** step, start by entering the provisioning passphrase for your grid.
2. Select **Download recovery package**.

You should download the current copy of the Recovery Package file before you upgrade the primary Admin Node. The Recovery Package file allows you to restore the system if a failure occurs.

3. When the file is downloaded, confirm that you can access the contents, including the `Passwords.txt` file.
4. Copy the downloaded 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.

5. Select **Run prechecks**, and wait for the prechecks to complete.
6. Review the details for each reported precheck and resolve any reported errors. See the [StorageGRID software upgrade resolution guide](#) for the StorageGRID 11.7 release.

You must resolve all precheck *errors* before you can upgrade your system. However, you don't need to address precheck *warnings* before upgrading.



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.

7. If you made any configuration changes to resolve the reported issues, select **Run prechecks** again to get updated results.

If all errors have been resolved, you are prompted to start the upgrade.

Start upgrade and upgrade primary Admin Node

When you start the upgrade, the upgrade prechecks are run again, and the primary Admin Node is automatically upgraded. This part of the upgrade can take up to 30 minutes.



You will not be able to access any other Grid Manager pages while the primary Admin Node is being upgraded. Audit logs will also be unavailable.

Steps

1. Select **Start upgrade**.

A warning appears to remind you will temporarily lose access to the Grid Manager.

2. Select **OK** to acknowledge the warning and start the upgrade.
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.

If the grid has another Admin Node that is online and ready, you can use it to monitor the status of the primary Admin Node. As soon as the primary Admin Node is upgraded, you can approve the other grid nodes.

4. As required, select **Continue** to access the **Upgrade other nodes** step.

Upgrade other nodes

You must upgrade all grid nodes, but you can perform multiple upgrade sessions and customize the upgrade sequence. For example, you might prefer to upgrade the nodes at site A in one session and then upgrade the nodes at site B in a later session. If you choose to perform the upgrade in more than one session, be aware that you can't start using the new features until all nodes have been upgraded.

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.



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, don't 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.

Steps

1. For the **Upgrade other nodes** step, review the Summary, which provides the start time for the upgrade as a whole and the status for each major upgrade task.

- **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 on each node.
- When the **Start upgrade service** task is complete, the **Upgrade other grid nodes** task starts, and you are prompted to download a new copy of the Recovery Package.

2. When prompted, enter your provisioning passphrase and download a new copy of the Recovery Package.



You should download a new copy of the Recovery Package file after the primary Admin Node is upgraded. The Recovery Package file allows you to restore the system if a failure occurs.

3. Review the status tables for each type of node. There are tables for non-primary Admin Nodes, Gateway Nodes, Storage Nodes, and Archive Nodes.

A grid node can be in one of these stages when the tables first appear:

- Unpacking the upgrade
- Downloading
- Waiting to be approved

4. When you are ready to select grid nodes for upgrade (or if you need to unapprove selected nodes), use these instructions:

Task	Instruction
Search for specific nodes to approve, such as all nodes at a particular site	Enter the search string in the Search field
Select all nodes for upgrade	Select Approve all nodes
Select all nodes of the same type for upgrade (for example, all Storage Nodes)	Select the Approve all button for the node type If you approve more than one node of the same type, the nodes will be upgraded one at a time.
Select an individual node for upgrade	Select the Approve button for the node
Postpone the upgrade on all selected nodes	Select Unapprove all nodes
Postpone the upgrade on all selected nodes of the same type	Select the Unapprove all button for the node type
Postpone the upgrade on an individual node	Select the Unapprove button for the node

5. Wait for the approved nodes to proceed through these upgrade stages:

- Approved and waiting to be upgraded
- Stopping services



You can't remove a node when its Stage reaches **Stopping services**. The **Unapprove** button is disabled.

- Stopping container
- Cleaning up Docker images
- Upgrading base OS packages



When an appliance node reaches this 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.

- Rebooting



Some appliance models might reboot multiple times to upgrade the firmware and BIOS.

- Performing steps after reboot
- Starting services
- Done

6. Repeat the [approval step](#) as many times as needed until all grid nodes have been upgraded.

Complete upgrade

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

Steps

1. As soon as the **Enable features** task is complete (which occurs quickly), you can 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 11.6 to 11.7 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** to complete.
4. When the **Final upgrade steps** have completed, the upgrade is done. The first step, **Select files**, is redisplayed with a green success banner.
5. Verify that grid operations have returned to normal:
 - a. Check that the services are operating normally and that there are no unexpected alerts.
 - b. Confirm that client connections to the StorageGRID system are operating as expected.

Troubleshoot upgrade issues

If something goes wrong when you perform an upgrade, you might be able to resolve the issue yourself. If you can't resolve an issue, gather as much information as you can and then contact technical support.

Upgrade does not complete

The following sections describe how to recover from situations where the upgrade has partially failed.

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.

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 aren't 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](#)

User interface issues

You might experience issues with the Grid Manager or the Tenant Manager during or after the upgrade.

Grid Manager displays multiple error messages during upgrade

If you refresh your browser or navigate to another Grid Manager page while the primary Admin Node is being upgraded, you might see multiple "503: Service Unavailable" and "Problem connecting to the server" messages. You can safely ignore these messages—they will stop appearing soon as the node is upgraded.

If these messages appear for more than an hour after you started the upgrade, something might have occurred that prevented the primary Admin Node from being upgraded. If you are unable to resolve the issue on your own, contact technical support.

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.

Message	Cause	Solution
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.

Copyright information

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

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

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

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

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

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

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

Trademark information

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