



**StorageGRID® 11.3**

# Upgrade Guide

November 2019 | 215-14194\_2019-11\_en-us  
[doccomments@netapp.com](mailto:doccomments@netapp.com)

 **NetApp®**



# Contents

<b>About StorageGRID 11.3 .....</b>	<b>4</b>
What's new in StorageGRID 11.3 .....	4
Removed or deprecated features .....	10
Changes to the Grid Management API .....	11
Changes to the Tenant Management API .....	11
<b>Upgrade planning and preparation .....</b>	<b>13</b>
Estimating the time to complete an upgrade .....	13
How your system is affected during the upgrade .....	15
Impact of an upgrade on groups and user accounts .....	16
Verifying the installed version of StorageGRID .....	17
Obtaining the required materials for a software upgrade .....	17
Web browser requirements .....	18
Downloading the StorageGRID upgrade file .....	19
Downloading the Recovery Package .....	20
Checking the system's condition before upgrading software .....	20
<b>Performing the upgrade .....</b>	<b>22</b>
Starting the upgrade .....	23
Upgrading grid nodes and completing the upgrade .....	27
Verifying the completion of your upgrade .....	30
<b>Troubleshooting upgrade issues .....</b>	<b>32</b>
User interface issues .....	32
“Docker image availability check” error messages .....	33
<b>Copyright .....</b>	<b>34</b>
<b>Trademark information .....</b>	<b>35</b>
<b>How to send comments about documentation and receive update notifications .....</b>	<b>36</b>

## About StorageGRID 11.3

---

Before starting an upgrade, review this section to learn about the new features and enhancements in StorageGRID 11.3, determine whether any features have been deprecated or removed, and find out about changes to StorageGRID APIs.

### What's new in StorageGRID 11.3

StorageGRID 11.3 introduces a new Load Balancer service, support for high availability node groups, new alerts functionality, improvements to Cloud Storage Pools, enhancements to object ingest and delete processing, new StorageGRID appliances, and more.

#### New Load Balancer service

A new Load Balancer service is included on Gateway Nodes and on all Admin Nodes. This service provides Layer 7 load balancing of S3 and Swift traffic from clients to Storage Nodes. The legacy Connection Load Balancer (CLB) service on Gateway Nodes is still supported; however, configuring endpoints for the new Load Balancer service is recommended (**Configuration > Load Balancer Endpoints**).

[Administering StorageGRID](#)

#### High availability groups

You can now create high availability (HA) groups of Admin Nodes and Gateway Nodes (**Configuration > High Availability Groups**). HA groups use virtual IP addresses to provide active-backup access to Gateway Node or Admin Node services. For example, you can create an HA group of Gateway Nodes and Admin Nodes to provide highly available data connections for S3 and Swift clients. Or, you can create an HA group of Admin Nodes to provide highly available connections to the Grid Manager and the Tenant Manager.

If required, you can achieve an active-active configuration by using round-robin DNS or a third-party load balancer and multiple HA groups.

[Administering StorageGRID](#)

#### Untrusted Client Network feature

You can use the Untrusted Client Network feature to secure the StorageGRID nodes on the Client Network from hostile attacks. The new feature allows you to specify that a given node only accept inbound connections on ports explicitly configured as load balancer endpoints (**Configuration > Untrusted Client Network**).

For example, you might want a Gateway Node to refuse all inbound traffic on the Client Network except for HTTPS S3 requests. Or, you might want to enable outbound S3 platform service traffic from a Storage Node, while preventing any inbound connections to that Storage Node on the Client Network.

[Administering StorageGRID](#)

#### New alerts functionality

A new alerts system is available to preview in StorageGRID 11.3. The alerts system is designed to be easier to use and more powerful than the legacy alarms system.

**Attention:** For StorageGRID 11.3, consider the alerts system to be a supplement to the alarms system, not a replacement for it. You must continue to use the alarms system as your primary tool for detecting and resolving any issues with your system.

Some of the benefits of the new alerts system include the following:

- Multiple alerts of the same type are reported in one email notification to reduce the number of emails received.
- The Alerts page provides a user friendly interface for viewing current problems across your StorageGRID system. You can expand and collapse groups of alerts and sort the listing by severity, location, or time triggered.
- Alerts use intuitive names and descriptions to help you understand quickly what the problem is, and they provide the recommended actions for resolving the alert.
- If you need to temporarily suppress the notifications at one or more severity levels, you can easily silence a specific alert rule for the entire grid, a single site, or a single node.
- You can create custom alert rules to target the specific conditions that are relevant to your situation and provide your own recommended actions. To define the conditions for different alert severities, you create expressions using the Prometheus metrics listed in the Metrics section of the Grid Management API.

**Note:** As part of this enhancement, the existing alarms and monitoring information was moved from the instructions for administering StorageGRID to the new instructions for monitoring and troubleshooting StorageGRID.

### *Monitoring and troubleshooting StorageGRID*

#### **Enhancements to Cloud Storage Pools**

In addition to using a Cloud Storage Pool to tier object data from StorageGRID to an external location, you can now use Cloud Storage Pools for backup. You can also configure more than one Cloud Storage Pool endpoint. Specifically:

- You can now back up an object to a Cloud Storage Pool while one or more copies of that object also exist in StorageGRID.
- You can now create Cloud Storage Pools that connect to Microsoft Azure Blob storage endpoints. You can also create Cloud Storage Pools within the AWS Secret Region.
- You can now create up to 10 Cloud Storage Pools, each with a unique cloud endpoint. However, you cannot store the same object in more than one Cloud Storage Pool at a time.
- You can now use ILM rules to restore objects from a Cloud Storage Pool back to StorageGRID. The ILM process automatically issues the required transition requests for objects that are in a non-retrievable state.
- The lifecycle policy on the external S3 bucket used for a Cloud Storage Pool can now include transitions to S3 Glacier Deep Archive.

### *Administering StorageGRID*

#### **Enhancements to object ingest**

When creating an ILM rule, you can now indicate whether you want the rule's placement instructions to be satisfied when the objects are ingested. Previously, StorageGRID used dual commit—it made two interim copies during ingest and evaluated ILM later.

On upgrade, existing ILM rules continue to use dual commit. After upgrade is complete, you can configure ILM rules to use the new ingest behavior by selecting one of these options: Balanced

(which attempts to make all required copies during ingest and performs dual commit if that is not possible) or Strict (which fails ingest if StorageGRID cannot immediately make all required copies).

The Balanced and Strict options cannot be used for some types of object placements. In addition, these options are not recommended for use with erasure-coded objects when objects are larger than 4 MB or if the erasure-coding scheme creates more than seven fragments. (That is, only the 2+1, 4+1, 4+2, and 6+1 erasure coding schemes are recommended.)

### *Administering StorageGRID*

#### **Enhancements to object deletion**

StorageGRID 11.3 improves delete performance and introduces synchronous deletion, which enables content to be removed from the grid more quickly in response to client requests.

In previous releases, StorageGRID always provided an immediate response to client delete requests and queued object copies for deletion later. With synchronous deletion, StorageGRID attempts to remove all object copies before providing a client response. This change means that clients might sometimes receive a slower response, even though objects are generally being removed more quickly than they were in the past.

In addition, when an S3 versioned object is deleted, StorageGRID now creates a delete marker as the current version of the object. This behavior matches AWS S3 behavior.

### *Administering StorageGRID*

#### **Enhancements to object capacity**

StorageGRID 11.3 optimizes database operations and metadata space allocations to increase the grid's object capacity. These changes significantly increase the number of objects per node that a StorageGRID deployment can support in many circumstances. The exact number depends on factors such as how many times ILM rules change object placements and how much user metadata and tags are stored per object.

As part of these changes, more space is now reserved for metadata on volume 0 of Storage Nodes that have 128 GB or more of RAM. When you upgrade, the size of the metadata reservation is automatically increased to 4 TB for these larger Storage Nodes, unless the Metadata Reserved Space (CAWM) setting has been changed from its default value of 3 TB (**Configuration > Storage Options > Overview**).

### *Administering StorageGRID*

#### **Changes to metadata usage reporting**

StorageGRID 11.2 and earlier under-reported the amount of metadata used by approximately 10%. After upgrade to StorageGRID 11.3, the reported metadata usage will increase and reflect the actual value. To see the value for used metadata, select **Nodes > Storage Node > Storage**, and hover over the Storage Used – Object Metadata graph. A pop-up displays Used (%), Used, and Total (allowed) values.

### *Monitoring and troubleshooting StorageGRID*

#### **Changes to ILM processing for Last Access Time**

Changing the Last Access Time for an object no longer adds the object to an ILM queue for immediate processing. Instead, the object's placements are re-evaluated during background ILM processing. If you use Last Access Time as a reference time for an ILM rule, you should check and update the time periods you have specified for object placements. Placements should typically last for more than one month.

### *Administering StorageGRID*

## Enhancements to the Grid Manager

- If you send AutoSupport messages using HTTP or HTTPS, you can now configure a non-transparent proxy server between Admin Nodes and technical support (**Configuration > Proxy Settings > Admin**).
- When configuring identity federation for the Grid Manager, you can now use LDAP over SSL (LDAPS) to secure communications between StorageGRID and the LDAP server. STARTTLS is still the recommended method for securing identity federation.
- You can now prevent grid administrators from being able to access the Tenant Manager by opening port 8443 on the external firewall and closing port 443. All Tenant Manager and internal traffic is rejected on port 8443.
- A new **Usage** button on the Tenants page allows you to monitor the storage usage for each tenant account, including which S3 buckets or Swift containers are consuming the most storage. If a quota was set for the tenant, you can see how much of that quota has been used.
- To ensure that operations are not disrupted by a failed server certificate, new alarms (and corresponding new alerts) are triggered if the Management Interface Server Certificate or the Object Storage API Service Endpoints Server Certificate is about to expire. To view the number of days until a server certificate expires, go to **Support > Grid Topology > primary Admin Node > CMN > Resources**.
- When adding Storage Nodes in an expansion, you can use the percentage complete estimate for the “Starting Cassandra and streaming data” stage to better estimate how much time this operation might take.
- You can now use the Grid Manager to delete quarantined objects and reset the count of quarantined objects to zero (**Support > Grid Topology > site > Storage Node > LDR > Verification > Configuration > Main**). Previously, quarantined objects could only be deleted by technical support.
- If you have internet access, you can now access the StorageGRID 11.3 Documentation Center directly from the Grid Manager (**Help > Documentation Center**).
- When working with technical support to troubleshoot an issue, you can use the new Metrics page (**Support > Metrics**) to review detailed metrics and charts for your StorageGRID system.
 

**Attention:** The tools available on the Metrics page are intended for use by technical support. Some features and menu items within these tools are intentionally non-functional.

### *Administering StorageGRID*

#### *Expanding a StorageGRID system*

#### *Monitoring and troubleshooting StorageGRID*

## Enhancements to the Tenant Manager

- When configuring identity federation for the Tenant Manager, you can now use LDAP over SSL (LDAPS) to secure communications between StorageGRID and the LDAP server. STARTTLS is still the recommended method for securing identity federation.
- You can now prevent tenants from being able to access the Grid Manager by opening port 9443 on the external firewall and closing port 443. All Grid Manager and internal traffic is rejected on port 9443.
- Several enhancements were made to S3 platform services, including the following:

- Basic HTTP authentication is now supported for connections to Elasticsearch clusters. You specify authentication credentials when creating an endpoint for the search integration platform service.
- The search integration service now includes the bucket's region in the metadata notifications that it sends to an endpoint.
- If you have internet access, you can now access the StorageGRID 11.3 Documentation Center directly from the Tenant Manager (**Help > Documentation Center**).

### *Using tenant accounts*

#### **Enhancements to S3 REST API support**

- Support for server-side encryption with customer-provided keys (SSE-C) has been added to object operations in the StorageGRID S3 REST API.
- S3 clients can now configure a bucket lifecycle to control how long their objects are retained. If an S3 bucket lifecycle exists, the Expiration action in the lifecycle will always override ILM rule settings. As part of this change, support has been added for the S3 DELETE, GET, and PUT Bucket lifecycle operations (Expiration and NoncurrentVersionExpiration actions only) and for the x-amz-expiration response header, which returns expiry-date and rule-id.
- Improvements were made to the validation and processing of the Content-Encoding field.
- Clients can now use the TLS\_ECDHE\_RSA\_WITH\_CHACHA20\_POLY1305\_SHA256 cipher when establishing a Transport Layer Security (TLS) session with StorageGRID.

### *Implementing S3 client applications*

#### *Administering StorageGRID*

#### **Enhancements to Swift REST API support**

- Clients can now use the TLS\_ECDHE\_RSA\_WITH\_CHACHA20\_POLY1305\_SHA256 cipher when establishing a Transport Layer Security (TLS) session with StorageGRID.

### *Implementing Swift client applications*

#### **Audit message changes**

- New fields were added to the SDEL, SGET, SPOS, and SPUT audit messages to provide additional auditing information for operations on S3 object and bucket subresources:
  - S3SR: Identifies the S3 subresource being operated on.
  - SRCF: Shows the new configuration. Included only in requests that set a non-sensitive subresource configuration:
    - PUT Bucket lifecycle
    - PUT Bucket policy
    - PUT Bucket compliance
    - PUT Bucket consistency
    - PUT Bucket last access time
    - PUT Bucket versioning
    - POST Object restore
  - CNCH: For requests with a Consistency-Control header, shows the value of the Consistency-Control header.
- Several fields were changed in the following audit messages:



- PUT Bucket compliance: ATYPE was changed from SUPD to SPUT. CMPS was replaced by SRCF, which is the whole request XML body.
- PUT Bucket versioning: VSST was replaced by SRCF. (VSST is still used in audit message where the bucket is versioned.)
- POST Object restore: RRDA (Restore Days) and RRTI (Restore Tier) were replaced by SRCF.
- PUT Object tagging: ATYP was changed from SUPD to SPUT.

### *Understanding audit messages*

## **Changes to the internal firewall**

The firewall service inside of StorageGRID has changed from UFW to nftables, and it has moved to inside the Docker container. This change allows for some firewall ports to be opened only when configured, such as the ports used by the new Load Balancer service. During the upgrade to StorageGRID 11.3, the open ports are reset to the default set.

**Note:** During the upgrade precheck process, any custom firewall ports that you might have opened are flagged. You must contact technical support before proceeding with the upgrade.

## **New StorageGRID appliances**

- The SG1000 services appliance can operate as a Gateway Node or an Admin Node to provide high availability grid administration and load balancing services.
- The all-flash SGF6024 storage appliance includes 24 flash drives, two EF570 storage controllers, and the same compute controller used for the SG6060 appliance.
- The SG6060 storage appliance can now optionally support one or two 60-drive expansion shelves, which must be installed during initial installation.

### *SG1000 appliance installation and maintenance*

### *SG6000 appliance installation and maintenance*

## **NAS Bridge enhancements**

- You can now define Secondary file systems. If you lose access to the data on the Primary file system, having a Secondary file system on standby allows you to recover the data by performing an “emergency takeover.”
- The NAS Bridge API includes new attributes for SMB file systems:
  - Mandatory attribute: `ntfs_file_attrs`. Set this attribute to `true` for access-control list (ACL) support and to track Windows NT file system (NTFS) style file attributes. Note that these tracking operations decrease performance.
  - Optional attributes: `allow_write` and `allow_read`. Use these attributes to specify user names or group names to restrict read or write access to the share.

### *Administering NAS Bridge*

### *Using the NAS Bridge Management API*

## Removed or deprecated features

Some features have been removed or deprecated in StorageGRID 11.3. You must review these items to understand whether you need to update client applications or modify your configuration before you upgrade.

### Temporary storage pools are deprecated

Designating a storage pool to use as a temporary location in an ILM rule has been deprecated in StorageGRID 11.3, due to changes in how objects are ingested into the grid. Any existing ILM rules that use temporary pools will continue to operate as they have in the past. However, new rules should not use temporary pools.

[Administering StorageGRID](#)

### CMS service removed

The CMS service (Content Management System), which was deprecated in StorageGRID 11.2, has now been removed. This service formerly managed object data that was queued by the legacy ILM system. Related to the removal of the CMS, port 1503 is no longer required for internal grid node communications.

### Removed audit message element

The following audit message element is obsolete and no longer appears in audit messages:

- SPAR (Security Partition ID) element of the ORLM (Object Rules Met) message

### Alarms removed

The following alarms have been removed:

- MMQS (Peak Message Queue Size): Both the attribute and the alarm have been removed.
- NSTA (NTP Sources Available): The NSTA attribute is still available, but the alarm on the attribute has been removed.

### Removed and deprecated cipher suites

TLS 1.1 ciphers are no longer supported.

All CBC and SHA1 ciphers are deprecated. Support for these ciphers has been removed from the S3 and Swift services in 11.3 and will be fully removed in a future release.

TLS\_RSA\_\* ciphers are also deprecated. Support for the following ciphers will be removed in a future release:

- TLS\_RSA\_WITH\_AES\_128\_GCM\_SHA256
- TLS\_RSA\_WITH\_AES\_256\_GCM\_SHA384

[Implementing S3 client applications](#)

[Implementing Swift client applications](#)

### SNMP configuration restricted

For information about using SNMP with StorageGRID 11.3, contact your NetApp account representative

## Changes to the Grid Management API

StorageGRID 11.3 uses version 3 of the Grid Management API. Version 3 deprecates version 2; however, version 1 and version 2 are still supported.

**Attention:** You can continue to use version 1 and version 2 of the management API with StorageGRID 11.3; however, support for these versions of the API will be removed in a future release of StorageGRID. After upgrading to StorageGRID 11.3, the deprecated v1 and v2 APIs can be deactivated using the PUT `/grid/config/management` API.

### New attributes in expansion-nodes section

The response body for the GET `/grid/expansion/nodes` and GET `/grid/expansion/nodes/{id}` endpoints has been updated to include three new attributes: `currentCassandraData`, `estimatedCassandraData`, and `stageProgress` (under the `current progress` attribute).

You can use `stageProgress` (calculated from `currentCassandraData` and `estimatedCassandraData`) to estimate how long the “Streaming Cassandra data” stage will take to complete for a new Storage Node.

### New Admin Node operations in Proxy section

The Proxy section now includes operations for configuring an Admin Node proxy for AutoSupport.

### Enhancements to SNMP agent address configuration

The GET `/grid/snmp` and PUT `grid/snmp` endpoints have been updated to include three new parameters: `protocol`, `network`, and `port`. You can use these parameters to configure the SNMP agent. Previously, these values defaulted to all interfaces and to the default SNMP UDP port.

For information about using SNMP with StorageGRID 11.3, contact your NetApp account representative.

### Character limit corrected for EC profile name in Swagger model

Previously, the `maxLength` for the EC profile name in the Swagger model for `/grid/ec-profiles` was incorrectly listed as 32 characters. The corrected limit is 64 characters.

### New attribute in the ILM section

The data submitted or returned for an ILM rule using PUT `/grid/ilm-rules` or GET `/grid/ilm-rules` now includes the `ingestBehavior` attribute. This attribute can take any of these values: `dual-commit`, `balanced`, or `strict`.

### Related information

[Administering StorageGRID](#)

## Changes to the Tenant Management API

StorageGRID 11.3 uses version 3 of the Tenant Management API. Version 3 deprecates version 2; however, version 1 and version 2 are still supported.

**Attention:** You can continue to use version 1 and version 2 of the management API with StorageGRID 11.3; however, support for these versions of the API will be removed in a future

release of StorageGRID. After upgrading to StorageGRID 11.3, the deprecated v1 and v2 APIs can be deactivated using the `PUT /grid/config/management` API.

### **Changes to endpoints for basic HTTP authentication**

To support the use of basic HTTP authentication for connections to Elasticsearch clusters, new `authType` and `basicHttpCredentials` properties have been added to endpoint operations. It is recommended that you submit a value for `authType` when creating an endpoint, although for backwards compatibility, StorageGRID can infer the correct value from the credentials supplied with the POST request. When updating an existing endpoint using PUT, you must supply the correct value for `authType` (`anonymous`, `accesskey`, or `basicHttp`).

### **Related information**

[\*Using tenant accounts\*](#)

## Upgrade planning and preparation

---

You must plan the upgrade of your StorageGRID system to ensure that the system is ready for the upgrade, and that the upgrade can be completed with minimal disruption.

### Steps

1. [Estimating the time to complete an upgrade](#) on page 13
2. [How your system is affected during the upgrade](#) on page 15
3. [Impact of an upgrade on groups and user accounts](#) on page 16
4. [Verifying the installed version of StorageGRID](#) on page 17
5. [Obtaining the required materials for a software upgrade](#) on page 17
6. [Downloading the StorageGRID upgrade file](#) on page 19
7. [Downloading the Recovery Package](#) on page 20
8. [Checking the system's condition before upgrading software](#) on page 20

## Estimating the time to complete an upgrade

When planning an upgrade to StorageGRID 11.3, 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 stages and lists the approximate time required for each stage. The steps after the table provide instructions you can use to estimate the upgrade time for your system.

**Note:** The upgrade from StorageGRID 11.2 to 11.3 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 or weeks, based on the amount of metadata in your system.

Upgrade stage	Description	Approximate time required	During this stage
Pre-upgrade validation	The grid's condition is validated.	3 minutes per grid node, unless validation errors are reported	The upgrade is not yet running. As required, you can perform this step before the scheduled upgrade maintenance window.
Primary Admin Node upgrade	The primary Admin Node is stopped, upgraded, and restarted.	30 minutes	You cannot access the primary Admin Node.

Upgrade stage	Description	Approximate time required	During this stage
Upgrade of all 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 for several minutes each.	15 to 45 minutes per node, with appliance Storage Nodes requiring the most time	<ul style="list-style-type: none"> <li>• Do not change the grid configuration.</li> <li>• Do not change the audit level configuration.</li> <li>• Do not update the ILM configuration.</li> <li>• Performing other maintenance procedures, including decommissioning and expansion, is not supported.</li> </ul>
Cassandra database update	Storage Nodes only. The upgrade process stops the Cassandra service, checks each node to verify that the Cassandra database does not need to be updated, and restarts the service.	3 minutes for each Storage Node	<ul style="list-style-type: none"> <li>• If you need to perform a recovery procedure, contact technical support.</li> </ul>
Restart services	Some grid node services are restarted. Affected grid nodes might be shown as Administratively Down.	15 minutes per node	
Complete final steps	The upgrade to the new release completes.	5 minutes	<p>When the final upgrade steps complete, you can:</p> <ul style="list-style-type: none"> <li>• Enable or disable new features.</li> <li>• Change the grid configuration.</li> <li>• Change the audit level configuration.</li> <li>• Update the ILM configuration.</li> <li>• Perform decommissioning, expansion, and recovery procedures.</li> </ul>

**Steps**

1. Multiply the number of nodes in your StorageGRID system by 45 minutes/node (average).
2. Multiply the number of Storage Nodes by 3 minutes for the Cassandra upgrade step, and add this number to the total.
3. Add 1 hour to this time to account for the time required to download the .upgrade file, upgrade the primary Admin Node, and complete the final steps.

**Example: Estimating the time to upgrade from StorageGRID 11.2 to 11.3**

Suppose your system has 14 grid nodes, including 9 Storage Nodes. The estimated time to upgrade all nodes is 12 hours.

```

630 minutes (14 × 45 minutes/node)
+ 27 minutes (9 × 3 minutes/Storage Node)
+ 60 minutes (Admin Node and final steps)
-----
717 minutes (~12 hours)

```

## How your system is affected during the upgrade

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

**Client applications might experience short-term disruptions**

The StorageGRID system can ingest and retrieve data from client applications throughout the upgrade process except for a short period of time when services are restarting and the client connections to individual Gateway Nodes or Storage Nodes are disrupted. Connectivity will be restored after the upgrade finishes and services resume on the individual nodes.

Every node in your StorageGRID system will be brought down one at a time for several minutes each during the upgrade. You might need to schedule downtime for the upgrade if loss of connectivity for a short period is not acceptable.

You must decide when to upgrade Gateway Nodes based on your grid's configuration. If your StorageGRID system has multiple Gateway Nodes, you must sequence the upgrade so that client applications are always directed to an available Gateway Node. If your StorageGRID system has only one Gateway Node, you must plan a downtime for the upgrade because client applications will not be able to access the system while the Gateway Node is being upgraded.

**Alarms might be triggered**

Alarms 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). In general, these alarms will clear when the upgrade completes.

**Many emails are generated**

When you upgrade grid nodes, email notifications are generated when the node is stopped and restarted. To avoid excessive emails, you can disable email notifications before upgrading the first node and re-enable notifications after the upgrade is completed.

**Configuration changes are restricted**

While you are upgrading StorageGRID:

- Do not make any grid configuration changes until the upgrade is complete.
- Do not change the audit level configuration until the upgrade is complete.
- Do not enable or disable any new features until the upgrade is complete.
- Do not update the ILM configuration until the upgrade is complete. Otherwise, you might experience inconsistent and unexpected ILM behavior.

### Metadata reserved space increases on larger-capacity Storage Nodes

During upgrade, the metadata reserved space for Storage Nodes that have 128 GB or more of RAM is automatically increased from 3 TB to 4 TB, unless the Metadata Reserved Space (CAWM) setting has been changed from its default value of 3 TB (3000000000000). If you have changed the value of Metadata Reserved Space, or if the Storage Node was installed when the default value for Metadata Reserved Space was different than 3 TB, the amount of space that is reserved for metadata on volume 0 is not changed during upgrade.

To see the current setting for Metadata Reserved Space (CAWM), go to **Configuration > Storage Options > Overview**. Note that because CAWM is a system-wide setting, its value does not change after the upgrade to reflect the new 4 TB metadata reservation for larger Storage Nodes.

To see the impact of the increased amount of metadata reserved space on larger Storage Nodes, select **Nodes > Storage Node > Storage** and look at the graph of **Storage Used - Object Metadata**. The value of the Metadata Used Space (Percent) attribute, or CDLP, decreases after a larger Storage Node is upgraded, due to the increased space reservation.

### Changes to open firewall ports

During the upgrade, open firewall ports are reset to the default set.

## 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 permissions for the Grid Manager

The following management permission has been changed in StorageGRID 11.3.

Permission	Description
Root Access	Required to perform these tasks: <ul style="list-style-type: none"> <li>• Access any of the new Alerts (preview) options from the <b>Alarms</b> menu</li> <li>• Enable a Storage proxy or an Admin proxy from <b>Configuration &gt; Proxies</b></li> <li>• Access <b>Configuration &gt; Load Balancer Endpoints</b></li> <li>• Access <b>Configuration &gt; High Availability Groups</b></li> <li>• Access <b>Configuration &gt; Untrusted Client Networks</b></li> </ul>

### Related information

[Administering StorageGRID](#)



## Verifying the installed version of StorageGRID

Before starting the upgrade, you must verify which version of the StorageGRID is currently installed.

### Steps

1. Sign in to the Grid Manager using a supported browser.
2. Select **Help > About**.
3. Verify that the **Version** is 11.2.x.

**Attention:** If you have an earlier version of the software, you must upgrade to version 11.2.x before proceeding with these steps.

### Related information

[Administering StorageGRID](#)

## Obtaining 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	<p>You must download one of the following sets of files to your service laptop:</p> <ul style="list-style-type: none"> <li>• Upgrade file for VMware</li> <li>• Upgrade file and RPM file (.zip or .tgz) for Red Hat Enterprise Linux or CentOS</li> <li>• Upgrade file and DEB file (.zip or .tgz) for Ubuntu or Debian</li> </ul>
Service laptop	<p>The service laptop must have:</p> <ul style="list-style-type: none"> <li>• Network port</li> <li>• SSH client (for example, PuTTY)</li> </ul>
Supported web browser	<p>You must confirm that the web browser on the service laptop is supported for use with StorageGRID 11.3.</p> <p>See “Web browser requirements.”</p> <p><b>Note:</b> Browser support has changed for StorageGRID 11.3. Confirm you are using a supported version.</p>
Recovery Package (.zip) file	<p>Before upgrading, you should download the most recent Recovery Package file in case any problems occur during the upgrade.</p> <p>After you upgrade the primary Admin Node, you must download a new copy of the Recovery Package file and save it in a safe location. The updated Recovery Package file allows you to restore the system if a failure occurs.</p> <p>See “Downloading the Recovery Package” for instructions.</p>

Item	Notes
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.
Provisioning passphrase	The passphrase is created and documented when the StorageGRID system is first installed. The provisioning passphrase is not listed in the Passwords.txt file.
Related documentation	<ul style="list-style-type: none"> <li>• Release notes for StorageGRID 11.3. Be sure to read these carefully before starting the upgrade.</li> <li>• Instructions for administering StorageGRID</li> <li>• If you are upgrading a Linux deployment, the StorageGRID installation instructions for your Linux platform.</li> <li>• Other StorageGRID documentation, as required.</li> </ul>

**Related tasks**

[Downloading the StorageGRID upgrade file](#) on page 19

[Downloading the Recovery Package](#) on page 20

**Related references**

[Web browser requirements](#) on page 18

**Related information**

[Administering StorageGRID](#)

[Red Hat Enterprise Linux or CentOS installation](#)

[Ubuntu or Debian installation](#)

[VMware installation](#)

[StorageGRID release notes](#)

**Web browser requirements**

You must use a supported web browser.

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

You should set the browser window to a recommended width.

Browser width	Pixels
Minimum	1024
Optimum	1280

## Downloading the StorageGRID upgrade file

You must download the upgrade file to a service laptop before you upgrade your StorageGRID system. If your StorageGRID system is deployed on Linux hosts, you must also download the StorageGRID installation files.

### Steps

1. Go to the NetApp Downloads page for StorageGRID.  
*NetApp Downloads: StorageGRID*
2. Click the button for downloading the latest release, or select another version from the pull-down menu and click **Go**.
3. Sign in using the username and password for your NetApp account.
4. Read and accept the End User License Agreement.  
The downloads page for the version you selected appears. The page contains columns for new installation files, upgrade files, and NAS Bridge.
5. In the **Upgrade** column, select and download the upgrade archive for your platform.
6. If your StorageGRID system is deployed on Linux hosts, complete these steps before you start the upgrade.
  - a. Download either the `.tgz` file or the `.zip` file for your Linux platform.

**Note:** Select 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"> <li>• <code>StorageGRID-WebScale-version-RPM-uniqueID.zip</code></li> <li>• <code>StorageGRID-WebScale-version-RPM-uniqueID.tgz</code></li> </ul>
Ubuntu or Debian	<ul style="list-style-type: none"> <li>• <code>StorageGRID-WebScale-version-DEB-uniqueID.zip</code></li> <li>• <code>StorageGRID-WebScale-version-DEB-uniqueID.tgz</code></li> </ul>

- b. Extract the RPM or DEB packages from the installation file.
- c. Install the RPM or DEB packages on all Linux hosts as described in “Installing StorageGRID host services” in the installation instructions for your Linux platform.
- d. For each containerized StorageGRID node on each Linux host, run the following commands in this order:

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

```
storagegrid node start <node-name>
```

Make sure each node boots correctly before taking the next node down.

### Related information

*Red Hat Enterprise Linux or CentOS installation*

*Ubuntu or Debian installation*

## Downloading the Recovery Package

You must download an updated copy of the Recovery Package file before and after making grid topology changes to the StorageGRID system and before and after upgrading the software. The Recovery Package file allows you to restore the system if a failure occurs.

### Before you begin

- You must be signed in to the Grid Manager using a supported browser.
- You must have the provisioning passphrase.
- You must have specific access permissions.

### Steps

1. Select **Maintenance > Recovery Package**.
2. Enter the provisioning passphrase, and click **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.  
**Attention:** The Recovery Package file must be secured because it contains encryption keys and passwords that can be used to obtain data from the StorageGRID system.

### Related information

[Administering StorageGRID](#)

## Checking the system's condition before upgrading software

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 browser.
2. Check for and resolve any active alarms.  
For information on specific alarms, see the monitoring and troubleshooting instructions.
3. Confirm that no conflicting grid tasks are active or pending.
  - a. Select **Support > 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.

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

4. Refer to the lists of internal and external ports in the 11.3 version of the installation instructions for your platform, and ensure that all required ports are opened before you upgrade.

**Attention:** 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.

StorageGRID 11.3 supports the use of three new ports:

- Required internal port. You must ensure that this port is not blocked between sites or nodes:
  - 7443: Used for internal HTTP traffic for maintenance procedures and error reporting.
- Optional external ports. If you want to separate Grid Manager communications from Tenant Manager communications, you can now use these ports between your browser and Admin Nodes:
  - 8443: Used by web browsers and management API clients for accessing the Grid Manager.
  - 9443: Used by web browsers and management API clients for accessing the Tenant Manager.

**Note:** Ports used in optionally configurable load balancer endpoints are not enabled by default. After upgrading to StorageGRID 11.3, you can choose whether you want to use load balancer endpoints.

#### Related information

*[Monitoring and troubleshooting StorageGRID](#)*

*[Administering StorageGRID](#)*

*[Recovery and maintenance](#)*

*[Red Hat Enterprise Linux or CentOS installation](#)*

*[Ubuntu or Debian installation](#)*

*[VMware installation](#)*

## Performing the upgrade

---

The Software Upgrade page guides you through the process of uploading the required file and upgrading all of the grid nodes in your StorageGRID system.

### Before you begin

You are aware of the following:

- You must upgrade all grid nodes for all data center sites from the primary Admin Node, using the Grid Manager.
- To detect and resolve issues, you can manually run the upgrade prechecks before starting the actual upgrade. The same prechecks are performed when you start the upgrade. Precheck failures will stop the upgrade process and might require technical support involvement to resolve.
- When you start the upgrade, the primary Admin Node is upgraded automatically.
- Shortly after the primary Admin Node has been upgraded, you can select which grid nodes to upgrade next.
- You must upgrade all grid nodes in your StorageGRID system to complete the upgrade, but you can upgrade individual grid nodes in any order. You can select individual grid nodes, groups of grid nodes, or all grid nodes. You can repeat the process of selecting grid nodes as many times as necessary, until all grid nodes at all sites are upgraded.
- When the upgrade starts on a grid node, the services on that node are stopped. Later, the grid node is rebooted. Do not approve the upgrade for a grid node unless you are sure that node is ready to be stopped and rebooted.
- When all grid nodes have been upgraded, the Cassandra database is upgraded, and the upgrade process completes.

**Note:** The upgrade from StorageGRID 11.2 to 11.3 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 or weeks, based on the amount of metadata in your system.

- You must complete the upgrade on the same hypervisor platform you started with.

### Steps

1. [Starting the upgrade](#) on page 23
2. [Upgrading grid nodes and completing the upgrade](#) on page 27
3. [Verifying the completion of your upgrade](#) on page 30

### Related tasks

[Estimating the time to complete an upgrade](#) on page 13

### Related information

[Administering StorageGRID](#)

## Starting the upgrade

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

### Before you begin

You have reviewed all of the considerations and completed all of the steps in “Upgrade planning and preparation.”

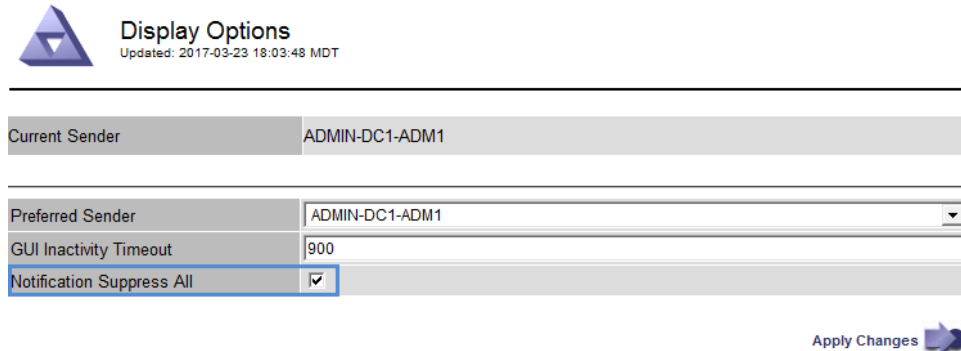
### Steps

1. Sign in to the Grid Manager using a supported browser.
2. Optionally, disable email notifications for alarms.


You can disable email notifications for alarms during the upgrade to avoid receiving excessive email notifications about node outages and upgrade processes.

- a. Select **Configuration > Display Options**.
- b. Select the **Notification Suppress All** check box.

All email notifications are suppressed when this check box is selected, including those unrelated to the upgrade, such as event-triggered AutoSupport email notifications.



Display Options	
Updated: 2017-03-23 18:03:48 MDT	
Current Sender	ADMIN-DC1-ADM1
Preferred Sender	ADMIN-DC1-ADM1
GUI Inactivity Timeout	900
Notification Suppress All	<input checked="" type="checkbox"/>

Apply Changes 

- c. Click **Apply Changes**.
3. Select **Maintenance > Software Upgrade**.

The Software Upgrade page appears. The date and time that the most recent upgrade completed are displayed, unless the primary Admin Node has been rebooted or the management API restarted since that upgrade was performed.

Software Upgrade

Before starting the upgrade process, you must confirm that there are no active alerts and that all grid nodes are online and available. Temporarily disable email notifications during the upgrade to avoid node outage warnings (Configuration > Display Options > Notification Suppress All).

After uploading the upgrade file, click the Run Prechecks button to detect problems that will prevent upgrade from starting. These prechecks also run when you start the upgrade.

Software upgrade completed at 2018-02-26 18:05:25 MST.

Upgrade file

Upgrade file

Upgrade Version No software upgrade file selected

Passphrase

Provisioning Passphrase

4. Select the .upgrade file you downloaded.
  - a. Click **Browse**.
  - b. Locate and select the file:  
NetApp\_StorageGRID\_version\_Software\_uniqueID.upgrade
  - c. Click **Open**.

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

Software Upgrade

Before starting the upgrade process, you must confirm that there are no active alerts and that all grid nodes are online and available. Temporarily disable email notifications during the upgrade to avoid node outage warnings (Configuration > Display Options > Notification Suppress All).

After uploading the upgrade file, click the Run Prechecks button to detect problems that will prevent upgrade from starting. These prechecks also run when you start the upgrade.

Upgrade file

Upgrade file  ✓ NetApp\_StorageGRID\_11.3.0\_Software\_20190719.2216.9a87d41

Upgrade Version StorageGRID® 11.3.0

Passphrase

Provisioning Passphrase

5. Enter the provisioning passphrase in the text box.  
The **Run Prechecks** and **Start Upgrade** buttons become enabled.



Software Upgrade

Before starting the upgrade process, you must confirm that there are no active alerts and that all grid nodes are online and available. Temporarily disable email notifications during the upgrade to avoid node outage warnings (Configuration > Display Options > Notification Suppress All).

After uploading the upgrade file, click the Run Prechecks button to detect problems that will prevent upgrade from starting. These prechecks also run when you start the upgrade.

**Upgrade file**

---

Upgrade file  ✔ NetApp\_StorageGRID\_11.3.0\_Software\_20190719.2216.9a87d41

Upgrade Version StorageGRID® 11.3.0

**Passphrase**

---

Provisioning Passphrase

6. If you want to validate the condition of your system before you start the actual upgrade, click **Run Prechecks**. Then, resolve any precheck errors that are reported.

**Note:** The same prechecks are performed when you click **Start Upgrade**. Clicking **Run Prechecks** allows you to detect and resolve issues before starting the upgrade.

**Attention:** 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. When you are ready to perform the upgrade, click **Start Upgrade**.

A warning box appears to remind you that your browser's connection will be lost when the primary Admin Node is rebooted. When the primary Admin Node is available again, you will need to clear your web browser's cache and reload the Software Upgrade page.

**⚠ Connection Will Be Temporarily Lost**

During the upgrade, your browser's connection to StorageGRID Webscale will be lost temporarily when the primary Admin Node is rebooted.

**Attention:** You must clear your cache and reload the page before starting to use the new version. Otherwise, StorageGRID Webscale might not respond as expected.

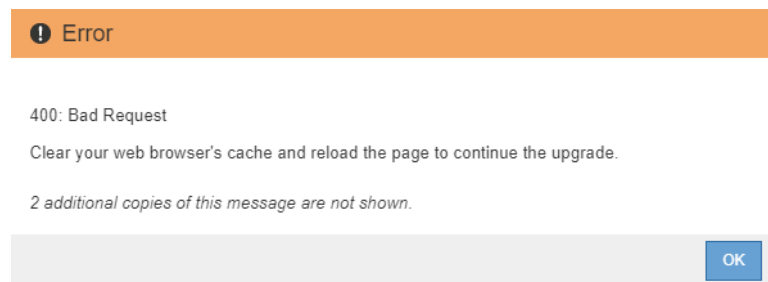
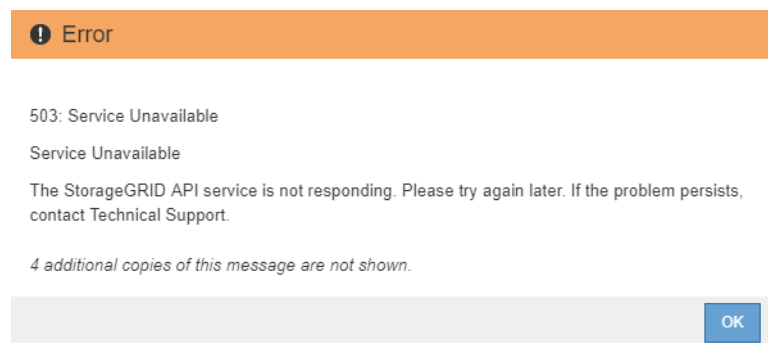
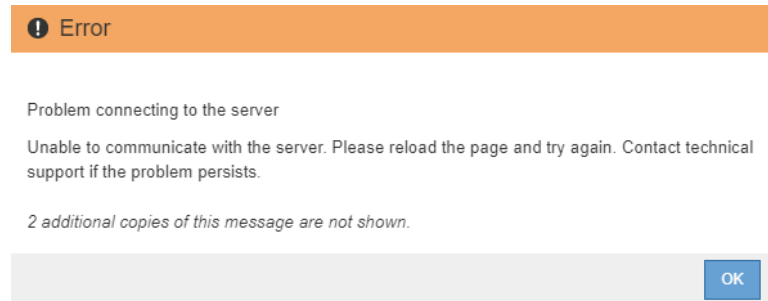
Are you sure you want to start the upgrade process?

8. Click **OK** to acknowledge the warning and start the upgrade process.

When the upgrade starts:

- a. The upgrade prechecks are run.
  - Note:** If any precheck errors are reported, resolve them and click **Start Upgrade** again.
- b. The primary Admin Node is upgraded, which includes stopping services, upgrading the software, and restarting services. You will not be able to 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.

**Note:** While the primary Admin Node is being upgraded, multiple copies of the following error messages appear:



9. After the primary Admin Node has been upgraded, clear your web browser's cache, sign back in, and reload the Software Upgrade page.

For instructions, see the documentation for your web browser.

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

#### Related concepts

[Upgrade planning and preparation](#) on page 13

## Upgrading grid nodes and completing the upgrade

After the primary Admin Node has been upgraded, you must upgrade all other grid nodes in your StorageGRID system. You can customize the upgrade sequence by selecting to upgrade individual grid nodes, groups of grid nodes, or all grid nodes.

### Steps

1. Review the Upgrade Progress section on the Software 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 Completed, 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 all grid nodes in your system.
2. 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.

**Attention:** 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.

3. Review the information in the Grid Node Status table. Grid nodes are arranged by type.

## 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.

Site	Name	Progress	Stage	Error	Action
Data Center 1	DC1-ADM1	<div style="width: 100%; height: 10px; background-color: green;"></div>	Done		

Site	Name	Progress	Stage	Error	Action
Data Center 1	DC1-S1	<div style="width: 25%; height: 10px; background-color: #00a0e3;"></div>	Waiting for you to approve		<input type="button" value="Approve"/>
Data Center 1	DC1-S2	<div style="width: 25%; height: 10px; background-color: #00a0e3;"></div>	Waiting for you to approve		<input type="button" value="Approve"/>
Data Center 1	DC1-S3	<div style="width: 25%; height: 10px; background-color: #00a0e3;"></div>	Waiting for you to approve		<input type="button" value="Approve"/>

A grid node can have 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
4. Optionally, sort the lists of nodes in ascending or descending order by **Site**, **Name**, **Progress**, **Stage**, or **Error**. Or, enter a term in the **Search** box to search for specific nodes.
  5. Approve the grid nodes you are ready to add to the upgrade queue.

**Attention:** When the upgrade starts on a grid node, the services on that node are stopped. Later, the grid node is rebooted. These operations might cause service interruptions for clients 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.

- Click one or more **Approve** buttons to add one or more individual nodes to the upgrade queue.
- Click the **Approve All** button within each section to add all nodes of the same type to the upgrade queue.

**Note:** Nodes of the same type are upgraded one at a time.

- Click the top-level **Approve All** button to add all nodes in the grid to the upgrade queue.

6. If you need to remove a node or all nodes from the upgrade queue, click **Remove** or **Remove All**.

As shown in the example, when the Stage reaches “Stopping services,” the **Remove** button is hidden and you can no longer remove the node.

Site	Name	Progress	Stage	Error	Action
Data Center 1	DC1-S1	<div style="width: 100%;"></div>	Stopping services		
Data Center 1	DC1-S2	<div style="width: 50%;"></div>	Queued		Remove
Data Center 1	DC1-S3	<div style="width: 50%;"></div>	Queued		Remove

7. Wait for each node to proceed through the upgrade stages, which include applying the upgrade, stopping services, upgrading the base operating system, rebooting, and starting services.

When all grid nodes have been upgraded, the Upgrade Grid Nodes task is shown as Completed, and the Upgrade Cassandra task starts. During this stage, the upgrade process checks each node to verify that the Cassandra database does not need to be updated.

**Note:** The upgrade from StorageGRID 11.2 to 11.3 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 or weeks, based on the amount of metadata in your system.

8. When the Upgrade Cassandra upgrade task has completed, wait a few minutes for the Final Upgrade Steps task to complete.

**Upgrade Progress**

Start Upgrade Service	Completed
Upgrade Grid Nodes	Completed
Upgrade Cassandra	Completed
Final Upgrade Steps	In Progress

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

**Related tasks**

[Downloading the Recovery Package](#) on page 20

## Verifying the completion of your upgrade

You must verify that the upgrade completed successfully and make any required configuration changes to ensure that your grid is operating optimally.

**About this task**

This procedure asks you to re-verify some items that you checked while the database upgrade was still in progress. You must ensure that the last stages of the upgrade completed successfully.

**Steps**

1. Sign in to the Grid Manager using a supported browser.
2. Confirm that the upgrade completed successfully.
  - a. Click **Help > About**, and confirm that the displayed version is what you would expect.
  - b. Select **Maintenance > Software Upgrade**.
  - c. Confirm that the green banner shows that the software upgrade was completed on the date and time you expected.

## Software Upgrade

Before starting the upgrade process, you must confirm that there are no active alerts and that all grid nodes are online and available. Temporarily disable email notifications during the upgrade to avoid node outage warnings (Configuration > Display Options > Notification Suppress All).

After uploading the upgrade file, click the Run Prechecks button to detect problems that will prevent upgrade from starting. These prechecks also run when you start the upgrade.

Software upgrade completed at 2019-07-24 06:27:42 MDT.

**Upgrade file**

Upgrade file

Upgrade Version No software upgrade file selected

**Passphrase**

Provisioning Passphrase

3. Verify that grid operations have returned to normal:
  - a. Check that the services are operating normally and that there are no new alarms.
  - b. Review all custom alarms to verify that they are still required and usable.
  - c. Confirm that client connections to the StorageGRID system are operating as expected.
4. Re-enable email notifications if you suppressed them for the upgrade.
  - a. Select **Configuration > Display Options**.

- b. Unselect the **Notification Suppress All** check box.
- c. Click **Apply Changes**.

## Troubleshooting upgrade issues

---

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 tasks

*[Checking the system's condition before upgrading software](#) on page 20*

## 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. For example, after you use the Grid Manager to acknowledge an alarm and click **Apply Changes**, the change might not be saved.

If you experience issues with the web interface:

- Make sure you are using a supported browser.
  - Note:** Browser support has changed for StorageGRID 11.3. Confirm you are using a supported version.
- 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.

**Related references**

[Web browser requirements](#) on page 18

## “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 <code>{file_path}</code> 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 <code>{file_path}</code> 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 <code>{node_name}</code> .	A critical file on the node is corrupt.	Contact technical support.
Connection error while attempting to list versions on <code>{node_name}</code>	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 <code>{node_name}</code> does not have StorageGRID <code>{upgrade_version}</code> image loaded. Images and services must be installed on the host before the upgrade can proceed.	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.  <b>Note:</b> This error only applies to nodes that are running as containers on Linux.	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.  For more information, see the installation instructions for your Linux platform.
Error while checking node <code>{node_name}</code>	An unexpected error occurred.	Wait a few minutes, and try again.
Uncaught error while running prechecks. <code>{error_string}</code>	An unexpected error occurred.	Wait a few minutes, and try again.

**Related information**

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

## Copyright

---

Copyright © 2019 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.

Data contained herein pertains to a commercial item (as defined in FAR 2.101) and is proprietary to NetApp, Inc. The U.S. Government has a non-exclusive, non-transferrable, non-sublicensable, 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).

## Trademark information

---

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

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

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

---

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

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

[\*doccomments@netapp.com\*](mailto:doccomments@netapp.com)

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

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

You can also contact us in the following ways:

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