



Upgrade software

StorageGRID 11.5

NetApp
August 30, 2024

Table of Contents

- Upgrade software 1
 - About StorageGRID 11.5 1
 - Upgrade planning and preparation 14
 - Performing the upgrade 25
 - Troubleshooting upgrade issues 37

Upgrade software

Learn how to upgrade a StorageGRID system to a new release.

- [About StorageGRID 11.5](#)
- [Upgrade planning and preparation](#)
- [Performing the upgrade](#)
- [Troubleshooting upgrade issues](#)

About StorageGRID 11.5

Before starting an upgrade, review this section to learn about the new features and enhancements in StorageGRID 11.5, determine whether any features have been deprecated or removed, and find out about changes to StorageGRID APIs.

- [What's new in StorageGRID 11.5](#)
- [Removed or deprecated features](#)
- [Changes to the Grid Management API](#)
- [Changes to the Tenant Management API](#)

What's new in StorageGRID 11.5

StorageGRID 11.5 introduces S3 Object Lock, support for KMIP encryption of data, usability improvements to ILM, a redesigned Tenant Manager user interface, support for decommissioning a StorageGRID site, and an appliance node clone procedure.

S3 Object Lock for compliant data

The S3 Object Lock feature in StorageGRID 11.5 is an object-protection solution that is equivalent to S3 Object Lock in Amazon Simple Storage Service (Amazon S3). You can enable the global S3 Object Lock setting for a StorageGRID system to allow S3 tenant accounts to create buckets with S3 Object Lock enabled. The tenant can then use an S3 client application to optionally specify retention and legal hold settings for the objects in those buckets.

S3 Object Lock lets tenant users comply with regulations that require certain objects to be retained for a fixed amount of time or indefinitely.

Learn more

- [Manage objects with ILM](#)
- [Use S3](#)
- [Use a tenant account](#)

KMS encryption key management

You can now configure one or more external key management servers (KMS) in the Grid Manager to provide encryption keys to StorageGRID services and storage appliances. Each KMS or KMS cluster uses the Key Management Interoperability Protocol (KMIP) to provide an encryption key to the appliance nodes at the

associated StorageGRID site. After the appliance volumes are encrypted, you cannot access any data on the appliance unless the node can communicate with the KMS.



If you want to use encryption key management, you must use the StorageGRID Appliance Installer to enable the **Node Encryption** setting for the appliance before you add the appliance to the grid.

Learn more

- [Administer StorageGRID](#)

Usability enhancements for information lifecycle management (ILM)

- You can now view the total capacity of a storage pool, including the amount of used and free space. You can also see which nodes are included in a storage pool and which ILM rules and Erasure Coding profiles use the storage pool.
- You can now design ILM rules that apply to more than one tenant account.
- When you create an ILM rule for erasure coding, you are now reminded to set the Object Size (MB) advanced filter to greater than 0.2 to ensure that very small objects are not erasure coded.
- The ILM policy interface now ensures that the default ILM rule will be always be used for any objects not matched by another rule. Starting in StorageGRID 11.5, the default rule cannot use any basic or advanced filters and is automatically placed as the last rule in the policy.



If your current ILM policy does not conform to the new requirements, you can continue to use it after you upgrade to StorageGRID 11.5. However, if you attempt to clone a non-conforming policy after you upgrade, you are prompted to select a default rule that does not include filters and you are required to place the default rule at the end of the policy.

- The stock All Storage Nodes storage pool is no longer selected by default when you create a new ILM rule or a new Erasure Coding profile. In addition, you can now remove the All Storage Nodes storage pool as long as it not used in any rule.



Using the All Storage Nodes storage pool is not recommended because this storage pool contains all sites. Multiple copies of an object might be placed on the same site if you use this storage pool with a StorageGRID system that includes more than one site.

- You can now remove the stock Make 2 Copies rule (which uses the All Storage Nodes storage pool) as long as it is not used in an active or proposed policy.
- Objects stored in a Cloud Storage Pool can now be deleted immediately (synchronous deletion).

Learn more

- [Manage objects with ILM](#)

Enhancements to the Grid Manager

- The redesigned Tenant Accounts page makes it easier to view tenant account usage. The tenant summary table now includes columns for Space Used, Quota Utilization, Quota, and Object Count. A new **View Details** button accesses an overview of each tenant as well as details about the account's S3 buckets or Swift containers. In addition, you can now export two `.csv` files for tenant usage: one containing usage values for all tenants and one containing details about a tenant's buckets or containers.

Related to this change, three new Prometheus metrics were added to track tenant account usage:

- storagegrid_tenant_usage_data_bytes
- storagegrid_tenant_usage_object_count
- storagegrid_tenant_usage_quota_bytes

- The new **Access Mode** field on the Admin Groups page (**Configuration > Access Control**) allows you to specify whether the management permissions for the group are read-write (default) or read-only. Users who belong to a group with read-write access mode can change settings and perform operations in the Grid Manager and the Grid Management API. Users who belong to a group with read-only access mode can only view the settings and features that are selected for the group.



When you upgrade to StorageGRID 11.5, the read-write access mode option is selected for all existing admin groups.

- The AutoSupport user interface was redesigned. You can now configure event-triggered, user-triggered, and weekly AutoSupport messages from a single page in the Grid Manager. You can also configure an additional destination for AutoSupport messages.



If AutoSupport has not been enabled, a reminder message now appears on the Grid ManagerDashboard.

- When viewing the **Storage Used - Object Data** chart on the Nodes page, you can now see estimates for the amount of replicated object data and the amount of erasure-coded data on the grid, site, or Storage Node (**Nodes > grid/site/Storage Node > Storage**).
- Grid Manager menu options were reorganized to make options easier to find. For example, a new **Network Settings** submenu was added to the **Configuration** menu and options in the **Maintenance** and **Support** menus are now listed in alphabetic order.

Learn more

- [Administer StorageGRID](#)

Enhancements to the Tenant Manager

- The appearance and organization of the Tenant Manager user interface has been completely redesigned to improve the user experience.
- The new Tenant Manager dashboard provides a high-level summary of each account: it provides bucket details and shows the number of buckets or containers, groups, users, and platform services endpoints (if configured).

Learn more

- [Use a tenant account](#)

Client certificates for Prometheus metrics export

You can now upload or generate client certificates (**Configuration > Access Control > Client Certificates**), which can be used to provide secure, authenticated access to the StorageGRID Prometheus database. For example, you can use client certificates if you need to monitor StorageGRID externally using Grafana.

Learn more

- [Administer StorageGRID](#)

Load balancer enhancements

- When handling routing requests at a site, the Load Balancer service now performs load aware routing: it considers the CPU availability of the Storage Nodes at the same site. In some cases, information about CPU availability is limited to the site where the Load Balancer service is located.



CPU awareness will be not enabled until at least two-thirds of the Storage Nodes at a site have been upgraded to StorageGRID 11.5 and are reporting CPU statistics.

- For added security, you can now specify a binding mode for each load balancer endpoint. Endpoint pinning lets you restrict the accessibility of each endpoint to specific high availability groups or node interfaces.

Learn more

- [Administer StorageGRID](#)

Object metadata changes

- **New Actual reserved space metric:** To help you understand and monitor object metadata space usage on each Storage Node, a new Prometheus metric is shown on the Storage Used - Object Metadata graph for a Storage Node (**Nodes** > **Storage Node** > **Storage**).

```
storagegrid_storage_utilization_metadata_reserved
```

The **Actual reserved space** metric indicates how much space StorageGRID has reserved for object metadata on a specific Storage Node.

- **Metadata space increased for installations with larger Storage Nodes:** The system-wide Metadata Reserved Space setting has been increased for StorageGRID systems containing Storage Nodes with 128 GB or more of RAM, as follows:
 - **8 TB for new installations:** If you are installing a new StorageGRID 11.5 system and each Storage Node in the grid has 128 GB or more of RAM, the system-wide Metadata Reserved Space setting is now set to 8 TB instead of 3 TB.
 - **4 TB for upgrades:** If you are upgrading to StorageGRID 11.5 and each Storage Node at any one site has 128 GB or more of RAM, the system-wide Metadata Reserved Space setting is now set to 4 TB instead of 3 TB.

The new values for the Metadata Reserved Space setting increase the allowed metadata space for these larger Storage Nodes, up to 2.64 TB, and ensure that adequate metadata space is reserved for future hardware and software versions.



If your Storage Nodes have enough RAM and sufficient space on volume 0, you can manually increase the Metadata Reserved Space setting up to 8 TB after you upgrade. Reserving additional metadata space after the StorageGRID 11.5 upgrade will simplify future hardware and software upgrades.

[Increasing the Metadata Reserved Space setting](#)



If your StorageGRID system stores (or is expected to store) more than 2.64 TB of metadata on any Storage Node, the allowed metadata space can be increased in some cases. If your Storage Nodes each have available free space on storage volume 0 and more than 128 GB of RAM, contact your NetApp account representative. NetApp will review your requirements and increase the allowed metadata space for each Storage Node, if possible.

- **Automatic cleanup of deleted metadata:** When 20% or more of the metadata stored on a Storage Node is ready to be removed (because the corresponding objects were deleted), StorageGRID can now perform an automatic compaction on that Storage Node. This background process only runs if the load on the system is low—that is, when there is available CPU, disk space, and memory. The new compaction process removes metadata for deleted objects sooner than in previous releases and helps to free up space for new objects to be stored.

Learn more

- [Administer StorageGRID](#)

Changes to S3 REST API support

- You can now use the S3 REST API to specify [S3 Object Lock](#) settings:
 - To create a bucket with S3 Object Lock enabled, use a PUT Bucket request with the `x-amz-bucket-object-lock-enabled` header.
 - To determine if S3 Object Lock is enabled for a bucket, use a GET Object Lock Configuration request.
 - When adding an object version to a bucket with S3 Object Lock enabled, use the following request headers to specify the retention and legal hold settings: `x-amz-object-lock-mode`, `x-amz-object-lock-retain-until-date`, and `x-amz-object-lock-legal-hold`.
- You can now use DELETE Multiple Objects on a versioned bucket.
- You can now use PUT, GET, and DELETE Bucket encryption requests to manage encryption for an existing S3 bucket.
- A minor change was made to a field name for the `Expiration` parameter. This parameter is included in the response to a PUT Object, HEAD Object, or GET Object request if an expiration rule in the lifecycle configuration applies to a specific object. The field that indicates which expiration rule was matched was previously named `rule_id`. This field was renamed to `rule-id` to match the AWS implementation.
- By default, the S3 GET Storage Usage request now attempts to retrieve the storage used by a tenant account and its buckets using strong-global consistency. If strong-global consistency cannot be achieved, StorageGRID attempts to retrieve the usage information using strong-site consistency.
- The `Content-MD5` request header is now correctly supported.

Learn more

- [Use S3](#)

Maximum size for CloudMirror objects increased to 5 TB

The maximum size for objects that can be replicated to a destination bucket by the CloudMirror replication service was increased to 5 TB, which is the maximum object size supported by StorageGRID.

Learn more

- [Use S3](#)

- [Use Swift](#)

New alerts added

The following new alerts were added for StorageGRID 11.5:

- Appliance BMC communication error
- Appliance Fibre Channel fault detected
- Appliance Fibre Channel HBA port failure
- Appliance LACP port missing
- Cassandra auto-compactor error
- Cassandra auto-compactor metrics out of date
- Cassandra compactions overloaded
- Disk I/O is very slow
- KMS CA certificate expiration
- KMS client certificate expiration
- KMS configuration failed to load
- KMS connectivity error
- KMS encryption key name not found
- KMS encryption key rotation failed
- KMS is not configured
- KMS key failed to decrypt an appliance volume
- KMS server certificate expiration
- Low free space for storage pool
- Node network reception frame error
- Services appliance storage connectivity degraded
- Storage appliance storage connectivity degraded (previously named Appliance storage connectivity degraded)
- Tenant quota usage high
- Unexpected node reboot

Learn more

- [Monitor & troubleshoot](#)

TCP support for SNMP traps

You can now select Transmission Control Protocol (TCP) as the protocol for SNMP trap destinations. Previously, only the User Datagram Protocol (UDP) protocol was supported.

Learn more

- [Monitor & troubleshoot](#)

Installation and networking enhancements

- **MAC address cloning:** You can now use MAC address cloning to enhance the security of certain environments. MAC address cloning enables you to use a dedicated virtual NIC for the Grid Network, Admin Network, and Client Network. Having the Docker container use the MAC address of the dedicated NIC on the host allows you to avoid using promiscuous mode network configurations. Three new MAC address cloning keys were added to the node configuration file for Linux-based (bare metal) nodes.
- **Automatic discovery of DNS and NTP host routes:** Previously, there were restrictions on which network your NTP and DNS servers had to connect to, such as the requirement that you could not have all of your NTP and DNS servers on the Client Network. Now, those restrictions are removed.

Learn more

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

Support for rebalancing erasure-coded (EC) data after Storage Node expansion

The EC rebalance procedure is a new command-line script that might be required after you add new Storage Nodes. When you perform the procedure, StorageGRID redistributes erasure-coded fragments among the existing and the newly expanded Storage Nodes at a site.



You should only perform the EC rebalance procedure in limited cases. For example, if you cannot add the recommended number of Storage Nodes in an expansion, you can use the EC rebalance procedure to allow additional erasure-coded objects to be stored.

Learn more

- [Expand your grid](#)

New and revised maintenance procedures

- **Site decommission:** You can now remove an operational site from your StorageGRID system. The connected site decommission procedure removes an operational site and preserves data. The new Decommission Site wizard guides you through the process (**Maintenance > Decommission > Decommission Site**).
- **Appliance node cloning:** You can now clone an existing appliance node to upgrade the node to a new appliance model. For example, you can clone a smaller-capacity appliance node to a larger-capacity appliance. You can also clone an appliance node to implement new functionality, such as the new **Node Encryption** setting that is required for the KMS encryption.
- **Ability to change the provisioning passphrase:** You can now change the provisioning passphrase (**Configuration > Access Control > Grid Passwords**). The passphrase is required for recovery, expansion, and maintenance procedures.
- **Enhanced SSH password behavior:** To enhance the security of StorageGRID appliances, the SSH password is no longer changed when you place an appliance into maintenance mode. In addition, new SSH host certificates and host keys are generated when you upgrade a node to StorageGRID 11.5.



If you use SSH to log in to a node after upgrading to StorageGRID 11.5, you will receive a warning that the host key has changed. This behavior is expected and you can safely approve the new key.

Learn more

- [Maintain & recover](#)

Changes to StorageGRID appliances

- **Direct access to SANtricity System Manager for storage appliances:** You can now access the E-Series SANtricity System Manager user interface from the StorageGRID Appliance Installer and from the Grid Manager. Using these new methods enables access to SANtricity System Manager without using the management port on the appliance. Users who need to access SANtricity System Manager from the Grid Manager must have the new Storage Appliance Administrator permission.
- **Node encryption:** As part of the new KMS encryption feature, a new **Node Encryption** setting was added to the StorageGRID Appliance Installer. If you want to use encryption key management to protect appliance data, you must enable this setting during the hardware configuration stage of appliance installation.
- **UDP port connectivity:** You can now test the network connectivity of a StorageGRID appliance to UDP ports, such as those used for an external NFS or DNS server. From the StorageGRID Appliance Installer, select **Configure Networking > Port Connectivity Test (nmap)**.
- **Automating installation and configuration:** A new JSON configuration upload page was added to the StorageGRID Appliance Installer (**Advanced > Update Appliance Configuration**). This page enables you to use one file to configure multiple appliances in large grids. Additionally, the `configure-sga.py` Python script has been updated to match the capabilities of the StorageGRID Appliance Installer.

Learn more

- [SG100 & SG1000 services appliances](#)
- [SG6000 storage appliances](#)
- [SG5700 storage appliances](#)
- [SG5600 storage appliances](#)

Changes to audit messages

- **Automatic cleanup of overwritten objects:** Previously, objects that were overwritten were not removed from disk in specific cases, which resulted in additional space consumption. These overwritten objects, which are inaccessible to users, are now automatically removed to save storage space. Refer to the LKCU audit message for more information.
- **New audit codes for S3 Object Lock:** Four new audit codes were added to the SPUT audit message to include [S3 Object Lock](#) request headers:
 - LKEN: Object Lock Enabled
 - LKLH: Object Lock Legal Hold
 - LKMD: Object Lock Retention Mode
 - LKRU: Object Lock Retain Until Date
- **New fields for Last Modified Time and Previous Object Size:** You can now track when an object was overwritten as well as the original object size.
 - The MTME (Last Modified Time) field was added to the following audit messages:
 - SDEL (S3 DELETE)
 - SPUT (S3 PUT)
 - WDEL (Swift DELETE)
 - WPUT (Swift PUT)

- The CSIZ (Previous Object Size) field was added to the OVWR (Object Overwrite) audit message.

Learn more

- [Review audit logs](#)

New nms.requestlog file

A new log file, `/var/local/log/nms.requestlog`, is maintained on all Admin Nodes. This file contains information about outgoing connections from the Management API to internal StorageGRID services.

Learn more

- [Monitor & troubleshoot](#)

StorageGRID documentation changes

- To make networking information and requirements easier to find and to clarify that the information also applies to StorageGRID appliance nodes, the networking documentation was moved from the software-based installation guides (RedHat Enterprise Linux/CentOS, Ubuntu/Debian, and VMware) to a new networking guide.

[Network guidelines](#)

- To make ILM-related instructions and examples easier to find, the documentation for managing objects with information lifecycle management was moved from the *Administrator Guide* to a new ILM guide.

[Manage objects with ILM](#)

- A new FabricPool guide provides an overview of configuring StorageGRID as a NetApp FabricPool cloud tier and describes the best practices for configuring ILM and other StorageGRID options for a FabricPool workload.

[Configure StorageGRID for FabricPool](#)

- You can now access several instructional videos from the Grid Manager. The current videos provide instructions for managing alerts, custom alerts, ILM rules, and ILM policies.

Removed or deprecated features

Some features were removed or deprecated in StorageGRID 11.5. You must review these items to understand whether you need to update client applications or modify your configuration before you upgrade.

Weak consistency control removed

The Weak consistency control was removed for StorageGRID 11.5. After you upgrade, the following behaviors will apply:

- Requests to set Weak consistency for an S3 bucket or Swift container will succeed, but the consistency level will actually be set to Available.
- Existing buckets and containers that use Weak consistency will be silently updated to use Available consistency.
- Requests that have a Weak consistency-control header will actually use Available consistency, if

applicable.

The Available consistency control behaves the same as the “read-after-new-write” consistency level, but only provides eventual consistency for HEAD operations. The Available consistency control offers higher availability for HEAD operations than “read-after-new-write” if Storage Nodes are unavailable.



Alarm for grid health deprecated

The `/grid/health/topology` API, which checks for active *alarms* on nodes, is deprecated. In its place, a new `/grid/node-health` endpoint was added. This API returns the current status of each node by checking for active *alerts* on nodes.

Compliance feature deprecated

The S3 Object Lock feature in StorageGRID 11.5 replaces the Compliance feature that was available in previous StorageGRID versions. Because the new S3 Object Lock feature conforms to Amazon S3 requirements, it deprecates the proprietary StorageGRID Compliance feature, which is now referred to as “legacy Compliance.”

If you previously enabled the global Compliance setting, the new global S3 Object Lock setting is enabled automatically when you upgrade to StorageGRID 11.5. Tenant users will no longer be able to create new buckets with Compliance enabled in StorageGRID; however, as required, tenant users can continue to use and manage any existing legacy Compliant buckets.

In the Tenant Manager, a shield icon  indicates a legacy Compliant bucket. Legacy Compliant buckets might also have a hold badge  to indicate that the bucket is under a legal hold.

[KB: How to manage legacy Compliant buckets in StorageGRID 11.5](#)

[Manage objects with ILM](#)

“S3 multipart part too small” alert removed

The **S3 multipart part too small** alert was removed. Previous, this alert was triggered if an S3 client attempted to complete a multipart upload with parts that did not meet Amazon S3 size limits. After the upgrade to StorageGRID 11.5, any multipart upload requests that do not meet the following size limits will fail:

- Each part in a multipart upload must be between 5 MiB (5,242,880 bytes) and 5 GiB (5,368,709,120 bytes).
- The last part can be smaller than 5 MiB (5,242,880 bytes).
- In general, part sizes should be as large as possible. For example, use part sizes of 5 GiB for a 100 GiB object. Since each part is considered a unique object, using large part sizes reduces StorageGRID metadata overhead.
- For objects smaller than 5 GiB, consider using non-multipart upload instead.

"Appliance link down on Grid Network" alerts removed

The following alerts were removed. If the Grid Network is down, the metrics that would trigger these alerts are not accessible:

- Services appliance link down on Grid Network
- Storage appliance link down on Grid Network

Support for fully qualified domain name removed from SNMP configuration

When configuring an SNMP server in the baseboard management controller (BMC) for the SG6000, SG100, or SG1000, you must now specify an IP address instead of a fully qualified domain name. If a fully qualified domain name was previously configured, change it to an IP address before upgrading to StorageGRID 11.5.

Legacy attributes removed

The following legacy attributes were removed. As applicable, equivalent information is provided by Prometheus metrics:

Legacy attribute	Equivalent Prometheus metric
BREC	storagegrid_service_network_received_bytes
BTRA	storagegrid_service_network_transmitted_bytes
CQST	storagegrid_metadata_queries_average_latency_milliseconds
HAIS	storagegrid_http_sessions_incoming_attempted
HCCS	storagegrid_http_sessions_incoming_currently_established
HEIS	storagegrid_http_sessions_incoming_failed
HISC	storagegrid_http_sessions_incoming_successful
LHAC	<i>none</i>
NREC	<i>none</i>
NTSO (Chosen Time Source Offset)	storagegrid_ntp_chosen_time_source_offset_milliseconds
NTRA	<i>none</i>
SLOD	storagegrid_service_load
SMEM	storagegrid_service_memory_usage_bytes
SUTM	storagegrid_service_cpu_seconds
SVUT	storagegrid_service_uptime_seconds
TRBS (Total bits per second received)	<i>none</i>

Legacy attribute	Equivalent Prometheus metric
TRXB	storagegrid_network_received_bytes
TTBS (Total bits per second transmitted)	<i>none</i>
TTXB	storagegrid_network_transmitted_bytes

The following related changes were also made:

- The `network_received_bytes` and `network_transmitted_bytes` Prometheus metrics were changed from gauges to counters because the values of these metrics only increase. If you are currently using these metrics in Prometheus queries, you should start using the `increase()` function in the query.
- The Network Resources table was removed from the Resources tab for StorageGRID services. (Select **Support > Tools > Grid Topology**. Then, select **node > service > Resources**.)
- The HTTP Sessions page was removed for Storage Nodes. Previously, you could access this page by selecting **Support > Tools > Grid Topology** and then selecting **Storage Node > LDR > HTTP**.
- The HCCS (Currently Established Incoming Sessions) alarm was removed.
- The NTSO (Chosen Time Source Offset) alarm was removed.

Changes to the Grid Management API

StorageGRID 11.5 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 11.5; however, support for these versions of the API will be removed in a future release of StorageGRID. After upgrading to StorageGRID 11.5, the deprecated v1 and v2 APIs can be deactivated using the `PUT /grid/config/management` API.

New client-certificates section

The new section, `/grid/client-certificates`, allows you to configure client certificates to provide secure, authenticated access to the StorageGRID Prometheus database. For example, you can monitor StorageGRID externally using Grafana.

Legacy compliance endpoints moved to new s3-object-lock section

With the introduction of StorageGRID S3 Object Lock, the APIs used to manage the legacy compliance settings for the grid were moved to a new section of the Swagger user interface. The **s3-object-lock** section includes the two `/grid/compliance-global` API endpoints, which now control the global S3 Object Lock setting. The endpoint URIs remain unchanged for compatibility with existing applications.

Swift-admin-password Accounts endpoint removed

The following Accounts API endpoint, which was deprecated in StorageGRID 10.4, has now been removed:

```
https://<IP-Address>/api/v1/grid/accounts/<AccountID>/swift-admin-password
```

New grid-passwords section

The **grid-passwords** section enables operations for grid password management. The section includes two `/grid/change-provisioning-passphrase` API endpoints. The endpoints allow users to change the StorageGRID provisioning passphrase and retrieve the status of the passphrase change.

storageAdmin permission added to Groups API

The `/grid/groups` API now includes the `storageAdmin` permission.

New parameter for Storage Usage API

The `GET /grid/accounts/{id}/usage` API now has a `strictConsistency` parameter. To enforce a strong-global consistency when retrieving storage usage information across Storage Nodes, set this parameter to `true`. When this parameter is set to `false` (default), StorageGRID attempts to retrieve usage information using strong-global consistency, but falls back to strong-site consistency if strong-global consistency cannot be met.

New Node Health API

A new `/grid/node-health` endpoint was added. This API returns the current status of each node by checking for active *alerts* on the nodes. The `/grid/health/topology` API, which checks for active *alarms* on nodes, is deprecated.

Change to "ApplianceStorageShelvesPowerSupplyDegraded" alert rule ID

The alert rule ID "ApplianceStorageShelvesPowerSupplyDegraded" has been renamed to "ApplianceStorageShelvesDegraded" to better reflect the alert's actual behavior.

Related information

[Administer StorageGRID](#)

Changes to the Tenant Management API

StorageGRID 11.5 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 11.5; however, support for these versions of the API will be removed in a future release of StorageGRID. After upgrading to StorageGRID 11.5, the deprecated v1 and v2 APIs can be deactivated using the `PUT /grid/config/management` API.

New parameter for tenant Storage Usage API

The `GET /org/usage` API now has a `strictConsistency` parameter. To enforce a strong-global consistency when retrieving storage usage information across Storage Nodes, set this parameter to `true`. When this parameter is set to `false` (default), StorageGRID attempts to retrieve usage information using strong-global consistency, but falls back to strong-site consistency if strong-global consistency cannot be met.

Related information

[Use S3](#)

[Use a tenant account](#)

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](#)
2. [How your system is affected during the upgrade](#)
3. [Impact of an upgrade on groups and user accounts](#)
4. [Verifying the installed version of StorageGRID](#)
5. [Obtaining the required materials for a software upgrade](#)
6. [Downloading the StorageGRID upgrade files](#)
7. [Downloading the Recovery Package](#)
8. [Checking the system's condition before upgrading software](#)

Estimating the time to complete an upgrade

When planning an upgrade to StorageGRID 11.5, you must consider when to upgrade, based on how long the upgrade might take. You must also be aware of which operations you can and cannot perform during each stage of the upgrade.

About this task

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

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



During the upgrade from StorageGRID 11.4 to 11.5, Cassandra database tables on Storage Nodes will be upgraded. The **Upgrade Database** task occurs in the background, but might require an extensive amount of time to complete. While the database is being upgraded, you can safely use new features, apply hotfixes, and perform node recovery operations. However, you might be prevented from performing other maintenance procedures.



If an expansion is urgently required, perform the expansion before upgrading to 11.5.

Upgrade task	Description	Approximate time required	During this task
Start Upgrade Service	Upgrade prechecks are run, the software file is distributed, and the upgrade service is started.	3 minutes per grid node, unless validation errors are reported	As required, you can run the upgrade prechecks manually before the scheduled upgrade maintenance window.
Upgrade Grid Nodes (primary Admin Node)	The primary Admin Node is stopped, upgraded, and restarted.	Up to 30 minutes	You cannot access the primary Admin Node. Connection errors are reported, which you can ignore.
Upgrade Grid Nodes (all other nodes)	The software on all other grid nodes is upgraded, in the order in which you approve the nodes. Every node in your system will be brought down one at a time for several minutes each.	15 to 45 minutes per node, with appliance Storage 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"> • Do not change the grid configuration. • Do not change the audit level configuration. • Do not update the ILM configuration. • Do not perform another maintenance procedure, such as hotfix, decommission, or expansion. <p>Note: If you need to perform a recovery procedure, contact technical support.</p>
Enable Features	The new features for the new version are enabled.	Less than 5 minutes	<ul style="list-style-type: none"> • Do not change the grid configuration. • Do not change the audit level configuration. • Do not update the ILM configuration. • Do not perform another maintenance procedure.

Upgrade task	Description	Approximate time required	During this task
Upgrade Database	Cassandra database tables, which exist on all Storage Nodes, are upgraded.	Hours or days, based on the amount of metadata in your system	<p>During the Upgrade Database task, the upgraded grid will operate normally; however, the upgrade will still be in progress. During this task, you can:</p> <ul style="list-style-type: none"> • Use the new features in the new StorageGRID version. • Change the audit level configuration. • Update the ILM configuration. • Apply a hotfix. • Recover a node. <p>Note: You cannot perform a decommission or expansion procedure until the Final Upgrade Steps 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 (consider all upgrade tasks except for **Upgrade Database**).
 - a. Multiply the number of nodes in your StorageGRID system by 30 minutes/node (average).
 - 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. Estimate the time required to upgrade the database.
 - a. From the Grid Manager, select **Nodes**.
 - b. Select the first entry in the tree (entire grid), and select the **Storage** tab.
 - c. Hover your cursor over the **Storage Used - Object Metadata** chart, and locate the **Used** value, which indicates how many bytes of object metadata are on your grid.
 - d. Divide the **Used** value by 1.5 TB/day to determine how many days will be needed to upgrade the database.

4. Calculate the total estimated time for the upgrade by adding the results of steps 1, 2, and 3.

Example: Estimating the time to upgrade from StorageGRID 11.4 to 11.5

Suppose your system has 14 grid nodes, of which 8 are Linux nodes. Also, assume that the **Used** value for object metadata is 6 TB.

1. Multiply 14 by 30 minutes/node and add 1 hour. The estimated time to upgrade all nodes is 8 hours.
2. Multiple 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.
3. Divide 6 by 1.5 TB/day. The estimated number of days for the **Upgrade Database** task is 4 days.



While the **Upgrade Database** task is running, you can safely use new features, apply hotfixes, and perform node recovery operations.

4. Add the values together. You should allow 5 days to complete the upgrade of your system to StorageGRID 11.5.0.

How your system is affected during the upgrade

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

StorageGRID upgrades are non-disruptive

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

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

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

Appliance firmware is upgraded

During the StorageGRID 11.5 upgrade:

- All StorageGRID appliance nodes are automatically upgraded to StorageGRID Appliance Installer firmware version 3.5.
- SG6060 and SGF6024 appliances are automatically upgraded to BIOS firmware version 3B03.EX and BMC firmware version BMC 3.90.07.
- SG100 and SG1000 appliances are automatically upgraded to BIOS firmware version 3B08.EC and BMC firmware version 4.64.07.

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). 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 11.5 but other nodes are still running StorageGRID 11.4.

In general, these alerts will clear when the upgrade completes.

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



During the upgrade to StorageGRID 11.5, the **ILM placement unachievable** alert might be triggered when Storage Nodes are stopped. This alert might persist for 1 day after the upgrade is completed successfully.

Many SNMP notifications are generated

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

Configuration changes are restricted

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

- Do not make any grid configuration changes.
- Do not change the audit level configuration.
- Do not enable or disable any new features.
- Do not update the ILM configuration. Otherwise, you might experience inconsistent and unexpected ILM behavior.
- Do not apply a hotfix or recover a grid node.

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

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

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 11.5, optionally select the following new permissions and options (**Configuration > Access Control > Admin Groups**).

Permission or option	Description
Storage Appliance Administrator	Required to access the SANtricity System Manager user interface from Grid Manager.
Access Mode	When managing groups, you can select Read-only for this new option to prevent users from changing the settings and features that are selected for the group. Users in groups with read-only access mode can view settings, but they cannot change them.

Related information

[Administer StorageGRID](#)

Verifying 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.

Steps

1. Sign in to the Grid Manager using a supported browser.
2. Select **Help > About**.
3. Verify that the **Version** is 11.4.x.y.

In the StorageGRID 11.4.x.y version number:

- The major release has an x value of 0 (11.4.0).
- A minor release, if available, has an x value other than 0 (for example, 11.4.1).
- A hotfix, if available, has a y value (for example, 11.4.0.1).



If you have an earlier version of StorageGRID, you must upgrade to any 11.4 version before upgrading to StorageGRID 11.5. You do not need to be at the highest 11.4 minor-version release to upgrade to StorageGRID 11.5.

4. If you are not at a StorageGRID 11.4 version you must upgrade to version 11.4, one release at a time, using the instructions for each release.

You must also apply the latest hotfix for each StorageGRID version before upgrading to the next level.

One possible upgrade path is shown in the example.

5. Once you are at StorageGRID 11.4, go to the NetApp Downloads page for StorageGRID and see if any hotfixes are available for your StorageGRID 11.4.x version.

[NetApp Downloads: StorageGRID](#)

6. Verify that your StorageGRID 11.4.x version has the latest hotfix applied.
7. If necessary, download and apply the latest StorageGRID 11.4.x.y hotfix for your StorageGRID 11.4.x version.

See the recovery and maintenance instructions for information about applying hotfixes.

Example: Preparing to upgrade to StorageGRID 11.5 from version 11.3.0.8

The following example shows the upgrade steps to prepare for an upgrade from StorageGRID version 11.3.0.8 to version 11.5. Before you can upgrade to StorageGRID 11.5, your system must have a StorageGRID 11.4 version installed with the latest hotfix.

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. (You do not need to install any 11.4.x minor releases.)
3. Apply the latest StorageGRID 11.4.0.y hotfix.

Related information

[Administer StorageGRID](#)

[Maintain & recover](#)

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	You must download the required files to your service laptop: <ul style="list-style-type: none">• All platforms: .upgrade file• Any node on Red Hat Enterprise Linux or CentOS: .upgrade file and RPM file (.zip or .tgz)• Any node on Ubuntu or Debian: .upgrade file and DEB file (.zip or .tgz)
Service laptop	The service laptop must have: <ul style="list-style-type: none">• Network port• SSH client (for example, PuTTY)
Supported web browser	You must confirm that the web browser on the service laptop is supported for use with StorageGRID 11.5. Web browser requirements Note: Browser support has changed for StorageGRID 11.5. Confirm you are using a supported version.

Item	Notes
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>Downloading the Recovery Package</p>
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"> • Release notes for StorageGRID 11.5. Be sure to read these carefully before starting the upgrade. • Instructions for administering StorageGRID • If you are upgrading a Linux deployment, the StorageGRID installation instructions for your Linux platform. • Other StorageGRID documentation, as required.

Related information

[Web browser requirements](#)

[Administer StorageGRID](#)

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

[Install Ubuntu or Debian](#)

[Install VMware](#)

[Downloading the StorageGRID upgrade files](#)

[Downloading the Recovery Package](#)

[Release notes](#)

Web browser requirements

You must use a supported web browser.

Web browser	Minimum supported version
Google Chrome	87

Web browser	Minimum supported version
Microsoft Edge	87
Mozilla Firefox	84

You should set the browser window to a recommended width.

Browser width	Pixels
Minimum	1024
Optimum	1280

Downloading the StorageGRID upgrade files

You must download the required files to a service laptop before you upgrade your StorageGRID system.

What you'll need

You must have installed all required hotfixes for the StorageGRID software version you are upgrading. See the hotfix procedure in the recovery and maintenance instructions.

About this task

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

Steps

1. Go to the NetApp Downloads page for StorageGRID.

[NetApp Downloads: StorageGRID](#)

2. Select the button for downloading the latest release, or select another version from the drop-down menu and select **Go**.

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

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

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

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

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

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

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

Every platform requires the `.upgrade` archive.

7. If any nodes are deployed on Linux hosts, also download the RPM or DEB archive in either `.tgz` or `.zip` format.

You must install the RPM or DEB archive on all Linux nodes before you start the upgrade.



No additional files are required for the SG100 or SG1000.



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`

Related information

[Linux: Installing the RPM or DEB package on all hosts](#)

[Maintain & recover](#)

Downloading the Recovery Package

The Recovery Package file allows you to restore the StorageGRID system if a failure occurs.

What you'll need

- 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.

About this task

Download the current Recovery Package file before making grid topology changes to the StorageGRID system or before upgrading software. Then, download a new copy of the Recovery Package after making grid topology changes or after upgrading software.

Steps

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

The download starts immediately.

3. When the download completes:
 - a. Open the `.zip` file.

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



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

Related information

[Administer 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 alerts.

For information about specific alerts, see the monitoring and troubleshooting instructions.

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 the lists of internal and external ports in the 11.5 version of the networking guidelines, and ensure that all required ports are opened before you upgrade.



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

Related information

[Monitor & troubleshoot](#)

[Administer StorageGRID](#)

[Maintain & recover](#)

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.

What you'll need

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.
- 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, new features are enabled and you can resume operations; however, you must wait to perform a decommission or expansion procedure until the background **Upgrade Database** task and the **Final Upgrade Steps** task have completed.
- You must complete the upgrade on the same hypervisor platform you started with.

Steps

1. [Linux: Installing the RPM or DEB package on all hosts](#)
2. [Starting the upgrade](#)
3. [Upgrading grid nodes and completing the upgrade](#)
4. [Increasing the Metadata Reserved Space setting](#)

Related information

[Administer StorageGRID](#)

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

Linux: Installing the RPM or DEB package on all hosts

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

What you'll need

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

StorageGRID.



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

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

Steps

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

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

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

[Install Ubuntu or Debian](#)

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

Starting the upgrade

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

What you'll need

You have reviewed all of the considerations and completed all steps in [Upgrade planning and preparation](#).

Steps

1. Sign in to the Grid Manager using a supported browser.
2. Select **Maintenance > System > Software Update**.

The Software Update page appears.

3. Select **StorageGRID Upgrade**.

The StorageGRID Upgrade page appears and shows the date and time of the most recently completed upgrade, unless the primary Admin Node has been rebooted or the management API restarted since that upgrade was performed.

4. Select the `.upgrade` file you downloaded.
 - a. Select **Browse**.

b. Locate and select the file: `NetApp_StorageGRID_version_Software_uniqueID.upgrade`

c. Select **Open**.

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

5. Enter the provisioning passphrase in the text box.

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

StorageGRID Upgrade

Before starting the upgrade process, you must confirm that there are no active alerts and that all grid nodes are online and available.

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

Upgrade file

Upgrade file ✔ NetApp_StorageGRID_11.5.0_Software_20210407.2135.8e126f1

Upgrade Version StorageGRID® 11.5.0

Passphrase

Provisioning Passphrase

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



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



The same prechecks are performed when you select **Start Upgrade**. Selecting **Run Prechecks** allows you to detect and resolve issues before starting the upgrade.

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

A warning appears to remind you that your browser's connection will be lost when the primary Admin Node is rebooted. When the primary Admin Node is available again, you 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 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 might not respond as expected.

Are you sure you want to start the upgrade process?

Cancel

OK

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

When the upgrade starts:

- a. The upgrade prechecks are run.



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



While the primary Admin Node is being upgraded, multiple copies of the following error messages appear, which you can ignore.

Error

Problem connecting to the server

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

2 additional copies of this message are not shown.

OK

! Error

503: Service Unavailable

Service Unavailable

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

4 additional copies of this message are not shown.

OK

! Error

400: Bad Request

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

2 additional copies of this message are not shown.

OK

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.



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

Related information

[Upgrade planning and preparation](#)

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 complete, the **Upgrade Grid Nodes** task starts.
 - c. While the **Upgrade Grid Nodes** task is in progress, the Grid Node Status table appears and shows the upgrade stage for each grid node in your system.
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.



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 in sections by type: Admin Nodes, API Gateway Nodes, Storage Nodes, and Archive Nodes.

Upgrade Progress

Start Upgrade Service	Completed
Upgrade Grid Nodes	In Progress

Grid Node Status

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

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

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

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

Approve All

Remove All

Admin Nodes

Search

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

◀ ▶

Storage Nodes

Approve All **Remove All**

Search

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

◀ ▶

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

- Done (primary Admin Node only)
- Preparing upgrade

- Software download queued
- Downloading
- Waiting for you to approve

4. Approve the grid nodes you are ready to add to the upgrade queue. Approved nodes of the same type are upgraded one at a time.

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. 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.

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

5. If you need to remove a node or all nodes from the upgrade queue, select **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.

Storage Nodes		Approve All		Remove All	
Site	Name	Progress	Stage	Error	Action
Data Center 1	DC1-S1	<div style="width: 50%; background-color: #007bff;"></div>	Stopping services		
Data Center 1	DC1-S2	<div style="width: 25%; background-color: #007bff;"></div>	Queued		Remove
Data Center 1	DC1-S3	<div style="width: 25%; background-color: #007bff;"></div>	Queued		Remove

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



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

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

- As soon as the **Enable Features** task is complete (which occurs quickly), you can start using the new features in the upgraded StorageGRID version.

For example, if you are upgrading to StorageGRID 11.5, you can now enable S3 Object Lock, configure a key management server, or increase the Metadata Reserved Space setting.

Increasing the Metadata Reserved Space setting

- Periodically monitor the progress of the **Upgrade Database** task.


During this task, the Cassandra database is upgraded on each Storage Node.



The **Upgrade Database** task might take days to complete. As this background task runs, you can apply hotfixes or recover nodes. However, you must wait for the **Final Upgrade Steps** task to complete before performing an expansion or decommission procedure.

You can review the graph to monitor the progress for each Storage Node.

Upgrade Progress



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

complete.

StorageGRID Upgrade

The new features are enabled and can now be used. While the upgrade background tasks are in progress (which might take an extended time), you can apply hotfixes or recover nodes. You must wait for the upgrade to complete before performing an expansion or decommission.

Status	In Progress
Upgrade Version	11.5.0
Start Time	2021-04-08 09:01:48 MDT

Upgrade Progress

Start Upgrade Service	Completed
Upgrade Grid Nodes	Completed
Enable Features	Completed
Upgrade Database	Completed
Final Upgrade Steps	In Progress 

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

10. Confirm that the upgrade completed successfully.
 - a. Sign in to the Grid Manager using a supported browser.
 - b. Select **Help > About**.
 - c. Confirm that the displayed version is what you would expect.
 - d. Select **Maintenance > System > Software Update**. Then, select **StorageGRID Upgrade**.
 - e. Confirm that the green banner shows that the software upgrade was completed on the date and time you expected.

StorageGRID Upgrade

Before starting the upgrade process, you must confirm that there are no active alerts and that all grid nodes are online and available.

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

Software upgrade completed at 2021-04-08 12:14:40 MDT.

Upgrade file

Upgrade file

Upgrade Version No software upgrade file selected

Passphrase

Provisioning Passphrase

11. 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.
12. Check the NetApp Downloads page for StorageGRID to see if any hotfixes are available for the StorageGRID version that you just installed.

[NetApp Downloads: StorageGRID](#)

In the StorageGRID 11.5.x.y version number:

- The major release has an x value of 0 (11.5.0).
 - A minor release, if available, has an x value other than 0 (for example, 11.5.1).
 - A hotfix, if available, has a y value (for example, 11.5.0.1).
13. If available, download and apply the latest hotfix for your StorageGRID version.

See the recovery and maintenance instructions for information about applying hotfixes.

Related information

[Downloading the Recovery Package](#)

[Maintain & recover](#)

Increasing the Metadata Reserved Space setting

After you upgrade to StorageGRID 11.5, you might be able to increase the Metadata Reserved Space system setting if your Storage Nodes meet specific requirements for RAM and available space.

What you'll need

- You must be signed in to the Grid Manager using a supported browser.
- You must have the Root Access permission or the Grid Topology Page Configuration and Other Grid Configuration permissions.
- You have started the StorageGRID 11.5 upgrade and the **Enable New Features** upgrade task has completed.

About this task

You might be able to manually increase the system-wide Metadata Reserved Space setting up to 8 TB after upgrading to StorageGRID 11.5. Reserving additional metadata space after the 11.5 upgrade will simplify future hardware and software upgrades.

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

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

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



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

For a detailed explanation of how the Metadata Reserved Space setting affects the allowed space for object metadata storage on a particular Storage Node, go to the instructions for administering StorageGRID and search for “managing object metadata storage.”

Administer StorageGRID

Steps

1. Sign in to the Grid Manager using a supported browser.
2. Determine the current Metadata Reserved Space setting.
 - a. Select **Configuration > System Settings > Storage Options**.
 - b. In the Storage Watermarks section, note the value of **Metadata Reserved Space**.
3. Ensure you have enough available space on storage volume 0 of each Storage Node to increase this value.
 - a. Select **Nodes**.
 - b. Select the first Storage Node in the grid.
 - c. Select the Storage tab.
 - d. In the Volumes section, locate the **/var/local/rangedb/0** entry.
 - e. Confirm that the Available value is equal to or greater than difference between the new value you want to use and the current Metadata Reserved Space value.

For example, if the Metadata Reserved Space setting is currently 4 TB and you want to increase it to 6 TB, the Available value must be 2 TB or greater.

- f. Repeat these steps for all Storage Nodes.
 - If one or more Storage Nodes do not have enough available space, the Metadata Reserved Space value cannot be increased. Do not continue with this procedure.
 - If each Storage Node has enough available space on volume 0, go to the next step.
4. Ensure you have at least 128 GB of RAM on each Storage Node.
 - a. Select **Nodes**.
 - b. Select the first Storage Node in the grid.
 - c. Select the **Hardware** tab.
 - d. Hover your cursor over the Memory Usage chart. Ensure that **Total Memory** is at least 128 GB.
 - e. Repeat these steps for all Storage Nodes.
 - If one or more Storage Nodes do not have enough available total memory, the Metadata Reserved Space value cannot be increased. Do not continue with this procedure.
 - If each Storage Node has at least 128 GB of total memory, go to the next step.
5. Update the Metadata Reserved Space setting.
 - a. Select **Configuration > System Settings > Storage Options**.
 - b. Select the Configuration tab.
 - c. In the Storage Watermarks section, select **Metadata Reserved Space**.
 - d. Enter the new value.

For example, to enter 8 TB, which is the maximum supported value, enter **800000000000** (8, followed by 12 zeros)

Storage Options

- Overview
- Configuration

Configure Storage Options
Updated: 2021-02-17 19:40:49 MST

Object Segmentation

Description	Settings
Segmentation	Enabled
Maximum Segment Size	1000000000

Storage Watermarks

Description	Settings
Storage Volume Read-Write Watermark	30000000000
Storage Volume Soft Read-Only Watermark	10000000000
Storage Volume Hard Read-Only Watermark	5000000000
Metadata Reserved Space	800000000000

Apply Changes

- e. Select **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 information

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

Troubleshooting user interface issues

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

Web interface does not respond as expected

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

If you experience issues with the web interface:

- Make sure you are using a supported browser.



Browser support has changed for StorageGRID 11.5. 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 information

[Web browser requirements](#)

“Docker image availability check” error messages

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

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

Message	Cause	Solution
Unable to determine upgrade version. Upgrade version info file {file_path} did not match the expected format.	The upgrade package is corrupt.	Re-upload the upgrade package, and try again. If the problem persists, contact technical support.
Upgrade version info file {file_path} was not found. Unable to determine upgrade version.	The upgrade package is corrupt.	Re-upload the upgrade package, and try again. If the problem persists, contact technical support.
Unable to determine currently installed release version on {node_name}.	A critical file on the node is corrupt.	Contact technical support.
Connection error while attempting to list versions on {node_name}	The node is offline or the connection was interrupted.	Check to make sure that all nodes are online and reachable from the primary Admin Node, and try again.
The host for node {node_name} does not have StorageGRID {upgrade_version} image loaded. Images and services must be installed on the host before the upgrade can proceed.	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. Note: 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 {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.

Related information

Install Red Hat Enterprise Linux or CentOS

Install Ubuntu or Debian

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.