



Product feature guides

StorageGRID solutions and resources

NetApp
December 12, 2025

Table of Contents

Product feature guides	1
Achieving zero RPO with StorageGRID - A Comprehensive Guide to Multi-Site Replication	1
StorageGRID Overview	1
Requirements for Zero RPO with StorageGRID	6
Synchronous Deployments across multiple sites	6
A Single Grid Multi-site deployment	6
A multi-site multi-grid deployment	10
Conclusion	12
Create Cloud Storage Pool for AWS or Google Cloud	12
Create Cloud Storage Pool for Azure Blob Storage	13
Use a Cloud Storage Pool for backup	14
Configure StorageGRID search integration service	15
Introduction	15
Create tenant and enable platform services	15
Search integration services with Amazon OpenSearch	16
Platform services endpoint configuration	20
Search integration services with on premises Elasticsearch	22
Platform services endpoint configuration	25
Bucket search integration service configuration	27
Where to find additional information	31
Node Clone	31
Node clone considerations	31
Node clone Performance estimates	31
Grid site relocation and site-wide network change procedure	34
Considerations before site relocation	34
Migrating object-based storage from ONTAP S3 to StorageGRID	38
Enabling enterprise-grade S3 by seamlessly migrating object-based storage from ONTAP S3 to StorageGRID	38
Enabling enterprise-grade S3 by seamlessly migrating object-based storage from ONTAP S3 to StorageGRID	38
Enabling enterprise-grade S3 by seamlessly migrating object-based storage from ONTAP S3 to StorageGRID	50
Enabling enterprise-grade S3 by seamlessly migrating object-based storage from ONTAP S3 to StorageGRID	62
Enabling enterprise-grade S3 by seamlessly migrating object-based storage from ONTAP S3 to StorageGRID	71

Product feature guides

Achieving zero RPO with StorageGRID - A Comprehensive Guide to Multi-Site Replication

This technical report provides a comprehensive guide to implementing StorageGRID replication strategies to achieve a Recovery Point Objective (RPO) of zero in the event of a site failure. The document details various deployment options for StorageGRID, including multi-site synchronous replication and multi-grid asynchronous replication. It explains how StorageGRID Information Lifecycle Management (ILM) policies can be configured to ensure data durability and availability across multiple locations. Additionally, the report covers performance considerations, failure scenarios, and recovery processes to maintain uninterrupted client operations. The goal of this document is to provide information to ensure that data remains accessible and consistent, even in the event of a complete site failure, by leveraging both synchronous and asynchronous replication techniques.

StorageGRID Overview

NetApp StorageGRID is an object-based storage system that supports the industry-standard Amazon Simple Storage Service (Amazon S3) API.

StorageGRID provides a single namespace across multiple locations with variable levels of service driven by information lifecycle management policies (ILM). With these lifecycle policies you can optimize where your data lives throughout its lifecycle.

StorageGRID allows for configurable durability and availability of your data in local and geo-distributed solutions. Whether your data is on premises or in a public cloud, integrated hybrid cloud workflows allow your business to leverage cloud services like Amazon Simple Notification Service (Amazon SNS), Google Cloud, Microsoft Azure Blob, Amazon S3 Glacier, Elasticsearch, and more.

StorageGRID scale

A minimal StorageGRID deployment consists of an Admin node and 3 Storage nodes in a single site. A single grid can grow up to 220 nodes.

StorageGRID can be deployed as a single site or extended to 16 sites.

The Admin node contains the management interface, a central point for metrics and logging, and maintains the configuration of the StorageGRID components. The Admin node also contains an integrated load balancer for S3 API access.


StorageGRID can be deployed as software-only, as VMware virtual machine appliances, or as purpose-built appliances.

A Storage node can be deployed as:

- A metadata only node maximizing object count
- An object storage only node maximizing object space
- A combined metadata and object storage node adding both object count and object space

Each Storage node can scale to multi-petabyte capacity for object storage allowing for a single namespace in the hundreds of petabytes. StorageGRID also provides an integrated load balancer for S3 API operations called a gateway node.

Delivery paths for any workload



StorageGRID consists of a collection of nodes placed into a site topology. A site in StorageGRID can be a unique physical location or reside in a shared physical location as other sites in the grid as a logical construct. A StorageGRID site should not span multiple physical locations. A site represents a shared local area network (LAN) infrastructure and failure domain.

StorageGRID and failure domains

StorageGRID contains multiple layers of failure domains to be considered in deciding how to architect your solution, how to store your data and where your data should be stored to mitigate the risks of failures.

- Grid level - A grid consisting of multiple sites can have site failures or isolation and the accessible site(s) can continue operating as the grid.
- Site level - Failures within a site may impact operations of that site but will not impact the rest of the grid.
- Node level - A node failure will not impact the operation of the site.
- Disk level - a disk failure will not impact operation of the node.

Object data and metadata

With object storage, the unit of storage is an object, rather than a file or a block. Unlike the tree-like hierarchy of a file system or block storage, object storage organizes data in a flat, unstructured layout. Object storage decouples the physical location of the data from the method used to store and retrieve that data.

Each object in an object-based storage system has two parts: object data and object metadata.

- Object data represents the actual underlying data, for example, a photograph, a movie, or a medical record.

- Object metadata is any information that describes an object.

StorageGRID uses object metadata to track the locations of all objects across the grid and to manage each object's lifecycle over time.

Object metadata includes information such as the following:

- System metadata, including a unique ID for each object, the object name, the name of the S3 bucket, the tenant account name or ID, the logical size of the object, the date and time the object was first created, and the date and time the object was last modified.
- The current storage location of each object's replicate copy or erasure-coded fragment.
- Any custom user metadata key-value pairs associated with the object.
- For S3 objects, any object tag key-value pairs associated with the object
- For segmented objects and multipart objects, segment identifiers, and data sizes.

Object metadata is customizable and expandable, making it flexible for applications to use. For detailed information about how and where StorageGRID stores object metadata, go to [Manage object metadata storage](#).

StorageGRID's Information lifecycle management (ILM) system is used to orchestrate the placement, duration, and ingest behavior for all object data in your StorageGRID system. ILM rules determine how StorageGRID stores objects over time using replicas of the objects or erasure coding the object across nodes and sites. This ILM system is responsible for the object data consistency within a grid.

Erasure coding

StorageGRID provides the ability to erasure code data at node level and the drive level. With StorageGRID appliances we erasure code the data stored on each node across all the drives within the node providing local protection against multiple disk failures causing data loss or interruptions. Rebuilds from drive failures are local to the node and do not require data replicated over the network.

Additionally, StorageGRID appliances use erasure coding schemes to store object data across the nodes within a site or spread across 3 or more sites in the StorageGRID system though StorageGRID's ILM rules protecting against node failure.

Erasure coding provides a storage layout that is resilient to node and site failures with a lower overhead than replication. All StorageGRID erasure coding schemes are deployable in a single site provided the minimum number of nodes required to store the data chunks are met. This means for an EC scheme of 4+2 there needs to be a minimum of 6 nodes available to receive the data.

Erasure-coding scheme ($k+m$)	Minimum number of deployed sites	Recommended number of Storage Nodes at each site	Total recommended number of Storage Nodes	Site loss protection?	Storage overhead
4+2	3	3	9	Yes	50%
6+2	4	3	12	Yes	33%
8+2	5	3	15	Yes	25%
6+3	3	4	12	Yes	50%
9+3	4	4	16	Yes	33%
2+1	3	3	9	Yes	50%
4+1	5	3	15	Yes	25%
6+1	7	3	21	Yes	17%
7+5	3	5	15	Yes	71%

Metadata consistency

In StorageGRID, metadata is typically stored with three replicas per site to ensure consistency and availability. This redundancy helps maintain data integrity and accessibility even in the event of a failure.

The default consistency is defined at a grid wide level. Users can change the consistency at the bucket level at any time.

The bucket consistency options available in StorageGRID are:

- **All:** Provides the highest level of consistency. All nodes in the grid receive the data immediately, or the request will fail.
- **Strong-global:**
 - **Legacy Strong Global:** Guarantees read-after-write consistency for all client requests across all sites.
 - This is the default behavior for all systems upgraded from 11.9 or older to 12.0 without manually changing to the new Quorum Strong Global.
 - **Quorum Strong-global:** Guarantees read-after-write consistency for all client requests across all sites. Offers consistency for multiple nodes or even a site failure if metadata replica quorum is achievable.
 - This is the default behavior for all systems newly installed at 12.0 or higher.
 - QUORUM consistency is defined as a quorum of Storage Node metadata replicas, where each site has 3 metadata replicas. It may be calculated as follows: $1+((N*3)/2)$ where N is the total number of sites
 - For example, a minimum of 5 replicas must be made from a 3-site grid with a maximum of 3 replicas within a site.
- **Strong-site:** Guarantees read-after-write consistency for all client requests within a site.
- **Read-after-new-write(default):** Provides read-after-write consistency for new objects and eventual

consistency for object updates. Offers high availability and data protection guarantees. Recommended for most cases.

- **Available:** Provides eventual consistency for both new objects and object updates. For S3 buckets, use only as required (for example, for a bucket that contains log values that are rarely read, or for HEAD or GET operations on keys that don't exist). Not supported for S3 FabricPool buckets.

Object data consistency

While metadata is automatically replicated within and across sites, object data storage placement decisions are up to you. Object data can be stored in replicas within and across sites, erasure coded within or across sites, or a combination of replicas and erasure coded storage schemes. ILM rules can apply to all objects, or be filtered to only apply to certain objects, buckets, or tenants. ILM rules define how objects are stored, replicas and/or erasure coded, how long objects are stored in those locations, if the number of replicas or erasure coding scheme should change, or locations should change over time.

Each ILM rule will be configured with one of three ingest behaviors for protecting objects: Dual commit, balanced or strict.

The dual commit option will make two copies on any two different storage nodes in the grid immediately and return the request to be successful to the client. The node selection will be tried within the site of the request but may use nodes of another site in some circumstances. The object is added to the ILM queue to be evaluated and placed according to the ILM rules.

The balanced option evaluates the object against the ILM policy immediately and places the object synchronously before returning the request to be successful to the client. If the ILM rule cannot be met immediately due to an outage or inadequate storage to meet the placement requirements, then dual commit will be used instead. Once the issue is resolved, ILM will automatically place the object based on the defined rule.

The strict option evaluates the object against the ILM policy immediately and places the object synchronously before returning the request to be successful to the client. If the ILM rule cannot be met immediately due to an outage or inadequate storage to meet the placement requirements, then the request will fail, and the client will need to retry.

Load balancing

StorageGRID can be deployed with client access through the integrated gateway nodes, an external 3rd party load balancer, DNS round robin, or directly to a storage node. Multiple gateway nodes can be deployed in a site and configured in high availability groups providing automated failover and fail back in the event of a gateway node outage. You can combine load balancing methods in a solution to provide a single point of access for all sites in a solution.

The gateway nodes will balance the load between the storage nodes in the site where the gateway node resides by default. StorageGRID can be configured to allow the gateway nodes to balance load using nodes from multiple sites. This configuration would add the latency between those sites to the response latency to the client's requests. This should only be configured if the total latency is acceptable to the clients.

Ensuring an RTO of zero can be achieved with a combination of local and global load balancing. Ensuring uninterrupted client access requires load balancing of client requests. A StorageGRID solution can contain many gateway nodes and high availability groups in each site. To provide uninterrupted access for clients in any site even in a site failure, you should configure an external load balancing solution in combination with StorageGRID Gateway nodes. Configure Gateway node high availability groups that manage the load within each site and use the external load balancer to balance the load across the high availability groups. The external load balancer must be configured to perform a health check to ensure requests are sent only to

operational sites. For more information on load balancing with StorageGRID please see the [StorageGRID load balancer technical report](#).

Requirements for Zero RPO with StorageGRID

To achieve zero Recovery Point Objective (RPO) in an object storage system, it is crucial that at the time of failure:

- Both metadata and object contents are in sync and considered consistent
- Object content remain accessible despite the failure.

For a multi-site deployment, Quorum Strong Global is the preferred consistency model to ensure metadata is synchronized across all sites, making it essential for meeting the zero RPO requirement.

Objects in the storage system are stored based on Information Lifecycle Management (ILM) rules, which dictate how and where data is stored throughout its lifecycle. For synchronous replication, one can consider between Strict execution or Balanced Execution.

- Strict execution of these ILM rules is necessary for zero RPO because it ensures that objects are placed in the defined locations without any delay or fallback, maintaining data availability and consistency.
- StorageGRID's ILM balance ingest behavior provides a balance between high availability and resiliency, allowing users to continue ingesting data even in the event of a site failure.

Synchronous Deployments across multiple sites

Multi-site solutions: StorageGRID allows you to replicate objects across multiple sites within the grid synchronously. By setting up Information Lifecycle Management (ILM) rules with balance or strict behavior, objects are placed immediately in the specified locations. Configuring bucket consistency level to Quorum Strong Global will ensure synchronous metadata replication as well. StorageGRID uses a single global namespace, storing object placement locations as metadata, so every node knows where all copies or erasure coded pieces are located. If an object can't be retrieved from the site where the request was made, it will be automatically retrieved from a remote site without needing failover procedures.

Once the failure is resolved, no manual failback efforts are required. The replication performance depends on the site with the lowest network throughput, highest latency, and lowest performance. A site's performance is based on the number of nodes, CPU core count and speed, memory, drive quantity, and drive types.

Multi-grid solutions: StorageGRID can replicate tenants, users, and buckets between multiple StorageGRID systems using Cross-Grid replication (CGR). CGR can extend select data to more than 16 sites, increase the usable capacity of your object store, and provide disaster recovery. The replication of buckets with CGR includes objects, object versions, and metadata, and can be bi-directional or one-way. The recovery point objective (RPO) depends on the performance of each StorageGRID system and the network connections between them.

Summary:

- Intra-grid replication includes both synchronous and asynchronous replication, configurable using ILM ingest behavior and metadata consistency control.
- Inter-grid replication is asynchronous only.


A Single Grid Multi-site deployment

In the following scenarios the StorageGRID solutions are configured with an optional external load balancer

managing requests to the integrated load balancer high availability groups. This will achieve an RTO of zero in addition to an RPO of zero. ILM is configured with Balanced ingest protection for synchronous placement. Each bucket is configured with the Quorum version of Strong Global consistency model for grids of 3 or more sites and the Legacy version of Strong Global consistency for 2 sites.


Scenario 1:

In a two site StorageGRID solution, there are at least two replicas of every object and 6 replicas of all metadata. Upon failure recovery, updates from the outage will synchronize to the recovered site/nodes automatically. With only 2 sites it is not likely to achieve a zero RPO in failure scenarios beyond a full site loss.



Scenario 2:

In a StorageGRID solution of three or more sites, there are at least 3 replicas or 3 EC chunks of every object and 9 replicas of all metadata. Upon failure recovery, updates from the outage will synchronize to the recovered site/nodes automatically. With three or more sites it is possible to achieve a zero RPO.



Multi-site failure scenarios

Failure	2-site Outcome Legacy Strong Global	3 or more sites outcome Quorum Strong Global
Single node drive failure	Each appliance uses multiple disk groups and can sustain at least 1 drive per group failing without interruption or data loss.	Each appliance uses multiple disk groups and can sustain at least 1 drive per group failing without interruption or data loss.
Single node failure in one site	No interruption to operations or data loss.	No interruption to operations or data loss.
Multiple node failure in one site	<p>Disruption to client operations directed to this site but no data loss.</p> <p>Operations directed to the other site remain uninterrupted and no data loss.</p>	Operations are directed to all other sites and remain uninterrupted and no data loss.

Failure	2-site Outcome Legacy Strong Global	3 or more sites outcome Quorum Strong Global
Single node failure at multiple sites	<p>No disruption or data loss if:</p> <ul style="list-style-type: none"> • At least one replicate copy exists in the grid • Sufficient EC chunks exist in the grid <p>Operations disrupted and risk of data loss if:</p> <ul style="list-style-type: none"> • No replicate copies exist • Insufficient EC chunks exist 	<p>No disruption or data loss if:</p> <ul style="list-style-type: none"> • At least a single replicate copy exists in the grid • Sufficient EC chunks exist in the grid <p>Operations disrupted and risk of data loss if:</p> <ul style="list-style-type: none"> • No replicate copies exist • Insufficient EC chunks exist to retrieve the object
Single site failure	<p>Some client operations will be interrupted until the failure is resolved.</p> <p>GET and HEAD operations will continue without interruption.</p> <p>Reduce bucket consistency to read-after-new-write or lower to continue operations uninterrupted in this failure state.</p>	No interruption to operations or data loss.
Single site plus single node failures	<p>Some client operations will be interrupted until either the failure is resolved.</p> <p>HEAD operations will continue without interruption.</p> <p>GET operations will continue without interruption if a replicate copy or sufficient EC chunks exist.</p> <p>Reduce bucket consistency to read-after-new-write or lower to continue operations uninterrupted in this failure state.</p>	<p>No interruption to operations or data loss.</p> <p>Possible data loss depending on the number of replicate copies.</p> <p>Local Erasure coding can prevent data loss.</p>


Failure	2-site Outcome Legacy Strong Global	3 or more sites outcome Quorum Strong Global
Single site plus a node from each remaining site	<p>Only two sites exist.</p> <p>See: Single site plus a single node.</p>	<p>Operations will be disrupted If metadata replica quorum cannot be met.</p> <p>Reduce bucket consistency to read-after-new-write or lower to continue operations uninterrupted in this failure state.</p> <p>Possible data loss for permanent failure depending on the number of replicate copies.</p> <p>Local Erasure coding can prevent data loss.</p>
Multi-site failure	<p>No operational sites remain.</p> <p>Data will be lost if at least one site cannot be recovered in its entirety.</p>	<p>Operations will be disrupted If metadata replica quorum cannot be met.</p> <p>Reduce bucket consistency to read-after-new-write or lower to continue operations uninterrupted in this failure state.</p> <p>Possible data loss for permanent failure if not enough erasure coded chunks remain.</p> <p>Local Erasure coding or replicate copies can prevent data loss.</p>
Network isolation of a site	<p>client operations will be interrupted until either the failure is resolved.</p> <p>Reduce bucket consistency to read-after-new-write or lower to continue operations uninterrupted in this failure state.</p> <p>No data loss</p>	<p>Operations will be disrupted for the isolated site, but no data loss.</p> <p>Reduce bucket consistency to read-after-new-write or lower to continue operations uninterrupted in this failure state.</p> <p>No disruption to operations in the remaining sites and no data loss.</p>

A multi-site multi-grid deployment

To add an extra layer of redundancy, this scenario will employ two StorageGRID Clusters and use cross-grid replication to keep them in sync. For this solution each StorageGRID clusters will have three sites. Two sites will be used for object storage and metadata while the third site will be used solely for metadata. Both systems will be configured with a balanced ILM rule to synchronously store the objects using erasure coding in each of the two data sites. Buckets will be configured with the Quorum Strong Global consistency model. Each grid will be configured with bi-directional cross-grid replication on every bucket. This provides the asynchronous replication between the regions. Optionally a global load balancer can be implemented to manage requests to

the integrated load balancer high availability groups of both StorageGRID systems to achieve a zero RPO.

The solution will use four locations equally divided into two regions. Region 1 will contain the 2 storage sites of grid 1 as the primary grid of the region and the metadata site of grid 2. Region 2 will contain the 2 storage sites of grid 2 as the primary grid of the region and the metadata site of grid 1. In each region the same location can house the storage site of the primary grid of the region as well as the metadata only site of the other regions grid. Using metadata only nodes as the third site will provide the consistency required for the metadata and not duplicate the storage of objects in that location.



This solution with four separate locations provides complete redundancy of two separate StorageGRID systems maintaining an RPO of 0 and will make use of both multi-site synchronous replication, and multi-grid asynchronous replication. Any single site can fail while maintaining uninterrupted client operations on both StorageGRID systems.

In this solution, there are four erasure coded copies of every object and 18 replicas of all metadata. This allows for multiple failure scenarios without client operations impact. Upon failure recovery updates from the outage will synchronize to the failed site/nodes automatically.

Multisite, multi-grid failure scenarios

Failure	Outcome
Single node drive failure	Each appliance uses multiple disk groups and can sustain at least 1 drive per group failing without interruption or data loss.
Single node failure in one site in a grid	No interruption to operations or data loss.
Single node failure in one site in each grid	No interruption to operations or data loss.
Multiple node failure in one site in a grid	No interruption to operations or data loss.

Failure	Outcome
Multiple node failure in one site in each grid	No interruption to operations or data loss.
Single node failure at multiple sites in a grid	No interruption to operations or data loss.
Single node failure at multiple sites in each grid	No interruption to operations or data loss.
Single site failure in a grid	No interruption to operations or data loss.
Single site failure in each grid	No interruption to operations or data loss.
Single site plus single node failures in a grid	No interruption to operations or data loss.
Single site plus a node from each remaining site in a single grid	No interruption to operations or data loss.
Single location failure	No interruption to operations or data loss.
Single location failure in each grid DC1 & DC3	Operations will be disrupted until either the failure is resolved, or the bucket consistency is lowered; each grid has lost 2 sites All data still exists at 2 locations
Single location failure in each grid DC1 & DC4 or DC2 & DC3	No interruption to operations or data loss.
Single location failure in each grid DC2 & DC4	No interruption to operations or data loss.
Network isolation of a site	Operations will be disrupted for the isolated site but no data will be lost No disruption to operations in the remaining sites or data loss.

Conclusion

Achieving zero Recovery Point Objective (RPO) with StorageGRID is a critical goal for ensuring data durability and availability in the event of site failures. By leveraging StorageGRID's robust replication strategies, including multi-site synchronous replication and multi-grid asynchronous replication, organizations can maintain uninterrupted client operations and ensure data consistency across multiple locations. The implementation of Information Lifecycle Management (ILM) policies and the use of metadata-only nodes further enhance the system's resilience and performance. With StorageGRID, businesses can confidently manage their data, knowing that it remains accessible and consistent even in the face of complex failure scenarios. This comprehensive approach to data management and replication underscores the importance of meticulous planning and execution in achieving zero RPO and safeguarding valuable information.

Create Cloud Storage Pool for AWS or Google Cloud

You can use a Cloud Storage Pool if you want to move StorageGRID objects to an external S3 bucket. The external bucket can belong to Amazon S3 (AWS) or Google Cloud.

What you'll need

- StorageGRID 11.6 has been configured.
- You have already set up an external S3 bucket on AWS or Google Cloud.

Steps

1. In the Grid Manager, navigate to **ILM > Storage Pools**.
2. In the Cloud Storage Pools section of the page, select **Create**.

The Create Cloud Storage Pool pop-up appears.

3. Enter a display name.
4. Select **Amazon S3** from the Provider Type drop-down list.

This provider type works for AWS S3 or Google Cloud.

5. Enter the URI for the S3 bucket to be used for the Cloud Storage Pool.

Two formats are allowed:

`https://host:port`

`http://host:port`

6. Enter the S3 bucket name.

The name you specify must exactly match the S3 bucket's name; otherwise, Cloud Storage Pool creation fails. You cannot change this value after the Cloud Storage Pool is saved.

7. Optionally, enter the Access Key ID and the Secret Access Key.
8. Select **Do Not Verify Certificate** from the drop-down.
9. Click **Save**.

Expected result

Confirm that a Cloud Storage Pool has been created for Amazon S3 or Google Cloud.

By Jonathan Wong

Create Cloud Storage Pool for Azure Blob Storage

You can use a Cloud Storage Pool if you want to move StorageGRID objects to an external Azure container.

What you'll need

- StorageGRID 11.6 has been configured.
- You have already set up an external Azure container.

Steps

1. In the Grid Manager, navigate to **ILM > Storage Pools**.
2. In the Cloud Storage Pools section of the page, select **Create**.

The Create Cloud Storage Pool pop-up appears.

3. Enter a display name.
4. Select **Azure Blob Storage** from the Provider Type drop-down list.
5. Enter the URI for the S3 bucket to be used for the Cloud Storage Pool.

Two formats are allowed:

`https://host:port`

`http://host:port`

6. Enter the Azure container name.

The name you specify must exactly match the Azure container name; otherwise, Cloud Storage Pool creation fails. You cannot change this value after the Cloud Storage Pool is saved.

7. Optionally, enter the Azure container's associated account name and account key for authentication.
8. Select **Do Not Verify Certificate** from the drop-down.
9. Click **Save**.

Expected result

Confirm that a Cloud Storage Pool has been created for Azure Blob Storage.

By Jonathan Wong

Use a Cloud Storage Pool for backup

You can create an ILM rule to move objects into a Cloud Storage Pool for backup..

What you'll need

- StorageGRID 11.6 has been configured.
- You have already set up an external Azure container.

Steps

1. In the Grid Manager, navigate to **ILM > Rules > Create**.
2. Enter a description.
3. Enter a criterion to trigger the rule.
4. Click **Next**.
5. Replicate the object to Storage Nodes.
6. Add a placement rule.
7. Replicate the object to the Cloud Storage Pool
8. Click **Next**.
9. Click **Save**.

Expected result

Confirm that the retention diagram shows the objects stored locally in StorageGRID and in a Cloud Storage

Pool for backup.

Confirm that, when the ILM rule is triggered, a copy exists in the Cloud Storage Pool and you can retrieve the object locally without doing an object restore.

By Jonathan Wong

Configure StorageGRID search integration service

This guide provides detailed instructions for configuring NetApp StorageGRID search integration service with either Amazon OpenSearch Service or on-premises Elasticsearch.

Introduction

StorageGRID supports three types of platform services.

- **StorageGRID CloudMirror replication.** Mirror specific objects from a StorageGRID bucket to a specified external destination.
- **Notifications.** Per-bucket event notifications to send notifications about specific actions performed on objects to a specified external Amazon Simple Notification Service (Amazon SNS).
- **Search integration service.** Send Simple Storage Service (S3) object metadata to a specified Elasticsearch index where you can search or analyze the metadata by using the external service.


Platform services are configured by the S3 tenant through the Tenant Manager UI. For more information, see [Considerations for using platform services](#).

This document serves as a supplement to the [StorageGRID 11.6 Tenant Guide](#) and provides step by step instructions and examples for the endpoint and bucket configuration for search integration services. The Amazon Web Services (AWS) or on-premises Elasticsearch setup instructions included here are for basic testing or demo purposes only.

Audiences should be familiar with Grid Manager, Tenant Manager, and have access to the S3 browser to perform basic upload (PUT) and download (GET) operations for StorageGRID search integration testing.

Create tenant and enable platform services

1. Create an S3 tenant by using Grid Manager, enter a display name, and select the S3 protocol.
2. On the Permission page, select the Allow Platform Services option. Optionally, select other permissions, if necessary.



3. Set up the tenant root user initial password or, if identity federation is enabled on the grid, select which federated group has root access permission to configure the tenant account.
4. Click Sign In As Root and select Bucket: Create and Manage Buckets.

This takes you to the Tenant Manager page.

5. From Tenant Manager, select My Access Keys to create and download the S3 access key for later testing.

Search integration services with Amazon OpenSearch

Amazon OpenSearch (formerly Elasticsearch) service setup

Use this procedure for a quick and simple setup of the OpenSearch service for testing/demo purposes only. If you are using on-premises Elasticsearch for search integration services, see the section [Search integration services with on premises Elasticsearch](#).



You must have a valid AWS console login, access key, secret access key, and permission to subscribe to the OpenSearch service.

1. Create a new domain using the instructions from [AWS OpenSearch Service Getting Started](#), except for the following:
 - Step 4. Domain name: sgdemo
 - Step 10. Fine-grained access control: deselect the Enable Fine-Grained Access Control option.
 - Step 12. Access policy: select Configure Level Access Policy, select the JSON tab to modify the access policy by using the following example:
 - Replace the highlighted text with your own AWS Identity and Access Management (IAM) ID and user name.
 - Replace the highlighted text (the IP address) with the public IP address of your local computer that you used to access the AWS console.
 - Open a browser tab to <https://checkip.amazonaws.com> to find your public IP.

```
{  
  
    "Version": "2012-10-17",  
    "Statement": [  
        {  
            "Effect": "Allow",  
            "Principal":  
                {"AWS": "arn:aws:iam:: nnnnnnn:user/xyzabc"},  
            "Action": "es:*",  
            "Resource": "arn:aws:es:us-east-1:nnnnnn:domain/sgdemo/*"  
        },  
        {  
            "Effect": "Allow",  
            "Principal": {"AWS": "*"},  
            "Action": [  
                "es:ESHttp*"  
            ],  
            "Condition": {  
                "IpAddress": {  
                    "aws:SourceIp": [ "nnn.nnn.nn.n/nn"  
                ]  
            }  
        },  
        "Resource": "arn:aws:es:us-east-1:nnnnnn:domain/sgdemo/*"  
    ]  
}
```

Fine-grained access control

Fine-grained access control provides numerous features to help you keep your data secure. Features include document-level security, field-level security, read-only users, and OpenSearch Dashboards/Kibana tenants. Fine-grained access control requires a master user. [Learn more](#)

Enable fine-grained access control

SAML authentication for OpenSearch Dashboards/Kibana

SAML authentication lets you use your existing identity provider for single sign-on for OpenSearch Dashboards/Kibana. [Learn more](#)

Prepare SAML authentication

 To use SAML authentication, you must first enable fine-grained access control.

Amazon Cognito authentication

Enable to use Amazon Cognito authentication for OpenSearch Dashboards/Kibana. Amazon Cognito supports a variety of identity providers for username-password authentication. [Learn more](#)

Enable Amazon Cognito authentication

Access policy

Access policies control whether a request is accepted or rejected when it reaches the Amazon OpenSearch Service domain. If you specify an account, user, or role in this policy, you must sign your requests. [Learn more](#)

Domain access policy

- Only use fine-grained access control
Allow open access to the domain.
- Do not set domain level access policy
All requests to the domain will be denied.
- Configure domain level access policy

Visual editor


JSON

Import policy

Access policy

```
3+ "Statement": [
4+   {
5+     "Effect": "Allow",
6+     "Principal": {
7+       "AWS": "arn:aws:iam::210111111111:user/sgdemo"
8+     },
9+     "Action": "es:*",
10+    "Resource": "arn:aws:es:us-east-1:210111111111:domain/sgdemo/*"
11+  },
12+  {
13+    "Effect": "Allow",
14+    "Principal": {
15+      "AWS": "*"
16+    },
17+    "Action": [
18+      "es:ESHttp*"
19+    ],
20+    "Condition": {
21+      "IpAddress": [
22+        "aws:SourceIp": [
23+          "216.210.111.24"
24+
25+        ]
26+      ],
27+      "Resource": "arn:aws:es:us-east-1:210111111111:domain/sgdemo/*"
28+    }
29+  }
30+ ]
```

2. Wait 15 to 20 minutes for the domain to become active.



Amazon OpenSearch Service > Domains > sgdemo

sgdemo [Info](#)

General information

Name	sgdemo	Domain status	Active	Version	Info
Domain ARN	arn:aws:es:us-east-1:123456789012:domain/sgdemo	Cluster health	Info	OpenSearch 1.1 (latest)	OpenSearch Dashboards URL
			Yellow	R20211203-PS (latest)	https://search-sgdemo-1.es.amazonaws.com/_dashboards
				Service software version	Domain endpoint
				Info	https://search-sgdemo-1.es.amazonaws.com

3. Click OpenSearch Dashboards URL to open the domain in a new tab to access the dashboard. If you get an access denied error, verify that the access policy source IP address is correctly set to your computer public IP to allow access to the domain dashboard.
4. On the dashboard welcome page, select Explore On Your Own. From the menu, go to Management → Dev Tools
5. Under Dev Tools → Console , enter `PUT <index>` where you use the index name 'sgmetadata' in the following example. Click the small triangle symbol to execute the PUT command. The expected result displays on the right panel as shown in the following example screenshot.



OpenSearch Dashboards

☰ Dev Tools

Console

History Settings Help

```
1 PUT sgmetadata ▶
```

```
1 {  
2   "acknowledged" : true,  
3   "shards_acknowledged" : true,  
4   "index" : "sgmetadata"  
5 }
```

6. Verify that the index is visible from Amazon OpenSearch UI under sgdomain > Indices.

Successfully updated to service software version 820211203-15

Amazon OpenSearch Service > Domains > sgdemo

sgdemo Info

General information

Name: sgdemo	Domain status: Active	Version info: OpenSearch 1.1 (latest)	OpenSearch Dashboards URL: https://search-sgdemo-1234567890.us-east-1.es.amazonaws.com/_dashboards
Domain ARN: arn:aws:es:us-east-1:211234567890:domain/sgdemo	Cluster health: Yellow	Service software version: 820211203-15 (latest)	Domain endpoint: https://search-sgdemo-1234567890.us-east-1.es.amazonaws.com

Indices [2]

Indices is the method by which search engines organize data by field names. Before you can search data, you must index it. Learn more [»](#)

Index	Document count	Size (byte)	Query total	Mapping type	Field mappings	Actions
arrows-es-test-1	1	5.08 kB	10	dynamic,meta,proper...		View
sgmedata	6	1.02 kB	0			View


Platform services endpoint configuration

To configure the platform services endpoints, follow these steps:

1. In Tenant Manager, go to STORAGE(S3) > Platform services endpoints.
2. Click Create Endpoint, enter the following, and then click Continue:
 - Display name **example aws-opensearch**
 - The domain endpoint in the **example** screenshot under Step 2 of the preceding procedure in the **URI field**.
 - The domain ARN used in Step 2 of the preceding procedure in the **URN field** and add `/<index>/_doc` to the end of ARN.

In this example, URN becomes `arn:aws:es:us-east-1:211234567890:domain/sgdemo/sgmedata/_doc`.

Create endpoint



3. To access the Amazon OpenSearch sgdomain, choose Access Key as the authentication type and then enter the Amazon S3 access key and secret key. To go the next page, click Continue.

Create endpoint

4. To verify the endpoint, select Use Operating System CA Certificate and Test and Create Endpoint. If verification is successful, an endpoint screen similar to the following figure displays. If verification fails, verify that the URN includes `<index>/_doc` at the end of the path and the AWS access key and secret key are correct.

Platform services endpoints

A platform services endpoint stores the information StorageGRID needs to use an external resource as a target for a platform service (CloudMirror replication, notifications, or search integration). You must configure an endpoint for each platform service you plan to use.

1 endpoint

[Create endpoint](#)

Delete endpoint	Display name	Last error	Type	URI	URN
<input type="checkbox"/>	aws-opensearch		Search	https://search-sgdemo-422715828337272.us-east-1.es.amazonaws.com/	arn:aws:es:us-east-1:21011100129:domain/sgdemo/sgmetadata/_doc

Search integration services with on premises Elasticsearch

On premises Elasticsearch setup

This procedure is for a quick setup of on premises Elasticsearch and Kibana using docker for testing purposes only. If the Elasticsearch and Kibana server already exists, go to Step 5.

1. Follow this [Docker installation procedure](#) to install docker. We use the [CentOS Docker install procedure](#) in this setup.

```
sudo yum install -y yum-utils
sudo yum-config-manager --add-repo
https://download.docker.com/linux/centos/docker-ce.repo
sudo yum install docker-ce docker-ce-cli containerd.io
sudo systemctl start docker
```

- To start docker after reboot, enter the following:

```
sudo systemctl enable docker
```

- Set the `vm.max_map_count` value to 262144:

```
sysctl -w vm.max_map_count=262144
```

- To keep the setting after reboot, enter the following:

```
echo 'vm.max_map_count=262144' >> /etc/sysctl.conf
```


2. Follow the [Elasticsearch Quick start guide](#) self-managed section to install and run the Elasticsearch and Kibana docker. In this example, we installed version 8.1.



Note down the user name/password and token created by Elasticsearch, you need these to start the Kibana UI and StorageGRID platform endpoint authentication.

Elasticsearch Service	Self-managed
Install and run Elasticsearch 1. Install and start Docker Desktop . 2. Run: <pre>docker network create elastic docker pull docker.elastic.co/elasticsearch/elasticsearch:8.1.0 docker run --name es-node01 --net elastic -p 9200:9200 -p 9300:9300 -it</pre> <p>When you start Elasticsearch for the first time, the following security configuration occurs automatically:</p> <ul style="list-style-type: none">• Certificates and keys are generated for the transport and HTTP layers.• The Transport Layer Security (TLS) configuration settings are written to <code>elasticsearch.yml</code>.• A password is generated for the <code>elastic</code> user.• An enrollment token is generated for Kibana. <p> NOTE You might need to scroll back a bit in the terminal to view the password and enrollment token.</p> <p>3. Copy the generated password and enrollment token and save them in a secure location. These values are shown only when you start Elasticsearch for the first time. You'll use these to enroll Kibana with your Elasticsearch cluster and log in.</p> <p> NOTE If you need to reset the password for the <code>elastic</code> user or other built-in users, run the <code>elasticsearch-reset-password</code> tool. To generate new enrollment tokens for Kibana or Elasticsearch nodes, run the <code>elasticsearch-create-enrollment-token</code> tool. These tools are available in the Elasticsearch bin directory.</p> <h3>Install and run Kibana</h3> <p>To analyze, visualize, and manage Elasticsearch data using an intuitive UI, install Kibana.</p> <ol style="list-style-type: none">1. In a new terminal session, run: <pre>docker pull docker.elastic.co/kibana/kibana:8.1.0 docker run --name kib-01 --net elastic -p 5601:5601 docker.elastic.co/kibana/kibana:8.1.0</pre><p>When you start Kibana, a unique link is output to your terminal.</p><ol style="list-style-type: none">2. To access Kibana, click the generated link in your terminal.<ol style="list-style-type: none">a. In your browser, paste the enrollment token that you copied and click the button to connect your Kibana instance with Elasticsearch.b. Log in to Kibana as the <code>elastic</code> user with the password that was generated when you started Elasticsearch.	

3. After the Kibana docker container has started, the URL link `https://0.0.0.0:5601` displays in the console. Replace 0.0.0.0 with the server IP address in the URL.
4. Log in to the Kibana UI by using user name `elastic` and the password generated by Elastic in the preceding step.
5. For first time login, on the dashboard welcome page, select **Explore On Your Own**. From the menu, select **Management > Dev Tools**.
6. On the Dev Tools Console screen, enter `PUT <index>` where you use this index for storing StorageGRID object metadata. We use the index name `sgmetadata` in this example. Click the small triangle symbol to execute the PUT command. The expected result displays on the right panel as shown in the following example screenshot.



The screenshot shows the Elasticsearch Dev Tools Console interface. The top navigation bar includes the elastic logo, a search bar, and tabs for Dev Tools (which is selected), Console, Search Profiler, Grok Debugger, and Painless Lab (BETA). Below the tabs are History, Settings, and Help buttons. The main area has a code editor with the following content:

```

1 PUT sgmetadata
2
3
4
5
6

```

On the right, the results panel shows the response:

```

1 { "acknowledged" : true,
2   "shards_acknowledged" : true,
3   "index" : "sgmetadata"
4 }
5
6

```

Platform services endpoint configuration

To configure endpoints for platform services, follow these steps:

1. On Tenant Manager, go to **STORAGE(S3) > Platform services endpoints**
2. Click **Create Endpoint**, enter the following, and then click **Continue**:
 - Display name example: `elasticsearch`
 - URI: `https://<elasticsearch-server-ip or hostname>:9200`
 - URN: `urn:<something>:es:::<some-unique-text>/<index-name>/_doc` where the index-name is the name you used on the Kibana console.
Example: `urn:local:es:::sgmd/sgmetadata/_doc`

Create endpoint

1 Enter details 2 Select authentication type 3 Verify server

Enter endpoint details

Enter the endpoint's display name, URI, and URN.

Display name ?

URI ?

URN ?

[Cancel](#)

[Continue](#)

3. Select Basic HTTP as the authentication type, enter the user name `elastic` and the password generated by the Elasticsearch installation process. To go to the next page, click Continue.

Authentication type ?

Select the method used to authenticate connections to the endpoint.

Basic HTTP

Username ?

Password ?




[Previous](#)

[Continue](#)

4. Select Do Not Verify Certificate and Test and Create Endpoint to verify the endpoint. If verification is

successful, an endpoint screen similar to the following screenshot displays. If the verification fails, verify the URN, URI, and username/password entries are correct.



Display name	Last error	Type	URI	URN
aws-opensearch		Search	https://search- sgdemo-fw223hpjv6lccxrpw3v3rte7i.us-east-1.es.amazonaws.com/	arn:aws:es:us-east-1:210811600158:domain/ sgdemo/sgmetadata/_doc
elasticsearch		Search	https:// 10.0.0.1:9200	urn:local:es::sgmd/sgmetadata/_doc

Bucket search integration service configuration

After the platform service endpoint is created, the next step is to configure this service at bucket level to send object metadata to the defined endpoint whenever an object is created, deleted, or its metadata or tags are updated.

You can configure search integration by using Tenant Manager to apply a custom StorageGRID configuration XML to a bucket as follows:

1. In Tenant Manager, go to STORAGE(S3) > Buckets
2. Click Create Bucket, enter the bucket name (for example, sgmetadata-test) and accept the default us-east-1 region.
3. Click Continue > Create Bucket.
4. To bring up the bucket Overview page, click the bucket name, then select Platform Services.
5. Select the Enable Search Integration dialog box. In the provided XML box, enter the configuration XML using this syntax.

The highlighted URN must match the platform services endpoint that you defined. You can open another browser tab to access the Tenant Manager and copy the URN from the defined platform services endpoint.

In this example, we used no prefix, meaning that the metadata for every object in this bucket is sent to the Elasticsearch endpoint defined previously.

```

<MetadataNotificationConfiguration>
  <Rule>
    <ID>Rule-1</ID>
    <Status>Enabled</Status>
    <Prefix></Prefix>
    <Destination>
      <Urn> urn:local:es:::sgmd/sgmetadata/_doc</Urn>
    </Destination>
  </Rule>
</MetadataNotificationConfiguration>

```

6. Use S3 Browser to connect to StorageGRID with the tenant access/secret key, upload test objects to sgmetadata-test bucket and add tags or custom metadata to objects.

File	Size	Type	Last Modified	Storage Class
Koala.jpg	762.53 KB	JPG File	3/19/2022 12:39:52 AM	STANDARD
Lighthouse.jpg	548.12 KB	JPG File	3/19/2022 12:39:52 AM	STANDARD
test1.txt	45 bytes	Text Document	3/19/2022 12:39:52 AM	STANDARD
test2.txt	35 bytes	Text Document	3/19/2022 12:39:52 AM	STANDARD

Tasks (14) Permissions Http Headers Tags Properties Preview Versions EventLog

URL: <https://10.193.204.106:10445/sgmetadata-test/Koala.jpg> Copy

Key	Value
date	01-01-2020
owner	testuser
project	test
type	jpg

+ Add | Edit | Delete | Default Tags.. ✓ Apply changes | Reload

7. Use the Kibana UI to verify that the object metadata was loaded to sgmetadata's index.
 - From the menu, select Management > Dev Tools.
 - Paste the sample query to the console panel on the left and click the triangle symbol to execute it.

The query 1 sample result in the following example screenshot shows four records. This matches number of objects in the bucket.

```
GET sgmetadata/_search
{
  "query": {
    "match_all": {}
  }
}
```

```
1 GET sgmetadata/_search
2 {
3   "query": {
4     "match_all": {}
5   }
6 }
```


```
1 {
2   "took": 1,
3   "timed_out": false,
4   "_shards": {
5     "total": 1,
6     "successful": 1,
7     "skipped": 0,
8     "failed": 0
9   },
10  "hits": {
11    "total": {
12      "value": 4,
13      "relation": "eq"
14    },
15    "max_score": 1.0,
16    "hits": [
17      {
18        "_index": "sgmetadata",
19        "_id": "sgmetadata-test_test1.txt",
20        "_score": 1.0,
21        "source": {
22          "bucket": "sgmetadata-test",
23          "key": "test1.txt",
24          "accountId": "18656646746705816489",
25          "size": 45,
26          "md5": "36b194a8ac536f09a7061f024b97211e",
27          "region": "us-east-1",
28          "metadata": {
29            "s3b-last-modified": "2017-04-29T10:24:29Z",
30            "sha256": "6bf96e898615852c94f8701580d9a0299487f4cbe4429a1e1d7d7f427ab10f51"
31          },
32          "tags": {
33            "owner": "testuser",
34            "project": "test"
35          }
36        },
37      },
38      {
39        "_index": "sgmetadata",
40        "_id": "sgmetadata-test_Koala.jpg",
41        "_score": 1.0,
42        "source": {
43          "bucket": "sgmetadata-test",
44          "key": "Koala.jpg",
45          "accountId": "18656646746705816489",
46          "size": 788831,
47          "md5": "2b04df3ecc1d94afddff082d139c6f15",
48          "region": "us-east-1",
49          "metadata": {
50            "s3b-last-modified": "2019-01-02T09:49Z",
51            "sha256": "84adda0e4c52c469acefe0c674a9144cd43eb2628c401c6b56b41242e2be4af1"
52          },
53          "tags": {
54            "date": "01-01-2020",
55            "owner": "testuser",
56            "project": "test",
57            "type": "jpg"
58          }
59        }
60      }
61    ]
62  }
63 }
```

The query 2 sample result in the following screenshot shows two records with tag type jpg.

```

GET sgmetadata/_search
{
  "query": {
    "match": {
      "tags.type": {
        "query" : "jpg" }
      }
    }
}

```



The screenshot shows the Elasticsearch Dev Tools Console interface. The top navigation bar includes 'Dev Tools' and 'Console'. The 'Console' tab is selected. The main area displays a search query and its results.

Search Query (Line 8):

```

GET sgmetadata/_search
{
  "query": {
    "match": {
      "tags.type": {
        "query" : "jpg" }
      }
    }
}

```

Search Results (Lines 10-59):

```

{
  "took" : 1,
  "timed_out" : false,
  "shards" : {
    "total" : 1,
    "successful" : 1,
    "skipped" : 0,
    "failed" : 0
  },
  "hits" : {
    "total" : 2,
    "value" : 2,
    "relation" : "eq"
  },
  "max_score" : 0.18232156,
  "hits" : [
    {
      "_index" : "sgmetadata",
      "_id" : "sgmetadata-test_koala.jpg",
      "_score" : 0.18232156,
      "_source" : {
        "bucket" : "sgmetadata-test",
        "key" : "Koala.jpg",
        "accountId" : "18656646746705016489",
        "size" : 288631,
        "md5" : "2b04df3ecc1d94afddff882d139c6f15",
        "region" : "us-east-1",
        "metadata" : {
          "s3b-last-modified" : "20190102T070549Z",
          "sha256" : "84a4da0e4c52c469ace6e0c674a9144cd43eb1628c001c0b56b41242e2be4af1"
        },
        "tags" : [
          "date" : "01-01-2020",
          "owner" : "testuser",
          "project" : "test",
          "type" : "jpg"
        ]
      }
    },
    {
      "_index" : "sgmetadata",
      "_id" : "sgmetadata-test_lighthouse.jpg",
      "_score" : 0.18232156,
      "_source" : {
        "bucket" : "sgmetadata-test",
        "key" : "Lighthouse.jpg",
        "accountId" : "18656646746705016489",
        "size" : 561276,
        "md5" : "896928874245120e07c3870287cce0ff3",
        "region" : "us-east-1",
        "metadata" : {
          "s3b-last-modified" : "20090714T051231Z",
          "sha256" : "ff86372ca435196075b6d8d29c98e9ccbe905da00ba057c8544fa001fa0d8e73"
        },
        "tags" : [
          "date" : "02-02-2022",
          "owner" : "testuser",
          "project" : "test",
          "type" : "jpg"
        ]
      }
    }
  ]
}

```

Where to find additional information

To learn more about the information that is described in this document, review the following documents and/or websites:

- [What are platform services](#)
- [StorageGRID 11.6 Documentation](#)

By *Angela Cheng*

Node Clone

Node clone considerations and performance.

Node clone considerations

Node clone can be a faster method for replacing existing appliance nodes for a tech refresh, increase capacity, or increase performance of your StorageGRID system. Node clone can also be useful for converting to node encryption with a KMS, or changing a storage node from DDP8 to DDP16.

- The used capacity of the source node is not relevant to the time required for the clone process to complete. Node clone is a full copy of the node including free space in the node.
- The source and destination appliances must be at the same PGE version
- The destination node must always have larger capacity than the source
 - Make sure the new destination appliance has a larger drive size than the source
 - If the destination appliance has the same size drives and is configured for DDP8, you can configure the destination for DDP16. If the source is already configured for DDP16 then node clone will not be possible.
 - When going from SG5660 or SG5760 appliances to SG6060 appliances be aware that the SG5x60's have 60 capacity drives where the SG6060 only has 58.
- The node clone process requires the source node to be offline to the grid for the duration of the cloning process. If an additional node goes offline during this time client services may be impacted.
- 11.8 and below: A storage node can only be offline for 15 days. If the cloning process estimate is close to 15 days or will exceed 15 days, use the expansion and decommission procedures.
 - 11.9: The 15 day limit has been removed.
- For a SG6060 or SG6160 with expansion shelves, you need to add the time for the correct shelf drive size to the time of the base appliance time to get the full clone duration.
- The number of volumes in a target storage appliance must be greater than or equal to the number of volumes in the source node. You cannot clone a source node with 16 object store volumes (rangedb) to a target storage appliance with 12 object store volumes even if the target appliance has larger capacity than the source node. Most storage appliances have 16 object store volumes, except the SGF6112 storage appliance that has only 12 object store volumes. For example, you cannot clone from a SG5760 to a SGF6112.

Node clone Performance estimates

The following tables contain calculated estimates for node clone duration. Conditions vary so, entries in **BOLD** may risk exceeding the 15 day limit for a node down.

DDP8

SG5612/SG5712/SG5812 → Any

Network Interface speed	4TB Drive size	8TB Drive size	10TB Drive size	12TB Drive size	16TB Drive size	18TB Drive size	22TB Drive size
10GB	1 Day	2 Days	2.5 Days	3 Days	4 Days	4.5 Days	5.5 Days
25GB	1 Day	2 Days	2.5 Days	3 Days	4 Days	4.5 Days	5.5 Days

SG5660 → SG5760/SG5860

Network Interface speed	4TB Drive size	8TB Drive size	10TB Drive size	12TB Drive size	16TB Drive size	18TB Drive size	22TB Drive size
10GB	3.5 Day	7 Days	8.5 Days	10.5 Days	13.5 Days	15.5 Days	18.5 Days
25GB	3.5 Day	7 Days	8.5 Days	10.5 Days	13.5 Days	15.5 Days	18.5 Days

SG5660 → SG6060/SG6160

Network Interface speed	4TB Drive size	8TB Drive size	10TB Drive size	12TB Drive size	16TB Drive size	18TB Drive size	22TB Drive size
10GB	2.5 Day	4.5 Days	5.5 Days	6.5 Days	9 Days	10 Days	12 Days
25GB	2 Day	4 Days	5 Days	6 Days	8 Days	9 Days	10 Days

SG5760/SG5860 → SG5760/SG5860

Network Interface speed	4TB Drive size	8TB Drive size	10TB Drive size	12TB Drive size	16TB Drive size	18TB Drive size	22TB Drive size
10GB	3.5 Day	7 Days	8.5 Days	10.5 Days	13.5 Days	15.5 Days	18.5 Days
25GB	3.5 Day	7 Days	8.5 Days	10.5 Days	13.5 Days	15.5 Days	18.5 Days

SG5760/SG5860 → SG6060/SG6160

Network Interface speed	4TB Drive size	8TB Drive size	10TB Drive size	12TB Drive size	16TB Drive size	18TB Drive size	22TB Drive size
10GB	2.5 Day	4.5 Days	5.5 Days	6.5 Days	9 Days	10 Days	12 Days
25GB	2 Day	3.5 Days	4.5 Days	5.5 Days	7 Days	8 Days	9.5 Days

SG6060/SG6160 → SG6060/SG6160

Network Interface speed	4TB Drive size	8TB Drive size	10TB Drive size	12TB Drive size	16TB Drive size	18TB Drive size	22TB Drive size
10GB	2.5 Day	4.5 Days	5.5 Days	6.5 Days	8.5 Days	9.5 Days	11.5 Days
25GB	2 Day	3 Days	4 Days	4.5 Days	6 Days	7 Days	8.5 Days

DDP16

SG5760/SG5860 → SG5760/SG5860

Network Interface speed	4TB Drive size	8TB Drive size	10TB Drive size	12TB Drive size	16TB Drive size	18TB Drive size	22TB Drive size
10GB	3.5 Day	6.5 Days	8 Days	9.5 Days	12.5 Days	14 Days	17 Days
25GB	3.5 Day	6.5 Days	8 Days	9.5 Days	12.5 Days	14 Days	17 Days

SG5760/SG5860 → SG6060/SG6160

Network Interface speed	4TB Drive size	8TB Drive size	10TB Drive size	12TB Drive size	16TB Drive size	18TB Drive size	22TB Drive size
10GB	2.5 Day	5 Days	6 Days	7.5 Days	10 Days	11 Days	13 Days
25GB	2 Day	3.5 Days	4 Days	5 Days	6.5 Days	7 Days	8.5 Days

SG6060/SG6160 → SG6060/SG6160

Network Interface speed	4TB Drive size	8TB Drive size	10TB Drive size	12TB Drive size	16TB Drive size	18TB Drive size	22TB Drive size
10GB	3 Day	5 Days	6 Days	7 Days	9.5 Days	10.5 Days	13 Days
25GB	2 Day	3.5 Days	4.5 Days	5 Days	7 Days	7.5 Days	9 Days

Expansion shelf (add to above SG6060/SG6160 for each shelf on source appliance)

Network Interface speed	4TB Drive size	8TB Drive size	10TB Drive size	12TB Drive size	16TB Drive size	18TB Drive size	22TB Drive size
10GB	3.5 Day	5 Days	6 Days	7 Days	9.5 Days	10.5 Days	12 Days
25GB	2 Day	3 Days	4 Days	4.5 Days	6 Days	7 Days	8.5 Days

By Aron Klein


Grid site relocation and site-wide network change procedure

This guide describes the preparation and procedure for StorageGRID site relocation in a multi-sites Grid. You should have a complete understand of this procedure and plan ahead to ensure smooth process and minimize interruption to clients.

If you need to change the Grid network of entire Grid, see [Change IP addresses for all nodes in grid](#).

Considerations before site relocation

- Site move should be completed, and all nodes online within 15 days to avoid Cassandra database rebuild. [Recover Storage Node down more than 15 days](#)
- If any ILM rule in active policy is using strict ingest behavior, consider changing it to balance or dual commit if customer wants to continue to PUT objects into the Grid during site relocation.
- For storage appliances with 60 drives or more, never move the shelf with disk drives installed. Label each disk drives and remove them from storage enclosure before pack/move.
- Change StorageGRID appliance Grid network VLAN can be performed remotely over admin network or client network. Or else plan to be onsite to perform the change before or after the relocation.
- Check if customer application is using HEAD or GET nonexistence object before PUT. If yes, change the bucket consistency to strong-site to avoid HTTP 500 error. If you are not sure, check S3 overview Grafana charts **Grid manager > Support > Metrics**, mouse over the 'Total Completed Request' chart. If there is very high count of 404 get Object or 404 head object, likely one or more applications are using head or get nonexistence object. The count is accumulative, mouse over different timeline to see the difference.



Procedure to change Grid IP address before site relocation

Steps

1. If new Grid network subnet will be used at the new location, [add the subnet to Grid network subnet list](#)
2. Log in to the primary Admin Node, use change-ip to make Grid IP change, must **stage** the change before shutdown the node for relocation.
 - a. Select 2 then 1 for Grid IP change

```
Editing: Node IP/subnet and gateway
```

```
Use up arrow to recall a previously typed value, which you can then edit
Use d or 0.0.0.0/0 as the IP/mask to delete the network from the node
Use q to complete the editing session early and return to the previous menu
Press <enter> to use the value shown in square brackets
```

```
=====
Site: LONDON
=====
```

```
LONDON-ADM1 Grid IP/mask [ 10.45.74.14/26 ]: 10.45.74.24/26
LONDON-S1 Grid IP/mask [ 10.45.74.16/26 ]: 10.45.74.26/26
LONDON-S2 Grid IP/mask [ 10.45.74.17/26 ]: 10.45.74.27/26
LONDON-S3 Grid IP/mask [ 10.45.74.18/26 ]: 10.45.74.28/26
=====
```

```
LONDON-ADM1 Grid Gateway [ 10.45.74.1 ]:
LONDON-S1 Grid Gateway [ 10.45.74.1 ]:
LONDON-S2 Grid Gateway [ 10.45.74.1 ]:
LONDON-S3 Grid Gateway [ 10.45.74.1 ]:
=====
```

```
=====
Site: OXFORD
=====
```

```
OXFORD-ADM1 Grid IP/mask [ 10.45.75.14/26 ]:
OXFORD-S1 Grid IP/mask [ 10.45.75.16/26 ]:
OXFORD-S2 Grid IP/mask [ 10.45.75.17/26 ]:
OXFORD-S3 Grid IP/mask [ 10.45.75.18/26 ]:
=====
```

```
OXFORD-ADM1 Grid Gateway [ 10.45.75.1 ]:
OXFORD-S1 Grid Gateway [ 10.45.75.1 ]:
OXFORD-S2 Grid Gateway [ 10.45.75.1 ]:
OXFORD-S3 Grid Gateway [ 10.45.75.1 ]:
=====
```

```
Finished editing. Press Enter to return to menu. █
```

- b. select 5 to show changes

```
=====
Site: LONDON
=====
LONDON-ADM1 Grid IP      [ 10.45.74.14/26 ]: 10.45.74.24/26
LONDON-S1  Grid IP      [ 10.45.74.16/26 ]: 10.45.74.26/26
LONDON-S2  Grid IP      [ 10.45.74.17/26 ]: 10.45.74.27/26
LONDON-S3  Grid IP      [ 10.45.74.18/26 ]: 10.45.74.28/26
Press Enter to continue
```

c. select 10 to validate and apply the change.

```
Welcome to the StorageGRID IP Change Tool.

Selected nodes: all

1:  SELECT NODES to edit
2:  EDIT IP/mask and gateway
3:  EDIT admin network subnet lists
4:  EDIT grid network subnet list
5:  SHOW changes
6:  SHOW full configuration, with changes highlighted
7:  VALIDATE changes
8:  SAVE changes, so you can resume later
9:  CLEAR all changes, to start fresh
10: APPLY changes to the grid
0:  Exit

Selection: 10
```

d. Must select **stage** in this step.

```
Validating new networking configuration... PASSED.
Checking for Grid Network IP address swaps... PASSED.

Applying these changes will update the following nodes:

LONDON-ADM1
LONDON-S1
LONDON-S2
LONDON-S3

The following nodes will also require restarting:

LONDON-ADM1
LONDON-S1
LONDON-S2
LONDON-S3

Select one of the following options:

apply: apply all changes and automatically restart nodes (if necessary)
stage: stage the changes; no changes will take effect until the nodes are restarted
cancel: do not make any network changes at this time

[apply/stage/cancel]> stage
```

e. If primary admin node is included in above change, Enter 'a' to restart primary admin node manually

```
10.45.74.14 - PuTTY
Validating new networking configuration... PASSED.
Checking for Grid Network IP address swaps... PASSED.

Applying these changes will update the following nodes:
LONDON-ADM1
LONDON-S1
LONDON-S2
LONDON-S3

The following nodes will also require restarting:
LONDON-ADM1
LONDON-S1
LONDON-S2
LONDON-S3

Select one of the following options:
apply: apply all changes and automatically restart nodes (if necessary)
stage: stage the changes; no changes will take effect until the nodes are restarted
cancel: do not make any network changes at this time

[apply/stage/cancel]> stage

Generating new grid networking description file... PASSED.
Running provisioning... PASSED.
Updating network configuration on LONDON-S1... PASSED.
Updating network configuration on LONDON-S2... PASSED.
Updating network configuration on LONDON-S3... PASSED.
Updating network configuration on LONDON-ADM1... PASSED.

Finished staging network changes. You must manually restart these nodes for the changes to take effect:
LONDON-ADM1 (has IP 10.45.74.14 until restart)
LONDON-S1 (has IP 10.45.74.16 until restart)
LONDON-S2 (has IP 10.45.74.17 until restart)
LONDON-S3 (has IP 10.45.74.18 until restart)

Importing bundles... PASSED.
*****
*          IMPORTANT          *
*          *          *
* A new recovery package has been generated as a result of the *          *
* configuration change. Select Maintenance > Recovery Package *          *
* in the Grid Manager to download it.          *
*****
Network Update Complete. Primary admin restart required. Select 'continue' to restart this node immediately, 'abort' to restart manually.
Enter a to abort, c to continue [a/c]> a
```

f. Press enter to return to previous menu and exit from change-ip interface.

```
Network Update Complete. Primary admin restart required. Select 'continue' to restart this node immediately, 'abort' to restart manually.
Enter a to abort, c to continue [a/c]> a
Restart aborted. You must manually restart this node as soon as possible
Press Enter to return to the previous menu.
```

3. From Grid Manager, download the new recovery package. **Grid manager > Maintenance > Recovery package**
4. If VLAN change is required on StorageGRID appliance, see the section [Appliance VLAN change](#).
5. Shutdown all nodes and/or appliances at the site, label/remove disk drives if necessary, unrack, pack and move.
6. If you plan to change admin network ip and/or client VLAN and ip address, you can perform the change after the relocation.

Appliance VLAN change

The procedure below assume you have remote access to StorageGRID appliance's admin or client network to perform the change remotely.

Steps

1. Before shutdown the appliance,
[place the appliance in maintenance mode](#).

2. Using a browser to access the StorageGRID appliance installer GUI using <https://<admin-or-client-network-ip>:8443>. Cannot use Grid IP as the new Grid IP already in place once the appliance is boot into maintenance mode.
3. Change the VLAN for Grid network. If you are accessing the appliance over client network, you cannot change Client VLAN at this time, you can change it after the move.
4. ssh to the appliance and shutdown the node using 'shutdown -h now'
5. After the appliances are ready at new site, access to the StorageGRID appliance installer GUI using <https://<grid-network-ip>:8443>. Confirm the storage are in optimal state and network connectivity to other Grid nodes using ping/nmap tools in the GUI.
6. If plan to change client network IP, you can change the client VLAN at this stage. The client network is not ready until you update the client network ip using change-ip tool in later step.
7. Exit maintenance mode. From the StorageGRID Appliance Installer, select **Advanced > Reboot Controller**, and then select **Reboot into StorageGRID**.
8. After all nodes are up and Grid shows no connectivity issue, use change-ip to update the appliance admin network and client network if necessary.

Migrating object-based storage from ONTAP S3 to StorageGRID

Enabling enterprise-grade S3 by seamlessly migrating object-based storage from ONTAP S3 to StorageGRID

Enabling enterprise-grade S3 by seamlessly migrating object-based storage from ONTAP S3 to StorageGRID

Migration Demo

This is a demonstration on migrating users and buckets from ONTAP S3 to StorageGRID.

Enabling enterprise-grade S3 by seamlessly migrating object-based storage from ONTAP S3 to StorageGRID

Enabling enterprise-grade S3 by seamlessly migrating object-based storage from ONTAP S3 to StorageGRID

Preparing ONTAP

For demonstration purposes we will create an SVM object store server, user, group, group policy and buckets.

Create the Storage Virtual Machine

In ONTAP System Manager, navigate to Storage VM's and add a new storage VM.



The screenshot shows the ONTAP System Manager interface. The left sidebar has a 'DASHBOARD' tab, followed by a 'STORAGE' section with 'INSIGHTS' and 'STORAGE' sub-options: Overview, Volumes, LUNs, Consistency groups, NVMe namespaces, Shares, Buckets, Qtrees, Quotas, Storage VMs (which is selected and highlighted in blue), and Tiers. The main content area is titled 'cluster1 Version 9.14.1P1'. It displays a 'Health' section with a green checkmark and the text 'Cluster is healthy'. Below this is a 'SIMBOX' icon with a lightning bolt, representing a storage virtual machine. The SIMBOX icon contains a small ONTAP logo and a grid of horizontal lines representing storage volumes.

Select the "Enable S3" and "Enable TLS" check boxes and configure the HTTP(S) ports. Define the IP, subnet mask and define the gateway and broadcast domain if not using the default or required in your environment.

Add storage VM

STORAGE VM NAME

Access protocol

SMB/CIFS, NFS, S3 iSCSI FC NVMe

Enable SMB/CIFS
 Enable NFS
 Enable S3

S3 SERVER NAME

Enable TLS

PORT

443

CERTIFICATE

Use system-generated certificate i
 Use external-CA signed certificate

Use HTTP (non-secure)

PORT

8080

DEFAULT LANGUAGE i

NETWORK INTERFACE

Use multiple network interfaces when client traffic is high.

onPrem-01


IP ADDRESS	SUBNET MASK	GATEWAY	BROADCAST DOMAIN AND PORT
192.168.0.200	24	Add optional gateway	Default

Storage VM administration

Enable maximum capacity limit
The maximum capacity that all volumes in this storage VM can allocate. [Learn More](#)

Manage administrator account

As part of the SVM creation a user will be created. Download the S3 keys for this user and close the window.



Once the SVM has been created, edit the SVM and add the DNS settings.

Services

NIS ✗ ⚙

Not configured

Name service switch ✗ ✎

Services lookup order ⓘ

HOSTS
Files, then DNS

GROUP
Files


NAME MAP
Files

NETGROUP
Files

DNS ✗ Configure

Not configured

Define the DNS name and IP.



Create SVM S3 User

Now we can configure the S3 users and group. Edit the S3 settings.

Protocols

NFS Not configured	 	SMB/CIFS Not configured	 	iS N
NVMe Not configured	 	S3 STATUS  Enabled TLS Disabled HTTP Enabled	 	

Add a new user.

Storage VMs

+ Add : More

Name
svm_demo

S3 All settings

Enabled

Server [Edit](#)


FQDN: s3portal.demo.netapp.com

TLS: Disabled TLS PORT: 443

HTTP: Enabled HTTP PORT: 8080

Users [Groups](#) [Policies](#)

User name	Access key	Key expiration time
root		-
sm_s3_user	34EH21411SMW1YOV3NQY	Valid forever



Input the user name and key expiration.

Storage VMs

+ Add : More

Name
svm_demo

S3 All settings

Enabled

Server [Edit](#)


FQDN: s3portal.demo.netapp.com

TLS: Disabled TLS PORT: 443

HTTP: Enabled HTTP PORT: 8080

Users [Groups](#) [Policies](#)

User name	Access key	Key expiration time
root		-
sm_s3_user	34EH21411SMW1YOV3NQY	Valid forever



Download the S3 keys for the new user.

Added user

**USER NAME**

demo_s3_user

ACCESS KEY

3TVPI142JGE3Y7FV2KC0

**SECRET KEY**

75a1QqKBU4quA132twl4g4lC4Gg5PP30ncy0sPE8

[Hide secret key](#)**KEY EXPIRATION TIME**

Valid forever

 The secret key won't be displayed again. Save this key for future use.

[Download](#)[Close](#)

USER NAME

3TVPI142JGE3Y7FV2KC0

VALID FOREVER

Create SVM S3 group

On the Groups tab of the SVM S3 settings, add a new group with the user created above and FullAccess permissions.

Add group

X

NAME

demo_s3_group

USERS

demo_s3_user X

POLICIES

FullAccess X

Cancel

Save



Create SVM S3 buckets

Navigate to the Buckets section and click the "+Add" button.



DASHBOARD

INSIGHTS

STORAGE ^

Overview

Volumes

LUNs

Consistency groups

NVMe namespaces

Shares

Buckets

Qtrees

Quotas

Storage VMs

Tiers

Buckets



Name

Stora

Enter a name, capacity, and deselect the "Enable ListBucket access..." check box. and click on the "More options" button.

Add bucket

X

NAME

bucket

CAPACITY

100



GiB



Enable ListBucket access for all users on the storage VM "svm_demo".
Enabling this will allow users to access the bucket.

More options

Cancel

Save

In the "More options" section select the enable versioning check box. and click the "Save" button.

Add bucket

X

NAME

FOLDER (OPTIONAL)

[Browse](#)

Specify the folder to map to this bucket. [Know more](#)

CAPACITY

Use for tiering

If you select this option, the system will try to select low-cost media with optimal performance for the tiered data.

Enable versioning


Versioning-enabled buckets allow you to recover objects that were accidentally deleted or overwritten. After versioning is enabled, it can't be disabled. However, you can suspend versioning.

PERFORMANCE SERVICE LEVEL

▼

Not sure? [Get help selecting type](#)

Repeat the process and create a second bucket without versioning enabled. Enter a name, the same capacity as bucket one, and deselect the "Enable ListBucket access..." check box. and click on the "Save" button.



By Rafael Guedes, and Aron Klein

Enabling enterprise-grade S3 by seamlessly migrating object-based storage from ONTAP S3 to StorageGRID

Enabling enterprise-grade S3 by seamlessly migrating object-based storage from ONTAP S3 to StorageGRID

Preparing StorageGRID

Continuing the configuration for this demo we will create a Tenant, user, security group, group policy, and bucket.

Create the tenant

Navigate to the "Tenants" tab and click on the "create" button

NetApp | StorageGRID Grid Manager

Search by page title Search

DASHBOARD

ALERTS 0

NODES

TENANTS

ILM

CONFIGURATION

MAINTENANCE

SUPPORT

Tenants

View information for each tenant account. Depending on the timing of ingests, network connectivity, and node status, the usage data shown might be out of date. To view more recent values, select the tenant name.

date Export to CSV Actions Search tenants by name or ID Search

No results

	Name ? ▲	Logical space used ? ▲	Quota utilization ? ▲	Quota ? ▲	Object count ? ▲	Sign in/Copy URL ?
No tenants found.						
Create						

Fill in the details for the tenant providing a tenant name, select S3 for the client type, and no quota is required. No need to select platform services or allow S3 select. You can choose to use own Identity source if you choose. Set the root password and click on the finish button.

Click on the tenant name to view the tenant details. **You will need the tenant ID later so copy it off.** Click on the Sign in button. This will bring you to the tenant portal login. Save the URL for future use.

Tenants


View information for each tenant account. Depending on the timing of ingests, network connectivity, and node status, the usage data shown might be out of date. To view more recent values, select the tenant name.

Create Export to CSV Actions Search tenants by name or ID Search Displaying one result

	Name ? ▲	Logical space used ? ▲	Quota utilization ? ▲	Quota ? ▲	Object count ? ▲	Sign in/Copy URL ?
<input type="checkbox"/>	tenant_demo	0 bytes	—	—	0	→ Copy

← Previous 1 Next →

This will bring you to the tenant portal login. Save the URL for future use, and enter the root user credentials.



Create the user

Navigate to the Users tab and create a new user.

Actions	Username	Full Name	Denied	Type
<input type="checkbox"/>	root	Root		Local

Optional

Enter user credentials

Create a new local user and configure user access.

Full name ?

Must contain at least 1 and no more than 128 characters

Username ?

Password

?

Must contain at least 8 and no more than 32 characters

Confirm password


?

Deny access

Do you want to prevent this user from signing in regardless of assigned group permissions?

Yes No

[Cancel](#) [Continue](#)



Now that the new user has been created, click on the users name to open the details of the user.

Copy the user ID from the URL to be used later.

⚠ Not secure | <https://192.168.0.80/ui/#/users/ebc132e2-cfc3-42c0-a445-3b4465cb523c>

Power Controls Accounts cluster1-mgmt cluster2-mgmt Blue XP

NetApp | StorageGRID Tenant Manager

Users > Demo S3 User

Overview

Full name:	Demo S3 User edit
Username:	demo_s3_user
User type:	Local
Denied access:	Yes
Access mode:	No Groups
Group membership:	None

>Password [edit](#) Access [edit](#) Access keys [edit](#) Groups [edit](#)

Change password

Change this user's password.

***** [Change Password](#)

To create the S3 keys click on the user name.

NetApp | StorageGRID Tenant Manager

DASHBOARD STORAGE (S3) My access keys Buckets Platform services endpoints ACCESS MANAGEMENT Groups Users Identity federation

Users


View local and federated users. Edit properties and group membership of local users.

2 users [Create user](#)

Actions	Username	Full Name	Denied	Type
<input type="checkbox"/>	root	Root		Local
<input type="checkbox"/>	demo_s3_user	Demo S3 User	✓	Local

← Previous 1 Next →

Select the "Access keys" tab and click on the "Create Key" button. There is no need to set an expiration time. Download the S3 keys as they cannot be retrieved again once the window is closed.



Create the security group

Now go to the Groups page and create a new group.

Create group

X

- 1 Choose a group type
- 2 Manage permissions
- 3 Set S3 group policy
- 4 Add users
Optional

Choose a group type ?

Create a new local group or import a group from the external identity source.

Local group

Federated group

Create local groups to assign permissions to any local users you defined in StorageGRID.

Display name

Demo S3 Group

Must contain at least 1 and no more than 32 characters

Unique name ?

demo_s3_group

Cancel

Continue



Set the group permissions to Read-Only. This is the Tenant UI permissions, not the S3 permissions.



Choose a group type



Manage permissions



Set S3 group policy

Add users
Optional

Manage group permissions

Select an access mode for this group and select one or more permissions.

Access mode [?](#)

Select whether users can change settings and perform operations or whether they can only view settings and features.



Read-write



Read-only

Group permissions [?](#)

Select the permissions you want to assign to this group.

[Root access](#)

Allows users to access all administration features. Root access permission supersedes all other permissions.

[Manage all buckets](#)

Allows users to change settings of all S3 buckets (or Swift containers) in this account.

[Manage endpoints](#)

Allows users to configure endpoints for platform services.

[Manage your own S3 credentials](#)

Allows users to create and delete their own S3 access keys.

[Previous](#)[Continue](#)

S3 permissions are controlled with the group policy (IAM Policy). Set the Group policy to custom and paste the json policy in the box. This policy will allow users of this group to list the buckets of the tenant and perform any S3 operations in the bucket named "bucket" or sub-folders in the bucket named "bucket".

```
{
  "Statement": [
    {
      "Effect": "Allow",
      "Action": "s3>ListAllMyBuckets",
      "Resource": "arn:aws:s3:::*"
    },
    {
      "Effect": "Allow",
      "Action": "s3:*",
      "Resource": ["arn:aws:s3:::bucket", "arn:aws:s3:::bucket/*"]
    }
  ]
}
```

Create group

1 Choose a group type — 2 Manage permissions — 3 Set S3 group policy — 4 Add users Optional

Set S3 group policy ?

An S3 group policy controls user access permissions to specific specific S3 resources, including buckets. Non-root users have no access by default.

No S3 Access

Read Only Access

Full Access

Custom


(Must be a valid JSON formatted string.)

```

    "Effect": "Allow",
    "Action": "s3>ListAllMyBuckets",
    "Resource": "arn:aws:s3:::*"
  },
  {
    "Effect": "Allow",
    "Action": "s3:*",
    "Resource": ["arn:aws:s3:::bucket", "arn:aws:s3:::bucket/*"]
  }
}
  
```

[Previous](#) [Continue](#)

Finally, add the user to the group and finish.



Add users

(This step is optional. If required, you can save this group and add users later.)

Select local users to add to the group **Demo S3 Group**.

Username	Full Name	Denied
demo_s3_user	Demo S3 User	✓

Previous Create group

Create two buckets

Navigate to the buckets tab and click on the Create bucket button.

The screenshot shows the 'Buckets' tab in the NetApp StorageGRID Tenant Manager. The left sidebar has sections for 'DASHBOARD', 'STORAGE (S3)' (selected), 'My access keys', 'Buckets' (selected), 'Platform services endpoints', 'ACCESS MANAGEMENT', 'Groups', 'Users', and 'Identity federation'. The main area is titled 'Buckets' and contains a sub-header 'Create buckets and manage bucket settings.' Below this is a table with the following columns: 'Name' (with a dropdown arrow), 'Region' (with a dropdown arrow), 'Object Count' (with a dropdown arrow), 'Space Used' (with a dropdown arrow), and 'Date Created' (with a dropdown arrow). A note below the table says '0 buckets'. At the bottom of the table is a blue button labeled 'Create bucket'. A note at the top right says 'Experimental S3 Console' with a link icon.

Define the bucket name and region.

Create bucket

1 Enter details

2

Manage object settings
Optional

Enter bucket details

Enter the bucket's name and select the bucket's region.

Bucket name [?](#)

bucket

Region [?](#)

us-east-1

[Cancel](#)

[Continue](#)

On this first bucket enable versioning.

Create bucket

1 Enter details

2

Manage object settings
Optional

Manage object settings Optional

Object versioning


Enable object versioning if you want to store every version of each object in this bucket. You can then retrieve previous versions of an object as needed.

Enable object versioning

[Previous](#)

[Create bucket](#)

Now create a second bucket without versioning enabled.



Enter bucket details

Enter the bucket's name and select the bucket's region.

Bucket name ?

sg-dummy


Region ?

us-east-1

Cancel

Continue

Do not enable versioning on this second bucket.



Previous

Create bucket

By Rafael Guedes, and Aron Klein

Enabling enterprise-grade S3 by seamlessly migrating object-based storage from ONTAP S3 to StorageGRID

Enabling enterprise-grade S3 by seamlessly migrating object-based storage from ONTAP S3 to StorageGRID

Populate the Source Bucket

Lets put some objects in the source ONTAP bucket. We will use S3Browser for this demo but you could use any tool you are comfortable with.

Using the ONTAP user s3 keys created above, configure S3Browser to connect to your ontap system.

 Add New Account

Add New Account

Enter new account details and click Add new account

Display name:

Bucket (original and post-migration)

Assign any name to your account.

Account type:

S3 Compatible Storage

Choose the storage you want to work with. Default is Amazon S3 Storage.

REST Endpoint:

s3portal.demo.netapp.com:8080

Specify S3-compatible API endpoint. It can be found in storage documentation. Example: rest.server.com:8080

Access Key ID:

3TVPI142JGE3Y7FV2KC0

Required to sign the requests you send to Amazon S3, see more details at <https://s3browser.com/keys>

Secret Access Key:

Required to sign the requests you send to Amazon S3, see more details at <https://s3browser.com/keys>

Encrypt Access Keys with a password:

Turn this option on if you want to protect your Access Keys with a master password.

Use secure transfer (SSL/TLS)

If checked, all communications with the storage will go through encrypted SSL/TLS channel

[advanced settings..](#)

 [Add new account](#)  [Cancel](#)

Now lets upload some files to the versioning enabled bucket.

S3 Browser 11.6.7 - Free Version (for non-commercial use only) - Bucket (original and post-migration)


Accounts Buckets Files Bookmarks Tools Upgrade to Pro! Help

New bucket + Add external bucket Refresh

Path: /

Name	Size	Type	Last Modified	Storage Class

Upload file(s) 

Upload folder(s) 

Upload Download Delete New Folder Refresh

Tasks (1) Permissions Headers Tags Properties Preview Versions Eventlog

Task	Size	%	Progress	Status	Speed

S3 Browser 11.6.7 - Free Version (for non-commercial use only) - Bucket (original and post-migration)

Open

Downloads This PC Downloads

Search Downloads

Organize New folder

Name	Date modified	Type	Size
9141P1_q_image.tgz	3/22/2024 1:25 AM	TGZ File	2,641,058 KB
cluster1_demo_s3_user_s3_user.txt	3/23/2024 11:04 PM	Text Document	1 KB
cluster1_svm_demo_s3_details (1).txt	3/23/2024 11:03 PM	Text Document	1 KB
cluster1_svm_demo_s3_details.txt	3/23/2024 11:01 PM	Text Document	1 KB
hfs.exe	3/22/2024 1:24 AM	Application	2,121 KB
hotfix-install-11.6.0.14	3/23/2024 11:55 AM	14 File	717,506 KB
putty	7/18/2020 6:39 PM	Shortcut	2 KB
s3browser-11-6-7.exe	3/23/2024 12:36 PM	Application	9,807 KB

File name: "s3browser-11-6-7.exe" "cluster1_demo_s3_user_s3_user.txt" "cluster1_svm_demo_s3_details (1).txt" "cluster1_svm_demo_s3_details.txt" "hfs.exe" "putty"

Open Cancel

S3 Browser 11.6.7 - Free Version (for non-commercial use only) - Bucket (original and post-migration)

Accounts Buckets Files Bookmarks Tools Upgrade to Pro! Help

New bucket Add external bucket Refresh

Path: /

Name	Size	Type	Last Modified	Storage Class
cluster1 дем...	157 bytes	Text Document	3/23/2024 11:23:25 PM	STANDARD
cluster1_svm...	211 bytes	Text Document	3/23/2024 11:23:25 PM	STANDARD
cluster1_svm...	211 bytes	Text Document	3/23/2024 11:23:25 PM	STANDARD
putty.exe	834.05 KB	Application	3/23/2024 11:23:25 PM	STANDARD
hfs.exe	2.07 MB	Application	3/23/2024 11:23:25 PM	STANDARD
s3browser-11...	9.58 MB	Application	3/23/2024 11:23:26 PM	STANDARD

Upload Download Delete New Folder Refresh

Tasks (1) Permissions Headers Tags Properties Preview Versions Eventlog

Now lets create some object versions in the bucket.

Delete a file.

S3 Browser 11.6.7 - Free Version (for non-commercial use only) - Bucket (original and post-migration)

Accounts Buckets Files Bookmarks Tools Upgrade to Pro! Help

New bucket Add external bucket Refresh

Path: /

Name	Size	Type	Last Modified	Storage Class
cluster1 дем...	157 bytes	Text Document	3/23/2024 11:23:25 PM	STANDARD
cluster1_svm...	211 bytes	Text Document	3/23/2024 11:23:25 PM	STANDARD
cluster1_svm...	211 bytes	Text Document	3/23/2024 11:23:25 PM	STANDARD
putty.exe	834.05 KB	Application	3/23/2024 11:23:25 PM	STANDARD
hfs.exe	2.07 MB	Application	3/23/2024 11:23:25 PM	STANDARD
s3browser-11...	9.58 MB	Application	3/23/2024 11:23:26 PM	STANDARD

Upload Download Delete New Folder


Confirm File Delete

Are you sure to delete 'putty.exe'?

Yes No

Tasks (1) Permissions Headers Tags Properties Preview Versions Eventlog

Upload a file that already exists in the bucket to copy the file over itself and create a new version of it.



In S3Browser we can view the versions of the objects we just created.

Name	Size	Type	Last Modified	Storage Class
cluster1_demo_s3_user_s3_user.txt	157 bytes	Text Document	3/23/2024 11:23:25 PM	STANDARD
cluster1_svm_demo_s3_details (1).txt	211 bytes	Text Document	3/23/2024 11:23:25 PM	STANDARD
hfs.exe	2.07 MB	Application	3/23/2024 11:23:36 PM	STANDARD
s3browser-11-6-7.exe	9.58 MB	Application	3/23/2024 11:23:26 PM	STANDARD

Key	Last Modified	ETag	Size	Storage Class	Owner	Version Id
cluster1_demo_s3_user_s3_user.txt	3/23/2024 11:23:25 PM	ac4fc9543e97ef3678b2b6ed6a60e1bc	157 bytes	STANDARD	Unknown (Unknown)	Mzg0MjQ1MDAwL...
revision # 1 (current)						
cluster1_svm_demo_s3_details (1).txt	3/23/2024 11:23:25 PM	40775b646aacfef19de71eebf5f04	211 bytes	STANDARD	Unknown (Unknown)	NDg0MjQ1MDAw...
revision # 1 (current)						
cluster1_svm_demo_s3_details.txt	3/23/2024 11:23:25 PM	17d06518569480a507af99ecc010e2	211 bytes	STANDARD	Unknown (Unknown)	NTU2NzMDAwL...
revision # 1 (current)						
hfs.exe	3/23/2024 11:23:36 PM	9e8557e98ed1269372fface91d63477	2.07 MB	STANDARD	Unknown (Unknown)	NzQ1OTE4MDAw...
revision # 2 (current)						
revision # 1	3/23/2024 11:23:25 PM	9e8557e98ed1269372fface91d63477	2.07 MB	STANDARD	Unknown (Unknown)	Njk2ODI0MDAwL...
putty.exe	3/23/2024 11:23:31 PM	54cb91395cdaad9d4782533c21f6e9	834.05 KB	STANDARD	Unknown (Unknown)	NjMzMDAwMC5...
revision # 2 (deleted)						
revision # 1	3/23/2024 11:23:25 PM	54cb91395cdaad9d4782533c21f6e9	834.05 KB	STANDARD	Unknown (Unknown)	NzE2NzEyMDAw...
s3browser-11-6-7.exe	3/23/2024 11:23:26 PM	ae36fb97054762962d6937c5df0820-2	9.58 MB	STANDARD	Unknown (Unknown)	NDY2ODcwMDEu...
revision # 1 (current)						

Establish the replication relationship

Lets start sending data from ONTAP to StorageGRID.

In ONTAP System Manager navigate to "Protection/Overview". Scroll down to "Cloud object stores". and click the "Add" button and select "StorageGRID".

Lets you select specific volumes for protection if you don't need to protect entire storage VMs.

Lets you select which volumes you want to be backed up to a cloud destination.

NetApp SnapCenter software simplifies backup, restore, and clone management for the applications hosted across ONTAP enabled platforms. [Use NetApp SnapCenter for application-consistent protection.](#)

Bucket protection

SnapMirror (local or remote)

0% 10% 20% 30% 40% 50% 60% 70% 80% 90% 100% **2 of the 2 buckets aren't protected.**

Back up to cloud

0% 10% 20% 30% 40% 50% 60% 70% 80% 90% 100% **2 of the 2 buckets aren't backed up to the cloud.**

Protect buckets

Lets you select specific buckets for setting up SnapMirror protection to either an ONTAP destination or a cloud destination.

StorageGRID **i**

ONTAP S3 **i**

AWS Amazon S3 **i**

Others **i**

StorageGRID is a cloud object store that provides a simple, efficient, and cost-effective way to store and manage data in the cloud. It offers a range of features, including data replication, data deduplication, and data compression, to help you manage your data more effectively. StorageGRID is designed to be highly available and reliable, with built-in redundancy and failover mechanisms.

+ Add

Input the StorageGRID information by providing a name, URL style (for this demo we will use Path-style URLs). Set the object store scope to "Storage VM".

Add cloud object store

NAME

URL STYLE

Path-style URL

OBJECT STORE SCOPE

Cluster Storage VM

USE BY **i**


SnapMirror ONTAP S3 SnapMirror

SERVER NAME (FQDN)

If you are using SSL, set the load balancer endpoint port and copy in the StorageGRID endpoint certificate

here. otherwise uncheck the SSL box and input the HTTP endpoint port here.

Input the StorageGRID user S3 keys and bucket name from the StorageGRID configuration above for the destination.




NODE	IP ADDRESS	SUBNET MASK	BROADCAST DOMAIN	GATEWAY
onPrem-01	192.168.0.113	24	Default	192.168.0.1

Use HTTP proxy

Save **Cancel**

Now that we have a destination target configured, we can configure the policy settings for the target. Expand "Local policy settings" and select "continuous".



Edit the continuous policy and change the "Recovery point objective" from "1 Hours" to "3 Seconds".

Policies Protection overview

Name	Description	Policy type	Scope
Continuous	Policy for S3 bucket mirroring.	(All)	Cluster
THROTTLE Unlimited	RECOVERY POINT OBJECTIVE 1 Hours		

Now we can configure snapmirror to replicate the bucket.

```
snapmirror create -source-path sv_demo: /bucket/bucket -destination-path sgws_demo: /objstore -policy Continuous
```

```
cluster1-mgmt
Using username "admin".
Using keyboard-interactive authentication.
Password:
Last login time: 3/24/2024 00:02:00
cluster1::> snapmirror create -source-path sv_demo:/bucket/bucket -destination-path sgws_demo:/objstore -policy Continuous
[Job 220] Job is queued: Create an S3 SnapMirror relationship between bucket "sv_demo:bucket" and bucket "objstore/sgws_demo"..
cluster1::>
```

The bucket will now show a cloud symbol in the bucket list under protection.

Name	Storage VM	Lifecycle rules	Capacity (available total)	Protection	Path
bucket	svm_demo	0	100 GiB 100 GiB	Cloud	-
ontap-dummy	svm_demo	0	100 GiB 100 GiB	Cloud	-

If we select the bucket and go to the "SnapMirror (ONTAP or Cloud)" tab we will see the snapmirror relationship status.

Buckets

The screenshot shows the 'Buckets' page in the NetApp StorageGRID web interface. On the left, a sidebar lists buckets: 'bucket' (selected) and 'ontap-dummy'. The main content area shows a 'SnapMirror (ONTAP or cloud)' tab. A table displays the replication configuration:

Source	Destination	Protection policy	Relationship health	State
svm_demo/bucket/bucket	sgws_demo/objstore	Continuous	Healthy	Mirrored

The replication details

We now have a successfully replicating bucket from ONTAP to StorageGRID. But what is actually replicating? Our source and destination are both versioned buckets. Do the previous versions also replicate to the destination? If we look at our StorageGRID bucket with S3Browser we see that the existing versions did not replicate and our deleted object does not exist, nor does a delete marker for that object. Our duplicated object only has 1 version in the StorageGRID bucket.

S3 Browser 11.6.7 - Free Version (for non-commercial use only) - Bucket (Migration Temp)

The screenshot shows the S3 Browser interface. The left sidebar shows buckets: 'bucket' (selected) and 'sg-dummy'. The main content area shows a table of objects in the 'bucket' bucket:

Name	Size	Type	Last Modified	Storage Class
cluster1_demo...	157 bytes	Text Document	3/24/2024 12:13:53 AM	STANDARD
cluster1_svm...	211 bytes	Text Document	3/24/2024 12:13:53 AM	STANDARD
cluster1_svm...	211 bytes	Text Document	3/24/2024 12:13:53 AM	STANDARD
hfs.exe	2.07 MB	Application	3/24/2024 12:13:53 AM	STANDARD
s3brd4-user11...	9.58 MB	Application	3/24/2024 12:13:53 AM	STANDARD

Below the table are buttons for 'Upload', 'Download', 'Delete', 'New Folder', and 'Refresh'. The 'Tasks' tab is selected, showing a table with one entry:

Key	Last Modified	ETag	Size	Storage Class	Owner	Version Id
hfs.exe	3/24/2024 12:13:53 AM	"9e0557e98ed1269372ff0ace91d6347"	2.07 MB	STANDARD	tenant_demo (27041610751...	NJUSRDRhCNDlR...

In our ONTAP bucket, let's add a new version to our same object that we used previously and see how it replicates.

Bucket contents:

Name	Size	Type	Last Modified	Storage Class
cluster1_demo...	157 bytes	Text Document	3/23/2024 11:23:25 PM	STANDARD
cluster1_svm...	211 bytes	Text Document	3/23/2024 11:23:25 PM	STANDARD
cluster1_svm...	211 bytes	Text Document	3/23/2024 11:23:25 PM	STANDARD
putty.exe	834.05 KB	Application	3/23/2024 11:23:25 PM	STANDARD
hfs.exe	2.07 MB	Application	3/24/2024 12:14:52 AM	STANDARD
s3browser-11...	9.58 MB	Application	3/23/2024 11:23:26 PM	STANDARD

Version history for hfs.exe:

Key	Last Modified	ETag	Size	Storage Class	Owner	Version Id
cluster1_demo_s3_user_s3_user.txt	3/23/2024 11:23:25 PM	acf4c9543e97e0678b2b6ed5a60e1bc	157 bytes	STANDARD	Unknown (Unknown)	Mzg0MjQ1MDAwL...
cluster1_svm_demo_s3_details (1).txt	3/23/2024 11:23:25 PM	407753b646a6cfef19de71eefb5ff04	211 bytes	STANDARD	Unknown (Unknown)	NDg0MjQ1MDAw...
cluster1_svm_demo_s3_details.txt	3/23/2024 11:23:25 PM	17d20651856480a587a039ecc10e2	211 bytes	STANDARD	Unknown (Unknown)	NTU2NzI0MDAwL...
hfs.exe	revision # 3 (current)	9e8557e98ed1269372fface91d63477	2.07 MB	STANDARD	Unknown (Unknown)	NTY0NDg0MDAw...
revision # 2	3/23/2024 11:23:36 PM	9e8557e98ed1269372fface91d63477	2.07 MB	STANDARD	Unknown (Unknown)	NzQ1OTE4MDAw...
revision # 1	3/23/2024 11:23:25 PM	9e8557e98ed1269372fface91d63477	2.07 MB	STANDARD	Unknown (Unknown)	NjK2ODI0MDAwL...
putty.exe	revision # 1 (current)	54cb91395daad9d7882533c21fe0e9	834.05 KB	STANDARD	Unknown (Unknown)	NzE2NzEyMDAwL...
s3browser-11-6-7.exe	revision # 1 (current)	ae36fb97054782962d6597c5d082d-2	9.58 MB	STANDARD	Unknown (Unknown)	NDY2ODcwMDEu...

If we look on the StorageGRID side we see that a new version has been created in this bucket too, but is missing the initial version from before the snapmirror relationship.

Bucket contents:

Name	Size	Type	Last Modified	Storage Class
cluster1_demo...	157 bytes	Text Document	3/24/2024 12:13:53 AM	STANDARD
cluster1_svm...	211 bytes	Text Document	3/24/2024 12:13:53 AM	STANDARD
cluster1_svm...	211 bytes	Text Document	3/24/2024 12:13:53 AM	STANDARD
putty.exe	834.05 KB	Application	3/24/2024 12:14:28 AM	STANDARD
hfs.exe	2.07 MB	Application	3/24/2024 12:14:56 AM	STANDARD
s3browser-11...	9.58 MB	Application	3/24/2024 12:13:53 AM	STANDARD

Version history for hfs.exe:

Key	Last Modified	ETag	Size	Storage Class	Owner	Version Id
hfs.exe	revision # 2 (current)	9e8557e98ed1269372fface91d63477	2.07 MB	STANDARD	tenant_demo (27041610751...)	OE14RjYNDgIRT...
revision # 1	3/24/2024 12:13:53 AM	9e8557e98ed1269372fface91d63477	2.07 MB	STANDARD	tenant_demo (27041610751...)	NjU5RDhCNDI0T...

This is because the ONTAP SnapMirror S3 process only replicates the current version of the object. This is why we created a versioned bucket on the StorageGRID side to be the destination. This way StorageGRID can maintain a version history of the objects.

By Rafael Guedes, and Aron Klein


Enabling enterprise-grade S3 by seamlessly migrating object-based storage from ONTAP S3 to StorageGRID

Enabling enterprise-grade S3 by seamlessly migrating object-based storage from ONTAP S3 to StorageGRID


Migrate S3 Keys

For a migration, most of the time you will want to migrate the credentials for the users rather than generate new credentials on the destination side. StorageGRID provides api's to allow s3 keys to be imported to a user.

Logging into the StorageGRID management UI (not the tenant manager UI) open the API Documentation swagger page.



Expand the "accounts" section, select the "POST /grid/account-enable-s3-key-import", click the "Try it out" button, then click on the execute button.



Now scroll down still under "accounts" to "POST /grid/accounts/{id}/users/{user_id}/s3-access-keys"

Here is where we are going to input the tenant ID and user account ID we collected earlier. fill in the fields and the keys from our ONTAP user in the json box. you can set the expiration of the keys, or remove the ", "expires": 123456789" and click on execute.

POST`/grid/accounts/{id}/users/{user_id}/s3-access-keys` Imports S3 credentials for a given user in a tenant account**Parameters**

Name

Description

id * requiredstring
(path)

ID of Storage Tenant Account

27041610751165610501

user_id * requiredstring
(path)

ID of user in tenant account.

ebc132e2-cfc3-42c0-a445-3b4465cb523c

body * required[Edit Value](#) | Model

(body)

```
{  
  "accessKey": "3TVPI142JGE3Y7FV2KC0",  
  "secretAccessKey": "75a10qKBU4quA132twI4g41C4Gg5PP30ncy0sPE8"  
}
```

Once you have completed all of your user key imports you should disable the key import function in "accounts" "POST /grid/account-disable-s3-key-import"

POST`/grid/account-disable-s3-key-import` Disables the Import S3 Credentials feature on this node.[Cancel](#)**Parameters**

No parameters

[Execute](#)**Responses**

Response content type

application/json



If we look at the user account in the tenant manager UI, we can see the new key has been added.

Overview

Full name: ?	Demo S3 User 
Username: ?	demo_s3_user
User type: ?	Local
Denied access: ?	Yes
Access mode: ?	Read-only
Group membership: ?	Demo S3 Group

- [Password](#)
- [Access](#)
- [Access keys](#)
- [Groups](#)

Manage access keys

Add or delete access keys for this user.

Create key	Actions 
<input type="checkbox"/>	Access key ID 
<input type="checkbox"/>	*****86TR 
<input type="checkbox"/>	*****2KC0 

The final cut-over

If the intention is to have a perpetually replicating bucket from ONTAP to StorageGRID, you can end here. If this is a migration from ONTAP S3 to StorageGRID, then its time to put an end to it and cut over.

Inside ONTAP system manager, edit the S3 group and set it to "ReadOnlyAccess". This will prevent the users from writing to the ONTAP S3 bucket anymore.

Edit group

X

NAME

demo_s3_group

USERS

demo_s3_user x

POLICIES

ReadOnlyAccess x

Cancel

Save




All that is left to do is configure DNS to point from the ONTAP cluster to the StorageGRID endpoint. Make sure your endpoint certificate is correct and if you need virtual hosted style requests then add the endpoint domain names in storageGRID

Endpoint Domain Names

Virtual Hosted-Style Requests

Enable support of S3 virtual hosted-style requests by specifying API endpoint domain names. Support is disabled if this list is empty. Examples: s3.example.com, s3.example.co.uk, s3-east.example.com



Endpoint 1 s3portal.demo.netapp.com + **Save**

Your clients will either need to wait for the TTL to expire, or flush DNS to resolve to the new system so you can test that everything is working. All that is left is to clean up the initial temporary S3 keys we used to test the StorageGRID data access (NOT the imported keys), remove the snapmirror relationships, and remove the ONTAP data.

By Rafael Guedes, and Aron Klein

Copyright information

Copyright © 2025 NetApp, Inc. All Rights Reserved. Printed in the U.S. No part of this document covered by copyright may be reproduced in any form or by any means—graphic, electronic, or mechanical, including photocopying, recording, taping, or storage in an electronic retrieval system—without prior written permission of the copyright owner.

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

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

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

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

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

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

Trademark information

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