

Product feature guides

StorageGRID solutions and resources

NetApp July 14, 2025

This PDF was generated from https://docs.netapp.com/us-en/storagegrid-enable/product-feature-guides/achieve-zero-rpo.html on July 14, 2025. Always check docs.netapp.com for the latest.

Table of Contents

Product feature guides	1
Achieving zero RPO with StorageGRID - A Comprehensive Guide to Multi-Site Replication	1
StorageGRID Overview	1
How to get to Zero RPO with StorageGRID.	5
Synchronous Deployments across multiple sites	6
A Single Grid Multi-site deployment	6
A multi-site multi-grid deployment	9
Conclusion	11
Create Cloud Storage Pool for AWS or Google Cloud	11
Create Cloud Storage Pool for Azure Blob Storage	12
Use a Cloud Storage Pool for backup	13
Configure StorageGRID search integration service	14
Introduction	14
Create tenant and enable platform services	14
Search integration services with Amazon OpenSearch	15
Platform services endpoint configuration	19
Search integration services with on premises Elasticsearch	21
Platform services endpoint configuration	24
Bucket search integration service configuration	26
Where to find additional information	30
Node Clone	30
Node clone considerations	30
Node clone Performance estimates	30
How to use port remap	33
Migrate S3 clients from CLB to NGINX with Port ReMap	33
Remap port 443 for client S3 access on an Admin node	37
Restore Databases and logs	41
Grid site relocation and site-wide network change procedure	43
Considerations before site relocation	43
Migrating object-based storage from ONTAP S3 to StorageGRID	48
Enabling enterprise-grade S3 by seamlessly migrating object-based storage from ONTAP S3 to	
StorageGRID	48
Enabling enterprise-grade S3 by seamlessly migrating object-based storage from ONTAP S3 to	
StorageGRID	48
Enabling enterprise-grade S3 by seamlessly migrating object-based storage from ONTAP S3 to	
StorageGRID	60
Enabling enterprise-grade S3 by seamlessly migrating object-based storage from ONTAP S3 to	
StorageGRID	72
Enabling enterprise-grade S3 by seamlessly migrating object-based storage from ONTAP S3 to	
StorageGRID	81

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's 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 the 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

StorageGRID can be deployed with as few as 3 storage nodes and a single grid can grow up to 200 nodes. A single grid can be deployed as a single site or extend to 16 sites. A minimal grid consists of an admin node and 3 storage nodes in a single site. 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 StorageGRID 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.

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 replica 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 multiple levels. With StorageGRID appliances we erasure code the data stored on each node across all the drives with RAID providing protection against multiple disk failures causing data loss or interruptions. Additionally, StorageGRID can 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.

Erasure coding provides a storage layout that is resilient to node failures with low overhead, while replication can do the same thing, with more overhead. 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: Guarantees read-after-write consistency for all client requests across all sites.
- **Strong-global V2**: 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. 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 or 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 is successful to the client. The node selection will try 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 is 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 is 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 requests. This should only be configured if the total latency is acceptable to the clients.

How to get to 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, Strong Global V2 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.

Optionally, 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.

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 Strong Global v2 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 strong global v2 consistency model for grids of 3 or more sites and strong Global consistency for less than 3 sites.

In a two site StorageGRID solution there are at least two replicas or 3 EC chunks 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.



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	3 or more sites 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.	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	Disruption to client operations directed to this site but no data loss. Operations directed to the other site remain uninterrupted and no data loss.	Operations are directed to all other sites and remain uninterrupted and no data loss.

Failure	2-site Outcome	3 or more sites outcome
Single node failure at multiple sites	No disruption or data loss if:	No disruption or data loss if:
	 At least a single replica exists in the grid 	 At least a single replica exists in the grid
	 Sufficient EC chunks exist in the grid 	 Sufficient EC chunks exist in the grid
	Operations disrupted and risk of data loss if:	Operations disrupted and risk of data loss if:
	 No replicas exist 	 No replicas exist
	 Insufficient EC chucks exist 	 Insufficient EC chucks exist to retrieve the object
Single site failure	client operations will be interrupted until either the failure is resolved, or the bucket consistency is lowered to strong site or lower to allow operations to succeed but no data loss.	No interruption to operations or data loss.
Single site plus single node failures	client operations will be interrupted until either the failure is resolved, or the bucket consistency is lowered to read-after-new-write or lower to allow operations to succeed and possible data loss.	No interruption to operations or data loss.
Single site plus a node from each remaining site	client operations will be interrupted until either the failure is resolved, or the bucket consistency is lowered to read-after-new-write or lower to allow operations to succeed and possible data loss.	Operations will be disrupted If metadata replica quorum cannot be met and possible data loss.
Multi-site failure	No operations sites remain data will be lost if at least 1 site cannot be recovered in its entirety.	Operations will be disrupted If metadata replica quorum cannot be met. No data loss as long as at least 1 site remains.
Network isolation of a site	client operations will be interrupted until either the failure is resolved, or the bucket consistency is lowered to strong site or lower to allow operations to succeed, but no data loss	Operations will be disrupted for the isolated site, but no data loss No disruption to operations in the remaining sites and no data loss

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 strong global v2 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.

Failure	Outcome
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.
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 identify 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:: nnnnn: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.n/nn"
                   ]
               }
        },
       "Resource": "arn:aws:es:us-east-1:nnnnnn:domain/sgdemo/*"
       }
   ]
}
```

ine-grain	and access control	
me-gran		
ne-grained a	cess control provides numerous features to help you keep your data secure. Features i	include document-level security, field-
ves security, r	tad-only users, and Upensearch Dashboards/Albana tenants. Pine-grained access com	troi requires a master user. Learn more
2		
Enable fi	ne-grained access control	
AML aut	hentication for OpenSearch Dashboards/Kibana	
AML authenti	ation lets you use your existing identity provider for single sign-on for OpenSearch D	ashboards/Kibana. Learn more 🖸
Prepare 5	AML authentication	
0.7	CALF - shortlestic	
10 10 US	2 SAME authentication, you must first enable fine-grained access control.	
	and a state of the	
mazon (ognito authentication	
able to use A	mazon Cognito authentication for OpenSearch Dashboards/Kibana. Amazon Cognito	supports a variety of identity provide
r üsername-ş	assword authentication. Learn more	
Enable A	mazon Cognite authentication	
Enable A	nazon cognito automication	
	Ren.	
access po	ucy	
ccess policies	control whether a request is accepted or rejected when it reaches the Amazon OpenSe	earch Service domain. If you specify a
count, user, o	r role in this policy, you must sign your requests. Learn more 🖸	
omain acce	is policy	
(Call	Res and a second s	
Allow one	nne-grained access control	
Allow ope	access to the domain.	
Do not se	t domain level access policy	
All request	s to the domain will be denied.	
Continue	a descente la colorada na lista	
Conngun	domain level access policy	
Visual ed	ISON	Import policy
visual eq	lor JSON	import poincy
start palls		
ccess pouc	/	
3+ 15	tatement": [
4 -		
5	"Effect": "Allow",	
6.°	TAUST: Tana and and an and a second second second	
8	have a second seco	
9	"Action": "es:*",	
18	"Resource": "arm:aws:es:us-east-1:2100000000000000000000000000000000000	
11		
12.*	"iffect": "Allow".	
14 *	"Principal": {	
15	"AuS": "**	
16	1,	
17 -	Action': [
18	as is smith	
	"Production" + 1	
28 -	CONDECTOR 2 C	
28 × 21 =	"IpAddress": {	
28 - 21 - 22 -	"lpAddress": { "aws:Sourcelp": {	
20 * 21 * 22 * 23 24	"lpAddress": { "aws:Sourcelp": { "216/24"	
20 * 21 * 22 * 23 24 25	"IpAddress": { "aws:Sourcelp": { "216	
28 - 21 - 22 - 23 24 25 26	"IpAddress": { "aws:Sourcelp": { "216	
28 - 21 - 22 - 23 24 25 26 27	"IpAddress": { "aws:Sourcelp": { "216.200000000000000000000000000000000000	

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

Amazon OpenSearch Service > Domains > sgdemo			
sgdemo anto			Delete Actions V
General information			
Name sgdemo Da	Domain status	Version Into OpenSearch 1.1 (latest)	OpenSearch Dashboards URL https://search-sgdemo-
Domain ARN	Cluster health Into	Service software version Info	1.es.amazonaws.com/_dashboards 🖸
arrcavesescus-east-1:3444444444444444444444444444444444444	Yellow	R20211203-P5 (latest)	https://search-sgleario-testinglishing and an //Lus-cost- 1 estamazanaws.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.
- 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 for storing StorageGRID object metadata. We 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.

Some and the second sec	
■ Dev Tools	
Console	
History Settings Help	
1 PUT sgmetadata D ≥	<pre>1* { 2 "acknowledged" : true, 3 "shards_acknowledged" : true, 4 "index" : "sgmetadata" 5* }</pre>

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

	2112/09/09			
Anazon OpenSearch Service 🗧 Domaine 🗦 ogdi	me			
sgdemo 🛶				Delete Actions v
General information				
Nome egdeno Domula AIN S arrawses is cart 12	Domosis status Statice Cluster health Yellow	aria	Version Iule OperSearch 1.1 (latent) Service software variation (me 920211203-05 (latent)	OpenSourch Deseboards URL https://www.hougdemo-d- east-tues.amazonaes.com/_datafocards.@ Domain endpoint https://www.houghemo-d- mad-1.ex.amazonaes.com 23
	en Cluster health	Instance Insultin Auto-Tune I	Logs Tags Connections Pa	ckages Notifications
Cluster configuration Security configuration	en. Sources constants			
Cluster configuration Security configuration	ige data for fact retrievel. Before	rybu can selarat data, yeu musi inetx it, usa	m more 😢	< 1 > @
Cluster configuration Security configuration Indices (2) Indices (2) Indices (2) Indices (2) Q Find half on Index Document conditions	nge sata for tast resnevel, Befor nit. w. Size (byte)	r pisu can search data, veu must innen it. Lea P Query total	m more 🕑	< 1 > 0
Cluster configuration Security coeffiguration Indices (2) Indices (2) Induces (2) Induces Index	nge sata for tait retrieval. Safer ont w Size (byte) 1	rybu can seleren data, yeu must index it, uta v Querry fotal 5.08 K/B	m more C • Mapping type Field mappings 10 dynamic_meta.oroper	< 1 > 0 v 0

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	
Enter details 2 Select authentication type Optional	③ Verify server Optional
Enter endpoint details	
Enter the endpoint's display name, URI, and URN.	
Display name 🔹	
aws-opensearch	
URI 😢	
https://search-sgdemo-/####################################	
URN Ø	
s:es:us-east-1:200010500000;domain/sgdemo/sgmetadata/_doc	
	Cancel Continue

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	
Enter details 2 Select authentication type	Verify server Optional
Authentication type 🔞	
Select the method used to authenticate connections to the endpoint.	
Access Key 🗸	
Access key ID 😮	
AKIA	
Secret access key 😨	
	Previous

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.

Pla	atforr	าร	servi	CE	es e	nc	points			
A platfon configure	m services end e an endpoint f	oint st ir each	ores the info platform ser	rmati vice y	on Storage ou plan to	eGRID n use.	eeds to use an external resource a	as a target for a platform servi	ce (CloudMirror replication, notifications, or s	search integration). You mu
1 endpoi	int									Create endpoint
Defete e	endpoint									
	Display name	•	Last error	•	Туре	¢ U	IRI Ø ≑		urn 😧 ≑	

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.



- 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.

🍪 elastic	Q Search Elastic	Q Search Elastic			
Dev Tools Console					
Console Search Profiler Grok	Debugger Painless Lab				
History Settings Help 1 PUT sgmetadata	► ئ ى	1* {			
		<pre>2 "acknowledged" : true, 3 "shards_acknowledged" : true, 4 "index" : "sgmetadata" 5 ^ }</pre>			
		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
 - o URN: urn:<something>:es:::<some-unique-text>/<index-name>/_doc where the indexname is the name you used on the Kibana console. Example: urn:local:es:::sgmd/sgmetadata/_doc

Create endpoint			
1 Enter details 2 Select authentication typ Optional	e 3	Verify server Optional	
Enter endpoint details			
Enter the endpoint's display name, URI, and URN.			
Display name 🥥			
elasticsearch			
https://10.0000000000000000000000000000000000			
URN 0			
urn:local:es:::sgmd/sgmetadata/_doc			
		Cancel Cont	inue

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.

Select the method used to authenticate connections to the endpoint.
Basic HTTP V
Username 🥝
elastic
Password ()

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.

Pla	atform servic	es en	dpoints	
A platfor configure	m services endpoint stores the inform e an endpoint for each platform servic	ation StorageGRI e you plan to use	D needs to use an external resource as a target for a platform service .	(CloudMirror replication, notifications, or search integration). You must
2 endpoi	nts			Create endpoint
Deleter	redpoint			
	Display name Last error	о ^{Тури} о	URI 🔮 🗢	URN 😌 🗢
	aws-opensearch	Search	https://search-sgdemo-fw223hpljv6lzcxrpw3v3rte7i.us-east- 1.es.amazonaws.com/	arrcaws:es:us-east- 1:210811600158:domain/sgdemo/sgmetadata/_doc
	elasticsearch	Search	https://10.	um: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
- Click Create Bucket, enter the bucket name (for example, sgmetadata-test) and accept the default useast-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>

</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.

New bucket	Path: /				1076
i sgmetadata-test	File	Size	Туре	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.bd	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
	Upload • 📄 Download	Delete 🕞 Ni	ew Folder	2	1 file (762,53 kB) selecte
Tasks (14) Permissions URL: https://10.193.204 Key	Upload • Download Http Headers Tags Propert 5.106:10445/sgmetadata-test/Koo Value	Delete Review Version	ew Folder Refresh is EventLog		1 file (762, 53 K8) selecte
Tasks (14) Permissions URL: https://10.193.204 Key date	Upload • Download Http Headers Tags Propert 1.106:10445/sgmetadata-test/Kor Value 01-01-2020	Delete 🥁 Na ies Preview Version sla.jpg	ew Folder Refresh Is EventLog		1 file (762.53168) selecte
Tasks (14) Permissions URL: https://10.193.204 Key date owner	Upload • Download Http Headers Tags Propert 1.106:10445/sgmetadata-test/Kor Value 01-01-2020 testuser	Delete Delete Na ies Preview Version ala.jpg	ew Folder Refresh Is EventLog		1 file (762.53 KB) selecte
Tasks (14) Permissions URL: https://10.193.204 Key date owner project	Upload • Download Http Headers Tags Propert 5.106:10445/sgmetadata-test/Kor Value 01-01-2020 testuser test	Delete Ra Na Na Preview Version	ew Folder Refresh is EventLog		1 file (762,53 k8) selecte

- 7. Use the Kibana UI to verify that the object metadata was loaded to sgmetadata's index.
 - a. From the menu, select Management > Dev Tools.
 - b. 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": { }
}
}
```



The query 2 sample result in the following screenshot shows two records with tag 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
 - $\,\circ\,$ 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 bellow: 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

$\textbf{SG5612/SG5712/SG5812} \rightarrow \textbf{Any}$

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

$\text{SG5660} \rightarrow \text{SG5760/SG5860}$

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

$\text{SG5660} \rightarrow \text{SG6060/SG6160}$

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

$\text{SG5760/SG5860} \rightarrow \text{SG5760/SG5860}$

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

$\textbf{SG5760/SG5860} \rightarrow \textbf{SG6060/SG6160}$

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

$\textbf{SG6060/SG6160} \rightarrow \textbf{SG6060/SG6160}$

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

$\text{SG5760/SG5860} \rightarrow \text{SG5760/SG5860}$

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

$\textbf{SG5760/SG5860} \rightarrow \textbf{SG6060/SG6160}$

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

$\text{SG6060/SG6160} \rightarrow \text{SG6060/SG6160}$

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

How to use port remap

You may have a need to remap an incoming or outbound port for multiple reasons. You may be moving from the legacy CLB load balancer service to the current nginx service load balancer endpoint and maintain the same port to reduce the impact to clients, wish to use port 443 for client S3 on an admin node client network, or for firewall restrictions.

Migrate S3 clients from CLB to NGINX with Port ReMap

In releases earlier than StorageGRID 11.3, the included Load Balancer service on the Gateway Nodes is the Connection Load Balancer (CLB). In StorageGRID 11.3, NetApp introduces the NGINX service as a feature rich integrated solution for load balancing HTTP(s) traffic. Because the CLB service remains available in the current release of StorageGRID, you cannot reuse port 8082 in the new load balancer endpoint configuration. To work around this, the 8082 inbound port is remapped to 10443. This makes all HTTPS requests coming into port 8082 on the gateway redirect to port 10443, bypassing the CLB service and instead connecting to the NGINX service. Although the following instructions are for VMware, the PORT_REMAP functionality exists for all installation methods, and you can use a similar process for bare metal deployments and appliances.

VMware virtual machine Gateway Node deployment

The following steps are for a StorageGRID deployment where the Gateway Node or Nodes are deployed in VMware vSphere 7 as VMs using the StorageGRID Open Virtualization Format (OVF). The process entails destructively removing the VM and redeploying the VM with the same name and configuration. Before you power on the VM, change the vAPP property to remap the port, then power on the VM and follow the node recovery process.

Prerequisites

- You are running StorageGRID 11.3 or later
- You have downloaded and have access to the installed StorageGRID version VMware install files.
- You have a vCenter account with permissions to power on/off VMs, change the settings of the VMs and vApps, remove VMs from vCenter, and deploy VMs by OVF.
- · You have created a load balancer endpoint
 - · The port is configured to the desired redirect port
 - The endpoint SSL certificate is the same as installed for the CLB service in the Configuration/Server Certificates/ Object Storage API Service Endpoints Server Certificate or the client is able to accept a change in certificate.



If your existing certificate is self-signed, you cannot reuse it in the new endpoint. You must generate a new self-signed certificate when creating the endpoint and configure the clients to accept the new certificate.

Destroy the first Gateway Node

To destroy the first Gateway Node, follow these steps:

1. Choose the Gateway Node to start with if the grid contains more than one.

- 2. Remove the node IPs from all DNS round-robin entities or load balancer pools, if applicable.
- 3. Wait for Time-to-Live (TTL) and open sessions to expire.
- 4. Power off the VM node.
- 5. Remove the VM node from the disk.

Deploy the replacement Gateway Node

To deploy the replacement Gateway Node, follow these steps:

- 1. Deploy the new VM from OVF, selecting the .ovf, .mf, and .vmdk files from the install package downloaded from the support site:
 - vsphere-gateway.mf
 - vsphere-gateway.ovf
 - NetApp-SG-11.4.0-20200721.1338.d3969b3.vmdk
- 2. After the VM has been deployed, select it from the list of VMs, select the Configure tab vApp Options.

Summary Monitor	Configure Permissions Datastores	Networks Snapshots Updates
Settings VM SDRS Rules	> Deployment	
vApp Options		
Alarm Definitions	OVF Settings View OVF ENVIRONMENT	í
Scheduled Tasks	OVF environment transport	VMware Tools
Policies Guest User Mappings	Installation boot	Disabled
	Properties	
	ADD EDIT SET VALUE DELETE	

3. Scroll down to the Properties section and select the PORT_REMAP_INBOUND property

Summary Monitor	Cor	nfigure Permissi	ons Datastores	Networks Snap	shots Updates		
Settings 🗸 🗸	0	ADMIN_IP	Primary Admin IP	10.193.204.110	0.0.0	Grid Network (eth0)	ip
VM SDRS Rules vApp Options	0	ADMIN_NETWO RK_ESL	Admin network ext ernal subnet list			Admin Networ k (eth1)	string
Alarm Definitions Scheduled Tasks	0	ADMIN_NETWO RK_IP	Admin network IP	10.193.174.112	0.0.0.0	Admin Networ k (eth1)	ip
Policies Guest User Mappings	0	NODE_TYPE	Node type		VM_API_Gate way	Grid Node Par ameters	string["VM_Storage_Node", "VM_ min_Node", "VM_API_Gateway", _Archive_Node"]
	0	CLIENT_NETWO RK_CONFIG	Client network IP c onfiguration	STATIC	DISABLED	Client Networ k (eth2)	string["DISABLED", "STATIC", "DH P"]
	۰	PORT_REMAP_I NBOUND	Inbound port rema pping specification			Advanced	string
	0	GRID_NETWORK	Grid network IP co	STATIC	STATIC	Grid Network	string["STATIC", "DHCP"]

4. Scroll to the top of the Properties list and click Edit


5. Select the Type tab, confirm that the User Configurable checkbox is selected, and then click Save.

General Type				
Static property				
Туре	String		4	
User configurable				
Length	0	0	65535	0
Default value			 S.	
) Dynamic property				
Macro	IP address			
Network	MGMT_564		×.	

6. At the top of the Properties list, with the "PORT_REMAP_INBOUND" property still selected, click Set Value.

Propertie	S		
ADD	EDIT	SET VALUE	DELETE

7. In the Property Value field, enter the network (grid, admin, or client), TCP, the original port (8082), and the new port (10443) with "/" in between each value as depicted following.

Set value	Inbound port remapping specification	×
Property value	grid/tcp/8082/10443	
	CANCEL	ок

8. If you are using multiple networks, use a comma (,) to separate the network strings, for example, grid/tcp/8082/10443,admin/tcp/8082/10443,client/tcp/8082/10443

Recover the Gateway Node

To recover the Gateway Node, follow these steps:

1. Navigate to the Maintenance/Recovery section of the Grid Management UI.

	Maintenance -	Support -	
	Maintenance Tasks	Network	System
	Expansion	Grid Network	Software Update
ery	Decommission	DNS Servers	License
	Recovery	NTP Servers	Recovery Package

2. Power on the VM node and wait for the node to appear in the Maintenance/Recovery Pending Nodes section of the Grid Management UI.

Recovery

Select the failed grid node to recover, enter your provisioning passphrase, and then click Start Recovery to begin the recovery procedure.

ending Nodes			
Name	IPv4 Address	11 State	1 Recoverable
For Fortunal IOGRA.			
i	For information and directi https://docs.netapp.com/sgw maint/GUID-7E22B1B9-4169-48 and Maintenance guide]	ons for node recover s-114/topic/com.neta 00-8727-75F25FC0FFB1	y, see the pp.doc.sg- .html[Recovery

3. After the node has been recovered, the IP can be included in all DNS round-robin entities, or load balancer pools, if applicable.

Now, any HTTPS sessions on port 8082 go to port 10443

Remap port 443 for client S3 access on an Admin node

The default configuration in the StorageGRID system for an admin node, or HA group containing an Admin node is for port 443 and 80 to be reserved for the management and tenant manager UI's and cannot be used for load balancer endpoints. The solution to this is to use the port remap feature and redirect inbound port 443 to a new port that will be configured as a load balancer endpoint. Once this completed Client S3 traffic will be able to use port 443, the Grid management UI will only be accessible through port 8443, and the Tenant management UI will only be accessible on port 9443. The remap port feature can only be configured at install time of the node. In order to implement a port remap of an active node in the grid, it must be reset to the pre-installed state. This is a destructive procedure that includes a node recovery once the configuration change has been made.

Backup logs and databases

Admin nodes contain audit logs, prometheus metrics, as well as historical information about attributes, alarms, and alerts. Having multiple admin nodes means you have multiple copies of this data. If you do not have multiple admin nodes in your grid, you should make sure to preserve this data to restore after the node has been recovered in the end of this process. If you have another admin node in your grid, you can copy the data from that node during the recovery process. If you do not have another admin node in the grid you can follow these instructions to copy the data before destroying the node.

Copy audit logs

- 1. Log in to the Admin Node:
 - a. Enter the following command: ssh admin@grid_node_IP
 - b. Enter the password listed in the Passwords.txt file.
 - c. Enter the following command to switch to root: su -

- d. Enter the password listed in the Passwords.txt file.
- e. Add the SSH private key to the SSH agent. Enter: ssh-add
- f. Enter the SSH Access Password listed in the Passwords.txt file.

When you are logged in as root, the prompt changes from `\$` to `#`.

2. Create the directory to copy all audit log files to a temporary location on a separate grid node lets use storage_node_01:

```
a.ssh admin@storage_node_01_IP
```

```
b. mkdir -p /var/local/tmp/saved-audit-logs
```

- 3. Back on the admin node, stop the AMS service to prevent it from creating a new log file: service ams stop
- 4. Rename the audit.log file so that it does not overwrite the existing file when you copy it to the recovered Admin Node.
 - a. Rename audit.log to a unique numbered file name such as yyyy-mm-dd.txt.1. For example, you can rename the audit log file to 2015-10-25.txt.1

```
cd /var/local/audit/export
ls -1
mv audit.log 2015-10-25.txt.1
```

- 5. Restart the AMS service: service ams start
- 6. Copy all audit log files: scp * admin@storage_node_01_IP:/var/local/tmp/saved-auditlogs

Copy Prometheus data



Copying the Prometheus database might take an hour or more. Some Grid Manager features will be unavailable while services are stopped on the Admin Node.

- 1. Create the directory to copy the prometheus data to a temporary location on a separate grid node, again we will user *storage_node_01*:
 - a. Log in to the storage node:
 - i. Enter the following command: ssh admin@storage_node_01_IP
 - ii. Enter the password listed in the Passwords.txt file.
 - iii. mkdir -p /var/local/tmp/prometheus`
- 2. Log in to the Admin Node:
 - a. Enter the following command: ssh admin@admin node IP
 - b. Enter the password listed in the Passwords.txt file.
 - c. Enter the following command to switch to root: su -

- d. Enter the password listed in the Passwords.txt file.
- e. Add the SSH private key to the SSH agent. Enter: ssh-add
- f. Enter the SSH Access Password listed in the Passwords.txt file.

When you are logged in as root, the prompt changes from `\$` to `#`.

- 3. From the Admin Node, stop the Prometheus service: service prometheus stop
 - a. Copy the Prometheus database from the source Admin Node to the storage node backup location Node: /rsync -azh --stats "/var/local/mysql_ibdata/prometheus/data" "storage_node_01_IP:/var/local/tmp/prometheus/"
- 4. Restart the Prometheus service on the source Admin Node.service prometheus start

Backup historical information

The historical information is stored in a mysql database. In order to dump a copy of the database you will need the user and password from NetApp. If you have another admin node in the grid, this step is not necessary and the database can be cloned from a remaining admin node during the recovery process.

- 1. Log in to the Admin Node:
 - a. Enter the following command: ssh admin@admin_node_IP
 - b. Enter the password listed in the Passwords.txt file.
 - c. Enter the following command to switch to root: su =
 - d. Enter the password listed in the Passwords.txt file.
 - e. Add the SSH private key to the SSH agent. Enter: ssh-add
 - f. Enter the SSH Access Password listed in the Passwords.txt file.

When you are logged in as root, the prompt changes from `\$` to `#`.

- 2. Stop StorageGRID services on Admin Node and startup ntp and mysql
 - a. Stop all services: service servermanager stop
 - b. restart ntp service: service ntp start ..restart mysql service: service mysql start
- 3. Dump mi database to /var/local/tmp
 - a. enter the following command: mysqldump -u username -p password mi > /var/local/tmp/mysql-mi.sql
- 4. Copy the mysql dump file to an alternate node, we will use *storage_node_01:* scp /var/local/tmp/mysql-mi.sql storage node 01 IP:/var/local/tmp/mysql-mi.sql
 - a. When you no longer require passwordless access to other servers, remove the private key from the SSH agent. Enter: ssh-add -D

Rebuild the Admin node

Now that you have a backup copy of all desired data and logs either on another admin node in the grid or stored in a temporary location it is time to reset the appliance so the port remap can be configured.

- 1. Resetting an appliance returns it to the pre-installed state where it only retains the host name, IP's and network configurations. All data will be lost which is why we made sure to have a backup of any important information.
 - a. enter the following command: sgareinstall

```
root@sg100-01:~ # sgareinstall
WARNING: All StorageGRID Webscale services on this node will be shut
down.
WARNING: Data stored on this node may be lost.
WARNING: You will have to reinstall StorageGRID Webscale to this
node.
After running this command and waiting a few minutes for the node to
reboot,
browse to one of the following URLs to reinstall StorageGRID Webscale
on
this node:
    https://10.193.174.192:8443
    https://10.193.204.192:8443
    https://169.254.0.1:8443
Are you sure you want to continue (y/n)? y
Renaming SG installation flag file.
Initiating a reboot to trigger the StorageGRID Webscale appliance
installation wizard.
```

- 2. After some time has passed the appliance will reboot and you will be able to access the node PGE UI.
- 3. Browse to the Configure Networking



4. Select the desired network, protocol, direction and ports then click the Add Rule button.



Remap of inbound port 443 on on the GRID network will break install, and expansion procedures. It is not recommended to remap port 443 on the GRID network.

otAp	p* StorageGRID* /	Appliance In	staller												
Home	e Configure Networ	king - Co	infigure Ha	rdware -	Monitor I	Installation	Advanced +								
Rem If req	lap Ports uired, you can remap the	internal ports o	on the appli	ance Stora	ge Node to di	fferent exter	mal ports. For exam	nple, you migh	t nee	d to remap por	ts because o	f a firewall is	sue		
.9	Femilye Selected Rule	+ Add Rule	Network	Grid y	Protocol	TCP v	Remap Direction	Inbound	÷	Original Port	1	101	Mapped-To Port	1	0
3	Network	+ Add Rule Protocol	Network	Grid y	Protocol temap Direct	TCP v	Remap Direction	Inbound Orig	y Jinal	Original Port	t.	C Mapp	Mapped-To Port	1	0

5. One the desired port remaps have been added, you can return to the home tab and click on the Start Installation button.

You can now follow the Admin node recovery procedures in the product documentation

Restore Databases and logs

Now that the admin node has been recovered, you can restore the metrics, logs, and historical information. If you have another admin node in the grid, follow the product documentation utilizing the *prometheus-clone-db.sh* and *mi-clone-db.sh* scripts. If this is your only admin node and you chose to backup this data, you can follow the below steps to restore the information.

Copy audit logs back

- 1. Log in to the Admin Node:
 - a. Enter the following command: ssh admin@grid node IP
 - b. Enter the password listed in the Passwords.txt file.
 - c. Enter the following command to switch to root: su -
 - d. Enter the password listed in the Passwords.txt file.
 - e. Add the SSH private key to the SSH agent. Enter: ssh-add
 - f. Enter the SSH Access Password listed in the Passwords.txt file.

When you are logged in as root, the prompt changes from `\$` to `#`.

- 2. Copy the preserved audit log files to the recovered Admin Node: scp admin@ grid node IP:/var/local/tmp/saved-audit-logs/YYYY* .
- 3. For security, delete the audit logs from the failed grid node after verifying that they have been copied successfully to the recovered Admin Node.
- 4. Update the user and group settings of the audit log files on the recovered Admin Node: chown amsuser:bycast *

You must also restore any pre-existing client access to the audit share. For more information, see the instructions for administering StorageGRID.

Restore Prometheus metrics



Copying the Prometheus database might take an hour or more. Some Grid Manager features will be unavailable while services are stopped on the Admin Node.

- 1. Log in to the Admin Node:
 - a. Enter the following command: ssh admin@grid_node_IP
 - b. Enter the password listed in the Passwords.txt file.
 - c. Enter the following command to switch to root: su -
 - d. Enter the password listed in the Passwords.txt file.
 - e. Add the SSH private key to the SSH agent. Enter: ssh-add
 - f. Enter the SSH Access Password listed in the <code>Passwords.txt</code> file.

When you are logged in as root, the prompt changes from `\$` to `#`.

- 2. From the Admin Node, stop the Prometheus service: service prometheus stop
 - a. Copy the Prometheus database from the temporary backup location to the admin node: /rsync -azh --stats "backup_node:/var/local/tmp/prometheus/" "/var/local/mysql ibdata/prometheus/"
 - b. verify the data is in the correct path and is complete ls /var/local/mysql ibdata/prometheus/data/
- 3. Restart the Prometheus service on the source Admin Node.service prometheus start

Restore historical information

- 1. Log in to the Admin Node:
 - a. Enter the following command: ssh admin@grid_node_IP
 - b. Enter the password listed in the Passwords.txt file.
 - c. Enter the following command to switch to root: su -
 - d. Enter the password listed in the Passwords.txt file.
 - e. Add the SSH private key to the SSH agent. Enter: ssh-add
 - f. Enter the SSH Access Password listed in the Passwords.txt file.

When you are logged in as root, the prompt changes from `\$` to `#`.

2. Copy the mysql dump file from the alternate node: scp grid_node_IP_:/var/local/tmp/mysql-mi.sql /var/local/tmp/mysql-mi.sql

- 3. Stop StorageGRID services on Admin Node and startup ntp and mysql
 - a. Stop all services: service servermanager stop
 - b. restart ntp service: service ntp start ..restart mysql service: service mysql start
- 4. Drop the mi database and create a new empty database: mysql -u username -p password -A mi -e "drop database mi; create database mi;"
- 5. restore the mysql database from the database dump: mysql -u username -p password -A mi < /var/local/tmp/mysql-mi.sql
- 6. Restart all other services service servermanager start

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.

	2024-04-24 19:49:00	1	
1.74 day	- 200 create_multipart_upload:		
	- 200 delete_bucket:		
1.16 day	200 delete_object:	12.9 K	
	200 delete_objects:	150	
13.9 hour	200 get_bucket_policy:	1	_
	= 200 get_object:	10.2 Mil	
ums	200 head_object:	1.52 Mil	
	= 200 list_buckets:	258 K	
sg6060-tme-04 delete_bucket	200 list_objects:	339 K	0-tm
🗕 sg6060-tme-05 get_usage 👝 sg6060-tme-05 head_object 👝 sg6060-tme-05 options 🥏 sg6060-tme-05 put_bucket 👝 sg6060-tme-05 put_object 🦲	= 200 options:	10	he-06
🗕 sg6060-tme-06 options 🗕 sg6060-tme-06 put_bucket 🗕 sg6060-tme-06 put_bucket_policy 🛥 sg6060-tme-06 put_object	200 other:	135 K	
	200 put_bucket:	- 3	
	= 200 put_bucket_tagging:	10791	
	= 200 upload part:	27.0 K	
4 Mil	= 204 put bucket policy:	1	
	- 400 get object:	8	
S MIL	- 400 list buckets:	108	
2 Mil	- 400 list objects:	16	
1 Mil	- 400 other:	8	
	- 400 put_bucket_policy:	4	
	- 403 delete_object:	5.72 K	10
19:19 19:20 19:21 19:22 19:23 19:24 19:25 19:20 19:27 19:26 19:29 19:30 19:31 19:32 19:33	- 403 get_object:	1.08 K	15
👝 sg6060-tme-04 delete_bucket 👝 sg6060-tme-04 delete_object 👝 sg6060-tme-04 get_object 👝 sg6060-tme-04 head_object 👝 sg6060-tme-04 options	- 403 head_bucket:	8	0-tm
😑 sg6060-tme-05 get_usage 👝 sg6060-tme-05 head_object 👝 sg6060-tme-05 options 🔤 sg6060-tme-05 put_bucket 👝 sg6060-tme-05 put_object 🛁	- 403 list_objects:	516	ne-06
🕳 sq6060-tme-06 options 📥 sq6060-tme-06 put_bucket 📥 sq6060-tme-06 put_bucket_policy 🕳 sq6060-tme-06 put_object	 404 get_bucket_policy: 		
	- 404 head_object:	1.05 K	
Table Completed Descents	- 404 other:	7.20 K	
iotal Completed Requests V	- 405 other:	242	
15 Mil	- 409 put_bucket:		
	- 411 other:	24	
10 Mil	6		
	5.50		
5 Mit	5		
	100		
	4.50		
19:20 19:22 19:24 19:26 19:26 19:30 19:32 19:34 19:36 19:38 19:40 19:44 19:46 19:48	4		
🗕 200 complete_multipart_upload 🗕 200 create_multipart_upload 🗕 200 delete_bucket 🗕 200 delete_object 🗕 200 delete_objects	3.50		
🗕 200 get_bucket_policy 😐 200 get_object 😑 200 head_object 😑 200 list_buckets 📥 200 list_objects 📥 200 options 💻 200 other	19:20 19:22 1	9:24	19:26
👝 200 put_bucket 🔤 200 put_bucket_tagging 🧫 200 put_object 💼 200 upload_part 🖨 204 put_bucket_policy 🛖 400 get_object 🕳 400 list_buckets	get object		

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]: 10.45.74.1]: LONDON-S1 Grid Gateway [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 1: OXFORD-S1 Grid Gateway [10.45.75.1]: 10.45.75.1]: OXFORD-S2 Grid Gateway [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 SELECT NODES to edit 1: 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 SAVE changes, so you can resume later 8: CLEAR all changes, to start fresh 9: 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-52
 LONDON-53
The following nodes will also require restarting:
 LONDON-ADM1
  LONDON-S1
 LONDON-52
 LONDON-53
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

```
PuTTY 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-53
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]>
```

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.
```

- From Grid Manager, download the new recovery package. Grid manager > Maintenance > Recovery package
- 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

- 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-networkip>: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 'shutdoown -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.

ONTAP System Manager								
DASHBOARD	cluster1 Version 9.14.1P1							
INSIGHTS STORAGE	Health \rightarrow \bigcirc Cluster is healthy							
Volumes LUNs Consistency groups NVMe namespaces Shares Buckets	SIMBOX							
Qtrees Quotas Storage VMs Tiers								

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

svm_demo

Access protocol

⊘ SMB/CIFS, NFS, S3	iSCSI FC	NVMe			
Enable SMB/CIFS					
Enable NFS					
Enable S3					
S3 SERVER NA	MF				
s3portal	.demo.netapp.con	n			
	LS				
ſ	443				
c					
Ċ	Use system-gener	ated certificate	D		
C) Use external-CA si	igned certificate			
	P (non-secure)				
Pr L					
l	8080				
DEFAULT LANGUAGE			-		
c.utf_8		~	J		
NETWORK INTERFA	CE				
Use multiple netwo	rk interfaces when	a client traffic is	high.		
onPrem-01					
IP ADDRESS		sk	GATEWAY	BROADCAST DOMAIN AND PORT	
192.168.0.200	24		Add optional gateway	Default	~
Storage VM adn	ninistration				
Enable maximum capaci The maximum capacity that	ty limit all volumes in this storag	je VM can allocate. Le	earn More [격		
Manage administrator a	count				
Can	.ei				

×

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

Added storage VM	>
storage vm svm_demo	s3 server NAME s3portal.demo.netapp.com
User details USER NAME sm_s3_user	
The secret key won't be displaye ACCESS KEY	ed again. Save this key for future use.
34EH21411SMW1YOV3NQY	þ
secret key Show secret key	
	Download Close

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

Services				
NIS Not configured	\$	©	Name service switch Services lookup order (i) HOSTS Files, then DNS GROUP Files NAME MAP Files NETGROUP	\$
DNS Not configured	Ļ			

51



Create SVM S3 User

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

Protocols				
NFS Not configured	<i>\$</i> ₽	SMB/CIFS Not configured	☆ ۞	iS N
NVMe Not configured	A Ø	S3 status Status Enabled TLS Disabled HTTP Enabled	*	

Add a new user.

Storage VMs			
+ Add : More			
Vame Name	S3 All settings		
✓ svm_demo	Enabled		
	Server FQDN s3portal.demo.netapp.com TLS Disabled HTTP Enabled Users Groups Policies	✓ Edit TLS PORT 443 HTTP PORT 8080	
	+ Add	Access key	Key expiration time
	root		-
	sm_s3_user	34EH21411SMW1YOV3NQY	Valid forever

Input the user name and key expiration.

Storage VMs			
+ Add : More			
Name	S3 All settings		
✓ svm_demo	Enabled		
	Server FGDN s3portal.demo.netapp.com TLS Disabled HTTP Enabled Users Groups Policies	<pre> Edit TLS PORT 443 HTTP PORT 8080 S</pre>	
	+ Add User name	Access key	Key expiration time
	root		•
	sm_s3_user	34EH21411SMW1YOV3NQY	Valid forever

Download the S3 keys for the new user.

S 3	All settings		
	Added user		×
	USER NAME demo_s3_user		
S	ACCESS KEY		
FC S	3TVPI142JGE3Y7FV2KC0	ļ	ð
D	SECRET KEY		
нт Ei	75a1QqKBU4quA132twl4g4lC4Gg5PP30ncy0sPE8	ļ	ð
r.	Hide secret key		
	KEY EXPIRATION TIME Valid forever		
	A The secret key won't be displayed again. Save this key for future use.		
	Download		se
12.		,	

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		×
NAME		
demo_s3_group		
USERS		
demo_s3_user ×		
POLICIES		
FullAccess ×		
	Cancel	Save

Create SVM S3 buckets

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

C2	■ ONTAP Sy	stem Manager	
	DASHBOARD	Buckets	
	INSIGHTS	+ 40	
	STORAGE ^	Name	Stor
	Overview		
	Volumes		
	LUNs		
	Consistency groups		
	NVMe namespaces		
	Shares		
	Buckets		
	Qtrees		
	Quotas		
	Storage VMs		
	Tiers		

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

NAME		
bucket		
CAPACITY 100 GiB ✓ Enable ListBucket access for all u Enabling this will allow users to	users on the storage VM ' access the bucket.	"svm_demo".

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

Add bucket

NAME
bucket
FOLDER (OPTIONAL)
Browse
Specify the folder to map to this bucket. Know more
CAPACITY 100 GiB V
Use for tiering
If you select this option, the system will try to select low-cost media with optimal performance for the tiered data.
S 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
Extreme ~
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 Q
DASHBO	DARD	
ALERTS	ø •	Tenants
NODES		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
TENAN	s	tenant name.
ILM		Leaster Export to CSV Actions ~ Search tenants by name or ID Q No results
CONFIG	URATION	Name 🕢 💠 Logical space used 🕢 💠 Quota utilization 🙆 💠 Quota 🕢 🗘 Object count 🌒 💠 Sign in/Copy URL 🥥
MAINTE	NANCE	
SUPPOI	रा	
		No tenante 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.

Ter	nants					
View info tenant na	rmation for each ter me.	nant account. Depending on the timing of ingest	s, network connectivity, and node s	tatus, the usage data shown	might be out of date. To view m	ore recent values, select the
Create	Export to CSV	Actions 🗸 Search tenants by name or ID	Q			Displaying one result
	Name 👔 🗘	Logical space used 👔 ≑	Quota utilization 👔 💠	Quota 🍘 ≑	Object count 🍘 💠	Sign in/Copy URL 👔
	tenant demo	0 bytes	-	-	0	→] []
						\leftarrow Previous 1 Next \rightarrow

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

$\leftarrow \ \ \rightarrow \ \ G$	A Not secure 192.168.0.80	/?accountId=270416	10751165610501					
🕙 Lab Status	S Power Controls S Accounts	cluster1-mgmt	S cluster2-mgmt	S Blue XP				
								NetApp Support NetApp
						Otorogo		
						StorageG	RID Tenant Manager	
						Recent	Optional	•
						Account ID	27041610751165610501	
						Account ib	2704101010100010001	
					NetApp	Username	root	
						Password		
							Sign in	

Create the user

Navigate to the Users tab and create a new user.

=	🗖 NetApp	orageGRID Tenant Manager	
DASHB	OARD		
STORA	GE (S3)	Users	
My aco	ess keys	View local and federated users. Edit properties and group membership of local users.	
Bucket	s		
Platfor	m services endpoints	1 user Creation Creatio Creation Creation Creation Creation Creation Creati	te user
ACCES	SS MANAGEMENT		
Groups	5	Actions ~	
Users			
Identit	y federation	Usernanie V Pur nanie V Peried V type V	
		root Root Local	
		\leftarrow Previous 1	Next \rightarrow

Demo S3 User		
fust contain at least 1 and no more than 128 chara	cters	
Jsername 😮		
demo_s3_user		
Password		
		0
fust contain at least 8 and no more than 32 charac	ters	
Confirm password		
		O
Deny access		

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.

A Not secure ht	ttps://192.168.0.80/ui/#/users/ebc132e2-cfc3-42c0-a4	445-3b4465cb523c		
Power Controls	Accounts 🖬 cluster1-mgmt 🔇 cluster2-mgmt 🤅	Blue XP		
NetApp Sto	rageGRID Tenant Manager			
^	Users > Demo S3 User			
	Overview			
	Full name: 🔗		Demo S3 User 🧪	
es endpoints	Username: 🝘		demo_s3_user	
GEMENT A	User type: 🥥		Local	
	Denied access: 🍘		Yes	
	Access mode: 👔		No Groups	
tion	Group membership: 🏈	•	None	
	Password Change password Change this user's passwo	Access 🔓 d ord.	Access keys	Groups
	*******		Change Password	

To create the S3 keys click on the user name.

≡	🗖 NetApp	StorageGRID Tenai	nt Manager				9
DASHB	OARD						
STORA	GE (S3)		Users				
Му асо	ess keys		View local and federated users. Edit properti	es and group membership of local users.			
Bucket	s		2.00000				
Platfor	m services endpoints		2 users				Create user
ACCES	S MANAGEMENT						
Groups			Actions V				
Users			Username 💠	Full Name 🗢	Denied 🗢	Туре 🗢	
Identit	yfederation						
			root	Root		Local	
			demo_a3_user	Demo S3 User	~	Local	
							\leftarrow Previous 1 Next \rightarrow

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 access key	×
Choose expiration time 2 Download access key	
Download access key To save the keys for future reference, select Download .csv , or copy and paste th	he values to another location.
(i) You will not be able to view the Access key ID or Secret access key after you close this	s dialog.
Access key ID	
7CT7L1X5MIO5091E86TR	
Secret access key	
RIJnC5N5FX9RSWgFdj6SQ7wMrfRZYu5bQLdNQT0c	
Download .csv	

Create the security group

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

C	Create group	×
1	1 Choose a group type Q Manage permissions Q Set S3 group policy Add use 0 Optional Optional Q	ers
(Choose a group type 🔞	
C	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 2	Manage permissions —— (3)	Set S3 group policy	Add users Optional
Manage group permissions Select an access mode for this group and sele Access mode ② Select whether users can change settings and Read-write ③ Read-only Group permissions ② Select the permissions you want to assign to the	ct one or more permissions. perform operations or whether they c this group.	an only view settings and featu	ires.
Allows users to access all administrat	ion features. Root access permission s	upersedes all other permission	s.
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 S credentials Allows users to create their own S3 access ke	33 and delete eys.
Previous Con			

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".



Create group	×	
Choose a group type — 🔗 M	lanage permissions 3 Set S3 group policy (4) Add users	
Set S3 group policy controls user access permise by default. No S3 Access Read Only Access Full Access Custom (Must be a valid JSON formatted string.)	sions to specific specific S3 resources, including buckets. Non-root users have no access "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.



Create two buckets

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

	■ NetApp	StorageGRID Tenant Manager	
DASHB	OARD		
STOR/	AGE (S3) ^	Buckets	
My acc	ess keys	Create buckets and manage bucket settings.	
Bucket	s m services endpoints	0 buckets Create Nicket	1
ACCES	S MANAGEMENT		1
Groups		Actions 🗸	Č
Users		Name 🗢 Region 🗢 Object Count 🚱 🗢 Space Used 🚱 🗢 Date Created 🗢	
Identit	y federation		-
		No buckets found	
		Create bucket	

Define the bucket name and region.

1 Enter details 2 Manage object settings Optional Enter bucket details Enter the bucket's name and select the bucket's region. Bucket name @	object settings on.
Enter bucket details Enter the bucket's name and select the bucket's region. Bucket name @	on.
Enter the bucket's name and select the bucket's region. Bucket name ②	n.
Bucket name 😧	
bucket	
Region 😢	
us-east-1 🗸	\sim

On this first bucket enable versioning.

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 v	rersioning		
Previous	Create bucket		

Now create a second bucket without versioning enabled.
Create bucket	
1 Enter details (Manage object settings Optional
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.

	Create bucket			×
(Senter 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		
	ĭ∞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 ...

Cancel

online help

X

~

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

Þ New bucket 🖶 Add external bucket 🛛 Refresh	Path: /					
ontap-dummy	Name		Size	Туре	Last Modified	Storage Class
	Upi	load file(s) load folder	(5)	~		
Tasks (1) Permissions Headers Tags Properties Preview Version	Uploa	d •	Download	Delete Rew	Folder Refresh	
Task	Size	%	Progress	Status	Speed	

📼 🌩 i

→ ~ ↑ 🔸 > Th	is PC > Downloads			✓ Ö Sea	rch Downloads
ganize 👻 New folde	er				
🖶 Downloads 👒 ^	Name	Date modified	Туре	Size	
🔮 Documents 🖈	9141P1_q_image.tgz	3/22/2024 1:25 AM	TGZ File	2,641,058 KB	
📰 Pictures 🛛 🖈	cluster1_demo_s3_user_s3_user.txt	3/23/2024 11:04 PM	Text Document	1 KB	
This PC	cluster1_svm_demo_s3_details (1).txt	3/23/2024 11:03 PM	Text Document	1 KB	
3D Objects	cluster1_svm_demo_s3_details.txt	3/23/2024 11:01 PM	Text Document	1 KB	
SU OBJECIS	🚔 hfs.exe	3/22/2024 1:24 AM	Application	2,121 KB	
Cloud Storage o	hotfix-install-11.6.0.14	3/23/2024 11:55 AM	14 File	717,506 KB	
E Desktop	🔗 putty	7/18/2020 6:39 PM	Shortcut	2 KB	
Documents	s3browser-11-6-7.exe	3/23/2024 12:36 PM	Application	9,807 KB	
Downloads					
b Music					
E Pictures					
Videos					
Local Disk (C:)					
		N			
		hr			

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

Accounts Bucke	ts Files	Bookmarks	Tools	Upgrade to Pro!	Help
----------------	----------	-----------	-------	-----------------	------

🖶 New bucket 🖶 Add external bucket 💋 Refresh	Path: /				
ontap-dummy	Name Cluster1_dem	Size 157 bytes 211 bytes	Type Text Document Text Document	Last Modified 3/23/2024 11:23:25 PM 3/23/2024 11:23:25 PM	Storage Class STANDARD STANDARD
	cluster1_svm	211 bytes 834.05 KB	Text Document Application	3/23/2024 11:23:25 PM 3/23/2024 11:23:25 PM 3/23/2024 11:23:25 PM	STANDARD STANDARD
	s3browser-11	9.58 MB	Application	3/23/2024 11:23:26 PM	STANDARD
	Upload -	Download	Delete Can New Fo	older Refresh	

📼 🌩 i 🛽

Tasks (1) Permissions Headers Tags Properties Preview Versions Event log

Now lets create some object versions in the bucket.

Delete a file.

ontap-dummy	Name	Size	Туре	LastModified	Storage Class
_ bucket	cluster1_dem	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
				0,20,2021112020111	
				Confirm File Delete	
	Upload -	Download	Delete 🧱 New Fo	lder 🥐 Are you sure	to delete 'putty.exe'?

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

→ * 个 🕹 > This	s PC > Downloads			✓ Ö Seard	h Downloads	9	
ganize 👻 New folder					EE ▼ III	?	Storage Clas
🕹 Downloads 🖈 ^	Name	Date modified	Туре	Size		И	STANDARD
🗄 Documents 🖈	9141P1 g image.tgz	3/22/2024 1:25 AM	TGZ File	2,641,058 KB		и	STANDARD
Nictures 🖈	cluster1_demo_s3_user_s3_user.txt	3/23/2024 11:04 PM	Text Document	1 KB		И	STANDARD
This DC	cluster1_svm_demo_s3_details (1).txt	3/23/2024 11:03 PM	Text Document	1 KB		и	STANDARD
	cluster1_svm_demo_s3_details.txt	3/23/2024 11:01 PM	Text Document	1 KB			
3D Objects	hfs.exe	3/22/2024 1:24 AM	Application	2,121 KB			
🔮 Cloud Storage o	hotfix-install-11.6.0.14	3/23/2024 11:55 AM	14 File	717,506 KB			
Desktop	🔗 putty	7/18/2020 6:39 PM	Shortcut	2 KB			
Documents	s3browser-11-6-7.exe	3/23/2024 12:36 PM	Application	9,807 KB			
Downloads							
Music							
Pictures							
Videos							
Local Disk (C:)							
~							
File pa	met hfe eve						
File ha	me: Infs.exe					~	

Tasks (1) Permissions Headers Tags Properties Preview Versions Event log

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

Accounts Buckets Files Bookmarks 100is Opgrade to Pro-	Petho (
New bucket 🖷 Add external bucket 😥 Kerresh	Path: /	-							
- intap-dummy	Name	Size	Туре	Last Modified	Storage Class				
- Ducket	cluster1_dem	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				
	lill cluster1_svm	211 bytes	Text Document	3/23/2024 11:23:25 PM	STANDARD				
	Infs.exe	2.07 MB	Application	3/23/2024 11:23:36 PM	STANDARD				
	S3Drowser-11	9.56 MB	Application	3/23/2024 11:23:26 PM	STANDARD				
									5 files (11.65 MR) and (
	Upload -	Download	Delete 📴 New F	older Refresh					5 files (11.65 MB) and 0
Tasks (1) Permissions Headers Tags Properties Preview	Versions Event log	Download	Delete Delete New F	older Refresh					5 files (11.65 MB) and C
Tasks(1) Permissions Headers Tags Properties Preview URL: http://bucket.s3portal.demo.netasp.com.8080/	Versions Event log	Download	Delete P	older Refresh					5 files (11.65 MB) and 0
Tasks(1) Permissions Headers Tags Properties Preview URL: http://buckt.s3portal.demo.netapp.com:8080/ Kev	Versions Event log	Download	Delete Rew F	older Refresh		Size	Storane Class	Quiner	5 files (11.65 MB) and 0
Tasks (1) Permissions Headers Tags Properties Preview URL: http://bucket.s3portal.demo.netapp.com:8080/ Key Dictated demo.s3 user 53 user 51	Versions Event log	Download	Delete g New F	older Refresh		Size	Storage Class	Owner	5 files (11.65 MB) and 0
Tasks (1) Permissions Headers Tags Properties Preview URL: http://bucket.s3portal.demo.netapp.com:8080/ Key Clustert demo_s3_user_s3_user.bd servicing #: M_custert	Versions Event log	Download	Delete Rew F	older Refresh	Dolhe	Size 157 butes	Storage Class	Owner Lickonaur (Lickonaur)	S files (11.65 MB) and G
Tasks (1) Permissions Headers Tags Properties Preview URI: http://bucket.s3portal.demo.netapp.com:8080/ Key Claster1_demo_s3_user_s3_usertM revision #1 (current) Cutater1_sem_demo_s3_datable (1) bt	Versions Event log Last Modified 3/23/2024 11:23.	Download	Delete Rew F	older Refresh	Oe1bc	Size 157 bytes	Storage Class STANDARD	Owner Unknown (Unknown)	5 files (11.65 MB) and 0
Tasks (1) Permissions Headers Tags Properties Preview URL: http://bucket.s3portal.demo.netapp.com:8080/ Key Cluster1_demo_s3_user_d3_usertM revision #1 (urrend) Cluster1_svm_demo_s3_detals (1) bt revision #1 (urrend)	Versions Event log	25 PM	Delete Rew F ETag ac4c 40775	older Refresh	Фе1bс #14	Size 157 bytes 211 bytes	Storage Class STANDARD STANDARD	Owner Unknown (Unknown) Liekonard Inknown)	5 files (11.65 MB) and 0
Tasks (1) Permissions Headers Tags Properties Preview URL: http://bucket.s3portal.demo.netapp.com:8080/ Key © cluster1_demo_s3_user_s3_user bt revision # 1 (current) © cluster1_sym.demo_s3_details (1) bt revision # 1 (current) © cluster1_sym.demo_s3_details (1) bt	Versions Eventlog Last Modified 3/23/2024 11.23 3/23/2024 11.23	Download 25 PM 25 PM	Delete Rew F	older eresh 9543e97ef3678b2b6ed6ad 33b646a6cfef19fde71eefb5	- Oe1bc #04	Size 157 bytes 211 bytes	Storage Class STANDARD STANDARD	Owner Unknown (Unknown) Unknown (Unknown)	5 files (11.65 HB) and c C Version Id Mtg0Mj01MDAwL. NDg0Mj01MDAw.
Tasks (1) Permissions Headers Tags Properties Preview URL: http://bucket.s3portal.demo.netapp.com:0880/ Key Cluster1_demo_s3_user_s3_uset.st revision #1 (current) Cluster1_svm_demo_s3_details (1) bt revision #1 (current) Cluster1_svm_demo_s3_details bt revision #1 (current)	Versions Event log	25 PM 25 PM 25 PM	Delete Rew F	older Refresh	0e1bc #04	Size 157 bytes 211 bytes 211 bytes	Storage Class STANDARD STANDARD STANDARD	Owner Unknown (Unknown) Unknown (Unknown)	5 files (11.63 HB) and (0 Version Id Mrg0Mj01MDAwL. NDg0Mj01MDAw. NTI (20.0MDA.)
Tasks (1) Permissions Headers Tags Properties Preview URL: http://bucket.s3portal.demo.netapp.com:8080/ Key ClusterT_demo_s3_user_s3_user bd revision # 1 (current) ClusterT_syme_demo_s3_details (1) bd revision # 1 (current) ClusterT_syme_demo_s3_details bd revision # 1 (current)	Versions Eventlog Last Modified 3/23/2024 11.23 3/23/2024 11.23 3/23/2024 11.23	25 PM 25 PM 25 PM 25 PM	Delete New F	older Refresh 9543e97ef3678b2b6ed6a 33b646a6cfef19fde71eefb5	00e1bc #04 10e2	Size 157 bytes 211 bytes 211 bytes	Storage Class STANDARD STANDARD STANDARD STANDARD	Owner Unknown (Unknown) Unknown (Unknown) Unknown (Unknown)	S files (11.65 MB) and C S files (11.65 MB) an
Tasks (1) Permissions Headers Tags Properties Preview URL: http://bucket.s3portal.demo.netspp.com:0880/ Key Cluster1_demo_s3_user_s3_user bt revision #? (current) Cluster1_svm_demo_s3_details (1) bt revision #? (current) Cluster1_svm_demo_s3_details bt revision #? (current) Testister #? (current)	Versions Event log Last Modified 3/23/2024 1123. 3/23/2024 1123. 3/23/2024 1123.	25 PM 25 PM 25 PM 25 PM	ETag actic 17d2	older pg Refresh 9543e97e19678b2b6ed6ad 3b646a6cfer191de71 eefb 56518564490a597arE94eccc0	0e1bc #04 10e2 65477	Size 157 bytes 211 bytes 211 bytes 211 bytes 217 bytes	Storage Class STANDARD STANDARD STANDARD STANDARD	Owner Unknown (Unknown) Unknown (Unknown) Unknown (Unknown)	5 files (11.65 He) and (2 Version Id Mrg0MjO1MDAwL, NDg0MjO1MDAw., NTU2NzI0MDAwL, NTU2NzI0MDAwL,
Tasks (1) Permissions Headers Tags Properties Preview URL: http://bucket.s3portal.demo.netapp.com:8080/ Key Choster1.gence.s3_user_s3_user.bit revision # 1 (current) Cutser1.gence.s3_details (1) bit revision # 1 (current) Cutser1.gence.s3_details bit revision # 1 (current) Finfs.com	Versions Eventlog Last Modified 3/23/2024 1123. 3/23/2024 1123. 3/23/2024 1123. 3/23/2024 1123.	25 PM 25 PM 25 PM 25 PM 36 PM 36 PM	ET ag acHc 40775 17d2	older Peffesh 9543e97erJ6678b2b6ed6ad 316646ae6clef19de71ee6b 16511566490a567atT9fec7c	0e1bc m04 10e2 63477 63477	Size 157 bytes 211 bytes 211 bytes 207 MB	Storage Class STANDARD STANDARD STANDARD STANDARD STANDARD	Owner Unknown (Unknown) Unknown (Unknown) Unknown (Unknown) Unknown (Unknown)	S files (11.65 HB) and C
Tasks (1) Permissions Headers Tags Properties Preview URL: http://bucket.s3portal.demo.netapp.com:0800/ Key Cluster1_demo_s3_user_s3_user bt revision # 1 (current) Cluster1_svm_demo_s3_details (1) bt revision # 1 (current) Tervision # 2 (current) Tervision # 1 Current) Tervision # 1	Versions Eventlog	25 PM 25 PM 25 PM 25 PM 36 PM 25 PM	ETag acHc 40775 17d2(9e855	older Press 9543997e13678b2b6e46a46 336646a8cfer19fefe71e0b5 9651856490.as57at534ece91 6651856490.as57at534ace91	00e1bc #04 10e2 63477 63477	Size 157 bytes 211 bytes 211 bytes 210 7 MB 207 MB	Storage Class STANDARD STANDARD STANDARD STANDARD STANDARD	Owner Unknown (Unknown) Unknown (Unknown) Unknown (Unknown) Unknown (Unknown)	5 files (11.65 MB) and CC CC Version Id Mrg0Mj01MDAwL, NDg0Mj01MDAwL, NTU23kJ0MDAwL, Ng1010F4MDAwL, Ng1020J3MDAwL,
Tasks (1) Permissions Headers Tags Properties Preview URL: http://bucket.s?portal.demo.netspp.com:8080/ Key Cluster1_demo_s3_user_s3_user.bd revision #.1 (current) Cluster1_svm_demo_s3_details (1) bd revision #.1 (current) Cluster1_svm_demo_s3_details bd revision #.1 (current) This.exe revision #.2 (current) revision #.2 (current) revision #.2 (current) revision #.2 (current) revision #.2 (current) revision #.2 (current)	Versions Eventlog Last Modified 3/23/2024 11.23. 3/23/2024 11.23. 3/23/2024 11.23. 3/23/2024 11.23. 3/23/2024 11.23. 3/23/2024 11.23.	25 PM 25 PM 25 PM 25 PM 25 PM 36 PM 25 PM 31 PM	ETag ac4c 40775 17d2 9e855	older Peffesh 9543e97et/3678b2b6ed6ad 336466a6cfef19fde71eefb 5651856490a657af39feccc 57e98ed1269372ff0ace91c	0e1bc #04 10e2 63477 63477	Size 157 bytes 211 bytes 211 bytes 207 MB 207 MB	Storage Class STANDARD STANDARD STANDARD STANDARD STANDARD	Owner Unknown (Unknown) Unknown (Unknown) Unknown (Unknown) Unknown (Unknown) Unknown (Inknown)	S files (11.65 HB) and C
Tasks (1) Permissions Headers Tags Properties Preview URL: http://bucket.s3portal.demo.netapp.com:0080/ Key Cluster1_demo_s3_user_s3_user bd revision # 1 (current) Cluster1_svm_demo_s3_details t01 bd revision # 2 (current) Tevision # 1 (current) Puptly ave revision # 2 (current) revision # 2 (current)	Versions Event log	25 PM 25 PM 25 PM 25 PM 36 PM 37 PM 31 PM 31 PM	Delete Rag acf4c 40775 17d2(9e655 9e655	older Py Refresh 9543957613678b2b6ed6ad 336646a6Cefr19(de 71 eebb 9651856480.0a587at394eccc 77698ed1269372870ace91c	00-1bc #04 10+2 63477 63477 150+9	Size 157 bytes 211 bytes 211 bytes 207 MB 207 MB 207 MB	Storage Class STANDARD STANDARD STANDARD STANDARD STANDARD	Owner Unknown (Unknown) Unknown (Unknown) Unknown (Unknown) Unknown (Unknown) Unknown (Unknown)	5 files (11.65 He) and C C Version Id Meg0Mj01MDAwL, NDg9Mj01MDAwL, NTU28420MDAwL, Ng120DBMDAWL, Ng120DB
Tasks (1) Permissions Headers Tags Properties Preview URL: http://bucket.s3portal.demo.netapp.com:0080/ Key Cluster1_demo_s3_user_s3_userbt revision # 1 (current) Cluster1_svm_demo_s3_details bt revision # 1 (current) Cluster1_svm_demo_s3_details bt revision # 1 (current) revision # 1 (current) revision # 2 (current) revision # 2 (cluster) revision # 1 Storgere=116-6 zep	Versions Event log Last Modified 3/23/2024 1123: 3/23/2024 1124: 3/23/2024 1124: 3/23/2024 1124: 3/23/2024 1124: 3/23/	25 PM 25 PM 25 PM 25 PM 36 PM 35 PM 31 PM 25 PM	ETag ac4c 40775 17d2 9e855 9e855	older Petersh 9543-897-815678b2b6e456a4 33646-6a6clef1916e71eeb5 165115564830a587at394eccc 77e98ed1269372t10acc91c 77e98ed1269372t10acc91c	0e1bc m04 10e2 63477 1fc0e9	Size Size 157 bytes 211 bytes 207 MB 2.07 MB 2.07 MB 2.07 MB	Storage Class STANDARD STANDARD STANDARD STANDARD STANDARD	Owner Unknown (Unknown) Unknown (Unknown) Unknown (Unknown) Unknown (Unknown) Unknown (Unknown) Unknown (Unknown)	S files (11.65 MB) and C
Tasks (1) Permissions Headers Tags Properties Preview URL: http://bucket.s3portal.demo.netspp.com:0080/ Key Cluster1_demo_s3_user_s3_user tot revision # 1 (current) Cluster1_svm_demo_s3_details tot revision # 1 (current) Phick revision # 1 (current) Phick revision # 1 (current) Publy revision # 1 (current) Publy revision # 2 (current) revision # 2 (current) revision # 2 (current) revision # 1 (current) Image:	Versions Event log	25 PM 25 PM 25 PM 25 PM 36 PM 36 PM 31 PM 25 PM 36 PM	Deter Rag actic 40775 17d2 9e655 9e655 54cb5	older Pys Refresh 9543957613678b2b6ed6ad 336646a6Cefr19(de71 eebb 9651856480.a657a1394eccc 74986ed1269372870ace91 c 74986ed1269372870ace91 c 11395cdaad9447882533c2	00-1bc m04 10-2 63477 16c0e9 16c0e9	Size 157 bytes 211 bytes 211 bytes 207 MB 2.07 MB 834.05 KB 9.88 MB	Storage Class STANDARD STANDARD STANDARD STANDARD STANDARD STANDARD	Owner Unknown (Unknown) Unknown (Unknown) Unknown (Unknown) Unknown (Unknown) Unknown (Unknown) Unknown (Unknown)	5 files (11.65 He) and C C Version Id Meg0Mj01MDAwL, NDg0Mj01MDAwL, NTU2Nz0MDAwL, NE2010TE4MDAWL, NE2010TE4MDAWL, NE2

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".

	tem Manager		Search actions, objects	and pages	۹		? ↔ 8
DASHBOARD	Peer cluster	Lets you select specific volumes don't need to protect entire stor	for protection if you rage VMs.	Lets you select where to a cloud destination	nich volumes y ntion.	you want to be backed up	Lets you protect a consistency group with a zero recovery time objective.
INSIGHTS		 NetApp SnapCenter software s 	simplifies backup, restore, and	clone management for t	the applications I	hosted across ONTAP enabled p	latforms. Use NetApp SnapCenter for application-consistent
STORAGE 🗸 🗸							
NETWORK ~							
EVENTS & JOBS 🛛 🗸		Bucket protection					
PROTECTION ^						2 of the 2 bucke	ts aren't protected.
Overview							
Relationships		Back up to cloud					
ноѕтѕ 🗸		0% 10% 20% 30				2 of the 2 bucke	ts aren't backed up to the cloud.
CLUSTER 🗸 🗸							
		Protect buckets					
		Lets you select specific buckets f	for setting up SnapMirror	protection to either a	in ONTAP dest	tination or a cloud destination	n.
		StorageGRID	0				
		ONTAP S3					
		aws Amazon S3					
		C Others					
		bj	ect store provider to acce	ss the data bucket.	alio data conta	ainer must be created with ti	ne object store provider. This assumes that the user has valid
		+ Add ~					

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

sgws_demo		
URL STYLE		
Path-style URL		~
OBJECT STORE SCOPE		
O Cluster	Storage VM	
USE BY 👔		

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.

ACCESS KEY				
7CT7L1X5MIO5	091E86TR			
SECRET KEY				
	•			
bucket				
Network fo	r cloud object sto	ore		Considerations
	r cloud object sto	ORE SUBNET MASK	BROADCAST DOMAIN	Considerations
Network fo NODE onPrem-01	r cloud object sto IP ADDRESS 192.168.0.113	O re SUBNET MASK 24	broadcast domain Default	Considerations GATEWAY 192.168.0.1
Network fo NODE onPrem-01	r cloud object sto IP ADDRESS 192.168.0.113	O re SUBNET MASK 24	broadcast domain Default	Considerations GATEWAY 192.168.0.1
Network fo NODE onPrem-01	r cloud object sto IP ADDRESS 192.168.0.113	O re SUBNET MASK 24	broadcast domain Default	Considerations GATEWAY 192.168.0.1
Network fo NODE onPrem-01	r cloud object sto IP ADDRESS 192.168.0.113	O re SUBNET MASK 24	broadcast domain Default	Considerations GATEWAY 192.168.0.1

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

🗧 🔳 ONTAP Sy	stem Manager	Search actions	, objects, and pages Q	Ø ↔ 8
DASHBOARD INSIGHTS STORAGE ~ NETWORK ~ EVENTS & JOBS ~		Back up to cloud C% 10% 20% 30% 40% Protect buckets Lets you select specific buckets for setting up Sn	2. 50% 60% 70% 80% 90% 180%	of the 2 buckets aren't backed up to the cloud. bud destination.
PROTECTION ^ Overview Relationships		^ Local policy settings ●		
CLUSTER ~		Protection policies Applicable when this cluster is the destina Agrichronoa Al Similae sast the hour, every hour AdomatedFaiCver No schedules Constructions No schedules Contrologies No schedules	→ Snapshot policies Applicable when this duster is the source actual 3-Shondwis actual:-weekly 3-Shondwis none Ne schedules	→ Schedules → Srin At 0, 13 15, 20, 25, 30, 35, 40, 45, 50, and 35 minutes past the Act 0, 13 15, 20, 25, 30, 35, 40, 45, 50, and 35 minutes past the Act 0, 13 15, 20, 25, 30, 35, 40, 45, 50, and 50 minutes past the Note: Act 0, 15, 20, 30, 40, and 50 minutes past the hour, every hour 10. Changy

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

Pol	icie	Protection	overview			
	Prote	ection policies	Snapshot policies			
	+ 4	Add				
		Name		Description	Policy type	Scope
		Continuous			(All) 🗸	
		Continuous	i	Policy for S3 bucket mirroring.	Continuous	Cluster
	T	THROTTLE Unlimited	Edit	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

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

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

В	uckets						
	+ Add				Q Search	⊥ Download © Show/hid	de 💙 🖙 Filter
	Name	Storage VM	Lifecycle rules	Capacity (available total)		Protection	Path
	bucket	svm_demo			100 GiB 100 GiB	80 <u>0</u>	-
	ontap-dummy	svm_demo			100 GiB 100 GiB	800	-

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

Buc	kets												
+ A0	d 🗊 Delete 🗘 Protect											Q Search	≡ Filter
	Name	buc	-ket All	All Ruckets de ca								More	
	ontap-dummy			Permissions		SnapM	lirror (ONTAP or cloud)						
			6	•			Destination		Backastlan anti-	Balatianakia kashk		N	
			Source			Destination		Protection policy	Relationship health	3	otate		
			svm_den	no:/bucket/bu	cket		sgws_demo:/objstore			(Healthy	6	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) - Buc	ket (Migration Temp)			📼 🌩 i	1				- 0	×
New bucket Add external bucket Refresh	Path: /								\$ / D Y	7 6
bucket	Name cluster1_dem cluster1_svm cluster1_svm in his exe sthut	Size 157 bytes 211 bytes 211 bytes 207 MB 9 58 MB	Type Text Document Text Document Text Document Application	Last Modified 3/24/2024 12:13:53 AM 3/24/2024 12:13:53 AM 3/24/2024 12:13:53 AM 3/24/2024 12:13:53 AM 3/24/2024 12:13:53 AM	Storage Class STANDARD STANDARD STANDARD STANDARD STANDARD					
			.,,							
	Upload -	Download	Delete 🧱 New F	older Refresh					1 file (2.07 MB) s	elect
URL: http://192.168.0.80:8080/bucket/hfs.exe	ew versions Eventing								j] o	ору
Key	Last Modified		ETag	1		Size	Storage Class	Owner	Version Id	
revision #: 1 (current)	3/24/2024 12:13:	53 AM	"9e85	57e98ed1269372ff0ace91	d63477"	2.07 MB	STANDARD	tenant_demo (27041610751	NjU5RDhCNDltRT	

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

S3 Browser 11.6.7 - Free Version (for non-commercial use only) - B	lucket (original and post-migration)			📼 🌩 i	1				- 0	×
Accounts Buckets Files Bookmarks Tools Upgrade to	Pro! Help									
🖶 New bucket 🖶 Add external bucket 🥰 Refresh	Path: /									1 🗈
- 🧐 ontap-dummy	Name	Size	Туре	Last Modified	Storage Class					
-e bucket	Cluster1_dem	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					
	Dutty.exe	834.05 KB	Application	3/23/2024 11:23:25 PM	STANDARD					
	This.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					
				-						
	Upload -	Download	Delete Delete New F	older Refresh					6 files (12.46 MB) and	folde
Tasks (1) Permissions Headers Tags Properties Pre	eview Versions Event log									
URL: http://bucket.s3portal.demo.netapp.com:8080/										ору
Key	Last Modified		ETag			Size	Storage Class	Owner	Version Id	
cluster1_demo_s3_user_s3_user.txt										
revision #: 1 (current)	3/23/2024 11:23:	25 PM	acf4c!	9543e97ef3678b2b6ed6a6	50e1bc	157 bytes	STANDARD	Unknown (Unknown)	Mzg0MjQ1MDAwL	
cluster1_svm_demo_s3_details (1).txt										
revision #: 1 (current)	3/23/2024 11:23:	25 PM	40775	3b646a6cfef19fde71eefb5	5ff04	211 bytes	STANDARD	Unknown (Unknown)	NDg0MjQ1MDAw	
cluster1 svm demo s3 details.txt										
revision #: 1 (current)	3/23/2024 11:23:	25 PM	17d20	651856f480a587af39fecco	:10e2	211 bytes	STANDARD	Unknown (Unknown)	NTU2NzI0MDAwL	
This exe										
revision # 3 (current)	3/24/2024 12:14	52 AM	9e855	7e98ed1269372ff0ace91c	163477	2.07 MB	STANDARD	Unknown (Unknown)	NTY0NDaxMDAw.	
revision 2	3/23/2024 11:23	36 PM	9e855	7e98ed1269372ffDace91c	63477	2.07 MB	STANDARD	Unknown (Unknown)	NzQ10TE4MDAw	
revision # 1	3/23/2024 11:23	25 PM	9e855	7e98ed1269372ff0ace91c	63477	2.07 MB	STANDARD	Linknown (Linknown)	Nik20DI3MDAwl n	
I putty eye	5/25/2024 11:25.		36035			2.01 110	ST. NOARD	Children (Children)	ingrade broth Drivelin	
revision # 1 (current)	3/23/2024 11:23-	25 PM	54cb9	1395cdaad9d47882533c3	P1fc0a9	834.05 KB	STANDARD	Linknown (Linknown)	NzE2NzEvMDAwl	
shrowson11.6.7 oxo	5/25/2024 11.25.	L	54003	1555566666564766255562		004.00 KD	GTANDARD	Chikitowii (Chikitowii)	Manager and Awe	
Subiovati i to read	2/22/2024 11:22	26 DM	204	-0705447920623-6203-5-4#	19:040 0	0.50 MD	STANDARD	Liebe sure (Liebe sure)	NDV20D	
revision #. I (current)	5/25/2024 11:23:	20 P/WI	ae3bt	uarount/023020033/C50R	10210-2	9.00 MD	STANDARD	Unknown (Unknown)	NDT20DCWMDEU	

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.

S3 Browser 11.6.7 - Free Version (for non-commercial use only) - Bucket (Migration Temp)			📼 🌩 i	1				- 6
Accounts Buckets Files Bookmarks 100is Upgrade to Pro:	Path: /								\$ I 🗆
sg dummy	Name Custer 1_dem_ Custer 1_svm_ Custer 1_svm_ Custer 1_svm_ Support	Size 157 bytes 157 bytes 211 bytes 834 05 KB 207 MB 9.58 MB	Type Tex Document Tex Document Application Application Application	Last Modified 3/24/2024 12:13:53 AM 3/24/2024 12:13:53 AM 3/24/2024 12:13:53 AM 3/24/2024 12:13:53 AM 3/24/2024 12:14:54 AM 3/24/2024 12:14:56 AM 3/24/2024 12:13:53 AM	Storage Class STANDARD STANDARD STANDARD STANDARD STANDARD STANDARD STANDARD				
	Upload -	Download	Delete 🥁 New Fo	older Refresh					1 file (2.07 MB)
Tasks (1) Permissions Headers Tags Properties Preview	Versions Event log								153
Ver N	Lost Medified		ETaa			Cizo	Staroas Class	Ounor	Version Id
This exe	Lastwoollied		Ling			0120	Glorage class	Owner	Veraiorria
revision #: 2 (current)	3/24/2024 12:14:5	56 AM	"9e85	57e98ed1269372ff0ace91	d63477"	2.07 MB	STANDARD	tenant_demo (27041610751	OEI4RjY4NDgtRT
revision #: 1	3/24/2024 12:13:5	53 AM	"9e85	57e98ed1269372ff0ace91	d63477*	2.07 MB	STANDARD	tenant_demo (27041610751	NjU5RDhCNDltRT

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.

=	NetApp	StorageGRID Grid Manager	Search by page title	৫ ় ^	💄 Root 🗸
DASH	IOARD		Documentation Ce	iter	
ALERT	s 🥝 🗸 🗸	Dashboard	API Documentation		
NODE			About		
TENA	тѕ	Health 😡	Available Storage 😡		
ILM		U License Status	Overall 📲		
CONFI	GURATION	1	Used		
MAINT	ENANCE	License			

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.

accounts	Operations on accounts		×
POST	/grid/account-enable-s3-key-import Enables the Import S3 Cred to have full access to tenant	entials feature on this node. Warning: Enabling this feat t data. This feature should be disabled immediately after	ture allows Grid Manager users with Change Tenant Root Password permission r use.
Parameters			Cancel
No parameters			
	Execute	h	Clear
			U.C.M

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 * required string (path)	ID of Storage Tenant Account
	27041610751165610501
<pre>user_id * required string (path)</pre>	ID of user in tenant account.
	ebc132e2-cfc3-42c0-a445-3b4465cb523c
body * required (body)	Edit Value Model
	<pre>{ "accessKey": "<u>3TVPI142JGE3Y7FV2KC0</u>", "secretAccessKey": "<u>75alQqKBU4quA132twI4g41C4Gg5PP30ncy0sPE8</u>" }</pre>

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.	â		
Parameters	Cappel		
No parameters			
Execute			
Responses	Response content type application/json v		

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

Users >	Demo	\$3	User
03013 -	DCIIIO	~~	0001

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 Manage access keys Add or delete access keys for this user. Create key Actions ∽	Access keys Groups	
Access key ID 💠		Expiration time 🗢
**********************************		None

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.



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

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.