



## **Product feature guides**

### **How to enable StorageGRID in your environment**

NetApp

March 05, 2024

This PDF was generated from <https://docs.netapp.com/us-en/storagegrid-enable/product-feature-guides/create-cloud-storage-pool-aws-google-cloud.html> on March 05, 2024. Always check docs.netapp.com for the latest.

# Table of Contents

- Product feature guides . . . . . 1
  - Create Cloud Storage Pool for AWS or Google Cloud . . . . . 1
  - Create Cloud Storage Pool for Azure Blob Storage . . . . . 1
  - Use a Cloud Storage Pool for backup . . . . . 2
  - Configure StorageGRID search integration service . . . . . 3
  - Node Clone . . . . . 19
  - How to use port remap . . . . . 22

# Product feature guides

## 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 11.6 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:: nnnnnn:user/xyzabc"},
      "Action": "es:*",
      "Resource": "arn:aws:es:us-east-1:nnnnnn:domain/sgdemo/*"
    },
    {
      "Effect": "Allow",
      "Principal": {"AWS": "*"},
      "Action": [
        "es:ESHttp*"
      ],
      "Condition": {
        "IpAddress": {
          "aws:SourceIp": [ "nnn.nnn.nn.n/nn"
          ]
        }
      },
      "Resource": "arn:aws:es:us-east-1:nnnnnn:domain/sgdemo/*"
    }
  ]
}

```

## Fine-grained access control

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



☐ Enable fine-grained access control

## SAML authentication for OpenSearch Dashboards/Kibana

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



☐ Prepare SAML authentication

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

## Amazon Cognito authentication

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



☐ Enable Amazon Cognito authentication

## Access policy

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



### Domain access policy

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

Visual editor

JSON

Import policy

### Access policy

```
3+  "Statement": [  
4+  {  
5+    "Effect": "Allow",  
6+    "Principal": {  
7+      "AWS": "arn:aws:iam::123456789012:user/ashley"  
8+    },  
9+    "Action": "es:*",  
10+   "Resource": "arn:aws:es:us-east-1:123456789012:domain/sgdemo/*"  
11+ },  
12+ {  
13+   "Effect": "Allow",  
14+   "Principal": {  
15+     "AWS": "*"   
16+   },  
17+   "Action": [  
18+     "es:ESHttp*"   
19+   ],  
20+   "Condition": {  
21+     "IpAddress": {  
22+       "aws:SourceIp": [  
23+         "216.240.240.0/24"  
24+       ]  
25+     }  
26+   },  
27+   "Resource": "arn:aws:es:us-east-1:123456789012:domain/sgdemo/*"  
28+ }
```




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



3. Click OpenSearch Dashboards URL to open the domain in a new tab to access the dashboard. If you get an access denied error, verify that the access policy source IP address is correctly set to your computer public IP to allow access to the domain dashboard.
4. On the dashboard welcome page, select Explore On Your Own. From the menu, go to Management → Dev Tools
5. Under Dev Tools → Console, enter `PUT <index>` where you use the index 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.



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



## 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

1

Enter details

2

Select authentication type  
Optional

3

Verify server  
Optional

[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  
Optional

✓ Verify server  
Optional

### Authentication type ?

Select the method used to authenticate connections to the endpoint.

Access Key

Access key ID ?

AKIA[REDACTED]UWO

Secret access key ?

.....

👁

Previous

Continue

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

## Platform services endpoints

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

1 endpoint Create endpoint

Delete endpoint

<input type="checkbox"/>	Display name ?	Last error ?	Type ?	URI ?	URN ?
<input type="checkbox"/>	aws-opensearch		Search	https://search-sgdemo-2-2021-11-24-1-us-east-1.es.amazonaws.com/	arn:aws:es:us-east-1:2[REDACTED]:domain/sgdemo/sgmetadata/_doc

## Search integration services with on premises Elasticsearch

### On premises Elasticsearch setup

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

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

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

- To start docker after reboot, enter the following:

```
sudo systemctl enable docker
```

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

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

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

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

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



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

## Install and run Elasticsearch

1. Install and start [Docker Desktop](#).
2. Run:

```
docker network create elastic
docker pull docker.elastic.co/elasticsearch/elasticsearch:8.1.0
docker run --name es-node01 --net elastic -p 9200:9200 -p 9300:9300 -it
```

When you start Elasticsearch for the first time, the following security configuration occurs automatically:

- [Certificates and keys](#) are generated for the transport and HTTP layers.
- The Transport Layer Security (TLS) configuration settings are written to `elasticsearch.yml`.
- A password is generated for the `elastic` user.
- An enrollment token is generated for Kibana.



You might need to scroll back a bit in the terminal to view the password and enrollment token.

3. Copy the generated password and enrollment token and save them in a secure location. These values are shown only when you start Elasticsearch for the first time. You'll use these to enroll Kibana with your Elasticsearch cluster and log in.



If you need to reset the password for the `elastic` user or other built-in users, run the `elasticsearch-reset-password` tool. To generate new enrollment tokens for Kibana or Elasticsearch nodes, run the `elasticsearch-create-enrollment-token` tool. These tools are available in the Elasticsearch `bin` directory.

## Install and run Kibana

To analyze, visualize, and manage Elasticsearch data using an intuitive UI, install Kibana.

1. In a new terminal session, run:

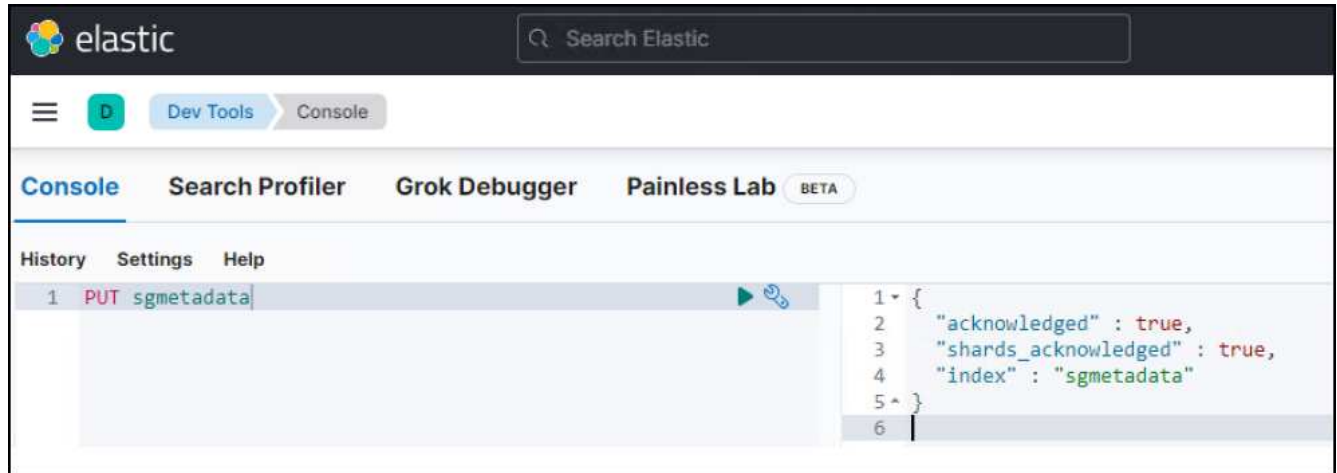
```
docker pull docker.elastic.co/kibana/kibana:8.1.0
docker run --name kib-01 --net elastic -p 5601:5601 docker.elastic.co/k
```

When you start Kibana, a unique link is output to your terminal.

2. To access Kibana, click the generated link in your terminal.

- a. In your browser, paste the enrollment token that you copied and click the button to connect your Kibana instance with Elasticsearch.
- b. Log in to Kibana as the `elastic` user with the password that was generated when you started Elasticsearch.

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



## Platform services endpoint configuration

To configure endpoints for platform services, follow these steps:

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

## Create endpoint


1 Enter details


2 Select authentication type  
Optional


3 Verify server  
Optional

### Enter endpoint details

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

Display name 

URI 


URN 


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


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


### Authentication type

Select the method used to authenticate connections to the endpoint.

Basic HTTP 

Username 

Password 

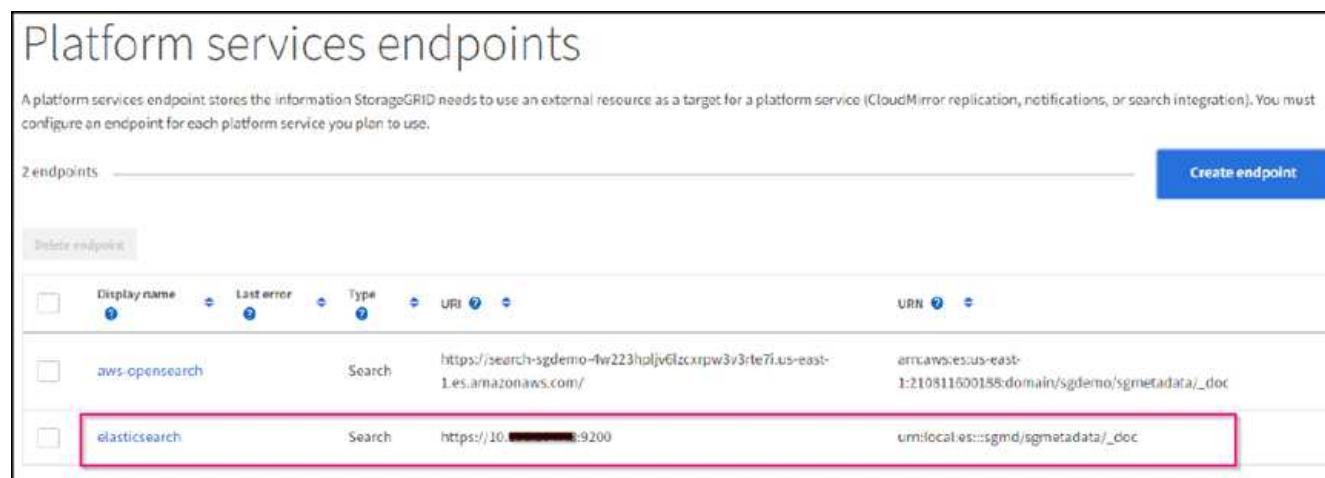
..... 

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

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



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



## Bucket search integration service configuration

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

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

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

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

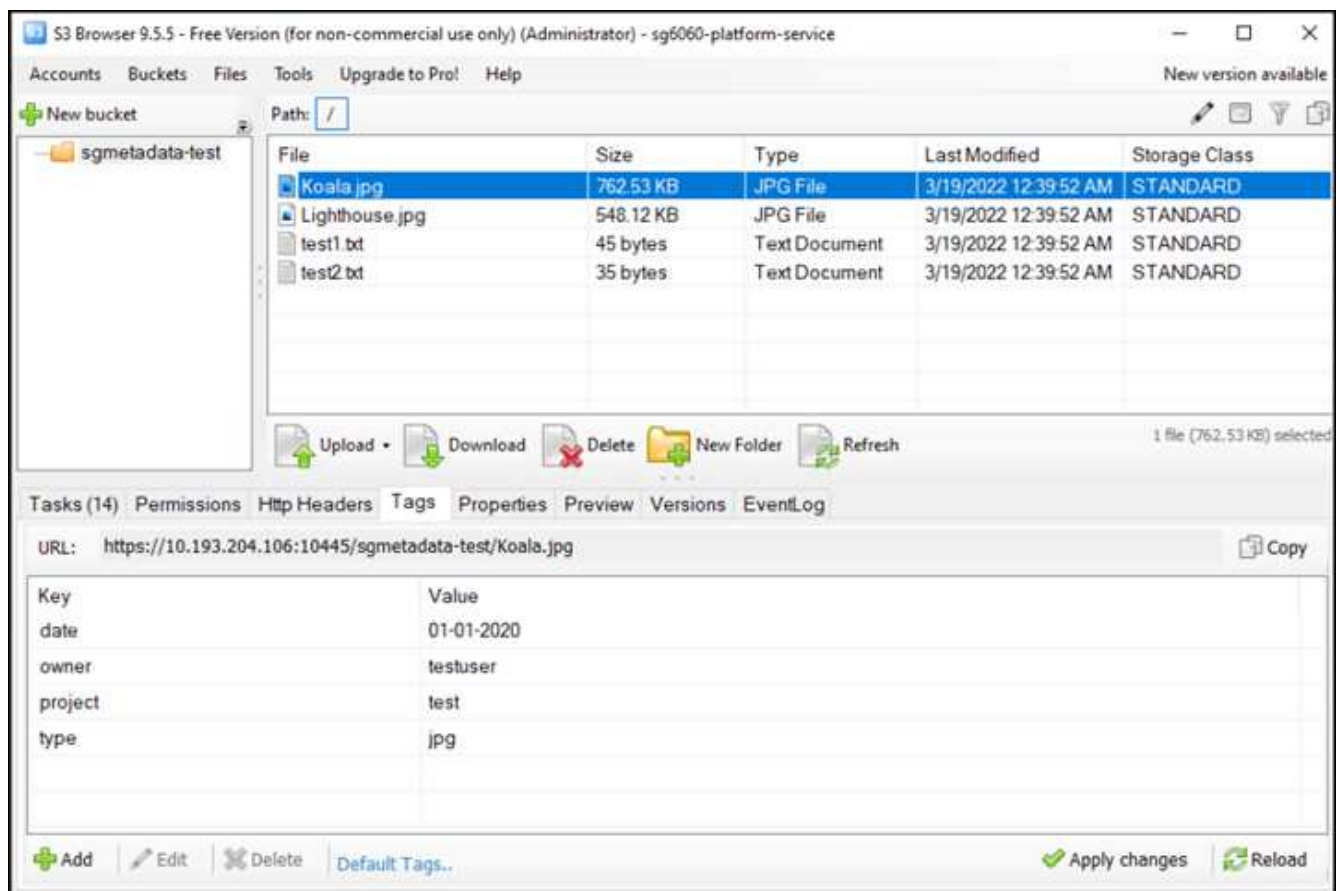
In this example, we used no prefix, meaning that the metadata for every object in this bucket is sent to the Elasticsearch endpoint defined previously.

```

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

```

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



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 screenshot shows the Elastic Search Console interface. On the left, the query is entered in the console: `GET sgmetadata/_search` with a body of `{ "query": { "match_all": { } } }`. On the right, the search results are displayed as a JSON array. The first record is for a file named `test1.txt` with a score of 1.0. The second record is for a file named `test_Koala.jpg` with a score of 1.0. Both records include metadata such as `bucket`, `key`, `accountId`, `size`, `md5`, `region`, `metadata`, and `tags`.

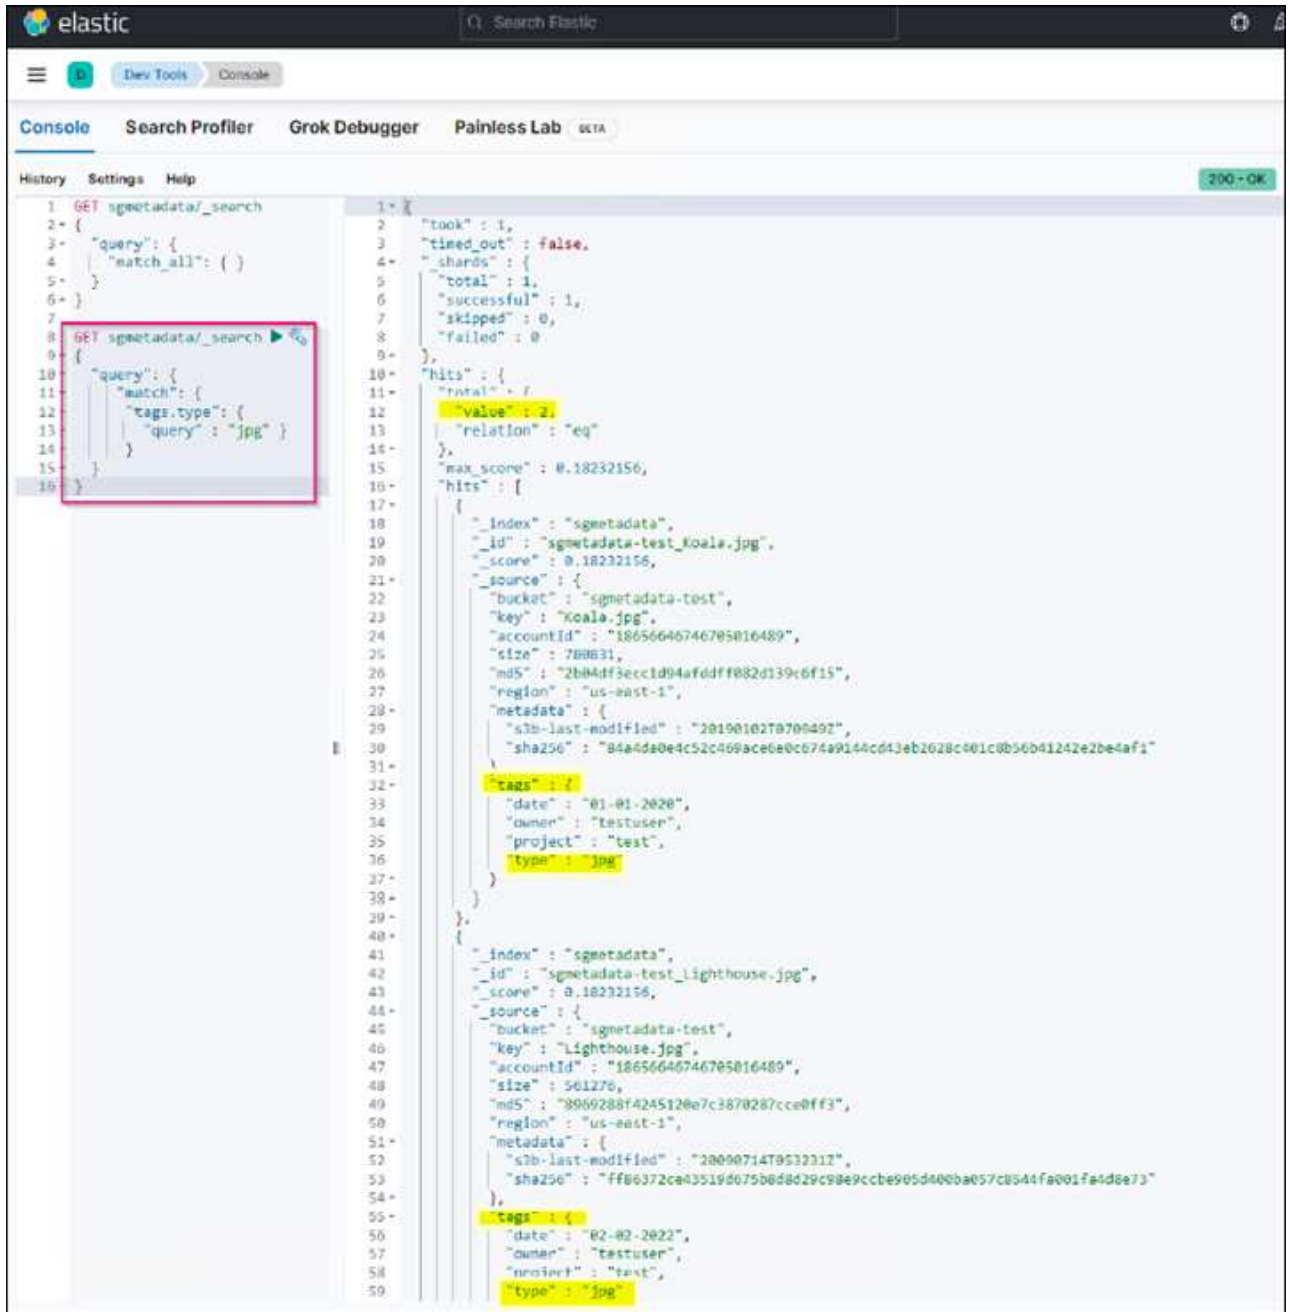
```
1 {
2   "took": 1,
3   "timed_out": false,
4   "_shards": {
5     "total": 1,
6     "successful": 1,
7     "skipped": 0,
8     "failed": 0
9   },
10  "hits": {
11    "total": {
12      "value": 4,
13      "relation": "eq"
14    },
15    "max_score": 1.0,
16    "hits": [
17      {
18        "_index": "sgmetadata",
19        "_id": "sgmetadata-test_test1.txt",
20        "_score": 1.0,
21        "_source": {
22          "bucket": "sgmetadata-test",
23          "key": "test1.txt",
24          "accountId": "18656646746705016489",
25          "size": 45,
26          "md5": "36b194a8ac536f09a7061f024b97211e",
27          "region": "us-east-1",
28          "metadata": {
29            "s3b-last-modified": "20170429T010249Z",
30            "sha256": "6bf95e898615852c94fa701580d9a0399487f4cbe4429e1a1d7d7f4270b10f51"
31          },
32          "tags": {
33            "owner": "testuser",
34            "project": "test"
35          }
36        }
37      },
38      {
39        "_index": "sgmetadata",
40        "_id": "sgmetadata-test_Koala.jpg",
41        "_score": 1.0,
42        "_source": {
43          "bucket": "sgmetadata-test",
44          "key": "Koala.jpg",
45          "accountId": "18656646746705016489",
46          "size": 780831,
47          "md5": "2b04df3ecc1d94afddff082d139c6f15",
48          "region": "us-east-1",
49          "metadata": {
50            "s3b-last-modified": "20190102T070949Z",
51            "sha256": "84adda0e4c52c409ace6e0c674a9144cd43eb2628c401c8b56b41242e2be4af1"
52          },
53          "tags": {
54            "date": "01-01-2020",
55            "owner": "testuser",
56            "project": "test",
57            "type": "jpg"
58          }
59        }
60      }
61    ]
62  }
63 }
```

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

```

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

```



The screenshot shows the Elastic Search Console interface. On the left, the 'Console' tab is active, displaying a search query. The query is highlighted with a red box:

```

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

```

On the right, the search results are displayed. The first result is highlighted with a yellow box:

```

{
  "_index": "sgmetadata",
  "_id": "sgmetadata-test_koala.jpg",
  "_score": 0.18232156,
  "_source": {
    "bucket": "sgmetadata-test",
    "key": "Koala.jpg",
    "accountId": "18656646746705016489",
    "size": 788631,
    "md5": "2b04df3eccc1d94afddff882d139c6f15",
    "region": "us-east-1",
    "metadata": {
      "s3b-last-modified": "20190102T070949Z",
      "sha256": "84a4da0e4c52c469ace0e0c674a9144cd13eb2628c001c0b56b41242e2be4af1"
    },
    "tags": {
      "date": "01-01-2020",
      "owner": "testuser",
      "project": "test",
      "type": "jpg"
    }
  }
}

```

The second result is also highlighted with a yellow box:

```

{
  "_index": "sgmetadata",
  "_id": "sgmetadata-test_lighthouse.jpg",
  "_score": 0.18232156,
  "_source": {
    "bucket": "sgmetadata-test",
    "key": "Lighthouse.jpg",
    "accountId": "18656646746705016489",
    "size": 561276,
    "md5": "8969288f4245120e7c3870287cce0ff3",
    "region": "us-east-1",
    "metadata": {
      "s3b-last-modified": "20090714T053231Z",
      "sha256": "ff06372ce43519d075b0d8d29c98e9ccbe965d400ba057c0544fa001fa4d8e73"
    },
    "tags": {
      "date": "02-02-2022",
      "owner": "testuser",
      "project": "test",
      "type": "jpg"
    }
  }
}

```

## Where to find additional information

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

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

*By Angela Cheng*

## Node Clone

Node clone considerations and performance.

### Node clone considerations

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

- The used capacity of the source node is not relevant to the time required for the clone process to complete. Node clone is a full copy of the node including free space in the node.
- The source and destination appliances must be at the same PGE version
- The destination node must always have larger capacity than the source
  - Make sure the new destination appliance has a larger drive size than the source
  - If the destination appliance has the same size drives and is configured for DDP8, you can configure the destination for DDP16. If the source is already configured for DDP16 then node clone will not be possible.
  - When going from SG5660 or SG5760 appliances to SG6060 appliances be aware that the SG5x60's have 60 capacity drives where the SG6060 only has 58.
- The node clone process requires the source node to be offline to the grid for the duration of the cloning process. If an additional node goes offline during this time client services may be impacted.
- 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.
- For a SG6060 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.

### Node clone Performance estimates

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

#### DDP8

SG5612 → Any

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

#### SG5712 → Any

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

#### SG5660 → SG5760

Network Interface speed	4TB Drive size	8TB Drive size	10TB Drive size	12TB Drive size	16TB Drive size	18TB Drive size
10GB	3 Day	6 Days	7 Days	8.5 Days	11.5 Days	<b>13 Days</b>
25GB	3 Day	6 Days	7 Days	8.5 Days	11.5 Days	<b>13 Days</b>

#### SG5660 → SG6060

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

#### SG5760 → SG5760

Network Interface speed	4TB Drive size	8TB Drive size	10TB Drive size	12TB Drive size	16TB Drive size	18TB Drive size
10GB	3 Day	6 Days	7 Days	8.5 Days	11.5 Days	<b>13 Days</b>
25GB	3 Day	6 Days	7 Days	8.5 Days	11.5 Days	<b>13 Days</b>

#### SG5760 → SG6060

Network Interface speed	4TB Drive size	8TB Drive size	10TB Drive size	12TB Drive size	16TB Drive size	18TB Drive size
10GB	2.5 Day	4.5 Days	5.5 Days	6.5 Days	9 Days	10 Days
25GB	1.5 Day	3 Days	3.5 Days	4.5 Days	6 Days	6.5 Days

#### SG6060 → SG6060

Network Interface speed	4TB Drive size	8TB Drive size	10TB Drive size	12TB Drive size	16TB Drive size	18TB Drive size
10GB	2.5 Day	4.5 Days	5.5 Days	6.5 Days	8.5 Days	9.5 Days
25GB	1.5 Day	3 Days	3.5 Days	4 Days	5.5 Days	6 Days

#### DDP16

#### SG5760 → SG5760

Network Interface speed	4TB Drive size	8TB Drive size	10TB Drive size	12TB Drive size	16TB Drive size	18TB Drive size
10GB	3.5 Day	6.5 Days	8 Days	9.5 Days	12.5 Days	<b>14 Days</b>
25GB	3.5 Day	6.5 Days	8 Days	9.5 Days	12.5 Days	<b>14 Days</b>

#### SG5760 → SG6060

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

#### SG6060 → SG6060

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

Expansion shelf (add to above SG6060 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
10GB	3.5 Day	5 Days	6 Days	7 Days	9.5 Days	10.5 Days
25GB	2 Day	3 Days	4 Days	4.5 Days	6 Days	7 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.

The screenshot shows the vSphere VM configuration interface. The 'Configure' tab is selected, and the 'vApp Options' sub-tab is active. The left sidebar lists various settings, with 'vApp Options' highlighted. The main content area shows 'OVF Settings' with a 'VIEW OVF ENVIRONMENT' button and an information icon. Below this, there are two rows of settings: 'OVF environment transport' with a value of 'VMware Tools', and 'Installation boot' with a value of 'Disabled'. At the bottom, there is a 'Properties' section with buttons for 'ADD', 'EDIT', 'SET VALUE', and 'DELETE'.

3. Scroll down to the Properties section and select the PORT\_REMAP\_INBOUND property

Summary	Monitor	Configure	Permissions	Datastores	Networks	Snapshots	Updates			
<div>Settings</div> <div>VM SDRS Rules</div> <div>vApp Options</div> <div>Alarm Definitions</div> <div>Scheduled Tasks</div> <div>Policies</div> <div>Guest User Mappings</div>		<input type="radio"/>	ADMIN_IP	Primary Admin IP	10.193.204.110		0.0.0.0	Grid Network (eth0)	ip	
		<input type="radio"/>	ADMIN_NETWORK_ESL	Admin network external subnet list					Admin Network (eth1)	string
		<input type="radio"/>	ADMIN_NETWORK_IP	Admin network IP	10.193.174.112		0.0.0.0		Admin Network (eth1)	ip
		<input type="radio"/>	NODE_TYPE	Node type			VM_API_Gateway		Grid Node Parameters	string["VM_Storage_Node", "VM_min_Node", "VM_API_Gateway", "_Archive_Node"]
		<input type="radio"/>	CLIENT_NETWORK_CONFIG	Client network IP configuration	STATIC		DISABLED		Client Network (eth2)	string["DISABLED", "STATIC", "DHCP"]
		<input checked="" type="radio"/>	PORT_REMAP_INBOUND	Inbound port remapping specification					Advanced	string
		<input type="radio"/>	GRID_NETWORK	Grid network IP configuration	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.

Edit property

Inbound port remapping specification... X

General

Type

Static property

Type

String

User configurable

☒

Length

0

-

65535

Default value

Dynamic property

Macro

IP address

Network

MGMT\_564


CANCEL

SAVE

- At the top of the Properties list, with the “PORT\_REMAP\_INBOUND” property still selected, click Set Value.



- 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.

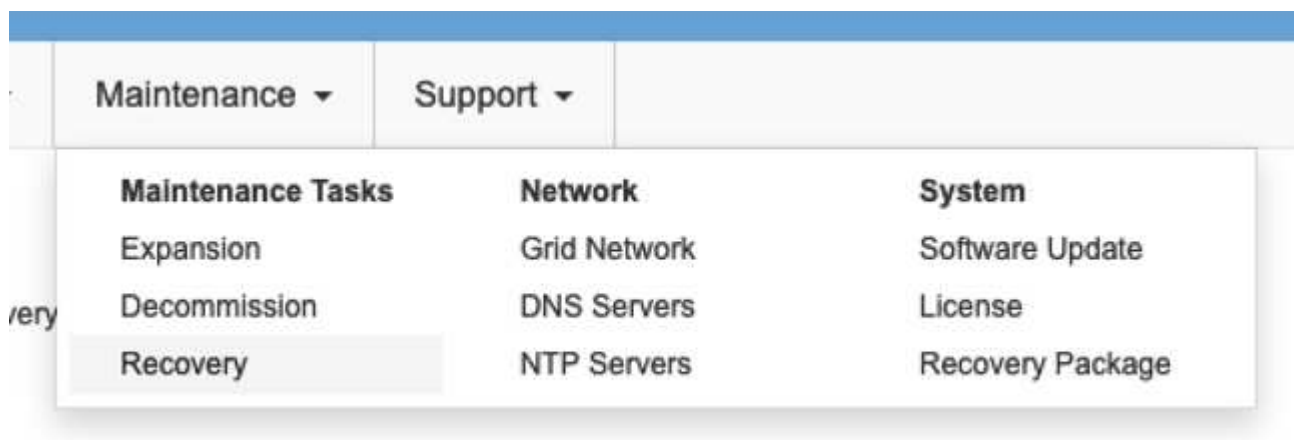


- If you are using multiple networks, use a comma (,) to separate the network strings, for example, grid/tcp/8082/10443,admin/tcp/8082/10443,client/tcp/8082/10443

#### Recover the Gateway Node

To recover the Gateway Node, follow these steps:

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



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.

#### Pending Nodes

Name	IPv4 Address	State	Recoverable
No results found.			



For information and directions for node recovery, see the <https://docs.netapp.com/sgws-114/topic/com.netapp.doc.sg-maint/GUID-7E22B1B9-4169-4800-8727-75F25FC0FFB1.html> [Recovery and Maintenance guide]

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 -l
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-audit-logs`

#### 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 use `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`

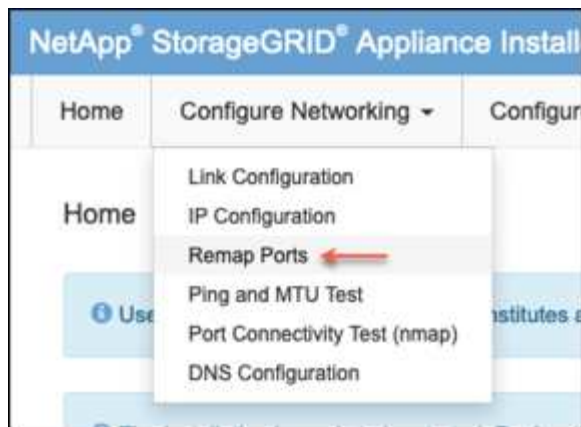
```
root@sg100-01:~ # sgareinstall
WARNING: All StorageGRID Webscale services on this node will be shut
down.
WARNING: Data stored on this node may be lost.
WARNING: You will have to reinstall StorageGRID Webscale to this
node.

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.
```

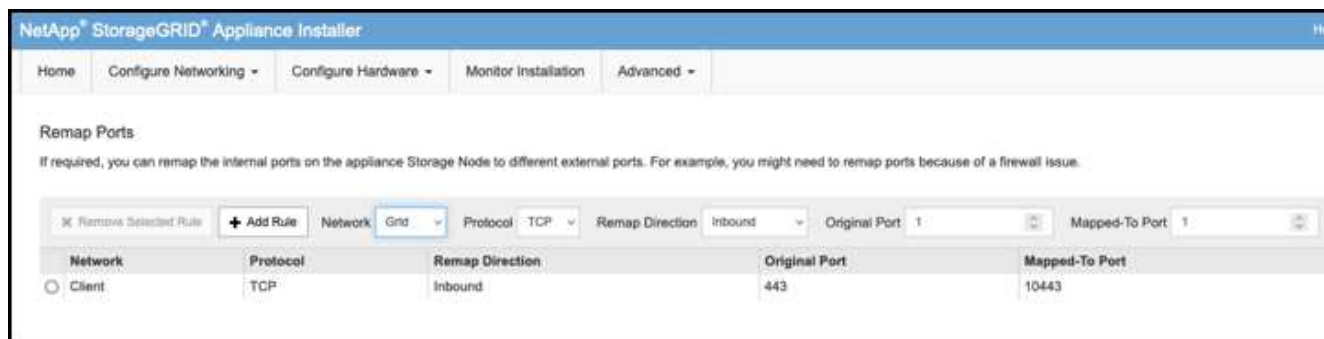
2. After some time has passed the appliance will reboot and you will be able to access the node PGE UI.
3. Browse to the Configure Networking



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



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



5. Once 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 ams-user: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 `Passwords.txt` 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*

## Copyright information

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

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

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

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

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

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

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

## Trademark information

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