



Technical Report

StorageGRID search integration service

Configuration guide

Angela Cheng, NetApp
April 2022 | TR-4922

Abstract

This technical report provides detailed instructions for configuring NetApp® StorageGRID® 11.6 search integration service with either Amazon OpenSearch Service or on-premises Elasticsearch.

TABLE OF CONTENTS

Introduction	3
Create tenant and enable platform services	3
Search integration services with Amazon OpenSearch	4
Amazon OpenSearch (formerly Elasticsearch) service setup.....	4
Platform services endpoint configuration.....	7
Search integration services with on premises Elasticsearch.....	9
On premises Elasticsearch setup	9
Platform services endpoint configuration.....	11
Bucket search integration service configuration	13
Where to find additional information	18
Version history.....	18

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.

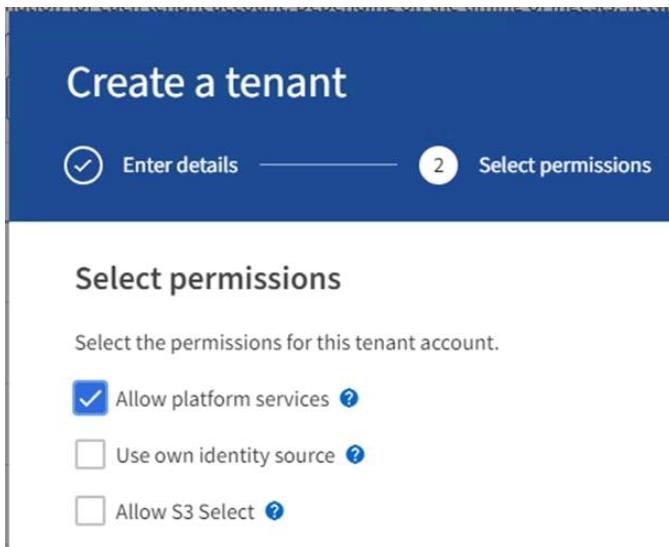
Platform services are configured by the S3 tenant through the Tenant Manager UI. For more information, see the section “Considerations for using platform services” in [Manage S3 platform services](#).

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

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

Create tenant and enable platform services

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



3. Set up the tenant root user initial password or, if identity federation is enabled on the grid, select which federated group has root access permission to configure the tenant account.
4. Click Sign In As Root and select Bucket: Create and Manage Buckets.
This takes you to the Tenant Manager page.
5. From Tenant Manager, select My Access Keys to create and download the S3 access key for later testing.

Search integration services with Amazon OpenSearch

Amazon OpenSearch (formerly Elasticsearch) service setup

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

Note: 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::210123456789:user/xyzabc"  
      },  
      "Action": "es:*",  
      "Resource": "arn:aws:es:us-east-1:210123456789:domain/sgdemo/*"  
    },  
    {  
      "Effect": "Allow",  
      "Principal": {  
        "AWS": "*"  
      },  
      "Action": [  
        "es:ESHttp*"  
      ],  
      "Condition": {  
        "IpAddress": {  
          "aws:SourceIp": [  
            "214.123.45.0/24"  
          ]  
        }  
      },  
      "Resource": "arn:aws:es:us-east-1:210123456789: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

i 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::222222222222:user/sgdemo"
8+     },
9+     "Action": "es:*",
10+    "Resource": "arn:aws:es:us-east-1:222222222222: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+          "216.222.222.222/24"
24+        ]
25+      }
26+    },
27+    "Resource": "arn:aws:es:us-east-1:222222222222:domain/sgdemo/*"
28+  }
```

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

Amazon OpenSearch Service > Domains > sgdemo

sgdemo [Info](#)

General information

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

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

OpenSearch Dashboards

☰ Dev Tools

Console

History Settings Help

1 `PUT sgmetadata` ▶ 🔍

```

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

```

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

The screenshot shows the Amazon OpenSearch Service console with the following details:

General information:

- Name: sgdemo
- Domain status: Active
- Domain ARN: arn:aws:es:us-east-1:211234567890:domain/sgdemo
- Cluster health: Yellow
- Service software version: R20211203-PS (latest)
- OpenSearch Dashboards URL: https://search-sgdemo-211234567890-us-east-1.es.amazonaws.com/_dashboards
- Domain endpoint: https://search-sgdemo-211234567890-us-east-1.es.amazonaws.com

Indices (2):

Index | Document count | Size (byte) | Query total | Mapping type | Field mappings

Index	Document count	Size (byte)	Query total	Mapping type	Field mappings
.kibana_1	1	5.08 KiB	10	dynamic,meta,proper...	0
sgmetadata	0	1.02 KiB	0		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/sgmetadata/_doc.

Create endpoint

1 Enter details 2 Select authentication type 3 Verify server

Enter endpoint details

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

Display name ?

aws-opensearch

URI ?

https://search-sgdemo-4...1.us-east-1.e

URN ?

ss:es:us-east-1:2...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

1 Enter details 2 Select authentication type 3 Verify server

Authentication type ?

Select the method used to authenticate connections to the endpoint.

Access Key

Access key ID ?

AKIA...UWO

Secret access key ?

.....



Previous

Continue

4. To verify the endpoint, select Use Operating System CA Certificate and Test and Create Endpoint. If verification is successful, an endpoint screen similar to the following figure displays. If verification

fails, verify that the URN includes /<index>/_doc at the end of the path and the AWS access key and secret key are correct.

Platform services endpoints

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



Display name	Last error	Type	URI	URN
aws-opensearch	Search	https://search-sgdemo-4299112111.us-east-1.es.amazonaws.com/	arn:aws:es:us-east-1:219011698108: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.

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

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.

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.



NOTE 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.

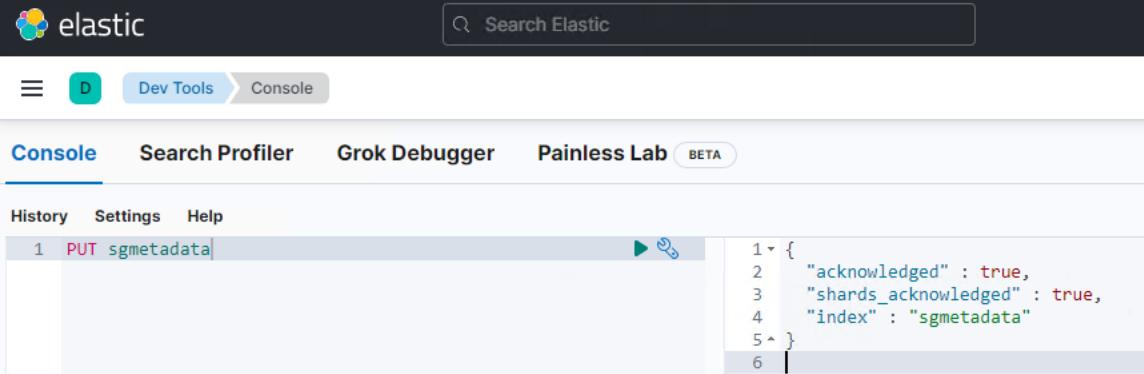
1. In a new terminal session, run:

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

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. Log in to the Kibana UI by using user name `elastic` and the password generated by Elastic in the preceding step.
4. For first time login, on the dashboard welcome page, select **Explore On Your Own**. From the menu, select **Management** → **Dev Tools**.
5. On the Dev Tools Console screen, enter `PUT <index>` where you use this index for storing StorageGRID object metadata. We use the index name `sgmetadata` in this example. Click the small triangle symbol to execute the `PUT` command. The expected result displays on the right panel as shown in the following example screenshot.



The screenshot shows the Elasticsearch Dev Tools Console interface. The top navigation bar includes the elastic logo, a search bar with 'Search Elastic', and tabs for 'Dev Tools' (which is selected and highlighted in green) and 'Console'. Below the tabs are 'Console', 'Search Profiler', 'Grok Debugger', and 'Painless Lab' tabs, with 'Painless Lab' being a BETA feature. The main area has tabs for 'History', 'Settings', and 'Help'. In the code editor, line 1 shows the command `PUT sgmetadata`. To the right, the response is displayed in a code block with line numbers 1 through 6. The response content is:

```

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

```

Platform services endpoint configuration

To configure endpoints for platform services, follow these steps:

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

Example: `urn:local:es:::sgmd/sgmetadata/_doc`



Enter endpoint details

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

Display name [?](#)

elasticsearch

URI [?](#)

https://10.10.10.10:9200

URN [?](#)

urn:local:es::sgmd/sgmetadata/_doc

[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 [?](#)

elastic

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.

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.



Display name	Type	URI	URN
aws-opensearch	Search	https://search-sgdemo-4w223hpljv6lzcxrpw3v3rte7i.us-east-1.es.amazonaws.com/	arn:aws:es:us-east-1:210811600188:domain/sgdemo/sgmetadata/_doc
elasticsearch	Search	https://10.0.0.10:9200/	urn:local:es:::sgmd/sgmetadata/_doc

Bucket search integration service configuration

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

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

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

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

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

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

Overview ^

Name:	sgmetadata-test
Region:	us-east-1
S3 Object Lock:	Disabled
Date created:	2022-03-18 20:09:35 EDT

[View bucket contents in Experimental S3 Console](#) ↗

Bucket options **Bucket access** **Platform services**

Replication Disabled ▼

Event notifications Disabled ▼

Search integration Enabled ^

Enable the search integration service to send object metadata to a destination search index whenever an object is created, deleted, or its metadata or tags are updated.

- Platform services must be enabled for your tenant account by a StorageGRID administrator.
- You must have already configured an endpoint for the search integration service.
- You must specify the URN of that endpoint in the search integration configuration XML for the bucket you want to index.

Enable search integration Clear

```

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

Save changes

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

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

URL: <https://10.193.204.106:10445/sgmetadatatest/Koala.jpg>

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

7. Use the Kibana UI to verify that the object metadata was loaded to sgmetadatas index.

- From the menu, select Management → Dev Tools.
- Paste the sample query to the console panel on the left and click the triangle symbol to execute it.

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

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

```

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

1 [
2   {
3     "took": 1,
4     "timed_out": false,
5     "_shards": {
6       "total": 1,
7       "successful": 1,
8       "skipped": 0,
9       "failed": 0
10    },
11    "hits": {
12      "total": {
13        "value": 4,
14        "relation": "eq"
15      },
16      "max_score": 1.0,
17      "hits": [
18        {
19          "_index": "sgmetadata",
20          "_id": "sgmetadata-test-test1.txt",
21          "_score": 1.0,
22          "source": {
23            "bucket": "sgmetadata-test",
24            "key": "test1.txt",
25            "accountId": "18656646746705016489",
26            "size": 45,
27            "md5": "36b194a8ac536f09a7061f024b97211e",
28            "region": "us-east-1",
29            "metadata": {
30              "s3b-last-modified": "20170429T010249Z",
31              "sha256": "6bf96e898615852c94fa701580d9a0399487f4cbe4429a1a1d7d7f427ab10f51"
32            },
33            "tags": {
34              "owner": "testuser",
35              "project": "test"
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": "84a4da0e4c52c469ace6e0c674a9144cd43eb2628c401c8b56b41242e2be4af1"
52              },
53              "tags": {
54                "date": "01-01-2020",
55                "owner": "testuser",
56                "project": "test",
57                "type": "jpg"
58              }
59            }
60          }
61        ]
62      }
63    }
64  ]
65 }

```

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"
      }
    }
  }
}

```

elastic Dev Tools Console

Console Search Profiler Grok Debugger Painless Lab BETA

History Settings Help 200 - OK

```
1 GET sgmetadata/_search
2 {
3   "query": {
4     "match_all": {}
5   }
6 }
7
8 GET sgmetadata/_search ►
9 {
10   "query": {
11     "match": {
12       "tags.type": {
13         "query" : "jpg"
14       }
15     }
16 }
```

17 {
18 "took" : 1,
19 "timed_out" : false,
20 "_shards" : {
21 "total" : 1,
22 "successful" : 1,
23 "skipped" : 0,
24 "failed" : 0
25 },
26 "hits" : {
27 "total" : 2,
28 "value" : 2,
29 "relation" : "eq"
30 },
31 "max_score" : 0.18232156,
32 "hits" : [
33 {
34 "_index" : "sgmetadata",
35 "_id" : "sgmetadata-test_Koala.jpg",
36 "_score" : 0.18232156,
37 "_source" : {
38 "bucket" : "sgmetadata-test",
39 "key" : "Koala.jpg",
40 "accountId" : "18656646746705016489",
41 "size" : 780831,
42 "md5" : "2b04df3ecc1d94afddff082d139c6f15",
43 "region" : "us-east-1",
44 "metadata" : {
45 "s3b-last-modified" : "20190102T070949Z",
46 "sha256" : "84a4da0e4c52c469ace6e0674a9144cd43eb2628c401c8b56b41242e2be4af1"
47 },
48 "tags" : {
49 "date" : "01-01-2020",
50 "owner" : "testuser",
51 "project" : "test",
52 "type" : "jpg"
53 }
54 },
55 {
56 "_index" : "sgmetadata",
57 "_id" : "sgmetadata-test_Lighthouse.jpg",
58 "_score" : 0.18232156,
59 "_source" : {
60 "bucket" : "sgmetadata-test",
61 "key" : "Lighthouse.jpg",
62 "accountId" : "18656646746705016489",
63 "size" : 561276,
64 "md5" : "8969288f4245120e7c3870287cce0ff3",
65 "region" : "us-east-1",
66 "metadata" : {
67 "s3b-last-modified" : "20090714T053231Z",
68 "sha256" : "ff86372ce43519d675b8d829c98e9ccbe905d400ba057c8544fa001fa4d8e73"
69 },
70 "tags" : {
71 "date" : "02-02-2022",
72 "owner" : "testuser",
73 "project" : "test",
74 "type" : "jpg"
75 }
76 }
77 }
78 }
79 }
80 }

Where to find additional information

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

- Manage S3 platform services—What are platform services?
<https://docs.netapp.com/us-en/storagegrid-116/tenant/what-platform-services-are.html>
- NetApp StorageGRID 11.6 Documentation
<https://docs.netapp.com/us-en/storagegrid-116/>

Version history

Version	Date	Document version history
Version 1.0	April 2022	Initial release— StorageGRID 11.6 search integration service.

Refer to the [Interoperability Matrix Tool \(IMT\)](#) on the NetApp Support site to validate that the exact product and feature versions described in this document are supported for your specific environment. The NetApp IMT defines the product components and versions that can be used to construct configurations that are supported by NetApp. Specific results depend on each customer's installation in accordance with published specifications.

Copyright information

Copyright © 2022 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—with prior written permission of the copyright owner.

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

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

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

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

Data contained herein pertains to a commercial item (as defined in FAR 2.101) and is proprietary to NetApp, Inc. The U.S. Government has a non-exclusive, non-transferable, non-sublicensable, worldwide, limited irrevocable license to use the Data only in connection with and in support of the U.S. Government contract under which the Data was delivered. Except as provided herein, the Data may not be used, disclosed, reproduced, modified, performed, or displayed without the prior written approval of NetApp, Inc. United States Government license rights for the Department of Defense are limited to those rights identified in DFARS clause 252.227-7015(b).

Trademark information

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.

TR-4922-0422