



Use search integration service

StorageGRID 11.8

NetApp
March 19, 2024

Table of Contents

- Use search integration service 1
- Configuration XML for search integration 1
- Configure the search integration service 5
- JSON generated by search integration service 7
- Object metadata included in metadata notifications 7

Use search integration service

The search integration service is one of the three StorageGRID platform services. You can enable this service to send object metadata to a destination search index whenever an object is created, deleted, or its metadata or tags are updated.

You can configure search integration by using the Tenant Manager to apply custom StorageGRID configuration XML to a bucket.



Because the search integration service causes object metadata to be sent to a destination, its configuration XML is referred to as *metadata notification configuration XML*. This configuration XML is different than the *notification configuration XML* used to enable event notifications.

See the [instructions for implementing S3 client applications](#) for details about the following custom StorageGRID S3 REST API operations:

- DELETE Bucket metadata notification configuration
- GET Bucket metadata notification configuration
- PUT Bucket metadata notification configuration

Related information

[Configuration XML for search integration](#)

[Object metadata included in metadata notifications](#)

[JSON generated by search integration service](#)

[Configure search integration service](#)

[Use S3 REST API](#)

Configuration XML for search integration

The search integration service is configured using a set of rules contained within `<MetadataNotificationConfiguration>` and `</MetadataNotificationConfiguration>` tags. Each rule specifies the objects that the rule applies to, and the destination where StorageGRID should send those objects' metadata.

Objects can be filtered on the prefix of the object name. For example, you could send metadata for objects with the prefix `images` to one destination, and metadata for objects with the prefix `videos` to another. Configurations that have overlapping prefixes aren't valid, and are rejected when they are submitted. For example, a configuration that includes one rule for objects with the prefix `test` and a second rule for objects with the prefix `test2` is not allowed.

Destinations must be specified using the URN of a StorageGRID endpoint that has been created for the search integration service. These endpoints refer to an index and type defined on an Elasticsearch cluster.

```

<MetadataNotificationConfiguration>
  <Rule>
    <ID>Rule-1</ID>
    <Status>rule-status</Status>
    <Prefix>key-prefix</Prefix>
    <Destination>
      <Urn>arn:aws:es:region:account-
ID:domain/mydomain/myindex/mytype</Urn>
    </Destination>
  </Rule>
  <Rule>
    <ID>Rule-2</ID>
    ...
  </Rule>
  ...
</MetadataNotificationConfiguration>

```

The table describes the elements in the metadata notification configuration XML.

Name	Description	Required
MetadataNotificationConfiguration	Container tag for rules used to specify the objects and destination for metadata notifications. Contains one or more Rule elements.	Yes
Rule	Container tag for a rule that identifies the objects whose metadata should be added to a specified index. Rules with overlapping prefixes are rejected. Included in the MetadataNotificationConfiguration element.	Yes
ID	Unique identifier for the rule. Included in the Rule element.	No
Status	Status can be 'Enabled' or 'Disabled'. No action is taken for rules that are disabled. Included in the Rule element.	Yes

Name	Description	Required
Prefix	<p>Objects that match the prefix are affected by the rule, and their metadata is sent to the specified destination.</p> <p>To match all objects, specify an empty prefix.</p> <p>Included in the Rule element.</p>	Yes
Destination	<p>Container tag for the destination of a rule.</p> <p>Included in the Rule element.</p>	Yes
Urn	<p>URN of the destination where object metadata is sent. Must be the URN of a StorageGRID endpoint with the following properties:</p> <ul style="list-style-type: none"> • es must be the third element. • The URN must end with the index and type where the metadata is stored, in the form domain-name/myindex/mytype. <p>Endpoints are configured using the Tenant Manager or Tenant Management API. They take the following form:</p> <ul style="list-style-type: none"> • arn:aws:es:region:account-ID:domain/mydomain/myindex/mytype • urn:mysite:es:::mydomain/myindex/mytype <p>The endpoint must be configured before the configuration XML is submitted, or configuration will fail with a 404 error.</p> <p>URN is included in the Destination element.</p>	Yes

Use the sample metadata notification configuration XML to learn how to construct your own XML.

Metadata notification configuration that applies to all objects

In this example, object metadata for all objects is sent to the same destination.

```

<MetadataNotificationConfiguration>
  <Rule>
    <ID>Rule-1</ID>
    <Status>Enabled</Status>
    <Prefix></Prefix>
    <Destination>
      <Urn>urn:myes:es:::sgws-notifications/test1/all</Urn>
    </Destination>
  </Rule>
</MetadataNotificationConfiguration>

```

Metadata notification configuration with two rules

In this example, object metadata for objects that match the prefix `/images` is sent to one destination, while object metadata for objects that match the prefix `/videos` is sent to a second destination.

```

<MetadataNotificationConfiguration>
  <Rule>
    <ID>Images-rule</ID>
    <Status>Enabled</Status>
    <Prefix>/images</Prefix>
    <Destination>
      <Urn>arn:aws:es:us-east-1:3333333:domain/es-
domain/graphics/imagetype</Urn>
    </Destination>
  </Rule>
  <Rule>
    <ID>Videos-rule</ID>
    <Status>Enabled</Status>
    <Prefix>/videos</Prefix>
    <Destination>
      <Urn>arn:aws:es:us-west-1:22222222:domain/es-
domain/graphics/videotype</Urn>
    </Destination>
  </Rule>
</MetadataNotificationConfiguration>

```

Related information

[Use S3 REST API](#)

[Object metadata included in metadata notifications](#)

[JSON generated by search integration service](#)

[Configure search integration service](#)

Configure the search integration service

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

Before you begin

- Platform services were enabled for your tenant account by a StorageGRID administrator.
- You have already created an S3 bucket whose contents you want to index.
- The endpoint that you intend to use as a destination for the search integration service already exists, and you have its URN.
- You belong to a user group that has the [Manage all buckets or Root access permission](#). These permissions override the permission settings in group or bucket policies when configuring the bucket using the Tenant Manager.

About this task

After you configure the search integration service for a source bucket, creating an object or updating an object's metadata or tags triggers object metadata to be sent to the destination endpoint. If you enable the search integration service for a bucket that already contains objects, metadata notifications aren't automatically sent for existing objects. You must update these existing objects to ensure that their metadata is added to the destination search index.

Steps

1. Use a text editor to create the metadata notification XML required to enable search integration.
 - See the information about configuration XML for search integration.
 - When configuring the XML, use the URN of a search integration endpoint as the destination.

```
<MetadataNotificationConfiguration>
  <Rule>
    <Status>Enabled</Status>
    <Prefix></Prefix>
    <Destination>
      <Urn>arn:aws:es:us-east-
1:1111111111111111:domain/mydomain/myindex/mytype</Urn>
    </Destination>
  </Rule>
</MetadataNotificationConfiguration>
```

2. In the Tenant Manager select **STORAGE (S3) > Buckets**.
3. Select the name of the source bucket.

The bucket details page appears.

4. Select **Platform services > Search integration**
5. Select the **Enable search integration** checkbox.
6. Paste the metadata notification configuration into the text box, and select **Save changes**.

Bucket options
Bucket access
Platform services

Replication
Disabled
▼

Event notifications
Disabled
▼

Search integration
Disabled
▲

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>
    <Status>Enabled</Status>
    <Prefix></Prefix>
    <Destination>
      <Urn>arn:aws:es:us-east-1:111111111111:domain/mydomain/myindex/mytype</Urn>
    </Destination>
  </Rule>
</MetadataNotificationConfiguration>

```

Save changes



Platform services must be enabled for each tenant account by a StorageGRID administrator using the Grid Manager or Management API. Contact your StorageGRID administrator if an error occurs when you save the configuration XML.

7. Verify that the search integration service is configured correctly:
 - a. Add an object to the source bucket that meets the requirements for triggering a metadata notification as specified in the configuration XML.

In the example shown earlier, all objects added to the bucket trigger a metadata notification.

- b. Confirm that a JSON document that contains the object's metadata and tags was added to the search index specified in the endpoint.

After you finish

As necessary, you can disable search integration for a bucket using either of the following methods:

- Select **STORAGE (S3) > Buckets** and clear the **Enable search integration** checkbox.
- If you are using the S3 API directly, use a DELETE Bucket metadata notification request. See the instructions for implementing S3 client applications.

Related information

[Understand search integration service](#)

[Configuration XML for search integration](#)

[Use S3 REST API](#)

[Create platform services endpoint](#)

JSON generated by search integration service

When you enable the search integration service for a bucket, a JSON document is generated and sent to the destination endpoint each time object metadata or tags are added, updated, or deleted.

This example shows an example of the JSON that could be generated when an object with the key `SGWS/Tagging.txt` is created in a bucket named `test`. The `test` bucket is not versioned, so the `versionId` tag is empty.

```
{
  "bucket": "test",
  "key": "SGWS/Tagging.txt",
  "versionId": "",
  "accountId": "86928401983529626822",
  "size": 38,
  "md5": "3d6c7634a85436eee06d43415012855",
  "region": "us-east-1",
  "metadata": {
    "age": "25"
  },
  "tags": {
    "color": "yellow"
  }
}
```

Object metadata included in metadata notifications

The table lists all the fields that are included in the JSON document that is sent to the destination endpoint when search integration is enabled.

The document name includes the bucket name, object name, and version ID if present.

Type	Item name and description
Bucket and object information	<code>bucket</code> : Name of the bucket
	<code>key</code> : Object key name
	<code>versionID</code> : Object version, for objects in versioned buckets
	<code>region</code> : Bucket region, for example <code>us-east-1</code>
System metadata	<code>size</code> : Object size (in bytes) as visible to an HTTP client
	<code>md5</code> : Object hash
User metadata	<code>metadata</code> : All user metadata for the object, as key-value pairs <code>key:value</code>
Tags	<code>tags</code> : All object tags defined for the object, as key-value pairs <code>key:value</code>



For tags and user metadata, StorageGRID passes dates and numbers to Elasticsearch as strings or as S3 event notifications. To configure Elasticsearch to interpret these strings as dates or numbers, follow the Elasticsearch instructions for dynamic field mapping and for mapping date formats. You must enable the dynamic field mappings on the index before you configure the search integration service. After a document is indexed, you can't edit the document's field types in the index.

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.