



Archive to the cloud through S3 API

StorageGRID 11.8

NetApp
March 19, 2024

Table of Contents

- Archive to the cloud through the S3 API 1
 - Configure connection settings for the S3 API 1
 - Modify connection settings for S3 API 3
 - Modify the Cloud Tiering Service state 4
 - Reset the Store Failure Count for S3 API connection 5
 - Migrate objects from Cloud Tiering - S3 to a Cloud Storage Pool 6

Archive to the cloud through the S3 API

You can configure an Archive Node to connect directly to Amazon Web Services (AWS) or to any other system that can interface to the StorageGRID system through the S3 API.

Support for Archive Nodes is deprecated and will be removed in a future release. Moving objects from an Archive Node to an external archival storage system through the S3 API has been replaced by ILM Cloud Storage Pools, which offer more functionality.



The Cloud Tiering - Simple Storage Service (S3) option is also deprecated. If you are currently using an Archive Node with this option, [migrate your objects to a Cloud Storage Pool](#) instead.

Additionally, you should remove Archive Nodes from the active ILM policy in StorageGRID 11.7 or earlier. Removing object data stored on Archive Nodes will simplify future upgrades. See [Working with ILM rules and ILM policies](#).

Configure connection settings for the S3 API

If you are connecting to an Archive Node using the S3 interface, you must configure the connection settings for the S3 API. Until these settings are configured, the ARC service remains in a Major alarm state as it is unable to communicate with the external archival storage system.

Support for Archive Nodes is deprecated and will be removed in a future release. Moving objects from an Archive Node to an external archival storage system through the S3 API has been replaced by ILM Cloud Storage Pools, which offer more functionality.



The Cloud Tiering - Simple Storage Service (S3) option is also deprecated. If you are currently using an Archive Node with this option, [migrate your objects to a Cloud Storage Pool](#) instead.

Additionally, you should remove Archive Nodes from the active ILM policy in StorageGRID 11.7 or earlier. Removing object data stored on Archive Nodes will simplify future upgrades. See [Working with ILM rules and ILM policies](#).

Before you begin

- You are signed in to the Grid Manager using a [supported web browser](#).
- You have [specific access permissions](#).
- You have created a bucket on the target archival storage system:
 - The bucket is dedicated to a single Archive Node. It can't be used by other Archive Nodes or other applications.
 - The bucket has the appropriate region selected for your location.
 - The bucket should be configured with versioning suspended.
- Object Segmentation is enabled and the Maximum Segment Size is less than or equal to 4.5 GiB (4,831,838,208 bytes). S3 API requests that exceed this value will fail if S3 is used as the external archival storage system.

Steps

1. Select **SUPPORT > Tools > Grid topology**.
2. Select **Archive Node > ARC > Target**.
3. Select **Configuration > Main**.

Target Type: Cloud Tiering - Simple Storage Service (S3)

Cloud Tiering (S3) Account

Bucket Name:

Region: Virginia or Pacific Northwest (us-east-1)

Endpoint: Use AWS

Endpoint Authentication:

Access Key:

Secret Access Key:

Storage Class: Standard (Default)

Apply Changes

4. Select **Cloud Tiering - Simple Storage Service (S3)** from the Target Type drop-down list.

Configuration settings are unavailable until you select a Target Type.

5. Configure the cloud tiering (S3) account through which the Archive Node will connect to the target external S3 capable archival storage system.

Most of the fields on this page are self-explanatory. The following describes fields for which you might need guidance.

- **Region:** Only available if **Use AWS** is selected. The region you select must match the bucket's region.
- **Endpoint** and **Use AWS:** For Amazon Web Services (AWS), select **Use AWS**. **Endpoint** is then automatically populated with an endpoint URL based on the Bucket Name and Region attributes. For example:

`https://bucket.region.amazonaws.com`

For a non-AWS target, enter the URL of the system hosting the bucket, including the port number. For example:

`https://system.com:1080`

- **End Point Authentication:** Enabled by default. If the network to the external archival storage system is trusted, you can clear the checkbox to disable endpoint SSL certificate and hostname verification for the targeted external archival storage system. If another instance of a StorageGRID system is the target archival storage device and the system is configured with publicly signed certificates, you can keep the checkbox selected.
- **Storage Class:** Select **Standard (Default)** for regular storage. Select **Reduced Redundancy** only for objects that can be easily recreated. **Reduced Redundancy** provides lower cost storage with less reliability. If the targeted archival storage system is another instance of the StorageGRID system, **Storage Class** controls how many interim copies of the object are made at ingest on the target system, if dual commit is used when objects are ingested there.

6. Select **Apply Changes**.

The specified configuration settings are validated and applied to your StorageGRID system. After the settings are applied, the target can't be changed.

Modify connection settings for S3 API

After the Archive Node is configured to connect to an external archival storage system through the S3 API, you can modify some settings should the connection change.

Before you begin

- You are signed in to the Grid Manager using a [supported web browser](#).
- You have [specific access permissions](#).

About this task

If you change the Cloud Tiering (S3) account, you must ensure that the user access credentials have read/write access to the bucket, including all objects that were previously ingested by the Archive Node to the bucket.

Steps

1. Select **SUPPORT > Tools > Grid topology**.
2. Select **Archive Node > ARC > Target**.
3. Select **Configuration > Main**.

Overview | Alarms | Reports | **Configuration**

Main | Alarms

 **Configuration: ARC (98-127) - Target**
 Updated: 2015-09-24 15:48:22 PDT

Target Type: Cloud Tiering - Simple Storage Service (S3)

Cloud Tiering (S3) Account

Bucket Name:

Region:


Endpoint: Use AWS

Endpoint Authentication:

Access Key:

Secret Access Key:

Storage Class:

Apply Changes 

4. Modify account information, as necessary.

If you change the storage class, new object data is stored with the new storage class. Existing object continue to be stored under the storage class set when ingested.



Bucket Name, Region, and Endpoint, use AWS values and can't be changed.

5. Select **Apply Changes**.

Modify the Cloud Tiering Service state

You can control the Archive Node's ability read and write to the targeted external archival storage system that connects through the S3 API by changing the state of the Cloud Tiering Service.

Before you begin


- You must be signed in to the Grid Manager using a [supported web browser](#).
- You have [specific access permissions](#).
- The Archive Node must be configured.

About this task

You can effectively take the Archive Node offline by changing the Cloud Tiering Service State to **Read-Write Disabled**.

Steps

1. Select **SUPPORT > Tools > Grid topology**.
2. Select **Archive Node > ARC**.
3. Select **Configuration > Main**.



4. Select a **Cloud Tiering Service State**.
5. Select **Apply Changes**.

Reset the Store Failure Count for S3 API connection


If your Archive Node connects to an archival storage system through the S3 API, you can reset the Store Failure Count, which can be used to clear the ARVF (Store Failures) alarm.

Before you begin

- You are signed in to the Grid Manager using a [supported web browser](#).
- You have [specific access permissions](#).

Steps

1. Select **SUPPORT > Tools > Grid topology**.
2. Select **Archive Node > ARC > Store**.
3. Select **Configuration > Main**.



4. Select **Reset Store Failure Count**.

5. Select **Apply Changes**.

The Store Failures attribute resets to zero.

Migrate objects from Cloud Tiering - S3 to a Cloud Storage Pool

If you are currently using the **Cloud Tiering - Simple Storage Service (S3)** feature to tier object data to an S3 bucket, you should migrate your objects to a Cloud Storage Pool instead. Cloud Storage Pools provide a scalable approach that takes advantage of all of the Storage Nodes in your StorageGRID system.

Before you begin

- You are signed in to the Grid Manager using a [supported web browser](#).
- You have [specific access permissions](#).
- You have already stored objects in the S3 bucket configured for Cloud Tiering.



Before migrating object data, contact your NetApp account representative to understand and manage any associated costs.

About this task

From an ILM perspective, a Cloud Storage Pool is similar to a storage pool. However, while storage pools consist of Storage Nodes or Archive Nodes within the StorageGRID system, a Cloud Storage Pool consists of an external S3 bucket.

Before migrating objects from Cloud Tiering - S3 to a Cloud Storage Pool, you must first create an S3 bucket and then create the Cloud Storage Pool in StorageGRID. Then, you can create a new ILM policy and replace the ILM rule used to store objects in the Cloud Tiering bucket with a cloned ILM rule that stores the same objects in the Cloud Storage Pool.



When objects are stored in a Cloud Storage Pool, copies of those objects can't also be stored within StorageGRID. If the ILM rule you are currently using for Cloud Tiering is configured to store objects in multiple locations at the same time, consider whether you still want to perform this optional migration because you will lose that functionality. If you continue with this migration, you must create new rules instead of cloning the existing ones.

Steps

1. Create a Cloud Storage Pool.

Use a new S3 bucket for the Cloud Storage Pool to ensure it contains only the data managed by the Cloud Storage Pool.

2. Locate any ILM rules in the active ILM policies that cause objects to be stored in the Cloud Tiering bucket.
3. Clone each of these rules.
4. In the cloned rules, change the placement location to the new Cloud Storage Pool.
5. Save the cloned rules.
6. Create a new policy that uses the new rules.

7. Simulate and activate the new policy.

When the new policy is activated and ILM evaluation occurs, the objects are moved from the S3 bucket configured for Cloud Tiering to the S3 bucket configured for the Cloud Storage Pool. The usable space on the grid is not affected. After the objects are moved to the Cloud Storage Pool, they are removed from the Cloud Tiering bucket.

Related information

[Manage objects with ILM](#)

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.