



Creating an ILM rule

StorageGRID 11.5

NetApp
August 30, 2024

Table of Contents

- Creating an ILM rule 1
 - Step 1 of 3: Define basics 2
 - Step 2 of 3: Define placements 7
 - Step 3 of 3: Define ingest behavior 14
- Creating a default ILM rule 15

Creating an ILM rule

ILM rules allow you to manage the placement of object data over time. To create an ILM rule, you use the Create ILM Rule wizard.

Before you begin

- You must be signed in to the Grid Manager using a supported browser.
- You must have specific access permissions.
- If you want to specify which tenant accounts this rule applies to, you must have the Tenant Accounts permission or you must know the account ID for each account.
- If you want the rule to filter objects on last access time metadata, Last Access Time updates must be enabled by bucket for S3 or by container for Swift.
- If you are creating replicated copies, you must have configured any storage pools or Cloud Storage Pools you plan to use.
- If you are creating erasure-coded copies, you must have configured an Erasure Coding profile.
- You must be familiar with the [data-protection options for ingest](#).
- If you need to create a compliant rule for use with S3 Object Lock, you must be familiar with the [requirements for S3 Object Lock](#).



To create the default ILM rule for a policy, use this procedure instead: [Creating a default ILM rule](#).

About this task

When creating ILM rules:

- Consider the StorageGRID system's topology and storage configurations.
- Consider what types of object copies you want to make (replicated or erasure coded) and the number of copies of each object that are required.
- Determine what types of object metadata are used in the applications that connect to the StorageGRID system. ILM rules filter objects based on their metadata.
- Consider where you want object copies to be placed over time.
- Decide which option to use for data protection option at ingest (Balanced, Strict, or Dual commit)

Steps

1. Select **ILM > Rules**.

The ILM Rules page appears, with the stock rule, Make 2 Copies, selected.

ILM Rules

Information lifecycle management (ILM) rules determine how and where object data is stored over time. Every object ingested into StorageGRID is evaluated against the ILM rules that make up the active ILM policy. Use this page to manage and view ILM rules. You cannot edit or remove an ILM rule that is used by an active or proposed ILM policy.

+ Create Clone Edit Remove

Name	Used In Active Policy	Used In Proposed Policy
Make 2 Copies	✓	

Make 2 Copies

Ingest Behavior: Dual commit
Reference Time: Ingest Time
Filtering Criteria:
Matches all objects.

Retention Diagram:

The diagram shows two horizontal bars representing retention periods. The top bar is labeled 'All Storage Nodes' and starts at 'Day 0'. The bottom bar is also labeled 'All Storage Nodes' and starts at 'Day 0'. Both bars extend to the right, ending at 'Forever'. The x-axis is labeled 'Duration'.



The ILM Rules page looks slightly different if the global S3 Object Lock setting has been enabled for the StorageGRID system. The summary table includes a **Compliant** column, and the details for the selected rule include a **Compliant** field.

2. Select **Create**.

Step 1 (Define Basics) of the Create ILM Rule wizard appears. You use the Define basics page to define which objects the rule applies to.

Related information

[Use S3](#)

[Use Swift](#)

[Configuring Erasure Coding profiles](#)

[Configuring storage pools](#)

[Using Cloud Storage Pools](#)

[Data-protection options for ingest](#)

[Managing objects with S3 Object Lock](#)

Step 1 of 3: Define basics

Step 1 (Define Basics) of the Create ILM Rule wizard allows you to define the rule's basic and advanced filters.

About this task

When evaluating an object against an ILM rule, StorageGRID compares the object metadata to the rule's filters. If the object metadata matches all filters, StorageGRID uses the rule to place the object. You can design a rule to apply to all objects, or you can specify basic filters, such as one or more tenant accounts or bucket

names, or advanced filters, such as the object's size or user metadata.

Create ILM Rule Step 1 of 3: Define Basics

Name

Description

Tenant Accounts (optional)

Bucket Name Value

[Advanced filtering...](#) (0 defined)

Steps

1. Enter a unique name for the rule in the **Name** field.

You must enter between 1 and 64 characters.

2. Optionally, enter a short description for the rule in the **Description** field.

You should describe the rule's purpose or function so you can recognize the rule later.

Name

Description

3. Optionally, select one or more S3 or Swift tenant accounts to which this rule applies. If this rule applies to all tenants, leave this field blank.

If you do not have either the Root Access permission or the Tenant Accounts permission, you cannot select tenants from the list. Instead, enter the tenant ID or enter multiple IDs as a comma-delimited string.

4. Optionally, specify the S3 buckets or Swift containers to which this rule applies.

If **matches all** is selected (default), the rule applies to all S3 buckets or Swift containers.

5. Optionally, select **Advanced filtering** to specify additional filters.

If you do not configure advanced filtering, the rule applies to all objects that match the basic filters.



If this rule will create erasure-coded copies, select **Advanced filtering**. Then, add the **Object Size (MB)** advanced filter and set it to **greater than 0.2**. The size filter ensures that objects that are 2 MB or smaller will not be erasure coded.

6. Select **Next**.

Step 2 (Define Placements) appears.

Related information

Using advanced filters in ILM rules

Advanced filtering allows you to create ILM rules that apply only to specific objects based on their metadata. When you set up advanced filtering for a rule, you select the type of metadata you want to match, select an operator, and specify a metadata value. When objects are evaluated, the ILM rule is applied only to those objects that have metadata matching the advanced filter.

The table shows the types of metadata you can specify in advanced filters, the operators you can use for each type of metadata, and the metadata values expected.

Metadata type	Supported operators	Metadata value
Ingest Time (microseconds)	<ul style="list-style-type: none">• equals• does not equal• less than• less than or equals• greater than• greater than or equals	Time and date the object was ingested. Note: To avoid resource issues when activating a new ILM policy, you can use the Ingest Time advanced filter in any rule that might change the location of large numbers of existing objects. Set Ingest Time to be greater than or equal to the approximate time when the new policy will go into effect to ensure that existing objects are not moved unnecessarily.
Key	<ul style="list-style-type: none">• equals• does not equal• contains• does not contain• starts with• does not start with• ends with• does not end with	All or part of a unique S3 or Swift object key. For example, you might want to match objects that end with <code>.txt</code> or start with <code>test-object/</code> .

Metadata type	Supported operators	Metadata value
Last Access Time (microseconds)	<ul style="list-style-type: none"> • equals • does not equal • less than • less than or equals • greater than • greater than or equals • exists • does not exist 	<p>Time and date the object was last retrieved (read or viewed).</p> <p>Note: If you plan to use last access time as an advanced filter, Last Access Time updates must be enabled for the S3 bucket or Swift container.</p> <p>Using Last Access Time in ILM rules</p>
Location Constraint (S3 only)	<ul style="list-style-type: none"> • equals • does not equal 	<p>The region where an S3 bucket was created. Use ILM > Regions to define the regions that are shown.</p> <p>Note: A value of us-east-1 will match objects in buckets created in the us-east-1 region as well as objects in buckets that have no region specified.</p> <p>Configuring regions (optional and S3 only)</p>
Object Size (MB)	<ul style="list-style-type: none"> • equals • not equals • less than • less than or equals • greater than • greater than or equals 	<p>The object's size in MB.</p> <p>To filter on object sizes smaller than 1 MB, type in a decimal value. For example, set the Object Size (MB) advanced filter to greater than 0.2 for any rule that makes erasure-coded copies. This setting ensures that erasure coding is not used for objects 200 KB or smaller.</p> <p>Note: Your browser type and locale settings control whether you need to use a period or a comma as the decimal separator.</p>
User Metadata	<ul style="list-style-type: none"> • contains • ends with • equals • exists • does not contain • does not end with • does not equal • does not exist • does not start with • starts with 	<p>Key-value pair, where User Metadata Name is the key and User Metadata Value is the value.</p> <p>For example, to filter on objects that have user metadata of <code>color=blue</code>, specify <code>color</code> for User Metadata Name, <code>equals</code> for the operator, and <code>blue</code> for User Metadata Value.</p> <p>Note: User-metadata names are not case sensitive; user-metadata values are case sensitive.</p>

Metadata type	Supported operators	Metadata value
Object Tag (S3 only)	<ul style="list-style-type: none"> contains ends with equals exists does not contain does not end with does not equal does not exist does not start with starts with 	<p>Key-value pair, where Object Tag Name is the key and Object Tag Value is the value.</p> <p>For example, to filter on objects that have an object tag of Image=True, specify Image for Object Tag Name, equals for the operator, and True for Object Tag Value.</p> <p>Note: Object tag names and object tag values are case sensitive. You must enter these items exactly as they were defined for the object.</p>

Specifying multiple metadata types and values

When you define advanced filtering, you can specify multiple types of metadata and multiple metadata values. For example, if you want a rule to match objects between 10 MB and 100 MB in size, you would select the **Object Size** metadata type and specify two metadata values.

- The first metadata value specifies objects greater than or equal to 10 MB.
- The second metadata value specifies objects less than or equal to 100 MB.

Advanced Filtering

Use advanced filtering if you want a rule to apply only to specific objects. You can filter objects based on their system metadata, user metadata, or object tags (S3 only). When objects are evaluated, the rule is applied if the object's metadata matches the criteria in the advanced filter.

Objects between 10 and 100 MB

Matches all of the following metadata:

Object Size (MB)	▼	greater than or equals	▼	10	⊞ ⊞
Object Size (MB)	▼	less than or equals	▼	100	⊞ ⊞
+ ×					

Cancel

Remove Filters

Save

Using multiple entries allows you to have precise control over which objects are matched. In the following example, the rule applies to objects that have a Brand A or Brand B as the value of the camera_type user metadata. However, the rule only applies to those Brand B objects that are smaller than 10 MB.

Advanced Filtering

Use advanced filtering if you want a rule to apply only to specific objects. You can filter objects based on their system metadata, user metadata, or object tags (S3 only). When objects are evaluated, the rule is applied if the object's metadata matches the criteria in the advanced filter.

Multiple filters

Matches all of the following metadata:

User Metadata	camera_type	equals	Brand A	+	x
---------------	-------------	--------	---------	---	---

+

+ x |

Or matches all of the following metadata:

User Metadata	camera_type	equals	Brand B	+	x
Object Size (MB)		less than or equals	10	+	x

+

+ x |

Cancel Remove Filters Save

Related information

[Using Last Access Time in ILM rules](#)

[Configuring regions \(optional and S3 only\)](#)

Step 2 of 3: Define placements

Step 2 (Define Placements) of the Create ILM Rule wizard allows you to define the placement instructions that determine how long objects are stored, the type of copies (replicated or erasure coded), the storage location, and the number of copies.

About this task

An ILM rule can include one or more placement instructions. Each placement instruction applies to a single period of time. When you use more than one instruction, the time periods must be contiguous, and at least one instruction must start on day 0. The instructions can continue either forever, or until you no longer require any object copies.

Each placement instruction can have multiple lines if you want to create different types of copies or use different locations during that time period.

This example ILM rule creates two replicated copies for the first year. Each copy is saved in a storage pool at a different site. After one year, a 2+1 erasure-coded copy is made and saved at only one site.

Configure placement instructions to specify how you want objects matched by this rule to be stored.

Example rule
 Two copies for one year, then EC forever

Reference Time Ingest Time

Placements Sort by start day

From day 0 store for 365 days Add Remove

Type replicated Location DC1 x DC2 x Add Pool Copies 2 + x

Specifying multiple storage pools might cause data to be stored at the same site if the pools overlap. See [Managing objects with information lifecycle management](#) for more information.

From day 365 store forever Add Remove

Type erasure coded Location DC1 (2 plus 1) Copies 1 + x

Retention Diagram Refresh

The diagram shows a timeline starting at 'Trigger' (Day 0) and ending at 'Forever'.

- At Day 0, two copies are created: DC1 (blue bar) and DC2 (orange bar).
- At Year 1, the DC1 and DC2 copies are replaced by a single erasure-coded copy (DC1 (2 plus 1), orange bar with a pencil icon).
- The duration of the rule is marked as '1 years' and 'Forever'.

Cancel Back Next

Steps

1. For **Reference Time**, select the type of time to use when calculating the start time for a placement instruction.

Option	Description
Ingest Time	The time when the object was ingested.
Last Access Time	The time when the object was last retrieved (read or viewed). Note: To use this option, updates to Last Access Time must be enabled for the S3 bucket or Swift container. Using Last Access Time in ILM rules

Option	Description
Noncurrent Time	<p>The time an object version became noncurrent because a new version was ingested and replaced it as the current version.</p> <p>Note: Noncurrent Time applies only to S3 objects in versioning-enabled buckets.</p> <p>You can use this option to reduce the storage impact of versioned objects by filtering for noncurrent object versions. See “Example 4: ILM rules and policy for S3 versioned objects.”</p>
User Defined Creation Time	A time specified in user-defined metadata.



If you want to create a compliant rule, you must select **Ingest Time**.

- In the **Placements** section, select a starting time and a duration for the first time period.

For example, you might want to specify where to store objects for the first year (“day 0 for 365 days”). At least one instruction must start at day 0.

- If you want to create replicated copies:

- From the **Type** drop-down list, select **replicated**.
- In the **Location** field, select **Add Pool** for each storage pool you want to add.

If you specify only one storage pool, be aware that StorageGRID can store only one replicated copy of an object on any given Storage Node. If your grid includes three Storage Nodes and you select 4 as the number of copies, only three copies will be made—one copy for each Storage Node.



The **ILM placement unachievable** alert is triggered to indicate that the ILM rule could not be completely applied.

If you specify more than one storage pool, keep these rules in mind:

- The number of copies cannot be greater than the number of storage pools.
- If the number of copies equals the number of storage pools, one copy of the object is stored in each storage pool.
- If the number of copies is less than the number of storage pools, the system distributes the copies to keep disk usage among the pools balanced, while ensuring that no site gets more than one copy of an object.
- If the storage pools overlap (contain the same Storage Nodes), all copies of the object might be saved at only one site. For this reason, do not specify the default All Storage Nodes storage pool and another storage pool.

Placements ⓘ ⇅ Sort by start day

From day store

Type Location Copies

Specifying multiple storage pools might cause data to be stored at the same site if the pools overlap. See [Managing objects with information lifecycle management](#) for more information.

c. Select the number of copies you want to make.

A warning appears if you change the number of copies to 1. An ILM rule that creates only one replicated copy for any time period puts data at risk of permanent loss. If only one replicated copy of an object exists during a time period, that object is lost if a Storage Node fails or has a significant error. You also temporarily lose access to the object during maintenance procedures such as upgrades.



To avoid these risks, do one or more of the following:

- Increase the number of copies for the time period.
- Click the plus sign icon **+** to create additional copies during the time period. Then, select a different storage pool or a Cloud Storage Pool.
- Select **erasure coded** for Type, instead of **replicated**. You can safely ignore this warning if this rule already creates multiple copies for all time periods.

d. If you specified only one storage pool, ignore the **Temporary location** field.



Temporary locations are deprecated and will be removed in a future release.

4. If you want to store objects in a Cloud Storage Pool:

- a. From the **Type** drop-down list, select **replicated**.
- b. In the **Location** field, select **Add Pool**. Then, select a Cloud Storage Pool.

When using Cloud Storage Pools, keep these rules in mind:

- You cannot select more than one Cloud Storage Pool in a single placement instruction. Similarly, you cannot select a Cloud Storage Pool and a storage pool in the same placement instruction.

If you want to use a Cloud Storage Pool, you must remove any other storage pools or Cloud Storage Pools from this placement instruction.

- You can store only one copy of an object in any given Cloud Storage Pool. An error message appears if you set **Copies** to 2 or more.

Type Location Add Pool Copies

The number of copies cannot be more than one when a Cloud Storage Pool is selected.


- You cannot store more than one object copy in any Cloud Storage Pool at the same time. An error message appears if multiple placements that use a Cloud Storage Pool have overlapping dates or if multiple lines in the same placement use a Cloud Storage Pool.

Placements Sort by start day

From day store for days Add Remove

Type <input type="text" value="replicated"/>	Location <input type="text" value="csp1"/> Add Pool	Copies <input type="text" value="1"/>	+ x
Type <input type="text" value="replicated"/>	Location <input type="text" value="csp2"/> Add Pool	Copies <input type="text" value="1"/>	+ x

A rule cannot store more than one object copy in any Cloud Storage Pool at the same time. You must remove one of the Cloud Storage Pools (csp1, csp2) or use multiple placement instructions with dates that do not overlap. **Overlapping days:** 0-10.
To see the overlapping days on the Retention Diagram, click Refresh.



- You can store an object in a Cloud Storage Pool at the same time that object is being stored as replicated or erasure coded copies in StorageGRID. However, as this example shows, you must include more than one line in the placement instruction for the time period, so you can specify the number and types of copies for each location.

Placements

From day store for days

Type <input type="text" value="replicated"/>	Location <input type="text" value="DC1"/> <input type="text" value="DC2"/> Add Pool	Copies <input type="text" value="2"/>
Type <input type="text" value="replicated"/>	Location <input type="text" value="testpool2"/> Add Pool	Copies <input type="text" value="1"/>

- If you want to create an erasure-coded copy:
 - From the **Type** drop-down list, select **erasure coded**.

The number of copies changes to 1. A warning appears if the rule does not have an advanced filter to ignore objects that are 200 KB or smaller.

Do not use erasure coding for objects that are 200 KB or smaller. Select **Back** to return to Step 1. Then, use **Advanced filtering** to set the Object Size (MB) filter to "greater than 0.2".



Do not use erasure coding for objects smaller than 200 KB to avoid the overhead of managing very small erasure-coded fragments.

b. If the object size warning appeared, follow these steps to clear it:

- i. Select **Back** to return to Step 1.
- ii. Select **Advanced filtering**.
- iii. Set the Object Size (MB) filter to “greater than 0.2”.

c. Select the storage location.

The storage location for an erasure-coded copy includes the name of the storage pool, followed by the name of the Erasure Coding profile.

6. Optionally, add different time periods or create additional copies at different locations:

- Click the plus icon to create additional copies at a different location during the same time period.
- Click **Add** to add a different time period to the placement instructions.




Objects are automatically deleted at the end of the final time period unless the final time period ends with **forever**.

7. Click **Refresh** to update the Retention Diagram and to confirm your placement instructions.

Each line in the diagram shows where and when object copies will be placed. The type of copy is represented by one of the following icons:

	Replicated copy
	Erasure-coded copy
	Cloud Storage Pool copy

In this example, two replicated copies will be saved to two storage pools (DC1 and DC2) for one year. Then, an erasure-coded copy will be saved for an additional 10 years, using a 6+3 erasure-coding scheme at three sites. After 11 years, the objects will be deleted from StorageGRID.



8. Click **Next**.

Step 3 (Define Ingest Behavior) appears.

Related information

[What ILM rule placement instructions are](#)

[Example 4: ILM rules and policy for S3 versioned objects](#)

[Why you should not use single-copy replication](#)

[Managing objects with S3 Object Lock](#)

[Using a storage pool as a temporary location \(deprecated\)](#)

[Step 3 of 3: Define ingest behavior](#)

Using Last Access Time in ILM rules

You can use Last Access Time as the reference time in an ILM rule. For example, you might want to leave objects that have been viewed in the last three months on local Storage Nodes, while moving objects that have not been viewed as recently to an off-site location. You can also use Last Access Time as an advanced filter if you want an ILM rule to apply only to objects that were last accessed on a specific date.

About this task

Before using Last Access Time in an ILM rule, review the following considerations:

- When using Last Access Time as a reference time, be aware that changing the Last Access Time for an object does not trigger an immediate ILM evaluation. Instead, the object's placements are assessed and the object is moved as required when background ILM evaluates the object. This could take two weeks or more after the object is accessed.

Take this latency into account when creating ILM rules based on Last Access Time and avoid placements that use short time periods (less than one month).

- When using Last Access Time as an advanced filter or as a reference time, you must enable last access time updates for S3 buckets. You can use the Tenant Manager or the Tenant Management API.



Last access time updates are always enabled for Swift containers, but are disabled by default for S3 buckets.



Be aware that enabling last access time updates can reduce performance, especially in systems with small objects. The performance impact occurs because StorageGRID must update the objects with new timestamps every time the objects are retrieved.

The following table summarizes whether the Last Access Time is updated for all objects in the bucket for different types of requests.

Type of request	Whether Last Access Time is updated when last access time updates are disabled	Whether Last Access Time is updated when last access time updates are enabled
Request to retrieve an object, its access control list, or its metadata	No	Yes
Request to update an object's metadata	Yes	Yes
Request to copy an object from one bucket to another	<ul style="list-style-type: none"> • No, for the source copy • Yes, for the destination copy 	<ul style="list-style-type: none"> • Yes, for the source copy • Yes, for the destination copy
Request to complete a multipart upload	Yes, for the assembled object	Yes, for the assembled object

Related information

[Use S3](#)

[Use a tenant account](#)

Step 3 of 3: Define ingest behavior

Step 3 (Define ingest behavior) of the Create ILM Rule wizard allows you to choose how the objects filtered by this rule are protected as they are ingested.

About this task

StorageGRID can make interim copies and queue the objects for ILM evaluation later, or it can make copies to meet the rule's placement instructions immediately.

Create ILM Rule Step 3 of 3: Define ingest behavior

Select the data protection option to use when objects are ingested:

- Strict
Always uses this rule's placements on ingest. Ingest fails when this rule's placements are not possible.
- Balanced
Optimum ILM efficiency. Attempts this rule's placements on ingest. Creates interim copies when that is not possible.
- Dual commit
Creates interim copies on ingest and applies this rule's placements later.

Cancel Back Save

Steps

1. Select the data protection option to use when objects are ingested:

Option	Description
Strict	Always uses this rule's placements on ingest. Ingest fails when this rule's placements are not possible.

Option	Description
Balanced	Optimum ILM efficiency. Attempts this rule's placements on ingest. Creates interim copies when that is not possible.
Dual commit	Creates interim copies on ingest and applies this rule's placements later.

Balanced offers a combination of data security and efficiency that is suitable in most cases. Strict or Dual commit are generally used to meet specific requirements.

See “What the data-protection options for ingest are” and “Advantages and disadvantages of each data-protection option” for more information.



An error message appears if you select the Strict or Balanced option and the rule uses one of these placements:

- A Cloud Storage Pool at day 0
- An Archive Node at day 0
- A Cloud Storage Pool or an Archive Node when the rule uses a User Defined Creation Time as a Reference Time

2. Click **Save**.

The ILM rule is saved. The rule does not become active until it is added to an ILM policy and that policy is activated.

Related information

[Data-protection options for ingest](#)

[Advantages, disadvantages, and limitations of the data-protection options](#)

[Example 5: ILM rules and policy for Strict ingest behavior](#)

[Creating an ILM policy](#)

Creating a default ILM rule

Every ILM policy must have a default rule that does not filter objects. Before creating an ILM policy, you must create at least one ILM rule that can be used as the default rule for the policy.

What you'll need

- You must be signed in to the Grid Manager using a supported browser.
- You must have specific access permissions.

About this task

The default rule is the last rule to be evaluated in an ILM policy, so it cannot use any filters. The placement instructions for the default rule are applied to any objects that are not matched by another rule in the policy.

In this example policy, the first rule applies only to objects belonging to Tenant A. The default rule, which is last, applies to objects belonging to all other tenant accounts.

+ Select Rules			
Default	Rule Name	Tenant Account	Actions
	Erasure Coding for Tenant A	Tenant A (94793396288150002349)	✘
✓	2 Copies 2 Data Centers	Ignore	✘

When you create the default rule, keep these requirements in mind:

- The default rule is automatically placed as the last rule in the policy.
- The default rule cannot use any basic or advanced filters.
- The default rule should create replicated copies.



Do not use a rule that creates erasure-coded copies as the default rule for a policy. Erasure-coding rules should use an advanced filter to prevent smaller objects from being erasure coded.

- In general, the default rule should retain objects forever.
- If you are using (or you plan to enable) the global S3 Object Lock setting, the default rule for the active or proposed policy must be compliant.

Steps

1. Select **ILM > Rules**.

The ILM Rules page appears.

2. Select **Create**.

Step 1 (Define Basics) of the Create ILM Rule wizard appears.

3. Enter a unique name for the rule in the **Name** field.
4. Optionally, enter a short description for the rule in the **Description** field.
5. Leave the **Tenant Accounts** field blank.

The default rule must apply to all tenant accounts.

6. Leave the **Bucket Name** field blank.

The default rule must apply to all S3 buckets and Swift containers.

7. Do not select **Advanced filtering**

The default rule cannot specify any filters.

8. Select **Next**.

Step 2 (Define Placements) appears.

9. Specify the placement instructions for the default rule.

- The default rule should retain objects forever. A warning appears when you activate a new policy if the default rule does not retain objects forever. You must confirm this is the behavior you expect.
- The default rule should create replicated copies.



Do not use a rule that creates erasure-coded copies as the default rule for a policy. Erasure-coding rules should include the **Object Size (MB) greater than 0.2** advanced filter to prevent smaller objects from being erasure coded.

- If you are using (or you plan to enable) the global S3 Object Lock setting, the default rule must be compliant:
 - It must create at least two replicated object copies or one erasure-coded copy.
 - These copies must exist on Storage Nodes for the entire duration of each line in the placement instructions.
 - Object copies cannot be saved in a Cloud Storage Pool.
 - Object copies cannot be saved on Archive Nodes.
 - At least one line of the placement instructions must start at day 0, using Ingest Time as the reference time.
 - At least one line of the placement instructions must be “forever.”

10. Click **Refresh** to update the Retention Diagram and to confirm your placement instructions.

11. Click **Next**.

Step 3 (Define Ingest Behavior) appears.

12. Select the data protection option to use when objects are ingested, and select **Save**.

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.