



Technical Report

## **NetApp StorageGRID with Rubrik CDM configuration guide**

### **Implementing NetApp StorageGRID as an archival location for Rubrik backups**

Jonathan Wong and Steve Pruchniewski, NetApp  
Ben Kendall, Rubrik  
September 2021 | TR-4812

In partnership with



## **Abstract**

This report describes the configuration, best practices, and guidelines for NetApp® StorageGRID® as an archive endpoint for Rubrik.

## TABLE OF CONTENTS

<b>Introduction .....</b>	<b>3</b>
Scope .....	3
Rubrik Cloud Data Management .....	3
NetApp StorageGRID .....	3
Reasons to choose StorageGRID .....	3
<b>Requirements .....</b>	<b>4</b>
Rubrik .....	5
StorageGRID .....	5
<b>Configuring StorageGRID with Rubrik .....</b>	<b>5</b>
Configuring an S3 tenant on StorageGRID .....	5
Configuring StorageGRID as an archival location .....	7
Configuring VM backup and restore .....	9
Creating an SLA .....	9
Selecting a VM to archive to StorageGRID .....	12
Creating an ILM rule to protect the Rubrik archive data .....	13
Multisite configuration (optional) .....	16
<b>Configuration recommendations .....</b>	<b>18</b>
<b>Where to find additional information .....</b>	<b>19</b>
<b>Version history .....</b>	<b>19</b>

## LIST OF FIGURES

Figure 1) NetApp StorageGRID with Rubrik as archival location.....	4
--	---

# Introduction

## Scope

Rubrik Cloud Data Management (CDM) is a backup and recovery solution for enterprise workloads. As customers set out to modernize their data management and data protection practices, Rubrik's CDM platform is an increasingly attractive solution. For longer-term backup retention, data can be archived from Rubrik to the public cloud or to object-based storage such as NetApp® StorageGRID®.

This document is a reference for configuring NetApp StorageGRID as an archival target. It also covers configuring a backup policy in Rubrik, configuring a S3 tenant in StorageGRID, and an information lifecycle management (ILM) policy in StorageGRID, as well as sizing and performance guidelines for StorageGRID with Rubrik.

The recommendations in this document are suggested guidelines. When designing a solution, it is important to consider as many aspects of the backup environment as possible.

This document is intended for NetApp employees and partners who are familiar with StorageGRID concepts and terminology, and who also have technical familiarity with backup solutions.

## Rubrik Cloud Data Management

Rubrik's CDM platform organizes and protects data on the premises, at the edge, and in the cloud. Users can easily automate backup jobs by configuring SLAs and archiving data to public or private clouds, including NetApp StorageGRID object-based storage.

Rubrik scales linearly, so customers are not hampered by forklift upgrades. Deduplication, compression, and other data services scale in line with the cluster to maximize efficiency and savings. Furthermore, Rubrik simplifies data recovery by allowing "Google-like" search functionality to recover only the database, applications, or files needed.

In addition, Rubrik is built on an API-first architecture that is fully featured and easy to learn, allowing customers to automate their data workflows.

## NetApp StorageGRID

NetApp StorageGRID is a software-defined object storage solution that supports industry-standard object APIs such as Amazon S3 API and OpenStack Swift API. StorageGRID uses intelligent, policy-driven data management to store, protect, and preserve data, enabling you to create metadata-driven object lifecycle policies to optimize durability, performance, cost, and location across multiple geographies.

In addition, NetApp StorageGRID is built as a scale-out, node-based architecture that gives you the agility to increase capacity and performance on demand across your sites. Scaling is asymmetrical and as simple as adding more storage nodes, which are available as an appliance, software only/bare metal, or virtual machine.

## Reasons to choose StorageGRID

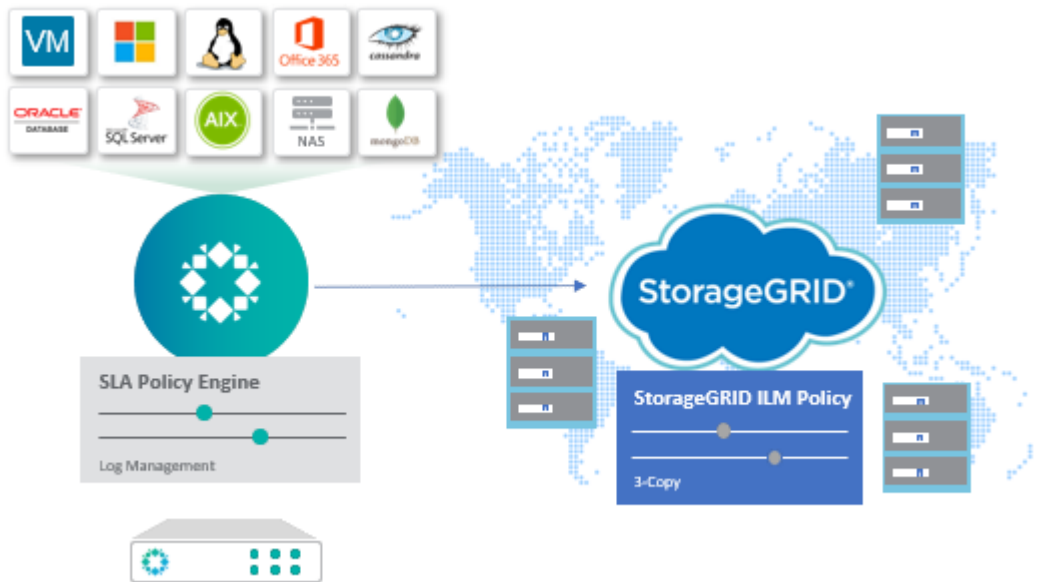
By reducing complexity and providing flexible, policy-driven data management, Rubrik delivers simplicity and efficiency. At the same time, given the popularity and ever-expanding use case for object storage, customers are choosing NetApp StorageGRID for similar reasons: policy-driven data management, simplicity, and flexibility.

NetApp and Rubrik are delivering newly integrated solutions that offer policy-based simplicity, cost-effective scale, and cloud mobility.

StorageGRID makes a superior archive endpoint in the following scenarios:

- You operate in a hybrid cloud model to optimize costs and avoid vendor lock-in. You can deploy locally or globally with no proprietary hardware lock in.
- You want a system that can grow organically. Refresh, expand, and migrate non-disruptively without the dependencies that other systems have.
- You require granular data protection to comply with data sovereignty and other regulatory compliance requirements. Dynamically balance data durability, performance, cost, and location with the industry's leading lifecycle policy engine.
- You want to archive data and leverage object storage for other use cases. Easily stand up other workflows with our multi-tenant capabilities and hybrid cloud capabilities.

**Figure 1) NetApp StorageGRID with Rubrik as archival location.**



## Requirements

Although StorageGRID and Rubrik have been implemented in the field as far back as StorageGRID 10.4 and Rubrik 2.3, this document validates the current General Availability (GA) versions:

- StorageGRID 11.5.0
- Rubrik CDM 6.0.0-p1-12566

As prerequisites, the following must be configured on StorageGRID:

- FQDN for S3 endpoint; for example, s3.company.com
- API service endpoint server certificates
- Commercial trusted SSL is preferred
- Self-signed is acceptable but requires Rubrik Support to enable
- S3 tenant created with S3 keys (Swift API is not supported for this solution.)
- Load balancer endpoint
- Disable grid encryption and compression (default setting). Rubrik already sends data encrypted and compressed.

## Rubrik

Rubrik can be deployed on a hardware appliance, software at the edge, software in the cloud, and on third-party industry platforms.

For this guide, a minimum configuration was deployed on VMware:

- One Rubrik Edge

## StorageGRID

StorageGRID can be deployed on a hardware appliance, software only/bare-metal servers, or a hypervisor. For this test, a minimum configuration was deployed on a VMware hypervisor:

- One admin node
- Three storage nodes

## HTTP and HTTPS support

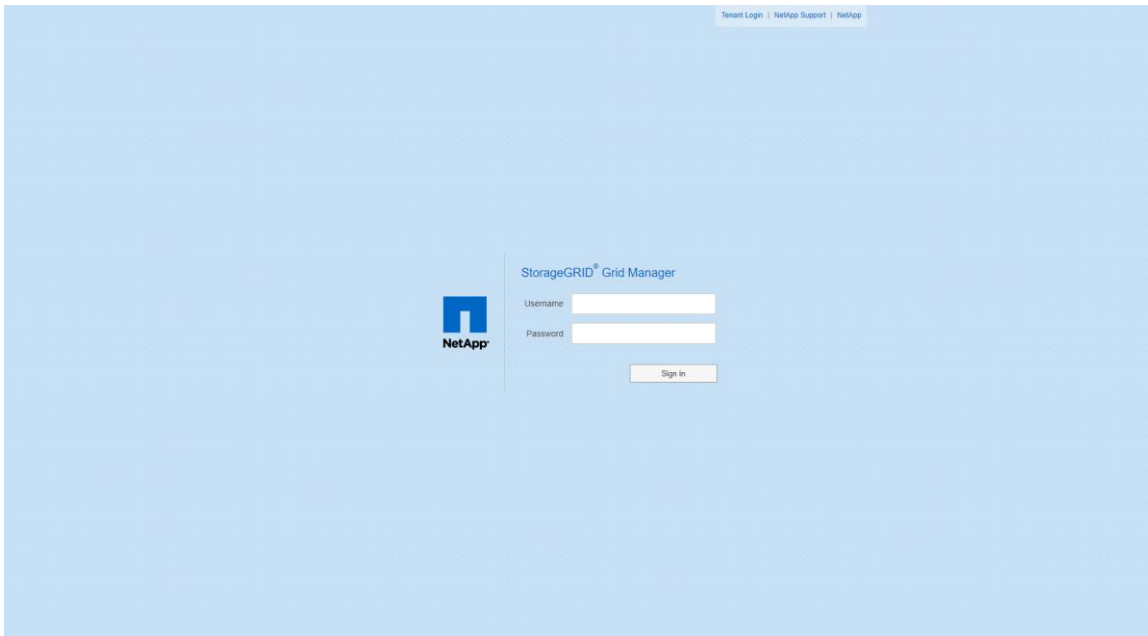
StorageGRID supports both HTTP and HTTPS. HTTPS is enabled by default. HTTP is recommended only for lab environments and is disabled by default.

Customers can choose to use standard ports with a third-party load balancer or create a load-balancer endpoint (11.3 and later). For details, see [Configuring Load Balancer Endpoints](#) in the [StorageGRID 11.5 Documentation Center](#).

# Configuring StorageGRID with Rubrik

## Configuring an S3 tenant on StorageGRID

1. Log in to the StorageGRID Grid Manager Interface.



2. Click Tenants on the top menu bar and select + Create.
3. Create a tenant account:

- a. Enter a display name.
- b. Select S3 as the protocol.
- c. Enter a password for the tenant's local root user.

### Tenant Account

**Details**

Display Name

Protocol  S3  Swift

Allow Platform Services

Storage Quota (optional)

---

**Authentication**

Determine how the tenant account will be accessed.

Uses Own Identity Source

---

Specify a password for the tenant's local root user.

Username

Password

Confirm Password


4. After the tenant account is created, select Sign In next to the tenant to access the Tenant Manager Interface (TMI) and log in.

5. In the TMI, go to S3 → My Access Keys, and click Create Key.
6. Step through the process and be sure to download your Access Key and Secret Access Key.

The StorageGRID S3 Access and Secret Access Keys are now ready to be used with Rubrik.

## Configuring StorageGRID as an archival location

1. Log in to the Rubrik Cluster using an account with Admin permission.



2. Click the cog symbol at the top right in the Rubrik GUI. Under System Configuration, select Archival Locations.
3. Click the + to create a new archival location and then follow these configuration steps:
  - a. Set Archival Type as Object Store and set Object Store Vendor as S3 Compatible (StorageGRID).

- b. Enter an S3 Access Key and Secret Access Key.
- c. Set Host Name to the FQDN of the S3 endpoint. If not using standard HTTPS port 443, enter the port of the load balancer endpoint:
- d. Enter a bucket prefix, for example, sg.
- e. Rubrik creates multiple buckets named `prefix-rubrik-x`.
- f. Select the number of buckets to be used.
- g. Rubrik stores all data for a source (VM, database, etc.) in a single bucket along with the metadata that allows Rubrik to validate and perform recovery. NetApp recommends creating no more than one bucket per Rubrik Archive.
- h. Generate an RSA key for encryption.
  - Run the following command on a secure computer that has the OpenSSL toolkit:

```
openssl genrsa -out rubrik_encryption_key.pem 2048
```

- Paste the RSA key into the window.
- Rubrik uses this key to encrypt the archival data.

- i. Click Add.

### Add Archival Location

Archive Type  
Object Store

Object Store Vendor  
S3 Compatible (StorageGRID, Cloudian, IBM COS, or other compatible object storage)

Access Key  
SYAQWPG7RB65CFMR0ZU

Secret Key  
.....

Host Name  
[REDACTED] 43:1044

Bucket Prefix  
rb

Number of Buckets  
1

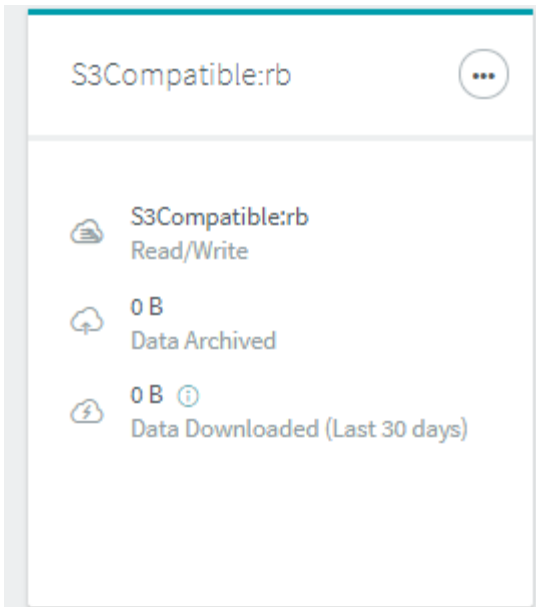
Archival Location Name  
S3Compatible:rb

RSA Key ⓘ  
-----BEGIN RSA PRIVATE KEY-----  
MIIEowIBAAKCAQEAyofxgBF2zbs4FXS8KzuoG4YQdW0EkPoU0YR2v5g++yIDR6hwy  
J3a9QglRlt5QFdy4VMPaKP1vHASG+6hloVkfzJarsGdjdtj9jQ9P03je6BvRv7tc  
aRM6ZnWvVwZl6yXzkMA8aPn94+xo5WXqfO1GelTrTmXoV5LB4o+5QxTYA4pi3Ua8  
0LJQPt9cxosjC+KIAAVQRtKB35kAL4ePGEbgellTolmsXf07EV7J7UUqSht2P42l  
szc6/0SD1pJorpFiSa7mJ2V7gCV4jqmK2Ay3awK+Uim9PvcvMVjKktmUUy2lbpyR  
SxMEE4N/2t8acyE0jpd83GjZDCphCJ/QYcOQwIDAQABAoIBABPHMs5O5kswWsm9  
KXMMF6B...  
-----END RSA PRIVATE KEY-----

Cancel Add

4. The StorageGRID S3 object store is added to the Rubrik Cluster as an archival location and is available to be used by Rubrik SLAs.





## Configuring VM backup and restore

Rubrik supports backup for different enterprise workloads such as databases and virtual machines. Rubrik also supports different hypervisors – VMware vSphere (ESXi), Microsoft Hyper-V, and Nutanix AHV. For this test, we connected Rubrik to a vCenter Server for VM backup.

1. Click the cog symbol at the top right in the Rubrik GUI to add the vCenter Servers.
2. In the Rubrik GUI click, the + to add vCenter and then follow these configuration steps:
  - a. Enter a vCenter IP or FQDN.
  - b. Enter a vCenter username and password.

3. When the vCenter has been added, Rubrik makes an inventory of VMs, hosts, and folders.

## Creating an SLA

Rubrik SLAs define protection levels for workloads composed of snapshot protection and retention, replication, and archiving. For this test, we configured an SLA targeted at archiving to StorageGRID.

1. From the left-side menu, select SLA Domains > Local Domains.
2. In the Rubrik GUI, click the + to create a new SLA domain and then follow these configuration steps:

- a. Enter an SLA domain name.
- b. Specify snapshot protection and retention.

### Create SLA Domain

1 Set Frequency and Retention
2 Set Archiving and Replication (Optional)
3 Summary

SLA Domain Name  
Archive

Continuous Data Protection

Service Level Agreement  
Choose how often we take snapshots and the length of time we keep them.

Advanced Frequencies

Take Snapshots:	Keep Snapshots:
Every (Hours) _____	For (Days) _____
Every (Days) <u>1</u> _____	For (Days) <u>7</u> _____
Every (Months) <u>1</u> _____	For (Months) <u>13</u> _____
Every (Years) <u>1</u> _____	For (Years) <u>2</u> _____

Local retention set to 2 years.

Snapshot Window

Take snapshots from: \_\_\_\_\_ : \_\_\_\_\_ to \_\_\_\_\_ : \_\_\_\_\_

Take first full between: First Opportunity at \_\_\_\_\_ : \_\_\_\_\_

Cancel
Next

3. Click Next.
4. Enable Archiving and select the configured StorageGRID archival target.

### Create SLA Domain

1 Set Frequency and Retention — 2 Set Archiving and Replication (Optional) — 3 Summary

Retention On Brik

0 2 years 2 years

Archiving

S3Compatible:rb  Enable Instant Archive ⓘ

Instant Archive sends policy-based snapshots to the archival location as soon as they are processed and retains them at the archival location for the maximum retention period specified by the associated policy.

Replication

A replication target has not been set up yet. Please [add a replication target](#) to configure retention.

Cancel Back Next

5. Verify the settings for SLA frequency, retention, and archiving.

## Create SLA Domain

✓ Set Frequency and Retention
 ✓ Set Archiving and Replication (Optional)
 3 Summary

---

SLA Domain Name
Archive

---

### Frequency and Retention

Take every	Retain for
1 day	7 days
1 month on the Last day of the month	13 months
1 year on the Last day of the year in January	2 years

---

Local Retention	2 years
-----------------	---------

---

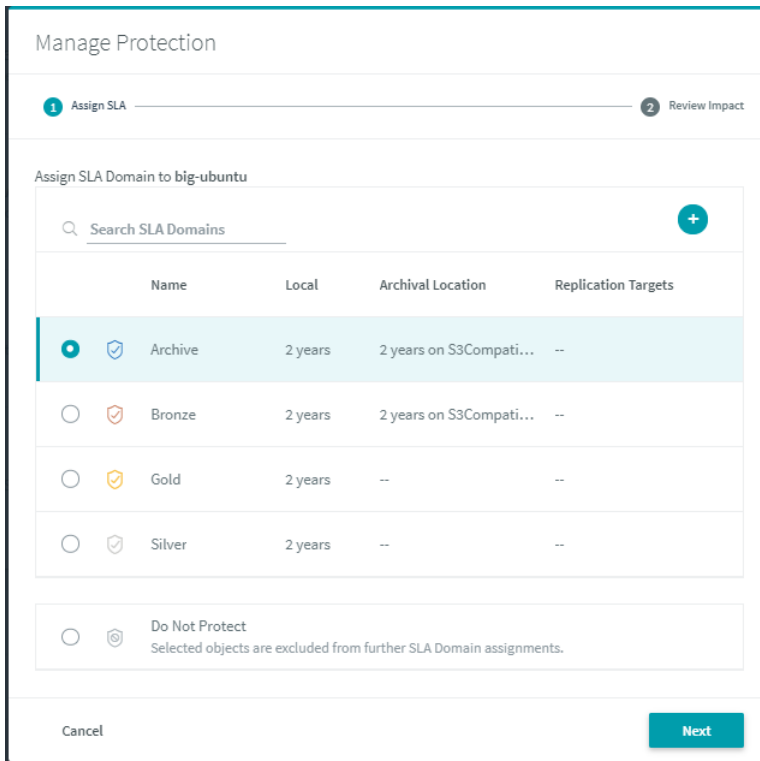
### Archiving

Location	S3Compatible:rb
Threshold	Instant
Retention	2 years

Cancel
Back
Create

## Selecting a VM to archive to StorageGRID

1. From the left-side menu, select Virtual Machines → vSphere VMs.
2. In the Rubrik UI, select a VM and click Manage Protection.
3. Select an SLA domain and click Next.



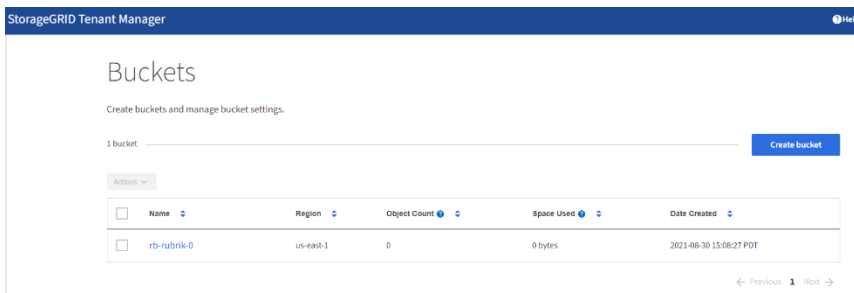
4. Review the protection settings and click Submit

## Creating an ILM rule to protect the Rubrik archive data

In StorageGRID, an ILM rule contains the instructions for placing objects in the system over time. Rules perform actions based on matching criteria such as bucket name or user-specified metadata fields.

In this test, we replicated our archival data to three sites and retained the replications for a year before transitioning to erasure coding.

1. Log in to the StorageGRID Grid Management Interface (GMI).
2. When StorageGRID is configured as a Rubrik archival location, buckets are automatically created based on the Rubrik configurations set in section 3.2. To see the buckets, follow these steps.
  - a. In the StorageGRID GMI, go to Tenants → Sign In (next to the Rubrik Tenant) to access the Tenant Manager.
  - b. Enter the Rubrik tenant root account and password.
  - c. Go to S3 → Buckets.



3. Create an ILM rule to manage the data. In the StorageGRID GMI, go to ILM → ILM Rules → + Create.
  - a. Enter a name.
  - b. Enter a description.
  - c. Specify criteria that the Tenant Account must be the configured Rubrik tenant.
  - d. Specify criteria that the Bucket Name must start with `prefix-rubrik`. (Prefix configured in section 3.2.)
  - e. Click Next.

Create ILM Rule Step 1 of 3: Define Basics

Name

Description

Tenant Accounts (optional)

Bucket Name

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

4. Example configuration of an ILM placement rule:
  - a. Replicate three copies across three data centers for geo distribution.
  - b. Transition to erasure coding after a year to reduce storage space.
  - c. Verify that the retention diagram is correct.
  - d. Click Save.

## Create ILM Rule Step 2 of 2: Define Placements

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

### Archive Tier

Data set from Rubrik

Reference Time

**Placements** Sort by start day

From day  store for  days Add Remove

Type  Location  Copies  + -

Specifying multiple storage pools might cause data to be stored at the same site if the pools overlap. See the Administrator Guide for more information.

From day  store forever Add Remove

Type  Location  Copies  + -

**Retention Diagram** Refresh

The diagram shows a timeline from Day 0 to Forever. At Day 0, three data centers (DC1, DC2, DC3) are triggered. DC1 is represented by a blue bar, DC2 by an orange bar, and DC3 by a grey bar. All three bars extend to Year 1. At Year 1, an erasure-coded pool (EC Pool (4plus2)) is triggered, represented by an orange bar with a right-pointing arrow, extending to Forever. The x-axis is labeled 'Duration' with markers for 'Day 0', 'Year 1', and 'Forever'.

Cancel Back Save

5. Go to ILM → Policies → + Create Proposed Policy:
  - a. Enter a name.
  - b. Enter a reason for the change.
  - c. Select the rules to apply.
  - d. Click Apply.
  - e. Select a default rule.
  - f. Click Save.

## Configure ILM Policy

Create a proposed policy by selecting and arranging rules. Then, save the policy and edit it later as required. Click Simulate to verify a saved policy using test objects. When you are ready, click Activate to make this policy the active ILM policy for the grid.

Name

Reason for change

### Rules

Select the rules you want to add to the policy. Drag and drop rows to reorder the rules. Rules are evaluated in order, starting at the top.

	Default	Rule Name	Tenant Account	Actions
	<input type="radio"/>	Archive Tier	Rubrik (81801872353536787046)	
	<input checked="" type="radio"/>	Make 2 Copies	Ignore	

## 6. Activate the policy.

Viewing Proposed Policy - Archive Tier from Rubrik

Errors in an ILM policy can cause irreparable data loss. Review and simulate the policy carefully before activating.

Review the rules in this policy. If this is a proposed policy, click Simulate to verify the policy and then click Activate to make the policy active.

Reason for change: Include Archive Tier ILM Rule

Rules are evaluated in order, starting from the top.

Rule Name	Default	Tenant Account
Archive Tier	<input type="radio"/>	Rubrik (81801872353536787046)
Make 2 Copies	<input checked="" type="radio"/>	Ignore

## Multisite configuration (optional)

It is a common deployment architecture to have a primary data center and a secondary data center both deploying Rubrik and StorageGRID. The secondary data center ensures backups can still be archived and restored during a full primary data center failover. To create a basic configuration for a two-site deployment, follow these steps:

1. Deploy a two-site StorageGRID grid with a primary Admin Node in the primary data center and a non-primary Admin Node in the secondary data center.



11.5 2sites test grid

SITE1

- ✓ SITE1-ADM1
- ✓ SITE1-S1
- ✓ SITE1-S2
- ✓ SITE1-S3
- ✓ SITE1-S4

SITE2

- ✓ SITE2-ADM1
- ✓ SITE2-S1
- ✓ SITE2-S2
- ✓ SITE2-S3
- ✓ SITE2-S4

The Admin Nodes act as the S3 endpoint to provide load balancing to the Storage Nodes. S3 endpoints can also be provided by dedicated Gateway Nodes.

- To enable connectivity to StorageGRID during data center failures, configure a [high availability group](#). This ties the two S3 endpoints into an active-backup group with a virtual IP (VIP).

High Availability Groups

High availability (HA) groups allow multiple nodes to participate in an active-backup group. HA groups maintain virtual IP addresses on the active node and switch to a backup node automatically if a node fails.

Name	Description	Virtual IP Addresses	Interfaces
○ site1-HA	site1-HA	10.193.205.43	SITE1-ADM1.eth2 (preferred Master) SITE2-ADM1.eth2

- The VIP is the IP address that Rubrik connects to when configuring the archival location. During a primary data center loss, the VIP reroutes to the S3 endpoint in the secondary data center transparently.
- Make sure that the backups archived to StorageGRID are replicated into both primary and secondary data centers. That way, during a data center failure, the object data is still available.
- Create two [storage pools](#); one for the primary data center and another for the secondary. Storage pools are logical groupings of nodes used to define object placement.
  - Go to ILM → Storage Pools → + Create
  - Follow the steps and create a storage pool for each data center.

Node Name	Site Name	Used (%)
SITE1-S3	SITE1	0.449%
SITE1-S4	SITE1	0.401%
SITE1-S2	SITE1	0.393%
SITE1-S1	SITE1	0.312%

Storage Pool Details - site2		
Nodes Included		ILM Usage
Number of Nodes: 4		
Storage Grade: All Storage Nodes		
Node Name	Site Name	Used (%) <span>?</span>
SITE2-S2	SITE2	0.382%
SITE2-S1	SITE2	0.417%
SITE2-S3	SITE2	0.434%
SITE2-S4	SITE2	0.323%

6. [Create an ILM rule](#) specifying object data to be stored in both sides. In this example, a one copy in each site rule was created.
  - a. Set Ingest Behavior to [Balanced](#) (default). This makes sure that when a site is lost, object data can still be written to StorageGRID.

1 copy per site

**Description:** 1 copy per site

**Ingest Behavior:** Balanced

**Reference Time:** Ingest Time

**Filtering Criteria:**

Matches all objects.

**Retention Diagram:**

The diagram shows two horizontal bars representing sites: 'site1' and 'site2'. Above each bar is a 'Trigger' icon and 'Day 0'. Below each bar is a 'Duration' icon and 'Forever'. A vertical line is positioned between the two sites.

7. The above configuration example results in the following:
  - a. A virtual S3 endpoint IP that during primary data center loss reroutes traffic to the secondary data center automatically.
  - b. When backups are archived to StorageGRID, one copy is made per data center. If the primary data center fails, a replicated copy still exists in the secondary data center ready for retrieval.
  - c. Both are transparent to Rubrik. The archival location continues to use the same IP and bucket name.

## Configuration recommendations

Rubrik already encrypts and compresses its data. Therefore, NetApp recommends disabling encryption and compression on StorageGRID.

Rubrik is not versioning aware, so versioning on StorageGRID buckets should be disabled.

Rubrik has not been validated with Cloud Storage Pools. If customers want to use capacity from the public cloud, you can configure Rubrik to support additional cloud storage targets.

StorageGRID performs better with larger objects, and you can increase part size to 128MB or more to increase throughput.

Rubrik, as of the current release tested in this document, performs nonexistent HEADs as part of the connectivity check. NetApp recommends changing the bucket consistency of the StorageGRID buckets to

available to enable greater availability during node or site loss. For more information, read about the [StorageGRID consistency controls](#). For information about how to change the consistency level, see the instructions in the StorageGRID documentation center.

## Where to find additional information

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

- NetApp StorageGRID documentation  
<http://docs.netapp.com/sgws-115/index.jsp>  
<https://www.netapp.com/us/media/ds-3613.pdf>
- TR-6773 StorageGRID Performance 11.5  
<https://fieldportal.netapp.com/content/205463>
- StorageGRID Solution Designer  
<https://fieldportal.netapp.com/content/593648>
- Rubrik resources  
<https://www.rubrik.com/resources/>
- Rubrik support articles (requires login)  
<https://support.rubrik.com/s/article/000003591>  
<https://support.rubrik.com/s/article/000003058>  
<https://support.rubrik.com/s/article/000002692>
- Rubrik SDK  
<https://build.rubrik.com/>  
<https://github.com/rubrikinc>

## Version history

Version	Date	Document version history
Version 1.0	December 2019	Initial release.
Version 1.1	May 2020	Updated sizing and performance section
Version 1.2	September 2021	<ul style="list-style-type: none"><li>• Updated to the latest template.</li><li>• Updated screenshots to latest Rubrik and StorageGRID version</li><li>• Included information on multisite deployments</li></ul> Updated "Configuration recommendations" section

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 © 2019–2021 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.

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-transferrable, 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.

Rubrik, the Rubrik graphic, Rubrik Polaris, Polaris GPS, Polaris Radar, Rubrik Envision, Rubrik Edge, Rubrik Mosaic, and Datos IO are trademarks or registered trademarks of Rubrik, Inc. in the U.S. and/or other countries. All other trademarks are the property of their respective owners.

TR-4812-0921