



Technical Report

# NetApp FabricPool with StorageGRID

## Recommendation guide

Rohan Thomas and Aron Klein | NetApp  
February 2023 | TR-4826

### Abstract

In this document, we discuss best practices and recommendations for deploying and sizing NetApp® StorageGRID® as a capacity tier for the NetApp ONTAP® software component FabricPool. This document also covers capabilities, requirements, implementation, and best practices when using StorageGRID.

## TABLE OF CONTENTS

<b>Overview .....</b>	<b>3</b>
<b>FabricPool policies .....</b>	<b>3</b>
Auto tiering policy .....	3
Snapshot-Only tiering policy .....	3
All tiering policy.....	4
<b>StorageGRID configuration.....</b>	<b>4</b>
<b>StorageGRID recommendations .....</b>	<b>5</b>
<b>StorageGRID performance and capacity sizing .....</b>	<b>5</b>
<b>StorageGRID performance and capacity sizing .....</b>	<b>5</b>
High-availability groups .....	6
Load balancer endpoints .....	6
Load balancer topology .....	6
Load balancer performance.....	6
Load balancer QoS.....	6
<b>Example StorageGRID 11.4 deployments .....</b>	<b>7</b>
Example 1: Auto tiering policy .....	7
Example 2: All tiering policy.....	7
<b>Where to find additional information .....</b>	<b>7</b>
<b>Version history.....</b>	<b>7</b>

## Overview

The NetApp StorageGRID object-based storage solution, which can store cold data from NetApp ONTAP to an object store, has supported NetApp FabricPool since the release of StorageGRID 10.3. This technical report focuses on best practices and recommendations for deploying and sizing StorageGRID for FabricPool, a component of ONTAP. We also cover the capabilities, requirements, implementation, and best practices for using StorageGRID.

Unlike public clouds that might set a maximum number of supported IOPS at the bucket or container level, StorageGRID performance scales with the number of nodes in a system. For acceptable performance targets, NetApp recommends using enough nodes to meet or exceed FabricPool connectivity requirements.

**Note:** This technical report uses the terms performance tier and all-flash storage interchangeably. It also uses capacity tier and object storage interchangeably.

## FabricPool policies

The main objective of using FabricPool with StorageGRID is to reclaim SSD capacity on the primary storage system by moving infrequently accessed data to low-cost object storage. This movement of cold data blocks is performed with FabricPool volume tiering policies.

FabricPool currently supports three tiering policies:

- `auto`
- `snapshot-only`
- `all`

The following subsections briefly cover considerations for using these FabricPool tiering policies with StorageGRID as the object storage.

### Auto tiering policy

The `auto` tiering volume policy has been available since ONTAP 9.4 and moves all cold blocks in the volume to the capacity tier, StorageGRID. The different operations that FabricPool performs on StorageGRID can be a combination of the following:

- Small random-range read GET operations (4KB, 8KB, or 32KB random-range reads)
- Large sequential-range read GET operations (256KB sequential-range reads)
- 4MB PUT operations (1000 x 4KB cold blocks)

Because the `auto` policy can perform these three operations at the same time, it requires high performance from the capacity tier. Therefore, you should emphasize performance sizing.

**Note:** When you use the `auto` tiering policy, NetApp recommends using the StorageGRID SG6060 appliance.

### Snapshot-only tiering policy

The `snapshot-only` volume tiering policy for FabricPool is an effective, intuitive way to reclaim storage space on SSDs. This policy tiers cold NetApp Snapshot™ volume blocks that are not shared with the active file system to StorageGRID. If read, cold data blocks on the capacity tier become hot and are moved back to the performance tier.

The `snapshot-only` tiering policy performs two operations:

- 4MB PUT operations (1000 x 4KB cold blocks)

- Large sequential-range read GET operations (256KB sequential range reads)

Unlike the `auto` tiering policy, the `snapshot-only` tiering policy requests either a GET or PUT operation at a given time (not both at the same time). Therefore, this policy doesn't require as much performance as the `auto` tiering policy does.

**Note:** When you use the `snapshot-only` tiering policy, NetApp recommends using the SG6060 appliance, the SG6060 appliance with one or two expansion shelves, the SG5760 appliance, or the SG5712 appliance.

## All tiering policy

The `all` volume tiering policy covers one of the most common uses of FabricPool. It lets you move entire volumes of data to StorageGRID. The `all` policy is primarily used with secondary data and data protection volumes.

The `all` tiering policy performs two operations:

- 4MB PUT operations (1000 x 4KB cold blocks)
- Large sequential-range read GET operations (256KB sequential range reads)

Unlike the `auto` tiering policy, the `all` tiering policy requests either a GET or PUT operation at a given time (not both at the same time). Therefore, this policy doesn't require as much performance as the `auto` tiering policy does.

**Note:** When you use the `all` tiering policy, NetApp recommends using the SG6060 appliance, the SG6060 appliance with one or two expansion shelves, the SG5760 appliance, or the SG5712 appliance.

## StorageGRID configuration

After the basic requirements for FabricPool have been met, attach an object store to an aggregate in NetApp ONTAP. The detailed steps for this process can be found in [TR-4598](#).

To create a bucket in NetApp StorageGRID, use StorageGRID Tenant Manager to complete the following steps:

1. Log in with your tenant account ID, user name, and password.
2. Select **S3**.
3. Select **Buckets**.
4. Click **Create Bucket**.
5. Provide a DNS-compliant name.
6. Click **Save**.

The screenshot shows a 'Create Bucket' dialog box. At the top, it says 'Create Bucket'. Below that is a section titled 'Bucket Details' with a help icon. There are two input fields: 'Name' with the value 'fabricpool789' and 'Region' with a dropdown menu showing 'us-east-1'. At the bottom right, there are two buttons: 'Cancel' and 'Save'.

**Note:** In versions earlier than StorageGRID 11.1, creating a bucket requires using a third-party Amazon S3 client such as S3 Browser.

After you create a bucket, you must perform additional steps in ONTAP, such as adding a cloud tier, adding a cloud tier as an aggregate, and setting volume tiering policies. For the detailed steps of this process, see [TR-4598](#).

## StorageGRID recommendations

FabricPool supports NetApp StorageGRID information lifecycle management (ILM) policies for data replication and erasure coding to protect FabricPool cloud tier data from failure. However, FabricPool does not support advanced ILM rules such as filtering based on user metadata or tags.

ILM can include various movement and deletion policies based on geography, storage class, retention, and other categories that would be disruptive to FabricPool cloud tier data. FabricPool has no knowledge of ILM policies or configurations set on external object stores, and misconfiguration of ILM policies can result in data loss. For example, FabricPool cloud tier data must not be expired, deleted, or tiered to cloud storage pools (Glacier).

StorageGRID uses two-copy replication as the default ILM rule for data protection. As of StorageGRID 11.2+, intrasite erasure coding using a 2+1 scheme is the recommended best practice for cost efficient data protection. Erasure coding uses more CPU, but significantly less storage capacity, than replication. 4+1 and 6+1 schemes will use even less capacity than 2+1, but at the cost of lower throughput and less flexibility when adding storage nodes during grid expansion.

**Note:** Single-copy replication is not recommended due to lowered system availability and data durability. Geographically dispersed erasure coding such as 4+2 or 6+3 over multiple physical sites is not recommended due to additional latencies. Geographically dispersed erasure coding such as 4+2 or 6+3 over multiple physical sites is not recommended due to additional latencies.

NetApp recommends that you do not enable platform services such as cloud mirroring or metadata search on the object store bucket because FabricPool cloud data is opaque and cannot be used further.

For FabricPool buckets, the recommended bucket consistency level is `read-after-new-write`. The `available` consistency level is not recommended and should not be used.

NetApp ONTAP storage efficiencies such as compression, deduplication, and compaction are preserved when you move data to the cloud tier. NetApp recommends disabling StorageGRID compression and encryption.

## StorageGRID performance and capacity sizing

Different FabricPool policies have different performance and capacity requirements for the capacity tier.

To view the FabricPool PUT and FabricPool range read performance numbers, see [TR-6773: StorageGRID Performance](#). Please note that these performance numbers are only applicable to `snapshot-only` and `all` tiering policies. Contact [ng-storagegrid-tme@netapp.com](mailto:ng-storagegrid-tme@netapp.com) for guidance on sizing StorageGRID for `snapshot-only` and `all` tiering policies.

Performance and capacity results for the `auto` tiering policy are included in the Fusion sizing tool for AFF with StorageGRID using FabricPool. For more information, see [fusion.netapp.com](https://fusion.netapp.com).

## StorageGRID performance and capacity sizing

This section covers configuration recommendations for NetApp StorageGRID load balancers used in FabricPool workloads.

## High-availability groups

In StorageGRID, high-availability (HA) groups use virtual IP addresses (VIPs) to provide active-backup access. HA groups can consist of a combination of admin nodes, gateway nodes, or both. Gateway nodes are dedicated load balancer nodes, whereas admin nodes run both management and Amazon Simple Storage Service (Amazon S3) load-balancing services. You assign one node in the group to be the active primary. When admin nodes are configured in an HA group, ports 443 and 80 cannot be configured for Amazon S3 access because they are reserved for the Grid Manager and Tenant Manager UIs.

## Load balancer endpoints

The load balancer service in StorageGRID provides Layer-7 load balancing. It also performs Transport Layer Security (TLS) termination of client requests, inspects the requests, and establishes new secure connections to the storage nodes. First, create a load balancer endpoint to be used for the FabricPool workload with the display name (for example, S3.netapp.com) and the port to be used (for example, 10443). When selecting HTTPS, you must upload or generate a certificate. Make sure to upload the certificate to the ONTAP cluster, as well as the root and any subordinate certificate authority (CA) certificates.

## Load balancer topology

A general best practice is to use two load-balancing nodes in each site in an HA group with at least one load balancer endpoint configured. Load balancers pass all data from the sources and load-balance the load to the targets. If you are implementing SG1000 appliances, bare metal, or virtual machines (VMs), make sure that there is adequate networking, hardware, or virtualization infrastructure.

For configurations of primary local performance tiers with the `auto` or `all` tiering policies, use a StorageGRID HA group of at least two nodes.

For configurations of primary local performance tiers with the `snapshot-only` tiering policy or non-primary local performance tiers (for example, disaster recovery locations or NetApp SnapMirror® destinations), you can use a single-node load balancer configuration.

## Load balancer performance

A FabricPool workload typically consists of 4MB PUT operations and range reads of 4KB to 256KB. A single SG1000 load balancer has the throughput to support many ONTAP clusters.

## Load balancer QoS

FabricPool might not be the only workload you want to configure for the use of the StorageGRID load balancers. To help a workload get most of the bandwidth, the load balancers can provide Layer-7 quality of service (QoS) based on IP. You can configure this setting through the grid management API.

For configurations of primary local performance tiers with the `auto` or `all` tiering policy, provide a QoS configuration that limits other workloads so that FabricPool workloads have the majority of bandwidth.

For configurations of primary local performance tiers with the `snapshot-only` tiering policy, or non-primary local performance tiers (for example, disaster recovery locations or SnapMirror destinations), prioritizing other primary workloads can be beneficial.

# Example StorageGRID 11.4 deployments

## Example 1: Auto tiering policy

If FabricPool cloud tier data sent to the capacity tier requires specific performance objectives, NetApp recommends a StorageGRID configuration based on the StorageGRID SG6060 appliance for improved read performance and the SG1000 appliance for highly available load balancing.

- Two SG1000 services appliances (a load balancer and an admin node configured with virtual IP)
- Six SG6060 StorageGRID appliances (storage node)
- Single-site, two-copy replication

## Example 2: All tiering policy

In other cases, the FabricPool cloud tier data sent to the capacity tier might have low SLAs. For example, SnapMirror backups are governed by the required recovery time objective (RTO) and (RPO), which can be met by the more cost-efficient SG5700-series StorageGRID appliance:

- Two VM gateway nodes with virtual IP
- One VM admin node
- Six SG5760 StorageGRID appliances (storage node)
- Single-site 4+1 erasure coding

## Where to find additional information

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

- StorageGRID Documentation Resources page  
<https://www.netapp.com/us/documentation/storagegrid.aspx>
- StorageGRID 11.4 Documentation Center  
<https://docs.netapp.com/sqws-114/index.jsp>
- NetApp Product Documentation  
<https://docs.netapp.com>

## Version history

Version	Date	Document version history
Version 1.0	April 2020	Initial version written by Rohan Thomas and Aron Klein.
Version 1.1	September 2020	Minor updates made by Rohan Thomas and Aron Klein.
Version 1.2	February 2023	Update to the StorageGRID Recommendation section with respect to the bucket consistency level.

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 © 2023 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, 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) (FEB 2014).

TR-4826-0223