



Configure StorageGRID for FabricPool

StorageGRID 11.5

NetApp
November 07, 2024

Table of Contents

- Configure StorageGRID for FabricPool 1
 - Configuring StorageGRID for FabricPool 1
 - Information needed to attach StorageGRID as a cloud tier 3
 - Using StorageGRID information lifecycle management with FabricPool data 14
 - Creating a traffic classification policy for FabricPool 17
 - Other best practices for StorageGRID and FabricPool 19

Configure StorageGRID for FabricPool

Learn how to configure StorageGRID as a NetApp FabricPool cloud tier.


- [Configuring StorageGRID for FabricPool](#)
- [Information needed to attach StorageGRID as a cloud tier](#)
- [Using StorageGRID information lifecycle management with FabricPool data](#)
- [Creating a traffic classification policy for FabricPool](#)
- [Other best practices for StorageGRID and FabricPool](#)

Configuring StorageGRID for FabricPool

If you use NetApp ONTAP software, you can use NetApp FabricPool to tier inactive, or cold, data to a NetApp StorageGRID object storage system.

Use these instructions to:

- Get an overview of configuring a StorageGRID object storage system for use with FabricPool.
- Learn how to obtain the information you provide to ONTAP when you attach StorageGRID as a FabricPool cloud tier.
- Learn the best practices for configuring the StorageGRID information lifecycle management (ILM) policy, a StorageGRID traffic classification policy, and other StorageGRID options for a FabricPool workload.



What you'll need

Before using these instructions:

- Decide which FabricPool volume tiering policy you will use to tier inactive ONTAP data to StorageGRID.
- Plan and install a StorageGRID system to meet your storage capacity and performance needs.
- Become familiar with StorageGRID system software, including the Grid Manager and the Tenant Manager.

Related information

- [TR-4598: FabricPool Best Practices for ONTAP 9.8](#)
- [ONTAP 9 Documentation Center](#)

What FabricPool is

FabricPool is an ONTAP hybrid storage solution that uses a high-performance flash aggregate as the performance tier and an object store as the cloud tier. Data in a FabricPool is stored in a tier based on whether it is frequently accessed or not. Using a FabricPool helps you reduce storage cost without compromising performance, efficiency, or protection.

No architectural changes are required, and you can continue managing your database and application environment from the central ONTAP storage system.

What object storage is

Object storage is a storage architecture that manages data as objects, as opposed to other storage architectures such as file or block storage. Objects are kept inside a single container (such as a bucket) and are not nested as files inside a directory inside other directories. Although object storage generally provides lower performance than file or block storage, it is significantly more scalable. StorageGRID buckets can hold petabytes of data.

Using StorageGRID as a FabricPool cloud tier

FabricPool can tier ONTAP data to a number of object store providers, including StorageGRID. Unlike public clouds that might set a maximum number of supported input/output operations per second (IOPS) at the bucket or container level, StorageGRID performance scales with the number of nodes in a system. Using StorageGRID as a FabricPool cloud tier allows you to keep your cold data in your own private cloud for highest performance and complete control over your data.

In addition, a FabricPool license is not required when you use StorageGRID as the cloud tier.

Using multiple ONTAP clusters with StorageGRID

These instructions describe how to connect StorageGRID to a single ONTAP cluster. However, you might want to connect the same StorageGRID system to multiple ONTAP clusters.

The only requirement for tiering data from multiple ONTAP clusters to a single StorageGRID system is that you must use a different S3 bucket for each cluster. Based on your requirements, you can use the same high availability (HA) group, load balancer endpoint, and tenant account for all clusters, or you can configure each of these items for each cluster.

Information needed to attach StorageGRID as a cloud tier

Before you can attach StorageGRID as an cloud tier for FabricPool, you must perform some configuration steps in StorageGRID and obtain certain values.

About this task

The following table lists the information you must provide to ONTAP when you attach StorageGRID as a cloud tier for FabricPool. The topics in this section explain how to use the StorageGRID Grid Manager and Tenant Manager to obtain the information you need.



The exact field names listed and the process you use to enter the required values in ONTAP depend on whether you are using the ONTAP CLI (storage aggregate object-store config create) or ONTAP System Manager (**Storage > Aggregates & Disks > Cloud Tier**).

For more information, refer to the following:

- [TR-4598: FabricPool Best Practices for ONTAP 9.8](#)
- [ONTAP 9 Documentation Center](#)

ONTAP field	Description
Object store name	Any unique and descriptive name. For example, <code>StorageGRID_Cloud_Tier</code> .
Provider type	StorageGRID (System Manager) or SGWS (CLI).
Port	The port that FabricPool will use when it connects to StorageGRID. You determine which port number to use when you define the StorageGRID load balancer endpoint. Creating a load balancer endpoint for FabricPool
Server name	The fully qualified domain name (FQDN) for the StorageGRID load balancer endpoint. For example, <code>s3.storagegrid.company.com</code> . Note the following: <ul style="list-style-type: none">• The domain name that you specify here must match the domain name on the CA certificate you upload for the StorageGRID load balancer endpoint.• The DNS record for this domain name must map to each IP address you will use to connect to StorageGRID. Configuring the DNS server for StorageGRID IP addresses
Container name	The name of the StorageGRID bucket you will use with this ONTAP cluster. For example, <code>fabricpool-bucket</code> . You create this bucket in the Tenant Manager. Note the following: <ul style="list-style-type: none">• The bucket name cannot be changed once the configuration is created.• The bucket cannot have versioning enabled.• You must use a different bucket for each ONTAP cluster that will tier data to StorageGRID. Creating an S3 bucket and obtaining an access key

ONTAP field	Description
Access key and secret password	<p>The access key and secret access key for the StorageGRID tenant account.</p> <p>You generate these values in the Tenant Manager.</p> <p>Creating an S3 bucket and obtaining an access key</p>
SSL	Must be enabled.
Object store certificate	<p>The CA certificate you uploaded when you created the StorageGRID load balancer endpoint.</p> <p>Note: If an intermediate CA issued the StorageGRID certificate, you must provide the intermediate CA certificate. If the StorageGRID certificate was issued directly by the Root CA, you must provide the Root CA certificate.</p> <p>Creating a load balancer endpoint for FabricPool</p>

After you finish

After you have obtained the required StorageGRID information, you can go to ONTAP to add StorageGRID as a cloud tier, add the cloud tier as an aggregate, and set volume tiering policies.

Best practices for load balancing

Before attaching StorageGRID as a FabricPool cloud tier, you use the StorageGRID Grid Manager to configure at least one load balancer endpoint.

What load balancing is

When data is tiered from FabricPool to a StorageGRID system, StorageGRID uses a load balancer to manage the ingest and retrieval workload. Load balancing maximizes speed and connection capacity by distributing the FabricPool workload across multiple Storage Nodes.

The StorageGRID Load Balancer service is installed on all Admin Nodes and all Gateway Nodes and provides Layer 7 load balancing. It performs Transport Layer Security (TLS) termination of client requests, inspects the requests, and establishes new secure connections to the Storage Nodes.

The Load Balancer service on each node operates independently when forwarding client traffic to the Storage Nodes. Through a weighting process, the Load Balancer service routes more requests to Storage Nodes with higher CPU availability.

Although the StorageGRID Load Balancer service is the recommended load balancing mechanism, you might want to integrate a third-party load balancer instead. For information, contact your NetApp account representative or refer to the following technical report:

[StorageGRID Load Balancer Options](#)



The separate Connection Load Balancer (CLB) service on Gateway Nodes is deprecated and no longer recommended for use with FabricPool.

Best practices for StorageGRID load balancing

As a general best practice, each site in your StorageGRID system should include two or more nodes with the Load Balancer service. For example, a site might include both an Admin Node and a Gateway Node or even two Admin Nodes. Make sure that there is adequate networking, hardware, or virtualization infrastructure for each load-balancing node, whether you are using SG100 or SG1000 services appliances, bare metal nodes, or virtual machine (VM) based nodes.

You must configure a StorageGRID load balancer endpoint to define the port that Gateway Nodes and Admin Nodes will use for incoming and outgoing FabricPool requests.

Best practices for the load balancer endpoint certificate

When creating a load balancer endpoint for use with FabricPool, you must use HTTPS as the protocol. You can then either upload a certificate that is signed by either a publicly trusted or a private Certificate Authority (CA), or you can generate a self-signed certificate. The certificate allows ONTAP to authenticate with StorageGRID.

As a best practice, you should use a CA server certificate to secure the connection. Certificates signed by a CA can be rotated nondisruptively.

When requesting a CA certificate for use with the load balancer endpoint, ensure that the domain name on the certificate matches the server name you enter in ONTAP for that load balancer endpoint. If possible, use a wildcard (*) to allow for virtual-host-style URLs. For example:

```
*.s3.storagegrid.company.com
```

When you add StorageGRID as a FabricPool cloud tier, you must install the same certificate to the ONTAP cluster, as well as the root and any subordinate certificate authority (CA) certificates.



StorageGRID uses server certificates for a number of purposes. If you are connecting to the Load Balancer service, you do not need to upload the Object Storage API Service Endpoints Server Certificate.

To learn more about the server certificate for a load balancing endpoint:

- [Managing load balancing](#)
- [Hardening guidelines for server certificates](#)

Best practices for high availability groups

Before attaching StorageGRID as a FabricPool cloud tier, you use the StorageGRID Grid Manager to configure a high availability (HA) group.

What a high availability (HA) group is

To ensure that the Load Balancer service is always available to manage FabricPool data, you can group the network interfaces of multiple Admin and Gateway Nodes into a single entity, known as a high availability (HA) group. If the active node in the HA group fails, another node in the group can continue to manage the workload.

Each HA group provides highly available access to the shared services on the associated nodes. For example,

an HA group consisting of all Admin Nodes provides highly available access to some Admin Node management services and to the Load Balancer service. An HA group that consists of only Gateway Nodes or of both Admin Nodes and Gateway Nodes provides highly available access to the shared Load Balancer service.

When creating an HA group, you select network interfaces belonging to the Grid Network (eth0) or the Client Network (eth2). All interfaces in an HA group must be within the same network subnet.

An HA group maintains one or more virtual IP addresses that are added to the active interface in the group. If the active interface becomes unavailable, the virtual IP addresses are moved to another interface. This failover process generally takes only a few seconds and is fast enough that client applications should experience little impact and can rely on normal retry behaviors to continue operation.

If you configure an HA group of load-balancing nodes, FabricPool connects to the virtual IP addresses of that HA group.

Best practices for high availability (HA) groups

The best practices for creating a StorageGRID HA group for FabricPool depend on the workload, as follows:

- If you plan to use FabricPool with primary workload data, you must create a HA group that includes at least two load-balancing nodes to prevent data retrieval interruption.
- If you plan to use the FabricPool snapshot-only volume tiering policy or non-primary local performance tiers (for example, disaster recovery locations or NetApp SnapMirror® destinations), you can configure an HA group with only one node.

These instructions describe setting up an HA group for Active-Backup HA (one node is active and one node is backup). However, you might prefer to use DNS Round Robin or Active-Active HA. To learn the benefits of these other HA configurations, see [Configuration options for HA groups](#).

Configuring the DNS server for StorageGRID IP addresses

After configuring high availability groups and load balancer endpoints, you must ensure that the domain name system (DNS) for the ONTAP system includes a record to associate the StorageGRID server name (fully qualified domain name) to the IP address that FabricPool will use to make connections.

The IP address you enter in the DNS record depends on whether you are using an HA group of load-balancing nodes:

- If you have configured a HA group, FabricPool will connect to the virtual IP addresses of that HA group.
- If you are not using a HA group, FabricPool can connect to the StorageGRID Load Balancer service using the IP address of any Gateway Node or Admin Node.

You must also ensure that the DNS record references all required endpoint domain names, including any wildcard names.

Creating a high availability (HA) group for FabricPool

When configuring StorageGRID for use with FabricPool, you can optionally create one or more high availability (HA) groups. An HA group consists of one or more network interfaces on Admin Nodes, Gateway Nodes, or both.

What you'll need

- You must be signed in to the Grid Manager using a supported browser.
- You must have the Root Access permission.

About this task

Each HA group uses virtual IP addresses (VIPs) to provide highly available access to the shared services on the associated nodes.

For details about this task, see [Managing high availability groups](#).

Steps

1. Select **Configuration > Network Settings > High Availability Groups**.
2. Select one or more of the network interfaces. The network interfaces must belong to the same subnet on either the Grid Network (eth0) or the Client Network (eth2).
3. Assign one node to be the Preferred Master.

The preferred Master is the active interface unless a failure occurs that causes the VIP addresses to be reassigned to a Backup interface.

4. Enter up to ten IPv4 addresses for the HA group.

The addresses must be within the IPv4 subnet shared by all of the member interfaces.

Create High Availability Group

High Availability Group

Name

Description

Interfaces

Select interfaces to include in the HA group. All interfaces must be in the same network subnet.

Select Interfaces

Node Name	Interface	IPv4 Subnet	Preferred Master
DC1-ADM1	eth0	10.96.98.0/23	<input checked="" type="radio"/>
DC1-G1	eth0	10.96.98.0/23	<input type="radio"/>

Displaying 2 interfaces.

Virtual IP Addresses

Virtual IP Subnet: 10.96.98.0/23. All virtual IP addresses must be within this subnet. There must be at least 1 and no more than 10 virtual IP addresses.

Virtual IP Address 1



Cancel

Save

Creating a load balancer endpoint for FabricPool

When configuring StorageGRID for use with FabricPool, you configure a load balancer endpoint and upload the load balancer endpoint certificate, which is used to secure the connection between ONTAP and StorageGRID.

What you'll need

- You must be signed in to the Grid Manager using a supported browser.
- You must have the Root Access permission.
- You have the following files:
 - Server Certificate: The custom server certificate file.
 - Server Certificate Private Key: The custom server certificate private key file.
 - CA Bundle: A single file containing the certificates from each intermediate issuing Certificate Authority

(CA). The file should contain each of the PEM-encoded CA certificate files, concatenated in certificate chain order.

About this task

For details about this task, see [Configuring load balancer endpoints](#).

Steps

1. Select **Configuration > Network Settings > Load Balancer Endpoints**.

Create Endpoint

Display Name

Port

Protocol HTTP HTTPS

Endpoint Binding Mode Global HA Group VIPs Node Interfaces

2. Select **Add endpoint**.
3. Enter the following information.

Field	Description
Display name	A descriptive name for the endpoint
Port	<p>The StorageGRID port you want to use for load balancing. This field defaults to 10433, but you can enter any unused external port. If you enter 80 or 443, the endpoint is configured only on Gateway Nodes, since these ports are reserved on Admin Nodes.</p> <p>Note: Ports used by other grid services are not permitted. See the list of ports used for internal and external communications:</p> <p>Network port reference</p> <p>You must provide this same port number to ONTAP when you attach StorageGRID as a FabricPool cloud tier.</p>
Protocol	Must be HTTPS .

Field	Description
Endpoint Binding Mode	<p>Use the Global setting (recommended) or restrict the accessibility of this endpoint to one of the following:</p> <ul style="list-style-type: none"> • Specific high availability (HA) virtual IP addresses (VIPs). Use this selection only if you require much higher levels of isolation of workloads. • Specific network interfaces of specific nodes.

4. Select **Save**.

The Edit Endpoint dialog box appears.

5. For **Endpoint Service Type**, select **S3**.

6. Select **Upload Certificate** (recommended) and then browse to your server certificate, certificate private key, and CA bundle.

Load Certificate

Upload the PEM-encoded custom certificate, private key, and CA bundle files.

Server Certificate

Certificate Private Key

CA Bundle

7. Select **Save**.

Creating a tenant account for FabricPool

You must create a tenant account in the Grid Manager for FabricPool use.

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

Tenant accounts allow client applications to store and retrieve objects on StorageGRID. Each tenant account has its own account ID, authorized groups and users, buckets, and objects.

You can use the same tenant account for multiple ONTAP clusters. Or, you can create a dedicated tenant account for each ONTAP cluster as required.



These instructions assume that you have configured single sign-on (SSO) for the Grid Manager. If you are not using SSO, use the instructions for [creating a tenant account if StorageGRID is not using SSO](#).

Steps

1. Select **Tenants**.
2. Select **Create**.
3. Enter a display name for the FabricPool tenant account.
4. Select **S3**.
5. Leave the **Allow Platform Services** check box selected to enable the use of platform services.

If platform services are enabled, a tenant can use features, such as CloudMirror replication, that access external services.

6. Leave the **Storage Quota** field blank.
7. In the **Root Access Group** field, select an existing federated group from the Grid Manager to have the initial Root Access permission for the tenant.
8. Select **Save**.

Creating an S3 bucket and obtaining an access key

Before using StorageGRID with a FabricPool workload, you must create an S3 bucket for your FabricPool data. You also need to obtain an access key and secret access key for the tenant account you will use for FabricPool.

What you'll need

- You must have created a tenant account for FabricPool use.

About this task

These instructions describe how to use the StorageGRID Tenant Manager to create a bucket and obtain access keys. You can also perform these tasks using the Tenant Management API or the StorageGRID S3 REST API.

To learn more:

- [Use a tenant account](#)
- [Use S3](#)

Steps

1. Sign in to the Tenant Manager.

You can do either of the following:

- From the Tenant Accounts page in the Grid Manager, select the **Sign in** link for the tenant, and enter your credentials.
- Enter the URL for the tenant account in a web browser, and enter your credentials.

2. Create an S3 bucket for FabricPool data.

You must create a unique bucket for each ONTAP cluster you plan to use.

- a. Select **STORAGE (S3) > Buckets**.
- b. Select **Create bucket**.
- c. Enter the name of the StorageGRID bucket you will use with FabricPool. For example, `fabricpool-bucket`.



You cannot change the bucket name after creating the bucket.

Bucket names must comply with these rules:

- Must be unique across each StorageGRID system (not just unique within the tenant account).
 - Must be DNS compliant.
 - Must contain at least 3 and no more than 63 characters.
 - Can be a series of one or more labels, with adjacent labels separated by a period. Each label must start and end with a lowercase letter or a number and can only use lowercase letters, numbers, and hyphens.
 - Must not look like a text-formatted IP address.
 - Should not use periods in virtual hosted style requests. Periods will cause problems with server wildcard certificate verification.
- d. Select the region for this bucket.

By default, all buckets are created in the `us-east-1` region.

A screenshot of a 'Create bucket' dialog box. The dialog has a dark blue header with the title 'Create bucket' and a close button (X) in the top right corner. Below the header, the section is titled 'Enter bucket details' with a subtitle 'Enter the bucket's name and select the bucket's region.' There are two input fields: 'Bucket name' with a help icon and a text input containing 'fabricpool-bucket'; and 'Region' with a help icon and a dropdown menu showing 'us-east-1'. At the bottom right, there are two buttons: 'Cancel' and 'Create bucket'.

- e. Select **Create bucket**.
3. Create an access key and a secret access key.
 - a. Select **STORAGE (S3) > My access keys**.

- b. Select **Create key**.
- c. Select **Create access key**.
- d. Copy the access key ID and the secret access key to a safe location, or select **Download .csv** to save a spreadsheet file containing the access key ID and secret access key.

You will enter these values in ONTAP when you configure StorageGRID as a FabricPool cloud tier.



If you create a new access key and secret access key in the future, remember to update the corresponding values in ONTAP immediately to ensure that ONTAP can store and retrieve data in StorageGRID without interruption.

Using StorageGRID information lifecycle management with FabricPool data

If you are using FabricPool to tier data to StorageGRID, you must understand the requirements for creating StorageGRID information lifecycle management (ILM) rules and an ILM policy to manage FabricPool data. You must ensure the ILM rules that apply to FabricPool data are not disruptive.



FabricPool has no knowledge of StorageGRID ILM rules or policies. Data loss can occur if the StorageGRID ILM policy is misconfigured.

To learn more: [Manage objects with ILM](#)

ILM guidelines for FabricPool data

Review these guidelines to ensure that your ILM rules and ILM policy are suitable for FabricPool data and your business requirements. If you are already using StorageGRID ILM, you might need to update your active ILM policy to meet these guidelines.

- You can use any combination of replication and erasure-coding rules to protect cloud tier data.

The recommended best practice is to use 2+1 erasure coding within a site for cost-efficient data protection. Erasure coding uses more CPU, but significantly less storage capacity, than replication. The 4+1 and 6+1 schemes use less capacity than 2+1, but at the cost of lower throughput and less flexibility when you add Storage Nodes during grid expansion.

- Each rule applied to FabricPool data must either use erasure coding or it must create at least two replicated copies.



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, 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.

- Do not use an ILM rule that will expire or delete FabricPool cloud tier data. Set the retention period in each ILM rule to "forever" to ensure that FabricPool objects are not deleted by StorageGRID ILM.
- Do not create rules that will move FabricPool cloud tier data out of the bucket to another location. You cannot use ILM rules to archive FabricPool data to tape using an Archive Node or use a Cloud Storage

Pool to move FabricPool data to Glacier.



Using Cloud Storage Pools with FabricPool is not supported because of the added latency to retrieve an object from the Cloud Storage Pool target.

- Starting with ONTAP 9.8, you can optionally create object tags to help classify and sort tiered data for easier management. For example, you can set tags only on FabricPool volumes attached to StorageGRID. Then, when you create ILM rules in StorageGRID, you can use the Object Tag advanced filter to select and place this data.

Example ILM policy for FabricPool data

Use this simple example policy as a starting point for your own ILM rules and policy.

This example assumes you are designing the ILM rules and an ILM policy for a StorageGRID system that has four Storage Nodes at a single data center in Denver, Colorado. The FabricPool data in this example uses a bucket named `fabricpool-bucket`.



The following ILM rules and policy are only examples. There are many ways to configure ILM rules. Before activating a new policy, simulate the proposed policy to confirm it will work as intended to protect content from loss.

To learn more: [Manage objects with ILM](#)

Steps

1. Create a storage pool named **DEN**. Select the Denver site.
2. Create an Erasure Coding profile named **2 plus 1**. Select the 2+1 erasure-coding scheme and the **DEN** storage pool.
3. Create an ILM rule that applies only to the data in `fabricpool-bucket`. This example rule creates erasure-coded copies.

Rule definition	Example value
Rule Name	2 plus 1 erasure coding for FabricPool data
Bucket Name	<code>fabricpool-bucket</code> You could also filter on the FabricPool tenant account.
Advanced Filtering	Object Size (MB) greater than 0.2 MB. Note: FabricPool only writes 4 MB objects, but you must add an Object Size filter because this rule uses erasure coding.
Reference Time	Ingest Time
Placement	From day 0 store forever

Rule definition	Example value
Type	Erasure coded
Location	DEN (2 plus 1)
Ingest Behavior	Balanced

4. Create an ILM rule that will create two replicated copies of any objects not matched by the first rule. Do not select a basic filter (tenant account or bucket name) or any advanced filters.

Rule definition	Example value
Rule Name	Two replicated copies
Bucket Name	<i>none</i>
Advanced Filtering	<i>none</i>
Reference Time	Ingest Time
Placement	From day 0 store forever
Type	Replicated
Location	DEN
Copies	2
Ingest Behavior	Balanced

5. Create a proposed ILM policy and select the two rules. Because the replication rule does not use any filters, it can be the default (last) rule for the policy.
6. Ingest test objects into the grid.
7. Simulate the policy with the test objects to verify the behavior.
8. Activate the policy.

When this policy is activated, StorageGRID places object data as follows:

- The data tiered from FabricPool in `fabricpool-bucket` will be erasure coded using the 2+1 erasure-coding scheme. Two data fragments and one parity fragment will be placed on three different Storage Nodes.
- All objects in all other buckets will be replicated. Two copies will be created and placed on two different Storage Nodes.
- The erasure-coded and replicated copies will be maintained in StorageGRID until they are deleted by the S3 client. StorageGRID ILM will never delete these items.

Creating a traffic classification policy for FabricPool

You can optionally design a StorageGRID traffic classification policy to optimize quality of service for the FabricPool workload.

What you'll need

- You must be signed in to the Grid Manager using a supported browser.
- You must have the Root Access permission.

About this task

The best practices for creating a traffic classification policy for FabricPool depend on the workload, as follows:

- If you plan to tier FabricPool primary workload data to StorageGRID, you should ensure that the FabricPool workload has the majority of bandwidth. You can create a traffic classification policy to limit all other workloads.



In general, FabricPool read operations are more important to prioritize than write operations.

For example, if other S3 clients use this StorageGRID system, you should create a traffic classification policy. You can limit network traffic for the other buckets, tenants, IP subnets, or load balancer endpoints.

- As a general rule, you should not impose quality of service limits on any FabricPool workload; you should only limit the other workloads.
- The limits placed on other workloads might need to be broad to account for the unknown behavior of those workloads. The limits imposed will also vary based on the sizing and capabilities of your grid and what the expected amount of utilization is.

To learn more: [Managing traffic classification policies](#)

Steps

1. Select **Configuration > Network Settings > Traffic Classification**.
2. Enter a name and a description.
3. In the Matching Rules section, create at least one rule.
 - a. Select **Create**.
 - b. Select **Endpoint**, and select the load balancer endpoint you created for FabricPool.

You can also select the FabricPool tenant account or bucket.

- c. If you want this traffic policy to limit traffic for the other endpoints, select **Inverse Match**.
4. Optionally, create one or more limits.



Even if no limits are set for a traffic classification policy, metrics are collected so you can understand traffic trends.


- a. Select **Create**.
- b. Select the type of traffic you want to limit and the limit to apply.

This example FabricPool traffic classification lists the types of network traffic you can limit and the types of values you can select. The traffic types and values for an actual policy would be based on your

specific requirements.

Edit Traffic Classification Policy "FabricPool"




Policy

Name  FabricPool

Description (optional) Limit traffic other than FabricPool

Matching Rules




Traffic that matches any rule is included in the policy.

Type	Inverse Match	Match Value
<input checked="" type="radio"/> Endpoint	<input checked="" type="checkbox"/>	FabricPool (https 10443)



Displaying 1 matching rule.

Limits (Optional)

Type	Value	Units
<input checked="" type="radio"/> Concurrent Read Requests	50	Concurrent Requests
<input checked="" type="radio"/> Concurrent Write Requests	15	Concurrent Requests
<input checked="" type="radio"/> Read Request Rate	100	Requests/Second
<input checked="" type="radio"/> Write Request Rate	25	Requests/Second
<input checked="" type="radio"/> Per-Request Bandwidth In	2000000	Bytes/Second
<input checked="" type="radio"/> Per-Request Bandwidth Out	10000000	Bytes/Second


Displaying 6 limits.

5. After creating the traffic classification policy, select the policy and then select **Metrics** to determine if the policy is limiting traffic as expected.

Traffic Classification Policies

Traffic classification policies can be used to identify network traffic for metrics reporting and optional traffic limiting.



Name	Description	ID
<input checked="" type="radio"/> FabricPool	Limit traffic other than FabricPool	587f53b2-7cf2-44b9-af5c-694ebbd4a2c5

Displaying 1 traffic classification policy.

Other best practices for StorageGRID and FabricPool

When configuring a StorageGRID system for use with FabricPool, you should avoid setting global options that might affect how your data is saved.

Object encryption

When configuring StorageGRID, you can optionally enable the global **Stored Object Encryption** setting if data encryption is required for other StorageGRID clients (**Configuration > System Settings > Grid Options**). The data that is tiered from FabricPool to StorageGRID is already encrypted, so enabling the StorageGRID setting is not required. Client-side encryption keys are owned by ONTAP.

Object compression

When configuring StorageGRID, do not enable the global **Compress Stored Objects** setting (**Configuration > System Settings > Grid Options**). The data that is tiered from FabricPool to StorageGRID is already compressed. Enabling **Compress Stored Objects** will not further reduce an object's size.

Consistency level

For FabricPool buckets, the recommended bucket consistency level is **Read-after-new-write**, which is the default setting for a new bucket. Do not edit FabricPool buckets to use **Available** or any other consistency level.

FabricPool tiering

If the StorageGRID node uses storage assigned from a NetApp AFF system, confirm that the volume does not have a FabricPool tiering policy enabled. For example, if a StorageGRID node is running on a VMware host, ensure the volume backing the datastore for the StorageGRID node does not have a FabricPool tiering policy enabled. Disabling FabricPool tiering for volumes used with StorageGRID nodes simplifies troubleshooting and storage operations.



Never use FabricPool to tier any data related to StorageGRID back to StorageGRID itself. Tiering StorageGRID data back to StorageGRID increases troubleshooting and operational complexity.

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.