



Technical Report

S3 in ONTAP best practices

ONTAP 9.9.1

John Lantz, NetApp
August 2021 | TR-4814

Abstract

This technical report describes best practices for using the Amazon Simple Storage Service (S3) with NetApp® ONTAP® software. We also cover capabilities and configurations for using ONTAP as an object store with native S3 applications or as a tiering destination for NetApp FabricPool.

TABLE OF CONTENTS

| | |
|----------------------------------------------------------------------------------|-----------|
| Overview | 4 |
| Primary use cases | 4 |
| Native S3 applications | 4 |
| FabricPool endpoints | 4 |
| Requirements | 5 |
| Platforms | 5 |
| Data LIFs | 5 |
| Cluster LIFs | 5 |
| S3 license | 5 |
| Architecture | 6 |
| Service policy | 6 |
| Object store server | 7 |
| Bucket | 7 |
| Users | 8 |
| Configuration for native S3 applications and remote cluster tiering | 8 |
| ONTAP System Manager | 8 |
| ONTAP CLI | 12 |
| Configuration for local cluster tiering | 15 |
| ONTAP System Manager | 16 |
| ONTAP CLI | 17 |
| Security | 19 |
| Local tier | 19 |
| Over the wire | 19 |
| Supported S3 actions | 20 |
| Buckets | 20 |
| Objects | 20 |
| Group policies | 20 |
| User management | 20 |
| ONTAP 9.9.1 | 21 |
| Interoperability | 21 |
| Where to find additional information | 22 |
| Version history | 22 |

| | |
|-------------------------|-----------|
| Contact us | 22 |
|-------------------------|-----------|

LIST OF TABLES

| | |
|---------------------------------------|----|
| Table 1) NetApp interoperability..... | 21 |
|---------------------------------------|----|

LIST OF FIGURES

| | |
|-------------------------------------------------------------------|----|
| Figure 1) The core elements of an S3 object storage in ONTAP..... | 6 |
| Figure 2) FlexGroup volume..... | 7 |
| Figure 3) Local cluster tiering..... | 15 |

Overview

NetApp ONTAP 9.8 software supports the Amazon Simple Storage Service (S3). ONTAP supports a subset of AWS S3 API actions and allows data to be represented as objects in ONTAP-based systems, including AFF, FAS, and ONTAP Select.

NetApp StorageGRID® software is, and will remain, the NetApp flagship solution for object storage. ONTAP complements StorageGRID by providing an ingest and preprocessing point on the edge, expanding the data fabric powered by NetApp for object data, and increasing the value of the NetApp product portfolio.

Primary use cases

The primary purpose of S3 in ONTAP is to provide support for objects on ONTAP-based systems. The ONTAP unified storage architecture now supports files (NFS and SMB), blocks (FC and iSCSI), and objects (S3).

Native S3 applications

An increasing number of customers need ONTAP to support objects using S3. Although well suited for high-capacity archival workloads, demand for native S3 applications is growing rapidly and includes:

- Analytics
- Artificial intelligence
- Edge-to-core ingest
- Machine learning

Customers can now use familiar manageability tools such as ONTAP System Manager to rapidly provision high-performance object storage for development and operations in ONTAP, taking advantage of ONTAP's storage efficiencies and security as they do so.

FabricPool endpoints

Starting in ONTAP 9.8, FabricPool supports tiering to buckets in ONTAP, allowing for ONTAP to ONTAP tiering. This is an excellent option for customers who wish to repurpose existing FAS infrastructure as an object store endpoint.

FabricPool supports tiering to ONTAP in two ways:

- **Local cluster tiering.** Inactive data is tiered to a bucket located on the local cluster using cluster LIFs.
- **Remote cluster tiering.** Inactive data is tiered to a bucket located on a remote cluster similarly to a traditional FabricPool cloud tier using IC LIFs on the FabricPool client and data LIFs on the ONTAP object store.

NetApp recommends using StorageGRID, the premier NetApp object store solution, when tiering more than 300TB of inactive data. A FabricPool license is not required when using ONTAP or StorageGRID as the cloud tier.

Requirements

Platforms

- **NetApp AFF storage system.** S3 is supported on all AFF platforms using ONTAP 9.8+.
- **FAS storage system.** S3 is supported on all FAS platforms using ONTAP 9.8+.
- **NetApp ONTAP Select.** S3 is supported on all platforms using ONTAP Select 9.8+.
- **Cloud Volumes ONTAP.** Starting in ONTAP 9.9.1, S3 is supported on Cloud Volumes ONTAP for Azure. S3 is not supported on other Cloud Volumes ONTAP providers.

Data LIFs

Storage virtual machines (SVMs) hosting object store servers require data LIFs to communicate with client applications using S3. When configured for remote cluster tiering, FabricPool is the client and the object store is the server.

Cluster LIFs

When configured for local cluster tiering, a local tier (also known as a storage aggregate in the ONTAP CLI) is attached to a local bucket. FabricPool uses cluster LIFs for intracluster traffic.

Note: Performance degradation might occur if cluster LIFs resources become saturated. To avoid this, NetApp recommends using two-node, or greater, clusters when tiering to a local bucket—the recommended best practice being an HA pair for the local tier and an HA pair for the local bucket. Tiering to local buckets on single node clusters is not recommended.

S3 license

As with other protocols such as FC, iSCSI, NFS, NVMe_oF, and SMB, S3 requires the installation of a license before it can be used in ONTAP. The S3 license is a zero-cost license, but it must be installed on systems upgrading to ONTAP 9.8.

New ONTAP 9.8 systems have the S3 license pre-installed.

The S3 license can be downloaded from the [Master License Keys page](#) on the NetApp support site.

Installation

To install the S3 license, run the following command in the ONTAP CLI:

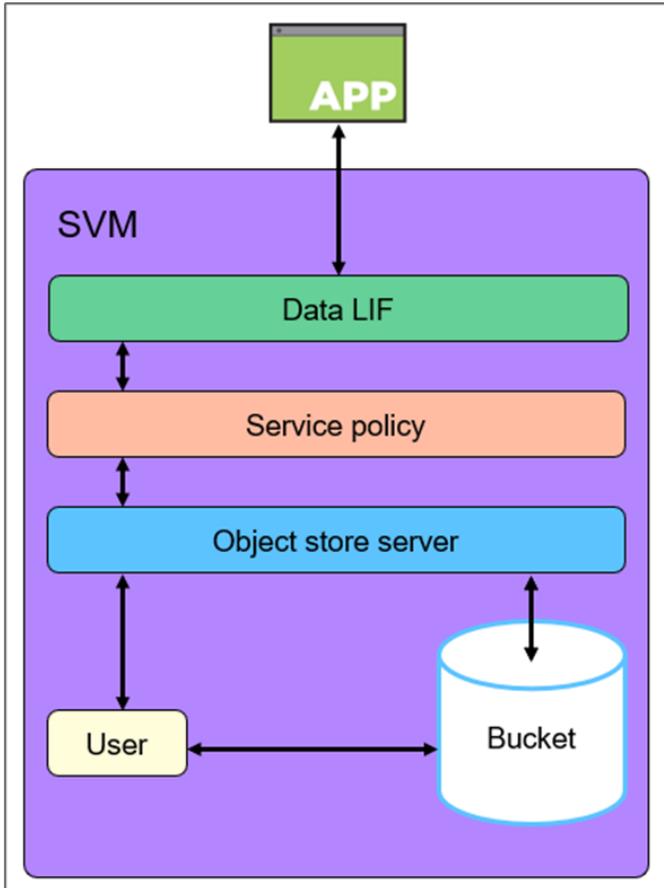
```
system license add <license_key>
```

Architecture

Object storage is an 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 might be less performative than file or block storage, it is significantly more scalable, and buckets containing petabytes of data are not uncommon.

Figure 1) The core elements of an S3 object storage in ONTAP.



Service policy

Data service policies are assigned to SVMs and provide a collection of network services required by data LIFs to support client application protocols. For example, data-nfs is used to support NFS traffic, data-iscsi is used to support iSCSI traffic, and so on.

New in ONTAP 9.8, the data-s3-server service, allows data LIFS to support client application traffic using S3.

Note: In addition to the data-s3-server service, the data-core service should be included in any service policy to ensure applications using the LIF work as expected.

Object store server

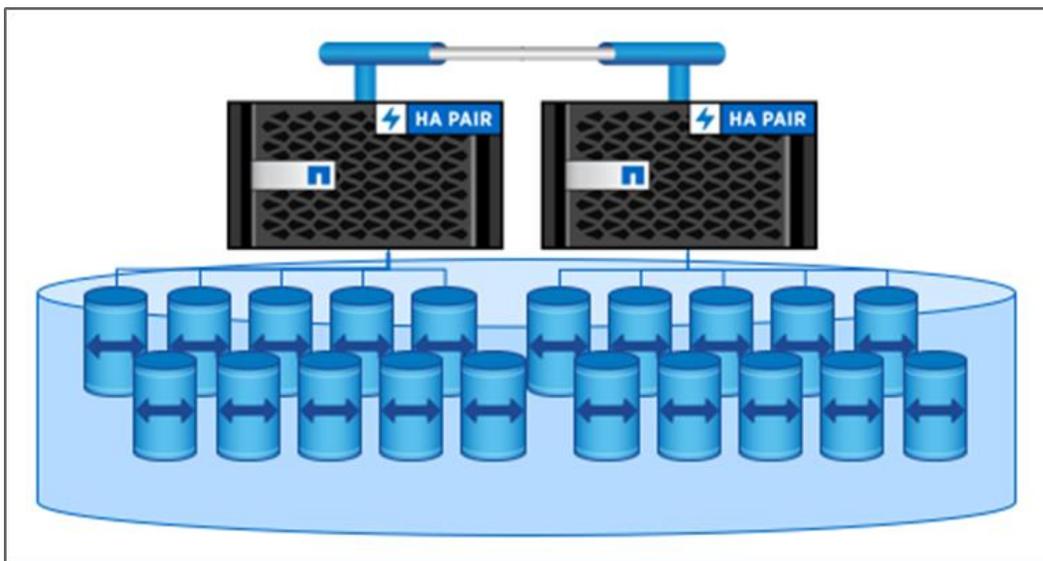
The SVM's object store server manages data as objects, as opposed to other storage architectures such as file or block storage. Management of bucket and user permission levels also takes place at the object store server level.

ONTAP S3 supports one object store server per SVM.

Bucket

In ONTAP, the underlying architecture for a bucket is a [FlexGroup volume](#)—a single namespace that is made up of multiple constituent member volumes but is managed as a single volume, as shown in Figure 2. Individual objects in a bucket are allocated to individual member volumes and are not striped across volumes or nodes. Individual buckets cannot be provisioned smaller than 96GB.

Figure 2) FlexGroup volume.



When used by buckets, FlexGroup volumes use elastic sizing, not volume autogrow. FlexGroup volume maximums are only limited by the physical maximums of the underlying hardware and have been tested to 20PB and 400 billion files in a 10-node cluster.

ONTAP S3 supports up to 12,000 buckets, although no more than 1,000 buckets should be created on a single FlexGroup volume.

The Amazon S3 maximum object size is 5TB. ONTAP S3 supports objects up to 16TB. Objects greater than 5TB might result in interoperability issues for clients that cannot exceed Amazon-defined maximum object sizes.

Note: Underlying architectural changes between ONTAP 9.7 buckets (one bucket per FlexGroup volume) and ONTAP 9.8 (multiple buckets per FlexGroup volume) cannot be made in place. Data must be migrated from preexisting buckets to ONTAP 9.8 buckets to take advantage of the new architecture.

Default bucket settings

Buckets that are not [manually configured](#) will use default settings for aggregate, FlexGroup, and bucket provisioning.

Aggregates

FlexGroup volumes supporting buckets are provisioned on aggregates by using the following priorities:

- Flash Pool aggregate
- HDD aggregate
- QLC SSD aggregate
- TLC SSD aggregate

FlexGroup volumes

The default FlexGroup size is large and provides significant room for expansion in most environments:

- 1.6PB in ONTAP
- 100TB in ONTAP Select

If a cluster does not have enough capacity to provision the default size, the size will be reduced by half until it can be provisioned in the existing environment. For example, in a 300TB environment, a FlexGroup volume would be automatically provisioned at 200TB. (1.6PB, 800TB, and 400TB FlexGroup volumes being too large for the environment.)

Buckets

The default bucket size is:

- 800GB in ONTAP
- 200MB in ONTAP Select

In order to provide capacity for bucket expansion, the total capacity of all buckets on the FlexGroup volume should be less than 33% of the FlexGroup volume capacity. If this cannot be met, the bucket being created will automatically be provisioned on a newly created FlexGroup volume.

Users

User authorization is required on all ONTAP object stores in order to restrict connectivity to authorized clients. Access to specific buckets or S3 actions can be allowed, denied, or made conditional at the user level.

ONTAP S3 supports 4,000 users per object store.

Configuration for native S3 applications and remote cluster tiering

External clients such as native S3 applications and FabricPool clients connect to the ONTAP object store using data LIFs. The easiest way to create an object store in ONTAP is by using ONTAP System Manager. Processes that require multiple steps when using the CLI are reduced to a few clicks using NetApp recommended best practices. Configuration with CLI is required for more custom configurations.

ONTAP System Manager

The easiest way to create an object store in ONTAP is by using the ONTAP System Manager, reducing multiple steps needed with the CLI to a few clicks. Object stores created using ONTAP System Manager allow for less customization, but they are created with NetApp recommended best practices by default. Configuration with the CLI is required for custom configurations.

To create an object store, bucket, and permission users using ONTAP System Manager, complete the following steps:

Configure the object store

To configure the object store, complete the following steps:

1. Launch ONTAP System Manager.
2. Click STORAGE.
3. Click Storage VMs.
4. Click Add. A new SVM is not necessary. S3 functionality can be added to existing SVMs using the SVM's Settings menu.
5. Name the SVM.
6. Select Enable S3 as an access protocol. The options Enable TLS (port 443) and Use System-Generated Certificate are selected by default. Using signed certificates from a third-party certificate authority is a recommended best practice.
7. Name the S3 server.
Note: The server name is used as the fully qualified domain name (FQDN) by client applications.
8. Enter network interfaces for the nodes.

Configure a bucket

To configure a bucket, complete the following steps:

1. Launch ONTAP System Manager.
2. Click STORAGE.
3. Click Buckets.
4. Click Add.
5. Name the bucket.
6. Select the SVM/object store that the bucket will be assigned to. This should be the same SVM/object store created earlier.
7. Click Save.

Add Bucket [X]

NAME
bucket-name

STORAGE VM
S3_object_store [v]

CAPACITY
2 PB [v]

[More Options] [Cancel] [Save]

More options

Use for tiering

If you select this option, ONTAP System Manager creates the bucket on the least expensive media, prioritizing HDD > QLC > TLC > NVMe.

Performance service level

Select the appropriate quality of service (QoS) for the bucket. Options include:

- **Extreme.** 50,000 IOPS; 1562MBps
- **Performance.** 30,000 IOPS; 937MBps
- **Value.** 15,000 IOPS; 468MBps
- **Custom.** Use an existing QoS policy or create a new one.

Note: Performance service levels are not selectable if the bucket is used for tiering. FabricPool does not support QoS minimums.

Permissions

Copy access permissions from an existing bucket or create new ones.

Note: Users and groups must be configured before they can be permissioned. See [Add Users and Groups](#).

To create new permissions, complete the following steps:

1. From the Add Bucket page, scroll down to Permissions and click Add.
2. Set principal users. Options include All users of the SVM (default), All public and anonymous users, and individual users associated with the SVM.
3. Set effect. Options include Allow (default) and Deny.
4. Set actions. Options include GetObject, PutObject, DeleteObject, ListBucket (default), GetBucketAcl, GetObjectAcl, ListBucketMultipartUploads, and ListMultipartUploadParts.
5. Set resources. bucket-name and bucket-name/* are used by default.
6. Set conditions.
7. Add conditions. Up to 10 conditional statements can be added. Each conditional statement is composed of a key, an operator, and one or more values.

New Permission

PRINCIPAL

All users of this stor... X

EFFECT

Allow

ACTIONS

ListBucket X

RESOURCES ?

bucket-name,bucket-name/*

Conditions ?

| KEY | OPERATOR | VALUE ? |
|------------|---------------|---------|
| delimiters | string_equals | |

+ Add

Add users and groups

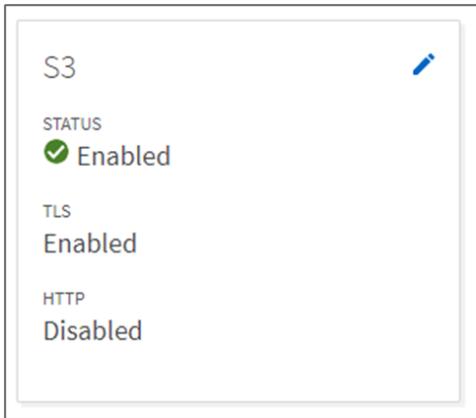
User authorization is required on all ONTAP object stores in order to restrict connectivity to authorized clients. Access to specific buckets or S3 actions can be allowed, denied, or made conditional at the user and group level using [permissions](#).

ONTAP S3 supports 4000 users per object store or SVM.

Note: A root user (UID 0) is created by default when the bucket is created. The root user has full access to all buckets and objects. Do not use the root user for client application access. Additional users must be created for client access.

To manage users and groups, complete the following steps:

1. Launch ONTAP System Manager.
2. Click STORAGE.
3. Click Storage VMs.
4. Select the SVM to add users and groups to.
5. Click the Edit icon on the S3 protocol box.



6. Select the Users or Groups tab.
7. Click Add.
8. Name the user or group.
9. Copy and/or download the access and secret key for future use.
Note: The secret key is not displayed again.
10. If you are configuring a group, assign users and policies.
11. If you are configuring a user, use the [permissions menu](#).

ONTAP CLI

Although the easiest way to create an object store in ONTAP is by using the ONTAP System Manager, object stores created using ONTAP system manager allow for less customization.

For example, ONTAP System Manager automatically selects the local tiers (aggregates) use by a bucket for storage. Although it uses recommended best practices to do so, for complex environments, the selected local tiers might not be the same ones an experienced storage administrator would use.

Configuration using the ONTAP CLI is required for custom configurations.

To create an object store, bucket, and permission users using ONTAP CLI, complete the following steps:

1. Create the service policy.
2. Create a data LIF to use S3.
3. Install a CA certificate.
4. Create the object store server.
5. Create the bucket.
6. Create a user.

Create the service policy

A service policy is required to enable S3 data traffic on the SVM LIFs.

To create the service policy by using the ONTAP CLI, run the following command:

```
network interface service-policy create
-vserver <name>
-policy <name>
-services data-s3-server, data-core
```

Note: In addition to the data-s3-server service, the data-core service should be included in any service policy to ensure applications using the LIF work as expected.

Create a data LIF to use S3

SVMs hosting object store servers require data LIFs to communicate with client applications using S3. NetApp recommends creating an S3 data LIF on all nodes as a best practice.

When configured for remote cluster tiering, FabricPool is the client and the object store is the server. Because FabricPool requires the object store to use a FQDN, all S3 DATA LIFs must be associated with the FQDN used by the Object Store Server.

Note: Creation of the DNS entry is external to ONTAP. NetApp recommends creating a single host entry that uses all S3 data LIF IP addresses.

The `dns-zone` setting is for ONTAP DNS load balancing. For more information, see [TR-4523: DNS Load Balancing in ONTAP](#).

To create a LIF to use the service policy using the ONTAP CLI, run the following command:

```
network interface create
-vserver <name>
-lif <name>
-service-policy <name>
-home-node <node>
-home-port <port>
-address <number>
-netmask <number>
-status-admin up
```

Install a CA certificate

Using CA certificates creates a trusted relationship between client applications and the ONTAP object store server. A CA certificate should be installed on ONTAP before using it as object store that is accessible to remote clients.

Although ONTAP can generate self-signed certificates, using signed certificates from a third-party certificate authority is the recommended best practice.

To install a CA certificate using the ONTAP CLI, run the following command:

```
security certificate install -type server -vserver <name> -type server-ca
```

Create the object store server

The ONTAP object store server manages data as objects, as opposed to other storage architectures such as file or block storage.

To create an object store server using the ONTAP CLI, run the following command:

```
vserver object-store-server create
-vserver <name>
-object-store-server <FQDN>
-certificate-name <name>
-secure-listener-port <443>
-is-http-enabled <false>
```

Note: FabricPool must resolve this name to all IP addresses used by S3 data LIFs through DNS.

Create the bucket

To create a bucket using the ONTAP CLI, run the following command:

```
vserver object-store-server bucket create
-vserver <name>
-bucket <name>
-aggr-list <aggregate name>,<aggregate name>
-aggr-list-multiplier <number of constituent volumes per aggregate> (default 4)
-size <size>
```

Create a user

User authorization is required on all ONTAP object stores in order to restrict connectivity to authorized clients.

Note: All S3 users with valid access and a secret key-pair can access all buckets and objects in the SVM.

To create a user by using the ONTAP CLI, run the following command:

```
vserver object-store-server user create
-vserver <name>
-user <name>
```

To view the user's access and secret key by using the ONTAP CLI, run the following command:

Note: Advanced privilege level is required.

```
object-store-server user show
```

Root user

A root user (UID 0) is created by default when the bucket is created. The root user has full access to all buckets and objects. Do not use the root user for client application access. Additional users must be created for client access.

The ONTAP administrator must run the `object-store-server users regenerate-keys` command to set the access key and secret key for this user.

Configuration for local cluster tiering

Starting in ONTAP 9.8, FabricPool supports tiering to buckets in ONTAP, allowing for ONTAP-to-ONTAP tiering. This is an excellent option for customers who wish to repurpose existing FAS infrastructure as an object store endpoint.

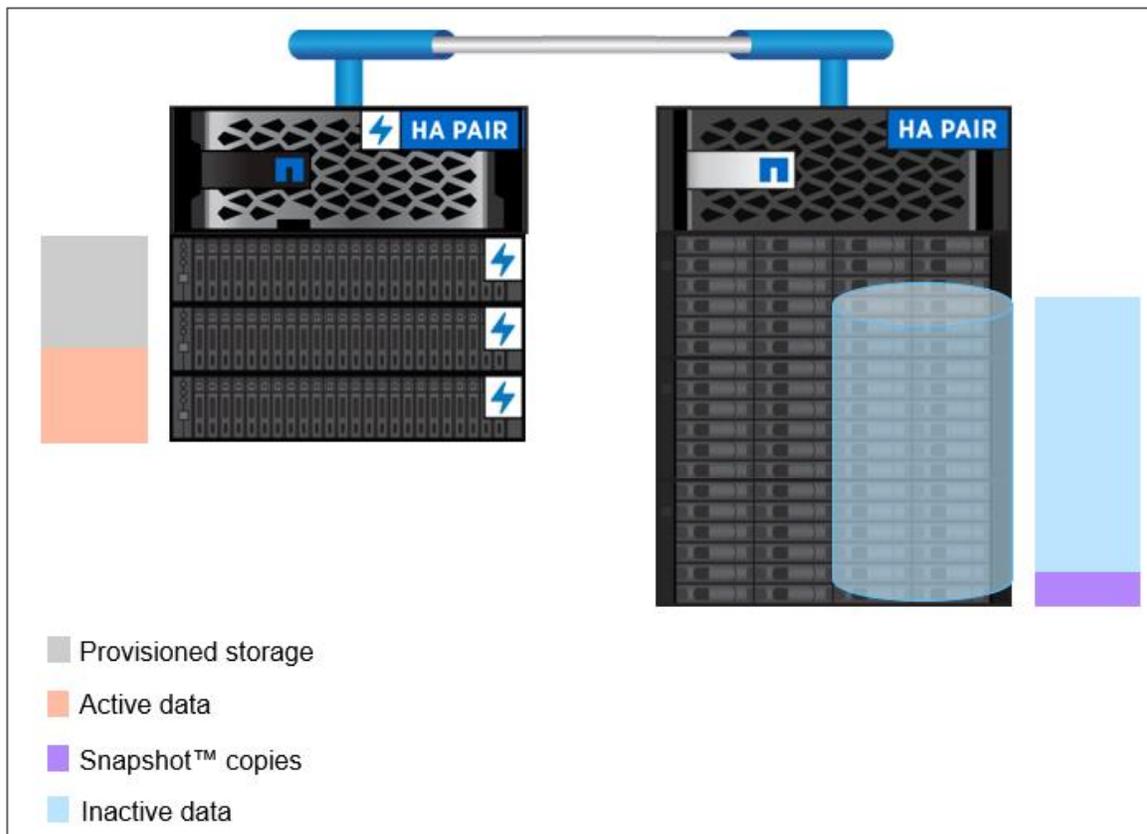
When configured for local cluster tiering, inactive data is tiered from local aggregates (typically SSD) to a local bucket (typically HDD) using cluster LIFs.

NetApp recommends using StorageGRID, NetApp's premier object store solution, when tiering more than 300TB of inactive data. A FabricPool license is not required when using ONTAP or StorageGRID as the cloud tier.

For more information on FabricPool, see [TR-4598: FabricPool Best Practices](#).

Note: Performance degradation might occur if cluster LIFs resources become saturated. To avoid this, NetApp recommends using two-node, or greater, clusters when tiering to a local bucket—the recommended best practice being an HA pair for the local tier and an HA pair for the local bucket. Tiering to local buckets on single-node clusters is not recommended.

Figure 3) Local cluster tiering.



ONTAP System Manager

The easiest way to create an object store for local tiering in ONTAP is by using ONTAP System Manager, which reduces multiple steps using the CLI to a few clicks. Object stores created using ONTAP system manager allow for less customization but use NetApp recommended best practices by default. Configuration via the CLI is required for custom configurations.

Configure the object store

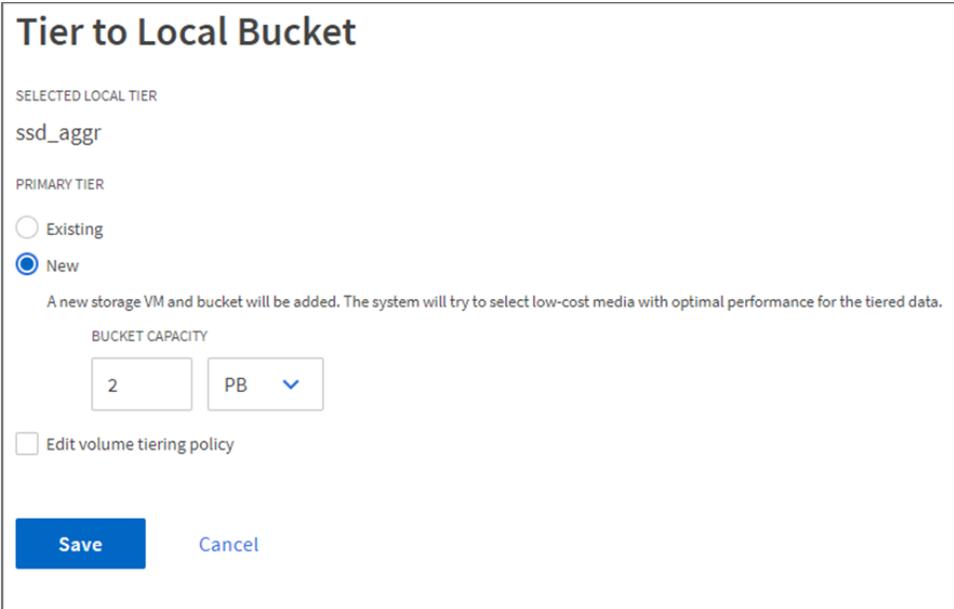
To create an object store used for local cluster tiering, complete the following steps:

1. Launch ONTAP System Manager.
2. Click STORAGE.
3. Click Tiers
4. Select a local tier.
5. Click More.
6. Select Tier to Local Bucket.
7. Select New if this is the first local bucket on the system.

A new SVM, object store server, and bucket are created. ONTAP System Manager creates the bucket on the least expensive media, prioritizing HDD > QLC > TLC > NVMe.

Select Existing if a local bucket has already been created.

Note: Attaching the same local bucket to all FabricPool local tiers in the cluster enables optimized volume moves. If a volume move's destination local tier uses the same bucket as the source local tier, data on the source volume that is stored in the bucket does not move back to the local tier. Optimized volume moves result in significant network efficiencies.



The screenshot shows a configuration window titled "Tier to Local Bucket". It includes the following elements:

- SELECTED LOCAL TIER:** A text field containing "ssd_aggr".
- PRIMARY TIER:** Two radio buttons: "Existing" (unselected) and "New" (selected).
- Message:** "A new storage VM and bucket will be added. The system will try to select low-cost media with optimal performance for the tiered data."
- BUCKET CAPACITY:** A numeric input field with "2" and a dropdown menu with "PB" selected.
- Checkbox:** "Edit volume tiering policy" (unchecked).
- Buttons:** "Save" (blue) and "Cancel" (grey).

8. Set bucket capacity.
9. Edit volume tiering policies (optional).
10. Click Save.

ONTAP CLI

Although the easiest way to create an object store for local tiering in ONTAP is by using ONTAP System Manager, object stores created using ONTAP System Manager allow for less customization.

For example, ONTAP System Manager automatically selects the local tiers (aggregates) used by a bucket for storage. Although ONTAP System Manager uses recommended best practices to do so, for complex environments, the selected local tiers might not be the same ones an experienced storage administrator would select.

Configuration using the ONTAP CLI is required for custom configurations.

To create an object store and bucket for local tiering using ONTAP CLI, complete the following steps:

1. Create the object store server on the Cluster SVM.
2. Create a bucket on a data SVM.
3. Create a user.
4. Add a cloud tier using the object store and bucket
5. Attach the cloud tier to a local tier

Create the object store server on the Cluster SVM

To create an object store server on the Cluster SVM using the ONTAP CLI, run the following command:

```
vserver object-store-server create
-vserver Cluster
-object-store-server <name> (This is the FGDN used by FabricPool)
-certificate-name <name>
```

Although installation and use of certificate authority (CA) certificates are recommended best practices, installation of CA certificates is not required. If not using a certificate, http must be enabled and https disabled:

```
vserver object-store-server create
-vserver Cluster
-object-store-server <name> (This is the FGDN used by FabricPool)
-is-http-enabled true
-is-https-enabled false
```

Set object store permissions

Permissions can be set at the object store level that will apply to all (or specified) buckets in the object store. To set an object store policy statement using the ONTAP CLI, run the following command:

```
vserver vserver object-store-server policy statement create
-vserver <data svm>
-policy <name>
-effect <allow/deny>
-action <*, GetObject, PutObject, DeleteObject, ListBucket, GetBucketAcl,GetObjectAcl,
ListBucketMultipartUploads, ListMultipartUploadParts>
-principal <S3 user or group> (maximum of 10 per policy)
-resource <bucket name>
```

Create a bucket on a data SVM

To create a bucket using the ONTAP CLI, run the following command:

```
vserver object-store-server bucket create
-vserver <name>
-bucket <name>
-aggr-list <aggregate name>,<aggregate name>
-aggr-list-multiplier <number of constituent volumes per aggregate> (default 4)
-size <size> (95GB minimum)
```

Note: Advanced privilege required to use `-aggr-list`.

Set bucket permissions

To set a bucket permission statement using the ONTAP CLI, run the following command:

```
vserver vserver object-store-server bucket policy add-statement
-vserver <data svm>
-bucket <name>
-effect <allow/deny>
-action <*, GetObject, PutObject, DeleteObject, ListBucket, GetBucketAcl,GetObjectAcl,
ListBucketMultipartUploads, ListMultipartUploadParts>
-principal <S3 user or group> (maximum of 10 per policy)
-resource <bucket name>
```

Create a user

User authorization is required on all ONTAP object stores in order to restrict connectivity to authorized clients.

Note: All S3 users with valid access and a secret key-pair can access all buckets and objects in the SVM.

To create a user by using the ONTAP CLI, run the following command:

```
vserver object-store-server user create
-vserver <name>
-user <name>
```

To view the user's access and secret key by using the ONTAP CLI, run the following command:

Note: Advanced privilege level is required.

```
object-store-server user show
```

User groups

User can be added to groups which can be associated with policy statements at the object store level or bucket level. To create a group policy and add users to using the ONTAP CLI, run the following command:

```
vserver vserver object-store-server group create
-vserver <data svm>
-name <group name>
-users <user1, user2, etc.
-policy <policy name>
```

Add a cloud tier using the object store and bucket

To add a cloud tier using the ONTAP CLI, run the following commands:

```
storage aggregate object-store config create
-object-store-name <name the cloud tier>
-provider-type ONTAP_S3
-server <name of the Cluster svm object store server>
-container-name <bucket-name>
-access-key <string>
-secret-password <string>
-ip-space Cluster
-ssl-enabled <true/false>
-is-certificate-validation-enabled true
-use-http-proxy false
-url-style <path-style/virtual-hosted-style>
```

Attach the cloud tier to a local tier

To attach the local bucket tier to a local tier (storage aggregate) by using the ONTAP CLI, run the following commands:

```
storage aggregate object-store attach
-aggregate <name>
-object-store-name <cloud tier name>
```

Note: Attaching a local bucket to a local tier is a permanent action. A local bucket cannot be unattached from a local tier after being attached. By using FabricPool Mirror, a different local bucket or cloud tier can be attached.

Security

Local tier

NetApp Storage Encryption (NSE), NetApp Volume Encryption (NVE), and NetApp Aggregate Encryption (NAE) work equally well for objects written to buckets in ONTAP. Neither NSE, NVE, nor NAE are required for S3 in ONTAP.

Over the wire

TLS/SSL encryption is enabled by default using a system-generated certificate. Using signed certificates from a third-party certificate authority is a recommended best practice.

Client-object store communication without TLS encryption (HTTP, Port 80) is supported but is not a recommended best practice.

Signature Version 4

S3 in ONTAP requires the use of Signature Version 4 (v4 signatures).

Note: Using v2 signatures result in a failure to connect. It is important to be aware of this because many client applications, including commonly used S3 browsers, use v2 signatures by default. Configure client applications to use v4 signatures to avoid connectivity errors.

Supported S3 actions

Buckets

Actions marked with an asterisk are supported by ONTAP, not S3 REST APIs.

- DeleteBucket*
- DeleteBucketPolicy*
- GetBucketAcl
- HeadBucket
- ListBuckets
- PutBucket*

Objects

- PutObject
- PutObjectTagging (9.9.1)
- GetObject
- GetObjectAcl
- GetObjectTagging (9.9.1)
- DeleteObject
- DeleteObjectTagging (9.9.1)
- HeadObject
- ListObjects
- ListObjectsV2
- ListParts
- UploadPart
- AbortMultipartUpload
- CompleteMultipartUpload
- CreateMultipartUpload
- ListMultipartUpload

Group policies

These operations are not specific to S3 and generally associated with Identity and Management (IAM). ONTAP supports these commands but does not use the IAM REST APIs.

- Create Policy
- AttachGroup Policy

User management

These operations are not specific to S3 and generally associated with IAM:

- CreateUser
- DeleteUser
- CreateGroup
- DeleteGroup

ONTAP 9.9.1

ONTAP 9.9.1 adds object metadata and tagging support to ONTAP S3.

- PutObject, CreateMultipartUpload, now include key-value pairs using `x-amz-meta-<key>`
For example: `x-amz-meta-project: ontap_s3`
- GetObject, and HeadObject now return user-defined metadata
- Tags can also be use with Buckets. Unlike metadata, tags can be read independently of objects using:
 - PutObjectTagging
 - GetObjectTagging
 - DeleteObjectTagging

Interoperability

The exceptions to normal interoperability that are listed in Table 1 are unique to ONTAP object stores.

Table 1) NetApp interoperability.

| Focus | Supported | Not supported |
|------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Data protection | <ul style="list-style-type: none"> • Cloud Sync | <ul style="list-style-type: none"> • Erasure coding • Information lifecycle management • NetApp MetroCluster™ • NDMP • NetApp SnapLock® technology • NetApp SnapMirror® technology • NetApp SyncMirror® technology • Object versioning • SMTape • SVM-DR • User-created snapshots • WORM |
| Encryption | <ul style="list-style-type: none"> • NetApp Aggregate Encryption (NAE) • NetApp Volume Encryption (NVE) • NetApp Storage Encryption (NSE) • TLS/SSL | <ul style="list-style-type: none"> • SLAG |
| Storage efficiency | <ul style="list-style-type: none"> • Deduplication • Compression • Compaction | Aggregate-level efficiencies |
| Storage virtualization | – | NetApp FlexArray® technology |
| QoS | QoS maximums (ceiling) QoS minimums (floors) | – |
| Additional features | – | <ul style="list-style-type: none"> • Audit • NetApp FPolicy software • Qtrees • Quotas |

Where to find additional information

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

- S3 Configuration Power Guide
<https://docs.netapp.com/ontap-9/topic/com.netapp.doc.pow-s3-cg/S3%20configuration.pdf>
- Provision Object Storage
https://docs.netapp.com/ontap/us-en/pdfs/sidebar/Provision_object_storage.pdf
- TR-4598: FabricPool Best Practices
<https://www.netapp.com/us/media/tr-4598.pdf>
- ONTAP 9 Documentation Center
<https://docs.netapp.com/ontap-9/index.jsp>
- ONTAP and ONTAP System Manager Documentation Resources
<https://www.netapp.com/us/documentation/ontap-and-oncommand-system-manager.aspx>
- NetApp Product Documentation
<https://www.netapp.com/us/documentation/index.aspx>

Version history

| Version | Date | Document version history |
|---------|--------------|------------------------------------------------------------------------------------------------------------------------------|
| 1.3 | August 2021 | Updated for 9.9.1. Support for object tagging. Added details regarding default provisioning capacities and permissioning. |
| 1.2 | March 2021 | Updated ONTAP CLI for local cluster tiering. |
| 1.1 | January 2021 | Updated supported S3 actions. |
| 1.0 | January 2021 | Initial release. |

Contact us

Let us know how we can improve this technical report.

Contact us at doccomments@netapp.com.

Include TR-4814: S3 in ONTAP in the subject line.

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 © 2020–2021 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.

RESTRICTED RIGHTS LEGEND: Use, duplication, or disclosure by the government is subject to restrictions as set forth in subparagraph (c)(1)(ii) of the Rights in Technical Data and Computer Software clause at DFARS 252.277-7103 (October 1988) and FAR 52-227-19 (June 1987).

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.

TR-4814-0821