



Technical Report

S3 in ONTAP best practices

ONTAP 9.13.1

John Lantz, NetApp
June 2023 | 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
S3 in multiprotocol NAS volumes	8
Configuration for native S3 applications and remote cluster tiering	9
ONTAP System Manager	9
ONTAP CLI	12
Configuration for local cluster tiering	15
ONTAP System Manager	16
ONTAP CLI	17
Configuration for S3 in multiprotocol NAS volumes	19
ONTAP System Manager	20
ONTAP CLI	22
Lifecycle rules	24
Expiration	24
Noncurrent Version Expiration	25
Abort Incomplete Multipart Upload	25
Security	26
Local tier	26
Over the wire	26
Signature Version 4	26

S3 SnapMirror	26
Snapshot copies	27
Protecting buckets using S3 SnapMirror	27
Requirements	27
Protection policies	28
Supported S3 actions	28
Buckets.....	28
Objects	28
Group policies.....	29
User management	29
Not supported in multiprotocol NAS volumes	29
S3 actions by release	30
ONTAP 9.13.1	30
ONTAP 9.12.1	30
ONTAP 9.11.1	30
ONTAP 9.10.1	31
ONTAP 9.9.1	31
Interoperability	32
Where to find additional information	33
Version history.....	34
Contact us	34
LIST OF TABLES	
Table 1) S3 interoperability.....	32
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

Beginning in ONTAP 9.8, ONTAP 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 the ONTAP storage efficiencies and security as they do so.

Beginning in ONTAP 9.12.1, the S3 protocol can also be enabled in multiprotocol NAS volumes that have been preconfigured to use NAS protocols. When the S3 protocol is enabled in multiprotocol NAS volumes, client applications can read and write data using S3, NFS, and SMB, which opens up a variety of additional use cases. One of the most common use cases is NAS clients writing data to a volume and S3 clients reading the same data and performing specialized tasks such as analytics, business intelligence, machine learning, and optical character recognition.

FabricPool endpoints

Beginning with 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**
 - S3 is supported on Cloud Volumes ONTAP for Azure using ONTAP 9.9+.
 - S3 is supported on Cloud Volumes ONTAP for AWS and Amazon FSx for NetApp ONTAP using ONTAP 9.11+.
 - S3 is supported on Cloud Volumes ONTAP for Google Cloud using ONTAP 9.12+.

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 four-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 HA pair 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:

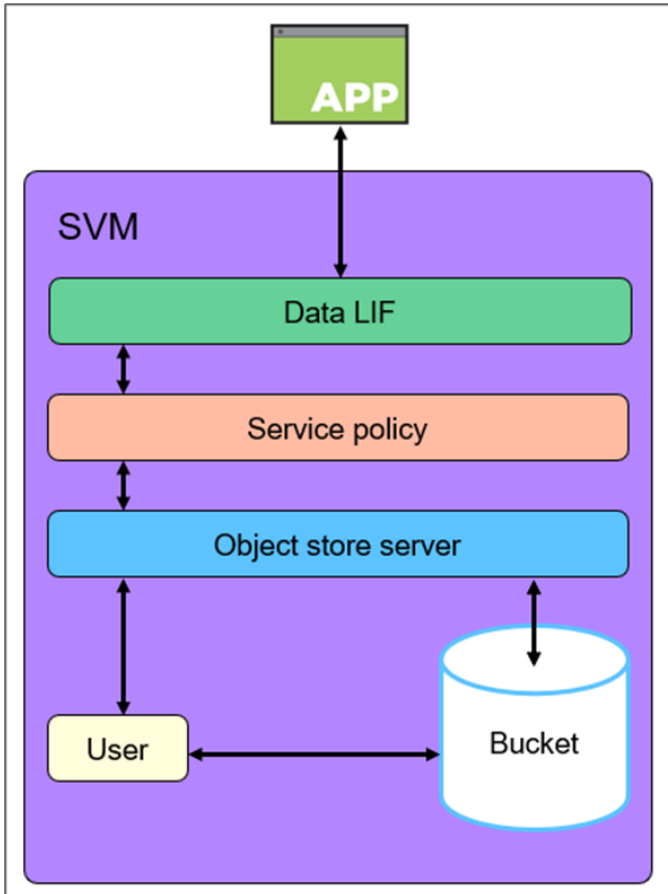
```
system license add <license_key>
```

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 make sure that 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.

For more information about FlexGroup volumes, see [TR-4557: NetApp ONTAP FlexGroup Volumes](#).

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](#) 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 is reduced by half until it can be provisioned in the existing environment. For example, in a 300TB environment, a FlexGroup volume is 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

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 is automatically provisioned on a newly created FlexGroup volume.

Users

User authorization is required on all ONTAP object stores 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.

S3 in multiprotocol NAS volumes

When S3 is used in multiprotocol NAS volumes (ONTAP 9.12.1+), it is mapped to existing NAS hierarchies. For example, buckets are mapped to a volume or a directory inside a volume. NAS security configurations, including file, directory, and user permissions, are preserved and mapped to S3 users in the same way that NFS and SMB configurations are mapped to each other.

See [TR-4887: Multiprotocol NAS in NetApp ONTAP](#) for more information on how names and permissions are mapped in multiprotocol volumes.

Objects are mapped to files and are presented to S3 clients using a naming scheme based on the underlying NAS hierarchy with folder/object corresponding to directory/file.

Note: Because the underlying architecture is file-based, not object-based, S3 in multiprotocol NAS volumes imposes NAS-related limits that might not exist when using native S3. For example, file and directory names are limited to 255 characters and 1024-byte paths, so corresponding object names are limited to 255 characters and 1024 bytes as well.

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 the CLI is required for more custom configurations.

ONTAP System Manager

To create an object store, bucket, and permission, users using ONTAP System Manager must 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 ✕

NAME

STORAGE VM

CAPACITY

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.
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

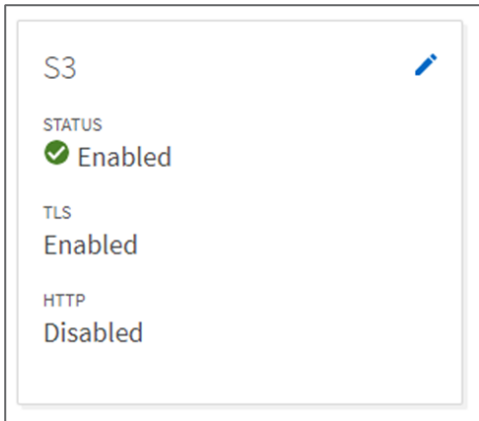
User authorization is required on all ONTAP object stores to restrict connectivity to authorized clients. Access to specific buckets or S3 actions can be allowed, denied, or made conditional at the user and group level using [permissions](#).

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) used 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 make sure that applications using the LIF work as expected.

Create a data LIF to use with 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 an 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 an 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
```

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 a user

User authorization is required on all ONTAP object stores 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.

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>
-type s3
-used-as-capacity-tier <true|false>
-aggr-list <aggregate name>, <aggregate name> (option for non-capacity tier)
-exclude-aggr-list <aggregate name>, <aggregate name> (option for capacity tier)
-aggr-list-multiplier <number of constituent volumes per aggregate> (default 4)
-size <size>
```

Beginning with ONTAP 9.11.1, ONTAP S3 supports bucket versioning. Enabling versioning allows for the creation of multiple versions of an object. Much like Snapshot copies, these objects can be retrieved and restored, enabling client applications to restore deleted objects or retrieve earlier versions of an object.

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

```
vserver object-store-server bucket modify
-vserver <name>
-bucket <name>
-versioning-state <disabled|enabled|suspend>
```

Note: The default versioning state is Disabled.

Configuration for local cluster tiering

Beginning with 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, 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.

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 when 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

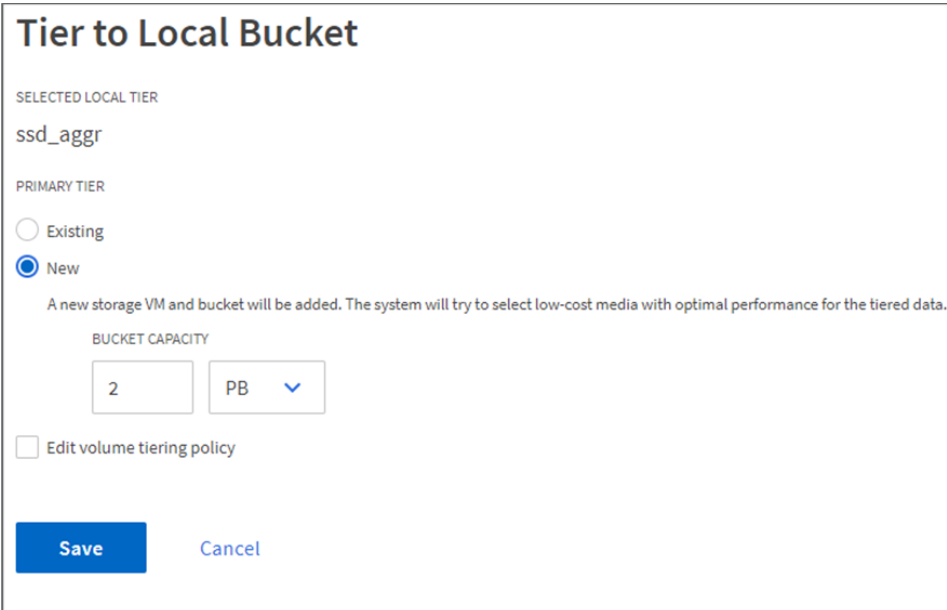
To create an object store used for local cluster tiering, complete the following steps:

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 local FabricPool 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 button options: "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)
-is-http-enabled true
-is-https-enabled false
-status-admin up
```

Although installation and use of certificate authority (CA) certificates are recommended best practices, installation of CA certificates is not required when tiering locally. If you are not using a certificate, HTTP must be enabled and HTTPS must be disabled.

Set object-store permissions

Permissions can be set at the object store level that 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, etc.>
-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>
-type s3
-used-as-capacity-tier true
-exclude-aggr-list <aggregate name>,<aggregate name>
-aggr-list-multiplier <number of constituent volumes per aggregate> (default 4)
-size <size> (95GB minimum)
```

Note: Advanced privileges are 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, etc.>
-principal <S3 user or group> (maximum of 10 per policy)
-resource <bucket name, bucket-name/*>
```

Note: To add anonymous access, a principal must be configured as `*`.

Create a user

User authorization is required on all ONTAP object stores 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

A user can be added to groups that can be associated with policy statements at the object store level or the bucket level. To create a group policy and add users to it 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 a FabricPool mirror, a different local bucket or cloud tier can be attached.

Configuration for S3 in multiprotocol NAS volumes

Beginning in ONTAP 9.12.1, S3 can be enabled in pre-existing NAS volumes that have been fully configured to serve NFS or SMB clients. Enabling the S3 protocol on a volume that supports NFS and/or SMB but has not been configured to serve data does not work. ONTAP must be able to map S3 users to pre-existing users created with Unix or Windows security styles. Enabling S3 on a NAS volume that has not been configured to serve NFS or SMB clients does not work.

To enable NAS protocols, see the following resources:

- [Provision NAS storage for Linux servers using NFS](#)
- [Provision NAS storage for Windows servers using SMB](#)

Note: Multiprotocol NAS volumes are NAS volumes that present NAS hierarchies and files as buckets and objects. Actions and capabilities associated with metadata, multipart objects, tags, and versioning are **not supported** when using S3 in multiprotocol NAS volumes. Clients that require these actions and capabilities should use native ONTAP S3.

ONTAP System Manager

The easiest way to enable S3 in a multiprotocol NAS volume in ONTAP is by using the ONTAP System Manager; this reduces the 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 enable S3 in a multiprotocol NAS volume using ONTAP System Manager, complete the following steps:

1. Enable S3 on the SVM.
2. Create the bucket.
3. Enable name mapping.
4. Add bucket permissions.

Enable S3 on the SVM

1. Launch ONTAP System Manager.
2. Click **Storage**.
3. Click **Storage VMs**.
4. Select a SVM configured to use NFS or SMB/CIFS protocols.
5. Click **Settings**.
6. Click the S3 gear icon.
7. Name the S3 server.

Note: The server name is used as the fully qualified domain name (FQDN) by client applications.

8. 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.
9. Enter network interfaces for the nodes.

Create the bucket

1. Click **Storage**.
2. Click **Buckets**.
3. Click **Add**.
4. Name the bucket.
5. Select the SVM/object store that the bucket will be assigned to. This should be the same S3 server created on the multiprotocol SVM earlier. Clicking **More Options** allows you to map the bucket to a specific folder inside the volume.

Note: S3 buckets in multiprotocol NAS volumes cannot be used for tiering as FabricPool cloud tiers.

6. Click **Save**.

Enable name mapping

User authorization is required on all ONTAP object stores 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. Click **Storage**.
2. Click **Storage VMs**.
3. Select the SVM to add users and groups to.
4. Click the **Settings** tab.
5. Click **Name Mapping**.
6. Select **S3 to Windows** or **S3 to Unix** (both can be used).
7. Click **Add**.
8. Select the pattern (S3) and replacement (Windows or Unix).

Add bucket permissions

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. Click **Storage**.
2. Click **Buckets**.
3. Select a bucket.
4. Click **Edit**.
5. Set principal users. Options include All users of the SVM (default), All public and anonymous users, and individual users associated with the SVM.
6. Set effect. Options include **Allow** (default) and **Deny**.
7. Set actions.
8. Set resources. bucket-name and bucket-name/* are used by default. NAS directories/folder paths can also be used.
9. Set conditions.
10. 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.

ONTAP CLI

To enable S3 in a multiprotocol NAS volume using ONTAP CLI, complete the following steps:

1. Add the S3 service policy.
2. Verify the data LIF.
3. Install a CA certificate.
4. Create the object store server.
5. Create a bucket.
6. Enable name mapping.

Add the S3 service policy

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

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

```
network interface service-policy add-service
-vserver <name>
-policy <name>
-services data-s3-server
```

Note: S3 in multiprotocol volumes requires a pre-existing SVM configured to serve NAS data using data-core, and data-nfs, and/or data-cifs services.

Verify the data LIF

SVMs hosting object store servers require data LIFs to communicate with client applications using NFS, SMB/CIFS, and S3. NetApp recommends using data LIFs on all nodes as a best practice.

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 verify the data LIF has already been configured to support client traffic, run the following command:

```
network interface show
-vserver <name>
```

Install 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
```

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>
```

Create a bucket

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

```
vserver object-store-server bucket create
-vserver <name>
-bucket <name>
-type nas
-nas-path <junction_path>
```

Note: Because S3 in multiprotocol NAS volumes uses pre-existing FlexVol or FlexGroup volumes, a new FlexGroup volume exclusively for S3 objects is not created. The volume already exists, so there is no need to define aggregates, constituent volumes, or size.

Note: ONTAP S3 supports object versioning in native S3 buckets. Object versioning is not supported in multiprotocol NAS volumes. Consider using SnapMirror instead.

Enable name mapping

User authorization is required on all ONTAP object stores to restrict connectivity to authorized clients. When using S3 in multiprotocol NAS volumes, ONTAP must be able to map S3 users to pre-existing users created with Unix or Windows security styles.

To map S3 users to existing Unix and/or Windows users, run the following command:

```
vserver name-mapping create
-vserver <name>
-direction <s3-win|s3-unix>
-position <1|2>
-pattern <S3 user>
-replacement <unix or windows user>
```

Create a bucket policy

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, etc.>
-principal <S3 user or group> (maximum of 10 per policy)
-resource <bucket name, bucket-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
```

Lifecycle rules

Beginning in ONTAP 9.13.1, ONTAP S3 supports expiration rules that can be used to provide bucket-level information lifecycle management (ILM) capabilities. Using expiration rules, retention policies can be created that apply to specific objects in ONTAP S3 buckets. Each expiration rule consists of:

- Metadata containing the rule ID and status indicating whether the rule is enabled or disabled.
- One or more expiration actions. Options include: Expiration, Noncurrent Version Expiration, and Abort Incomplete Multipart Upload.
- Filters used to match the set of objects that needs to be deleted. Filters include object prefix, tags, object size, age, etc. If no filters are set, the expiration rule will be applied to all objects in the bucket.

After a bucket lifecycle rule has been created, the expiration rule will be added to the header of all new objects put in the bucket.

Note: ONTAP S3 does not support transition rules.

Expiration

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

```
vserver object-store-server bucket lifecycle-management-rule create
-vserver <name>
-bucket <name>
-rule-id <name>
-index <#>
-is-enabled <true|false>
-action Expiration
-obj-age-days <#>
-obj-exp-date <"MM/DD/YYYY HH:MM:SS">
-expired-obj-del-marker <true|false>
-prefix <name>
-tags <name, name> (maximum of 4)
-obj-size-greater-than <#[KB|MB|GB|TB|PB]>
-obj-size-less-than <#[KB|MB|GB|TB|PB]>
```

Examples

Expire objects starting with 'test' after 30 days

```
vserver object-store-server bucket lifecycle-management-rule create
-vserver svml
-bucket mybucket
-rule-id rule1
-index 1
-is-enabled true
-action Expiration
-prefix testobj
-obj-age-days 30
```

Expire objects tagged "proj1=test" on January 1st, 2025

```
vserver object-store-server bucket lifecycle-management-rule create
-vserver svml
-bucket mybucket
-rule-id rule2
-index 2
-is-enabled true
-action Expiration
-tags proj1=test
-obj-exp-date "2025-01-01T00:00:00"
```


Expire objects ranging between 100MB and 1GB after 365 days

```
vserver object-store-server bucket lifecycle-management-rule create
-vserver svml
-bucket mybucket
-rule-id rule3
-index 3
-is-enabled true
-action Expiration
-obj-size-greater-than 100MB
-obj-size-less-than 1GB
-obj-age-days 365
```

Noncurrent Version Expiration

To create a bucket noncurrent version expiration rule using the ONTAP CLI, run the following command:

```
vserver object-store-server bucket lifecycle-management-rule create
-vserver <name>
-bucket <name>
-rule-id <name>
-index <#>
-is-enabled <true|false>
-action NonCurrentVersionExpiration
-new-non-curr-versions <#>
-non-curr-days <#>
-prefix <name>
-tags <name, name> (maximum of 4)
-obj-size-greater-than <#[KB|MB|GB|TB|PB]>
-obj-size-less-than <#[KB|MB|GB|TB|PB]>
```

Examples

Expire non-current versions of objects after 30 days, retaining up to 10 non-current versions

```
vserver server object-store-server bucket lifecycle-management-rule create
-vserver svml
-bucket mybucket
-rule-id rule4
-index 4
-action NoncurrentVersionExpiration
-is-enabled true
-non-curr-days 30
-new-non-curr-versions 10
```

Abort Incomplete Multipart Upload

To create a bucket Abort Incomplete Multipart Upload rule using the ONTAP CLI, run the following command:

```
vserver object-store-server bucket lifecycle-management-rule create
-vserver <name>
-bucket <name>
-rule-id <name>
-index <#>
-is-enabled <true|false>
-action AbortIncompleteMultipartUpload
-after-initiation-days <#>
-prefix <name>
-obj-size-greater-than <#[KB|MB|GB|TB|PB]>
-obj-size-less-than <#[KB|MB|GB|TB|PB]>
```

Examples

Abort incomplete multipart uploads after 7 days

```
vserver object-store-server bucket lifecycle-management-rule create
-vserver svml
-bucket mybucket
-rule-id rule4
-index 4
-action AbortIncompleteMultipartUpload
-is enabled true
-after-initiation-days 7
```

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

Prior to ONTAP 9.11.1, S3 in ONTAP did not support Signature Version 2 (v2 signatures) and required the use of v4 signatures.

Note: Prior to ONTAP 9.11.1, using v2 signatures results 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. NetApp recommends that client applications use v4 signatures when possible.

S3 SnapMirror

Beginning with ONTAP 9.10.1, data in ONTAP S3 buckets can be protected by S3 SnapMirror. SnapMirror allows you to define synchronization schedules to meet specific recovery point objectives (RPOs). SnapMirror can also be used to create a variety of data protection relationships including source-to-destination, fan-out, and cascading. Fan-in relationships are not supported.

S3 SnapMirror has two primary use cases:

- Backup and recovery, where the objective is to restore from the destination bucket to the source bucket with no intention of failing over to the destination bucket. If the S3 SnapMirror relationship is broken, objects in the destination bucket remain read-only.
- Disaster recovery (DR) and failover, where the objective is to be able to serve data to client applications from the destination bucket in the event a disaster event takes place. If the S3 SnapMirror relationship is broken, the destination bucket supports reads and writes. ONTAP is the only destination target that supports DR and failover operations.

Note: S3 SnapMirror is exclusively for the protection of native S3 objects. NAS and SAN data tiered by FabricPool to ONTAP S3 buckets is protected as normal using SnapMirror or other data protection applications—not S3 SnapMirror.

Snapshot copies

Although ONTAP S3 supports object versioning, because object storage is not transactional like file or block storage, S3 SnapMirror does not make use of Snapshot copies that capture the state of a file at a specific point in time and act as highly efficient point-in-time deltas.

Protecting buckets using S3 SnapMirror

The full set of instructions for mirroring bucket data, setting data protection policies restoring data, and performing takeover operations is found in [Protecting buckets with S3 SnapMirror](#).

Requirements

S3 SnapMirror requires ONTAP 9.10.1 or later. Prior to ONTAP 9.10.1, data protection could be achieved by using NetApp Cloud Sync.

Destination targets

- NetApp:
 - ONTAP
 - StorageGRID
 - Cloud Volumes ONTAP for AWS
 - Cloud Volumes ONTAP for Azure
 - Amazon FSx for NetApp ONTAP
- Third party:
 - Amazon S3

Note: When using ONTAP as a destination target, S3 SnapMirror supports creating a data protection relationship between source and destination buckets in both same- and remote-cluster relationships. Same-cluster relationships do not protect data from cluster or site-wide disaster events. NetApp recommends using S3 SnapMirror with destination targets outside the local cluster.

License

Enabling S3 SnapMirror requires the use of the Data Protection Bundle. Both the Data Protection Bundle and the Hybrid Cloud Bundle are required when using S3 SnapMirror to replicate data to third-party object stores such as Amazon S3.

Certificate authority certification

When using TLS, S3 SnapMirror must be configured to use the destination's CA certificates on both the source and the destination.

Although CA certificates are not required, ONTAP S3 uses TLS and self-signed certificates by default. Using signed certificates from a third-party certificate authority is the recommended best practice.

Cluster peer relationship

A cluster peer relationship must be established before using a different ONTAP cluster as an S3 SnapMirror target destination. For more information, see [Prepare for mirroring and vaulting](#) and [Create a cluster peer relationship](#).

Cloud object store

A cloud object store—such as StorageGRID, Amazon S3, or Microsoft Azure Blob Storage—must be identified by ONTAP before it can be used as a S3 SnapMirror target destination. For more information, see [Add a Cloud Object Store](#).

Protection policies

S3 SnapMirror creates a data protection relationship that replicates data in a source bucket to a destination bucket. Replication of data is based on the protection policy selected when a bucket is protected. The S3 SnapMirror default protection policy, *Continuous*, replicates data continuously to the destination bucket, using a one-hour RPO, and does not throttle data.

Protection policies can be created and saved for use when protecting one or more buckets. Customizable parameters include the following:

- **Policy type.** S3 SnapMirror protection policies must use the Continuous policy type. Asynchronous and Synchronous policies can be created using the Add Protection Policy menu but cannot be selected as a protection policy when protecting a bucket.
- **Throttle.** Set the maximum bandwidth allowed to attain the RPO. The default of zero does not set any throttle.
- **RPO.** Set a delay between the time a change is made in the source bucket and when that change is pushed to the destination bucket. The default is one hour.

Supported S3 actions

Buckets

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

- CreateBucket (9.11.1)
- DeleteBucket (9.11.1)
- DeleteBucketLifecycleConfiguration (9.13.1)
- DeleteBucketPolicy (9.12.1)
- GetBucketAcl
- GetBucketLifecycleConfiguration (9.13.1)
- GetBucketLocation (9.10.1)
- GetBucketPolicy (9.12.1)
- GetBucketVersioning (9.11.1)
- HeadBucket
- ListBuckets
- ListBucketVersioning (9.11.1)
- PutBucket*
- PutBucketLifecycleConfiguration (9.13.1)
- PutBucketPolicy (9.12.1)
- PutBucketVersioning (9.11.1)

Objects

- AbortMultipartUpload
- CompleteMultipartUpload

- CopyObject (9.12.1)
- CreateMultipartUpload
- DeleteObject
- DeleteObjects (9.11.1)
- DeleteObjectTagging (9.9.1)
- GetObject
- GetObjectAcl
- GetObjectTagging (9.9.1)
- HeadObject
- ListMultipartUpload
- ListObjectVersions (9.11.1)
- ListObjects
- ListParts
- PutObject
- PutObjectTagging (9.9.1)
- HeadObject
- UploadPart
- UploadPartCopy (9.12.1)

Group policies

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

ONTAP S3 groups can have a maximum of 10 attached policies. Group policies can have a maximum of five statements. Each statement can have a maximum of 10 resources.

- Create Policy
- AttachGroup Policy

User management

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

- CreateUser
- DeleteUser
- CreateGroup
- DeleteGroup

Not supported in multiprotocol NAS volumes

Actions and capabilities associated with metadata, multipart objects, tags, and versioning are **not supported** when using S3 in multiprotocol NAS volumes. This includes the following:

- Key-values pairs using `x-amz-meta-<key>` are not saved, and request headers using `x-amz-meta` are ignored.
- Requests to update tags are rejected, and headers using `x-amz-tagging` are ignored.
- AbortMultipartUpload
- CompleteMultipartUpload
- CreateMultipartUpload

- DeleteObjectTagging
- GetBucketVersioning
- GetObjectTagging
- PutBucketVersioning
- PutObjectTagging
- ListBucketVersioning
- ListMultipartUpload
- ListObjectVersions

S3 actions by release

ONTAP 9.13.1

ONTAP 9.13.1 adds bucket lifecycle configuration.

- DeleteBucketLifecycleConfiguration
- GetBucketLifecycleConfiguration
- PutBucketLifecycleConfiguration

ONTAP 9.12.1

ONTAP 9.12.1 adds bucket policies and the ability to copy objects.

- DeleteBucketPolicy
- GetBucketPolicy
- PutBucketPolicy
- CopyObject
- UploadPartCopy

ONTAP 9.11.1

ONTAP 9.11.1 adds versioning, presigned URLs, chunked uploads, and support for common S3 actions such as creating and deleting buckets using S3 APIs.

- ONTAP S3 now supports chunked uploads signing requests using `x-amz-content-sha256:STREAMING-AWS4-HMAC-SHA256-PAYLOAD`
- ONTAP S3 now supports client applications using presigned URLs to share objects or allow other users to upload objects without requiring user credentials.
- CreateBucket
- DeleteBucket
- GetBucketVersioning
- ListBucketVersioning
- PutBucketVersioning
- DeleteObjects
- ListObjectVersions

Note: Because the underlying FlexGroup is not created until the first bucket is, a bucket must first be created in ONTAP before an external client can create a bucket using CreateBucket.

ONTAP 9.10.1

ONTAP 9.10.1 adds support for S3 SnapMirror and GetBucketLocation.

- GetBucketLocation

ONTAP 9.9.1

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

- PutObject and 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 used with buckets. Unlike metadata, tags can be read independently of objects using:
 - PutObjectTagging
 - GetObjectTagging
 - DeleteObjectTagging

Interoperability

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

Table 1) S3 interoperability.

Focus	Supported	Not supported
Data protection	<ul style="list-style-type: none"> • Cloud Sync • S3 SnapMirror (9.10.1) • Unmirrored NetApp MetroCluster™ aggregates (9.12.1) 	<ul style="list-style-type: none"> • Erasure coding • Mirrored NetApp MetroCluster aggregates • NDMP • NetApp SnapLock® technology • NetApp SnapMirror® technology • NetApp SyncMirror® technology • SMTape • SVM-DR • WORM
Encryption	<ul style="list-style-type: none"> • NetApp Aggregate Encryption (NAE) • NetApp Storage Encryption (NSE) • NetApp Volume Encryption (NVE) • TLS/SSL 	<ul style="list-style-type: none"> • SLAG
Storage efficiency	<ul style="list-style-type: none"> • Compaction • Compression • Deduplication • Temperature Sensitive Storage Efficiency (TSSE) 	<ul style="list-style-type: none"> • Aggregate-level efficiencies
Storage virtualization	–	<ul style="list-style-type: none"> • NetApp FlexArray® technology
QoS	<ul style="list-style-type: none"> • QoS maximums (ceiling) • QoS minimums (floors) 	–
Additional features	<ul style="list-style-type: none"> • Audit • Bucket Lifecycle Management (9.13.1) • FabricPool cloud tier (native S3 only) • FabricPool local tier (NAS volumes only) 	<ul style="list-style-type: none"> • FabricPool cloud tier (NAS volumes only) • FabricPool local tier (native S3 only) • 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 overview
<https://docs.netapp.com/us-en/ontap/s3-config/index.html>
- Protect buckets with S3 SnapMirror
https://docs.netapp.com/us-en/ontap/pdfs/sidebar/Protect_buckets_with_S3_SnapMirror.pdf
- S3 object storage management
<https://docs.netapp.com/us-en/ontap/object-storage-management/index.html>
- TR-4598: FabricPool Best Practices
<https://www.netapp.com/us/media/tr-4598.pdf>
- TR-4887: Multiprotocol NAS in NetApp ONTAP
<https://www.netapp.com/pdf.html?item=/media/27436-tr-4887.pdf>
- ONTAP 9 Documentation Center
<https://docs.netapp.com/ontap-9/index.jsp>

Version history

Version	Date	Document version history
1.7	June 2023	Updated for 9.13.1. Support for bucket lifecycle management.
1.6	March 2023	Updated for 9.12.1. Support for S3 in multiprotocol NAS volumes. Support for DeleteBucketPolicy, GetBucketPolicy, PutBucketPolicy, CopyObject, and the UploadPartCopy S3 actions. ONTAP S3 now supported on unmirrored MetroCluster aggregates and NetApp Cloud Volumes ONTAP for Google Cloud.
1.5	August 2022	Updated for 9.11.1. Support for CreateBucket, DeleteBucket, GetBucketVersioning, ListBucketVersioning, PutBucketVersioning, DeleteObjects, and the ListObjectVersions S3 actions. ONTAP S3 now supported on NetApp Cloud Volumes ONTAP for AWS and Amazon FSx for NetApp ONTAP.
1.4	February 2022	Updated for 9.10.1. Support for S3 SnapMirror and the GetBucketLocation S3 action. Updated ONTAP CLI for configuration cluster tiering.
1.3	August 2021	Updated for 9.9.1. Support for object tagging. Added details regarding default provisioning capacities and permissioning. ONTAP S3 now supported on NetApp Cloud Volumes ONTAP for Azure.
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 best practices 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 © 2020–2023 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.

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-4814-0323