



Add storage capacity to an S3-enabled SVM

ONTAP 9

NetApp
February 12, 2026

Table of Contents

Add storage capacity to an S3-enabled SVM	1
Create an ONTAP S3 bucket	1
Create S3 buckets with the ONTAP CLI	2
Create S3 buckets with System Manager	3
Increase or decrease the ONTAP S3 bucket size	5
Create an ONTAP S3 bucket on a mirrored or unmirrored aggregate in a MetroCluster configuration	5
Process to create buckets	6
Create an ONTAP S3 bucket lifecycle management rule	10
Manage lifecycle management rules with the CLI	11
Manage lifecycle management rules with System Manager	12
Create an ONTAP S3 user	14
Create or modify ONTAP S3 user groups to control access to buckets	15
Regenerate ONTAP S3 keys and modify their retention period	16

Add storage capacity to an S3-enabled SVM

Create an ONTAP S3 bucket

S3 objects are kept in *buckets*. They are not nested as files inside a directory inside other directories.

Before you begin

A storage VM containing an S3 server must already exist.

About this task

- Beginning with ONTAP 9.14.1, automatic resizing has been enabled on S3 FlexGroup volumes when buckets are created on them. This eliminates excessive capacity allocation during bucket creation on existing and new FlexGroup volumes. FlexGroup volumes are resized to a minimum required size based on the following guidelines. The minimum required size is the total size of all the S3 buckets in a FlexGroup volume.
 - Beginning with ONTAP 9.14.1, if an S3 FlexGroup volume is created as part of a new bucket creation, the FlexGroup volume is created with the minimum required size.
 - If an S3 FlexGroup volume was created prior to ONTAP 9.14.1, the first bucket created or deleted subsequent to ONTAP 9.14.1 resizes the FlexGroup volume to the minimum required size.
 - If an S3 FlexGroup volume was created prior to ONTAP 9.14.1, and already had the minimum required size, the creation or deletion of a bucket subsequent to ONTAP 9.14.1 maintains the size of the S3 FlexGroup volume.
- Storage service levels are predefined adaptive Quality of Service (QoS) policy groups, with *value*, *performance*, and *extreme* default levels. Instead of one of the default storage service levels, you can also define a custom QoS policy group and apply it to a bucket. For more information about storage service definitions, see [Storage service definitions](#). For more information about performance management, see [Performance management](#). Beginning with ONTAP 9.8, when you provision storage, QoS is enabled by default. You can disable QoS, or choose a custom QoS policy during the provisioning process or at a later time.
- If you are configuring local capacity tiering, you create buckets and users in a data storage VM, not in the system storage VM where the S3 server is located.
- For remote client access, you must configure buckets in an S3-enabled storage VM. If you create a bucket in a storage VM that is not S3-enabled, it will only be available for local tiering.
- Beginning with ONTAP 9.14.1, you can [create a bucket on a mirrored or unmirrored aggregate in a MetroCluster configuration](#).
- For the CLI, when you create a bucket, you have two provisioning options:
 - Let ONTAP select the underlying aggregates and FlexGroup components (default)
 - ONTAP creates and configures a FlexGroup volume for the first bucket by automatically selecting the aggregates. It will automatically select the highest service level available for your platform, or you can specify the storage service level. Any additional buckets you add later in the storage VM will have the same underlying FlexGroup volume.
 - Alternatively, you can specify whether the bucket will be used for tiering, in which case ONTAP tries to select low-cost media with optimal performance for the tiered data.
 - You select the underlying aggregates and FlexGroup components (requires advanced privilege command options): You have the option to manually select the aggregates on which the bucket and

containing FlexGroup volume must be created, and then specifying the number of constituents on each aggregate. When adding additional buckets:

- If you specify aggregates and constituents for a new bucket, a new FlexGroup will be created for the new bucket.
- If you do not specify aggregates and constituents for a new bucket, the new bucket will be added to an existing FlexGroup. See [FlexGroup volumes management](#) for more information.

When you specify aggregates and constituents when creating a bucket, no QoS policy groups, default or custom, are applied. You can do so later with the `vserver object-store-server bucket modify` command.

Learn more about `vserver object-store-server bucket modify` in the [ONTAP command reference](#).

Note: If you are serving buckets from Cloud Volumes ONTAP, you should use the CLI procedure. It is strongly recommended that you manually select the underlying aggregates to ensure that they are using one node only. Using aggregates from both nodes can impact performance, because the nodes will be in geographically separated availability zones and hence susceptible to latency issues.

Create S3 buckets with the ONTAP CLI

1. If you plan to select aggregates and FlexGroup components yourself, set the privilege level to advanced (otherwise, admin privilege level is sufficient): set `-privilege advanced`
2. Create a bucket:

```
vserver object-store-server bucket create -vserver <svm_name> -bucket <bucket_name> -size [integer{KB|MB|GB|TB|PB}] [-comment text] [additional_options]
```

The storage VM name can be either a data storage VM or Cluster (the system storage VM name) if you are configuring local tiering.

If you want ONTAP to create a bucket based on performance or usage, use one of the following options:

- service level

Include the `-storage-service-level` option with one of the following values: `value`, `performance`, or `extreme`.

- tiering

Include the `-used-as-capacity-tier true` option.

If you want to specify the aggregates on which to create the underlying FlexGroup volume, use the following options:

- The `-aggr-list` parameter specifies the list of aggregates to be used for FlexGroup volume constituents.

Each entry in the list creates a constituent on the specified aggregate. You can specify an aggregate multiple times to have multiple constituents created on the aggregate.

For consistent performance across the FlexGroup volume, all of the aggregates must use the same disk type and RAID group configurations.

- The `-aggr-list-multiplier` parameter specifies the number of times to iterate over the aggregates that are listed with the `-aggr-list` parameter when creating a FlexGroup volume.

The default value of the `-aggr-list-multiplier` parameter is 4.

3. Add a QoS policy group if needed:

```
vserver object-store-server bucket modify -bucket bucket_name -qos-policy  
-group qos_policy_group
```

4. Verify bucket creation:

```
vserver object-store-server bucket show [-instance]
```

Example

The following example creates a bucket for storage VM `vs1` of size 1TB and specifying the aggregate:

Learn more about the commands described in this procedure in the [ONTAP command reference](#).

```
cluster-1::>*> vserver object-store-server bucket create -vserver  
svm1.example.com -bucket testbucket -aggr-list aggr1 -size 1TB
```

Create S3 buckets with System Manager

1. Add a new bucket on an S3-enabled storage VM.

- a. Click **Storage > Buckets**, then click **Add**.

- b. Enter a name, select the storage VM, and enter a size.

- If you click **Save** at this point, a bucket is created with these default settings:
 - No users are granted access to the bucket unless any group policies are already in effect.



You should not use the S3 root user to manage ONTAP object storage and share its permissions, because it has unlimited access to the object store. Instead, create a user or group with administrative privileges that you assign.

- A Quality of Service (performance) level that is the highest available for your system.
- Click **Save** to create a bucket with these default values.

Configure additional permissions and restrictions

You can click **More Options** to configure settings for object locking, user permissions, and performance level when you configure the bucket, or you can modify these settings later.

If you intend to use the S3 object store for FabricPool tiering, consider selecting **Use for tiering** (use low-cost media with optimal performance for the tiered data) rather than a performance service level.

If versioning is enabled on a bucket, Object Lock retention time can be placed on specific versions of an object using S3 clients. Locking a specific version of an object does not prevent other versions of the object from being deleted. If you want to enable versioning for your objects for later recovery, select **Enable Versioning**. Versioning is enabled by default if you are enabling object locking on the bucket. For information about object versioning, see the [Using versioning in S3 buckets for Amazon](#).

Beginning with 9.14.1, object locking is supported on S3 buckets. S3 Object Lock must be enabled when a bucket is created. Object Lock cannot be enabled on preexisting buckets. Object Lock can only be used in native S3 use cases. Multiprotocol NAS volumes configured to use the S3 protocol should use SnapLock to commit data to WORM storage. S3 object locking requires a standard SnapLock license. This license is included with [ONTAP One](#).

Prior to ONTAP One, the SnapLock license was included in the Security and Compliance bundle. The Security and Compliance bundle is no longer offered but is still valid. Although not currently required, existing customers can choose to [upgrade to ONTAP One](#). If you are enabling object locking on a bucket, you should [verify that a SnapLock license is installed](#). If a SnapLock license is not installed, you must [install](#) it before you can enable object locking.

When you have verified that the SnapLock license is installed, to protect objects in your bucket from getting deleted or overwritten, select **Enable object locking**. Locking can be enabled on either all or specific versions of objects, and only when the SnapLock compliance clock is initialized for the cluster nodes. Follow these steps:

1. If the SnapLock compliance clock is not initialized on any node of the cluster, the **Initialize SnapLock Compliance Clock** button appears. Click **Initialize SnapLock Compliance Clock** to initialize the SnapLock compliance clock on the cluster nodes.
2. Select **Governance** mode to activate a time-based lock that allows *Write once, read many (WORM)* permissions on the objects. Even in *Governance* mode, the objects can be deleted by administrator users with specific permissions.
3. Select **Compliance** mode if you want to assign stricter rules of deletion and update on the objects. In this mode of object locking, the objects can be expired only on the completion of the specified retention period. Unless a retention period is specified, the objects remain locked indefinitely.
4. Specify the retention tenure for the lock in days or years if you want the locking to be effective for a certain period.



Locking is applicable to versioned and non-versioned S3 buckets. Object locking is not applicable to NAS objects.

You can configure protection and permission settings, and performance service level for the bucket.



You must have already created user and groups before configuring the permissions.

For information, see [Create mirror for new bucket](#).

Verify access to the bucket

On S3 client applications (whether ONTAP S3 or an external third-party application), you can verify your access to the newly created bucket by entering the following:

- The S3 server CA certificate.
- The user's access key and secret key.

- The S3 server FQDN name and bucket name.

Increase or decrease the ONTAP S3 bucket size

When necessary, you can increase or decrease the size of an existing bucket.

Steps

You can use System Manager or the ONTAP CLI to manage the bucket size.

System Manager

1. Select **Storage > Buckets** and locate the bucket you want to modify.
2. Click  next to the bucket name and select **Edit**.
3. In the **Edit bucket** window, change the capacity for the bucket.
4. **Save**.

CLI

1. Change the bucket capacity:

```
vserver object-store-server bucket modify -vserver <SVM_name>
-bucket <bucket_name> -size {<integer>[KB|MB|GB|TB|PB]}
```

Create an ONTAP S3 bucket on a mirrored or unmirrored aggregate in a MetroCluster configuration

Beginning with ONTAP 9.14.1, you can provision a bucket on a mirrored or unmirrored aggregate in MetroCluster FC and IP configurations.

About this task

- By default, buckets are provisioned on mirrored aggregates.
- The same provisioning guidelines outlined in [Create a bucket](#) apply to creating a bucket in a MetroCluster environment.
- The following S3 object storage features are **not** supported in MetroCluster environments:
 - SnapMirror S3
 - S3 bucket lifecycle management
 - S3 object lock in **Compliance** mode



S3 object lock in **Governance** mode is supported.

- Local FabricPool tiering

Before you begin

An SVM containing an S3 server must already exist.

Process to create buckets

CLI

1. If you plan to select aggregates and FlexGroup components yourself, set the privilege level to advanced (otherwise, admin privilege level is sufficient): set `-privilege advanced`
2. Create a bucket:

```
vserver object-store-server bucket create -vserver <svm_name> -bucket <bucket_name> [-size integer[KB|MB|GB|TB|PB]] [-use-mirrored-aggregates true/false]
```

Set the `-use-mirrored-aggregates` option to `true` or `false` depending on whether you want to use a mirrored or unmirrored aggregate.



By default, the `-use-mirrored-aggregates` option is set to `true`.

- The SVM name must be a data SVM.
- If you specify no options, ONTAP creates an 800GB bucket with the service level set to the highest level available for your system.
- If you want ONTAP to create a bucket based on performance or usage, use one of the following options:
 - service level

Include the `-storage-service-level` option with one of the following values: `value`, `performance`, or `extreme`.

- tiering

Include the `-used-as-capacity-tier true` option.

- If you want to specify the aggregates on which to create the underlying FlexGroup volume, use the following options:
 - The `-aggr-list` parameter specifies the list of aggregates to be used for FlexGroup volume constituents.

Each entry in the list creates a constituent on the specified aggregate. You can specify an aggregate multiple times to have multiple constituents created on the aggregate.

For consistent performance across the FlexGroup volume, all of the aggregates must use the same disk type and RAID group configurations.

- The `-aggr-list-multiplier` parameter specifies the number of times to iterate over the aggregates that are listed with the `-aggr-list` parameter when creating a FlexGroup volume.

The default value of the `-aggr-list-multiplier` parameter is 4.

3. Add a QoS policy group if needed:

```
vserver object-store-server bucket modify -bucket <bucket_name> -qos-policy -group <qos_policy_group>
```

4. Verify bucket creation:

```
vserver object-store-server bucket show [-instance]
```

Example

The following example creates a bucket for SVM vs1 of size 1TB on a mirrored aggregate:

```
cluster-1::*> vserver object-store-server bucket create -vserver
svm1.example.com -bucket testbucket -size 1TB -use-mirrored-aggregates
true
```

System Manager

1. Add a new bucket on an S3-enabled storage VM.
 - a. Click **Storage > Buckets**, then click **Add**.
 - b. Enter a name, select the storage VM, and enter a size.

By default, the bucket is provisioned on a mirrored aggregate. If you want to create a bucket on an unmirrored aggregate, select **More Options** and uncheck the **Use the SyncMirror tier** box under **Protection** as shown in the following image:

Add bucket

⚠ To use this bucket from a remote cluster, configure S3 service on storage VM "vs1".

FOLDER (OPTIONAL)

Browse

Specify the folder to map to this bucket. [Know more](#)

CAUTION

Size
GB

Use for tiering
If you select this option, the system will try to select low-cost media with optimal performance for the tiered data.

Enable versioning
Versioning-enabled buckets allow you to recover objects that were accidentally deleted or overwritten. After versioning is enabled, it can't be disabled. However, you can suspend versioning.

PERFORMANCE SERVICE LEVEL

Value
▼

Not sure? [Get help selecting type](#)

Permissions

Copy access permissions from an existing bucket

Principal	Effect	Actions	Resources	Conditions
All users of this stor...	allow	ListBucket	-	-

[+ Add](#)

Object locking

Enable object locking
Object locking utilizes the "Write Once, Read Many" (WORM) model in which objects or their versions are protected from being deleted or overwritten during the specified retention period.

Protection

Use the SyncManager

[Save](#) [Cancel](#)

- If you click **Save** at this point, a bucket is created with these default settings:
 - No users are granted access to the bucket unless any group policies are already in effect.

i

You should not use the S3 root user to manage ONTAP object storage and share its permissions, because it has unlimited access to the object store. Instead, create a user or group with administrative privileges that you assign.

- A Quality of Service (performance) level that is the highest available for your system.
- You can click **More Options** to configure user permissions and performance level when you configure the bucket, or you can modify these settings later.
 - You must have already created user and groups before using **More Options** to configure their permissions.

- If you intend to use the S3 object store for FabricPool tiering, consider selecting **Use for tiering** (use low-cost media with optimal performance for the tiered data) rather than a performance service level.

2. On S3 client apps (another ONTAP system or an external 3rd-party app) verify access to the new bucket by entering the following:

- The S3 server CA certificate.
- The user's access key and secret key.
- The S3 server FQDN name and bucket name.

Create an ONTAP S3 bucket lifecycle management rule

Beginning with ONTAP 9.13.1, you can create lifecycle management rules to manage object lifecycles in your S3 buckets. You can define deletion rules for specific objects in a bucket, and through these rules, expire those bucket objects. This enables you to meet retention requirements and manage overall S3 object storage efficiently.

 If object locking is enabled for your bucket objects, the lifecycle management rules for object expiration will not be applied on locked objects. For information about object locking, see [Create a bucket](#).

Before you begin

- An S3-enabled SVM containing an S3 server and a bucket must already exist. See [Create an SVM for S3](#) for more information.
- Bucket lifecycle management rules are not supported when using S3 in multiprotocol NAS volumes, or when using S3 in MetroCluster configurations.

About this task

When creating your lifecycle management rules, you can apply the following deletion actions to your bucket objects:

- Deletion of current versions - This action expires objects identified by the rule. If versioning is enabled on the bucket, S3 makes all expired objects unavailable. If versioning is not enabled, this rule deletes the objects permanently. The CLI action is `Expiration`.
- Deletion of non-current versions - This action specifies when S3 can permanently remove non-current objects. The CLI action is `NoncurrentVersionExpiration`.



A non-current version is based on the current version's creation or modification time. Delayed removal of non-current objects can be helpful when you accidentally delete or overwrite an object. For example, you can configure an expiration rule to delete non-current versions five days after they become non-current. For example, suppose that on 1/1/2014 at 10:30 AM UTC, you create an object called `photo.gif` (version ID 111111). On 1/2/2014 at 11:30 AM UTC, you accidentally delete `photo.gif` (version ID 111111), which creates a delete marker with a new version ID (such as version ID 4857693). You now have five days to recover the original version of `photo.gif` (version ID 111111) before the deletion is permanent. On 1/8/2014 at 00:00 UTC, the Lifecycle rule for expiration runs and permanently deletes `photo.gif` (version ID 111111), five days after it became a non-current version.

- Deletion of expired delete markers - This action deletes expired object delete markers. In versioning-enabled buckets, objects with a delete markers become the current versions of the objects. The objects are not deleted, and no action can be performed on them. These objects become expired when there are no current versions associated with them. The CLI action is `Expiration`.
- Deletion of incomplete multipart uploads - This action sets a maximum time (in days) that you want to allow multipart uploads to remain in progress. Following which, they are deleted. The CLI action is `AbortIncompleteMultipartUpload`.

The procedure you follow depends on the interface that you use. With ONTAP 9.13.1, you need to use the CLI. Beginning with ONTAP 9.14.1, you can also use System Manager.

Manage lifecycle management rules with the CLI

Beginning with ONTAP 9.13.1, you can use the ONTAP CLI to create lifecycle management rules to expire objects in your S3 buckets.

Before you begin

For the CLI, you need to define the required fields for each expiration action type when creating a bucket lifecycle management rule. These fields can be modified after initial creation. The following table displays the unique fields for each action type.

Action type	Unique fields
NonCurrentVersionExpiration	<ul style="list-style-type: none"> • <code>-non-curr-days</code> - Number of days after which non-current versions will be deleted • <code>-new-non-curr-versions</code> - Number of latest non-current versions to be retained
Expiration	<ul style="list-style-type: none"> • <code>-obj-age-days</code> - Number of days since creation, after which current version of objects can be deleted • <code>-obj-exp-date</code> - Specific date when the objects should expire • <code>-expired-obj-del-markers</code> - Cleanup object delete markers
AbortIncompleteMultipartUpload	<ul style="list-style-type: none"> • <code>-after-initiation-days</code> - Number of days of initiation, after which upload can be aborted

In order for the bucket lifecycle management rule to only be applied to a specific subset of objects, admins must set each filter when creating the rule. If these filters are not set when creating the rule, the rule will be applied to all objects within the bucket.

All filters can be modified after initial creation *except* for the following:

- `-prefix`
- `-tags`
- `-obj-size-greater-than`
- `-obj-size-less-than`

Steps

1. Use the `vserver object-store-server bucket lifecycle-management-rule create` command with required fields for your expiration action type to create your bucket lifecycle management rule.

Example

The following command creates a NonCurrentVersionExpiration bucket lifecycle management rule:

```
vserver object-store-server bucket lifecycle-management-rule create
-vserver <svm_name> -bucket <bucket_name> -rule-id <rule_name> -action
NonCurrentVersionExpiration -index <lifecycle_rule_index_integer> -is
-enabled {true|false} -prefix <object_name> -tags <text> -obj-size-greater
-than {<integer>[KB|MB|GB|TB|PB]} -obj-size-less-than
{<integer>[KB|MB|GB|TB|PB]} -new-non-curr-versions <integer> -non-curr
-days <integer>
```

Example

The following command creates an Expiration bucket lifecycle management rule:

```
vserver object-store-server bucket lifecycle-management-rule create
-vserver <svm_name> -bucket <bucket_name> -rule-id <rule_name> -action
Expiration -index <lifecycle_rule_index_integer> -is-enabled {true|false}
-prefix <object_name> -tags <text> -obj-size-greater-than
{<integer>[KB|MB|GB|TB|PB]} -obj-size-less-than
{<integer>[KB|MB|GB|TB|PB]} -obj-age-days <integer> -obj-exp-date
<"MM/DD/YYYY HH:MM:SS"> -expired-obj-del-marker {true|false}
```

Example

The following command creates an AbortIncompleteMultipartUpload bucket lifecycle management rule:

```
vserver object-store-server bucket lifecycle-management-rule create
-vserver <svm_name> -bucket <bucket_name> -rule-id <rule_name> -action
AbortIncompleteMultipartUpload -index <lifecycle_rule_index_integer> -is
-enabled {true|false} -prefix <object_name> -tags <text> -obj-size-greater
-than {<integer>[KB|MB|GB|TB|PB]} -obj-size-less-than
{<integer>[KB|MB|GB|TB|PB]} -after-initiation-days <integer>
```

Manage lifecycle management rules with System Manager

Beginning with ONTAP 9.14.1, you can expire S3 objects by using System Manager. You can add, edit, and delete lifecycle management rules for your S3 objects. Additionally, you can import a lifecycle rule created for one bucket and use it for the objects in another bucket. You can disable an active rule and enable it later.

Add a lifecycle management rule

1. Click **Storage > Buckets**.

2. Select the bucket for which you want to specify the expiration rule.
3. Click the  icon and select **Manage lifecycle rules**.
4. Click **Add > Lifecycle rule**.
5. On the Add a lifecycle rule page, add the name of the rule.
6. Define the scope of the rule, whether you want it to apply to all the objects in the bucket or on specific objects. If you want to specify objects, add at least one of the following filter criteria:
 - a. **Prefix**: Specify a prefix of the object key names to which the rule should apply. Typically, it is the path or folder of the object. You can enter one prefix per rule. Unless a valid prefix is provided, the rule applies to all the objects in a bucket.
 - b. **Tags**: Specify up to three key and value pairs (tags) for the objects to which the rule should apply. Only valid keys are used for filtering. The value is optional. However, if you add values, ensure that you add only valid values for the corresponding keys.
 - c. **Size**: You can limit the scope between the minimum and maximum sizes of the objects. You can enter either or both the values. The default unit is MiB.
7. Specify the action:
 - a. **Expire the current version of objects**: Set a rule to make all current objects permanently unavailable after a specific number of days since their creation, or on a specific date. This option is unavailable if the **Delete expired object delete markers** option is selected.
 - b. **Permanently delete non-current versions**: Specify the number of days after which the non-current version is deleted, and the number of versions to retain.
 - c. **Delete expired object delete markers**: Select this action to delete objects with expired delete markers, that is delete markers without an associated current object.
-  This option becomes unavailable when you select the **Expire the current version of objects** option that automatically deletes all objects after the retention period. This option also becomes unavailable when object tags are used for filtering.
- d. **Delete incomplete multipart uploads**: Set the number of days after which incomplete multipart uploads are to be deleted. If the multipart uploads that are in progress fail within the specified retention period, you can delete the incomplete multipart uploads. This option becomes unavailable when object tags are used for filtering.
- e. Click **Save**.

Import a lifecycle rule

1. Click **Storage > Buckets**.
2. Select the bucket for which you want to import the expiration rule.
3. Click the  icon and select **Manage lifecycle rules**.
4. Click **Add > Import a rule**.
5. Select the bucket from which you want to import the rule. The lifecycle management rules defined for the selected bucket appear.
6. Select the rule that you want to import. You have the option to select one rule at a time, with the default selection being the first rule.
7. Click **Import**.

Edit, delete, or disable a rule

You can only edit the lifecycle management actions associated with the rule. If the rule was filtered with object tags, then the **Delete expired object delete markers** and **Delete incomplete multipart uploads** options are unavailable.

When you delete a rule, that rule will no longer apply to previously associated objects.

1. Click **Storage > Buckets**.
2. Select the bucket for which you want to edit, delete, or disable the lifecycle management rule.
3. Click the  icon and select **Manage lifecycle rules**.
4. Select the required rule. You can edit and disable one rule at a time. You can delete multiple rules at once.
5. Select **Edit**, **Delete**, or **Disable**, and complete the procedure.

Create an ONTAP S3 user

Create an S3 user with specific permissions. User authorization is required on all ONTAP object stores to restrict connectivity to authorized clients.

Before you begin.

An S3-enabled storage VM must already exist.

About this task

An S3 user can be granted access to any bucket in a storage VM. When you create an S3 user, an access key and a secret key are also generated for the user. They should be shared with the user along with the FQDN of the object store and bucket name.

For added security, beginning with ONTAP 9.15.1, access keys and secret keys are only displayed at the time the S3 user is created and cannot be displayed again. If the keys are lost, [new keys must be regenerated](#).

You can grant specific access permissions to S3 users in a bucket policy or an object server policy.



When you create a new object store server, ONTAP creates a root user (UID 0), which is a privileged user with access to all buckets. Rather than administering ONTAP S3 as the root user, NetApp recommends that an admin user role be created with specific privileges.

CLI

1. Create an S3 user:

```
vserver object-store-server user create -vserver svm_name -user user_name  
-comment [-comment text] -key-time-to-live time
```

- Adding a comment is optional.
- Beginning with ONTAP 9.14.1, you can define the period of time for which the key will be valid in the `-key-time-to-live` parameter. You can add the retention period in this format, to indicate the period after which the access key expires:

`P[<integer>D]T[<integer>H][<integer>M][<integer>S] | P<integer>W`

For example, if you want to enter a retention period of one day, two hours, three minutes, and four seconds, enter the value as `P1DT2H3M4S`. Unless specified, the key is valid for an indefinite period of time.

The below example creates a user with name `sm_user1` on storage VM `vs0`, with a key retention period of one week.

```
vserver object-store-server user create -vserver vs0 -user  
sm_user1 -key-time-to-live P1W
```

2. Be sure to save the access key and secret key. They will be required for access from S3 clients.

System Manager

1. Click **Storage > Storage VMs**. Select the storage VM to which you need to add a user, select **Settings** and then click  under S3.
2. To add a user, click **Users > Add**.
3. Enter a name for the user.
4. Beginning with ONTAP 9.14.1, you can specify the retention period of the access keys that get created for the user. You can specify the retention period in days, hours, minutes, or seconds, after which the keys automatically expire. By default, the value is set to 0 that indicates that the key is indefinitely valid.
5. Click **Save**. The user is created, and an access key and a secret key are generated for the user.
6. Download or save the access key and secret key. They will be required for access from S3 clients.

Next steps

- [Create or modify S3 groups](#)

Create or modify ONTAP S3 user groups to control access to buckets

You can simplify bucket access by creating groups of users with appropriate access authorizations.

Before you begin

S3 users in an S3-enabled SVM must already exist.

About this task

Users in an S3 group can be granted access to any bucket in an SVM but not in multiple SVMs. Group access permissions can be configured in two ways:

- At the bucket level

After creating a group of S3 users, you specify group permissions in bucket policy statements and they apply only to that bucket.

- At the SVM level

After creating a group of S3 users, you specify object server policy names in the group definition. Those policies determine the buckets and access for the group members.

System Manager

1. Edit the storage VM: click **Storage > storage VMs**, click the storage VM, click **Settings** and then click  under S3.
2. Add a group: select **Groups**, then select **Add**.
3. Enter a group name and select from a list of users.
4. You can select an existing group policy or add one now, or you can add a policy later.

CLI

1. Create an S3 group:

```
vserver object-store-server group create -vserver svm_name -name group_name -users user_name\(\s\) [-policies policy_names] [-comment text\]
```

The **-policies** option can be omitted in configurations with only one bucket in an object store; the group name can be added to the bucket policy.

The **-policies** option can be added later with the `vserver object-store-server group modify` command after object storage server policies are created.

Regenerate ONTAP S3 keys and modify their retention period

Access keys and secret keys are automatically generated during user creation for enabling S3 client access. You can regenerate keys for a user if a key is expired or compromised.

For information about generation of access keys, see [Create an S3 user](#).

System Manager

1. Click **Storage > Storage VMs** and then select the storage VM.
2. In the **Settings** tab, click  in the **S3** tile.
3. In the **Users** tab, verify that there is no access key, or the key has expired for the user.
4. If you need to regenerate the key, click  next to the user, then click **Regenerate Key**.
5. By default, generated keys are valid for an indefinite amount of time. Beginning with 9.14.1, you can modify their retention period, after which the keys automatically expire. Enter the retention period in days, hours, minutes, or seconds.
6. Click **Save**. The key is regenerated. Any change in the key retention period takes effect immediately.
7. Download or save the access key and secret key. They will be required for access from S3 clients.

CLI

1. Regenerate access and secret keys for a user by running the `vserver object-store-server user regenerate-keys` command.
2. By default, generated keys are valid indefinitely. Beginning with 9.14.1, you can modify their retention period, after which the keys automatically expire. You can add the retention period in this format: `P[<integer>D]T[<integer>H][<integer>M][<integer>S] | P<integer>W`
For example, if you want to enter a retention period of one day, two hours, three minutes, and four seconds, enter the value as `P1DT2H3M4S`.

```
vserver object-store-server user regenerate-keys -vserver svm_name
-user user -key-time-to-live 0
```

3. Save the access and secret keys. They will be required for access from S3 clients.

Copyright information

Copyright © 2026 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—with prior written permission of the copyright owner.

Software derived from copyrighted NetApp material is subject to the following license and disclaimer:

THIS SOFTWARE IS PROVIDED BY NETAPP “AS IS” AND WITHOUT ANY EXPRESS OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE, WHICH ARE HEREBY DISCLAIMED. IN NO EVENT SHALL NETAPP BE LIABLE FOR ANY DIRECT, INDIRECT, INCIDENTAL, SPECIAL, EXEMPLARY, OR CONSEQUENTIAL DAMAGES (INCLUDING, BUT NOT LIMITED TO, PROCUREMENT OF SUBSTITUTE GOODS OR SERVICES; LOSS OF USE, DATA, OR PROFITS; OR BUSINESS INTERRUPTION) HOWEVER CAUSED AND ON ANY THEORY OF LIABILITY, WHETHER IN CONTRACT, STRICT LIABILITY, OR TORT (INCLUDING NEGLIGENCE OR OTHERWISE) ARISING IN ANY WAY OUT OF THE USE OF THIS SOFTWARE, EVEN IF ADVISED OF THE POSSIBILITY OF SUCH DAMAGE.

NetApp reserves the right to change any products described herein at any time, and without notice. NetApp assumes no responsibility or liability arising from the use of products described herein, except as expressly agreed to in writing by NetApp. The use or purchase of this product does not convey a license under any patent rights, trademark rights, or any other intellectual property rights of NetApp.

The product described in this manual may be protected by one or more U.S. patents, foreign patents, or pending applications.

LIMITED RIGHTS LEGEND: Use, duplication, or disclosure by the government is subject to restrictions as set forth in subparagraph (b)(3) of the Rights in Technical Data -Noncommercial Items at DFARS 252.227-7013 (FEB 2014) and FAR 52.227-19 (DEC 2007).

Data contained herein pertains to a commercial product and/or commercial service (as defined in FAR 2.101) and is proprietary to NetApp, Inc. All NetApp technical data and computer software provided under this Agreement is commercial in nature and developed solely at private expense. The U.S. Government has a non-exclusive, non-transferrable, nonsublicensable, worldwide, limited irrevocable license to use the Data only in connection with and in support of the U.S. Government contract under which the Data was delivered. Except as provided herein, the Data may not be used, disclosed, reproduced, modified, performed, or displayed without the prior written approval of NetApp, Inc. United States Government license rights for the Department of Defense are limited to those rights identified in DFARS clause 252.227-7015(b) (FEB 2014).

Trademark information

NETAPP, the NETAPP logo, and the marks listed at <http://www.netapp.com/TM> are trademarks of NetApp, Inc. Other company and product names may be trademarks of their respective owners.