



Add storage capacity to an S3-enabled SVM

ONTAP 9

NetApp
March 08, 2024

Table of Contents

- Add storage capacity to an S3-enabled SVM 1
 - Create a bucket 1
 - Create a bucket on a mirrored or unmirrored aggregate in a MetroCluster configuration 4
 - Create a bucket lifecycle management rule 9
 - Create an S3 user 13
 - Create or modify S3 groups 14
 - Regenerate keys and modify their retention period 15
 - Create or modify access policy statements 16
 - Enable client access to S3 object storage 26

Add storage capacity to an S3-enabled SVM

Create a 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.

See [vserver object-store-server bucket modify](#) for more information.

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 specify no options, ONTAP creates a 5GB 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 storage VM `vs1` of size 1TB and specifying the aggregate:

```
cluster-1::*> vserver object-store-server bucket create -vserver  
svml.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 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. If you want 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.

Create a 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:
 - S3 SnapMirror
 - 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 a 5GB 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  
svml.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 ✕

NAME

 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](#)

CAPACITY

Size

GB

▼

☐ Use 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 S3x3l0n0r4t10n

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.



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 a 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.

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`.
- 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 utilize 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 noncurrent versions:** Specify the number of days after which the version becomes non-current, and thereafter can be 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 S3 user

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. An S3 users' keys can be viewed with the `vserver object-store-server user show` command.

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 S3 groups

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 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](#).



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.

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.

Create or modify access policy statements

About bucket and object store server policies

User and group access to S3 resources is controlled by bucket and object store server policies. If you have a small number of users or groups, controlling access at the bucket level is probably sufficient, but if you have many users and groups, it is easier to control access at the object store server level.

Modify a bucket policy

You can add access rules to the default bucket policy. The scope of its access control is the containing bucket, so it is most appropriate when there is a single bucket.

Before you begin

An S3-enabled storage VM containing an S3 server and a bucket must already exist.

You must have already created users or groups before granting permissions.

About this task

You can add new statements for new users and groups, or you can modify the attributes of existing statements. For more options, see the `vserver object-store-server bucket policy man` pages.

User and group permissions can be granted when the bucket is created or as needed later. You can also modify the bucket capacity and QoS policy group assignment.

Beginning with ONTAP 9.9.1, if you plan to support AWS client object tagging functionality with the ONTAP S3 server, the actions `GetObjectTagging`, `PutObjectTagging`, and `DeleteObjectTagging` need to be allowed using the bucket or group policies.

The procedure you follow depends on the interface that you use—System Manager or the CLI:

System Manager

Steps

1. Edit the bucket: click **Storage > Buckets**, click the desired bucket, and then click **Edit**.

When adding or modifying permissions, you can specify the following parameters:

- **Principal**: the user or group to whom access is granted.
- **Effect**: allows or denies access to a user or group.
- **Actions**: permissible actions in the bucket for a given user or group.
- **Resources**: paths and names of objects within the bucket for which access is granted or denied.

The defaults **bucketname** and **bucketname/*** grant access to all objects in the bucket. You can also grant access to single objects; for example, **bucketname/*_readme.txt**.

- **Conditions** (optional): expressions that are evaluated when access is attempted. For example, you can specify a list of IP addresses for which access will be allowed or denied.



Beginning with ONTAP 9.14.1, you can specify variables for the bucket policy in the **Resources** field. These variables are placeholders that are replaced with contextual values when the policy is evaluated. For example, If `${aws:username}` is specified as a variable for a policy, then this variable is replaced with the request context username, and the policy action can be performed as configured for that user.

CLI

Steps

1. Add a statement to a bucket policy:

```
vserver object-store-server bucket policy add-statement -vserver svm_name
-bucket bucket_name -effect {allow|deny} -action object_store_actions
-principal user_and_group_names -resource object_store_resources [-sid
text] [-index integer]
```

The following parameters define access permissions:

-effect	The statement may allow or deny access
-action	You can specify * to mean all actions, or a list of one or more of the following: <code>GetObject</code> , <code>PutObject</code> , <code>DeleteObject</code> , <code>ListBucket</code> , <code>GetBucketAcl</code> , <code>GetObjectAcl</code> , <code>ListBucketMultipartUploads</code> , and <code>ListMultipartUploadParts</code> .

-principal	<p>A list of one or more S3 users or groups.</p> <ul style="list-style-type: none"> • A maximum of 10 users or groups can be specified. • If an S3 group is specified, it must be in the form <code>group/group_name</code>. • * can be specified to mean public access; that is, access without an access-key and secret-key. • If no principal is specified, all S3 users in the storage VM are granted access.
-resource	<p>The bucket and any object it contains. The wildcard characters * and ? can be used to form a regular expression for specifying a resource. For a resource, you can specify variables in a policy. These are policy variables are placeholders that are replaced with the contextual values when the policy is evaluated.</p>

You can optionally specify a text string as comment with the `-sid` option.

Examples

The following example creates an object store server bucket policy statement for the storage VM `svm1.example.com` and `bucket1` which specifies allowed access to a `readme` folder for object store server user `user1`.

```
cluster1::> vservers object-store-server bucket policy statement create
-vserver svm1.example.com -bucket bucket1 -effect allow -action
GetObject,PutObject,DeleteObject,ListBucket -principal user1 -resource
bucket1/readme/* -sid "fullAccessToReadmeForUser1"
```

The following example creates an object store server bucket policy statement for the storage VM `svm1.example.com` and `bucket1` which specifies allowed access to all objects for object store server group `group1`.

```
cluster1::> vservers object-store-server bucket policy statement create
-vserver svm1.example.com -bucket bucket1 -effect allow -action
GetObject,PutObject,DeleteObject,ListBucket -principal group/group1
-resource bucket1/* -sid "fullAccessForGroup1"
```

Beginning with ONTAP 9.14.1, you can specify variables for a bucket policy. The following example creates a server bucket policy statement for the storage VM `svm1` and `bucket1`, and specifies `${aws:username}` as a variable for a policy resource. When the policy is evaluated, the policy variable is replaced with the request context username, and the policy action can be performed as configured for that user. For example, when the following policy statement is evaluated, `${aws:username}` is replaced with the user performing the S3 operation. If a user `user1` performs the operation, that user is granted access to `bucket1` as `bucket1/user1/*`.

```
cluster1::> object-store-server bucket policy statement create -vserver  
svml -bucket bucket1 -effect allow -action * -principal - -resource  
bucket1,bucket1/${aws:username}/*##
```

Create or modify an object store server policy

You can create policies that can apply to one or more buckets in an object store. Object store server policies can be attached to groups of users, thereby simplifying the management of resource access across multiple buckets.

Before you begin

An S3-enabled SVM containing an S3 server and a bucket must already exist.

About this task

You can enable access policies at the SVM level by specifying a default or custom policy in an object storage server group. The policies do not take effect until they are specified in the group definition.



When you use object storage server policies, you specify principals (that is, users and groups) in the group definition, not in the policy itself.

There are three read-only default policies for access to ONTAP S3 resources:

- FullAccess
- NoS3Access
- ReadOnlyAccess

You can also create new custom policies, then add new statements for new users and groups, or you can modify the attributes of existing statements. For more options, see the `vserver object-store-server policy` [command reference](#).


Beginning with ONTAP 9.9.1, if you plan to support AWS client object tagging functionality with the ONTAP S3 server, the actions `GetObjectTagging`, `PutObjectTagging`, and `DeleteObjectTagging` need to be allowed using the bucket or group policies.

The procedure you follow depends on the interface that you use—System Manager or the CLI:

System Manager

Use System Manager to create or modify an object store server policy

Steps

1. Edit the storage VM: click **Storage > storage VMs**, click the storage VM, click **Settings** and then click  under S3.
2. Add a user: click **Policies**, then click **Add**.
 - a. Enter a policy name and select from a list of groups.
 - b. Select an existing default policy or add a new one.

When adding or modifying a group policy, you can specify the following parameters:

- Group: the groups to whom access is granted.
- Effect: allows or denies access to one or more groups.
- Actions: permissible actions in one or more buckets for a given group.
- Resources: paths and names of objects within one or more buckets for which access is granted or denied.

For example:

- * grants access to all buckets in the storage VM.
 - **bucketname** and **bucketname/*** grant access to all objects in a specific bucket.
 - **bucketname/readme.txt** grants access to an object in a specific bucket.
- c. If desired, add statements to existing policies.

CLI

Use the CLI to create or modify an object store server policy

Steps

1. Create an object storage server policy:

```
vserver object-store-server policy create -vserver svm_name -policy policy_name [-comment text]
```

2. Create a statement for the policy:

```
vserver object-store-server policy statement create -vserver svm_name -policy policy_name -effect {allow|deny} -action object_store_actions -resource object_store_resources [-sid text]
```

The following parameters define access permissions:

-effect	The statement may allow or deny access
---------	--

<code>-action</code>	You can specify <code>*</code> to mean all actions, or a list of one or more of the following: <code>GetObject</code> , <code>PutObject</code> , <code>DeleteObject</code> , <code>ListBucket</code> , <code>GetBucketAcl</code> , <code>GetObjectAcl</code> , <code>ListAllMyBuckets</code> , <code>ListBucketMultipartUploads</code> , and <code>ListMultipartUploadParts</code> .
<code>-resource</code>	The bucket and any object it contains. The wildcard characters <code>*</code> and <code>?</code> can be used to form a regular expression for specifying a resource.

You can optionally specify a text string as comment with the `-sid` option.

By default, new statements are added to the end of the list of statements, which are processed in order. When you add or modify statements later, you have the option to modify the statement's `-index` setting to change the processing order.

Configure S3 access for external directory services

Beginning with ONTAP 9.14.1, services for external directories have been integrated with ONTAP S3 object storage. This integration simplifies user and access management through external directory services.

You can provide user groups belonging to an external directory service with access to your ONTAP object storage environment. Lightweight Directory Access Protocol (LDAP) is an interface for communicating with directory services, such as Active Directory, that provide a database and services for identity and access management (IAM). To provide access, you need to configure LDAP groups in your ONTAP S3 environment. After you have configured access, the group members have permissions to ONTAP S3 buckets. For information about LDAP, see [Overview of using LDAP](#).

You can also configure Active Directory user groups for fast bind mode, so that user credentials can be validated and third-party and open-source S3 applications can be authenticated over LDAP connections.

Before you begin

Ensure the following before configuring LDAP groups and enabling the fast bind mode for group access:

1. An S3-enabled storage VM containing an S3 server has been created. See [Create an SVM for S3](#).
2. A bucket has been created in that storage VM. See [Create a bucket](#).
3. DNS is configured on the storage VM. See [Configure DNS services](#).
4. A self-signed root certification authority (CA) certificate of the LDAP server is installed on the storage VM. See [Install the self-signed root CA certificate on the SVM](#).
5. An LDAP client is configured with TLS enabled on the SVM. See [Create an LDAP client configuration](#) and [Associate the LDAP client configuration with SVMs for information](#).

Configure S3 access for external directory services

1. Specify LDAP as the *name service database* of the SVM for the group and password to LDAP:

```
ns-switch modify -vserver <vserver-name> -database group -sources  
files,ldap  
ns-switch modify -vserver <vserver-name> -database passwd -sources  
files,ldap
```

For more information about this command, see the [vserver services name-service ns-switch modify](#) command.

2. Create an object store bucket policy statement with the `principal` set to the LDAP group to which you want to grant access:

```
object-store-server bucket policy statement create -bucket <bucket-name>  
-effect allow -principal nasgroup/<ldap-group-name> -resource <bucket-  
name>, <bucket-name>/*
```

Example: The following example creates a bucket policy statement for `buck1`. The policy allows access for the LDAP group `group1` to the resource (bucket and its objects) `buck1`.

```
vserver object-store-server bucket policy add-statement -bucket buck1  
-effect allow -action  
GetObject,PutObject,DeleteObject,ListBucket,GetBucketAcl,GetObjectAcl,Li  
stBucketMultipartUploads,ListMultipartUploadParts,  
ListBucketVersions,GetObjectTagging,PutObjectTagging,DeleteObjectTagging  
,GetBucketVersioning,PutBucketVersioning -principal nasgroup/group1  
-resource buck1, buck1/*
```

3. Verify that a user from the LDAP group `group1` is able to perform S3 operations from the S3 client.

Use LDAP fast bind mode for authentication

1. Specify LDAP as the *name service database* of the SVM for the group and password to LDAP:

```
ns-switch modify -vserver <vserver-name> -database group -sources  
files,ldap  
ns-switch modify -vserver <vserver-name> -database passwd -sources  
files,ldap
```

For more information about this command, see the [vserver services name-service ns-switch modify](#) command.

2. Ensure that an LDAP user accessing the S3 bucket has permissions defined in the bucket-policies. For

more information, see [Modify a bucket policy](#).

3. Verify that a user from the LDAP group can perform the following operations:
 - a. Configure the access key on the S3 client in this format:
"NTAPFASTBIND" + base64-encode(user-name:password)
Example: "NTAPFASTBIND" + base64-encode(ldapuser:password), which results in
NTAPFASTBINDbGRhcHVzZXI6cGFzc3dvcmQ=



The S3 client might prompt for a secret key. In the absence of a secret key, any password of at least 16 characters can be entered.

- b. Perform basic S3 operations from the S3 client for which the user has permissions.

Enable LDAP or domain users to generate their own S3 access keys

Beginning with ONTAP 9.14.1, as an ONTAP administrator, you can create custom roles and grant them to local or domain groups or Lightweight Directory Access Protocol (LDAP) groups, so that the users belonging to those groups can generate their own access and secret keys for S3 client access.

You have to perform a few configuration steps on your storage VM, so that the custom role can be created and assigned to the user that invokes the API for access key generation.

Before you begin

Ensure the following:

1. An S3-enabled storage VM containing an S3 server has been created. See [Create an SVM for S3](#).
2. A bucket has been created in that storage VM. See [Create a bucket](#).
3. DNS is configured on the storage VM. See [Configure DNS services](#).
4. A self-signed root certification authority (CA) certificate of the LDAP server is installed on the storage VM. See [Install the self-signed root CA certificate on the SVM](#).
5. An LDAP client is configured with TLS enabled on the storage VM. See [Create an LDAP client configuration](#) and .
6. Associate the client configuration with the Vserver. See [Associate the LDAP client configuration with SVMs and vservice services name-service ldap create](#).
7. If you are using a data storage VM, create a management network interface (LIF) and on the VM, and also a service policy for the LIF. See the [network interface create](#) and [network interface service-policy create](#) commands.

Configure users for access key generation

1. Specify LDAP as the *name service database* of the storage VM for the group and password to LDAP:

```
ns-switch modify -vserver <vserver-name> -database group -sources  
files,ldap  
ns-switch modify -vserver <vserver-name> -database passwd -sources  
files,ldap
```

For more information about this command, see the [vserver services name-service ns-switch modify](#) command.

2. Create a custom role with access to S3 user REST API endpoint:

```
security login rest-role create -vserver <vserver-name> -role <custom-role-name> -api "/api/protocols/s3/services/*/users" -access <access-type>
```

In this example, the `s3-role` role is generated for users on the storage VM `svm-1`, to which all access rights, read, create, and update are granted.

```
security login rest-role create -vserver svm-1 -role s3role -api  
"/api/protocols/s3/services/*/users" -access all
```

For more information about this command, see the [security login rest-role create](#) command.

3. Create an LDAP user group with the security login command and add the new custom role for accessing the S3 user REST API endpoint. For more information about this command, see the [security login create](#) command.

```
security login create -user-or-group-name <ldap-group-name> -application  
http -authentication-method nsswitch -role <custom-role-name> -is-ns  
-switch-group yes
```

In this example, the LDAP group `ldap-group-1` is created in `svm-1`, and the custom role `s3role` is added to it for accessing the API endpoint, along with enabling LDAP access in the fast bind mode.

```
security login create -user-or-group-name ldap-group-1 -application http  
-authentication-method nsswitch -role s3role -is-ns-switch-group yes  
-second-authentication-method none -vserver svm-1 -is-ldap-fastbind yes
```

For more information, see [Use LDAP fast bind for nsswitch authentication](#).

Adding the custom role to the domain or LDAP group allows users in that group a limited access to the ONTAP `/api/protocols/s3/services/{svm.uid}/users` endpoint. By invoking the API, the domain or LDAP group users can generate their own access and secret keys to access the S3 client. They can generate the keys for only themselves and not for other users.

As an S3 or LDAP user, generate your own access keys

Beginning with ONTAP 9.14.1, you can generate your own access and secret keys for accessing S3 clients, if your administrator has granted you the role to generate your own keys. You can generate keys for only yourself by using the following ONTAP REST API endpoint.

HTTP method and endpoint

This REST API call uses the following method and endpoint. For information about the other methods of this endpoint, see the reference [API documentation](#).

HTTP method	Path
POST	/api/protocols/s3/services/{svm.uuid}/users

Curl example

```
curl
--request POST \
--location "https://$FQDN_IP /api/protocols/s3/services/{svm.uuid}/users "
\
--include \
--header "Accept: */*" \
--header "Authorization: Basic $BASIC_AUTH"
--data '{"name": "_name_"}'
```

JSON output example

```
{
  "records": [
    {
      "access_key":
"Pz3SB54G2B_6dsXQPrA5HrTPcf478qoAW6_Xx6qyqZ948AgZ_7YfCf_9nO87YoZmskxx3cq41
U2JAH2M3_fs321B4rkzS3a_oC5_8u7D8j_45N8OsBCBPWGD_1d_ccfq",
      "_links": {
        "next": {
          "href": "/api/resourcelink"
        },
        "self": {
          "href": "/api/resourcelink"
        }
      },
      "name": "user-1",
      "secret_key":
"A20_tDhC_cux2C2BmtL45bXB_a_Q65c_96FsAcOdo14Az8V31jBKDTc0uCL62Bh559gPB8s9r
rn0868QrF38_1dsV2u1_9H2tSf3qQ5xp9NT259C6z_GiZQ883Qn63X1"
    }
  ],
  "num_records": "1"
}
```

Enable client access to S3 object storage

Enable ONTAP S3 access for remote FabricPool tiering

For ONTAP S3 to be used as a remote FabricPool capacity (cloud) tier, the ONTAP S3 administrator must provide information about the S3 server configuration to the remote

ONTAP cluster administrator.

About this task

The following S3 server information is required to configure FabricPool cloud tiers:

- server name (FQDN)
- bucket name
- CA certificate
- access key
- password (secret access key)

In addition, the following networking configuration is required:

- There must be an entry for the remote ONTAP S3 server's hostname in the DNS server configured for the admin SVM, including the S3 server's FQDN name and the IP addresses on its LIFs.
- Intercluster LIFs must be configured on the local cluster, although cluster peering is not required.

See the FabricPool documentation about configuring ONTAP S3 as a cloud tier.

[Managing Storage Tiers By Using FabricPool](#)

Enable ONTAP S3 access for local FabricPool tiering

For ONTAP S3 to be used as a local FabricPool capacity tier, you must define an object store based on the bucket you created, and then attach the object store to a performance tier aggregate to create a FabricPool.

Before you begin

You must have the ONTAP S3 server name and a bucket name, and the S3 server must have been created using cluster LIFs (with the `-vserver Cluster` parameter).

About this task

The object-store configuration contains information about the local capacity tier, including the S3 server and bucket names and authentication requirements.

An object-store configuration once created must not be reassociated with a different object-store or bucket. You can create multiple buckets for local tiers, but you cannot create multiple object stores in a single bucket.

A FabricPool license is not required for a local capacity tier.

Steps

1. Create the object store for the local capacity tier:

```
storage aggregate object-store config create -object-store-name store_name
-ipSPACE Cluster -provider-type ONTAP_S3 -server S3_server_name -container
-name bucket_name -access-key access_key -secret-password password
```

- The `-container-name` is the S3 bucket you created.
- The `-access-key` parameter authorizes requests to the ONTAP S3 server.

- The `-secret-password` parameter (secret access key) authenticates requests to the ONTAP S3 server.
- You can set the `-is-certificate-validation-enabled` parameter to `false` to disable certificate checking for ONTAP S3.

```
cluster1::> storage aggregate object-store config create
-object-store-name MyLocalObjStore -ipspace Cluster -provider-type
ONTAP_S3 -server s3.example.com
-container-name bucket1 -access-key myS3key -secret-password myS3pass
```

2. Display and verify the object store configuration information:

```
storage aggregate object-store config show
```

3. Optional: To see how much data in a volume is inactive, follow the steps in [Determining how much data in a volume is inactive by using inactive data reporting](#).

Seeing how much data in a volume is inactive can help you decide which aggregate to use for FabricPool local tiering.

4. Attach the object store to an aggregate:

```
storage aggregate object-store attach -aggregate aggr_name -object-store-name
store_name
```

You can use the `allow-flexgroup` **true** option to attach aggregates that contain FlexGroup volume constituents.

```
cluster1::> storage aggregate object-store attach
-aggregate aggr1 -object-store-name MyLocalObjStore
```

5. Display the object store information and verify that the attached object store is available:

```
storage aggregate object-store show
```

```
cluster1::> storage aggregate object-store show
```

Aggregate	Object Store Name	Availability State
-----	-----	-----
aggr1	MyLocalObjStore	available

Enable client access from an S3 app

For S3 client apps to access the ONTAP S3 server, the ONTAP S3 administrator must provide configuration information to the S3 user.

Before you begin

The S3 client app must be capable of authenticating with the ONTAP S3 server using the following AWS signature versions:

- Signature Version 4, ONTAP 9.8 and later
- Signature Version 2, ONTAP 9.11.1 and later

Other signature versions are not supported by ONTAP S3.

The ONTAP S3 administrator must have created S3 users and granted them access permissions, as an individual users or as a group member, in the bucket policy or the object storage server policy.

The S3 client app must be capable of resolving the ONTAP S3 server name, which requires that ONTAP S3 administrator provide the S3 server name (FQDN) and IP addresses for the S3 server's LIFs.

About this task

To access an ONTAP S3 bucket, a user on the S3 client app enters information provided by the ONTAP S3 administrator.

Beginning with ONTAP 9.9.1, the ONTAP S3 server supports the following AWS client functionality:

- user-defined object metadata

A set of key-value pairs can be assigned to objects as metadata when they are created using PUT (or POST). When a GET/HEAD operation is performed on the object, the user-defined metadata is returned along with the system metadata.

- object tagging

A separate set of key-value pairs can be assigned as tags for categorizing objects. Unlike metadata, tags are created and read with REST APIs independently of the object, and they implemented when objects are created or any time after.



To enable clients to get and put tagging information, the actions `GetObjectTagging`, `PutObjectTagging`, and `DeleteObjectTagging` need to be allowed using the bucket or group policies.

For more information, see the AWS S3 documentation.

Steps

1. Authenticate the S3 client app with the ONTAP S3 server by entering the S3 server name and the CA certificate.
2. Authenticate a user on the S3 client app by entering the following information:
 - S3 server name (FQDN) and bucket name
 - the user's access key and secret key

Copyright information

Copyright © 2024 NetApp, Inc. All Rights Reserved. Printed in the U.S. No part of this document covered by copyright may be reproduced in any form or by any means—graphic, electronic, or mechanical, including photocopying, recording, taping, or storage in an electronic retrieval system—without prior written permission of the copyright owner.

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

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

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

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

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

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

Trademark information

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