Protect S3 data with snapshots

ONTAP 9

NetApp
December 21, 2024

Table of Contents

Protect S3 data with snapshots

S3 snapshot overview

Beginning with ONTAP 9.16.1, you can use ONTAP snapshot technology to generate read-only, point-in-time images of your ONTAP S3 buckets.

Using the S3 snapshots feature, you can manually create snapshots or automatically generate them through snapshot policies. S3 snapshots are presented as S3 buckets to S3 clients. You can browse and restore the content from the snapshots through S3 clients.

In ONTAP 9.16.1, S3 snapshots only capture the current versions of the objects in S3 buckets. The non-current versions of versioned buckets are not captured in the S3 snapshots. Also, the point-in-time object tags are not captured in the snapshots if the object tags are modified after the snapshots are taken.

S3 snapshots rely on the cluster time. You should configure the NTP server in your cluster to synchronize the time. For more information, refer to Manage the cluster time.

Quota and space usage

Quotas track the number of objects and logical size used in an S3 bucket. When S3 snapshots are created, the objects captured in the S3 snapshots are counted towards the bucket object count and size used, until the snapshots are deleted from the file system.

Multipart objects

For multipart objects, only the final objects are captured in snapshots. Partial uploads of multipart objects are not captured in snapshots.

Snapshots on versioned and non-versioned buckets

You can create snapshots on both versioned and non-versioned buckets. The snapshot contains only the current object versions at a time when the snapshot is captured.

Versioned buckets and snapshots

In buckets with object versioning enabled, a snapshot retains the content of the most recent object version after which the snapshot was taken. It excludes non-current versions in the bucket.

Consider this example: In a bucket where object versioning is enabled, object `obj1` has versions v1, v2, v3, v4, v5. You created a snapshot `snap1` from `obj1` v3 (the most recent version at the point of capture). When browsing `snap1`, `obj1` will appear as an object with content created at v3. Content of the previous versions will not be returned.

The non-current versions are retained in the filesystem, until the snapshots are deleted.

Non-versioned buckets and snapshots

In non-versioned buckets, S3 snapshots preserve the content of the latest commits prior to the snapshot creation.

Consider this example: In a bucket where object versioning is unavailable, object `obj1` has been overwritten several times at (t1, t2, t3, t4, and t5). You created an S3 snapshot `snap1` sometime between t3 and t4. When browsing `snap1`, `obj1` will appear with the content created at t3.

Object expiration and snapshots

ONTAP S3 object expiration and S3 snapshots feature function independently of each other. ONTAP object expiration feature expires object versions according to the lifecycle management rules defined for the S3 bucket. S3 snapshots are static copies of the bucket objects at a point in time when the snapshot is created.

If object versioning is enabled in a bucket, when a specific version of an object is deleted due to an expiration rule defined for that bucket, the content of the expired object version continues to remain in the filesystem if the version has been captured as a current version in one or more S3 snapshots. That object version will cease to exist in the file system only when that snapshot is deleted.

Similarly, in a bucket in which versioning is disabled, if an object is deleted based on an expiration rule, but the object is still captured in some existing S3 snapshots, the object will be retained in the file system. The object will be permanently removed from the file system when the snapshots capturing it are deleted.

For information about S3 object expiration and lifecycle management, refer to Create a bucket lifecycle management rule.

Limitations with S3 snapshots

Note the following feature exclusions and scenarios in ONTAP 9.16.1:

- You can generate up to 1023 snapshots for an S3 bucket.
- It is necessary to delete all the S3 snapshots and metadata from all the buckets in a cluster before reverting the cluster to an ONTAP version earlier than ONTAP 9.16.1.
- If you need to delete an S3 bucket containing objects with snapshots, ensure that you have deleted all the corresponding snapshots of all the objects in that bucket.
- S3 snapshots are not supported in these configurations:
  - On buckets in a SnapMirror relationship
  - On buckets where object-locking is enabled
  - On NetApp BlueXP
  - On System Manager
  - In ONTAP MetroCluster configurations

Create S3 snapshots

You can either manually generate S3 snapshots or set up snapshot policies to automatically create S3 snapshots for you. Snapshots serve as static copies of objects that you use for data backup and recovery. For determining the tenure of snapshot retention, you can create snapshot policies that facilitate automatic snapshot creation at specified intervals.

S3 snapshots help you protect your object data in S3 buckets with or without object versioning enabled.

Snapshots can be especially useful in establishing data protection when object versioning is not enabled in an S3 bucket, because they act as point-in-time records that you can use for restore operations when a previous object version is not available.

About this task

- The following naming rules apply to snapshot (for both manual and automatic snapshots):
  - S3 snapshot names can be up to 30 characters
  - S3 snapshot names can consist only of lowercase letters, numbers, dots (.), and hyphens (-)
  - S3 snapshot names must end with a letter or number
  - S3 snapshot names cannot contain substring `s3snap`
- In the context of the S3 protocol, the bucket naming restrictions limit a bucket name to 63 characters. Because ONTAP S3 snapshots are presented as buckets through the S3 protocol, similar restrictions apply to the snapshot bucket names. By default, the original bucket name is used as the base bucket name.
- To make it easier to identify which snapshot belongs to which bucket, the snapshot bucket name consists of the base bucket name, along with a special string, `-s3snap-`, that's prefixed to the snapshot name. The snapshot bucket names are formatted as `<base_bucket_name>-s3snap-<snapshot_name>`.

  For example, running the following command to create `snap1` on `bucket-a` creates a snapshot bucket with name `bucket-a-s3snap-snap1`, which is accessible to you through S3 clients if you have permissions to access the base bucket.

  ```
  vserver object-store-server bucket snapshot create -bucket bucket-a
  -snapshot snap1
  ```

- You cannot create a snapshot that results in a snapshot bucket name with more than 63 characters.
- The automatic snapshot name contains the policy schedule name and the timestamp, which is similar to the naming convention for the traditional volume snapshots. For example, the scheduled snapshot names can be `daily-2024-01-01-0015` and `hourly-2024-05-22-1105`.

Manually create S3 snapshots

You can manually create an S3 snapshot by using the ONTAP CLI. The procedure creates a snapshot on the local cluster only.

Steps

1. Create an S3 snapshot:

   ```
   vserver object-store-server bucket snapshot create -vserver <svm_name>
   -bucket <bucket_name> -snapshot <snapshot_name>
   ```

   The following example creates a snapshot named `pre-update` on the `vs0` storage VM and `website-data` bucket:

   ```
   vserver object-store-server bucket snapshot create -vserver vs0 -bucket
   website-data -snapshot pre-update
   ```

Assign an S3 snapshot policy to a bucket

When you configure snapshot policies at the S3 bucket level, ONTAP creates scheduled S3 snapshots for you automatically. Like traditional snapshot policies, up to five schedules can be configured for S3 snapshots.

A snapshot policy typically specifies the schedules to create snapshots, the number of copies to retain for each schedule, and the schedule prefix. For example, a policy can create one S3 snapshot every day at 12:10 AM, retain the two most recent copies, and name them `daily-<timestamp>`.

The default snapshot policy preserves:

- Six hourly snapshots
- Two daily snapshots
- Two weekly snapshots

Before you begin

- A snapshot policy must have been created before assigning it to the S3 bucket.

> Policies for S3 snapshots follow the same rules as other ONTAP snapshot policies. However, a snapshot policy with a retention period configured in any of the snapshot schedules cannot be assigned to an S3 bucket.

For more information about creating snapshot policies for autogenerating snapshots, refer to Configure custom snapshot policies overview.

Steps

1. Assign the snapshot policy on your bucket:

```
vserver object-store-server bucket create -vserver <svm_name> -bucket
<bucket_name> -snapshot-policy <policy_name>
```

or

```
vserver object-store-server bucket modify -vserver <svm_name> -bucket
<bucket_name> -snapshot-policy <policy_name>
```

> If you need to revert a cluster to an ONTAP version earlier than ONTAP 9.16.1, ensure that the value for `snapshot-policy` for all the buckets is set to `none` (or `-`).

Related information

S3 snapshot overview

View and restore S3 snapshots

The ONTAP S3 snapshot feature enables you to view and browse the S3 snapshot content for your buckets from S3 clients. In addition, you can restore a single object, a set

of objects, or a whole bucket on an S3 client from an S3 snapshot.

Before you begin

For viewing, browsing, and restoring ONTAP S3 snapshots on your buckets, the snapshots should have been created and the S3 base bucket should be accessible to you through the S3 protocol client.

List and view S3 snapshots

You can view the S3 snapshot details, compare them, and identify errors. Using the ONTAP CLI, you can list all the snapshots created on your S3 buckets.

Steps

1. List S3 snapshots:

   ```
   vserver object-store-server bucket snapshot show
   ```

   You can view the snapshot names, storage VMs, buckets, creation time, and `instance-uuid` of the S3 snapshots created for all your buckets on the cluster.

2. You can also specify a bucket name to view the names, creation time, and `instance-uuid` of all the S3 snapshots created for that specific bucket.

   ```
   vserver object-store-server bucket snapshot show -vserver <svm_name>
   -bucket <bucket_name>
   ```

Browse S3 snapshots content

If you notice any failures or issues in your environment, you can browse the content of the S3 bucket snapshots to identify the errors. You can also browse the S3 snapshots to determine the error-free content to restore.

S3 snapshots are presented as snapshot buckets to the S3 clients. The snapshot bucket name is formatted as `<base_bucket_name>-s3snap-<snapshot_name>`. You can see all the snapshot buckets in a storage VM using the `ListBuckets` S3 API operation.

The S3 snapshot bucket inherits the access policies of the base bucket, and supports only read-only operations. If you have permissions to access the base bucket, you can also perform read-only S3 API operations on the S3 snapshot bucket, such as `HeadObject`, `GetObject`, `GetObjectTagging`, `ListObjects`, `ListObjectVersions`, `GetObjectAcl`, and `CopyObject`.

> The `CopyObject` operation is supported on an S3 snapshot bucket only if it is a snapshot copy of the source bucket, not if it is the storage destination of the snapshot.

For more information about these operations, refer to ONTAP S3 supported actions.

Restore content from S3 snapshots

You can perform a restore operation on an S3 client to recover a single object, a set of objects, or an entire

bucket by copying content from a snapshot bucket to the original or a different bucket. You can browse snapshots to determine which snapshot content you should copy.

You restore the entire bucket, objects with a prefix, or a single object by using the `aws s3 cp` command.

Steps

1. Take a snapshot of the base S3 bucket.

```
vserver object-store-server bucket snapshot create -vserver <svm_name>
-bucket <base_bucket_name> -snapshot <snapshot_name>
```

2. Restore the base bucket using the snapshot:
   - Restore an entire bucket. Use the snapshot bucket name in the format `<base_bucket_name>-s3snap-<snapshot_name>`.

```
aws --endpoint http://<IP> s3 cp s3://<snapshot-bucket-name>
s3://<base-bucket>  --recursive
```

   - Restore objects in a directory with the prefix `dir1`:

```
aws --endpoint http://<IP> s3 cp s3://<snapshot-bucket-name>/dir1
s3://<base_bucket_name>/dir1 --recursive
```

   - Restore a single object named `web.py`:

```
aws --endpoint http://<IP> s3 cp s3:// <snapshot-bucket-name>/web.py
s3://<base_bucket_name>/web.py
```

Delete S3 snapshots

You can delete S3 snapshots that you no longer require, and free up storage space in your buckets. You can manually remove S3 snapshots, or modify the snapshot policies attached to the S3 buckets to change the number of snapshots to be retained for a schedule.

Snapshot policies for S3 buckets follow the same deletion rules as the traditional ONTAP snapshot policies. For more information about creating snapshot policies, refer to Create a snapshot policy.

About this task

- If an object version (in a versioned bucket) or an object (in a non-versioned bucket) is captured in multiple snapshots, the object will be removed from the file system only after the last snapshot protecting it is deleted.
- If you need to delete an S3 bucket containing objects with snapshots, ensure that you have deleted all the

snapshots of all the objects in that bucket.

- If you need to revert a cluster to an ONTAP version earlier than ONTAP 9.16.1, ensure that you have deleted all the S3 snapshots for all the buckets. You might also need to run the `vserver object-store-server bucket clear-snapshot-metadata` command to remove the snapshot metadata for an S3 bucket. For information, refer to Clear S3 snapshots metadata.
- When you delete snapshots in batches, you can remove a large number of objects captured in several snapshots, effectively freeing up more space than individual snapshot deletion would cause. As a result, you can reclaim more space for your storage objects.

Steps

1. To delete a specific S3 snapshot, run this command:

```
vserver object-store-server bucket snapshot delete -vserver <svm_name>
-bucket <bucket_name> -snapshot <snapshot_name>
```

2. To remove all the S3 snapshots in a bucket, run this command:

```
vserver object-store-server bucket snapshot delete -vserver <svm_name>
-bucket <bucket_name> -snapshot *
```

Clear S3 snapshots metadata

With S3 snapshots, snapshot metadata is also generated in a bucket. The snapshot metadata continues to be in the bucket even if all the snapshots are removed from it. The presence of snapshot metadata blocks the following operations:

- Cluster revert to an ONTAP version earlier than ONTAP 9.16.1
- Configuration of SnapMirror S3 on the bucket

Before performing these operations, you should clear all snapshot metadata from the bucket.

Before you begin

Ensure that you have removed all the S3 snapshots from a bucket before you start clearing the metadata.

Steps

1. To clear the snapshot metadata from a bucket, run this command:

```
vserver object-store-server bucket clear-snapshot-metadata -vserver
<svm_name> -bucket <bucket_name>
```

Copyright information

Trademark information