Veritas Access Appliance Cloud Storage Tiering Solutions Guide

7.4.3 Linux

December 2021

VERITAS™

Access Appliance Cloud Storage Tiering Solutions Guide

Last updated: 2021-12-13

Legal Notice

Veritas Technologies LLC
2625 Augustine Drive
Santa Clara, CA 95054

http://www.veritas.com

Technical Support

Technical Support maintains support centers globally. All support services will be delivered in accordance with your support agreement and the then-current enterprise technical support policies. For information about our support offerings and how to contact Technical Support, visit our website:

https://www.veritas.com/support

You can manage your Veritas account information at the following URL:

https://my.veritas.com

If you have questions regarding an existing support agreement, please email the support agreement administration team for your region as follows:

Worldwide (except Japan)	CustomerCare@veritas.com
Japan	CustomerCare_Japan@veritas.com

Documentation

Make sure that you have the current version of the documentation. Each document displays the date of the last update on page 2. The latest documentation is available on the Veritas website:

https://www.veritas.com/content/support/en_US/dpp.Appliances.html

Documentation feedback

Your feedback is important to us. Suggest improvements or report errors or omissions to the documentation. Include the document title, document version, chapter title, and section title of the text on which you are reporting. Send feedback to:

APPL.docs@veritas.com

You can also see documentation information or ask a question on the Veritas community site:

http://www.veritas.com/community/

Veritas Services and Operations Readiness Tools (SORT)

Veritas Services and Operations Readiness Tools (SORT) is a website that provides information and tools to automate and simplify certain time-consuming administrative tasks. Depending on the product, SORT helps you prepare for installations and upgrades, identify risks in your datacenters, and improve operational efficiency. To see what services and tools SORT provides for your product, see the data sheet:

https://sort.veritas.com/data/support/SORT_Data_Sheet.pdf

Contents

Chapter 1

Introduction to cloud storage tiering

This chapter includes the following topics:

About moving Access Appliance on-premises data to different cloud services

The cloud as a tier feature for a file system lets you move data to different cloud services. The data is always written to the on-premises storage tier and then data can be moved to the cloud tier using a tiering mechanism. The cloud is used as a storage tier in a file system. File metadata including any attributes set on the file resides on-premises even though the file is moved to the cloud. The cloud as a tier feature is best used for moving infrequently accessed data to the cloud.

Once you configure cloud tiering in Access Appliance, data that is stored in the file system can be intelligently moved between the on-premises tier and the cloud tier.

Access Appliance moves the data from on-premises to the following cloud providers based on automated policy management:

- Amazon Glacier - for storing data that is rarely accessed, and retrieval latency of several hours is acceptable (costs more to retrieve data).
  Amazon Glacier is part of the Amazon Web Services (AWS) suite of cloud computing services, and is designed for long-term storage of infrequently used data.
  See "Amazon Glacier considerations and retrieval options" on page 7.
- AWS S3 - for storing a variety of objects, mostly images and videos.
  Amazon S3 (Simple Storage Services) stores unstructured data in the form of objects. Objects are organized into buckets (each owned by an AWS account) and identified within each bucket by a unique, user-assigned key.
  See "AWS S3 considerations" on page 8.
- S3-compatible - for storing a variety of objects, mostly images and videos.
  S3-compatible is any third-party implementation of the Amazon S3 APIs.
  See "S3-compatible considerations" on page 8.
- AWS GovCloud (US) - for storing sensitive data and regulated workloads, helping customers support their U.S. government compliance requirements.
- Microsoft Azure
  See "Moving Access Appliance on-premises data to the Microsoft Azure cloud" on page 6.
- Google cloud.

Moving Access Appliance on-premises data to the Microsoft Azure cloud

You can move your Access Appliance on-premises data to Azure to save on storage costs and for access to Azure storage services.

You can have multiple Azure storage accounts configured with a file system cloud tier so that you can have storage beyond 500 TB available to that cloud tier.

Access Appliance includes the following features for an Azure cloud tier:

- Ability to attach multiple Azure storage accounts so that you can store data beyond Azure's limit of 500 TB per storage account.
- Ability to store 100 PB of data per Azure subscription in a single Azure tier.
- Ability to store data that is encrypted at rest using Azure's storage account encryption feature.
  Use the Azure Portal to enable or disable the storage service encryption of a storage account used in the file system cloud storage tier.
- Ability to use Azure's cool access tier for cost savings for cold data.

Use the Azure Portal to switch the access tier between hot and cool access tiers.

Support for multiple cloud tiers

You can have up to 30 cloud tiers for a file system. You can configure moving data between these different cloud tiers. For example, moving data from Azure to Glacier.

Note: You can add multiple cloud tiers using CLISH only. Only one additional tier can be added using the GUI.

Amazon Glacier considerations and retrieval options

Adding an Amazon Glacier type tier to a file system creates a vault in Amazon Glacier.

The Amazon Glacier tier is an offline tier. Read, write, and truncation file operations fail with an EIO error for files moved to the Amazon Glacier tier.

When files are moved to an Amazon Glacier tier, an archive is created per file.

If you want to read or modify the data that is moved to Amazon Glacier, move the data back to on-premises using policies.

See the Amazon Glacier website for storage and retrieval costs.

The maximum file size for moving files to Amazon Glacier is 4 GB.

The Amazon Glacier cloud tier usage statistics are not immediately reflected.

Amazon Glacier archive retrieval options:

- **Expedited** - Retrievals typically complete within 1-5 minutes.
  The expedited option is expensive and you should use it conservatively.
  Files moved from the Amazon Glacier tier with the expedited option might return the following error:

  ```
  InsufficientCapacityException (503 service unavailable)
  ```

  This error occurs if there is insufficient capacity to process the expedited request. This error only applies to expedited retrievals and not to standard or bulk retrievals.
- **Standard** - Retrievals typically complete within 3-5 hours.
- **Bulk** - Retrievals typically complete within 5-12 hours.

Note: The retrieval operation is asynchronous whereas the upload operation is synchronous for the Amazon Glacier cloud. When a file is moved to the Amazon Glacier cloud, all the data in the file is written as one archive in a glacier vault. When a file is retrieved from the Amazon Glacier cloud, a job gets initiated and the output gets downloaded after the job is completed.

Note: If you execute a policy to move files from Amazon Glacier to the primary file system, you should check if all the files are moved back using the `storage tier listfiles` command. If any of the files do not get listed, Veritas recommends that you re-run the policy.

AWS S3 considerations

Adding an S3 type tier to a file system creates a bucket in S3. When files are moved to an S3 tier, data is chunked to 64 MB sizes and objects are created for each chunk. So for a 1 GB file on-premises, when moved to S3, will have 16 objects of 64 MB size. An S3 tier can be removed only if the bucket corresponding to it is empty, so either delete all the files that were moved to S3 or move the files back to on-premises. When Amazon S3 or any S3-compatible cloud storage provider is used as the cloud tier, the data present on S3 can be accessed any time (unlike in Amazon Glacier). An EIO error is returned if you try to write, or truncate the files moved to the S3 tier.

If you want to modify the data, move the data to on-premises using policies.

See the Amazon S3 website for storage and retrieval pricing.

S3-compatible considerations

S3-compatible is any third-party implementation of Amazon S3 APIs. If you want to add any S3-compatible storage, it can be qualified with Access Appliance and used.

Access Appliance supports either AWS signature version 2 or version 4 authentication for an S3-compatible service. Access Appliance does not differentiate between S3 and S3-compatible when storing and retrieving data.

Configuring cloud storage tiering

You can configure cloud storage tiering using the following interfaces:

- GUI

File systems panel for creating a file system and for adding a cloud tier.
Settings > **Cloud storage registration** for adding a cloud subscription.

- Veritas Access command-line interface (CLI)
  See the `storage cloud(1)` and the `storage tier(1)` manual pages for more information.
  See the *Veritas Access Command Reference Guide* for a PDF version of the manual pages.
- RESTful APIs
  See the *Veritas Access RESTful API Guide* for more information.

Chapter 2

Configuring the cloud gateway

This chapter includes the following topics:

About the cloud gateway

You can configure Access Appliance as a gateway to cloud storage. You can register cloud subscriptions to your Access Appliance cluster. Multiple cloud subscriptions can be attached, so you need to assign a service name to each subscription. You can then use the service name to attach the subscription to the file system as a storage tier.

The following clouds can be added as storage tiers for a file system:

- Amazon S3
- AWS Glacier
- AWS Gov Cloud (US)
- Azure
- Google
- S3-Compatible

If you want to add any other S3-compatible storage, it can be qualified with Access Appliance and used.

The cloud as a tier feature lets you have hybrid storage that uses both on-premises storage and public or private cloud storage. After the gateway and tier are configured,

you can use the cloud as a tier feature to move data between the cloud and the on-premises storage. The files in the cloud, like the files in the on-premises storage, are accessible using the NFS, S3, and CIFS protocols. Access to the data present in the cloud tier is transparent to the application.

Before you provision cloud storage, you set up cloud subscriptions. To set up the cloud gateway, you attach cloud subscriptions to your Access Appliance cluster. You need to have the subscription credentials to add the cloud subscriptions. You need different subscription credentials based on your cloud provider.

Configuring the cloud gateway

A cloud gateway enables you to register one or more cloud services to the cluster. You can then add the cloud service as a cloud tier for a file system. Whenever you add a tier to a file system, you need to add a cloud service for the cloud tier.

You can add more than one cloud tier to the same file system and the cloud tiers can be associated with different services

---

Note: Before adding the cloud service, make sure that the pool is created on the setup because of the dependency on shared location which gets created during the first pool creation.

To add the cloud service

- Add the cloud service.

  ```
  Storage> cloud addservice service_name service_provider=AWS
  ```

  You are prompted to provide Amazon S3 subscription credentials.

To display the cloud services

- Display the added cloud services.

  ```
  Storage> cloud listservice service_name
  ```

To remove the cloud service

- Remove the cloud service.

  ```
  Storage> cloud removeservice service_name
  ```

Chapter 3

Configuring the cloud as a tier using CLISH

This chapter includes the following topics:

Configuring the cloud as a tier

You can configure the following clouds as a storage tier with a file system:

- Amazon S3
- Amazon Glacier
- AWS GovCloud (US)
- Azure
- Google cloud
- S3-compatible

If you want to add any other S3-compatible storage, it can be qualified with Access Appliance and used.

The data is always written to the on-premises storage tier and then data can be moved to the cloud tier by setting up policies. The policies can be configured based on access time, modification time, and pattern of the files. File metadata including any attributes set on the file resides on-premises even though the file is moved to

the cloud. The cloud as a tier feature is best used for moving infrequently accessed data to the cloud.

Warning: When any cloud service is used as a cloud tier for a file system, Access Appliance exclusively owns all the buckets and the objects created by Access Appliance. Any attempt to tamper with these buckets or objects outside of Access Appliance corrupts the files represented by those modified objects.

See the `storage_cloud(1)` man page for detailed examples.

See the `storage_tier(1)` man page for detailed examples.

Warning: Access Appliance cannot add a cloud tier if the clock on the Access Appliance system is more than 15 minutes out-of-sync with the actual time. To ensure that the Access Appliance clock time is accurate, configure an NTP server or use the `System> clock set` command.

To configure the cloud as a tier for a file system

1 Display the available cloud services.

```
Storage> cloud listservice service_name
```

If the cloud service is not listed, you may need to add the cloud subscription to the cluster.

See "About the cloud gateway" on page 10.

2 Add the cloud as a tier to a file system.

```
Storage> tier add cloud 'S3 | Glacier' fs_name tier_name
service_name region interface
```

interface is an optional argument. It takes the name of the public data network interface and uses this interface during cloud data transfer.

Note:

- If the cloud tier is configured over a bond which is created over two public data interfaces, and later the bond is broken, then the data transfer will fail for that cloud tier.
- If the cloud tier is configured over two public data interfaces and then a bond is created with both the data interfaces, then data transfer will fail for that cloud tier.

Amazon AWS has standard regions defined. Based on the region you choose in Access Appliance, AWS storage is served through that region. To get better performance, always select the closest geographic region.

3 Verify that the cloud tier is configured on the specified file system.

```
Storage> tier list fs_name
```

To remove the cloud tier

- Remove the cloud tier.

```
Storage> tier remove fs_name tier_name
```

If there are any files present in the cloud tier, the remove cloud tier operation fails. Move the files back to the primary tier before removing the cloud tier.

Creating and scheduling a policy for a file system

A file system can have multiple on-premises tiers for storage. You can add a cloud service as an additional tier. After a cloud tier is configured, you can move data between the tiers of the file system as needed.

Use policies to define a set of data movement rules for the file system. Each file system can include a policy for deletion and a policy for data movement between tiers.

Be careful when specifying the criteria for moving files. Conflicting policies may cause data to move from one tier to another tier. A best practice is to use policies with a smaller data set first before applying those policies to file systems using a schedule.

A data movement policy can use the following criteria to indicate which files or directories to move between tiers:

- pattern
- atime
- mtime

You can also perform a dry run of a policy.

See the `storage_fs(1)` policy section of the manual page for detailed examples.

Creating a policy

To create a policy

1. Create a data movement policy *policy1*, for file system *fs1*, to move the files with file name extensions of .txt from the primary tier (disk tier) to tier1 (cloud tier), which did not get accessed or modified for the last two days.

   ```
   Storage> fs policy add operation=move policy1 fs1 primary tier1
   *.txt atime >2d mtime >2d
   ACCESS policy SUCCESS V-288-0 Policy policy1 for fs fs1 added
   successfully.
   ```

2. Retrieve data from Amazon Glacier. Create a policy *pol1*, to move all the files with the file name extension of .txt from Amazon Glacier to the primary tier using the Bulk retrieval option.

   Files are copied to on-premises and then deleted from Amazon Glacier. The time when the files become available on-premises depends on the type of retrieval option selected.

   ```
   Storage> fs policy add operation=move pol1 gfs2 gtier primary
   retrieval_option=Bulk *.txt
   ```

3. Create a data deletion policy *policy2*, for file system *fs1* to move the files with file name extensions of .txt from tier1 (cloud tier), which did not get accessed or modified for the last two days.

   ```
   Storage> fs policy add operation=delete policy2 fs1 tier1 \*.txt
   atime >2d mtime >2d
   ACCESS policy SUCCESS V-288-0 Policy policy2 for fs fs1 added
   successfully.
   ```

4. Modify data movement policy *policy1*, for file system *fs1* to move the files with the file name extension of .doc, which did not get accessed or modified for the last three days.

   ```
   Storage> fs policy modify policy1 \*.doc atime >3d mtime >3d
   ACCESS policy SUCCESS V-288-0 Policy policy1 modified
   successfully.
   ```

5 List all the policies.

```
Storage> fs policy list
Name     FS name  Action   Source Tier     Destination Tier
=======  ======== =======  =============== =================
policy2  fs1      delete   tier1           -
policy1  fs1      move     primary         tier1

Retrieval Option  Pattern     Atime   Mtime   State
================= =========   ======  ======  =========
Standard          *.txt       >2d     >2d     not running
Standard          *.doc       >3d     >3d     running
```

6 List all the policies set for file system *fs1*.

```
Storage> fs policy list
Name     FS name  Action   Source Tier     Destination Tier
=======  ======== =======  =============== =================
policy2  fs1      delete   tier1           -
policy1  fs1      move     primary         tier1

Retrieval Option  Pattern     Atime   Mtime   State
================= =========   ======  ======  =========
Standard          \*.txt      >2d     >2d     running
Standard          \*.doc      >3d     >3d     not running
```

7 Delete policy *policy1*, set for file system *fs1*.

```
Storage> fs policy delete policy1 fs1
ACCESS policy SUCCESS V-288-0 Policy policy1 for fs fs1 deleted
successfully.
```

8 Rename *policy2* to *policy3*.

```
Storage> fs policy rename policy2 policy3
ACCESS policy SUCCESS V-288-0 Policy policy2 renamed to policy3.
```

9 Show the status of policy run for the policy *Policy1*.

```
Storage> fs policy status Policy1
Policy Name:                    Policy1
================================================
Policy Run Type:                unknown
Policy Run Status:              completed
Moved Data:                     600.0 MB
Total Moved Files:              6
Total Deleted Files:            3
```

10 Abort the currently running policy *Policy1*.

```
Storage> fs policy abort Policy1
ACCESS policy INFO V-288-0 Policy Policy1 aborted successfully.
```

11 Start a dry run of the policy *Policy1*.

```
Storage> fs policy dryrun Policy1
ACCESS policy INFO V-288-0 Policy Policy1 dryrun started in background,
please check 'fs policy status' for progress.
```

12 Pause the currently running policy *Policy1*.

```
Storage> fs policy pause Policy1
ACCESS policy INFO V-288-0 Policy Policy1 paused successfully.
```

13 Run the currently paused policy *Policy1*.

```
Storage> fs policy run Policy1
Policy Policy1 is not running currently, as it was killed/paused.
Would you like to start new run (y/n): y
ACCESS policy INFO V-288-0 Policy Policy1 run started in background,
please check 'fs policy status' for progress.
```

14 Resume the currently paused policy *Policy1*.

```
Storage> fs policy resume Policy1
ACCESS policy INFO V-288-0 Policy Policy1 resume started in background,
please check 'fs policy status' for progress.
```

Scheduling a policy

Creating a policy schedule

1 To create a policy schedule:

```
storage> fs policy schedule create fs1 23 20
ACCESS policy_schedule SUCCESS V-493-10-2945
Schedule create for file system fs1 done successfully.
```

2 To list a policy schedule:

```
storage> fs policy schedule list fs10
File System Name Minute Hour Day of Month Month Day of Week
================ ====== ==== ============ ===== ===========
fs1              23     20   *            *     *
byos1> storage fs policy schedule remove fs1
ACCESS policy_schedule SUCCESS V-493-10-2947
Schedule for file system fs1 removed successfully.
```

3 To remove a policy schedule:

```
storage> fs policy schedule remove fs1
ACCESS policy_schedule SUCCESS V-493-10-2947
Schedule for file system fs1 removed successfully.
```

About policies for file systems

When a file system includes a cloud tier, you can use policies to control the movement of data between the on-premises storage and the cloud tier. A policy is a set of rules defined for a file system for deleting data or moving data. If you want to specify repeatable rules for maintaining data on the tiers, you can set up a policy for the file system.

Each rule defines the following criteria:

- What action should be taken (move or delete)
- When the data should be moved or deleted based on the access time or modified time of the file
- Which data should be moved based on the pattern matching of the files and directories.
  See “About pattern matching for data movement policies” on page 19.

Each file system can have more than one rule, though you should be careful not to create rules that conflict or cause looping.

To move or delete the data, you run the policy. When you run a policy, the rules are applied to the data at the time the policy is run. By default, the policy does not run automatically. You can attach a schedule to a policy to have it run automatically at the specified times.

About pattern matching for data movement policies

Within a policy, you can use a pattern to specify that the rule applies to file names or directory names that match the pattern. By using a pattern, you do not need to know the exact file name in advance; the files that match the pattern are selected dynamically.

A pattern uses special characters, which are case sensitive. There are the following types of patterns:

- Directory patterns
  A pattern that ends with a slash (/) is matched only against directories.
- File patterns
  A pattern that does not end with a slash (/) is matched only against files.

The following is a list of supported special characters and their meanings:

* (asterisk)	Matches any character any number of times.
? (question mark)	Matches any single character.

** (two asterisks)	Matches across child directories recursively. The pattern `fs1/*/*.pdf` will match `.pdf` file names present after first sub-directory of `fs1/`. For example, if the following files exist: fs1/dir1/a.pdf fs1/dir2/b.pdf fs1/dir3/dir4/c.pdf dir5/d.pdf e.pdf then the pattern `fs1/*/*.pdf` will match only `a.pdf` and `b.pdf`. It will not match `c.pdf`, `d.pdf` and `e.pdf`. The pattern `fs1/**/*.pdf` will match `.pdf` files in any directory after `fs1`. For the above file list, it will match all of the files: `a.pdf`, `b.pdf`, and `c.pdf`. The pattern `**/*.pdf` will match all the .pdf files in the whole file system. For the above file list, it will match all of the files: `a.pdf`, `b.pdf`, `c.pdf`, `d.pdf` and `e.pdf`.
[ ] (square brackets)	Matches either range or set of characters. [0-5] will match any character in range of 0 to 5. [a-g] will match any character in range of a to g. [abxyz] will match any one character out of a,b,x,y,z.
! (exclamation point)	Can be used as the first character in a range to invert the meaning of the match. [!0-5] will match any character which is not in range of 0 to 5.
\ (backslash)	Can be used as an escape character. Use this to match for one of the above pattern matching characters to avoid the special meaning of the character.

About schedules for running policies

A schedule is specified in a format similar to the UNIX `crontab` format. The format uses five fields to specify when the schedule runs:

minute	Enter a numeric value between 0-59, or an asterisk (*), which represents every minute. You can also enter a step value (*/x), or a range of numbers separated by a hyphen.

hour	Enter a numeric value between 0-23, or an asterisk (*), which represents every hour. You can also enter a step value (*/x), or a range of numbers separated by a hyphen.
day_of_the_month	Enter a numeric value between 1-31, or an asterisk (*), which represents every day of the month. You can also enter a step value (*/x), or a range of numbers separated by a hyphen.
month	Enter a numeric value between 1-12, or an asterisk (*), which represents every month. You can also use the names of the month. Enter the first three letters of the month (you must use lower case letters). You can also enter a step value (*/x), or a range.
day_of_the_week	Enter a numeric value between 0-6, where 0 represents Sunday, or an asterisk (*), which represents every weekday. You can also enter the first three letters of the week (you must use lower case letters). You can also enter a step value (*/x), or a range.

A step value (*/x) specifies that the schedule runs at an interval of x. The interval should be an even multiple of the field's range. For example, you could specify */4 for the hour field to specify every four hours, since 24 is evenly divisible by 4. However, if you specify */15, you may get undesired results, since 24 is not evenly divisible by 15. The schedule would run after 15 hours, then 7 hours.

A range of numbers (two values separated by a hyphen) represents a time period during which you want the schedule to run.

Examples: To run the schedule every two hours every day:

```
0 */2 * * *
```

To run the schedule on 2:00 a.m. every Monday:

```
* 2 * * 1
```

To run the schedule at 11:15 p.m. every Saturday:

```
15 23 * * 6
```

Obtaining statistics on data usage in the cloud tier

You can find the following information for data stored in the cloud tier:

- Storage utilization in the cloud

- Number of the objects that are stored in the cloud
- Number of the files that are stored in the cloud

See the `storage_tier` man page for detailed examples.

To show the total data usage in the cloud tier

- Show the total data usage in the cloud tier.

  ```
  Storage> tier stats usage fs_name
       tier_name
  ```

  Unlike the `Storage> tier stats show` command, these statistics are persistent across reboots.

  Example:

  ```
  Storage> tier stats usage fs1 cloudtier
  Storage Utilized   :223.1GB
  Number of objects  :488
  Number of files    :231
  ```

  This example shows that 223.1 GB is used in the cloud tier. Based on the size of the file, each file is chunked to multiple objects when moved to the cloud, so 231 files were stored as 488 objects in the cloud.

To display vendor specific cloud usage

- Show the vendor specific cloud usage.

  ```
  Storage> cloud usage
  ```

  Example:

  ```
  Storage> cloud usage
  SCLOUD          Files  Objects  Size(Bytes)
  ==============  =====  =======  ===========
  AZURE           6      6        98304
  AWS             34     34       39556096
  GOOGLE          60     60       2891776
  GOVCLOUD        45     45       876544
  S3-COMPATIBLE   6      6        33660928
  ```

Chapter 4

Configuring the cloud as a tier using GUI

This chapter includes the following topics:

Workflow for adding a cloud tier

By default, a file system has a single primary tier, which is the on-premises storage for the file system. You can add only one additional cloud tier using the GUI. After a cloud tier is configured, you can move data between the tiers of the file system.

Prerequisites:

- Ensure that the cloud vendor and file system are created.
- You have to configure the NTP server to add a cloud vendor.

To perform a cloud storage registration

1. Navigate to **Settings** > **Cloud storage registration**. Click **Add**.
2. Enter the details in the **Add cloud registration** form. Select the cloud service provider. Enter the required input fields for the selected cloud service provider.
3. Click **Add** to add a cloud subscription for the selected cloud provider.

To create a file system

1. Navigate to **File systems** and click **Create** to create a file system.
2. In the **Create file system** form, enter the name and size of the file system. You can click **Advanced** to specify other details such as layout, and block size.
3. Select a storage pool for the on-premises storage and click **Create**.

You can now add a cloud tier.

To add a cloud tier

1. Navigate to **File systems**. Click on the name of the file system. Select the **Storage** tab to add the cloud tier.
2. Select **Add cloud tier**. In the **Add cloud tier** form, specify a name for the cloud tier and provide other details, such as cloud subscription, service provider and region name.
   - For **AWS**: Select the tier type as Glacier or S3.
   - For **AWS GovCloud (US)**: Select the tier type as S3.
   - For **Azure**: You can add additional Azure storage accounts to the same cloud tier to expand the capacity of the cloud tier beyond 500 TB. Simply repeat this procedure with a different Azure storage account, and specify the same cloud tier name to add an additional storage account to the cloud tier. You provide the storage account information when you add the cloud subscription. Storage account information is contained in the cloud subscription.
   - For **Google**: The access credentials are present in the .json file. Select the tier type: Coldline, Multi-Regional, Nearline, Regional
   - For **S3-compatible**: Enter the access key, secret key, and the REST endpoint of the S3 server.
3. Click **Add** to add the cloud tier.
4. View the **View details** panel for the status of the task.

You can remove an existing cloud tier. Right-click the ellipses icon next to the cloud tier and select **Remove cloud tier**. This operation deletes associated policies and schedules which are configured on the cloud tier.

Adding a secondary tier

You can add a secondary tier to move aged data to secondary storage.

To add a secondary tier

1. In the **File Systems** tab, select the file system for which you want to add the secondary tier.
2. Select the **Storage** tab, and select the **Add secondary tier** button.
3. Specify the tier name size.
4. Click **Advanced** to specify the layout and other details.
5. Select the storage pool for storing your on-premises data. Click **Next**.
6. Click **Finish**.
7. Check the **View details** panel for the status of the task.

Creating policy rules using GUI

For each file system, you can create policy rules. The file system policy, which may consist of multiple rules, is run at the designated time.

Note: To create a policy rule, a secondary or cloud tier must already exist on the file system.

To configure a policy schedule

1. Navigate to **File systems**. Click on the name of the file system. Select the **Storage** tab. In the **Configured policy rules** section, click on the **+** icon.
2. In the **Add policy** form, enter all the required details. Choose **Data movement** if you want to move data. Choose **Data deletion** if you want to delete data. Enter a name for the new policy.

   Enter the schedule frequency. Specify the hours and minutes that represent the time of day according to a 24-hour clock.
3. Click **Add** to add the new policy.
4. You can see the **View details** panel for the status of the task. The newly created policy is listed in the **Configured policy rules** section.

It is possible to edit or remove a policy rule.

Click the **Edit** icon to modify the details of an existing policy. You can enter your changes in the **Edit policy** form.

Click the **Delete** icon to remove an existing policy.

Creating schedules using GUI

For each policy created for the file system, you can create a schedule. The schedule is run at the designated time.

To create a policy schedule

1. **Prerequisites:**

   Ensure that the policy is already configured.

2. Create or modify the schedule by selecting the **Edit** icon.

   In the **Configuration schedule** form, select one of the following:

   - Daily
   - Weekly. Specify the day of the week on which to run the schedule.
   - Monthly. Specify the day of the month on which to run the schedule..

   For the Start time, specify the hours and minutes that represent the time of day according to a 24-hour clock.

3. Click **Next**. The task is initiated.
4. Click **Finish**.
5. View the **Recent Activity** panel for the status of the task.

The same wizard is used to create, modify, or remove the policy schedule.

Viewing information about the secondary tier

Access Appliance provides two types of tiers: primary tier and a secondary tier. You can add a secondary tier to the file system for moving aged data to secondary storage.

You can view the following information about the secondary tier:

- Tier size
- Tier used size
- Subscription name