Veritas Access Cloud Storage Tiering Solutions Guide

Linux

7.4

VERITAS™

Veritas Access Cloud Storage Tiering Solutions Guide

Last updated: 2018-07-24

Document version: 7.4 Rev 0

Legal Notice

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 document version appears on page 2 of each guide. The latest documentation is available on the Veritas website:

https://sort.veritas.com/documents

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:

doc.feedback@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

Use cases for cloud storage tiering

This chapter includes the following topics:

- About moving Veritas Access on-premises data to different cloud services
- Benefits of moving Veritas Access on-premises data to the Microsoft Azure cloud

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

The cloud as a tier feature for a scale-out 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 scale-out 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 Veritas Access, data that is stored in a scale-out file system can be intelligently moved between the on-premises tier and the cloud tier.

Veritas Access 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 10.

- 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 11.

- AWS 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 11.

- AWS GovCloud (US) - for storing sensitive data and regulated workloads, helping customers support their U.S. government compliance requirements.

- Microsoft Azure
  See "Benefits of moving Veritas Access on-premises data to the Microsoft Azure cloud" on page 6.

- Alibaba

- Google cloud

- IBM Cloud Object Storage

Support for multiple cloud tiers

You can have up to eight cloud tiers for a scale-out file system. You can configure moving data between these different cloud tiers, moving data from Azure to Glacier for example.

Benefits of moving Veritas Access on-premises data to the Microsoft Azure cloud

You can move your Veritas Access on-premises data to Azure for saving on storage costs and for access to Azure storage services. You can have multiple Azure storage accounts configured with a scale-out file system cloud tier so that you can have storage beyond 500 TB available to that cloud tier.

Veritas Access 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 a scale-out 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.

Moving infrequently used data to the cloud

This chapter includes the following topics:

- Configuring cloud storage tiering
- Workflow for adding a cloud tier
- Amazon Glacier considerations and retrieval options
- AWS S3 considerations
- S3-compatible considerations

Configuring cloud storage tiering

You can configure cloud storage tiering using the following Veritas Access interfaces:

- GUI
  File Systems panel for creating a scale-out file system and for adding a cloud tier.
  Settings > Cloud Storage Registration for adding a cloud subscription.

- CLISH
  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.

Workflow for adding a cloud tier

By default, a scale-out file system has a single primary tier, which is the on-premises storage for the scale-out file system. 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 scale-out file system.

See "Amazon Glacier considerations and retrieval options" on page 10.

See "Benefits of moving Veritas Access on-premises data to the Microsoft Azure cloud" on page 6.

To add a cloud tier (Alibaba, AWS, AWS GovCloud (US), Azure, or Google, or S3-compatible), perform the actions in the order listed (GUI operations order):

1. Add a cloud subscription for the selected cloud provider.

2. Create a scale-out file system with a minimum file system size of 10 GB.

3. Select a storage pool for the on-premises storage.

4. Specify that you want to add cloud storage and specify a name for the cloud tier.

5. Select the cloud storage provider and other fields.

   Alibaba

   Select the cloud service.

   Select a region.

   AWS

   Select the cloud service.

   Select the tier type, Glacier or S3

   Select the region.

   AWS GovCloud (US)

   Select the cloud service.

   Select the tier type, Glacier, S3, or S3 (FIPS)

   Azure

   Select the cloud service.

Note: 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.

Google

Select the cloud service.

A .json file is required. The .json file is contained in the cloud service.

Select the tier type, Coldline, Multi-Regional, Nearline, Regional

Select a region.

IBM Cloud Object Storage

Select the cloud service.

Select a region.

S3-compatible

Select the cloud service.

The access key, secret key, and the REST endpoint of the S3 server are contained in the cloud service.

Amazon Glacier considerations and retrieval options

Adding an Amazon Glacier type tier to a scale-out 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 back the data to on-premises using `Storage> tier move start` or 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.

AWS S3 considerations

Adding an S3 type tier to a scale-out 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 back the files 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 `Storage> tier move start` or 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. You can add an S3-compatible storage service as a cloud tier. Veritas Access supports either AWS signature version 2 or version 4 authentication for an S3-compatible service. Veritas Access does not differentiate between S3 and S3-compatible when storing and retrieving data.