Enterprise Vault™ Migrating Data Using the Simple Storage Service (S3) API

11.0.1 or later

VERITAS™

Enterprise Vault™: Migrating Data Using the Simple Storage Service (S3) API

Last updated: 2021-03-23.

Legal Notice

Veritas Technologies LLC
2625 Augustine Drive
Santa Clara, CA 95054

https://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

Before you contact Technical Support, run the Veritas Quick Assist (VQA) tool to make sure that you have satisfied the system requirements that are listed in your product documentation. You can download VQA from the following article on the Veritas Support website:

https://www.veritas.com/support/en_US/vqa

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/docs/100040095

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:

evdocs@veritas.com

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

https://www.veritas.com/community

Contents

Chapter 1 Introducing the Enterprise Vault S3 API migrator 5

About the Enterprise Vault S3 API migrator 5
Requirements for using the Enterprise Vault S3 API migrator 6
Configuring the Enterprise Vault S3 API migrator 6

Chapter 2 Configuring the Enterprise Vault S3 API migrator 7

Getting the S3 access credentials 7
Configuring Simple Storage Service (S3) API as secondary storage 7
Simple Storage Service (S3) API server properties 8

Appendix A Troubleshooting 15

Using DTrace to view diagnostic logs 15

Chapter 1

Introducing the Enterprise Vault S3 API migrator

This chapter includes the following topics:

- About the Enterprise Vault S3 API migrator
- Requirements for using the Enterprise Vault S3 API migrator
- Configuring the Enterprise Vault S3 API migrator

About the Enterprise Vault S3 API migrator

The Enterprise Vault Simple Storage Service (S3) API storage migrator lets you migrate archived data to and retrieve it from any S3-compliant storage solution. You can use the S3 API to connect to cloud storages and hosting providers and use them as secondary storage to store infrequently accessed data.

The Enterprise Vault S3 API storage migrator is installed as part of the Enterprise Vault 11.0.1 Cumulative Hotfix 4 or later installation. It moves CAB files, created by the Enterprise Vault file collection software, to S3-compatible storage.

This guide shows how to configure the Enterprise Vault S3 API storage migrator to migrate data to the S3-compatible storage.

Caution: Using S3-compliant cloud storage as a secondary storage can make some Enterprise Vault operations, such as retrieval, take a long time to process.

Requirements for using the Enterprise Vault S3 API migrator

The following are the requirements for using the Enterprise Vault S3 API migrator:

- Veritas Enterprise Vault 11.0.1 Cumulative Hotfix 4 or later.
- The Access Key ID and Secret Access Key to connect to the storage provider.
- At least one S3 bucket to store data. Buckets represent the actual storage in the S3 compliant cloud storage.

Note: This guide assumes a working knowledge of Enterprise Vault tasks such as creating and configuring vault store partitions, and an understanding of the storage solution that is using the S3 API.

Configuring the Enterprise Vault S3 API migrator

The following table outlines the steps with which you can configure the Enterprise Vault S3 API migrator.

Table 1-1 Configuring the Enterprise Vault S3 API migrator

Step	Action	Description
Step 1	Get the S3 access credentials.	See "Getting the S3 access credentials" on page 7.
Step 2	For each vault store partition, configure Simple Storage Service (S3) API as secondary storage.	See "Configuring Simple Storage Service (S3) API as secondary storage" on page 7.

Chapter 2

Configuring the Enterprise Vault S3 API migrator

This chapter includes the following topics:

- Getting the S3 access credentials
- Configuring Simple Storage Service (S3) API as secondary storage
- Simple Storage Service (S3) API server properties

Getting the S3 access credentials

Before you migrate files to the S3-compliant cloud storage, you must have an S3 Access Key ID and a Secret Access Key. You need the Access Key ID and Secret Access Key to connect to the storage provider.

Note: The S3 API supports Signature V2 for authentication.

Configuring Simple Storage Service (S3) API as secondary storage

Use the Enterprise Vault Administration Console to enable and configure Simple Storage Service (S3) API as secondary storage.

For each vault store partition for which you want to configure Simple Storage Service (S3) API as secondary storage, follow these steps.

To configure Simple Storage Service (S3) API as secondary storage

1. Start the Enterprise Vault Administration Console.
2. In the left pane of the Administration Console, expand the Enterprise Vault site hierarchy until the vault store partition from which you want to migrate data is visible.
3. Right-click the vault store partition and click **Properties**. You can also double-click the vault store partition to open the vault store partition properties page.
4. In the **Collection** tab, do the following:
   - Select **Use collection files** to enable collection.
   - Select **Enterprise Vault** as the file collection software.
   - Specify the daily collection times and the amount of time that must elapse since items were archived before they are eligible for collection.

   You can also enable and configure collection when creating a new vault store partition using the **New Partition** wizard.
5. In the **Migration** tab, do the following:
   - Select **Migrate files** to enable migration.
   - Select **Simple Storage Service (S3) API** as the migrator software.
   - Specify the amount of time that must elapse since a collection file was created before it is eligible for migration.
   - Choose whether to remove the collection files from primary storage after they have been migrated. Specify the amount of time that must elapse since the migration before the files are deleted.

   You can also enable and configure migration when creating a new vault store partition using the **New Partition** wizard.
6. In the **Advanced** tab, configure the S3 API storage server properties.

Simple Storage Service (S3) API server properties

After you select Simple Storage Service (S3) API as the secondary storage for a vault store partition, configure the properties for your S3-complaint cloud storage. Use the **Advanced** tab of the vault store partition properties page to configure the storage server properties.

Most of the configuration parameters are populated with their default values. You need to manually set the following parameters:

- Secure access key ID
- Shared secret
- Bucket name

Table 2-1 S3-complaint cloud storage server properties

Option	Description	Default value
Service host name	The fully qualified host name of the computer that hosts the S3 service.	
Http port	The http port number on which the S3 service is configured.	80
Https port	The https port number on which the S3 service is configured.	443
Use SSL for control	When set to **Yes**, Enterprise Vault uses the TLS 1.2 protocol to establish a secure connection to the S3 storage server. Note that the S3 API does not support self-signed certificates.	Yes
Use SSL for data r/w	When set to **Yes**, Enterprise Vault uses the TLS 1.2 protocol to establish a secure connection for read and write operations to the S3 storage server.	Yes
Access Key ID	The secure access key ID, also known as the username, provided by the S3-complaint cloud storage vendor.	None
Secret Access Key	The account shared secret, also known as the password, provided by the S3-complaint cloud storage vendor.	None

Table 2-1 S3-complaint cloud storage server properties *(continued)*

Option	Description	Default value
Supported regions	The geographical regions that are supported by the cloud storage solution for storing data. Regions must be specified in the format *identifer, displayname, servicehostname*. Use semicolons to separate multiple region names. For example, us-west-1,US West,s3-us-west-1.amazonaws.com; eu-west-1,EU West,s3-eu-west-1.amazonaws.com. Note the following: ■ Each region should contain 3 parts: *identifier*, *display name* and *service host*. The identifier must be a location constraint of the region in cloud storage solution. ■ Specify the default region for the cloud storage first. The first region specified in the supported region list is treated as the default region. ■ If the cloud storage solution has region support you must specify all regions. ■ Do not add whitespace in the identifier and service host. ■ None of the fields of the region list should be blank. ■ Multiple regions should be separated with semicolon. ■ Each field in the region details should be separated with a comma. ■ The S3 API supports Signature V2 for authentication.	

Table 2-1 S3-complaint cloud storage server properties *(continued)*

Option	Description	Default value
	■ The region list should end with a semicolon. ■ You can specify only English characters in the region list. ■ Do not specify duplicate region names. ■ All regions should have similar SSL configuration, either enabled or disabled. Mixed SSL configuration is not supported.	
Bucket name	The name of the bucket. The bucket name must be unique across all existing bucket names in the S3 storage solution. To ensure that you use a unique name you could prefix your bucket names with your company's name. There are other requirements that you need to take care of while naming the buckets. Check the appropriate storage solution documentation for bucket naming requirements and guidelines.	
Bucket region	The geographical location where the bucket is created.	
Bucket access type	Specifies whether the URL is virtual-hosted-style URL or path-style URL. In a virtual-hosted-style URL, the bucket name is part of the domain name in the URL. For example: http://bucket.s3.amazonaws.com. In a path-style URL, the bucket name is not part of the domain. For example: http://s3.amazonaws.com/bucket.	Virtual

Table 2-1 S3-complaint cloud storage server properties *(continued)*

Option	Description	Default value
Write buffer size	The buffer size, in megabytes, Enterprise Vault uses for data uploads. Ensure that this value is greater than the **Maximum collection file size** setting on the **Collections** tab of the vault store partitions page. Set this option to zero (0) to disable the use of buffers.	20
Read buffer size	The buffer size, in megabytes, Enterprise Vault uses for data downloads.	20
Log level	The amount of detail to include in the log file. You can select from the following: ■ No logging ■ Errors only ■ Errors, Warnings ■ Errors, Warnings, Info ■ Everything **Note:** If you choose **No logging**, Enterprise Vault does not log cURL messages even if `Log CURL messages` is set to Yes.	Errors, Warnings
Log CURL messages	Specifies whether to log cURL activity. cURL is a command line tool for sending or receiving files using URL syntax. Enterprise Vault uses the cURL library to transfer data to the S3 cloud.	No
CURL connect timeout	The maximum amount of time, in seconds, the Enterprise Vault S3 API Migrator waits to connect to the S3 storage server. This only limits the connection phase, not the session time.	300

Table 2-1 S3-complaint cloud storage server properties *(continued)*

Option	Description	Default value
CURL operation timeout	The maximum amount of time, in seconds, the Enterprise Vault S3 Migrator waits to transfer data to and retrieve data from the S3 storage server.	900
CURL proxy type	The CURL proxy types. Proxy types are HTTP, SOCKS, SOCKS4, SOCKS5, SOCKS4A.	None
CURL proxy IP	The CURL proxy IP. By default, no proxy is used.	<your_proxy_ip>
CURL proxy port	The CURL proxy port number.	0
User wait timeout	If an Enterprise Vault user's request to retrieve an archived item from the S3 storage server takes longer than normal, specifies the number of seconds after which to present the user with the following message: "The archived item is being retrieved from a slow device. Try again later." Enterprise Vault continues to retrieve the item in the background until the **System wait timeout** period has elapsed. Enterprise Vault then abandons the attempt to retrieve the item, and the user must submit the retrieval request again. The recommended value is 40 seconds.	40
System wait timeout	If an attempt to retrieve an archived item from the S3 storage server takes an excessively long time, specifies the number of seconds after which to abandon the attempt and remove the requested item from the retrieval queue. The recommended value is 900 seconds.	900

Table 2-1 S3-complaint cloud storage server properties *(continued)*

Option	Description	Default value
Recalled file cache period	The number of days, since the last accessed date, that Enterprise Vault should retain recalled files in the cache. The collection process deletes the recalled files when the cache period has elapsed.	7
Migrate all files	If the value is set to Yes, Enterprise Vault forces all eligible files to be collected and migrated. Setting this value to Yes may cause Enterprise Vault to create a large number of collection files. If the value is set to No, Enterprise Vault may leave some saveset files uncollected and thus unmigrated.	No

Appendix A

Troubleshooting

This appendix includes the following topics:

- Using DTrace to view diagnostic logs

Using DTrace to view diagnostic logs

If you encounter problems when you store or retrieve archived data with the Enterprise Vault S3 API migrator, you can run the DTrace utility to help you identify their cause. DTrace let you monitor multiple services simultaneously, write the trace to a file, filter for specific words, and trigger tracing based on filters.

The following table lists the migrator processes for which you can get diagnostic logs with DTrace.

Note: We recommend that you set the monitoring level to Verbose in all cases.

Table A-1 Troubleshooting with DTrace logs

Monitor this process	To do this
EVStgOfflineOpns.exe	To get diagnostic logs for the restore or recall of the archived files that are located in the S3 cloud.
StorageFileWatch.exe	To get diagnostic logs for the following operations: ■ The upload of data to the S3 cloud through the Enterprise Vault collection and migration processes. ■ The deletion of empty collection files from the primary storage.

Table A-1 Troubleshooting with DTrace logs *(continued)*

Monitor this process	To do this
StorageManagement.exe	To view the information that is logged when Enterprise Vault checks whether the migrator software is installed and registered correctly and validates the configured settings. Enterprise Vault logs this information when you click the **Test** button on the **Advanced** tab of the vault store properties page. You can set a filter on the “OST Streamer” keyword to view logs specific to the S3 API storage migrator.

For more information on DTrace, see the *Utilities* guide.