



Cloud Storage for Enterprise Vault

Provided by
Business Critical Services

Table of Contents

Contributors.....	3
Revision History	3
Introduction	4
1. Primary partition on Cloud Storage	5
1.1 Dell EMC Elastic Cloud Storage (ECS)	5
1.2 IBM Cloud Object Storage	10
2. Secondary Storage (Migrator)	17
2.1 Common Terminology	17
2.2 Amazon Web Service S3 (Simple Storage Service)	19
2.3 Microsoft Azure Data Blob storage	25
2.4 Google Cloud Storage	30
2.5 Rackspace Cloud Files	35
Troubleshooting	42

Contributors

<i>Who</i>	<i>Contribution</i>
Pradeep Papnai, Business Critical Engineer	Author
Rajesh Nagarkar, Business Critical Engineer	Reviewer
May Ting, Business Critical Account Manager	Reviewer
Patti Rogers, Business Critical Engineer	Reviewer

Revision History

<i>Version</i>	<i>Date</i>	<i>Change</i>
1.0	June 2018	Initial Version

Introduction

This document is designed to assist VERITAS customers and partners who wish to configure Enterprise Vault with cloud based primary OR secondary storage for vault store partitions, with examples of the most commonly used cloud service providers such as EMC, IBM, AWS, and Azure.

Readers of this document should already have basic knowledge of Enterprise Vault and reviewed Veritas and cloud vendor supplied documentation. This document does not cover the theoretical concepts; rather, it is a 'How to do' guide.

The steps and screenshots of this document are taken from the latest available version of Enterprise Vault (12.3.0) and cloud providers that may vary with previous OR future versions of the products. Please always review the latest compatibility guide for Enterprise Vault before implementing cloud storage in a production environment.

This presentation is provided for informational purposes only and is not intended as advertising of VERITAS or any other cloud service provider. The information in this document is subject to change without notice.

Customers who purchase Veritas or third party vendor's offerings should make their purchase decision based upon features available at the time of purchase.

If you have any feedback or questions about this document, please email them to ii-tec@veritas.com stating the document title.

1. Primary partition on Cloud Storage

Vault store partitions can be placed on different physical disks and on various types of storage medium. Following two examples shows how cloud storage can be used to store vault partition.

1.1 Dell EMC Elastic Cloud Storage (ECS)

Dell EMC provides streamer based Vault store partition that allows EV to write files on S3 bucket associated with Namespace created on EMC cluster.

- Access the ECS Test Drive portal at <https://portal.ecstestdrive.com/account/register> and complete the registration process. Once you are registered for an ECS Test Drive account, log in and click the **CREDENTIALS** link at the top of the page

The screenshot shows the Dell EMC ECS Test Drive portal. At the top, there is a navigation bar with links for HOME, FAQ, SUPPORT, CREDENTIALS, and LOG OUT. Below the navigation bar, the page title "ECS TEST DRIVE" is displayed, followed by links for MANAGE SECRET KEYS, SWIFT PASSWORD, and CAS PEA F. The main content area is titled "All Credentials" and lists the following credential details:

AWS S3
Endpoint: <https://object.ecstestdrive.com>
Public Endpoint: [http://131693042649205091.public.ecstestdrive.com/\[bucket_name\]/\[key_name\]](http://131693042649205091.public.ecstestdrive.com/[bucket_name]/[key_name])
Access Key: 131693042649205091@ecstestdrive.emc.com
Secret Key1: fbHlum2QY3A5xSr7Vlx635+USGw3O1ULrHS9jmom
Secret Key2: DYjW8jwXmWPACQVaQRTOIXc2/6yk1JFIDU3ELI

Atmos
Endpoint: <https://atmos.ecstestdrive.com>
Subtenant ID: 1f0d6e48874d4a3f87aadf7800cefa03
UID: 131693042649205091@ecstestdrive.emc.com
Secret Key1: fbHlum2QY3A5xSr7Vlx635+USGw3O1ULrHS9jmom
Secret Key2: DYjW8jwXmWPACQVaQRTOIXc2/6yk1JFIDU3ELI

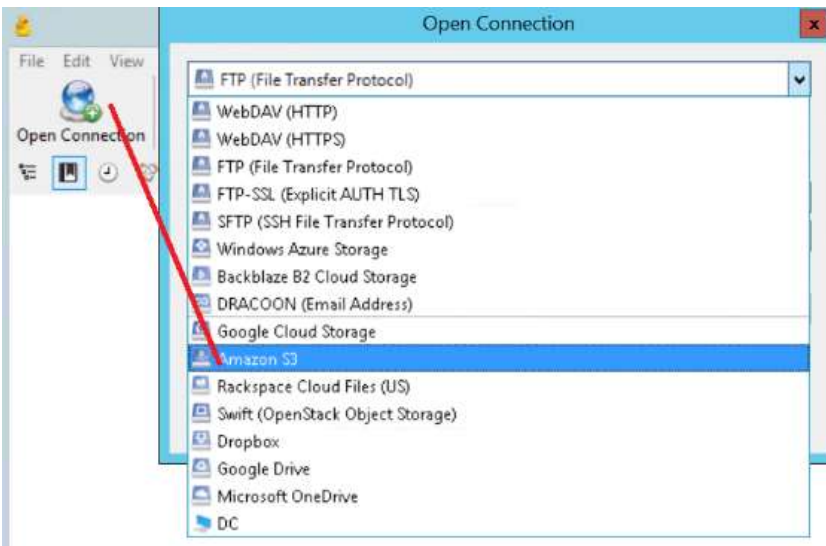
Swift / Swift-Keystone
Endpoint (swauth): <https://swift.ecstestdrive.com/auth/v1.0>
Endpoint (keystone): <https://swift.ecstestdrive.com/v2.0>
Username: 131693042649205091@ecstestdrive.emc.com
Tenant ID: 131693042649205091

Centera CAS
Endpoint: http://cas.ecstestdrive.com/your_pea_file.pea

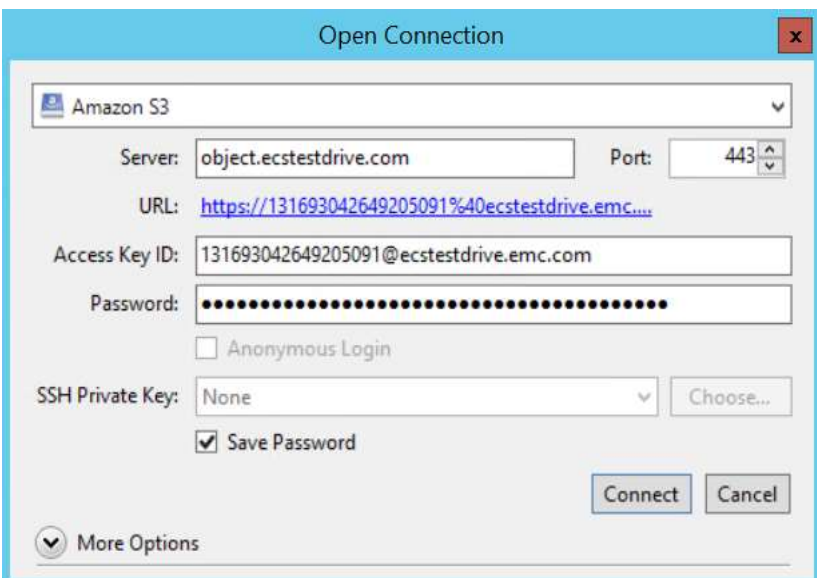
ECS Management
Endpoint: <https://portal.ecstestdrive.com>
Replication Group ID: urn:storageec:ReplicationGroupInfo:104b3728-fba1-41b3-8055-4592348f1d24:global
Namespace: 131693042649205091
Username: 131693042649205091-admin
Password: N2Q4NGU3Mjg3ZDNIbG9mYmVxYzYwMWUyZmE1NzVmMzk=

- In **AWS S3** section, copy **Endpoint, Access Key, Secret key**
- Use one of S3 compliant client application such as CyberDuck, cloudberry OR S3Browser to create bucket for Enterprise vault partition. Following steps shows the bucket creation using **CyberDuck**


- Click on **Open Connection**
- Choose **Amazon S3** from the first dropdown list



- Use information such as **server name**, **Access Key ID**, **password (secret) key** from the credential page then click on **connect**




- Click on **New Folder**, give an appropriate name to the bucket, eg. *'storage-ecs-part1'*. The bucket name must be unique and lower case



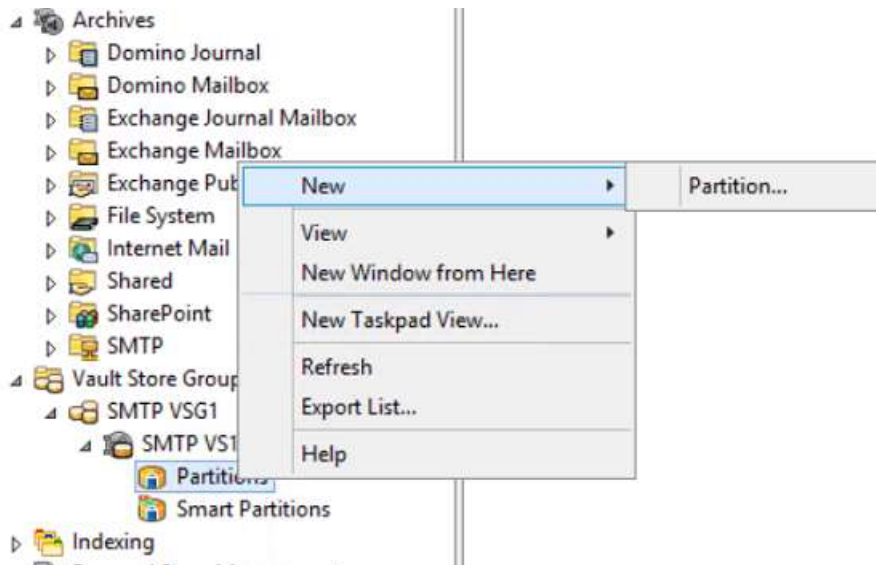
Creating Enterprise Vault Partition

- Download and install ECS Streamer driver <http://support.emc.com> on EV server

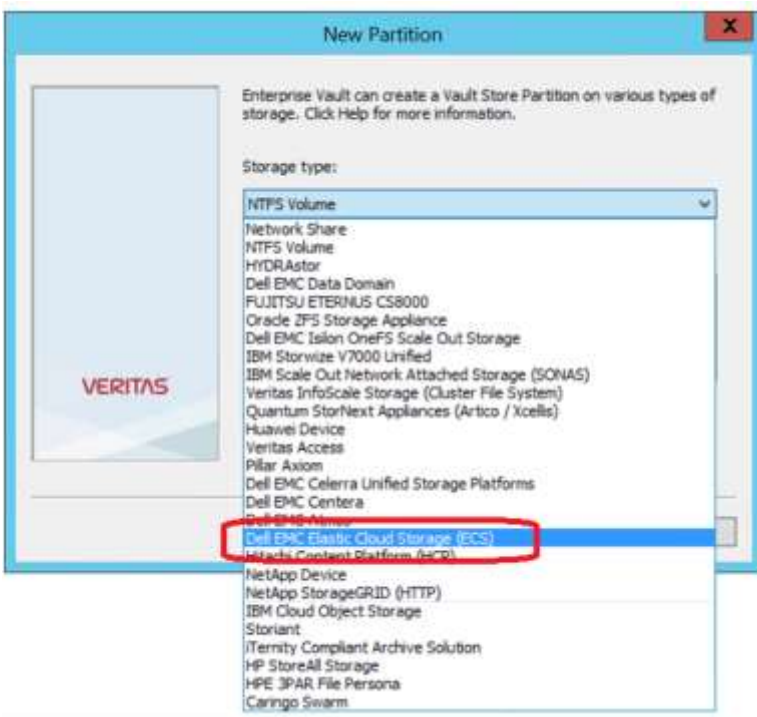


(Follow the installation steps as per the wizard)

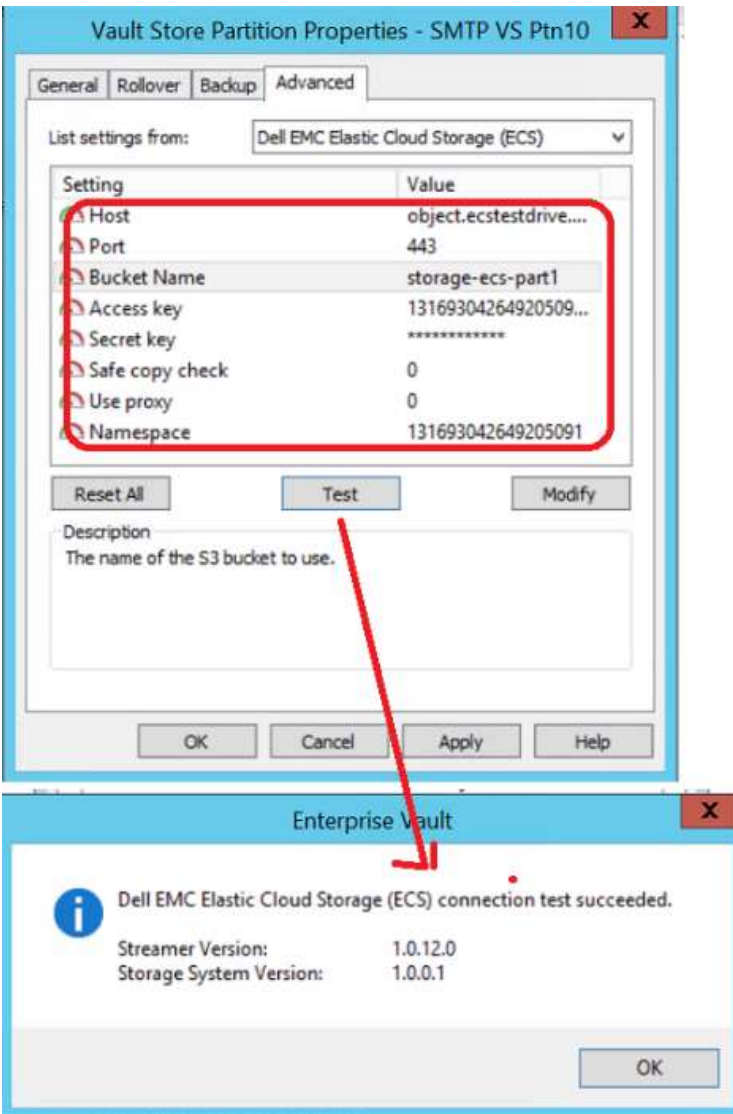
- Create a new **Partition**



- Select **Dell EMC Elastic cloud storage (ECS)**



- Fill in details of **hostname, port no., bucket name, access key, secret key, namespace (copy from ECS credential page)** of ECS configuration then click on **Test**. A *"Dell EMC Elastic Cloud Storage (ECS) connection test succeeded"* message should display



Once the partition setup completes successfully you can start archiving on Dell-EMC ECS partition.

Additional reference material provided by VERITAS:

How to configure EMC Elastic Cloud Storage (ECS) as an Enterprise Vault Partition

https://www.veritas.com/support/en_US/article.100039067

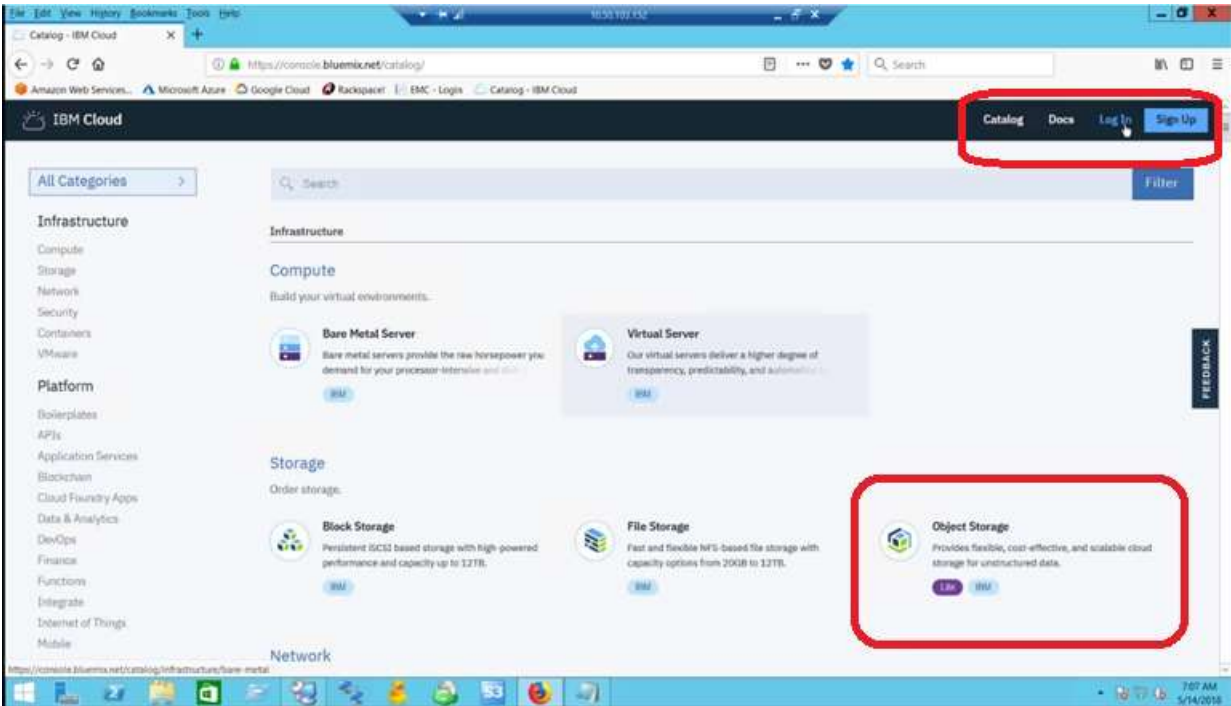
Additional reference material provided by Dell-EMC:

<https://www.emc.com/collateral/white-papers/h15309-veritas-vault-emc-elastic-cloud-storage.pdf>

1.2 IBM Cloud Object Storage

IBM Cloud Object Storage (COS) provide streamer based Vault store partition that allows EV to write files on S3 compliant bucket created on IBM cloud.

- Open IBM COS URL <https://console.bluemix.net/catalog/>
- Select **Object Storage** from **Storage** container



- Give an appropriate name to **Service name** then click on **Create**

Cloud Object Storage

Looking for our infrastructure or Swift Object Storage offerings? [Compare Versions](#)

IBM Cloud Object Storage is a highly scalable cloud storage service, designed for high durability, resiliency and security. Store, manage and access your data via our self-service portal and RESTful APIs. Connect applications directly to Cloud Object Storage use other IBM Cloud Services with your data.

[Like](#) [App](#)

[View Docs](#) [Terms](#)

AUTHOR IBM
PUBLISHED 05/11/2018
TYPE Service

Service name: storage project 1

Select a resource group: Default

Features

- Storage for the IBM Cloud**
IBM Cloud Object Storage provides unstructured data storage for cloud applications. Libraries and SDKs support a common set of S3 API functions for connecting new applications to scalable cloud storage and integrating your data into other services on the IBM Watson and Cloud Platform.
- Encryption management**
- IAM Policies - Bucket level access management**
IBM Identity and Access Management (IAM) integration allows for granular access control at the bucket level using role-based policies.
- Regional and Cross Region resiliency options**

[Need Help?](#) [Contact IBM Cloud Sales](#) [Estimate Monthly Cost](#) [Cost Calculator](#) [Create](#)

- Once the service is ready, either click on **Create Bucket** OR **Create your first bucket** as highlighted below

Storage / storage project 1

Resource Group: Default

Buckets

Q Type the first 3 characters of the bucket name(s)

[Create bucket](#)

NAME	LOCATION	CLASS	ADVANCED
Your Service Instance is empty. Create your first bucket.			

- Give an appropriate **Name** to bucket
- Select the appropriate option for **Resiliency** and **location**. **Storage Class** must be **Standard**
- Click on **Create**

Create a bucket

Name: 📘

storagebucket1

Resiliency: 📘

Cross Region

Location: 📘

us-geo

Standard

Vault

Cold Vault

Flex

Standard

ADVANCED CONFIGURATION

☐ Add Key Protect Keys 📘

Key Protect is not available in the selected location. To enable, choose another location.

Cancel Create

- Select **Endpoint**
- Copy the endpoint which will be used by the Enterprise vault to access bucket created on IBM cloud storage

Getting started

Buckets

Endpoint

Service credentials

Connections

Usage details

Plan

Storage / storage project 1

Resource Group: Default

Endpoints

Service endpoints 📘

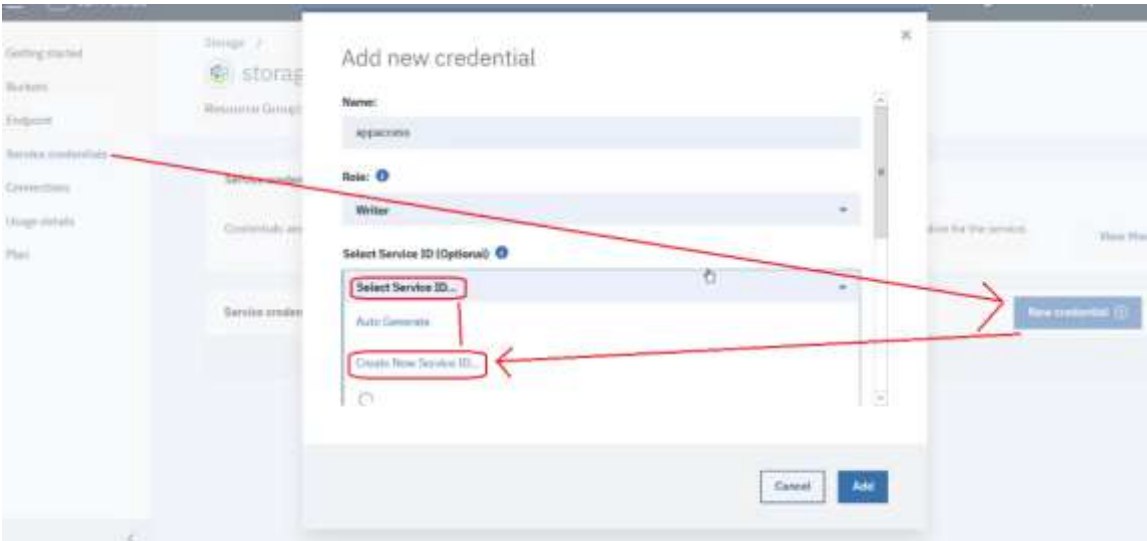
Select resiliency: Cross Region

Select location: us-geo

PUBLIC		PRIVATE	
us-geo:	s3-ap.us-geo.objectstorage.softlayer.net	us-geo:	s3-ap.us-geo.objectstorage.service.networklayer.com
Dallas:	s3-ap.dal-us-geo.objectstorage.softlayer.net	Dallas:	s3-ap.dal-us-geo.objectstorage.service.networklayer.com
Washington:	s3-ap.wdc-us-geo.objectstorage.softlayer.net	Washington:	s3-ap.wdc-us-geo.objectstorage.service.networklayer.com
San Jose:	s3-ap.sjc-us-geo.objectstorage.softlayer.net	San Jose:	s3-ap.sjc-us-geo.objectstorage.service.networklayer.com

- From the **Service Credential** page, click on **New credential**
- Give an appropriate **Name** to credential
- If you already created Service ID, then click **Select Service ID**

- If not, select **Create New Service ID**, (Note: If you are creating it for the first time, then create a new Service ID (do not click on **Add**)



- Give an appropriate name to the **New Service ID**
- In **Add inline configuration parameter** paste `{"HMAC":true}` (this will generate **Access Key ID & Secret key id** for API access.)
- Click on **Add**

Add new credential

Create New Service ID...

New Service ID Name

appaccess

New Service ID Description (Optional)

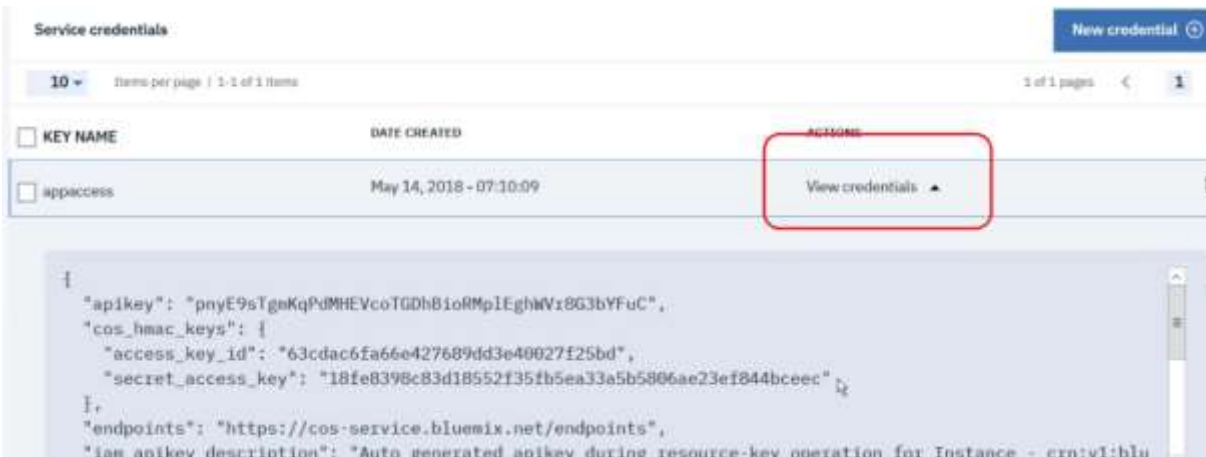
Enter Service ID Description

Add Inline Configuration Parameters (Optional): ⓘ

`{"HMAC":true}`

Cancel Add

- Click on **View Credential** then note down **access_key_id** & **secret_access_key**



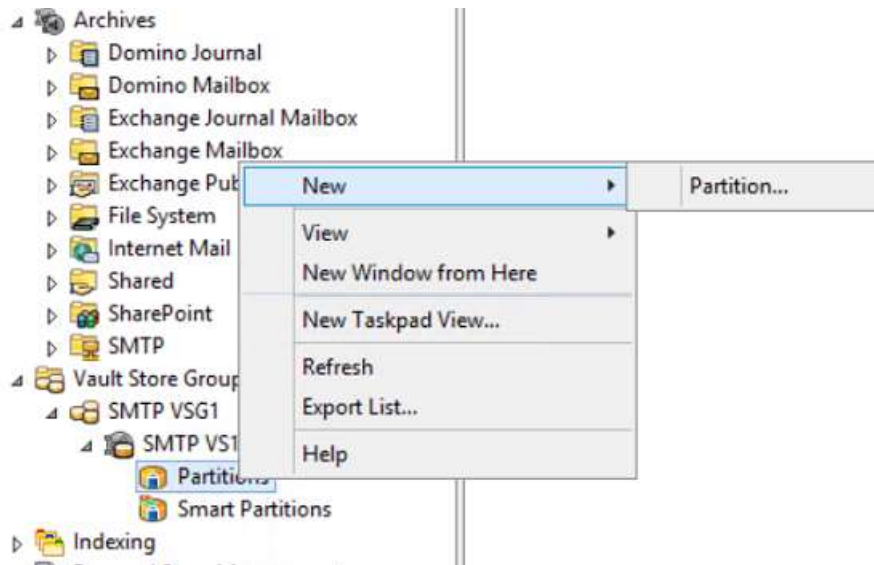
Configuration of Enterprise Vault partition with primary location on IBM COS.

- Install IBM COS streamer (provided by IBM support)

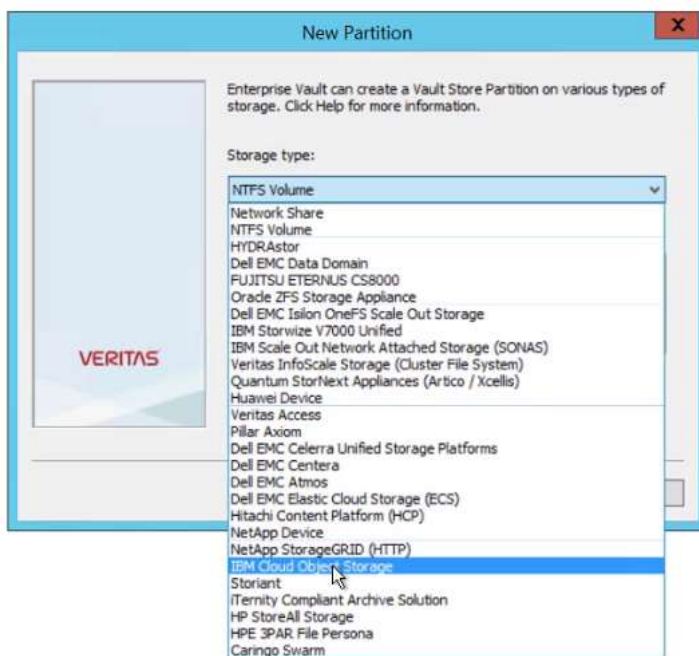


(Follow the installation steps as per the wizard)


- Create a **New Partition**



- Select type as **IBM Cloud object storage**



- Enter configuration details such as **Accessor IP (Endpoint)**, **vault name (Bucket)**, **authentication type (AWS2 OR AWS4)**, **access key**, **secrete key**, **disable peer authentication (true)**. Keep **server authentication** as blank



- Click on **Test**. “IBM Cloud Object Storage connection test succeeded” should appear if there are no issues encountered



Once partition setup completes successfully you can start archiving on IBM COS partition.

Additional reference material provided by IBM:

<https://www-01.ibm.com/common/ssi/cgi-bin/ssialias?htmlfid=KUO12405USEN>

2. Secondary Storage (Migrator)

Where vault store partitions are held on non-WORM devices, you can configure and schedule the collection and migration of the files that are stored in the partition.

Collection involves collecting multiple small files into much larger collection files (.cab files). Collection may give you a significant improvement in backup times.

Migration involves moving the collection files onto longer term storage devices. For example, you may want to migrate older collections to cheaper, slower storage. If you choose to use collection files you can configure the collection criteria, and optionally provide details of how and when to migrate the collection files to secondary storage.

Following example will show how to setup AWS S3, Azure, Rackspace & google cloud storage that can be used as a secondary storage.

2.1 Common Terminology

Storage Server name

The name of the server used by Cloud provider.

Access Key ID

The secure access Key ID name that cloud vendor provides.

Secret key ID

The account shared secret that cloud vendor provides.

Bucket (Container in Azure & google cloud)

The name of the bucket created on cloud storage. The bucket name must be unique across all existing bucket names in the cloud provider. 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. Refer to the Storage Service provider's documentation for bucket naming requirements and guidelines.

Bucket Region

The geographical location where the bucket is created.

Write buffer size

The buffer size, in megabytes, Enterprise Vault uses for data uploads. Ensure this value is greater than the Maximum Collection File Size setting on the Collections tab of the vault store partitions. Set this option to zero (0) to disable the use of buffers.

Read buffer size

The buffer size, in megabytes, Enterprise Vault uses for data downloads.

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 cloud.

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.

User wait timeout

Specify the number of seconds after a retrieval request, after which the user is presented with the 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.

System wait timeout

If an attempt to retrieve an archived item from the Amazon 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.

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.

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 un-migrated.

2.2 Amazon Web Service S3 (Simple Storage Service)

The Enterprise Vault Amazon S3 storage migrator lets you migrate archived data to and retrieve it from Amazon Simple Storage Service (Amazon S3). You can use Amazon S3 as a secondary storage location in the cloud to store infrequently accessed data. The Enterprise Vault Amazon S3 storage migrator is installed as part of the Enterprise Vault 10.0.1 or later installation. It moves CAB files, created by the Enterprise Vault file collection software, to the Amazon S3 storage.

Following steps shows the configuration of S3 Bucket using AWS console and then use it as migrator for vault store partition.


- Sign in with root a/c credential to AWS console
<https://console.aws.amazon.com/console/home>



- Select **IAM** under **Security, identity and compliance** container



- Click on **Add user** under **Users**



- Give appropriate **User name**
- Select **Access Type**

Please note: *selection of both access types is not recommended for production use.*

Set user details

User name*

[+ Add another user](#)

Select how these users will access AWS. Access keys and autogenerated passwords are provided in the last step. [Learn more](#)

^a Required

- Add permissions to evuser

Use IAM policies to grant permissions. You can assign an existing policy or create a new one.

https://www.veritas.com/support/en_US/article.100032260

- Review the setting and then click on **Create user**

Review

Review your choices. After you create the user, you can view and download the autogenerated password and access key.

User details

User name	AppUser
AWS access type	Programmatic access and AWS Management Console access
Console password type	Custom
Require password reset	No

Permissions summary

The following policies will be attached to the user shown above.

Type	Name
Managed policy	AmazonS3FullAccess

Cancel Previous **Create user**

- Note down details of the **user name**, **access Key ID**, **Secret Access Key** and **Sign-in URL**. You can additionally **download** this information in CSV format

Add user

1

2

3

4

Success

You successfully created the users shown below. You can view and download user security credentials. You can also email users instructions for signing in to the AWS Management Console. This is the last time these credentials will be available to download. However, you can create new credentials at any time.

Users with AWS Management Console access can sign in at: <https://signin.aws.amazon.com/console>

Download .csv

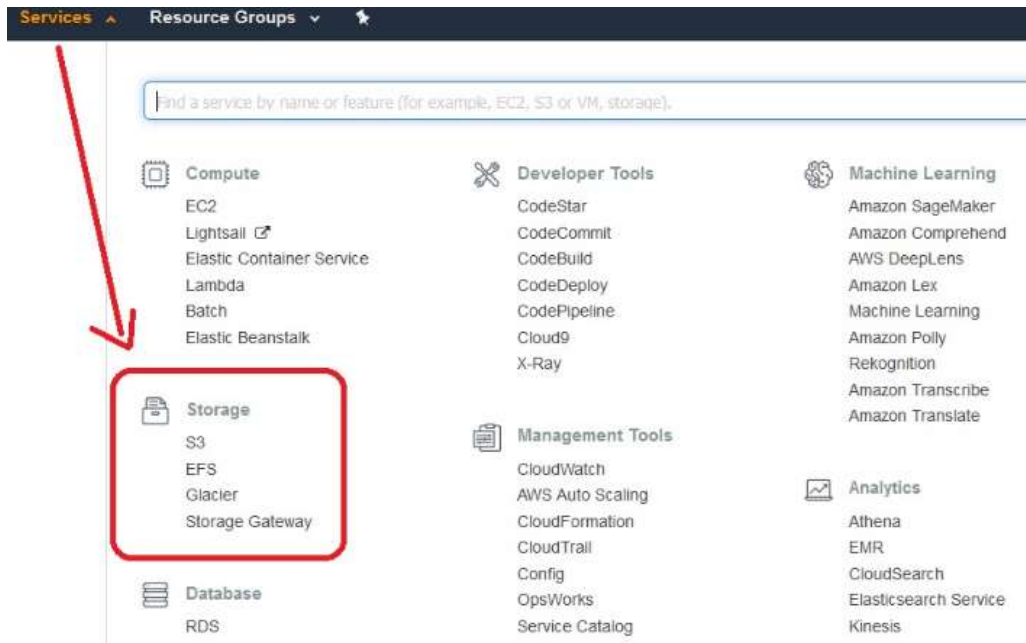
User	Access key ID	Secret access key
user1	AKIAJ2NWMQWQWTS7SA	wjAhtGacw9Fy2y3Bkuz+uG8RT8Cqm ar1QY view

Close

- Go to S3 Service page. Select **S3** from **Storage** section

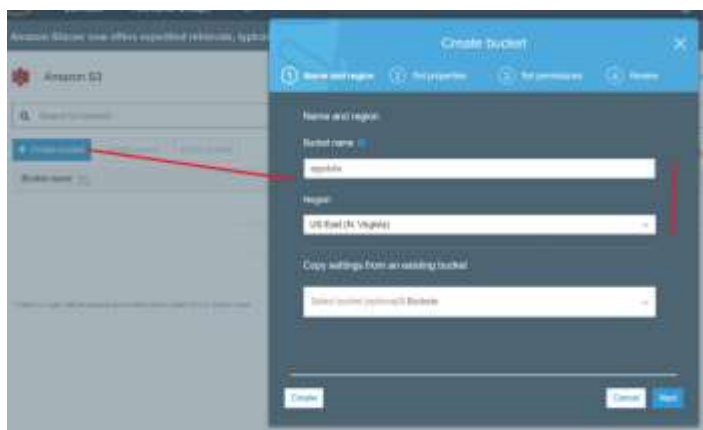
Cloud Storage for Enterprise Vault

22



- Click on **Create Bucket**
- Give an appropriate name to **Bucket**
- Select **Region** (note the bucket name should be unique in your AWS infrastructure)
- Click **Create**

Please note remaining criteria such as version, permission and website related items have been skipped as they are not required for EV configuration.



Configuring AWS S3 storage as a Migrator (Secondary) storage for Enterprise Vault

- Select **Amazon Simple Storage Service** as migrator type



- Enter **Access Key, Secret key, Bucket Name, Bucket Region information** as specified during AWS S3 bucket creation
- Click on **Test** (Migrator Configuration Test succeeded message appears)



Once AWS S3 migrator configured successfully. Enterprise vault collection process will collect then migrate EV files from primary partition to secondary AWS S3 storage.

More details can be found at:

Enterprise Vault™ Migrating Data Using the Amazon S3 Storage Migrator

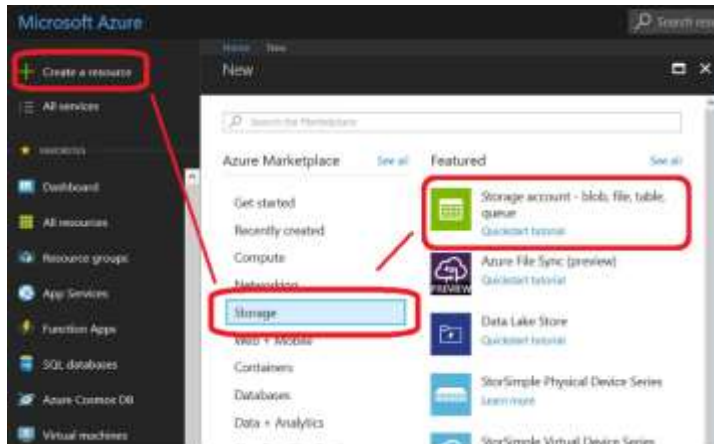
https://www.veritas.com/content/support/en_US/doc/67282638-129299793-0/v66079963-129299793

2.3 Microsoft Azure Data Blob storage

The Microsoft Azure Blob Storage migrator lets you migrate archived data to and retrieve it from Azure Blob storage. You can use Azure Blob storage as a secondary storage location to store archived data in the cloud. The Microsoft Azure Blob Storage migrator is installed automatically when you install Enterprise Vault 12.2 or later.

Following steps shows the configuration of cloud container using Azure portal and then use it as migrator for vault store partition.

- Open Azure web portal - <https://portal.azure.com/>
- Create a blob storage resource



- Give appropriate name for **Storage account**, remember this should be unique in AWS infrastructure, and select storage kind as **Blob Storage**
- Select **New** in **Resource group** section. (If you have already created the resource group you wish to use, select it.)
- Select **Pin to dashboard** to easily access it directly from console

Create storage account

The cost of your storage account depends on the usage and the options you choose below. [Learn more](#)

* Name **appazurestorage1** ✓
 .core.windows.net

Deployment model **Resource manager** Classic

Account kind **Blob storage**

* Location **East US**

Replication **Read-access geo-redundant storage (R...)**

Performance **Standard** Premium

Access tier (default) **Cool** Hot

* Secure transfer required **Disabled** Enabled

* Subscription **Free Trial**

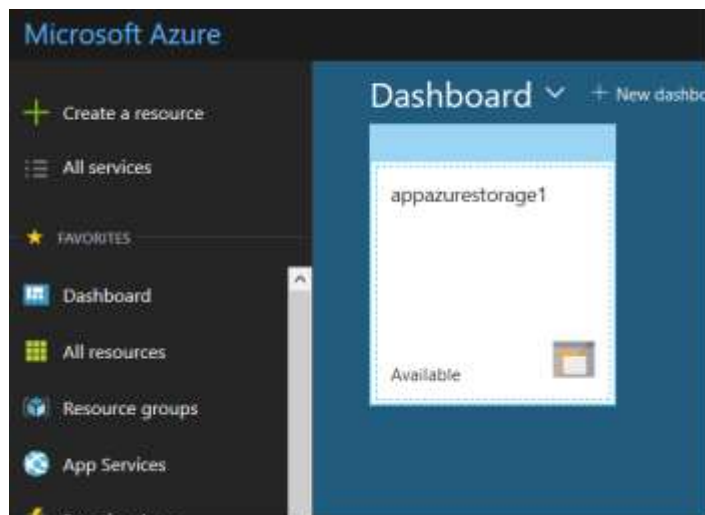
* Resource group **MyAppStorage** ✓
☒ Create new ☐ Use existing

Virtual networks **Disabled** Enabled

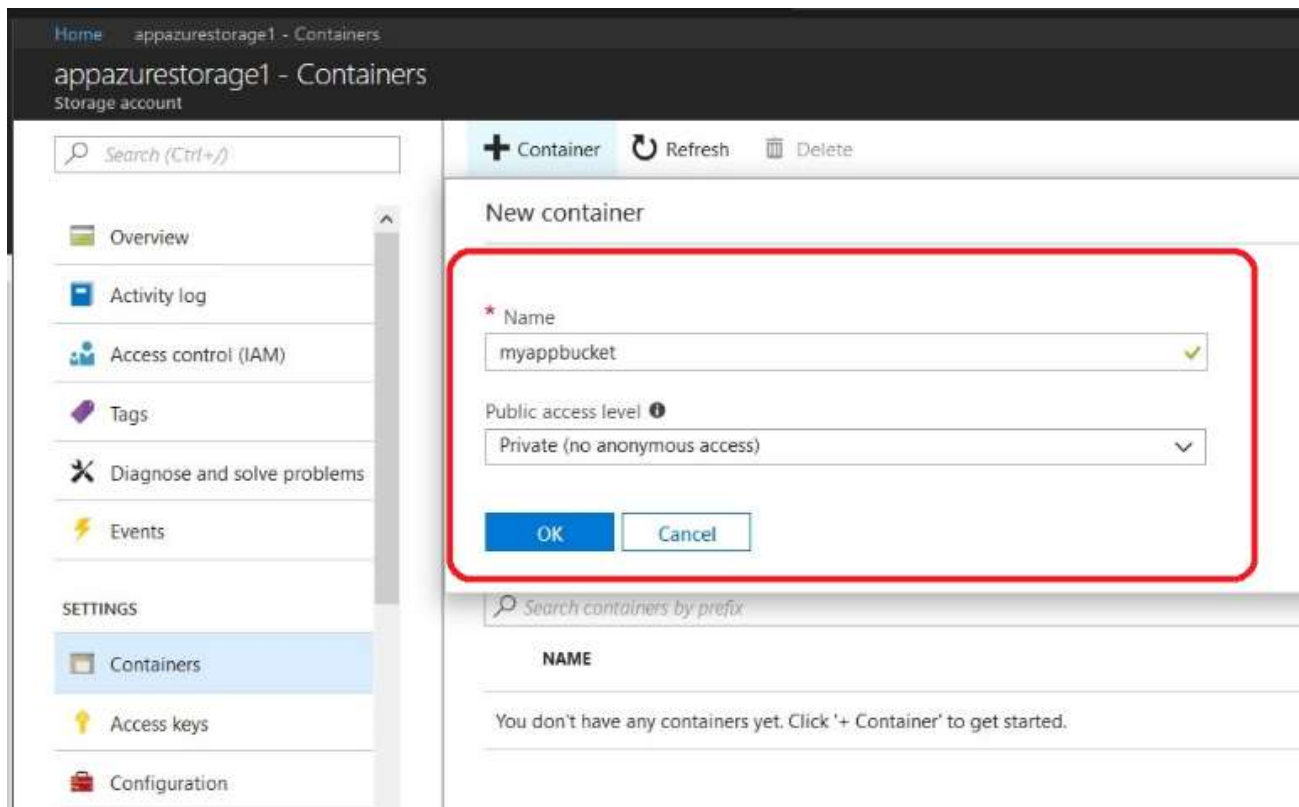
☒ Pin to dashboard

Create [Automation options](#)

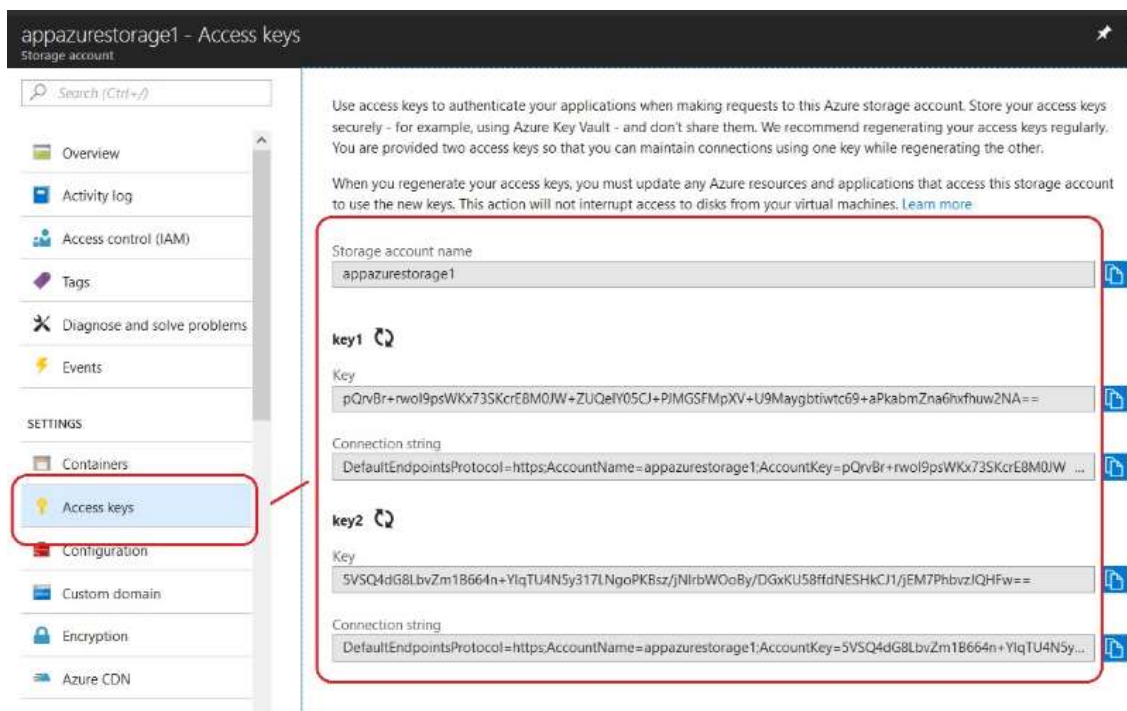
- From the dashboard, click on storage account **appazurestorage1** which was created in an earlier step



- Create a new **Container**, keep the access level as **Private(no anonymous access)**

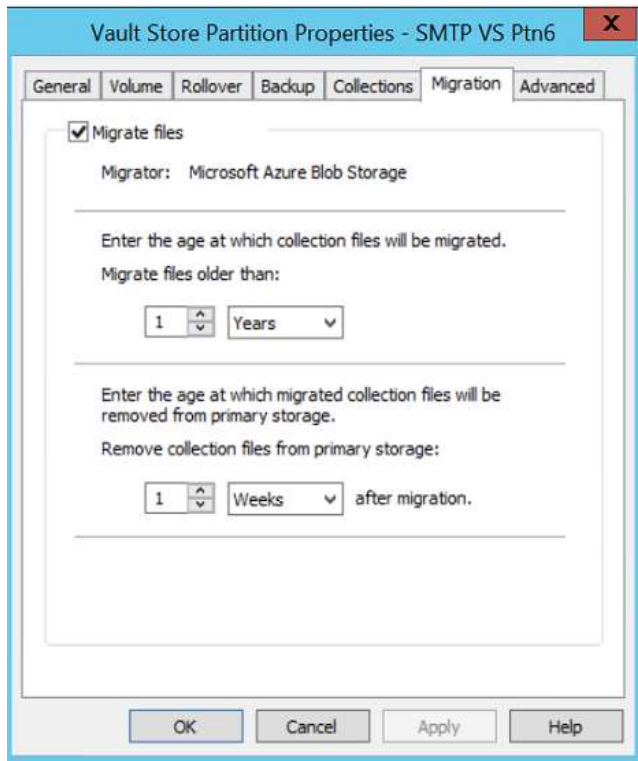


- Click on **Access key** tab and note the **storage account name & key**

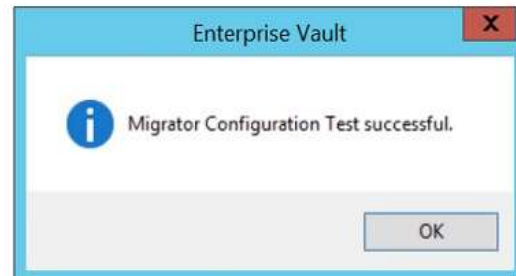
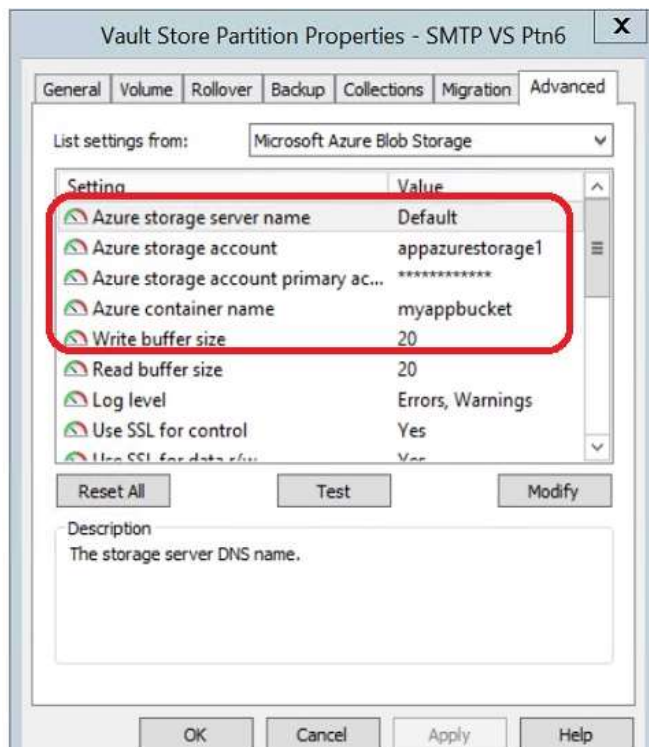


Configuring Azure Data blob as secondary storage for Enterprise vault.

- In the properties of vault store partition, select **Migration** tab
- Select **Migrate files**
- Select **Microsoft Azure blob storage**



- In **Advanced** tab, specify the **storage account, primary access key & container name (bucket)** information
- Click on **Test** connection. "*Migrator Configuration Test successful*" message appears if there are no errors



Once Azure migrator has been configured successfully, the Enterprise Vault collection process will collect then migrate EV files from primary partition to secondary AWS S3 storage.

More information can be found at:

Enterprise Vault™ Migrating Data Using the Microsoft Azure Blob Storage Migrator
https://www.veritas.com/support/en_US/doc/125282611-125282615-0/index

2.4 Google Cloud Storage

The Google Cloud Storage migrator lets you migrate archived data to and retrieve it from Google Cloud Storage. You can use Google Cloud Storage as a secondary storage location to store archived data in the cloud. The Google Cloud Storage migrator is installed automatically when you install Enterprise Vault 12.2 or later.

Following steps shows the configuration of bucket using Google portal and then use it as migrator for vault store partition.

- Open <https://cloud.google.com/>
- Open **console** and create a new **Project**

New Project

Project Name *

mystorageproject 

Project ID: mystorageproject-202809. It cannot be changed later. [EDIT](#)

Location *

 No organisation [BROWSE](#)

Parent organisation or folder

CREATE

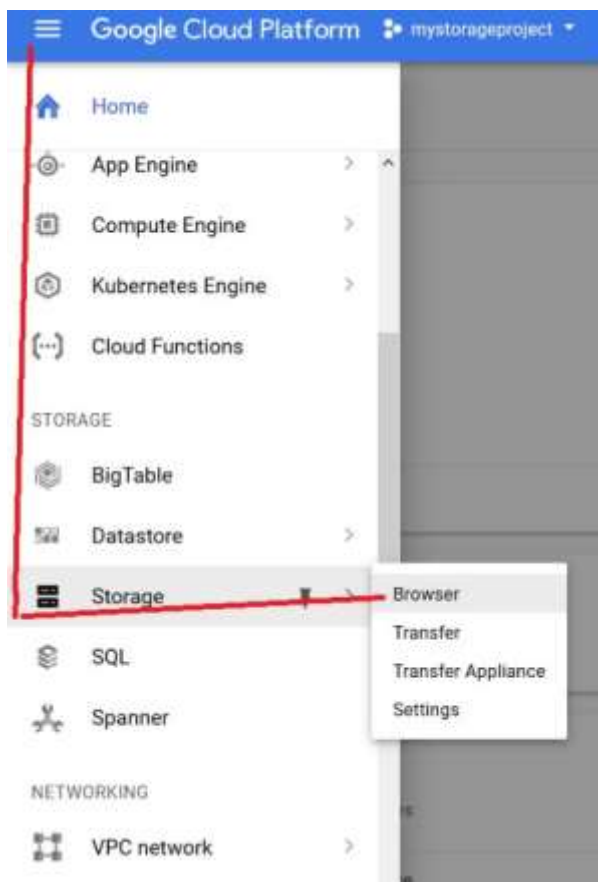
CANCEL

- Select the correct project if you have access to multiple projects

Select



- Select storage from **Product and Services** menu



- Click on **Create Bucket** and populate the **Name**, **storage class** and **location** fields

Cloud Storage
Buckets

Cloud Storage lets you store unstructured objects in containers called buckets. You can serve static data directly from Cloud Storage, or you can use it to store data for other Google Cloud Platform services.

Create bucket

or

Take the quickstart tutorial

[←](#) Create a bucket

Name [?]
Must be unique across Cloud Storage. If you're [serving website content](#), enter the website domain as the name.

Default storage class [?]
[Compare storage classes](#)
☒ Multi-Regional
☐ Regional
☐ Nearline
☐ Coldline

Location

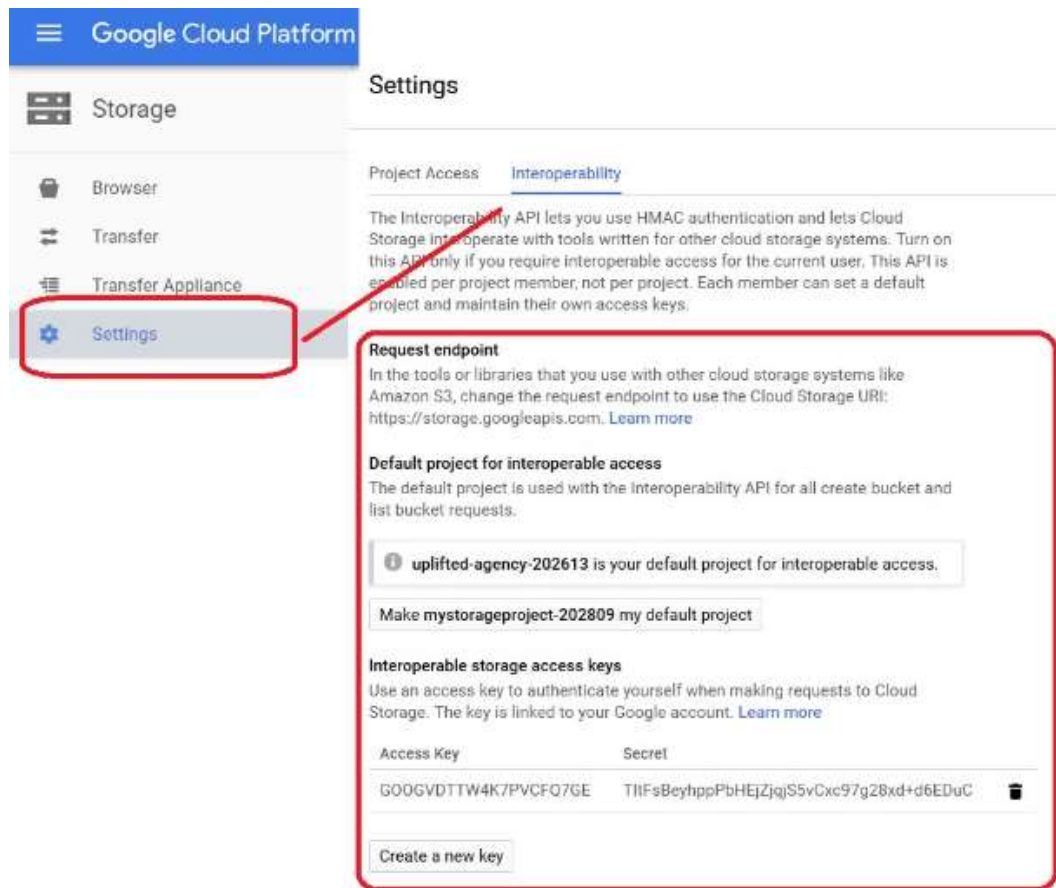
Storage cost	Retrieval cost	Class A operations [?]	Class B operations [?]
\$0.026 per GB-month	Free	\$0.005 per 1,000 ops	\$0.0004 per 1,000 ops

[Show advanced settings](#)

Create

Cancel

- Click on **Settings**
- Select **Interoperability** (enable if it is not already)
- Note down **cloud storage URL**, **access key** and **secret key** details (If secret key is not present, click on **Create a new Key**)

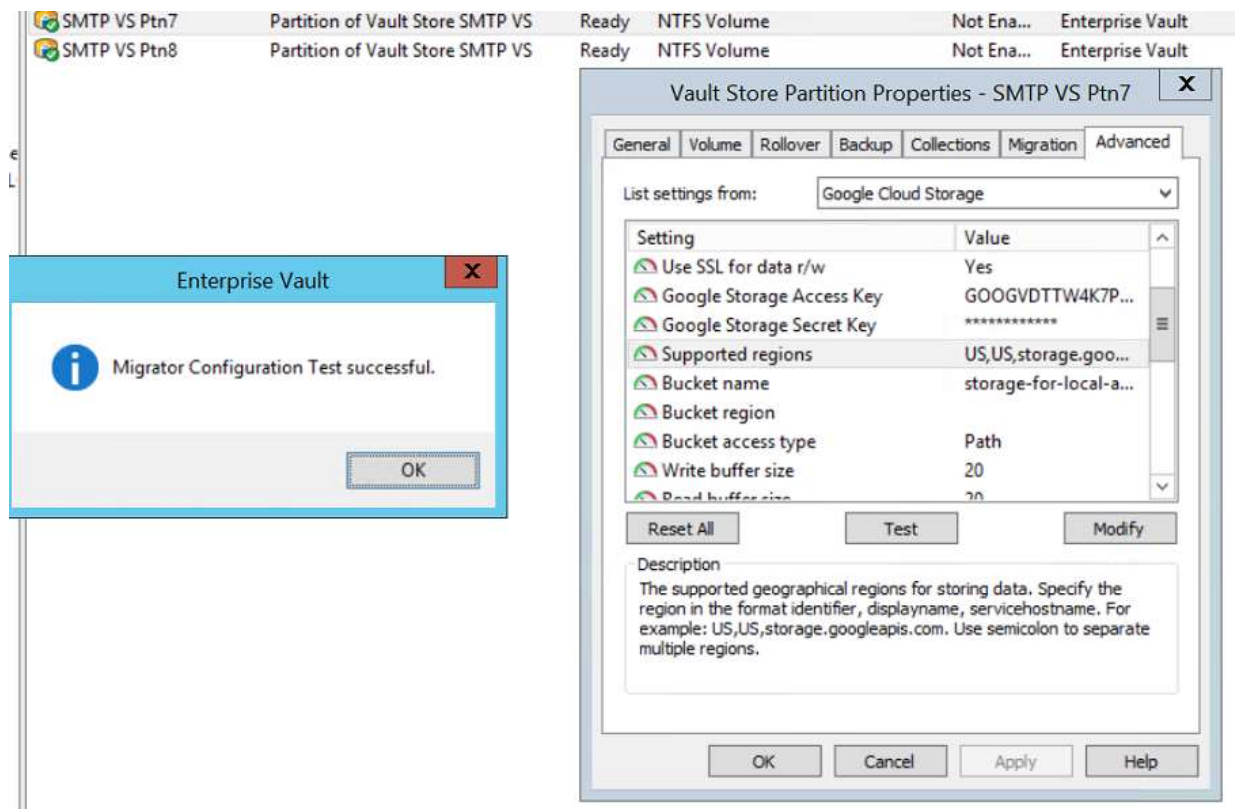


Configuration of Google cloud as Migrator for Enterprise Vault partition.

- In the properties of the EV partition, select **Migration** tab
- Select **Google Cloud storage** as migrator



- Select **Advanced** tab, based on configuration
- Specify **Storage server name** (storage.googleapis.com), **access key**, **secret key**, **Support region** (US,US,storage.googleapis.com)
- Specify **Bucket Name**
- Click on **Test** connection, “*Migrator Configuration Test successful*” message appears if there are no errors



Once the Google migrator configured successfully, the Enterprise Vault collection process will collect and then migrate EV files from the primary partition to secondary Google storage.

Additional reference material provided by VERITAS:

How to configure Enterprise Vault 12.1 with Google Cloud Storage

https://www.veritas.com/support/en_US/article.100039053

https://www.veritas.com/content/support/en_US/doc/125452905-125453076-0/v125452521-125453076

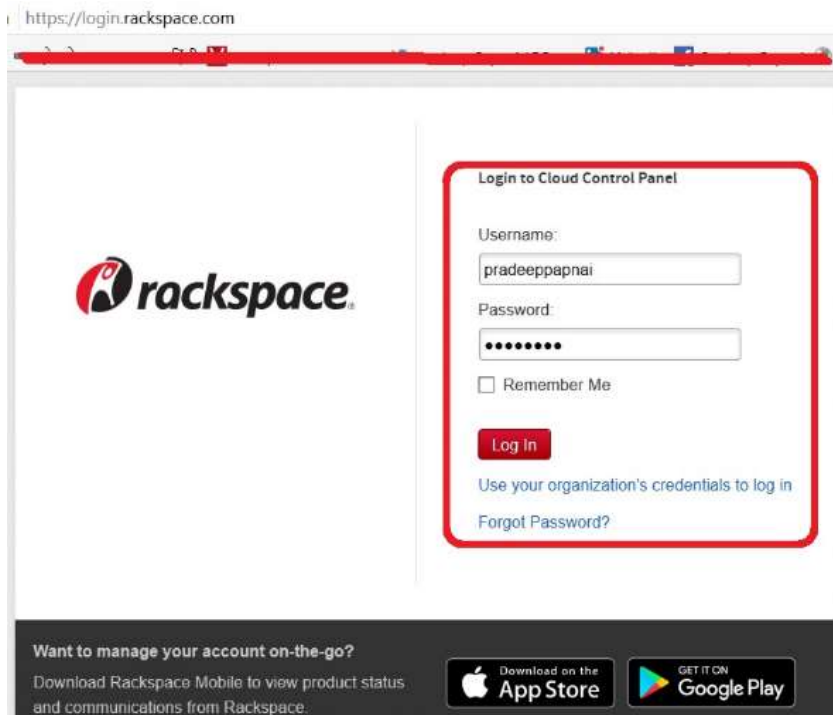
2.5 Rackspace Cloud Files

The Enterprise Vault Rackspace Cloud Files storage migrator lets you migrate archived data to and retrieve it from Rackspace Cloud Files storage. You can use Rackspace Cloud Files as a secondary storage location in the cloud to store infrequently accessed data.

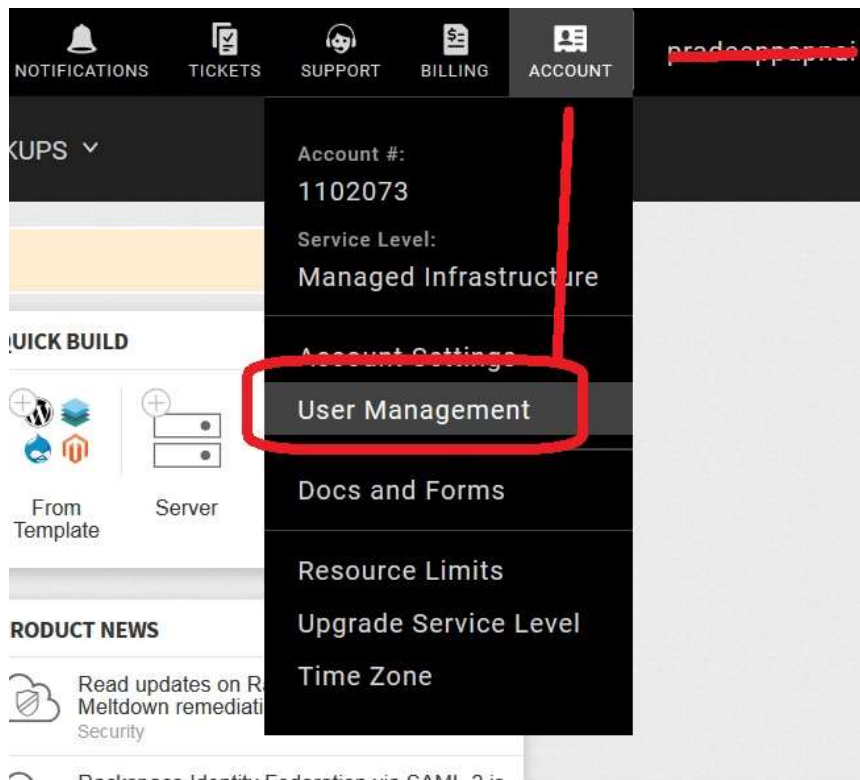
The Enterprise Vault Rackspace Cloud Files storage migrator is installed as part of Enterprise Vault 10.0.1 or later. It moves CAB files, created by the Enterprise Vault file collection software, to the Rackspace Cloud Files storage.

Following steps shows the configuration of Rackspace and then use it as migrator for vault store partition.

- Signup to Rackspace cloud <http://www.rackspacecloud.com/signup>
- Go to Rackspace control panel <https://login.rackspace.com/>
- Provide the root user account you configured during signup process



- Create a new user **Account** for API access
- Go to **User Management** from **Account** tab



- Click on **Create user**, this console will give you list of all users created so far

User Management



- Give user details, contact type must be **Technical**
- Select appropriate permission on **Rackspace cloud**

User Information

First Name

Last Name

Username
The username cannot be changed.

Password
At least 8 characters with 1 uppercase, 1 lowercase, and 1 number

Confirm Password

Email Address

Phone Number

Contact Type: ☒ Technical ☐ Administrative

Product Permissions

Fanatical Support for AWS Rackspace Cloud

Global Permissions

- Product Access
- ☒ Full Access (All Products) ⓘ
 - ☐ Read Only Access (All Products) ⓘ
 - ☐ No Access (No Product Access)
 - ☐ Custom (Per Product Access) ⓘ

Create User Cancel

- Once user is created successfully go to the properties of user account and copy the **Rackspace API key**

Security Settings

Password

Multi-Factor Authentication Disabled • [Enable...](#)

Secret Question What is the location of a dream vacation? • [Reset...](#)

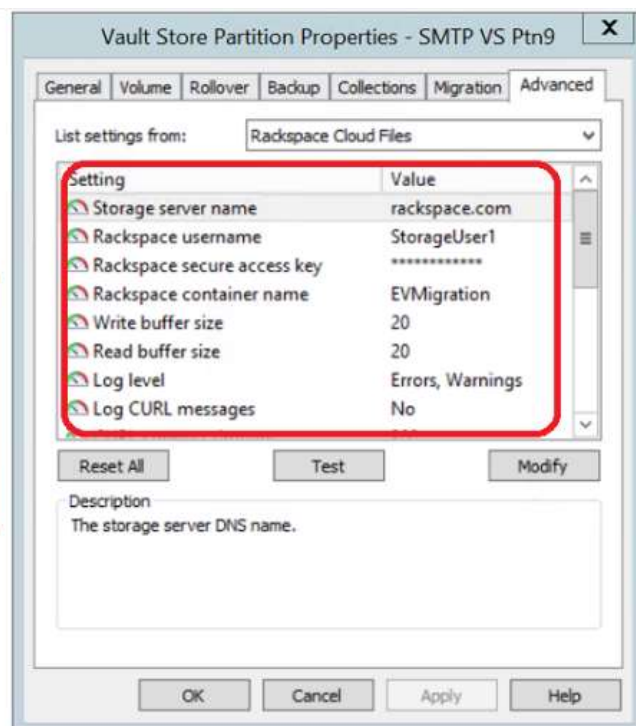
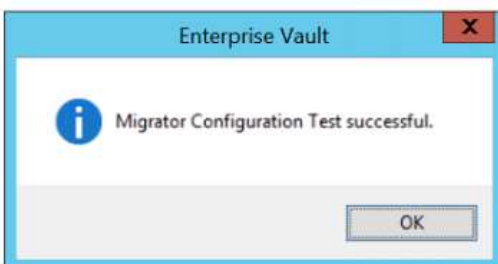
Rackspace API Key 368fba0be46d41d199fce07ccc2252ac [Hide](#) • [Reset](#)

Configuring Rackspace cloud file storage as secondary (migrator) for Enterprise vault

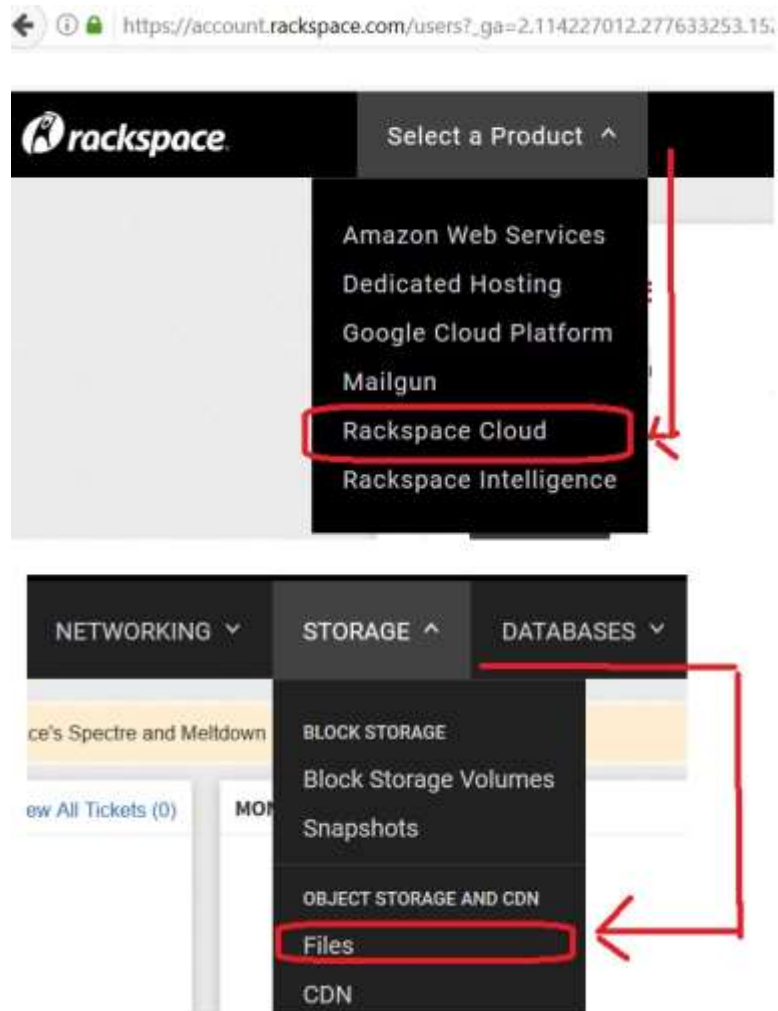
- In the vault store partition property, select **Rackspace cloud file storage** as migrator



- In **Advanced** tab, provide details such as **user name**, **secure access key (API key)**, **container name**
- Click on **Test** (**Please note:** *Container should not be created from Rackspace control pane. EV migrator will create it automatically. It should be a unique name in Rackspace Cloud.*)







- Once configured, go to the **control panel**
- Select **Rackspace cloud** from the product list
- Select **Files** from **Storage** list



- After clicking on **Test** connection, EV process creates a bucket automatically in Rackspace cloud

Cloud Files / Containers

Create Container		Northern Virginia (IAD) ▾		
	Container Name	Region	Files	Size
 	EVMigration	Northern Virginia (IAD)	0	0 B
 	StorageAccess	Northern Virginia (IAD)	0	0 B
Show More				

Once the Rackspace migrator has been configured successfully, the Enterprise Vault collection process will collect then migrate EV files from primary partition to secondary Rackspace cloud storage.

More information can be found at:

Enterprise Vault™ Migrating Data Using the Rackspace Cloud Files Migrator
https://www.veritas.com/support/en_US/doc/72538174-72538188-0/index

Troubleshooting

The following 3rd party utilities can be used to verify connectivity and permissions issue with a bucket (or container) created on cloud storage.

- CloudBerry
- CyberDuck
- S3Browser
- TNTDrive

Additionally the cloud provider may also provide command line utilities that can be used to isolate configuration issues.

The following processes can be enabled for Enterprise Vault debugging using Dtrace utility.

- **StorageArchive** Responsible for writing EV objects to storage when primary partition is configured on cloud.
- **EVStgOfflineOpns.exe** Retrieve/Recall request to secondary storage.
- **StorageFileWatch.exe** Responsible for collection, migration to cloud, deletion of objects on cloud and validation of safety copies.
- **StorageManagement.exe** Verification of library and connection test.

Additional information about the Dtrace utility:

How to run Dtrace to help diagnose issues with Enterprise Vault

https://www.veritas.com/support/en_US/article.100038975

List of Dtrace processes and descriptions of each process responsibility

https://www.veritas.com/support/en_US/article.100001741

The Veritas logo is displayed in a bold, red, sans-serif font. The word "VERITAS" is in all caps. A small trademark symbol (TM) is located at the top right of the letter "S".

VERITAS™

© 2017 Veritas Technologies LLC. All rights reserved.
Veritas and the Veritas Logo are trademarks or registered
trademarks of Veritas Technologies LLC or its affiliates in
the U.S. and other countries. Other names may be
trademarks of their respective owners.