

Understanding AWS Storage Gateway-VTL

<https://campus.barracuda.com/doc/58490951/>

The AWS Storage Gateway is a service connecting an on-premises software appliance with cloud-based storage to provide seamless and secure integration between an organization's on-premises IT environment and the AWS storage infrastructure. The service enables you to securely store data to the AWS cloud for scalable and cost-effective storage. The AWS Storage Gateway supports industry-standard storage protocols that work with your existing applications. It provides low-latency performance by maintaining frequently accessed data on-premises while encrypting and storing all of your data in Amazon Simple Storage Service (Amazon S3) or Amazon Glacier.

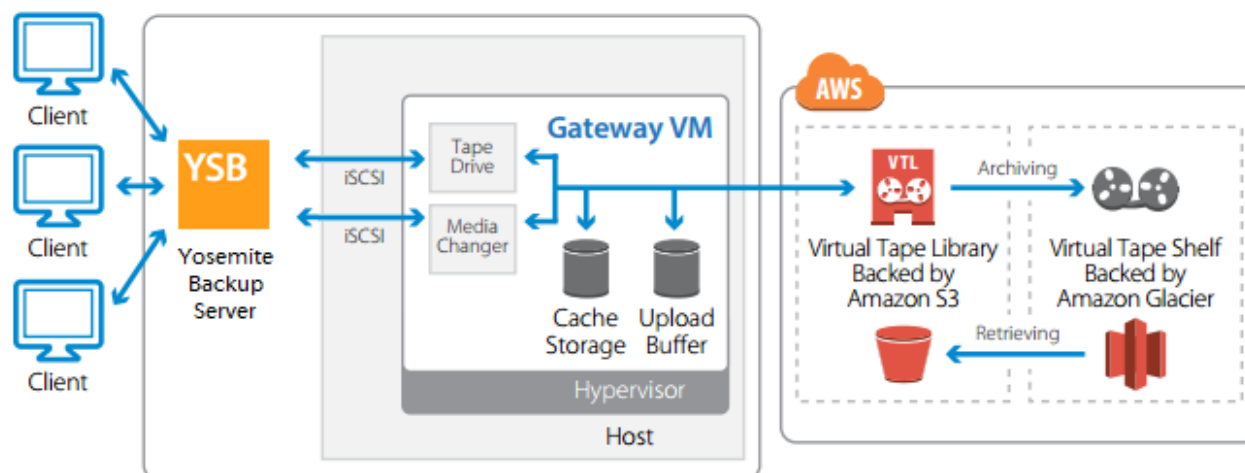
There are three types of AWS Storage Gateway; Yosemite Server Backup supports the Storage Gateway-Virtual Tape Library (VTL). With Gateway-VTL you have a limitless collection of virtual tapes. Each virtual tape can be stored in a Virtual Tape Library backed by Amazon S3 or a Virtual Tape Shelf (VTS) backed by Amazon Glacier. The Virtual Tape Library exposes an industry standard iSCSI interface which provides Yosemite Server Backup online access to the virtual tapes. When you no longer require immediate or frequent access to data contained on a virtual tape, use Yosemite Server Backup to move it from its Virtual Tape Library to your VTS to further reduce your storage costs.

AWS Storage Gateway-VTL with Yosemite Server Backup

AWS Storage Gateway-VTL with Yosemite Server Backup is designed for customers want to offload certain data sets from their own archival devices to AWS Glacier.

Figure 1 illustrates a typical environment using Yosemite Server Backup and AWS Storage Gateway-VTL.

Figure 1. Yosemite Server Backup and AWS Storage Gateway-VTL Architecture.



Infrastructure

The AWS Storage Gateway's software appliance is available for download as a virtual machine (VM) image that you install on a host in your datacenter. Once you install your gateway and associated it with your AWS Account through Amazon's activation process, you can use the AWS Management Console to create either Gateway-Cached or Gateway-Stored storage volumes or Gateway-VTL virtual tape libraries that you can mount as iSCSI devices using your on-premises applications.

Gateway-VTL enables you to seamlessly replace your physical tape infrastructure with a virtual tape infrastructure. Each Gateway-VTL presents Yosemite Server Backup with an industry-standard iSCSI-based VTL consisting of a virtual media changer and tape drives. Create virtual tapes in your VTL using the AWS Management Console. Each VTL can hold up to 1,500 virtual tapes with a maximum aggregate capacity of 1 PB. Virtual tapes are discovered by Yosemite Server Backup using its standard media inventory procedure. Virtual tapes in your VTL are available for immediate access and are backed by Amazon S3. Yosemite Server Backup can read data from or write data to virtual tapes by mounting them to virtual tape drives using the virtual media changer.

Hardware and Software Requirements

The AWS Storage Gateway VM must be either installed on a host in your datacenter running supported versions of VMware ESXi or Microsoft Hyper-V, or as an AMI running on an EC2 instance. The gateway VM must be deployed with a minimum set of hardware resources. The AWS Storage Gateway currently supports Microsoft Windows, Red Hat Enterprise Linux, and VMware ESXi, iSCSI initiators. For additional information, refer to the [AWS Storage Gateway User Guide](#).

Long Term Retention

For cost-effective long term retention of data requiring infrequent access, you can use Yosemite Server Backup to move virtual tapes from one or more of your VTLs to your VTS backed by Amazon Glacier. Your VTS is automatically created when you activate your first Gateway-VTL. Store virtual tapes that need to be accessed frequently in a VTL. Data that does not need to be retrieved frequently can be archived to your VTS. Access to virtual tapes in your VTL is immediate while virtual tapes in your VTS must be retrieved and loaded into a VTL before it can be accessed. You can retrieve virtual tapes from your VTS using the AWS Management Console. Virtual tapes retrieved from your VTS take approximately 24 hours before they are automatically loaded into your VTL.

Restore Data

You can retrieve a virtual tape using the AWS Management Console. First select the virtual tape, then choose the VTL into which you want to load the virtual tape. It takes about 24 hours for the retrieved virtual tape to be available in the selected VTL. Once the virtual tape is available in the VTL, use the Yosemite Server Backup to make use of the virtual tape to restore data.

Performance

As the Gateway-VTL sits between your application, AWS, and underlying on-premises storage, the performance you experience is dependent upon a number of factors including the speed and configuration of your underlying local disks, provisioned host processor and memory, network bandwidth between your iSCSI initiator and gateway VM, the amount of local storage allocated to the gateway VM, and the bandwidth between the gateway VM and AWS. For Gateway-VTL, to provide predictable write performance to Yosemite Server Backup, it is important that you allocate enough local cache disk storage to durably buffer data uploaded to AWS. For additional information, refer to the [AWS Storage Gateway User Guide](#) for guidance on sizing cache. Amazon recommends configuring your disks in a redundant array of independent disks (RAID) configuration to improve performance and to protect against disk failures.

Security

The AWS Storage Gateway encrypts all data in-transit to and from AWS via SSL. All volume and snapshot data stored in AWS using Gateway-Stored Volumes and Gateway-Cached Volumes, and all virtual tape data stored in AWS using Gateway-VTL, is encrypted at-rest using Advanced Encryption

Standard (AES) 256, a secure symmetric-key encryption standard using 256-bit encryption keys. Additionally, Yosemite Server Backup has optional built-in encryption and compression features.

Storage

Table 1. Storage.

Feature	Description
Data Reduction	The AWS Storage Gateway performs compression of data in-transit and at-rest, which can reduce both data transfer and storage charges. All data transfer between the AWS Storage Gateway VM and AWS, and all data stored in AWS, is compressed. In addition, the AWS Storage Gateway VM only uploads data that has changed, minimizing the amount of transferred data.
Bandwidth Throttling	Use the AWS Management Console to restrict the bandwidth between your gateway and AWS based on the selected rate. You can specify individual rates for inbound and outbound traffic.
Data Storage on a Virtual Tape	When creating a virtual tape, select from the following sizes: 100 GB, 200 GB, 400 GB, 800 GB, 1.5 TB, and 2.5 TB
Data Storage in VTL	Each VTL can store up to 1500 virtual tapes with a maximum aggregate capacity of 1 PB.
Data Storage on a VTS	There is no limit to the amount of data you can store on a VTS.

Access Data

The virtual tape containing your data must be stored in a VTL before you can access it, however, access to the virtual tapes is instantaneous. If the virtual tape containing your data is in your VTS, you must first retrieve the virtual tape from your VTS. You can retrieve the virtual tape using the AWS Management Console. First select the virtual tape, then choose the VTL into which you want to load the virtual tape. It takes about 24 hours for the retrieved virtual tape to be available in the selected Virtual Tape Library. Once the virtual tape is available in the Virtual Tape Library, you can use the Yosemite Server Backup to make use of the virtual tape to restore data.

Get Started with AWS Storage Gateway

To sign up for the AWS Storage Gateway, go to the AWS Storage Gateway detail page, and create an account or log in with an existing account. Once logged in, you can begin the AWS Storage Gateway

sign-up process. After you sign up, you can begin setting up and activating your gateway by visiting the AWS Management Console. For details, refer to the Amazon Web Services article [Getting Started with AWS Storage Gateway](#).

Billing

Note: All fees quoted here are controlled by Amazon and were current at the time of publication. For additional storage pricing, refer to the [AWS Storage Gateway Pricing](#).

Barracuda Networks does not charge a fee to use the Amazon Storage Gateway. Amazon charges for the Storage Gateway, the cost to store data in their cloud, along with a \$125 fee per activated virtual gateway per month.

- You are billed a monthly fee for each of gateway; this fee is prorated daily. Billing for a gateway begins upon activation and continues until you delete the gateway from the AWS Management Console or via the API.
- You are billed for the virtual tape data you store in AWS. This fee is prorated daily. You are only billed for the portion of virtual tape capacity that you use, not for the size of the virtual tape you create.
- You are billed for Internet data transfer for each GB downloaded from AWS to your gateway. All data transfer for uploading to AWS is free.
- You are charged for the peak concurrent data retrieval in a month.

Table 2. Cost to Store Data in AWS.

VTs Storage	\$0.007 per GB per month of data stored
VTL Storage	\$0.030 per GB per month of data stored
VTs Retrieval	\$0.30 per GB ⁽¹⁾
Virtual Tape Deletion	Free ⁽²⁾

Notes:

⁽¹⁾ Pay only for the highest peak concurrent data retrieval in a month.

⁽²⁾ There are no charges for deleting virtual tapes from your VTL and VTsf. However, there is a prorated charge of \$0.021 per gigabyte for virtual tapes deleted prior to 90 days from your VTs.

Concurrent Data Retrieval Charges

Concurrent data retrieval is calculated when a retrieval from VTs is initiated. The concurrent data retrieval is the sum of the size of all virtual tape retrievals initiated concurrently or concurrently in progress. If the concurrent data retrieval is more than any prior concurrent data retrieval in the month you are charged for the difference between this new monthly peak and the prior monthly peak. If the concurrent data retrieval is less than or equal to any prior concurrent data retrieval in the month, the retrieval is free.

- **Example 1** – You initiate retrieval of a 100 GB virtual tape from your VTS in US East (Northern Virginia) Region. This is the first retrieval of the month. The prior peak concurrent data retrieval for the month was 0 GB. When the retrieval is initiated the concurrent data retrieved is 100 GB. 100 GB is your new monthly peak concurrent data retrieval. The charge for the retrieval is the difference between the prior and the new peak concurrent data retrieval, that is, $(100 \text{ GB} - 0 \text{ GB}) \times \$0.30/\text{GB} = \$30.00$.
- **Example 2** – You initiate retrieval of one virtual tape containing 500 GB of data from your VTS in US East (Northern Virginia) Region. This is the first retrieval of the month. The prior peak concurrent data retrieval for the month was 0 GB. When the retrieval is initiated the concurrent data retrieved is 500 GB. 500 GB is your new monthly peak concurrent data retrieval. You are charged the difference between the prior and new peak concurrent data retrieval, that is, $(500 \text{ GB} - 0 \text{ GB}) \times \$0.30 / \text{GB} = \$150$. Twelve hours after initiating the retrieval of the 500 GB virtual tape you initiate retrieval of a virtual tape containing 600 GB of data. When the second retrieval is initiated you have two virtual tape retrievals concurrently in progress as any virtual tape retrieval takes about 24 hours to complete. The concurrent data retrieved is $500 \text{ GB} + 600 \text{ GB} = 1100 \text{ GB}$. The new peak concurrent data retrieval for the month is 1100 GB and you are now charged the difference between the new peak concurrent data retrieval and the existing peak concurrent data retrieval, that is, $(1100 \text{ GB} - 500 \text{ GB}) \times \$0.30 = \$180$.
- **Example 3** – You initiate retrieval of one 500 GB virtual tape from the VTS in US East (Northern Virginia) Region. This is the first retrieval of the month. The prior peak concurrent data retrieval for the month was 0 GB. When the retrieval is initiated the concurrent data retrieved is 500 GB which is your new peak concurrent data retrieved. The charge for the retrieval is the difference between the prior and the new peak concurrent data retrieval, that is, $(500 \text{ GB} - 0 \text{ GB}) \times \$0.30/\text{GB} = \$150.00$. The next day, you initiate a retrieval of one 500 GB virtual tape after the first tape retrieval is complete. Because 500 GB is equal to the prior peak concurrent data retrieval of 500 GB, the second retrieval is free.

Data Charges for Deleting Data

VTS is designed for use cases where data is retained for months, years, or decades. Deleting virtual tapes from a VTS is free if the virtual tape being deleted has been stored for three months or longer. If a virtual tape is deleted within three months of being archived, you are charged an early deletion fee. If you delete data from your VTS that is less than three months old, in the US East (Northern Virginia) Region, you are charged a prorated early deletion fee of \$0.03 per GB deleted within three months. For example, if you delete one virtual tape containing 1 GB of data one month after uploading it, you are charged a \$0.02 early deletion fee. If, instead, you delete the same virtual tape after two months, you would be charged a \$0.01 early deletion fee.

Figures

1. AWSglacier.png

© Barracuda Networks Inc., 2024 The information contained within this document is confidential and proprietary to Barracuda Networks Inc. No portion of this document may be copied, distributed, publicized or used for other than internal documentary purposes without the written consent of an official representative of Barracuda Networks Inc. All specifications are subject to change without notice. Barracuda Networks Inc. assumes no responsibility for any inaccuracies in this document. Barracuda Networks Inc. reserves the right to change, modify, transfer, or otherwise revise this publication without notice.