



Leveraging AWS Volume Gateway for Disaster Recovery

A summary of AWS Storage Gateway offerings, Volume Gateway use cases, and disaster recovery implementation

By Chris Madison

Introduction

Amazon Web Services (AWS) introduced the Storage Gateway software appliance in 2012 and has expanded that offering as recently as 2016. Additionally, in 2017, the Storage Gateway was added to the AWS list of Health Insurance Portability and Accountability Act (HIPAA) compliant services that are covered under its HIPAA Business Associate Addendum (BAA). The Storage Gateway is a hybrid cloud solution that allows on-premises systems to seamlessly back up volumes, files, or tape backups into AWS's extremely durable Simple Storage Service (S3). However, the use of S3 is hidden by the Storage Gateway.

The Storage Gateway comes in different options as outlined below:

Volume Gateway

This provides point-in-time, volume-based backups. On-premises systems mount Storage Gateway volumes as iSCSI devices. There are two types of volumes provided by Volume Gateways.

Stored volumes: The entire data set is stored locally on the Volume Gateway volume and in the AWS cloud. The stored volume option should be used when an application requires low latency access to its entire data set. Additionally, stored volumes provide an excellent choice for disaster recovery use cases.

Cached volumes: Cached volumes cache the most frequently accessed data on-premises, while maintaining the full data set in the AWS cloud.

File Gateway

The most recent addition to the Storage Gateway family, the File Gateway, is an NFS-based solution. On-premises systems mount the File Gateway offered NFS file systems for storage of application data. The files are maintained locally and backed up into S3.

Virtual Tape Library

Instead of creating tape backups, shipping them offsite, and paying for storage, leverage the Virtual Tape Library solution offered through the Storage Gateway. The virtual tape solution leverages iSCSI virtual tape drives and media changer and integrates with existing backup software to provide seamless and secure media storage.

The rest of this report will focus on the Volume Gateway, its uses, and where it fits architecturally in the enterprise. This report will conclude with a disaster recovery simulation.

Use Cases

The Volume Gateway is ideally suited for several use cases that large organizations grapple with on a daily basis. The following use cases cover standard backup and recovery scenarios, hybrid cloud solutions, and disaster recovery.

Backup and Restore: The primary premise for this use case is that point-in-time snapshots of your application data is securely stored in S3. For example, if you have several Windows-based applications whose hosts have been infected with ransomware, the Volume Gateway solution would allow for you to restore to a previous application state prior to infection.

The Volume Gateway hides the complexity of cloud integration and S3 storage through an industry-standard iSCSI interface. On-premises systems mount the iSCSI drives, and applications interact with the volume as a normal file system. However, the Volume Gateway is storing (or caching) data locally and sending asynchronous, encrypted snapshots to AWS.

Shared Workloads: For enterprises in a hybrid cloud environment, the Volume Gateway is one solution to share data between on-premise and cloud-based applications. Data written to on-premise volumes are replicated to the Elastic Block Storage (EBS) snapshots and stored on S3. The volume snapshots may be copied and directly leveraged by cloud-based applications.

One scenario for a shared workload for the Volume Gateway is a cognitive application feedback loop. For example, a cognitive application for image recognition or a question and answer system might record user feedback that corresponds to hits and misses of the cognitive application. Additionally, this feedback would have logged to a local file. It is conceivable to leverage the Volume Gateway to capture this log information over time. Snapshots of the data volume may be used by an image to retrain the cognitive application to increase its accuracy and precision.

Cloud Migration: On-premises applications that are migrating to the cloud may take advantage of the Volume Gateway solution. Application data is stored in AWS by virtue of using the Volume Gateway. In a 'shift-and-lift' model, the application may be migrated directly to AWS EC2 complete with a new volume based on the latest volume snapshot.

In this case, assume an organization has a Windows or Red Hat-hosted application running on-premise. The organization has a mandate to move this application to the cloud. To satisfy the organizational imperative, the IT team may first leverage the Storage Gateway to get data volumes into the cloud.

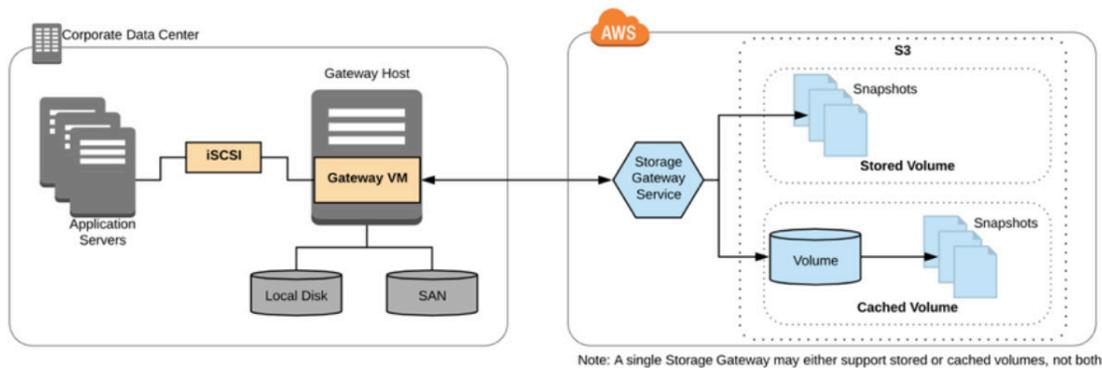
In the context of the Storage Gateway, the data volume snapshots are block storage devices that are directly mountable by Windows or Red Hat EC2 images. After data volume snapshots have been established, the application may be rehosted on corresponding EC2 images. Those images would then leverage the snapshots as data volumes. The application is now running in the cloud with the latest data available—migration complete.

Disaster Recovery: Cloud-based disaster recovery is a great use case for the Volume Gateway. If your local environment fails, you can quickly spin up an AWS-based disaster recovery environment using EC2. The EC2 images will have full access to the volume snapshots from the on-premises systems. If not leveraging EC2, a new gateway can be set up in a disaster recovery environment and volume snapshots may be made available to your systems and applications in that environment.

Hypothetically, physical infrastructure in data centers can be damaged or wiped out by natural disasters like hurricanes, floods, and earthquakes. These incidents have the potential to decimate entire businesses. However, prudent use of Storage Gateway can mitigate the impact of these events and provide continuity in the event of disaster. In this case, the Volume Gateway can be leveraged to keep hourly snapshots of data volumes in AWS. Your disaster recovery environment can be provisioned with minimal operational impact.

Architecture

The following diagram illustrates how the various components of the Volume Gateway fit together. For expediency, the diagram contains both the stored and cached volume options, but a single Volume Gateway may only support one or the other volume types, not both.



In the corporate data center, the **Volume Gateway** appliance is deployed as a virtual machine on a host. The gateway host may have directly attached and network-based storage capability, which may be used as local stored volumes or for cache. The Volume Gateway offers iSCSI devices for Application Servers to mount and leverage as block storage devices.

Within the AWS Cloud, the **Storage Gateway Service** manages the communication with the remote Volume Gateway and associated volumes and snapshots. If the Volume Gateway is hosting stored volumes, then the Storage Gateway Service maintains snapshots of the on-premise volumes. However, if the Volume Gateway is hosting cached instances, then the Storage Gateway Service

Considerations

Recovery Point Objective

The organization's Recovery Point Objective defines the acceptable amount of time that data may be lost due to service interruption. Generally, the Volume Gateway Service provides the capability for automated snapshots to occur as frequently as one hour.

Host Sizing

The Volume Gateway may be configured to host stored or cached volumes, but not both. The implication of volume type selection is the sizing of host machine. The following table lists the different memory requirements for the Volume Gateway image based on volume type.

Volume Type	Cache (Min)	Cache (Max)	Upload Buffer (Min)	Upload Buffer (Max)
Stored	N/A	N/A	150 GB	2 TB
Cached	150 GB	16 TB	150 GB	2 TB

Hypervisor Support

The Volume Gateway software appliance may be hosted on either VMware ESXi (4.1 to 6.5) or Microsoft Hyper-V (2008 R2, 2012, or 2012 R2) hypervisor platforms.

iSCSI Initiator Support

The Volume Gateway supports the following iSCSI Initiators:

- Windows Server 2012 and 2012 R2
- Windows Server 2008 and 2008 R2
- Windows 7
- Red Hat Enterprise Linux 5, 6, and 7
- VMware ESXi Initiato

Volume Size Support

A single gateway may support up to 32 volumes. Using cached volumes, each volume may be up to 32 TB in size for a maximum of 1 PB support for the cached-volume configured gateway. Using stored volumes, each volume may be scaled to 16 TB for a gateway maximum of 512 TB.

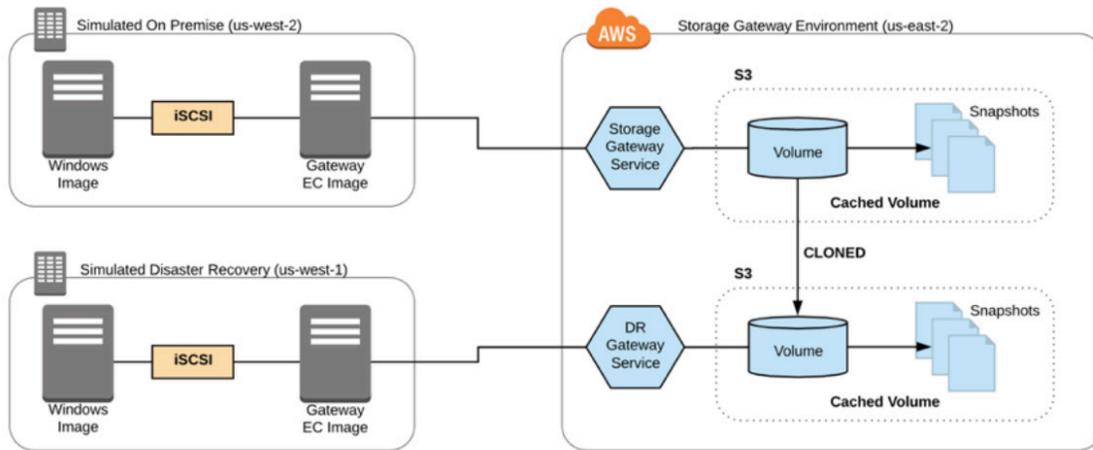
Pricing

Cost is an important consideration of finding a solution that fits business needs and budget. At this time, leveraging the us-east-2 region (Ohio), the cost breakdown is listed in the following table.

Service	Cost (us-east-2)
Volume Storage	\$0.023 per GB-month of data stored
EBS Snapshot Storage	\$0.05 per GB-month of data stored
Data Written	Maximum of \$125 a month
Data Transfer In	FREE
Data Transfer Out	Tiered, max of \$0.09 per GB, first 1 GB Free

Simulation

To demonstrate the Volume Gateway’s disaster recovery capabilities, the following environment was created in AWS.



In us-west-2, a Windows image and a Volume Gateway were created. The Volume Gateway shared a cached volume with the Windows image. The Volume Gateway was connected to the Storage Gateway Service. Test data, in the form of images, was added to cached volume.

A similar environment for disaster recovery was set up in us-west-1. The main exception is that a cached volume was not created.

There are two Service Gateway services now running. The AWS Console screenshot is below.

<input type="checkbox"/>	Name	Gateway ID	Status	Type
<input type="checkbox"/>	cmm-DluPGJBC-ServiceGateway	sgw-E0BE5B89	Running	Volume cached
<input type="checkbox"/>	cmm-DluPGJBC-DR-ServiceGateway	sgw-E2BE5B8B	Running	Volume cached

And the cached volume has data stored on it.

Volume ID	Status	Type	Used	Size	Gateway
vol-0d103f46633bc3c91	Available	Cached	4%	1 GiB	cmm-DluPGJBC-ServiceGateway

The next step is to clone the storage volume from the Storage Gateway Service to the DR Gateway Service. The next image illustrates the interface to create a clone. Note that one could have cloned from an existing EBS snapshot. In this case, the iSCSI target is named drvolume.

Create volume ✕

Gateway: cmm-DluPGJBC-DR-ServiceGateway

Capacity: GiB

Volume contents:

- New empty volume
- Based on EBS snapshot
- Clone from last volume recovery point [Learn more](#)

Source volume: vol-0d103f46633bc3c91

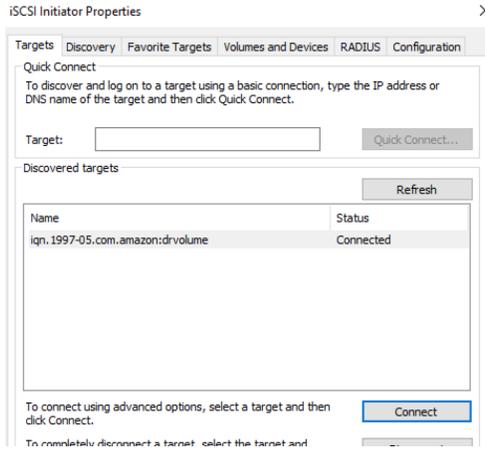
iSCSI target name:
1 to 200 characters including a-z, 0-9, and '-', '-'

Cancel Create volume

There are now two volumes available, each one managed by the two gateway services.

Volume ID	Status	Type	Used	Size	Gateway
vol-0685a59fd043a6617	Available	Cached	4%	1 GiB	cmm-DluPGJBC-DR-ServiceGateway
vol-0d103f46633bc3c91	Available	Cached	4%	1 GiB	cmm-DluPGJBC-ServiceGateway

On the disaster recovery Windows host, the next step is to mount the iSCSI target. As shown below, the disaster recovery volume, drvolume, is available.



Once the volume is mounted, the images are available for access.

This PC > New Volume (E:)

Name	Date modified	Type	Size
images	5/31/2018 3:41 PM	JPEG image	10 KB
images-2	5/31/2018 3:41 PM	JPEG image	9 KB
images-3	5/31/2018 3:41 PM	JPEG image	12 KB

At this point, setup for the AWS Volume Gateway service is complete.

Summary

This report presented a summary of the AWS Storage Gateway offerings, a deeper dive on the Volume Gateway use cases and architecture, and a brief simulation outlining the ease of implementing disaster recovery solution. The important takeaway is that the Storage Gateway is a flexible hybrid cloud storage solution that requires little effort to be effective.

If you are interested in learning more about the AWS Storage Gateway and how it would fit into your enterprise environment, please feel free to contact the Level Cloud Team at cloud@level.io.

Resources

CloudFormation templates—the templates used to build out the simulation environment: <https://github.com/GetLevel/cloud-cftemplates/blob/master/storage-gateway/readme.md>

AWS Storage Gateway documentation: https://aws.amazon.com/storagegateway/?sc_medium=AW_PAC_Adopt_StorageGateway_b_030118sc_content=vtl_p?s_kwcid=AL!4422!3!270968207905!p!!g!!storage%20gateway&s_kwcid=AL!4422!3!270968207905!p!!g!!storage%20gateway&ef_id=WwcSTAAABZdhy@K6:20180524192812:s

About Level

Level helps clients transform their business with strategic consulting and technical execution services. We work with your IT organization, product groups, and innovation teams to design and deliver on your technical priorities.

Level's cloud experts combine decades of traditional architecture, development, security, and infrastructure experience with a complete mastery of available and emerging cloud offerings. Our client-centric approach focuses first on understanding your business needs and goals, then selecting the right cloud technology to make you efficient, agile, and scalable.

We tailor custom solutions to fit within your business processes, simultaneously reducing TCO and downtime while increasing productivity, security, ROI, and speed to market.

For more information, contact us at hello@level.io.