

Exploring AWS Storage Gateway: the perfect bridge between the On-prem and the AWS Cloud

13 October 2023 - 7 min. read

[AWS Storage Gateway](#)

[Hybrid Cloud](#)

Introduction

We are noticing an increasing number of customers contacting us to discuss migration projects that involve moving their on-premise infrastructures to Cloud environments. Whether the project will be hosted on hybrid infrastructures or entirely in the Cloud, the main concern (among many others) is identifying the most suitable AWS storage solution for their needs and infrastructure. By infrastructure, we refer to the set of applications, files, and data in general that will depend on the most reliable storage possible.

It is not only easy to give a detailed answer before clearly understanding the context and requirements of the infrastructure to be migrated. This article also attempts to answer a complex question, but to do so, it is necessary to change the point of view of the question itself. I'll start by citing some words from the famous speech that inspired entire generations, inviting them to change their perspective and approach in everyday situations. John F. Kennedy, during his inauguration as President of the United States, said: "Ask not what your country can do for you - ask what you can do for your country".

Likewise, we can paraphrase those words and finally imagine the question from the true point of view: "Don't ask which AWS storage is suitable for your infrastructure - ask your infrastructure which AWS storage will best adapt to".

Amazon Web Services (AWS) Storage Gateway

This managed AWS service is often underestimated or underused. We will realize how the great flexibility and multiple conditions of use will favor the creation, in the first instance, of

hybrid infrastructures, which will subsequently be able, with some well-defined steps, to complete their transition and remain entirely in the Cloud.

AWS Storage Gateway works as a bridge between the local infrastructure and the AWS cloud services. It also allows us to choose between a series of storage protocols and interfaces that will give access to data management in the Cloud. AWS Storage Gateway supports the most popular storage access protocols, including NFS (Network File System), SMB (Server Message Block), and iSCSI (Internet Small Computer System Interface). The flexibility of the AWS Storage Gateway will also allow us to range from Windows to AWS S3 (Simple Storage Service) as a backend for our data, favoring the integration of a series of applications and workloads resident in on-premise infrastructures.

Below and for completeness, you will find the four configuration modes as destination backend, related to Storage Gateway.

- **S3 File Gateway:** This mode supports access via NFS and SMB to files stored in Amazon S3. It is ideal for companies that want to store files in the Cloud without changing existing applications.
- **FSX File Gateway:** this mode provides access to the on-premise infrastructure, allowing us to access the shares managed in Amazon FSx for Windows File Server.
- **Tape Gateway:** This type of storage gateway provides a virtual tape library (VTL) interface that allows you to replace physical tape libraries with virtual tapes stored in Amazon S3 and Glacier. It is useful for backup and archiving purposes.
- **Volume Gateway:** Provides block storage volumes mounted as iSCSI devices on local servers.

Why use AWS Storage Gateway

If you come from a legacy on-premise environment, the best way to get familiar with AWS-managed services is to start thinking about hybrid solutions. AWS Storage Gateway perfectly adapts to this need, especially with file gateways. Once you have completed the registration steps and verified your AWS account, you can access the AWS management console to create resources and services and use the various tools offered by AWS, including Storage Gateway.

To keep your AWS account secure, remember to keep your login credentials, including your username and password, in a safe place. In addition, it is crucial to increase the security level of your AWS account by using features like two-factor authentication.

Before configuring AWS Storage Gateway, you need to define the connection method between your infrastructure hosted in an on-premise data center and the AWS cloud. The most popular ways for interconnecting the two environments involve the configuration of a VPN (site-to-site), for which AWS provides a double channel where the encrypted connection tunnels can be certified, which is the recommended choice, or through the public Internet network, which is sometimes not recommended due to governance and compliance policies required in specific business scenarios.

Using the AWS console, you can define which gateway to use, such as Amazon S3 file gateway or Amazon FSx File Gateway, and subsequently, the virtualization technology (hypervisor) used in your on-premise infrastructure. Supported hypervisors include VMware ESXi, Microsoft Hyper-V, and Linux KVM. After configuring the virtual appliance in your on-premise environment, you can share data with AWS services. Cloud-native and legacy applications can access data and work with S3, for example, without necessarily having to rewrite code.

However, AWS Storage Gateway is not always applicable. For example, suppose you have managed your file servers in a scenario where authentication and authorizations are provided through Active Directory, with many ACLs on the various shares. In that case, you will be faced with a limit that implies a maximum of 10 Windows ACLs per file or folder. You will encounter similar limitations when dealing with shared folders, as AWS Storage Gateway supports up to 50 shared folders. Therefore, it is essential to consider the context before deciding to use Storage Gateway.

AWS Storage Gateway won't groove with SAMBA (SMBv1): Exploring its usage limits.

Many companies are still using legacy industrial applications or machinery that require outdated and deprecated technologies, such as SMB1 (Server Message Block 1). Despite its limitations, some organizations hesitate to abandon it due to their desire to innovate while preserving what they can. However, SMB1 is no longer recommended for use in modern environments due to various issues such as security vulnerabilities, poor performance, lack of advanced features, and complex maintenance.

To ensure the safety and efficiency of modern networks and operating systems, it is highly recommended to disable SMB1 and migrate to newer versions of the protocol, such as SMB2 or SMB3, which are fully supported by AWS Storage Gateways and offer greater security, better performance, and advanced features. It is also important to keep operating systems and applications up-to-date to avoid security risks associated with outdated protocols like SMB1.

If you are still dancing SMB1, your infrastructure is likely not keeping up with the modern rhythm of AWS, rather than say that Storage Gateway cannot dance to the new music.

Use cases

AWS Storage Gateway is a flexible solution that offers various practical applications for managing data and storage resources in hybrid environments, combining on-premises storage with Amazon Web Services (AWS) Cloud.

Backup and Long-Term Archiving

One of the main applications is **Backup and Long-Term Archiving**. With AWS Storage Gateway, you can securely and cost-effectively store data and long-term backups on services like Amazon S3 or Amazon Glacier. This is especially useful for organizations needing a low-cost and secure storage solution for their backup and archival data.

File and NAS Storage

Another key application is **File and NAS Storage**. AWS Storage Gateway allows you to transform an on-premises storage system into a file server or network-attached storage (NAS) system, synchronizing data directly with Amazon S3. This enables businesses to expand their on-premises storage space without additional hardware investments.

Data Migration

AWS Storage Gateway is also an excellent choice for streamlining **Data Migration** to AWS. Initially, data can be copied to Amazon S3 via Storage Gateway, allowing for a gradual and planned migration to the AWS environment.

Caching

Another crucial application is the use of caching to optimize **Cloud Application Performance**. By using the caching mode, AWS Storage Gateway accelerates access to data frequently used by cloud applications, thus improving overall application efficiency and performance.

Backup for Cloud Applications

Finally, AWS Storage Gateway can be used for **Backups for Cloud Applications**, enabling the backup of cloud applications in an on-premises or cloud-based storage environment. This ensures reliable data protection for cloud applications.

Conclusions

AWS Storage Gateway is a highly flexible solution that provides a wide range of integration options with Amazon Web Services' cloud storage infrastructure. It is an essential technology to consider when looking for innovative and updated solutions. This service acts as a **critical link between the legacy world and the Cloud**, enabling us to take advantage of the great scalability and resilience of Amazon S3 without needing to rewrite applications or make significant changes to existing infrastructures.

In summary, AWS Storage Gateway provides a comprehensive and versatile solution for data storage and management challenges in hybrid environments, enabling businesses to fully leverage AWS cloud storage's benefits.

Did you have the chance to use it in one of the mentioned Use Cases or in different ones? Share your experience!

See you in 14 days with a new blog post on **Proud2beCloud!**

About Proud2beCloud

Proud2beCloud is a blog by **beSharp**, an Italian APN Premier Consulting Partner expert in designing, implementing, and managing complex Cloud infrastructures and advanced services on AWS. Before being writers, we are Cloud Experts working daily with AWS services since 2007. We are hungry readers, innovative builders, and gem-seekers. On Proud2beCloud, we regularly share our best AWS pro tips, configuration insights, in-depth news, tips&tricks, how-tos, and many other resources. Take part in the discussion!



Gianluca Laino

Technical Project Manager @ beSharp. After over twenty years of on-premise projects, I finally understood what I was looking at in the sky as a child: the “clouds”. From the Commodore 64 to the EC2 with an eye towards the future. My only analog passion? football!
