IJIRSET
International Journal of Innovative Research in
SCIENCE | ENGINEERING | TECHNOLOGY

e-ISSN: 2319-8753 | p-ISSN: 2347-6710

INTERNATIONAL JOURNAL OF INNOVATIVE RESEARCH

IN SCIENCE | ENGINEERING | TECHNOLOGY

Volume 10, Issue 12, December 2021

INTERNATIONAL STANDARD SERIAL NUMBER INDIA

Impact Factor: 7.569

9940 572 462 | 6381 907 438 | ijirset@gmail.com | www.ijirset.com

| e-ISSN: 2319-8753, p-ISSN: 2347-6710| www.ijirset.com | Impact Factor: 7.569|

|| Volume 10, Issue 12, December 2021 ||

| DOI:10.15680/IJIRSET.2021.1012067|

Security and Privacy of AWS S3

Shubham Choudhari[1], Aditya Gupta[2], Nitin Kamble[3]

Student, School of Computer Science and Engineering, Ajeenkya D Y Patil University, Pune, India

Student, School of Computer Science and Engineering, Ajeenkya D Y Patil University, Pune, India

Sr. Faculty, School of Computer Science and Engineering, Ajeenkya D Y Patil University, Pune, India

ABSTRACT: Many people use cloud computing services, which are provided by companies like Amazon, Microsoft, and Google. The storage service is one of the most regularly used services. Users can store and access data online from any location using the service. However, data that is stored remotely and managed by a service provider raises issues about the data's security and privacy. The security and privacy of Amazon AWS S3 and are reviewed in this study. The comparison will be based on their security and privacy compliance processes, as well as how the access control handles data protection for users.

KEYWORD: Cloud computing services, AWS, security, privacy, compliance, data storage, access control

I. INTRODUCTION

Amazon S3 (Simple Storage Service) is Amazon Web Services' online storage web service. In March 2006, Amazon introduced S3, its first publicly available online service, in the United States, then in November 2007 in Europe.
Amazon charged end users $0.15 per gigabyte-month when it first launched, with additional fees for bandwidth consumed in transmitting and receiving data, as well as a per-request (get or put) fee.
Pricing was changed to tiers on November 1, 2008, with end users holding more than 50 terabytes receiving a discount.
Amazon claims that S3 is built on the same scalable storage technology that powers Amazon.com's global e-commerce network.
As of December 2011, Amazon S3 was said to be storing over 762 billion things. Web hosting, image hosting, and backup storage are all examples of S3 applications. S3 has a monthly uptime guarantee of 99.9%, which translates to about 43 minutes of downtime each month.
Amazon S3 is an online storage service. It's a straightforward storage solution that provides software developers with a cost-effective, highly scalable, dependable, and low-latency data storage infrastructure.
It has a simple web services interface that allows you to save and retrieve unlimited amount of data from anywhere on the internet at any time. Developers can simply create applications that utilise Internet storage by using this web service. Because Amazon S3 is highly scalable and you only pay for what you use, developers can start small and scale up as needed without sacrificing performance or reliability. It's built to be extremely versatile: store whatever type and amount of data you want; read the same piece of data a million times or only for emergency disaster recovery; create a basic FTP application or a sophisticated web application like Amazon.com's retail website. Developers may focus on innovation rather than figuring out how to store their data with Amazon S3.
You can, for example, use Amazon S3 to host your complete static website. Customers frequently use Amazon S3 to store photographs, video, and other information for their websites. To host your full website from your Amazon S3 bucket, you may take use of root and index document support, as well as custom HTML errors. You can store almost any type of data in any format you want.

Figure 1. Security and privacy relationship

| e-ISSN: 2319-8753, p-ISSN: 2347-6710| www.ijirset.com | Impact Factor: 7.569|

|| Volume 10, Issue 12, December 2021 ||

| DOI:10.15680/IJIRSET.2021.1012067|

The total amount of data and things you can store is limitless. Individual AmazonS3 objects might be anywherefrom 1 byte and 5 gigabytes in size. The maximum size of an object that can be uploaded in a single PUT is 5GB. Customers should use the Multipart Upload functionality for objects larger than 100 megabytes.

II. SECURING S3

Data Protection

Redundancy

Reduced Redundancy Storage (RRS) is an Amazon S3 storage option that allows users to store noncritical, repeatable data with less redundancy than standard Amazon S3 storage. It provides a highly accessible option for publishing or sharing content that has been archived elsewhere, as well as for storing thumbnails, transcoded media, and other easily reproducible processed data. RRS stores things across numerous devices across various sites, offering 400 times the durability of a regular disc drive, however it does not replicate objects as frequently as standard Amazon S3 storage.

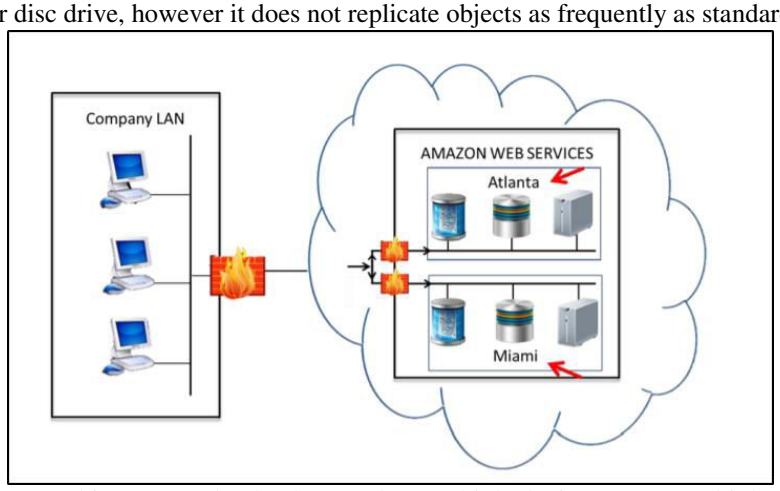

Figure 2: Objects are redundantly stored on multiple devices across multiple facilities.

Versioning

Every version of every object in an Amazon S3 bucket may be saved, retrieved, and restored using versioning. When you enable Versioning for a bucket, Amazon S3 ensures that existing objects are preserved if you PUT, POST, COPY, or DELETE them. GET requests, by default, return the most

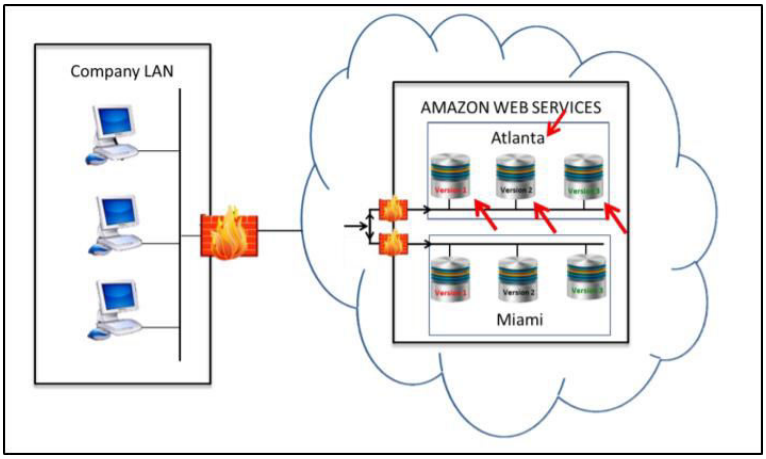

Figure 3: You can use versioning to preserve, retrieve, and restore every version of every object stored in your Amazon S3 bucket.

recently written version. By mentioning a version in the request, older versions of an overwritten or deleted item can be retrieved. Versioning adds another layer of security by allowing customers to recover objects that have been mistakenly

| e-ISSN: 2319-8753, p-ISSN: 2347-6710| www.ijirset.com | Impact Factor: 7.569|

|| Volume 10, Issue 12, December 2021 ||

| DOI:10.15680/IJIRSET.2021.1012067|

overwritten or deleted. This makes it simple to recover from unintentional user activities and application faults.

Checksums

Both the x-amz-sha256-tree-hash and x-amz-content-sha256 headers must be included when uploading an archive. The checksum of the payload in your request body is contained in the x-amz-sha256-tree-hash header. The x-amz-sha256-tree-hash header is calculated in this topic. The x-amz-content-sha256 header is required for authorization because it is a hash of the full payload. See Example Signature Calculation for Streaming API for additional information.
Your request's payload can be one of the following:

Entire archive - The complete archive is sent in the request body when using the Upload Archive API to upload an archive in a single request. In this scenario, you must include the complete archive's checksum.

Archive part - When utilising the multipart upload API to upload an archive in parts, you only provide a portion of the archive in the request body. In this scenario, the archive part's checksum is included. After you've uploaded all of the parts, you'll need to make a Complete Multipart Upload request, which must include the complete archive's checksum.

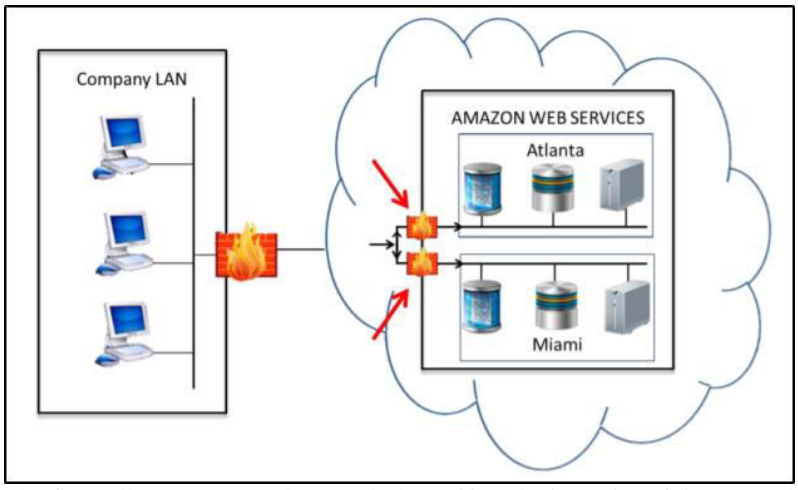

Figure 4: At the firewall Amazon S3 also regularly verifies the integrity of data stored using checksums.

Data Encryption

Server Side Encryption (SSE) or a client library like the Amazon S3 Encryption Client can be used to encrypt data. If you utilise a client library, you keep control of the encryption keys. Some customers appreciate having more control over their keys, while others would rather not have to deal with the administrative burden of managing and protecting keys. AWS manages key management and key security for you if you use SSE. If you prefer AWS to manage your keys, SSE is the way to go. SSE employs the 256-bit Advanced Encryption Standard, one of the most powerful block cyphers available (AES-256). The maximum key size defined for AES is 256 bits. For objects stored in Standard Storage and Reduced Redundancy Storage, both client-side encryption and server-side encryption are supported (RRS)

Server-Side Encryption

Amazon S3 encrypts your data as it writes it to discs in its data centres and decrypts it for you when you access it, which is known as server-side encryption. There is no difference in how you access encrypted or unencrypted items as long as you authenticate your request and have access permissions. Amazon S3 takes care of encryption and decryption on your behalf. If you share your objects with a pre-signed URL, for example, the pre-signed URL works for both encrypted and unencrypted items.

Strong multi-factor encryption is used in Amazon S3 Server Side Encryption. Amazon S3 uses a unique key to encrypt each object. It encrypts the key with a master key that it rotates on a regular basis as an additional protection. To protect your data, Amazon S3 Server Side Encryption uses a 256-bit Advanced Encryption Standard (AES-256), which is one of the most powerful block cyphers available. The object level allows you to specify data encryption. As a result, both

| e-ISSN: 2319-8753, p-ISSN: 2347-6710| www.ijirset.com | Impact Factor: 7.569|

|| Volume 10, Issue 12, December 2021 ||

| DOI:10.15680/IJIRSET.2021.1012067|

encrypted and unencrypted objects could be found in your bucket.

Client-Side Encryption

Important: Because your private encryption keys and unencrypted data are never delivered to AWS, it's critical that you keep your encryption keys safe. You won't be able to decrypt your data if you lose your encryption keys.

AWS employs a technique known as envelope encryption. In envelope encryption, you give the Amazon S3 encryption client your encryption key, and the client takes care of the rest. The procedure is as follows:

1) The Amazon S3 encryption client generates a one-time-use symmetric key for encrypting your data.

2) Using your private encryption key, the client encrypts the envelope symmetric key.

3) The client then uploads your encrypted data and the encrypted envelope key to Amazon S3.

The encryption described in the previous slide is reversed when retrieving and decrypting client-side encrypted data from Amazon S3.

Your private encryption key can be an asymmetric key pair (made of public and private keys) or a symmetric key that the client employs in the envelope encryption process. Asymmetric keys are recommended by Amazon because they provide a higher level of security in the envelope encryption process.

Access Control

Identity and Access Management (IAM)

You can use AWS Identity and Access Management (IAM) to safeguard your users' access to AWS services and resources. IAM allows you to establish and manage users in AWS' identity management system (referred to as "IAM users"), as well as offer access to AWS services to users in your corporate directory who are controlled outside of AWS (referred to as "federated users"). When utilising AWS, IAM provides more security, flexibility, and control. IAM makes it simple to grant safe access to your AWS account and resources to numerous users. You can use IAM to:

1) Manage IAM users and their access - In AWS' identity management system, you can create users, assign them individual security credentials (such as Access Keys, passwords, and Multi Factor Authentication devices), or request temporary security credentials to grant them access to AWS services and resources. You can control which operations a user can conduct by managing permissions.

2) You can manage federated user access - You can request security credentials with configurable expirations for users in your corporate directory, allowing you to provide secure access to resources in your AWS account without having to create an IAM user for them. To govern which operations a user can execute, you establish permissions for these security credentials.

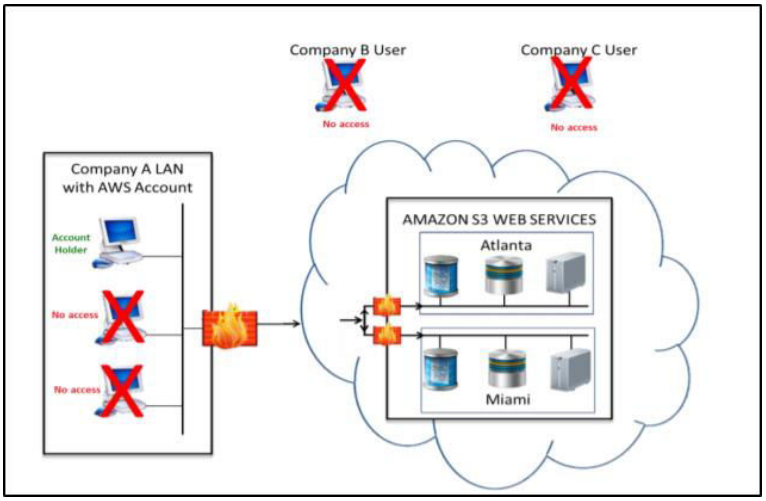

Figure 5: Notice in this figure only the Account Holder has access to the cloud

IAM Policies

IAM policies are used to manage permissions. These policies are linked to users, allowing you to manage permissions for all users in your AWS account from a single location.

IAM enables you to:

1) Create IAM user identities and Add IAM users to your AWS account.

| e-ISSN: 2319-8753, p-ISSN: 2347-6710| www.ijirset.com | Impact Factor: 7.569|

|| Volume 10, Issue 12, December 2021 ||

| DOI:10.15680/IJIRSET.2021.1012067|

2) Group IAM users - Create groups to handle rights for many IAM users in your AWS account more effectively.
Control which actions your users can execute, such as accessing certain AWS service APIs and resources, by centralising user access control.
4) Add restrictions to control how your users can use AWS, such as the time of day, their originating IP address, or if they're using SSL.
5) Create and assign security credentials - Create and assign security credentials to your IAM users, then rotate and/or revoke them as needed.
6) Create temporary security credentials - For your IAM users, federated users, or applications, request temporary security credentials with configurable expiration and permissions.
IAM allows you to create, delete, and list IAM users, manage group membership, manage user security credentials, and assign rights using a variety of methods. IAM APIs, Command Line Tools, and the IAM console, which provides a point-and-click web-based interface, can be used to create and manage users, groups, and permissions. You can also develop policies with the AWS Policy Generator.

Resource-based Permissions on Objects and Buckets

Using access control lists (ACLs) and bucket policies, Amazon S3 supports resource-based permissions on objects and buckets. Buckets and objects contain ACLs and bucket policies that define which AWS accounts have access and what type of access they have. You can only give other AWS Accounts access to your Amazon S3 resources using ACLs.

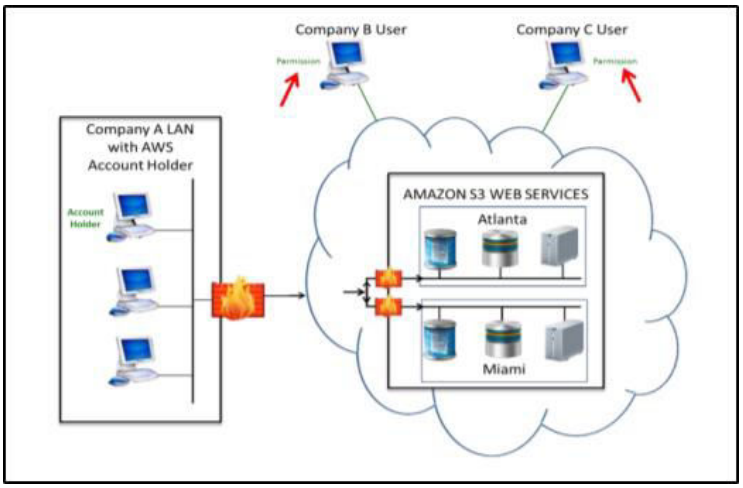

Figure 6: With ACLs, you can only grant other AWS Accounts access to your Amazon S3 resources

Bucket Policies
You can do both with Bucket Policies. You can provide users in your AWS Account access to your Amazon S3 resources, and you can also give users in other AWS Accounts access to your Amazon S3 resources.

| e-ISSN: 2319-8753, p-ISSN: 2347-6710| www.ijirset.com | Impact Factor: 7.569|

|| Volume 10, Issue 12, December 2021 ||

| DOI:10.15680/IJIRSET.2021.1012067|

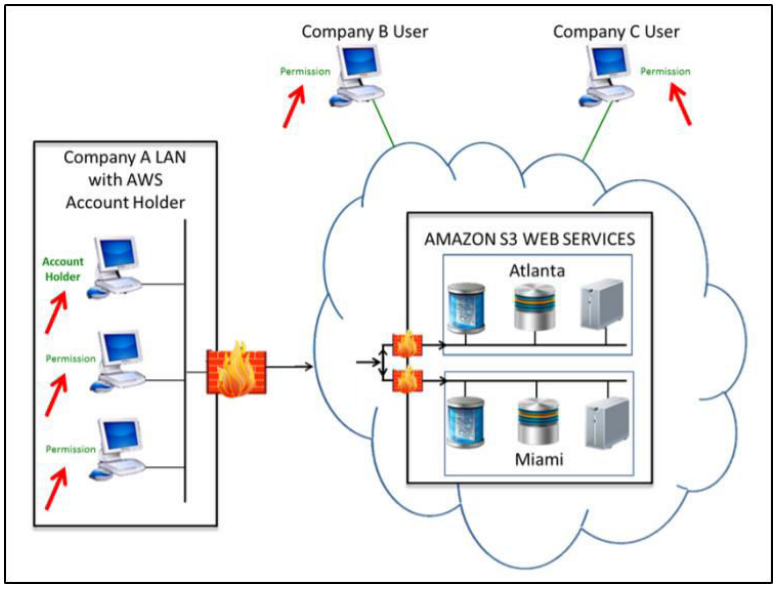

Figure 7: With Bucket Policies, you can do both.

The main challenge with any shared storage system is whether unwanted users can access information on purpose or by accident. Amazon S3 APIs provide both bucket- and object-level access restrictions, with defaults that only allow authenticated access by the bucket and/or object creator, to ensure clients have the most flexibility in determining how, when, and to whom they desire to disclose the information they store in AWS. An Access Control List (ACL) linked with the bucket controls write and delete permissions. An ACL controls permission to edit the bucket ACLs, and it defaults to creator-only access. As a result, the customer retains complete control over who has access to their information. Access to Amazon S3 can be allowed based on an AWS Account ID, a Depay Product ID, or it can be granted to anyone.

A clean root account is the starting point for an Amazon S3 account. Buckets are created at the root level. The root folder in Amazon S3 is referred to as a bucket. You can make numerous buckets, and you can put your folders and photographs inside of them.

Amazon S3 has released a set of APIs and developers around the world that let your Amazon S3 account to communicate with your local computer, allowing you to handle all of your file uploading, synchronisation, and backup, among other things.

III. FUTURE PROSPECTS OF S3

"Everything we've added over time, from lifecycle management to security properties to methods you can add to an object, and so on, is far from simple." However, we started small to ensure that the core API interfaces would be sufficient years down the road. APIs are eternal. It's not something you can later decide you don't want after you've made them for your servers. They are the foundation of many businesses. However, things change over time, and we've progressed from people posting thumbnails of stuff on eBay to storing literally petabytes of video, image, and other data and archiving it."

IV. CONCLUTION

Amazon S3 is Amazon's cloud storage system that underpins most of the modern internet. S3 is flexible, scalable, extremely dependable, and accessible because to its object-based storage paradigm. It's also inexpensive. Instead of paying a hefty flat-rate price for data storage, businesses simply pay for what they use.

A web-based console, APIs, and a vast number of Amazon products and third-party software are all available to interact with your storage. It's an excellent choice if you need to store a significant quantity of data, want secure file uploads and downloads, or require a cloud storage solution that can grow with your business.

| e-ISSN: 2319-8753, p-ISSN: 2347-6710| www.ijirset.com | Impact Factor: 7.569|

|| Volume 10, Issue 12, December 2021 ||

| DOI:10.15680/IJIRSET.2021.1012067|

REFRENCES

[1] Palinka, Mayur R., et al. "Amazon S3 for science grids: a viable solution." Proceedings of the 2008 international workshop on Data-aware distributed computing. 2008.
[2]Saeed, Iman, Sarah Baras, and Hassan Hajjdiab. "Security and privacy of AWS S3 and Azure Blob storage services." 2019 IEEE 4th International Conference on Computer and Communication Systems (ICCCS). IEEE, 2019.
[3]Shu, Chiao-Fe, et al. "IBM smart surveillance system (S3): a open and extensible framework for event based surveillance." *IEEE Conference on Advanced Video and Signal Based Surveillance, 2005.*. IEEE, 2005.
[4]Garfinkel, Simson. "An evaluation of Amazon's grid computing services: EC2, S3, and SQS." (2007).
[5]Subashini, Subashini, and Veeraruna Kavitha. "A survey on security issues in service delivery models of cloud computing." *Journal of network and computer applications* 34.1 (2011): 1-11.
[6]Oh, Seung-hee, Daehee Seo, and Byunggil Lee. "S3 (secure ship-to-ship) information sharing scheme using ship authentication in the e-navigation." *International Journal of Security and its Applications* 9.2 (2015): 97-110.
[7]Brantner, Matthias, et al. "Building a database on S3." *Proceedings of the 2008 ACM SIGMOD international conference on Management of data*. 2008.

INNO SPACE
SJIF Scientific Journal Impact Factor

Impact Factor:
7.569

doi cross ref

ISSN INTERNATIONAL STANDARD SERIAL NUMBER INDIA

निस्केयर NISCAIR

INTERNATIONAL JOURNAL OF INNOVATIVE RESEARCH

IN SCIENCE | ENGINEERING | TECHNOLOGY

9940 572 462 6381 907 438 ijirset@gmail.com

www.ijirset.com

Scan to save the contact details