**IOSR Journal of Engineering (IOSRJEN)** **www.iosrjen.org**
ISSN (e): 2250-3021, ISSN (p): 2278-8719
PP 41-46

# An Effective Data Disaster Recovery Services Using AWS Cloud

Pradeep Fale[1], Punam Wadhi[2], Smita Chauhan[3], Chetana Ghonge[4], Ashwini Uikey[5], Diksha Lautre[6]

[1]*Assistant Professor, Priyadarshini College of Engineering, Nagpur, India, 440019*
[2]*B.E. student, Dept. of Information Technology, Priyadarshini College ofEngineering, Nagpur, India, 440019*
[3]*B.E. student, Dept. of Information Technology, Priyadarshini College ofEngineering, Nagpur, India, 440019*
[4]*B.E. student, Dept. of Information Technology, Priyadarshini College of Engineering, Nagpur, India, 440019*
[5]*B.E. student, Dept. of Information Technology, Priyadarshini College ofEngineering, Nagpur, India, 440019*
[6]*B.E. student, Dept. of Information Technology, Priyadarshini College of Engineering, Nagpur, India, 440019*

***Abstract –*** *Nowadays cloud computing is most demanding technology. While dealing with the data in cloud we faces some problems like downtime, storage, security and privacy, vulnerabilities to attack, limited control and flexibility, Vendor lock- in, costs. AWS provides the facility to overcome this problem. Amazon Web Services (AWS) is a comprehensive, evolving cloud computing platform provided by Amazon. An advantage of the AWS cloud is that it allows customers to scale and innovate, while maintaining a secure environment. Customers pay only for the services they use, meaning that you can have the security you need, but without the upfront expenses, and at a lower cost than in an on-premises environment. This paper describes to implementing a proven and cost-effective disaster recovery solution for file servers that can minimize data loss and provide fast, automatic recovery of file services running on the AWS cloud.*
***Keywords****- Cloud computing, Data replication, AWS cloud, Disaster recovery.*

## I. Introduction

Cloud computing is the delivery of computing services—servers, storage, databases, networking, software, analytics, intelligence and more—over the Internet ("the cloud") to offer faster innovation, flexible resources and economies of scale. You typically pay only for cloud services you use,helping lower your operating costs, run your infrastructure more efficiently and scale as your business needs change.

Amazon Web Services provides a comprehensive set of services and tools for deploying Microsoft Windows-based workloads on its reliable and secure cloud infrastructure. Active Directory Domain Services (AD DS) and Domain Name Server (DNS) are core Windows services that provide the foundation for many enterprise class Microsoft based solutions including Microsoft Share Point, Microsoft Exchange, and .NET applications.

The aim of a disaster recovery site is to keep business operations active in case of any damage or hazard occurs. Disaster recovery helps in restoring applications, data, and hardware quickly for business continuity.

Businesses are using the AWS cloud to enable faster disaster recovery of their critical IT systems without incurring the infrastructure expense of a second physical site. The AWS cloud supports many popular disaster recovery (DR) architectures from "pilot light" environments that may be suitable for small customer workload data center failures to "hot standby" environments that enable rapid fail over at scale. With data centers in Regions all around the world, AWS provides a set of cloud-based disaster recovery services that enable rapid recovery of your IT infrastructure and data.

## II. Methodology

**AWS Storage Services**

**S3 (**Simple Storage Service)

Amazon Simple Storage Service is storage for the Internet. It is designed to make web-scale computing easier for developers. Amazon S3 has a simple web services interface that you can use to store and retrieve any amount of data, at any time, from anywhere on the web. It gives any developer access to the same highly scalable, reliable, fast, inexpensive data storage infrastructure that Amazon uses to run its own global network of web sites. The service aims to maximize benefits of scale and to pass those benefits on to developers.

This guide explains the core concepts of Amazon S3, such as buckets and objects, and how to work with these resources using the Amazon S3 application programming interface (API).

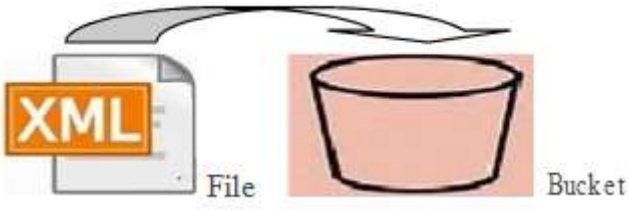


**Fig. S3 (Simple Storage Service) Cloud Front**

Amazon Cloud Front is a web service that speeds up distribution of your static and dynamic web content, such as .html, .css, .js, and image files, to your users. Cloud Front delivers your content through a worldwide network of data centers called edge locations. When a user requests content that you're serving with Cloud Front, the user is routed to the edge location that provides the lowest latency (time delay), so that content is delivered with the best possible performance.

If the content is already in the edge location with the lowest latency, Cloud Front delivers it immediately.

If the content is not in that edge location, Cloud Front retrieves it from an origin that you've defined—such as an Amazon S3 bucket, a Media Package channel, or an HTTP server (for example, a web server) that you have identified as the source for the definitive version of your content.

As an example, suppose that you're serving an image from a traditional web server, not from Cloud Front. For example,you might serve an image,sunsetphoto. png, using the URL http:// example.com/ sunsetphoto.png.

Your users can easily navigate to this URL and see the image. But they probably don't know that their request was routed from one network to another— through the complex collection of interconnected networks that comprise the internet—until the image was found.

Cloud Front speeds up the distribution of your content by routing each user request through the AWS backbone network to the edge location that can best serve your content. Typically, this is a Cloud Front edge server that provides the fastest delivery to the viewer. Using the AWS network dramatically reduces the number of networks that your users' requests must pass through, which improves performance. Users get lower latency— the time it takes to load the first byte of the file— and higher data transfer rates.

You also get increased reliability and availability because copies of your files (also known as objects) are now held (or cached) in multiple edge locations around the world.

**Elastic Block Storage**

Amazon Elastic Block Store (Amazon EBS) provides persistent block storage volumes for use with Amazon EC2 instances in the AWS Cloud. Each Amazon EBS volume is automatically replicated within its Availability Zone to protect you from component failure, offering high availability and durability. Amazon EBS volumes offer the consistent and low-latency performance needed to run your workloads. With Amazon EBS, you can scale your usage up or down withinminutes – all while paying a low price for only what you provision.

Amazon EBS is designed for application workloads that benefit from fine tuning for performance, cost and capacity. Typical use cases include Big Data analytics engines (like the Hadoop/HDFS ecosystem and Amazon EMR clusters), relational and NoSQL databases (like Microsoft SQL Server and MySQL or Cassandra and MongoDB), stream and log processing applications (like Kafka and Splunk), and data warehousing applications (like Vertica and Teradata).

**Glacier**

Amazon Simple Storage Service Glacier, that is Amazon S3 Glacier (Glacier), is a storage service optimized for infrequently used data, or "cold data."

Glacier is an extremely low-cost storage service that provides durable storage with security features for data archiving and backup. With Glacier, customers can store their data cost effectively for months, years, or even decades. Glacier enables customers to offload the administrative burdens of operating and scaling storage to AWS, so they don't have to worry about capacity planning, hardware provisioning, data replication, hardware failure detection and recovery, or time-consuming hardware migrations.

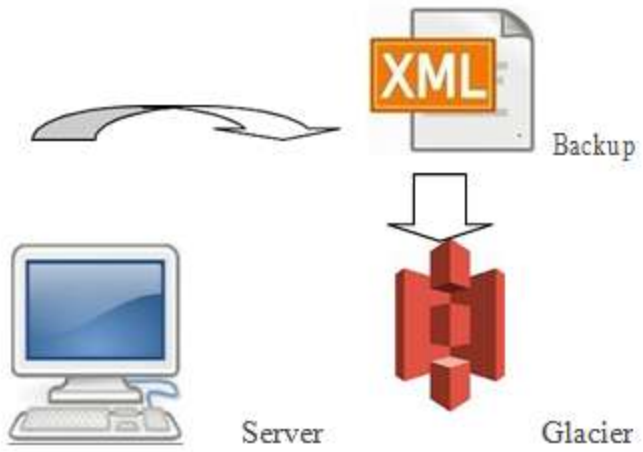


**Fig. Glacier(**Data Archiving service offering low price storage.)

**Snowball**

Snowball is a petabyte-scale data transport solution that uses devices designed to be secure to transferlarge amounts of data into and out of the AWS Cloud. Using Snowball addresses common challenges with large-scale data transfers including high network costs, long transfer times, and security concerns. Customers today use Snowball to migrate analytics data, genomics data, video libraries, image repositories, backups, and to archive part of data center shutdowns, tape replacement or application migration projects. Transferring data with Snowball is simple, fast, more secure, and can be as little as one-fifth the cost of transferring data via high-speed Internet.

With Snowball, you don't need to write any code or purchase any hardware to transfer your data. Simply create a job in the AWS Management Console ("Console") and a Snowball device will be automatically shipped to you. Once it arrives, attach the device to your local network, download and run the Snowball Client ("Client") to establish a connection, and then use the Client to select the file directories that you want to transfer to the device. The Client will then encrypt and transfer the files to the device at high speed. Once the transfer is complete and the device is ready to be returned, the E Ink shipping label will automatically update and you can track the job status via Amazon Simple Notification Service (SNS), text messages, or directly in the Console.

**Storage Gateway**

AWS Storage Gateway connects an on-premises software appliance with cloud-based storage to provide seamless integration with data security features between your on-premises IT environment and the AWS storage infrastructure. You can use the service to store data in the AWS Cloud for scalable and cost-effective storage that helps maintain data security.

AWS Storage Gateway offers file-based, volume- based, and tape-based storage solutions:

**File Gateway** – A file gateway supports a file interface into Amazon Simple Storage Service (Amazon S3) and combines a service and a virtual software appliance. By using this combination, you can store and retrieve objects in Amazon S3 using industry-standard file protocols such as Network File System (NFS) and Server Message Block (SMB). The software appliance, or gateway, is deployed into your on-premises environment as a virtual machine (VM) running on VMware ESXi or Microsoft Hyper-V hypervisor.

**Volume Gateway** – A volume gateway provides cloud-backed storage volumes that you can mount as Internet Small Computer System Interface (iSCSI) devices from your on-premises application servers.

**Tape Gateway** – With a tape gateway, you can cost-effectively and durably archive backup data in Glacier. A tape gateway provides a virtual tape infrastructure that scales seamlessly with your business needs and eliminates the operational burden of provisioning, scaling, and maintaining a physical tape infrastructure.

You can run AWS Storage Gateway either on- premises as a VM appliance, or in AWS as an Amazon Elastic Compute Cloud (Amazon EC2) instance. You deploy your gateway on an EC2 instance to provision iSCSI storage volumes in AWS. Gateways hosted on EC2 instances can be used for disaster recovery, data mirroring, and providing storage for applications hosted on Amazon EC2.

## III. Design

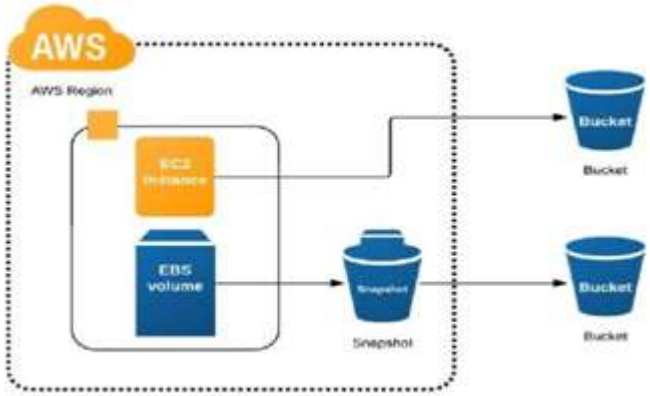


**Fig. AWS architecture**

This is the simple structure of **AWS EC2**, where **EC2** stands for Elastic Compute Cloud. EC2 allow users to use virtual machines of different configurations as per their requirement. It allows various configuration options, mapping of individual server, various pricing options, etc.

**Load balancing** simply means to hardware or software load over web servers, that improver's the efficiency of the server as well as the application. Following is the diagrammatic representation of AWS architecture with load balancing.
Hardware load balancer is a very common network appliance used in traditional web application architectures.
AWS provides the Elastic Load Balancing service, it distributes the traffic to EC2 instances across multiple available sources, and dynamic addition and removal of Amazon EC2 hosts from the load-balancing rotation.

**Amazon Cloud-front**

It is responsible for content delivery, i.e. used to deliver website. It may contain dynamic, static, and streaming content using a global network of edge locations. Requests for content at the user's end are automatically routed to the nearest edge location, which improves the performance.

Amazon Cloud-front is optimized to work with other Amazon Web Services, like Amazon S3 and Amazon EC2. It also works fine with any non- AWS origin server and stores the original files in a similar manner.

In Amazon Web Services, there are no contracts or monthly commitments. We pay only for as much or as little content as we deliver through the service.

**Elastic Load Balancer**

Elastic Load Balancer is basically is used to distribute your work load among a number of instances and equally among the instances, so that the work is done efficiently and also the work is consistent.
Ex. Joy is using Lobo website and Lobo are using same website, Lobo and joy will experience the same kind of time, because it being distributed equally among the instances and the instances are busy on the same level.

Elastic Load Balancing can dynamically grow and shrink the load-balancing capacity as per the traffic conditions.

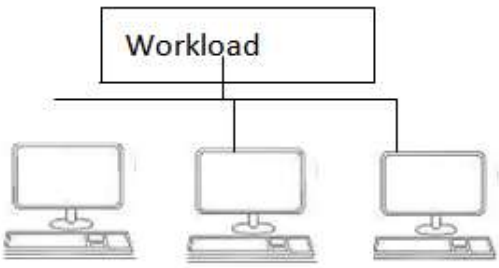


**Fig. Elastic Load Balancer Security Management**

Amazon's Elastic Compute Cloud (EC2) provides a feature called security groups, which is similar to an inbound network firewall, in which we have to specify the protocols, ports, and source IP ranges that are allowed to reach your EC2 instances.

Each EC2 instance can be assigned one or more security groups, each of which routes the appropriate traffic to each instance. Security groups can be configured using specific subnets or IP addresses which limits access to EC2 instances.

**Elastic Caches**

Amazon Elastic Cache is a web service that manages the memory cache in the cloud. In memory management, cache has a very important role and helps to reduce the load on the services, improves the performance and scalability on the database tier by caching frequently used information.

**Amazon RDS**

Amazon RDS (Relational Database Service) provides a similar access as that of MySQL, Oracle, or Microsoft SQL Server database engine. The same queries, applications, and tools can be used with Amazon RDS.

It automatically patches the database software and manages backups as per the user's instruction. Italso supports point-in-time recovery. There are no up-front investments required, and we pay only for the resources we use.

**Hosting RDMS on EC2 Instances**

Amazon RDS allows users to install RDBMS (Relational Database Management System) of your choice like MySQL, Oracle, SQL Server, DB2, etc. on an EC2 instance and can manage as required.

Amazon EC2 uses Amazon EBS (Elastic Block Storage) similar to network-attached storage. All data and logs running on EC2 instances should be placed on Amazon EBS volumes, which will be available even if the database host fails.

Amazon EBS volumes automatically provide redundancy within the availability zone, which increases the availability of simple disks. Further if the volume is not sufficient for our databases needs, volume can be added to increase the performance for our database.

Using Amazon RDS, the service provider manages the storage and we only focus on managing the data.

**Server Capacity**

AWS is providing best server capacity. All AWS competitors combined say server capacity of x and alone AWS server capacity is 6x.

**Storage & Backups**

AWS cloud provides various options for storing, accessing, and backing up web application data and assets. The Amazon S3 (Simple Storage Service) provides a simple web-services interface that can be used to store and retrieve any amount of data, at any time, from anywhere on the web.

Amazon S3 stores data as objects within resources called **buckets**. The user can store as many objects as per requirement within the bucket, and can read, write and delete objects from the bucket.

Amazon EBS is effective for data that needs to be accessed as block storage and requires persistence beyond the life of the running instance, such as database partitions and application logs.

Amazon EBS volumes can be maximized up to 1 TB, and these volumes can be striped for larger volumes and increased performance. Provisioned IOPS volumes are designed to meet the needs of database workloads that are sensitive to storage performance and consistency.

Amazon EBS currently supports up to 1,000 IOPS per volume. We can stripe multiple volumes together to deliver thousands of IOPS per instance to an application.

**AWS pricing**

AWS offers you a pay-as-you-go approach for pricing for over 120 cloud services. With AWS you pay only for the individual services you need, for as long as you use them, and without requiring long-term contracts or complex licensing. AWS pricing is similar to how you pay for utilities like water and electricity. You only pay for the services you consume, and once you stop using them, there are no additional costs or termination fees.

**Acknowledgment**

The authors thank all reviewers for improving this paper.

## IV. Conclusion

This paper explained about the concept of cloud computing. Amazon Web Service (AWS) is a cloud that provides secure and efficient data storage system. This paper proposed the different services provided by the AWS cloud for the data storage and replication system. These services have many advantages like it speed up the distribution of web content, it provides durable storage with security features for data archiving and backup, etc. Along with this the architecture of AWS cloud is also explained with the help of different components.

## References


[1]. CLOUD COMPUTING SECURITY:AMAZON WEB SERVICE Saakshi NarulaM.tech Student Department of CSE & IT ITM University Gurgaon ,India

[2]. Exploring the Cloud from Passive Measurements: the Amazon AWS CaseIgnacio Bermudez, Stefano Traverso, Marco Mellia, Maurizio Munafˆ o DET, Politecnico di Borislav S. Ĉorÿeviü, Slobodan P. Jovanoviü and Valentina V. Timþenko, Member, IEEE Torino, Italy – {lastname}@tlc.polito.it

[3]. Microsoft and Amazon A comparison of approaches to cloud security Golnoosh Tajadod School of Information Technology Deakin University Melbourne, Australia gtaj@deakin.edu.au

[4]. Cloud Computing in Amazon and Microsoft Azure platforms: performance and service comparison Ibrahim Ejdayid A.Mansour Dept. of Computing & Informatics Bournemouth University Bournemouth, UK imansour@bournemouth.ac.uk

[5]. Cloud Computing Services and Applications B.Sai.Karthik*M. Nalini Sri ECM&K L University, India.

[6]. Database Security Management for Healthcare SaaS in the Amazon AWS ClouFabio Bracci, Antonio Corradi, Luca Foschini Dipartimento di Elettronica, Informatica e Sistemistica (DEIS) University of Bologna, Italy {fabio.bracci, antonio.corradi, luca.foschini}@unibo.it d

[7]. Comparative study of Amazon EC2 and Microsoft Azure cloud architecture Prof Vaibhav A Gandhi Research Scholar, Dept of Computer Science, Sau Uni, & Associate Professor, Dept of MCA,B H Gardi College of Engg & Tech, Rajkot, India

[8]. gandhi.vaibhav@gmail.com

[9]. CLOUD COMPUTING CHALLENGES WITH EMPHASIS ON AMAZON EC2 AND WINDOWS AZURE Azzam Sleit*, Nada Misk1, Fatima Badwan1, Tawfiq Khalil2

[10]. A Security Analysis of Amazon's Elastic Compute Cloud Service – Long Version –

[11]. Marco Balduzzi EURECOM marco.balduzzi@madlab.it

[12]. "CLOUD COMPUTING-INFRASTRUCTURE AS SERVICEAMAZON EC2"Gurudatt Kulkarni1, Ramesh Sutar 2 Jayant Gambhir 3