Is Amazon Web Service (AWS) cloud supports best cost effective & high performance modern disaster recovery.

Name: Rajib Das

Employee ID- tcs 231462

ISU-CSP

Email-rajib2.d@tcs.com

Introduction:

Disaster recovery (DR) is about preparing for and recovering from a disaster. Any event that has a negative impact on a company's business continuity or finances could be termed a disaster. This includes hardware or software failure, a network outage, a power outage, physical damage to a building like fire or flooding, human error, or some other significant event. To minimize the impact of a disaster, companies invest time and resources to plan and prepare, to train employees, and to document and update processes. The amount of investment for DR planning for a particular system can vary dramatically depending on the cost of a potential outage. Companies that have traditional physical environments typically must duplicate their infrastructure to ensure the availability of spare capacity in the event of a disaster. The infrastructure needs to be procured, installed, and maintained so that it is ready to support the anticipated capacity requirements. During normal operations, the infrastructure typically is under-utilized or over-provisioned. With Amazon Web Services (AWS), companies can scale up its infrastructure on an as-needed, pay-as-you-go basis. AWS also gives the flexibility to quickly change and optimize resources during a DR event, which can result in significant cost savings.

Cloud infrastructure is gaining in popularity for organizations that have very real business needs that rely on information technology (IT) infrastructure but struggle with the costs —both capital and operational—of expanding their data centers. A computing model in which the equipment—including servers, storage, and networking components—used to support an organization's IT operations is hosted by a service provider and made available to customers over a network, typically the Internet. The service provider owns the equipment and is responsible for housing, running, and maintaining it, with the client typically paying on a per-use basis.

Today, more than three-quarters (82%) of organizations have plans to leverage cloud-based services to some extent over the next five years. DR is an ideal use case for taking advantage of the cloud. While many organizations remain cautious about placing production services in the cloud, they are often more comfortable testing those waters for DR—especially since the cloud alters the economics of DR so radically. Some organizations implement DR only for their most critical applications to minimize risk and keep expenses in check. Using the cloud enables companies to extend DR services to additional workloads, further reducing their exposure to business interruption. Others that currently operate failover sites are finding the costs skyrocketing because of continual data growth. But for many (particularly smaller) organizations, the cloud actually makes DR possible for the first time.

A traditional approach to DR involves different levels of off-site duplication of data and infrastructure. Critical business services are set up and maintained on this infrastructure and tested at regular intervals.

Infrastructure required to support the duplicate environment should include

- Facilities to house the infrastructure which including power and cooling.
- Physical protection of assets.
- Repairing, replacing, and refreshing the infrastructure support.
- Contractual agreements with an Internet service provider (ISP) for Internet connectivity which can sustain bandwidth utilization for the environment under a full load.
- Network infrastructure components such as firewalls, routers, switches, and load balancers.
- Capacity planning to run all mission-critical services, including storage appliances for the supporting data, and servers to run applications and backend services such as user authentication, Domain Name System (DNS), Dynamic Host Configuration Protocol (DHCP), monitoring, and alerting.

AWS supports Infrastructure-as-a-Service (IaaS) which is a collection of modular cloud service components, that is well known for scalability, reliability, and security. These modular services can be used independently or combined to meet specific computing and storage requirements. AWS services make an option to seed the cloud repository is to import large amounts of data into AWS using portable storage devices transported via third-party logistics. Data can be exported in a similar fashion too. Ongoing data movement thereafter occurs using network links.
Customer can find that cost reduction is achieved in a number of ways. There are no fees for inbound data transfer, which eliminates costs associated with moving data to a secondary and/or tertiary site. Also AWS' pay-as-you-go pricing reduces monthly costs since only resources actually used are paid for.

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 failover 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.

Major features and benefits of using AWS cloud for disaster recovery.

- **Fast Performance:** Fast disk-based storage and retrieval of files.
- **No Tape:** Eliminate costs associated with transporting, storing, and retrieving tape media and associated tape backup software.
- **Compliance:** Fast retrieval of files allows you to avoid fines for missing compliance deadlines.
- **Elasticity:** Add any amount of data, quickly. Easily expire and delete without handling media.
- **Secure:** Secure and durable cloud disaster recovery platform with industry-recognized certifications and audits.
- **Partners:** AWS solution provider and system integration partners to help with your deployment.

AWS Services Essential for Disaster Recovery

- **Regions**

Amazon Web Services are available in multiple regions around the globe, customer can choose the most appropriate location for your DR site, in addition to the site where system is fully deployed. AWS has multiple general purpose regions in the Americas, EMEA, and Asia Pacific that anyone with an AWS account can access.

- **Storage**

Amazon Simple Storage Service (Amazon S3) provides a highly durable storage infrastructure designed for mission-critical and primary data storage. Objects are redundantly stored on multiple devices across multiple facilities within a region, designed to provide a durability of 99.999999999% (11 9s). AWS provides further protection for data retention and archiving through versioning in Amazon S3, AWS multi-factor authentication (AWS MFA), bucket policies, and AWS Identity and Access Management (IAM).

Amazon Glacier provides extremely low-cost storage for data archiving and backup. Objects (or archives, as they are known in Amazon Glacier) are optimized for infrequent access, for which retrieval times of several hours are adequate. Amazon Glacier is designed for the same durability as Amazon S3.

Amazon Elastic Block Store (Amazon EBS) provides the ability to create point-in-time snapshots of data volumes. You can use the snapshots as the starting point for new Amazon EBS volumes, and you can protect your data for long-term durability because snapshots are stored within Amazon S3. After a volume is created, you can attach it to a running Amazon EC2 instance. Amazon EBS volumes provide off-instance storage that persists independently from the life of an instance and is replicated across multiple servers in an Availability Zone to prevent the loss of data from the failure of any single component.

AWS Import/Export bypasses the Internet and transfers your data directly onto and off of storage devices by means of the high-speed internal network of Amazon. For data sets of significant size, AWS Import/Export is often faster than Internet transfer and more cost effective than upgrading your connectivity. Customer can use AWS Import/Export to migrate data into and out of Amazon S3 buckets and Amazon Glacier vaults or into Amazon EBS snapshots.

AWS Storage Gateway is a service that connects an on-premises software appliance with cloud-based storage to provide seamless and highly secure integration between your on-premises IT environment and the storage infrastructure of AWS.

No-Fee Data Import

Some of the biggest expenses of disaster recovery are those associated with physical secondary sites. When transitioning to cloud-based DR, moving backup data can represent a large up-front cost. However, AWS has eliminated fees for data import into Amazon S3. This can represent a significant cost savings at the onset, as well as long-term.

Amazon S3 is a cloud-based object store available that is through Web services interfaces. It is used as a cloud storage container for backup data and images. Organizations can write, read, and delete virtually an unlimited number of objects containing from one byte to 5 TB of data each. Amazon S3 is similar to a traditional on-premise SAN or NAS device. Amazon S3 is designed to deliver “11 nines” of durability per year—this is accomplished by automatically making redundant copies in multiple Availability Zones, reducing the chance of data loss to one in 150 billion. Amazon S3 is also designed to offer 99.99% availability of objects, equal to just under one hour of *yearly* downtime.

Compute

Amazon Elastic Compute Cloud (Amazon EC2) provides resizable compute capacity in the cloud. Quick way to create Amazon EC2 instances, which are virtual machines over which you have complete control. In the DR, the ability

to rapidly create virtual machines that user can control is critical.
Amazon Machine Images (AMIs) are preconfigured with operating systems, and some preconfigured AMIs might also include application stacks. Anyone can also configure your own AMIs. In the context of DR, it is recommend to configure and identify user created own AMIs so that they can launch as part of your recovery procedure. Such AMIs should be preconfigured with your operating system of choice plus appropriate pieces of the application stack.
Availability Zones are distinct locations that are engineered to be insulated from failures in other Availability Zones. They also provide inexpensive, low-latency network connectivity to other Availability Zones in the same region. By launching instances in separate Availability Zones, application can be protected from the failure of a single location.

Networking
In disaster recovery AWS offers several services and features that enable you to manage and modify network settings.
Amazon Route 53 is a highly available and scalable Domain Name System (DNS) web service. It gives developers and businesses a reliable, cost-effective way to route users to Internet applications. Amazon Route 53 includes a number of global load-balancing capabilities (which can be effective when you are dealing with DR scenarios such as DNS endpoint health checks) and the ability to failover between multiple endpoints and even static websites hosted in Amazon S3.
Elastic IP addresses are static IP addresses designed for dynamic cloud computing.

Elastic Load Balancing automatically distributes incoming application traffic across multiple Amazon EC2 instances. It enables you to achieve even greater fault tolerance in your applications by seamlessly providing the load-balancing capacity that is needed in response to incoming application traffic. Just as you can pre-allocate Elastic IP addresses, you can pre-allocate your load balancer so that its DNS name is already known, which can simplify the execution of your DR plan.
Amazon Virtual Private Cloud (Amazon VPC) lets you provision a private, isolated section of the AWS cloud where you can launch AWS resources in a virtual network that you define. You have complete control over your virtual networking environment, including selection of your own IP address range, creation of subnets, and configuration of route tables and network gateways. This enables you to create a VPN connection between your corporate data center and your VPC, and leverage the AWS cloud as an extension of your corporate data center. In the context of DR, you can use Amazon VPC to extend your existing network topology to the cloud; this can be especially appropriate when recovering enterprise applications that are typically on the internal network.
Amazon Direct Connect makes it easy to set up a dedicated network connection from your premises to AWS. In many cases, this can reduce your network costs, increase bandwidth throughput, and provide a more consistent network experience than Internet-based connections.

Databases
Amazon Relational Database Service (Amazon RDS) makes it easy to set up, operate, and scale a relational database in the cloud. You can use Amazon RDS either in the preparation phase for DR to hold your critical data in a database that is already running, or in the recovery phase to run your production database. When you want to look at multiple regions, Amazon RDS gives you the ability to snapshot data from one region to another, and also to have a read replica running in another region.

Amazon DynamoDB is a fast, fully managed NoSQL database service that makes it simple and cost-effective to store and retrieve any amount of data and serve any level of request traffic. It has reliable throughput and single-digit, millisecond latency. You can also use it in the preparation phase to copy data to DynamoDB in another region or to Amazon S3. During the recovery phase of DR, you can scale up seamlessly in a matter of minutes with a single click or API call.

Amazon Redshift is a fast, fully managed, petabyte-scale data warehouse service that makes it simple and cost-effective to efficiently analyze all your data using your existing business intelligence tools. You can use Amazon Redshift in the preparation phase to snapshoot your data warehouse to be durably stored in Amazon S3 within the same region or copied to another region. During the recovery phase of DR, you can quickly restore your data warehouse into the same region or within another AWS region.

AWS Import/Export
Large amounts of data can be imported into AWS (and exported from it) using subscriber's portable storage devices transported via third-party logistics via the AWS Import/Export option. AWS transfers data directly into Amazon S3 using AWS high-speed internal network and bypassing the Internet. For large data sets, this is often more rapid than Internet transfer and could help avoid the need to upgrade connectivity.

Security and Compliance

Security concerns are often a top barrier for those considering cloud implementations. Organizations depend on their data and applications to run businesses, and their level of trust that their data is secure from intrusion, as well as fully protected from loss, are critical.

AWS starts with shared responsibility between the subscriber and AWS, enabling flexibility and the levels of customer control required for certain industry-specific compliance requirements. AWS operates and manages the physical infrastructure and its security, as well as host operating systems and the virtualization layer. Customers assume the responsibility for guest operating systems and application software, including updates and security patches. Customers are also responsible for configuring the AWS-provided firewall. Subscribers can enhance security as they choose to meet more stringent compliance requirements with additional host-based firewalls and intrusion detection features, as well as encryption and encryption key management.
AWS services are built on an environment with extensive and validated security and controls, including:

1. Service Organization Controls 1 (SOC 1) Type 2 report[10] (formerly SAS 70[11] Type II report), with periodic independent audits to confirm security features and controls that safeguard customer data.
2. ISO 270001 Certification, an internationally-recognized security management standard that specifies leading practices and comprehensive security controls following the ISO 27002 best practice guidelines.
3. PCI DSS compliance, an independent validation of the platform for the secure use of processing, transmitting, and storing credit card data.
4. Relevant government agency and public sector compliance qualifications, such as an ITAR-[13]compliant environment.
5. To support customers with FIPS 140-2[14] requirements, the Amazon VPC VPN endpoints and SSL-terminating load balancers in AWS GovCloud (US) operate using FIPS 140-2-validated hardware.

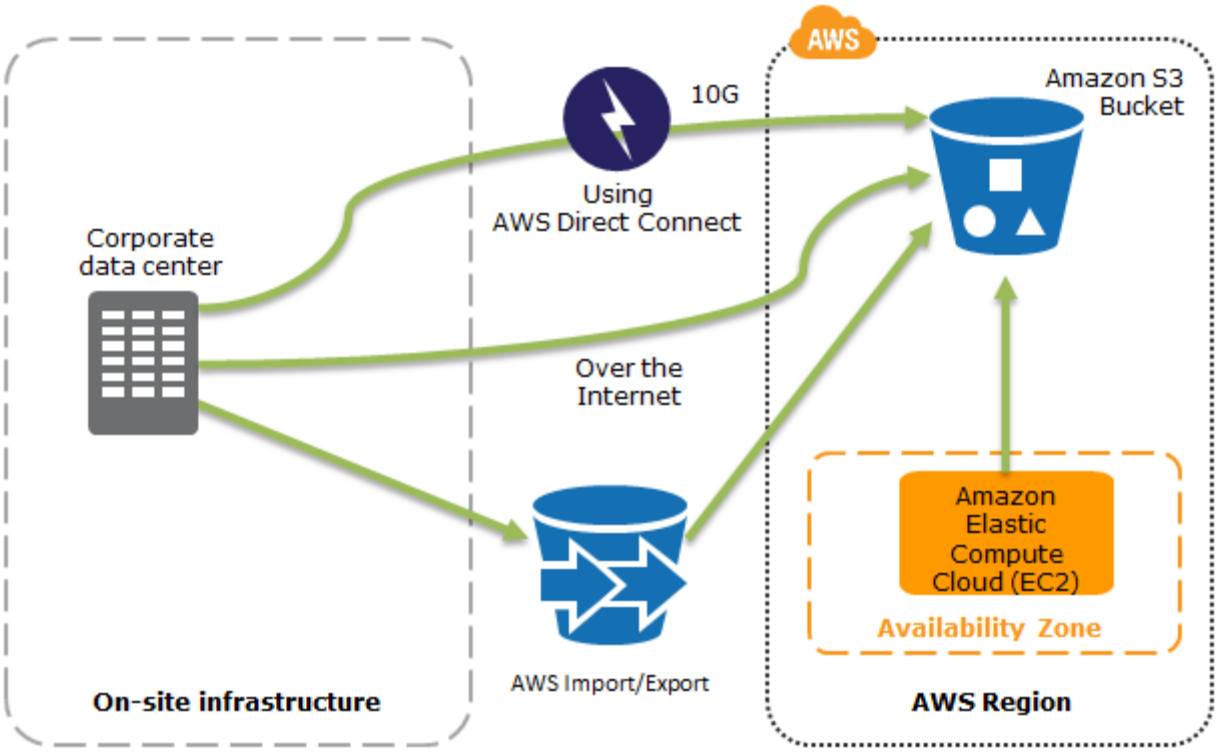

The following figure shows data backup options to Amazon S3, from either on-site infrastructure or from AWS

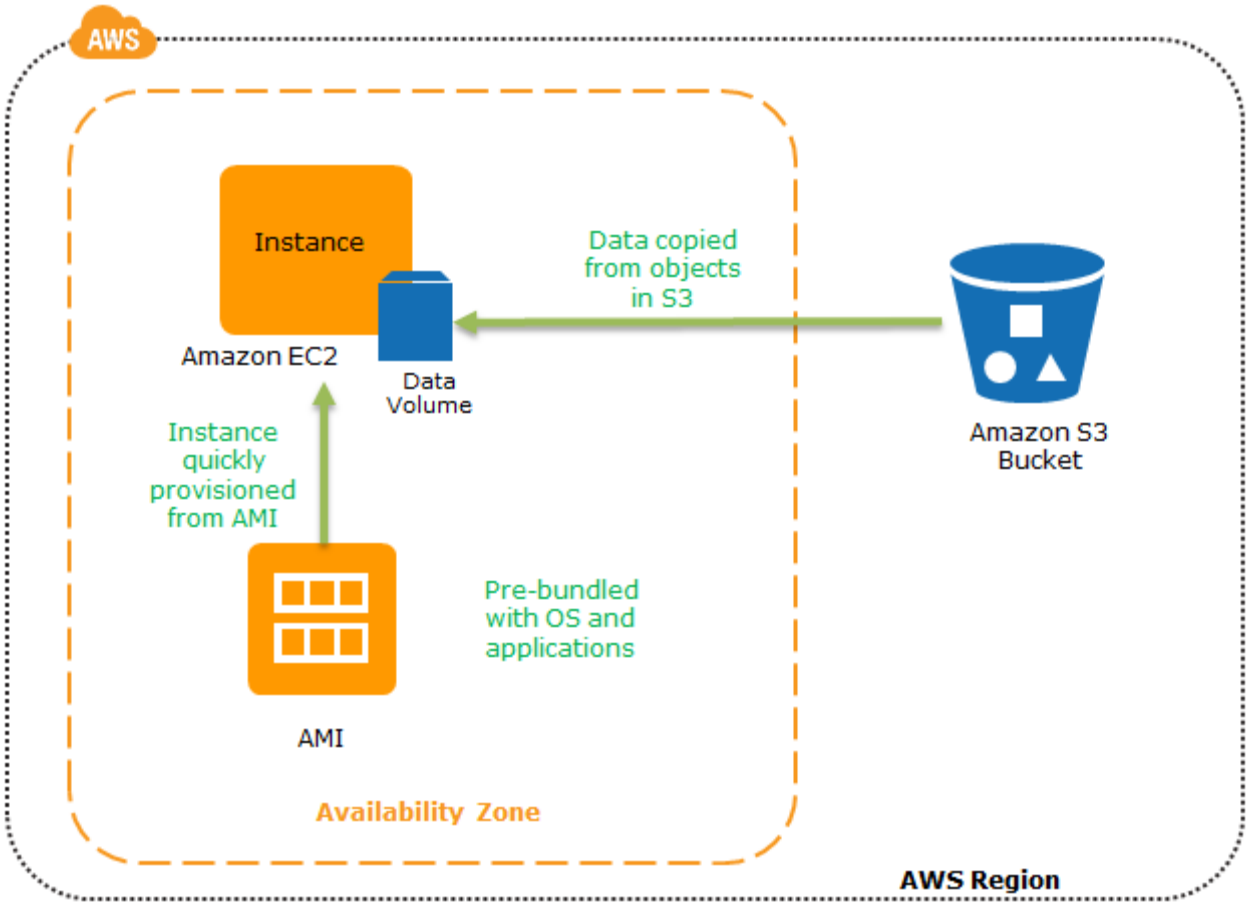

This figure shows how you can quickly restore a system from Amazon S3 backups to Amazon EC2.

Improving Your DR Plan

After your DR solution is in place, it needs to be tested. You can test frequently, which is one of the key advantages of deploying on AWS. “Game day” is when you exercise a failover to the DR environment, ensuring that sufficient documentation is in place to make the process as simple as possible should the real event take place. Spinning up a duplicate environment for testing your game-day scenarios is quick and cost-effective on AWS, and you typically don’t need to touch your production environment. You can use AWS Cloud Formation to deploy complete environments on AWS. This uses a template to describe the AWS resources and any associated dependencies or runtime parameters that are required to create a full environment.
Differentiating your tests is key to ensuring that you are covered against a multitude of different types of disasters. The following are examples of possible game-day scenarios:

- Power loss to a site or a set of servers
- Loss of ISP connectivity to a single site
- Virus impacting core business services that affects multi-sites
- User error that causes the loss of data, requiring a point-in-time recovery

Monitoring and alerting
You need to have regular checks and sufficient monitoring in place to alert you when your DR environment has been impacted by server failure, connectivity issues, and application issues. Amazon Cloud Watch provides access to metrics about AWS resources, as well as custom metrics that can be application–centric or even business-centric. You can set up alarms based on defined thresholds on any of the metrics and, where required, you can set up Amazon SNS to send alerts in case of unexpected behavior.

Backups
After you have switched to your DR environment, you should continue to make regular backups. Testing backup and restore regularly is essential as a fall-back solution.
AWS gives you the flexibility to perform frequent, inexpensive DR tests without needing the DR infrastructure to be "always on."
User access
To have secure access to resources in your DR environment by using AWS Identity and Access Management (IAM). With IAM, customer can create role-based and user-based security policies that segregate user responsibilities and restrict user access to specified resources and tasks in your DR environment.

Software Licensing and DR

AWS provides a variety of models to make licensing easier for manage. For example, "Bring Your Own License" is possible for several software components or operating systems. Alternately, there is a range of software for which the cost of the license is included in the hourly charge. This is known as "License included."
"Bring your Own License" enables you to leverage your existing software investments during a disaster. "License included" minimizes up-front license costs for a DR site that doesn't get used on a day-to-day basis.

AWS Storage Gateway
The new AWS Storage Gateway is a service that connects an on-premise software appliance with cloud-based storage. The appliance is downloaded from the Amazon website and launched on an on-premise virtualization host.
The AWS Storage Gateway enables customers to use existing applications to store data on Amazon S3 by exposing a standard iSCSI interface to the customer's on-premise application server. The customer points the application to the Gateway, which will store a primary copy locally on their on-premise storage (SAN or DAS) and also store a snapshot in the cloud on Amazon S3, as an Amazon EBS snapshot. Subsequent snapshots will store only differential changes.
Should a disaster occur, customers can restore their applications using Amazon EC2, pull their snapshots into Amazon EBS, and have the application up and running right away. Now, instead of paying for servers

and infrastructure that sits idle in a disaster recovery site, customers only pay for snapshots stored in S3. If they want to recover the application in the cloud using Amazon EC2 and Amazon EBS, those subscription-based fees are only incurred when needed in the case of a disaster.

No-Fee Data Import

Some of the biggest expenses of disaster recovery are those associated with physical secondary sites. When transitioning to cloud-based DR, moving backup data can represent a large up-front cost. However, AWS has eliminated fees for data import into Amazon S3. This can represent a significant cost savings at the onset, as well as long-term.

Conclusion

The AWS services are available on-demand, and you pay only for what you use. This is a key advantage for DR, where significant infrastructure is needed quickly, but only in the event of a disaster.