DIGITAL & TECHNOLOGY SERVICES

Cloud Disaster Recovery

Looking Beyond Regional Deployments

Introduction

Enterprise Disaster Recovery strategy has shifted in terms of how to meet the demand and availability of an information technology (IT) infrastructure. As more and more companies are adopting cloud, chief information officers (CIOs) and IT leaders are looking at cloud regional deployment as a way of achieving their disaster recovery (DR) requirements and meeting their recovery service level agreements (SLAs).

When creating a DR strategy, careful consideration must be given to ensure business continuity. An IT organization must evaluate their recovery time objective (RTO) and recovery point objective (RPO) requirements and then ensure that the cloud architecture can match these requirements across all disaster types within the required recovery SLAs of the organization.

This paper evaluates key considerations that an organization must make when building an IT/application architecture on public cloud and debunks some common myths surrounding cloud availability and the SLAs provided by cloud providers, enabling organizations to understand some of the key issues that are often missed in all of the marketing speak.

Rising Cloud Adoption

There is no denying that cloud adoption is on the rise. Worldwide end-user spending on public cloud services is forecasted to grow 20.7 percent to $591.8 billion in the current year, up from $490.3 billion in 2022, according to the latest forecast from Gartner, Inc.[1] It is estimated that more than 85 percent of organizations will embrace cloud-first principle by 2025. Cloud offers a way to scale fast, reduce costs, increase productivity, and reduce operations and management overhead. While moving to cloud is practically inevitable, it will necessitate companies rethink their disaster recovery approach and finetune their IT architecture to handle outages.

Cloud Myths Debunked

The trend of cloud adoption is rapidly growing due to the many advantages that cloud computing offers over traditional on-premises approaches. However, not all clouds are created equal and not every cloud has a silver lining.

"My IT deployment is on the public cloud so I don't need disaster recovery."

Moving to the cloud is not a panacea. If there are concerns about data integrity and availability risks, moving to the cloud will not instantly mitigate them. The reality is that no cloud is immune to downtime, and most organizations are less protected than they think.

As part of a cloud evaluation process, it is essential to deploy a parallel DR solution that meets an organization's RTP/RPO needs. Otherwise, the organization is responsible for any issues that may arise — not the cloud service provider (CSP).

CSPs do not typically include resilient DR as part of their stack. Instead, they recommend that organizations implement out-of-region DR for any business-critical applications. This is very important and often lost in glossy marketing speak. Many organizations that are already running their IT infrastructure on cloud are blissfully unaware of how vulnerable they really are, specifically if their business demands 24/7 availability to service their end customers. This leads to the next problematic statement.

"My IT is safe because I have regional high availability (multi-AZ deployments)."

1. Gartner, *Inflationary Pressures Creating a Push and Pull Effect for Cloud Spending, 2022*

Most of the CSPs highlight features like "in-region high availability" as a solution to most disaster recovery requirements. While in-region high availability can offer a good degree of resilience, the protection it offers is far from robust. So much so, that if you dig into the small print, most of the CSPs recommend an out-of-region standby environment for business-critical infrastructure and applications to meet RTO/RPO requirements.

Why? The truth is that big is not always better, and it is not uncommon for an entire region to go down rendering in-region high availability null and void.

Even instances where only a particular service goes offline within the region can impact application availability.

For a recent example, one only needs to look to Microsoft Azure Cloud's lengthy June 2022 outage. For 12 hours, customers had trouble connecting to US East 2 region. The reason provided was "an unplanned power oscillation in one of our datacenters within one of our availability zones in the East US 2 region," according to a Microsoft report.[2]

Other recent examples of public cloud outages impacting complete regions and bringing down many services are listed below.

November 2021
Google Cloud experienced two hours of downtime caused by a networking glitch.

December 2021
Amazon's entire US East Region experienced a major outage for five hours due to a networking issue.

March 2022
Users of Google's Traffic Director tool experienced an elevated service error for approximately three hours.

June 2022
Azure's East 2 Region experienced a cooling failure caused by a power oscillation, resulting in a 12-hour outage of storage and compute resources.

July 2022
AWS US East 2 region had regional connectivity issues that brought down services for 20 minutes.

December 2022
AWS US East 2 region had regional connectivity issues that brought down services for 75 minutes.

While CSPs offer "guaranteed uptime and availability," what they are actually guaranteeing is the provision of small rebates in an organization's monthly bill if SLAs are not met. It is important to understand the differentiation here: CSPs are not guaranteeing that a service will be available or ensuring access to an organization's data; they are simply offering modest refunds should downtime or a lack of availability happen.

In addition, quality service and uptime do not mean the same thing. While a region may be available, a host of technical reasons could mean that the service/user experience is temporarily degraded. If AWS S3 service faces an outage/degraded service, it could significantly impact an organization's business functionality. Furthermore, organization-specific networking and internet issues can and do frequently interrupt or degrade service in ways that materially impact the business.

A&M's recommendation is for the organization to evaluate their RTP/RPO requirements, and fully understand the cost of any downtime or service degradation to the business.

2. CRN, *The 10 Biggest Cloud Outages Of 2022 (So Far), 2022*

What approach is recommended by CSPs?

As previously discussed, while multi-AZ deployments improve resiliency by providing high-availability capabilities to protect from a data center outage, it does not protect from a region-wide outage or performance issues. A host of factors cause these outages – from natural disasters to cloud vendor misconfiguration to malicious attacks and others. For this reason, most CSPs recommend out-of-region replication as the basis of best-practice DR.

The graphics below look at two of the major CSPs and their approach to DR.

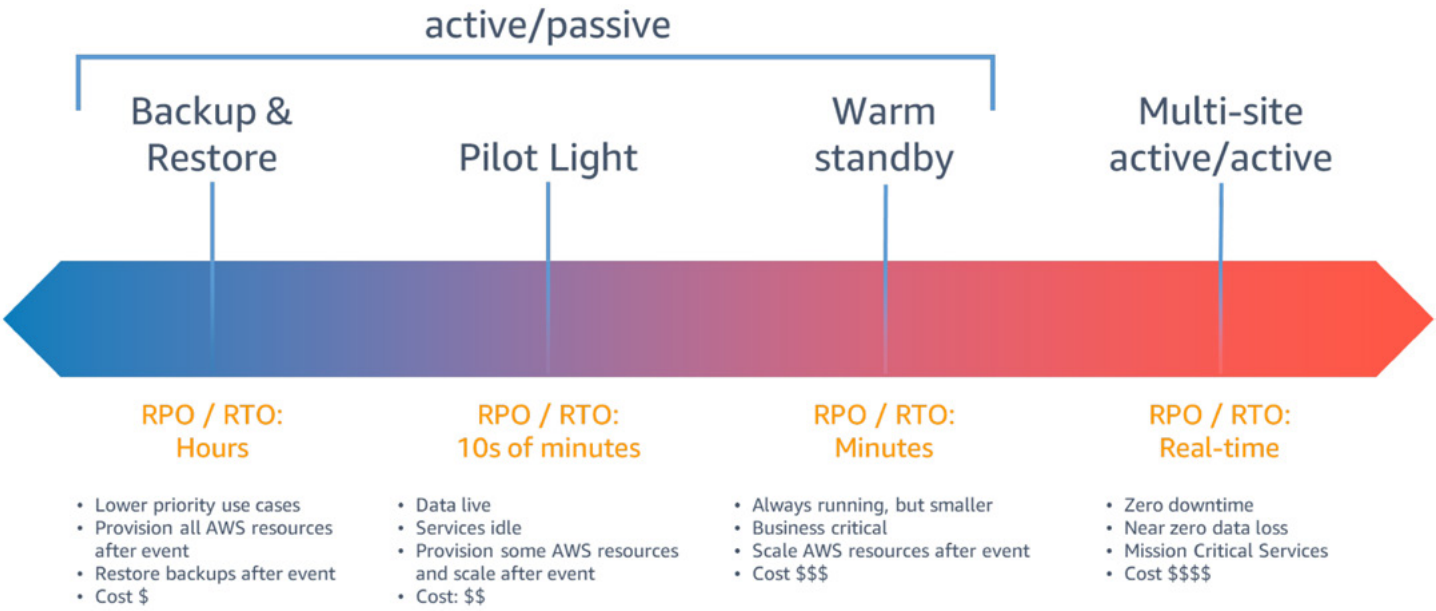

Source: AWS – Out of Region DR options in Cloud

Strategy	Description	Impacts	Speed to Recovery	Complexity to Execute	Complexity to Implement	Impact to Customers	Above line OPEX Cost
Wait for Microsoft	Wait for Microsoft to complete the recovery of services.	The solution would be completely offline until Microsoft recovers the impacted region and all required services/components.	●	○	●	●	●
Redeploy on Disaster	The solution is redeployed from scratch, post the disaster event.	The solution would be offline while Contoso completes failover activities into the secondary region, procuring new component instances , deploying the code base, and bringing the system up to date.	●	●	●	●	●
Warm Spare	A secondary hosted solution is created in an alternate region, and roles are deployed to guarantee minimal capacity; however, the roles don't receive production traffic.	The solution would be offline for a short period while Contoso completes failover activities into the secondary region and bringing the system up to date.	●	●	●	●	●
Hot Spare	The solution is hosted in an active/active setup across multiple regions, with both receiving, processing and serving data out.	The solution remains in service for customers with no impact resulting from the regional outage.	●	●	●	○	●

Source: Azure – Out of Region DR options in Cloud

Evaluating RTO/RPO Requirements

> It is important to understand an organization's RTO/RPO requirements and design the IT infrastructure and applications accordingly.
>
> **RTO: Recovery Time Objective –** refers to how much time an application can be down without causing any substantial damage to business.
>
> **RPO: Recovery Point Objective –** refers to company's loss tolerance: the amount of data that can be lost before causing significant harm to business.

When it comes to providing DR capabilities for environments with short RTO/RPO timelines, CSPs identify cross-region replication as a best practice. A highly available environment within the region can be sufficient for most business needs, but for critical 24/7 business services and applications, designing the right solution with data replication and system redundancy is the single most important piece of IT design for cloud. Achieving a near-zero RTO/RPO is not possible (and would come at an exorbitant cost). Businesses must set realistic recovery goals based on their budget, resources and business priorities.

Addressing Key Concerns of CIOs

With cloud, there are various options and services that enable an organization to build resilient infrastructure across region and be efficient both with cost and operations. Let's look at some of the challenges that CIOs and IT leaders face and the answers/solutions that cloud provides.

Do I need a like-for-like environment in another region?

Yes and no. Yes, in terms of application architecture and application configuration. No, for the size and quantity of services. With cloud offering auto-scaling and elasticity, you can build a smaller, trimmed down version of the production environment and scale out, as required, during a DR event.

Does the application require a high degree of changes to meet multi-region deployments?

The amount of required application reconfiguration depends on various factors like architecture of the application, cloud services that the application is using and others. For example, one of A&M's client's deployments only required changes.

Will my average monthly bill/cost double?

No. With cloud there are various services that come with zero additional cost while setting up as a multi-regional service. Also, with auto-scaling features and functions, there is no need to configure a 100 percent match of the production environment. One caveat: if the business requirement is an active/active configuration then yes, a higher cost is probable.

Do I need specialized skills to manage a multi-region environment?

Not usually. An organization should be able to continue to manage the ongoing operations of their DR environment. However, a trusted partner who has expertise and experience in helping an organization design, architect, build and manage the entire implementation can accelerate and de-risk execution.

FACILITY SERVICES PROVIDER

Implementing a Robust Regional Cloud Disaster Recovery Strategy on AWS for Tier-1 Applications

About the Customer and the Project

A privately owned service industry company had 95 percent of its business applications running on Amazon Web Services (AWS) with multi-AZ deployment. However, due to recent AWS regional outages, the company's online business was severely challenged during these disruptions, losing customers and revenue.

The company turned to A&M for strategic and technical support on building and implementing a regional disaster recovery (DR) strategy for their applications, which used multiple AWS services including Kubernetes, Lambda, KMS, Secrets, Databases, ECS, EC2, Route 53 and others.

Our Approach

A&M conducted a detailed analysis of the client's business applications and, after interviewing several key stakeholders in the company, created a map of the applications in different tiers and analyzed the critical Tier-1 applications experiencing negative impacts on revenue and customer satisfaction. These Tier-1 applications were selected by A&M to be included in the DR project, while any associated applications were kept out of scope to reduce implementation and AWS cost.

The proposed strategy was to build a regional DR model meeting two critical factors:

1. The ability to move applications from one region to another during a regional AWS outage with RPO of >= 5 min and RTO >= 30 minutes, and
2. The ability to recover quickly in case of a security breach and data loss in a primary region with RPO of >= 30 min and RTO >= 4 hours.

Our Challenges and Solution

One of the key challenges that A&M faced initially was understanding the complex nature of the applications' inter-connectivity and their associated AWS services, which were built by mainly using serverless architectures and native AWS services, including databases (Aurora and Postgre-SQL). Once the application architectures were untangled, our team mapped the individual AWS services and their associated restrictions to move from a single region model to multi-region service offerings. Each individual AWS service forming an application set was then targeted and migrated to the multi-regional model without impacting business operations. Stage 1 was comprised of setting and migrating AWS services for staging environment, building/modifying the services and testing for any issues in data replication, service recovery, recovery time objective and recovery point objective, making necessary adjustments to design and architecture to resolve those issues.

Following the successful deployment in the staging environment, Stage 2 began with moving production applications to the multi-regional model. A&M's team conducted production DR testing to demonstrate successful recovery of critical applications during an AWS regional outage.

Outcomes

A&M conducted multiple DR tests and achieved a 100 percent success rate. This was done in collaboration with client application teams for any application configuration changes and testing.

Additionally, A&M's engagement produced the following benefits to the client:

- Attained zero business impact to production during DR build and implementation;
- Secured the ability for the business and IT to continue their operations 24 hours a day and provide services to its customer continuously;
- Minimized the overall, ongoing AWS cost by leveraging free tiers, auto scaling groups and optimizing application architecture; and
- Improved the business's ability to improve the security posture of the AWS environment by setting up DR strategies for recovery from both regional outages and security breaches.

Conclusion

The cloud is not immune to disasters. With so much of an organization's business value tied up in its data, it is critical that this data be protected. While multi-AZ strategy might represent a "good enough" solution for customers to sustain and continue running their business, its multi-region deployments give complete recovery capabilities.

Contact us today to learn more about how A&M can help you deploy cloud services in a multi-regional model with disaster recovery services.

Contacts

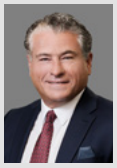

Joe Melfi
Managing Director

+1 917 743 4987
jmelfi@alvarezandmarsal.com

Amit Patel
Director

+1 818 770 2136
amit.patel@alvarezandmarsal.com

Follow A&M on:

ABOUT ALVAREZ & MARSAL

Companies, investors and government entities around the world turn to Alvarez & Marsal (A&M) for leadership, action and results. Privately held since its founding in 1983, A&M is a leading global professional services firm that provides advisory, business performance improvement and turnaround management services. When conventional approaches are not enough to create transformation and drive change, clients seek our deep expertise and ability to deliver practical solutions to their unique problems.

With over 7,000 people across five continents, we deliver tangible results for corporates, boards, private equity firms, law firms and government agencies facing complex challenges. Our senior leaders, and their teams, leverage A&M's restructuring heritage to help companies act decisively, catapult growth and accelerate results. We are experienced operators, world-class consultants, former regulators and industry authorities with a shared commitment to telling clients what's really needed for turning change into a strategic business asset, managing risk and unlocking value at every stage of growth.

To learn more, visit: AlvarezandMarsal.com