ActualTech MEDIA

SIOS

Ensuring High Availability for Essential Applications in a Multi-Cloud World

Lawrence Miller

CONTENTS

IN THIS PAPER

More and more organizations have deployed their mission-critical workloads across multiple public cloud providers. For distributed systems, such as databases and directory services, multi-cloud DR can be far more challenging. This paper discusses those unique challenges and how to address high availability and disaster recovery in multi-cloud environments.

Highlights include:

- How multi-cloud environments and strategies have evolved and become widely adopted
- The unique challenges organizations face with multi-cloud
- The importance of working with a trusted, cloud-agnostic partner to help design and implement a multi-cloud deployment using a holistic approach

Cloud computing has become ubiquitous over the last decade with 99% of organizations using at least one public or private cloud according to the Flexera 2021 State of the Cloud Report. While Amazon Web Services (AWS), Microsoft Azure, and Google Cloud Platform (GCP) are the top three public cloud providers today, many organizations—whether by design or by accident—have adopted a multi-cloud strategy that allows them to pick and choose which cloud services are most compelling and best suited to their unique business requirements. According to the Flexera report, 92% of enterprises today have a multi-cloud strategy and use an average of 2.6 public and 2.7 private clouds, including Software-as-a-Service (SaaS), Platform-as-a-Service (PaaS), and Infrastructure-as-a-Service (IaaS) offerings.

The Evolution (and Wide Adoption) of Multi-Cloud as a Strategy

A *multi-cloud* environment consists of a combination of any two or more public or private cloud offerings including SaaS, PaaS, and IaaS (see **Figure 1**). Thus, an organization's multi-cloud strategy may consist of an enterprise workload running on Amazon Elastic Cloud Compute (EC2) and using Microsoft 365 for email and back-office applications. Or an organization may connect a custom database hosted in a private cloud to Salesforce, a public cloud SaaS offering.

A *hybrid cloud* environment consists of a mix of on-premises, private cloud, and public cloud environments. According to the Flexera report, 80% of enterprises have a hybrid cloud strategy.

Multi-Cloud vs. Hybrid Cloud

A multi-cloud is simply an environment that consists of two or more public and/or private clouds (including SaaS, PaaS, and IaaS). The different services in a multi-cloud environment may interoperate (in which case it might be a hybrid cloud) or may not necessarily interoperate (essentially operating as separate cloud silos). Remember, although all hybrid clouds are multi-clouds, not all multi-clouds are hybrid clouds.

Multi-cloud environments often evolve as a result of shadow IT, in which different departments procure cloud services to meet their individual needs without necessarily consulting a centralized IT department. For example, your marketing team may have started using Salesforce long

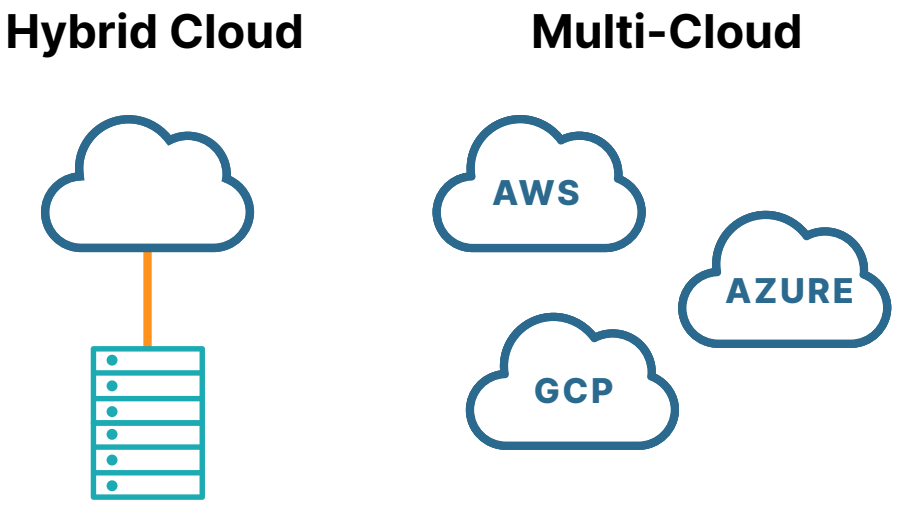

Figure 1: Hybrid cloud vs. multi-cloud

before IT deployed its first workload in AWS, while your HR and finance departments were busy adding Workday and Concur to the mix of SaaS applications that your organization now depends on. Or perhaps you have application development teams that work on different projects across the globe. One development team may prefer Azure DevOps, whereas another team may prefer the open source tools in AWS. Thus, your multi-cloud strategy may have evolved purely by accident—which isn't necessarily a bad thing.

Your different departments are empowered to select best-of-breed solutions to meet their needs while your app dev teams can maximize productivity and reduce time-to-market working in their preferred development environments.

Multi-cloud environments also evolve by design, for example, due to regulatory requirements, mergers and acquisitions, or to implement high availability (HA) and disaster recovery (DR) strategies.

> **Multi-cloud environments often evolve as a result of shadow IT, in which different departments procure cloud services to meet their individual needs without necessarily consulting a centralized IT department.**

Regulatory language can be vague and confusing. For example, the Financial Conduct Authority (FCA) regulations on outsourcing IT state that firms must be able to "know how they would transition to an alternative service provider and maintain business continuity." This statement implies that regulated firms need to at least plan for a secondary cloud environment. Given the risk-averse nature of many heavily regulated firms, these types of issues have led many to adopt a multi-cloud strategy.

Integrating IT systems and consolidating data centers and cloud environments after a merger or acquisition is a significant challenge. There are a number of factors that can complicate this challenge, including existing contracts with cloud providers or co-location providers. Similar to consolidating physical data centers, consolidating cloud workloads can be a major effort that doesn't deliver significant business value, so it's frequently delayed for higher-priority projects.

> **For distributed systems, such as databases and directory services (for example, Active Directory), multi-cloud DR can be far more challenging.**

Finally, multi-cloud strategies are often adopted to support HA and DR requirements. In evaluating major public cloud outages across AWS and Azure, most outages are typically limited to a single cloud region at a time (and are most commonly software-related).

More and more organizations (34% according to the Flexera report) have taken the added step of deploying their mission-critical workloads across multiple public cloud providers. This can be much easier for static workloads, such as websites and applications that can run independently of one another. For distributed systems, such as databases and directory services (for example, Active Directory), multi-cloud DR can be far more challenging.

Understanding Unique Challenges in Multi-Cloud Environments

Multi-cloud environments are more complex and thus more challenging to manage than single cloud deployments. Some unique challenges in multi-cloud environments include:

- **End-to-end visibility:** Ensuring complete visibility is a challenge in any IT environment—and it's exponentially more complex and challenging in a highly dynamic multi-cloud environment. However, end-to-end visibility is critical to troubleshooting performance

issues and bottlenecks, securing your digital footprint, and identifying single points of failure in mission-critical systems and applications.

- **Security and identity management:** Ransomware and other cybersecurity threats are top of mind for every IT leader today. While moving to a public cloud platform generally improves the security posture of an organization by shifting certain security responsibilities (such as data center and physical security) to the public cloud provider and providing on-demand access to services like encryption and network segmentation, it can also make it easier to make costly mistakes. For example, network misconfigurations can be common—thousands of data breaches have been caused by improperly configured AWS S3 storage buckets. Identity management is yet another challenge. For example, Azure Active Directory may be quite familiar to organizations that have previously used Active Directory in their on-premises environments, but extending identity management beyond Azure to AWS, GCP, and SaaS offerings (such as Salesforce, ServiceNow, Workday, and others) can introduce new challenges.

> The ability to dynamically move applications and data across different public cloud platforms in a hybrid (multi-cloud) environment is key to many multi-cloud strategies.

- **Application and data portability:** The ability to dynamically move applications and data across different public cloud platforms in a hybrid (multi-cloud) environment is key to many multi-cloud strategies. Although public cloud providers don't necessarily build their services to restrict application and data portability, they don't necessarily work together to facilitate this capability and there may be costs involved. Different cloud providers also use different technologies for their various service offerings.
- **Multi-cloud silos:** If organizations don't plan and design their multi-cloud deployments for application and data portability, they can end up with siloed applications and storage, essentially re-creating a common problem in traditional on-premises data center environments, across multiple cloud platforms. At the very least, organizations need multi-cloud security and management tools that allow them to effectively manage their risks and usage/costs across different cloud platforms.

Security and Cloud Spend Challenges

According to the Flexera 2021 State of the Cloud Report, 81% of organizations cite security as the top challenge in their cloud deployments, followed by managing cloud spend (79 %). Yet only 42% of organizations use multi-cloud cost management tools and only 38% use multi-cloud security tools.

Addressing High Availability and Disaster Recovery in Multi-Cloud Environments

While there are many challenges to multi-cloud deployments, they can provide additional availability, especially in the event of a major cloud outage, and DR. If your organization is pursuing a multi-cloud strategy, you should work with a trusted, cloud-agnostic partner to help you design and implement your multi-cloud deployment using a holistic approach.

> At the very least, organizations need multi-cloud security and management tools that allow them to effectively manage their risks and usage/costs across different cloud platforms.

For HA and DR, you also need a cloud-agnostic technology solution that spans your multi-cloud environment, irrespective of the cloud platforms you use. You always want to avoid a scenario where your HA solution causes more downtime in your environment than a standalone solution. Early versions of SQL Server clustering presented this conundrum—to add disk space, you had to incur downtime that wouldn't have occurred on a standalone solution.

> **If your organization is pursuing a multi-cloud strategy, you should work with a trusted, cloud-agnostic partner to help you design and implement your multi-cloud deployment using a holistic approach.**

While failing over something like a static website can be trivial, moving a multi-tier application stack is extremely complicated in terms of networking and data synchronization. You also need to avoid failing over to a less secure cloud environment that has potentially been misconfigured due to a lack of understanding the nuances between different security solutions across cloud providers.

Finally, in every public cloud, there are a handful of services that can increase costs quickly. These services are charged according to usage-based pricing and can mean steep cost increases after only a few days. One way to mitigate this risk is to ensure you're taking advantage of the cost monitoring services and alerts that are in each of your cloud platforms.

Understanding Your Hurdles

While multi-cloud deployments aren't for all organizations, many will go down this path. Understanding networking and security are among your biggest technical hurdles, and managing governance and costs are key functional challenges. Testing is critical to ensure your multi-cloud failback solution works. It's important to use an HA clustering solution that enables simple switchover and switchback and to understand how each of your applications will work with failover, and most importantly to regularly test that failover to understand any networking or data hurdles.

Learn more and schedule a personalized demonstration of SIOS at https://us.sios.com.