rubrik

Protecting Hybrid and Multicloud Data

How to bridge the gaps in data protection for seamless backup and recovery

Hybrid and Multicloud Environments: The New Normal

As enterprises pursue digital transformation, they are adopting the cloud for the agility, flexibility, and scale it offers. That often means attempting to choose between a variety of public clouds, while maintaining some on-premises infrastructure.

Today, hybrid or multicloud architectures are almost unavoidable. They occur sometimes by choice, but most often develop organically over time. Enterprises choose cloud providers for specific purposes based on their use cases, budget, technical and business requirements, geographic location, and other reasons. It's typical for an organization to have workloads across multiple clouds—preferring Google Cloud Platform for development and testing, let's say, or AWS for business analytics, or Microsoft Azure for disaster recovery.

Industry experts say hybrid and multicloud are fast becoming the new normal for enterprises.

81%

of respondents were already working with two or more providers*

93%

of organizations had adopted a multicloud approach**

87%

were taking a hybrid approach (using both public and private clouds)**

*Gartner, 2019 **Flexera, *State of the Cloud Report*, 2020

The move to hybrid and multicloud environments has brought with it a number of benefits, including:

- Flexibility to choose best-in-class providers for specific workloads or use cases
- Ability to find the most competitive pricing
- Avoidance of vendor lock-in
- Enhanced resilience
- Improved risk management, in the event one provider is compromised, for example

The new normal has given rise to new challenges, too—challenges that have significantly complicated the way enterprises protect their data, especially when it comes to backup and recovery.

Data Fragmentation Complicates Management and Protection

Simply put, the biggest challenge to data management and protection in hybrid and multicloud environments is **data fragmentation**.

Today, a preferred architecture has evolved in most organizations that looks primarily like this:

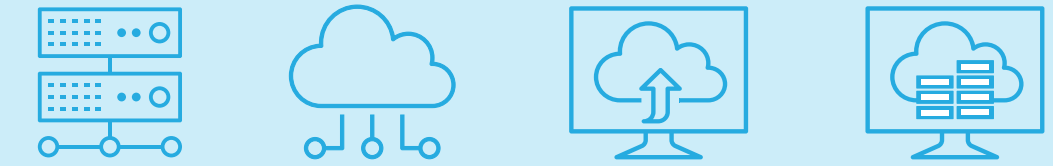

On-premises servers, storage, and networks	**CLOUD Software as a Service**	**CLOUD Platform as a Service (PaaS)**	**CLOUD Infrastructure as a Service (IaaS)**
High-value or high-risk workloads, such as financial data or intellectual property	Common business processes, such as CRM, marketing, and HR	Rapid app devlopment and testing before they are moved to best execution venue	Rapid provisioning of compute, storage, and network resources

Typically, these on-prem and cloud services operate independently of each other.

Source: Carl Lehmann, "Why is a multi-cloud approach gaining such popularity?" TechBeacon

Each cloud platform, as well as on-premises infrastructure, is siloed and that complicates the way you protect, manage, and secure your data, operations, and business. It can make governance, security, and compliance infinitely more complex.

That's a big problem. Even several years ago, some in the industry identified this very real challenge. A study by IBM in 2018 discovered that while 98% of organizations planned to adopt multicloud architectures by 2021, only 41% had a multicloud management strategy in place and just 38% had the required procedures and tools to operate a multicloud environment.

What makes data fragmentation such an issue is that these siloed platforms each have their own tools and processes for data management and protection. Even within each environment, there can be nuances. For example, the way you protect on-premises VMware workloads might be significantly different from the way you protect your on-premises databases—and both of these would be significantly different from protecting cloud workloads.

You need to know that all of your workloads have the right levels of protection. You need to be able to easily tell what has been implemented and what is actively functional across the entire ecosystem. And when disaster hits, you need to be able to quickly recover applications, databases, and other workloads regardless of whether they're in a public cloud, private cloud, or on-prem.

Legacy Tools Aren't Sufficient

Managing and protecting data in the cloud is very different from what IT teams are used to on-premises. APIs, agility, serverless concepts, security—they're all vastly different from legacy practices. While some enterprises might think they can simply "lift and shift" legacy tooling into the cloud, that approach can seriously affect the elasticity and agility that drew them to the cloud in the first place. You can't assume that a cloud platform is basically a new type of data center and expect to use it the same way.

Native Cloud Tools Don't Work Across the Entire Ecosystem

Each cloud provider offers proprietary tools to help you manage and protect your data on that platform. When you have multiple public or private clouds, however, you're still managing and protecting the data on each platform in a siloed way.

The way you protect data on AWS is drastically different from how you protect data on Microsoft Azure, for instance, in terms of operator experience, policy model, capabilities, and limitations. Some of the processes can be rudimentary, and involve manual tasks. With multiple clouds, your IT team has to learn the processes and procedures for each cloud, and ensure that the data on each one is adequately protected. The way you monitor and report differs from cloud to cloud, which also differs from how you do it on-premises.

Even if your IT team could become experts on each of the native toolsets, it would take more time and effort to manage and protect data—and recover it quickly when needed—across multiple clouds.

Creating an In-house Solution Is Extremely Expensive

Some enterprises aim to solve the hybrid/multicloud data fragmentation problem by hiring a global cloud architect, who can create an in-house solution that works across platforms. That's one way to do it, but it can cost a lot. A cloud architect with the right skills is one of the most expensive resources in the technology organization today—precisely because they're solving such complicated problems. For organizations that don't have that kind of budget, they must look elsewhere for solutions.

WHY SHOULD I BUY A BACKUP SOLUTION WHEN NATIVE TOOLING IS INCLUDED FOR FREE?

The question to ask yourself is, are you confident that you can meet your RPO and RTO when restoring data and applications in each of these disparate environments? And can you do it at scale?

Say you have workloads on AWS, Azure, Google Cloud Platform, and on-premises. Each one uses different tools to restore data. You have 10 accounts in a region, and data for each account that goes across all the platforms.

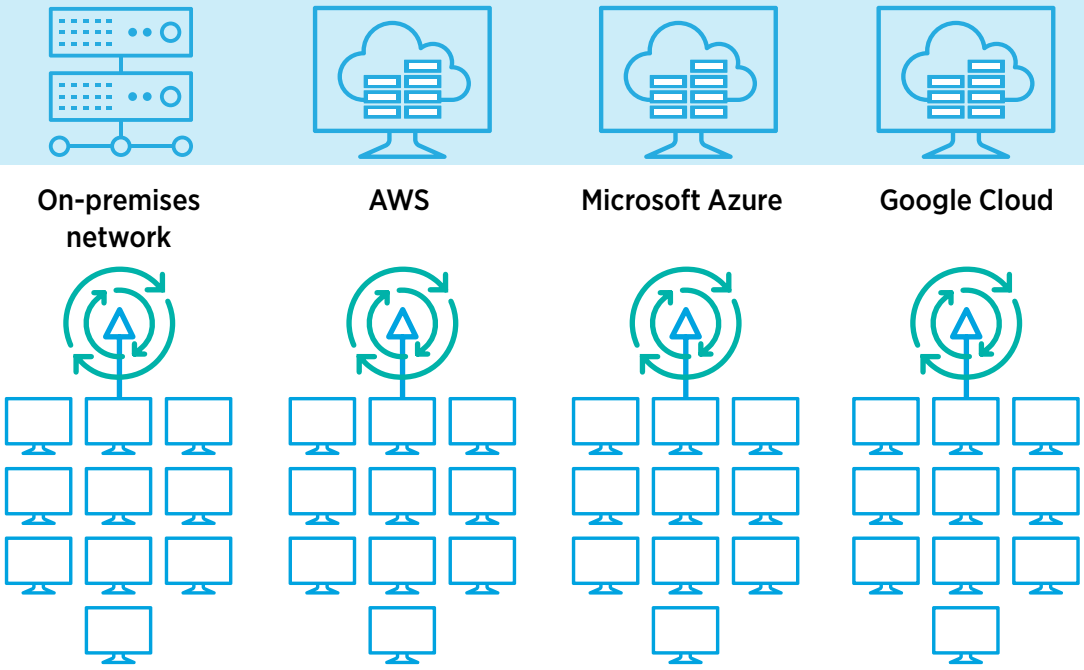

When disaster hits, are you really going achieve your RPO and RTO using a range of recovery protocols for 10 accounts across four platforms? (And that's assuming your team has the procedures down pat.)

Haven't well-established cloud providers such as AWS already figured out how to protect my data and applications?

All the major cloud providers offer tools for data management and protection on their platforms. They don't work cross-platform, as we've already mentioned, and they do have some limitations.

The most important thing to know is that public cloud providers operate under a **Shared Responsibility** model, which means while they are responsible for the protection and availability *of* the cloud, it is still your responsibility to protect your resources *in* the cloud.

Cloud providers each have their own slightly different definitions of shared responsibility, but they largely include the same main categories. The bottom line is this: Everything you previously had to protect on-premises, you still have to protect in the cloud. And that includes creating data backups. Just because your applications or data resides in distributed cloud storage, that doesn't mean it's backed up. You're still responsible for your own RPOs and RTOs, and for protecting your data against unauthorized access, faulty configurations, insecure APIs, and other risks.

There's still an assumption—or lack of knowledge—in many enterprises today about who is responsible for what.

A survey by Veritas showed that 85% of respondents said the cloud service provider is responsible for protecting their data. That's just not true.

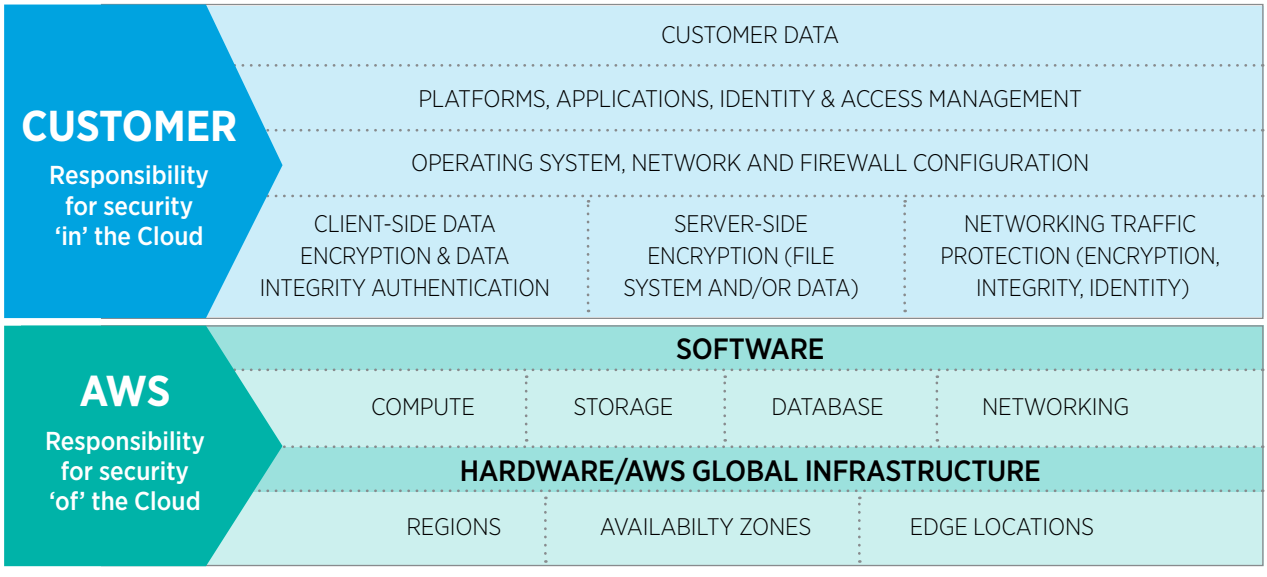

Source: AWS

A More Innovative Approach Is Needed

When it comes to data protection in the cloud, legacy tooling isn't enough and cloud-provider tooling only works on their specific platform. Enterprises need something more to break down silos and eliminate complexity. They need a solution that will allow them to easily build policies that match their business goals and then assign those policies quickly and seamlessly to workloads across the entire hybrid or multicloud ecosystem.

The ideal solution would integrate with public and private cloud platforms as well as on-premises infrastructure and provide a single control pane to manage, back up, and restore data across the whole environment. It would enable automated backups, make backup management easy (even at scale), and make it simple to recover any workload across the ecosystem.

Set Your Business Up for Success in the Cloud

A big part of simplifying backup and recovery in hybrid and multicloud environments is building a smart cloud adoption strategy from the start. Here are a few tips to keep in mind:

- **Start with visibility.** Discover and document your infrastructure and your applications. You won't be successful without it.
- **Create a plan to adopt cloud services.** Any plan is better than no plan. Check out the various cloud adoption frameworks available on the Internet.
- **Build a team of stakeholders focused on cloud adoption, strategy, and governance.** You can find various models online for creating a Cloud Center of Excellence (CcoE) specifically for this purpose.
- **Strive to be well architected, even outside of the cloud.** There are many published cloud architecture frameworks that can help drive the decision-making process. This is critical.
- **Find low-hanging fruit.** Look for options to ease capacity woes, service bursty workloads, or deploy new applications in the public cloud.
- **Don't be afraid to experiment.** Good governance is key, and once in place, enables you to iterate and evolve quickly.

Rubrik Delivers Cloud-Native Protection Across Your Entire Ecosystem

Achieve a consistent level of data protection and peace of mind by eliminating the complexity of hybrid and multicloud environments. With Rubrik, you get simple, intuitive solutions that allow you to implement data protection policies across public and private clouds and on-premises infrastructure with ease. Automatically discover, protect, organize, and manage all of your data and applications on multiple clouds from a single management pane. Recover data in place cleanly and simply.

Rubrik data protection solutions for hybrid and multicloud environments integrate with AWS, Azure, and Google Cloud Platform to protect native workloads. You can automate backup policy management to eliminate time-consuming scripting and manual job scheduling—and access near-zero RTOs, fast object-level recovery, and global management and reporting across your entire ecosystem.

TO LEARN MORE, VISIT

https://www.rubrik.com/en/solutions/cloud-native-protection

rubrik