ActualTech Media

CLUMIO

How to Protect Data in an All Cloud World

Protect Your Enterprise Data—Wherever It Needs to Be

James Green

IN THIS PAPER

Architecting IT systems has always been complicated and time-consuming. But the recent wave of innovation in cloud computing services, the rise of edge computing models, and the proliferation of software-as-a-service (SaaS) offerings have made architecture more complex than ever before.

As enterprises lean heavily on "cloud first" strategies, data is increasingly and necessarily scattered. In contrast to legacy architectures which centralize data storage, enterprise data is now spread across tens or hundreds of independent working sets. The data dispersion that is underway brings both significant challenges and exciting possibilities for cloud architects.

PROTECTING DATA IN THE NEW, DISTRIBUTED ENTERPRISE

This paper focuses on *protecting* enterprise data across an array of silos, and highlights some critical opportunities for simplifying data protection, containing cloud spending, and easing the burdens of security and compliance.

The key takeaways from this discussion:

- Modern data protection can and should be simple, even across clouds and SaaS
- Cloud data protection platforms require elasticity and must scale on demand
- Security and reliability are of paramount importance as data dispersion takes hold
- Costs must be predictable to realize a cloud-first vision
- Fast and efficient backup and recovery are crucial to modern data protection, as well as achieving optimal disaster recovery and compliance status

Cloud-first strategies need not pose a threat to your desire for simple and potent data protection. This paper will show how to unleash the full power of cloud to deliver secure backup and recovery for your data—wherever it needs to be.

Modern cloud architecture as a discipline is not for the faint of heart. Individuals in this role are responsible for bridging complex business problems with technical, cloud-based solutions. They need to identify and vet technical enablers for business goals such as:

- Reducing operational complexity, even as actual complexity grows
- Honoring cloud-first initiatives, while at the same time avoiding cloud lock-in
- Keeping a watchful eye on cloud spending and avoiding unexpected cloud expenses
- Pinpointing opportunities to drive business breakthroughs with cloud technology

And there's a long list of second-order architecture concerns, such as ensuring that enterprise data is safe from ransomware, and ensuring that compliance, governance, and data sovereignty considerations are always accounted for.

> While the premise behind a cloud-first strategy is widely accepted and is generally sound, it does create serious challenges when it comes to data.

Nowhere does the burgeoning menu of cloud options frustrate cloud architects more than when it comes to confidently managing and protecting data.

Cloud-First Initiatives Result in Data Dispersion

The pace and innovation in cloud technology leads many enterprises to implement "cloud-first" and "SaaS-first" strategies (See **Figure 1**). Essentially, leaders believe that cloud providers are better suited to deliver certain foundational IT building blocks than their own organization. They also like the consumption-based expense model of the cloud versus the capital investment in on-premises data centers and the army of staff needed to run them.

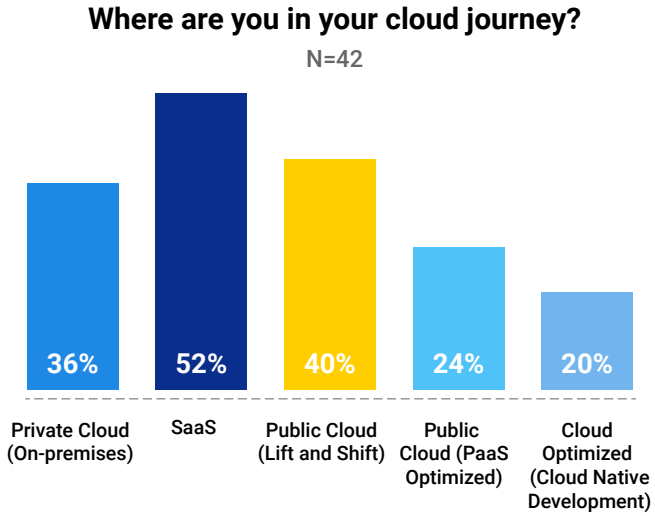

Figure 1: Although organizations vary widely on the stage of their cloud journey, it is clear that most of them are adopting cloud technologies in some form.

So they issue a decree that cloud is to be the default option for any new project—architects should only fall back on another option when a design constraint precludes a cloud-based deployment.

While the premise behind a cloud-first strategy is widely accepted and is generally sound, it does create serious challenges when it comes to data. With each new cloud service adopted, a new silo of data is erected. The result is data dispersion—the scattering of enterprise data across many independent locations. In effect, one problem is solved while a new one is created.

And if it were only one new silo, it wouldn't be such a big problem. But the reality is that enterprises are adopting new applications at a record pace. Okta's 2020 Business @ Work report indicates that the average Okta customer has integrated 88 separate enterprise applications.

The Arduous Task of Protecting Dispersed Data

Protecting data in one place is hard enough. Enterprises employ teams of people and complex, multi-pronged approaches to back up and recover data from purely on-premises infrastructure. It's easy to see how involving multiple cloud providers turns this already thorny problem into an overwhelming mess.

DATA DISPERSION CHALLENGES

Protecting enterprise data strewn about the world in tens or hundreds of data silos threatens the sanity of even the most level-headed cloud architects. Consider the following challenges:

- **Complexity.** With each additional data silo, guaranteeing that data is protected becomes exponentially more difficult. Each workload comes with its own set of native data protection mechanisms. Ensuring that you're protected beyond those native protections is up to you. Further, keeping documentation up to date and helping other team members understand the architecture becomes increasingly hard to do.
- **Hidden and unpredictable costs.** Cloud service billing is notoriously confusing and fraught with opportunities to ruin your budgetary numbers with even a slight misconfiguration. Base configurations may not include adequate data retention, and adding sufficient data retention can easily double costs if done inefficiently.
- **Security concerns.** The good old days of a security "perimeter" to speak of are gone. In the era of dispersed data, every new data silo increases your attack surface and provides one more opportunity for misconfiguration. And misconfiguration is no peripheral threat. IBM's latest Cost of a Data Breach Report shows that an alarming 24% of data breaches have human error to blame.
- **Compliance.** A major factor in regulatory compliance is knowing *what* data exists in the first place. Each additional data silo makes knowing what you have harder. And confirming for auditors that it is protected according to regulatory guidelines gets harder as well.

Those are just passive challenges that data dispersion introduces by its nature. What about active security threats such as the scourge of ransomware impacting enterprises today? Or the possibility of disgruntled employees destroying company data on their way out the door? It's safe to say that each new data silo makes protecting your enterprise data against these threats harder, too.

In a time where "cloud-first" thinking permeates the collective culture of enterprise IT, could it be that cloud is also the answer to the data protection problem?

Data Protection for the Distributed Enterprise

Sometimes, it can feel like a top-down initiative like a "cloud-first strategy" is an imposition which makes it harder to do things that should be easy. The good news is that in the case of data protection, the opposite turns out to be true. By embracing the cloud-first paradigm and leaning into it, protecting data across a plethora of enterprise data sources becomes refreshingly simple.

> In a time where "cloud-first" thinking permeates the collective culture of enterprise IT, could it be that cloud is also the answer to the data protection problem?

Secure backup as a service brings the same windfalls you know and love from line-of-business SaaS applications to data protection. True backup as a service is not simply a legacy data protection solution repackaged with a subscription model to masquerade as SaaS, though. To realize the benefits of data protection as a service, it needs to be authentic SaaS, complete with the experience you've come to expect from any cloud service:

- Onboard and protect data in minutes
- No infrastructure to maintain
- Proactive support
- Over-the-air updates with no downtime

Further, an effective service protects data across private cloud, public cloud, and SaaS. Let's examine some of the qualities of effective backup as a service and how such a platform can help overcome data dispersion challenges.

IT'S SIMPLE, EVEN ACROSS CLOUDS

Most enterprises need to protect data in cloud-based infrastructure like Amazon Elastic Block Store (EBS). Further, a majority of enterprises need to protect VMware vSphere data on-premises or in VMware Cloud on AWS, or in a hybrid scenario which spans both environments. And in many cases, there is yet more data to protect in SaaS platforms like Microsoft 365.

One of the biggest nuisances in protecting a myriad of data sources is managing each one independently. Effective backup as a service must solve this for you by providing a single user interface for managing multiple business-critical enterprise data sources.

Because of this, data protection should be policy-based, and a single policy should apply to multiple data sources. That way, applying data protection rules across multiple data silos is as easy as creating a single policy. Beyond simplicity, a unified policy ensures consistency which reduces the possibility that something important gets overlooked. This ultimately impacts your security and compliance posture.

> A modern backup as a service approach can leverage a cloud-native serverless architecture and scale up instantly on demand.

And while legacy backup solutions might leave you grasping at straws when things go haywire, a service-based approach should feature proactive support. That means that a support ticket may well be opened and resolved before you're even aware a problem exists. It's what you'd expect from your SaaS CRM; it's what you should expect from backup as a service, too.

ELASTIC AND BUILT TO SCALE

A primary driver for cloud adoption is the desire to manage less infrastructure. Enterprise backup is one of the last major infrastructure holdouts, accounting for a significant data center footprint in most enterprises, and requiring the constant attention of IT administrators.

> While legacy backup solutions might leave you grasping at straws when things go haywire, a service-based approach should feature proactive support.

Even many contemporary data protection solutions still employ an infrastructure-centric design where organizations deploy and manage physical appliances. And although appliance-based data protection solutions employ some principles borrowed from the cloud such as scale-out architecture, you're still stuck managing infrastructure at the end of the day. And with it come all the "usual suspect" burdens like capacity planning, facilities management, and hardware troubleshooting.

In contrast, backup as a service must look and feel like a service. There can be no infrastructure to manage if it is "as a service."

Another shortcoming of legacy data protection architecture is that it scales like infrastructure. In a best-case scenario, scaling out requires procuring more boxes, racking and cabling them, and doing an initial configuration—easily a weeks-long process. In an even less desirable scenario, you could be looking at a forklift upgrade.

A modern backup as a service approach, on the other hand, can leverage a cloud-native serverless architecture and scale up instantly on demand. In the cloud, the number of VMs spun up to fill a specific need isn't limited by physical infrastructure.

If, for instance, you need to immediately quadruple your backup demand, there's **literally** nothing for you to do but adjust the backup job—the serverless back end automatically handles the scaling for you. It could not be more simple, easy, or fast. That is the power of public cloud computing.

THE SERVICE IS SECURE AND RELIABLE

Ideally, service-based data protection is deceptively simple to use. But under the hood, it should employ a security-first mindset which follows robust, cloud-native security best practices and then some.

In the spirit of a true service-based technology, the service provider should take on the burden of ensuring the security and compliance of your backup data, rather than leaving you holding the bag as some cloud data protection solutions do.

At a minimum:

- Your backup data needs to be encrypted both in-flight and at-rest.
- The service should automatically handle key rotation to keep things secure.
- Backup data should be kept in an immutable, air-gapped data store which ensures that it's safe even if the primary data source is compromised (like sending tapes to a vault).
- It should regularly perform background integrity checking of the data to be sure it's safe from degradation over time.

To validate the security posture of the platform, the service provider should proactively obtain security certifications such as:

- ISO/IEC 27001:2013
- AICPA SOC 2 Type I

All of this is critical to ensuring that your backup data stored within the service is safe and sound, and will be there when you need it most.

> In the spirit of a true service-based technology, the service provider should take on the burden of ensuring the security and compliance of your backup data.

DATA PROTECTION COSTS ARE PREDICTABLE

Erratic and unpredictable cloud spending is a common reason for repatriation—that is, the movement of workloads out of the cloud and back on-premises. Why? Unknowingly projecting "on-premises thinking" onto cloud implementations. The underlying principles that give cloud its edge operate on a fundamentally different set of assumptions.

Take, for example, protecting data stored in EBS by leveraging EBS-native volume snapshots. Through an on-premises lens, this seems like a logical approach. But achieving long-term retention of data with EBS snapshots overlooks cloud-native ways of protecting data that can be done at a fraction of the cost and with greater durability.

> Erratic and unpredictable cloud spending is a common reason for repatriation—that is, the movement of workloads out of the cloud and back on-premises.

You should expect a backup as a service provider to take all of the unpredictability out of data protection and leverage cloud-native thinking to protect enterprise data with no infrastructure.

Like any SaaS, costs should be straightforward and predictable. You don't want to be ambushed by egress fees or artificially limited on how much data you can restore.

BACKUP AND RESTORE: FAST AND EFFICIENT

When the unthinkable happens and data is lost, enterprise data protection systems are the backstop that keeps the whole business from going belly up. The speed and efficiency with which data protection systems operate is intrinsically linked to business continuity. Keeping recovery time low is especially crucial for many enterprises.

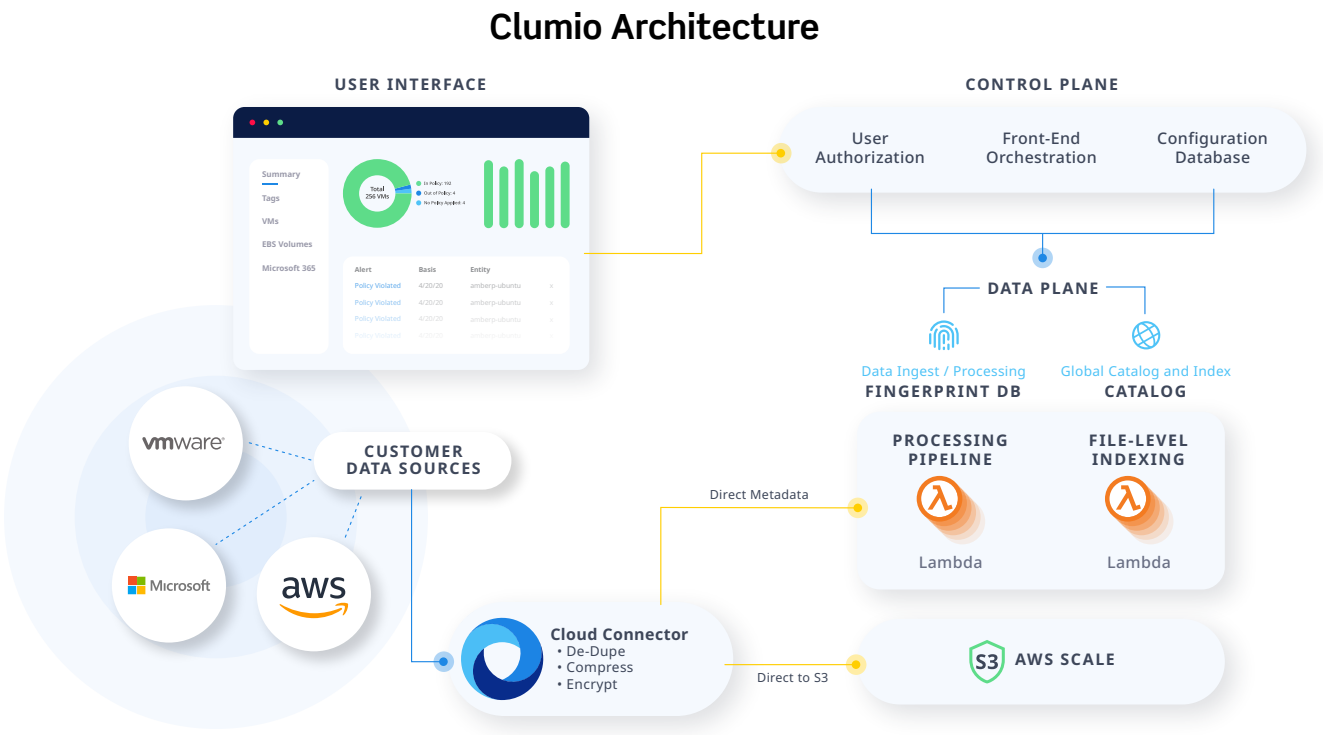

Figure 2: A serverless architecture leverages the full power of the cloud, and decoupled control and data planes maximize efficiency

Modern, cloud-based backup should continue to leverage tried-and-true tactics like source-side deduplication and forever incrementals to keep backup job times low. But beyond that, it should leverage cloud-inspired principles like decoupled control and data planes to allow the control plane to send instructions and then get out of the way and let the data plane do what it's good at. In this case, that means streaming data very fast to highly durable cloud storage.

Backup as a service solutions shouldn't force you to make sacrifices when it comes to recovery. Recovery times from the cloud to on-premises will necessarily have limitations caused by physics. However, the service should innovate to deliver recovery times that rival on-premises restore times wherever possible, and should be able to easily meet your recovery time objectives.

Clumio: SaaS Data Protection for an All Cloud World

An instantiation of the preceding ideals, Clumio takes these notions about what enterprise backup should be and delivers on them.

- **It's Simple, Even Across Clouds.** Clumio provides a single UI to protect enterprise data across clouds, and data protection is policy-driven. And the Clumio team takes care of the hard stuff. In fact, 68% of Clumio support tickets are opened by the Clumio support team on behalf of customers.
- **Scale on Demand.** The Clumio service is designed to scale by leveraging a serverless architecture to dynamically meet resourcing needs instantly and accurately (see **Figure 2**).
- **Always-On Security.** Clumio safeguards data from end to end with always-on encryption, automatic key management, and ongoing integrity checking. Clumio has achieved ISO 27001 and SOC 2 Type I certifications to prove that the platform is secure.

- **Predictable costs.** Clumio customers pay a straightforward per-VM, per-protected-TB, or per-seat fee depending on the workload. No hidden fees and no mathematicians needed.
- **Fast and efficient backup and restore.** Clumio backup as a service builds upon proven data protection techniques by adding game-changing features like Rapid Recovery. This proprietary technology leverages a mechanism called Reverse Changed Block Tracking in combination with the serverless architecture shown in Figure 2 to radically accelerate restores. Clumio can locate, rehydrate, and restore only the changed blocks required for the restore at lightning speed.

Taming Data Dispersion

The widespread adoption of SaaS and other clouds has driven the mass dispersion of enterprise data, and threatens to increase complexity to a tipping point. To make matters worse, data privacy regulations and governance mandates are growing stricter every day. As enterprise data gets spread around, ensuring compliance and securing data is harder than ever.

> Clumio is an authentic SaaS solution built in the cloud to address enterprise data protection needs brought about by the cloud.

In the face of overwhelming complexity, however, the same cloud-centric paradigm poised to cause a data protection problem is the one that can solve it. Clumio is an authentic SaaS solution built in the cloud to address enterprise data protection needs brought about by the cloud. To see just how easy it is to begin protecting data with Clumio, sign up for a demo at https://clumio.com/schedule-a-demo/.