Commvault®   CLUMIO
A Commvault Company

Redefining Data Protection for Amazon S3

Commvault Clumio simplifies S3 backup and recovery at any scale

REDEFINING DATA PROTECTION FOR AMAZON S3

Understand why Amazon S3 data needs protection

Amazon S3 is used to build websites, mobile apps, big data analytics, and dev/test environments by millions of customers. Its ease of use, reliability, and low cost make it very popular with both developers and cloud operations teams. Although S3 is highly available, it is critical to recognize that infrastructure or service durability is not the same as data durability. If a customer uses S3 to store data, the validity, security, accessibility, and usability of that data is the customers' responsibility. This is called the shared responsibility model, and is very important to understand for any AWS customer. If data on S3 gets deleted, compromised, corrupted, or ransomed, AWS is not responsible for restoring it.

Most S3 data loss happens due to user errors, and cybersecurity tools alone can't solve this problem

While AWS provides several security tools, and there is a plethora of cybersecurity vendors to reduce malicious atttacks, most S3 data loss is actually from user errors such as accidental deletions, misconfigured buckets, weak creden- tials, sensitive data exposure, and falling prey to phishing. In fact, according to Gartner, "through 2025, 99% percent of cloud security failures will be the customer's fault."

Common causes of user-driven S3 data loss:

- Erroneous data lifecycle policies
- Mistakenly deleting data
- Leaving sensitive S3 buckets open to the public internet
- Inadvertently publishing private keys on public repositories
- Software-based expunges
- Insider threats

> Even though we have a number of native data protection features built into S3, deletions are customer driven, or they can be machine driven, and we don't know the intent of a customer request to delete an object. Is that a malicious delete? Is that an accidental delete, or is that a correct delete? We have to honor them regardless. So that's why we're starting to see the real need to protect this data.
>
> **Principal Product Manager**
> Amazon S3

S3 is your innovation engine, and it needs to stay up

Organizations use S3 to power their most innovative applications. It is, therefore, important to secure what's in it, including customer data, user-facing media, sensitive records, and personally identifiable information.

Losing critical data from Amazon S3 can have material consequences, including operational disruptions and downtime, lost revenue, fines for non-compliance with industry regulations, loss of customer trust and reputational damage, and overall business fragility.

Build or Buy?

AWS provides its customers several powerful ways of creating copies of S3 data, including versioning, replication, and AWS Backup. However, it can take several full-time staff and millions of dollars annually to build and maintain a comprehensive data protection platform from these tools.

And even then, simple misconfigurations can create irreversible errors. For example, a replicated bucket can get deleted or corrupted with the original, a threat actor could delete all versions of a bucket, or a retention policy could make your data out of compliance. This is particularly pronounced for large S3 environments with millions of objects across hundreds of accounts. Moreover, tools cannot (and are not meant to) offer platform-level protection. For example, they cannot scale to protect beyond a few billion objects and are unable to offer an out-of-enterprise air gap. Importantly, recoveries

can get slowed down significantly as the user has to find the right copy to recover from many versions. Finally, protection policies can be set only at the level of buckets, which might increase costs where only object level protection is required. A solution that is built from composite pieces also leads to challenges in observability across different services and widens the surface area for data loss. For critical data with zero margin of error, this is a risky stance.

The consequences? Important data might get lost, recoveries could get delayed, and cloud costs can balloon. It is extremely difficult to build a robust, cost-effective data protection solution, especially one that can be relied upon to protect S3 environments at scale, and recover them quickly.

> The biggest benefit to our end customers has been [meeting] our SLAs by minimizing recovery time and improving our resiliency, ensuring we're able to meet our obligations even faster than we otherwise could if we built a solution ourselves.
>
> **Lead Cloud Security Engineer, Brightloom**
> Clumio customer

CLUMIO—SIMPLE PROTECTION AND RAPID RECOVERY FOR AMAZON S3

SaaS simplicity

Clumio is an agentless all-SaaS solution, without any of the management, upgrades, and installation hassles of traditional data protection. Updates are rolled out without the need for any intervention allowing customers to backup a slew of data sources within minutes, without worrying about asset size or capacity constraints.

Intuitive protection

Clumio's protection groups relieves customers from having to manage thousands of object versions across accounts, prefixes, and buckets. It provides a mechanism to intuitively classify data and apply policies to related or interdependent datasets, and recover them quickly when needed.

Flexible recovery

Unlike patched-together tooling, Clumio can help customers restore data to any point in time at the level of objects, buckets, prefixes, or protection groups, to any uncompromised account or region. Customers can simply pick the right recovery point-in-time on the Clumio calendar, or use the advanced filters to quickly find the exact object, bucket, prefix, or protection group they need to recover.

> With Clumio we can selectively provide the right levels of protection for our Amazon S3 data via global policies and protection groups. Global search enables quick recovery of any bucket or object for compliance needs, and optimizations to reduce small object overhead deliver a low TCO.
>
> **Senior Director, Engineering Enablement at Cox Automotive**
> Clumio customer

Commvault®

WHAT MAKES CLUMIO DIFFERENT?

INFINITE SCALE

Clumio is a serverless data processing engine and has no dependencies on legacy backup constructs such as minimum number of nodes, on-prem hardware, or object limits. Clumio has the highest scale for any backup technology today, and can safely protect exabytes of data consisting of tens of billions of objects per bucket.

COST-OPTIMIZED

Clumio does not force customers into purchasing large 4-node clusters or minimum commits, and instead meters for exactly the data that is protected, down to the byte. Clumio also provides insights into hidden data protection costs, intelligently estimates spend and proposes ways to reduce cloud bills.

INDUSTRY-LEADING PERFORMANCE

Clumio can get your entire S3 data estate protected in minutes. During a backup workflow, Clumio microservices inventory data sources, orchestrate, read, reduce, encrypt, and transfer data in parallel streams. Indexing and verification are also auto-scaled and enable an industry-low RPO of 15 minutes. During a restore operation, Clumio scales out rehydration while running parallel I/O operations across restore blocks to get customers back to a fully operational state in minutes. Clumio routinely clocks the fastest RTO for AWS workloads in the industry.

UNDELETABLE

Data on Clumio is air gapped outside enterprise access controls, so even if a threat actor gets hold of a customer's credentials and breaches their system, their data is safe. Clumio's resident data is outside of the customer's security sphere, and cannot be altered or deleted by anyone. There is literally no delete button. Each customer's data is encrypted and confined to an exclusively dedicated, immutable storage arena in AWS and access-controlled by MFA. Customers have the flexibility to bring their own keys, and define custom organizations and role based access controls.

> "
> Some of the things we were looking for in our data protection solution were really around immutability, ransomware protection, cloud-native, as well as API-first. These features and functions were paramount to our solution decision making and things we really only found in Clumio.
>
> **Managing Director, Defiance Digital**
> Clumio customer

CLUMIO—THE WORLD'S MOST ADVANCED DATA PROTECTION SOLUTION FOR S3
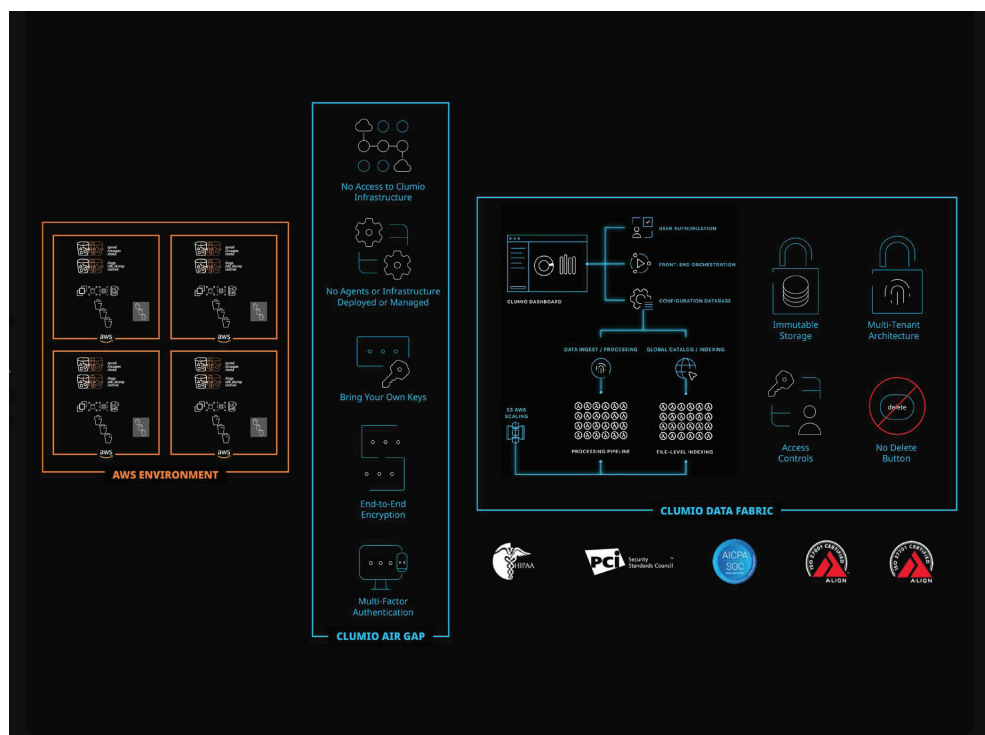
| Radically simple, even at exabyte scale | Air gapped peace of mind | Industry Leading RTO and RPO performance | Cost-optimized cloud data protection |

REQUEST A DEMO

commvault.com/clumio-request-demo

To learn more, visit **commvault.com**