

eBOOK

The Definitive Guide To Overcoming Challenges in AWS BACKUP

Contents

UNLOCKING THE FULL VALUE OF THE CLOUD

- Protecting Data in AWS
- Is AWS Backup the Answer?

THE CHALLENGES OF AWS BACKUP

- No Air Gap
- Long Recovery Times
- Zero Visibility
- Additional Complexity Protecting Data in AWS
- The Challenge of Third-Party Snapshot Managers
- Even More Hidden Costs
- Custom AWS Backup Scripts

COMMVAULT'S APPROACH TO AWS DATA PROTECTION

- Best-in-Class Security
- Rapid Recovery
- Better Visibility
- Simpler Management
- Lower TCO

Unlocking the Full Value of the Cloud

For the past several years, organizations have embarked on the digital transformation journey to modernize their services and improve customer experience. The year 2020 further challenged enterprises of all sizes, forcing them to re-think the way they operate. Digital transformation has moved from an option to stay competitive to an imperative for success.

Being shackled to a physical site to run day-to-day operations is no longer an option, as the new world requires the ability to efficiently run 100% from any location. Public clouds, such as AWS, have long been an option to achieve this goal, but in 2020 the velocity increased dramatically.

The driving forces to move to the cloud include access to the latest and greatest technology, being able to innovate faster, and not having to manage IT infrastructure. These benefits enable organizations to focus on their core business and grow at a much faster pace. The cloud is here to stay, and almost every organization will have a cloud presence.

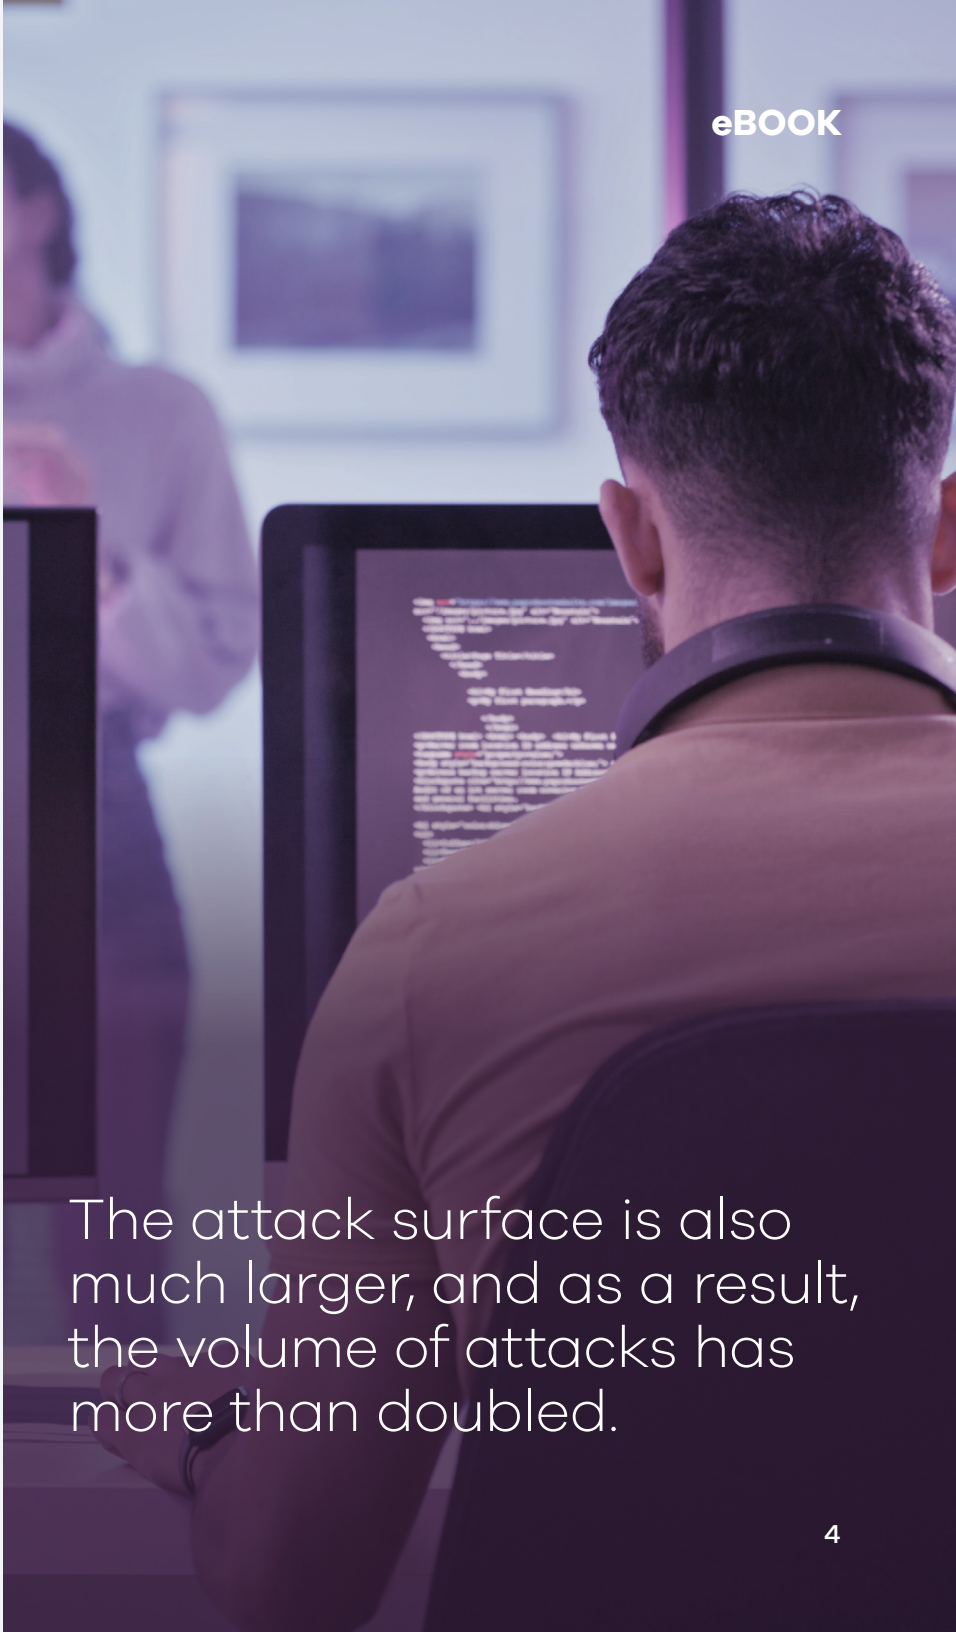
With the rapid pace of innovation in their services offering, AWS has been the leading cloud choice for many organizations. However, along with all of the benefits of AWS, there exist significant challenges that hinder unlocking the cloud's full potential.

Protecting Data in AWS

Protecting company and customer data in the cloud is one such key challenge. Having a sound backup and recovery strategy in AWS is as important as in the on-premises world, if not more. In many ways, it's also more challenging.

With the increased velocity of innovation enabled by the cloud, organizations are deploying and scaling their applications at a much faster pace, resulting in massive volumes of production data being generated that require protection. The data in the cloud is also more scattered—across applications, accounts, and regions as well as different public clouds.

The attack surface is also much larger, and as a result, the volume of attacks is also increasing. Finally, from a data protection point of view, organizations are finding themselves in an environment that is very decentralized, compared to the on-premises world.

A person is seen from behind, sitting at a desk in a server room. They are looking at a computer monitor that displays a complex interface with many lines of text, possibly a code editor or a system log. The room is dimly lit, with blue light emanating from the screens. In the background, another person is partially visible, also working at a computer. The overall atmosphere is professional and technical.

The attack surface is also much larger, and as a result, the volume of attacks has more than doubled.

Is AWS Backup the Answer?

Many organizations begin in the cloud with “shadow IT” projects or through a digital transformation initiative that includes cloud migration as part of that strategy. AWS makes it easy for engineers to spin up some EC2 instances and start putting some database workloads in RDS.

Once they are well underway, they realize that there is no sound data protection strategy for all these new cloud workloads—no solid backup and recovery plan. After all, backup is not typically in an engineer’s vernacular, and they assume it’s all taken care of. When it comes to data protection it can be difficult to know where to start in the cloud.

A convenient path at this point is to leverage AWS native backup services. One of the most common ways organizations start implementing data protection for AWS is using its snapshot management service.

At first, snapshots seem reasonably priced at \$0.05/GB per month. Businesses assume they have a painless way to take snapshots using the AWS Management Console. In fact, to make sure engineers don’t get burned, some organizations jump the gun and take some snapshots on all EBS and RDS volumes with a few clicks of the mouse.

Data protection appears under control. But things are not always what they seem to be.

Data protection appears under control, but things are not always what they seem to be.

The Challenges of AWS Backup

Unfortunately, it doesn't take long for organizations to start seeing the limitations of using snapshots to protect their critical company and customer data on AWS.

While snapshots provide basic data protection capabilities—such as recovering from operational errors—they fall short on delivering a comprehensive data protection solution that is simple and cost effective at the same time. It turns out that snapshots are rudimentary at best and require complex time-consuming processes to perform basic day-to-day tasks—defeating the whole purpose of embracing the public cloud.

Ultimately, cloud providers, including AWS, guarantee the safety of the infrastructure but not the data. Customers are responsible for managing their own data. Therefore, to take full advantage of AWS, it is critical for enterprises to get cloud protection done right.

Let's take a close look at what makes using snapshots for AWS backup and recovery so problematic.


No Air Gap

Ransomware attacks are skyrocketing, and the fear of criminals holding enterprise data hostage is turning out to be the worst nightmare for many CIOs and IT decision makers. According to Cyber security Ventures, ransomware is predicted to cost victims around \$275 billion annually by 2031, with a new attack every 2 seconds. Therefore, it is of paramount importance to have the right cyber resilience solution to help thwart such attacks.

The right solution recommended by experts such as CISA (Cybersecurity and Infrastructure Security Agency) is to have air-gapped backup data that is independently secured and saved, isolated from an organization's security sphere. This prevents hackers from finding the backup copy even after they have gained entry into your cloud account.

As shown, when organizations use AWS Backup to protect data sources such as EC2, EBS, RDS, etc., the snapshots are created in the same account as the primary data sources. The problem with this approach is that there is no separation or air gap between the primary data and the snapshots. While AWS has introduced logically air-gapped vaults for some resources, these are still within your enterprise security sphere and may be vulnerable to compromise.

If a hacker or bad actor gets access to the primary account, they will first compromise the snapshots before compromising the primary data. Now there is no valid usable backup copy that exists and hence no way to recover the primary data. This is precisely the situation organizations do not want to find themselves in and is a big limitation in data protection mechanisms in the cloud.



Long Recovery Times

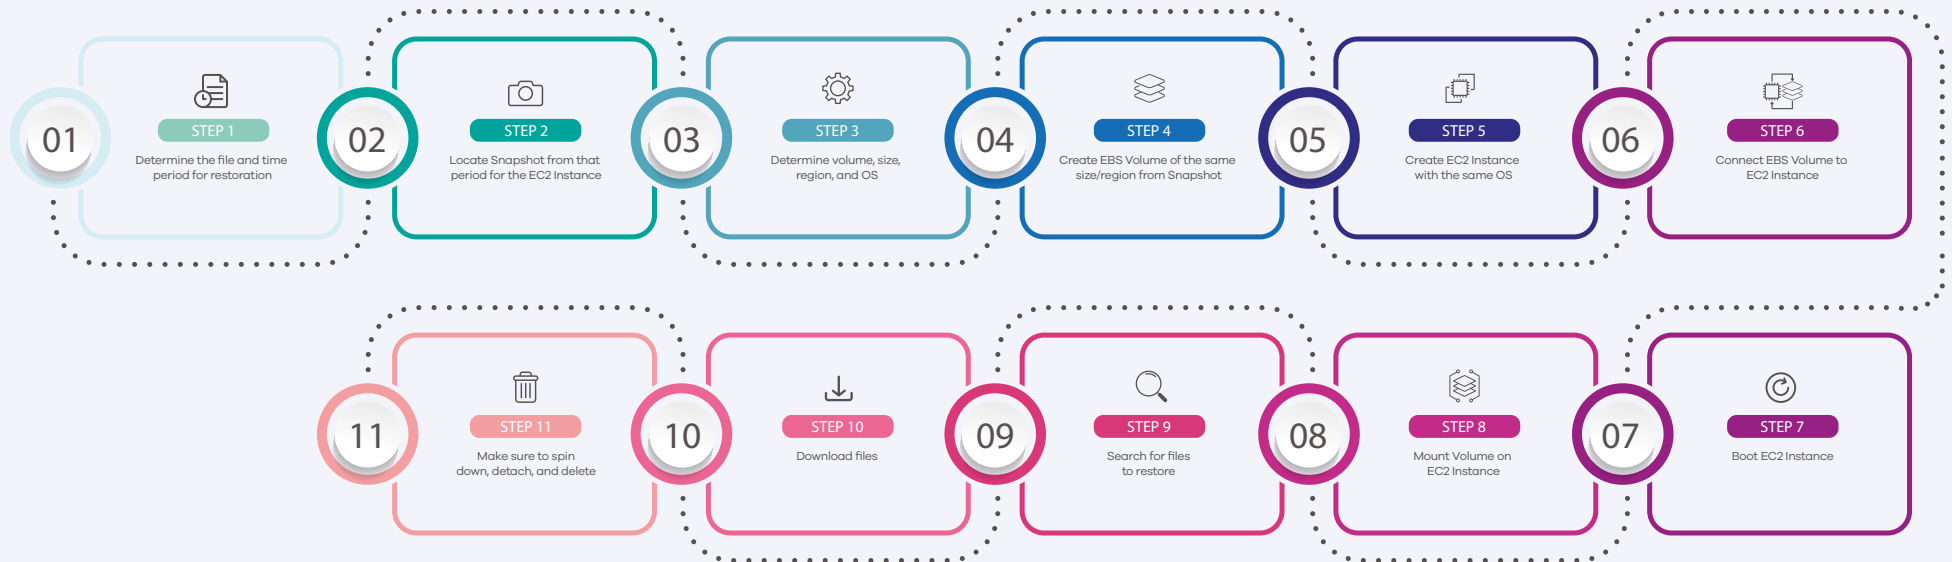
Data protection comprises two fundamental functions: backup and recovery. While it is important to make sure all critical data is backed up securely, it is equally important to have a solution that enables fast recovery in the event of a data failure or compromise. Not being able to recover business data in time during such events results in disruption to business continuity, poor customer experience, and, in extreme cases, can even put organizations out of existence.

Recovery is not all about just getting access to data, it is about the time it takes to get access to the data in a granular fashion. Recovering data from snapshots created using AWS snapshots can take several hours if not days.

The recovery process is like having numerous boxes of files in a warehouse. With snapshots there is almost no visibility as to what is inside each box, forcing an organization to open them up and rummage through every single file. The organization has to go through the boxes trying to figure out which file is the one they want. When they find the file that they think they want, they begin the restoration process and then have to wait for the whole thing to finish. This alone can take a significant amount of time.

eBOOK

Recovery is not all about just getting access to data, it is about the time it takes to get access to the data in a granular fashion.



Then they have to mount it to a server. After that is done, they have to load it and see if it truly is the data they want. Often they look at it and realize that the needed file is not there, or it isn't the correct version, and then they have to go back to the boxes and do the whole thing over again. This whole process can easily take several hours.

Let's look at a typical scenario to illustrate this point. An organization is using AWS Backup to protect its EBS volumes and now needs to recover a specific file from one of the volumes that is compromised. They will need to perform the following complicated steps in AWS to recover the file.

Depending upon how long it takes to find the right snapshot in Step 2, identify the right OS, volume size, region, etc. in Steps 3 to 5, and finally create an EC2 instance that matches the original instance, this whole process can easily take several hours.

And these complex steps need to be performed for every file that needs to be recovered. RDS is even worse as it requires the restore of an entire RDS instance to get access to the data, even if they only need a record or two. Finally, they need to make sure that the resources that were spun up in the cloud to recover the file are spun down in order to avoid unwanted expenses.

Zero Visibility

Setting data protection policies across all applications is not a daily task. In fact, they should be set once with the right attributes and not tampered with frequently to maintain compliance requirements. With this mode of operation, it is not easy for IT admins to remember every policy, backup history, and compliance requirement for each data source that is protected in their environment.

However, it is important to have this information at their fingertips when needed. For example, they should be able to:

- Quickly prove compliance during audits
- Easily find the right snapshot during the restore process
- Select the right policy to protect a new application or data source that gets added
- Do all the above across 100s of accounts

Launch

EC2 Image Builder

Actions

Owned by me

Filter by tags and attributes or search by keyword

1 to 50 of 100

Name	AMI Name	AMI ID	Source	Owner	Visibility	Status
woontest	AwsBackup_j-0f67d0c6b3a0be...	ami-0253e1ca0d57f95a6	786578629570/AwsBackup_j-0f67d0c6b3...	786578629570	Private	Available
woontest	AwsBackup_j-0f67d0c6b3a0be...	ami-0e1a9523d0733dcd9	786578629570/AwsBackup_j-0f67d0c6b3...	786578629570	Private	Available
woontest	AwsBackup_j-0f67d0c6b3a0be...	ami-02d0477ac5ce9f9ee0	786578629570/AwsBackup_j-0f67d0c6b3...	786578629570	Private	Available
woontest	AwsBackup_j-0f67d0c6b3a0be...	ami-0ea22f73e259b8e2	786578629570/AwsBackup_j-0f67d0c6b3...	786578629570	Private	Available
	clumio-app-logs	ami-bc441504	786578629570/clumio-app-logs	786578629570	Private	Available
	Clumio-App-Systemd1	ami-40acd538	786578629570/Clumio-App-Systemd1	786578629570	Private	Available
clumio-esx	clumio-esx	ami-8faee8fe	786578629570/clumio-esx	786578629570	Private	Available
clumio-esx-2	clumio-esx-2	ami-56a6ec2e	786578629570/clumio-esx-2	786578629570	Private	Available
	clumio-gpvn-1547577160	ami-083d4c81ecd017ee	786578629570/clumio-gpvn-1547577160	786578629570	Private	Available
	clumio-staging-test	ami-2309985b	786578629570/clumio-staging-test	786578629570	Private	Available
	clumio-test-vm	ami-5c084324	786578629570/clumio-test-vm	786578629570	Private	Available
clumio-staging-test-v2	clumio-test-vm-v2	ami-9f7c12e7	786578629570/clumio-test-vm-v2	786578629570	Private	Available
clumio vapp builder v1.3	clumio-vapp-builder-3	ami-8da08e5	786578629570/clumio-vapp-builder-3	786578629570	Private	Available
clumio vapp builder v1.4	clumio-vapp-builder-v1.4	ami-4cafd534	786578629570/clumio-vapp-builder-v1.4	786578629570	Private	Available
clumio-vapp-builder-v1.5	clumio-vapp-builder-v1.5	ami-3e413646	786578629570/clumio-vapp-builder-v1.5	786578629570	Private	Available
	clumio-vapp-builder-v1.6	ami-6b4739f3	786578629570/clumio-vapp-builder-v1.6	786578629570	Private	Available
	clumio-vapp-builder-v1.7	ami-7e89d006	786578629570/clumio-vapp-builder-v1.7	786578629570	Private	Available
	clumio-vapp-builder-v1.8	ami-d08f6a8	786578629570/clumio-vapp-builder-v1.8	786578629570	Private	Available
clumio-vapp-builder-2.0	clumio-vapp-builder-v2.0	ami-bc55d404	786578629570/clumio-vapp-builder-v2.0	786578629570	Private	Available
	clumioapp	ami-9ba3f4e3	786578629570/clumioapp	786578629570	Private	Available
ClumioApp-1.1	ClumioApp-1.1	ami-0c85229d7550061b6	786578629570/ClumioApp-1.1	786578629570	Private	Available
ClumioApp-1.2	ClumioApp-1.2	ami-0096ec7ee3cd5c22d	786578629570/ClumioApp-1.2	786578629570	Private	Available
	ClumioApp-1.3	ami-0f7a4f4f	786578629570/ClumioApp-1.3	786578629570	Private	Available

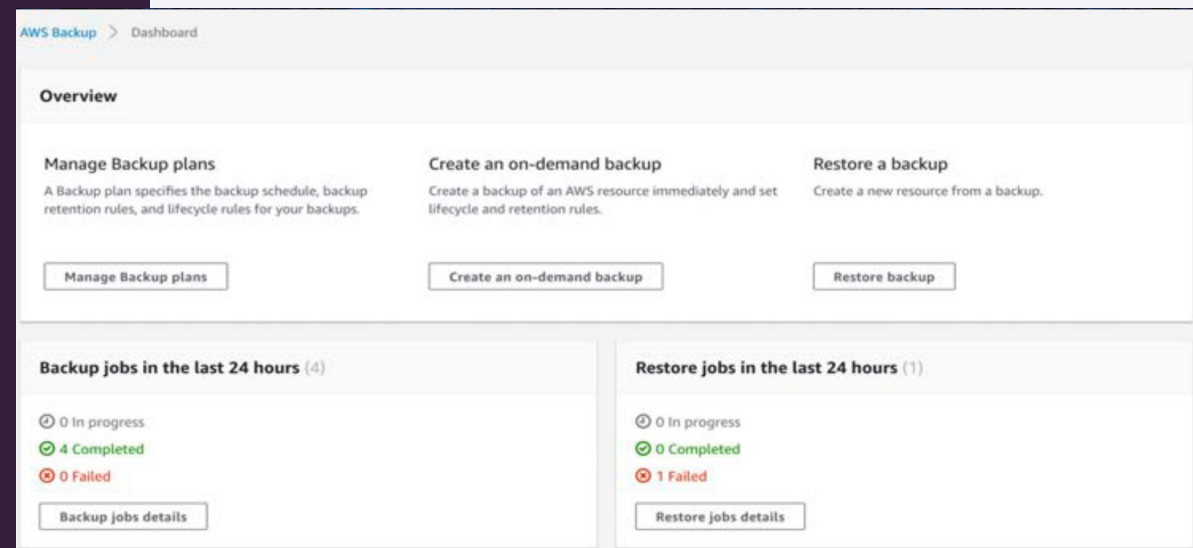
Lots of Snapshots to Scroll Through

Complicated and time consuming to find a snapshot to restore

AWS Backup does not provide the above functionalities, which results in limited visibility into your data protection plan.

Let's take a look at a real-world example that demonstrates the inherent dangers this could have on a business. IT departments often get audited to prove that they're meeting various compliance standards. An insurer, for example, might require a business to prove that they have 30 days of backups.

To do so with AWS Backup, an organization would have to write code to prove it or show it manually in response to an auditor request. And utilizing a manual process then introduces the potential for errors and risks in protecting the data and to being in compliance. The IT department likely will have to go through this painful process multiple times a year, depending on the compliance schedule, spending upwards of a week running reports for auditors on top of everything else their job requires.

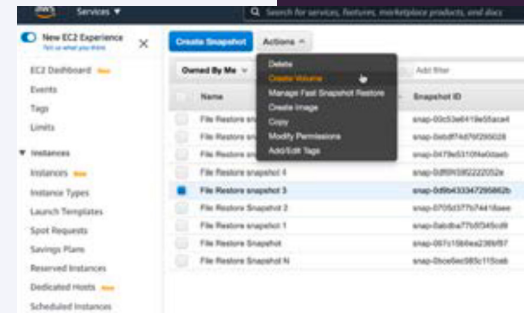


No visibility on compliance

Additional Complexity

From the gaps in AWS backup discussed so far, it should be clear that snapshots provide very basic capabilities. Organizations that start using snapshots to back up their AWS data soon run into these shortcomings and find themselves forced to respond.

They end up writing complex scripts to add missing functionalities into their data protection solution. They have to dedicate valuable IT resources to develop and maintain these scripts on an ongoing basis rather than having them focus on their core business. This is not in line with leveraging the cloud to add agility and enable faster innovation into your business.



Instance Restore Only

Limited instance restore in the same account only

Unavailable

Browse and Restore

Unavailable

Global Search

Unavailable

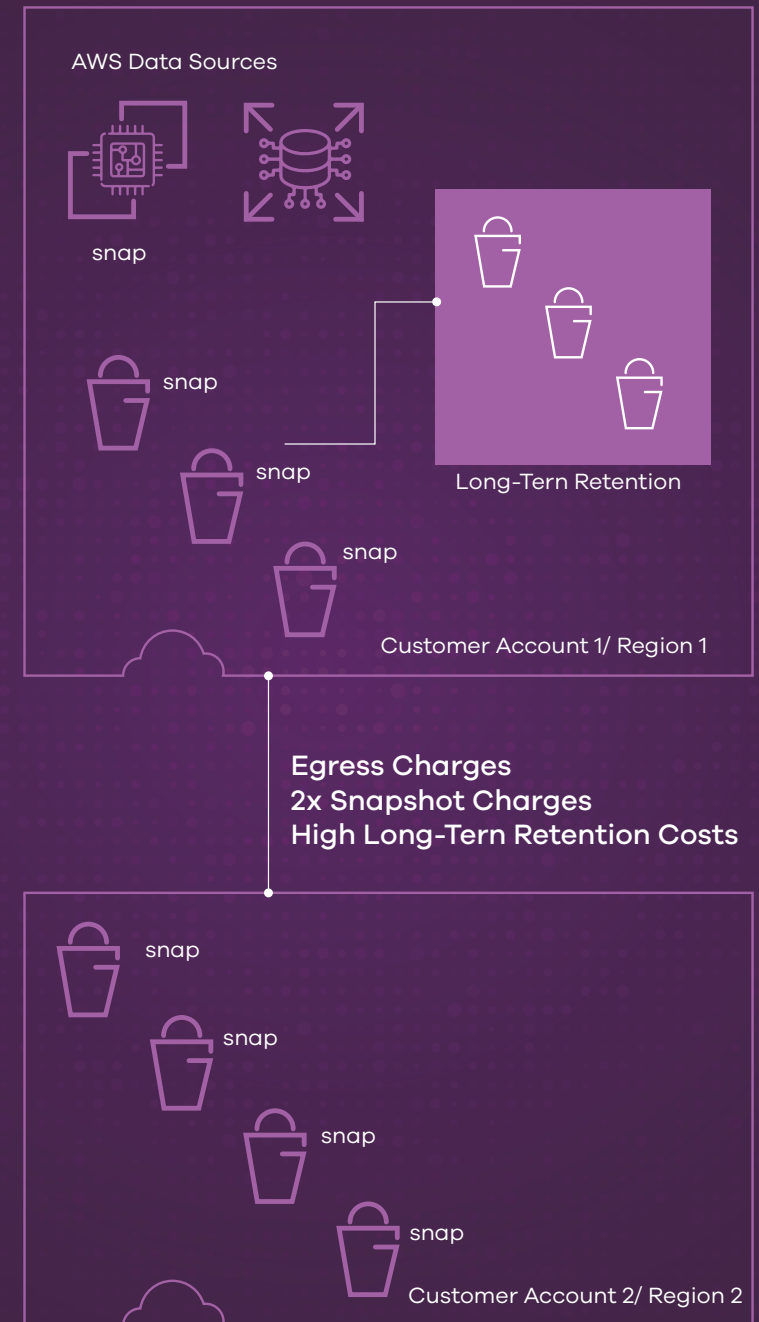
Granular Record Restore

Protecting Data in AWS

As part of a data protection strategy, it is typical for organizations to set up long-term retention for production data as well as protection against account compromises from attacks such as ransomware. To accomplish these two objectives via AWS Backup, organizations do the following:

- Create snapshot based long-term retention that results in creating tons of snapshots per account in high-tier storage
- Replicating snapshots from one account to another account to protect against account compromises that results in doubling the number of snapshots and thus the backup costs. There also are additional egress charges for the cross-account transfers


In a typical scenario, a few months go by and an organization is starting to realize a concerning trend. Their AWS bill has steadily increased and doesn't seem to be showing signs of leveling out. Nor do they have any idea what's driving their AWS backup costs. Meanwhile data protection with AWS can be costing them up to 50% more.



The Challenge of Third-Party Snapshot Managers

As organizations grapple with managing the cost of their AWS backup, they start to look at third-party snapshot managers. These snapshot managers talk about helping reduce organizational costs significantly by being able to tier to S3.

The cost seems reasonable at a glance, perhaps presented as a per-instance license fee. However, the total cost is actually higher, since it includes the additional cost that the organization has to incur to store the data in AWS and often, for the compute resources running their backup management. For example, if you consider a typical EBS volume with daily changes, you're still paying AWS its \$0.05/GB/month for the underlying snapshot storage, and that license fee is added on top. This creates an immediate premium. On top of that, these solutions often run their own in-account EC2 backup manager instances and temporary worker EC2 instances, which add further, blended usage costs to your AWS bill.



The cost per protected EC2 instance per month may seem reasonable, but the total cost is actually higher

Even More Hidden Costs

But do organizations actually save money on the tiering-to-S3 feature? In order to offer this capability, third-party snapshot managers spin up temporary EC2 instances that run for longer periods of time and add a hidden cost to your bill. The costs are not obvious because they are blended in with your own EC2 bill. The temporary EC2 instances have to traverse through a company's EBS snapshots and copy chunks over to S3.

The break-even on just moving it to S3 tends to be around the three-month mark due to all these hidden costs. This means if a business has retention periods less than three months for their daily backups, they could actually be spending more than if they had just left it in EBS snapshots. Add in the licensing costs we talked about earlier and the cost is even higher.

And because EBS snapshots are incremental, it's hard for the S3 tiering to effectively only move what's needed to S3. This is due to the fact that daily snapshots (that they didn't want to move to S3) may point to older blocks that are part of a yearly backup that were already moved to S3. Due to this complexity, it is likely the backup manager has a lot of your data in both EBS snapshots and S3.



The costs are not obvious because they are blended in with your own EC2 bill

Custom AWS Backup Scripts

Management sees just how much is being spent on AWS backups, especially the third-party manager, which convinces them to dedicate a shared engineering resource to help develop custom backup scripts that can handle the organization's backup requirements.

The organization eventually gets what it needs up and running, and they are license free and are finally feeling they have a handle on their backup policies. But as AWS changes their APIs and business needs adjust (like needing to backup cross-region or protect backups from ransomware), so do scripts. The shared engineering resource becomes a full-time role. And now the business has a not-so-hidden cost for managing, refining, and creating custom backup scripts.

Now the business has a not-so-hidden cost for managing, refining, and creating custom backup scripts.

Commvault's Approach to AWS Data Protection

From all the challenges discussed in the previous section—long recovery times, no true air gap, lack of visibility, mounting costs—it is easy to conclude that effective data protection in AWS is nearly impossible. Yet the advantages of AWS for a business as it continues its journey to the cloud are too significant to ignore.

It is increasingly critical that organizations find and implement the right data protection solution to truly leverage the benefits of AWS. To make sure it is done right, organizations need to implement a solution that addresses each of the key challenges of AWS backup today.

Organizations need to take a look at Clumio. We have taken a long look at the limitations of AWS Backup and created a solution that builds on top of native snapshots and addresses the challenges. Let's take a look at how we do it.


Best-in-Class Security

To make sure that backups are valid and usable to initiate a recovery process when primary data is compromised, it is necessary for backups to be saved outside of the security sphere of the primary data. This is called having air-gapped backup data. By air gapping the backups, hackers or bad actors cannot get to it, thereby leaving open the path to a successful recovery from any account compromises. Given the rise of ransomware attacks, combined with the large and decentralized attack surface of the cloud, organizations should pay extra attention to the security posture of their cloud data protection solution.

The solution should deliver:

- Air gap backup
- Immutable backups, so that the backup copies cannot be modified even if bad actors somehow get access to it
- No “Delete” option for backup data. This combined with immutable backups keeps the backup data well secured
- End-to-end encryption of user data, in transit and at rest

Long-Term Retention



Air-Gap + Long-Term Retention

Ransomware and bad actor protection

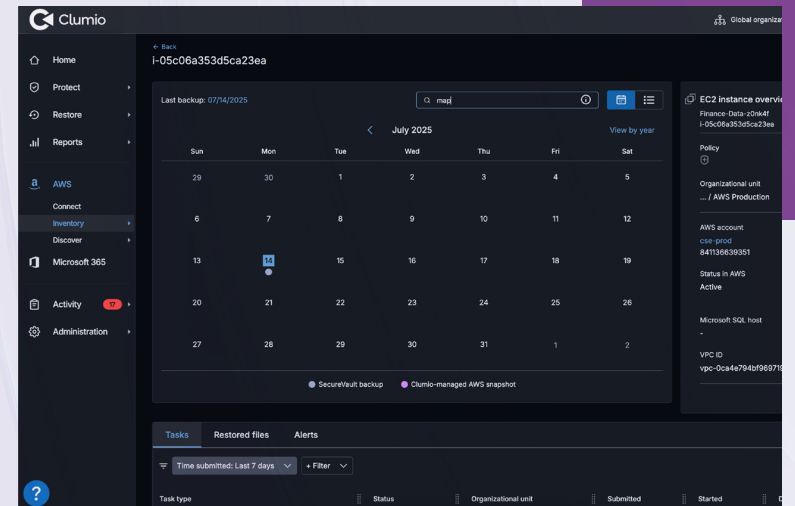
Immutable backups

No delete button

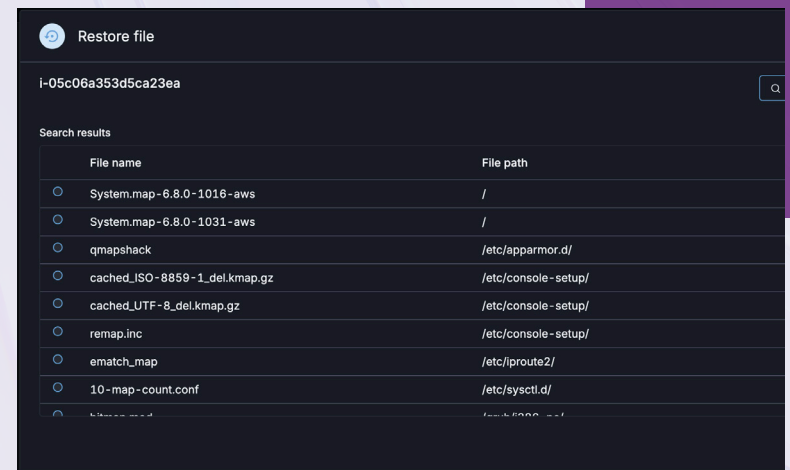
Lifecycle management delivering cost savings

Rapid Recovery

When it is time to recover from a data failure, one should be able to do so quickly to maintain business continuity. The right data protection solution must provide a quick way to find the data (snapshots, instances, files, records, etc.) that needs to be recovered and then restore it. Here is how Clumio enables rapid recovery for files in Amazon EC2.



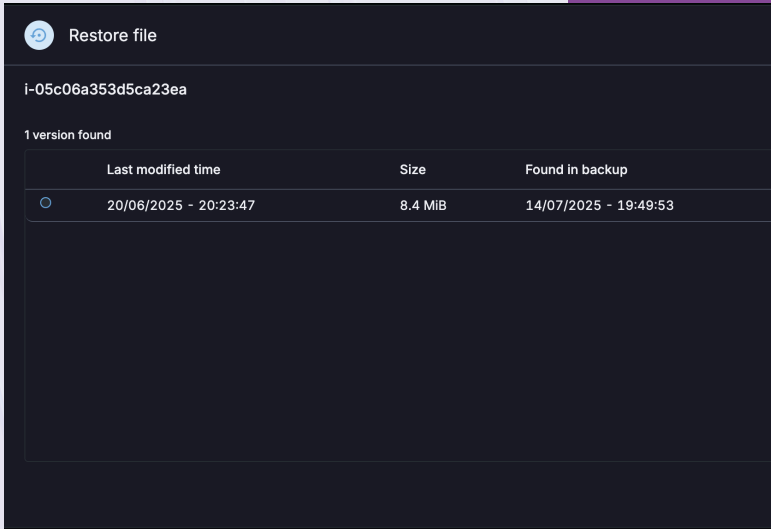
Step 1: Type in a search term for the file you want to recover



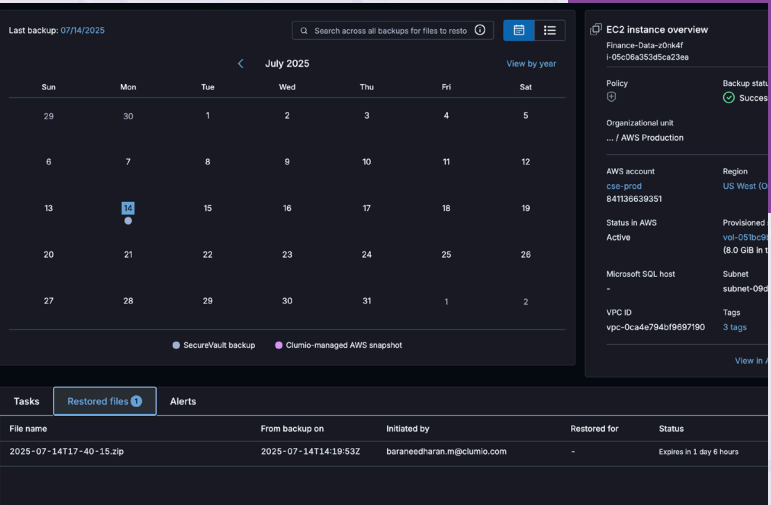
Step 2: Select the file you want to restore

In contrast to the low visibility AWS Backup process, Clumio's calendar view allows an organization to efficiently browse the entire file system. Instead of having to load or restore a file and then confirm it is the correct one, an organization can identify the needed file just by typing in data parameters and conducting a search. The user will see all the different versions stored, with time stamps, and can easily identify and restore whatever file or files are needed. Clumio significantly reduces recovery time essentially by letting an organization search its entire warehouse of file boxes without having to open them up first.

Such an easy-to-use rapid recovery process has helped Clumio customers reduce average recovery times from over 4 hours to as little as 10 minutes.



Step 3: Select a version you wish to download

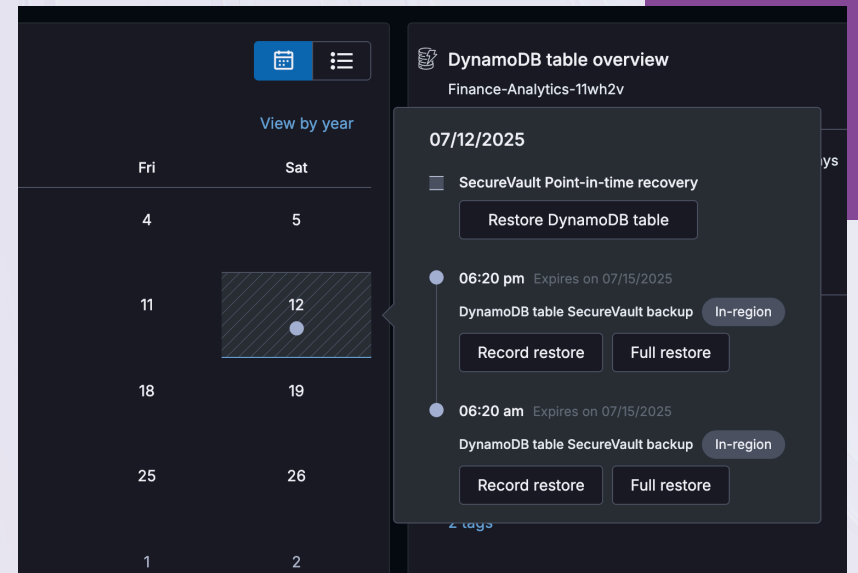


Done! Download the file

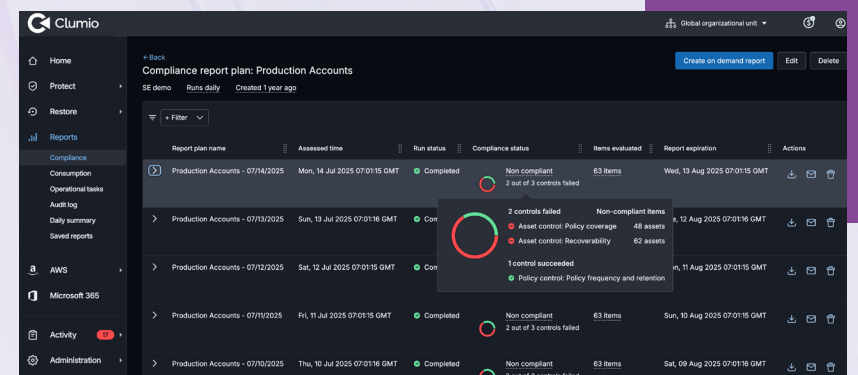
Better Visibility

A key requirement for a good data protection solution is the ability to provide the necessary information about protected assets in a simple to understand manner. If the user is required to develop custom reports and dashboards to get a view of assets protected by data sources, policies in place, global cross-account status or to understand if they are meeting compliance requirements, this is a recipe for introducing errors and an approach that will end up consuming valuable resources on an ongoing basis. Instead, these tasks should be built into the data protection solution.

The calendar view in Clumio presents a global understanding of all the snapshots and backups created for an AWS data source (EC2, EBS, RDS, etc.). You can then easily perform a point in time restore of the entire data source or view individual files/records to perform selective restores. Clumio's environment dashboard and the compliance report provide real-time compliance status for selected accounts or the entire environment. This allows the security team to stay on top of their data governance requirements and be audit-ready when the need arises.



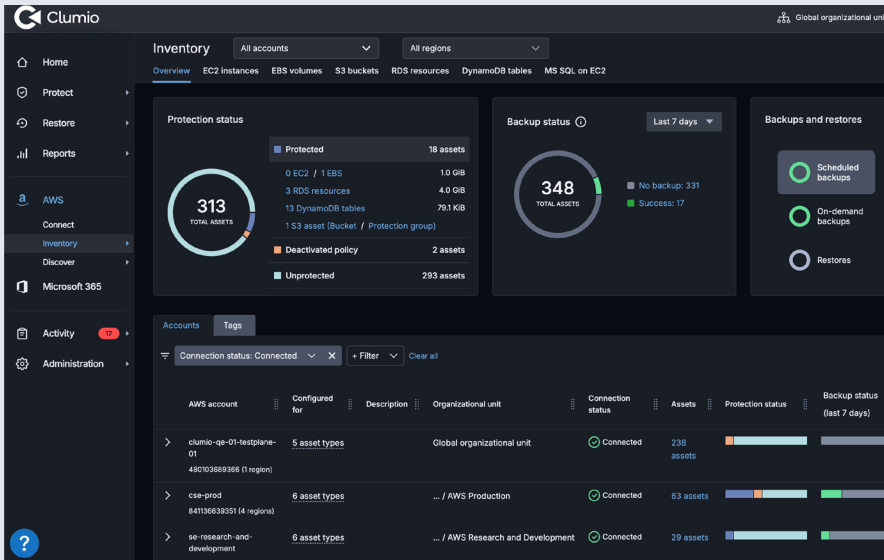
Backup History Calendar View:
Find any backup quickly for fast restore at any point in time




Global Compliance Reporting:
Single source of truth for audits and compliance

Simpler Management

Organizations should be looking for cloud data protection solutions that simplify and automate backup orchestration. It should enable simple onboarding, quick backups, global policy setting, and ease of recovery.




With a SaaS data protection service such as Clumio, it takes as little as 10 minutes to start protecting AWS data sources, and you get fine-grained control into your backup and recovery operations. This simplifies and automates day-to-day tasks for the IT and Operations teams, freeing up time to focus on their core business.




Instance Restore

Restore an entire Instance to any account



Browse and Restore

Browse the file system to restore the file



Granular Record

Scheme browser with record recovery w/SQL query



Global Search

Search for any file across any instance or volume



Global Policies

Simplify policy setting across multiple accounts

Lower TCO

One of Clumio's significant TCO advantages is in how it implements its air gap. In a typical scenario, if an organization had \$100,000 worth of snapshots sitting in their account and wanted to copy these files into an "air-gapped" account (basically saving a copy into an account with a different password), the total investment would now double to \$200,000.

With Clumio, the process is automatic and fully managed. There is no tiering that the customer needs to worry about; it is all automatically taken care of, and the customer simply pays on one easy to understand per GB model. In as little as 10 minutes, an organization can click a button and back up their data in an air-gapped environment. A business doesn't need to set up scripts, manage the process, or make mirror copies of their snapshots—thereby saving work hours and reducing backup costs.

The chart on the next page shows that by using a data protection service such as Clumio that has built-in air gap protection as well as lifecycle management for long term retention, organizations can avoid inefficient ways of backing up their data and 30% or more on their AWS backup costs on average. Considering all the other benefits that are delivered by this solution, this is icing on the cake.

On top of that, we help our customers understand what's driving their backup costs. And for customers that want to easily protect their backups in a secure air-gapped vault, customers can easily move their snapshots into Clumio at a lower cost than AWS native snapshots. It's no wonder that customers are quickly moving their AWS backups into Clumio.



THE DEFINITIVE GUIDE TO AWS BACKUP

By moving to Clumio, you get all the benefits of moving to a custom solution without the ongoing cost of developing and maintaining it: no hidden license fees, easier administration, lower costs than AWS native snapshots, and better reporting.



Ready to simplify backup for AWS?

Request a 1:1 demo:

www.clumio.com/demo

commvault.com | 888.746.3849 | get-info@commvault.com



© 2025 Commvault. See [here](#) for information about our trademarks and patents. 07_25