druva

White paper

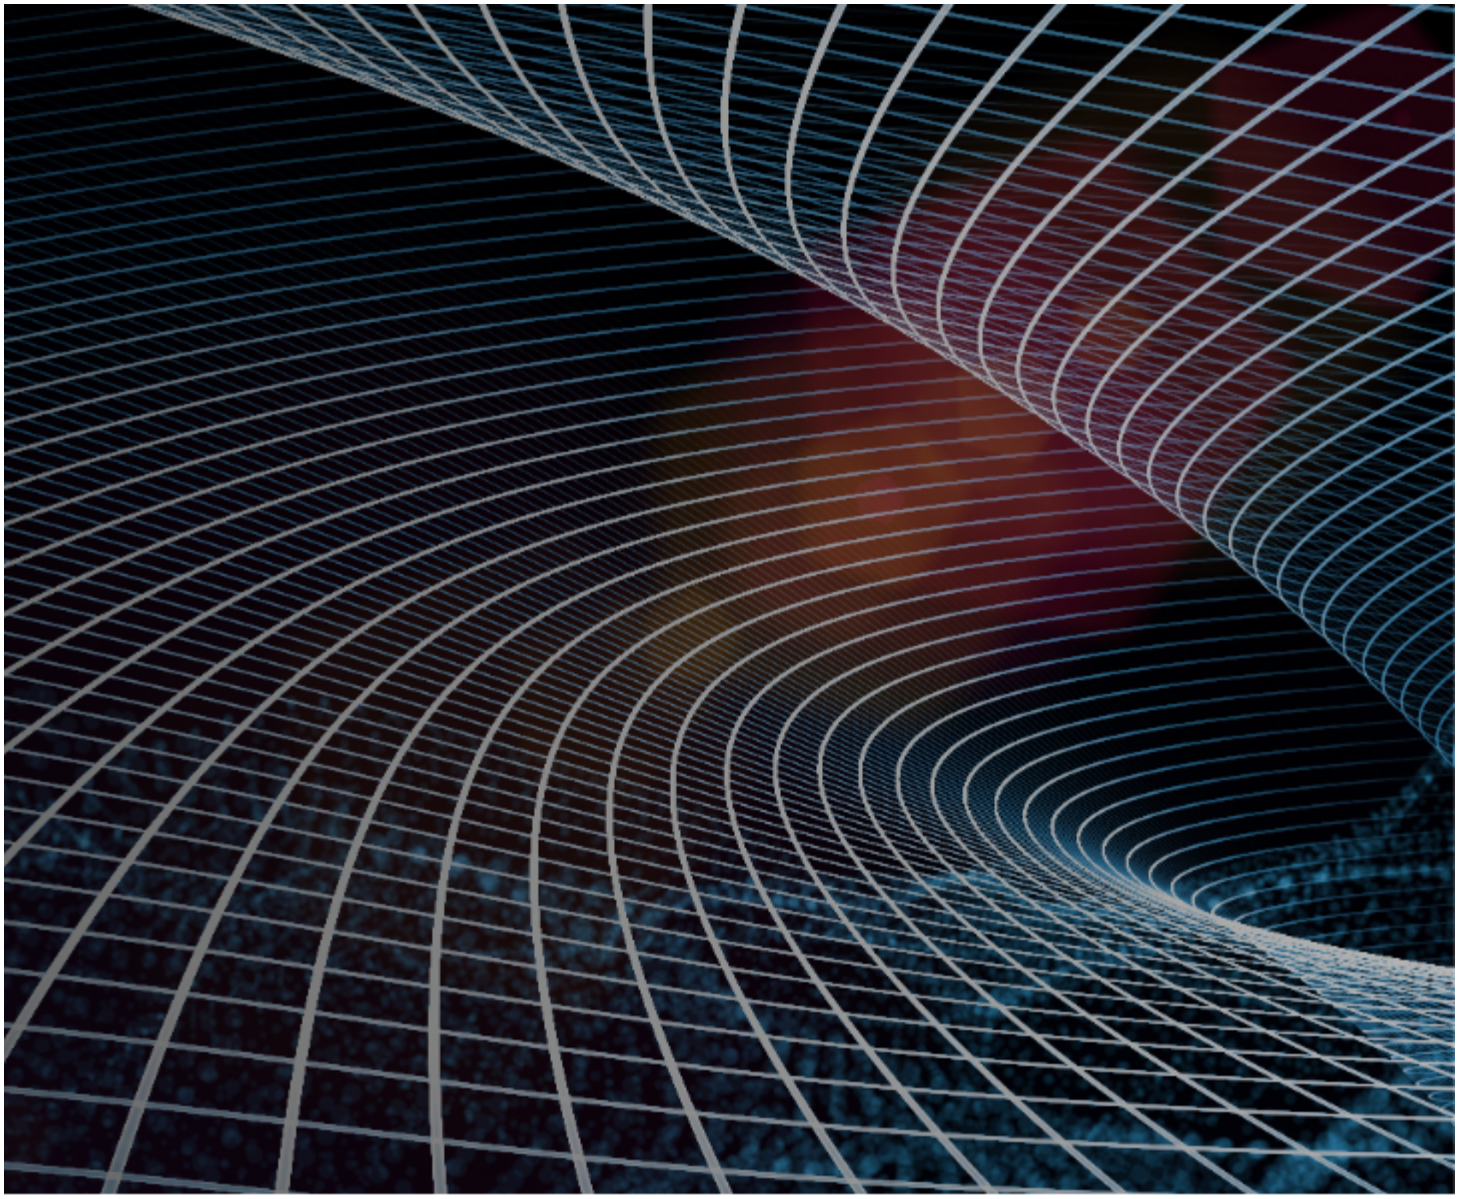

How to Secure Your AWS Data From Ransomware Without Breaking the Bank

Q422-20339

Executive Summary

Your business runs on AWS. To date, you believe you have been successfully protecting your applications with snapshots; but are snapshots alone the best approach? The most secure and cost-efficient?

For many companies, it's time to start thinking about more than just snapshots to protect AWS resources such as Amazon EC2. While such snapshots will always be the foundation of any operational and disaster recovery plan for AWS resources, they are far from a complete solution. This paper reviews what AWS snapshots are and the benefits they provide. It then explains the limitations of AWS snapshots and the challenges of using them as a complete data resilience and disaster recovery solution. Finally, it describes a new service from Druva designed to address those data security challenges for EC2 and EBS backup while delivering exceptional cost savings for your business — up to 50%.

How do snapshots differ from backups for AWS resources?

What are AWS snapshots?

An AWS snapshot is an *image copy* of the protected resource (e.g. EC2 instance, EBS volume) stored in object storage in your AWS account. A full or baseline snapshot is a byte-for-byte copy of the protected resource from a single point in time. Once the first snapshot is created, you can take incremental snapshots of any blocks that have changed since the last snapshot was taken — as many as you would like. This makes them efficient from a storage consumption perspective as each subsequent snapshot contains only the blocks that have changed since the previous.

AWS snapshots are image copies stored in a different storage system and, as a complete copy, can be used to restore a damaged or deleted resource. In fact, they will most likely be the first thing that you reach for if something damages the primary copy. For example, if you accidentally delete an EC2 instance or EBS volume, that does not damage the snapshot. You can quickly recover the deleted instance or volume from the snapshot.

This means that an AWS snapshot provides the basic building block to create a good backup: an independent copy of the data stored in a different storage system. They are by far the easiest and quickest way to recover a damaged resource in AWS.

If you copy a snapshot to a different region, it can also be used for disaster recovery. By using snapshots to recover your environment into the region they were copied, your recovery will be as fast as if you had kept them locally. This is perhaps one of the most beautiful things about using a cloud provider such as AWS; all the resources you need are at your fingertips — you simply need a snapshot to recover your environment.

What is a backup?

It's a simple question, but it's not as easy to answer as you'd think. The most basic definition of a backup is an isolated copy of data created for the purposes of restoring a resource after some type of damage. Backups are perfect for long-term data storage since they are not designed to override each other. Ideally, backups should be used to restore older backup files, mitigate the potential risk of data corruption or cyber threats, enable point-in-time recovery, etc. They can also be

compressed and/or deduplicated copies that are retained for potentially lengthy periods of time.They can be kept for months and even years, providing the ability to recover files that may subsequently have been deleted, corrupted or simply need to be re-accessed.

A backup system should be designed based on the agreed-upon needs of the organization. These needs will dictate how quickly a given resource needs to be restored, which will determine your recovery time objective (RTO). They will also dictate how much data you are allowed to lose in a restore, which in turn will determine your recovery point objective (RPO). Most organizations decide on an RTO and RPO for each application, and that dictates what type of backup system to use and how often to run backups.

Once your organization has agreed to an RTO and RPO for each application, the job of the backup system is to assure that these requirements are met across the entire organization. Once configured, monitor the system and test your backups occasionally by doing either a partial or full recovery of your environment.

A well-run backup and recovery system shouldn't be something that you have to spend a lot of time managing. Monitoring the backups is an important job and should be easy to do. The more centralized the backup system and fewer backup products used to build your backup system, the easier it will be to accomplish this important task.

Following the 3-2-1 rule

This rule refers to at least three copies or versions of data, stored on at least two different types of media, one of which should be stored off-site (i.e. in another region).

The moment you make an AWS snapshot, you're storing a copy of your data in a different storage system. However, the part of the 3-2-1 rule most difficult to comply with when using AWS snapshots is the last part — storing at least one copy off-site. It is possible to copy AWS snapshots to another region, but this can be difficult to automate and expensive. Copying data from one region to another creates egress charges and a less efficient second copy that can be quite costly. It can also be difficult to consistently apply this concept across all AWS accounts with centralized management and reporting. It is possible, but not easy without help.

The reason most AWS experts recommend copying your snapshots to another region and account is a result of the business-critical risks if a region is damaged or your account becomes compromised. This is why you must put as much distance as possible between the protected resource and the protection copy — a process typically referred to as creating an "air gap." Copying the snapshot to another region and another account has historically been the best way to accomplish this for AWS resources. But the Druva solution is both easier and less expensive, with air-gapped, automated backups.

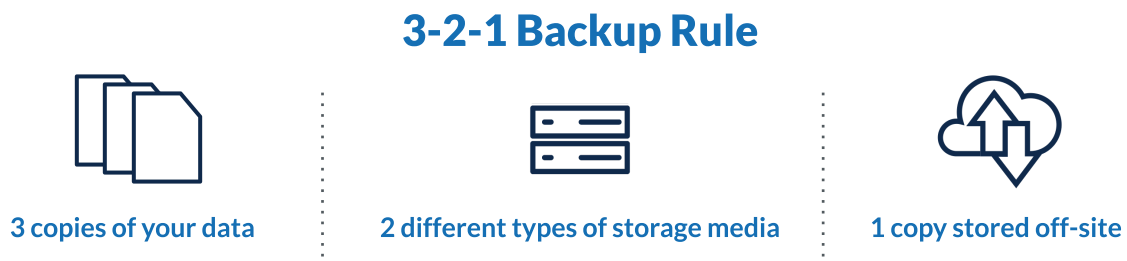

The 3-2-1 rule of backups provides a solid foundation for cyber resiliency.

Limitations of AWS snapshots

Snapshots are a tool — nothing more. As mentioned in the previous section, they are the best tool for restoring AWS resources that have been damaged or deleted for whatever reason. But like any tool, they require knowledge and skill to

use. Failure to use them properly can result in significant cost overruns or — even worse — the complete loss of your organization due to a misconfiguration or actions of a bad actor.

Ransomware is a growing risk

As AWS snapshots are created in the same AWS account, and region as their source, they are at increased risk of being infected by ransomware, tampered with by malicious insiders, or accidentally deleted. AWS themselves recommend leveraging third-party, offsite backup for further protection, saying:

> *"Some newer, smarter ransomware variants are designed to search for stored backups and encrypt or delete them to disrupt recovery efforts. Multiple copies of backups should exist and they should be stored in isolated, offline locations."*
>
> — AWS eBook: "Securing your AWS Cloud environment from ransomware," 2020

In a 2021 study[1], researchers found a large majority of AWS accounts were vulnerable to ransomware.

90%

of S3 buckets in an AWS account are at risk as a result of vulnerabilities common in the platform.

70%+

environments feature machines directly exposed to public internet with permissions exploitable by ransomware.

And attacks on AWS-based resources are ramping up. As a recent example, Portugal's largest media group, Impresa, fell victim to a ransomware attack over New Year's weekend 2022. The largest ransomware attack in the nation's history, attackers were able to gain access and take the site down by manipulating vulnerabilities in Impresa's AWS account.

Increased operational complexity

If you have a single AWS account and relatively few resources, creating an automated backup system with consistent policies is a relatively simple thing to do. However, if you have many AWS accounts in many different regions, enforcing consistent data protection policies across all of these accounts — while monitoring all of their backups — can be incredibly difficult.

In addition, there are some applications that require a more sophisticated data protection mechanism than snapshots. Some databases and Kubernetes are good examples. While snapshots may be part of the overall design, relying solely on snapshots might not allow you to meet your recovery objectives.

Soaring AWS bills

Then there is the issue of ensuring that all snapshots are copied to a different account in a different region. Copying thousands of snapshots across regions is very expensive. While it is lower than the typical egress charge, AWS charges you for copying data between regions. Finally, while AWS snapshots are storage efficient, they can get quite costly if you store

them for any length of time. Amazon EBS snapshot standard storage pricing is significantly costlier than other options for long-term retention ($0.05 per GB of data stored).

Druva: Delivering radical storage efficiency and air-tight security

For applications that need something more than snapshots, Druva's cutting-edge, source-side global deduplication is the most efficient way to back up these applications. Source-side deduplication creates the lowest impact on the system you're backing up, while also sending and storing the least necessary amount of data to dramatically reduce costs — storage cost reductions of up to 50% in some cases. Built-in and automated cold-storage tiering for long-term retention is a cost-efficient storage option for infrequently accessed data. All data is encrypted in transit and at rest, and the customer always maintains all the encryption keys. This zero-trust architecture means no Druva employee ever has access to your data.

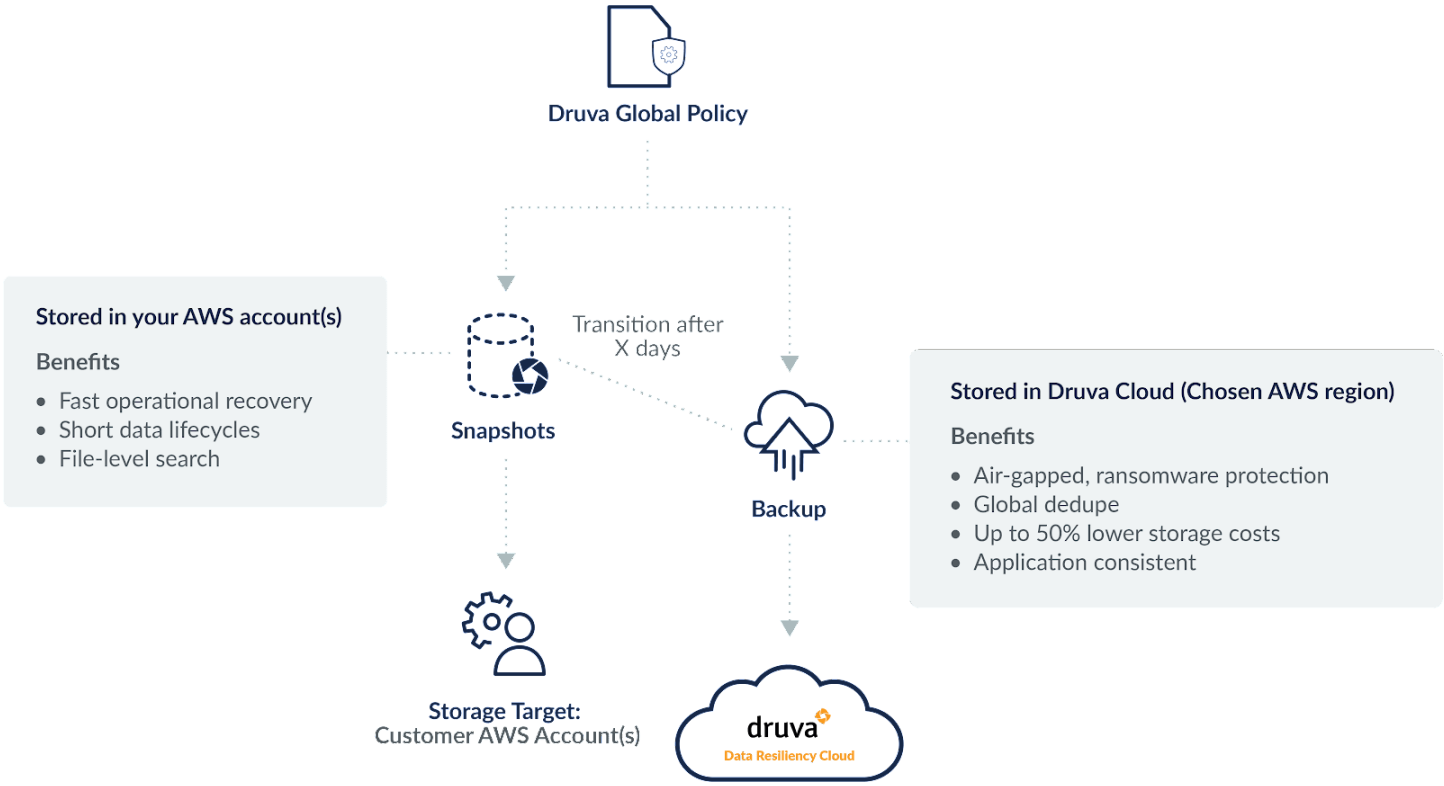

For increased security, all copies of your data stored in the Druva Data Resiliency Cloud are air-gapped from the original AWS production environment. Ransomware attacking your primary environment will find no route between your protected resources and the encrypted protected copy stored in the cloud.

Best-in-class AWS data protection that doesn't break the bank

Druva has combined these two technologies to give you snapshots plus an air-gapped, deduplicated, truly immutable copy stored in the Druva Cloud. By eliminating unnecessary cross-account and cross-region snapshot copies, you receive all the benefits of secure backup without the complexity or cost. Druva estimates 30-50% savings for the typical customer.

Explore example cost savings scenarios in the following diagrams.

Year 1 Savings $789K

Year 1 % Savings 46%

Breakdown, After 1 Year

Total cost with Druva	$906,347.26
Original cost with AWS	$1,696,050.03
Savings	**$789,702.77**
Savings %	**46.56%**

Retention

Daily	7
Weekly	4
Monthly	12
Yearly	3

EC2 instances and source data

Total Instances	1000 EC2
Total source	500TB

3 Year Cost Comparison

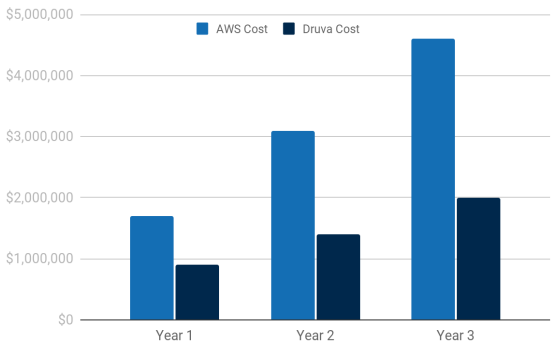

Benefits of EC2 Backup

	with Druva	Snapshots only
Fast RTO	✓	✓
Compression	✓	⊘
Long Term Retention	✓	⊘
Deduplication	✓	⊘
Ransomware Protection	✓	⊘

*Change rate and data growth used in the TCO are on the basis of assumptions as per average industry standards

***Note:** Get free Amazon EC2 snapshot orchestration licenses for every instance you backup to Druva Cloud for additional savings!*

Visibility and control from one easy-to-use console

Create organization-wide policies, authenticate Druva with your various AWS accounts, and Druva will enforce those policies across your entire environment. Druva provides a simple and unified management and reporting console to configure and monitor all of your AWS backups across your entire environment. As a result, customers receive the following benefits:

- Single-click enablement enables restore and back up in seconds — across regions and accounts
- Agentless approach removes the need for any hardware or software
- Unified view of all snapshot and backup jobs in one place simplifies management
- Visibility into the status of all updates or changes through the *Audit* page

Customize to meet your needs

Druva has expanded its capabilities for Amazon EC2 backup and recovery, and now offers two extra options: back up to the Druva Cloud, and immutability. Customers determine how long to retain snapshots as well as how long to keep backups in the Druva Cloud. If the immutable option is checked, even an authorized user will not be able to delete backups prior to

their designated retention period, nor will they be able to reduce the retention period for a given backup once it has been created.

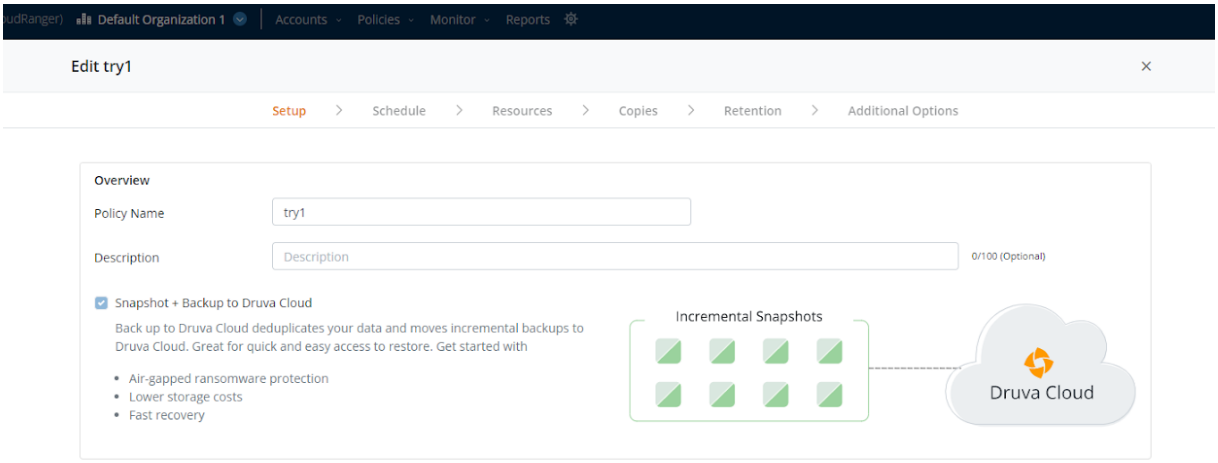

If a customer initiates a restore, Druva uses the local snapshot copy if available for fast RTO, otherwise, the deduplicated backup copy stored in Druva Cloud will be restored.

This snapshot plus air-gapped, deduplicated, immutable copy gives you the best of both worlds. You get AWS snapshots for easy, operational recovery while using the Druva Cloud copy to protect your backups from those things that would damage them. Have your cake and eat it too.

Having your cake and eating it too

Protection of AWS resources should start with AWS snapshots, and Druva makes them incredibly easy to use. The Druva Data Resiliency Cloud is the industry's first 100% SaaS platform to truly combine backups, snapshots, and disaster recovery into a single solution, at enterprise scale. Druva removes the need for cross-account and cross-region snapshots and provides secure, air-gapped backups for ransomware recovery readiness with storage cost savings up to 50% compared to the native AWS backup solution.

Next steps

Visit the Druva site to learn more about how we provide comprehensive data resilience for your Amazon EC2 environments and schedule a free live demo to experience Druva for yourself.

druva Sales: +1 888-248-4976 | sales@druva.com

Americas: +1 888-248-4976
Europe: +44 (0) 20-3750-9440
India: +91 (0) 20 6726-3300
Japan: japan-sales@druva.com
Singapore: asean-sales@druva.com
Australia: anz-sales@druva.com

Druva is the industry's leading SaaS platform for data resiliency, and the only vendor to ensure data protection across the most common data risks backed by a $10 million guarantee. Druva's innovative approach to backup and recovery has transformed how data is secured, protected and utilized by thousands of enterprises. The Druva Data Resiliency Cloud eliminates the need for costly hardware, software, and services through a simple, and agile cloud-native architecture that delivers unmatched security, availability and scale. Visit druva.com and follow us on LinkedIn, Twitter, and Facebook.

[1] Ermetic, "Ermetic Finds Majority of AWS Accounts Surveyed are Vulnerable to Ransomware Due to Misconfigurations," Published 7 October 2021.