druva

White paper

How to add disaster recovery to your AWS environment

Protect vital AWS workloads with Druva disaster-recovery-as-a-service

Q221-20169

Introduction

Enterprises need to prepare for potential failures to their cloud infrastructure due to user errors, malicious external and insider threats, unforeseen disasters, operational mishaps, as well as site outages. For example, ransomware attacks and associated costs continue to annually rise world-wide, impacting various industries and sensitive corporate data, including AWS workloads. According to MIT Technology Review, "The potential cost of ransomware in the United States last year was over $7.5 billion."[1] Cybersecurity firm, Emsisoft, reported 113 governments and agencies, 764 health-care providers, and up to 1,233 individual schools affected by ransomware in America. Big cities including Baltimore and New Orleans were both struck by ransomware attacks last year.[2] In the case of AWS workloads and data protection, it's critical that IT managers create a comprehensive disaster recovery (DR) plan in order to actively protect against unpredictable threats and events. An AWS Public Sector leader comments, "All these customers I talk to are running on data centers, unpatched, they talk about having a DR strategy — but guess what, they don't have that."[3]

That's why enterprises must also prepare across availability zones, regions, or accounts within AWS. Systems, sites, and networks can fail no matter where they are located or who is managing them, and mission-critical service availability and protection of corporate data are essential no matter where services reside. The AWS Shared Responsibility Model outlines the difference between Security of the Cloud, AWS' job, and Security in the Cloud — the customer's responsibility. Yet Amazon is clear: "Workloads that are migrated or created in AWS are not implicitly protected with disaster recovery (DR) capabilities."[5]

> *Organizations with critical workloads in public cloud IaaS and PaaS must still perform disaster recovery (DR) planning, and many are not aware of the cloud differences, limitations, and risk.*
>
> — Gartner[4]

This white paper provides guidance and recommendations to cloud ops admins or cloud architects — you'll learn how to ensure your AWS workloads are protected in the event of disasters and threats such as a ransomware or a cybersecurity attack, a loss which happens more often than you might think. You'll also gain insights into how SaaS-based data protection, like Druva, can benefit AWS users, and how you can configure Druva disaster-recovery-as-a- service (DRaaS) for an AWS environment.

> *An organization's DR plan should be continually reviewed and enhanced. Any gaps identified during DR testing need to be prioritized, remediated and incorporated into the revised DR plan.*
>
> — Gartner[6]

[1] MIT Technology Review, Ransomware may have cost the US more than $7.5 billion in 2019, 2020
[2] EMSISOFT, The State of Ransomware in the US: Report and Statistics 2019
[3] ZDNet, Avoid ransomware by moving to the cloud, says AWS Public Sector boss, 2019
[4] Gartner, Implementing Disaster Recovery for Public-Cloud Workloads, Brian Adler, May 2020
[5] AWS Shared Responsibility Model
[6] Gartner, Implementing Disaster Recovery for Public-Cloud Workloads, Brian Adler, May 2020

DR for AWS workloads: evaluate and consider the challenges

When evaluating disaster recovery solutions for their various AWS environments, enterprises should look for the following capabilities for maintaining business continuity throughout disasters and service disruptions:

- End-to-end automated recovery procedures
- TCO of DR solutions and comprehensive support for AWS native services and resources
- Built-in orchestration and runbook execution
- Automation of failover compliance testing and efficient failback
- Centralized reporting and auditing

One of the biggest challenges with DR planning is the lack of adequate testing and verification. In this situation, you will have no idea as to whether or not you'll actually be able to recover from a disaster and whether your RTOs and RPOs are valid. Infrequent testing of backup environments puts businesses at substantial risk when outages eventually occur.

Small enterprises have a higher adoption rate of cloud technology, with 93 percent of companies using it. This is compared to 82 percent of mid-sized businesses and 81 percent of large businesses. Using cloud backup offers a number of advantages including ease-of-access and affordability. Currently, 36 percent of businesses use this, and a further 23 percent plan to add the technology within the next year.[8]

> *84 percent of all businesses store data or backups in the cloud, with a further eight percent planning to do so within the next year … cloud-based Disaster Recovery-as-a-Software (DRaaS) will be used by 59 percent of businesses by 2021.*
>
> — Unitrends[7]

A successful cloud DR solution should make it simple to set up and update DR plans with automation and take advantage of all the merits of the AWS cloud. In addition, thought should be given to these factors:

- **Flexibility and availability —** While AWS offers 20+ regions, 200+instance types, and 100+ services, not all regions have the same number of availability zones or support all instance types. As an example, Northern Virginia has 6 Availability Zones (AZs) while Ohio has 3 AZs.
- **Cost considerations —** An AWS subscription model requires you to only pay for services used but there are transfer (egress) charges when data is moved out of a region. Likewise, not all regions cost the same.
- **Compliance and controls —** It is important to consider data residency requirements and operational issues resulting from incorrect security settings in your DR sites. AWS Identity and Access Management (IAM) roles are account-specific and AWS Key Management Service (KMS) keys used for data encryption are region-specific.

The importance of cloud-native DR

Legacy DR solutions often leverage storage hardware replication that are not appropriate cloud storage targets. And for any storage capability working in an AWS context, the difference between cloud-enabled and cloud-native is considerable. Cloud-enabled DR solutions are typically built by taking traditional data center applications and adding-on features for cloud access. A cloud-native application is built, from the ground up, with principles of multi-tenancy, microservices, elastic scaling, and easy integration and administration in mind.

[7] Unitrends, Data Protection, Cloud, and Proof DraaS Delivers, 2019
[8] Unitrends, Data Protection, Cloud, and Proof DraaS Delivers, 2019

Cloud-native DR adds clear advantages for AWS users:

- Ability to recover VMs in the cloud within minutes
- Lower cost as subscription models replace capital expenditures
- Expanded DR coverage for applications/workloads
- Workload mobility across cloud regions

Druva is here for you

Druva disaster-recovery-as-a-service (DRaaS) is itself built on AWS and integrates natively with a user's AWS environment to provide a secure operating environment, available on demand, for comprehensive DR operations. You'll be able to:

- Seamlessly create cross account DR sites based on source sites by cloning VPCs and their dependents
- Set up backup policies to automatically create and copy snapshots of EC2 and RDS instances to DR sites based on RPO requirements
- Set up SLO-based DR plans with options to schedule automated test of DR plans and ensure compliance
- Monitor execution of DR plans from the console
- Generate compliance reports for DR failover and test execution

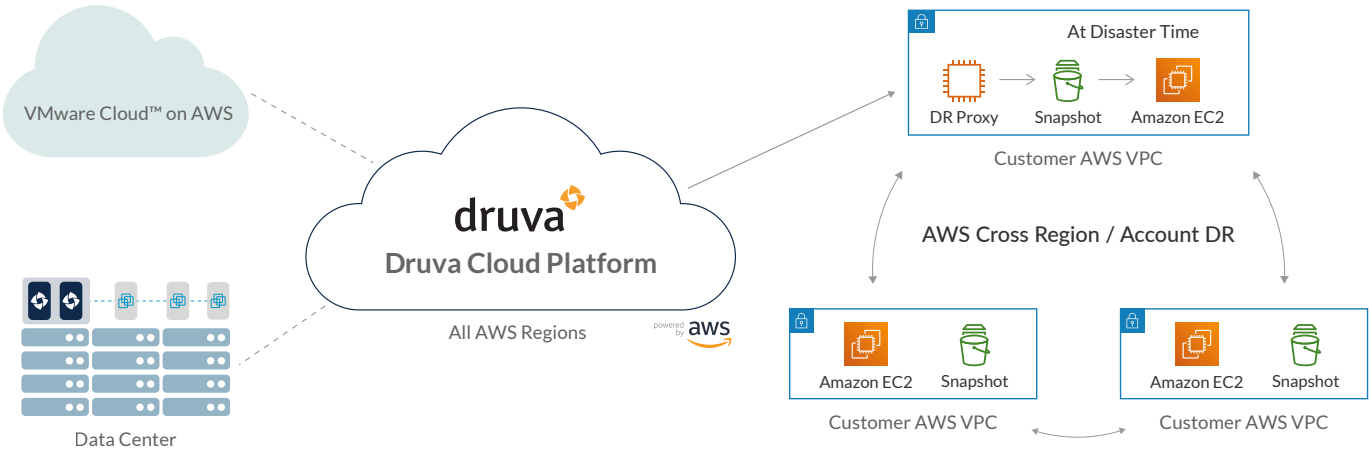

Features also include automated runbook execution, tight AWS integration, and simplified orchestration and testing, helping protect and recover data at scale. Druva provides the flexibility to adopt evolving infrastructure across geographic locations, adhere to the regulatory footprint, and recover workloads more quickly following disasters. This unified solution offers optimized recovery time objectives (RTO) and supports taking snapshots as low as every five minutes, also improving recovery point objectives (RPO).

Druva helps lower costs by eliminating traditional administration and maintenance of storage hardware and software, upgrades, patches, and integrations. And it transfers data efficiently by deduplicating content and incrementally backing up changes only rather than every data write.

How to configure Druva DRaaS

This workflow shows the ease with which you can configure Druva DRaaS in an AWS environment:

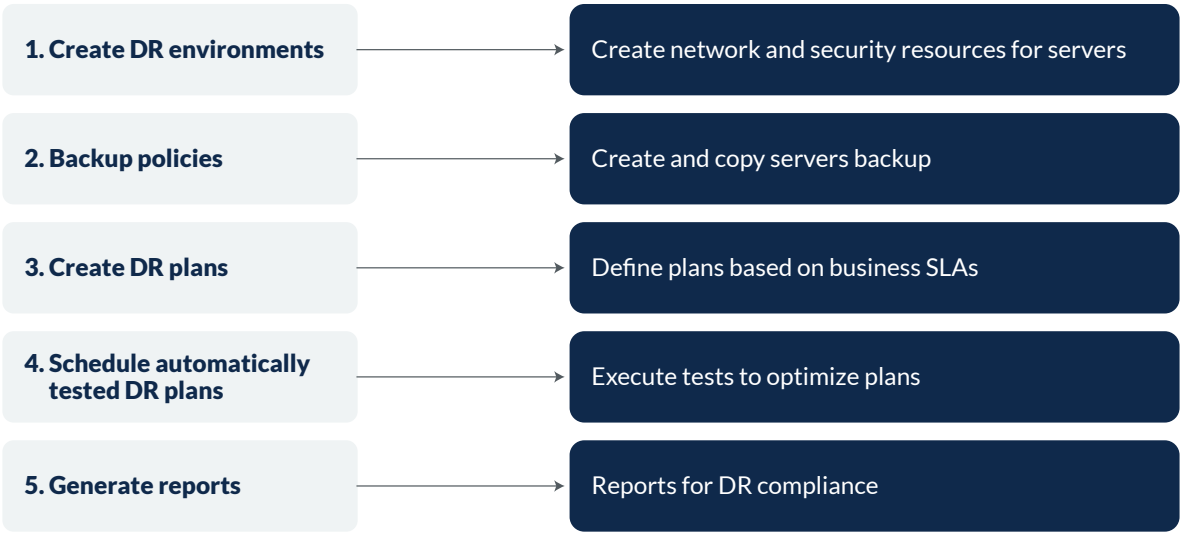

1. **Create DR environments**

   In the Druva control panel, create a source environment by selecting network and security resources such as Amazon Virtual Private Networks (VPNs) and security groups of the Amazon EC2 and Amazon RDS instances that are part of your application.

   Use the One-Click Clone option to clone the environment to another AWS account and region to automatically create a target DR environment based on production configurations. Restored instances will be in an identical environment despite being in a different region or account. This is ideal for when the infrastructure does not currently exist in your chosen destination.

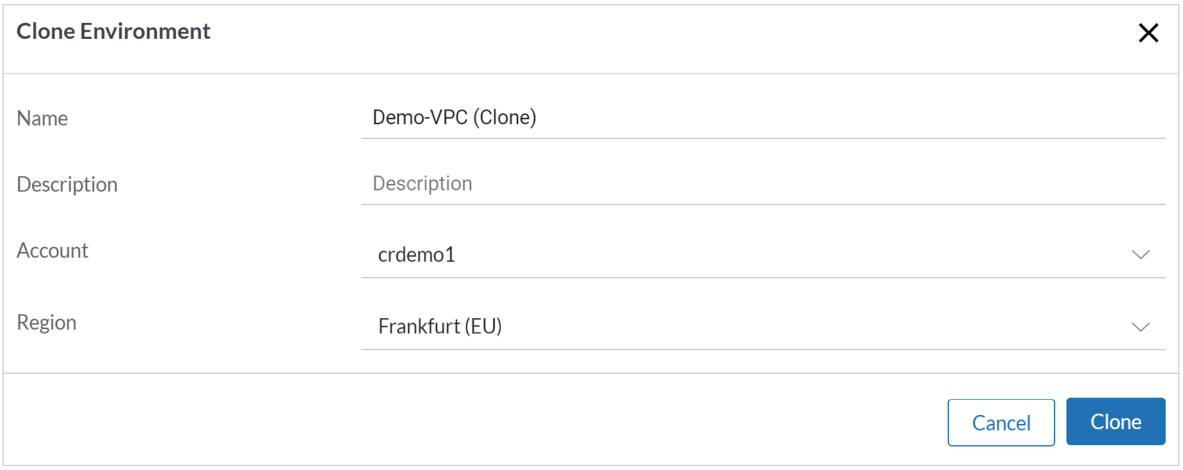

The environment cloning process captures all the areas listed in the following table.

		Clone settings
Amazon VPC	✓	CIDR ranges preserved
Subnets	✓	CIDR ranges preserved, AZs allocated in round robin
Route tables	✓	Routing preserved
Internet gateways	✓	Routing preserved
Egress only internet gateways	✓	Routing preserved
DHCP options sets	✓	Options preserved
NAT gateways	✓	Routing preserved
Elastic IPs	✓	New addresses allocated and assigned to Amazon VPCs for NAT Gateways and pre-allocated for instances with EIPs
Security	✓	Security groups
Network ACLs	✓	Rules and associations preserved
Security groups	✓	Ingress and egress rules preserved

2. **Create backup policies**

   Create a backup policy to backup and copy Amazon Machine Instances (AMIs) and snapshots based on the RPO to the target AWS account and region.

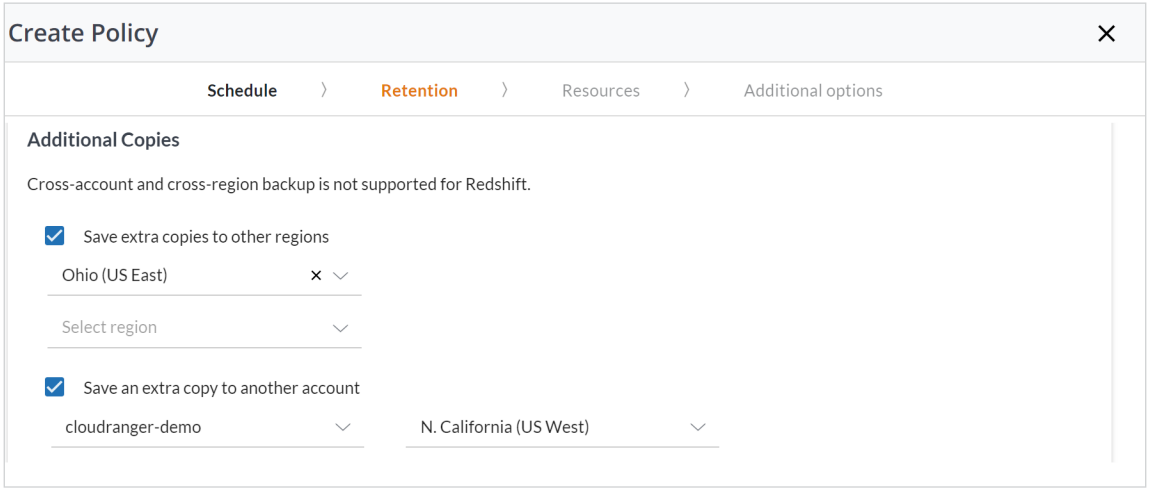

3. Create DR plans

Select source and target environments:

- Set up the DR network and security mapping for resource instances in the event of a failover
- Select appropriate Amazon EC2 and Amazon RDS instances based on IDs or AWS TAGS

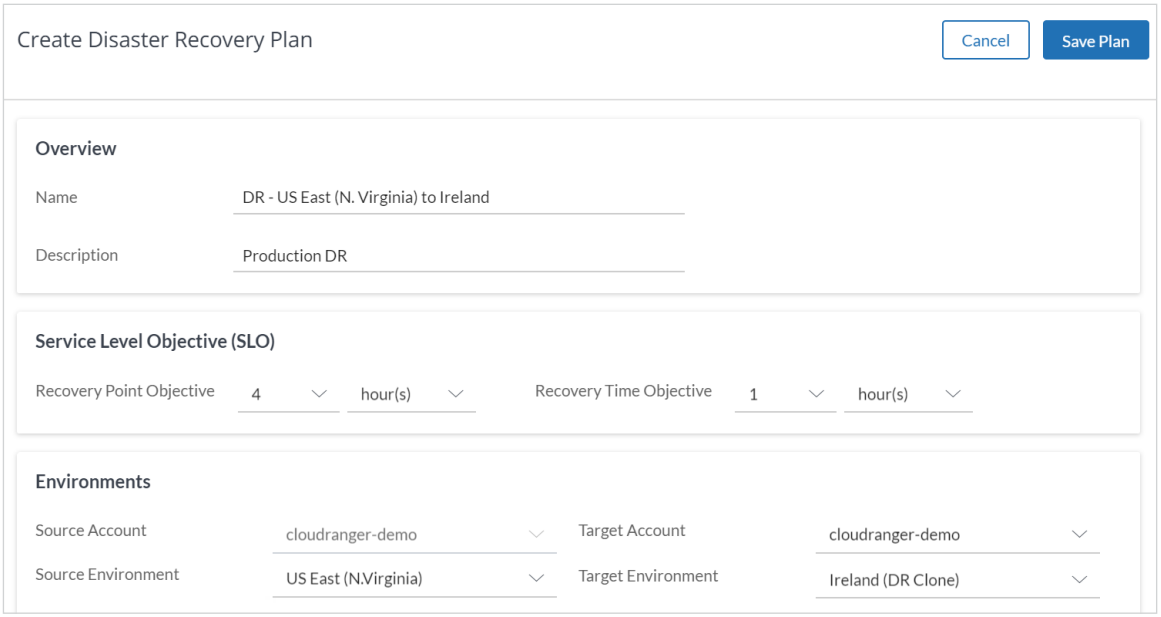

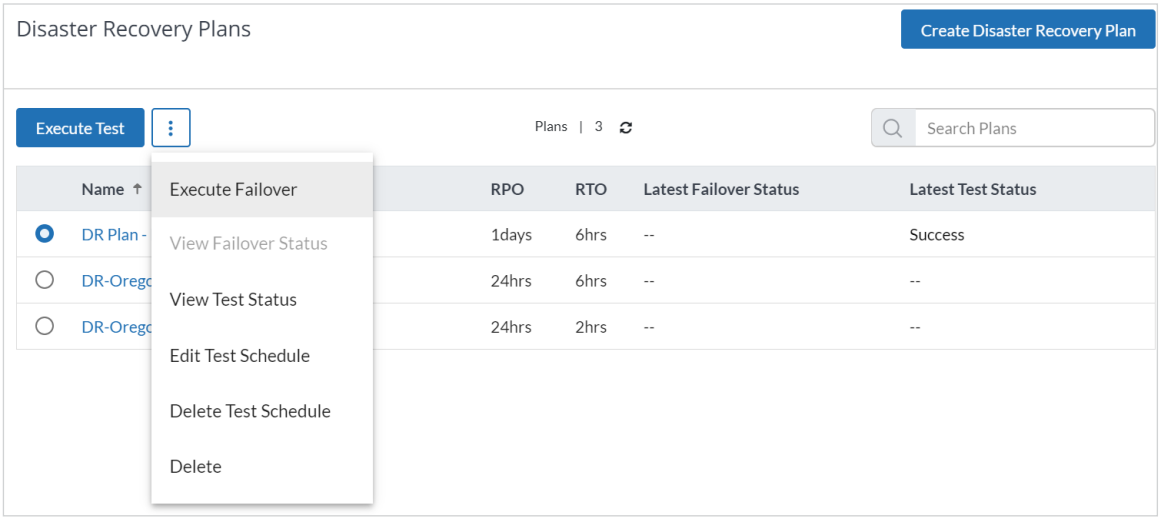

4. **Schedule automatically tested DR plans**

Define a recurring schedule to automatically test DR plans and verify if plans meet RPO and RTO requirements. The test options offer retention settings to clean up the instances after testing to minimize AWS DR test costs.

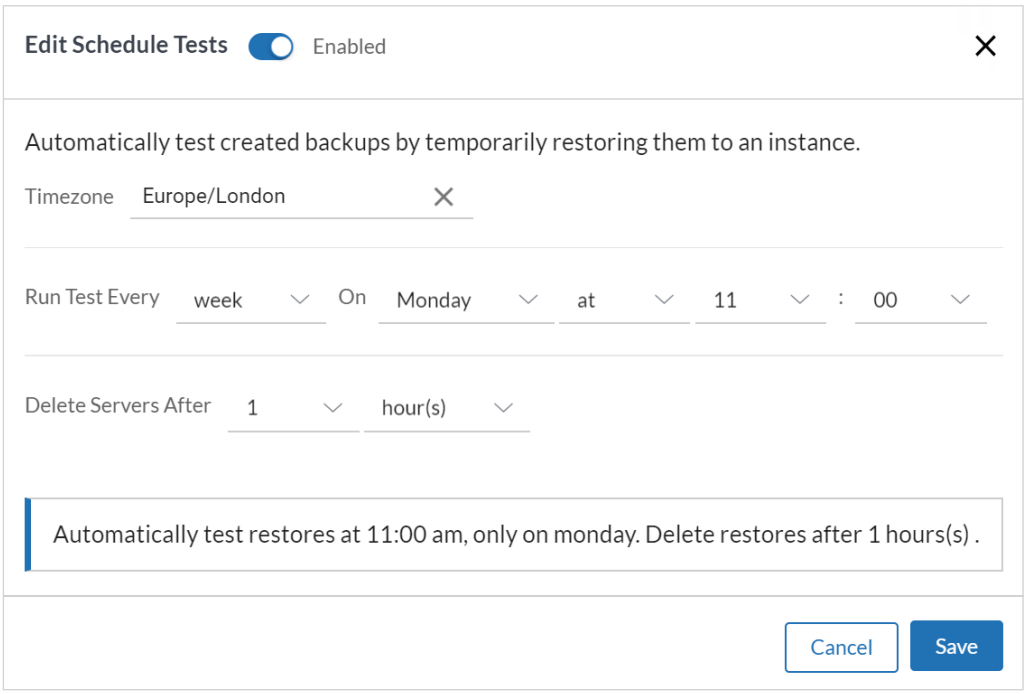

5. **Generate reports**

View and select DR plan test reports that can be shared to meet compliance requirements.

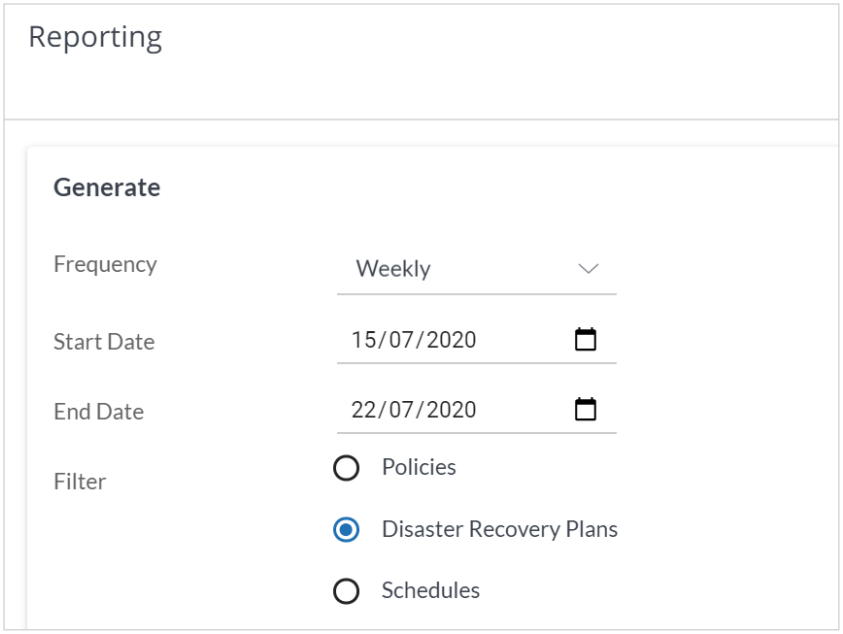

How Druva DRaaS works

Druva DRaaS backs up on-premises and cloud workloads to either the Druva Cloud Platform or assigned targets. It converts VMs to AWS EBS snapshots that are kept-at-the-ready inside a customer's AWS VPC for immediate spin-up of AWS EC2 instances at the time of a failover. Users can recover on-premises (failback) or in the cloud (failover) across any AWS region or accounts.

Users manage all aspects of data protection for AWS workloads through a single control panel:

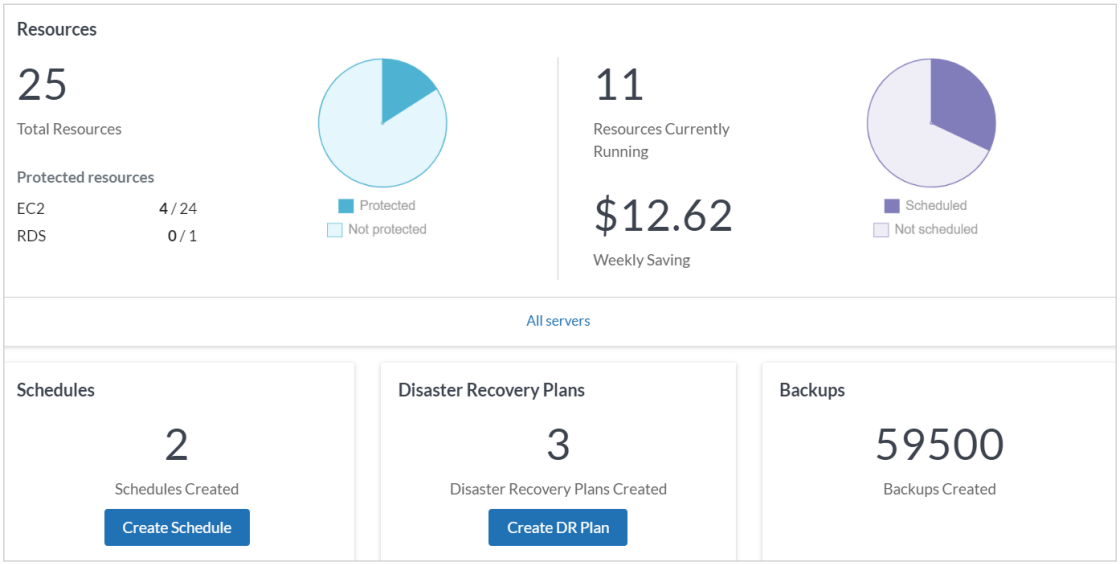

The central dashboard shows all administrative activities such as policy management, replication, and reporting, and provides visibility into the organization's entire AWS footprint. It strictly adheres to compliance and regulatory requirements such as legal holds and eDiscovery enablement.

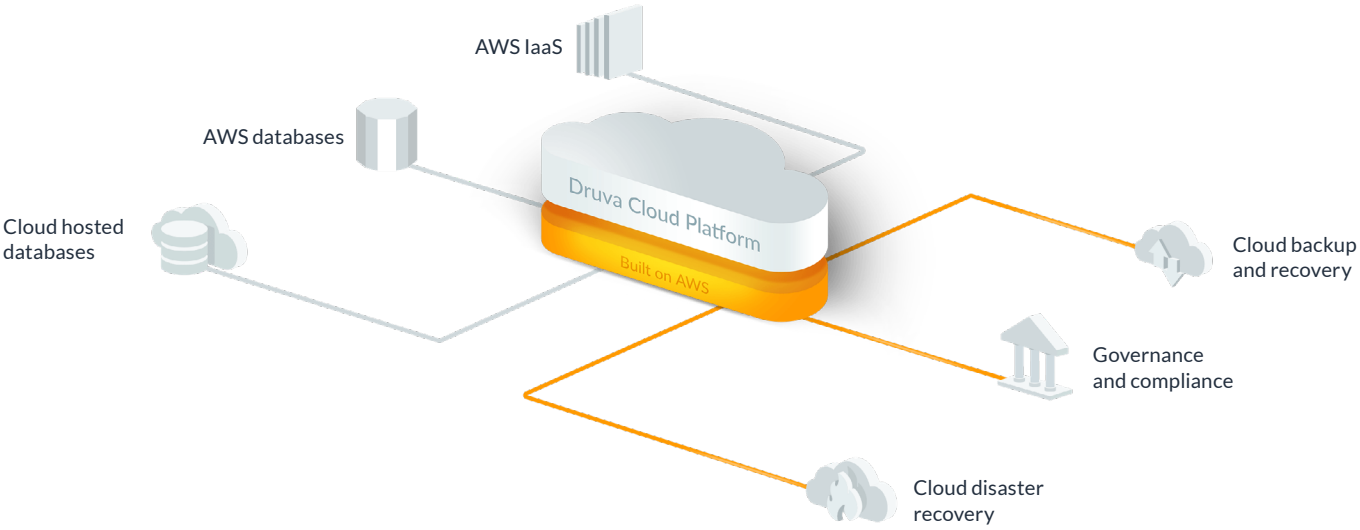

Access to AWS accounts via AWS-native APIs and microservices is controlled by role-based Amazon Identity and Access Management (IAM). This enables centralized snapshot policy orchestration and disaster recovery across an organization's AWS accounts and regions for Amazon Elastic Compute Cloud (EC2), Amazon Elastic Block Store (EBS), Amazon Relations Database Service (RDS), Amazon DynamoDB, Amazon Redshift, Amazon Document DB and Amazon Neptune. Druva also supports hybrid workload failback to VMware on AWS or on-premises data centers to align with enterprise compliance.

Conclusion

Doing business on AWS is a remarkably dependable and safe proposition. But just as Microsoft warns their Microsoft 365 customers to implement third-party data protection, AWS holds the customer fully responsible for any losses resulting from service outages. Fortunately, the easy-to-deploy, cost-efficient, cloud-native Druva Cloud Platform eliminates any potential loss by providing comprehensive DRaaS capabilities in any AWS environment.

Ensure your AWS workloads are protected with DRaaS druva.com/products/aws-backup/

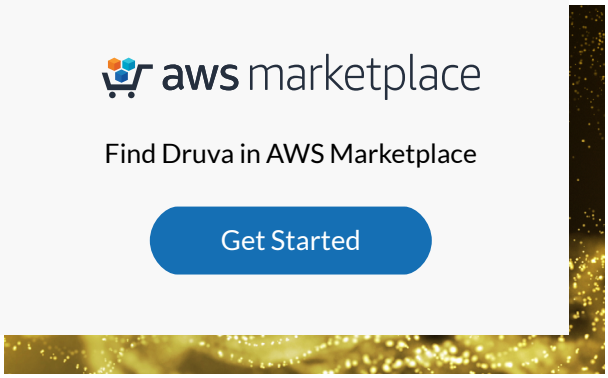

druva

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: +81-3-6890-8667
Singapore: +65 3158-4985
Australia: +61 1300-312-729

Druva™ delivers data protection and management for the cloud era. Druva Cloud Platform is built on AWS and offered as-a-Service; customers drive down costs by up to 50 percent by freeing themselves from the burden of unnecessary hardware, capacity planning, and software management. Druva is trusted worldwide by over 4,000 companies at the forefront of embracing cloud. Druva is a privately held company headquartered in Sunnyvale, California and is funded by Sequoia Capital, Tenaya Capital, Riverwood Capital, Viking Global Investors, and Nexus Partners. Visit Druva and follow us @druvainc.