Druva delivers cloud disaster recovery

The challenge

In the era of cloud, business continuity has become one of the top five key topics for leadership to reduce revenue impact of any event that disrupts service. Traditional disaster recovery solutions are struggling to follow data into the cloud because they were designed and optimized for client/service environments. Your business needs a solution that capitalizes on the benefits of cloud while ensuring your business is always up and serving customers.

The solution

Druva's SaaS solution eliminates the cost, provisioning and management tasks associated with legacy data protection solutions. Layered on top of Amazon Web Services (AWS), Druva delivers globally available, single-click recovery for on premise and cloud-based workloads. A single license combines Druva innovation with AWS' highly available infrastructure, eliminating ingress/egress charges and enabling your business to harness the speed and scale of cloud-based disaster recovery from a unified management console.

Business challenges

- **Achieving service level objectives.** Recovery times acceptable to customers have been continuing to decrease, to the point that 75% of all business applications have a recovery time objective of just 15 minutes to 1 hour.[1]
- **Using tape for disaster recovery use cases.** While organizations have traditionally turned to tape to meet their disaster recovery (DR) needs, in today's 24x7 landscape companies are finding that tape does not provide an acceptable recovery time for business critical apps.
- **Maintaining the disaster recovery plan/runbook.** Paper-based runbooks contain the actual sequence, dependencies and user connectivity for each recovery step, but they are difficult to keep up-to-date and expensive to verify recovery processes work as designed. Recovering applications in 15 minutes to 1 hour requires automated workflows that orchestrate each recovery steps to bring your business back online.

Benefits

- Powered by AWS, Druva works with Amazon Simple Storage Service (Amazon S3) storage family to automatically store data that needs to be retained longer than 90 days for compliance and legal hold purposes.
- Druva's cloud-native solution delivers recovery time objectives of 15-20 minutes and recovery point objects of 1 hour to meet the needs of today's business applications.
- Druva's automation and orchestration delivers one-click recovery and setup of business applications to ensure your business is up and running as quickly as possible after a disaster has been declared.

How it works

The first step for most customers is to backup their data into the Druva Cloud Platform. Using source deduplication, Druva works in tandem with Amazon DynamoDB to reduce duplicate data by slicing new data into chunks. A 120-bit value, known as a hashtag, is established by running the data through SHA-1.

[1] ESG, 2018

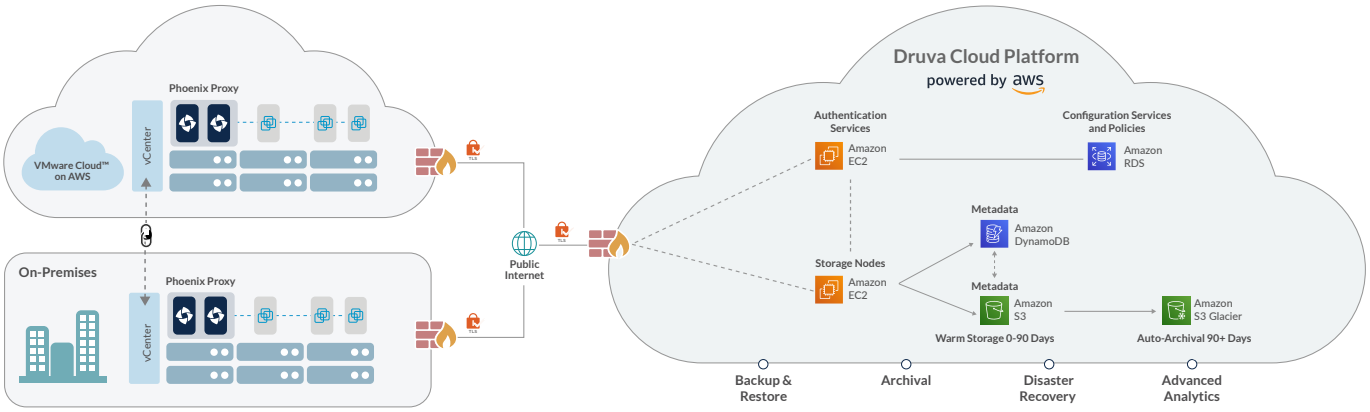

Figure 1

Druva sends the hashtag to Amazon DynamoDB where a comparison against the existing hash table is completed by answering a simple question: "Is this a new or duplicate hashtag?" If the answer is "no," Druva dedupes and encrypts the data and transmits it to the cloud. If the answer is "yes," then the hashtag is recorded in Amazon DynamoDB.

To automate disaster recovery, customers must have their own Amazon Virtual Private Cloud (VPC) instance and pre-defined disaster recovery workflows, which can be created from the Druva unified management console. The workflow defines the application components being recovered, size of the AWS recovery environment, and specific order for recovering data.

The Druva proxy inside of the customer's VPC injects wrappers and drivers into the identified virtual machines, which converts them into Amazon Machine Images (AMIs) so they can be recognized by Amazon Elastic Compute Cloud (Amazon EC2). Lastly, Druva takes a snapshot of the AMI, providing customers with a lower-cost alternative to storing the Amazon EC2 image.

When a disaster is declared, a single click begins the process of recovering all virtual machines defined in the DR workflow within 15-20 minutes regardless of their quantity or size. Customers can replicate virtual machines, clone full VPCs, and relocate them across multiple AWS regions and accounts to accommodate test and development use cases.

Summary

As applications become more complex, the number of interdependencies between components and other applications increases - further complicating your business continuity strategy. Druva's automated, cloud-native disaster recovery solution streamlines DR processes, ensuring your business is up and running quickly after a disruptive event.

Learn more about Druva at: https://www.druva.com/solutions/cloud-disaster-recovery/

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.