



CUSTOMER STORY

LEADING GLOBAL CRUISE LINE

AWS Application Resilience with Rapid Failover and Failback for a Leader in the Cruise Industry

Industry

Travel and Leisure

Key Metrics

- Rapid failover and failback of AWS application environments
- 2030 total resources protected
- Failover and failback RTO < 15 minutes

BACKGROUND

A leading global cruise holding company with a portfolio of three award-winning brands and a vision to be the vacation of choice for everyone around the world. It provides cruise experiences for travelers with itineraries in North America, the Mediterranean, the Baltic, Central America, Bermuda, and the Caribbean. It also offers an entirely inter-island itinerary in Hawaii.



First time in over multiple years of working across multiple recovery solutions, I am happy to be part of the test exercise that successfully failed over and failed back the application in running state with such ease."

Lead Cloud Engineer

A Leading Global Cruise Line

CHALLENGE

A leading global cruise holding company wanted their business-critical environment running on AWS to be resilient against any and all kinds of failures and was looking for a one stop solution to ensure their cloud environment is completely protected. The company just had migrated a lot of its on-prem business critical applications including the one that allows their customers to book their luxury cruises, to AWS in the US-East-1, North Virginia region. The IT team, coming from the datacenter mindset, their definition of success criteria for a DR solution was to successfully failover and then failback to the original environment. So, they wanted their cloud applications to failover and failback to make sure they can move back to the original production environment after a failover event.

Adding to the complexity, recovery of the applications, cloud services, dependencies and the data should be to a pre-created VPC in the DR region, in this case Oregon US-west-1 region. Recovering into an existing VPC is important as their production environments had some central systems running in their datacenters and both the primary and DR sites in the cloud needed to communicate to their DC based central systems in order to operate during a partial disaster scenario.

Moreover, the complete failover was expected to be fully automated along with the internal DNS configured in the AWS Route 53. Failback was expected to revert the Route53 DNS settings back to their original state so that the application endpoints talk to each other automatically without any manual intervention.

"First time in over multiple years of working across multiple recovery solutions, I am happy to be part of the test exercise that successfully failed over and failed back the application in running state with such ease."

Lead Cloud Engineer, A Leading Global Cruise Line.

SOLUTION**Completely Automated Failover and Failback**

The cloud engineering team, through one of the common cloud partners, reached out to Commvault Cloud Rewind, to enable protection of the entire environment with all the cloud services, configurations and the applications data. The success criteria was to protect the complete environment end to end, along with the ability to failover and failback in the cloud environment rapidly with complete automation.

Commvault Cloud Rewind worked along with the cloud engineering team to understand their DNS challenges and allowed a backup of DNS configurations along with entire environments cloud services, configurations and the application data based on the application RPO and RTO requirements. The customer's critical applications were set to 1 hour RPO and other regular applications were set to a 24 hour RPO by configuring hourly and daily policies respectively.

Commvault Cloud Rewind allows failover and failback customizations using programmable webhooks. Commvault Cloud Rewind and cloud engineering teams setup webhook with AWS lambda functions within their secure cloud environment. These webhooks automatically back up the DNS configurations in a pre-recovery webhook call.

After a failover event another webhook updates the Route 53 with the recovered instance details like recover EIP address, and RDS endpoints. After a successful failover using Commvault Cloud Rewind's single-click Recover operation that automatically creates the entire environment along with the application dependencies and the data, failed over instances were protected in a new Cloud Assembly in the recovered region. The failed over application environment was then run as a production by cutting off the current running production environment. The new production environment was kept running for 45 days before failing back to the original region. Before failing back, appropriate cleanup was done to the original production instances as it is now 45 days old. After failing back from the DR site to the original production site, a similar post recovery DNS update was run to re-configure EC2 instances and RDS end points along with application verification.

Overall, Commvault Cloud Rewind delivered a robust failover and failback that was practically impossible in the company's data center world: to do a failover; run the application in the failover site for 45 days; and failback without doing any manual work except for login and verifying that the application is working as if nothing happened in the backend.

To learn more, visit commvault.com