WHITE PAPER

Cloud Rewind Cloud Resilience Copilot Technical Overview

WHITE PAPER

Commvault

Table of Contents

CONTENTS

WHITE PAPER

Commvault

Introduction to Cloud Rewind Resilience Platform

Cloud Rewind Cloud Resilience is a comprehensive solution designed to ensure availability and recovery of cloud-based applications. Cloud Rewind helps organizations achieve resilience in the face of disruptions and outages, such as ransomware attacks, cloud infrastructure failures, software failures, security breaches, or natural disasters.

Cloud Rewind ensures resilience by continuously discovering and mapping dependencies of cloud resources of distributed systems. Cloud Rewind protects all the discovered cloud configurations, dependencies, and application data with the policies, per organization RPO, using a patented Dual-vault Cloud Time Machine technology. After an application outage or disruption, organizations can invoke Cloud Rewind Recovery-as-Code capabilities to rapidly recover applications or even rebuild the entire application environment with one-click in another region or another
cloud account.

Cloud Rewind holds a third-party AICPA audited SOC (Service Organization Control) Type II certification, which validates the effectiveness of its security controls and safeguards, and the availability of the SaaS platform, providing assurance to customers about the security and reliability of services.

USE CASES OF CLOUD REWIND

1. **Cloud Infrastructure Backup:** Backup your cloud configurations and dependencies continuously away from your production cloud for recovery and rebuilds.

2. **Application Data Backup and Restore:** Backup your distributed cloud application datain cloud-native format for rapid point-in-time restores without sacrificing data residencyand sovereignty. Cloud Rewind automates backup of all the databases and data services platforms from an application-centric perspective to avoid organizational risk of not backing up applications data as cloud resources dynamically change.

3. **Cloud-Native Data Replication:** Cloud Rewind continuously replicates your application datafor point-in-time recoveries across other regions or cloud tenant accounts per your RPO.

4. **One-Click Rebuild from Ransomware Attacks:** Rebuild entire distributed systems, cloud resources, dependencies, application images, and application data away from the affected regions or accounts with one-click.

5. **Cloud-Native Disaster Recovery:** Rapidly recover partial or full distributed applicationsor entire cloud account resources from disasters or other disruptions.

6. **On-Demand Cloud Spaces:** Automatically create sandboxed cloud spaces for cyberthreat scanning or dev/test or fault injection testing without affecting production at anypoint-in-time in any region.

7. **Control Cloud Costs:** Avoid multi-region cloud architecture or pilot light to control cloudcosts. Remove development, maintenance and operations costs with Cloud Rewind on-demand cloud space rebuilds.

PRIMARY OPERATIONS OF CLOUD REWIND

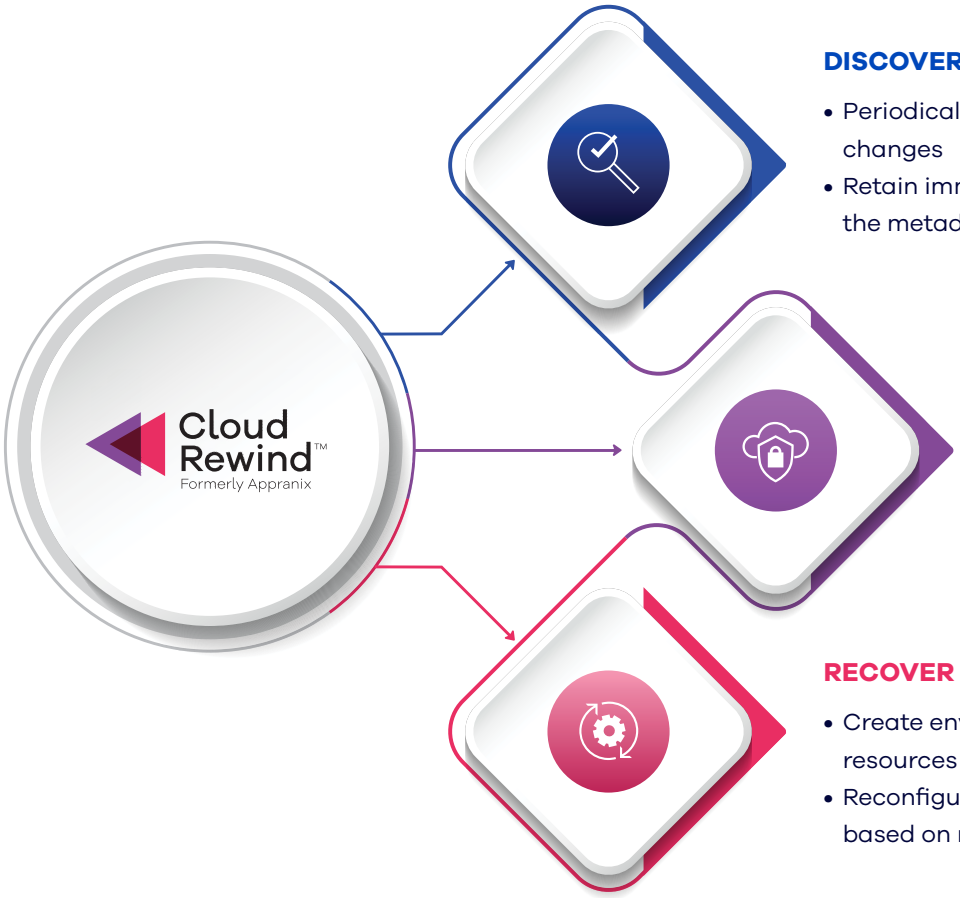

DISCOVER

- Periodically discover changes
- Retain immutable copy of the metadata

PROTECT

- Create Snapshot based on policy
  - Storage Snapshots
  - Config/Metadata Snapshots
- Replicate across regions
- Retain/Prune based on policy
- All captured information are stored in an immutable timeline

RECOVER

- Create environments/ resources from protections
- Reconfigure environment based on recovery requirement

DISCOVER

Cloud Rewind discovery process involves analyzing cloud application environments to gain comprehensive visibility. It examines infrastructure, configurations, dependencies, and interactions within the distributed application system to prepare recovery with dependencies after a disruption or an outage.

PROTECT

Cloud Rewind patented Dual-vault Cloud Time Machine offers some of the best protections against various disruptions for cloud-native and cloud-enabled applications. Cloud Rewind splits cloud configurations backup and replication, and application data backup and replication into two different vaults to avoid any form of compromise that could risk organizations recovery from outages.

CLOUD CONFIGURATION VAULT

Cloud Rewind backs up cloud configurations and dependencies continuously over a 256-bit encrypted channel as point-in-time snapshots away from the production cloud to provide a level of security that is not available in the native clouds. This immutable cloud configuration vault is secured with 256-bit encryption at rest with specific organization controls that are only accessible to customer authenticated users based on their SSO and MFA controls. This allows organizations to recover their environments even if their production cloud regions are not accessible and in certain occasions even if their cloud accounts have been compromised and not accessible.

CLOUD-NATIVE APPLICATION DATA VAULT FOR DATA RESILIENCE

Cloud Rewind takes a unique model to application data backup and replication. Cloud Rewind does not take customers' proprietary data to its cloud. Cloud Rewind also does not modify the application backup or replication data copy to its common format. Cloud Rewind leverages cloud-native mechanisms to make copies at a point-in-time and vaults them using customers cloud object storage such as S3, or other cloud-native object stores for faster backup and replication so they are immutable. This overcomes some of the common problems with current backup and replication mechanisms available along with following key benefits:

a. **No proprietary data lock-in:** As Cloud Rewind does not convert and move the data using a common backup and replication format that lock-in customers data with a proprietary backup format. Customers have complete control over their data so they can use cloud-platforms data residency and sovereignty easily.

b. **Higher performance:** Take advantage of all the cloud-native performance improvements. Cloud Rewind also takes advantage of individual cloud services data copy mechanism without dictating a common denominator model which slows down backup and replication. This is particularly important as ransomware attacks are increasingly becoming sophisticated, and organizations are forced to reduce RPO windows to be able to recover clean copies without sacrificing too much data loss faster.

c. **Wider support matrix:** Cloud Rewind offers backup and replication across various services including, multi-cloud compute services, container services, several PaaS databases, serverless objects, and key vaults and much more. As hyperscale providers add more and more data services, Cloud Rewind can readily take advantage of those services and provide data resilience at a much faster rate compared to the traditional common denominator model.

d. **Rapid recoveries:** As there is no data conversion, data recoveries are much faster in any region at any point in time. This is crucial for larger distributed applications across various data services or even a single data service. It is also very helpful when organizations try to rebuild their business applications after a ransomware attack.

RECOVER

Cloud Rewind patented system uses Recovery-as-Code to drastically reduce the risk and recovery time across the entire distributed application environment. As Cloud Rewind knows all the cloud services and their dependencies at the time of an outage, it can rapidly reconstruct the services for rebuilds at any point of time in any region of the cloud where the data copies reside. This model eliminates the need for customers to write complicated infrastructure-as-code for a particular cloud at a particular application recovery point-in-time in-sync with application data copies to guarantee application recoveries. This model also allows organizations to cut down the recovery time significantly, especially after a cyber disaster like a ransomware attack.

HOW CLOUD REWIND CONNECTS TO THE CUSTOMER CLOUD ACCOUNT

Cloud Rewind SaaS runs on GCP and accesses AWS through secure 256bit encrypted AWS APIs. Cloud Rewind has an AWS account for authentication and authorization on AWS and AWS cross-account ARN for inter-account access.

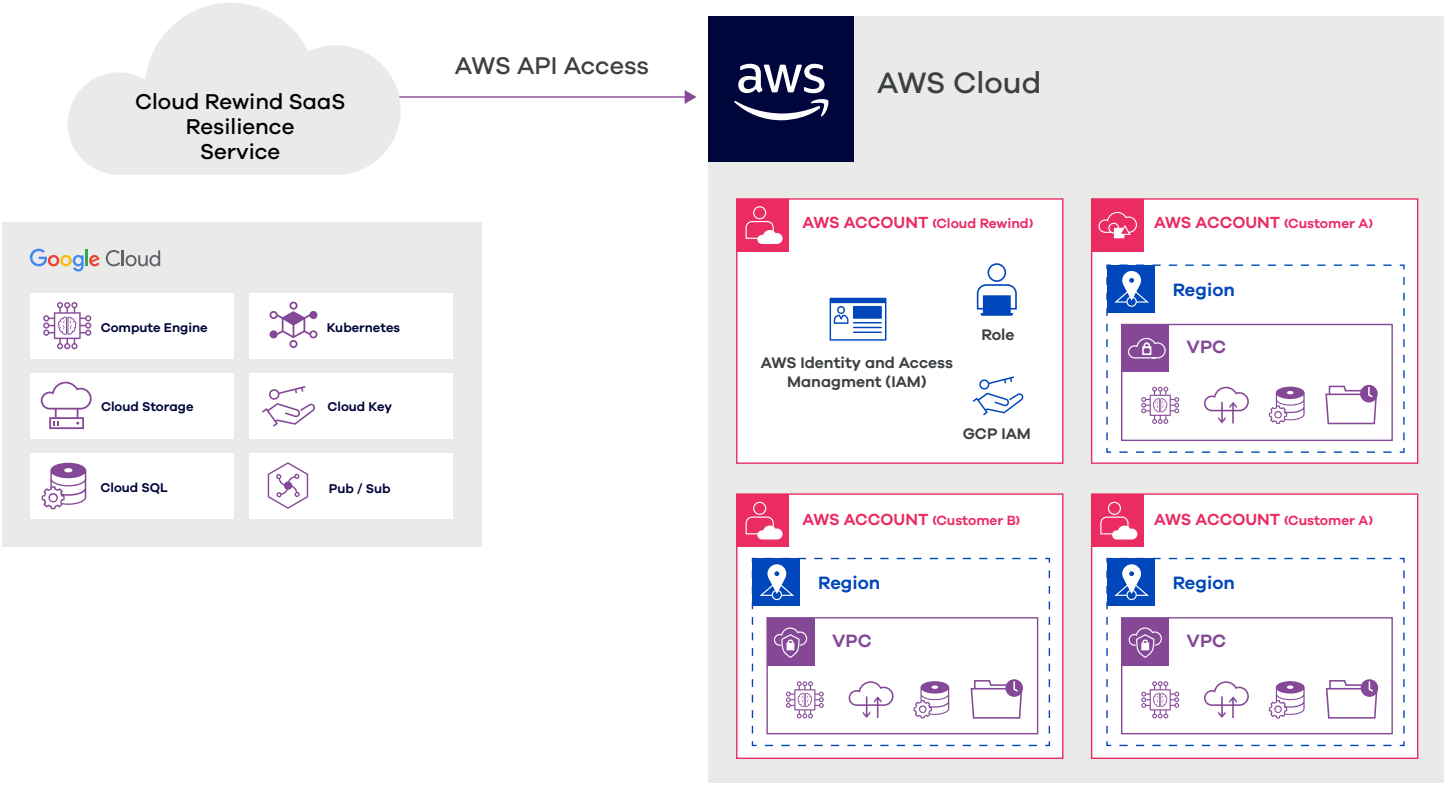

Customers create Roles and Policies for third party SaaS access in the required account and grant access to the dedicated Cloud Rewind account specific to a particular customer to assume the role. Cloud Rewind uses Assume Role with AWS STS token and performs the required operations for secure access. Cloud Rewind does not use or store AWS API keys and access keys of the customer environment, as advised by AWS. Cloud Rewind security is validated by the stringent AWS APN/Marketplace Foundational Technical Review process.

Read more about the cross-role permissions here:
https://docs.aws.amazon.com/IAM/latest/UserGuide/tutorial_cross-account-with-roles.html

PURE SAAS, AGENT-LESS, AND NO SOFTWARE INSTALLATIONS REQUIRED

Cloud Rewind does not use any agents, nor any proprietary software installations in the customer account and it is fully SaaS, making it easy to use securely. Cloud Rewind only requires AWS IAM permissions and can be onboarded using an account creation request or through the AWS marketplace. Cloud Rewind restricts itself from accessing the internals of customer environments by never installing any software which would have access to customers' data.

PERMISSIONS REQUIRED BY CLOUD REWIND

The following permissions are provided as Policies and are attached to the Cloud Rewind Role. This can be managed externally or through Cloud Rewind.

Operation	Permission Required
Discover (Mandatory)	*Resource List and Describe* permission to collect the metadata periodically. This is read-only permission, enabled for each service. *KMS Key Describe* is required to map resources to keys in cross-region for dependency mapping (Note: Cloud Rewind does not require permission to encrypt or decrypt using the Key, but only describe permission to know the key name and resources mapped for the key)
Protect (Granted only when protection is enabled)	<u>Protection based on Policy</u> *Create snapshot, Backup* permission based on the resource types. <u>Pruning after Retention Period</u> *Delete snapshots and backups* created by Cloud Rewind after the retention period. <u>Replication to cross-account and cross-region</u> Permission to *Copy snapshots and backups* to other regions (enabled regions only) and accounts (if enabled for cross-account) and delete the same after the retention period is over.
Recover (Can be granted only during the recovery process or always based on need)	<u>Region-based Permission</u> Permission to create resources during recovery. Permission to delete resources created by Cloud Rewind on reset. (Customers can use Policy attach and detach for highly sensitive environments as an internal operating procedure)

REVOKING ACCESS TO CLOUD REWIND

RECOVERY AND RESET PERMISSION REVOKE

Policies provided by Cloud Rewind, which can be attached and detached.

Policy name	Type	Used as	Description
AppranixProtectServicePermissions-20200226- -AppranixManagedDiscoveryPolicy-MPI X9	Customer managed	Permissions policy (1)	Appranix protect service discovery Policy
AppranixProtectServicePermissions-20200226- -AppranixManagedRecoveryPolicy-10 QM	Customer managed	Permissions policy (1)	Appranix protect service recovery Policy

Use the AWS IAM to add and remove Permissions using attach and detach policies.

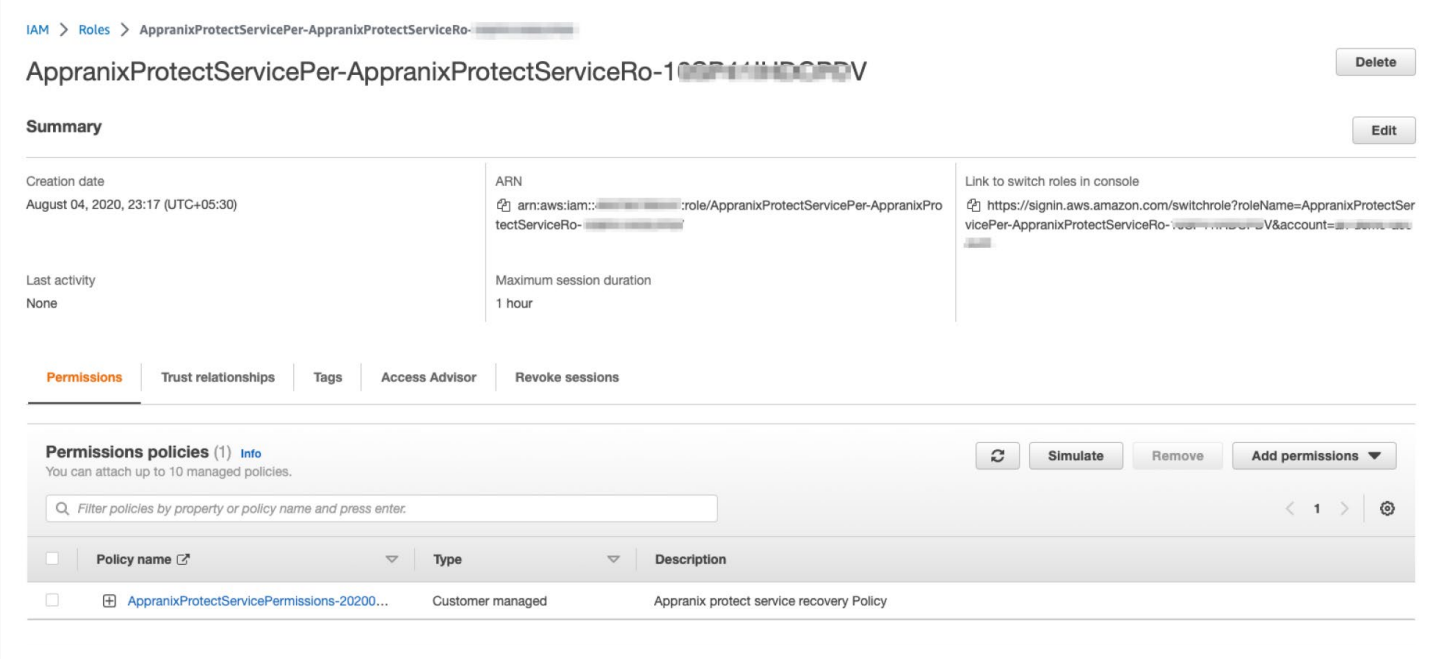

COMPLETE PERMISSION REVOKE

It is recommended to delete the Cloud Assemblies and Cloud Connections from Cloud Rewind first before revoking access provided to the Cloud Rewind role in the specific AWS account. Deleting Cloud Assemblies and Cloud Connections before deleting revoking permissions will help clean snapshots created by Cloud Rewind using an asynchronous process through the cloud API.

DATA TYPES AND STORAGE LOCATION

Data Type	What is it?	Where is it stored?
Metadata	Information of the resource provided by the cloud-provider. Example: • For a VM, the size, name, image and IP, NIC, Subnet and VPC are few of the meta-data collected. • For a EBS volume, the size, disk type, IOPS, name and attached VM are few of the metadata collected. Cloud Rewind Discovery collects the metadata of the resources based on the permission provided.	Stored in Cloud Rewind SaaS based on the region preferred by the customer. By default, Global replication for higher availability. As of June 2023 — Yet to be operational in India (Delhi and Mumbai).
Application Data	Data created by the customer's application. Example: • EBS volume data, RDS server data, etc., Usually snapshots of the EBS and RDS instances. • Cloud Rewind never reads nor requests access to store these data. It only creates a copy of these data in the customer account using AWS snapshots and backup permissions.	Always stored in customer-managed accounts, on the regions selected by the customer for replication. Data never leaves the customer accounts and stays within the cloud-provider environments.

HOW DOES THE CLOUD REWIND DUAL-VAULT CLOUD TIME MACHINE WORK?

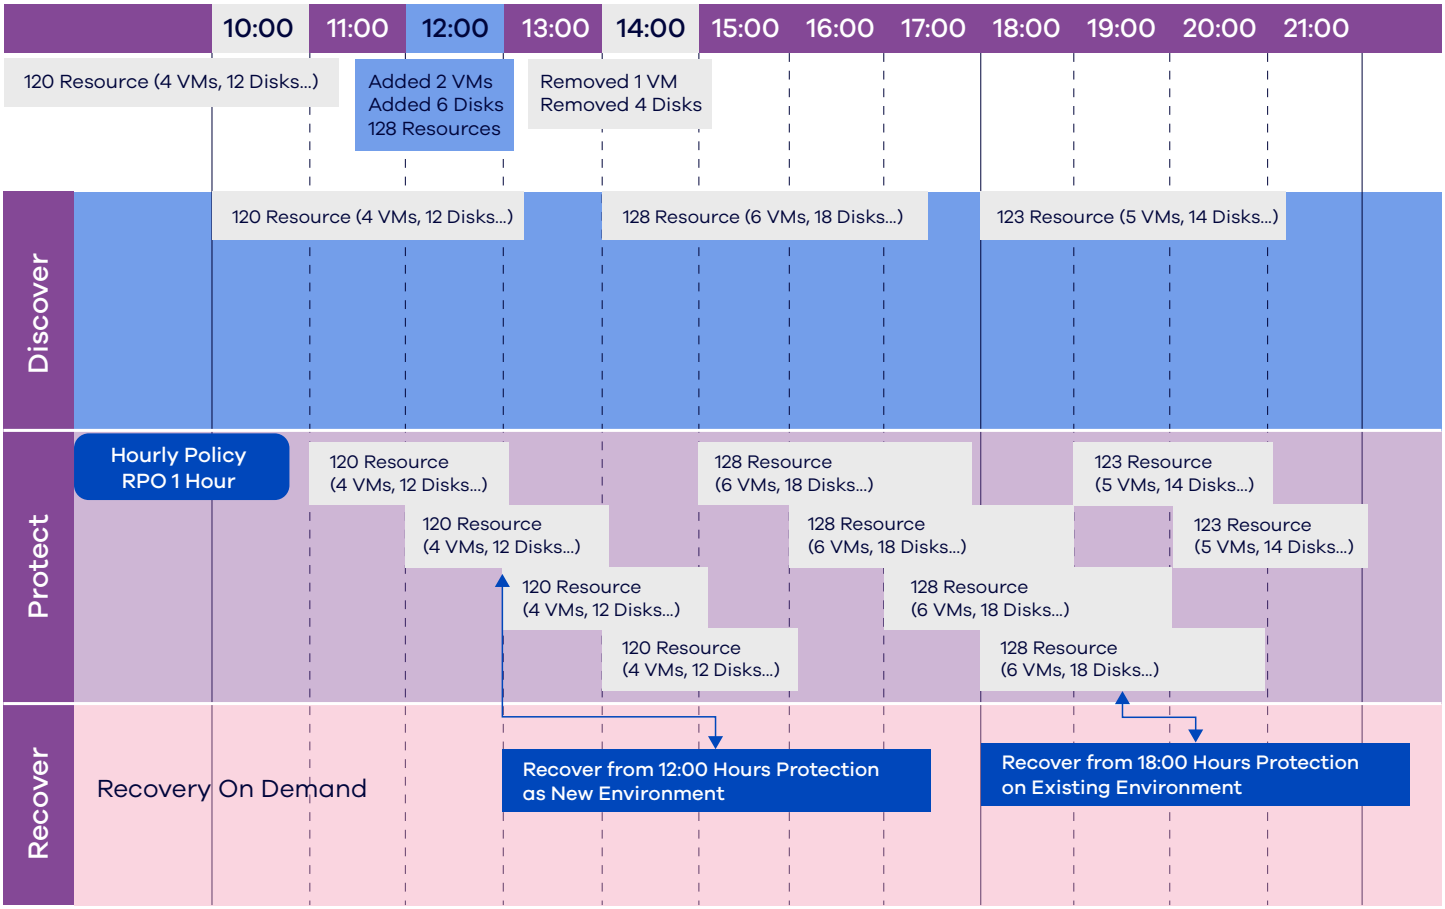

Based on the policy both configuration and application data changes are captured at periodic intervals. The above example uses a one-hour protection policy. Cloud Rewind has several options including recovery in isolated network environments if required for security reasons or in an existing customer-created network without affecting production network.

INTEGRATION WITH CLOUD REWIND

Certain use cases require internal applications of the customer environment to require configuration changes before or after recovery. These integrations are performed using Webhooks as post-recovery process. Cloud Rewind provides the recovery information to the application through the payload references. The webhooks application remains inside the customer environment and is owned by the customer, while Cloud Rewind only invokes them to provide the information required. These Webhooks can be written in any Lambda function and attached to the Cloud Assemblies as one time work.

SAAS TENANT ISOLATION AND DATA SECURITY

Cloud Rewind isolates tenant data at various levels to ensure the multi-tenant architecture of the service is available for enterprise customers. Cloud Rewind uses customer isolated KMS Keys for the data in the buckets and isolated database access for each tenant making the application access to databases with unique credentials for each customer.

All data considered are encrypted at transit making all communication through SSL and TLS. The data stored in the buckets and disks are encrypted with different levels of encryption for each tenant. All of the customers data are immutable by design.

Once a tenant account is removed, the KMS keys related to the account are removed after the grace recovery period and the data is hard deleted. Ensuring the data cannot be accessed.

CLOUD REWIND AVAILABILITY

Cloud Rewind uses GCP cloud to protect customers AWS accounts making the region entirely different from the customer operational regions. AWS regions and zones are in different geographical locations compared to GCP regions and zones making it protected from regional disasters. Cloud Rewind currently operates in GCP Iowa which is away from all the AWS regions.

AWS India works in Mumbai and Hyderabad, while GCP operates from Delhi and Mumbai, Cloud Rewind will operate from GCP Delhi for India to avoid geographical disasters in Mumbai and will use Mumbai as a secondary region to Delhi.

Note: As of June 2023, Cloud Rewind India regions are not operational yet.