

WHITE PAPER

Cloud Rewind Cloud Resilience Copilot Technical Overview

Table of Contents

CONTENTS

Introduction to Cloud Rewind Resilience Platform	3
Use Cases of Cloud Rewind	4
Primary Operations of Cloud Rewind	5
Discover	5
Protect	5
1 Cloud Configuration Vault	5
2 Cloud-native Application Data Vault for Data Resilience	6
Recover	6
How Cloud Rewind Connects to the Customer Cloud Account	7
Pure SaaS, Agent-less, and No Software Installations Required	7
Permissions Required by Cloud Rewind	8
Revoking Access to Cloud Rewind	9
Recovery and Reset Permission Revoke	9
Complete Permission Revoke	9
Data Types and Storage Location	9
How Does the Cloud Rewind Dual-vault Cloud Time Machine Work?	10
Integration with Cloud Rewind	10
Cloud Rewind Availability	11

Introduction to Cloud Rewind Resilience Platform

Cloud Rewind Cloud Resilience is a comprehensive solution designed to ensure the 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.

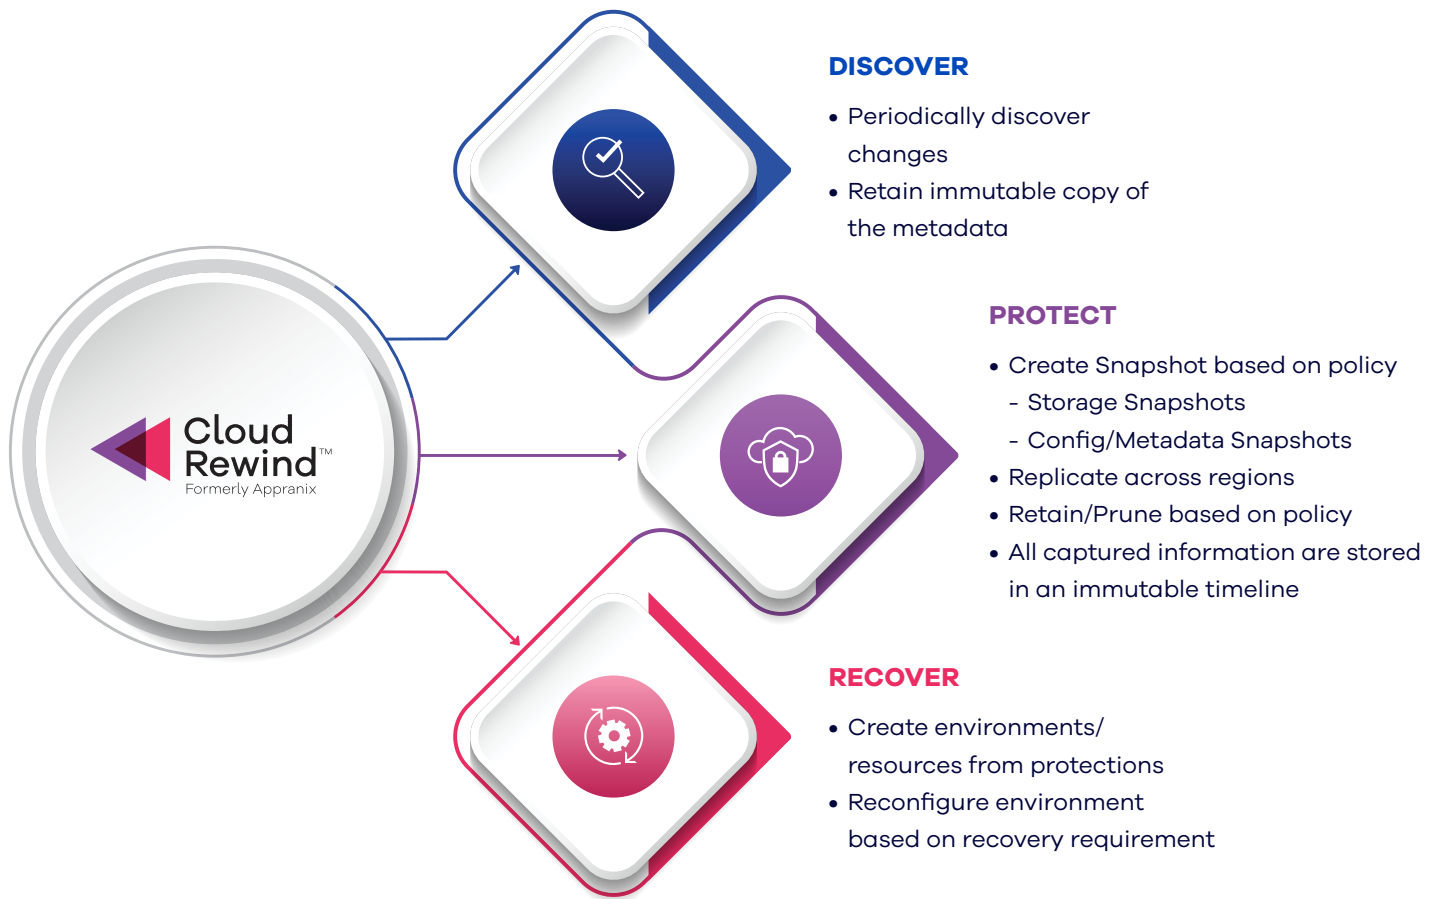
Cloud Rewind ensures resilience by continuously discovering and mapping dependencies of cloud resources of distributed systems. Cloud Rewind protects all the discovered cloud configurations, dependencies, and application data with the policies, per organization RPO, using a patented Dual-vault Cloud 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 the desired DR region.

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:** Back up 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 data in a cloud-native format for rapid point-in-time restores without sacrificing data residency and sovereignty. Cloud Rewind automates the backup of all the databases and data services platforms from an application-centric perspective to avoid the organizational risk of not backing up application data as cloud resources dynamically change.
- 3 Cloud-Native Data Replication:** Cloud Rewind continuously replicates your application data for point-in-time recoveries across other regions as 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 with one click.
- 5 Cloud-Native Disaster Recovery:** Rapidly recover partial or fully distributed applications or entire primary region resources from disasters or other disruptions.
- 6 On-Demand Cloud Spaces:** Automatically create sandboxed cloud spaces for cyber threat scanning or dev/test or fault injection testing without affecting production at any point in time in any region.
- 7 Control Cloud Costs:** Avoid multi-region cloud architecture or pilot light to control cloud costs. Remove development, maintenance, and operations costs with Cloud Rewind on-demand cloud space rebuilds.

PRIMARY OPERATIONS OF CLOUD REWIND



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’s 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.

2. CLOUD-NATIVE APPLICATION DATA VAULT FOR DATA RESILIENCE

Cloud Rewind takes a unique model for 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 snapshot mechanisms to make copies at a point in time and vaults them using customers' cloud storage such as Resource Group and Storage Account 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 the following key benefits:

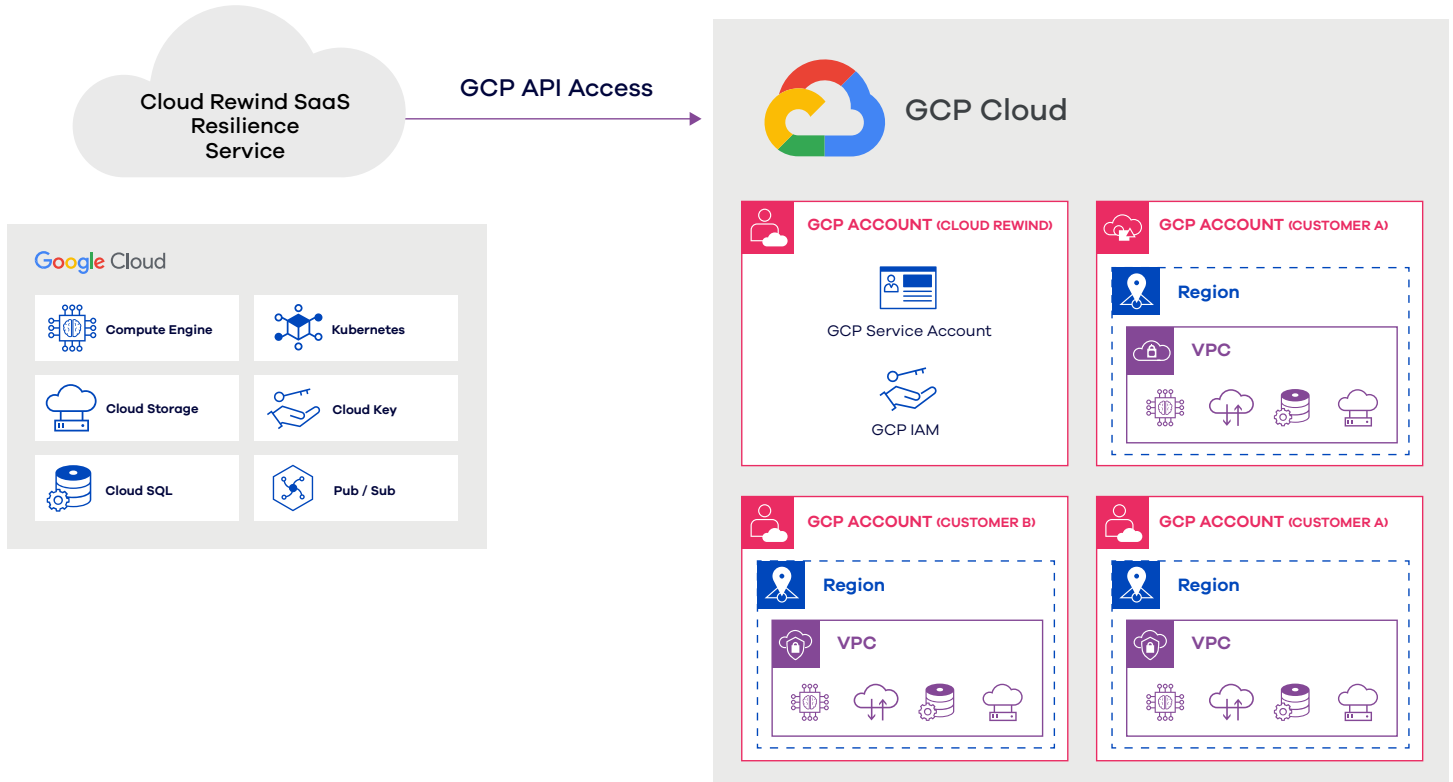
- a **No proprietary data lock-in:** As Cloud Rewind does not convert and move the data using a common backup and replication format that locks in the customer's data with a proprietary backup format. Customers have complete control over their data so they can use cloud-platform data residency and sovereignty 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.
- 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, 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's 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 in 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 your GCP through secure 256-bit encrypted GCP Cloud APIs. Cloud Rewind uses the Appianix enterprise application for authentication and authorization on GCP projects.



Customers register the Cloud Rewind app as an enterprise application within their GCP projects and assign permissions for the Cloud Rewind Enterprise Application in the registered GCP account. By leveraging the role, Cloud Rewind performs all operations within the registered GCP projects through the Cloud Rewind Enterprise Application. Cloud Rewind App registration and role assignments can be easily done through the Cloud Rewind SAAS tool.

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 a few GCP permissions assigned to the Cloud Rewind Enterprise Application and can be onboarded using an account creation request or through the GCP marketplace. Cloud Rewind restricts itself from accessing the internals of customer environments by never installing any software that would have access to customers' data.

PERMISSIONS REQUIRED BY CLOUD REWIND

The following permissions are provided to the Cloud Rewind Enterprise Application. This can be managed externally or through Cloud Rewind.

Operation	Permission Required
Discover	<i>Resource List and Describe</i> permission to collect the metadata periodically. This is read-only permission, enabled for each service.
Protect	<p><u>Protection based on Policy</u> <i>Create snapshot, Backup</i> permission based on the resource types.</p> <p><u>Pruning after Retention Period</u> <i>Delete snapshots and backups</i> created by Cloud Rewind after the retention period.</p> <p><u>Replication to cross-account and cross-region</u> Permission to <i>Copy snapshots and backups</i> to other regions (enabled regions only) and delete the same after the retention period is over.</p>
Recover	<p><u>Region-based Permission</u> Permission to create resources during recovery.</p> <p>Permission to delete resources created by Cloud Rewind on reset.</p>

REVOKING ACCESS TO CLOUD REWIND

RECOVERY AND RESET PERMISSION REVOKE

Permissions provided by the Cloud Rewind Enterprise Application can be attached to and detached from the GCP console.

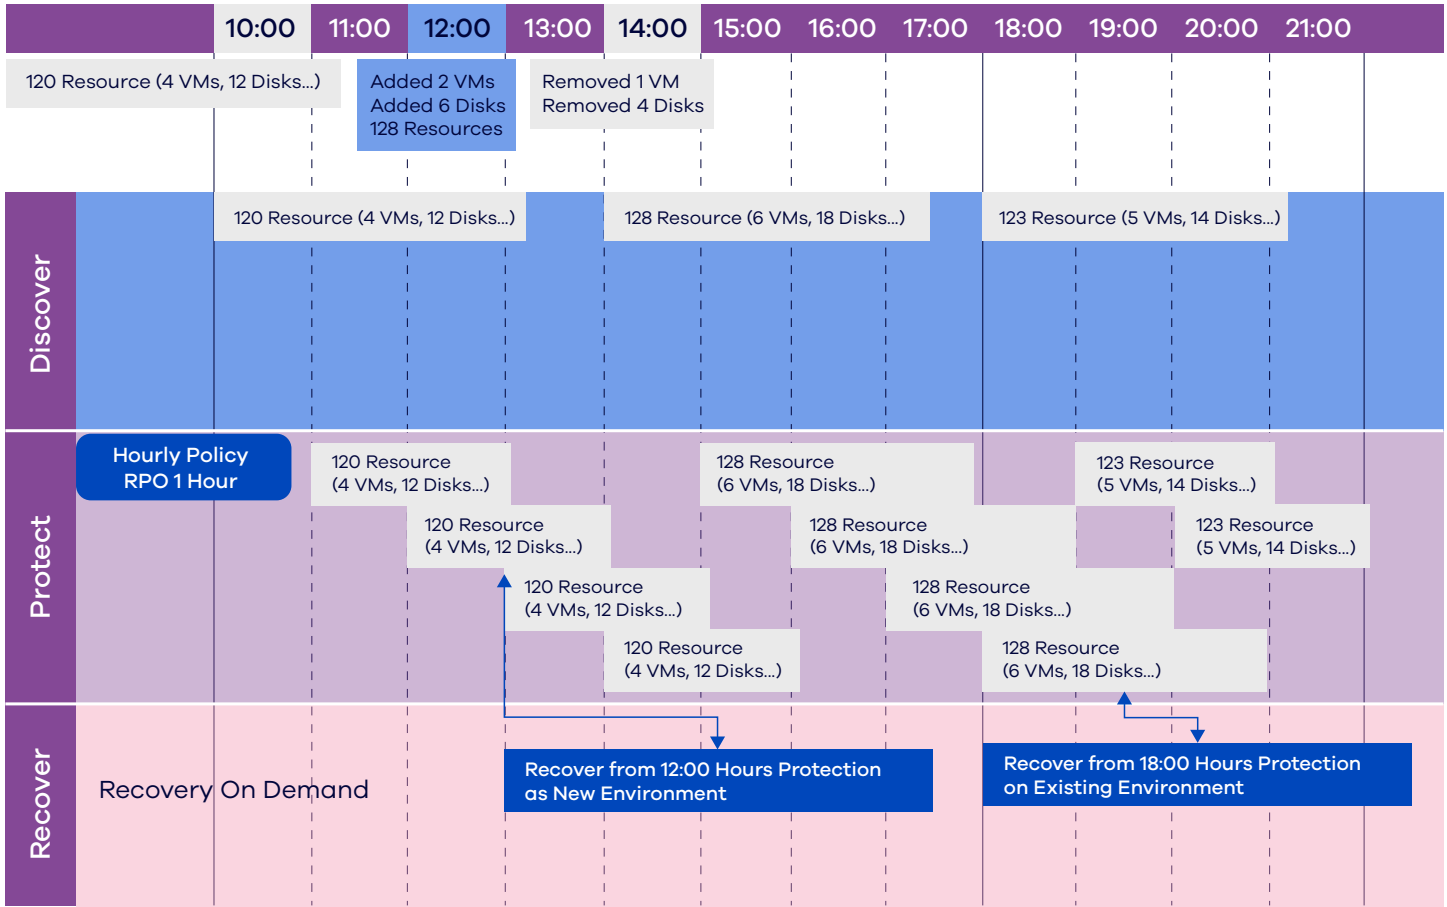
COMPLETE PERMISSION REVOKE

It is recommended to delete the Cloud Assemblies and Cloud Connections from Cloud Rewind first before revoking access provided to the Cloud Rewind Enterprise Application in the specific GCP 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	<p>Information of the resource provided by the cloud provider. Example:</p> <ul style="list-style-type: none"> For a VM, the size, name, image IP, NIC, Subnet, and VNet are a few of the meta-data collected. For Managed Disk, the size, disk type,, name, and attached VM are a few of the metadata collected. <p>Cloud Rewind Discovery collects the metadata of the resources based on the permission provided.</p>	<p>Stored in Cloud Rewind SaaS based on the region preferred by the customer.</p> <p>By default, Global replication for higher availability.</p> <p>As of June 2023 — Yet to be operational in India (Delhi and Mumbai).</p>
Application Data	<p>Data created by the customer's application. Example:</p> <ul style="list-style-type: none"> Managed Disk, etc., usually snapshots of the Managed Disk. Cloud Rewind never reads nor requests access to store these data. It only creates a copy of these data in the customer account using GCP snapshots and backup permissions. 	<p>Always stored in customer-managed accounts, on the regions selected by the customer for replication.</p> <p>Data never leaves the customer accounts and stays within the cloud-provider environments.</p>

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 the 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 a 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 GCP Cloud Function and attached to the Cloud Assemblies as one-time work.

CLOUD REWIND AVAILABILITY

Cloud Rewind uses Google Cloud Platform (GCP) to keep its clients' data safe in different zones. GCP's setup helps prevent big problems in case of disasters, making everything more secure and dependable. Right now, Cloud Rewind works from GCP's Iowa, which keeps things safe from any possible issues.

In India, even though GCP mostly works from Delhi and Mumbai, Cloud Rewind will set up its base in GCP's Delhi. This is to lower the chance of problems from disasters happening, with Mumbai as a backup to Delhi, making sure things keep running smoothly and stay strong.

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

To learn more, visit commvault.com