Commvault

Cloud Rewind
By Commvault

WHITE PAPER

Commvault Cloud Rewind Security Architecture

Table of Contents

CONTENTS

Introduction

This document is intended for enterprise cloud security and operations teams, especially those overseeing compliance, resiliency, and secure access to cloud services. It is written for Chief Information Security Officers (CISOs), security architects, cloud infrastructure leads, and risk management professionals who demand precision, control, and auditability.

The purpose of this paper is to explain how Commvault Cloud Rewind securely connects to customer AWS, Azure, and GCP environments to discover resources, protect configurations and data, and recover entire environments without impacting customer's security posture.

This white paper presents an architecture that aligns with Zero Trust principles, tenant isolation, least privilege access, and data sovereignty enforcement.

To learn more about Commvault Cloud compliance, security and Privacy, please visit https://www.commvault.com/legal/trust-center.

CLOUD INTEGRATION AND DISCOVERY

Connecting Cloud Rewind to AWS, Azure and GCP

Cloud Rewind connection to AWS, Azure and GCP is based on the principles of least privileged access. In AWS, Cloud Rewind leverages CloudFormation template to automatically create an AWS cross-account roles using STS tokens and IAM policies. In Azure, Cloud Rewind can connect using a managed Commvault App or Customers can bring their own Azure App Registration along with delegated fine grained RBAC roles. In GCP, Cloud Rewind uses service accounts with IAM roles.

Discovering Cloud Resources in AWS, Azure, and GCP

To discover cloud resources in AWS, Azure or GCP, Cloud rewind requires on read permissions. In AWS, Cloud Rewind uses fine grained read permissions only. In Azure, Cloud Rewind uses delegated RBAC roles with read-only permissions, limited to supported resources within the Cloud Rewind scope. In GCP, Cloud Rewind uses service accounts with IAM roles restricted to the viewer level.

Agentless Metadata Discovery and Permissions Model

No agents are installed. No software runs inside the customer's cloud perimeter. Customers have full control over what is discovered and when. Permissions are revocable at any time without needing to uninstall anything. Data collected during discovery is metadata only—never application data—and stored securely in Cloud Rewind's isolated configuration vault on GCP. Cloud Rewind performs read-only discovery using secure identity delegation, not credentials.

The discovery process collects configuration metadata, not application data. Metadata includes virtual machine sizes, instance names, network interface associations, volume types, encryption key references, and other cloud service descriptors.

Cloud Rewind connects to customer cloud accounts using short-lived security tokens and does not store persistent API keys or secrets. The connection is one-directional and outbound from the Cloud Rewind platform. Customers retain control over the scope of permissions granted. Revoking permissions is straightforward and can be performed by detaching policies or deleting the IAM role assigned to Cloud Rewind. These design decisions reduce the attack surface and limit Cloud Rewind from operating outside the scope of the customer-defined policy.

SECURITY AND DATA PROTECTION

Immutable Configuration Vault with Global and Regional Availability

All discovered cloud configurations are versioned and stored in an air-gapped configuration vault. Customer metadata is isolated with tenant specific encryption in transit and at rest using AES-256. For each customer's configuration data is isolated with tenant-specific encryption using Google Cloud KMS. Configuration snapshots are encrypted in transit using TLS 1.2+ and at rest using AES-256. The vault supports global replication for redundancy, but respects customer-chosen data residency requirements.

Each tenant has logically separate database schemas and KMS keys. Access to vault data requires customer-authenticated SSO and MFA, enforced through SAML 2.0 federation. No Commvault personnel have access to customer configuration data. Audit logs track every action and are retained for compliance review.

The configuration vault contains only metadata. No customer application data is ingested or processed. Each tenant has its own database schema and encryption context. This separation helps tenant isolation remain at the data, compute, and storage levels. All access to configuration data within the vault is controlled by the customer's administrative users via secure login. Authentication is managed through Single Sign-On (SSO) and Multi-Factor Authentication (MFA). Authorization policies follow least privilege access principles on role-based access controls defined by the customer.

Protecting Your Application Data

Cloud Rewind does not store application data or backups in its infrastructure. Instead, it leverages the native backup and snapshot APIs provided by cloud platforms such as AWS, Azure and GCP. When the customer enables protection policies, Cloud Rewind uses the assumed IAM role to trigger snapshots or backups of services such as EBS, RDS, or S3. The resulting data copies are created and retained within the customer's own cloud account. The data never leaves the customer's cloud environment. In AWS, this includes EBS, RDS, and DynamoDB snapshot APIs and so on. In Azure, Cloud Rewind uses Snapshot APIs for managed disks and Backup APIs for VMs, SQL, and blob data. In GCP, it uses persistent disk snapshot APIs and Cloud SQL snapshot and export mechanisms.

All snapshot data remains in the customer account. Customers choose regions for replication, retention duration, and deletion policies. Cloud Rewind executes snapshot creation, replication, and deletion using permissions that can be attached or detached dynamically. These operations are logged in customer cloud-native audit systems such as AWS CloudTrail, Azure Monitor, and GCP Cloud Audit Logs for customer auditability.

All snapshots are created using the native cloud platform's mechanisms and encryption. For example, AWS snapshots are encrypted using AWS KMS. Cloud Rewind does not decrypt or access the contents of any application data. Cloud Rewind's role in this process is limited to initiating and monitoring the execution of snapshots based on the customer's defined protection policies.

Cloud Rewind supports multi-region, multi-account replication within and across cloud providers. This enables point-in-time recovery in isolated environments. Snapshots are immutable, encrypted using cloud-native KMS systems, and never processed by Cloud Rewind. Customers maintain full ownership and control of their application data.

Customers can configure the regions where application data is replicated. Snapshots can be replicated to secondary regions or secondary accounts to meet recovery point objectives (RPO). Once the retention period expires, Cloud Rewind uses the assigned permissions to delete the expired snapshots. This lifecycle is fully under the customer's control and can be audited through their native cloud platform logs.

RECOVERY AND OPERATIONS

Recovery-as-Code Execution Model

During a recovery operation, Cloud Rewind reconstructs infrastructure using previously captured configuration metadata. The platform generates recovery blueprints from these snapshots and executes them using cloud-native APIs. The execution occurs in the customer's designated target region or account, enabling recovery even if the primary environment is compromised. This approach avoids reliance on infrastructure-as-code templates that may be out-of-date or unavailable during a disaster.

Cloud Rewind reconstructs cloud environments by replaying configuration snapshots through native APIs. This Recovery-as-Code model eliminates the need for infrastructure-as-code templates or manual intervention. In AWS, the system uses CloudFormation or direct API calls to recreate EC2 instances, VPCs, subnets, security groups, and IAM configurations. In Azure, it uses ARM templates or API equivalents to rebuild VMs, VNets, NSGs, and Managed Identities. In GCP, it orchestrates resource re-creation through Deployment Manager or direct API automation.

The recovery process includes the re-creation of compute instances, networking resources, load balancer configurations, and data storage links. Cloud Rewind helps establish infrastructure dependencies are restored in the correct order and with consistent configurations. This eliminates manual errors during incident response and reduces the mean time to recover (MTTR).

Customers may optionally configure isolated recovery networks to validate rebuilds without affecting production. This is useful for periodic disaster recovery drills or validation of recovery point integrity. Cloud Rewind supports webhook integrations so that post-recovery workflows, such as application configuration or DNS updates, can be triggered within the customer environment. These integrations do not require Cloud Rewind to execute commands within the customer's application layer. All webhook payloads are issued securely and initiated by the customer.

GOVERNANCE AND COMPLIANCE

Access Control, Authentication Architecture and Access Monitoring

Cloud Rewind enforces strong authentication practices. The SaaS portal supports user login using SSO, backed by SAML 2.0 or other federated identity providers. MFA is required for administrative access. Customers can define user roles and assign permissions to different operational tasks. These include discovery, protection, recovery, and reporting.

All user actions within the Cloud Rewind portal are logged and timestamped. These logs are retained according to SOC 2 requirements and can be exported upon request. Customers can monitor and audit all operations initiated by Cloud Rewind in their cloud environments through the logs and metrics provided by their own cloud platforms. This includes CloudTrail in AWS and Activity Logs in Azure.

Access to production systems within the Cloud Rewind SaaS platform is limited to authorized personnel within Commvault operations. Access is governed by strict role-based controls, and all activities are logged and reviewed. No Commvault personnel can access customer application data under any circumstances.

No persistent secrets are used. All API access to customer environments uses temporary tokens and roles. Cloud Rewind logs all access attempts and API calls within its platform and respects the audit boundaries of the customer's cloud provider. Customers can monitor actions initiated by Cloud Rewind through their own monitoring systems.

No access is allowed to application data or customer credentials. The principle of least privilege is enforced at all layers—UI, backend, and API invocation. Terminated users lose access immediately, and access can be revoked on-demand at both the Cloud Rewind and cloud provider levels.

Encryption and Key Management

Cloud Rewind encrypts all metadata in transit using TLS 1.2 or higher. Data at rest is encrypted using customer-isolated AES-256 keys managed by Google Cloud KMS. Each tenant has its own encryption keys and key access policies. The key lifecycle includes automatic rotation and periodic auditing. For application data, encryption is enforced using the customer's cloud-native KMS service, such as AWS KMS or Azure Key Vault. Cloud Rewind never stores or accesses these keys.

Snapshot operations are executed under the authority of the customer-provided IAM role. As a result, only the customer's cloud platform has visibility and control over how application data is encrypted or decrypted. This separation of responsibilities reduces the risk of privilege escalation or key compromise.

Cloud Rewind enforces end-to-end TLS encryption and signs all internal communications between microservices. Service boundaries are strictly controlled through Google's service mesh infrastructure. No shared services cross tenant boundaries. Once a customer account is terminated, all data is securely deleted and KMS keys are purged.

If a customer terminates their use of Cloud Rewind, all metadata stored in the configuration vault is securely deleted following a defined retention and purge policy. This includes the destruction of encryption keys and metadata entries. Customers can also delete their data on-demand before revoking IAM access.

Tenant Isolation by Design

Cloud Rewind uses a dual-vault architecture. One vault stores configuration metadata. The other handles references to cloud-native application data snapshots, which are never pulled into the SaaS environment. Each tenant has logically separated database instances, encryption keys, and access credentials.

The platform is designed for high availability with built-in redundancy and health checks across all major services. Metadata is replicated across availability zones within the hosting region. Additional copies of the encrypted metadata are stored in GCS buckets for retention with high reliability across many regions to avoid disruptions to the resiliency services provided by Cloud Rewind. Continuous monitoring helps to detect and investigate anomalies. Customers are notified of any potential degradation or incident through the service status page and support channels.

Commvault Cloud Rewind maintains comprehensive disaster recovery plans for its own infrastructure. This includes automated failover, regular recovery drills, and business continuity planning. These procedures are reviewed as part of the SOC 2 Type II audit process.

Operational Model and Platform Availability

Cloud Rewind operates from GCP's secure, globally redundant regions. For US federal workloads, this includes FedRAMP High-certified GCP zones. All control plane operations occur in isolated VPCs. The system uses Kubernetes for stateless services and Google Compute Engine for stateful components. Monitoring and alerting are handled through Google Stackdriver.

Backup and recovery orchestration for AWS and Azure workloads occurs without requiring the customer to mirror those workloads into GCP. Instead, Cloud Rewind connects to AWS and Azure using APIs over the public internet with TLS. The platform is entirely outbound and does not require any ingress firewall changes.

Availability is guaranteed through automated failover and regional redundancy. Recovery operations are decoupled from the SaaS control plane, so even if the SaaS layer is down, customers can initiate manual recoveries using the metadata previously exported.

Audit Controls, Incident Response, and Compliance

Cloud Rewind maintains a detailed audit trail of all user and system actions. These logs include login attempts, role changes, protection policy updates, recovery initiations, and snapshot lifecycle actions. Logs are timestamped and stored in immutable storage. Customers can access these logs via the platform interface or export them to SIEM systems.

Incident response procedures are defined and reviewed annually. In the event of a security incident, customers are notified through secure communication channels. Cloud Rewind conducts root cause analysis and shares findings with affected customers. The incident response plan is tested annually and reviewed by independent auditors.

Cloud Rewind's operational policies are aligned with NIST 800-53, ISO 27001, and the AICPA Trust Services Criteria. All system controls are tested as part of the SOC 2 Type II audit. Physical security, logical access, change management, and monitoring controls have been independently validated. Cloud Rewind inherits additional security from Google Cloud's FedRAMP High certification, in order to establish all infrastructure components meet stringent government-grade security standards.

Cloud Rewind supports data sovereignty maintaining all data and metadata remain in the regions chosen by the customer. It does not replicate data outside those boundaries. Deletion of snapshots and configuration data follows secure erasure policies and customer-defined retention schedules.

Customer Responsibilities and Control Points

Customers are responsible for defining protection policies, including which resources to discover, protect, and recover. Customers control all permissions granted to Cloud Rewind through their cloud provider IAM console. They can grant or revoke permissions at any time. Customers are also responsible for verifying that internal stakeholders are assigned the correct roles within the Cloud Rewind platform.

Customers must be responsible for maintaining their own IAM policies are configured securely and rotated periodically. They are responsible for enforcing MFA for their users and maintaining compliance with internal access control policies. Cloud Rewind provides guidelines for creating minimal permission policies, but customers are responsible for validating and approving these roles internally.

Cloud Rewind relies on the customer cloud platform's availability and performance. Customers need to properly configure their cloud platform accounts for service limits, region availability, and resource quotas to avoid disruptions during backup or recovery operations.

CONCLUSION

Commvault Cloud Rewind delivers an agentless, secure, and verifiable cloud resource protection solution across AWS, Azure, and GCP. It separates configuration metadata from application data. It respects cloud-native encryption, never stores secrets, and never ingests customer data. All operations occur with customer-granted, temporary permissions that can be revoked any time.

Built on Google Cloud's FedRAMP High infrastructure and audited for SOC 2 Type II compliance, Cloud Rewind enforces strict isolation, encryption, and access controls at every layer. The platform operates with operational transparency, customer-owned snapshots, and zero proprietary data formats. Recovery is fast, deterministic, and audit-ready.

Commvault Cloud Rewind provides a platform for discovering and protecting cloud infrastructure and application data using a secure, agentless approach. The system operates entirely within the customer's security perimeter and uses cloud-native APIs with restricted IAM roles. Metadata is encrypted and stored in a tenant-isolated configuration vault, while application data remains within the customer's cloud account.

No proprietary formats are used. No customer application data is read, stored, or transferred. The recovery process reconstructs cloud environments based on previously captured metadata, reducing RTO and enabling secure, regionally-isolated recoveries. Cloud Rewind meets the requirements of SOC 2 Type II and operates on Google Cloud regions certified under FedRAMP High.

Customers retain full control over access, encryption, and policy enforcement. Cloud Rewind enables a defensible, auditable, and secure method to build cyber-resilient cloud operations without increasing the attack surface or requiring invasive installations. Security operations and IT leaders can adopt the platform with confidence that it will not compromise internal security practices while strengthening recovery readiness against ransomware and outages.

This white paper documents implementation specifics based on verified operational patterns. No customer data is used for model training or analytics. All metadata is processed under strict access controls. Commvault Cloud Rewind is architected for security-first enterprises that require full visibility and recovery readiness in their multi-cloud environments.

To learn more, visit **commvault.com**