Commvault

Cloud Rewind™
Formerly Appranix

eBOOK

Cloud-Native Data Replication Management for Application Resilience

High Performer FALL 2022

AICPA SOC aicpa.org/soc4so

Commvault

Table of Contents

CLOUD-NATIVE DATA REPLICATION MANAGEMENT FOR APPLICATION RESILIENCE

As an increasing number of enterprises migrate their applications to the cloud, the requirement for efficient and effective data replication management grows. The fact that cloud-native apps are frequently distributed, dynamic, and auto-scaled makes data replication a challenging operation. However, automating the process of data replication in the cloud can help ensure that your data is always accessible, even in the face of unexpected outages or other disruption such as ransomware attacks.

In this paper, we will explore the various approaches to automating data replication management in the cloud. We will also discuss the importance of using multiple cloud regions and cross-tenant cloud accounts for better application resilience, and how point-in-time replicated data copies can help users recover from ransomware attacks and other data loss scenarios. Before we delve into the specifics of automating data replication management, it is important to understand the basics of data replication in the cloud. Data replication involves creating copies of data across multiple locations, either within a single region across zones or across different regions. This ensures that data is available even if one location becomes unavailable due to a disaster or other issue.

Organizations are no longer limited to dark-fiber based replication distances. Cloud changes the replication performance assumptions.

CLOUD DATA REPLICATION MANAGEMENT IS DIFFERENT FROM TRADITIONAL DATA CENTER DATA REPLICATION

Data centers are inherently heterogeneous, having been populated over time with various types of storage and host replication systems. To meet the needs of applications that demanded better RPOs, data replication necessitated the use of specific storage technologies. Furthermore, the distance between datacenters, as well as the speed and type of networking, would affect the performance of replication technologies. To avoid managing multiple storage systems for replications, one could augment and abstract various storage technologies at the host or hypervisor level, but performance would vary to meet RPO requirements.

Furthermore, in order to have multiple copies of data replicated to different locations, you must use the same vendor storage or hypervisor clusters. The entire setup becomes a very expensive configuration exercise, and upgrades become extremely time consuming and cumbersome. Now, managing replication across these various technologies typically entails the use of vendor-specific management tools. There are abstraction layers designed to improve replication management across storage, hypervisor host, database, file system, and virtual machine layers. The most difficult challenge for application resilience is understanding the connections between all of these layers and achieving the required RPOs.

RE-ARCHITECTED TRADITIONAL APPLICATIONS ON CLOUD PLATFORMS DEMAND A DIFFERENT TYPE OF REPLICATION MANAGEMENT

When traditional datacenter applications migrate to cloud platforms, they are re-architected to take advantage of the cloud platform storage services. For example, if an application like Sharepoint relied on SQL server clustering to achieve high availability for the entire Sharepoint farm in the on-premises data center using shared storage, the same SQL server must be re-architected using dual cloud zone architecture in the cloud due to the lack of clustering mechanisms on hyperscale platforms. Since cloud storage provisioning has changed, so has data replication and replication management.

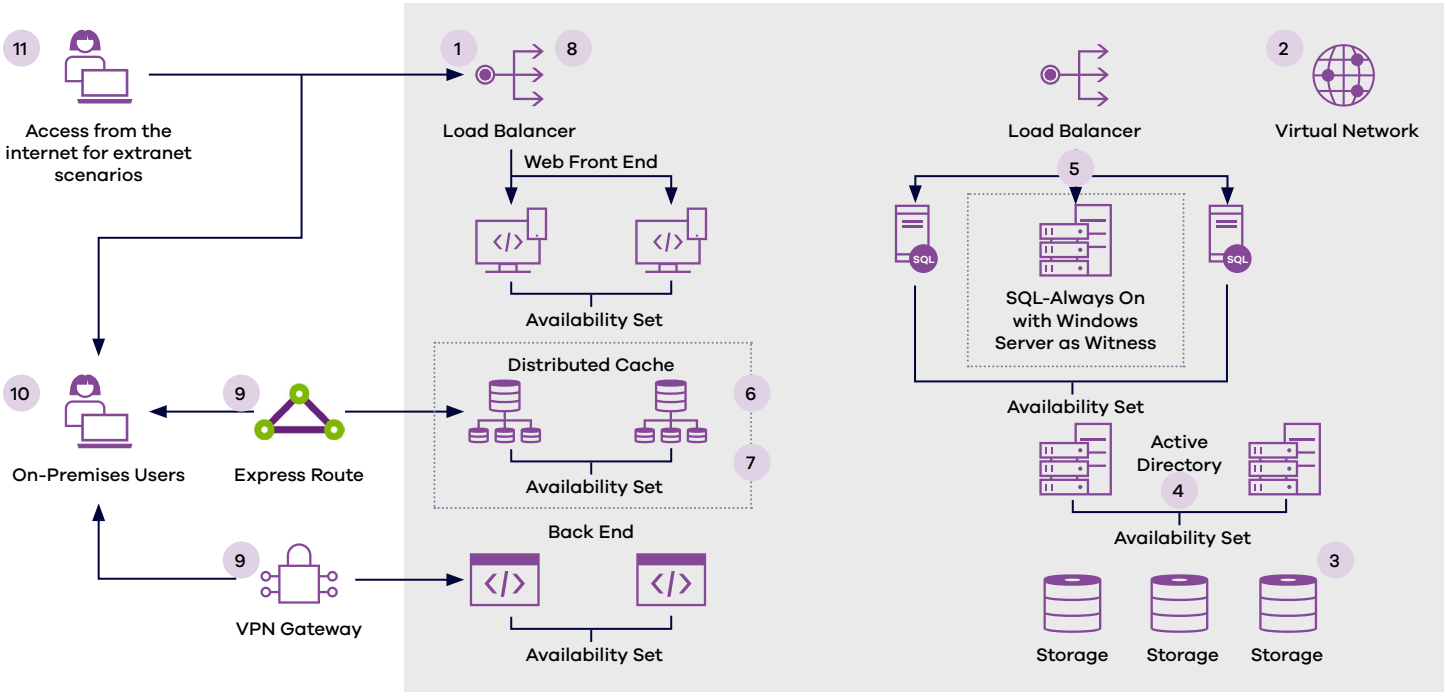

The number of resources that must be tracked to ensure replication SLAs increases the complexity of replication management within the same Sharepoint farm. Furthermore, as changes to the infrastructure become more frequent, these applications will further transform using evolving cloud infrastructure services.

More dynamism introduced by frequent provisioning or de-provisioning of underlying storage volumes adds new challenges to replication management. The cloud platform has the advantage of abstracting replication control from the underlying physical storage systems, allowing replication management to be controlled via the API. However, due to the asynchronous nature of the replication controls, new challenges emerge in tracking the volumes or databases being replicated to meet RPO requirements.

FOREVER INFINITELY INCREMENTAL CLOUD DATA REPLICATION

In the last decade, data backup and replication have introduced the possibility of forever incremental copies using snapshot-based mechanisms. Take a snapshot from the hypervisor, VM, or storage system level, and calculate the difference between the previous and current snapshots, moving only the delta. It is now possible to implement solutions that provide a unified snapshot and replication mechanism for heterogeneous host and storage systems. However, it continues to suffer from the same issues of source systems going offline due to a variety of factors such as operating system upgrades or hardware changes, unplanned downtime due to power outages, or simply human error.

When the source systems go offline and back online, incrementals may be lost, necessitating the need to capture a full again. So, while the forever incremental promise is theoretically possible, it is practically impossible to implement consistently.

Enter hyperscale cloud platforms, where the nature of allocating storage to specific virtual machine instances on an as-needed basis, combined with the unified nature of storage systems, has enabled forever incremental delta-based snapshot and replication. Changes to instance types are now possible at will, and operating system upgrades are now possible without losing deltas to continuously capture data and move it quickly to another region or even across another subscription/tenant. Or even to move to multiple regions at the same time for much better data resilience.

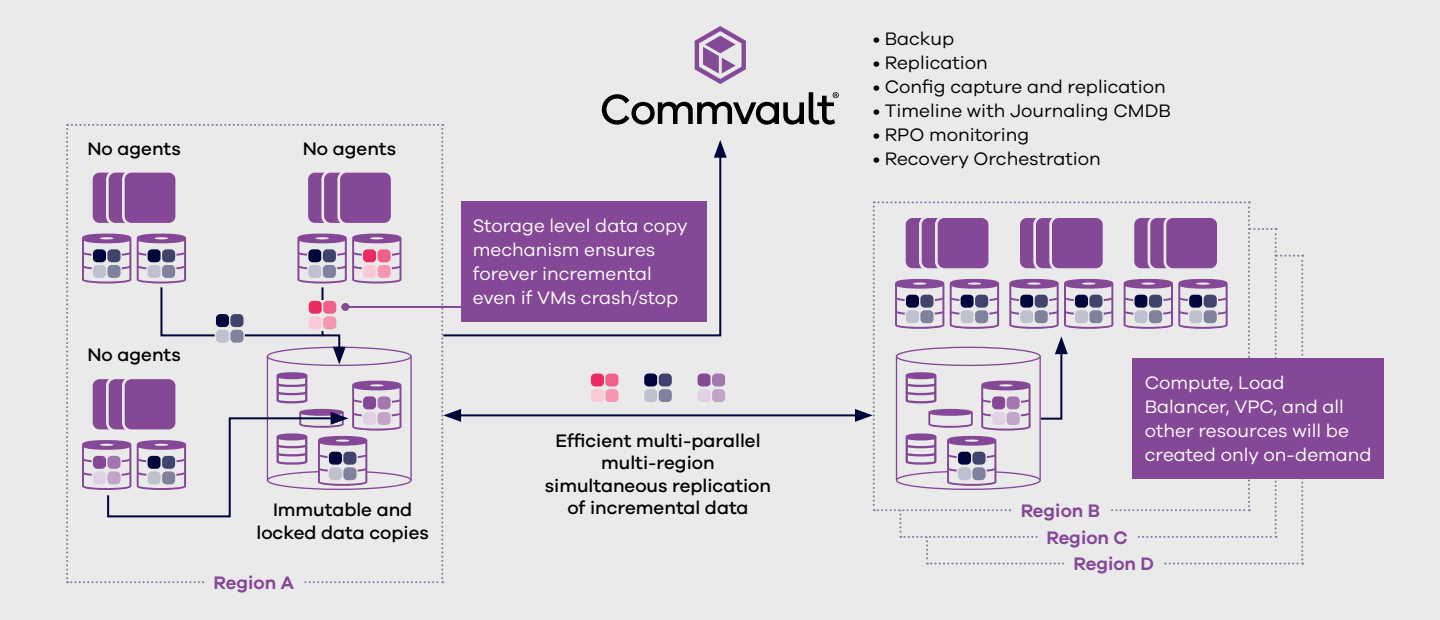

Continuous incremental replication across multiple cloud regions in parallel ensures much better resilience for applications

TAKE ADVANTAGE OF THE POWER OF THE CLOUD TO ACHIEVE BETTER DATA SECURITY

Cloud data security resilience is the ability to protect distributed and dynamic cloud application systems' data, spread across various cloud services with immutable data copies, at air-gapped locations, far away from the production environments. Importantly, a good cloud data resilience platform has the capacity to get back all your cloud services data rapidly at a particular point-in-time in a cloud region of your choice in an isolated environment so you can resume your business operations from clean data copies even after a ransomware attack.

Cloud Rewind™ can copy and move the data to various regions of the cloud really fast and keep them safe as immutable copies until someone really needs them. Cloud Rewind takes advantage of the hyperscale clouds network performance, and on-demand copying capabilities to achieve better replication SLAs for users.

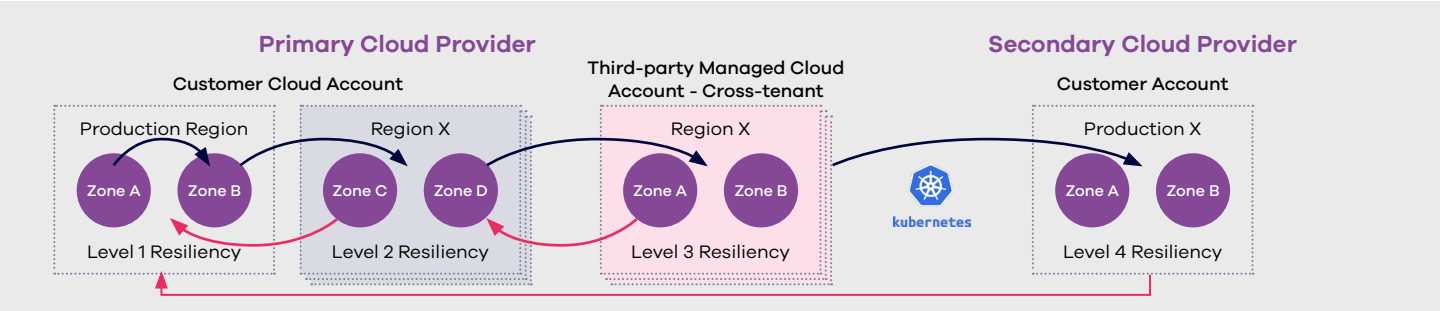

Infinitely incremental forever backup and replication for cloud cost optimization

CROSS-TENANT REPLICATION: REPLICATING DATA ACROSS TO A SERVICE PROVIDER CLOUD ACCOUNT FOR DATA SECURITY

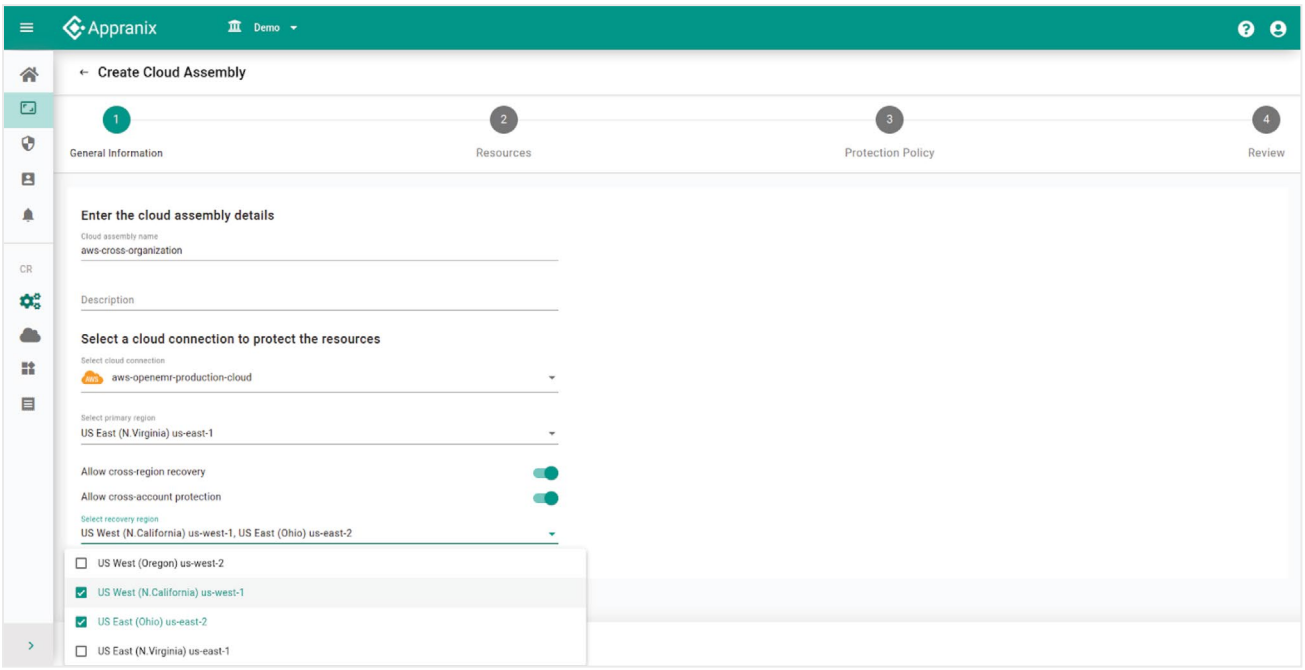

Configuring cross-tenant replication in Cloud Rewind with multiple cloud region selection as simple of selecting a few checkboxes

Hyperscale cloud platforms offer another level of protection with the capabilities to copy data across to another cloud account. This cloud account could also belong to another organization such as a service provider so that even if a customers' entire cloud account has been taken over by cyberattacks, there is still a possibility to recover critical data and business applications. Not only across another account but also across to a different region so even if the entire cloud region experiences outage along with account takeover scenario.

In addition to replicating data across multiple regions, it is also important to keep point-in-time copies of replicated data.These copies allow users to go back in time and identify clean, unaffected copies of data after a ransomware attack or other data loss event. Even more important is not only to protect the data remotely across another tenant but also to be able recover (hyper)fast all the data volumes and attach them across the distributed applications infrastructure such as VMs and database instances along with other dependencies in an environment.

REPLICATING TO AN ISOLATED SERVICE PROVIDER CLOUD ACCOUNT FOR AIR-GAPPED IMMUTABLE DATA SECURITY

Using clean point-in-time data copies, enterprises can recover from ransomware attacks by air-gapping cloud data copies in an offline location separate from the production environment. Using hyperscale globally distributed cloud infrastructure and replicating data to a completely separate cloud account from the production account increases data resiliency and offline security without sacrificing faster recoverability.

The traditional two-step copying mechanism is not required for cloud-enabled distributed dynamic workloads with a variety of backend cloud data services. The data is copied from the data service to another data management platform, and then from the data management platform to another location in the same cloud for offline access, which not only slows backup but also has a significant impact on recoverability.

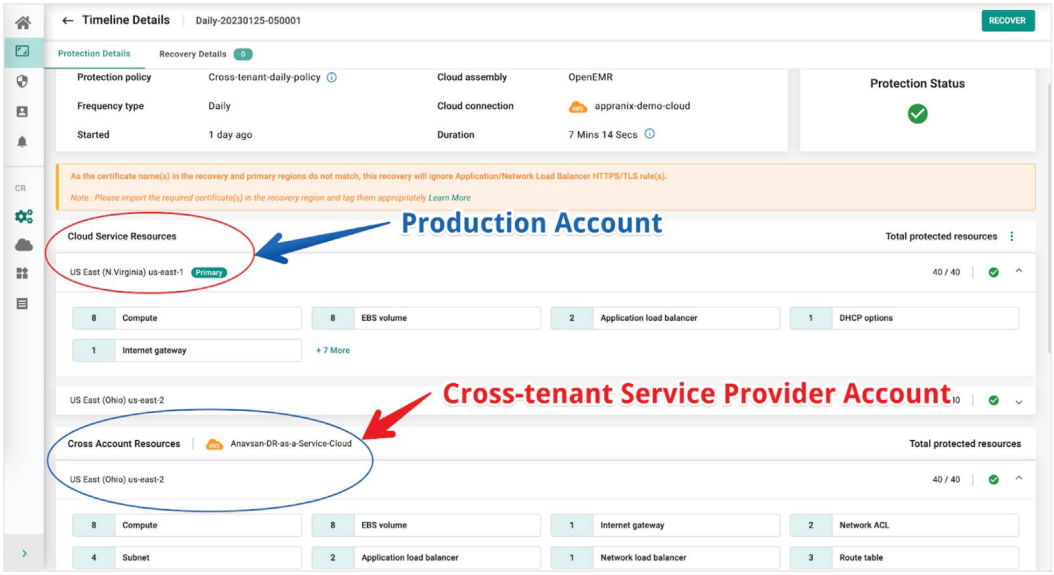

Cloud Rewind takes care of multi-parallel multi-region including cross-tenant replication automatically per customer SLAs

Cloud platform users should take advantage of fast and improving global networks and copying capabilities for faster recoveries while leveraging air-gapped data protection at a significantly lower cost. Cloud Rewind makes it much easier for organizations to protect their data by transferring it to another cloud account via cross-tenant sharing and replication, which is fully automated with policies.

INCREASED DATA SECURITY WITH DIFFERENT ENCRYPTIONS ACROSS REGIONS AND TENANTS

One of the most important ways to secure data copies for additional protection is to replicate data with two different encryptions. Customers and service providers can use their separate Key Management Systems in their cloud accounts to keep two different keys for storage volume level encryptions.

Customers can freely change the keys in accordance with their key rotation policies, and service providers can rotate their keys to protect specific customer data in their own KMS within their own cloud accounts dedicated to the specific customer.

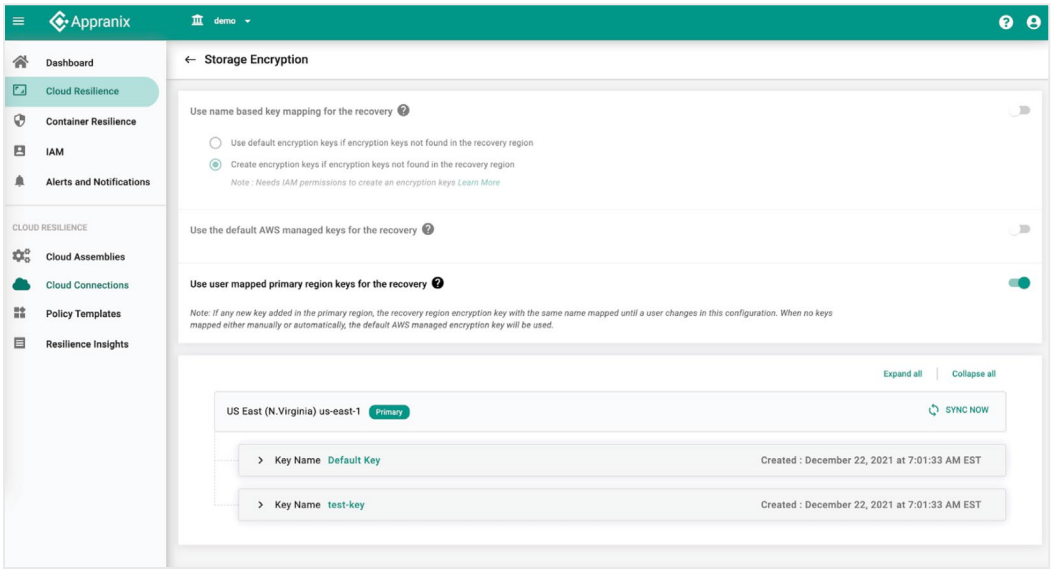

Setting up different encryptions is straightforward and completely automated to ensure data security

MANAGING ACTIVE-ACTIVE DATA REPLICATION WITH BACKUP FOR POINT-IN-TIME RECOVERY

Active-active replication, which involves replicating data to multiple zones and regions while keeping all regions online and accessible to users, is one method for improving application resilience. This method necessitates a more complex setup and configuration, but it provides the highest level of resilience and availability. This type of replication configuration is typically used at the database level to ensure that there is no data loss and that applications have the highest level of availability.

However, backup copies must be made in the active-active replication scenario so that applications can be recovered into the same configuration from a different point in time in the event of a ransomware attack.

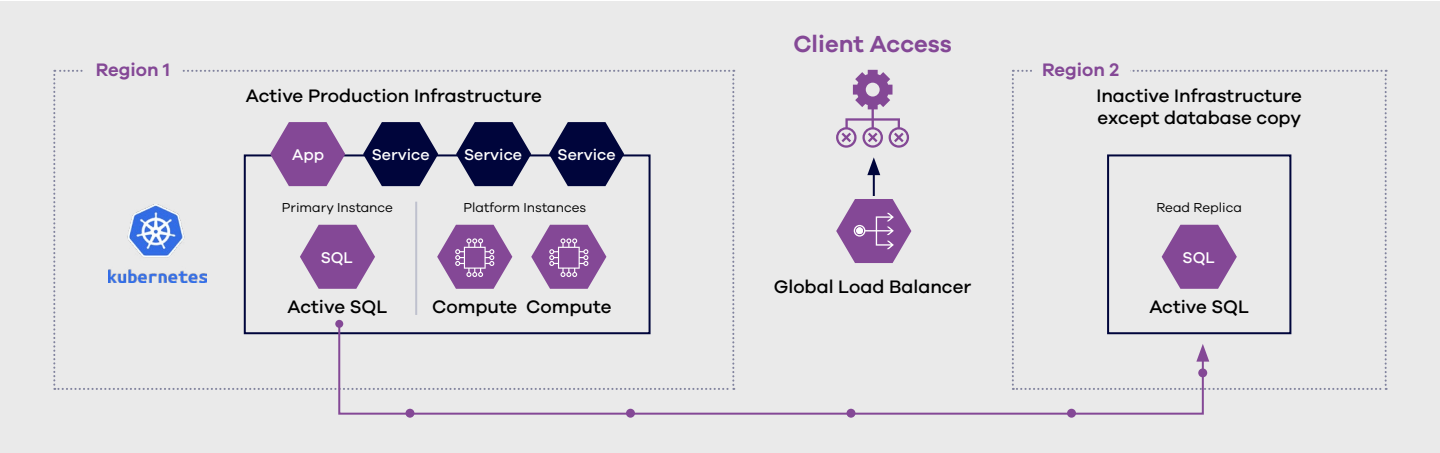

Cloud Rewind manages active-active PaaS databases with backup seamlessly to allow no data loss scenario along with recoveries from ransomware attacks

MANAGING LOCAL REGION + REMOTE REGION + CROSS-TENANT LOCAL REGION + CROSS-TENANT REMOTE-REGION REPLICATION COPIES IN ONE SIMPLE SETUP

The cloud has simplified many IT operations. It also allows organizations to investigate options that were previously unavailable. Take, for example, data replication. To have reliable replication within the same data center, across to another data center, and to a DR-as-a-Service provider data center, you need several products and, depending on the complexity and scale, several teams. Not to mention the management tools' complexity.

Hyperscale clouds have enabled a level of resilience not previously possible, but in order to fully leverage the capabilities, you must have a solid understanding of the cloud platforms as well as a clear understanding of your application requirements.

Cloud Rewind has simplified all replication operations by orchestrating data natively to regions or accounts using cloud platform capabilities, with a very simple application-centric setup. As applications evolve and their data gravity shifts, Cloud Rewind gives you the confidence to not only capture the data incrementally in the same region, but also across to another region and to another cloud account belonging to a service provider, with unprecedented ease of use. You can now easily handle all combinations of replication setups, including backups for point-in-time recovery, and recover your applications in-sync with data copies in whichever cloud account and regions at whatever point-in-time.

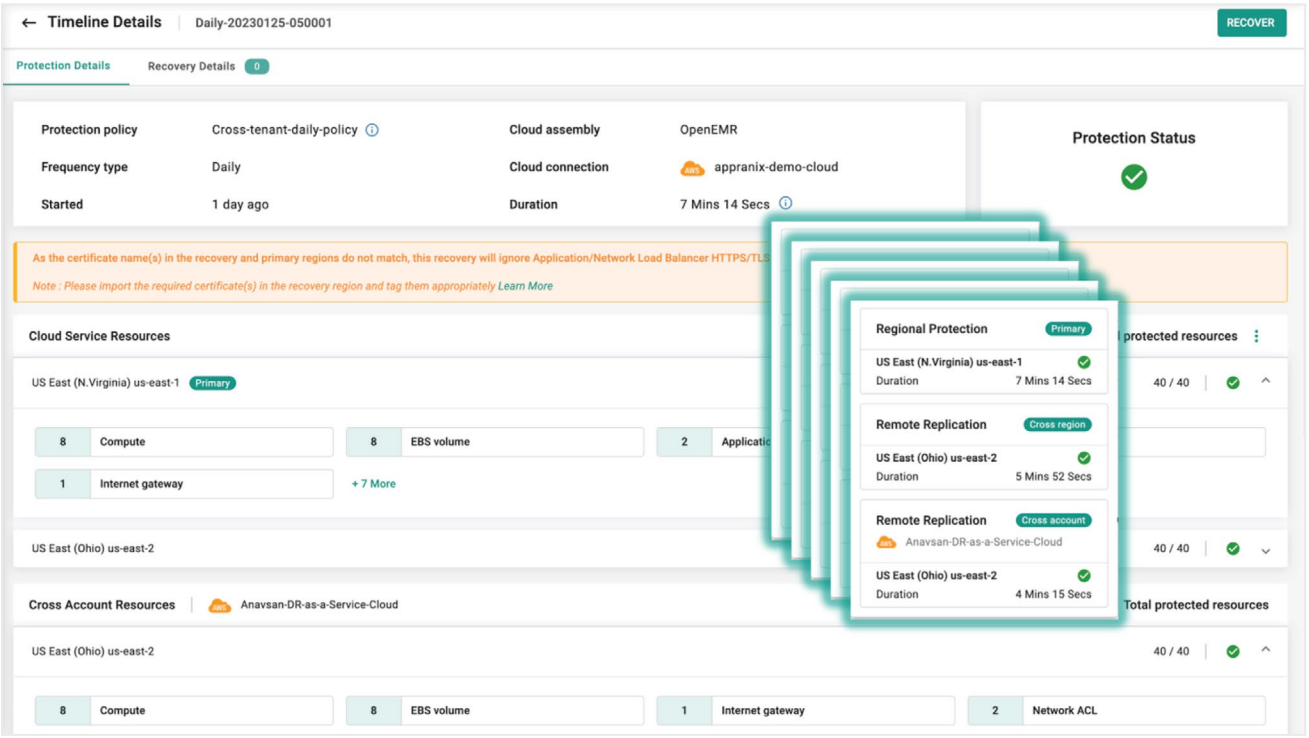

Tracking replication SLAs along with backup for all the policies is very simple over a period of time.

SUMMARY

Organizations migrate to hyperscale clouds not only for agility, but also for significantly improved application resilience. Taking advantage of everything that cloud platforms have to offer, as well as keeping up with the rapid changes in data services, is an important consideration for cloud-enabled organizations. The most important aspect is how dynamic, auto-scaled, distributed applications with evolving services or microservices architectures with multiple database sources or storage options can be resilient in the event of an outage and recover quickly. Cloud Rewind provides a simple yet effective SaaS solution that not only eliminates risk but also offers a guaranteed approach to application recoveries with significant data resilience built on hyperscale cloud platforms.

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