



eBOOK

Achieve Cloud Application Resilience With Hyperfast Recoveries


Table of Contents

Hyperfast Recovery of Cloud Application Environments	3
Why a Traditional Datacenter BCDR Model Is Not Suitable for Cloud Applications	3
Proliferating Ransomware Changes Application Recovery	4
Continuous Discovery of Cloud Resources Is Key to Better Resilience	4
Backup Your Cloud Environment State With a Continuously Learning System	5
Hyperfast Entire Environment Recoveries With DR-as-Code	5
Cloud-Native Data Copy Management	5
Cloud Application Environment Time Machine	6
Summary	6
About Cloud Rewind™	6

HYPERFAST RECOVERY OF CLOUD APPLICATION ENVIRONMENTS

Cloud-enabled organizations have rapidly shifted to a decentralized operating model for their applications and services. Software architectures have also become more distributed, making use of readily available cloud resources across cloud zones. Site Reliability Engineers have adopted more dynamic and faster release cycles through the use of DevOps practices in order to meet increased customer demands. Furthermore, programmable cloud resources have enabled environments to scale automatically in order to meet the performance requirements of critical business applications. On the negative side, though, all of these changes have created massive challenges for the shared operations service teams that manage resilience, security, and cost. The most pressing question now, especially when cloud environments are prone to increased cyberattacks, is how these dynamic, auto-scaled application environments can recover quickly from downtimes using cloud-native infrastructure so promised business SLAs can be maintained.

WHY A TRADITIONAL DATACENTER BCDR MODEL IS NOT SUITABLE FOR CLOUD APPLICATIONS



Applications no longer rely on a few servers or a single set of critical databases. Consider the following example of a simple three-tier cloud application with an auto-scaling consisting of two virtual machines and a database. It is made up of at least twenty (20) distinct cloud resource types and instances. Traditional backup and recovery systems were designed to just protect virtual machines disks, databases and file systems. In order to deliver resilience for the entire cloud application, all the cloud resources need to be protected so you can recover at any point in time in any region of the cloud. Legacy backup and recovery systems were not built to protect all these cloud resources used by dynamic, distributed auto-scaling apps that rely on software-defined cloud infrastructure.

PROLIFERATING RANSOMWARE CHANGES APPLICATION RECOVERY

As ransomware attacks proliferate and become increasingly sophisticated, cloud environment recoveries become harder and harder. What's more important is that new ransomware attacks target the backup products and their management consoles. As most BCDR products are installed in the primary domain's cloud account as that of their production systems, if a ransomware attack takes over the entire cloud account, it is not even possible to reach the consoles of the backup and recovery systems to be able to recover the application environments. This is perhaps one of the critical models that organizations need to rethink when re-architecting for application resilience as opposed to just the data backup and recovery.

As cloud application systems are made up of multiple cloud infrastructure services, users performing recoveries from ransomware attacks need a tremendous understanding to be able to piece together virtual machines, databases, networks, multitudes of cloud services and associated cloud configurations to recover them properly. Typically, key components of the application environments such as virtual private networks (VPCs), load balancers, gateways, security groups, database parameter groups, etc, need to be manually assembled ahead of time by the cloud operations teams before even engaging with BCDR systems for data recovery.

CONTINUOUS DISCOVERY OF CLOUD RESOURCES IS KEY TO BETTER RESILIENCE

Dynamic and auto-scaled cloud environments introduce enormous challenges for the operations teams to keep them secure and resilient. As multiple development teams self-service the majority of the cloud infrastructure resources, cloud application environments expand at a faster pace than a traditional data center model. These programmable ever-changing environments need a system to continuously discover all the resources that belong to an application. Organizations also have many cloud accounts to isolate their development, production, and test environments depending on their business needs. The number of cloud accounts in the 100s is not uncommon now-a-days.

The complexity of many cloud accounts along with fast-changing environments makes it hard for centralized teams to rely on traditional, non-application-centric protection and recovery systems as these systems simply rely on users to pick the right resources and manually apply protection for their applications. On the other hand, application developers do not keep track of all the cloud infrastructure resources used for their applications so they are unable to help SREs at the critical time of recovery. Typically several DevOps pipelines modify the central cloud environments making it harder for SREs to recover at the time of dire need. You need a system that continuously discovers cloud resources and is application-centric, and should have the capability to understand the system resources using automated dependency mapping to properly protect all the relevant cloud resources. It is then possible to rapidly recover or failover the applications, data, configurations, state, and dependencies to keep up with application uptime requirements.

BACKUP YOUR CLOUD ENVIRONMENT STATE WITH A CONTINUOUSLY LEARNING SYSTEM

Gartner estimates that a typical cloud environment goes through 50+ configuration changes per day. It is important to build an immutable cloud configuration meta-data repository of all critical cloud application environments. It is also very crucial to host cloud configuration meta-data away in a cloud application resilience system hosted on a different cloud to achieve additional levels of resilience. These configuration meta-data vaults need to be segmentable by application services, and should be journaled for any point-in-time recovery in any region of the cloud. They should be granular enough to allow operations teams to ask for a single resource at any point-in-time so a particular cloud service instance could be recovered quickly upon a failure.

A continuously learning system is key to keeping track of the changes so in the event of failure, distributed cloud application systems can be recreated based on what was in the production environment. A continuous discovery and meta-data learning system completely eliminates the need for manual assessment and risks associated with disconnected meta-data during the recovery process.

HYPERFAST ENTIRE ENVIRONMENT RECOVERIES WITH DR-AS-CODE

The most complex part of recovery is identifying the right compute, storage, PaaS, networking infrastructure resources corresponding to a set of applications and sequencing them for an orchestrated recovery. This is called a “Technical DR Plan”, TDP for short. There is also a non-technical aspect of the DR plan which concerns bringing human and other organizational resources for application validation after the recoveries.

TDPs are typically multiple pages and need several operational people working together to identify what runs in production, in terms of configurations, dependencies, sequencing, and scripting. Organizations that have used traditional BCDR products will tell you how complex TDPs are and why they don’t run recovery tests often.



It is now possible to eliminate manual TDPs completely with an automated infrastructure-as-code (IaC) model. In particular, for guaranteed recoveries, it is important to use cloud-native IaC, instead of a cloud-neutral IaC so the responsibility of running large system recoveries shifts over to the cloud provider with dynamically scalable resources to be able to complete the recoveries successfully during a downtime.

CLOUD-NATIVE DATA COPY MANAGEMENT

Cloud platforms have enough data management capabilities to be able to make much faster data copies for backup, replication, and recovery. There is no need to add any additional data management capabilities from third-party vendors. There is no need to change the data storage format from native application storage data to a common data backup format and go through the lengthy process of import and export into the neutral backup file system.

It is possible to make incremental data consistent copies from virtual machines and databases to reduce the cost of backup, and DR. Serverless services have enough built-in data management capabilities to avoid costly copying to and from data management platforms bolted onto a cloud environment.

Hyperscale cloud platforms really opened the doors for much better resilience compared to the datacenter infrastructure model. Global organizations can literally replicate the incremental data from one region of the cloud to another within minutes. This not only increases data resilience but also cheaper multiple copies across the globe enable for much better levels of application resilience in the event of a failure.

CLOUD APPLICATION ENVIRONMENT TIME MACHINE

Cloud Environment Time Machine is a simple concept in which an automated system can assemble all the app-centric cloud resources’ meta-data from a vault and application from an immutable repository and application data from storage and databases for a synchronized point-in-time recovery. You could imagine these time machines as journaled CMDBs that are automatically refreshed from an application-centric perspective using all the cloud-native capabilities.

However, the most important difference between a Cloud Time Machine compared to the older CMDBs is that it knows about the point-in-time data copies for the applications. Over time a cloud time machine becomes invaluable for organizations as multiple groups within an organization can readily tap into it for various rollbacks, recoveries and failovers. Traditional BCDR systems never gathered meta-data of systems to be useful beyond simple data backup requirements.

SUMMARY

The dynamic nature, complexity and speed of changes to the cloud applications really need a new application-centric resilience model as opposed to the legacy protection and recovery or disaster recovery models predominately created during the datacenter era. Whether the applications were migrated to the cloud or natively created on the cloud-platforms, this new model not only helps recovery of the entire application environments from multitudes of downtimes rapidly but also drastically reduces operations nightmare, especially when fewer operations teams manage a magnitude more resources compared to the last decade.

ABOUT CLOUD REWIND™

Cloud Rewind delivers cloud application resilience with entire cloud environment backup and recovery of all resources, services and dependencies at any point-in-time in any cloud region.

To learn more, visit commvault.com