

Bringing Enterprise-Grade Capabilities to the Public Cloud for Secondary Data

Public cloud expansion will continue to see double-digit growth over the next five years, according to IDC¹. Making the public cloud interoperable with on-premises solutions needs to be part of your hybrid cloud strategy for secondary data workloads. Legacy solutions, however weren't architected with the cloud in mind. They require the use of cloud gateways and don't scale well in the cloud.


Introducing a Better Way with Cohesity

Cohesity DataPlatform is a modern architecture for secondary data and apps designed to be cloud-first. Native integration with public clouds helps you take advantage of their scalability and cost effectiveness for multiple use cases, all while staying in control of your secondary data.

Cohesity integrates with Amazon Web Services, Microsoft Azure, Google Cloud Platform and Cohesity Powered service provider clouds to extend your on-premises data protection solution to these public cloud services. With Cohesity Platform CloudEdition, you can deploy a Cohesity cluster in the public cloud and enable replication from the data center to the public cloud.

"We saw clear value in adopting Cohesity to consolidate and simplify our secondary storage infrastructure. Cohesity allows us to consolidate data protection and file storage, while providing simple integration with the public cloud for long-term archival."


BRIAN SWEENEY,
Principal Engineer,
Manhattan Associates



¹IDC Worldwide Semiannual Public Cloud Services Spending Guide


Use Cases

Our built-in cloud capabilities give you flexibility to use the cloud for multiple use cases:




Long-Term Retention

Archive data to the cloud – Enable policy-based archival of older snapshots on-premises to the cloud for long-term retention. Reduce reliance on tape and lower TCO. Provide an easy way to retrieve data back to on-premises from the cloud and recover data to a different site. CloudArchive™ is the capability for archiving older local snapshots in the Cohesity cluster to the public cloud for long-term retention. CloudArchive dedupe optimizes data stored in the cloud.




Storage Tiering

Move cold data to the cloud – Utilize policy-based thresholds to move cold data to the public cloud. Leverage cloud storage as another tier. Tier data back to an on-premises cluster. CloudTier™ uses cloud as an extension to Cohesity’s built-in storage to tier the data between a Cohesity cluster and the cloud.




Test/Dev

Provision new applications or move on-premises applications to the cloud – Use your backup data as more than an insurance policy. Spin up new VMs with or without Cohesity in the cloud. CloudSpin is the Cohesity cloud capability that enables you to spin up on-premises VMs in the public cloud for test/dev. It allows policies to be set up to schedule backups to be converted and stored in the cloud without the need for running Cohesity in the cloud. Upon demand, VM in the cloud can be activated.



Disaster Recovery


Use the cloud for cost-effective disaster recovery (DR) – Replicate to the public cloud from on-premises for cost-effective disaster recovery. Spin up cloud instances of VMs with automatic format conversion. Enable disaster recovery. CloudReplicate™ allows you to replicate data from an on-premises Cohesity cluster to a Cohesity instance in the cloud. With CloudRetrieve, you can archive data from one cluster to an external target. Register another cluster with same external target and retrieve data from there to an alternate cluster.



NEW **Cloud Native Backup**

Automate backup of cloud applications – Protect cloud applications with API integration. Achieve best in class operational efficiency by leveraging integration with public cloud snapshot APIs. The Cohesity platform features native API integration so that you can integrate with public cloud snapshot APIs for native backup of cloud applications.

On-Premises Datacenter



Direct Backup to Cloud

Backup directly to the cloud – Backup on-premises applications and send backup data directly to the cloud without requiring backup software and target storage on-premises.

Key Platform Differentiators for the Cloud

Cohesity consolidates all of your secondary workloads, including backup and recovery, files and objects, test/dev and analytics in a single, cloud-native solution.

- **Web-scale:** Software-defined and web-scale platform provides unlimited scale, non-disruptive upgrades, and a pay-as-you-grow consumption model.
- **Global deduplication:** Solution maximizes space efficiency with global deduplication in the cloud. Deduplication can be configured to be variable, inline, or post-process.
- **Unlimited snapshots and clones:** Cohesity SnapTree takes an unlimited number of snapshots and clones without impacting performance. Each snapshot is instantly available as a fully hydrated copy of the file.
- **Automated global indexing:** Automated global indexing powers Google-like search, enabling instant wildcard searches for any VM, file, or object that is ingested. Granular recovery optimizes data egress from cloud.
- **Security:** Cohesity enables encryption of data at rest and in-flight. Provide role-based access control that enables mapping of Active Directory or local users to roles with segregation of duties.
- **Quality of Service:** The solution allows you to control Quality of Service at a granular level to ensure low-latency, high throughput to mission critical workloads.

Cohesity DataPlatform provides simple connectivity to public cloud services as an extension of the data center infrastructure for long-term retention, tiering, disaster recovery, test/dev and replication. By utilizing Cohesity you can make your hybrid cloud strategy for secondary data a reality.

Learn more at www.cohesity.com.