

# Cohesity Integrates the Cloud Into On-Premises Secondary Storage

## Situation

Use cases such as mobility, big data initiatives, social media and the Internet of Things (IoT) are driving exponential data growth, resulting in a need to double storage capacity every two years. IT organizations (and vendors) place an emphasis on the evaluation and acquisition of primary storage, which is the repository for mission-critical and other production-oriented applications. However, dark or secondary data represents 80% of enterprise data and imposes an enormous burden on enterprise IT budgets. The lower level of visibility leaves secondary storage prone to insufficient management and low utilization.

Secondary data originated as primary data, but it is no longer used for production or operational purposes. This fact does not lessen the business value of secondary data, as it is used to fulfill a number of essential IT processes.

Secondary data consists of multiple copies of data that companies use for analytics, business continuance, and Test/Dev. Businesses can have ten copies or more of data to satisfy these use cases. Another contributor of secondary data is shared files that distribute or provide access to digital media, such as computer programs, multimedia (audio, images, and video), and documents within a business or organization. These file shares are not copies but also consume a significant amount of storage.

The service level required by secondary data is less than that of primary data and is often satisfied with lower-cost secondary storage. Secondary storage consists of point solutions for the different use cases or older, cascaded storage, resulting in a fragmented and inefficient collection of resources. Secondary storage is a diverse patchwork of equipment with various services, performance levels, management tools, connectivity, and features. This diversity results in storage that is difficult to manage, inefficient to provision, and complex to protect.

## Event

On November 23, 2016 Cohesity, a provider of hyperconverged secondary storage, announced the DataPlatform Cloud Edition. The DataPlatform Cloud Edition extends the benefits of its on-premises consolidated DataPlatform secondary storage solution to public cloud infrastructure. Cohesity has headquarters in Santa Clara, California. For more information, visit <http://cohesity.com/> and follow @Cohesity on Twitter.



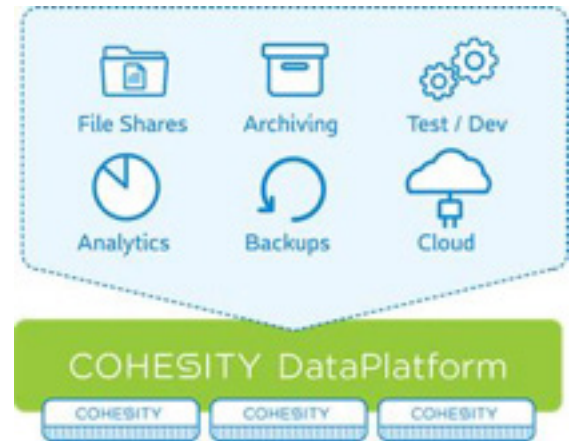
**80% of Data is Secondary**

## Cohesity DataPlatform On-Premises Secondary Storage

Cohesity has been shipping the on-premises DataPlatform since early 2016. The DataPlatform consolidates silos of on-premises secondary storage into a single operating environment. The solution enables the consolidation of file shares, archiving, test/dev, analytics, backups, and objects into a single storage infrastructure. With data consolidated into a common source, administrators can create clones from backup data for analytics and test/dev processes.

In addition to saving capacity, the DataPlatform simplifies administration by eliminating unnecessary copies and preventing unsynced versions used for essential testing or decision-making. The single repository eliminates the stranded, unused capacity that cannot be shared across the multiple point products. Cohesity DataProtect provides a complete backup and recovery solution fully integrated with DataPlatform. The single data protection scheme unifies end-to-end data protection infrastructure, including target storage, backup, replication, disaster recovery, and cloud tiering. It simplifies backup infrastructure and eliminates the need to run separate backup software, proxies, media servers, and archive software.

The DataPlatform is a software-driven architecture purpose-built for secondary storage. The DataPlatform is designed to scale to sizes that enable consolidation of all secondary data. The solution utilizes a distributed scale-out file system that resides on 12TB and 24TB nodes. The nodes are extensible into four node blocks, scaling out performance and capacity simultaneously to reduce complex configuration planning. Connectivity includes 1 GbE and 10 GbE, and the supported protocols are NFS V3, SMB 3.0, and REST. Storage services include those found on primary tier storage, including global deduplication, snapshots, analytics, encryption, intuitive management, and a RESTful API.

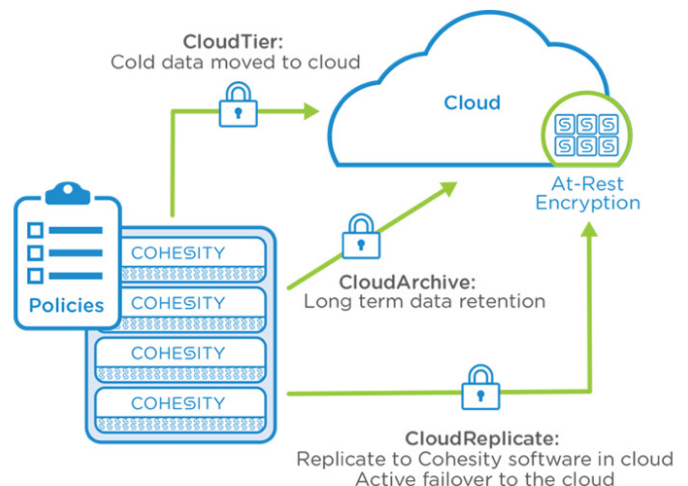


## Cohesity DataPlatform Cloud Edition

IT organizations and businesses owners are attracted by the self-service, on-demand, and elastic provisioning characteristics of the cloud. However, many IT organizations have issues related to compliance, security, and performance with a multi-tenant public cloud. Most enterprises take a hybrid approach to the cloud, utilizing both on-premise infrastructure and private cloud software as well as public cloud service providers. Cohesity extends its on-premises DataPlatform solution to leverage the cloud with the DataPlatform Cloud Edition.

The Cloud Edition provides seamless connectivity to public cloud storage and services. Supported targets include Google Cloud Storage, Microsoft Azure, and Amazon S3/Glacier.

The Cloud Edition takes a comprehensive approach to cloud adoption with the CloudArchive, CloudTier, and CloudReplicate features. CloudArchive gives customers the ability to seamlessly archive data for extended retention from the Cohesity Data Platform through pre-built integrations with cloud service providers. CloudTier gives



customers the ability to use the public cloud as an extension of their on-premise storage. It dynamically increases local storage capacity by moving less active data into the cloud. The CloudReplicate feature duplicates local storage instances to remote public or private cloud services for a low-cost disaster recovery solution. Since data is de-duplicated on-premises with the DataPlatform, using the DataPlatform Cloud Edition reduces network resource requirements and speeds up the transfer of data. Inherent encryption protects data in-flight and isolates data between tenants in the cloud.

## EMA Perspective

Cohesity's DataPlatform solution makes secondary storage dramatically more efficient and easier to manage by consolidating backup, archive, test and development, file services, and replication on a single scalable platform. In addition to lowering capital expenses and increasing operational efficiency, the single platform simplifies version control when repurposing data for multiple use cases. Business analytics or test development can leverage the existing data protection source without making additional copies. Also, with the DataPlatform Cloud Edition, Cohesity adds cost-saving opportunity for businesses. What sets Cohesity apart from other solutions is that there is a single user interface for management, protection, and data movement for data located on-premises or in the cloud. EMA applauds Cohesity for delivering a solution that further improves data center efficiency for secondary storage.

### About EMA

Founded in 1996, Enterprise Management Associates (EMA) is a leading industry analyst firm that provides deep insight across the full spectrum of IT and data management technologies. EMA analysts leverage a unique combination of practical experience, insight into industry best practices, and in-depth knowledge of current and planned vendor solutions to help EMA's clients achieve their goals. Learn more about EMA research, analysis, and consulting services for enterprise line of business users, IT professionals, and IT vendors at [www.enterprisemanagement.com](http://www.enterprisemanagement.com) or [blogs.enterprisemanagement.com](http://blogs.enterprisemanagement.com). You can also follow EMA on [Twitter](#), [Facebook](#), or [LinkedIn](#).

3503.122916