COHESITY

WHITE PAPER

Solve Hybrid Cloud Challenges for Secondary Data and Apps

IT organizations are undergoing a significant transformation as they move toward consolidating data centers, and migrate many workloads and their data to the cloud. Today, hybrid cloud is increasingly the norm, and enterprises are challenged with ways to have visibility, manage and make use of all this data—both on-premises and in the cloud. While much attention has been given to primary data affecting mission-critical workloads, data that lives in the secondary realm—backup, archiving, test/dev, and disaster recovery, to name a few—have become siloed the same way application data has, leading to multiple point solutions to manage an increasing amount of data.

This white paper looks at the evolution of these challenges and offers practical advice on ways to store, manage and move secondary data in hybrid cloud architectures while extracting the hidden value it can provide.

The Cloud and the Data Deluge

In today's data-driven age, enterprises require visibility into the data they already have, and must be able to properly mobilize that data—providing it to the people and organizations that need it in a way that furthers business goals in a timely, secure manner.

For many organizations, public cloud providers like AWS, Azure and GCP are helping them to increase their agility and also deal with the accelerating growth in the variety and amount of data. Furthering this trend is what's happening on the application front, as apps and their associated data are increasingly born in the cloud, or rapidly moving toward the cloud. Unfortunately, as enterprises adopt a hybrid approach, legacy backup and recovery architectures and solutions just can't keep up.

Traditional backups are unreliable, with one out of four restore attempts failing.[1] Homegrown scripts and manual backup processes often have gaps in protection as applications come and go, and workloads move from on-premises to the cloud. What's worse, enterprises that amass a variety of point solutions for each data protection use case end up creating the same data silos that IT has wrestled with for decades.

Challenges with Building a Hybrid Cloud

A recent Gartner poll indicated that developing a hybrid cloud strategy is the top challenge[2] for IT leaders, as data increasingly comes from all quarters and business units seek new applications to make use of it. Increasingly, hybrid cloud also means multi-cloud, as businesses turn to two or more public cloud providers for the sake of business continuity. The end result is more silos, as organizations that do more in the cloud often find data increasingly fragmented and difficult to extract value from.

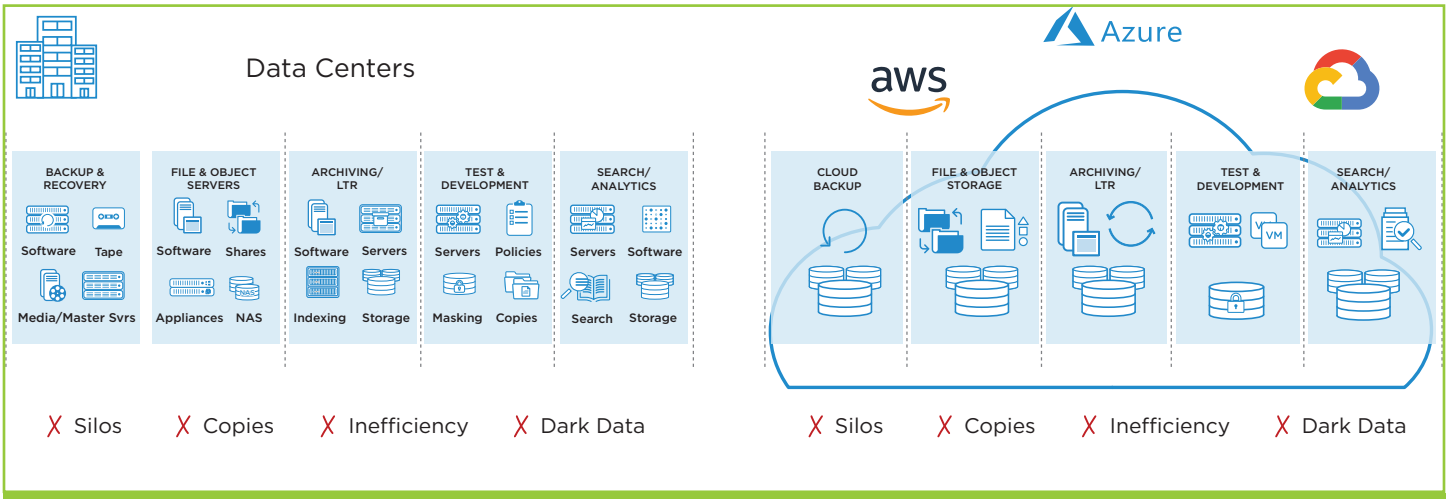

On-premises storage products present one set of silos, cloud storage presents yet another set.

Legacy storage products were not designed to manage data in hybrid cloud environments. On-premises data center storage consists of point solutions such as NAS appliances, deduplication appliances, archive storage and analytics clusters. Cloud storage presents another, different set of silos including object, file and cold archive storage.

Each of these silos is optimized for a limited set of use cases and is managed through its own dedicated user interfaces or APIs. Furthermore, traditional storage products aren't designed to manage and transfer data between on-premises and public cloud infrastructure—such transfers require a bolt-on cloud gateway, yielding yet another silo. The resulting infrastructure is complex, expensive, inefficient and can't keep up with the rapidly changing needs of the hybrid cloud.

1 "Survey: Users lose data despite backup," Kroll Ontrack LLC, March 2017.
2 "Top 4 Challenges Facing IT Infrastructure Leaders," Gartner, April 2018.

Enterprise Impact

There are a number of ways this fragmentation across a hybrid cloud environment hinders business outcomes:

- **Increased complexity:** Multiple point vendors translates to more complexity—cloud gateways, integration with a variety of storage providers and types—and all of this translates into an increased IT workload required to manage multiple products, vendors and interfaces.
- **Cost:** Enterprises may find they must be more thoughtful about which types of workloads go to which cloud provider, taking into consideration both storage and bandwidth costs, among others.
- **Unpredictable SLAs:** Service level agreements (SLAs) vary from provider to provider, leading to unpredictability and potential loss of access.
- **Lack of easy data mobility:** Multiple protocols and formats means moving data from cloud back to on-premises, or between clouds. Both of these are challenging and time-consuming. In some cases, data transfers are so large that you may have to resort to shipping physical media such as AWS Snowball, Azure Data Box and Google Transfer Appliance.
- **Lack of visibility and control:** It's difficult to keep track of data across your on-premises data center, and it's now being further fragmented in public cloud environments. Where is your data stored? Who owns the data? What is it being used for? What must be retained for how long? Do I have data that puts me at risk for compliance or regulatory violations? All these questions become exponentially harder to answer as your data is spread across clouds.
- **Limited security:** Prior to transferring data to another location, it should be encrypted and secured. This raises a number of challenges related to managing encryption across sites, rotating encryption keys, integrating with external key management systems, etc.

Adopt a Cloud-First Approach

Since legacy solutions by nature aren't "cloud-forward," they can't meet these challenges. What's needed is a holistic solution that spans multi-cloud and on-premises data for secondary workloads.

Your Cloud Initiative Can Begin Anywhere

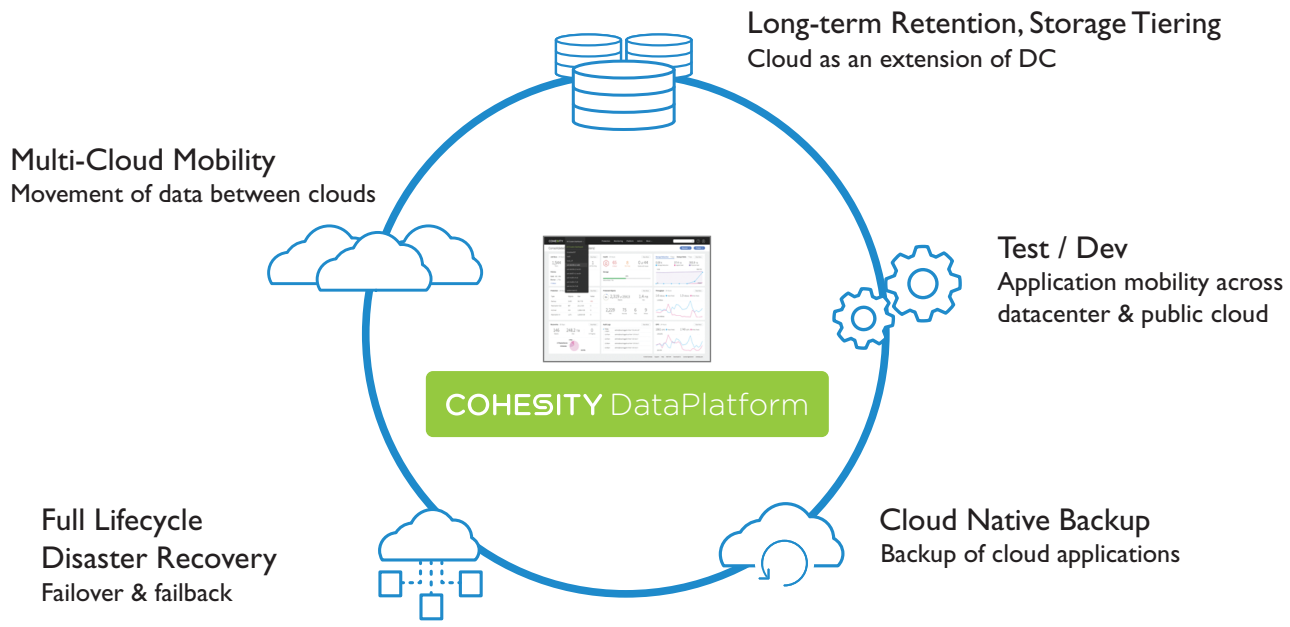

Your cloud initiative can start anywhere and, for many organizations, a good place to start is with one of the biggest problems for secondary data—namely, data protection, in particular:

- Long term retention and archival of data for regulatory or governance purposes
- Storage tiering to ensure frequently accessed data in the most performant storage devices
- Data mobility between on-premises and public cloud environments

Increasingly, public cloud has become a backup destination, with Gartner predicting that one in five enterprises will utilize the public cloud for backup by 2020.[3]

Built with native cloud integration, Cohesity DataPlatform addresses all these challenges and more. It integrates with public cloud platforms, not as a bolted-on afterthought but from the ground up.

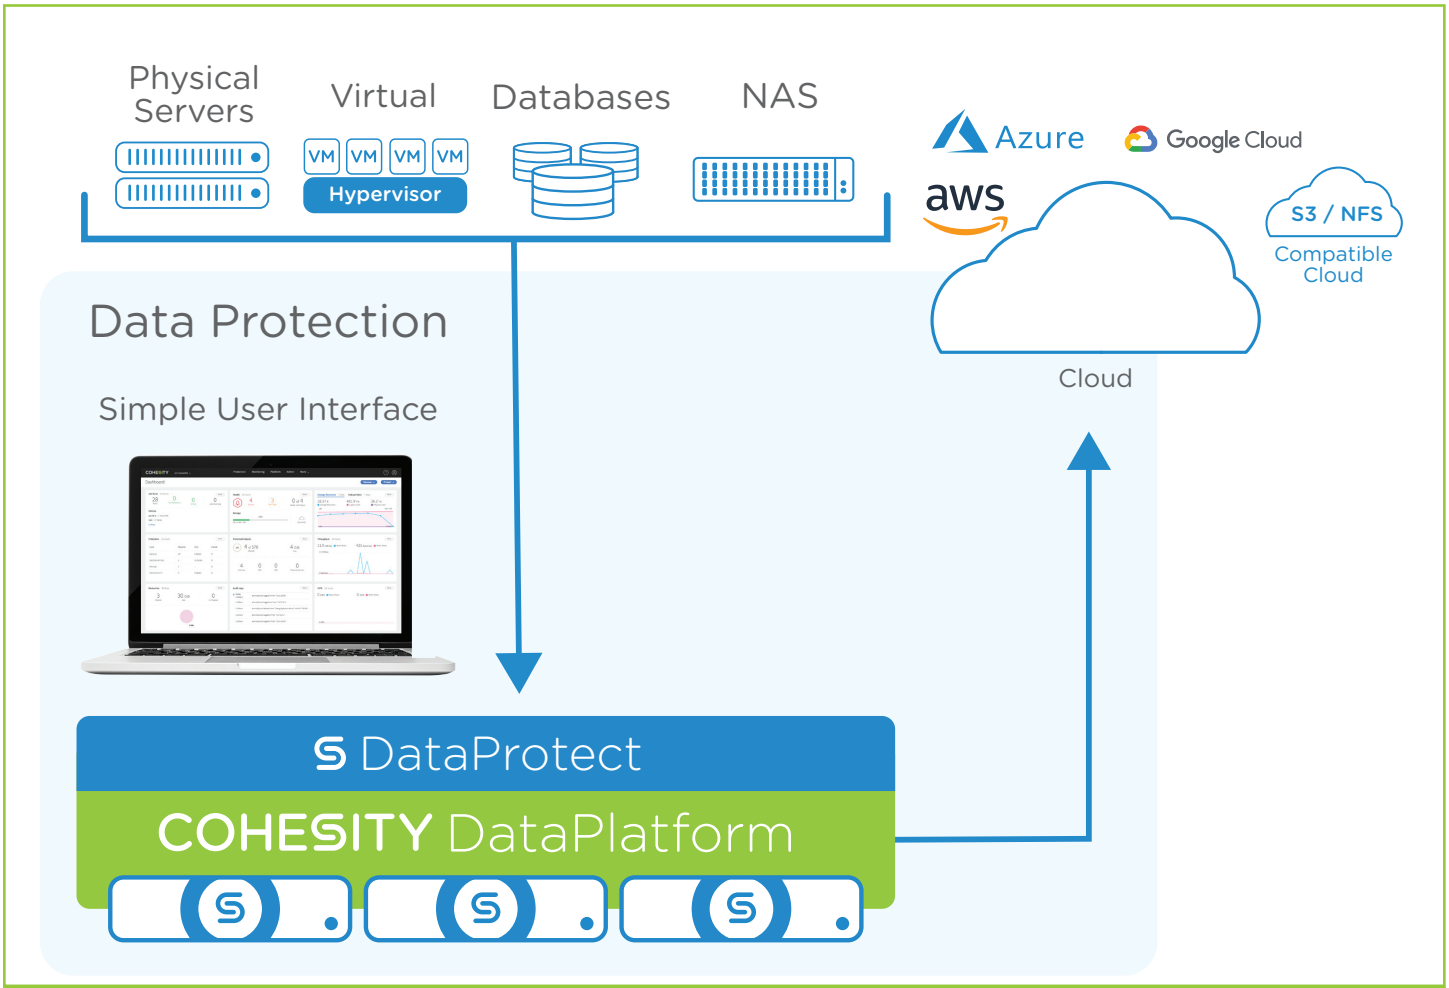

Native integration with Microsoft Azure, AWS and Google Cloud Platform allows you to extend on-premises secondary storage use cases to the cloud

3 "How to Leverage Public Cloud IaaS for Backup & Recovery," Gartner, August 2017.

Cohesity's cloud-agnostic design offers broad support for all major public cloud platforms and features a distributed file system that spans on-premises and cloud data, regardless of where it resides. The result is a platform that supports multiple secondary data use cases, both on-premises and in the cloud. This translates to unique benefits including:

- Faster recovery time for archived and tiered data
- Ability to perform instant mass restores
- Security of end-to-end data encryption, both at rest and in motion
- Ability to handle both on-premises and cloud-based data with equal agility

Build for Beyond Backup and Recovery

Accelerate Test & Development in the Cloud

Application owners and developers typically have to wait days or even weeks to get data that they can use for test/dev purposes. But wouldn't it benefit everyone to gain access to backup data faster? Cohesity helps turn backups into much more than an insurance policy by offering a complete set of cloud and data integrations to help unlock the value in cloud data.

Now, data for test/dev purposes can be simply provisioned, converted for use in the cloud and moved back with ease. Test/dev workloads are easily spun up when needed and perhaps more importantly, spun down when no longer required to save on cloud expenditures.

As a result, enterprises can eliminate the need for expensive copies of data, and ease application mobility between cloud and on-premises for test/dev, enabling otherwise idle backup data to be utilized while protecting the enterprise.

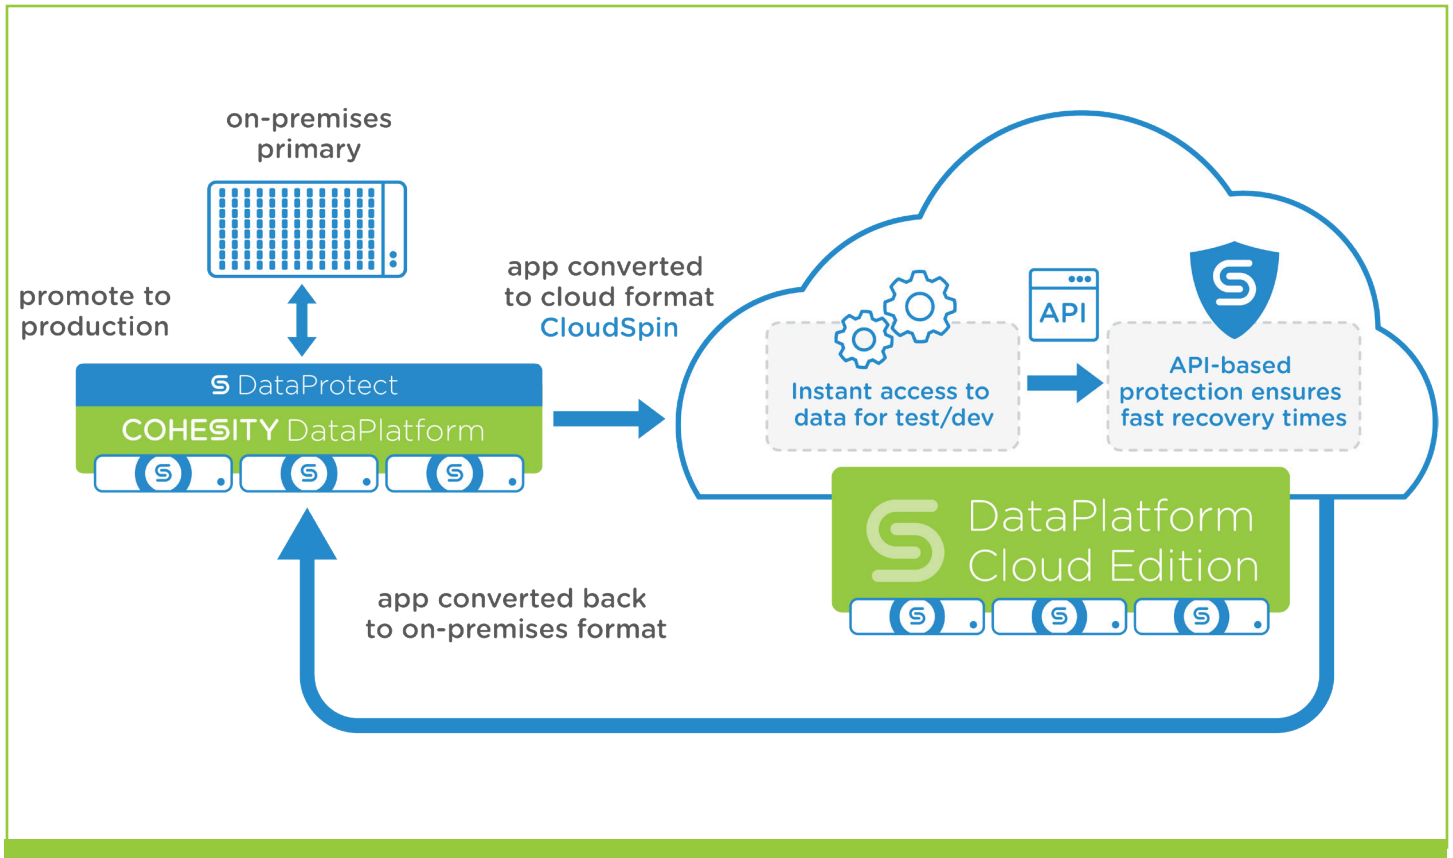

Accelerate application test and dev between on-premises and cloud

Simplifying Disaster Recovery

There is only one reason why businesses spend time and energy on data protection: to keep the business running. Increasingly, business continuity and disaster recovery (BCDR) is as important as backup, as natural disasters and cyber-attacks can otherwise leave businesses without access to data, with potentially devastating consequences.

Since all data is not created equal, having the flexibility to apply the right BCDR strategy—at a granular level—is critical, especially in today's hybrid environments.

Cohesity offers a range of data protection solutions including CloudRetrieve, CloudReplicate and CloudSpin, which, in combination, enable a variety of BCDR implementation strategies including:

- Replication from one on-premises site to another—or to the public cloud to mitigate against complete failure, and simple failback of data when failure occurs.
- Enabling data mobility from and to anywhere, leveraging policy-based automation for replication for hybrid clouds.
- Creating virtual clusters in the cloud, with the ability to replicate data from on-premises to cloud clusters.
- Converting on-premises VMs to the cloud for BCDR, spinning up and down instances as needed with automatic format conversions.

Three Ways to Implement Disaster Recovery

CloudRetrieve

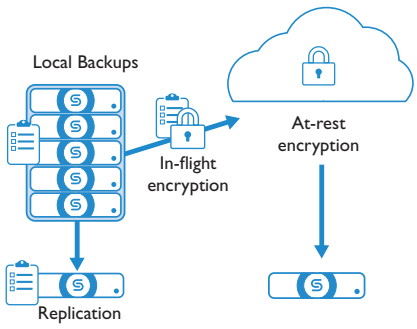

Use CloudRetrieve to recover archived data from the cloud.

CloudSpin

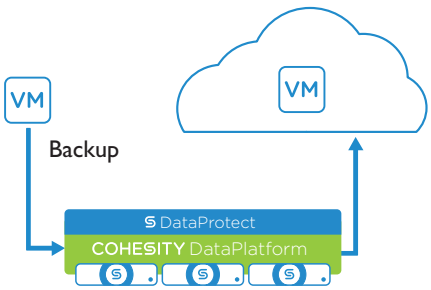

Convert on-premises VMs to cloud format and easily spin them up and down.

Cloud Edition with Replicate

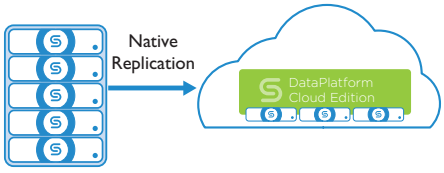

Replicate data from on-premises to the cloud.

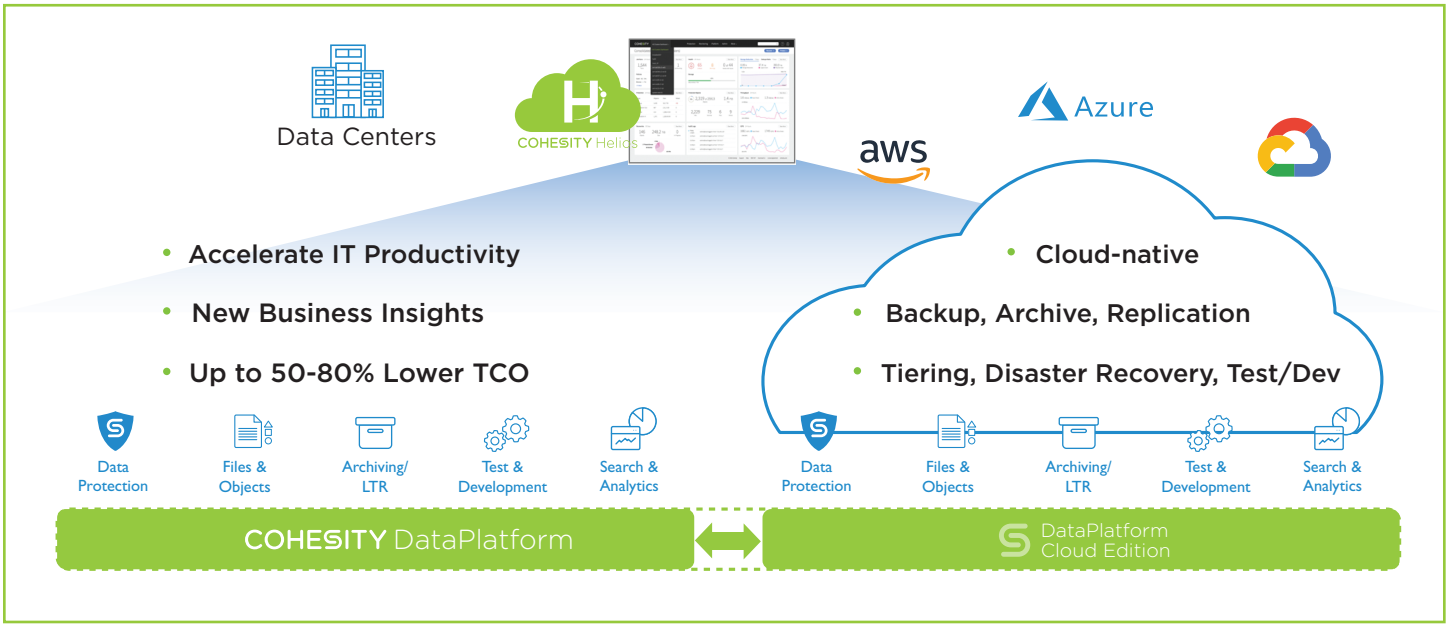

Why Now?

Nearly every enterprise is moving toward a more hybrid and complex data environment. Cohesity helps to reduce the number of data silos, providing a single platform for archival, backup, BCDR, data tiering, and support of test/dev data provisioning, all on the way to accelerating the journey to a multi-cloud, hybrid IT environment.

IT resources are precious, and organizations are always seeking ways to improve productivity. Cohesity offers policy-based automation that further reduces the IT management burden and frees up resources to focus on other priorities. The result is visible in cost savings, improved data availability, a higher return on investment and secondary data being more valuable than previously thought.

Conclusion

The need to smoothly and transparently span data and workloads from on-premises to one or more cloud providers and back is already apparent. Cohesity DataPlatform is the industry's first true web-scale platform for all secondary data and apps on-premises and in the public cloud. With the advent of a hybrid cloud environment, the Cohesity DataPlatform allows enterprises to improve efficiency and reduce total cost of ownership (TCO) across the data center, cloud and edge.

Click here **to find out more about how Cohesity can help you make the most out of your secondary data and apps.**

About Cohesity

Cohesity delivers web-scale simplicity for secondary data that eliminates silos and puts data to work. Enterprises rely on our software-defined, hyperconverged platform to radically streamline backup and data protection, converge file and object services, quickly deliver test/dev instances, and provide analytic functions on a single, global data store. Customers and partners, including Global 1000 companies and federal agencies, tired of data center complexity are modernizing and scaling secondary data protection with Cohesity.

Cohesity, Inc.
Address 300 Park Ave., Suite 800, San Jose, CA 95110
Email contact@cohesity.com www.cohesity.com