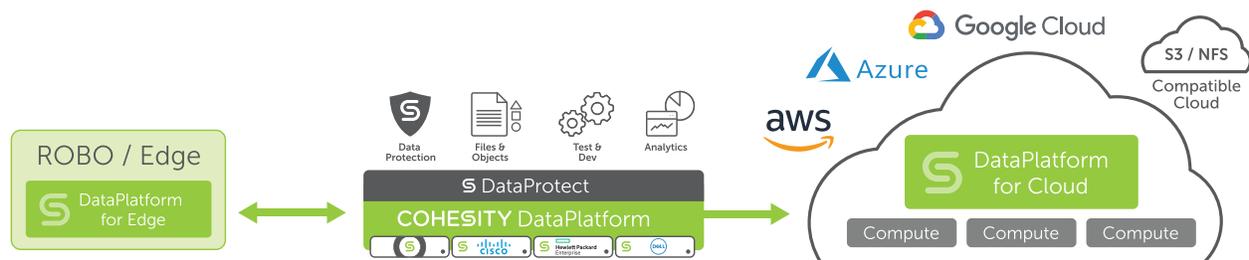


Cohesity DataPlatform

Cohesity DataPlatform™ consolidates secondary data and applications including backups, files, objects, test/dev, and analytics on a single, software-defined platform. Inspired by web-scale architecture, DataPlatform is a scale-out solution based on a unique distributed file system, SpanFS®. DataPlatform modernizes and simplifies secondary data and application management by providing one platform for multiple secondary workloads. Although most organizations begin their journeys to overcoming mass data fragmentation by simplifying data protection, DataPlatform's flexible architecture allows easy expansion to additional use cases, further increasing operational simplicity and improved TCO. Because DataPlatform is a software solution, it works equally efficiently on-premises on qualified Cisco, HPE, Dell, or Cohesity hyperconverged platforms as well as in the public cloud.



Web-Scale Simplicity and Efficiency

The flexible, resilient, and on-demand web-scale benefits pioneered by hyperscalers such as Google are now available to enterprises with DataPlatform. The solution eliminates unnecessary data silos and spans workloads as it converges secondary data on a single platform. It accommodates data growth by delivering scale-out file and object storage in a limitless, pay-as-you-grow model without disruptive upgrades.

Easy, Cloud-Ready Deployment

DataPlatform is cloud-ready, ensuring seamless and rapid data and application mobility with leading public cloud providers including Amazon Web Services (AWS), Microsoft Azure, and Google Cloud Platform. Take advantage of public cloud elasticity and economics using policy-based automation for archiving, tiering, and replication to the cloud. Move virtual machines (VMs) to the cloud for test/dev and then back again when ready.

One Platform, Endless Possibilities

DataPlatform is powered by a patented distributed file system architecture, making it much more than a modern converged target for backup data. It can be used to recover applications instantly from any point in time to meet business SLAs. Native applications also simplify file data search, speed eDiscovery, and streamline compliance. Built to support an API-first approach, DataPlatform integrates with leading automation solutions and custom business scripts to increase operational efficiency.

Specifications

Hyperconverged Architecture	<ul style="list-style-type: none"> • Cohesity SpanFS, a distributed file system • Web-scale architecture with unlimited scalability • Non-disruptive upgrades • Unlimited, distributed snaps and clones with no performance impact with Cohesity SnapTree
Multiprotocol Support	<ul style="list-style-type: none"> • Distributed volumes with multiprotocol NFS, SMB, and S3 access • Support for simultaneous sequential and random IOs
Storage Efficiency	<ul style="list-style-type: none"> • Global deduplication across all nodes, variable length, inline, or post-process • Compression, including Zstandard • Erasure coding 2:1, 4:2, and 5:2, and replication factor
Data Resiliency	<ul style="list-style-type: none"> • Strict consistency across nodes within a cluster, ensuring data resiliency and that reads always return the most recent data • Two node and two disk failure tolerance support
Use Cases	<ul style="list-style-type: none"> • Backup and recovery • Target storage • Files and object storage • Security and compliance • Analytics • Test/dev • Long-term retention • Disaster recovery
Supported Applications	<ul style="list-style-type: none"> • Cohesity DataProtect – DataPlatform is the foundation for DataProtect’s delivery of comprehensive data protection for all workloads, anywhere on a single platform that converges backup, recovery, replication, and target storage with fast RPOs and near-instant RTOs • Cohesity Helios – SaaS-based global secondary data and apps management for Cohesity environments • Pattern Finder – Searches for multiple patterns, such as \d{3}-\d{2}-\d{4} at once across a wide variety of files and formats—from text and office files to PDF and Zip formats—simplifying discovery • Password Detector – Searches for files with alphanumeric strings, raising the probability that a string is also a password • Video Compressor – Compresses video file formats, for example .mpeg, .avi, .mp4, .flv, .wmv, .mov, reducing retained file storage sizes • Custom Apps – Develop custom apps via Java archive files
Global Search	<ul style="list-style-type: none"> • Automated global indexing • Wildcard searches for any VM, file, or object on the platform
Flexible Deployment	<ul style="list-style-type: none"> • Software-defined solution for on-premises, public cloud, and edge sites • Private cloud/data center – Hyperconverged secondary storage appliances, supporting qualified Cisco, HPE, and Cohesity hardware • Virtual Infrastructure - DataPlatform for virtualized infrastructure (VMware vSphere) • Public cloud – DataPlatform Cloud Edition for a variety of use cases in leading public clouds • Remote/edge sites – Simplified secondary data and applications management with DataPlatform for Edge
Public Cloud Integration	<ul style="list-style-type: none"> • Native cloud integration with all leading public clouds including AWS, Microsoft Azure, Google Cloud Platform, and any S3 compatible clouds, as well as AWS Govcloud and Microsoft Azure GovCloud
Global Management	<ul style="list-style-type: none"> • Cohesity Helios for multi-cluster, global active management • vRealize • ServiceNow
API Integration	<ul style="list-style-type: none"> • API-first architecture • RESTful API • OpenAPI standard
Encryption	<ul style="list-style-type: none"> • Software-defined AES-256, FIPS compliant encryption of data in flight and at rest

Overcome mass data fragmentation by simplifying and unifying all secondary data and applications on Cohesity DataPlatform. Learn more at [Cohesity.com/products/data-platform](https://cohesity.com/products/data-platform).

CAPABILITIES	DATAPLATFORM BASE	DATAPLATFORM STANDARD
Enterprise-class Storage		
Unlimited Scalability (web-scale)	✓	✓
NFS, SMB, S3, Multi-protocol Support	✓	✓
Non-disruptive Software Upgrades	✓	✓
Non-disruptive Add/Remove Nodes/ Blocks	✓	✓
Unlimited, Distributed Snaps	✓	✓
Multi-Tenancy	✓	✓
App, Infrastructure test/dev	✓	✓
User Quotas, Storage quotas	-	✓
Capacity Optimization		
Global Variable-Length Dedupe	-	✓
Compression	-	✓
Unlimited Clones	-	✓
Data Integrity		
Data Resiliency w/ Strict Consistency	✓	✓
Two Node Failure Tolerance Support	✓	✓
Two Disk Failure Support	✓	✓
Automated Data/App Mobility		
Archive to Cloud/On-premises	✓	✓
Automated Data Tiering	✓	✓
Make Data Productive		
Cohesity Native Apps (AWB)	✓	✓
Security		
Datalock - WORM for Security	-	✓
Encryption AES-256 (In-flight / at rest)	-	✓
Global Active Management		
Multi-cluster Single Sign-on	✓	✓
Multi-cluster Dashboard	✓	✓
Global Search	Read Only	Actionable
Support Automation	✓	✓
Global Pre-Built Reports	✓	✓
Global Customizable Reports	-	✓
Global Policy	-	✓
Orchestrated Cluster Upgrade	-	✓
Long Term Retention (10 years) of meta- data	-	✓
Proactive Wellness	-	✓
Self-Serving	-	✓
Capacity Prediction	-	✓
What-if Analysis	-	✓
Performance Balancing Recommendations	-	✓
Supported Environments		
Public Cloud	✓	✓

Physical (on certified appliances)	✓	✓
Virtualized	-	✓
Edge/Robo	-	✓
Supported Workloads		
Backup Target	✓	✓
IOT Big Data	✓	-
Video Surveillance	✓	-
Healthcare Pacs	✓	-
Hadoop and Splunk Cold Buckets	✓	-
Archive	✓	✓
Backup and Recovery	-	✓
Disaster Recovery	-	✓
Home Directories	-	✓
File Shares	-	✓
Web-Scale NAS	-	✓
Test and Development	-	✓
Data Analytics	✓	✓