

Scale-out File and Object Storage, from Edge to Cloud

Is the growth of unstructured data putting pressure on your legacy storage silos?

According to Gartner, data will grow by 800% over the next five years, of which 80% will be unstructured in form of file shares, backups, archives, logs, media files, test/dev and analytics. Traditional network attached storage (NAS) was designed more than 10 years ago for these use cases, however no one solution addresses all of these workloads. Different vendor solutions are optimized for a specific subset of unstructured data workloads and have limited scale. This results in multiple storage silos, running on proprietary hardware, with different software and licensing.

“By 2021, more than 80% of enterprise data will be stored in scale-out storage systems in enterprise and cloud data centers, up from 30% today.

- Gartner

Secondary Storage

These workloads - backups, archives, file shares, logs, media files, test/dev and analytics - are not mission critical, yet they consume about 80% of enterprise storage capacity. This set of workloads can be categorized as secondary storage for two reasons. One because the data is not in the production environment and two, these workloads typically don't have strict SLA requirements. Yet this secondary storage imposes a huge burden on enterprise IT budgets, typically greater than what is needed for primary storage! In addition, because the infrastructure is fragmented, complex to manage, and inefficient, operational costs are high.

Rethink Files and Objects for the Cloud Era

This problem won't be solved by yet another point solution. A new architecture is needed that takes into account the radical changes that have occurred in the last 10 years since NAS was popularized. One that accommodates the exponential growth in data and web-scale requirements. And second, one that seamlessly integrates with public and private clouds.

Cohesity does that with an innovative new Google-inspired architecture. Cohesity DataPlatform simplifies secondary storage by consolidating all files and objects on a single, web-scale platform.

Key Benefits

Extensibility

- Consolidate files and objects silos on a single platform
- Manage all your data through the protocol of your choice with multiprotocol access (NFS, SMB and S3)

Web-scale

- Scale-out seamlessly to address data growth
- Expand easily with a pay-as-you-grow model
- Eliminate forklift upgrades with non-disruptive software and hardware upgrades
- Leverage commodity economics with software-defined storage that can be deployed on industry-standard hyperconverged nodes

Cloud-Ready

- Deploy in your data center, private cloud, edge, or Microsoft Azure, AWS or Google Cloud
- Leverage on premises benefits in the cloud with native protocol access

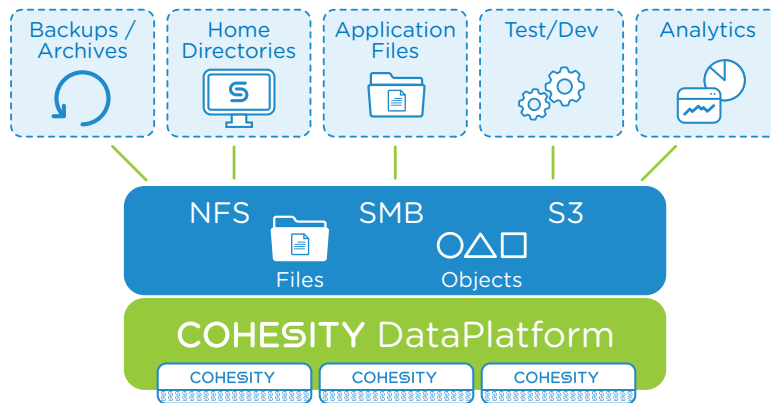
Storage Efficiency

- Gain unparalleled storage efficiency with true global dedupe, even across nodes
- Compress your cold or streaming data by 5x to 10x
- Make your storage environment elastic to address changing business requirements instantly

Data Productivity

- Accelerate application development by instantly provisioning zero-cost clones
- Gain greater insight into your data with in-place analytics





Cohesity DataPlatform is powered by SpanFS™, a purpose-built distributed file system that's been optimized to manage unstructured data at scale. SpanFS is the only distributed file system that combines global deduplication, strict consistency, unlimited snapshots and clones with SnapTree, and native integration with the public cloud.

Easily Extensible

Most organizations have multiple vendors for their file and object storage - some specialized on corporate directories, a different one for media content, and yet other for analytics. Cohesity consolidates all these workloads all on a single, scale-out platform. Cohesity DataPlatform allows you to read and/or write to the same data volume with simultaneous multiprotocol access for NFS, SMB and S3, from data center to edge to the cloud.

Web-scale

Legacy storage doesn't allow you to provision just what you need, or share the compute resources when under utilized. Cohesity is a true web-scale solution that scales limitlessly. With DataPlatform's pay-as-you-grow model, you can start as small as three nodes and simply add more nodes to address growing business requirements, with automatic usage optimization.

Cohesity DataPlatform is the only software-defined storage solution that guarantees data resiliency at scale with strict consistency.

Cloud-ready

Legacy storage systems weren't designed for our cloud-first world. Bolt-on cloud gateways are expensive and add complexity. Cohesity natively integrates with Microsoft Azure, Amazon and Google Cloud, and allows you to leverage the economics and elasticity of the public cloud.

Storage Efficiency

Legacy systems offer limited storage efficiency with deduplication only at the node or the block level. This drives up storage costs significantly. Cohesity offers true global variable-length deduplication, along with amazing compression and erasure coding to maximize every dollar.

Data Productivity

Unlike legacy solution, Cohesity makes your data productive. Cohesity DataPlatform allows you to instantly provision zero-cost clones and run test/dev environment on Cohesity to accelerate application development. This helps to eliminate unnecessary data copies and reduces data center footprint.

Software Defined for Flexible Deployment

As true software-defined storage, you have the flexibility to deploy Cohesity DataPlatform in a variety of ways to meet your needs. Choose Cohesity C2000 or C3000 hyperconverged appliances that are pre-configured and tested for your data center. Or gain the same capabilities on pre-qualified Cisco UCS and HPE Proliant servers. For branch or remote offices, use Cohesity Virtual Edition for rapid, remote deployment. , And, as you need the flexibility of the public cloud, choose Cohesity Cloud Edition.

Key Features

Features	Description
NFSv3, CIFS, SMB2.x, SMB 3.0, and S3 APIs	Multiprotocol access to same data allows support of applications across all major enterprise operating systems including Microsoft Windows, Linux and S3 API
Strict Consistency	Guaranteed data resiliency at scale
SnapTree® snapshots and clones	Limitless and fully-hydrated snapshots for granular Cohesity Views (file systems) as well as writable snapshot clones that provide instant creation, testing and development of view-based datasets
Web-scale File System	Limitless scalability, always-on availability, non-disruptive upgrades, pay-as-you-grow model
Hyperconverged Secondary Storage	Single platform for data protection, files, objects, test/dev, and analytics
Global deduplication and compression	Unparalleled storage efficiency with global deduplication and compression across all nodes of the cluster that significantly reduces data center footprint
Erasur coding	Data is protected against any individual node failure with erasure coding across nodes
Global indexing and search	File and object metadata is indexed upon ingest, enabling Google-like search across all files in a cluster
Mix-mode permission mapping	Cohesity manages the permission mapping and also natively integrates with Centrify. Centrify allows Cohesity to directly access the ID mapping information stored in Centrify's AD. This eliminates the need for LDAP proxy and simplifies the user experience.
Windows Active Directory and Kerberos Integration with Role-Based Access Control (RBAC)	Simplify user and group access to data utilizing credentials and permissions with Windows AD and Kerberos mechanisms. Create and manage custom Cohesity cluster administration roles for domain users and groups
External KMS integration	Snap-in for the Microsoft management console, which allows Cohesity file shares to be managed MMC
Quotas	Easily establish user and file system quotas with audit logs
Policy-based backup protection	Integrated data protection software and SnapTree technology is available to allow simplified data protection of objects with fully-hydrated snapshots
QoS	QoS policies are provided that optimize performance for different types of workloads
Encryption	Cohesity solution provides data-at-rest as well as data-in-flight encryption using industry standard 256-bit Advanced Encryption Standard (AES) algorithm. The platform is also FIPS 140-2 compliant
Write Once Read Many (WORM)	Enables long-term retention of data that have compliance controls mandating a policy that objects cannot be modified during the lock time
Replication for Disaster Recovery	Built-in, granular, and secure replication services for geo redundancy
Cloud integration (CloudArchive, CloudTier, CloudReplicate)	Archive into public cloud services for long-term retention. Utilize cloud tiering for transparent capacity expansion into the cloud. Replicate into the cloud for disaster recovery and test/dev