



Unlimited Scale-Out File and Object Storage

Unlimited Scale-Out Files and Objects

Struggling to Manage Unstructured Data at Scale?

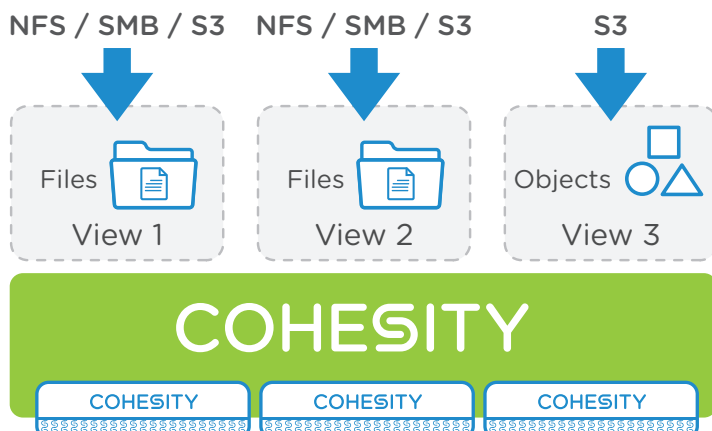
Unstructured data is growing 30% to 40% a year. IDC estimates that enterprises were managing a total of 6 ZB in 2016, and sees that number increasing to 93 ZB by 2025, with 80% being unstructured data in the form of files and objects. Enterprises have traditionally been using NAS appliances to store files. But traditional NAS is expensive and doesn't scale out – not a great fit for managing rapidly increasing data volumes. More recently, enterprises have started adopting scale-out NAS such as Dell EMC Isilon. But Isilon is delivered as expensive proprietary hardware appliances, with scaling limits (144 nodes max), and complicated InfiniBand requirements. Software-defined object storage is starting to make inroads into the enterprise, but most enterprise workloads weren't designed to use object storage, making adoption slow and difficult.

Enterprises need a more scalable and efficient solution to manage unstructured data. That's the problem that Cohesity is setting out to solve, with a web-scale platform designed to consolidate all secondary storage, including files and objects.

Cohesity Consolidates Files and Objects at Scale, with Global Deduplication, Global Search, and Multiprotocol Access

Cohesity provides globally distributed NFS, SMB and S3 object storage on Cohesity DataPlatform, a web-scale secondary storage platform. NFS, SMB and S3 volumes are provisioned as Views on Cohesity and benefit from all the DataPlatform capabilities. Views are globally deduplicated and distributed across all the nodes in the cluster for performance and capacity. DataPlatform indexes all the file and object metadata to enable global Google-like search across an entire cluster. It also provides software-based encryption, replication, snapshot-based data protection, pay-as-you-grow scalability, and always-on availability. Quotas can be established for volumes, file shares, and object buckets. Cohesity also integrates natively with the public cloud for archival, tiering and replication of files and objects in the public cloud.

Views provide multi-protocol access to data and support simultaneous access through NFS, SMB and S3 protocols. Data security is maintained using protocol specific permission and access controls. Active Directory and Kerberos authentication integration provide user and group directory and credential management. Additionally, QoS policies can be created that prioritize workloads across the cluster and advanced, integrated data protection can be enabled to protect file and object data.



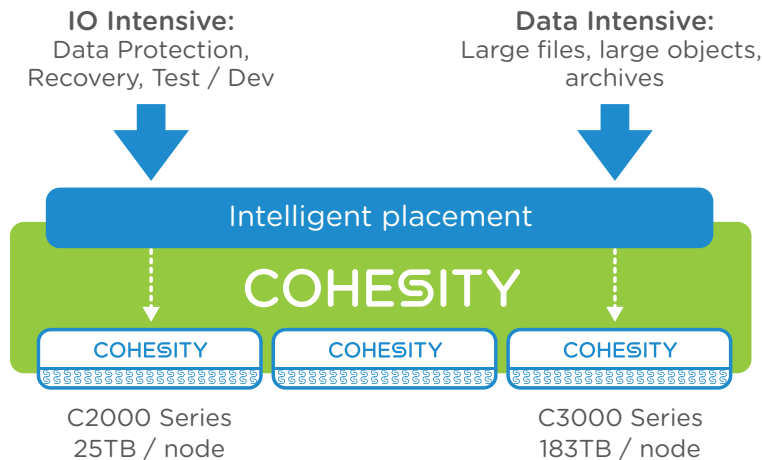
Key Benefits

- NFS, SMB, and S3 compatible object storage with multiprotocol access supporting NFSv3, CIFS, SMB2.x, SMB 3.0, and S3 APIs
- Converged on Cohesity DataPlatform which can be used to consolidate files, objects, data protection, and test/dev copies
- Web-scale platform with unlimited scalability, non-disruptive upgrades, pay-as-you-grow scalability
- Global deduplication and compression
- Global Google-like search on all file and object metadata
- User and file system quotas with audit logs
- C3000 capacity-optimized storage node optimized for large files and objects
- Public cloud integration for archival, tiering and replication

C3000 Dense Storage Node

Cohesity DataPlatform is a software-defined solution deployed on hyperconverged nodes. Cohesity offers two series of hyperconverged nodes: the C2000 and C3000. The C2000 series is designed for more CPU intensive workloads including data protection and recovery, and the C3000 series is capacity-optimized to consolidate large files and objects. The C3000 series offers almost twice the storage density of the C2000 series, with up to 183TB of raw capacity per node in a 2U form factor.

C2000 and C3000 series can be mixed in a single Cohesity cluster. DataPlatform provides intelligent placement of data based on QoS and IO profiles. More CPU and IO intensive data is preferentially placed on C2000, while less CPU and IO intensive data is preferentially placed on the C3000.



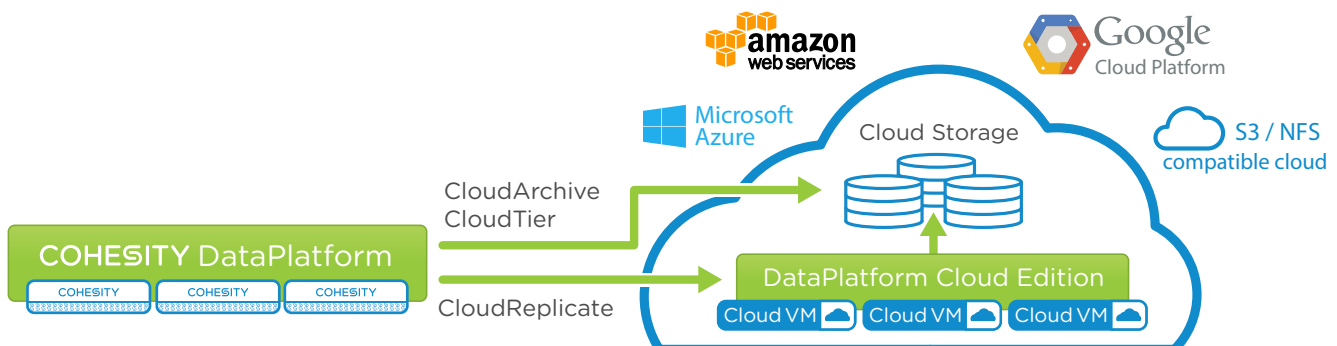
Cloud Archival, Tiering and Replication

Cohesity provides native integration with all the leading public cloud providers including Google Cloud Storage Nearline, Microsoft Azure, Amazon S3 and Glacier. Traditional NAS appliances don't provide this integration natively, and a separate storage gateway must be bolted on to leverage the public cloud. Cohesity can leverage the cloud for the following use cases:

CloudArchive: Replace tape rotations and vaulting with long-term archival to the cloud. Enable Google-like search on your local Cohesity cluster to quickly identify data archived in the cloud. Recover to the original cluster or to a new cluster.

CloudTier: Leverage compelling cloud economics with native cloud integration for data tiering. Colder data is automatically tiered to the cloud for capacity bursting to the cloud.

CloudReplicate: Replicate your data to a Cohesity cluster running in the cloud for disaster recovery, test/dev and analytics.



Advantages of Cohesity for File and Object Consolidation

Feature	Benefit
NFSv3, CIFS, SMB2.x, SMB 3.0, and S3 APIs	Multi-protocol access to data allows support of applications across all major enterprise operating systems including Microsoft Windows, Linux and S3 API
Web-Scale File System	Unlimited scalability, always-on availability, non-disruptive upgrades, pay-as-you-grow scalability
Converged on Cohesity DataPlatform	Store files on the same platform used for data protection, files, objects, and test/dev copies
Global deduplication and compression	Global deduplication and compression across all nodes of the cluster significantly reduces physical storage footprint
Global indexing and search	File and object metadata is indexed upon ingest, enabling Google-like search across all files in a cluster
SnapTree® snapshots and clones	Limitless and fully-hydrated snapshots for granular Cohesity Views (file systems) as well as writeable snapshot clones that provide instant creation, testing and development of view-based datasets
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.
Quotas	Easily establish user and file system quotas with audit logs
Erasur coding	Data is protected against any individual node failure with erasure coding across nodes
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

Summary

Cohesity's NFS, SMB, and S3 object storage provides advanced, enterprise-class data storage management that leverages Cohesity's DataPlatform to provide highly available, secure and efficient hyperconverged secondary storage.