

Cohesity SmartFiles: Unified File and Object Services for the Hybrid Cloud



Key Benefits

- Manage unstructured data on-premises, at the edge, or in the cloud.
- Software defined flexibility.
- Simple, efficient, powerful data management and insights.
- Massively scalable platform for data consolidation.
- Lower TCO with data consolidation and dramatic data reduction up to 96x or more.
- Multilayered cyber resiliency and ransomware protection to defend data and applications.
- Comply with regulatory mandates for data privacy and retention.

Unstructured Data Growth Challenges

Corporate data is growing at a rate as high as 60% per year, most of which is unstructured in the form of file shares, backups, archives, logs, media files, dev/test and analytics. As much as 90% of that data will be infrequently or never accessed. Traditional network attached storage (NAS) and object storage solutions were designed over a decade ago for traditional use cases, and without serious architectural consideration for hybrid cloud IT environments. These file and object silos create multiple challenges related to scalability, data management and security that drain precious staff resources, puts sensitive data a risk for exposure, and hinders the pace of business acceleration and innovation. Related costs also often destroy storage budgets. That's why we created Cohesity SmartFiles: to go beyond traditional file and object platforms and address escalating and evolving data management pain with reliable and secure on-premises, hybrid, and multicloud solutions.

What is Cohesity SmartFiles

Cohesity SmartFiles modernizes and simplifies data and application management by providing a modern economical converged target for consolidating data silos and securely managing unstructured content and application data. It is an enterprise-class, software-defined, datacentric, multiprotocol file and object solution that transcends traditional offerings in terms of manageability, scale, cybersecurity, efficiency, and multi-tiered data management. Integrated data services, Marketplace apps, and machine-learning assisted management uniquely provide intelligence, enhanced threat mitigation, and ease-of-administration that separate SmartFiles from the rest. SmartFiles is based on the Cohesity Data Cloud platform, which provides a patented distributed architecture for scale-out storage in a pay-as-you-grow model. To maximize availability, Cohesity's unique distributed file system provides resiliency and consistency while eliminating disruptive upgrades.

"

"We provide the back end for leading companies all over the world, and as an example, just one of our customers has more than 2.5 billion files, so it was imperative to enable a scale-out, software-defined file and object solution across our environment."

JOHN WHITE, chief innovation officer, Expedient

Why Is SmartFiles Smart?

SmartFiles provides modern data management that uniquely combines intelligent data management, hybrid cloud flexibility, and active threat defense, along with multiple features and options to extract value from your data. Unlike traditional NAS and object storage, the Cohesity Data Cloud platform with SmartFiles is data-centric, not storage or infrastructure centric. This means there is no need to move data to apps and add-on data services as we deliver them in the same platform. This removes the complexity and cost of integrating multiple ecosystem apps to enable file and object data management. Cohesity Marketplace apps and machine-learning assisted management also provide intelligence and ease-of management.

The bottom line is that SmartFiles is smart based on its deployment flexibility, broad compatibility, stress-free management, integrated security, visibility and control of unstructured data across the entire geographic enterprise, and lower storage cost.

Freedom of Choice

Software-defined means customers have the freedom to choose the most appropriate hardware for any file or object workload. It's the freedom to match capacity, performance, and cost with workload requirements. But this is not just about today, it's about tomorrow. SmartFiles future-proofs unstructured data workloads on the assurance that workloads are not bound to proprietary hardware—or even the data center. SmartFiles runs at the edge, in the data center, and in the cloud. Unified NFS and SMB protocols support file management on-premises and the cloud. SmartFiles is also an easy on-ramp to the cloud through its native S3 protocol.

Broad Compatibility

SmartFiles has you covered if you have a mixed OS environment or hybrid cloud environment. File permissions across NFS and SMB environments are mapped for transparent file access from either protocol. Files and objects can be accessed simultaneously, and all data is always available as S3. Hybrid environments are easier as SmartFiles speaks the language of Windows/Linux/Unix and the cloud. Seamless data interoperability helps make complex hybrid and multicloud environments easy.

Stress-Free Operation

It begins with an intuitive UI and “click-and-go” ease. It ends with a bit more time to focus on strategic matters—or just extra time to relax. Machine-learning driven management acts as a hidden smart assistant to remove complexity and save time. Data migrations are easy using Cohesity's backup solution—which removes the complexity of other migration tools. Automate operations through

rich APIs. Cold data on costly Tier 1 storage automatically moves itself to a cost-optimized tier or the cloud. Utilize policies to ensure consistent management. Manage across multiple clusters and locations from a single management pane. Cloud or on-premises. Core or edge. Physical or virtual. Data movement is always easy and transparent to applications. Cohesity SmartFiles alleviates the stress and complexity of traditional scale-out NAS.

Integrated Apps

Integrate an entire ecosystem on a single platform with apps from the Cohesity Marketplace. That means anti-virus protection, file audit, and content search all run on the Cohesity Data Cloud platform with no additional infrastructure required. Detect anomalous file accesses from file logs. Index and search content for compliance and security reasons. The SmartFiles solution eliminates the cost of infrastructure for an application ecosystem while removing the management effort required to manage ecosystem hardware.

Eliminate Security Risks

SmartFiles integrated cybersecurity makes security, including ransomware protection, easy to deploy and manage. It's a multi-layered security approach to protect, detect, and recover your data:

- **Prevent** – Immutable file system, software encryption, over the wire encryption, multi-factor and quorum authentication, DataLock (WORM), and adherence to multiple compliance standards.
- **Detect** – Detect infected files and prevent files from being compromised. Observe user and behavior analytics for security and compliance reasons.
- **Recover** – Rapidly clone your last good snapshot and connect to it for instant access to data.

“Google-like” Enterprise Search

It's not just simple, it's deep and data-centric. Get fast search results across file metadata and file content as both are indexed. Unlike traditional solutions, use results to drive data-driven services and make decisions on actual data contents, and not just data about the data. Data is no longer dark. Rather, it's a renewable resource you need for insights and smart operations. Conduct a single simple search across the entire enterprise. Search any data—anywhere. Across silos, VMs, backups, data centers, remote sites, and multiple clouds.

Low storage cost


Being software-defined, SmartFiles customers have a variety of hardware price points to choose from. SmartFiles provides as much as 96x data reduction or more through deduplication across data center volumes, compression, and small file optimization.

Cost of ownership can be reduced by as much as 51% through consolidation of silos and moving data from Tier 1 storage to a cost-optimized Cohesity tier. Let's explore this technology.







- **Data deduplication** – Based on advanced sliding window variable dedupe technology. This means that duplicate data that has been slightly changed is still deduped—which is beyond traditional dedupe capabilities.
- **Compression** – Zstandard compression algorithm for high compression ratios.
- **Small file optimization** – Eliminated the small file amplification penalty common to traditional NAS file systems. Small files are stored efficiently and without the need for triple-mirroring inefficiency. Put simply, 50TB of small files under SmartFiles consumes about 50TB of capacity.

- **Cross-volume dedupe** – Unlike many traditional NAS appliances, data is deduped across storage volumes and application silos across the data center. Duplicate data that can't be reduced by traditional silo-by-silo dedupe is eliminated by the Cohesity SmartFiles.
- **Data Insights** – SmartFiles provides the built-in capability to scan third-party NAS solutions to analyze file access trends. A policy can be created to transparently tier or migrate older inactive data from third-party and generic NAS solutions, including NetApp, Dell-EMC Isilon, and Pure Storage, to Cohesity. Data can be seamlessly accessed on SmartFiles while NAS storage costs are instantly and drastically reduced.

There are dozens of workloads and use cases in which SmartFiles provides benefit. Amongst these use cases are:

 <p>Content Management</p> <ul style="list-style-type: none"> • Secure file repository • Scalable libraries • File and Object 	 <p>Splunk Cold Buckets</p> <ul style="list-style-type: none"> • Cost-effective cold archiving • on-premises, cloud or both • Frozen data to cloud 	 <p>Corporate Video</p> <ul style="list-style-type: none"> • Corporate communication • In-house productions • Education and training 	 <p>Digital Archives</p> <ul style="list-style-type: none"> • Digital documents • Audio and video files • Active and passive archives • Backup datasets 	 <p>Video Surveillance</p> <ul style="list-style-type: none"> • Security cameras • Body and mobile cams • Rapid retrieval from large amounts of data 	 <p>Backup</p> <ul style="list-style-type: none"> • Databases backup • Target for third party backup applications
--	---	---	---	---	---

There are many industries that depend on Cohesity for file and object services. These industries include:

 <p>Financial Services</p> <ul style="list-style-type: none"> • Images and documents • Call logs and correspondence 	 <p>Law Firms</p> <ul style="list-style-type: none"> • Documents and research • Collaboration • eDiscovery and legal hold 	 <p>State, Local and Education</p> <ul style="list-style-type: none"> • Collaboration and research • Safety and security • Innovation and technology 	 <p>Healthcare</p> <ul style="list-style-type: none"> • PACS archives • VNA archives 	 <p>Life Sciences and Medical Research</p> <ul style="list-style-type: none"> • Genomic data archive • Security • Compliance 	 <p>Manufacturing and Technology</p> <ul style="list-style-type: none"> • Design files, images, documents • IOT data • Security and compliance
---	--	---	--	---	---

Feature	Description
Data management platform	Single platform for data protection and security, files, objects, dev/test, analytics, and cloud integration.
API-first architecture design	Maximum flexibility and ease of automation. Cohesity integrates with leading third-party apps and solutions you already use. Easily integrate into your existing on-prem environment and create bonds instead of divisions between operational teams.
NFS3, CIFS, SMB 2.x, SMB 3.0, S3, and OpenStack Swift APIs	Simultaneous multiprotocol access to the same data with unified permissions. Supports NFS, SMB, and S3 APIs.
MMC integration	Snap-in for the Microsoft management console, which allows Cohesity file shares to be managed by the MMC.
Fully distributed shared-nothing file system	Limitless scalability, always-on availability, non-disruptive upgrades, pay-as-you-grow model.
Global deduplication and compression	Unparalleled storage efficiency with global deduplication, compression, and small file capacity optimization across all cluster nodes to significantly reduce data center cost and footprint.
Quality of service (QoS)	QoS policies are provided that optimize performance for different types of workloads.
Global indexing and search	File and object metadata is indexed upon ingest, enabling Google-like search across all files in a cluster. Indexed search of file contents through Insight integrated Marketplace app. Search across geographic locations and clouds via Cohesity Helios UI.
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.
Multilayered authentication	Multi-factor authentication. Quorum authentication provides an extra layer of protection by requiring approval by a predefined group of administrators for all sensitive and privileged operations.
Encryption and key management	Support for data-at-rest as well as data-in-flight encryption using the industry standard 256-bit Advanced Encryption Standard (AES) algorithm. The platform is also FIPS 140-1 and 140-2 compliant. Use the built-in key management system, or a third-party solution.
Write Once Read Many (WORM) immutability	Enables long-term retention of data that has compliance controls mandating a policy that objects cannot be modified during the lock time.
Compliance certifications	SEC, FINRA, Common Criteria, PCI-DSS
Quotas	Easily establish user and file system quotas with audit logs.
Public cloud integration	Native 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
Cloud use case flexibility (CloudArchive, CloudTier, CloudReplicate)	Archive into leading 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 dev/test.
Strict consistency	Guaranteed consistent data regardless of 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 data sets.

Feature	Description
Policy-based backup protection	Integrated data protection software and SnapTree technology is available to allow simplified data protection of files and objects with fully-hydrated snapshots.
Replication for disaster recovery	Built-in, granular, and secure replication services for geo redundancy.

Learn more at Cohesity.com/products/smartfiles



© 2023 Cohesity, Inc. All rights reserved.

Cohesity, the Cohesity logo, SnapTree, SpanFS, DataPlatform, DataProtect, Helios, the Helios logo, DataGovern, SiteContinuity, and other Cohesity marks are trademarks or registered trademarks of Cohesity, Inc. in the US and/or internationally. Other company and product names may be trademarks of the respective companies with which they are associated. This material (a) is intended to provide you information about Cohesity and our business and products; (b) was believed to be true and accurate at the time it was written, but is subject to change without notice; and (c) is provided on an "AS IS" basis. Cohesity disclaims all express or implied conditions, representations, warranties of any kind.