COHESITY

Version 1.1

April 2024

Cohesity Storage Efficiency

Maximize Your Storage Capacity with Cohesity

ABSTRACT

The bulk of data being generated today by organizations is unstructured data. The unstructured data is primarily stored on NAS file systems. Over the years, tools such as deduplication and compression have helped organizations consume less space. Cohesity builds on these tools and implements the technology differently than traditional NAS vendors. As a result, the same amount of data requires less space to be stored on Cohesity. The storage-efficiency tools in Cohesity, combined with its snapshot and cloning capabilities, allow organizations to dramatically reduce their storage footprint.

Table of Contents

Figures

Introduction to Storage Efficiency

Given Gartner's prediction of 800% growth in enterprise data within the next five years, storage space efficiency capabilities like deduplication and compression become more critical to enterprise storage platforms every day. To deal with the expected data growth, enterprises need storage products capable of delivering the highest level of space efficiency — at optimal cost.

In addition to space efficiency, an important aspect of storage efficiency is time efficiency. Cohesity enables customers to efficiently protect data using snapshots while also leveraging those snapshots with cloning, producing significant time savings.

With inline deduplication, built-in compression, small file efficiency, and other key storage technologies, Cohesity enables organizations to deliver optimal storage efficiency and reduced infrastructure costs.

Global Space Efficiency

The Cohesity portfolio of storage efficiency capabilities consist of several technologies that can work together or individually to achieve maximum space efficiency, including:

- Deduplication
- Compression
- Small file storage efficiency
- Cohesity file system storage efficiency
- Cohesity file system and metadata resiliency

Deduplication — Fixed-length vs Cohesity's Variable-length

Cohesity uses a unique, variable-length data deduplication technology that spans an entire cluster, generating significant savings across the entire storage footprint. With variable-length deduplication, the block size is not fixed. Instead, our algorithm divides the data into chunks of varying sizes to achieve the maximum possible matching of duplicate data chunks.

With conventional fixed-length deduplication:

- Markers are placed at fixed-length intervals.
- Small file changes result in loss of efficiency.

With Cohesity's variable-length deduplication:

- Only new or changed data is stored and the remainder of the file is not affected.
- Data is divided into blocks based on the characteristics of the data itself and not an arbitrary block size.

Figure 1: Fixed- vs Variable-length Deduplication Example

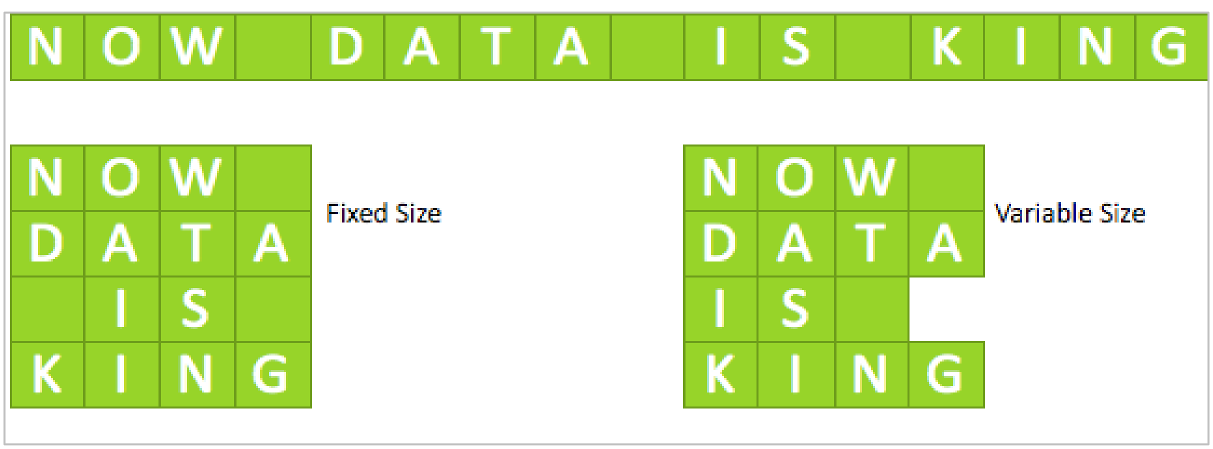

What happens when the data changes? In our example, what happens when 'NOW DATA IS KING' changes by a single character, like 'KNOW DATA IS KING'?

With fixed-block deduplication, unless the changes are exactly a multiple of the fixed-block size, all of the data past the first change is shifted. The shift changes subsequent blocks in the file with respect to the fixed-block boundaries. As a result, the changes look new to the fixed-block deduplication engine.

Variable-block deduplication divides data into blocks based on the characteristics of the data itself, not an arbitrary block size. As a result, this makes variable-block deduplication flexible when data changes. Only the new or changed data is stored and the remainder of the file is not affected, which allows for near-maximum data storage efficiency.

Figure 2: Impact of Changes on Deduplicated Data

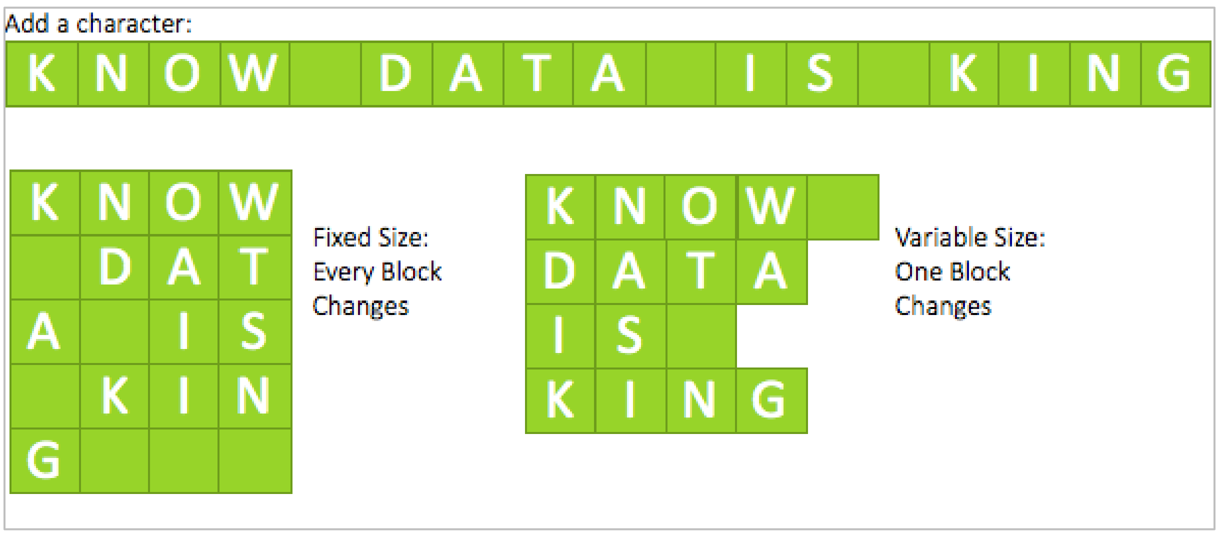

NOTE: The benefits of Cohesity's deduplication technology apply across all of Cohesity's replication and cloud-storage features: Replication, CloudArchive, and Cloud Tier.

Compression

Cohesity uses the Zstandard (zstd) lossless data compression algorithm. Zstandard offers high compression along with high performance. Compression reduces the storage capacity you consume by reducing the size of bit or byte strings in a data stream.

NOTE: Cohesity ensures that deduplication always precedes data compression.

Log in to Cohesity's platform to toggle inline compression:

- **On** (default). Compress inline for HDD and flash.
- **Off**. Turns inline compression off on the flash tier. Data in the flash tier is compressed only after it has been down-tiered to HDD.

  NOTE: In a few scenarios — such as data that is already stored in a compressed file format (H.264 videos, MP3 files, or zip files) or data that is encrypted — we recommend you disable Cohesity compression.

NOTE: The benefits of Cohesity's compression technology apply across all of Cohesity's replication and cloud-storage features: Replication, CloudArchive, and Cloud Tier.

Small File Storage Efficiency

Cohesity provides small file efficiency and avoids the small file amplification penalty common to traditional NAS file systems, which traditionally have a fixed block size of either 4K or 8K. On traditional NAS file systems, any files that are smaller than the fixed block size cause extra space to go unused.

Take, for instance, a storage vendor who employs a fixed block size of 8K. If you store a 1K file on the system, the system writes an 8K block, so 7K of space is unused. What's more, to protect that file, it is written multiple times, compounding the inefficient use of that block.

On SpanFS™, Cohesity's file system, any file smaller than 8 megabytes is considered small and stored together with multiple other compressed, small files in a Cohesity minion blob. A Cohesity minion blob is a coalesced disk-based object that is eligible for deduplication and compression.

Figure 3: Small File Storage Efficiency — Traditional NAS vs Cohesity File System

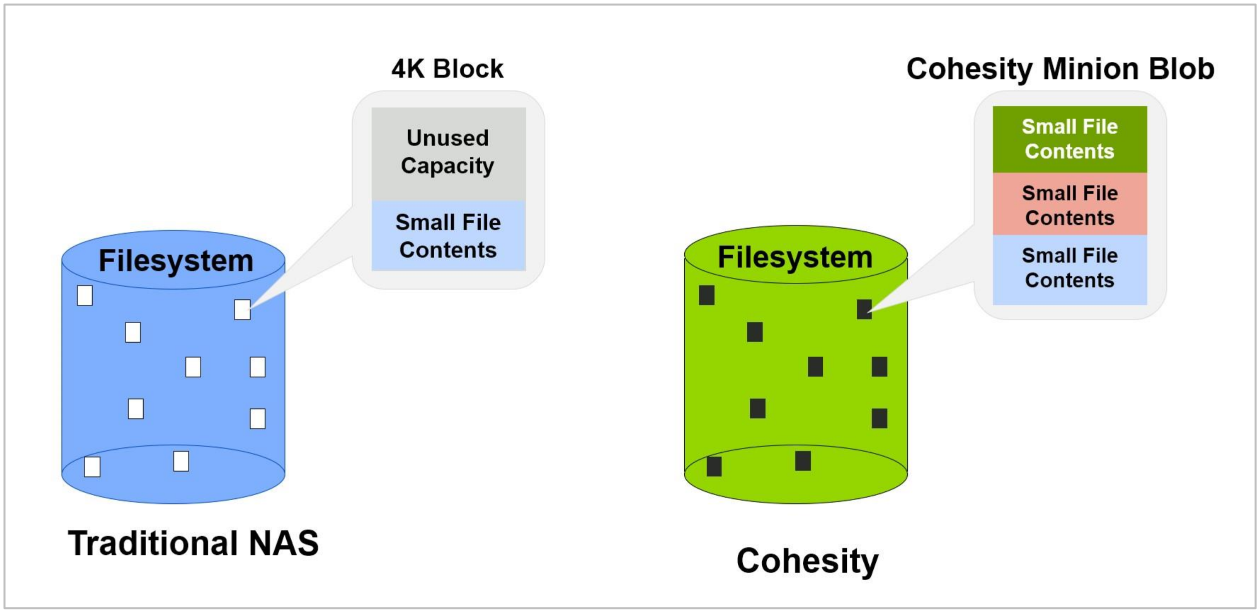

Standard Data Storage Efficiency Ratios

Across the leading vendors in the storage industry, typical storage efficiency ratios vary by workload. A use case of home directories or software development will benefit from deduplication and compression, while, by contrast, video and image files do not benefit much from compression, as the image codecs have already maximized compression of the images. Similarly, encrypted data does not benefit from traditional storage-efficiency technologies.

As demonstrated in Cohesity Storage Efficiency vs Scale-out NAS Vendors below, testing by ESG demonstrates that Cohesity's storage-efficiency technologies outperform these industry numbers by significant margins.

Cohesity File System and Metadata Resiliency

Even while delivering significant storage efficiencies, Cohesity delivers robust file system and metadata resiliency by employing erasure coding or replication factor technologies.

Erasure Coding (EC) refers to a scheme where data is broken into stripes, encoded and stored across available nodes or disks. A single code stripe unit can protect against one data (or code) stripe failure, and two code stripe units can protect against two data (or code) stripe unit failures.

The following diagram illustrates erasure coding 4:2. A erasure coding of 4:2 will have 4 data stripes, identified as D1, D2, D3, D4 and 2 code stripes, identified as C1 and C2. The first code stripe is calculated using an exclusive OR (XOR). The second code stripe is calculated using a large finite field identified as GF(2^n).

Figure 4: Erasure Coding Example of EC Level 4:2

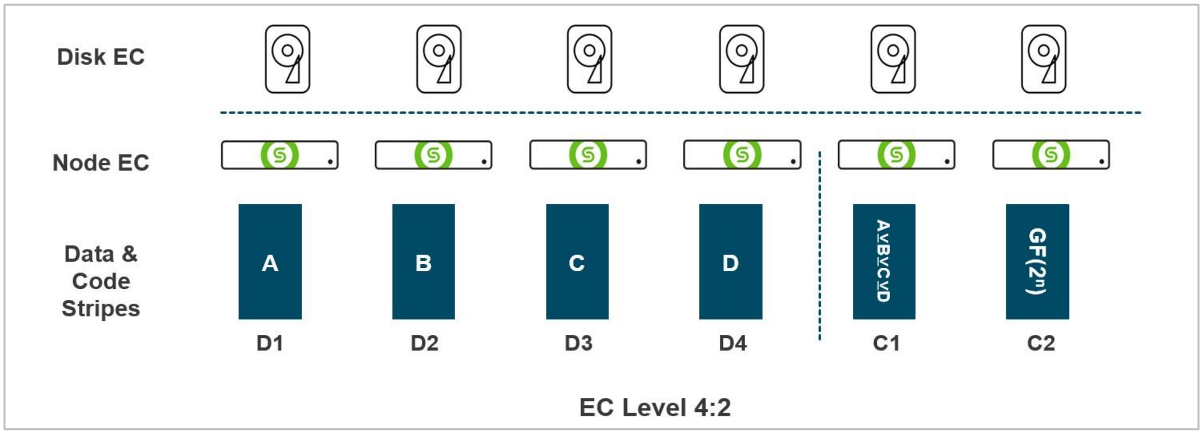

Replication Factor (RF) refers to the number of replicas of a unit of data. The unit of replication is a chunk file, and a chunk file is mirrored into either one or two other nodes depending on the Replication Factor number chosen. An RF2 mechanism provides resilience against a single data unit failure, and a RF3 provides resilience against two data unit failures.

Figure 5: Replication Factor Example

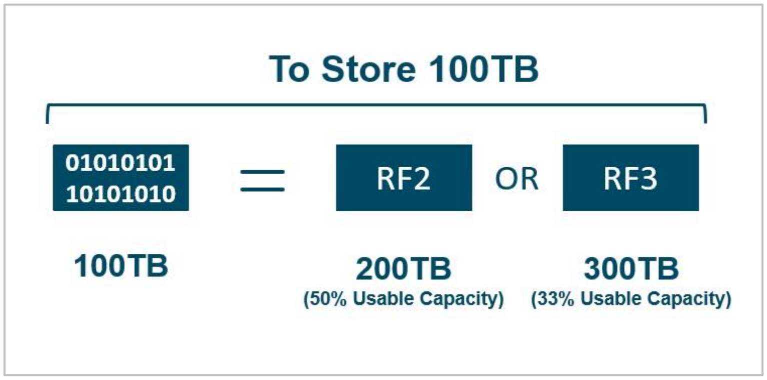

For more on Cohesity's data resiliency, see the *Cohesity Fault Tolerance - Data Integrity for Modern Web-scale Environments* white paper.

Cohesity Storage Efficiency vs Scale-out NAS Vendors

The Enterprise Strategy Group (ESG) examined the data reduction features and small file efficiency of various vendors in the storage domain. ESG compared Cohesity's global variable-length deduplication and small file optimization to a system from another vendor that uses fixed-length deduplication and does not deduplicate globally across Views.

ESG used synthetic datasets generated specifically for these tests:

- To test deduplication, a 1TB unstructured dataset with a duplication factor of 2:1.
- To test small file efficiency, a set of 1 million 1KB files.

Erasure Code (EC) redundancy was set to 2:1 in both systems.

In the first series of tests, the 1TB unstructured dataset was copied to two Views in Cohesity and two Views in the other vendor's system. As illustrated in Figure 6 below, Cohesity's global deduplication reduced the data footprint by 60.7% while the non-global deduplication system was only able to achieve 6.4% data reduction.

Figure 6: Cohesity Variable-length Deduplication vs Fixed-length Deduplication

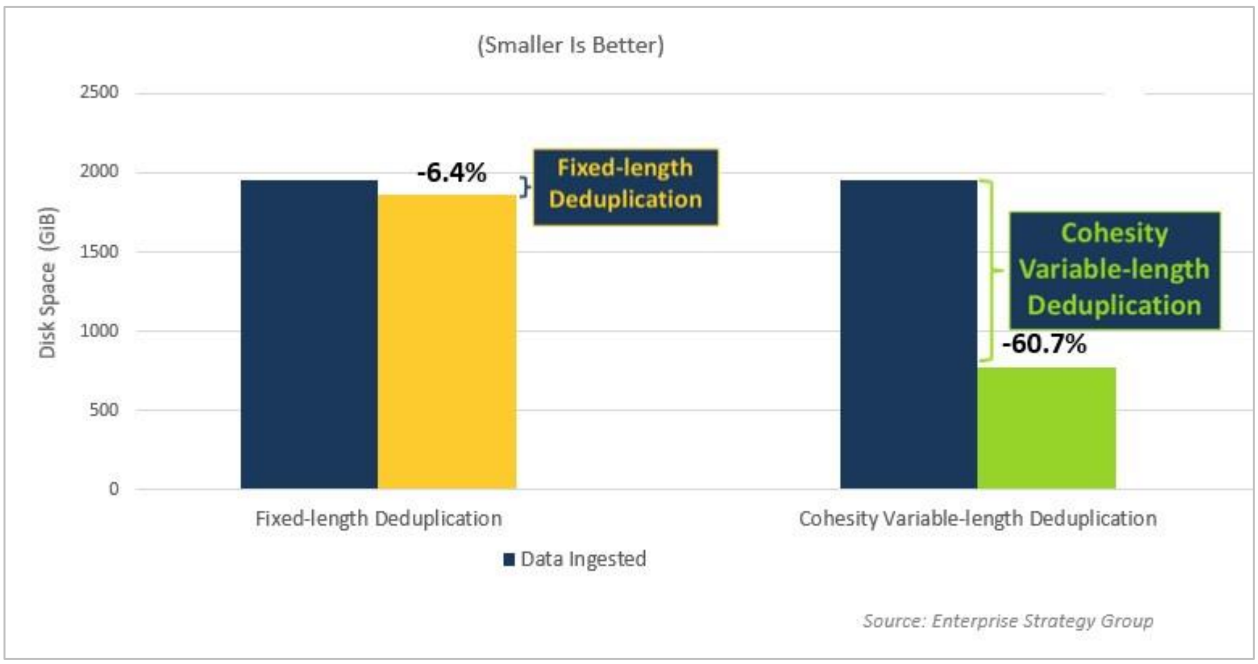

ESG also looked at small file efficiency. ESG copied 1 million 1KB files to a View on Cohesity and a View on the other vendor's system. To isolate the effect of small file optimization, Cohesity disabled deduplication for this test and set EC redundancy to 2:1 in both systems. Due to its 8KB block size and requirement to store multiple copies of the data for redundancy, the other vendor's approach resulted in significant storage amplification and inflated the data to 25.6GB — see Figure 7 below. Cohesity was able to store the data on disk using only 1.5GB of capacity.

NOTE: This test was designed to highlight the small file efficiency of Cohesity and does not, by itself, represent a real-world application. Organizations that have a sizable number of small (sub-8KB) files can benefit the most, depending on the workload and how the files are written.

Figure 7: Small File Efficiency

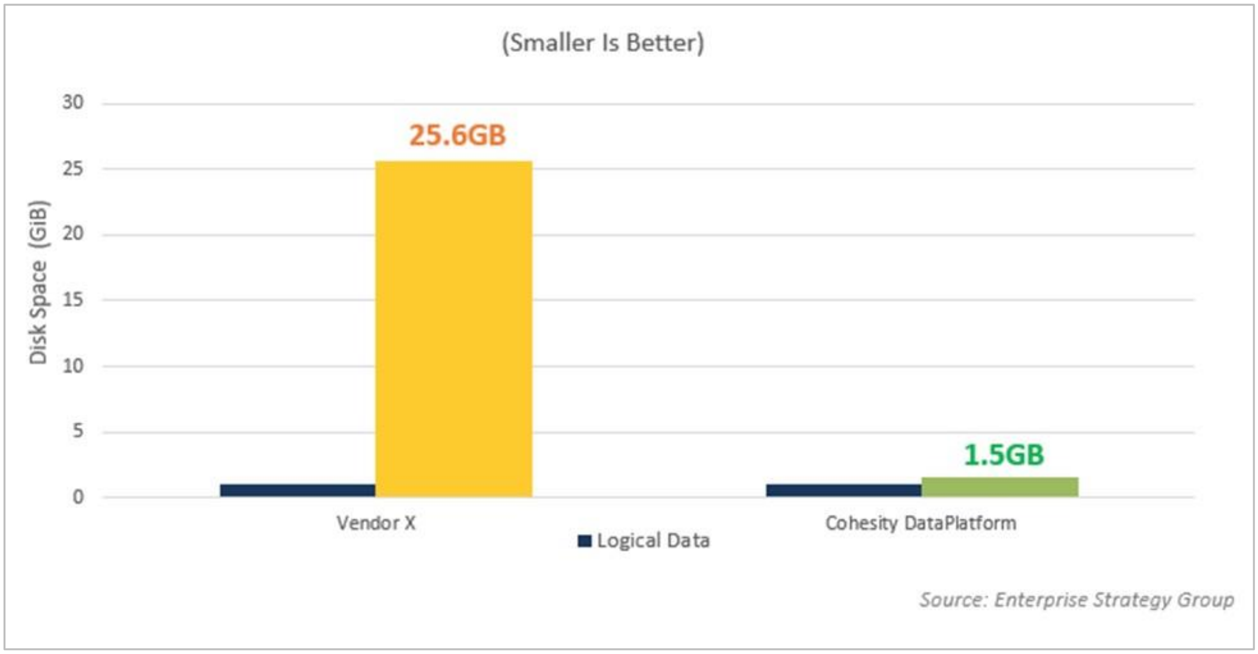

Cohesity Snapshots

Cohesity has brought innovation to data protection through its SnapTree™ technology. SnapTree is a *Distributed Redirect-on-write* (DROW) snapshot mechanism that provides speed and scalability, in addition to the inherent benefits of Redirect-on-write (ROW) snapshot technology. The design is optimized for write performance, so any changes are redirected to new blocks. What's more, all nodes participate in this process, leveraging the scalability of the Cohesity cluster.

Cohesity's SnapTree technology allows businesses to take a large number of snapshots at any time interval with uncapped retention policies, without ever affecting performance or consuming additional space. Cohesity's DROW implementation keeps track of changes by writing the changed data to new blocks.

SnapTree creates a tree of pointers that limits the number of traverses it takes to retrieve blocks of data, regardless of the number of snapshots that have been taken. SnapTree uses a B+ tree data structure (Figure 8 below) such that access to any point in the tree takes a fixed number of snapshots, no matter how many snapshots there are, without having to rebuild chain linkages.

Figure 8: Cohesity SnapTree with Distributed Redirect-on-write Snapshot Technology

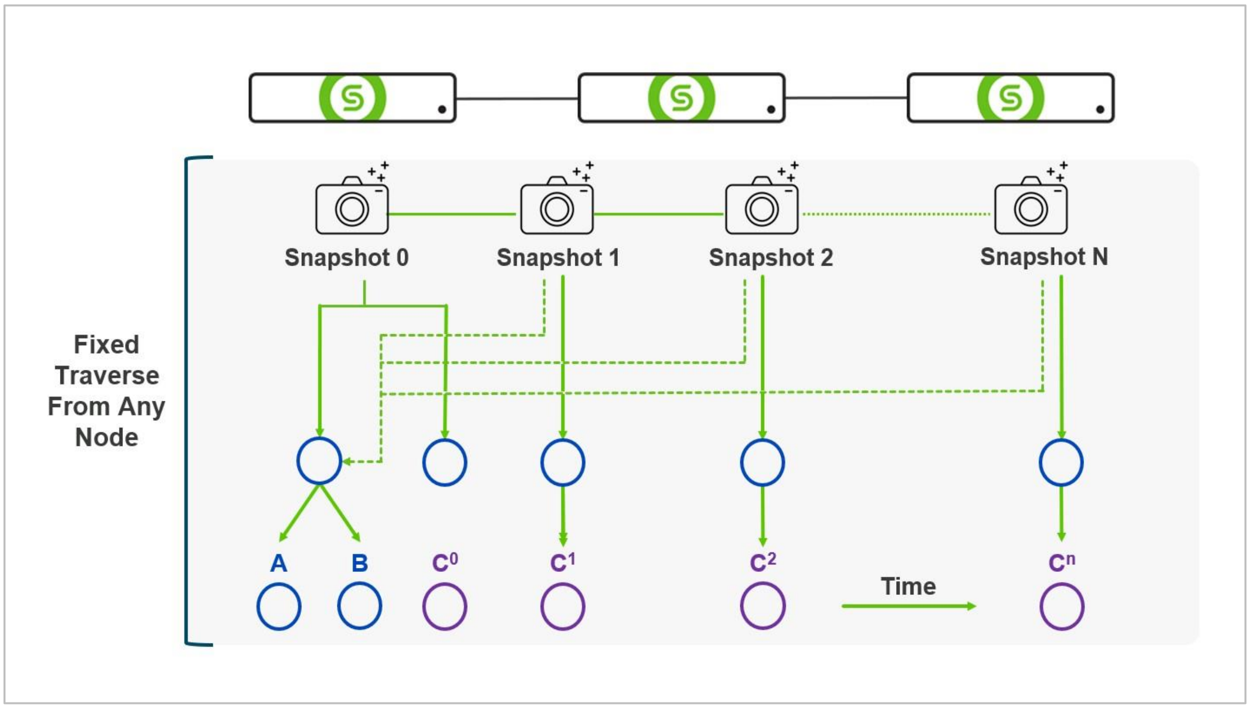

Because SnapTree is implemented on a distributed file system (Figure 9 below), every node sees the same nested structure of the chain with a fixed depth, independent of where the actual data is stored in the cluster. Keeping the snapshots fully hydrated improves the recovery times of any particular snapshot from **t0** to **t*n***, because it does not incur the time penalty of traversing the entire chain of changes. Each of these snapshot clones is fully hydrated so that businesses can achieve fast RTOs (Recovery Time Objectives) and near-continuous RPOs (Recovery Point Objectives).

Figure 9: Cohesity Distributed File System Enables Fast Recovery

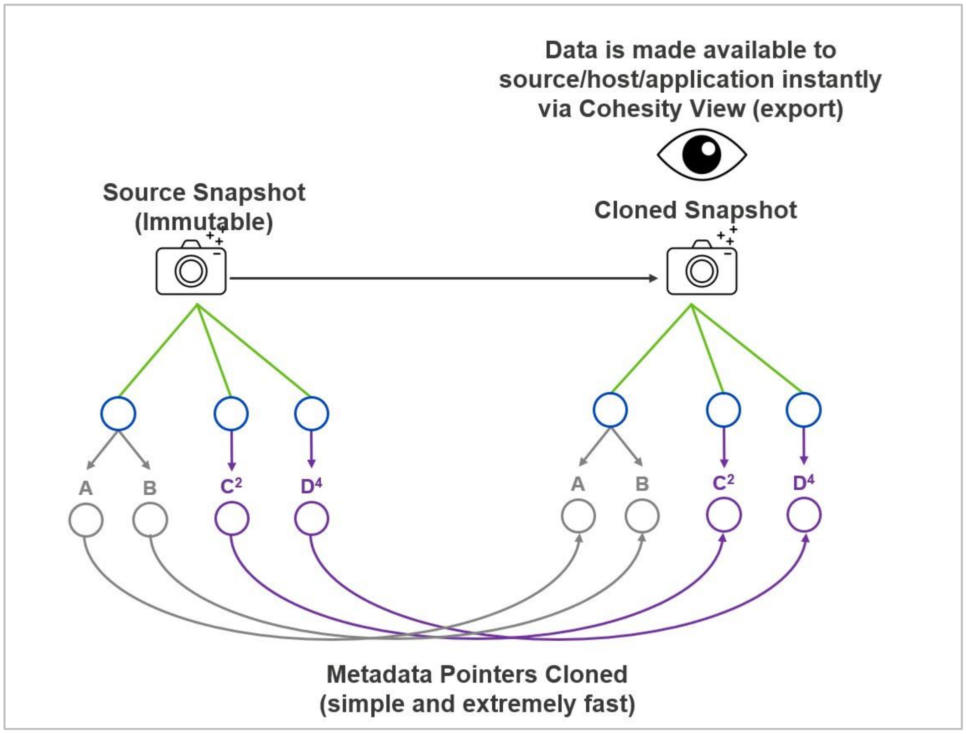

Cohesity Clones

With Cohesity SpanFS, customers can instantly provision clones of backup data, files, objects, or entire Views. They can then present those clones to support a variety of use cases.

In Cohesity, snapshots and clones are very efficient. Snapshots don't consume space until data is modified, and at that point, clones only need to store the changes from the original copy of the data. The clones can be created instantly without having to move data between storage devices. This is in stark contrast to the inefficiency of traditional storage, where full copies of data are created between storage silos, wasting lots of storage capacity, time, and IO bandwidth.

Cohesity clones:

- **Are time efficient.** Traditional copies can take many minutes or hours to make. With Cohesity's clone technology, even the largest Views can be cloned instantly.
- **Are space efficient.** A clone only consumes additional space as data is changed or added.
- **Lower costs.** Cohesity clone technology can dramatically reduce the storage you need for development and testing.
- **Efficiently test your DR environment.** Cohesity clone technology makes it possible to test disaster recovery processes fully without threatening your production environment.

Resources

This paper is based on research and testing by the Enterprise Strategy Group. For their original findings, see their ESG Technical Review: Cohesity Data Management report.

Your Feedback

Was this document helpful? Send us your feedback!

About the Authors

Scott Owens is a Technical Marketing Engineer at Cohesity. In his role, Scott focuses on file services.

Other essential contributors included:

- Karthick Radhakrishnan, Manager, Technical Marketing
- Vibhor Gupta, Product Management, File Services

Document Version History

VERSION	DATE	DOCUMENT HISTORY
1.1	April 2024	Cohesity rebranding updates
1.0	June 2020	First release

ABOUT COHESITY

Cohesity is a leader in AI-powered data security and management. Aided by an extensive ecosystem of partners, Cohesity makes it easier to protect, manage, and get value from data – across the data center, edge, and cloud. Cohesity helps organizations defend against cybersecurity threats with comprehensive data security and management capabilities, including immutable backup snapshots, AI-based threat detection, monitoring for malicious behavior, and rapid recovery at scale. Cohesity solutions are delivered as a service, self-managed, or provided by a Cohesity-powered partner. Cohesity is headquartered in San Jose, CA, and is trusted by the world's largest enterprises, including six of the Fortune 10 and 42 of the Fortune 100.

Visit our website and blog, follow us on Twitter and LinkedIn and like us on Facebook.