

Version 2.1

September 2021

Architecture Matters

Cohesity vs Dell EMC PowerScale

(previously Dell EMC Isilon) NAS Scale-Out Limits

ABSTRACT

This white paper looks at the scale-out capability of Cohesity and the benefits of a scale-out architecture without the design limitations of Dell EMC PowerScale (previously Dell EMC Isilon) scale-out solutions. This is important to consider for consolidation projects, large scale deployments, and for ensuring that an architecture has scale to address future strategic plans.

Table of Contents

Overview	3
Need for Limitless Scaling	4
<i>Video Surveillance</i>	4
<i>Home Directory</i>	4
<i>High File-count Applications</i>	4
SpanFS: File System Designed for Storage at Web Scale	5
Data Management at Scale	7
Tested Limits	8
Appendix A - Configuration Details	9
Document Version History	10
Your Feedback	10

Tables

Table 1: Tested Limits — Cohesity vs Dell EMC PowerScale	8
--	---

Overview

According to projections from the International Data Corporation (IDC), 80 percent of worldwide data will be unstructured by 2025¹. For many companies, the unfettered growth of data has reached critical mass already. Unstructured data creates a unique challenge for organizations wishing to use their information for analysis. The data can't easily be stored in a database, and has attributes that make it a challenge to search for, edit, and analyze. The challenges of managing unstructured data are part of the reason why this is such an important topic.

If your organization is struggling to manage their unstructured data now, the problem is going to get exponentially harder over time. Unstructured data comes in the form of file shares, backups, archives, logs, media files, IoT sensor data and more. Traditional network-attached, or scale-up, storage was designed more than 10 years ago for some of these use cases. As a result, existing NAS storage is optimized for only a specific subset of unstructured data workloads and has limited scale. This results in multiple storage silos, often running on proprietary hardware, with different software and licensing models.

With the advancement in applications, traditional NAS platforms are not able to keep up with the needs of ever-growing data. Although some of the traditional NAS platforms can scale-out in capacity, they are still limited by their architecture. [Cohesity SpanFS®](#) is designed for web-scale architecture, where all the components within Cohesity scale linearly. Cohesity SmartFiles is architected to meet your current and future data needs.

¹ Seagate (November 2018). The Digitization of the World from Edge to Core
<https://www.seagate.com/www-content/our-story/trends/files/idc-seagate-dataage-whitepaper.pdf>

Need for Limitless Scaling

There are countless use cases where arbitrary limits affect the design and capabilities of storage systems. There is also a need to provide for unknown scalability requirements in the future, so having greater-than-necessary scale today provides future-proofing for tomorrow. Three such cases are [video surveillance](#), [home directories](#), and high file-count applications.

Video Surveillance

Companies rely on 24x7 video streams with higher and higher resolution and more frames per second to better monitor their business. Coupled with broader camera deployment, the enhanced video quality of such streams gives security personnel a more detailed, comprehensive view of the areas under surveillance.

The higher resolution and frame rates of such 24x7 video streams can significantly increase the capacity requirements. Just one continuous video stream at 30fps/1920x1080 image resolution can fill a 16 TiB file in only 16 days with H.264 compression. On Dell EMC PowerScale the max file size is 16 TiB² regardless of the size of the cluster, so one would need to account for this file limitation as part of the design.

Home Directory

Dell EMC PowerScale has a limit of 80,000 SMB shares per cluster³. In an environment with a large number of home directories, such as a university or hospital, a new cluster will need to be deployed when nearing the 80,000 SMB share limit. The new Dell EMC PowerScale cluster will add additional management overhead, as well as contribute to mass data fragmentation. In addition, the new cluster will be unable to take advantage of any storage efficiency associated with data stored on the original cluster.

High File-count Applications

Companies have a variety of applications that result in a large number of files being stored in a single directory. For instance, an engineering group could set up an NFS export for scratch space for shared development efforts. Engineers would leverage the export to store temporary files. On Dell EMC PowerScale's solution, the best practice is to not have more than 1M files in a single directory. As a result, additional management would be required to watch the shared scratch space and to create jobs to constantly ensure the number of files stays below the limit.

² Dell EMC, (June 2020). Dell EMC PowerScale OneFS: A Technical Overview, 15, Retrieved from <https://www.dell EMC.com/en-me/collaterals/unauth/technical-guides-support-information/products/storage/h10719-isilon-onefs-technical-overview-wp.pdf>

³ Dell EMC, (May 2019). Dell EMC Isilon OneFS: Technical Specifications Guide, 6, Retrieved from <https://www.dell EMC.com/ja-jp/collaterals/unauth/technical-guides-support-information/products/storage/docu94152.pdf>

SpanFS: File System Designed for Storage at Web Scale

SpanFS provides the foundation for Cohesity's SmartFiles NAS and Object features. Legacy storage has evolved in two very different directions. Traditional enterprise storage focused on providing standardized file interfaces (NFS and SMB) on 'scale-up' hardware, and snapshots for resiliency. And cloud storage, developed by hyperscale companies like Google and Amazon, focused on delivering scale-out solutions on commodity hardware, with strong resilience to hardware failures, but relying on proprietary protocols and APIs for data access.

Span Everything

To enable enterprises to take back control of their data at scale, Cohesity has built a completely new file system: SpanFS. SpanFS is designed to effectively consolidate and manage all data, including backups, files, objects, test/dev, and analytics data, on a web-scale platform that spans from the edge to the cloud.

SpanFS is designed to span everything, including physical and logical constructs:

- Physical:
 - **Nodes:** SpanFS provides unlimited scale across many hyperconverged nodes. SpanFS is completely distributed and doesn't have a master node. It scales linearly and dynamically rebalances data as nodes are added or removed. It provides always-on availability, non-disruptive upgrades, and a pay-as-you-grow consumption model.
 - **Storage Tiers:** SpanFS spans across tiers of persistent storage technologies, including SSD, HDD, and even remote cloud storage. SpanFS effectively utilizes the most appropriate tier, based on IO profiles.
 - **Geographic Locations and Cloud:** SpanFS interconnects remote offices, branch offices, core data centers, and public clouds into one data fabric. SpanFS can be used to build a multi-cloud data fabric and span data across multiple private and public clouds.
- Logical:
 - **Workloads:** SpanFS supports data protection, files, objects, test/dev copies, and analytics data. It supports all the key capabilities required by these use cases, including globally distributed NFS, SMB, and S3 storage, unlimited snapshots, global deduplication, encryption, replication, global indexing and search, and good performance, for both sequential and random IO.
 - **Namespaces:** SpanFS can divide physical storage pools into separate shared namespaces (or Storage Domains) that have common data reduction, availability, and/or archive policies.
 - **Tenants:** SpanFS supports multiple tenants (or partitions) with strong QoS capabilities, data isolation between tenants, separate encryption keys, and role-based access control.

SnapTree®: In legacy storage, snapshots form a linked chain, with each link containing the changes from the prior snapshot. Every time a new snapshot is done, an additional link is added to the chain. As the chain grows, the performance overhead required to access the data increases proportionally because the file system must traverse the chain to access the data. Hence snapshots introduce performance overhead and are limited in scale.

[SnapTree](#) introduces a completely new approach to managing metadata at scale and enables SpanFS to provide unlimited snapshots with no performance overhead. SnapTree is based on a B+ tree metadata structure but adds multiple innovations, including:

- Distributes the tree across nodes.
- Provides concurrent access from multiple nodes.
- Supports the creation of instantaneous clones and snapshots.
- Garbage collects unreferenced nodes in the background, using Map-Reduce.
- Ensures consistent access performance regardless of the number of snapshots and clones.
- Stores only one value per leaf node, as opposed to multiple values in traditional B+ trees. This avoids unnecessary snapshotting of multiple values.
- Supports a variable fan-out factor that increases further down the tree. This avoids making any given sub-part of the tree too hot, while at the same time keeping tree-balancing costs low

With SnapTree, Views (volumes) and files are represented by a tree of pointers to the underlying data. The root node represents the View or individual files. The root node points to some intermediary nodes, which in turn point to the leaf nodes, which contain the location of the chunks of data. Customers can take snapshots of entire Views (volumes) or individual files within the Views. As snapshots are taken, the number of hops from the root to the leaves does not increase. **Unlike legacy NAS vendors, customers can take snapshots as frequently as desired – without ever experiencing performance degradations.**

While many traditional NAS vendors claim to be scale-out, they have limitations that contradict what it means to provide limitless scale-out.

Data Management at Scale

Cohesity was created to deliver value beyond scale-out NAS in many ways:

- **Scale without boundaries.** Unlimited everything: resource pools, snapshots, capacity, open and active connections, file size/name/path length, and performance linearity at scale.
- **Storage efficiency.** Sliding-window variable deduplication, compression, small file capacity optimization, and data deduplication across multiple storage silos.
- **Multi-tiered data management.** Transparent data tiering — on-premises to the cloud. Also, tiering from other vendor storage, too, and to the cloud.
- **Enterprise search.** A simple “Google-like” search across VMs, backups, data centers, clouds, and multiple sites.
- **Enterprise richness.** NFS, SMB, and S3 with unified permissions – same data accessed by all protocols, writable snapshots, API-driven, failover/failback, DR (disaster recovery) automation, and WORM ([write once read many](#)).
- **Powerful cybersecurity.** Software-based encryption, multi-factor authentication, FIPS 140-1 and 140-2, integrated anti-virus protection, and audit log analysis with anomalous access detection.

With this data-centric storage platform, IT teams are no longer required to deploy costly external infrastructure for anti-virus (AV) protection, as this is native within the Cohesity. There is also no additional management burden as the integrated AV environment automatically grows with the Cohesity platform.

Cohesity is the future of enterprise file and object services. It's not just the ability to run anti-virus, file audit, and content search on the data platform – it's also boundless scale, leading capacity efficiency, robust cyber-security, and transparent heterogeneous data tiering that transcends traditional scale-out NAS appliances.

Tested Limits

The following table compares the tested limits between Cohesity and Dell EMC PowerScale.

Table 1: Tested Limits — Cohesity vs Dell EMC PowerScale

ITEM	COHESITY	DELL EMC POWERSCALE	LIMITATION IMPACT
Number of Files Per Directory	3.2 Million	1 Million ⁴	Dell EMC PowerScale states a tested limit of 1 million files per directory.
Max Path Length	131,092 B	1,023 B ⁵	Dell EMC PowerScale limits the path length to 1023 bytes. In the case of multi-byte character sets, the path limit is reduced. This limit in bytes forces customers to consider directory structure.
Max File Size (non-Sparse)	70 TiB	16 TiB ⁶	On Cohesity, the max file size is based on the physical size of the cluster. In this test, the cluster had 4 nodes. If the cluster had had 8 nodes, the max file size would scale to 140 TiB.
Max Number of Views/Shares	800,000	80,000 (SMB) 40,000 (NFS) ⁷	Customer environments that require provisioning a high number of SMB shares or NFS exports would need to deploy multiple Dell EMC PowerScale clusters. Multiple clusters create silos of data, resulting in lower storage efficiency and a higher TCO.

NOTE: These tests were run on a Cohesity cluster with 4 nodes. On Cohesity, the max file size is only limited by the physical size of the cluster.

⁴ Dell EMC, (May 2019). Dell EMC Isilon OneFS: Technical Specifications Guide, 10, Retrieved from <https://www.dell EMC.com/ja-jp/collaterals/unauth/technical-guides-support-information/products/storage/docu94152.pdf>

⁵ Dell EMC, (May 2019). Dell EMC Isilon OneFS: Technical Specifications Guide, 11, Retrieved from <https://www.dell EMC.com/ja-jp/collaterals/unauth/technical-guides-support-information/products/storage/docu94152.pdf>

⁶ Dell EMC, (June 2020). Dell EMC PowerScale OneFS: A Technical Overview, 15, Retrieved from <https://www.dell EMC.com/en-me/collaterals/unauth/technical-guides-support-information/products/storage/h10719-isilon-onefs-technical-overview-wp.pdf>

⁷ Dell EMC, (May 2019). Dell EMC Isilon OneFS: Technical Specifications Guide, 4,6, Retrieved from <https://www.dell EMC.com/ja-jp/collaterals/unauth/technical-guides-support-information/products/storage/docu94152.pdf>

Appendix A - Configuration Details

Provided below are the specific components and configuration details of the Cohesity test environment.

Storage Hardware:

- Cohesity C2500: 4 nodes

Storage Software:

- Cohesity 6.1

Storage Configuration:

- Redundancy: Erasure Coding 2:1
- Deduplication: Post: On; Inline On
- Compression: Post: On; Inline On
- Encryption: Off
- Network: 10GbE per node

Your Feedback

Was this document helpful? [Send us your feedback!](#)

Document Version History

VERSION	DATE	DOCUMENT HISTORY
1.0	Oct 2019	Original document
2.0	Sept 2020	Updated document
2.1	Sept 2021	Cohesity rebranding updates

ABOUT COHESITY

[Cohesity](#) radically simplifies data management. We make it easy to protect, manage, and derive value from data — across the data center, edge, and cloud. We offer a full suite of services consolidated on one multicloud data platform: backup and recovery, disaster recovery, file and object services, dev/test, and data compliance, security, and analytics — reducing complexity and eliminating [mass data fragmentation](#). Cohesity can be delivered as a service, self-managed, or provided by a Cohesity-powered partner.

Visit our [website](#) and [blog](#), follow us on [Twitter](#) and [LinkedIn](#) and like us on [Facebook](#).

© 2021. Cohesity, Inc. All Rights Reserved.

Cohesity, the Cohesity logo, SnapTree, SpanFS, DataProtect, Helios, 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.

2000025-004-EN