

SFD18 - Cohesity

*Chris Colotti, Principal Technologist
300% Revenue YoY growth, 70% enterprise customers

Disruption I the consumer space

- consolidate multiple devices into one
- single UI to manage all apps and data
- machine learning for efficiency
- apps for added functionality

Problem - Complex DC Infrastructure

- Fragmented
- Inefficient
- Dark

“Mass Data Fragmentation”

Solution

Single platform for all workloads and data

Single UI for simplified management

Machine learning for efficiency

Run apps to deliver insight

*Apurv Gupta, Chief Architect

Under the covers: SpanFS

Distributed Lock Manager to achieve consensus

Distributed key value store - you want this to be fully consistent

Cluster management - helps with new nodes, hardware, etc

Cohesity-specific stuff

SnapTree - allows them to do very efficient cloning

Transactional Metadata Store - it's not good enough if your filesystem is eventually consistent

Data Repository - filesystem natively understands cloud

Data Journal - helps with random writes

Largest customer cluster is 3PB, tested 256 nodes in cloud (around 6PB)

Need at least 3 DCs to make a stretched cluster work, or 2 with a cloud witness

SpanFS ties this all together, supports SMB, NFS, S3

Add Applications on top of this. What do you do with this data?

Application layer needs to be distributed and fault-tolerant

Run in Docker containers, orchestrated by Kubernetes

- Splunk
- Sentinel One
- other 3rd party
- Cohesity Insight

- DataProtect

*Anand Nadathur, Senior Director of Product Management
And Hildebrand

Deep dive into Cohesity Scale-out NAS

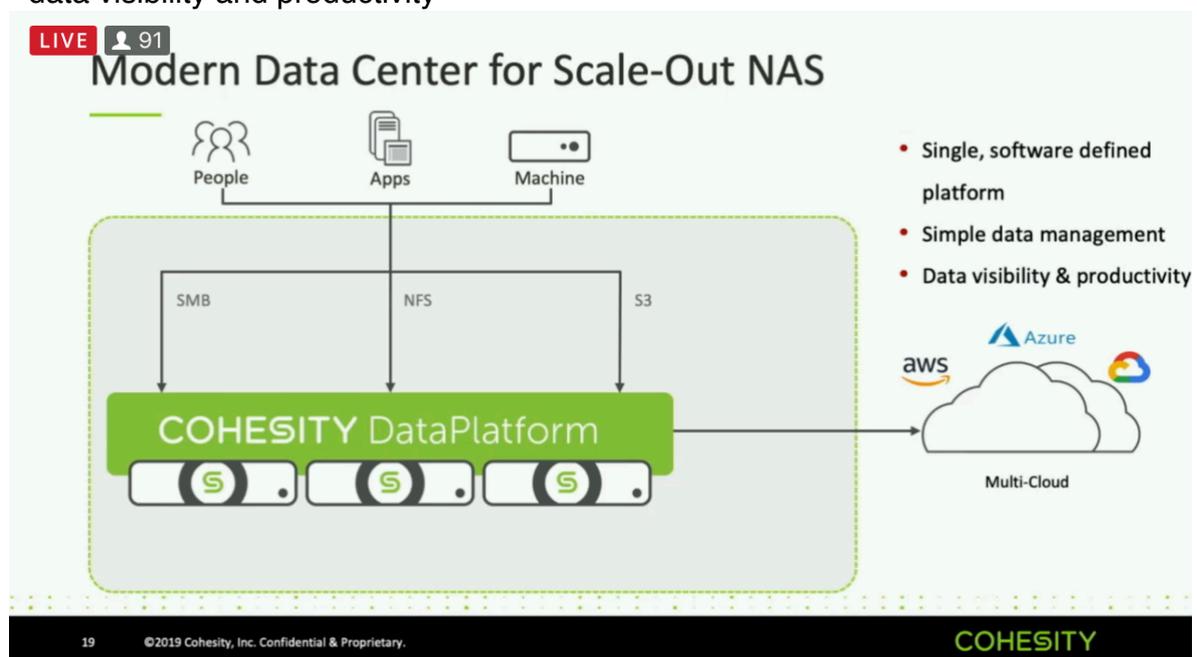
“Designed to eliminate silos of infrastructure that customers put in for files and objects”

Legacy DC Challenges

- Fragmented
- Inefficient
- Dark data

New DC

- single, software defined platform
- simple data management
- data visibility and productivity



Cohesity Difference

- Global space efficiency
- data mobility
- data resiliency & compliance
- instant mass restore
- apps integration

Global Space efficiency

Data Reduction and Small File Optimisation

- Global deduplication
- modes: inline, post-process
- archival to cloud is also deduped
- compression

- Zstandard algorithm
- Small file optimisation
- better performance for reads and writes
- benefits from dedupe and compression

Data Mobility

- data portability across clouds
- multi-cloud replication and archival (1:many)
- integrated indexing and search across locations

Simultaneous, multi-protocol access
Comprehensive file permissions

*Demo by Jon

Application development converting from SMB to NFS

Flexible Archive Solutions

- archive either on-premises or to cloud
- policy driven archival schedule for long term data retention
- data can be retrieved to same or different Cohesity cluster
- archived data is subject to further deduplication

Data Resiliency and Compliance

- ensures data integrity
- erasure coding
- highly available
- DataLock and legal hold

Achieving Compliance with File-level Datalock

- DataLock enables WORM functionality at a file level
- DataLock adheres to regulatory acts
- Can automatically lock a file after a period of inactivity
- Files can be locked manually by setting file attributes
- Minimum and maximum retention times can be set
- Cohesity provides a unique RBAC role for Data Security administration

*Chris Colotti

DataLock on Backups

- DataLock enables WORM functionality
- Prevent changes by locking Snapshots
- Applied via backup policy
- Operations performed by Data Security administrators

Ability to look within Helios for Ransomware

Prevent. Detect. Respond

Prevent

- Immutable file system
- DataLock (WORM)
- Multi-factor authentication

Detect

- machine-driven anomaly detection (backup data, unstructured data)
- automated alert

Respond

- scalable file system to store years worth of backup copies
- Google-like global actionable search
- Instant mass restore

*Sidarth Mishra, Product Manager

“Leveraging your backup and unstructured data”

Cohesity MarketPlace

- download and run apps directly on the platform
- From Cohesity and 3rd party ecosystem
- SDK for easy custom app development

Apps Integration

- bring compute to the data
- Marketplace
- 3rd party apps
- universal data access