

SFD15 - Cohesity

*Mohit Aron

Mission: Redefine secondary storage

Founded 2013 by Mohit (Co-founder of Nutanix, lead on Google file system)

Hundreds of customers (70% enterprise)

600% revenue growth YoY

50% headcount growth since September

The bulk of the work happening in the DC is boring - the secondary storage

Primary storage - Apps

Secondary - backups, archival, analytics, test/dev, file shares

This space is very fragmented. Eg with backups - one vendor for software, one for vendor, etc

Opportunity for modernisation?

“Clean up the mess left by enterprise backup products”

“Apple-like UI”, “Google-like scale”

Analogy is that if your house is messy, people will move to hotels (?)

Consolidate workloads on Cohesity DataPlatform then take advantage of copy data management

“Move the compute to the data - not data to the compute”

“Hyperconverged Secondary storage”

Cohesity as a virtual cluster in the cloud, then connect to physical appliances on-premises

You only think about plumbing if the plumbing is bad

SpanFS

- Unlimited scalability (256 nodes tested on Azure)
- Strict consistency (versus eventual consistency)
- Multiprotocol (NFS / SMB / S3)
- Global dedupe
- SnapTree limitless snaps and clones
- Self healing
- Automated Tiering
- Multicloud (Azure, Google, NFS or S3-compliant)
- Sequential and random IO
- multi-tenancy with QoS
- global indexing and search

From DC to edge to cloud

Edge - DP VE

The Cohesity Difference

Global Space Efficiency

- Variable length dedupe
- erasure coding

QoS

- multi workload isolation
- noisy neighbour prevention

Instant Mass Restore

- any point in time
- highly available

Data Resiliency

- strict consistency
- ensures data integrity

Cloud/Apps Integration

- multiprotocol
- universal access

*Ganesh

Challenge: 2 Storage architectures

Enterprise Stack

- scale up
- Reliable hardware
- NFS/SMB

Cloud Stack

- Scale out
- commodity hardware
- new protocols

CDP

Result

- distributed NFS/SMB
- pay-as-you-grow
- seamless upgrades

Challenge: Workload Consolidation [slide]

SpanFS Architecture

Distributed Lock Manager

Distributed Key Value Store

Cluster Management

SnapTree

Transactional Metadata Store

Data Repository

Data Journal

All tied together with SpanFS - SMB, NFS, S3, iSCSI, DataProtect

Healer runs a MapReduce for the metadata to get a global view of the system

QoS - every IO is tagged with QoS

Replication, Archival and CloudTier

Unlimited clones with low overhead

Traditional Clones - Performance overhead, limited clone frequency

SnapTree Clone - Minimal overhead, no clone frequency limitations

Lack of consistency with today's scale-out solutions

Vs

Guaranteed consistency for no loss of data

- write from client
- write propagated through cluster
- acknowledgement

*Gaurav Garg

Object Store Overview

Table Stake

- unlimited buckets
- store billions of objects
- AWS S3 compatible APIs
- buckets and object level ACLs
- Object versioning
- multi-part uploads

Global storage efficiency

Multiprotocol access

Instant bucket cloning

*Demo - simplicity of setting up and managing bucket for multi protocol access

*Demo - simultaneous multiprotocol access

*Demo - Accelerate Test/Dev with Instant Clones

*Sai Mukundan, Radhani Guturi

"Cohesity Anywhere: Driving to the Cloud"

Test/Dev & Data Recovery

- recover data to cloud for test/dev and DR to meet RPO/RTO SLAs

Long-term Retention

- Archive backup data to cloud for LTR
- granular recovery of data as needed

Create Hybrid Cloud

- seamless data movement across on-premises and cloud

Test/Dev

- policy drive backup to on-premises Cohesity

- Trigger VM conversion to cloud format
- deploy VM in cloud

LTR

- LTR data retention in the cloud
- encryption, dedupe, and compression
- granular recovery fo data (VM/File)

CloudRetrieve

- retrieve data to a new Cohesity Platform (on-premises / Virtual Edition / Cloud Edition)
- search and download metadata/data
- optimise data egress from cloud through granular recovery

*Mohit wrap-up

<https://www.sdxcentral.com/articles/news/cohesity-adds-file-object-sd-storage-data-backup-platform/2018/02/>

<http://www.futuriom.com/articles/news/cohesity-expands-into-file-and-object-storage/2018/02>

<http://www.computerweekly.com/news/252435045/Cohesity-adds-file-and-object-to-hyper-converged-hybrid-cloud>

http://www.channelbuzz.ca/2018/02/cohesity-making-push-into-file-and-object-storage-market-24693/?utm_source=website&utm_campaign=wordtwit&utm_medium=web

<http://searchdatabackup.techtarget.com/news/252435170/Cohesity-DataPlatform-now-spans-file-object-storage>