

SFD15 - Datrium

DataDomain and VMware heritage

*Rex Walters, VP of Product
Tier 1 HCI

Converging T1 HCI, Scale out backup, Cloud DR

DVX

Distributed Virtual:

- Compute
- Primary Storage
- Secondary storage
- data management
- data protection
- WAN optimisation
- Security and encryption
- scale-out filesystem
- cluster management
- disaster recovery
- cache management
- cloud software

DVX Components

Virtualized Compute

- vSphere, KVM or containers, with a global storage pool, unified management of the distributed system, and complete independence between servers. Any mix of choices all running together.

Primary Data Management

- Flash Performance
- Global Encryption
- Instant Cloning
- Always-on dedupe and compression
- quick, efficient “cold cache” refresh upon server failures
- one cache tier, zero parity overhead

Secondary Data Protection

- Always-on distributed erasure coding
- Choice of high density or high throughput flash copy data storage

Off-site data protection

- It's not enough to protect from failures within a DC. DVX can replicate data safely and efficiently off-site. Global dedupe and encryption makes retrieval equally efficient and safe

Built for Scale, Performance, Resilience

Performance Tier

1-128 compute nodes

- stateless hosts
- Up to 200GB/s 32K Random Reads
- Up to 18M 4K Random Read IOPS

Protection Tier

1-10 Data Nodes

- up to 1.7PB Capacity
- Up to 16GB/s 32K Random writes
- up to 1.2M Snapshots

Datrium Compute nodes

Also support third party in the primary tier

Compute - application CPU and memory, plus compute for data management

If you have a need for encryption - you need global encryption

Global data management: safe, efficient “pool” of storage

End-to-end integrity verification w/encryption

End-to-end distributed systems management, auto support

Always-on features - no tuning / decisions

Live data protection: distributed EC w/2-fail protection

Snapstore: on-premises

Elastic replication

Cloud DVX Software - <https://www.datrium.com/resources/cloud-dvx-datasheet/>

Largest HCI-like system ever measured

128 compute nodes / 10 data nodes

DVX: full scale, tier 1 validation (in partnership with Dell)

Dual-socket E5-2697 v4 (Broadwell) CPUs

2x1.92TB Samsung PM863a SATA SSDs

10 F12x2 Data Nodes

Hero Number? 4KB Random Read: 12.3M IOPS, 48GBps

32KB Random Read: 5.2M IOPS, 163GBps

8KB Random 70R/30W: 4.1M IOPS, 32GBps

Real world performance - <http://bit.ly/2AlfArl>

*Sazzala Reddy, CTO and Co-founder, Mike McLaughlin, Dir of Tech Marketing
Converging Primary and backup

Demo

Datrium Host Software - uses about 8GB RAM, 20% of host CPU resources, “Insane Mode” uses 40%

Every core yields about 10000 IOPS

Tier 1 HCI - Performance and resiliency

Scale-out backup - high data efficiency

Cloud DR - Fast Data Mobility

Common Technology Foundations

- Log-structured filesystem
- Global deduplication
- Backup catalog

SSDs across LAN have Little Law latency problem

HDDs and SSDs are not good at random writes

Dealing with variable sized blocks is hard for storage systems

Deduplication locally and across WAN

Finding needle in a haystack with millions of objects

Foundation 1 - Log-structured Filesystem

LFS Paper by Mendel & Ousterhout from Berkely in 1992

- Mendel later founded VMware, and is an investor in Datrium

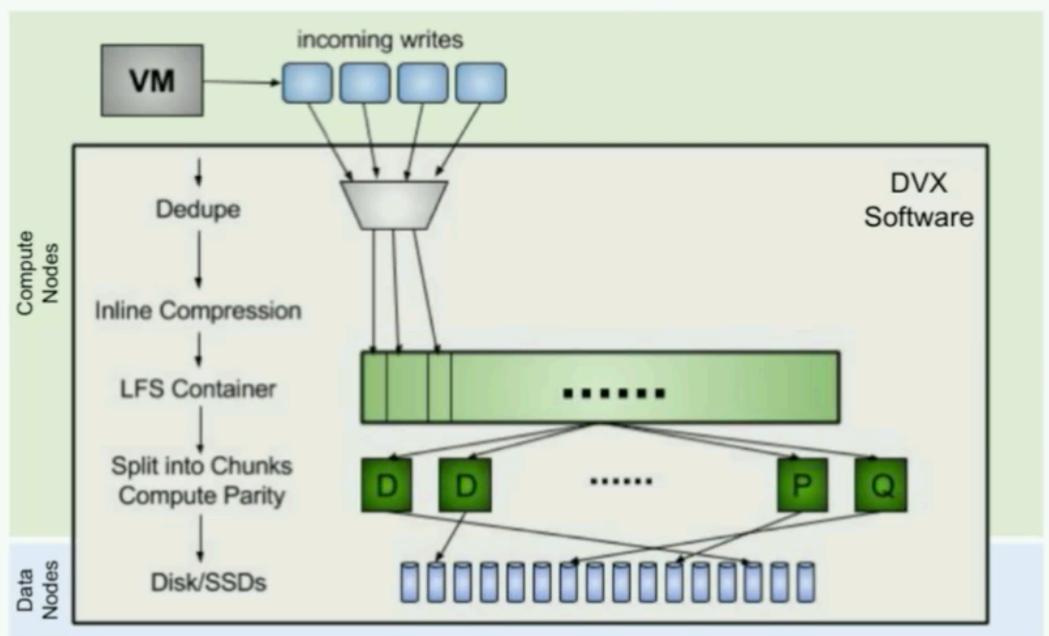
Main theory in original paper:

- Convert random writes into large sequential log on disk
- Commodity HDDs and SSDs are very good at large sequential writes

Datrium's Design Enhancements

- "Distributed" LFS design for PB scale-out object store
- Inline compression
- Inline erasure coding
- large scale algorithms

Datrium Simplified Write IO



8MB Containers

Erasure Coding

Devices Very good at large writes

Average deduplication and compression ratio is 4.4 based on customer data

Radical New Architecture

Compute Nodes

- Deduped local flash runs entire load
- LFS enables commodity SSDs
- Solves random reads problem

Data Nodes

- Durable scale-out object store
- LFS for commodity HDDs or SSDs
- Solves random writes
- Solves variable sized compressed blocks

DVX Log based filesystem solves more things!

High Performance

- Full stripe writes
- No read-modify-writes

Data Integrity

- No overwrites
- Avoids corruption

Cost Efficient

- Commodity HDDs and SSDs
- Inline erasure coding

Dedupe Index

- LSM Trees
- Fast key-value DB

ROW Snaps

- Write in new place
- no performance penalty

Designed for Cloud

- AWS S3 prefers full writes
- LFS maps to S3 objects

Foundation 2: Global Deduplication

Global Dedupe Aware

- faster movement of data
- Crypto hashes
- content-addressable system
- Blockchain-like verification

Foundation 3: Backup Catalog

Easy VM Policies

- Always consistent
- Automated membership

SnapStore

- Backup catalog
- Search & Operate

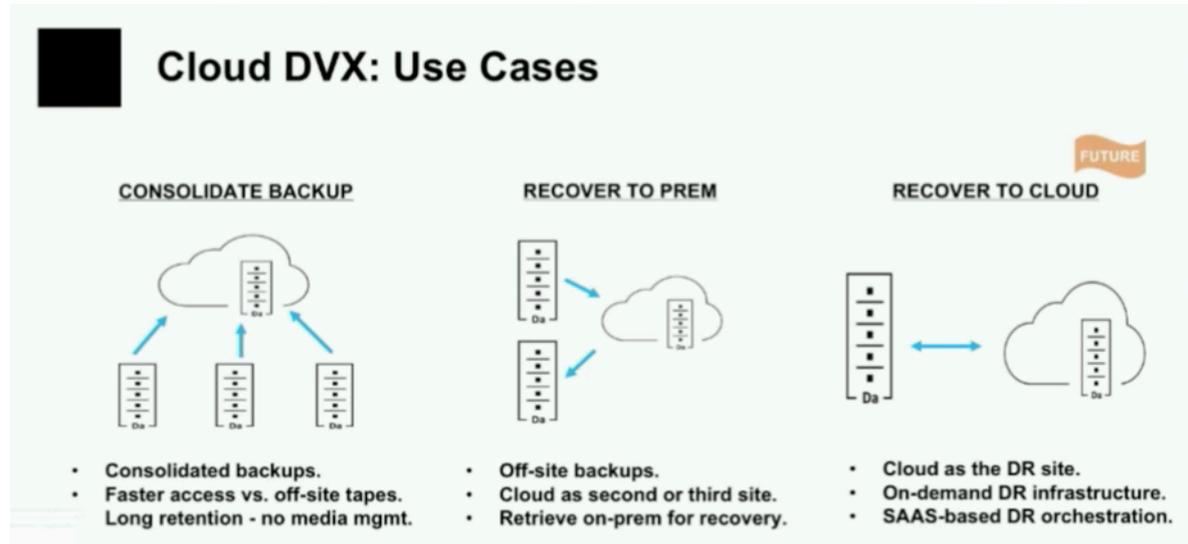
Granular Operations

- Snap, Clone, revert, replicate
- VMs, vDisks, Files, Containers

*Tushar Agrawal, Dir of Product, John Cho, MTS and Cloud Architect

Cloud DVX

Want an iPhone iCloud-like experience
As-a-service simple (zero management)
High efficiency - global deduplication
Faster recovery - direct to primary



Easy as 1-2-3

Customer

1. Create AWS account
2. Put AWS keys in on-premises DVX
3. Set protection policy

Datrium

- spin up AWS resources
- pair with on-premises DVX
- self-heal from failures
- auto upgrades
- sends proactive telemetry

Demo with John

Today one on-premises DVX can send data to one cloud DVX

Future? On-demand DR / SaaS

Leverage global dedupe and cloud DVX efficiency for near-zero RTO DR in AWS

- on-demand DR orchestration service
- on-premises and AWS
- eliminate second site for DR
- Instantiate data for test/dev/analytics