

SFD10

@DateraInc
<http://datera.io>

*Marc Fleischmann (CEO, Co-founder)

Bring hyper scale operations and economics to private and public clouds
Founded in 2013
Launched in April 2016
50+ employees, based in Sunnyvale

Elastic data fabric (<http://datera.io/product/#elastic-data-fabric>)
hyper scale API-driven elastic block storage (EBS)

shipped 2PB raw capacity in first revenue quarter
4 customers announced, more than 10 not announced

cloud not so much about the location of the data - it's really about the operational model

Scale-out storage platform

Application intent

mix and match nodes on standard hardware (NVMe and hard drives)
can have different application flavours

not all tenancy features have been implemented

“no more data migration orgies”

iSCSI connector, REST api

Summary

Operations

- API-driven
- designed to scale
- infrastructure-as-code

Nodes

- commodity parts
- software-controlled
- performance-optimised

Benefits

- elastic economics
- value-matched
- future-ready

Software available with qualified hardware (prescriptive, currently SuperMicro)
Can be licensed as software-only as well

2 SKUs - 50TB or 100TB

Replication, clones, thin provisioning

*Docker/Swarm Demo

Bill Borsari

@billborsari / <https://www.linkedin.com/in/billborsari>

DT5005 nodes

2 RU chassis

12 drives

2 NVMe drives

2 boot drives

1GbE mgmt

10 and 40 for data

Datera - software defined storage appliance that takes over the hardware

*Architecture Deep Dive

Raghu Krishnamurthy

<https://www.linkedin.com/in/rakrishnamurthy>

Whiteboard session

What makes Datera Unique?

Intent defined - templates, intelligent placement

economic flexibility - heterogeneous nodes (capacity, performance, media type)

API first Dev/Ops model - infrastructure as code, programmable/composable

Multi-tenant - network isolation, QoS

Infrastructure awareness - auto-forming, optimal allocation of infrastructure resources

primary interface can be 1/10GbE for access / front-end

secondary interface - backend network (distributed protocol, management)

12 x 4TB or 8TB disks (48 or 96TB), flash tier (capacity and performance tiers)

Other nodes in the system can be different, i.e. all flash

[SFD10_DateraControlPlane.jpg]

Datera control plane runs on every storage node

cluster mgmt

configuration mgmt - replicated to another set of storage nodes

node platform - what is it capable of?

optimiser - optimising the placement of data

policy engine - responsible for intent-based provisioning

UI and REST support

Volumes in a consistency group for snapshots

Top-down approach instead?

Intent provisioning

A way of describing what your application wants and then letting the system allocate the data

application template

- policies for management (e.g. QoS) - data redundancy, data protection, data placement
- storage template - defines how many volumes you want, size you want
- pools of resources

don't use RAID, currently use 1->5 replication (synchronous) within the cluster

snapshots are copy on write (at an application intent level)

*Bill Borsari - OpenStack Demo

doesn't matter which node you log into

You can use a VIP if you want

Can support IOPS as well as bandwidth QoS

Datera supports Python SDK