

SFD10

Primary Data (@primary_data)

*Lance Smith - CEO

exited stealth Nov 19 2014

DataSphere platform unveiled at VMworld 2015 US

CEO Lance, Chief Scientist Steve Woz, CTO David Flynn, CMO Rick White

HQ in Los Altos, CA

“Greater Data Lakes” at EMC World 2016 US

Global data space - Flash <-> Block <-> File <-> Object

“Different storage types naturally create data silos”

“31 flavours: good for ice cream, complex for storage management”

Enterprises have plenty of storage - dynamic data mobility finally puts all to work

Data mobility - change cars mid-race without application disruption

Virtualising data across a global namespace

Benefits?

- lower cost of over provisioning
- increase performance on demand
- break out of vendor lock in
- reduce cost and down time of maintenance (upgrades, scale, data migration)
- lower cost of archiving to object store (single namespace and online)

*David Flynn

Technology Vision

Data Virtualisation

- any piece of data
- on any storage
- movable at any time
- without interruption

What VMware did for compute is what they're doing for storage

Data access is direct and at full performance

Separating metadata and control plane

Datasphere - metadata

DSX - client (data path) - portal, mover and store

NFS as the control plane

Out-of-band, storage agnostic approach [SFD10_Pd_DavidFlynn.jpg]

Supports anything that can be presented to Linux as a block device ...

You could think of DSX as SDS (ppl are dropping DSX on a box to make the drives a

storage node)

Question about VSAN - Pd can sit on VSAN as a block storage pool, or just present those drives to DSX

HyperConverged?

Confluence of hyper converged and scale out NAS

Treats the storage as a BLOB store

Data Migration vs Data Mobility

[SFD10_Pd_DataMigrationvsMobility.jpg]

Resource visualisation and allocation

- identify and control what data consumes physical resources through telemetry and throttling capabilities at the client, object granular service levels

Single, universal VASA provider

Dynamic data mobility

- align data's needs (objectives) with storage capabilities (service Levels) through automated mobility, arbitrated by economic value and reported as compliance

Every storage container is a landlord with floorspace to lease and utilities available (capacity and performance)

SLOs - Aspects

Protection

- Durability
- Availability
- Recoverability
- Security
- Priority
- Sovereignty
- Access Control

Performance

- IOPS / Bandwidth / Latency
- Read / Write
- Sustained / Burst

Price / Penalty

- Per File
- Per Byte
- Per Operation

Automatic

Smart Objectives

- automatically optimise Objectives for varied and ongoing of entire application

environments using community extensible predicates and analytics

Predicates

Built-in

- File type
- name / path
- owner / group
- size / bytes used

Analytics

- frequency / recovery of access
- duty cycle
- type of access (direct / buffered, seq / random, read / write)
- Latency sensitivity

Extensible

- User defined
- admin defined
- application defined
- 3rd party tool defined

Capacity Planning

- simulate changes to physical infrastructure or amounts of data at specific SLAs to determine “what if”

“Data virtualisation obviates point solutions”

- caching: in-server, appliance
- auto-tiering
- cloud gateways
- hyper-converged, software defined storage
- scale-out, clustered NAS
- storage virtualisation
- data reduction
- copy data management

*Demo and Use Cases - Kaycee Lai - SVP Products and Sales

DataSphere Use Cases

- data migration
- scale out NAS [SFD10_Pd_ScaleOut_Deployment_Options.jpg]
- data lifecycle management

Going beyond SPBM to Objective Based Data Management

DataSphere Overview and Demo

