

SFD6 - Pure Storage

*Matt Kixmoeller, VP of Products

@mattkix

storage market in the midst of disruption

Tier 1 - going to AFAs

Tier 2 - hybrid disk tiering arrays, hyper-converged disk tiering appliances, commodity cloud

What if

- Flash could be everyday storage?
- Storage was fundamentally simpler?
- We could build a better storage company?

Flash Revolution?

- Disruptive performance and economics
- Built to share and scale
- Always on
- Simple and invisible

Simplicity is the manifesto

End-to-end guarantee

All-inclusive software pricing

No required training or PS (although they use partners for install and optimisation)

Fanatically proactive support

A better approach to the storage array lifecycle

Ask the right questions about AFA? Not

- Can my array do 1 million IOPS?
- Arch swordplay
- Raw price per gigabyte
- How do I adjust the Tiering? RAID? Caching? Disk layout? etc

How does the app perform?

How do I scale capacity and performance?

Usable storage cost for my applications?

How simple is this for my admins, DBAs, etc?

*Neil Vachharajani - Principal Architect

How can we make flash affordable? Available? Reliable?

What does this mean for performance?

Commodity x86 servers talking to each other over infiniband - between controllers.

SAS between controllers and Flash shelves

cMLC Flash (consumer grade flash)

RAID is segment-based

Host generates data -> goes to NVRAM - is duplicated - ACK is then sent back to host

Data reduction Round 1: inline dedupe, medium weight compression

SSD segments RAID-3D

Two options then

1 - background dedupe (Data Reduction Round 2A) - then write down metadata (garbage creating process).

2- garbage collection (Data reduction Round 2B) - heavy weight compression

What problem are you trying to solve with scale up vs scale out?

Capacity, Performance and Reliability in a single array. All while being worried about cost.

PPL get uncomfortable with around 0.5 - 1PB in a single array, regardless of the resiliency

*Vaughn

Asking the right questions

Simplicity

How simple is this for my storage admin?

Is storage invisible for my app owner?

No more app alignment - invisible

No more storage tuning - simple

And automated ...

Operate & Automate

- Web-based GUI

- CLI

- RESTful API

Integrated with your stack

- OpenStack

- VMware VAAI

- vCenter Plugin

- Log Insight plugin

- MSFT VSS Provider

Pure Storage joins OpenStack Foundation (code contributor and corporate sponsor)

- Pure Storage Cinder Driver - Standard Cinder deployment

- Python Automation Toolkit - semi-custom implementation

- RESTful API - fully customised

Demo - Project codename SPOG (single pane of glass) - simple, scalable, multi-array management

SaaS-based, part of CloudAssist

vVOLs Demo - VM-Granular Flash Storage

Pure have migration from VMFS to vVOLs sorted

Scalability

Stateless controllers provide CPU, memory, IO

- NVRAM write log in first shelves
- controllers mirror memory contents

Cluster has "virtual persona"

- WWNs
- IQNs
- Eth address
- Replication
- Performance history

Scalability Results:

FlashArray Today - Flexible Scale

Federated Management

Federated Scale

Performance

Performance claims versus real world

Average IO size is 40.6K

"4K Vanity Zone"

Do you know the I/O size of your application?

Adaptive metadata enables scale & multi-application consolidation

Adaptive metadata fabric

- Scale-optimized (no virtual space limit)
- Fine-grained, variable addressing (better dedupe, auto-align to all FS / VM / DB, perf optimised for real-world, large block datasets)
- protected against failure (dual-parity + log metadata protection, persisted to flash)

Don't confuse thin provisioning with data reduction