

SFD10

@Tintri

[Comment: Tintri are clear about what they offer: VM-aware storage. Not app-aware, but VM-aware]

*Kieran Harty, CTO and Co-founder
Product Vision and Overview

1000 customers worldwide

Restored the simple relationship between storage and the applications that are using storage

Support VMware, Hyper-V, OpenStack, RedHat

Tintri is VM-aware Storage

- VM-level QoS
- VM-level analytics
- VM data management
- VM-level automation with PowerShell and REST
- Multi-hypervisor

IO Blender = contention, lack of visibility vs VM aware storage (VAS) performance isolation

VMstore assigns every VM its own lane to optimise performance

visual guidance on how/when to throttle IOPS for each individual VM

- immediate feedback on the latency/contention impact

VM-level metrics for every VM all the time

Replicate only the VMs you care about

- one to one, bi-directional and many to one topologies

- data protection and DR in remote VMstores

- set policy for individual VMs

Deep Hypervisor integration

vCenter plugin 2.0

vRealize Operations Manager

New Stuff?

Modern Scale-out

set of arrays can be combined together (using scale-out software which makes decisions about which workloads to run where)

similar to storage DRS (but "done right")

New models announced as well

17TB - 308TB effective

Predictive analytics announced

VM Scale-out

Versatile

- loosely coupled scaling
- heterogeneous nodes
- hypervisor neutral

Intelligent

- based on deep VM workload analysis
- optimised distribution across hardware platforms / capabilities

Automated

- recommendation / fully-automated modes
- policy stickiness

*Justin Lauer, Global Evangelist
VAS and QoS Demo

Dashboard - rollup of all the per-VM stats

End to end latency - host / network / storage (contention and flash) / throttle latency

Space savings at 7.7x - 87% (including thin provisioning 14.6x - 93%)

Capacity gas gauge (107 days till full), performance gas gauge (when you'll run out of performance)

Customers live in the Virtual Machines tab

Synchronous replication is on the roadmap for this year (currently at 1 minute)

Current interface is Flash-based

*Mark Gritter, Principal Engineer and co-founder
VM Scale-out demo

Tintri Global Center

Each box is still self-contained

*Brandon Salmon, Principal Engineer
Tintri Analytics Demo

Already show real-time analytics on TGC today

Build models from data

Delivered as SaaS

Data is coming from phone-home service data

Understanding what applications need from your infrastructure

You need capacity

And performance

And working set (whether hybrid or all-flash)

Average VM needs

- capacity
- performance
- working set

*Christopher Slater - Technical Marketing Engineer
SyncVM Demo