

SFD6 - Wednesday PM

Tegile

*Narayan (CMO)

Business transformation through storage transformation

First product shipped Feb 2011

Close to 700 customers

Focused on enterprise mid-range

1200 systems in the field, 300+ partners

Early stages of the flash revolution

Dial up or down the flash in the system depending on requirements

Same OS, feature set or uxp whether you're using hybrid or AFA

HA2100, HA2130, HA2300, HA2400, T3400, T3600, T3700, T3800

All models can run AFA or hybrid as required

All software features bundled in

Support eMLC and spinning disk, can support SLC. TLC - the economics and resilience aren't there yet. PCIe Flash support is coming.

Support

Performance

metadata acceleration

r/w/meta in DRAM/SSD cache

pin data sets in all SSD pools

sustained low-latency

Economics

inline dedupe

inline compression

zero elimination

thin provisioning

automatic block reclamation

Protection

Dual a/a controllers

multiple RAID options

instantaneous thin snapshots, thin clones

block-based remote replication (asynchronous)

Multi-protocol

Block and file (FC, iSCSI, NFS & SMB)

Currently 2-controller, scale up architecture. Working on scale out architecture (12 - 18 months)

They are a systems company - you can't uncouple the hardware

IntelliCare - Cloud Analytics (sits on AWS), IntelliCare Portal and Proactive Customer Support

Value proposition:

Server virtualisation,

5 - 10x perf gain

3:1 - 5:1 data reduction

VM-consistent backup and recovery

desktop virtualisation,

no boot storms

10x better uxp

8:1 - 10:1 data reduction

VM-centric mgmt

databases

3x - 20x reduction in latency

sub ms response times

2:1 - 3:1 data reduction

integrated backup and DR

dedupe and compression can be turned on and off on a per volume or per pool basis

Dedupe is typically done on a pool basis

standard x86 architecture, SAS back-end, no custom silicon

tested up to 96000 snaps per system

Self-encrypting drives, SSDs

entry level 50000 IOPS up to 250 - 300000 IOPS

Design Premises

Flash densities doubling and cost halving every 18 - 24mos

Flash densities increasing faster than HD densities

latency and economics will drive adoption of new media types

continued innovations in incumbent media tech ensure longevity

Design Goals

architecture must support a perf and one or more capacity layers

data mgmt @ speed and scale of the fastest medium

comprehensive data mgmt regardless of media or access methods

abstraction of media types for rapid incorporation of new media

*Rajesh (CTO)

IntelliFlash

Physical Media - Performance (DRAM and eMLC today) and a capacity layer (dense eMLC, MLC or HDD)
DRAM, Perf Flash, Capacity Flash and Capacity HDD

Media Optimization - Media Management

Media aligned writes

- aligned along sector and page boundaries
- avoids page fragmentation
- adjust to specific drive models (e.g. SSDs that advertise a 512 byte sector size)

Even wear across flash pages

- segment flash into high-low I/O ranges
- over-provision high I/O ranges

Fewer writes extend media life

- compress and dedupe writes
- variable block size support
- zero block avoidance

Vendors provide them with “performance” firmware from drive vendors

Media Optimization - Storage Pooling

Pooled Storage

- Media virtualised and aggregated
- provision luns and shares from pool
- simplified on-demand provisioning
- optimal space utilisation

Dynamic Stripe Widths

- variable block sizes
- each logical block is its own stripe
- all writes are full stripe writes
- no read-modify-write media wear

Coalesced writes to media

- writes in DRMA aggregated
- Coalesced and mostly sequential writes to data layer
- Efficient media handling
- Consistent performance

Media Optimization - Resiliency

Redundant Controllers

- dual redundant controllers
- active i/o for handling for local store
- standby i/o handling for peer store

continuous data access

Redundant Paths

- dual ported SSDs and HDDs
- No SPoF
- Redundant I/O fabrics

Live swaps and Upgrades

- online capacity expansion
- hot swap controllers and media
- non-disruptive sw upgrades

Metadata
information about data

Metadata Acceleration: Aggregation and Placement

- Metadata abstracted from data
- aggregated metadata placed on logical metadata devices
- optimal I/O paths
- increases locality of metadata and ease of metadata access
- enables high-perf inline dedupe and accelerates all storage operations

Typically 10 - 15% of total capacity used for metadata (not advertised as useable space)

Metadata is accessible by both controllers, but lives in a pool, controlled by one controller

Multiple copies of metadata stored in the system in multiple media

Metadata Acceleration: Caching and Scaling

Metadata abstraction and aggregation

- intelligent real-time caching across Hdd, SSD and DRAM
 - adaptive algorithms opt fo multi io patterns
 - cachine engine adapts to block size, i/o patterns and media sizes
- automated expansion of metadata space
- threshold-based alerts for metadata space expansion

Data Services - Data Reduction, Protection, Recovery, Block and File Layers

Data Services: Data Reduction

in-line dedupe

- block-level
- dedupe across media
- perf multiplier

inline compression

- block level
- turn on/off at LUN / share level
- alorithm - LZ4, LZJB, GZIP
- perf multiplier

thin provisioning

- array-level thin
- for LUNs and shares
- supports VMware VAAI "STUN"
- JIT storage provisioning

Compress then dedupe

Data services: Data Protection

Instantaneous thin snaps

- PIT copies of data
- space allocated only for changed blocks

- no reserve space for snaps
- unlimited number of snaps
- VM-consistent and app-consistent

Instantaneous thin clones

- mountable copies
- R/W able copies
- PIT copies
- Space allocated only for deltas

Detect and correct silent corruption

- checksums all data blocks
- data and checksum in separate locations
- match data/checksum for integrity
- corrupt / mismatched data fixed using blocks from mirrored copy

Data Services: Data Recovery

Instantaneous stable Recovery

- data-consistent VM snaps
- hypervisor integrated
- MSFT VSS co-ordinated data-consistent snaps
- VM-consistent and app-consistent snaps

Intelligent data reconstruction

- no need to rebuild entire drive
- only portion of data rebuilt
- accelerated metadata accelerates rebuilds

WAN-optimized replication

- snap-based site-to-site replication
- no need to replicate to multiple updates to a block within repo interval
- minimizes bandwidth usage
- one to many / many to one replication

Data Services: Block and File Services

Block Layer

- logical block volumes
- thin / thick provisioning
- vol expansion / contraction
- optimised I/O services (block copy, clone, etc)

File Services Layer

- file system folders
- Rich ACLs
- Quotas / Reservations
- Optimized file I/O services (file copy, extended attributes, etc)

Protocol Choice

Block

iSCSI 1Gb 10Gb

FC - 4/8Gb

File Access

- NFS, CIFS and SMB 3.0 (coming)
- Autohome directories for CIFS

- AD integration

*Rajiev talks about Management

Web UI

REST API

VMware vCenter

MSFT SCVMM support coming soon

Call Home