Flash Storage

Roles & Opportunities

L.A. Hoffman/Ed Delgado
CIO & Senior Storage Engineer
Goodwin Procter L.L.P.

i/+^ 2013 the catalyst

Overview & Background

- Flash Storage history
  - First smartphones, then laptops, then onboard cards, and now – Enterprise Class Flash storage arrays to compete with spinning disk
- Flash as a technology solution
  - Enterprise Class features
    - Management & HA features
    - De-duplication
    - Snapshots, Thin Provisioning, Clones, etc.

Flash Chip Components

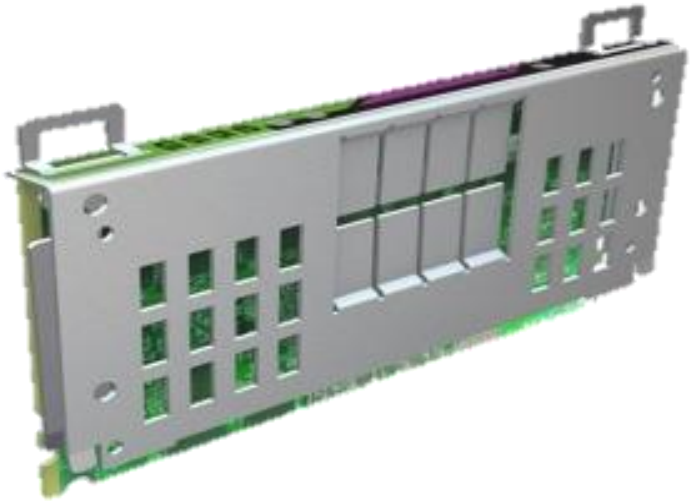

Violin Array internal architecture

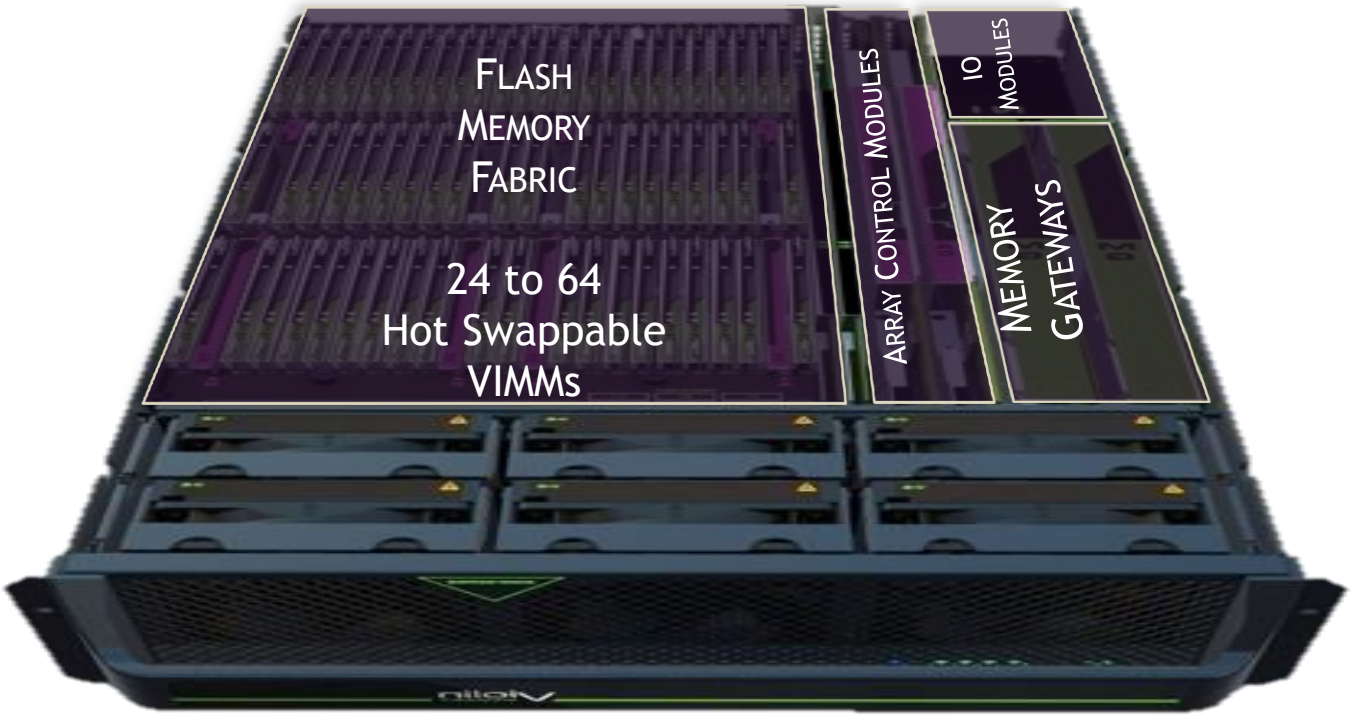

Flash Performance Examples Autonomy IDOL

- Autonomy IDOL Indexing Benchmark – Single Engine – 250gb

Media	Detail	Unused space	File Write (MB/s)	File Read (MB/s)	Time
Spinning Disk Raid Group	8x 300gb 15k RPM drives, RAID10	818gb	187.9	423.6	7m 30s
Flash based array	1x 250gb LUN	0gb	189.5	446.1	7m 12s

- Spinning disk results match the flash array for a single engine
- To match performance however requires lots of wasted space
- Spinning disk array cache nears saturation during file write affecting other workloads (Exchange, SQL) on array

Flash Performance Examples Autonomy IDOL

- Autonomy IDOL Indexing Benchmark – 4 Engines – 250gb

Media	Detail	Unused space	File Write (MB/s)	File Read (MB/s)	Time
Spinning Disk Raid Group	8x 300gb 15k RPM drives, RAID10	68gb	88.6	371.2	11m 38s
Flash based array	4x 250gb LUN	0gb	165.3	407.8	9m 17s

- Flash performance drops modestly, 9% for writes, 13% reads
- Spinning disk performance drops by 53% for writes, 13% reads
- Spinning disk array cache fully saturated. Writes have to wait on disks to complete requests before telling host to proceed
- 4 Engines running on one Raid Group to maximize $/TB

Flash Performance Examples
Autonomy IDOL

Autonomy IDOL Deployment

Spinning Disks

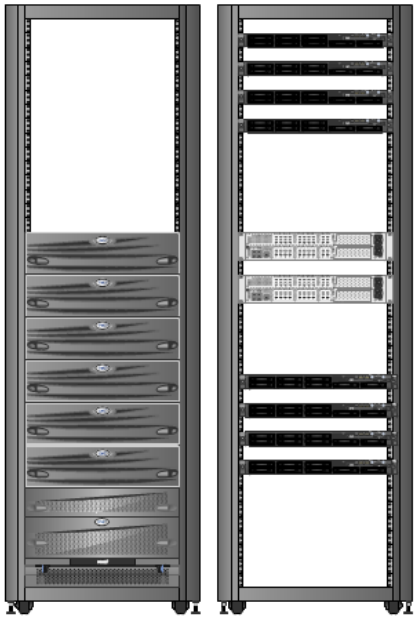

Flash

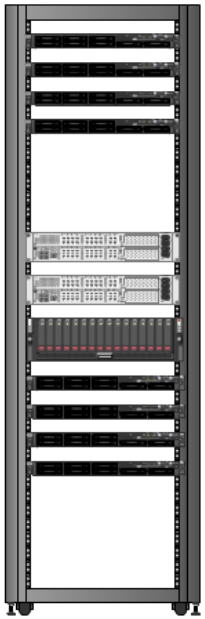

Proposed IDOL Configuration

- 2 Clusters of 4 Servers
- 5 Content Engines per server
- 250gb per Content Engine
- 40 Content Engines total
- 5TB per Cluster, 10 TB Total

i/+∧ 2013 the catalyst

Flash Performance Examples Autonomy IDOL

	Spinning Disks	Flash
Media Used	80 x 300gb 15,000 RPM Drives	Choice of SSDs or memory-based Flash
Array Size	10.3TB Usable with RAID 10	13.7TB Usable, Flash optimized RAID
Engine Layout	10 total Raid Groups, 4 Engines per	No Raid Groups
Number of Engines	40	40
Array Rack Space	Dedicated Cabinet - 25U in use	In Existing Rack - 3U
Array Power	Additional 30Amp Circuit Pair	Approximately 7 Amps
Estimated Cost of Array	$165,000	$135,000
Est. Cost of Array Power	$1,200 / month	$0 additional (existing rack)
Space Available for other projects	360gb	3700gb
Estimated Price per Usable GB	$15.45	$9.50

Flash & VDI

- Short history of VDI
  - Long on promise; short on delivery – thus far. Weak adoption rates to date
  - Why?
    - Expensive licensing regimes
    - Immature VDI software functionality & support
    - Difficult storage workload – 20/80 read/write and highly burstable due to boot storms, antivirus, etc.
    - IO performance demands require vast amounts of conventional disk storage, thus killing ROI projections

What Could Flash Drive VDI to Deliver?

- In an ideal state, VDI implementations potentially offer the prospects of:
  - Lower cost and faster OS and package deployment – hours to deploy images from a console, versus weeks or months via manual methods – lower Opex
  - Extended life of existing desktop assets – lower Capex

VDI Costs

- VDI Operational Benefits must be weighed against:
  - Computing and memory costs are pushed to the data center, so the effective leverage ratio, i.e. VDI images-to-core lifecycle costs, is critical
  - Storage is also moved to the Data center – Two models of VDI storage
    - Traditional spinning disk – not extensible at reasonable cost
    - Flash arrays – address extreme I/O performance requirements, and more recently, at a lower Opex

Flash Performance Examples VDI

- VDI Image Cloning Benchmark – Single Desktop

	Spinning Disks	Flash
Media Used	5 x 300gb 15,000 RPM Drives	Memory-based Flash
LUN Size and Layout	1024GB LUN	1024GB LUN
RAID Layout	RAID 5 (for cost)	Flash optimized RAID
VM Template Size	57gb	57gb
Time to clone a single VM	281sec	116sec
Bandwidth cloning single VM	197.1 MB/s	477.5 MB/s
Boot up time (to login screen)	66.2sec	22.7sec

Flash Performance Examples VDI

- VDI Provisioning Benchmark – 10 Desktops
  - Using VMWare Horizon View Suite – Creation and Configuration

	Spinning Disks	Flash
Media Used	5 x 300gb 15,000 RPM Drives	Memory-based Flash
LUN Size and Layout	1024GB LUN	1024GB LUN
RAID Layout	RAID 5 (for cost)	Flash optimized RAID
VM Template Size	54gb	54gb
Time until last VM powered up	48min 50sec	8min 2sec
Time until last VM was configured	1 Hour 7min	12min 20sec

Flash Performance Examples VDI

- VDI Basic Operations Benchmark – 10 Desktops
  - VMWare Tools used to initiate actions simultaneously (Guest)
  - Simultaneous Power On simulates a refresh of the VDI image
  - Guest reboot simulates a boot storm (usually IT initiated)
    - Patches, application deployments, etc - all could force reboot

	Spinning Disks	Flash
Time to Guest shut down	42sec	23sec
Time to power on (to login screen)	3min 54sec	52sec
Time to Guest reboot (boot storm)	3min 39sec	42sec

Flash Enables a Game Changer for VDI

- The speed of Flash allows for in-line deduplication
- Data reduction of at least 10:1 is achievable in most environments for the primary OS Image
- Requires careful planning as User Data achieves a much smaller Deduplication ratio (closer to 2½:1)

	Spinning Disks	Flash	In-Line Dedupe Flash
Size of 10 VMs (Base OS only)	540gb	540gb	540gb
Amount written to array	540gb	540gb	36gb