



# Flash Memory Arrays in Enterprise Applications

Ken Ow-Wing, Senior Product Line Manager  
Violin Memory, Inc.

685 Clyde Ave, Mountain View, CA 94043

Office: 650-396-1603 Mobile: 415-608-7773

- Enterprise Customer Requirements
- New Product Category
- Enterprise Use Cases
- Business Benefits
  
- Appendix
  - Economics
  - Array Characteristics



# Enterprise Environments: Requirements

- Flash Performance
- Consistent Low Response Time
- Reliability
- Availability
- Serviceability
- Scalability
- Manageability
- Resource Utilization

# Evolution of Use of Flash

Flash Memory Array



## PURPOSE-BUILT ENTERPRISE SOLUTION

- Networked/shared storage
- Sustained R/W throughput
- 7x24x365 operation

3<sup>RD</sup> GENERATION



2<sup>ND</sup> GENERATION

- Workstation/Gaming
- Memory extension/cache
- Limitations for High End Data Center Usage



1<sup>ST</sup> GENERATION

- Direct drive replacement
- Cost sensitive
- Limitations for High End Data Center Usage

# Flash Memory Arrays




Flash Chips  
4GB



10,400  
Chips

Flash Package  
32GB




1344 Packages

Capacity  
VIMMs  
512GB



84 memory modules

Flash vRAID Group  
2560GB



16 Groups

Capacity Flash Systems  
40TB in 3U




Flash Memory Array

*Data Center packaging reduces capital cost, space, power and operations costs. Infrastructure Consolidation*



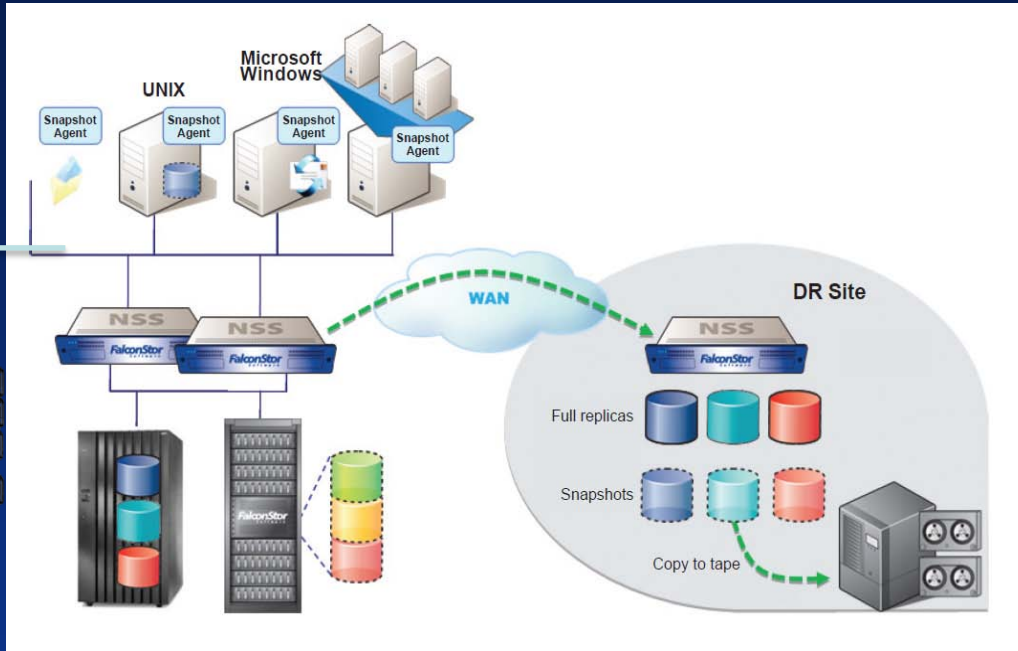
# Flash Memory Storage – 2PB

Silicon Virtualized Data Center



# Flash Memory Arrays

Available by  
the rack



Available as shelves



Fits in  
Virtualized  
Environments





# Database Appliance – 20,000 users



## High Performance Database Solution for OLTP

Architecture View  
512GB memory  
15 TB max DB size  
100M OLTP trans / hr

- Fits in OEM system

Production Database  
In Flash Memory Array

Disk)



2 x 10GbE Switches

Production Server

App/Test/Dev Server

Software  
Production Data Base  
Storage mgmt





# What's Different about Flash Memory Arrays?

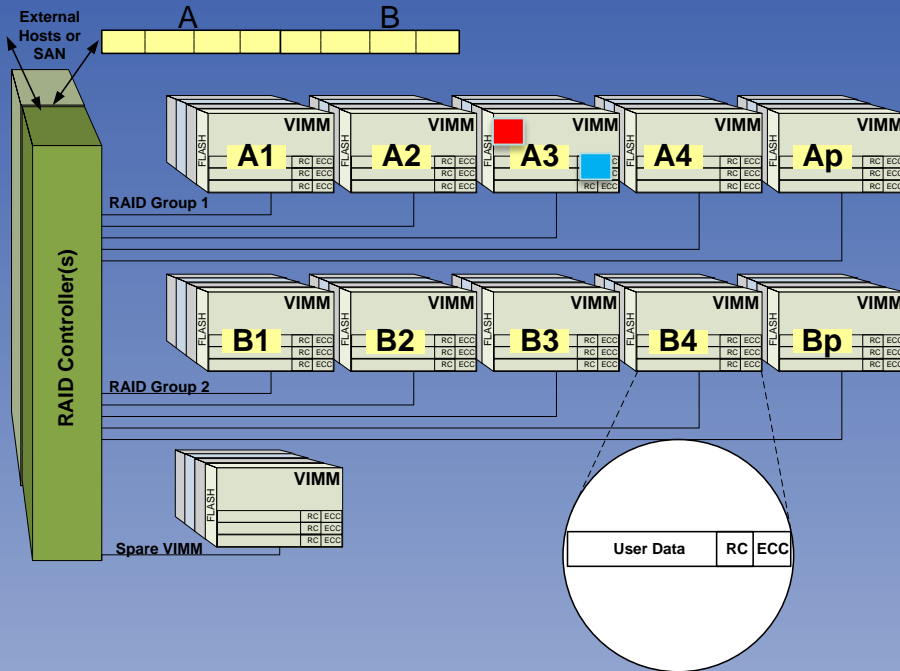
Compared to SSD

| Difference                        | Benefit                            |
|-----------------------------------|------------------------------------|
| * No support for rotating media   | Optimum performance with flash     |
| * Distributed Garbage Collection  | Sustained Writes, no "Write Cliff" |
| * Purpose Built "vRAID" for Flash | Sustained Writes, no "R/M/W"       |
| * vRAID not blocked by erasures   | Significant Latency reduction      |
| * vRAID protects flash devices    | No replacement on flash failure    |
| * Flash Packaging                 | Density > 10TB per RU              |

\* Flash Memory Arrays are different from SSD and/or flash cards

# Hardware Flash RAID

1<sup>st</sup> Purpose Built RAID for  
Flash Memory Arrays



\* Violin Intelligent Memory Module

Failure Handling Result:

- Data rebuilt on same VIMM
- VIMM stays in service
- No data loss
- Increases MTBF 4X

Details – Example

- Flash chip fails (Red)
- vRAID rebuilds data on same VIMM (Blue)
- Garbage collection avoided, performance maintained
- Rebuilt data on extra NAND
- HW RAID in Controller



# What's Different about Flash Memory Arrays?

## Compared to PCIe Card with Flash

| Difference                        | Benefit                            |
|-----------------------------------|------------------------------------|
| * No support for rotating media   | Optimum performance with flash     |
| * Distributed Garbage Collection  | Sustained Writes, no "Write Cliff" |
| * Purpose Built "vRAID" for Flash | Sustained Writes, no "R/M/W"       |
| * vRAID not blocked by erasures   | Significant Latency reduction      |
| * vRAID protects flash devices    | No replacement on flash failure    |
| * Hot swappable components        | No outage or data loss             |
| * Shareability                    | Max utilization by many servers    |
| * Scalability                     | Lg. dataset w/simplicity           |
| * Flash Packaging                 | Density > 10TB per RU              |

\* Flash Memory Arrays are different from SSD and/or flash cards

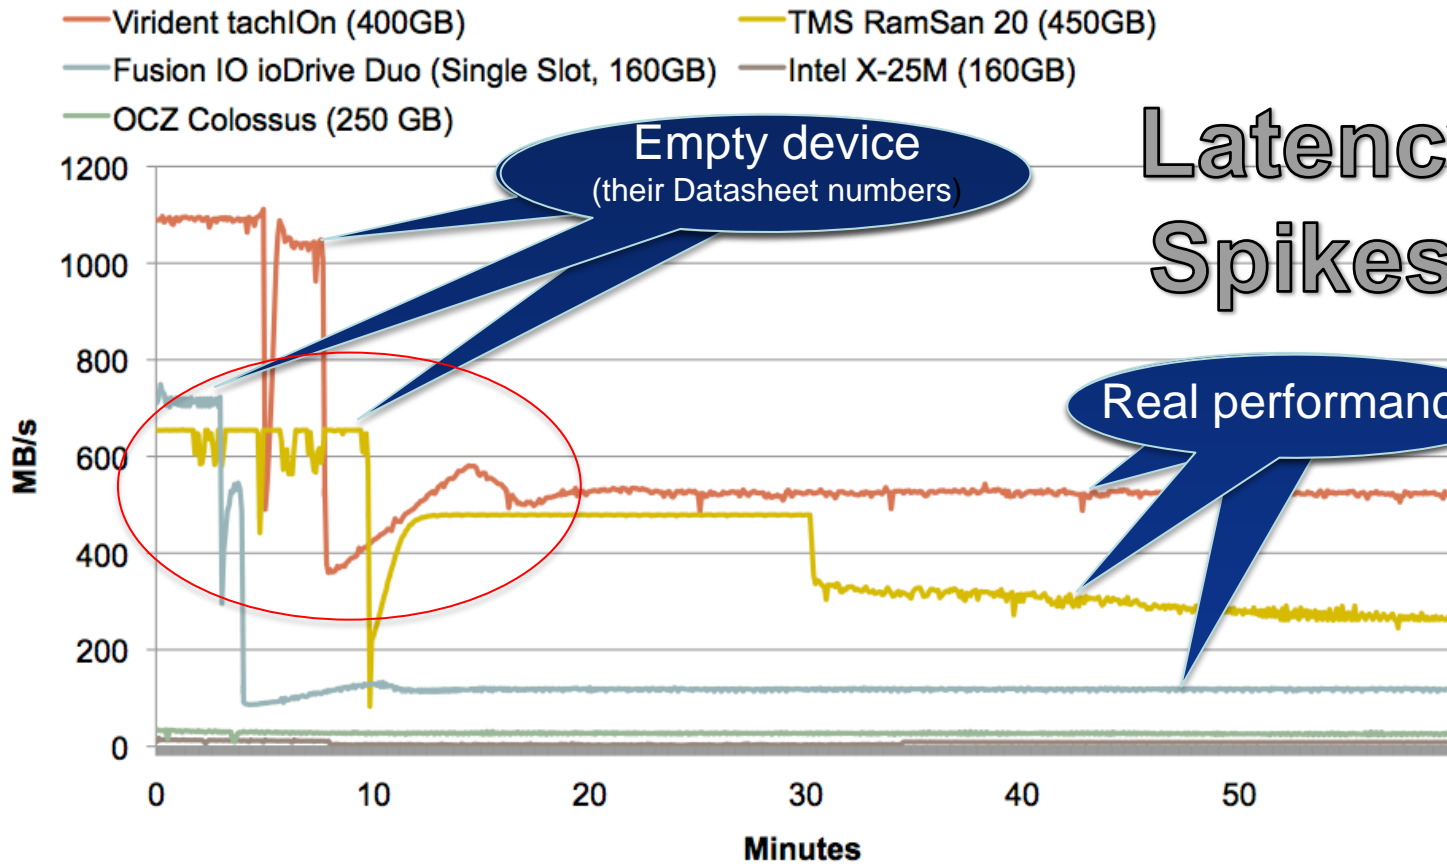
# The Infamous SSD "Write Cliff"

The elephant in the room everyone tries to ignore

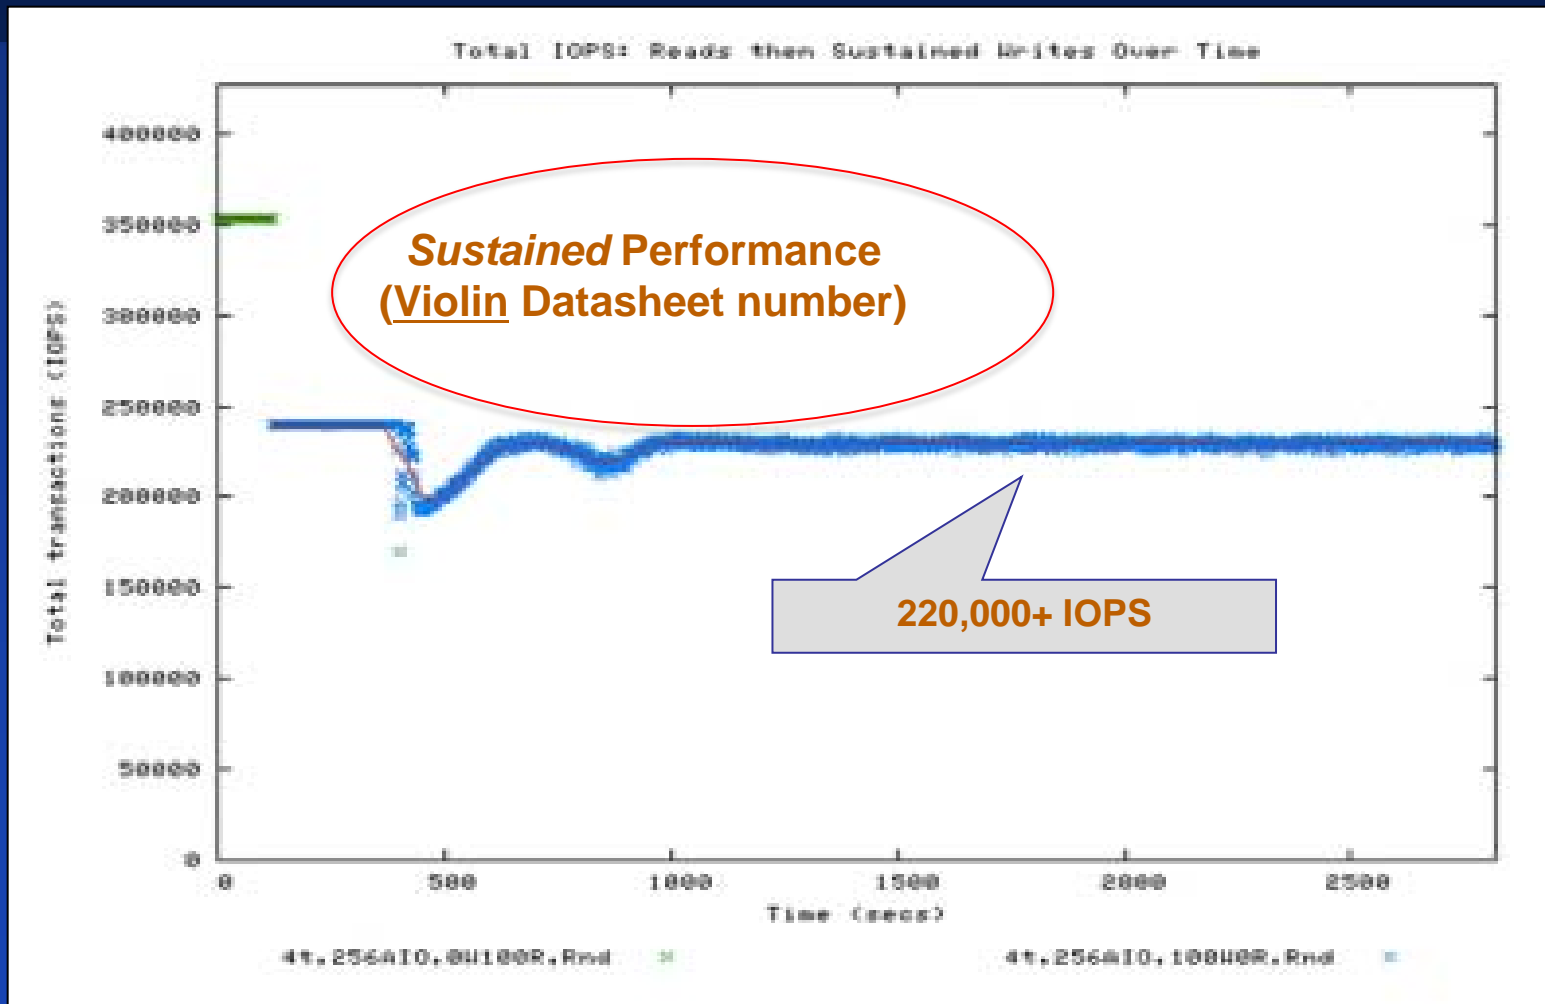
Flash Memory Summit



## Degradation - Bandwidth



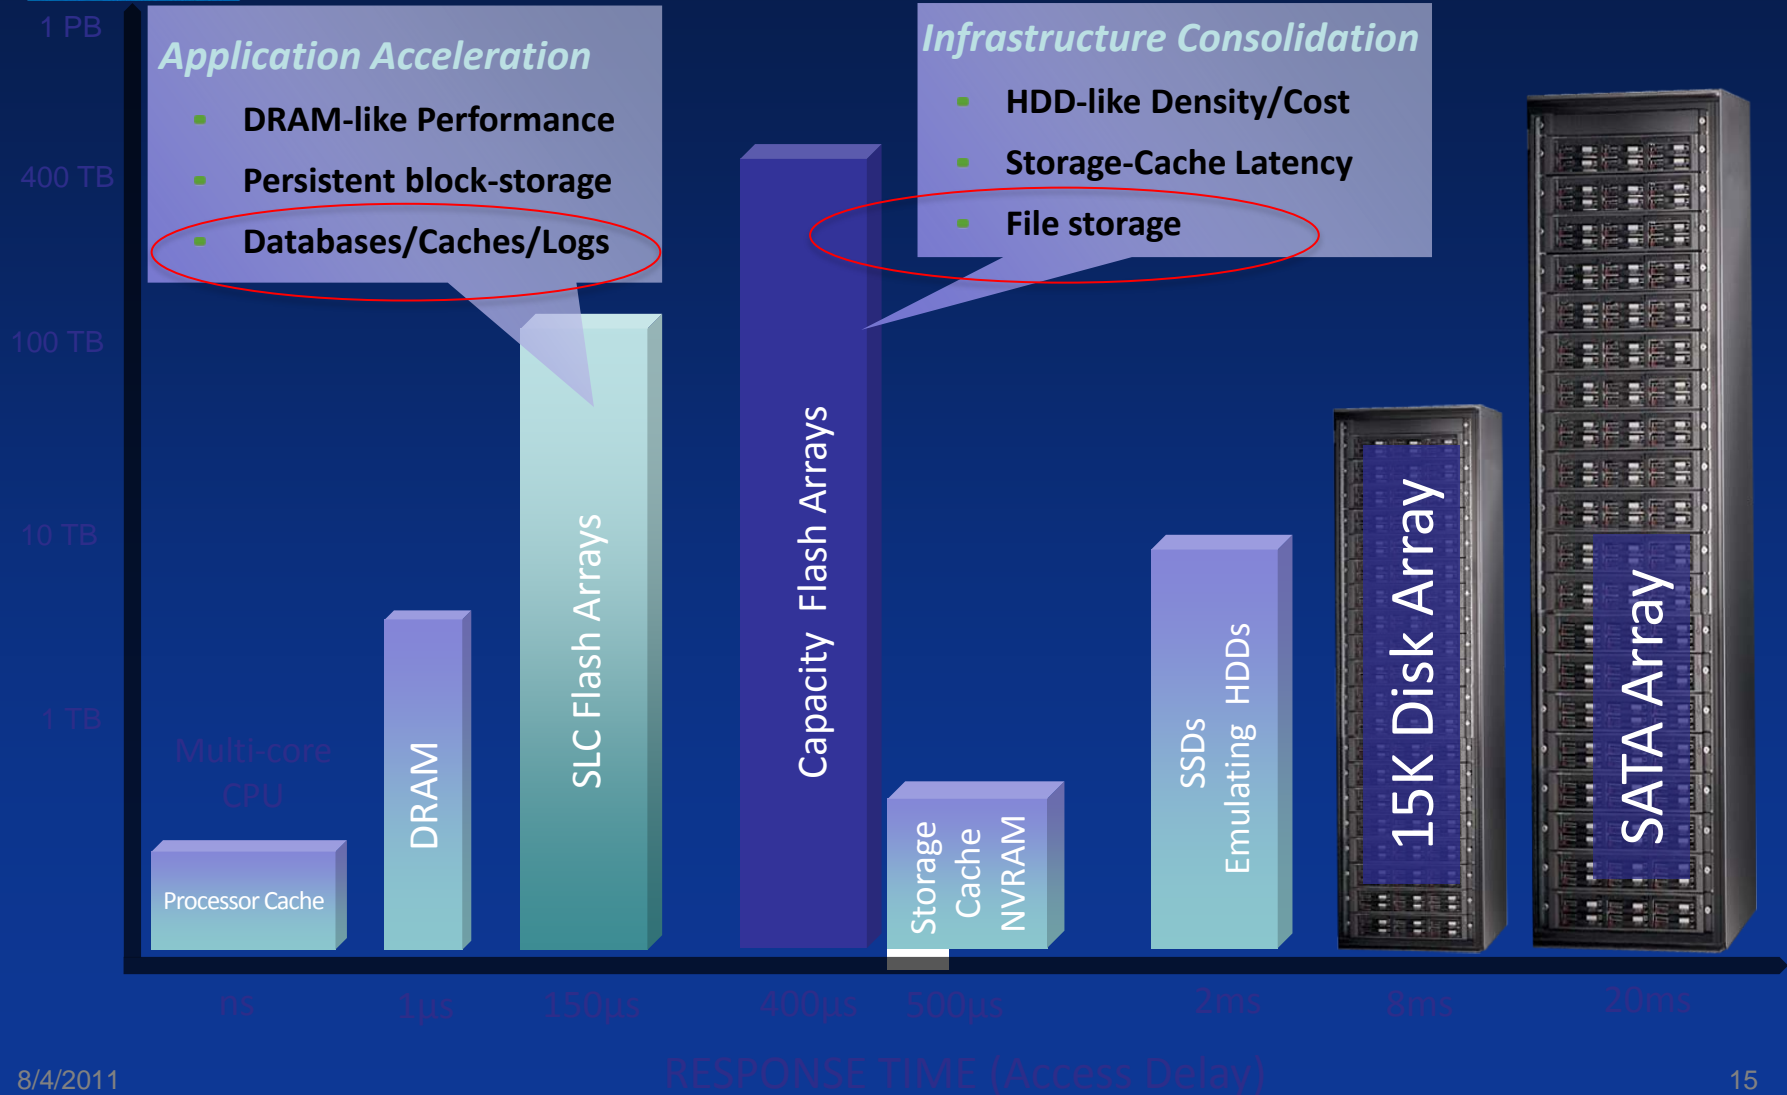
# Violin – Sustained performance





# Enterprise Use Cases

# Tiered Storage 2.0





# Transaction Processing

ORACLE®

Microsoft®  
SQL Server™ 2008

IBM® DB2®



MySQL™



Co-exist with Legacy HDD Systems



Co-exist with Legacy HDD Systems






100's of servers



Tape Archive

146 GB 15K    600 GB 15K    2 TB SATA 15K

# Transaction Processing



Flash  
Memory  
Arrays



Short-Stroke  
146-600GB 15K  
FC disk

400-600 GB  
FC disk

400 to 600  
GB FC disk

2-4 TB SATA/  
SAS disk

60 GB  
tape

Move high  
performance  
transactions  
to Flash  
Memory  
Arrays

High IOPs  
Low Latency  
>Server Utilization  
> IOPs/sq. foot

Fully  
Utilize  
Disk  
Capacity

OLTP

DW/ODS


Nearline

Archive

# Multi-Tenancy

- Max Availability, Isolation, Utilization


Big Host



Little  
Hosts



Each customer  
gets their own  
partition




Combine  
containers:  
Max HA & I/O

Or Container  
Level Isolation

2 partitions  
are HA with 2  
PCI-E each

# OLTP, DW, ODS

- Net Benefit: Analytics For Big Data





# Extending the Use of Flash.....

## Facilitates:

- Movement to High End Commercial Data Center usage
- Next evolutionary step beyond capabilities of SSD and Flash PCIe boards
- Extend Benefits of Flash beyond current performance and latency benefits

## Enablers:

- Scalability
- Share-ability
- Manageability
- I/O
  
- Sustained Writes
- Hot Swap
- HA
- RAID
- Fail-in-place
- Remote mgmt.
- Partitions

## SNMP

- Interface - System and network mgmt
- Ex: HP NNM and IBM Tivoli tools

## Array mgmt

- Wear mgmt
- 5 Year MLC lifetime under std maintenance agreement

## REST API

- Interface to proprietary provisioning systems
- XML interface to management systems

## Remote Admin

- Single Web GUI & CLI
- XML API & SNMP
- Email alerts
- Single multi-PB image




# Business Benefits





# Application Acceleration w/HP

|   |  |  |
|---|--|--|
|  |  |  |
| <h2>OLTP Results November, 2010</h2>  |  |  |
| Total System Cost:  | Transactions/Min                           | Price/Performance                          |
| \$2,126,304<br>(\$900,000 = Oracle SW)  | 3,388,535                                  | <b>\$0.63</b> (per transaction per minute) |
| Processors/Cores  | Database Manager                           | Operating System                           |
| 8/64<br>Intel Xeon 2.26 GHz   | Oracle Database<br>11g Rel 2<br>Enterprise | Oracle Linux Basic<br>TUXEDO 11gR1         |

**70% Reductions**

- ✓ Cost
- ✓ Rack space
- ✓ Power
- ✓ Response time



+



=

**Open Architecture**  
*Scales Linearly*

HP ProLiant DL980 G7

**Database Options:**

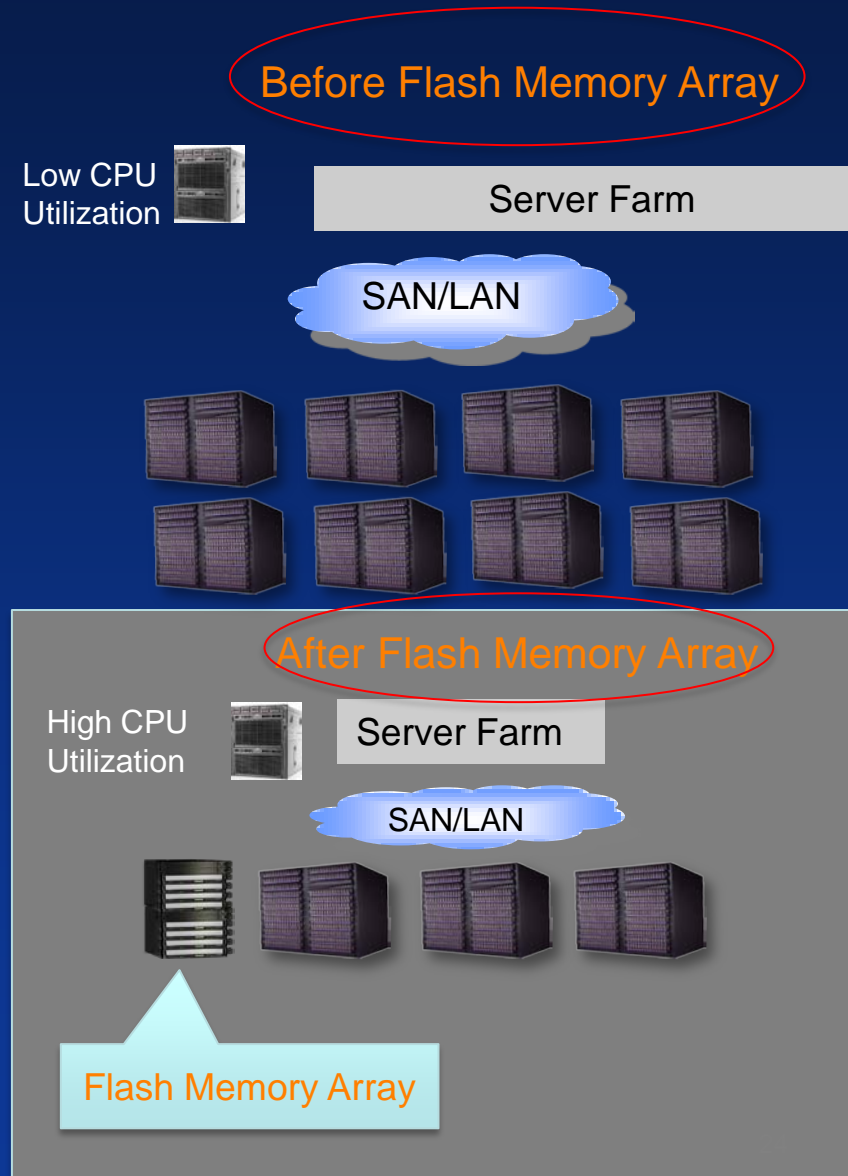
- Oracle 8/9/10/11/RAC
- MS SQL Server
- Sybase + Others

Flash Memory Array

**\$0.63 with Flash RAID**  
vs. \$2.40 (Oracle Exadata 2)  
or \$1.01 without RAID  
(Oracle SuperCluster 2011)

# Key Business Benefits

- **Application Acceleration**
  - Meet & Exceed SLAs
  - Simpler System Architectures
  - Deploy new apps faster
  - Reduce tuning costs
- **Infrastructure Consolidation**
  - Reduce CapEx and OpEX
  - Fewer Spindles, licenses, servers
  - Less Power, space, service
  - Leverage existing infrastructure
  - Enable Virtualization
- **Lower \$ per Application**



# Data Center Transformation

*“The transition from spinning to solid-state storage is already underway.”*

Steve O'Donnell, ESG



- Resource utilization
- OpEx Reduction
- Reliability
- Availability
- Serviceability
- Power
- Space
- Cooling



# Key Take Always

- Flash Memory Arrays:
  - Suitable for High End Enterprise Applications
  - Meet Enterprise Application requirements\*\*

\*\*Summary of requirements:

Flash Performance  
Reliability,  
Serviceability  
Manageability

Consistent low response time  
Availability,  
Scalability  
Resource Utilization



# Appendix



# Flash Memory Array Characteristics

| Category      | Characteristic (8 racks)   | Uses                                   |
|---------------|--|--|
| Scalability*  | 2 + PB   | Large Active Data Sets                 |
| IOPS**        | 64,000,000   | Migrate from short-stroked 15K FC HDD  |
| Bandwidth**   | 400 GB/sec read<br>256 GB/sec write  | Excellent ingest and data distribution |
| Latency       | 25 $\mu$ s write<br>75 $\mu$ s read  | Max server utilization                 |
| Availability  | HA and RAID  | High end applications                  |
| Manageability | XLM/SNMP interfaces  | High end applications                  |
| Protocols     | FC, iSCSI, IB (Q3), NFS  | Multiple environments                  |
| I/O           | (512) 8 Gbit FC ports or<br>(512) 10 GbE ports<br>(64) 40 GB/sec IB ports (Q3) | Max resource utilization               |

\* Raw \*\* Theoretical



# Compelling Economics

| Performance Per Rack | Flash Memory Arrays | Conventional HDD Arrays | HDD/SDD Combination |
|----------------------|---------------------|-------------------------|---------------------|
| IOPS                 | 2,000,000*          | 24,000                  | 40,000              |
| Latency              | 200 $\mu$ sec       | 5000 $\mu$ sec          | 2000 $\mu$ sec      |

| Cost per Application | Flash Memory Arrays | SATA/SAS | FC      |
|----------------------|---------------------|----------|---------|
| \$/IOPS (4K)         | \$1.00              | \$17.00  | \$20.00 |

| Cost per GB Flash | Flash Memory Arrays | RAID-1 SSDs in Array | PCIe Flash in Mirrored Systems |
|-------------------|---------------------|----------------------|--------------------------------|
| \$/GB with RAID   | \$22.00             | \$100 - \$200        | \$60.00                        |

\* Based on one rack with 8 memory arrays





# Flagship Customer

## 600+ Terabytes and counting

**Problem:** ORACLE Ad Server Reporting only met 8 hour SLA twice in 6 months

**Goal:** consistent sustainable IO performance to meet SLA under EMC's Enterprise Storage management tools

**Result:** On Violin Arrays without any tuning, haven't missed SLAs

AOL is now able to further enhance their Ad Campaign Reporting

- Reinforcing what works, pruning what doesn't
- Potential for positive revenue impact going forward

## AOL was one of EMC's VPLEX key launch customers

- Global production prior to official launch by EMC
- Significant amount of VPLEX support matrix was validated @ AOL
- Violin 3200 Memory Array certified under EMC VPLEX
  - Winning combination of consistent sustainable performance under world-class enterprise management system
- VPLEX certification enables Violin's products to be seamlessly used in EMC environments



## HP & Microsoft - Best of Breed



### TPC-E Blade server world Record – June 2010

- This is the first use of non-HP storage in an HP TPC benchmark
- Flash Memory Arrays only operating at 35% utilization
- Other HP benchmarks due shortly
- The TPC-E benchmark simulates the OLTP workload of a brokerage firm. The focus of the benchmark is the central database that executes transactions related to the firm's customer accounts. Although the underlying business model of TPC-E is a brokerage firm, the database schema, data population, transactions, and implementation rules have been designed to be broadly representative of modern OLTP systems.



# Thank You