Memory Channel Storage™ (MCS™) Demystified

Jerome McFarland
Principal Product Marketer

diablo technologies™ | PROBLEM SOLVED.

AGENDA

+ INTRO AND ARCHITECTURE
+ PRODUCT DETAILS
+ APPLICATIONS

THE COMPUTE-STORAGE **DISCONNECT**

+ Compute And Data Have **Far Outgrown Storage Advancements**

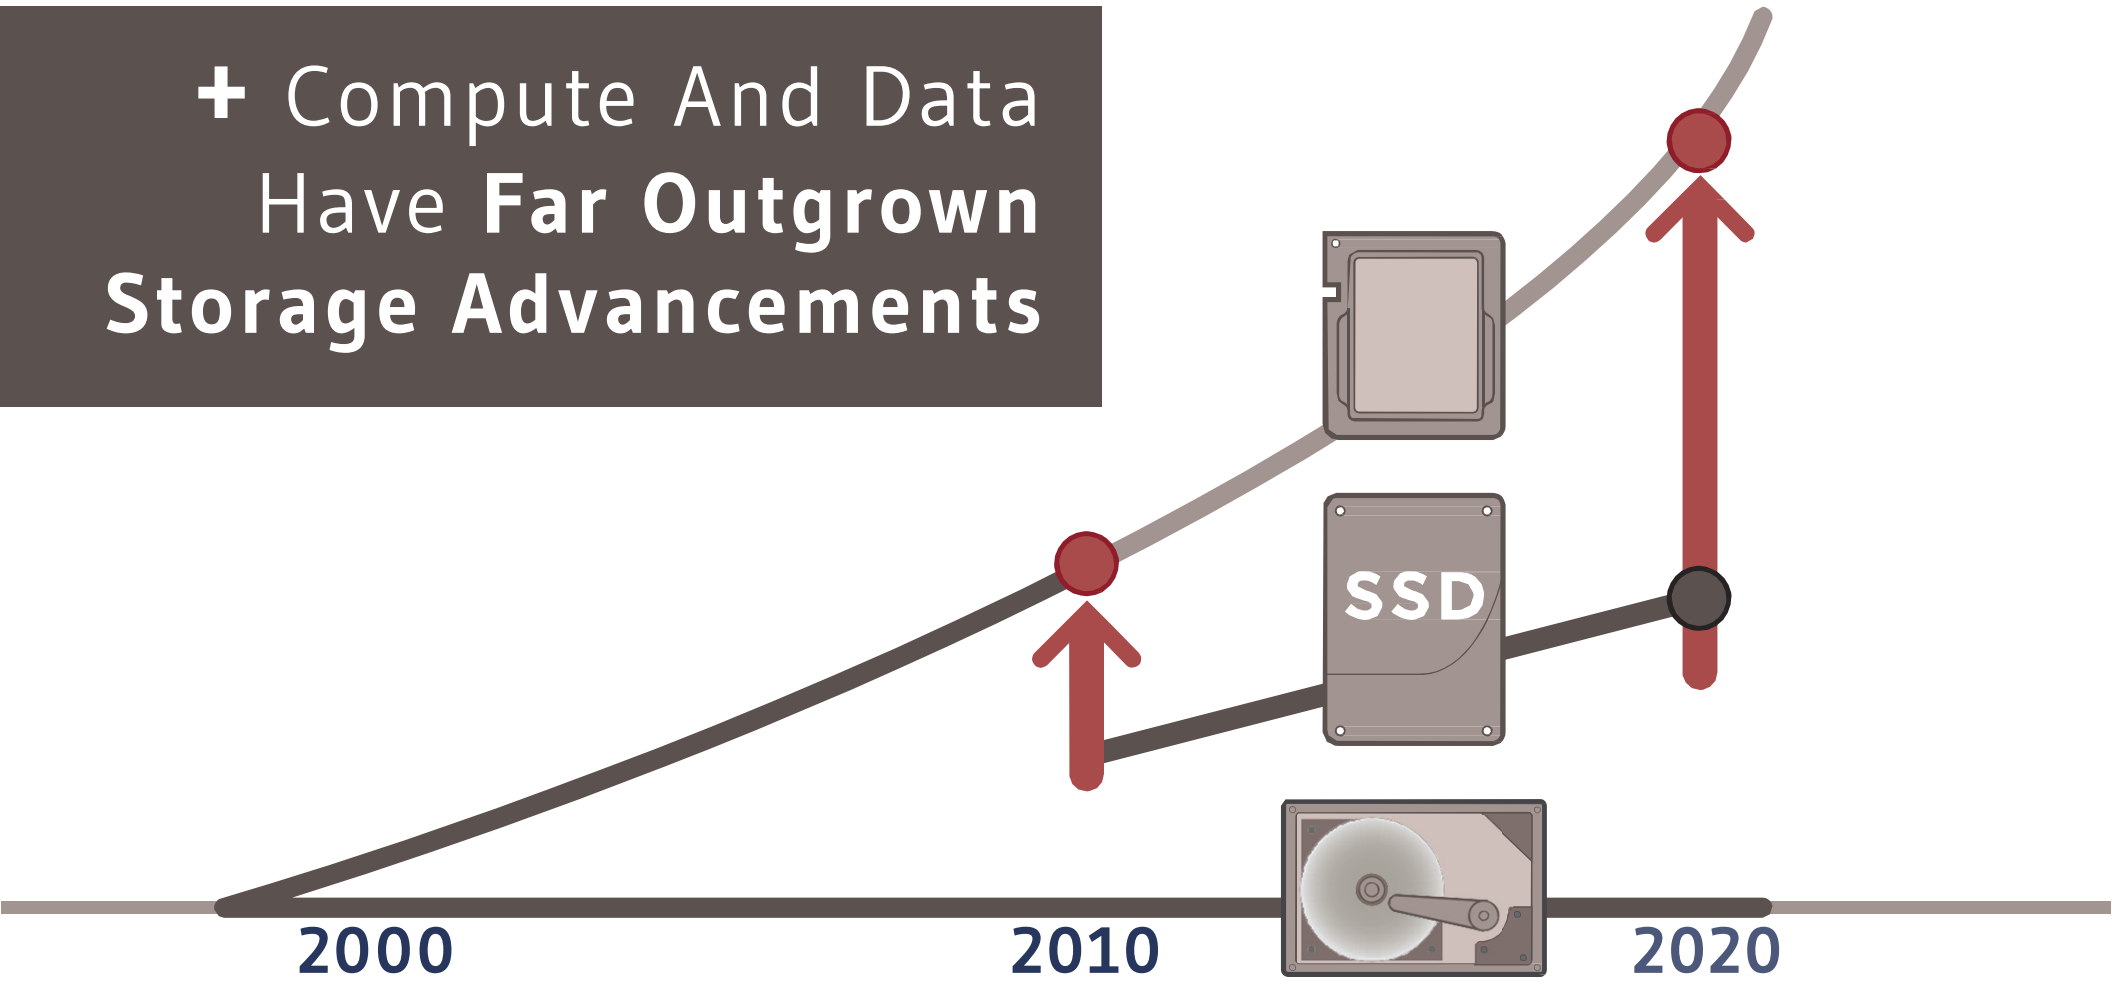

Enterprises need a solution to close the gap…

FLASH STORAGE EVOLUTION THUS FAR

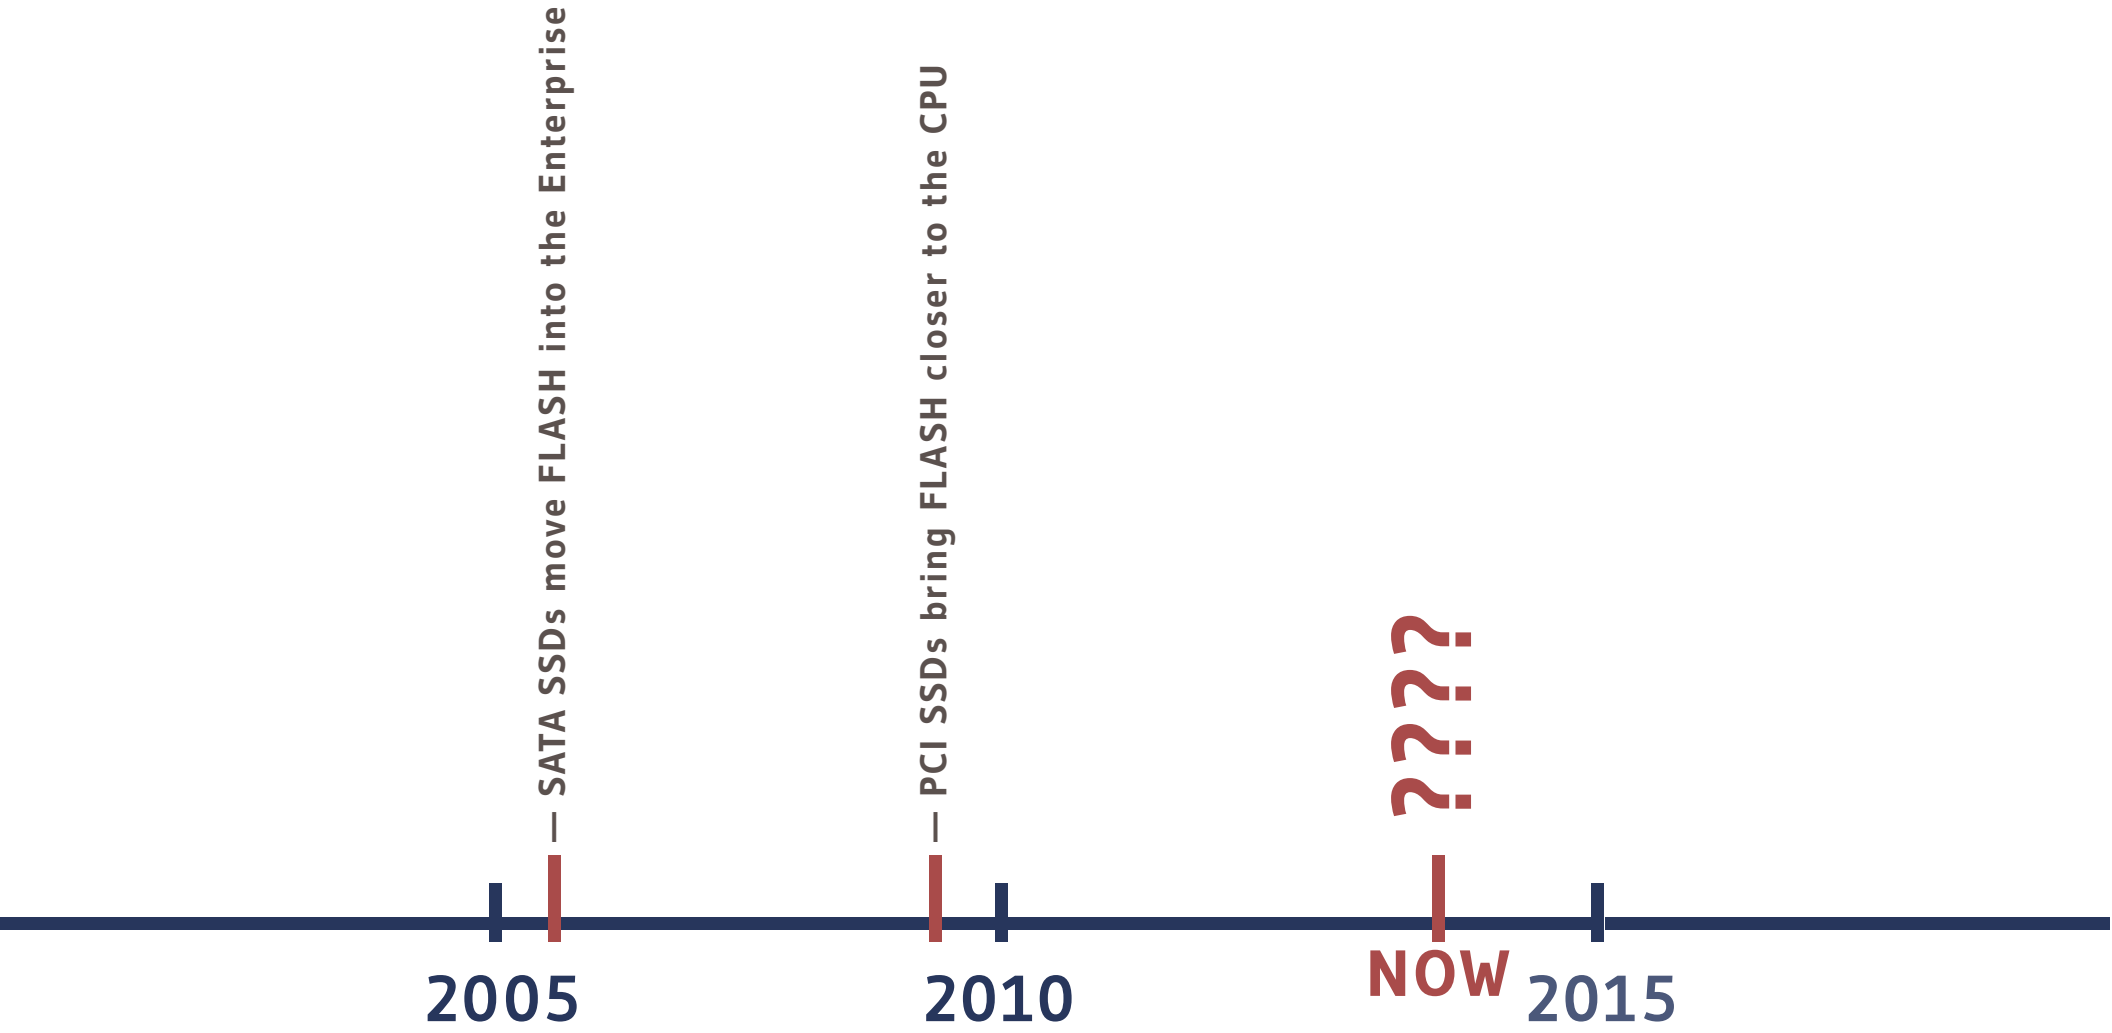

TODAY'S SOLUTION…

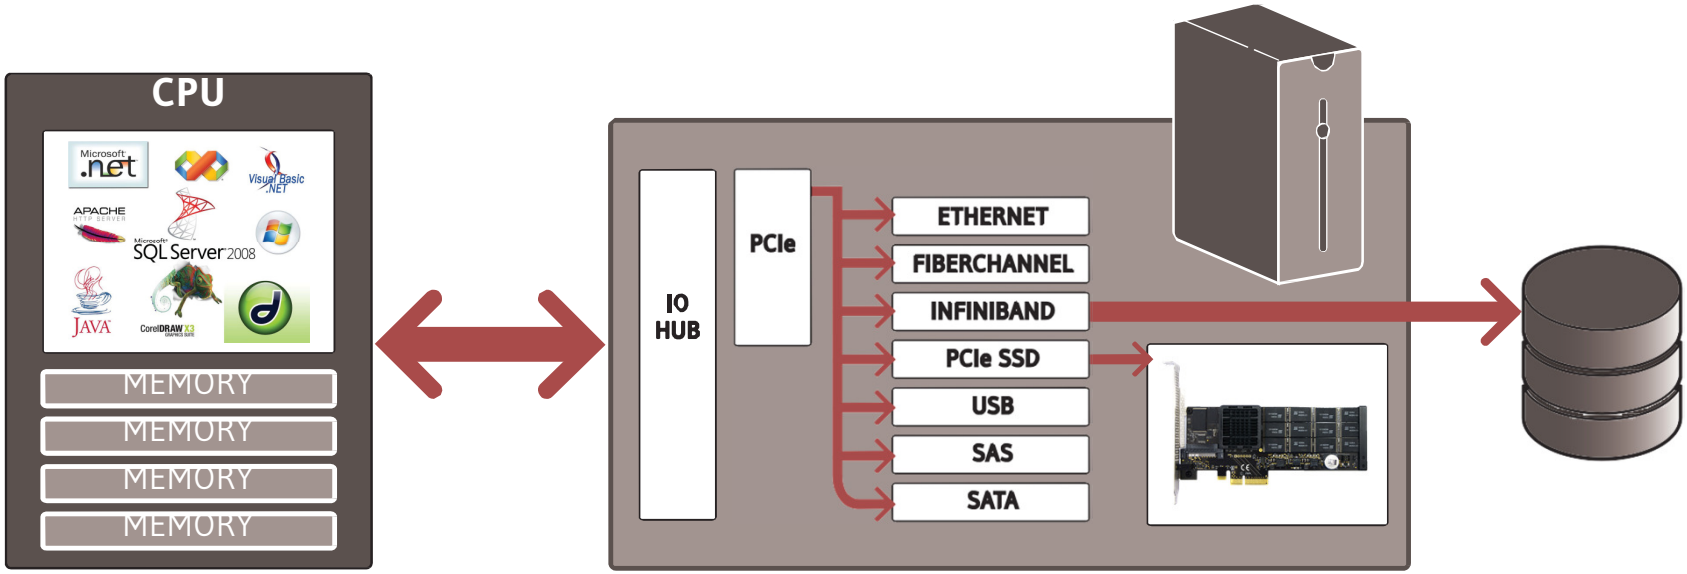

+ **Processor, memory and applications are tightly coupled**
+ **Storage not local to processor and application execution**
+ **SSDs deployed to minimize the disconnect, but critical issues remain:**
  - Long trips required for data retrieval
  - Resource contention
  - Response Time (Latency) suffers

WHAT IS THE IMPACT OF RESPONSE TIME (A.K.A. LATENCY)?

RESPONSE TIME → USER EXPERIENCE

SLOW RESPONSE TIME → APPLICATION LAG

FAST RESPONSE TIME → CRITICAL BUSINESS ADVANTAGES

DETERMINISTIC RESPONSE TIME → QUALITY OF SERVICE

THE PERFORMANCE TRADE-OFF

+ Traditionally customers have faced a suboptimal trade-off in storage system design:

OPTIMIZE LATENCY → SACRIFICE IOPS

OR

OPTIMIZE IOPS → SACRIFICE LATENCY

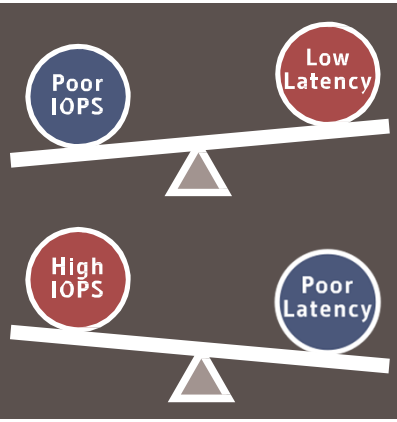

A Painful Workaround…

+ When SSD "IOPS vs. Latency" trade-offs are unacceptable, adding expensive RAM is a traditional recourse.

+ However, adding RAM can create an imbalance between incremental performance requirements and rapidly growing solution cost.

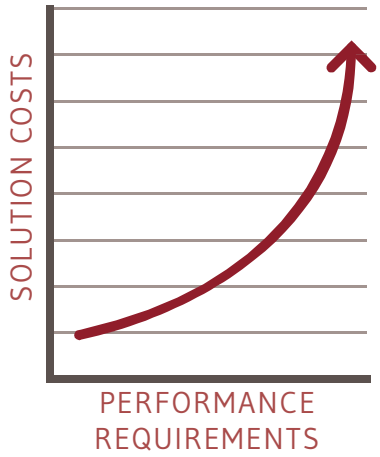

MEMORY CHANNEL STORAGE SOLUTION

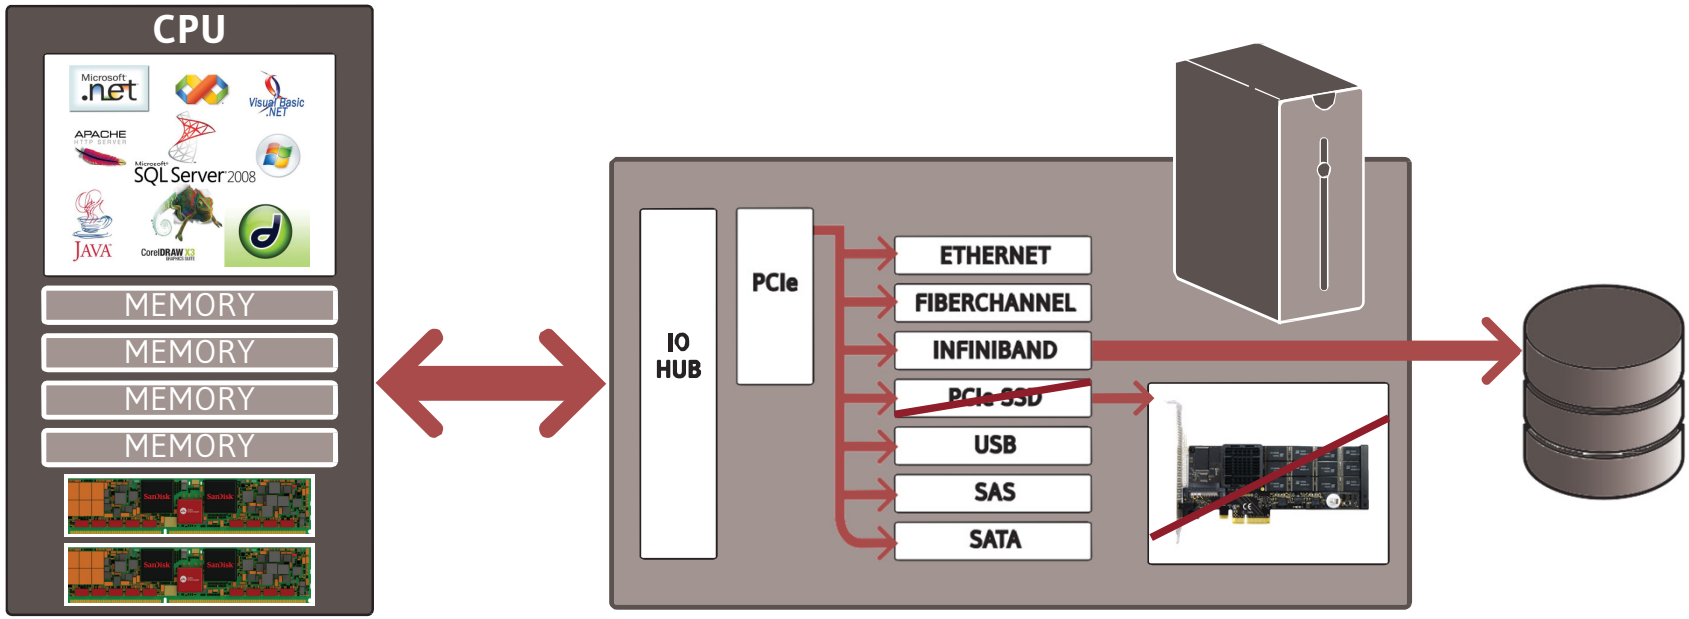

+ MCS is coupled with the processor, application, and system memory
+ Distance and contention issues are eliminated
+ Data stays within memory subsystem for local access
  + **Achieves ultra-low latency**
  + **Enables linear "performance vs. cost' scalability**
+ Enabled by a unique architectural approach

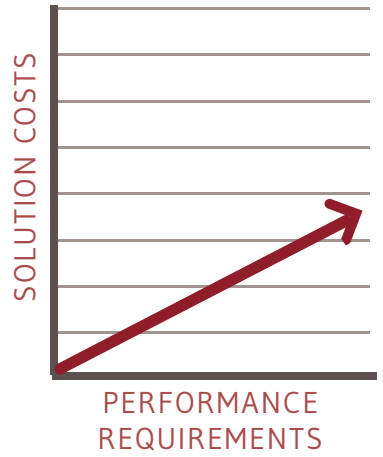

MCS SYSTEM VIEW

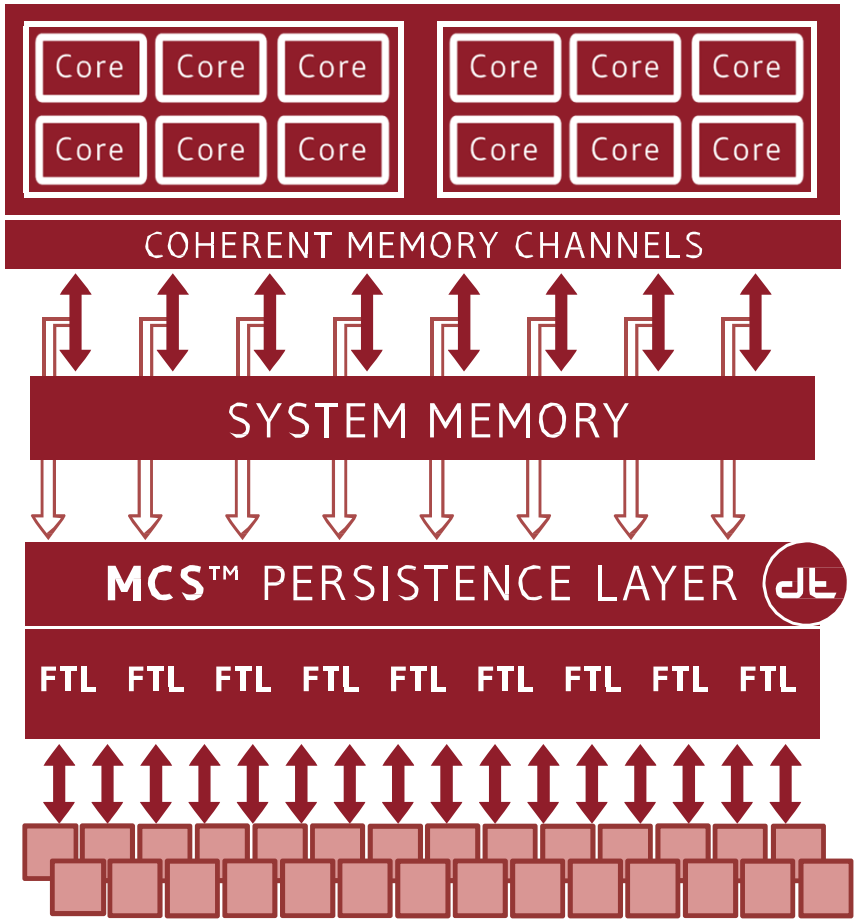

Leveraging the **Power of Parallelism...**

+ Massive Flash capacity exposed through the low-latency memory subsystem.

SOFTWARE ARCHITECTURE

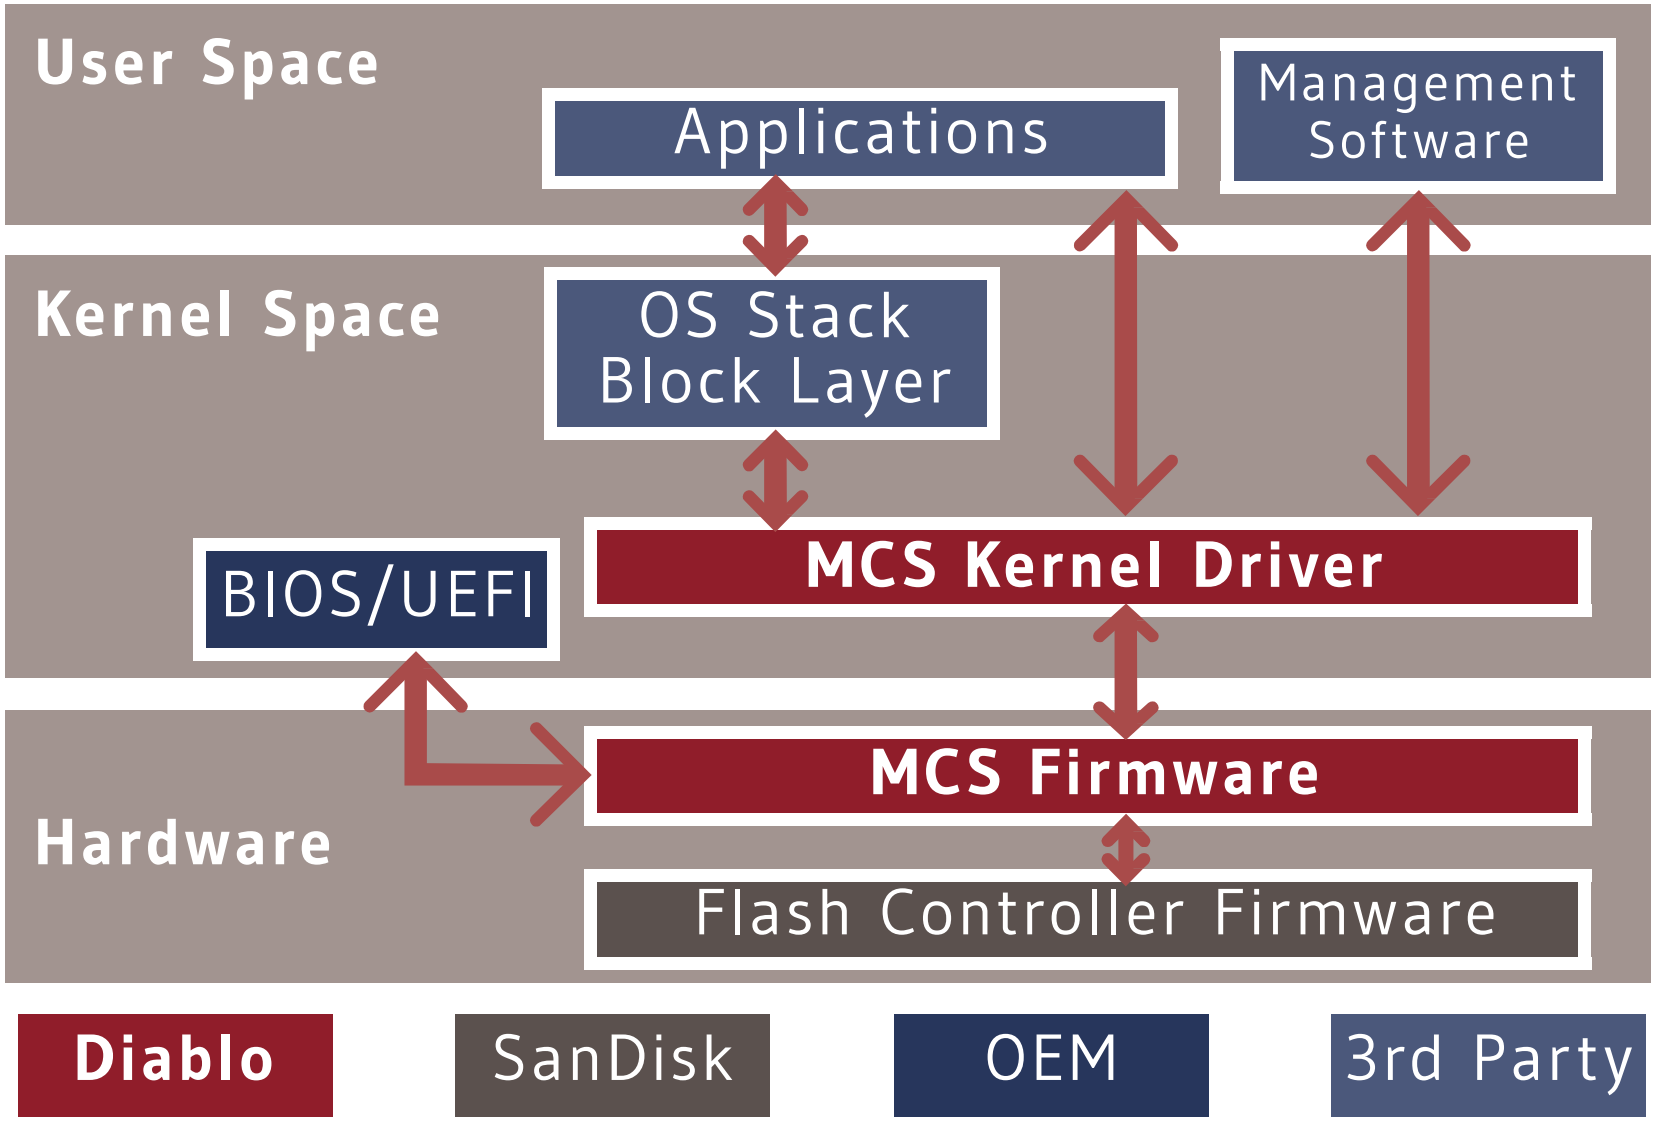

HARDWARE ARCHITECTURE

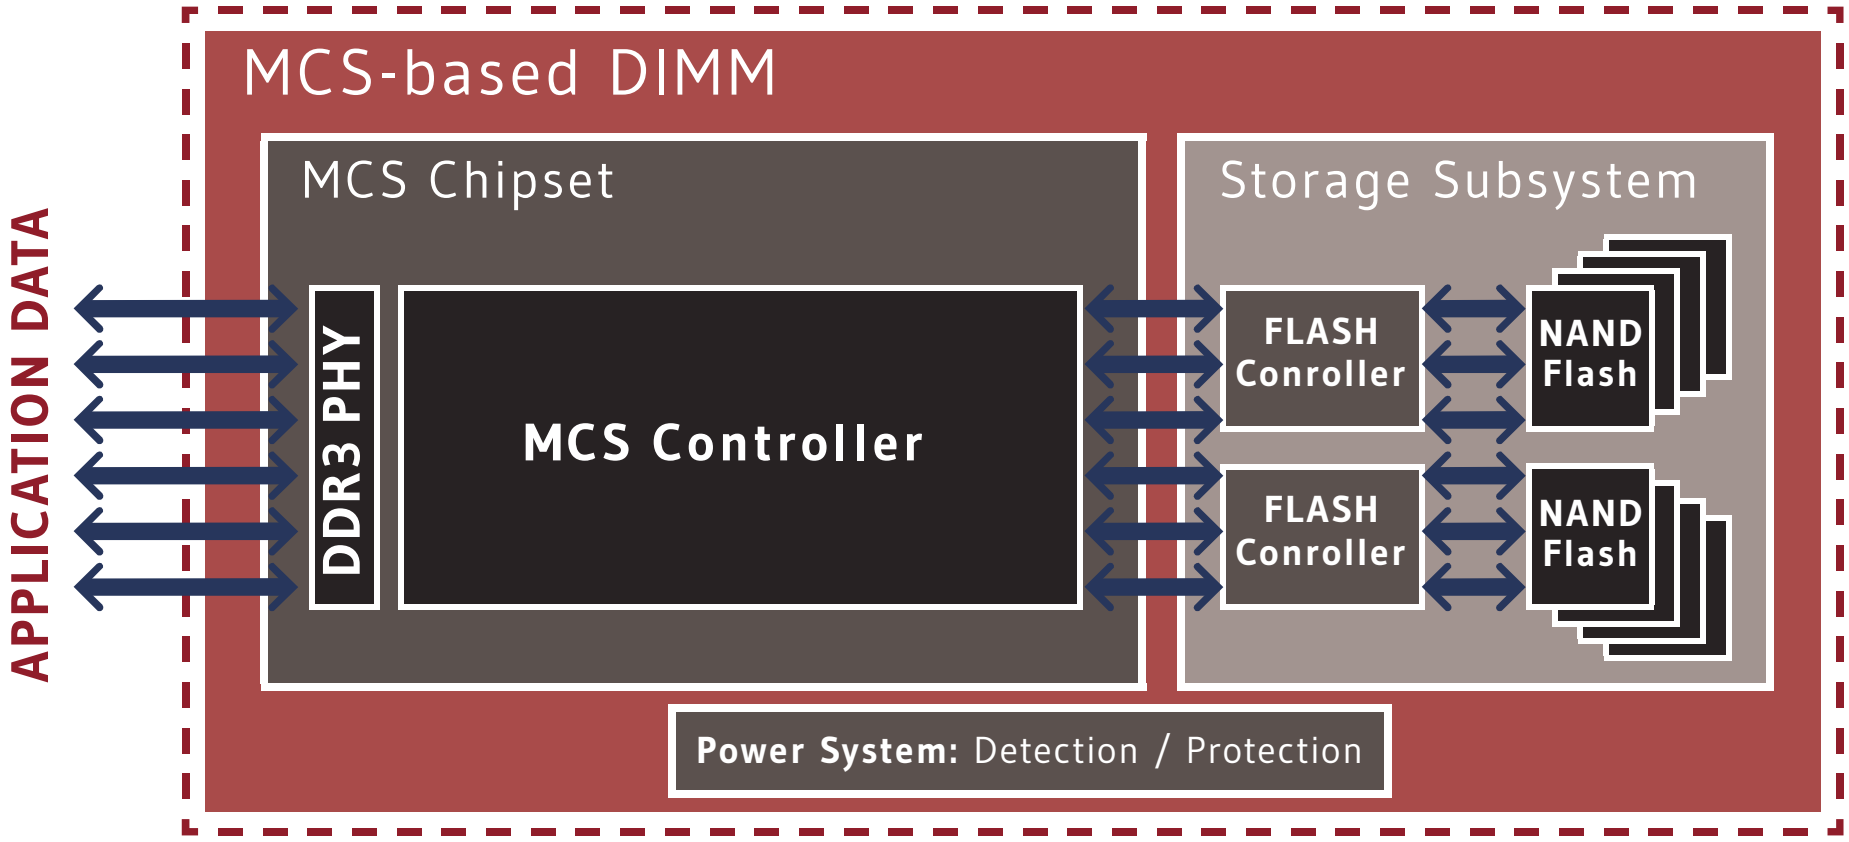

PERSISTENT MEMORY
WITHIN NUMA

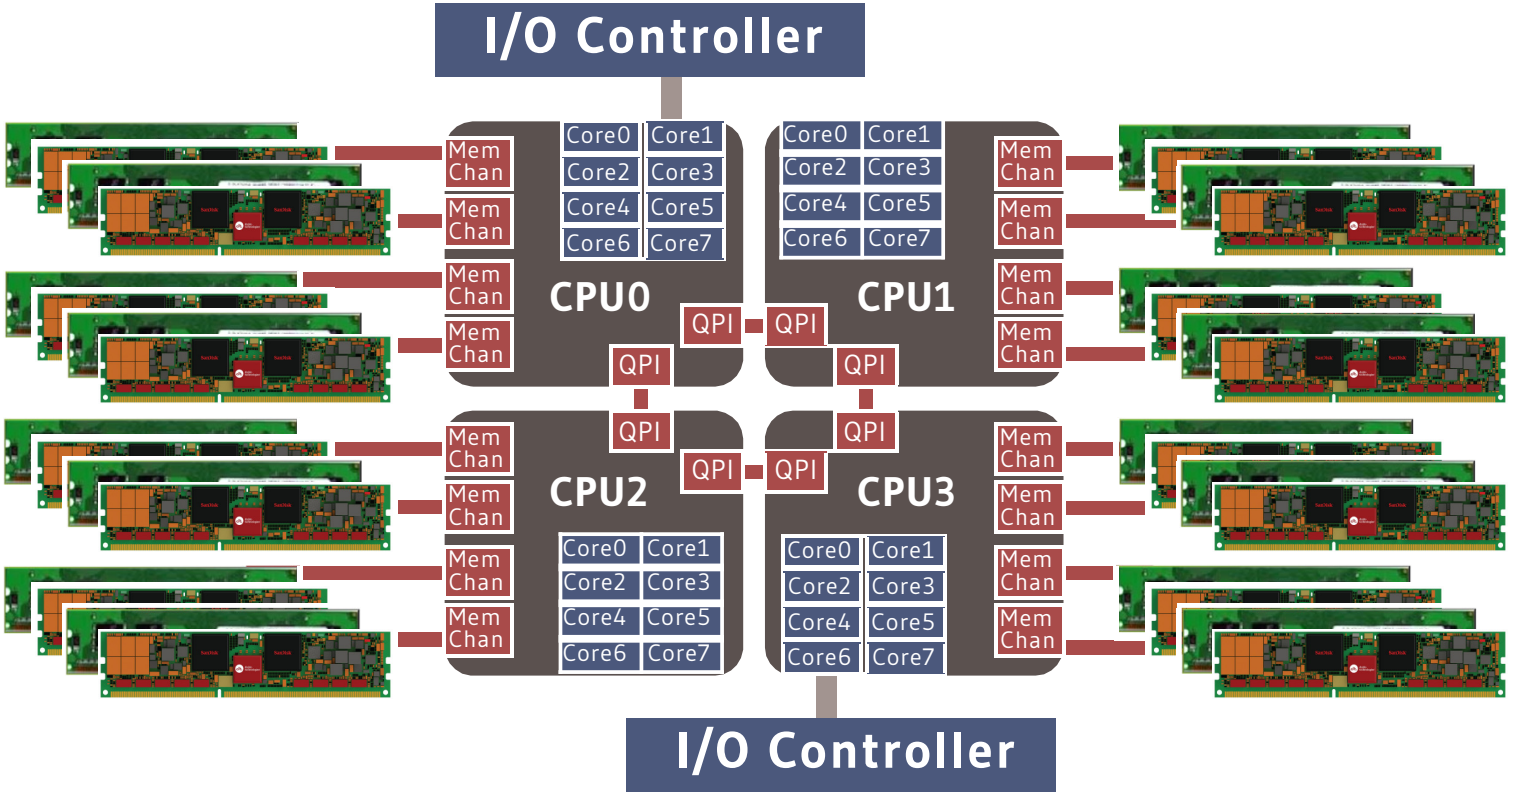

+ The best location for application data is within the NUMA architecture
+ Highly parallel (threaded) applicatons running on parallel processors and cores, with highly parallel memory and storage access

SO, WHAT IS **MEMORY CHANNEL STORAGE**?

- An Architecture (not a single product)
  - Enables Flash Storage to Directly Interface on the Memory Channel
- Presents as a Block I/O Device
  - Can be Managed just like Existing Storage Devices
- DDR3 Interface, Standard RDIMM Physical Form Factor
  - Plugs into Standard DIMM Slots
  - Self-contained, No External Connections Required

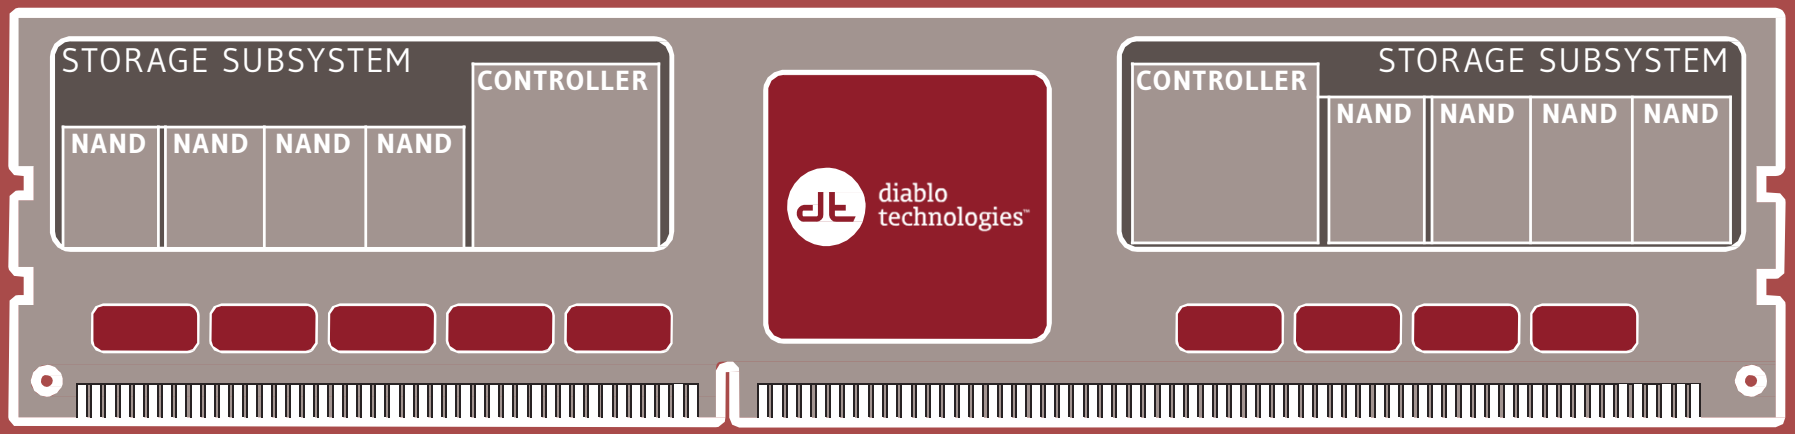

SYSTEM REQUIREMENTS & COMPATIBILITY

+ Hardware and BIOS Requirements

- + Server enabled with MCS UEFI BIOS modifications
- + DDR3-compatible processor
  - + MCS is compatible with standard JEDEC-compliant 240-pin RDIMMs
  - + Supports DDR3-800 through DDR3-1600
- + 8GB of standard memory (RDIMM) installed in the system
- + MCS follows standard server DIMM population rules

+ Initial OS Support

- + Linux (RHEL, SLES)
- + Windows Server
- + VMware ESXi

vmware

Microsoft

TECHNOLOGY COLLABORATION TO CREATE **THE FIRST MCS-ENABLED PRODUCT**

diablo technologies™

+ Reference architecture design
+ DDR3 to SSD ASIC/firmware
+ Kernel and application level software development
+ OEM System Integration and enterprise application domain knowledge

+

SanDisk®

+ Guardian Technology for enterprise applications
+ SSD controller & FTL firmware development and test
+ Supply Chain and Manufacturing with flash partner
+ System Validation

SanDisk™ "ULLtraDIMM™" POWERED BY **MCS**

MEMORY CHANNEL INTERFACE

+ DDR3 PROTOCOL
+ CONFIGURABLE AS BLOCK DEVICE
+ STANDARD RDIMM FORM FACTOR

GUARDIAN™ TECHNOLOGY

+ 19nm MLC
+ 10 DRIVE WRITES PER DAY
+ 5 YEAR WARRANTY

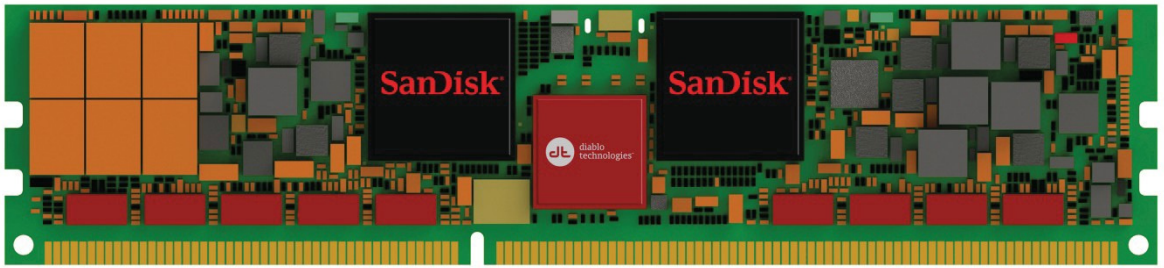

ENTERPRISE CLASS RELIABILITY

+ BACKUP POWER CIRCUITRY
+ END-TO-END DATA PROTECTION
+ 2M HOURS MTBF

SCALABLE & COST EFFECTIVE MEDIA

+ 200, 400GB CAPACITIES
+ SCALABLE ARCHITECTURE
+ 19nm MLC NAND

ADDITIONAL FEATURES

+ S.M.A.R.T. MONITORING
+ SUPPORTS TRIM
+ MAINTENANCE TOOLS

SUPERIOR PERFORMANCE
ACROSS WORKLOADS

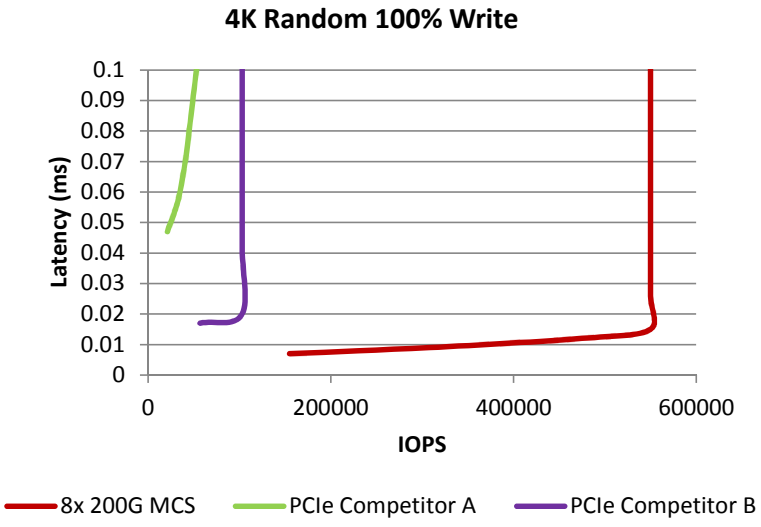

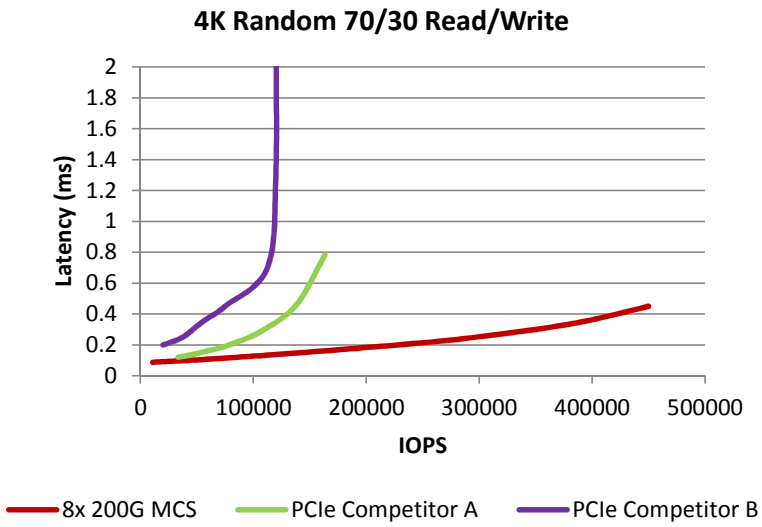

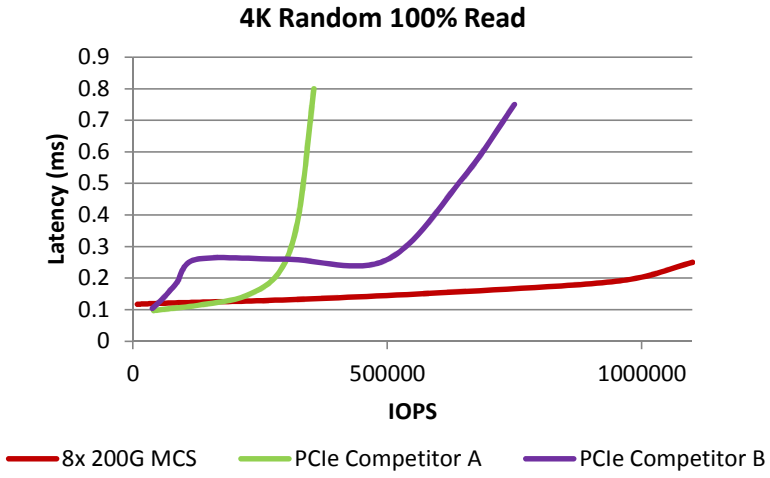

+ Enables standardization on a flexible, low-latency **platform**

Tested using MCS prototype modules.

IT'S ALL ABOUT THE **APPLICATIONS!**

LOW LATENCY APPLICATIONS

+ 3X MESSAGE RATE
+ 40X REDUCTION OF MAX LATENCY

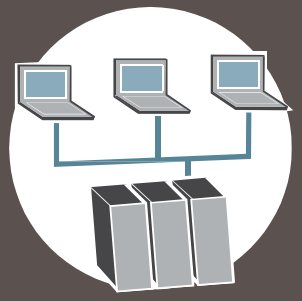

VIRTUAL DESKTOP

+ LOW σ = CONSISTENT USER EXPERIENCE
+ MEET QoS/SLA REQUIREMENTS

DATABASE/ CLOUD

+ 7X TPSE IMPROVEMENT
+ 3X REDUCTION OF AVERAGE LATENCY

BIG DATA ANALYTICS

+ MINIMIZE QUERY TIMES
+ EXTEND WORKING SET BEYOND RAM ALLOCATION

SERVER VIRTUALIZATION

+ 2X VMs PER NODE…
+ …USING 1/6 THE RAM PER VM

REDUCED LATENCY
ENABLES REAL-TIME ANALYTICS

60East TECHNOLOGIES + diablo technologies™

15% Read Mix

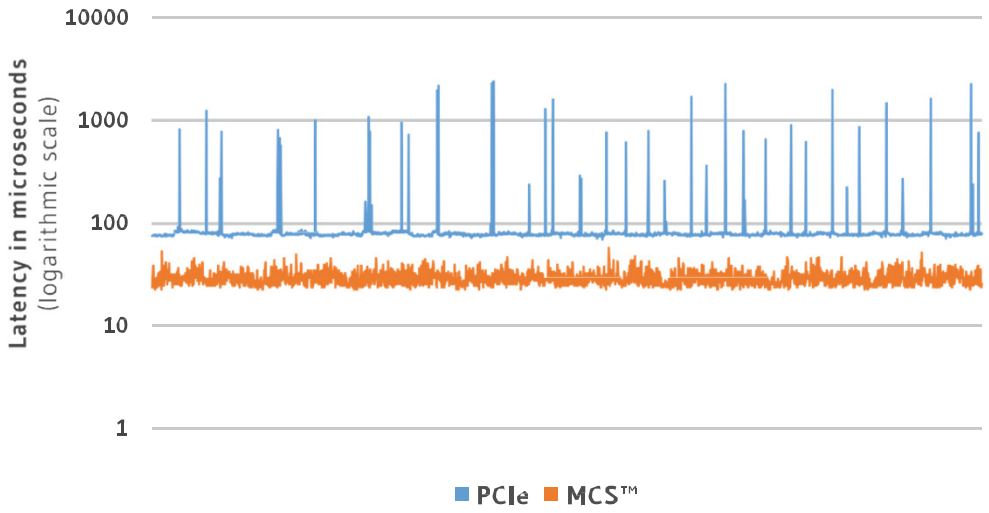

15% Read/Write Ratio Overview

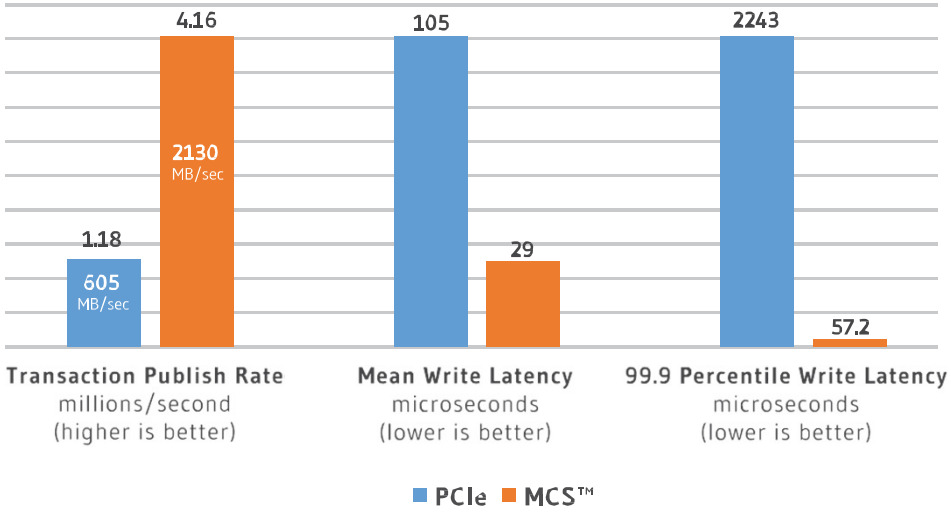

+ THE APPLICATION HAS BECOME THE BOTTLENECK IN E-TRADING

VIRTUAL DESKTOPS
SCALABLE ARCHITECTURE

+ 200 SERVERS, 8x SAS SSD
+ 3 GB RAM / VM

10,000 VIRTUAL DESKTOP DEPLOYMENT

+ 100 SERVERS, 4x MCS Modules
+ 1/2 GB RAM / VM

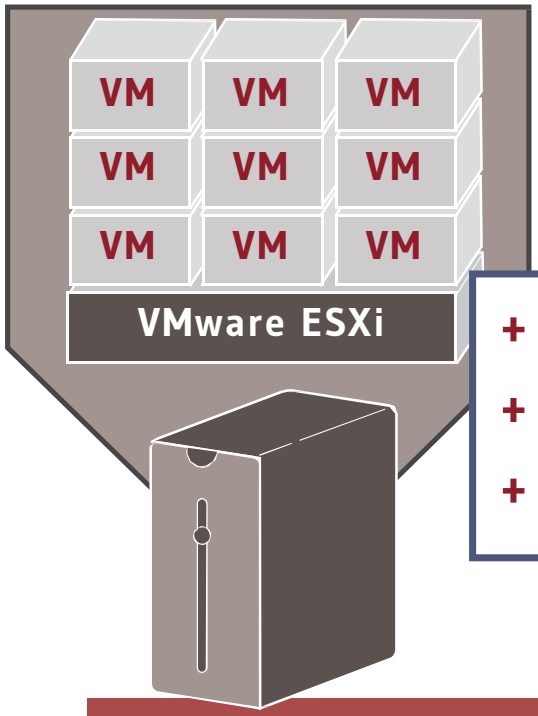

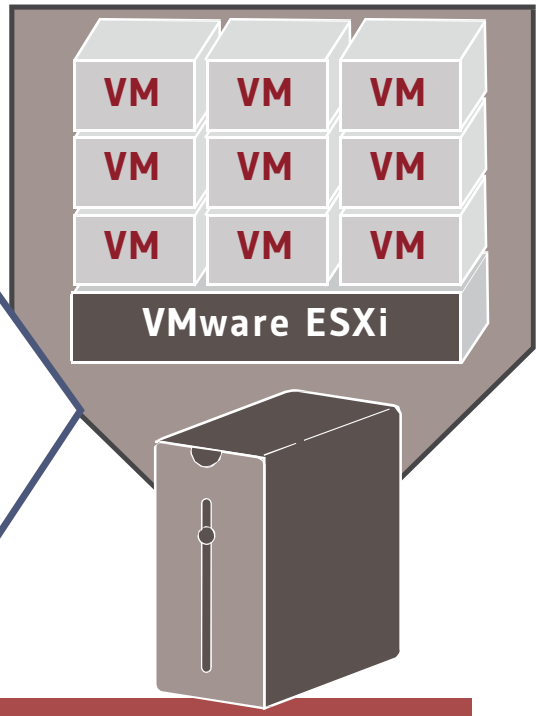

+ **2X VIRTUAL DESKTOPS** PER HOST
+ **75% DRAM REDUCTION** PER HOST
+ **MAINTAIN WORKLOAD PERFORMANCE**

+ IT'S NOW EASY AND COST EFFECTIVE TO ADD USERS AND SCALE INFRASTRUCTURE

MEMORY MAPPED **I/O ACCELERATION**

10 million records (20GB mmap) using synchronous msync calls

mmap Random Write: Write Latency Histogram

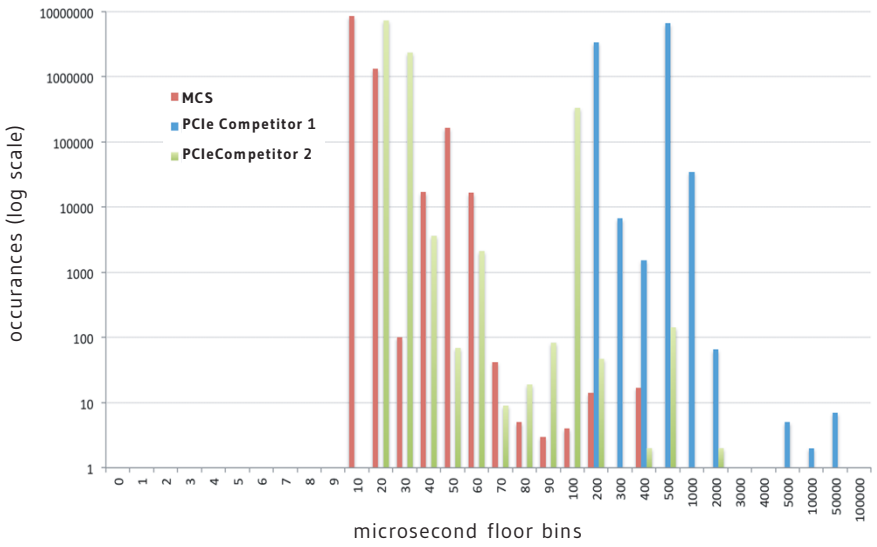

mmap Random Write: Write Latency Percentiles

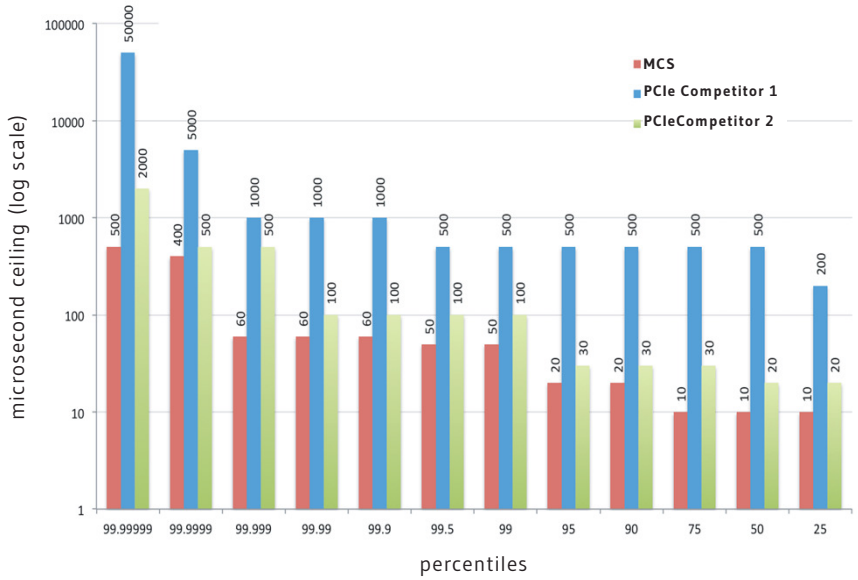

+ MCS 99th-percentile latency is **2x lower than Competitor 2** and **10x lower than Competitor 1**

+ MCS has the **tightest latency distribution**

TODAY'S
FRAGMENTED STORAGE SOLUTIONS

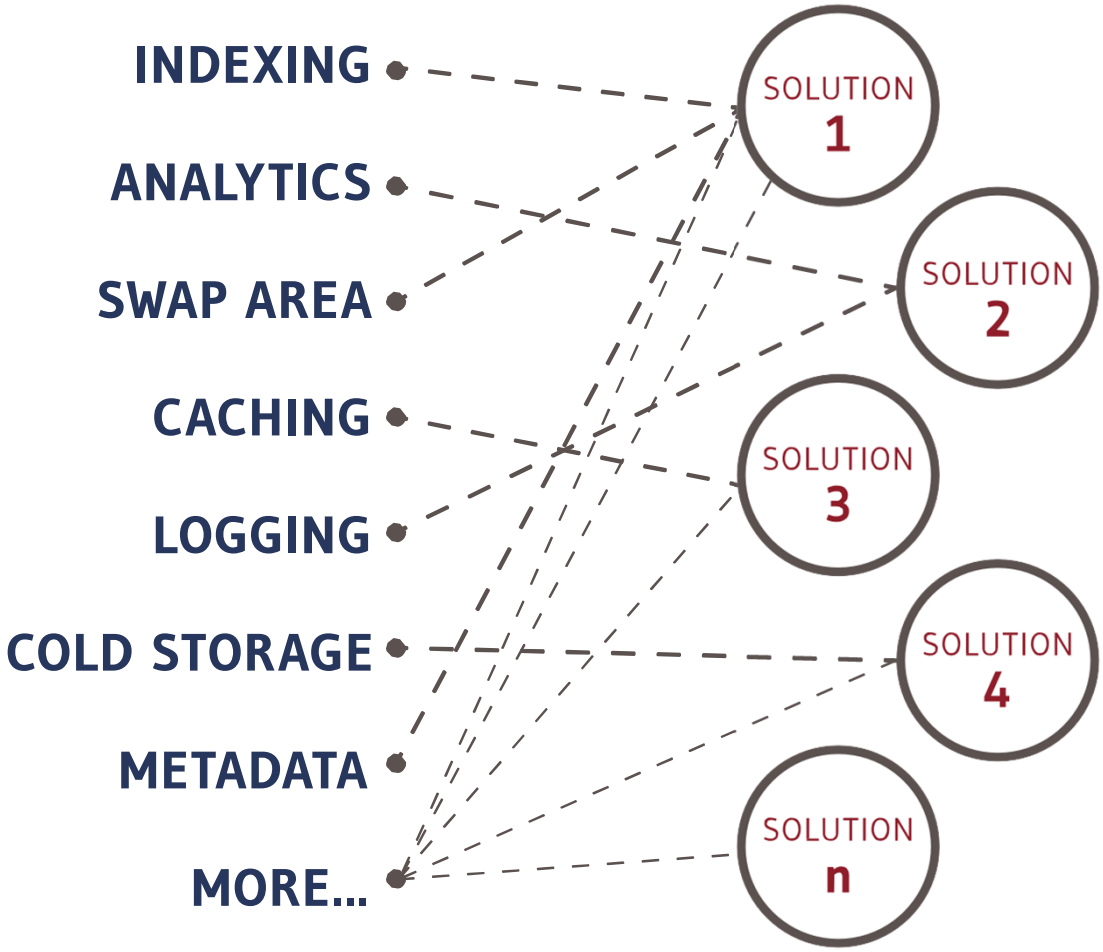

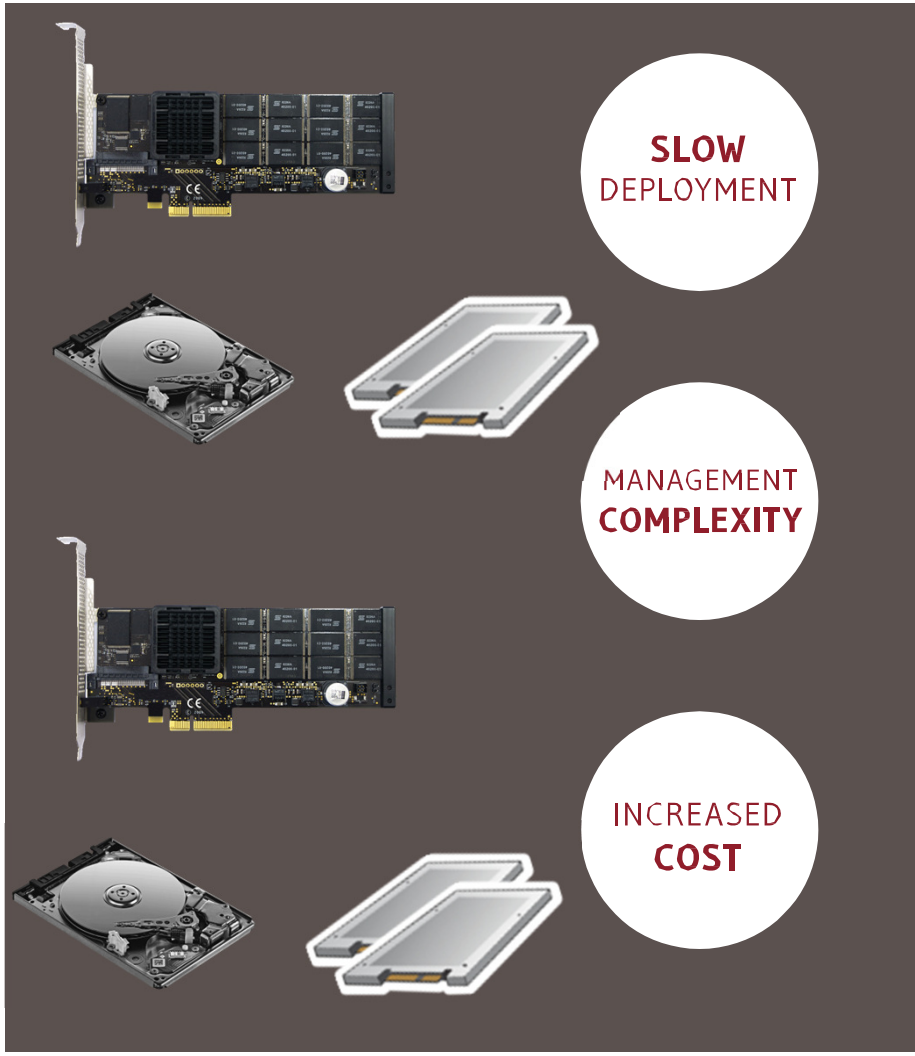

MCS-ENABLED HOMOGENOUS PLATFORM

- INDEXING
- ANALYTICS
- SWAP AREA
- CACHING
- LOGGING
- COLD STORAGE
- METADATA
- MORE...

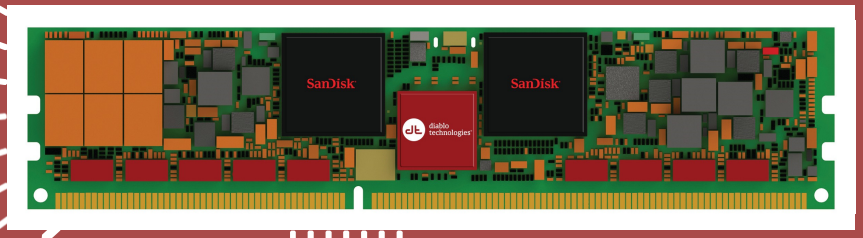

SanDisk® + diablo technologies™

- SIMPLIFIED DEPLOYMENT
- MANAGEMENT CONSISTENCY
- ECONOMIES OF SCALE

SUMMARY

Memory Channel Storage

+ Leverages parallelism and scalability of the memory channel

+ Significantly reduces data persistence latencies and improves single thread throughput

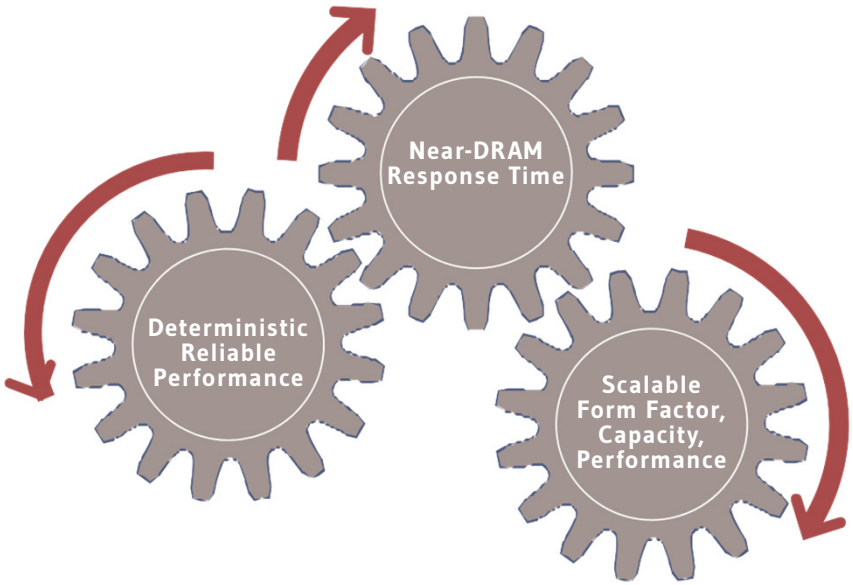

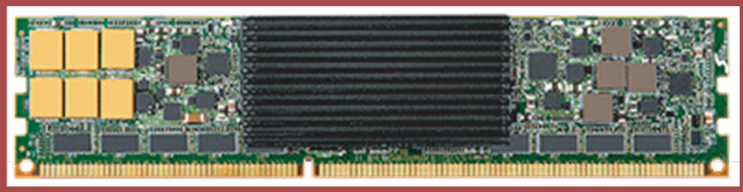

Benefits of MCS

+ 200GB to tens of TB's of flash in standard DIMM form factor and DDR3-CPU interface

+ Disruptive performance accelerates existing applications and enables new flash use cases

+ Scalability facilitates economic, "right-sized" system solutions

+ Form factor enables high-performance flash in servers, blades, and storage arrays

+ Future proofed with ability to utilize NAND-flash and future non-volatile memories

Thank You!

Jerome McFarland
jmcfarland@diablo-technologies.com

diablo technologies™ | PROBLEM SOLVED.

Q&A

1. What is the difference between Memory Channel Storage and NVDIMMs?

2. What is the difference between Memory Channel Storage and SATADIMMs?

3. Though being on the memory channel will reduce latency to the CPU, isn't I/O performance still limited by the IOPS & BW of the SSD NAND flash and controller?

4. How do Enterprise customers purchase Memory Channel Storage?

5. Will Memory Channel Storage support DDR4?

6. Will the Linux driver be open sourced?

7. Is there a limit to the number of MCS modules that can be deployed in single system?

8. What is the cost per GB for a Memory Channel Storage device?

Thank You!

Jerome McFarland
jmcfarland@diablo-technologies.com

diablo technologies™ | PROBLEM SOLVED.