



Diablo Technologies

Memory1™

Maher Amer

CTO

Diablo Technologies - Highly Confidential

➤ Big Memory for Big Data

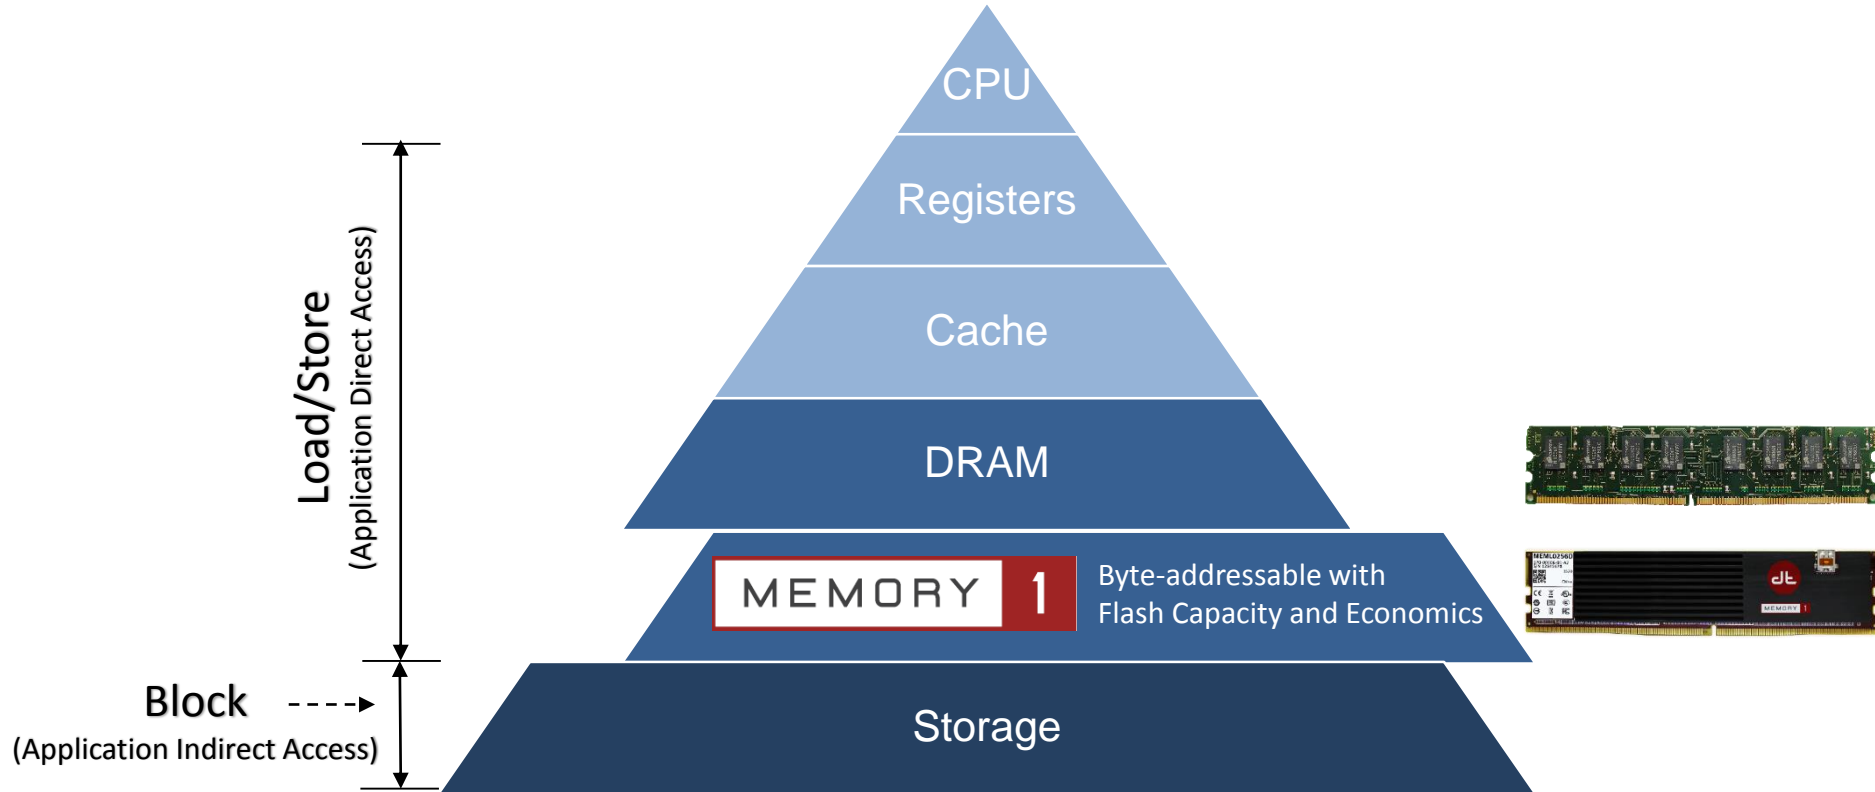
THE NEED: Big Data applications need Big Memory




**DRAM capacity is limited
while cost is prohibitive**

THE SOLUTION: Expand system memory with Memory1

➤ Two Tier Memory Sub-Systems are Here to Stay



➤ Memory1 At A Glance



- DDR4 memory DIMM with up to 128 GB per module
- Economically expands application memory by terabytes
 - No changes required to servers, applications
- Industry standard, JEDEC-compliant LRDIMM/RDIMM

➤ How It Works: High-Performance Hardware Solution

➤ Software Intelligently Manages Memory Access


- Acts as extension of OS Virtual Memory Manager
- Implements intelligent paging algorithms
- Optimizes performance and extends Flash endurance

➤ Modules Plug Into DDR4 Memory Slots

- Memory bus is highest performing interface to CPU
- Over 17GB/s per memory channel @ 2133 MT/s
- Lower latency than PCIe/NVMe

➤ Innovative All-Flash DDR4 DIMM Hardware

- Deployed in parallel with standard DRAM
- Leverages Flash capacity, power, and cost advantages



➤ Target Applications and Workloads

Application data doesn't fit in one machine



DATABASE

- Distributed database
- In-memory database
- Relational database

Application data doesn't fit in one machine



CLOUD

- Caching
- Paid search
- Key-value lookup

Application data doesn't fit in one machine but DRAM constrained



BIG DATA PROCESSING

- Real-time analysis
- Distributed caching
- Predictive analytics

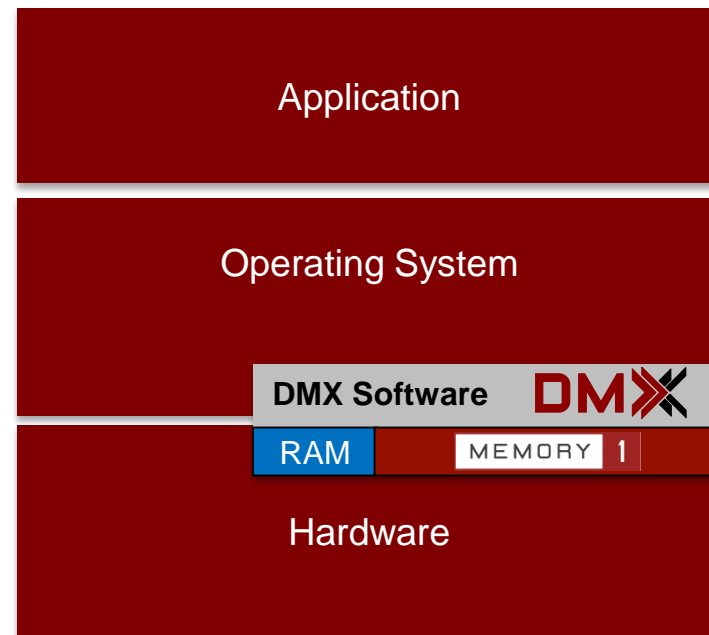


diablo
technologies™

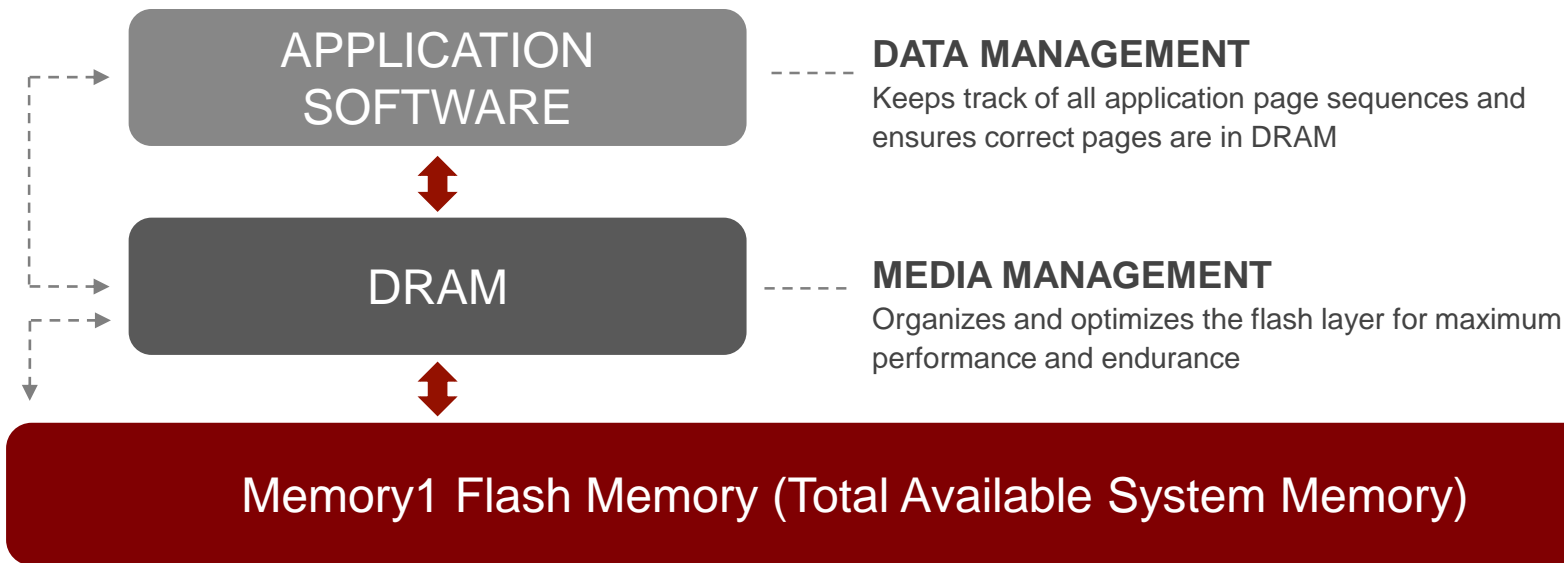
Diablo Memory Expansion

➤ Diablo Memory Expansion (DMX) Software

- **Combination of firmware and software**
 - Intelligently manages application memory access
 - Leverages CPU hardware and special statistics
 - Manages performance and endurance
- **No Application Changes**
 - Loads like a driver
- **Two Major Components**
 - Data Management
 - Media Management



➤ DMX Software Intelligently Manages Expanded Memory



- DMX Software loads and operates as an OS-level driver
- Application Software requires no changes

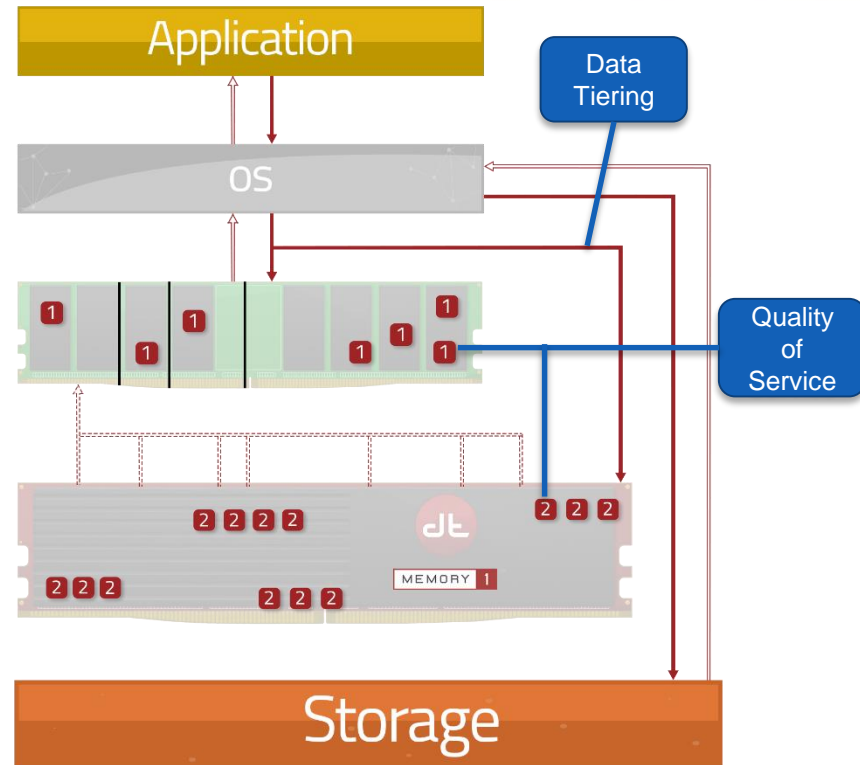
➤ Data Management Details

➤ Data Tiering

- DMX keeps hot data in DRAM
- High priority data maintained in DRAM
- Cold data evicted to flash

➤ Quality of Service


- Priority Associated Data Placement
- Additional DRAM allocated per application increasing hot data in DRAM
- Keeps data with response time requirements in DRAM



➤ Data Management Details

➤ Learning Engine

- Application Profiling and Analytics
 - Monitors application data access behaviors
- Data access prediction
 - Predicts next or additional pages required
- Smart Data pre-fetch
 - Pre-fetches pages to DRAM based on profiling, history, and data access patterns

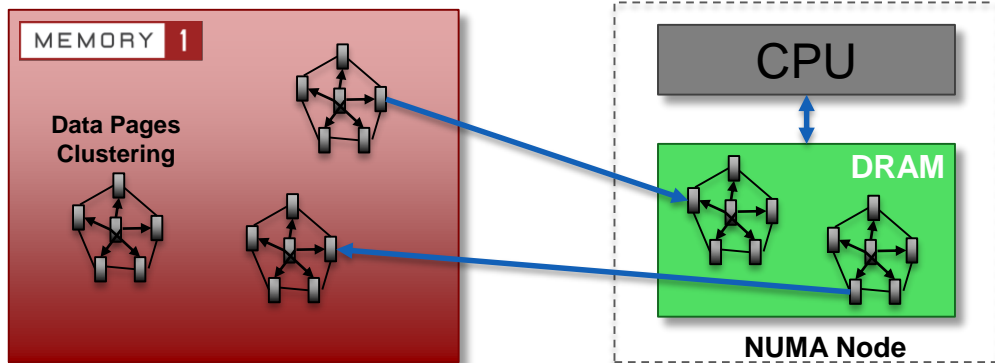


➤ Clustered Pages

- Prefetches grouped data typical in many applications

➤ Data Locality

- Movement between DRAM and Memory1 ensures data local to associated node



➤ Amortized Page Faults

- Groups page requests together, fully leveraging page fault

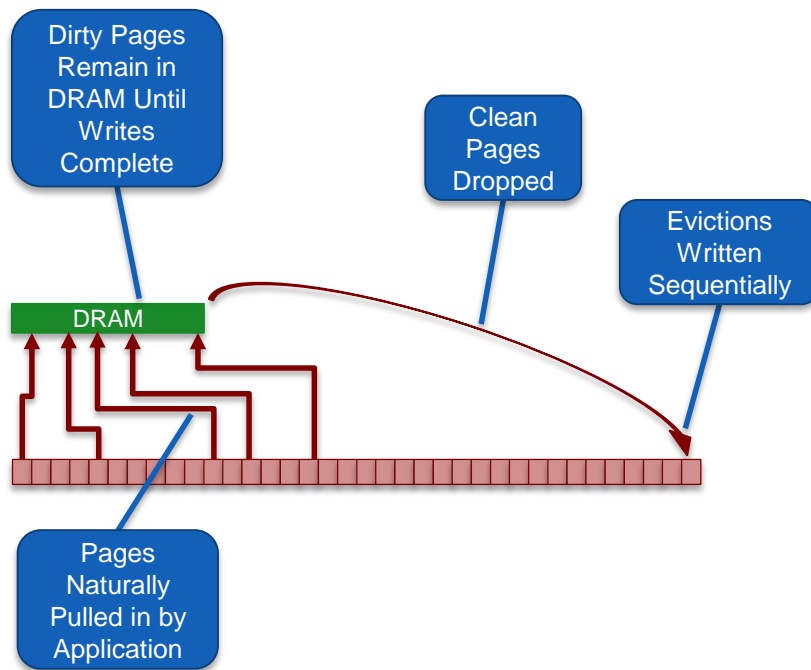
➤ Media Management Details

➤ Flash Management

- Low-level media management (handled in firmware)
- SoftFTL – Adaptation for 4K pages
- Tuneable cache ratio (4:1, 8:1)
- Device Striping

➤ Intelligent Traffic Management

- Dirty Page Writes
 - Avoids premature writes to flash for frequently written pages
 - Minimizes Read/Modify/Write operations
- Traffic Sequentialization
 - Pages evicted are written sequentially to flash





diablo
technologies™

Application Benchmarks


➤ Kdb+ Time Series Software: Memory1 Use Case

[Kdb+ time series data software | Kx Systems](https://kx.com/software.php)

The world's most powerful number cruncher, **kdb+** offers unparalleled performance for time-series data and analytics

<https://kx.com/software.php>

Stock Ticker Analysis and Regression:



300GB of in-memory data generated (after Garbage Collection)

2HRs of Execution Time on 2S EP Platform w/ 512GB DRAM

➤ Kdb+: 2TB Memory1 Configuration

Memory1 2TB KDB+ Appliance



CPU: Haswell/Broadwell

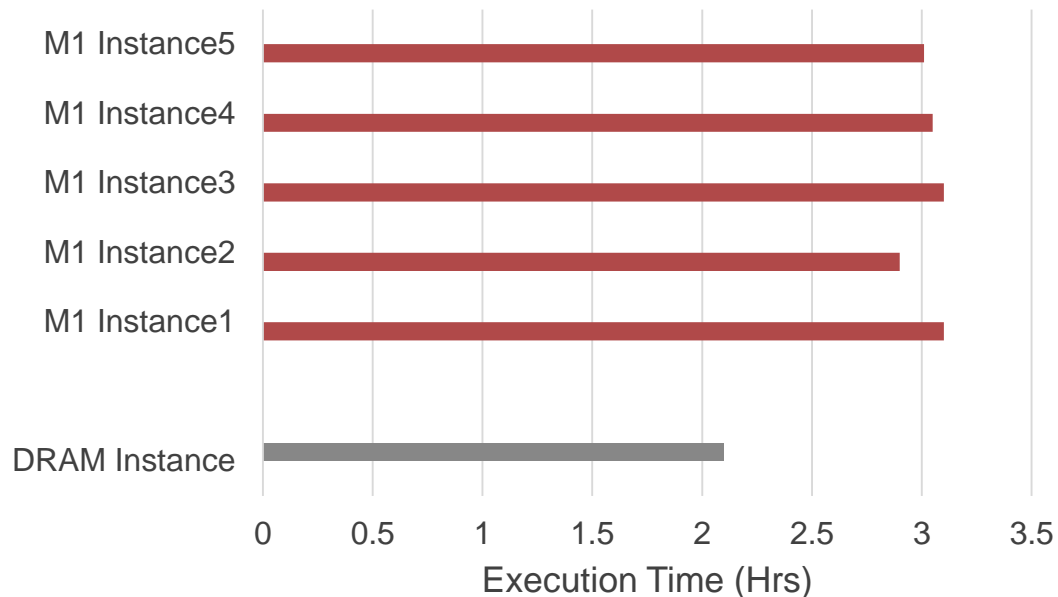
Cores: 14-18C/socket

DRAM: 256GB

MEMORY1: 1900GB

NIC: built-in GE + Add-on 10GE

Storage: N x HDD

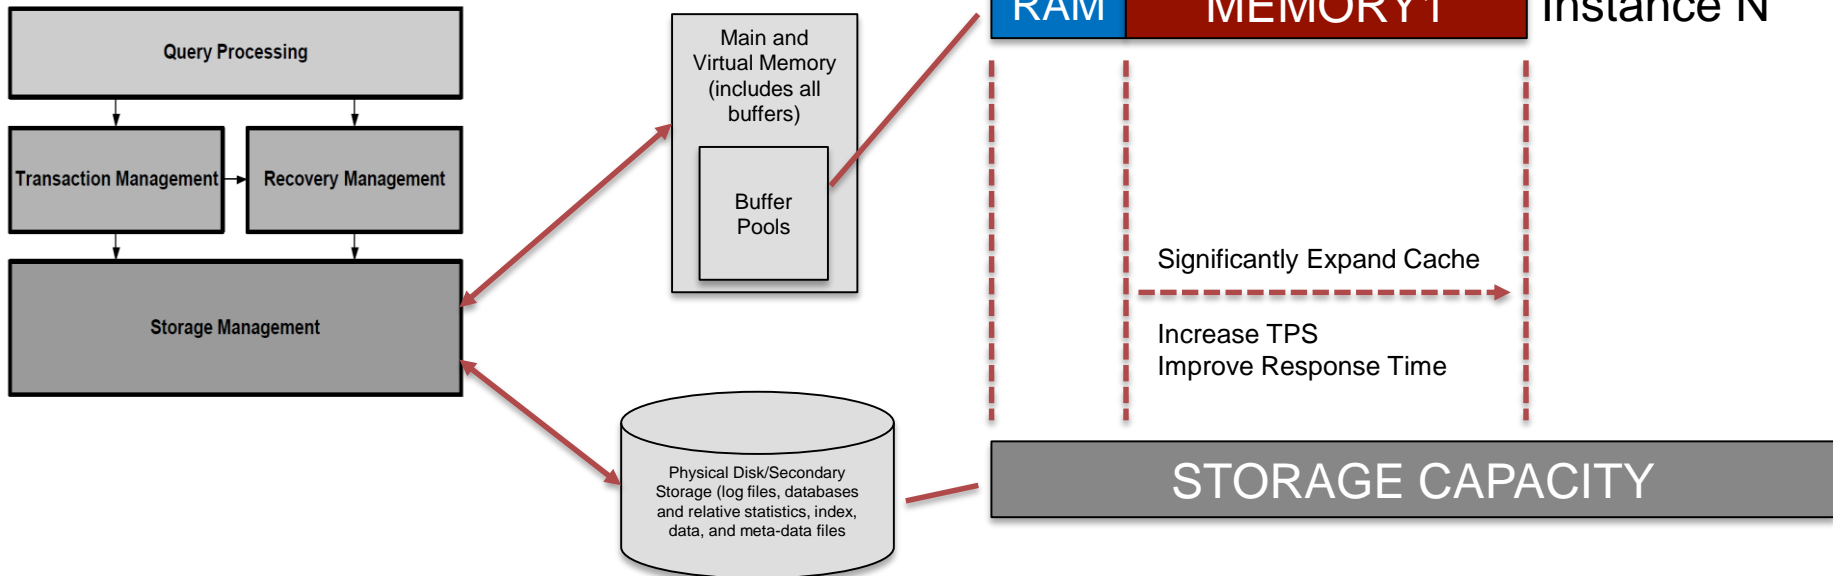


Increase dataset size and instances per machine by 5X and avoid EX Platforms


➤ Memory1 Advantage: Economically Increases MySQL Buffer Pool

zGB MySQL Instance Buffer Pool

MySQL Architecture




> Test Configuration



Increasing Buffer Pool size with MEMORY1 should:

- > Remove Bottleneck imposed by Storage Performance
- > Increase TPS and Reduce Response Time

Random, Normal Workload Distribution



➤ Hardware Configuration

All Servers Include: 10GE Optical NIC CentOS 7

Client (Load Generator)

Sysbench



10GE Optical Switch



DRAM / NVMe Server

Dual 14-Core CPUs
128GB DRAM
NVMe SSD as storage



Memory1 Server

Dual 14-Core CPUs
128GB DRAM
1TB Memory1

Random, Normal Workload Distribution



RAM

MEMORY1


STORAGE CAPACITY

Increasing Buffer Pool size with Memory1 will:

- Remove bottleneck imposed by storage performance
- Increase TPS and reduce response time
- Allow top performance with any storage solution


> Benchmark Scenarios

DMX:
6x Scenario A
Total RAM used = 128G
Sysbench Mixed Traffic



Scenario "A"


RAM:
6x Scenario B
Total RAM used = 128G
Sysbench Mixed Traffic



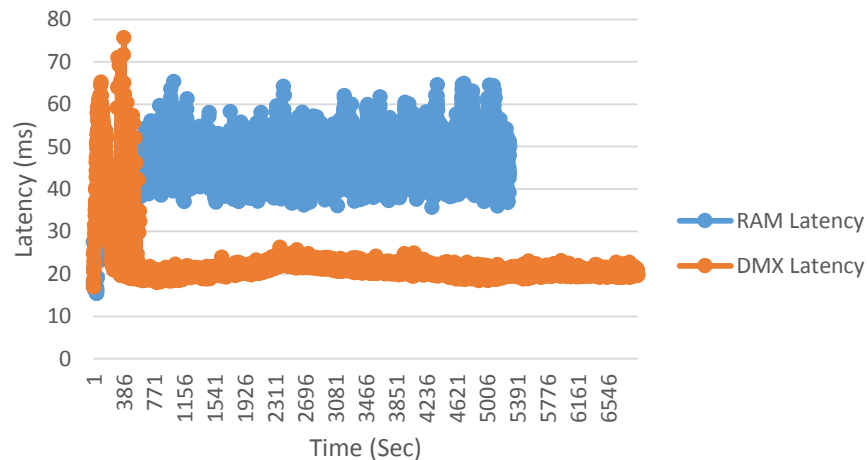
Scenario "B"

➤ DMX vs RAM

TPS Comparison



Latency Comparison




➤ Storage Read BW

sysbench / 1 / reads bandwidth per second

DMX / DB 6 x 75 / BP 64 x 122880M / LFS 8192M / RND uniform


bandwidth_type=reads
device=md0



sysbench / 1 / reads bandwidth per second

DMX / DB 6 x 75 / BP 16 x 16G / LFS 8192M / RND uniform

bandwidth_type=reads
device=md0





diablo
technologies™

Thank You!

Maher Amer
CTO, Diablo Technologies
mamer@diablo-technologies.com