diablo technologies

Diablo Technologies

Run Your Flash at DRAM Speeds

Aug. 13th, 2013

Riccardo Badalone, CEO

Problem solved.

Why are Writes So Important?

- ‘Persisting’ data is a critical operation within Databases, File Systems, and Messaging Systems
- Writes are usually synchronous and performed on a single thread (or at least very few threads), so they are typically slow and block other transactions
- To some degree, write performance in IOPS is synonymous with ‘transactions per second’, so it’s an important performance parameter in evaluating persistence technology
- The game is changing in terms of how data is written to Flash…

ULLtra DIMM System Solution

A Modular Flash System using the Memory Channel

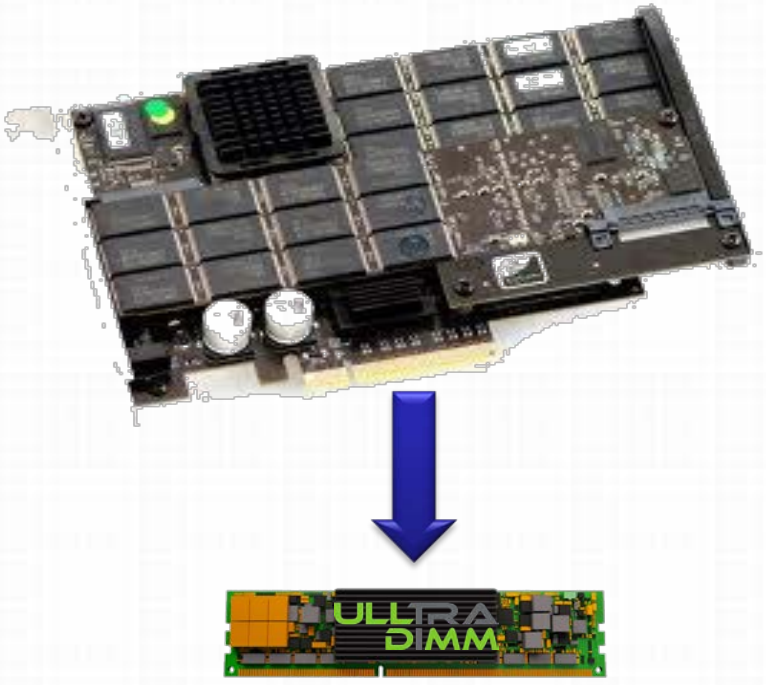

200/400GB MLC Flash

Solution features:

- Software Drivers for major Operating Systems
- Memory Channel Storage™ (MCS™) enables persistence accessed through DDR3 sub-system of CPU
  - IO Acceleration
  - Paging Acceleration

MCS System View

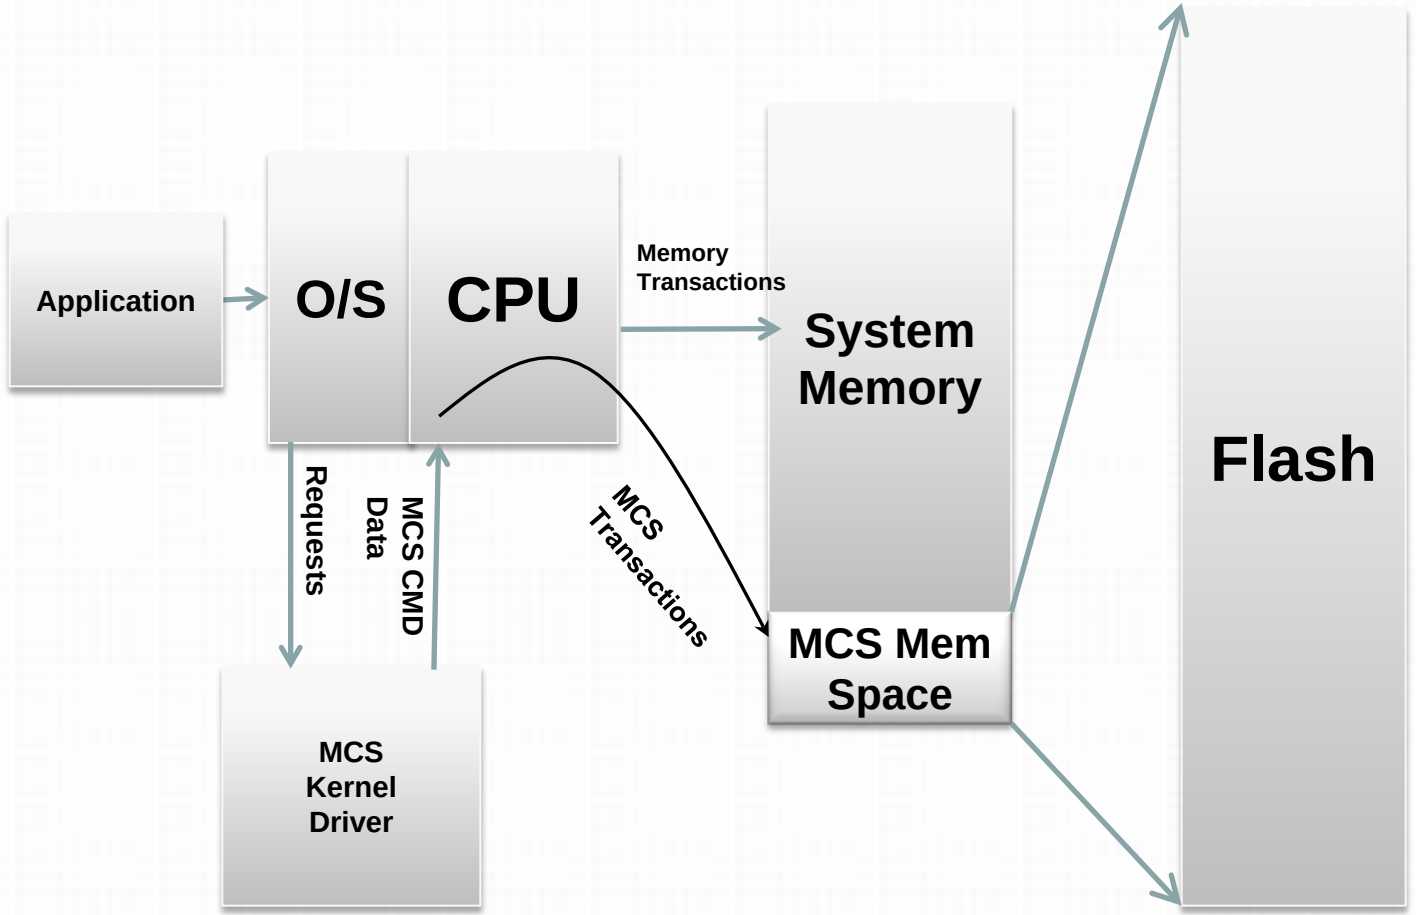

Read/Write Data remains within the memory sub-system

Example Deployment with Dual Socket Enterprise Processor

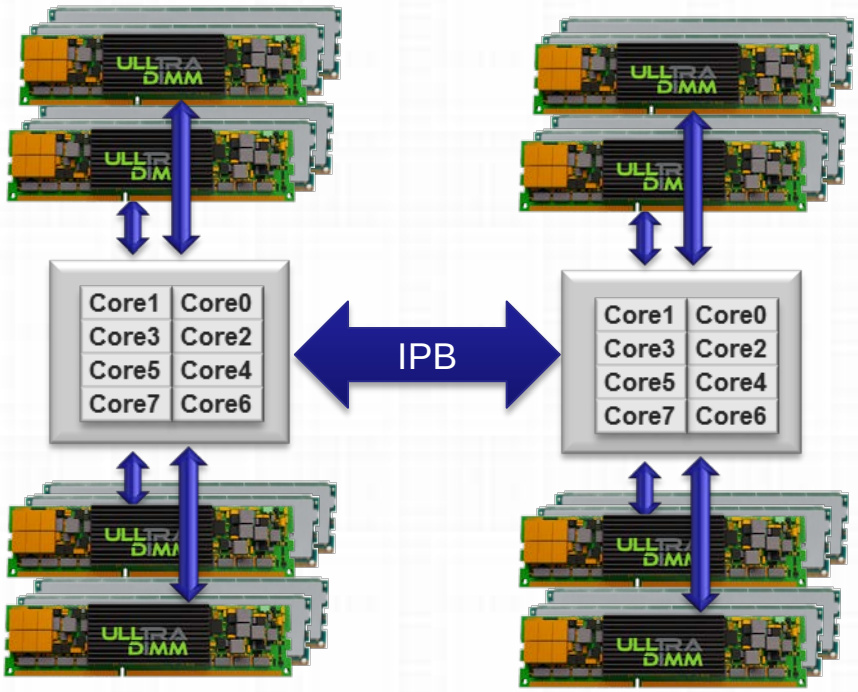

- Create distributed persistence with multiple entry points
- Leverage Transfer B/W of 8X concurrent DDR Channels for page transfer into persistence layer
- Increase # of parallel channels and FTL MIPS

Total DIMM slots	Total ULLtraDIMMs	Total RDIMMs	Total Provisioned Flash Capacity	DRAM Capacity Total (w/ 16GB DR)	Read IOPS	Write IOPS
24	8 x 200G	16	1.6TB	256GB	1.2M	520K

Burst Write Performance

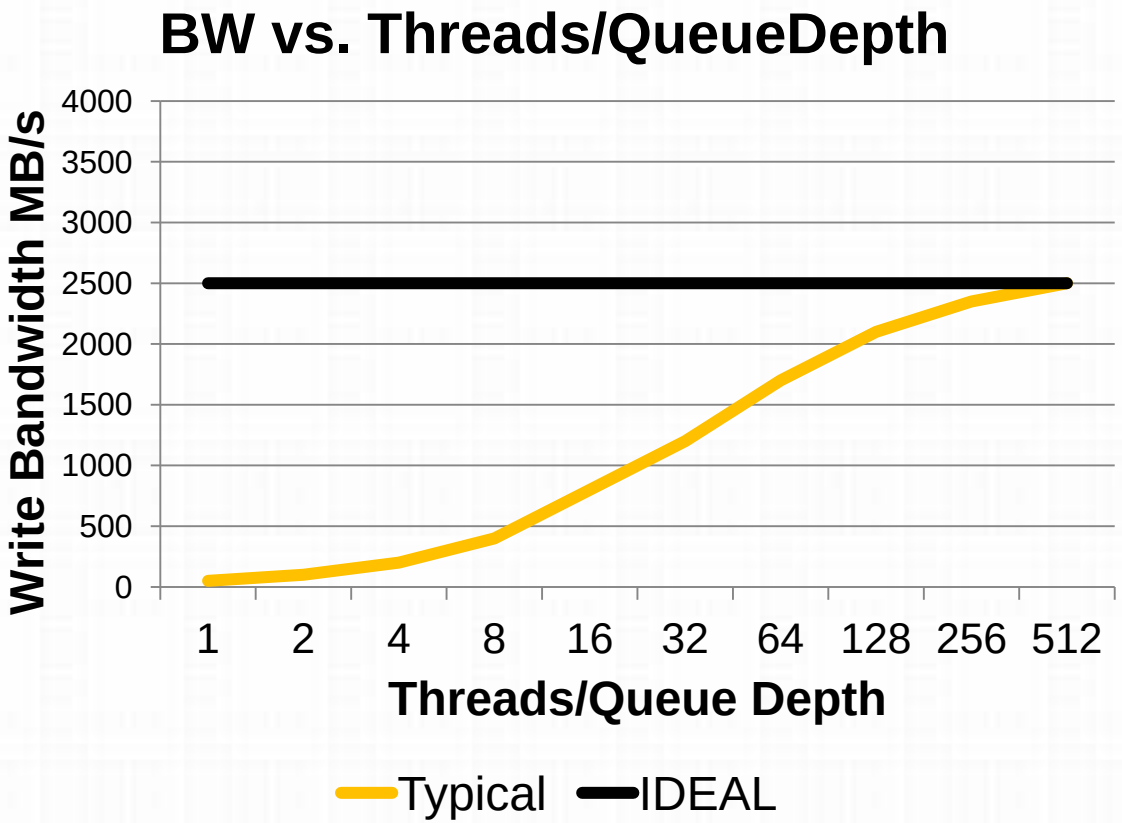

- **Typically, maximum throughput is achieved at very high thread count or queue depth**
- **Ideally, the maximum write throughput would be available to a single threaded process**

ULLtraDIMM Random Write Performance Scaling

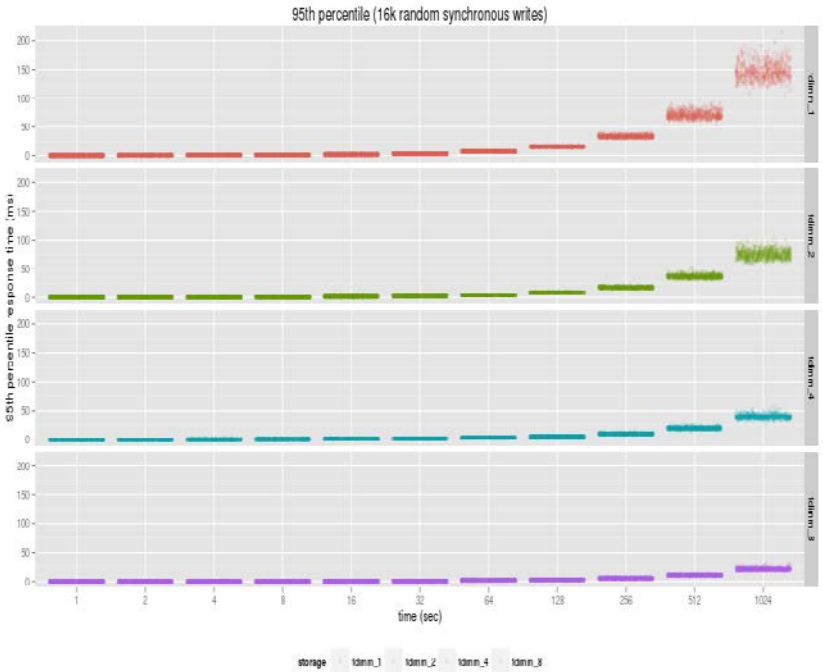

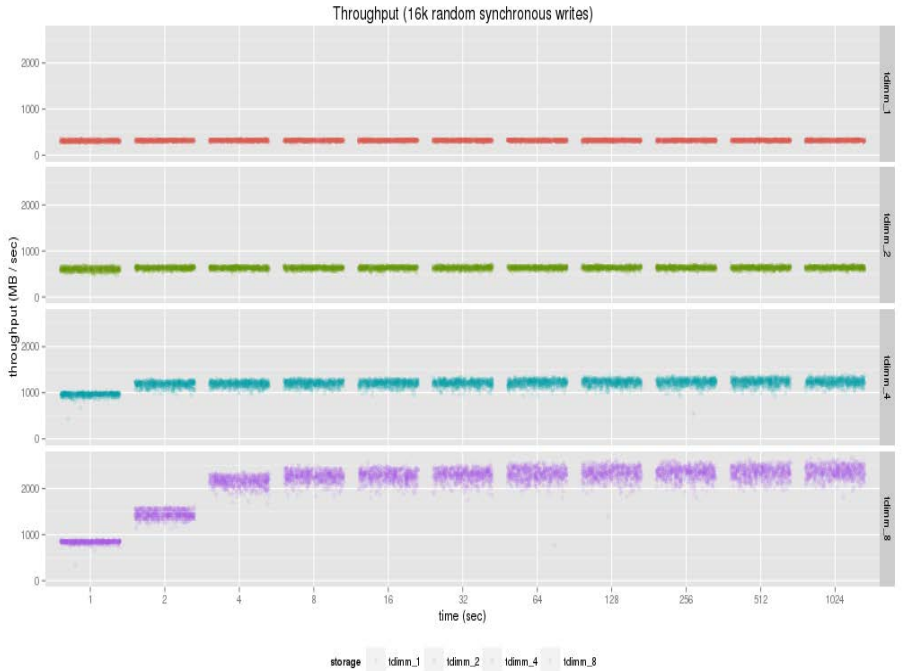

- 16K Random Writes on a Single Thread
  - 950MB/s @ 7us (thread limited)
- 16K Random Writes on Four Threads
  - 2.4GB/s @ 20us (max throughput achieved)

Summary

- ‘Persisting’ data is a critical operation within Databases, File Systems, Messaging Systems
- The memory sub-system offers high transfer rates and tremendous parallelism for persistence layers
- ULLtraDIMM scalability provides more synchronous, single threaded throughput than can be currently handled by most benchmarks
- The game is changing in terms of how data is written to Flash…

THANK YOU

RIC@DIABLO-TECHNOLOGIES.COM

613.277.7425