Accelerating Database Performance with Cisco UCS 785GB MLC Fusion ioDrive2

Cisco UCS B200 M3 B-Series Blade Servers with Cisco UCS 785GB MLC Fusion ioDrive2 Power MS SQL Server 2012 Databases

PERFORMANCE BRIEF

March 2013

The Cisco UCS B200 M3 Blade Server with a Cisco UCS 785GB MLC Fusion ioDrive2 was tested to assess its performance with an MS SQL Server 2012 database. The following tests were conducted:

1. Block-level tests: SQLIO was used to showcase IOPS and throughput for various block sizes, for both random and sequential workloads.
2. Transactional tests: HammerDB was used to run a TPC-C-like workload to determine application-level performance. These tests take into account the interplay of CPU, memory, and SQL server buffer management.
3. Database functionality tests: Certain I/O-intensive databases functions, such as backups, restores, integrity checks, and index rebuilds, were run to determine performance aspects.

CISCO UCS B200 M3 BLADE SERVER

The Cisco Unified Computing System™ (Cisco UCS®) B200 M3 Blade Server addresses the broadest set of workloads, from IT and web infrastructure to distributed databases. Building on the success of the Cisco UCS B200 M2 Blade Server, the enterprise-class Cisco UCS B200 M3 Blade Server further extends the capabilities of the Cisco UCS portfolio in a half-width blade form factor. It harnesses the power of the latest Intel® Xeon® processor E5-2600 product family, with up to 768 GB of RAM (using 32-GB DIMMs), two disk drives, and up to dual 4 x 10 Gigabit Ethernet throughput. In addition, Cisco UCS does not have to power and cool excess switches in each blade chassis, an architectural advantage. With a larger power budget per blade server, Cisco proves it can design uncompromised expandability and capabilities as shown by the UCS B200 M3, with its leading memory slot and drive capacity.

The Cisco UCS B200 M3 provides:

HIGHLIGHTS

Industry-leading Performance and Scalability

Storing full database structures on Cisco UCS 785GB MLC Fusion ioDrive2s in the Cisco UCS B200 M3 Blade Server allows customers to deliver consistent, low-latency performance to applications, reduce IT costs through consolidation, manage more data on less hardware, and make better business decisions in real time.

Increase Server Uptime

Cisco UCS 785GB MLC Fusion ioDrive2s eliminate disk aggregation and failure point proliferation, while improving performance significantly without database tuning.

Reduce Costs

Reduced server and storage footprint reduces hardware, maintenance, floor space, and energy costs.

- Two, multi-core, Intel® Xeon® processor E5-2600 CPUs, with up to 16 processing cores
- 24 DIMM slots for industry-standard double-data-rate 3 (DDR3) memory running up to 1600 MHz and up to 768GB of total memory (using 32GB DIMMs)
- Two optional, hot-pluggable SAS or SATA hard disk drives (HDDs) or solid-state drives (SSDs)
- Industry-leading 80Gbps throughput bandwidth
- Remote management through a Cisco Integrated Management Controller (CIMC) that implements policy established in Cisco UCS Manager
- Out-of-band access by remote keyboard, video, and mouse (KVM) device, Secure Shell (SSH) Protocol, and virtual media (vMedia), as well as the Intelligent Platform Management Interface (IPMI)
- A half-width blade. Up to eight of these high-density, two-socket blade servers can reside in the 6RU Cisco UCS 5108 Blade Server Chassis, offering one of the highest densities of servers per rack unit in the industry.

CISCO UCS 785GB MLC FUSION-IO IODRIVE2 ADAPTER

Fusion-io builds server-resident Peripheral Component Interconnect Express (PCIe) flash storage for applications requiring high input/output operations per second (IOPS) with low latency. Fusion's industry-leading, low-latency ioMemory eliminates I/O bottlenecks to dramatically improve server workload capability.

The Cisco UCS 785GB MLC Fusion ioDrive2 significantly reduces MS SQL Server 2012 database latency with a high-capacity, high-performance flash memory tier. The product offers a persistent storage option for any architecture. For maximum performance, store an entire database of up to 785GB on ioMemory or accelerate database performance significantly by storing key, I/O-intensive database structures on low-latency flash. By moving the I/O burden off the storage area network (SAN) to ioMemory within database servers, organizations improve the performance of database servers, while freeing SAN resources for other tasks.

SPECIFICATIONS FOR PERFORMANCE TESTING

Cisco UCS B200 M3 Blade Server

- Two Intel Xeon processor E5-2690 product family CPUs
- 256GB RAM
- Two mezzanine adapter slots (one dedicated to the Cisco UCS VIC 1240 for total throughput bandwidth of 80Gbps)
- Up to two optional front-access SAS or SATA HDDs or SSDs
- Up to 2TB of maximum internal storage
- An LSI SAS 2004 Integrated RAID controller
- RAID 0 and 1 support
- Two slots for 16-GB SD flash memory cards (future enablement via software update)

Cisco UCS 785GB MLC Fusion ioDrive2

- 785GB Multi-Level Cell (MLC) flash capacity
- 1.5GB/s bandwidth (1MB read blocks)
- 1.1GB/s bandwidth (1MB write blocks)
- 141,000 IOPS (512B random read)
- 535,000 IOPS (512B random write)
- 15µs write latency, 68µs read latency
- Hardware supported: All M3 Blades
- Software supported: Cisco UCS Manager 2.1

Test Transactional Database

- 150GB size
- Eight data files
- One Log file
- 4GB Ram allocated to SQL Server

BLOCK-LEVEL TESTS

We used SQLIO to demonstrate the block-level performance of the Cisco UCS 785GB MLC Fusion ioDrive2. The tool is available for download at the Microsoft Download Center (http://www.microsoft.com/en-us/download/details.aspx?id=20163). SQL Server 2012 performs its operations at different block size ranges as shown below (8KB is the most frequently used block size):

Workload	**Type**	Block Size
Log writes	Sequential	Up to 60KB
Checkpoint / Lazy Writes	Random	Up to 256KB
Index Seeks	Random	8KB or 64KB (for read-aheads)
Table / Range Scan	Sequential	64KB – 512KB
Bulk Load	Sequential	64KB – 256KB
Backup Operations	Sequential	64KB – 4MB

These tests highlight the maximum performance that the Cisco UCS 785GB MLC Fusion ioDrive2 provides for various block sizes, for both random and sequential workloads. The performance itself is measured using two metrics: Input/Output Operations per Second (IOPS) and throughput measured in MB/s (megabytes per second).

IOPS Results Across Various Block Sizes

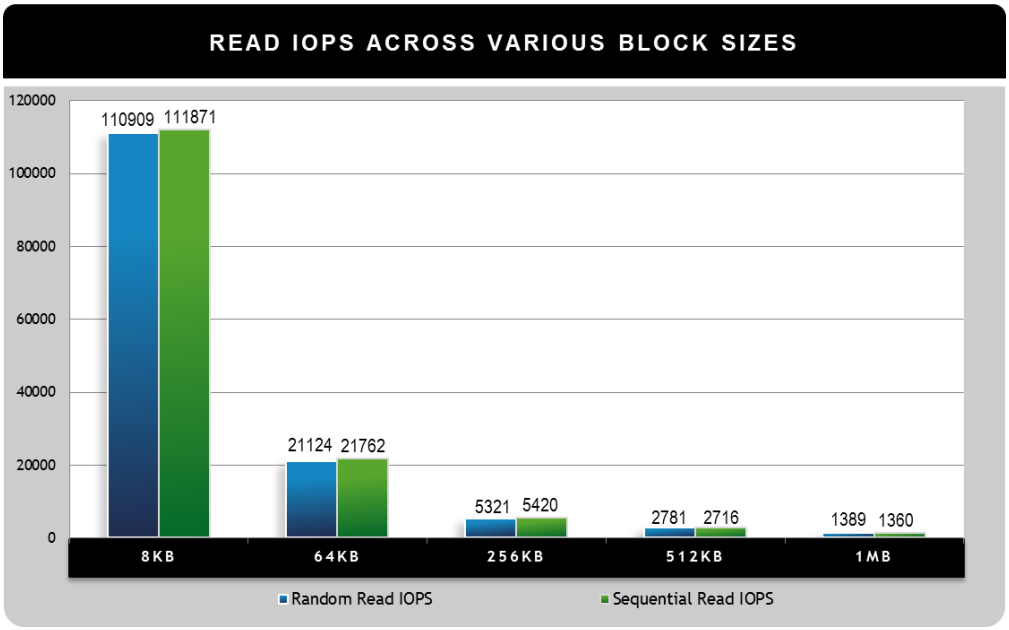

Figure 1: Random and Sequential Read IOPS are comparable across various block sizes.

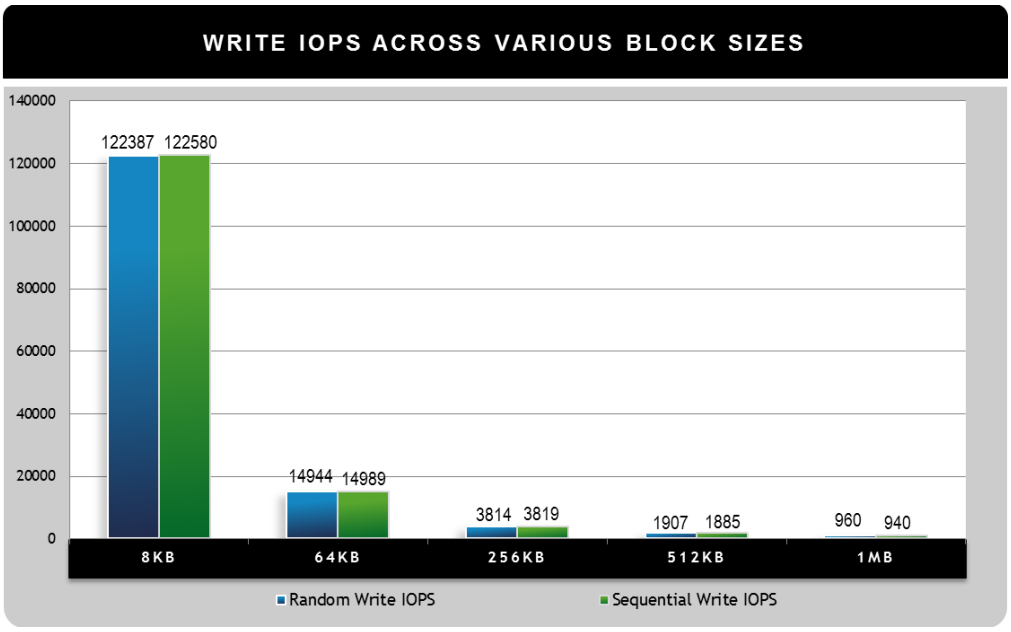

Figure 2: Random and Sequential Write IOPS are consistent across block sizes.

Throughput Results Across Various Block Sizes

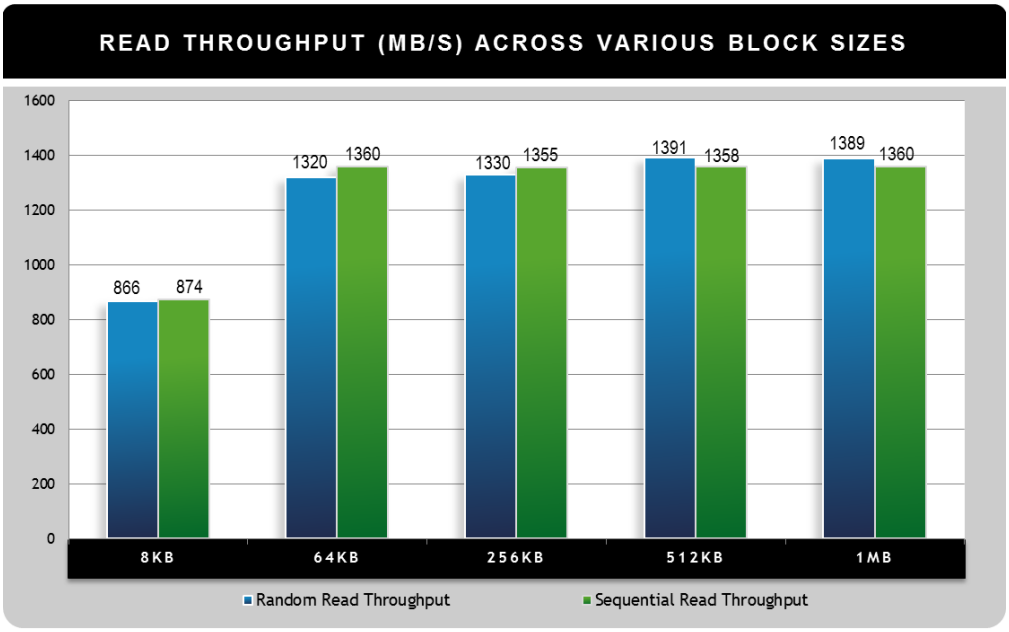

Figure 3: Consistent Read Throughput shows across various block sizes.

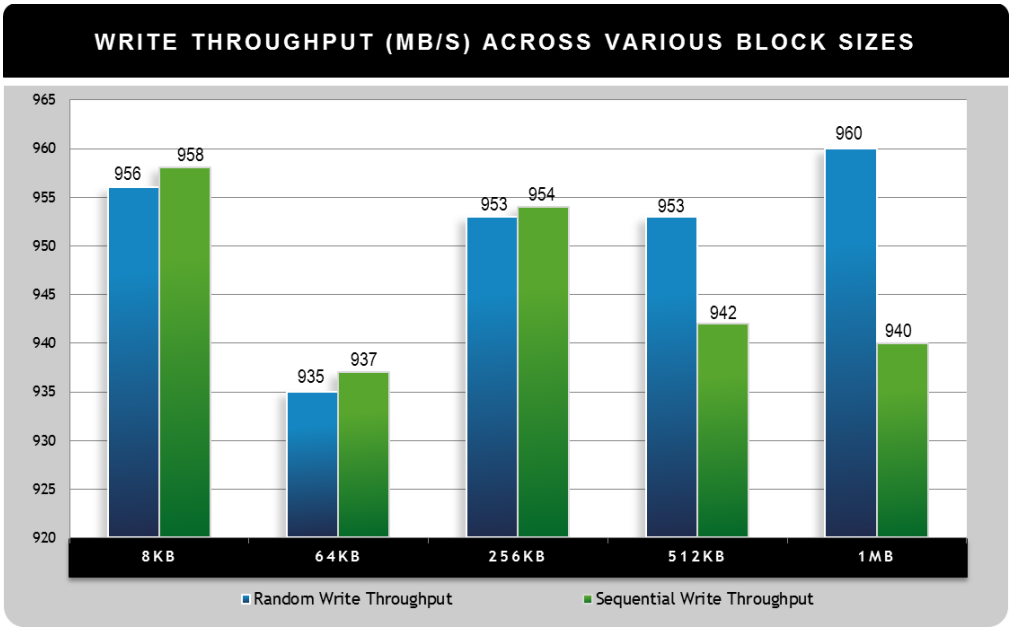

Figure 4: Sequential and Random Write Throughput is consistently high across block sizes.

Summary of Block-Level Tests

The Cisco UCS 785GB MLC Fusion ioDrive2 in the Cisco Blade provides:

- Over 122,000 IOPS for an 8KB block size (the most relevant block size for MS SQL Server). IOPS is a key indicator of transactional workload performance.
- Over 1.3 GB/s throughput for large block sizes. Bandwidth is a key indicator for data-warehousing performance.
- Unlike spinning disks, the random and sequential performance of the Cisco UCS 785GB MLC Fusion ioDrive2 is very similar. This makes it is universal solution for both transactional as well as data-warehousing workloads.

TRANSACTIONAL TESTS

We used HammerDB to run a TPC-C like workload on the Cisco UCS 785GB MLC Fusion ioDrive2. The tool is open source and available at http://hammerora.sourceforge.net/index.html. The test measures how well the SQL Server 2012 database can leverage the Cisco UCS 785GB MLC Fusion ioDrive2. This test demonstrates the full interplay of CPU, memory, and the SQL Server Buffer Management system. For comparison, we ran the same tests on a similar-sized database residing on a local hard drive (spinning media referred to as "disk" from here on) within the Cisco Blade Server. We ran the tests by scaling up the number of users running the transactional workload simultaneously.

Users	TPM (ioMemory)	TPM (Disk)
1	1,923,708	112,434
2	3,139,056	108,348
4	5,466,660	188,988
8	8,641,248	241,628
16	13,557,276	466,000

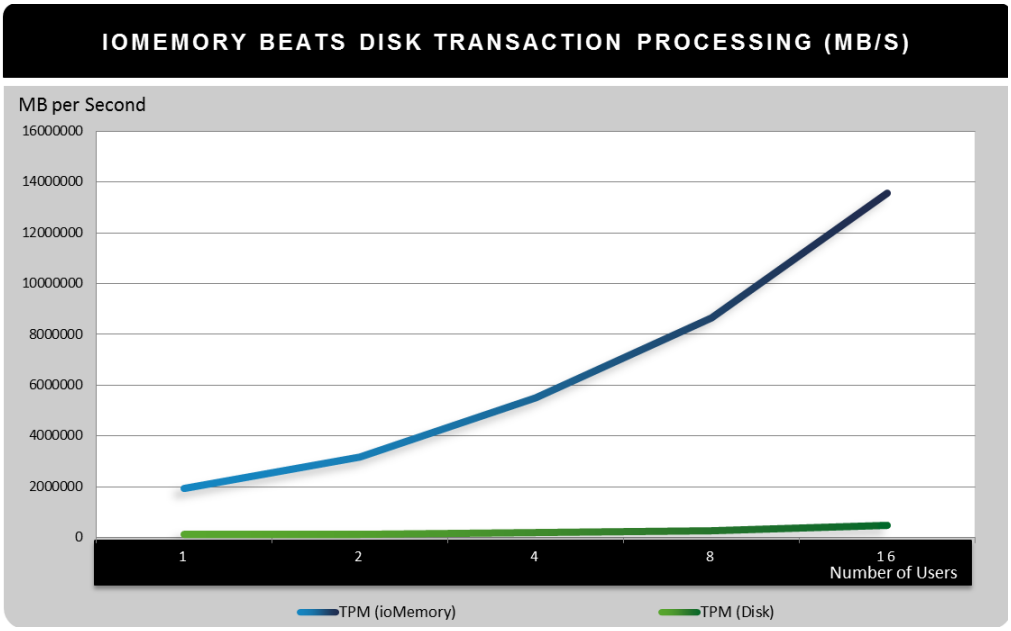

Figure 5: ioMemory transaction processing beats disk transaction processing, up to 36 times faster.

Results of Average Disk Access Latency for the Application Workload

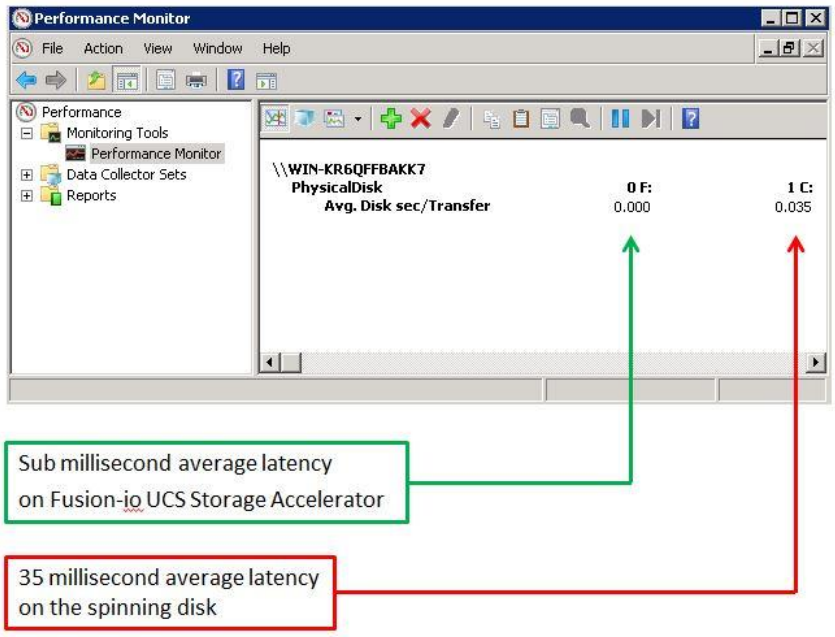

Figure 6: ioMemory latency so low it is not registering, while disk latency measures 35 milliseconds.

Summary of Transactional Tests

The Cisco UCS 785GB MLC Fusion ioDrive2 can support over 13.5 million transactions per minute for a TPC-C-like workload on MS SQL Server 2012 in the Cisco Blade Server, 29 times more than disk.

The Cisco UCS 785GB MLC Fusion ioDrive2 is able to deliver a very high transactions-per-minute rate by means of reducing the application-level, disk-access latency to less than 1 millisecond.

DATABASE FUNCTIONAL TESTS

These tests reflect the time a DBA might need to perform common, day-to-day database tasks. For comparison, we ran the tests on a database stored on a Cisco UCS 785GB MLC Fusion ioDrive2 and the same tests again with a database stored on a local spinning disk.

Database Operations	Time to Complete on Cisco UCS 785GB MLC Fusion ioDrive2 (minutes)	Time to Complete on Disk (minutes)	Cisco UCS 785GB MLC Fusion ioDrive2 is faster by
Back up database[1]	7	76	Over 10X
Restore database	9	362	Over 40X
DBCC Ceckdb(), integrity check	18	62	Over 3x
Rebuild indexes[2]	10	77	Over 7X

SUMMARY OF DATABASE FUNCTIONAL TESTS

The I/O limitations of hard disk are made clear in these tests. Solving the I/O bottleneck by storing the database on a Cisco UCS 785GB MLC Fusion ioDrive2 enables the powerful Cisco Blade Server processor to be used more effectively. This allows a DBA to perform routine operations much faster, and without affecting the production systems.

CONCLUSION

These tests show that running MS SQL Server 2012 on flash memory in a Cisco UCS 785GB MLC Fusion ioDrive2 deployed on a Cisco UCS B200 M3 Blade Server results in dramatic performance improvements.

For smaller databases, moving the full database to the Fusion-io product enables performance to scale linearly in terms of users and transactions. For larger databases, moving specific database structures to Fusion ioMemory reduces storage requirements in the SAN while accelerating performance. This solution results in significant overall performance improvements, reduced mechanical failures, and reduced power and cooling requirements for the SAN infrastructure.

[1] Database size is 150GB.
[2] Index size is 30GB.