



Dell EMC® PowerEdge™ MX750 Servers and MySQL Application Performance Gains with PCIe® 4.0 Technology

Dell EMC PowerEdge MX750 servers are based on the current PCIe 4.0 interface and the latest 3rd Gen Intel® Xeon® Ice Lake¹ PCIe 4.0 scalable processors. Generationally, servers with PCIe 4.0 enable twice the bandwidth versus the previous PCIe 3.0 technology so peripheral devices, such as SSDs, GPUs and NICs can access data faster than ever before. The speed upgrade is well-suited for data-intensive and computational applications such as cloud computing, databases, data analytics, artificial intelligence, machine learning, container orchestration and media streaming. A faster PCIe interface enables today's powerful CPUs, such as Ice Lake CPUs, to be continually fed with data.

To validate application performance and productivity gains that can be achieved with the PCIe 4.0 interface, KIOXIA Corporation, a leader in PCIe 4.0 SSD storage, compared the performance of a MySQL database application running on an MX750 server with PCIe 4.0 SSD technology versus an MX740 server with PCIe 3.0 SSD technology. The test process and results are presented.

System and Application Test Scenario

The tests utilized an operational, high-performance MySQL database workload that was based on comparable TPC-C™ benchmarks created by HammerDB software². The MySQL database is commonly used in hyperscale and enterprise environments, and widely deployed on PowerEdge servers. It supports key applications such as webserver, online transactional processing (OLTP), e-commerce and data warehousing, and is the most widely deployed open source database globally (ranked number two overall³).

The tests on each server platform were conducted with PCIe 4.0 and PCIe 3.0 SSDs that measured transactions per minute (TPM), average read/write latency and CPU utilization. For the MX750 server, four (4) KIOXIA CM6 Series PCIe 4.0 enterprise NVMe® SSDs were deployed. For the MX740 server, four (4) PCIe 3.0 specification-compliant enterprise NVMe SSDs were deployed.

The test results provide a real-world scenario of TPM performance, average read/write latency and CPU utilization when running a MySQL database application using comparable equipment and performing queries against it. Supporting details include a description of the test criteria, the set-up and associated test procedures, a visual representation of the test results, and a test analysis.

Test Criteria:

The hardware and software equipment used for these tests included:

- **Server Configurations:**

- **Server Setup 1:** One (1) Dell EMC PowerEdge MX750 dual socket server with two (2) Intel Xeon Ice Lake PCIe 4.0 CPUs, featuring 28 processing cores, 2.0 GHz frequency, and 512 gigabytes⁴ (GB) of DDR4 RAM

- **Server Setup 2:** One (1) Dell EMC PowerEdge MX740 dual socket server with two (2) Intel Xeon Cascade Lake⁵ PCIe 3.0 CPUs, featuring 24 processing cores, 2.2 GHz frequency, and 384 GB of DDR4 RAM

- **Operating System:** CentOS™ v8.3

- **Application:** MySQL v8.0 (database size of 440 GB)

- **Test Software:** Comparable TPC-C benchmark tests generated through HammerDB v3.3 test software
- **Storage Devices (Table 1):**
SSD Setup 1: Four (4) KIOXIA CM6-V Series (3 DWPD⁶) PCIe 4.0 enterprise NVMe SSDs with 3.2 terabyte⁴ (TB) capacities
SSD Setup 2: Four (4) PCIe 3.0 specification-compliant (3 DWPD) enterprise NVMe SSDs with 3.2 TB capacities

Specifications	CM6-V Series	PCIe 3.0-compliant
Interface	PCIe 4.0 NVMe (U.3)	PCIe 3.0 NVMe (U.2)
Capacity	3.2 TB	3.2 TB
Form Factor	2.5-inch ⁷ (15mm)	2.5-inch (15mm)
NAND Flash Type	BiCS4 3D NAND (96-layer)	V-NAND
Drive Writes per Day (DWPD)	3 (5 years)	3 (5 years)
Power	18W	20W
DRAM Allocation	96 GB	96 GB

Table 1: SSD specifications and set-up parameters

PLEASE NOTE: The MySQL database was limited to 96 GB of RAM. Although the total capacity available to the test system differs (512 GB for server setup 1 versus 384 GB for server setup 2), the actual amount of capacity used by each server setup was expected to be the same and not contribute to any performance advantage⁸.

Set-up & Test Procedures

Set-up: The MX750 and MX740 servers were setup with the CentOS v8.3 operating system and MySQL v8.0 software. The MySQL database was set to a maximum of 96 GB of DRAM and was placed into a RAID10 set. RAID10 was selected because it is commonly used in data center environments. Once each SSD array was initialized (PCIe 4.0 SSDs and PCIe 3.0 SSDs), the RAID10 set was formatted to the XFS file system. A 440 GB database was then loaded into each server setup (MX750/MX740) using HammerDB test software. Once the database was loaded, it was then backed up. Before each test run, the 440 GB test database was restored to the exact same state for each run to control the test inputs and database size.

Test Procedures: The first set of tests (TPM, latency and CPU utilization) were run on the MX750 server. The comparable TPC-C workload utilized HammerDB software to run the test. The four (4) KIOXIA CM6-V Series SSDs were then placed into a RAID10 set and the tests were conducted. Multiple iterations of each test were run to determine an optimal configuration of virtual users. A configuration of 480 virtual users delivered the highest TPM performance numbers. See Test Results section.

The first set of tests were run on the MX740 server. The comparable TPC-C workload utilized HammerDB software to run the test. The four (4) PCIe 3.0 specification-compliant SSDs were then placed into a RAID10 set and the tests were conducted. Multiple iterations of each test were run and a configuration of 480 virtual users delivered the highest TPM performance numbers. See Test Results section.

Test Results

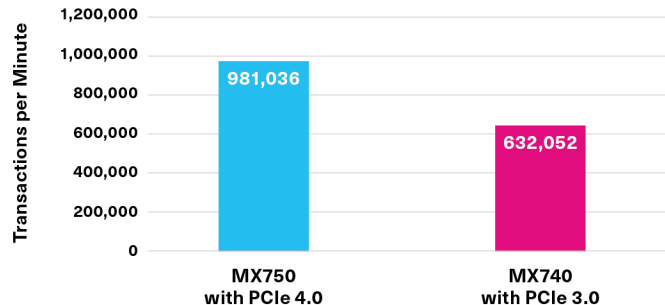
The tests were conducted with the MX750 server (Ice Lake CPUs with PCIe 4.0 technology) versus an MX740 server (Cascade Lake with PCIe 3.0 technology), with the results recorded. For the TPM result, the higher the test value, the better the result. For the latency and CPU utilization results, the lower the test value, the better the result.

Transactions Per Minute

In an Online Transaction Processing (OLTP) database environment, TPM is a measure of how many transactions in the TPC-C transaction profile that are being executed per minute. HammerDB software, executing the TPC-C transaction profile, randomly performs new order transactions and randomly executes additional transaction types such as payment, order status, delivery and stock levels. This benchmark simulates an OLTP environment where there are a large number of users that conduct simple, yet short transactions that require sub-second response times and return relatively few records. The TPM test results:

TPM: MySQL Comparable TPC-C Workload	MX750 (PCIe 4.0)	MX740 (PCIe 3.0)
TPM (higher is better)	981,036	632,052
Advantage	+55%	-

Transactions per Minute

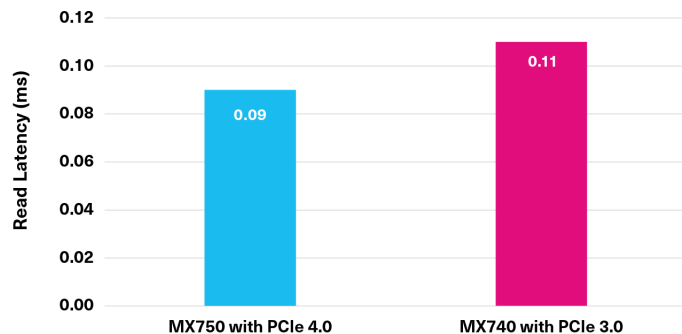


Average Read and Write Latency

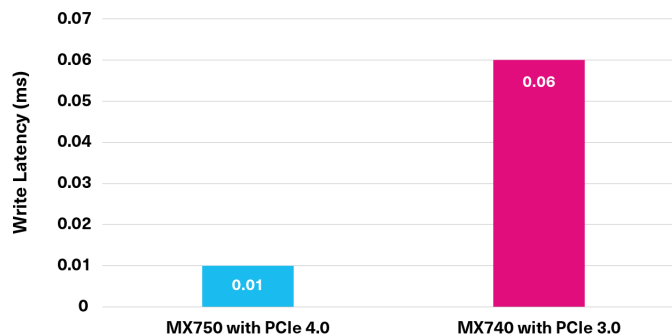
Latency is the delay in time before a storage device completes a data transaction following an instruction from the host for that request, which can greatly affect application performance and the user experience. Application response time often has built-in latency between the user and the server, so maintaining low-latency within the server will usually translate into an overall better user experience. The read and write latency test results:

Latency: MySQL Comparable TPC-C Workload	MX750 (PCIe 4.0)	MX740 (PCIe 3.0)
Average Read Latency (lower is better)	0.09 ms	0.11 ms
Advantage	-18%	-
Average Write Latency (lower is better)	0.01 ms	0.06 ms
Advantage	-83%	-

Average Read Latency



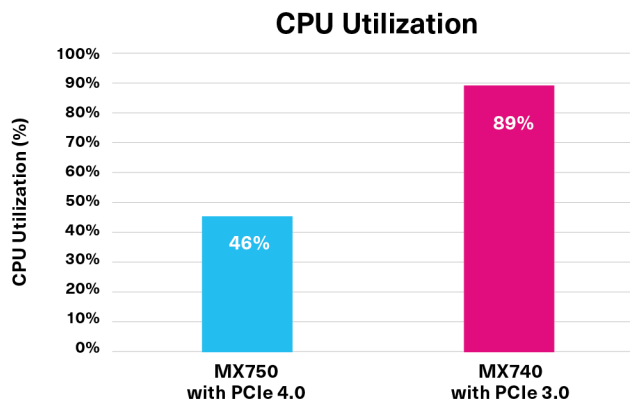
Average Write Latency



CPU Utilization

In general, CPU utilization represents a percentage of the total amount of computing tasks that are handled by the CPU, and is another estimation of system performance. For these tests, CPU utilization was measured to determine the unused CPU that is available for additional tasks, and where a lower result is better. The Dell MX750 server with the Ice Lake CPU has a 20% instructions per cycle (IPC) uplift over the previous CPU generation, which enables it to complete more instructions per clock cycle. The combination of higher transistor density (with the 10 nanometer fabrication process) and PCIe 4.0 technology results in faster storage transactions.

In this test scenario, the Ice Lake CPU was less taxed when running the test workload versus the MX740 server with the Cascade Lake CPU delivering about 46% utilization versus 89% utilization with the Cascade Lake CPU, as evident from the results:



Test Analysis

The test results validated that a PowerEdge MX750 server with PCIe 4.0 technology, such as KIOXIA CM6-V Series SSDs, delivers about 55% better TPM performance versus the previous generation server with PCIe 3.0 technology. Almost a million transactions per minute were delivered within this PCIe 4.0 server/storage platform, which in turn, enables systems and applications based on the PCIe 4.0 interface to run at higher performance. In addition, PCIe 4.0 technology delivers faster access to data per the latency results, enabling data transactions to be completed faster.

Though the total capacity available to each test system differed (512 GB for server setup 1 versus 384 GB for server setup 2), the actual amount of capacity used by each server setup for the MySQL database was the same and did not contribute to any performance advantage.

It should also be noted that the MX740 server with the Cascade Lake CPU was slightly faster in frequency than the MX750 server with the Ice Lake CPU (2.20 GHz vs 2.0GHz), but had less processing cores (24 vs 28). As it related to the testing process, system performance was roughly comparable between the two test systems. Therefore, the combination of higher PCIe 4.0 performance in both the Intel Ice Lake CPU and the KIOXIA CM6-V Series SSDs enabled significantly more transactions per minute and better read/write latency.

At 46% CPU utilization, the MX750 server with the Ice Lake CPU and PCIe 4.0 technology was less taxed when running the test workload versus the MX740 server. At 89% CPU utilization, the MX740 server with Cascade CPU and PCIe 3.0 technology was nearing the maximum workload leaving very little room for the CPU to address additional tasks.

CM6 Series SSD Overview

The CM6 Series is KIOXIA's 3rd generation enterprise-class NVMe SSD product line that features significantly improved performance from PCIe Gen3 to PCIe Gen4, 30.72 TB maximum capacity, dual-port for high availability, 1 DWPD for read-intensive applications (CM6-R Series) and 3 DWPD for mixed use applications (CM6-V Series), up to a 25-watt power envelope and a host of security options – all of which are geared to support a wide variety of workload requirements.

Summary

The test results presented validate that a Dell EMC PowerEdge MX750 PCIe 4.0 enabled server with KIOXIA CM6-V Series SSDs effectively delivered much faster TPM performance for MySQL database workloads than a comparable PCIe 3.0 system/server configuration.

CM6 Series SSDs

***PCIe 4.0 and NVMe 1.4
Specification Compliant***

High-Performance

SeqRead = up to 6,900 MB/s

RanRead = up to 1.4M IOPS

SeqWrite = up to 4,200 MB/s

RanWrite+ up to 350K IOPS

Configurable Flexibility

1 and 3 DWPD options

800 GB - 30,720 GB

NOTES:

¹ Ice Lake is the codename for Intel Corporation's 3rd generation Xeon scalable server processors.

² HammerDB is benchmarking and load testing software that is used to test popular databases. It simulates the stored workloads of multiple virtual users against specific databases to identify transactional scenarios and derive meaningful information about the data environment, such as performance comparisons. TPC Benchmark C is a supported OLTP benchmark that includes a mix of five concurrent transactions of different types, and nine types of tables with a wide range of record and population sizes and where results are measured in transactions per minute.

³ Source: <https://db-engines.com/en/ranking>, January 2021.

⁴ Definition of capacity - KIOXIA Corporation defines a megabyte (MB) as 1,000,000 bytes, a gigabyte (GB) as 1,000,000,000 bytes and a terabyte (TB) as 1,000,000,000,000 bytes. A computer operating system, however, reports storage capacity using powers of 2 for the definition of 1Gbit = 2^{30} bits = 1,073,741,824 bits, 1GB = 2^{30} bytes = 1,073,741,824 bytes and 1TB = 2^{40} bytes = 1,099,511,627,776 bytes and therefore shows less storage capacity. Available storage capacity (including examples of various media files) will vary based on file size, formatting, settings, software and operating system, and/or pre-installed software applications, or media content. Actual formatted capacity may vary.

⁵ Cascade Lake is the codename for Intel Corporation's 2nd generation Xeon scalable server processors.

⁶ Drive Write(s) per Day (DWPD): One full drive write per day means the drive can be written and re-written to full capacity once a day, every day, for the specified lifetime. Actual results may vary due to system configuration, usage, and other factors.

⁷ 2.5-inch indicates the form factor of the SSD and not the drive's physical size.

⁸ Modifying DRAM usage can impact performance. The purpose of this testing was to stress SSD performance. 96GB of DRAM was allocated to prevent the database from being cached into DRAM as that would reduce stress on the storage devices.

Trademarks:

CentOS is a trademark of Red Hat, Inc. in the United States and other countries. Dell, Dell EMC and PowerEdge are either registered trademarks or trademarks of Dell Inc. Intel and Xeon are registered trademarks of Intel Corporation or its subsidiaries in the U.S. and/or other countries. MySQL is a registered trademark of Oracle and/or its affiliates. NVMe is a registered trademark of NVM Express, Inc. PCIe is a registered trademark of PCI-SIG. TPC-C is a trademark of the Transaction Processing Performance Council. All company names, product names and service names may be the trademarks of their respective companies.

Disclaimers:

© 2021 KIOXIA America, Inc. All rights reserved. Information in this application brief, including product specifications, tested content, and assessments are current and believed to be accurate as of the date that the document was published, but is subject to change without prior notice. Technical and application information contained here is subject to the most recent applicable KIOXIA product specifications.