EXTREME SCALE INFRASTRUCTURE

TAKE COMPUTING TO THE EDGE

Building a high-performance enterprise infrastructure using **KIOXIA CM6 Series enterprise NVMe™ SSDs**

DELL Technologies

KIOXIA

EXTREME SCALE INFRASTRUCTURE

For years, the IT industry has been talking about an impending "data tsunami"—an unprecedented wave of data that will need to be stored and analyzed. This idea was primarily driven by an explosion in internet-connected devices; i.e., the internet of things (IoT)—and the desire to consume massive amounts of data from these devices in order to enhance services.

Today, the tsunami is upon us. Enterprise IT organizations must sustain critical application data flow with steadily increasing volumes of data arriving at greater velocity. This data flow is not limited to structured databases, but also includes semi-structured data, such as email, and unstructured data consisting of such items as videos, images, documents and sensor data, among others.

NVMe® SSD storage has brought enterprise IT performance to new levels. Simply put, the benefit of higher IOPS and throughput, along with lower latency cannot be ignored, even with SSD costs significantly higher than that of traditional spinning disks. As evidenced by rapid NVMe SSD adoption, the cost advantage of spinning HDD storage has been mitigated by NVMe SSD performance benefits in meeting the demands of critical enterprise application workloads.

In today's IT environment, many new terms are being heard, such as "digital transformation", "speed time to market", and "elastic compute." In reality, none of these concepts would be possible without NVMe SSDs consuming legions of data arriving from multiple locations simultaneously.

PCIe® 4.0 and NVMe technologies address the needs of enterprise IT organizations by enhancing the ability to quickly capture, consume, visualize, and publish business insights.

Building a DATA-DRIVEN Culture

Organizations are adapting to today's data-driven society. A recent Harvard Business Review survey found that organizations embracing data-driven cultures improved revenue by four times after unlocking rich insights, driving meaningful business decisions, and improving customer satisfaction through the use of data mining and analytics.

To translate a data-driven culture into a data-driven IT infrastructure, IT must anticipate the needs of the organization, including the ability to derive timely and accurate business insights across all types of data. This means that the IT infrastructure is no longer bound by applications, data types, processing engines, organizational boundaries, or data skills in order to transform data into insights.

As such, servers, CPUs, and storage devices must be adaptable in a data-driven IT infrastructure. For example, Dell PowerEdge servers are powered by the AMD EPYC™ CPU, which was designed with more Peripheral Component Interconnect Express (PCIe) lanes and additional Non-Uniform Memory Access (NUMA) nodes than traditional processors.

The significance of these powerful processors, when combined with high density PCIe 4.0 and NVMe SSD solutions, becomes clear during the sustained ingest of massive amounts of data. Dell PowerEdge servers can operate in conjunction with enterprise databases or machine learning applications in the backend, or within any number of IoT functions that require large amounts of data to feed real-time analytics, such as self-driving vehicles, satellite imagery, and weather telemetry.

In this paper, the performance of PCIe 4.0 NVMe SSDs will be evaluated and compared with PCIe 3.0 NVMe SSDs. The importance of throughput, IOPS, and latency within Online Transaction Processing (OLTP) databases and real-time analytics applications will also be discussed. Specifically, the performance of KIOXIA CM6 Series enterprise NVMe SSDs will be compared to the Intel P4510 PCIe 3.0 NVMe SSD.

About KIOXIA

KIOXIA America, Inc. (formerly Toshiba Memory America, Inc.) is a leading worldwide supplier of flash memory and solid-state drives (SSDs). From the original invention of flash memory to today's innovative 3D flash memory technology, BiCS FLASH™, KIOXIA continues to pioneer cutting-edge memory solutions and services that enrich people's lives and expand society's horizons. Focused on SSD end-to-end vertical integration, KIOXIA is a leader in providing outstanding quality with exceptional support.

KIOXIA NVMe SSDs

CM6 Series enterprise NVMe SSDs are part of KIOXIA enterprise SSD family. These enterprise SSDs are suitable for high-performance Tier 0 computing, enabling server and storage systems that require high levels of performance and reliability. KIOXIA enterprise SSDs offer high reliability, with data protection incorporating power-loss-protection (PLP) and encryption technology to support enterprise environments and applications.

Built on PCIe 4.0 and NVMe 1.4 technology, CM6 Series SSDs deliver excellent performance up to 1.4M IOPS (random read), with maximum power consumption of 25 W. Featuring KIOXIA Corporation's 96-layer BiCS FLASH™ 3D TLC memory, the CM6 Series of enterprise NVMe™ SSDs deliver 1 DWPD (Drive Writes Per Day) of endurance and support storage capacities up to 30.72 TB, making them ideally suited for read-intensive enterprise applications.

Please click **here** for detailed specifications of the KIOXIA SSD used in the **Dell PowerEdge Express Flash NVMe PCIe SSD**.

PCIe 3.0 vs. PCIe 4.0 NVMe SSD Benchmark Tests

In order to determine baseline NVMe SSD performance, the Flexible IO command line tool, or FIO, was utilized. FIO generates various user-defined IO workloads, including random read/write and sequential read/write. As such, FIO is an easy and versatile tool to quickly perform benchmark tests on storage devices. Key performance indicators of FIO include the speed of IO processing, measured as IOPS (input/output operations per second), and data throughput, measured in Gb/sec (Gigabytes per second).

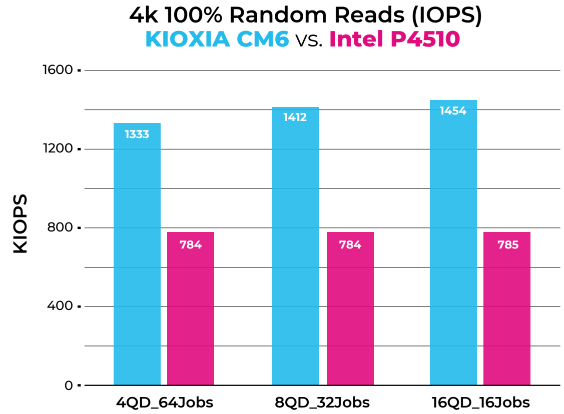

85% more read IOPS

BENCHMARK SUMMARY

The results of the FIO performance benchmarking confirmed that the KIOXIA PCIe 4.0 NVMe SSD showed a remarkable improvement over the PCIe 3.0 NVMe SSD. A maximum of 1.45 Million IOPS was observed with throughput up to 6.8 GB/sec on the PCIe 4.0 NVMe SSD, compared to a maximum 785K Million IOPS with throughput up to 2.53 GB/sec on the PCIe 3.0 NVMe SSD. This performance is consistent with published PCIe 4.0 NVMe specifications.

The charts on this page show results of the FIO benchmarks. Across all types and sizes of baseline tests, the PCIe 4.0 NVMe SSD demonstrated superior storage performance. When benchmarked against PCIe 3.0 NVMe SSDs, the test results validated that the PCIe 4.0 NVMe SSD achieved:

85% more 4K random read IOPS

120% more 4K random write IOPS

168% more sequential read throughput

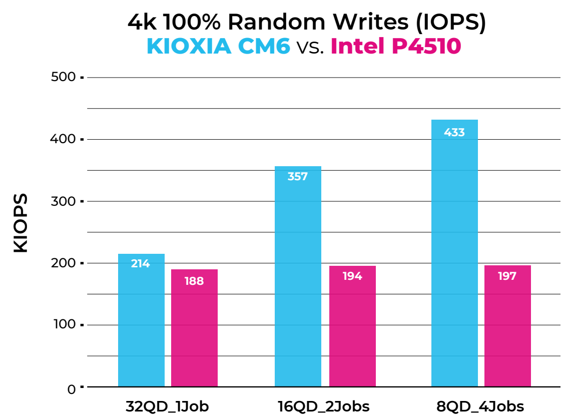

120% more write IOPS

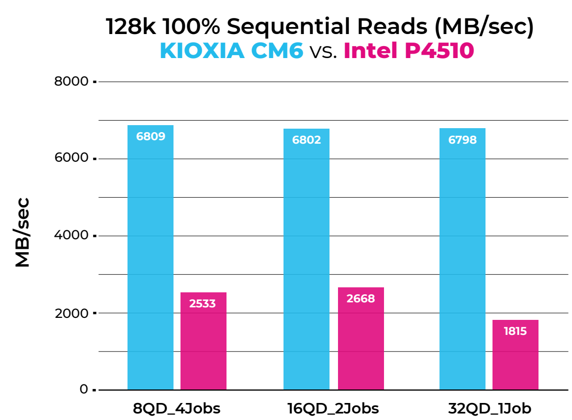

168% more throughput

PCIe 3.0 vs. PCIe 4.0 NVMe SSD Workload Tests

Once baseline SSD performance was established, workload performance tests were performed to further validate PCIe 4.0 SSD performance gains under OLTP-type workloads. HammerDB was used to conduct TPC-C benchmark tests. A popular industry standard open source TPC-C and TPC-H benchmarking tool, HammerDB simulates the workload of multiple virtual users against SQL databases for both transactional and analytic scenarios.

A MySQL 8.x database with its default InnoDB storage engine was deployed on a Dell PowerEdge R6525 server containing dual AMD EPYC 7402 processors and four NVMe SSDs. This configuration was chosen as it simulates a real-world production environment. HammerDB performance was measured by running an OLTP database workload with a "TPC-C-like" parameters. To assure consistency and accuracy, four identical test passes were completed on the above configuration using both PCIe 3.0 and PCIe 4.0 NVMe SSDs.

The key performance indicators provided by the HammerDB tests were TPM (Transactions per Minute), NOPM (New Order Transactions per Minute) and Latency (measured in milliseconds).

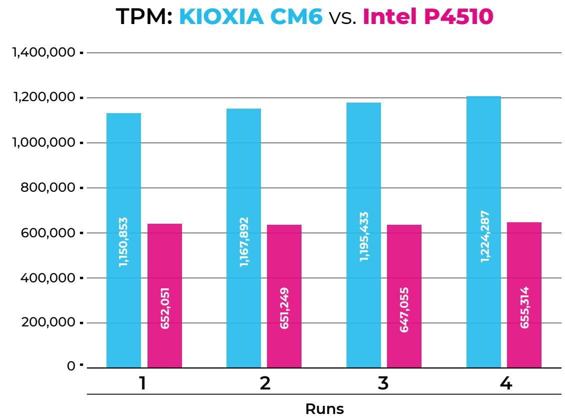

86% more TPM

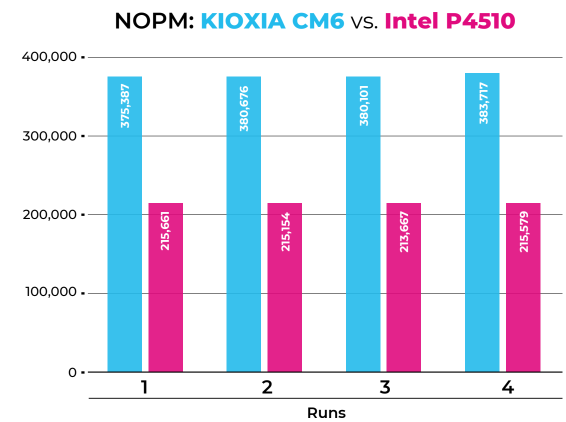

78% more NOPM

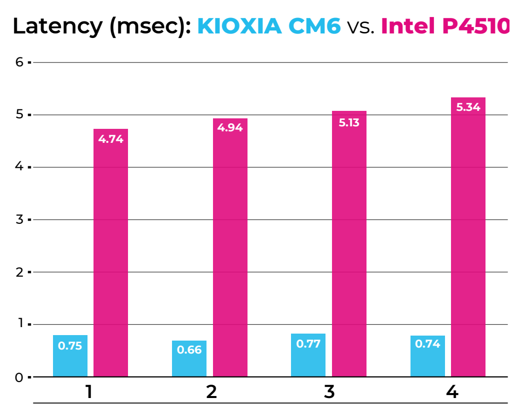

86% reduced latency

WORKLOAD SUMMARY

When database performances for OLTP workloads are compared, the PCIe 4.0 NVMe SSD achieved and demonstrated a substantial storage performance advantage over PCIe 3.0 NVMe SSD.

86% more transactions per minute (TPM)

78% more new orders per minute (NOPM)

86% reduction in latency

CONCLUSION

Dell Technologies PowerEdge servers using KIOXIA PCIe 4.0 NVMe SSDs handle diverse read and write workloads very efficiently, empowering administrators to fully support their IO-intensive enterprise applications.

In baseline performance tests, the PCIe 4.0 NVMe SSD recorded results in line with published specifications. With 1.5 million read IOPS, 433K write IOPS, and 6.8 GB/sec sustained read throughput, PCIe 4.0 NVMe SSDs far surpass PCIe 3.0 performance. This performance gain is due to the fact that PCIe 4.0 NVMe SSDs leverage the latest PCIe technology, and are not constrained by legacy protocols such as SAS and SATA. As such, PCIe 4.0 NVMe SSDs brings true storage-class memory (SCM) to IT organizations.

In the workload tests, the PCIe 4.0 NVMe SSD again proved its mettle. Demonstrating 1.2 million TPM, 384K NOPM, and ultra-low 0.74 msec sustained latency, PCIe 4.0 NVMe SSD brings critical advantages to real-world application scenarios, including:

1. **PCIe 4.0 NVMe SSDs optimize virtualized environments by increasing the number of virtual machines (VMs) supported per server. Virtual environments are often limited by the latency or IOPS ceiling of each server, leading to higher costs as additional servers are deployed to optimize VM workloads. PCIe 4.0 NVMe SSDs enable virtual environments to boost network speed and server performance *without* the need for additional servers or complex virtual partitioning.**

2. For many industries, such as healthcare, finance and telecom, any slowdown in data flow has devasting results. These companies exist and compete because of extremely fast servers with high-performance, low latency storage. Leveraging PCIe 4.0 NVMe SSDs for complex workloads keeps these companies competitive by nearly eliminating processor wait times when reading data from storage.

3. **OLTP databases and big data analytic applications also benefit from high-performance PCIe 4.0 NVMe SSDs. In databases, businesses can use NVMe SSDs as storage-class memory to pin metadata and indexes without slowing query speeds, ultimately improving database performance. In analytics, business analysts can make real-time decisions with quick, accessible, data since their data intensive workloads are no longer inhibited by performance bottlenecks.**

Based on the findings in this paper, if your real-time data is growing exponentially and requires the ultimate in IO performance, the platform utilized for the tests described in this paper—Dell PowerEdge family with KIOXIA PCIe 4.0 CM6 NVMe SSDs—is an optimal solution for critical enterprise OLTP, data analytics and distributed database applications.

BENCHMARK TEST NOTES

1

For all tests, a dual socket, 48-core, Dell PowerEdge R6525 server was utilized with the following parameters:

System CPU (2)	AMD EPYC™ 7402 (24 core)
Total number of cores	48
Processor Core Speed	2.80 GHz
Processor Bus Speed	16 GT/s
Family-Model-Stepping	17-31-0
Level 2 Cache	24x512 KB per CPU
Level 3 Cache	128 MB per CPU
PCIe 3.0 SSDs	(4) Dell NVMe® P4510 4TB SFF (Intel)
PCIe 4.0 SSDs	(4) Dell NVMe® CM6 RI 3.84TB (KIOXIA)

2

Before running any SSD tests, it is important to prepare the SSD by "preconditioning". The preconditioning process uses three steps to ensure that benchmarking results are accurate and repeatable. It is recommended to run the following workloads with twice the advertised capacity of the SSD to guarantee that all available memory is filled with data, including the factory provisioned area:

1. Secure erase the SSD
2. Fill SSD with 128k sequential data
3. Fill SSD with 4k random data

To better understand the importance of SSD preconditioning, please click HERE.

3

For HammerDB tests, the MySQL Innodb engine was optimized as follows:

innodb_pagesize=16K (Default)
innodb_log_file_size=1024M
innodb_log_files_in_group=32
innodb_buffer_pool_size=16G
innodb_buffer_pool_instances=16

innodb_log_buffer_size=64M
innodb_io_capacity=1000000
innodb_io_capacity_max=2000000
innodb_read_io_threads=16
innodb_write_io_threads=16

DELL Technologies   KIOXIA