KIOXIA

Application Brief

Accelerating Database Workloads with Ampere® Altra® Processors and New KIOXIA CD8 Series Data Center NVMe™ SSDs

Today's data centers collect an overwhelming amount of data on premises and at the edge of the network. A good portion of this data, such as sales transactions, customer/financial/product/vendor information, etc., is stored, maintained and accessed in relational databases represented in tables comprised of rows. To better address larger datasets, IT teams are scaling storage in the data center and the databases that are contained on them to deliver fast storage device performance to end-users. To meet this objective, IT teams are transitioning to NVMe SSDs given their very fast read and write Input/Output (I/O) capabilities.

Additionally, servers equipped with Arm® based processors are gaining popularity with hyperscalers and may soon spread more widely in the data center industry[1]. When properly implemented, a server equipped with Arm technology can deliver greater processing power while using less energy and requiring less cooling than similar technology based on x86-class processors[2], which are important considerations for most enterprise data centers.

The focus of this application brief is how database workloads can be accelerated using servers equipped with an Arm based processor and NVMe SSDs. For this brief, KIOXIA CD8 Series NVMe data center SSDs were deployed in a database server and powered by an Ampere Altra 80C processor. When tested against NVMe data center SSDs from another vendor, database acceleration was improved even further. Presented herein are the test objectives, setup, procedures and results of PCIe® 4.0 NVMe SSDs deployed on a server equipped with an Arm based processor.

The tests also monitored CPU power consumption. When paired with KIOXIA CD8 Series SSDs, the Ampere Altra 80C processor consumed only 142 watts while delivering 1.1 million database transactions per minute (TPM).

Test Objectives

- *Create a MySQL® database server on a server platform using an Ampere Altra processor running an Ubuntu® operating system (Linux® kernel)*
- *Use MySQL, a very commonly used database application, and store its data and log files on NVMe SSDs in a software RAID10 set through mdadm[3] — the buffer pool was limited to 64 gigabytes[4] (GB) of RAM in order to best showcase the performance that the SSDs themselves can contribute to database throughput (versus DRAM doing most of the work)*
- *Test MySQL database performance using HammerDB[5] benchmark software with a TPROC-C[6] workload on KIOXIA CD8 Series NVMe data center SSDs, as well as competitor SSDs*
- *Collect the performance (and associated CPU) metrics covering: Database Throughput; Drive Read and Write Latencies; Transaction Latencies (i.e. New Order, Payment; Delivery; Stock Level and Order Status), and CPU Utilization*

Products Tested

<u>*Ampere® Altra® Processor:*</u>

Ampere Altra processors include up to eighty processing cores and use Arm technology. They are typically deployed in high-performance data centers and power-constrained edge environments including use cases that require high-performance, scalable and sustainable cloud native solutions including data analytics, artificial intelligence (AI), database storage, telco stacks, edge computing and web hosting.

More information: https://amperecomputing.com/processors/ampere-altra/.

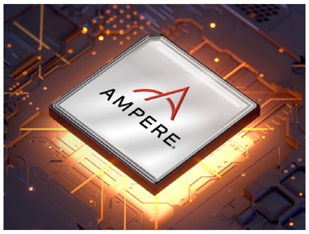

Used with permission from Ampere Computing

KIOXIA CD8 Series SSD[9]

KIOXIA CD8 Series NVMe Data Center SSDs:

KIOXIA recently launched its latest generation CD8 Series NVMe data center SSDs positioned for scale-out and cloud applications. The series is available in a 2.5-inch[7] form factor with supported capacities up to 15.36 terabytes[4] (TB) and in two configurations:

1. *CD8-R Series: Read-intensive supporting 1 DWPD[8]*
2. *CD8-V Series: Mixed-use supporting 3 DWPD*

More information: https://americas.kioxia.com/en-us/business/ssd/data-center-ssd.html.

MySQL Database Application

A MySQL relational database management system (RDBMS) was used for this comparison as it is versatile in both hyperscale and enterprise environments. It supports key applications such as online transaction processing (OLTP), e-commerce and data warehousing. It is often used for backend webserver data storage. MySQL is the most widely deployed open source database in the world today (ranked number two overall[10]).

Test Configuration

The hardware and software equipment used to run the tests described in this brief includes:

Database Server Information	
Model	Gigabyte R272-P30-JG
Number of CPU Sockets	1
CPU	Ampere® Altra® 80C
No. of CPU Cores	80
CPU Frequency	3.0 GHz
Total Memory	256 GB DDR4 DRAM
Memory Frequency	2.933 GHz
Operating System Information	
Operating System	Ubuntu
Version	22.04.1
Kernel	5.15.0-52-generic aarch64
mdadm	v4.2

Load Generator Server Information	
Model	HPE® ProLiant® DL385 Gen10 Plus v2
Number of CPU Sockets	2
CPU	AMD EPYC™ 7543
No. of CPU Cores	32
CPU Frequency	2.8 GHz
Total Memory	256 GB DDR4 DRAM
Memory Frequency	3.2 GHz
Operating System Information	
Operating System	Ubuntu
Version	22.04.1
Kernel	5.15.0-52-generic x86_64
mdadm	v4.2

Test Software Information	
Software	HammerDB
Version	4.5
Number of Virtual Users	64
User Repeat	1 ms
User Delay	1 ms

SSD Information		
Model	KIOXIA CD8-R Series	Vendor A
Number of Devices	4	4
Form Factor	2.5-inch	2.5-inch
Interface	PCIe 4.0	PCIe 4.0
Capacity	3.84 TB	3.84 TB
Drive Write(s) per Day	1 (5 years)	1 (5 years)
Active Power	14 W	16.6 W

***NOTE:** For further information on the **MySQL parameters** used for this comparison, refer to the **Addendum**.*

Test Setup and Procedures

The Gigabyte R272-P30-JG server was configured with an Ubuntu 22.04.1 Linux operating system, MySQL RDBMS version 8.0.31 and four CD8-R Series 3.84 TB SSDs. Each drive was configured in a software RAID 10[11] set for storing the MySQL database.

An HPE ProLiant DL385 Gen10 Plus v2 server was then configured as a load generator to drive database transactions to the MySQL (Gigabyte) server. HammerDB test software was deployed on this server and used to create a database containing 1,000 data warehouses[12] on the KIOXIA CD8-R Series SSD RAID 10 set. A TPROC-C workload from HammerDB software was then run against the MySQL database.

Metrics were recorded from HammerDB software and the TPROC-C workload that included database throughput (transactions per minute), drive read and write latency, transaction latencies (i.e., new order, payment, delivery, stock level and order status), and CPU utilization.

Once the workload completed, HammerDB software was then used to create a database on the Vendor A competitive SSDs RAID 10 set and the TPROC-C workload was then run against the MySQL database. Metrics covering the identical tests were recorded and compared to the KIOXIA CD8-R Series SSD results.

Test Results

Database Throughput (in TPM):
Database throughput measures how many transactions in the TPROC-C workload are being executed per minute. The HammerDB software, executing the TPROC-C transaction profile, randomly performed new order, payment, order status, delivery and stock level transactions. This benchmark simulates an OLTP environment where there are a large number of users that conduct simple short transactions that require sub-second response times and return relatively few records. *For these tests, the higher result is better.*

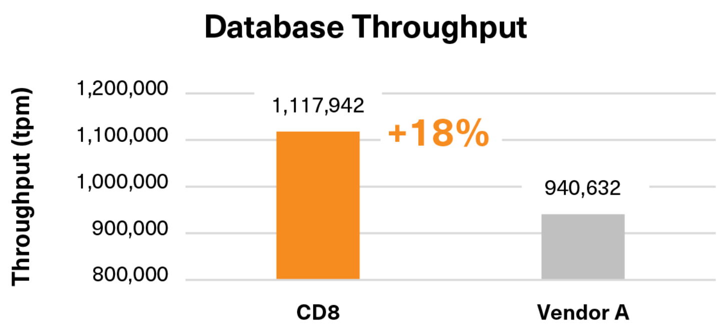

Drive Read Latency (in ms):
Drive read latency represents the time it takes to perform a drive read operation. It includes the average time it takes for the read operation to successfully complete once the drive has received the read operation request issued from the workload generator. This metric can positively impact database performance and application response times that can translate into a better user database experience. *For these tests, the lower result is better.*

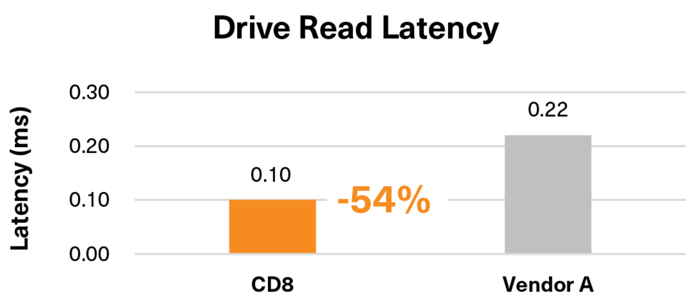

Drive Write Latency (in ms):
Drive write latency represents the time it takes to perform a drive write operation and is the same description as Drive Read Latency except for write operations. *For these tests, the lower result is better.*

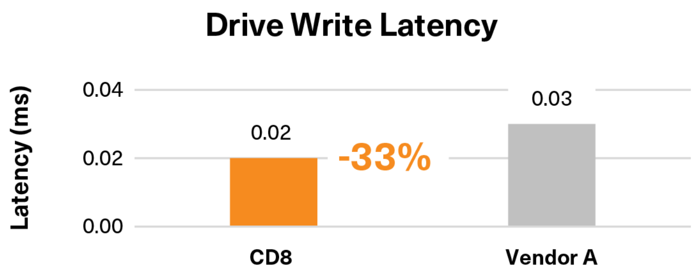

New Order Transaction Latency[13] (in ms):
"The New-Order business transaction consists of entering a complete order through a single database transaction. It represents a mid-weight, read-write transaction with a high frequency of execution and stringent response time requirements to satisfy online users. This transaction is the backbone of the workload. It is designed to place a variable load on the system to reflect online database activity as typically found in production environments."
For these tests, the lower result is better.

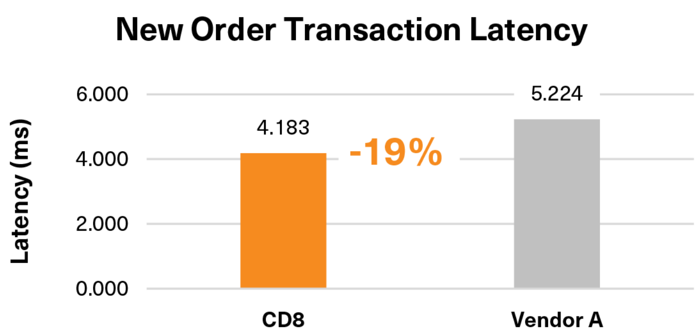

Payment Transaction Latency[13] (in ms):

"The Payment business transaction updates the customer's balance and reflects the payment on the district and warehouse sales statistics. It represents a light-weight, read-write transaction with a high frequency of execution and stringent response time requirements to satisfy online users." *For these tests, the lower result is better.*

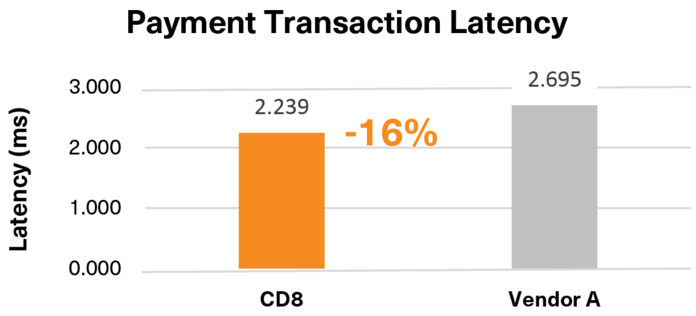

Delivery Transaction Latency[13] (in ms):

"The Delivery business transaction consists of processing a batch of 10 new (not yet delivered) orders. Each order is processed (delivered) in full within the scope of a read-write database transaction. The number of orders delivered as a group (or batched) within the same database transaction is implementation specific. The business transaction, comprised of one or more (up to 10) database transactions, has a low frequency of execution and must complete within a relaxed response time requirement." *For these tests, the lower result is better.*

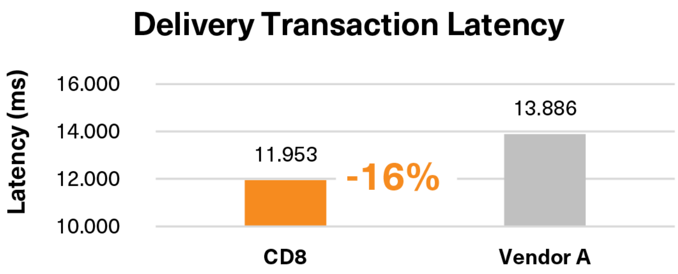

Stock Level Transaction Latency[13] (in ms):

"The Stock-Level business transaction determines the number of recently sold items that have a stock level below a specified threshold. It represents a heavy read-only database transaction with a low frequency of execution, a relaxed response time requirement, and relaxed consistency requirements." *For these tests, the lower result is better.*

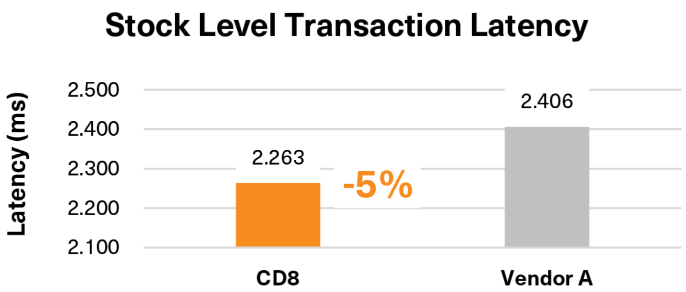

Order Status Transaction Latency[13] (in ms):
"The Order-Status business transaction queries the status of a customer's last order. It represents a mid-weight read-only database transaction with a low frequency of execution and response time requirement to satisfy online users." *For these tests, the lower result is better.*

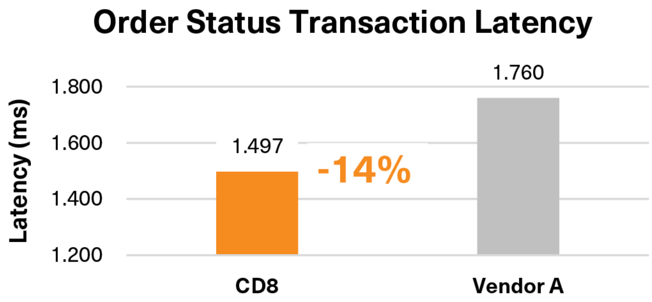

CPU Utilization (in percentages):
This metric represents a percentage of the CPU cycles that are being used for a given workload and measured to ensure that the Ampere Altra CPU was efficiently processing the database workloads. Low utilization means that the CPU was not being used efficiently, which could result in an underutilization of server capabilities and stranded compute resources. Efficient CPU utilization indicates there are fewer bottlenecks limiting application performance, which in turn enables the application to utilize more CPU cycles to deliver higher performance. For the KIOXIA CD8-R Series tests, the Ampere Altra processor delivered an efficient CPU utilization of 69.4%.

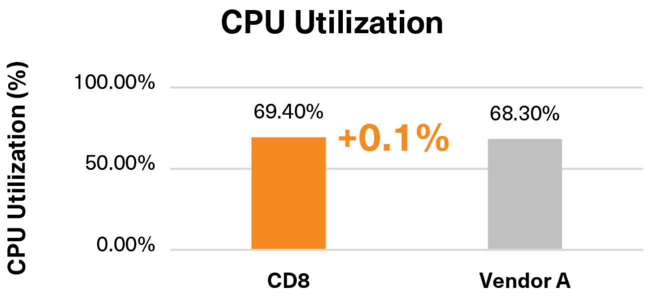

Test Analysis

The test results indicate that the KIOXIA CD8-R Series SSD configuration, when driven by an Ampere Altra 80C processor, accelerated database performance in all of the tested performance areas as follows:

Test Software Information		CD8-R Series 64 virtual users	Vendor A 64 virtual users	CD8-R Series Advantages
Database Throughput	*(higher is better)*	1,117,942 tpm	940,632 tpm	18%
Drive Read Latency	*(lower is better)*	0.10 ms	0.22 ms	54%
Drive Write Latency	*(lower is better)*	0.02 ms	0.03 ms	33%
New Order Transaction Latency	*(lower is better)*	4.183 ms	5.224 ms	19%
Payment Transaction Latency	*(lower is better)*	2.239 ms	2.695 ms	16%
Delivery Transaction Latency	*(lower is better)*	11.953 ms	13.886 ms	13%
Stock Level Transaction Latency	*(lower is better)*	2.263 ms	2.406 ms	5%
Order Status Transaction Latency	*(lower is better)*	1.497 ms	1.760 ms	14%
CPU Utilization	*(higher is better)*	69.40%	68.30%	0.1%

The Ampere Altra 80C processor delivered a high CPU utilization enabling the MySQL database application to utilize 69.4% of the CPU cycles for application performance. Additional monitoring showed that the Ampere Altra 80C processor only drew 142 watts in the data center, which is a great metric when compared to x86 CPUs that can typically draw 250 to 275 watts. The cost savings related to server wattage when scaled out to any sized data center could be significant.

Summary

The combination of KIOXIA CD8 Series SSDs and an Ampere Altra 80C processor accelerated database performance over data center SSDs from another vendor. The test results demonstrated low database latencies and high workload throughput that can enable very dense storage solutions while delivering high application performance. These results can also help to meet critical business needs and can heighten user experiences. This configuration also showed that the Ampere Altra CPU was efficiently processing the database workloads, which in turn can reduce data center costs without sacrificing application performance.

NOTES:

[1] Source: Data Center Knowledge, 'Arm Chips Gaining in Data Center, But Still in Single Digits, Maria Korolov, August 29, 2022, https://www.datacenterknowledge.com/arm/arm-chips-gaining-data-centers-still-single-digits.

[2] Source: Tech Target, Search Data Center Definitions, 'ARM Server,' Stephen Bigelow, https://www.techtarget.com/searchdatacenter/definition/advanced-RISC-machine-ARM.

[3] mdadm is a standard Linux operating system tool that manages and monitors software RAID sets and devices.

[4] Definition of capacity - KIOXIA Corporation defines a kilobyte (KB) as 1,000 bytes, a megabyte (MB) as 1,000,000 bytes, a gigabyte (GB) as 1,000,000,000 bytes, a terabyte (TB) as 1,000,000,000,000 bytes, and a petabyte (PB) as 1,000,000,000,000,000 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, 1TB = 2^{40} bytes = 1,099,511,627,776 bytes, and 1PB = 2^{50} bytes = 1,125,899,906,842,624 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] 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.

[6] TPROC-C is the OLTP workload implemented in HammerDB derived from the TPC-C™ specification with modification to make running HammerDB straightforward and cost-effective on any of the supported database environments. The HammerDB TPROC-C workload is an open source workload derived from the TPC-C Benchmark Standard and as such is not comparable to published TPC-C results, as the results comply with a subset rather than the full TPC-C Benchmark Standard. TPROC-C means Transaction Processing Benchmark derived from the TPC "C" specification.

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

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

[9] Product image may represent a design model.

[10] Source: https://db-engines.com/en/ranking, January 2023.

[11] RAID 10 is a RAID configuration that combines disk mirroring and disk striping to protect data. It requires a minimum of four drives and stripes data across mirrored pairs.

[12] The 1,000 data warehouses, in combination with the 64G buffer pool size, represents a database configuration that creates a typical database size.

[13] Source: TPC Benchmark™C Standard Specification, Revision 5.11, Clause 2: Transaction and Terminal Profiles, Sections 2.4-2.8, February 2010, https://www.tpc.org/tpc_documents_current_versions/pdf/tpc-c_v5.11.0.pdf.

ADDENDUM

MySQL Parameters

MySQL Parameters	
Version	8.0.31-0ubuntu0.22.04.1
innodb_adaptive_flushing	1
innodb_flush_neighbors	0
innodb_write_io_threads	64
innodb_read_io_threads	64
innodb_thread_concurrency	0
innodb_purge_threads	4
innodb_adaptive_hash_index	0
innodb_file_per_table	
innodb_log_file_size	1024M
innodb_log_files_in_group	32
innodb_open_files	4000
innodb_buffer_pool_size	64G
innodb_buffer_pool_instances	16
innodb_log_buffer_size	64M
innodb_doublewrite	0
innodb_thread_concurrency	0
innodb_flush_log_at_trx_commit	2
innodb_max_dirty_pages_pct	90
innodb_max_dirty_pages_pct_lwm	10
join_buffer_size	32K
sort_buffer_size	32K
innodb_use_native_aio	1
innodb_stats_persistent	1
innodb_spin_wait_delay	6
innodb_max_purge_lag_delay	300000
innodb_max_purge_lag	0
innodb_flush_method	O_DIRECT_NO_FSYNC
innodb_checksum_algorithm	none
innodb_io_capacity	760000
innodb_io_capacity_max	800000
innodb_lru_scan_depth	9000
innodb_change_buffering	none
innodb_read_only	0
innodb_page_cleaners	4
innodb_undo_log_truncate	off
innodb_log_compressed_pages	OFF