Micron® SSDs and Ampere® eMAG™ CPUs Drive Efficient Data Analytics

Architecting high-value solutions begins with efficient, high-performance QLC SSDs and ARM® technology

Overview

As cloud computing becomes mainstream, more computing power is being concentrated in larger, centralized data centers to increase efficiency at scale. This places power and hardware costs front-and-center for the data center. And while customers are looking for more efficient alternatives that enable continued growth of their solutions, they have reached a breaking point with existing compute architectures and slow, power-hungry spinning media.

With the availability of Micron's quad-level cell (QLC)-based 5210 ION SSD, flash storage has reached a point where the economics of SSDs now challenge those of HDD solutions in mainstream applications. Providing massive performance in throughput, latency, and I/O per device in a cost-effective package, QLC SSDs like the 5210 ION are a low-cost capacity choice of professionals around the world.

Likewise, the increasing growth of edge computing—using devices in disparate geographical locations numbering in the thousands—and its dependence on smaller form factors that must fit within increasingly restrictive physical and power requirements, also demands efficient, well-designed solutions that provide more compute and storage than has been available in the past.

Simultaneously, edge and cloud increasingly demand better efficiency, resulting in a new trend in computer architectures that involves alternatives to the legacy x86 compute as well as alternatives to spinning media-based HDD storage. But how effective are these new options?

While x86-based compute has been the standard in the data center for more than 30 years, the world of mobile computing has brought a new player into the data center market in the form of ARM processors. Ampere is extending ARM processor technology into the data center with their eMAG processor; making a new case for low-power, high-performance compute.

This brief highlights what is possible when Micron QLC SSDs and Ampere eMAG CPUs are combined to create a low-power, cost-efficient computing alternative to the status quo. Offering more operations per watt than legacy solutions, all-flash solutions using ARM technology-based servers become a compelling option for those concerned about long-term TCO.

All Flash vs. HDD and eMAG ARM vs. x86 Servers

This brief looks at the value of QLC SSDs vs. legacy 10K RPM HDDs and how an eMAG server-based solution compares to an equivalent x86 server-based configuration when running a common cloud-centric application—the Apache Cassandra® NoSQL database scale-out solution. To make these comparisons, we use the industry accepted Yahoo! Cloud Serving Benchmark (YCSB) family of tests.

Using this framework, we make two comparisons:

- Micron SATA SSDs using QLC vs. 10K RPM SATA HDDs
- Ampere eMAG ARM processor-based servers vs. x86 servers

QLC NAND Brings Competition to Legacy HDD-based Solutions

Building solutions that minimize TCO traditionally involve using lower-cost HDDs as the primary storage. The challenge with this is, lower TCO is not always attainable due to the low performance of individual HDDs.

HDD-based solutions increase performance by adding numerous drives to boost overall I/O performance to acceptable levels. This results in two things: an increase in the number of servers that are required to host the larger number of HDDs and an increase in power and excess capacity, both of which negate the lower cost of the HDDs and thus minimize the TCO advantages.

With the release of the latest in flash technology, QLC NAND, Micron has changed the TCO game. QLC SSDs offer I/O performance orders of magnitude better than HDDs, enabling deployment of fewer drives, better utilizing higher capacity, and thus reducing higher server demand. The result is improved overall solution performance and lower operating cost—better TCO.

To illustrate this, we executed several YCSB benchmark tests against identical Cassandra database clusters that differed only in the storage solution:

- 4X 1.92TB Micron 5210 ION SATA SSDs using QLC NAND
- 4X 2.4TB enterprise hybrid SAS HDDs using 10K RPM speeds

The following figures show database operations per second and average operations per WATT of power consumed.

Reviewing the results, we expected the SSD configuration would be much faster than the HDD configuration, and that was validated here. Depending on the workload profile, we experienced up to a 39X improvement in operations completed per second with an equivalent level of power efficiency. Because of the much higher performance of SSDs, the power efficiency—the number of operations per WATT of power consumed—goes up significantly. The results show SSD I/O performance an order of magnitude faster than HDDs, enabling you to get more work done with fewer servers, fewer drives, and thus less power.

Figure 1 and tables 1 and 2 illustrate the results, organized by YCSB workload. (A brief workload description and use cases are presented later in this brief.)

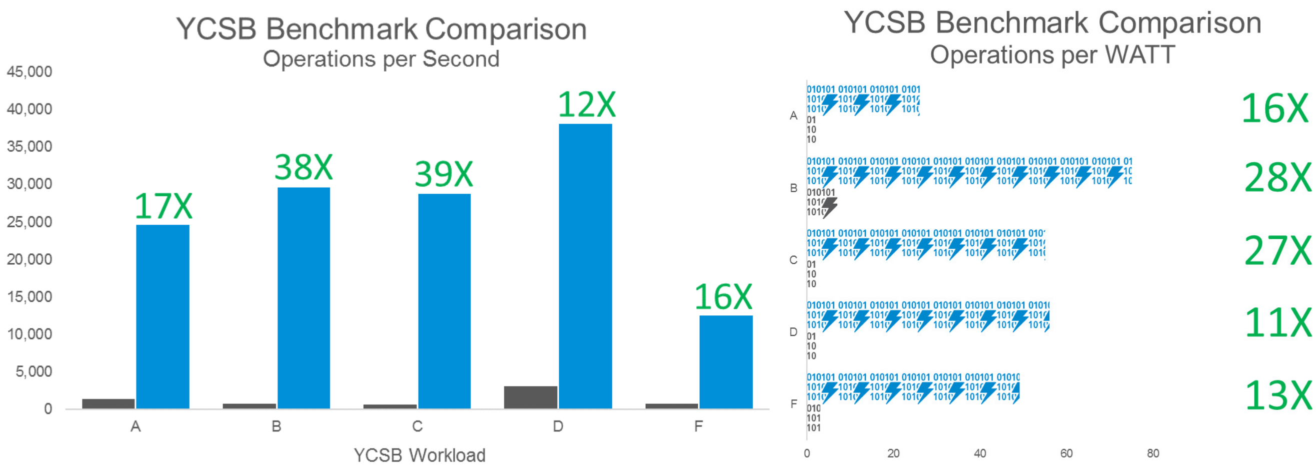

Figure 1a and 1b: Micron 5210 SSD (Blue) vs. 10K RPM HDDs (Gray) for Cloud-centric Workloads Comparison

YCSB Workload	10K RPM HDD	Micron 5210 ION
Workload A	1436	24,766
Workload B	776	29,763
Workload C	737	28,908
Workload D	3167	38,219
Workload F	767	12,580

Table 1: Total Database Operations for Cloud-centric Workloads: HDD vs. SSD

YCSB Workload	10K RPM HDD	Micron 5210 ION
Workload A	3 Ops/Watt	49 Ops/Watt
Workload B	1.5 Ops/Watt	56 Ops/Watt
Workload C	1.5 Ops/Watt	55 Ops/Watt
Workload D	1.7 Ops/Watt	75 Ops/Watt
Workload F	1.5 Ops/Watt	26 Ops/Watt

Table 2: Operations per Watt for Cloud-centric Workloads: HDD vs. SSD

Ampere eMAG Offers Efficient, Cost-Effective Performance for the Cloud

While storage performance and efficiency are major contributors to overall solution performance and cost, an equally impactful component of any solution is the server. To understand this, we next looked at how newly available Ampere eMAG CPU-based servers might impact the overall success of a cloud infrastructure. Ampere eMAG provides new opportunities to build efficient, lower-cost platforms for cloud-centric solutions.

The ARM processor-based eMAG and legacy x86 servers in our testing used an identical all-flash 5210 ION storage configuration. We used a Lenovo® HR350A server, powered by the Ampere eMAG processor, which provides 32 single-threaded 64-bit ARMv8 cores in a single socket design. For comparison, the x86 servers used a common dual-socket implementation, providing a total of 28 hyper-thread 64-bit cores for a total of 56 threads per server. Each cluster consisted of four Cassandra database nodes. (For more details, refer to the "How We Tested" section at the end of this brief.)

As the data below shows, Ampere eMAG CPUs provide the performance and efficiency demanded by enterprise-class analytics solutions when compared to legacy x86 based solutions of similar design in terms of database operations per second, database operations per CPU thread, and database operations per watt (W) of power consumed.

Figure 2 illustrates that eMAG processor-based servers can provide similar performance to that of a legacy x86-based solution. The x86 cluster exceeded the Ampere eMAG server performance on three YCSB workloads, but the eMAG server cluster outperformed the x86 solution on workloads that were heavily dependent on reads vs. writes or updates.

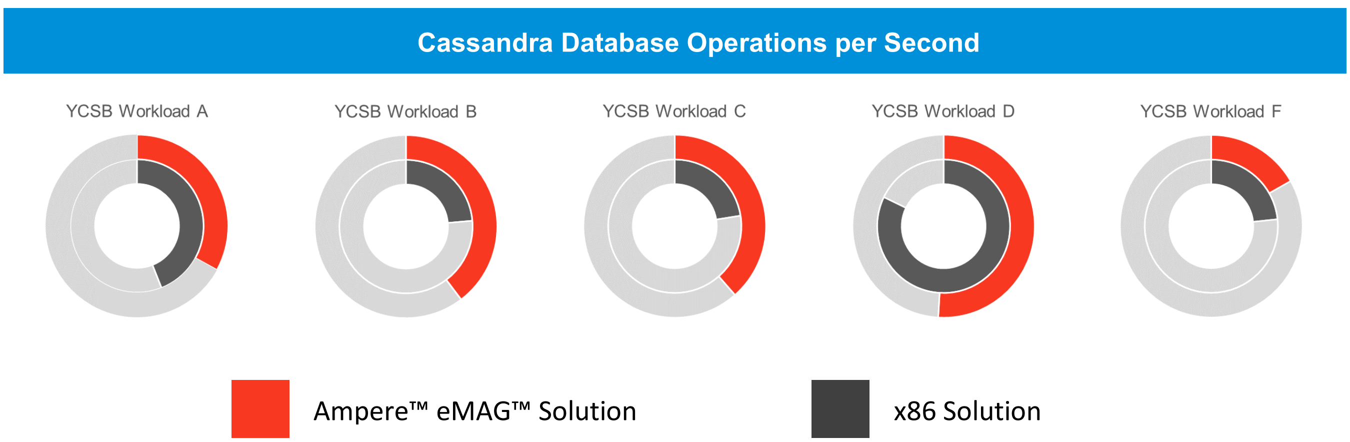

Figure 2: Cassandra Database Performance Comparison of eMAG vs. x86 Architectures for Cloud-centric Workloads

Where eMAG really shows its value is operational efficiency. When we compare database operations executed on a per-thread basis, the ARM processor-based solution exceeded that of the x86 solution for every workload. Operations per thread was improved using eMAG by up to 3X depending on workload (Figure 3).

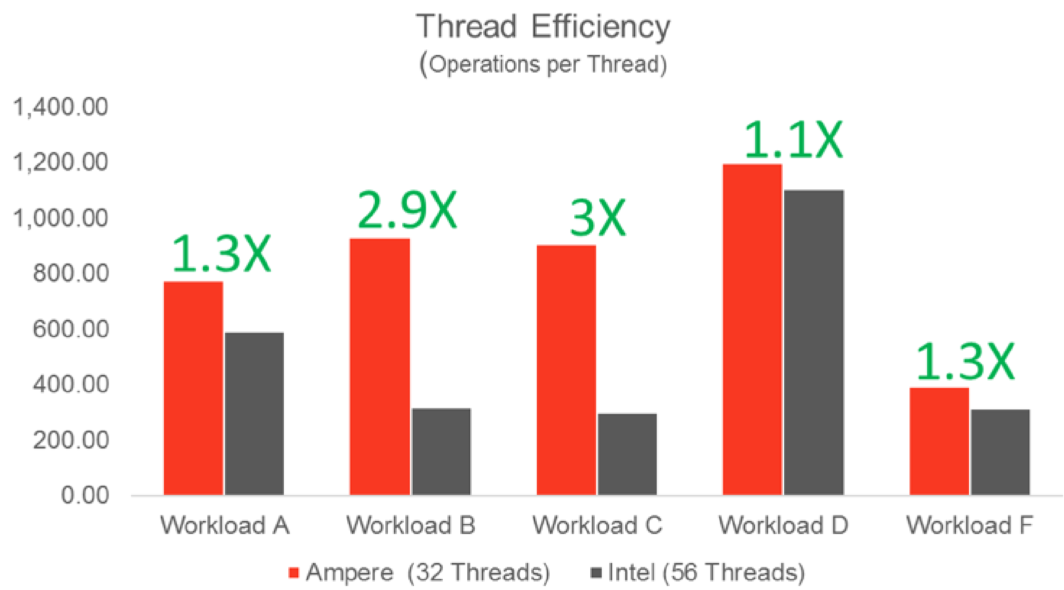

Figure 3: Compute Efficiency Comparison of eMAG vs. x86 Architectures for Cloud-centric Workloads

The amount of processing achieved per watt of power consumed speaks well of the eMAG processor's low-power design. As shown in Figure 4, each YCSB workload showed significantly better power efficiency for ARM vs. x86-based solutions. As dependence on reads increased (Workloads B and C), eMAG showed a 3X gain in Cassandra operations per watt of power consumed.

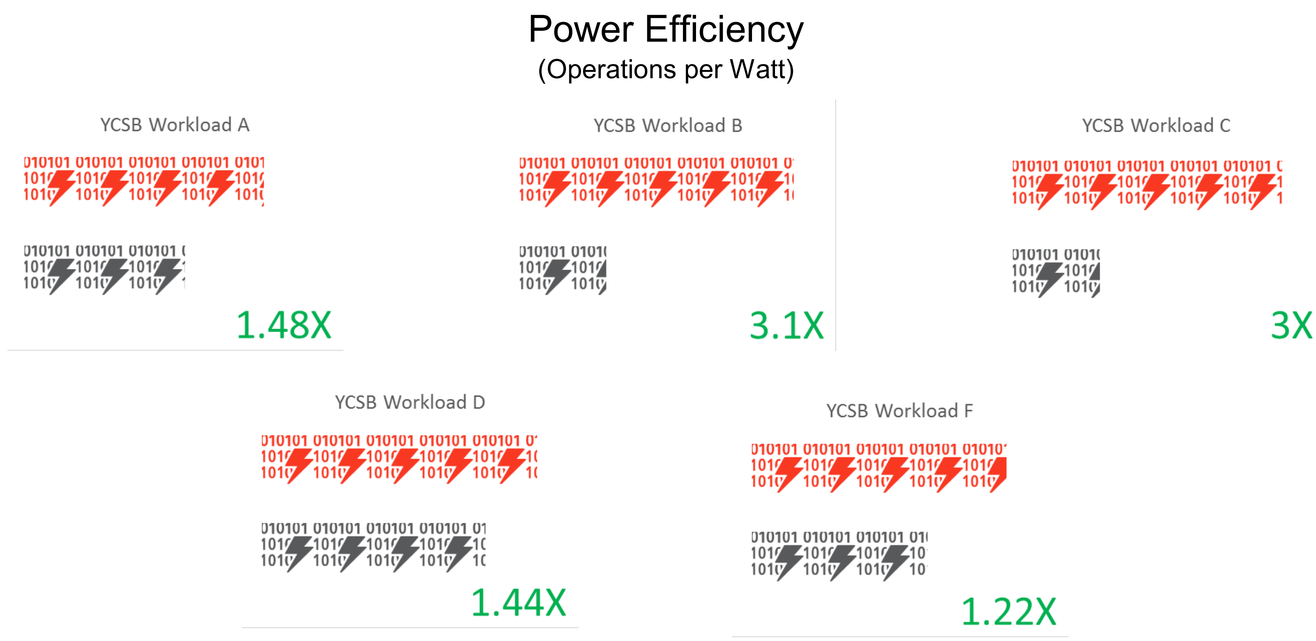

Figure 4: Power Efficiency Comparison of eMAG (Red) vs. x86 Architectures (Gray) for Cloud-centric Workloads

YCSB Workload	Database Operations per Second		Database Operations per Thread	
	x86	Ampere eMAG	x86 (56 Threads)	Ampere eMAG (32 Threads)
Workload A	33,013	24,766	589.52	773.94
Workload B	17,701	29,763	316.09	930.09
Workload C	16,854	28,908	300.96	903.38
Workload D	61,647	38,219	1100.84	1194.34
Workload F	17,506	12,580	312.61	393.13

Table 3: Total Database Operations for Cloud-centric Workloads: x86 vs. eMAG

YCSB Workload	x86	Ampere eMAG
Workload A	33 Ops/W	49 Ops/W
Workload B	18 Ops/W	56 Ops/W
Workload C	18 Ops/W	55 Ops/W
Workload D	61 Ops/W	75 Ops/W
Workload F	18 Ops/W	26 Ops/W

Table 4: Operations per Watt for Cloud-centric Workloads: x86 vs. eMAG

The Bottom Line

As more businesses rely on big data analytics and build out cloud-centric solutions, which can be heavily dependent on read disk operations, solutions built using ARM processor-based technology—such as the Ampere eMAG CPU—can have an distinct advantage. Ampere eMAG servers provide the expansion and performance of legacy x86 solutions while lowering power consumption to create a very compelling solution for big data and cloud-centric solutions. When combined with the additional benefits of low-cost, high-performance SSDs—such as the Micron 5210 ION, designed specifically for read-centric solutions—the benefits become even greater.

This brief demonstrates the ARM processor-based Ampere eMAG processor provides better efficiency for read-centric workloads. For workloads where x86 raw performance exceeded that of ARM, the power efficiency and thread efficiency benefit provided by eMAG CPUs compensate for potentially having to deploy additional nodes to come to parity with x86 for those workloads.

Learn more about the Micron 5210 ION and the complete line of Micron NVMe SSDs at www.micron.com.

Learn more about the Ampere eMAG family of solutions at www.amperecomputing.com.

How We Tested

The testing methodology for this brief attempts to represent a real world deployment and usage scenario for a Cassandra database.

- Four database nodes in cluster configuration
- 1TB database size distributed across all nodes of cluster resulting in the configuration shown in Table 5
- Replication factor for the database set to 3 (3TB of data stored across all nodes in cluster)

Server	Capacity	Tokens	Percent Owned
node01	752.66 GB	256	73.3%
node02	772.52 GB	256	75.2%
node03	783.19 GB	256	76.3%
node04	771.34 GB	256	75.1%

Table 5: Data Distributions for Cassandra Database Nodes

The database was initially created by utilizing the load parameter of YCSB workload A, generating a dataset of approximately 3TB when including replication. The database is then backed up to a separate location over an NFS mount point for quick reload of data between test runs. For each configuration under test, we restored the database from this backup to start every test from a consistent state.

All SSDs were restored to fresh out of box (FoB) state and preconditioned before measurement.

For each cluster configuration (ARM and x86), we tested a variety of thread count settings to determine where each workload reached maximum performance. We then chose a thread count to run longer tests for each workload. For this brief, the best performance, as measured by average latency vs. YCSB throughput, was used.

The tests for each of these workloads ran three times for one hour each. The results of these three test runs were then averaged to obtain the final results.

Test measurement used `dim_STAT` to capture statistics on the server running Cassandra. It captures `IOStat`, `VMStat`, `mpstat`, network load, processor load, and several other statistics. `dim_STAT` was configured to capture statistics on a 10-second interval.

About YCSB Benchmark

The Yahoo Cloud Service Benchmark (YCSB) framework was originally designed to facilitate performance comparisons between various cloud data serving systems for transaction-processing workloads. It includes a set of core workloads that describe common cloud-centric data management operations typically used to manipulate data. Six core workloads are predefined within the benchmark. For testing, five workloads are traditionally used for analyzing cloud-centric applications: Workloads A-D and Workload F. (Workload E is not used in testing as it requires a different database structure than the other five workloads.)

Full descriptions and source code for the YCSB benchmark workloads are available from GitHub (https://github.com/brianfrankcooper/YCSB/wiki/Core-Workloads). Where appropriate, we've used the definitions from GitHub as written.

Table 6 shows the I/O profiles for tested YCSB workloads (additional details are available at YCSB Core Workloads).

Name	Type	Use Case	I/O Profile
Workload A	Update heavy	Session store recording recent actions	50% Read, 50% Write
Workload B	Read mostly	Photo tag reading and adding	95% Read, 5% Write
Workload C	Read-only	User profile cache using profiles constructed elsewhere	100% Read, 0% Write
Workload D	Read latest	User status updates	95% Read, 5% Insert
Workload F	Read-modify-write (R/M/W)	User database where records are read, modified by the user, then written back to the database	50% Read, 50% R/M/W

Table 6: Database Threads for Each YCSB Workload for Each Architecture

micron.com

AMPERE