



MICRON® 6500 ION SSD DELIVERS BREAKOUT NoSQL DATABASE PERFORMANCE¹ WITHOUT BREAKING THE BUDGET

The Micron® 6500 ION SSD is built for applications that demand immense capacity and best-of-breed performance without compromising on application responsiveness.² This technical brief uses the common YCSB benchmark³ workloads to highlight NoSQL database performance and latency differences between the Micron 6500 ION SSD (a high-capacity TLC SSD) and the Solidigm™ D5-P5316 (a QLC SSD). Both tested SSDs are 30.72TB.⁴

How can we compare a TLC SSD to a QLC SSD? Simple. Micron's advancements in NAND technology (including 200+ layers) make it possible for the Micron 6500 ION to be offered at a comparable price to the Solidigm D5-P5316.⁵ When cost is similar, choices are typically influenced by features, endurance, lower power consumption, and real-world results.

Test results show that the Micron 6500 ION routinely demonstrates higher performance (operations per second, noted throughout as ops/sec) and better (lower) 99.99% latency compared to the Solidigm D5-P5316 QLC SSD.

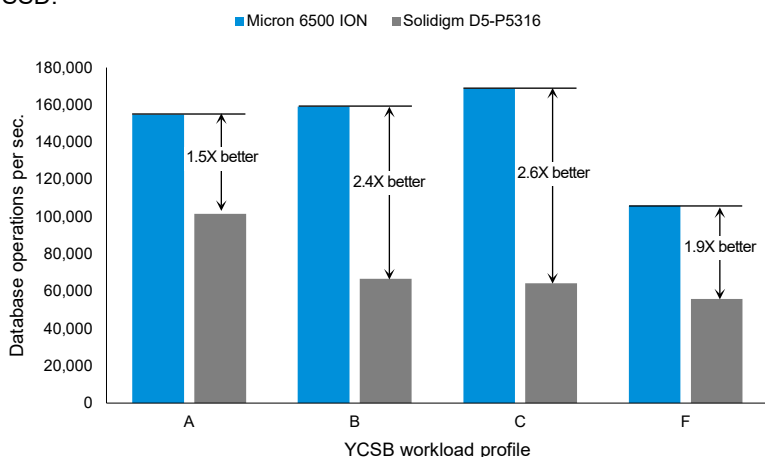


Figure 1: Cassandra maximum performance summary by workload

Fast Facts

The capacity-focused Micron 6500 ION SSD delivers 30.72TB per drive. This high-capacity NVMe™ SSD with 20% lower maximum power consumption enables innovative design opportunities and performance thresholds not found in 30.72TB QLC SSDs, such as the Solidigm D5-P5316.

Yahoo! Cloud Serving Benchmark (YCSB) workloads A–C and F⁶ compare single-node Cassandra results for these two drives. Both are 30.72TB SSDs.

We found that the cluster using the Micron 6500 ION offers:

Workload A: Recording User Sessions

1.5X better performance
7.0X lower latency

Workload B: Adding Metadata

2.4X better performance
4.3X lower latency

Workload C: Reading Profiles

2.6X better performance
3.0X lower latency

Workload F: Recording User Activity

1.9X better performance
9.2X lower latency

1. In this document, the terms performance and database operations per second interchangeably.

2. In this document, application responsive means 99.99% latency, meaning that 99.99% of storage accesses complete within the stated time value. Thus, the terms application responsiveness, responsiveness consistency, and 99.99% latency are used interchangeably in this document.

3. Additional details on YCSB are available from <https://github.com/brianfrankcooper/YCSB>.

4. Unformatted capacity. 1GB = 1 billion bytes, formatted capacity is less. TLC = three data bits per cell. QLC = four data bits per cell.

5. Comparison to the 30.72TB Solidigm D5-P5316 based on public information available at the time of this document's publication. 20% less power based on Micron 6500 ION default 4KB, 100% random, 100% read power = 20 watts. Solidigm P5316 default power consumption for 4KB, 100% random, 100% read = 25 watts. 56% greater power efficiency. 4KB, 100% random, 100% read IOPS per watt comparison.

6. We did not test Workload D (read latest) as its record updates result in a storage profile similar to Workload B (5% write). The major difference between these is a record's age. See <https://github.com/brianfrankcooper/YCSB/blob/master/doc/coreworkloads.html> for additional details. We did not test Workload E as it is not supported in all NoSQL databases.

A Closer Look at Micron 6500 ION Performance in Apache Cassandra™

The following figures show YCSB workload performance (operations per second) and 99.99% latency results for four common NoSQL workloads. Performance is shown on the x-axis (farther to the right is better) and 99.99% latency (in milliseconds) is shown on the y-axis (lower is better). Each point on the figure represents workload performance with the thread count scaled from 8, 16, 32, 64, and 128.

Workload A

Workload A is an update-heavy workload where approximately 50% of all the storage I/O is written and 50% is read. An example of this workload can be seen when user sessions are recorded.




Figure 2: Workload A performance vs. latency

Performance Analysis

Figure 2 shows that Micron 6500 ION performance is consistently higher (farther to the right) than the Solidigm QLC SSD at every tested thread count, with an improvement of 1.5X (154,815 ops/sec versus 101,593 ops/sec) at a thread count of 128.

Latency Analysis

The Micron 6500 ION performance versus latency curve is much flatter than the Solidigm QLC curve. This indicates that the Micron 6500 ION 99.99% read latency remains more consistent with only a slight increase as queue depth scales. In contrast, Solidigm QLC response times show an increase at higher queue depths, indicating that this workload is difficult for the Solidigm QLC SSD.

Workload B

This read-mostly workload comprises approximately 95% read and 5% write storage I/O. An example of this workload includes adding metadata to existing data (tagging) where most of the tags are read (as only a few tags need to be written or rewritten).

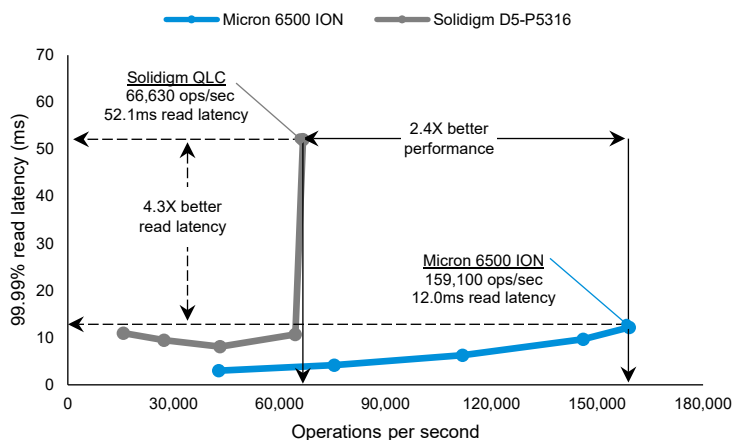


Figure 3: Workload B performance vs. latency

Performance Analysis

Figure 3 shows that Micron 6500 ION performance is consistently higher (farther to the right) than Solidigm QLC at every tested thread count. At a thread count of 128, the improvement observed is 2.4X.

Latency Analysis

The Micron 6500 ION performance versus latency curve gradually increases as operations per second increase, showing limited increases over all thread counts. The Solidigm QLC SSD shows significantly different responsiveness versus ops/sec characteristics. In contrast, the Solidigm QLC SSD shows a dramatic increase in latency at a thread count of 128, while the Micron 6500 ION latency remains lower at every thread count and is more consistent as the workload demand increases.

Workload C

This workload is 100% read (data does not change). An example includes reading immutable data for user authentication or reading a profile cache (for example, when a user or system profile was created elsewhere).

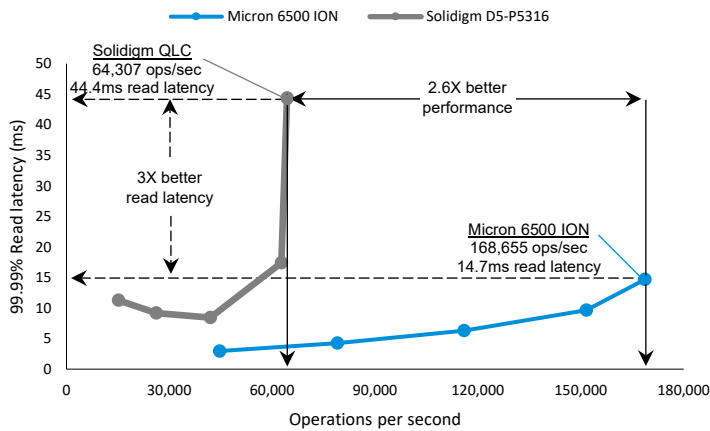


Figure 4: Workload C performance vs. latency

Performance Analysis

Workload C results seen in Figure 4 again show that the Micron 6500 ION SSD performance increases as the thread count increases (left to right) with no latency spikes.

The Solidigm SSD results are different. Its highest thread count performance data point shows little performance improvement over the prior data point suggesting that this SSD will exhibit no additional performance benefits as thread counts grow further.

Latency Analysis

Latency results again show that the Micron 6500 ION SSD latency increases gradually as its performance increases. The Solidigm SSD maximum performance reflects extremely high latency (the line is nearly vertical at this point). This behavior aligns well to the behavior seen in Figures 2 and 3 for this SSD — extremely high latency with little performance improvement.

Workload F

In this workload, the client reads a record, modifies it, and writes back the changes. Application examples include a user database where user records are read and modified by the user and written back. This workload is also used to record user activity.




Figure 5: Workload F performance vs. latency

Performance Analysis

Workload F results in Figure 5 again show that Micron 6500 ION performance increases as thread counts increase (left to right). Its resultant curve again shows no abrupt latency spikes.

The Solidigm QLC SSD results are similar to those observed in previous figures. Its farthest right data point shows a small performance improvement over the prior data point, and its latency spike again suggests that this SSD is nearing its performance limit at a thread count of 128.

Latency Analysis

Latency results again show that Micron 6500 ION latency increases very gradually as its performance increases. The Solidigm QLC SSD performance spike reflects extremely high read latency at its highest observed performance level.

Conclusion

Results show that the 30.72TB Micron 6500 ION SSD consistently demonstrated higher peak performance and better (lower) 99.99% read latency than the comparable 30.72TB Solidigm QLC SSD. The Micron 6500 ION SSD responded faster and more consistently relative to a competing, value-class SSD on the market.

Performance improvements ranged from 1.5X in Workload A (recording use sessions), to a maximum of 2.6X in Workload C (where immutable data like user profiles was used for authentication). Latency improvements range from 3.0X in Workload C, to a maximum of 9.2X (recording user sessions) in Workload F.

These improvements across a broad range of common NoSQL workloads will often have a significant impact on data center performance, making the Micron 6500 ION SSD the preferred high-capacity SSD for Cassandra and other NoSQL database deployments.

Test Configuration

| Hardware Configuration | | | |
|---------------------------|---|-------------------------------|--|
| Database Server | Supermicro® AS -1114S-WN10RT | Load Generation Server | Supermicro AS-1114S-WN10RT |
| CPUs | 1x AMD EPYC™ 74F3 CPU, 3.20Ghz 24 core per socket (total 48 cores in the server) | CPU | 1x AMD EPYC 74F3 CPU, 3.20Ghz 24 core per socket (total 48 cores in the server) |
| Memory | 256GB Micron DDR4-3200 | Memory | 256GB Micron DDR4-3200 |
| Server SSD Storage | Micron 6500 ION configuration: 1x 30.72TB Solidigm D5-P5316 configuration: 1x 30.72TB | Server SSD Storage | N/A |
| Boot Drive | Micron 7300 PRO: 1x 3.84TB M.2 NVMe SSD | Boot Drive | Micron 7300 PRO: 1x 3.84TB M.2 NVMe SSD |
| Network Adapter | NVIDIA® ConnectX®-6 | Network Adapter | NVIDIA ConnectX-6 |
| Operating System | AlmaLinux 8.7 Kernel: 4.18.0-425.3.1.el8.x86_64 Sysctl.conf: vm.swappiness=1 vm.max_map_count=1048575 | Operating System | AlmaLinux 9.0 Kernel: 5.14.0-70.30.1.el9_0.x86_64 |
| Database Version | Apache Cassandra 4.0.7 | YCSB Version | 0.17.0 |

Table 1: Server configuration

The [Yahoo! Cloud Service Benchmark](#) framework was originally designed to facilitate performance comparisons between various cloud data serving systems for transaction-processing workloads.

The core workloads provided by YCSB are listed below.

| Workload | Use Case | IO Type | Ratio |
|----------|-------------------------|-------------------|----------------------------------|
| A | Recording User Sessions | Update heavy | 50% read / 50% write |
| B | Tagging Existing Assets | Read mostly | 95% read / 5% write |
| C | Caching User Profiles | Read only | 100% read / 0% write |
| F | Users Modifying Records | Read-modify-write | 50% read / 50% read-modify-write |

Table 2: Workload details

Workload Modifications

YCSB workloads A, B, C, and F default to a Zipfian distribution for selecting keys. This distribution is meant to work on what is considered hot data, a side effect of which is stressing memory more than storage. We run these workloads using a uniform distribution to place more stress on the storage subsystem.

The default record size in the YCSB workloads is 1000 bytes (10 fields, 100 bytes per field). We used 4096-byte fields (4 fields, 1024 bytes per field), which allows us to use larger data sets due to the record count limit in YCSB. Currently YCSB has issues when using record counts above 2 billion (about 2TB of data with the default record size). The larger record size also allows us to fill larger-capacity SSDs.

micron.com/6500ION