



THE MICRON 7450 SSD DELIVERS IMPRESSIVE PERFORMANCE¹ FOR COMPLEX, MIXED WORKLOADS

The Micron 7450 SSD with NVMe™ provides leading latency quality of service (QoS) to enable responsive and predictable workload performance. It offers the industry's broadest range of PCIe® Gen4 SSD form factors² and enables several storage use cases, including boot, cache and main data storage unlocking performance in a wide variety of data center workloads.

In this brief we focus on the impact that low, consistent SSD latency has on application performance using RocksDB as an example. Our testing centered on 99.99% QoS read latency using the RocksDB database and a mixed (read and write) workload. Our testing compared the Micron 7450 SSD with three open market PCIe Gen4 NVMe SSDs.²

The 7450 SSD showed substantial QoS improvement compared to the first-generation SSD.




Figure 1: Micron 7450 SSD

Fast Facts

The Micron 7450 SSD demonstrated consistently low latency and high performance, eclipsing the low, consistent read latency capability of first-generation, PCIe Gen4 SSD with NVMe.

Why Micron chose RocksDB for this comparison

RocksDB is an embedded storage engine and key-value store database that is used in several backend systems at Facebook and in other data center environments:³

- It replaced another internal storage engine called Centrifuge in the Facebook newsfeed's backend.
- It is one of the many components used in ZippyDB (a distributed key-value store service used by Facebook products that rely on RocksDB).

Due to its performance, RocksDB is well-suited to diverse applications such as:

- A user-facing application that stores the viewing history and state of users of a website
- A spam detection application that needs fast access to big data sets
- A graph-search query that needs to scan a data set in real time
- Metadata management used by BlueStore in RedHat Ceph Storage to store objects to block locations on a disk
- A message queue that supports a high number of inserts and deletes

A 50% read, 50% write workload is extremely demanding

We used this workload because it is analogous to one of the most challenging use cases — recording user sessions.

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

2. We used a Micron 7450 SSD, two current, competitive PCIe Gen4 SSDs with NVMe and a popular first-generation U.2 PCIe Gen4 SSD with NVMe (ca. 2021) for this comparison. All SSDs tested used 7.68TB capacities.

3. Additional information is available here: [GitHub - facebook/rocksdb](https://github.com/facebook/rocksdb) and here: <https://docs.ceph.com/en/latest/rados/configuration/storage-devices/>

4. Retrieved from rocksdb.org/ at the date of this technical brief's publication.

An Introduction to Quality of Service

Many data center applications rely on real time results — they use, process and respond to data that must be delivered quickly, reliably and consistently. This includes active databases (SQL and NoSQL), voice, video and a host of other workloads. For these delay-sensitive applications, data must be delivered on time, without delay or interruption, as this can be detrimental to overall application performance.

Quality of service (QoS) is a measurement that describes latency consistency for a group of deliveries, like SSD read requests or network packets sent. QoS is often expressed in a number of “nines” (99.999% is “five nines”) and a threshold (a time value, such as “50 milliseconds”).

What this means

To understand this better, let's look at a simple example. If our test yields 50 milliseconds at “four-nines” or 99.99% QoS, then out of 10,000 events, the worst event is measured at the 50ms value.

How we measure QoS for SSDs

Testing methodologies for determining QoS values can be a bit more complex than our simple example. Our testing for four-nines calculates QoS in the following manner. For every 1 million data points, we identify the 100 highest latencies, the worst latencies, and report the lowest latency within that set, the best latency in that poor-performing set. Small QoS differences makes large differences at data center scale.

While the difference between the number of nines latency may not seem like much, perspective changes when this is scaled to the data center level. We will use 1 million 4KB random reads per second to show the difference. At this speed, 10 IOs will measure at or above 50ms at five-nines QoS, and 100 IOs will measure at or above 50ms at four-nines — every second.

QoS is also a measure of scalability. Suppose we have a workload that pushes an SSD to 1 million 4KB random read IOPS, this single SSD will have 10 of its read requests at or above the latency threshold for five-nines latency.

If we install 10 of these SSDs into a single server, 100 read requests per second would be at or above 50ms. If we have 20 servers in a rack, 2,000 read requests per second would be at or above 50ms. If we have 20 racks, then 40,000 read requests would be at or above 50ms — every second! For workloads that wait for all outstanding IOs to complete (analytics, for example), these laggard completions — these slowest responders — may cause the analytics process to wait.

QoS Latency Has Changed: A Close Look at Mixed IO Workload Latency

QoS can be a major focus in cloud and other multi-tenant environments. In large deployments, QoS values may frequently be reached, requiring that the SSDs provide fast and consistent response times.

The importance of QoS becomes clear when we analyze how close the Micron 7450 SSD and its competitors are (and see the stark differences between these leading SSDs and the first-generation model).

We will investigate key performance and 99.99% read latency in a 50% read / 50% update, random workload and then analyze:

- Read latency improvements for a set of operations per second target
- Performance improvements for a set service level
- Close performance results among leaders

Micron 7450 SSD: 2.3X lower latency at similar performance




Figure 2: 99.99% latency vs. operations per second – 40,000 ops/sec

Figure 2 shows 99.99% read latency and database operations per second. Database operations per second increase along the x-axis (left to right). Read latency, in milliseconds, is shown increasing along the y-axis. Since lower latency is desirable, lower values on the y-axis indicate lower latency (better results).

To understand latency improvements, we start by fixing a performance value and then comparing latencies.

Fixing a performance value: Start by locating a specific performance value on the x-axis. For our example, we chose 40,000 operations per second. Similar analysis can be done for other performance values as needed.

Micron 7450 SSD: When we trace the vertical 40,000 operations per second line to its intersection with Micron 7450 SSD performance (black) line and then look left, we see that its 99.99% read latency is ~10ms (value indicated on the y-axis).

First-gen SSD: Likewise, when we trace the same 40,000 operations per second vertical line upward to the first-gen SSD performance (grey) line and again look left, we see that the first-gen SSD latency is ~23ms.

Other SSDs: The Micron 7450 has a similar, in some cases better, latency than the leading competitors in this space (competitors A and B).

Results: Comparing these latency values shows that the Micron 7450 SSD offers a 2.3X 99.99% latency improvement (reduction) at 40,000 operations per second compared to the first-gen SSD.

Micron 7450 SSD: 95% operations per second improvement at similar latency




Figure 3: 99.99% latency vs. operations per second – 15 milliseconds

To understand performance improvements, we start by fixing a latency value, a threshold, and then compare performance.

Fixing a latency value: In this comparison, start by locating a desired service level (y-axis value). Here we chose 15ms 99.99% read service level. Similar analysis can be done for other service levels by moving the horizontal line to the desired y-axis value.

Setting the threshold to 15ms enables us to compare each SSD's operations per second at this threshold.

Micron 7450 SSD: When we follow the 15ms service level to the right until it intersects the Micron 7450 performance line and then trace down, we see that the Micron 7450 SSD reaches 45,000 operations per second at 15ms.

First-gen SSD: If we repeat this process for the first-gen SSD, we see that it reaches 23,000 operations per second at the same latency threshold.

Other SSDs: The Micron 7450 has a similar, in some cases better, latency than competitors A and B.

Results: Comparing these performance values at a 15ms service level shows that the Micron 7450 SSD reaches 95% higher performance than the first-gen SSD.

Latency leaders show similar results




Figure 4: 99.99% latency vs. operations per second – leaders

Finally, when we compare the 99.99% latency values for the Micron 7450 to only Competitor A and Competitor B (Figure 4 does not show the first-gen SSD), we see that the latency and operations-per-second values among these three SSDs are tightly grouped.

This indicates that these three SSDs exhibit similar read latency values at multiple performance levels.

Conclusion

Micron chose the RocksDB database because it is a high-performance key-value store. It is capable of supporting a broad variety of cloud and hyperscale solutions.

RocksDB is optimized for flash so it can fully exploit SSDs with NVMe and is capable of extreme storage throughput and database operations per second.⁵

Latency values show a major improvement over a first-generation SSD with NVMe. These differences reflect Micron's continued efforts to improve all aspects of our data center SSDs, and the small differences seen in this test among the leaders.

The Micron 7450 SSD with NVMe enables not only RocksDB, but a wide variety of workloads for flexible deployment in cloud, OEM, data center and system integrator designs. It is the right SSD where predictable and reliable responsiveness is needed for data center workloads.

5. As noted by RocksDB: <https://rocksdb.org/>

How We Tested

Hardware Configuration	
Server	Dell PowerEdge R7525
CPU	AMD EPYC 7713 64-Core Processor
Memory	512GB Micron DDR4-3200
Server Storage	1x Micron 7450 PRO 7.68TB SSD 1x Competitor A 7.68TB NVMe SSD 1x Competitor B 7.68TB NVMe SSD 1x 7.68TB First Generation NVMe SSD
Boot Drive	Micron 960GB M.2 NVMe SSD
YCSB Version	YCSB 0.17.0 (HSE fork — https://github.com/hse-project/hse-ycsb)
RocksDB Version	6.22
OS	CentOS Linux 8
Kernel	4.18.0-305.12.1.el8_4.x86_64

Table 1: RocksDB Server Configuration

System/Software Configuration	
File System	XFS
Mount Options	noatime, nodiratime, norelatime
Mount Point	/var/lib/rocksdb

Table 2: Device Configuration

The testing methodology was designed to show the performance of a single device to store and access a RocksDB database. Single-device performance differences help characterize their suitability as a storage building block for a broad range of node and cluster configurations. Actual results may vary based on configuration specifics.

The database was initially created by using the YCSB load parameter, which generated a dataset of approximately 1TB. After completing the load, the test sequence, in the order listed, was executed using the following parameters (Table 3). We used a uniform distribution to cause more stress on the storage subsystem.

Parameters	Value	Description
Number of threads (#Threads)	4, 8, 16, 32	Number of threads to generate
Fieldcount	4	4KB record size
Fieldlength	1,024	
Recordcount	250 million	Number of records in the database
Operationcount	200 million	Dataset size within database
Distribution	Uniform	Distribution of record accesses in database for each workload

Table 3: Testing Parameters

micron.com/7450

© 2022 Micron Technology, Inc. All rights reserved. All information herein is provided on an "AS IS" basis without warranties of any kind. Products are warranted only to meet Micron's production data sheet specifications. Products, programs and specifications are subject to change without notice. Micron Technology, Inc. is not responsible for omissions or errors in typography or photography. Micron, the Micron logo and all other Micron trademarks are the property of Micron Technology, Inc. All other trademarks are the property of their respective owners.
Rev. A 03/2022 CCM004-676576390-11598