



THE MICRON® 9400 NVMe™ SSD ENABLES CLASS-LEADING AEROSPIKE DATABASE PERFORMANCE¹

Turning data into insights as rapidly as possible is imperative in today's business environment. Architectures that provide high throughput while maintaining application responsiveness are essential to support these demands.

Aerospike Database is an open source NoSQL database that is optimized for flash storage. It is one of the fastest NoSQL databases underpinning time-critical web applications like fraud detection, recommendation engines, real-time payment processing, and stock trading.²

This technical brief compares the Micron 9400 NVMe SSD to three leading performance-focused NVMe SSDs³ using Aerospike Database and four common workloads similar to the Yahoo! Cloud Serving Benchmark (YCSB).⁴ It shows workload performance and responsiveness as demand is increased (increasing thread count). Aside from different SSDs, each configuration was the same.

Test results show that the Micron 9400 SSD demonstrates higher peak performance and superior application responsiveness across all tested workloads.

Test results are organized by workload, with a brief workload description and one example use case (these are based on the YCSB workload descriptions from GitHub).



Micron 9400 SSD

Fast Facts

The Micron 9400 SSD delivers superior Aerospike Database performance and rapid, consistent results compared to the leading competitors' performance-focused NVMe SSDs.

The highlights below illustrate performance benefits of the Micron 9400 NVMe SSD in Aerospike deployments. This technical brief provides additional performance details and in-depth comparisons versus other PCIe® Gen4 SSDs.

Workload A: Recording User Sessions

2X Higher peak performance

Workload B: Tagging Existing Assets

1.9X Higher peak performance

Workload C: Caching User Profiles

1.4X Higher peak performance

Workload F: Users Modifying Records

2.1X Higher peak performance

micron.com/9400

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

2. Additional Aerospike Database use cases are available from <https://aerospike.com>

3. Competitors as noted in SSD Insights Q4/22 (analyst firm Forward Insights).

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

A Closer Look at Micron 9400 SSD Aerospike Performance

The Micron 9400 SSD is built for applications that demand high throughput while maintaining responsiveness. This NVMe SSD is well-suited for broadly deployed, mixed read-write NoSQL database workloads that thrive on performance.

The following figures show YCSB workload Aerospike Database operations per second on the x-axis (farther to the right is better) and application responsiveness on the y-axis (lower is better). Each point on the figure represents workload performance, scaled until adding additional threads yielded negligible performance improvement. Performance values are stated as operations per second and are abbreviated as “OPS/s”.

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.

Workload A: Performance vs Application Responsiveness

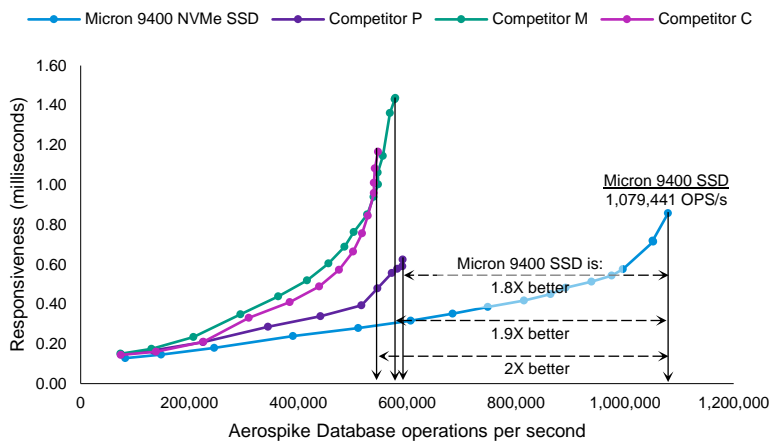


Figure 1

Performance Analysis: Figure 1 shows that the Micron 9400 SSD performance is consistently higher (farther to the right) than all three competitors at every thread count.

The Micron 9400 SSD performance response curve is much flatter than the competitors' curves. This indicates that the Micron 9400 SSD application responsiveness remains extremely consistent as demand increases.

Tested SSD	Max. performance (OPS/s)	Micron 9400 SSD Advantage
Micron 9400 SSD	1,079,441	–
Competitor P	592,150	1.8X
Competitor M	577,361	1.9X
Competitor C	546,483	2X

Table 1: Workload A performance differences

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 (only a few tags are written or rewritten).

Workload B: Performance vs Application Responsiveness

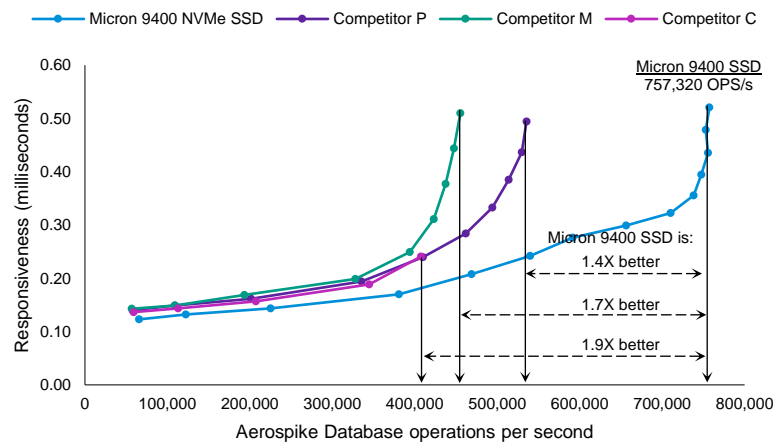


Figure 2

Performance Analysis: Figure 2 shows that the Micron 9400 SSD performance is consistently higher than all three competitors.

The Micron 9400 SSD reached 757,320 peak operations per second compared to just 535,471 operations per second for the Competitor P, 455,254 operations per second for Competitor M, and 407,534 operations per second for Competitor C.

Tested SSD	Max. performance (OPS/s)	Micron 9400 SSD advantage
Micron 9400 SSD	757,320	–
Competitor P	535,471	1.4X
Competitor M	455,254	1.7X
Competitor C	407,534	1.9X

Table 2: Workload B performance differences

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 (when a user or system profile was created elsewhere).

Workload C: Performance vs Application Responsiveness

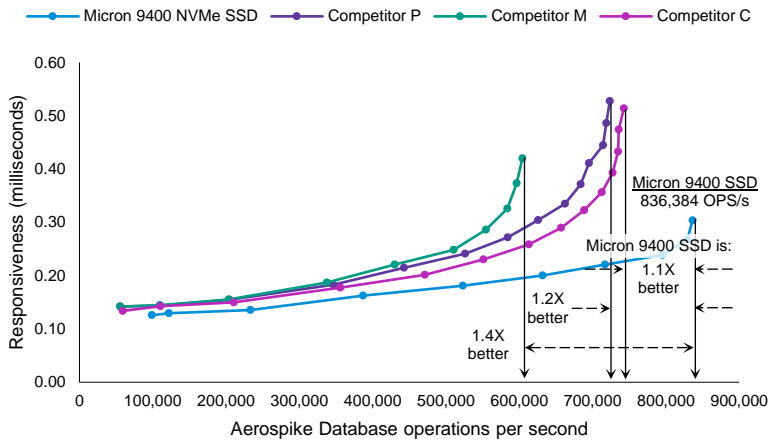


Figure 3

Performance Analysis: Figure 3 shows that the Micron 9400 SSD performance is consistently and significantly higher than all three competitors (summarized in Table 3).

The Micron 9400 SSD is also 67% faster than its nearest competitor at the lowest thread count (farthest left point). This difference shows that the Micron 9400 SSD offers superior performance even when the demand is extremely low.

Tested SSD	Max. performance (OPS/s)	Micron 9400 SSD advantage
Micron 9400 SSD	836,384	-
Competitor P	723,533	1.2X
Competitor M	604,297	1.4X
Competitor C	742,696	1.1X

Table 3: Workload C performance differences

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.

Workload F: Performance vs Application Responsiveness

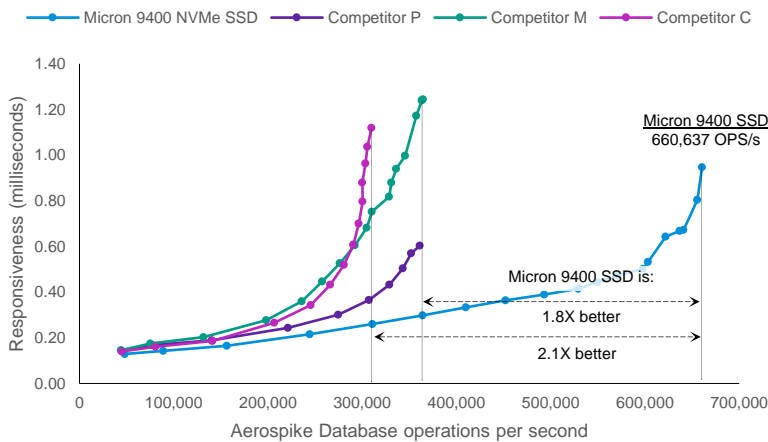


Figure 4

Performance Analysis: Figure 4 shows that the Micron 9400 SSD performance is consistently higher than all three competitors (summarized in Table 4).

The performance values for Competitors P and M are sufficiently close such that their performance indicators in Figure 4 overlap as indicated by an asterisk. This overlap is due to the resolution in Figure 4. Actual performance values are shown in Table 4.

Tested SSD	Max. performance (OPS/s)	Micron 9400 SSD advantage
Micron 9400 SSD	660,637	-
Competitor P	361,210	1.8X
Competitor M	364,704	1.8X
Competitor C	309,888	2.1X

Table 4: Workload F performance differences

Conclusion

The Micron 9400 NVMe SSD consistently demonstrated higher peak performance in several common NoSQL use cases while also providing the best workload responsiveness. Performance improvements ranged from a minimum of 1.1X in Workload C (where immutable data for user authentication or reading a profile cache is read) to a maximum of 2.1X in Workload F (a user database where user records are read, modified, and written back).

These improvements across a broad range of common NoSQL use cases will often have a real-world impact in the data center, making the Micron 9400 NVMe SSD the clearly preferred storage building block in Aerospike Database deployments.

How We Tested

Hardware Configuration	
Aerospike Database Server	Dell PowerEdge R750
CPUs	2x Intel Xeon Platinum 8358 CPUs @2.60Ghz 32 cores per socket, Hyper Threading enabled (total 128 cores in the server) High performance mode enabled Intel VT disabled
Memory	1TB Micron DDR4-3200
Server Storage	Micron 9400 Configuration: 1x Micron 9400 PRO 7.68TB NVMe SSD Competitors P, M, and C: 7.68TB performance-focused NVMe SSD All NVMe SSDs configured with 16 namespaces of 480GB each
Boot, Applications Drive	Micron 5200 PRO 480GB SATA SSD
Network Adapter	NVIDIA ConnectX-6 – 100Gbps
Operating System	AlmaLinux 8.6 Kernel: 4.18.0-372.19.1.el8_6.x86_64 Sysctl.conf: vm.swappiness=0 vm.min_free_kbytes=1310720 Disable transparent hugepages
Aerospike Database version	Enterprise Edition build 6.1.0.1
Load Generator	Supermicro AS-1114S-WN10RT
CPU	1x AMD EPYC 7F72 24 Cores, SMT enabled (total 48 cores)
Memory	256GB Micron DDR4-3200
Boot Drive	Micron 5200 PRO 480GB SATA SSD
Network Adapter	NVIDIA ConnectX-6 – 100Gbps
Operating System	Ubuntu 20.04 LTS Kernel: 5.4.0-99-generic

Table 5 Server Configuration

Each of the 16 namespaces is presented to Aerospike Database to increase parallelism and achieve better performance. Thread count was scaled (8, 16, 32, 64, etc.) until further increases resulted in negligible performance increase.

When starting a test, each SSD is securely erased and preconditioned. The namespaces or partitions are recreated on the device to match the testing configuration. For each test, the following process is followed:

- Stop database service
- Clear caches
- Start database service
- Start data capture
- Execute test (20 minutes)
- Stop data captures
- Collect data from data captures

We evaluated each thread count listed on Workloads A, B and F. Workload C is run separately to prevent Aerospike defragmentation from previous workloads, which may affect the results of the 100% read test.

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 6: 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 in order 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/9400