Open Source Storage Performance With Micron HSE Compared to RocksDB

HSE realizes improved throughput, reduced latency and improved write- and read- amplification when compared to RocksDB

Introducing Micron® Heterogenous-Memory Storage Engine (HSE), one of the first heterogeneous-memory storage engines designed from the ground up to accelerate Linux® workloads using flash-based storage and storage class memory. Legacy storage engines from the hard disk drive era were not optimized for modern technologies like solid-state drives (SSDs) and storage class memory (SCM).

Micron's unique open source HSE was designed to maximize the capabilities of these new technologies by intelligently and seamlessly managing data among multiple storage classes. The result is significantly improved performance, increased drive endurance and decreased latency — even under the crushing burden of massive-scale deployments.

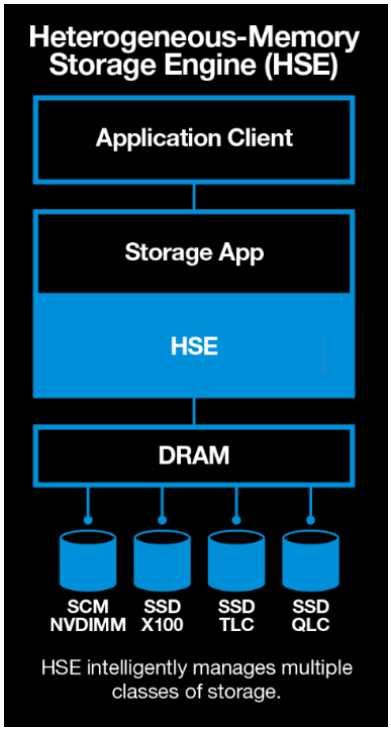

Figure 1: HSE Block Diagram

HSE Features

Unleash SSDs and storage-class memory

Many storage engines were written for hard drives and never optimized for SSDs, flash-based technologies or SCM. HSE offers a way to maximize those legacy applications.

Drive storage applications to massive scale

Many applications struggle to be open enough or scalable enough for today's largest workloads. HSE is designed with massive-scale databases in mind, such as billions of key counts, terabytes of data and thousands of concurrent operations. These capabilities make HSE ideal for even the largest data challenges.

Build on leading-edge storage and memory technologies

Few enterprise applications optimize across the complex stack of different nonvolatile memory hardware. HSE transparently exploits multiple classes of media — from QLC 3D NAND flash to 3D XPoint™ technology.

HSE was also designed for next-generation NVMe™, which is becoming a data center standard for fast, low-latency SSD connectivity. Designed to accommodate emerging technologies — such as NVMe sets, zoned namespaces, LBA streams and open channel — HSE keeps our customers at the forefront of storage technologies.

Tested Configurations

Tests were performed using standard, x86 dual-socket servers. Table 1 shows the tested hardware configurations and types and quantities of storage devices used. Table 2 shows software and version numbers used. Table 3 shows the Yahoo! Cloud Serving Benchmark (YCSB) benchmark configuration used for testing.[1]

Table 1: Hardware Details

Server Platforms	
Server Platform	Intel® based (dual-socket)
Processor	2x Intel E5-2690 v4
Memory	256GB DDR4 (16x16GB)
SSDs	4x Micron 9300 SSDs with NVMe
Storage Configuration	LVM striped logical volume

Table 2: Software Versions Used

Software Details	
Operating System	RHEL 8.1
HSE version	1.7.0
RocksDB version	6.6.4

Table 3: Benchmark[1]

YCSB Benchmark Configuration	
Dataset	2TB (2-billion 1,000-byte records)
Client threads	96
Operations	2 billion per workload

Results

The test results are organized by testing phase, with each YCSB workload result documented. The figures show database operations per second (higher is better) and tail latency (lower is better). These figures show that HSE enables far greater throughput with substantially lower tail (99.9%) latency for these YCSB workloads.

A brief description of each phase accompanies each workload's test results. Additional details on example YCSB workloads and their descriptions are available from GitHub.

Load Phase

YCSB starts by populating the dataset (database), a 100% insert workload. For this workload, HSE delivered more than 8x the throughput compared to RocksDB — reducing total workload duration by 87% (Table 4).

Table 4: Load Phase Results

Load Metric	HSE	RocksDB
Duration (minutes)	34	271
Inserts / second	986,009	122,899
Insert 99.9% latency (ms)	1.6	19.5

Run Phase

Run phase throughputs for HSE and RocksDB are shown in Figure 2. For these YCSB workloads, HSE delivered up to nearly 6x more throughput and up to 91% lower write tail latency than RocksDB. Figures 2 and 3 show these results.

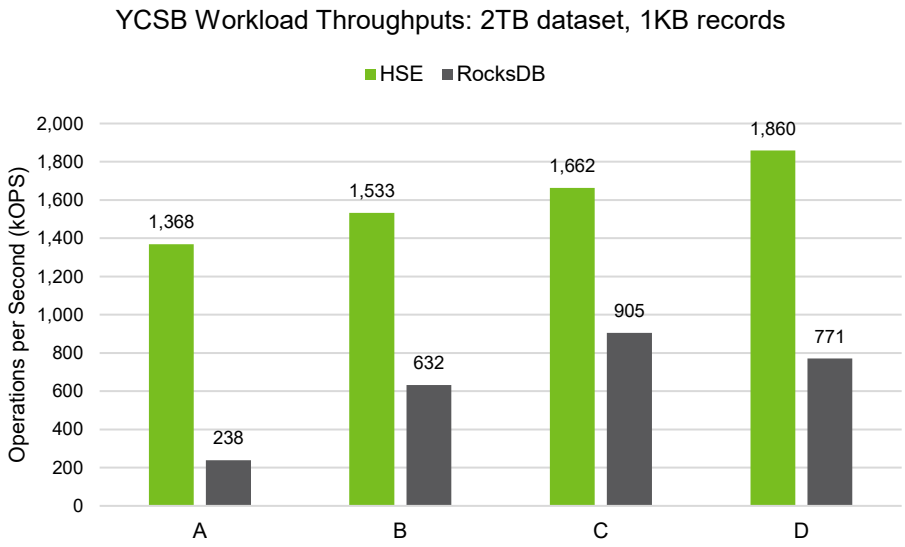

Figure 2: Run Phase Throughput

In delivering high throughput, HSE also demonstrated good 99.9% tail latency, as shown in Figure 3. For these YCSB workloads, HSE reduced read tail latency up to 44% and write (insert or update) tail latency by as much as 91%.

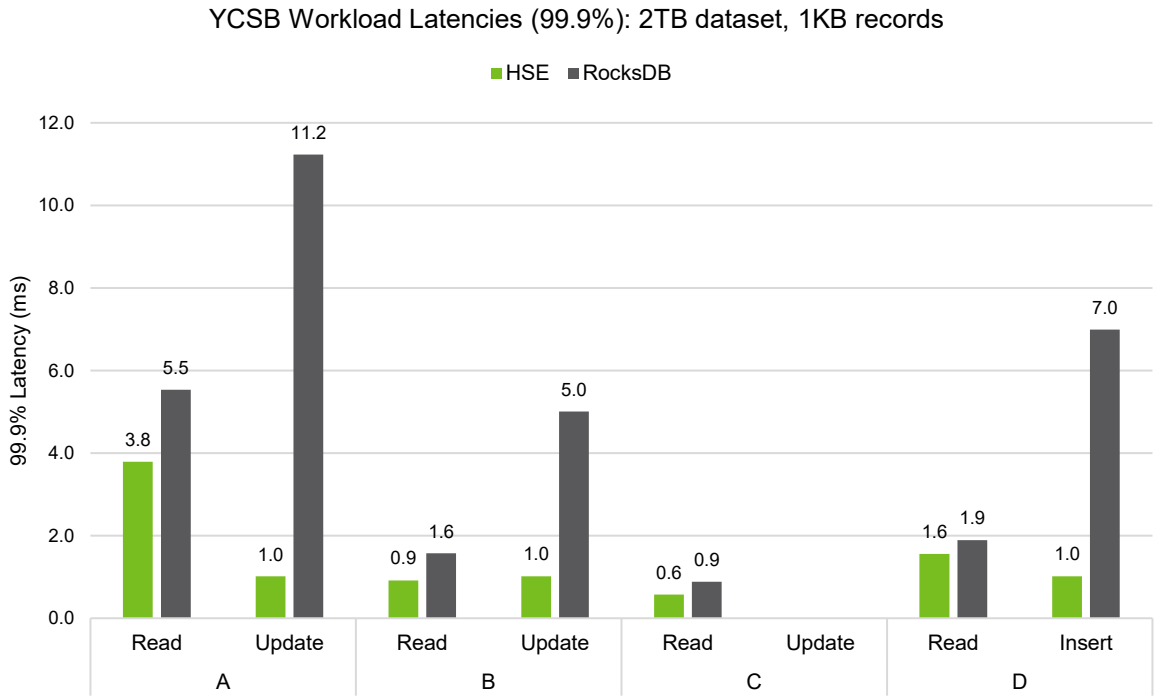

Figure 3: Run Phase Tail (99.9%) Latency

Data Written	Figure 4 shows the total bytes written to or read from the drives during workload execution. For these YCSB workloads, HSE reduced bytes read up to 64% and bytes written by as much as 84%. By reducing write-amplification, SSDs will see a commensurate endurance improvement.

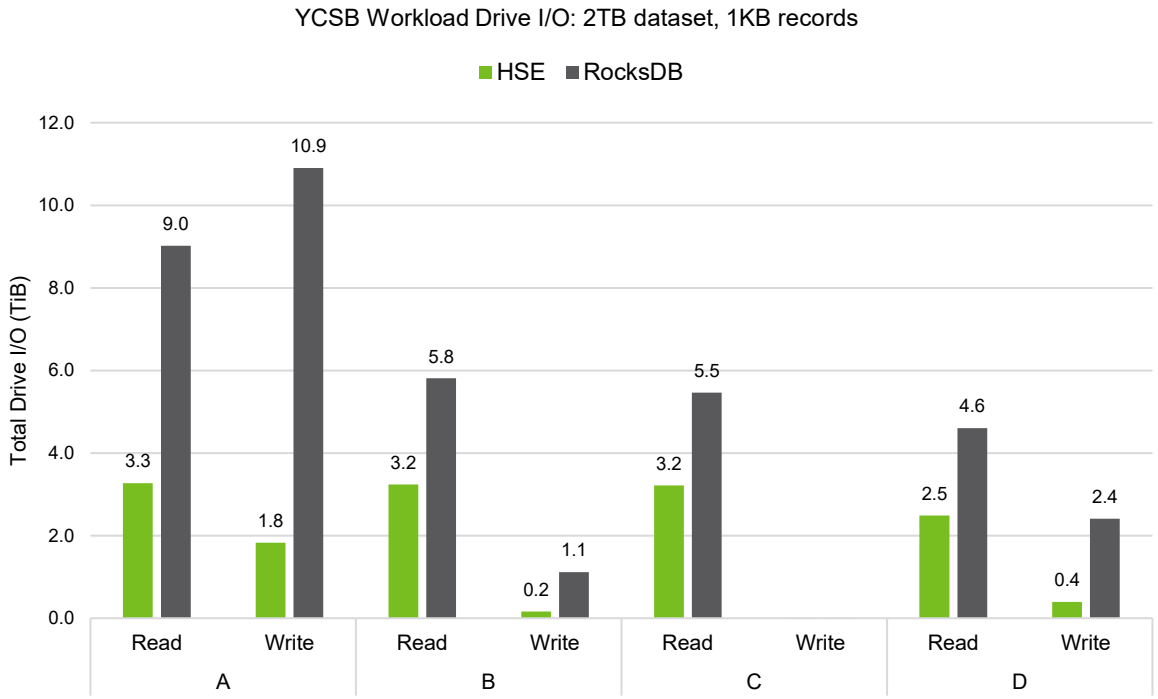

Figure 4: I/O Reduction

The Bottom Line

With this new open source heterogeneous-memory storage engine (HSE), data-focused software can now more effectively use modern storage performance in an open source world. Developers can build on and deploy HSE to accelerate massive-scale workloads. Designed to optimize across the complex stack of different nonvolatile storage and memory hardware, HSE is proven to dramatically increase throughput, reduce latency and maximize I/O efficiency for extended SSD endurance. HSE is designed with massive-scale databases in mind, those with billions of key counts. terabytes of data and thousands of concurrent operations.

These capabilities make HSE ideal for even the largest data challenges.

How to Get Started

Learn more at www.micron.com/hse

Get started developing at www.github.com/hse-project

1. Different versions of software, datasets, threads and other configuration parameters may show different results. Configurations described were used for all tests.

Why Micron for Storage Innovation

Why align with Micron for open source software? We're meeting our customers' needs today and tomorrow with the industry's broadest, most innovative memory and storage portfolio. Micron is one of the few silicon-to-system vendors in the world. Through our globally distributed network of supply chain, operations sales offices and customer labs, we deliver comprehensive collaboration, support and quality throughout a product's lifecycle.

Our global reach also extends farther than our physical locations. Micron has a history of building trusted relationships with the world's most influential technology ecosystems**,** including industry consortiums, leading research labs and academic institutions.

We keep innovating to make it easy for our customers to get the best performance from their systems. Micron's vertical integration allows us to build solutions by now adding software to our systems expertise.

micron.com

Rev. B 9/2020, CCM004-676576390-11449