

Open-Source OLTP Never Looked So Good

NVMe™ SSDs Deliver Amazing Results With MySQL® and OLTP

Overview

Open-source relational database management systems (RDBMS) like MySQL® are broadly deployed¹ and very flexible with extensive community support. MySQL powers some of the largest high-growth websites while remaining free.²

MySQL's licensing model shifts the way we think about storage. Compared to traditional closed-source RDBMSs with heavy licensing costs, MySQL can give us a different focus. In this brief, we discuss the configuration needed to achieve MySQL system responsiveness (average time for the database to respond to an inquiry), response consistency and power efficiency (database operations per watt of system-level power consumed). The goal is to generate more performance³ and more value for MySQL.

We also explore supporting more users and throughput when the system is heavily loaded⁴, which is an important consideration for I/O-intensive applications using scripting languages like PHP with a database backend — a common use for MySQL.

In our testing, we used standardized OLTP performance metrics and a data set that exceeds available system memory (to test storage system I/O) to compare performance, average response time, response time consistency and power efficiency for three storage configurations used with MySQL Community Edition 5.7:

- **NVMe 1:** 2x 3.8TB SSDs, RAID 1 (LVM)
- **NVMe 2:** 4x 3.8TB SSDs, RAID 10 (LVM)
- **Legacy:** 8x 300GB 15K RPM HDDs, hardware RAID 10 (baseline configuration for comparison)

We found that NVM Express® (NVMe™) SSD configurations demonstrated far higher performance, with quicker response times and more consistency, bringing more value to OLTP workloads on MySQL compared to our HDD configuration.

Fast Facts

- More orders, more fulfillments, more value to your bottom line with NVMe
- Data can't wait — NVMe enables amazing results for MySQL
- Release your data's potential like never before:
 - 2x NVMe SSDs: >103,000 new orders per minute
 - 8x 15K HDDs: ~1,800 new orders per minute
- Applications are highly responsive and more consistent with NVMe



1. Deployment data from <https://www.mysql.com/why-mysql/>
 2. MySQL GPL licensing: <https://www.mysql.com/about/legal/>
 3. Throughout this paper "performance" means new orders per minute (NOPM)
 4. Heavily loaded is defined as maximum user count without reaching one or more stop conditions (see the How We Tested section)

NVMe Provides Much Higher Performance

To test platform performance, we measured new orders per minute (NOPM) for all three configurations across all supported user counts and averaged the results. (See the How We Tested section for details on the range of user counts for each configuration.) This gives a good view of each configuration's performance.

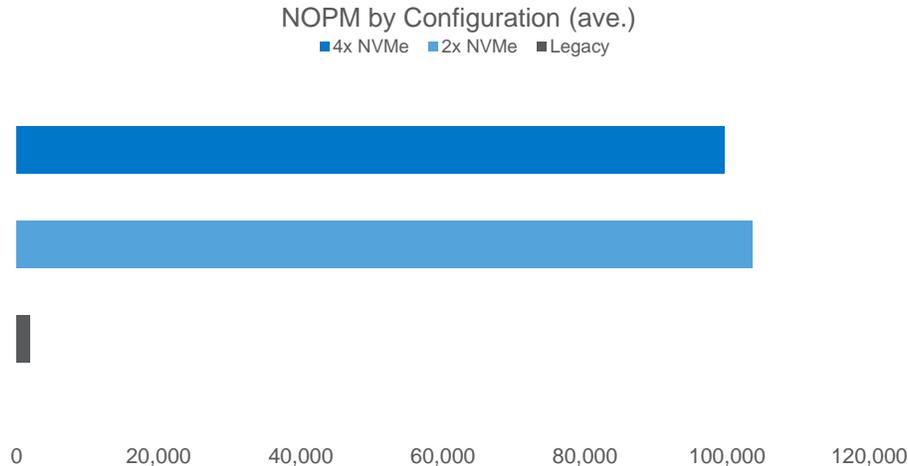


Figure 1: Average NOPM

As shown in Figure 1, both NVMe configurations showed very high performance compared to the legacy HDD configuration. The 2x NVMe configuration supported slightly higher performance (just over 100,000 NOPM) than the 4x NVMe configuration (99,000 NOPM). This was somewhat unexpected and may be due to MySQL performance not scaling well for very high-performance storage. As a result, a 4x NVMe configuration is a good option for high-performance, high-capacity platforms.



Tip: The 4x NVMe configuration is a great choice for high-performance, high-capacity MySQL platforms. 4x NVMe doubles storage capacity with excellent performance.

NVMe Responds Faster, More Consistently

For many OLTP deployments, performance is the design driver, translating to more revenue and greater benefit. However, this is not always the case.

For some, database responsiveness (mean response time) drives configuration choice. For still others, database response time consistency (90% transaction response time is a good measure of latency consistency) may be a primary configuration driver (this may be common in high performance OLTP systems that can't tolerate an occasional outlier during which the database response takes much longer than its average).

As shown in Figures 2a and 2b, we compared mean and 90% transaction response time for all three configurations. We found that the NVMe configurations responded very quickly and with excellent consistency.

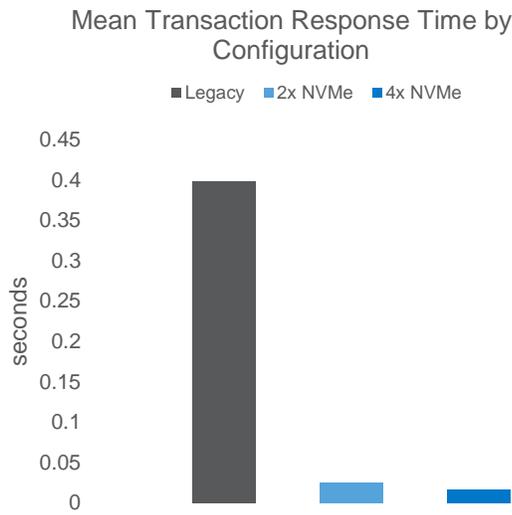


Figure 2a: Mean response time

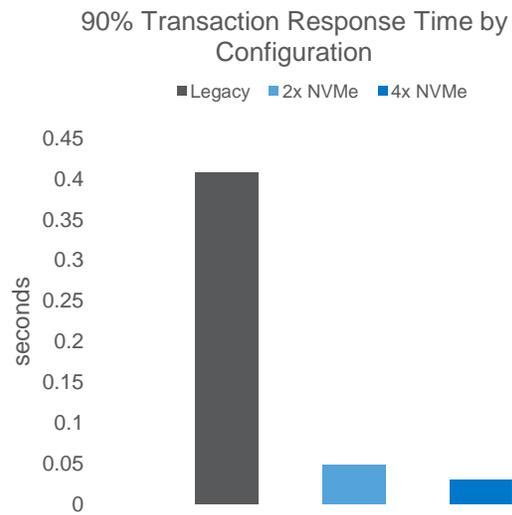


Figure 2b: 90th percentile response time

Figures 2a and 2b show that the NVMe configurations have dramatically lower mean response times and are more consistent than the legacy reference configuration (this result is not surprising). The legacy configuration shows high mean transaction response time, indicating its overall response times are high. It also shows a high 90% latency response time, indicating poor response time consistency.

Configuration	Mean Response Time Reduction (vs. Baseline)
4x NVMe	96% lower
2x NVMe	94% lower
Legacy	baseline

Table 1: Mean Response Times

Configuration	90% Response Time Improvement (vs. Baseline)
4x NVMe	93% lower
2x NVMe	88% lower
Legacy	baseline

Table 2: 90th Percentile Response Times



Tip: For high-capacity deployments that need lower, more consistent responses, consider a 4x NVMe configuration.

While both NVMe configurations showed low and consistent response times, with double-digit improvements over the baseline configuration, the 4x NVMe configuration showed lower average response time with more consistency than the 2x configuration.

NVMe Is More Power-Efficient

In densely packed, high-I/O deployments, power efficiency may be more important than raw performance or than faster, more consistent response time. To gauge power efficiency, we divided each configuration's performance (in NOPM) by the system power draw at that load to calculate performance per watt, as shown in Figure 3. Note that we used the system power draw (in watts) at the NOPM values shown in Figure 1 for this calculation.

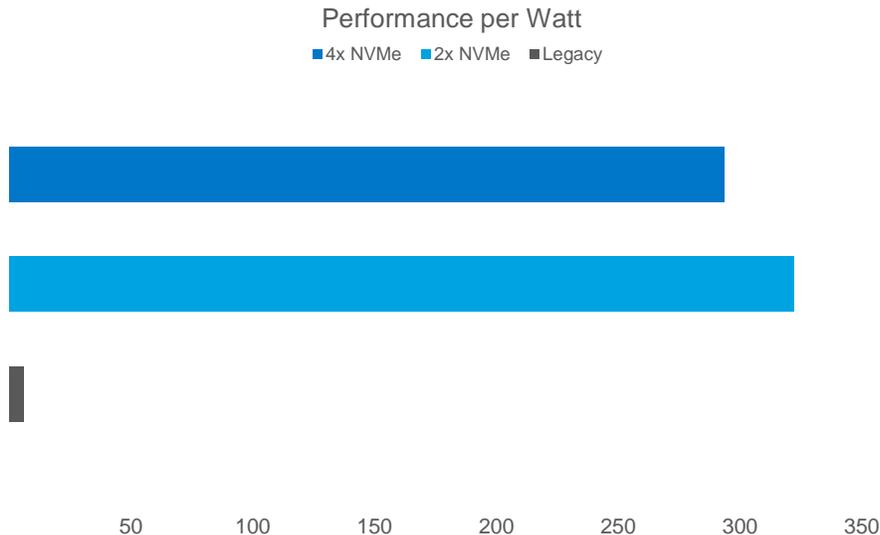


Figure 3: NOPM per watt

The 2x NVMe configuration showed the highest NOPM per watt, indicating greater power efficiency compared to the 4x NVMe configuration. This is because the 4x NVMe configuration has slightly lower NOPM and slightly higher power draw compared to the 2x NVMe configuration.

The legacy configuration NOPM is very low while its power draw is close to both NVMe configurations. This drives its calculated power efficiency much lower than both NVMe configurations.



Tip: The 2x NVMe supports slightly higher performance than the 4x NVMe configuration and draws slightly less power.

The Bottom Line

SSDs with NVMe bring amazing results to MySQL and OLTP. We used a standardized OLTP performance metric and a data set that exceeded available system memory to compare two NVMe configurations to one HDD-based legacy baseline configuration.

We've shown how SSDs with NVMe can be a great fit for MySQL. They show very high performance with fast and consistent responses and excellent power efficiency per operation. A legacy HDD configuration simply could not keep up, even with many more drives.

Micron offers a broad portfolio of [SSDs with NVMe](#). You can estimate the benefits of an all-flash data center using our [Move2SSD TCO Tool](#).

How We Tested

To ensure a fair assessment of each configuration's capabilities, we took a configuration-specific approach. We wanted to measure each configuration's performance over a broad range of user counts, as opposed to an arbitrary fixed number of users.

We started testing with 32 users and increased the user count until the platform was saturated. Prior to testing, we established stop conditions, as shown in Table 1. During testing, we increased the user count until reaching a stop condition.

Limit	Stop Condition
CPU utilization	80%
90 th percentile average transaction response time	See Table 2
Response time	Sharp increase (a "kink")
NOPM plateau	When NOPM fails to increase with user count

Table 1: Stop conditions^{5,6}

We set the 90th percentile transaction response time to the values in Table 2, which reflect common tolerance limits. When we reached a stop condition, we described the user count as the platform "saturation point."

Transaction	90 th Percentile Response Time
New order	5 seconds
Payment	5 seconds
Order status	5 seconds
Delivery	5 seconds
Stock level	20 seconds

Table 2: Threshold limits

5. We set the stop condition for CPU utilization at 80%. Many IT organizations plan for platform upgrade when CPU utilization reaches 50% and implement that plan when it reaches 80%.

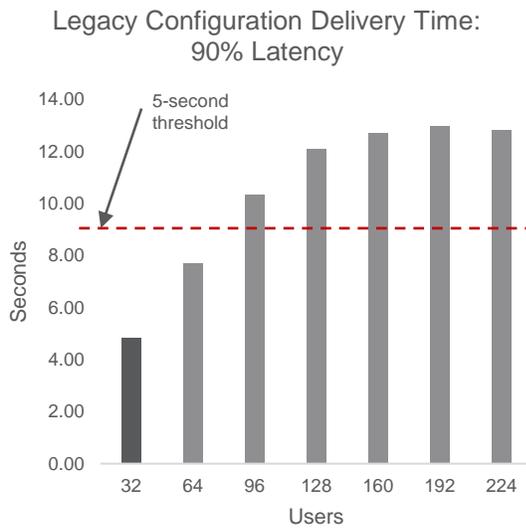
6. We sized the data set to ensure it was large enough to exercise storage I/O (data set size about 2X the memory size) but did not occupy more than 80% storage capacity.

Determining User Count by Configuration

This section shows the test condition(s) that established the maximum user count for each configuration.

Legacy Configuration Stop Condition: Delivery Latency

Figure 4 shows the legacy configuration’s average latency for delivery and the 90% delivery latency 5-second threshold (dashed line). When the user count exceeds 32, the 90% latency exceeded 5 seconds (the preset stop condition). As a result, we used 32 users for all legacy configuration data. Note, Figure 4 shows results beyond 32 users to provide additional detail. This data is shaded in gray to indicate that these values were not used in our comparisons).



2x NVMe Configuration Stop Condition: NOPM Plateau

Figure 5 shows the 2x NVMe configuration performance by user count. Performance plateaus at 96 users. As a result, we used 96 users for all 2x NVMe configuration data. Note, Figure 5 shows results beyond 96 users to provide additional detail. This data is shaded in gray to indicate that these values were not used in our comparisons.

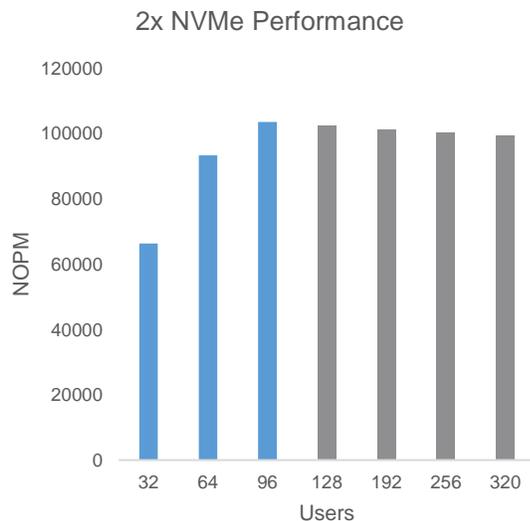


Figure 5: 2x NVMe stop condition

4x NVMe Configuration Stop Condition: NOPM Plateau

Figure 6 shows the 4x NVMe configuration performance by user count. Performance plateaus at 64 users. As a result, we used 64 users for all 4x NVMe configuration data. Note, Figure 6 shows results beyond 64 users to provide additional detail. This data is shaded in gray to indicate that these values were not used in our comparisons.

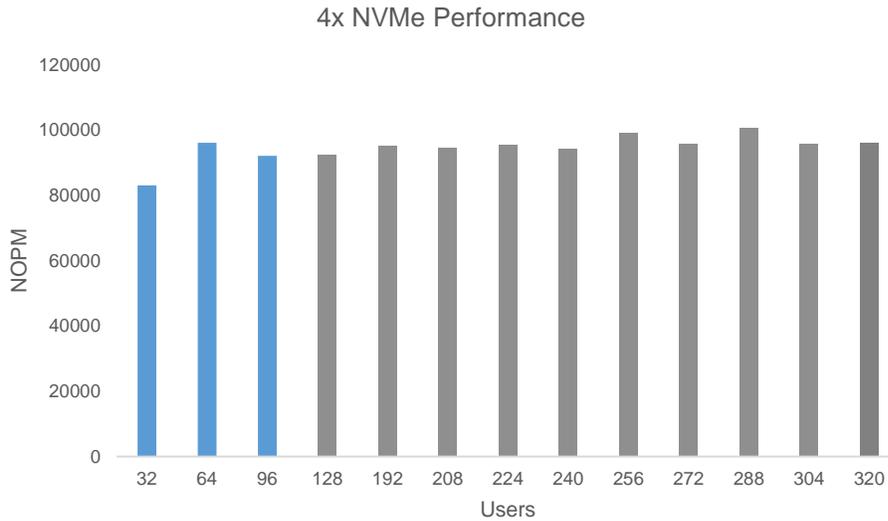


Figure 6: 4x NVMe stop condition

Table 3 shows the hardware platform configuration used for all testing. We used MySQL Community Edition 5.7. on CentOS 7.4 for all tests.

Component	Description
Server	2U, 2-socket (Intel based)
CPUs	Intel Xeon E5-2690 v4 (x2)
NVMe SSDs	Micron 9200 PRO 3.8TB (x2, x4)
NVMe Configuration	LVM: 2x NVMe, RAID 1; 4x NVMe, RAID 10
HDDs	15K RPM, 300GB (x16)
HDDs Configuration	RAID 10 (hardware)

Table 3: Hardware configuration

micron.com

This technical brief is published by Micron and has not been authorized, sponsored, or otherwise approved by <<company name>>. Products are warranted only to meet Micron's production data sheet specifications. Products, programs and specifications are subject to change without notice. Dates are estimates only. ©2017 Micron Technology, Inc. All rights reserved. All information herein is provided on an "AS IS" basis without warranties of any kind. Micron and the Micron logo are trademarks of Micron Technology, Inc. All other trademarks are the property of their respective owners. Rev. A 12/17, CCM004-676576390-10936