NVMe™ SSDs – A New Benchmark for OLTP Performance

OLTP Platforms Thrive on Fast, Consistent Results Delivered by NVMe SSDs

Overview

A recent DB-Engines ranking shows that Microsoft® SQL Server® is one of the most popular database management systems in the world. Top database platforms manage high-capacity, high-bandwidth transaction-based applications for online transaction processing (OLTP)-type workloads like order entry and fulfillment, real-time data analysis, e-commerce and real-time order management systems. These systems give fast access to mission-critical data, enabling transaction processing with ultra-low and consistent latency, where access delays can be extremely costly.

This technical brief discusses how we used standardized OLTP performance metrics and a data set that exceeds available system memory (to test storage system I/O) to compare new orders per minute (NOPM), database average response time and response time consistency on a SQL Server using SSD versus legacy HDD storage.

We used the same base hardware (server, CPUs and DRAM) with all three storage configurations:

- **Legacy:** 16x 300GB 15K RPM HDD configured RAID 10 (baseline configuration for comparison)
- **NVMe 1:** 2x 3.8TBin Mirrored Storage Spaces
- **NVMe 2:** 4x 3.8TB in Mirrored Storage Spaces

We found that the NVMe SSDs generated more NOPM to bring more value to OLTP workloads on Microsoft SQL Server. A pair of NVMe SSDs supported 730,000 NOPM and four supported over 900,000 NOPM. At these rates, responses were incredibly quick and consistent.

Fast Facts

- Complete more transactions with NVMe — more orders, more fulfillment, more to your bottom line
- Data can't wait — SSDs with NVMe release your data's potential like never before
- Compared to the baseline 16x HDD configuration:
  - 2x NVMe SSDs supported 216X more New Orders per Minute
  - 4x NVMe SSDs supported 276X more New Orders per Minute
- Databases hosted on NVMe response times were 77% to 90% lower and much more consistent

More Orders per Minute Brings More Value

SSDs are a mainstay of high-performance, low-latency IT systems. High-capacity, ultra-performance NVMe SSDs drive those systems farther and faster, processing more data and bringing more value.

More and more OLTP platforms are moving to NVMe, and the differences between the capabilities of NVMe today and what we used to think of as a performance HDD configuration are greater than ever, with the legacy configuration's 15K RPM HDDs being painfully slow in camparison. In OLTP systems, additional transactions can represent more orders, more fulfillment or more detailed analysis — all bringing more value.

The magnitude of the difference between current NVMe designs and legacy standards is evident in Figure 1, which shows each configuration's relative NOPM at a system load just before the test reached a stop condition. (See the How We Tested section for stop condition details.)

The NVMe configurations' NOPM are extremely high. The 2x NVMe achieved 216X more NOPM than the baseline and the 4x NVMe achieved 276X more NOPM than the baseline.

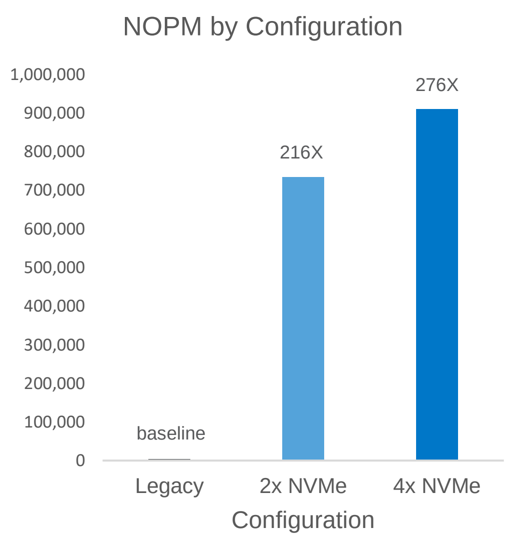

Figure 1: Relative NOPM

NVMe Brings Faster, More Consistent Responses

Many applications require high NOPM while quick, consistent database response (low latency) may be more important for applications that are very time-sensitive.

As shown in Figures 2a and 2b below, we calculated and compared the mean latency and the 99.9th percentile latency (a good indicator of latency consistency) at system load just before the test reached a stop condition (see the How We Tested section for stop condition details) for the three storage configurations. We used the same metrics, database and test conditions for each configuration.

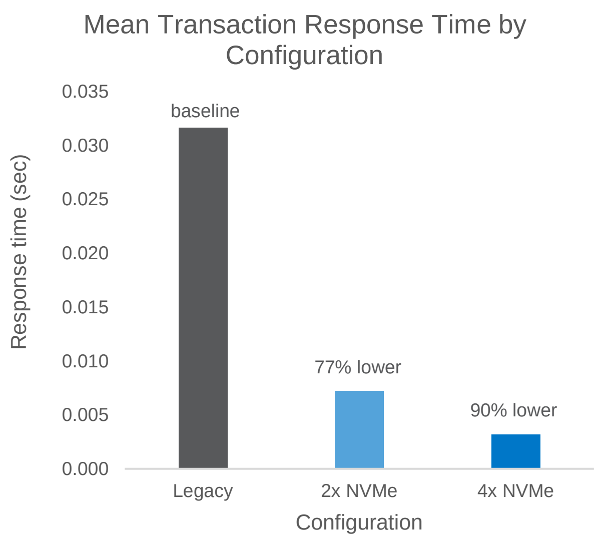

Figure 2a: Mean transaction time

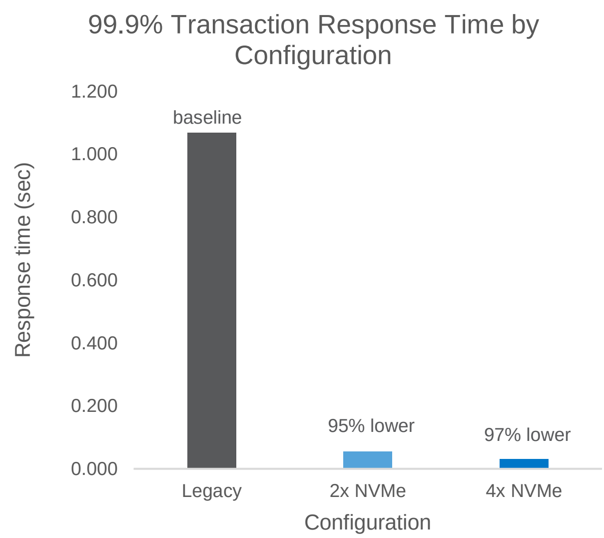

Figure 2b: 99.9% transaction time

Both NVMe configurations show low mean transaction times compared to the baseline configuration. Figure 2a shows that the 2x NVMe configuration measured 77% lower and the 4x NVMe configuration measured 90% lower than baseline.

Figure 2b shows that the NVMe configuation response times were more consistent than the baseline configuration (lower 99.9% transaction response time).

The two comparisons indicate that both NVMe configurations respond more quickly and more consistently than the baseline configuration. The legacy configuration shows the opposite — much higher mean transaction response time that is far less consistent.

The Bottom Line

Mission-critical data can't wait. Access delays or inconsistency can be extremely costly. Using NVMe SSDs enables fast transaction processing and fast, consistent response times.

In our testing, NVMe SSDs demonstrated tremendous benefits and new capabilities for one of the most popular database management systems and most challenging workloads — Microsoft SQL Server and OLTP. Supporting immense loading, far greater NOPM with lower and more consistent latency means more orders and more transactions completed faster and more consistently.

Learn more about NVMe SSDs and their transformative effect on your business at micron.com.

How We Tested

To ensure a fair assessment of the expected maximum NOPM of each configuration, we took a configuration-specific approach. We measured each configuration's NOPM at the maximum load the platform could reasonably support, as opposed to comparing NOPM and latency at an arbitrary load.

Prior to testing, we established stop conditions (Tables 1 and 2). As we tested, we increased the load until the test a reached a stop condition after which we stopped increasing the load and used the NOPM and latency values recorded when we reached the stop condition.

Limit	Stop Condition
CPU utilization	80%
90th percentile average transaction response time	See Table 2
Average log disk (partition) write latency	5ms
NOPM plateau	When NOPM fails to increase with higher load

Table 1: Stop conditions [1,2]

We set the 90th percentile transaction response time to the values in Table 2, which each reflect common tolerance limits.

Transaction	90th 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

1. We set the stop condition for CPU utilization at 80%. Many IT organizations plan for a platform upgrade when CPU utilization reaches 50% and implement that plan when it reaches 80%.
2. We sized the data set to ensure it was large enough to ensure storage I/O (data set size about 2X the memory size) but did not occupy more than 80% storage capacity.

Determining Maximum Load by Configuration

This section shows the test condition(s) that established each configuration's maximum load.

Legacy Configuration Stop Condition: Average Log Write Latency

Figure 3 shows the legacy configuration's average log disk (partition) write latency by load. The stop condition is shown in red (at which point the average log disk [partition] write latency significantly exceeds our 5ms stop condition).

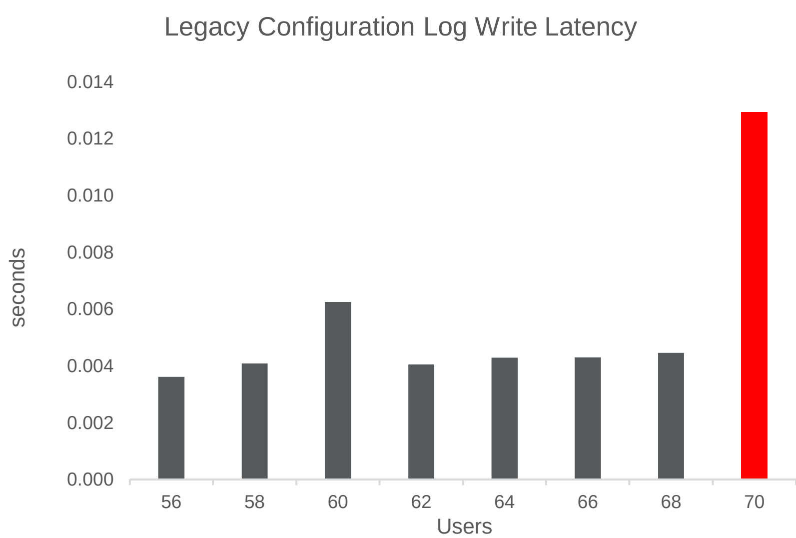

Figure 3: Legacy configuration stop condition

2x NVMe Configuration Stop Condition: NOPM Plateau

Figure 4 shows the NOPM (vertical bars) and 99.9% response time (line) by load for the 2x NVMe configuration. There is a NOPM plateau just before the far right of the loading range (the loading value when NOPM drops is shown in red). Note too there is a significant jump in 99.9% response time at the same loading point.

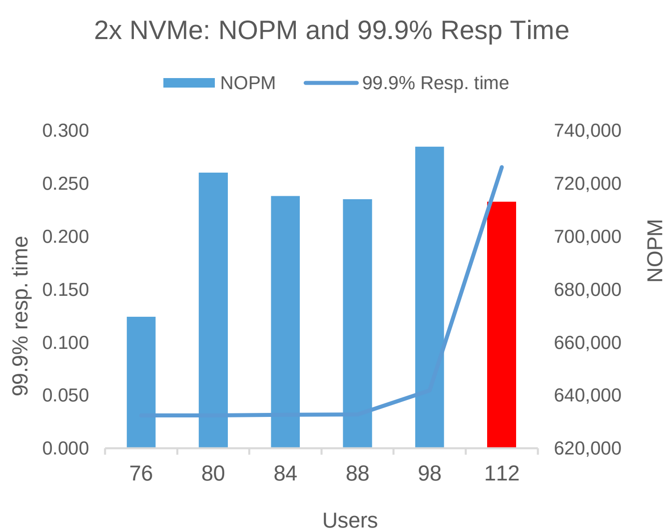

Figure 4: 2x NVMe stop condition

4x NVMe Configuration Stop Condition: %CPU Utilization

Figure 5 shows the %CPU utilization versus loading for the 4x NVMe configuration. %CPU utilization reaches 80% on the right side of Figure 5. We used a loading point where %CPU utilization exceeds 80% and the loading level aligns to two logical users per thread.

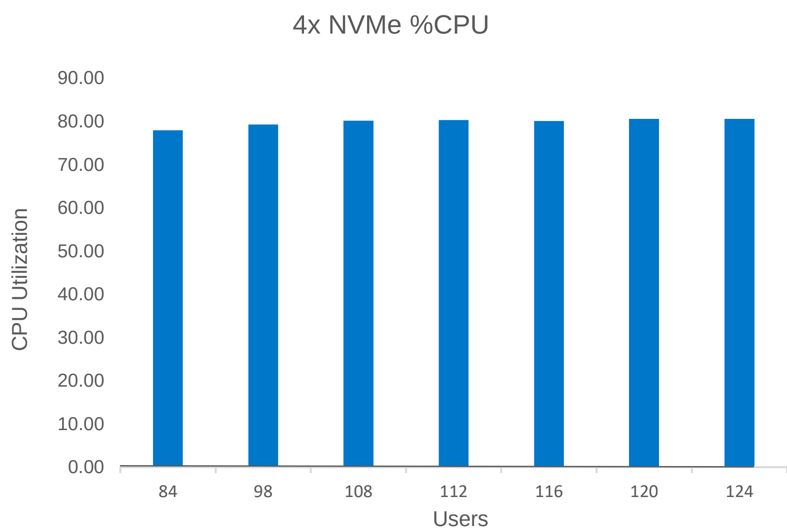

Figure 5: 4x NVMe stop condition

Table 3 shows the hardware configuration details used in all testing.

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	Mirrored Storage Spaces
HDDs	15K RPM, 300GB (x16)
HDDs Configuration	RAID 10

Table 3: Hardware configuration