Scaling High-Demand OLTP Workloads: Micron 9300 NVMe SSDs

OLTP Lives on Fast, Consistent Results. Micron NVMe SSDs Deliver

Overview

Starting with SQL Server 2017, Microsoft enables users to run SQL on their platform of choice, Windows or Linux. According to Microsoft, it is the same SQL Server engine and is very similar regardless of which platform you choose[1].

In this technical brief, we show how Microsoft SQL Server 2017 with Micron 9300 NVMe SSDs scales on Linux under a heavy demand OLTP workload. We compare business throughput and power efficiency (measured as mW per NOPM).

We found that MS SQL Server 2017 on Linux scales extremely well – showing increased throughput latency and higher power efficiency as we scaled.

We use two storage configurations to test scaling: one with a pair of NVMe SSDs and second with four. Each configuration used the same base hardware (SSDs, server, CPUs and DRAM):

NVMe # 1: 2x 9300 PRO 3.8TB in Mirrored Storage Spaces (baseline configuration)

NVMe #2: 4x 9300 PRO 3.8TB in Mirrored Storage Spaces ('scaled' configuration)

We found that a four NVMe SSD configuration demonstrates real scaling with significant benefits for demanding OLTP workloads with Microsoft SQL Server on Linux, generating a 35% higher transaction rate and lowering power consumed 18% (per NOPM).

Fast Facts

Scaling 9300 NVMe SSDs using MS SQL Server on Linux enables

35% higher business throughput.

Increased New Orders Per Minute can translate into increased transaction rates and numbers of orders fulfilled.

18% Less power per transaction

Drive increasing transaction rates and power efficiency

Micron 9300 series of SSDs with NVMe

1. See https://docs.microsoft.com/en-us/sql/linux/sql-server-linux-overview?view=sql-server-2017 for additional details on Microsoft SQL Server on Linux.

NVMe SSDs are a mainstay of high performance, low latency IT systems. Performance-focused NVMe SSDs like the 9300 series drive OLTP systems farther and faster, scaling to process more data and bring more value. When you need the ability to scale your business, the 9300 delivers.

The advantage of scaling up to four Micron 9300 PRO SSDs from two can clearly be seen in Figure 1. Each configuration's relative transaction rates is shown at a system load just before the test reached a stop condition (see "How We Tested" for stop condition details).

The scaled configuration with four 9300 SSDs delivered 35% more business throughtput.

Figure 1 shows NOPM for each storage configuration (up to saturation point, see "How We Tested" for additional details on saturation).

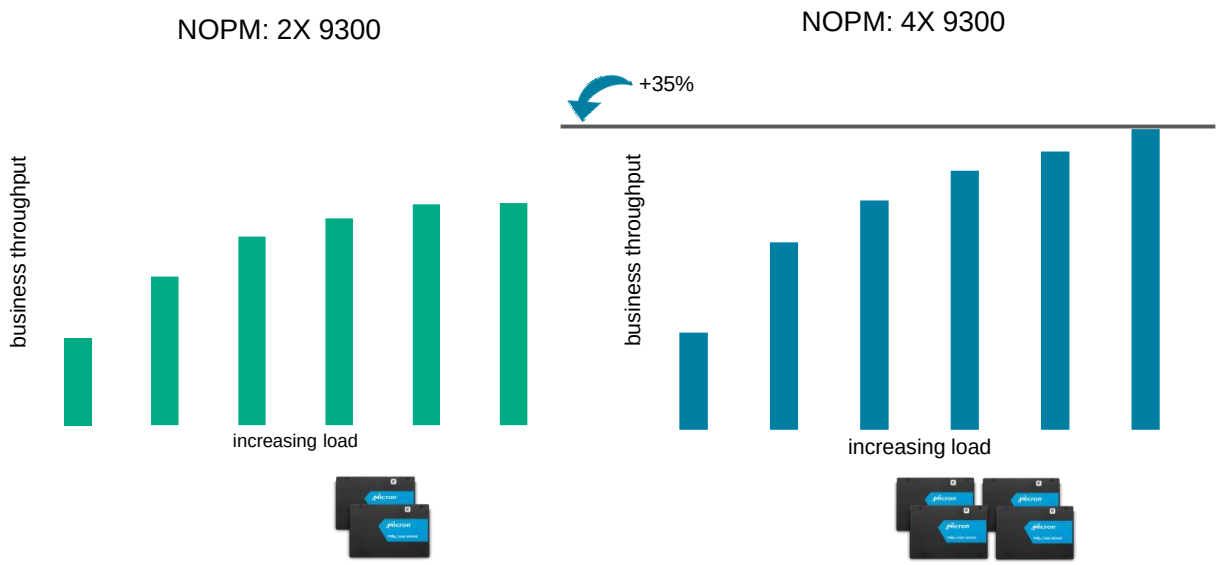

Figure 1: Relative NOPM by Storage Configuration (Taller is Better)

At system saturation, the scaled configuration delivered 35% more NOPM than the standard configuration (presaturation values shown for completeness).

Power Stingy, Scaled Efficiency

Power consumption per unit of work (NOPM) is a good metric to determine if power efficiency scales like the other metrics considered.

We measured instantaneous system power consumption and NOPM for each configuration (at the point of platform saturation) and divided power in mW by NOPM. This helps determine if power efficiency also scaled (we expected total power consumption for the system to increase with the scaled configuration simply due to the increased number of SSDs).

Figure 2 shows that the scaled (4X 9300) configuration consumed 18% less power per NOPM (a good measure of power efficiency).

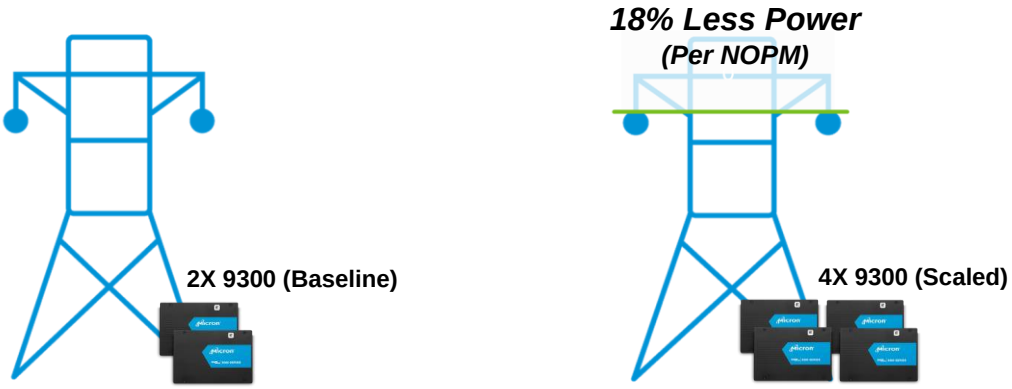

Figure 2: Power Consumption per NOPM (Lower is Better)

When we consider Figures 1 and 2 together, the 9300 scaled configuration delivers more throughput with better power efficiency.

The Bottom Line

We compared two NVMe SSD configurations supporting Microsoft SQL Server on Linux with an OLTP workload. The first, baseline configuration used a pair of Micron 9300 3.84TB PRO SSDs, the scaled configuration extended that to four of the same SSDs (each configuration used Storage Spaces Direct).

During testing we saw excellent results. Scaling NVMe storage generated 35% higher transaction rates (NOPM) accompanied by an 18% decrease in power consumption.

Mission-critical data at scale can't wait. Access delays or inconsistency can be extremely costly. Using NVMe SSDs enables fast transaction processing and fast, consistent response times. Learn more about NVMe SSDs and their transformative effect on your business at micron.com.

micron.com

How We Tested

To ensure a fair assessment of each configuration's capabilities – the maximum transaction rates one could expect – we took a configuration-specific approach. We wanted to ensure we measured each configuration's NOPM at the maximum load the platform could reasonably support (as opposed to comparing these rates and latency at an arbitrary load).

Before testing, we established stop conditions. As we tested, we increased the load until the test 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.

Determining Maximum Load by Configuration (Stop Conditions)

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

Stop Condition: Performance Plateau

Figure 3 shows each configuration's performance (NOPM) versus transaction latency. The circled points represent points at which increasing the system load resulted in a performance increase that did not show a commensurate increase in performance. The test data shown in this brief reflects this maximal loading for the two 9300 configurations: 2X 9300 baseline: 192 users; 4X 9300 scaled configuration: 408 users.

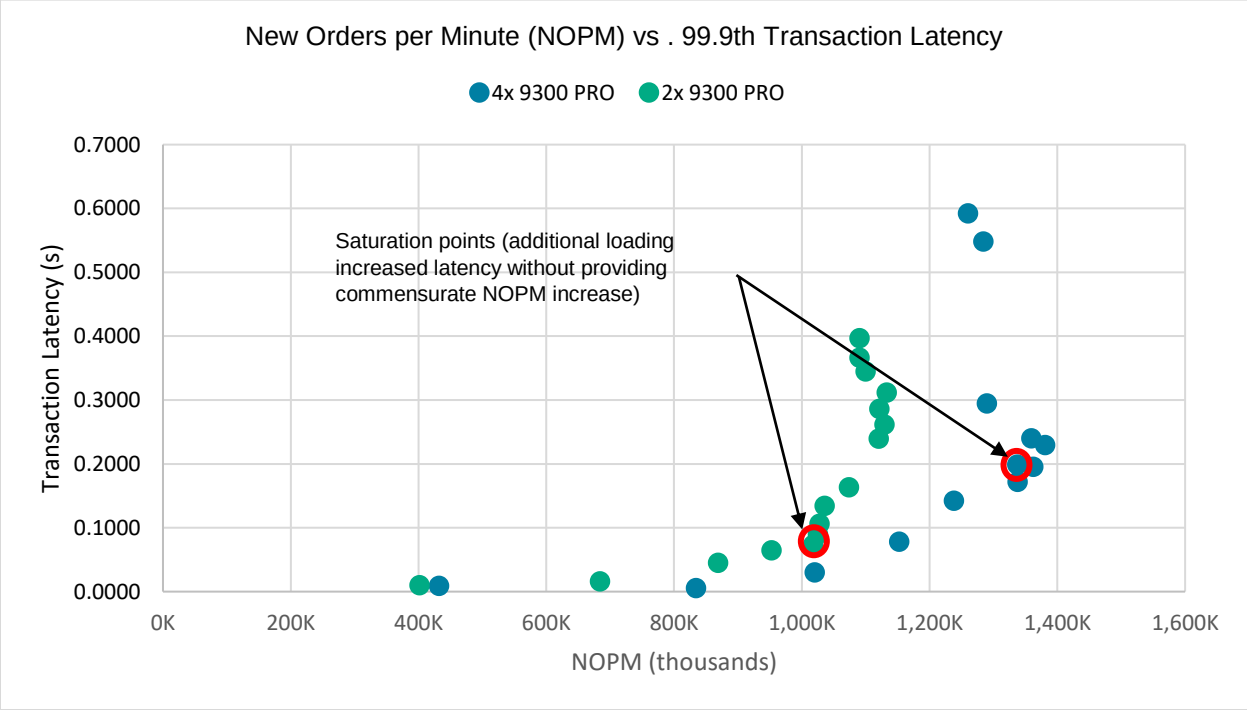

Figure 3: 9300 Stop Condition

Tested Server Configuration Details

Tables 1 and 2 show the hardware and software configuration details.

Component	Description
Server	2U, 2-socket (Intel based)
CPUs	Intel Xeon Platinum 8168 24-core (x2)
Memory	384 GB Micron DRAM
NVMe SSDs	Micron 9300 PRO 3.8TB (X2 baseline, X4 scaled)
Storage Configuration	NVMe: Mirrored Storage Spaces
Operating System	CentOS 7.6 – 3.10.0-957.12.1.el7.x86_64
SQL Server	Microsoft SQL Server 2017 Enterprise Edition (x64)

Table 1: Hardware and Software Configuration

Component	Description
Server BIOS	V2.2.10 High performance mode: Enabled Intel VT: Disabled CPU Hyper-threading: Enabled
OS Storage	240 GB Micron M.2 SSD (X2) RAID (Dell BOSS)
Database Storage	2X 9300 (baseline); 4X 9300 (scaled) TempDB: 500 (Stripe) Log: 500GB Data: 3 TB – 6.24 TB 9300 FW: 11300B20
OS Settings	kernel.numa_balancing=0 vm.max_map_count=262144 kernel.sched_min_granularity_ns=10000000 kernel.sched_wakeup_granularity_ns=15000000 vm.dirty_ratio=40 vm.dirty_background_ratio=10 ATTR{bdi/read_ahead_kb}="2048" (udev rule for drive)
File System	XFS
SQL Server Settings	Target Recovery Time: 60 seconds Recovery Model: Simple Maximum Degree of Parallelism: 1 Lock Pages in Memory: Enabled Number of database files: 48 Number of tempdb files: 8

Table 2: Configuration