MAINTAINING HIGH OLTP PERFORMANCE (→) IN MEMORY CONSTRAINED ENVIRONMENTS

SANDISK™

MSSQL Benchmark testing on the SANDISK® SN861 SSD

ALEX HAYEK
GOPAL SHARMA

"WHITE PAPER"

SEPTEMBER.2025

MAINTAINING HIGH OLTP PERFORMANCE (→) IN MEMORY CONSTRAINED ENVIRONMENTS

MSSQL Benchmark testing on the SANDISK® SN861 SSD

Executive Summary

Today's database servers are being tasked with more than ever before, with more consideration being put into how many resources these servers are given. Microsoft SQL Server (MSSQL) is one of the most popular databases in the world, many enterprises rely on it for **mission-critical** transaction processing at the core of their products and services. Not all database servers are created equal, and to show excellent performance even when limited on CPU and DRAM resources SANDISK conducted benchmarks using HammerDB, an open-source benchmarking tool, to generate an online transactional processing (OLTP) load for MSSQL. The SANDISK SN861 NVMe SSD rises to the challenge and is an ideal choice for database environments facing constraints, or for a game-changing level of performance in unconstrained environments.

The Tools and Software

To conduct these benchmarks SANDISK chose the following tools and software to highlight the performance of the SANDISK SN861 E3.S SSD

Microsoft Server 2022 Datacenter 21H2 20348.3807 as the host operating system, which gives us a native environment to test Microsoft SQL Server 2022 and allows for easy configuration for the multiple tests that were run, as well as direct metric collection for system-level metrics such as disk performance.

FIO or Flexible IO Tester is a free and open-source tool that allows for simulating a given I/O workload in a pre-defined and configurable way. SANDISK uses FIO to pre-condition the drives to achieve steady state before tests.

HammerDB is a free and open-source database benchmarking application, built for all platforms and capable of testing Microsoft SQL Server, IBM Db2, Oracle, PostgreSQL, MariaDB, and MySQL. The benchmarks TPROC-C and TPROC-H are fair-use implementations of the popular TPC-C and TPC-H benchmarks. While results can be interpreted similarly, they are not directly comparable.

Microsoft SQL Server 2022 is the latest release of a relational database management system (RDBMS) written by Microsoft. This version of SQL Server includes advancements for faster query performance plus improved query tuning with features like parameter-sensitive plan optimization.

The Setup

For this benchmark, the hardware used for the database was a Single Socket SuperMicro ASG-1115S-NE316R, a 1U server with 16 front-loading PCIe Gen 5 E3.S NVMe slots. This was paired with the AMD EPYC 9534 64 Core 128 Thread Processor, 256GB DDR5-4800 Registered ECC 16GB x16, and 100GbE networking. The base operating system was installed on a separate PCIe Gen 4 M.2 NVMe boot drive. The drive under test is a single SANDISK SN861 E3.S 16TB (15.36TB usable) PCIe Gen 5 NVMe SSD (SDS6BA1A1PBPAX3).

Server	MFG/Model	Motherboard/ Model #	CPU	DRAM	OS	Primary Drive	Drive Under Test (DUT)	System Variable during test	NIC
System Under Test	ASG-1115S-NE316R	Super H13SSF	Genoa 9534 DP/UP 64C/128T 2.45G 256M 280W SP5*	16x 16GB DDR5-4800 1Rx8 LP (16Gb) ECC REG DIMM, HF, RoHS**	Windows Server 2022 Datacenter 21H2 20348.3807	SSD M.2 NVMe PCIe4 480GB 1DWPD TLC D, 80mm	Sandisk SDS6BA1A1PBPAX3	See asterisks	AIOM 2-port 100GbE QSFP28,Mellanox CX-6 DX
Load Generation	AS -2125HS-TNR		Genoa 9354 DP/UP 32C/64T 3.25G 256M 280W SP5	16x 16GB DDR5-4800 1RX8 LP (16Gb) ECC RDIMM	Windows Server 2022 Datacenter 21H2 20348.3807	SSD M.2 480GB NVMe PCIe 4.0 M.2 22x80mm 3D TLC	N/A	N/A	MCX516A-CCAT PCIe 2-port 100GbE QSFP28 Gen3.0 x16 CX-5
	*Reduced to 32 Cores w/o SMT for all tests, except Server max								
	**Varied available memory in each run of test in 32Gb increments								

For the HammerDB Client, we used another server with similar specifications, also leveraging 100GbE networking. This client was in the same location and connected directly via a TOR switch to the database server.

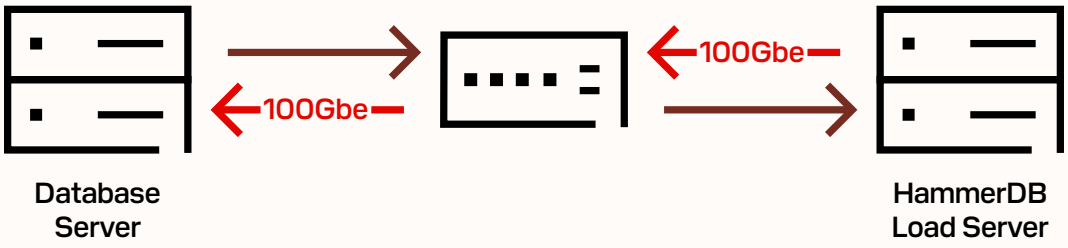

To prepare the server for the workload, pre-conditioning was performed on the drive with a full sequential drive write followed by two random drive writes to achieve a steady operating state. Once in steady operating state, HammerDB was used to create and populate the schema for the tests with the TPROC-C SQL Server version of the benchmark with 2500 Warehouses, resulting in a database file with size on disk of ~240GB. This was backed up using the simple SQL Server 2022 backup and restored for each individual run. The intent of sizing the database at 240GB was to ensure it could fit in memory in its entirety.

The Benchmark

To best reflect the transactional workload in many enterprise deployments, we chose the TPROC-C benchmark implemented by HammerDB. This focuses on OLTP; it simulates a system to fulfill orders from customers and tracks both New Orders Per Minute (NOPM) and Transactions Per Minute (TPM). For this benchmark, we focus on the raw performance of TPM, which includes transactions that are not included in the New Order transaction and that are solely counted for NOPM. TPM represents the entirety of work that the database conducts during the benchmark.

Transactions can consist of one of the following types and their typical share of count: New Order (45%), Payment (43%), Delivery (4%), Order-status (4%), Stock-level (4%).

To establish where we achieve CPU and Memory saturation prior to disk saturation, we first conducted a fully unconstrained test, then scaled back hardware availability. In our case, the server had half of its 64 cores disabled (32) and SMT disabled to ensure that the CPU would reach maximum utilization with 601 Virtual Users and 2500 Warehouses. This means that the CPU was averaging over 97% busy throughout each test run. A good way to think of this is maxing out available CPU resources as the control for changing other variables in this benchmark. We included the fully unconstrained results as a data point of how the drive is fully capable of performing in an unconstrained environment, but it is worth noting that these are not directly comparable, nor the focus of this benchmark.

HammerDB	Warehouses	Test Duration	Virtual Users	VU Use All DB	Test Variable
TPROC-C SQL Server Default unless specified ->	2500 (240GB)	10 mins	601	Yes	None

After establishing the correct number of CPU cores necessary to maximize disk load we applied pressure to the disk by running 8 tests starting with the database being configured to only use 32GB of DRAM and ending with the full 256GB of DRAM available increasing each test by 32GB. By limiting the amount of DRAM available to the database application, more data must be read from disk as there is less space to cache data in memory. This demonstrates the drive under test's overall ability to maintain high performance even when being tasked with the majority of reads in addition to all the writes.

SQL Server	Backup	AutoGrow Files	Test Variable	Backup / Restore
SQL Server 2022 16.0.1000.6	Simple	Yes	SQL Server Memory from 32GB to 256GB in 32GB increments	Backup/restore DB after each memory variation test result

The Results

As we mentioned previously, an initial fully unconstrained system resource run was conducted on the SANDISK SN861. That benchmark showed a TPM of over 5.19 million, and saw average write and read performance of ~3,700MB/s in writes and 251MB/s in reads for the 10-minute test. With 128 CPU threads at 48% utilization and 384GB of DRAM, these results are some of the upper limits of performance for a single SN861.

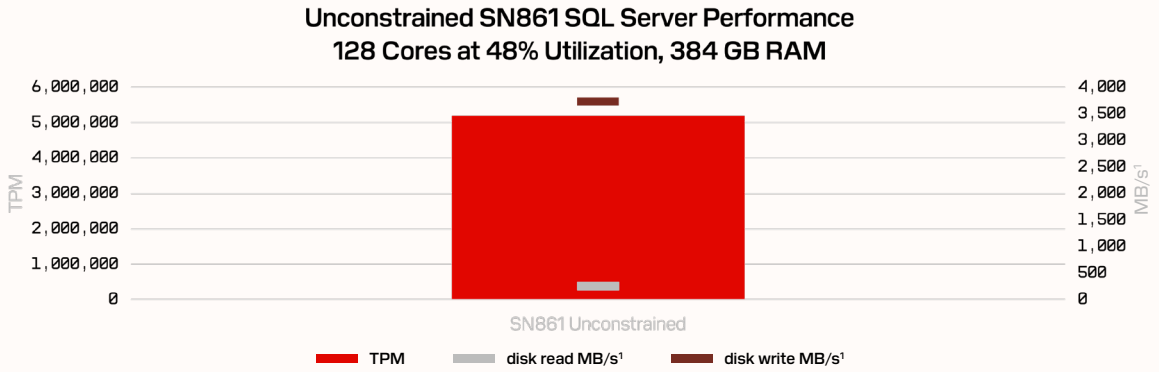

In our constrained test results, the SANDISK SN861 was able to maintain 72.4% performance while DRAM was constrained to just 32GB or 12.5% of the 256GB allocation. We also observe a steady increase in performance as we incrementally add DRAM until we reach 160GB or 62.5% of the 256GB allocation, which reaches 96.7% of the performance level when compared to the run with the DRAM allocation of 256GB. An important note is that for all these tests, CPU utilization remained above 97% for the duration of each test apart from a small ramp time in the first two minutes. These are promising results that show when limited on CPU and on DRAM, the SANDISK SN861 maintains a high level of transactions per second.

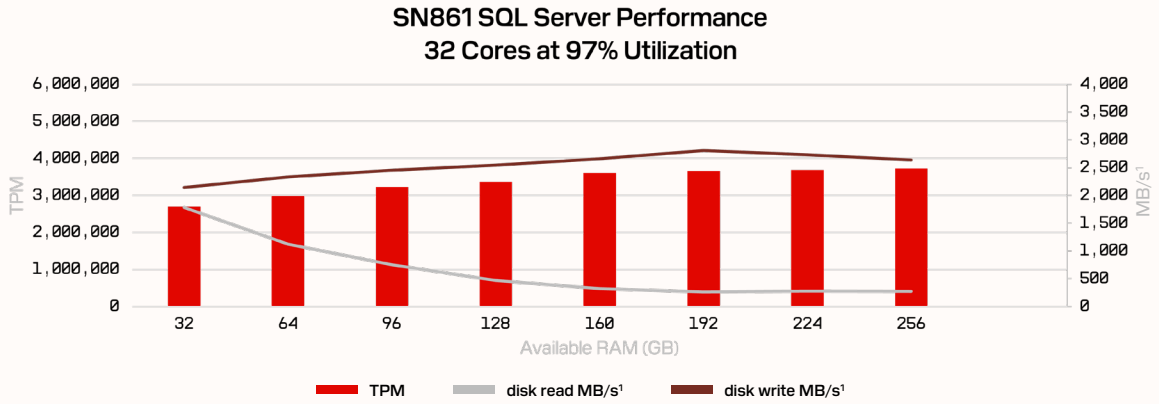

When we look at the disk read and write activity through the tests we see that read pressure comes down as memory is added to the database, and write pressure increases due to the addition of more transactions. In addition, we see beyond 160GB of DRAM allocation that the disk read throughput plateaus, with most indexes and table data being cached in memory.

Conclusion

The SANDISK SN861 is an excellent choice for infrastructure teams looking to maximize their performance in constrained environments. The high performance nature of the SANDISK SN861 allows for databases to perform well even with a lack of system resources whether that be CPU or DRAM. In addition we can extrapolate the environment constraints to large servers with multiple VMs and resources made available to each VM – being able to have more database servers per node, while maintaining high levels of performance with the SANDISK SN861 is something we will explore in the future so stay tuned!

1. One megabyte (MB) is equal to one million bytes, one gigabyte (GB) is equal to 1,000MB (one billion bytes), one terabyte (TB) is equal to 1,000GB (one trillion bytes), and one petabyte (PB) is equal to 1,000TB. Actual user capacity may be less due to operating environment.up to 60°C. Non-operational storage temperature does not guarantee data retention.

951 SANDISK DRIVE
MILPITAS, CA 95035, USA
WWW.SANDISK.COM