

Real-World Microsoft SQL Server™ Performance on Datrium DVX



Highlights

A single VM running Microsoft SQL Server 2014 TPC-C database achieved over 2 Million SQL Server transactions per minute (TPM) with HammerDB. The SQL Server TPC-C benchmark read average latency was 0.9 milliseconds and the flash read hit rate was 100%. The average application write latency was slightly north of 2 milliseconds.

Introduction: Microsoft SQL Server Performance on Datrium DVX

The purpose of this technical note is to demonstrate the real-world production performance of a virtualized Microsoft SQL Server™ on a Datrium DVX platform.

IT organizations virtualizing Microsoft SQL Server with Datrium DVX find that the solution is fast, efficient, and most importantly, predictable. With Datrium DVX, enterprises can scale application performance independently from storage capacity due to an innovative Open Converged approach. Furthermore, with Datrium, IT admins can choose to use existing x86 servers, deploy new x86 servers, or deploy a Datrium DVX Rackscale platform to build an elastic private cloud infrastructure for literally any application workload.

This document assumes that the reader has some familiarity with Microsoft SQL Server administration, HammerDB, and the Datrium DVX platform. This document is not intended to portraint a Datrium DVX reference architecture.

Datrium DVX

For this Microsoft SQL Server benchmark, a DVX CN2000 server was used as the Datrium DVX compute node. In this document, a Datrium DVX compute node uses the VMware ESXi hypervisor.

The server configuration chosen for this benchmark performance test is a Datrium DVX Compute Node. Additionally, the Datrium DVX compute node was configured with 2TB of usable Flash to service all active data for the Microsoft SQL Server databases. Figuring out what features should be enabled in an environment with 1000s of VMs and applying settings on a per-VM or per-some-group-of-objects basis is simply too cumbersome for the IT admin and unnecessary. Datrium DVX was designed from the ground up to support all data services turned ON all the time.

Important: The benchmark presented in this document uses in-line checksumming, compression, deduplication, and distributed Erasure Coding to fully protect the system and databases against failures.

Configurations

Datrium DVX System Information

table 1

Datrium DVX Compute Node (n1667):	Description:
Hypervisor	VMware ESXi 6.0
CPU	Intel Xeon® E5-2680 v4 @2.4GHz
CPU Sockets	2
CPU Cores per Socket	14
Logical Processors	56
RAM	512GB
Usable Flash	2TB
Datrium DVX Data Node:	Description:
Model	DVX CN2000
Software version	2.0.2.0

Virtual Machines

A virtual machine (VM) was configured for the Microsoft SQL Server performance testing on the Datrium DVX compute node. The Windows 2012 virtual machine had the following vDisks created to separate Operating System and Microsoft SQL files:

- Operating System disk for:
 - vDisk 0 - Windows 2012 R2 Data Center
- Multiple data disks for:
 - vDisk 1 - database
 - vDisk 2 - database log
 - vDisk 3 - TempDB
 - vDisk 4 - TempDB log

All vDisks for the Microsoft SQL Server VM were Thin Provisioned, and there is no tuning required for the ESXi datastore in use. In fact, vDisk Thin Provisioning is a Datrium best practice. The Datrium DVX storage is presented as an NFS mount to the ESXi server and is natively optimized for performance. For a deeper dive into Datrium DVX architecture and how the system is natively optimized for application performance, review the [Datrium DVX Architecture technical whitepaper](#).

Virtual Machine System Information

table 2

Virtual machine resource (dom-winsql1)	Description:
Number of CPUs	12
RAM	256GB
Network Adapter	VMXNET3
OS	Windows 2012 R2 Data Center
Microsoft SQL Server	Microsoft SQL Server 2014 Enterprise Edition

Datrium recommends the reading of the “Best Practices for Deploying SQL Server Using vSphere” section in the [VMware’s Microsoft SQL Server Best Practices Guide](#) for additional best practices and guidelines on:

1] Host Configuration

2] CPU Configuration

3] Memory Configuration

4] Network Configuration

The Microsoft SQL Server VM – in this test – was not setup for HA at the VM or application layer. For recommendations and guidelines on HA and DR, review Datrium’s [Microsoft SQL Server HA and DR with Datrium DVX white paper](#).

HammerDB

[HammerDB](#) is a free database load testing and benchmarking tool for databases. For this test, we have chosen the TPC-C schema to build a roughly 500+ GB database (to start) of 5,000 warehouses and used more than 500 concurrent users to drive the database transactions with no think-time enabled.

[TPC-C](#) simulates online transaction processing (OLTP) workloads that are typically identified by a database receiving both requests for data and multiple changes to the data from user transactions. The HammerDB workload simulated roughly a 70:30 split of read/write transactions at 8KB block size in most of the test runs.

Multiple test runs were executed as we tuned the SQL Server application parameters to get the maximum SQL Server transactions per minute (TPM). We chose to perform our tests with a Microsoft SQL Server VM that has 12 vCPU and 256 GB RAM.

Microsoft SQL Server 2014

The Microsoft SQL Server 2014 database was built and populated using HammerDB. At the start of the test run, the initial size of the TPC-C (OLTP) database was 500+GB, and the Microsoft SQL Server data disks are later extended as the Microsoft SQL Server database grew during the performance tests.


Review the “SQL Server and In-Guest Best Practices” section in the [VMware’s Microsoft SQL Server Best Practices Guide](#) for additional recommendations on optimization at the SQL Server and Windows Server side.

Architecture

For Datrium's best practices for VMware and applications on Datrium DVX, existing Datrium customers can access the current [Datrium DVX Handbook](#).

High Level Overview

Figure 1 shows a high-level overview of the Microsoft SQL Server 2014 setup on a Datrium DVX.




Windows 2012 SMB File Share Clients were used as a target for Microsoft SQL Server native backups. Additionally, during the tests Datrium DVX snapshots were used to protect the Microsoft SQL Server, allowing us to perform tests as normal Datrium DVX operations and to be able to revert the environment to a known good state, if necessary.

The platform used for this benchmark uses a single compute node and a single data node. While the minimum DVX configuration is one compute node and one data node, and that is what is being tested in this benchmark, DVX Split Provisioning (see figure 2) allows scalability up to 128 compute nodes and 10 data nodes in a single system. The system-wide performance and capacity is up to 200 gigabytes per second (GB/s) of read bandwidth and 18 million IOPS, 10 GB/s write throughput, and 1.7 petabytes (PB) of effective capacity.

Datrium DVX Scale-Out Matrix

figure 2



SQL Tests

As part of this performance benchmark exercise the following tests were performed:

- 1] HammerDB SQL Server TPC-C benchmark with Datrium DVX
- 2] HammerDB SQL Server TPC-C benchmark with Datrium DVX during Steady-State

During the tests, VM level performance metrics in the Datrium DVX was utilized to monitor and tune parameters at the application layer. As with any performance testing, it is not recommended to tweak all available parameters simultaneously. It is best to tune one parameter at a time and monitor if performance is improved before making additional changes. Review [Microsoft's recommendation and guidelines for the "max degree of parallelism" configuration option](#) for more information on Microsoft SQL Server tuning.

Monitoring


SQL Server TPC-C benchmark performance

Figure 3 demonstrates the HammerDB SQL Server performance numbers achieved with Datrium DVX in Fast Mode. Fast mode is the regular operation mode for a Datrium DVX compute node, whereas a maximum of 20% host CPU utilization is allocated to the storage IO and data services. In this benchmark test run (Figure 3), the Microsoft SQL Server 2014 TPC-C database achieved over 2 Million SQL Server transactions per minute (TPM) with HammerDB TPC-C workloads.

Datrium DVX also provides the ability to enable Insane Mode allowing compute nodes to utilize up to 40% of host CPU to improve storage IO operations. *This benchmark is not using Insane Mode.*

HammerDB SQL Server TPC-C Benchmark

figure 3



The SQL Server TPC-C benchmark average read latency was 0.8 milliseconds and the flash read hit rate was 100%. The average write latency was slightly above 2 milliseconds (see figure 4).

Datrium DVX DAVG Latency

figure 4

DAVG latency

● 0.8 ms read ● 1.9 ms write

SQL Server TPC-C benchmark performance in Steady-State

In this document, we define Steady-State when space reclamation and snapshots are enabled and in use to protect VMs and applications. Generally, during Steady-State, space reclamation processes would kick in when the data node cluster reaches 75% capacity utilization.

A Datrium DVX protection group was created during the test to protect the Microsoft SQL Server VM with multiple native DVX snapshots being created during the performance benchmark. The objective of this test is to determine if application performance is affected during Steady-State.

Figure 5 shows real-time statistics captured while space reclamation (manually started) and snapshots were in use to protect the Microsoft SQL Server VM. In this test run (see figure 5), the HammerDB transaction counter achieved north of 1.8 million TPM and the Microsoft SQL Server VM performance is not much affected during Steady-State.

HammerDB SQL Server TPC-C benchmark during Steady-State

figure 5




Figure 6 shows the DAVG latency during DVX Steady-State. The SQL Server TPC-C benchmark average read latency was 0.9 milliseconds and the flash read hit rate was 100%. The average write latency was slightly above 2 milliseconds.

Datrium DVX DAVG Latency during Steady-State

figure 6

DAVG latency

● 0.9 ms read ● 2.1 ms write

Conclusion


The results demonstrate that Datrium DVX can host Microsoft SQL Server 2014 databases with sub-millisecond latencies for read operations. A single Microsoft SQL Server VM - with 12 vCPU and 256 GB of RAM – can achieve 2 Million TPM before reaching 100% VM vCPU utilization.

During Steady-State, the HammerDB counters demonstrated that the Microsoft SQL Server VM could nevertheless produce over 1.8 Million TPM, showing that snapshot chaining and space reclamation processes do not seriously affect application performance.

With the Microsoft SQL Server VM in this configuration, it is possible to keep the transactions consistently above 1.7 Million TPM but that means the Microsoft SQL Server VM busy is always at 100% vCPU utilization (Figure 7).

SQL Server VM CPU and Memory Performance

figure 7



While the minimum DVX configuration is one compute node and one data node, and that is what is being tested in this benchmark, DVX Split Provisioning allows scalability up to 128 compute nodes and 10 data nodes in a single system. The overall system-wide performance and capacity is up to 200 gigabytes per second (GB/s) of read bandwidth and 18 million IOPS, 10 GB/s write throughput, and 1.7 petabytes (PB) of effective capacity. It is important to note that during the HammerDB TPC-C benchmark performance tests, the average read latencies was sub-milliseconds with more than 1.8 Million TPM - even during Datrium DVX Steady-State.

Important: This benchmark exercise was not designed to drive maximum IOPs or maximum HammerDB TPC-C TPM consistently and it is not a reference architecture paper. In production environments, VMs consistently running at 100% vCPU utilization would trigger alarms at the host and VM levels.

About the Author

Dominic Cheah is a Technical Marketing Engineer on the Solutions and Alliances group at Datrium, Inc. In his role, Dominic focuses on business-critical applications, DevOps, databases, and data protection.