



VMware Virtual SAN™ Performance with Microsoft SQL Server™

Performance Study

TECHNICAL WHITE PAPER

Table of Contents

Executive Summary	3
Introduction.....	3
Experiment Setup.....	3
Hardware Configuration.....	3
ESXi/Virtual SAN Configuration.....	3
Workload/Virtual Machine Configuration	4
Test Methodology	5
Results and Analysis.....	5
Workload Scaling with Virtual SAN.....	5
Conclusion.....	7
Appendix A. Hardware Configuration.....	7
Appendix B. Virtual Machine Configuration.....	8
Appendix C. Workload Configuration.....	8
References.....	8

Executive Summary

Microsoft SQL Server workload performance on VMware Virtual SAN is studied in this paper. The DVD Store benchmark [1] on Microsoft SQL Server performs well when using Virtual SAN storage. The workload scales linearly with respect to the number of virtual machines per node in the Virtual SAN cluster and also with respect to the number of nodes in the Virtual SAN cluster. Experiments show the Virtual SAN storage sub-system is never the bottleneck and the workload saturates the host CPU while the I/O latency remains constant. The workload also exhibited a stable and consistent performance on Virtual SAN storage. The variability over time is minimal. An aggregated “orders per minute” of 77,206 across 12 DVD Store instances and 148,099 across 24 DVD Store instances, was achieved in a 4-node and an 8-node Virtual SAN clusters respectively.

These tests show that VMware vSphere 5.5 with Virtual SAN is an effective storage platform for transactional database workloads like Microsoft SQL Server.

Introduction

Virtual SAN is a distributed layer of software that runs natively as part of the ESXi hypervisor. Virtual SAN aggregates local or direct-attached storage disks in a host cluster and creates a single storage pool that is shared across all hosts of the cluster. Virtual SAN eliminates the need for external shared storage and simplifies storage configuration and virtual machine provisioning operations. In addition, Virtual SAN supports VMware features that require shared storage such as high availability (HA), vMotion, and distributed resource scheduling (DRS) for failover. More information on Virtual SAN can be obtained from the [Virtual SAN Design and Sizing Guide](#) [2].

This paper demonstrates the performance of Microsoft SQL Server with VMware Virtual SAN on VMware vSphere 5.5. This paper focuses on one of the key aspects in Virtual SAN performance, namely workload scale-out of a widely-used transactional database benchmark, DVD Store.

Note: Hosts in a Virtual SAN cluster are also called nodes. The terms “host” and “node” are used interchangeably in this paper.

Experiment Setup

Hardware Configuration

“[Appendix A](#)” provides detailed hardware configuration. Experiments used 4 or 8 hosts depending on the intended size of Virtual SAN cluster. Each of these hosts contained a dual-socket Intel Xeon Processor E5-2650 v2 (16 cores, 32 threads, @2.6GHz), 128GB memory, an LSI 9207-8i controller hosting one 400GB Intel S3700 SSD, and seven 1.1TB, 10K RPM SAS drives. The Virtual SAN traffic was configured to use a 10GbE network (over a 10GbE switch).

ESXi/Virtual SAN Configuration

VMware vSphere 5.5 U2 was used in this study. Two setups were used for the experiments: a 4-node Virtual SAN cluster and an 8-node Virtual SAN cluster. Both of these setups had homogenous hardware and software configurations. The workload virtual machine was entirely stored on the Virtual SAN datastore. The Virtual SAN datastore was configured with the default policy of HostFailuresToTolerate=1 and StripeWidth=1.

Workload/Virtual Machine Configuration

The open-source DVD Store version 2.0 was used as the benchmark tool for this performance study. DVD Store simulates an online ecommerce DVD store, where customers log in, browse, and order products. The benchmark tool is designed to utilize a number of advanced database features, including transactions, stored procedures, triggers, and referential integrity. The primary performance metric of DVD Store is orders per minute (OPM).

The DVD Store benchmark driver outputs a moving average of orders per minute and also a cumulative number of transactions every 10 seconds. Because a moving average result would not show fine-grained changes in performance during scenarios like Virtual SAN failover, the absolute orders per minute were computed from the cumulative number of transactions that was printed every 10 seconds.

The entire DVD Store benchmark tools, including the query generator and the database backend, were encapsulated in a single virtual machine, which ran the Microsoft Windows Server 2008 R2 operating system and Microsoft SQL Server 2008. The virtual machine was configured with 4 virtual CPUs (vCPUs), and 4GB of memory. The virtual machine was configured with 3 virtual disks: a 50GB boot disk containing Windows Server 2008 R2 and Microsoft SQL Server 2008, a 200GB database disk, and a 10GB log disk. Figure 1 illustrates the experiment setup in detail.

The DVD Store workload used a database size of 100GB with 200 million customers and 10 million products. A more detailed workload configuration can be found in “Appendix C.”

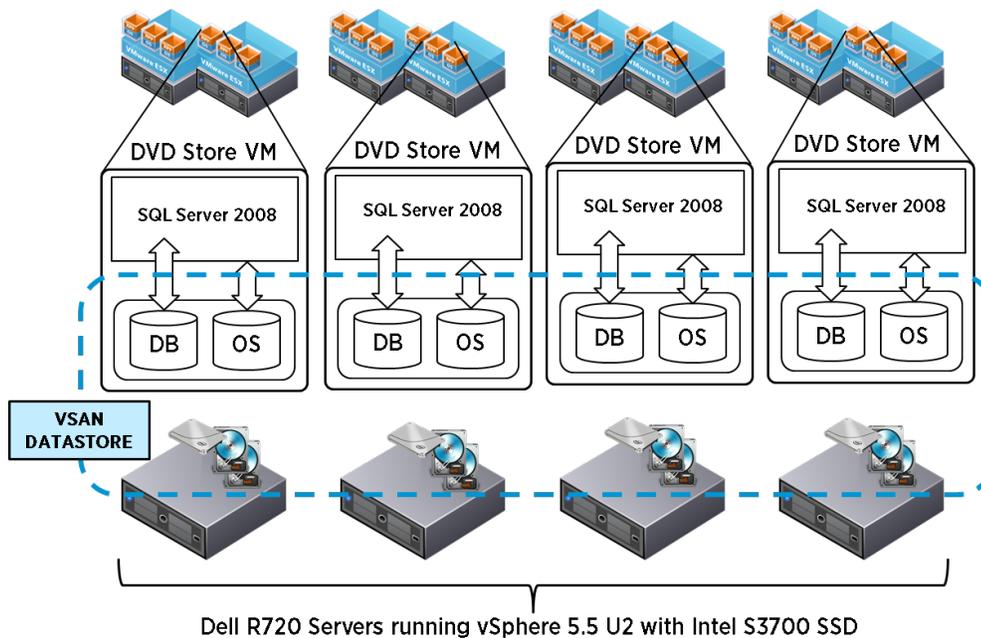


Figure 1. Experiment setup

Test Methodology

The benchmark DVD Store was run for 3 hours for every test case. After the initial 20-30 minutes of the experiment, the “Orders Per Minute” metric stabilized throughout the entire run. In order to stress the Virtual SAN storage, the number of virtual machines was increased per Virtual SAN node (a host) stepwise from 1 to 3. This set of experiments provides an idea of the impact of the working set size on the overall performance. Experiments were also conducted to measure the scalability of Virtual SAN storage with respect to increasing the number of nodes in the Virtual SAN cluster. Numbers were reported from a 4-node cluster and an 8-node cluster.

Because the amount of data that is cached in the SSD is vital to achieve good performance, it was ensured that the results across runs were reliable by flushing out the Virtual SAN cache (SSD contents) at the end of every run. Because of this, every experiment went through an initial warm-up phase where the data was staged to cache and then stabilized over time. To avoid the impact of database caches, the virtual machine was power cycled between runs. The entire database was also restored from a backup before every run in order to achieve the same initial state for every run.

Results and Analysis

Workload Scaling with Virtual SAN

This section presents the results of a scalability study with the DVD Store workload on Virtual SAN. Two cases were examined:

- The number of nodes in the Virtual SAN cluster were scaled up from 4 to 8.
- The number of DVD Store instances (virtual machines) per node were scaled up in a Virtual SAN cluster from 1 to 3.

Figures 2 and 3 show the DVD Store benchmark performance in terms of “Orders Per Minute” application metric, which is a measure of application throughput, for 4- and 8-node Virtual SAN clusters. The graphs show that the application performance on Virtual SAN is stable over time irrespective of the I/O load handled by the Virtual SAN storage.

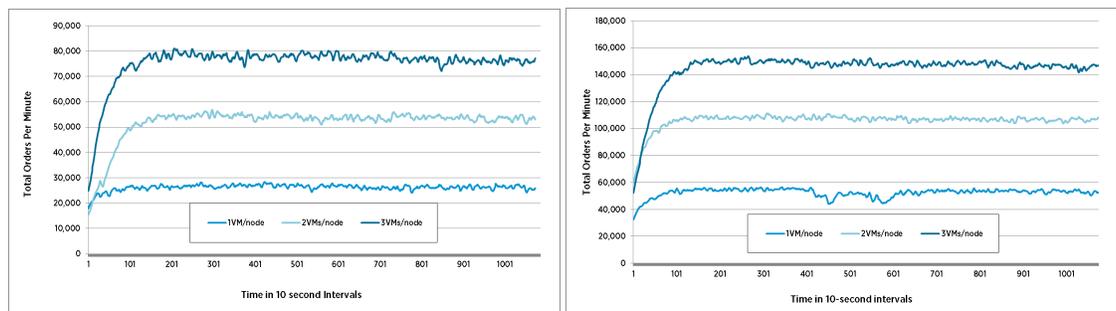


Figure 2. DVD Store performance on 4-node VSAN cluster Figure 3. DVD Store performance on 8-node VSAN cluster

Figures 4 and 5 show the linear scaling of application performance in Virtual SAN with respect to an increasing number of DVD Store instances per Virtual SAN node. According to the results, the application scales linearly up to 3 virtual machines per node.

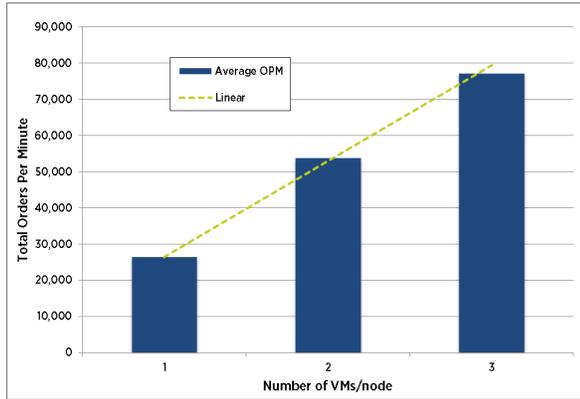


Figure 4. Scaling wrt. number of VMs/node – 4 node VSAN cluster




Figure 5. Scaling wrt. number of VMs/node – 4 node VSAN cluster

The experiments were conducted up to 3 DVD Store virtual machines per node and not more because at 3 virtual machines per node, the total CPU utilization across the 16 physical cores in every node was close to 100%. This means that Virtual SAN is able to handle the DVD Store I/Os comfortably until the point when the application saturates its CPU and can't scale further. Figures 6, 7 and 8 show the per-node CPU usage patterns for different numbers of virtual machines on a 4-node Virtual SAN cluster. Note that nodes in the Virtual SAN cluster were homogenous and each of them had 16 physical cores. The values plotted in these graphs are the sum of "CPU Used %" from each physical core.

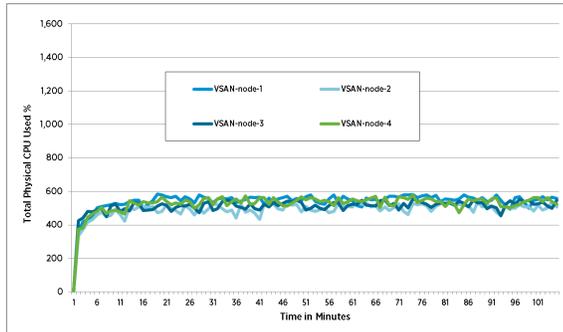


Figure 6. Total CPU used – 1VM/node

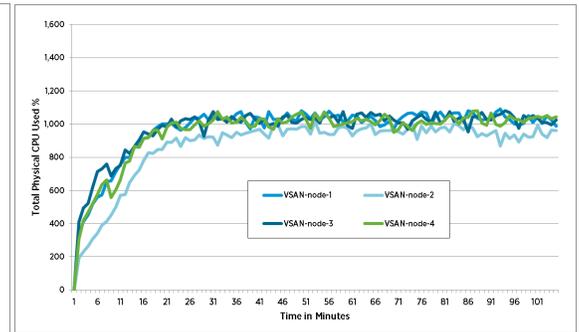


Figure 7. Total CPU used – 2VMs/node

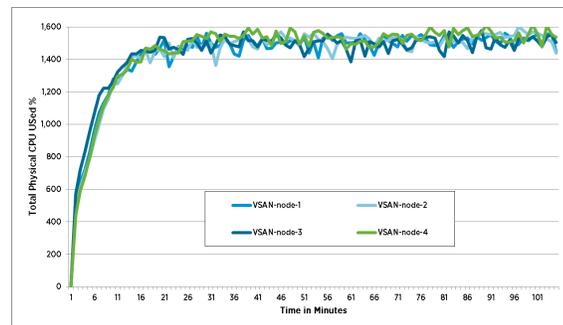


Figure 8. Total CPU used – 3VMs/node

DVD Store on Virtual SAN also scales linearly with respect to increasing the number of nodes in the Virtual SAN cluster. To illustrate this, we ran the DVD Store benchmark on two different Virtual SAN setups: with 4 nodes and 8 nodes. Figure 9 shows the “Orders Per Minute” metric per DVD Store virtual machine on the primary y-axis and the total “Orders Per Minute” across all the virtual machines on the secondary y-axis for these two configurations. The graph shows that the application scales almost linearly from a 4- to an 8-node Virtual SAN cluster.

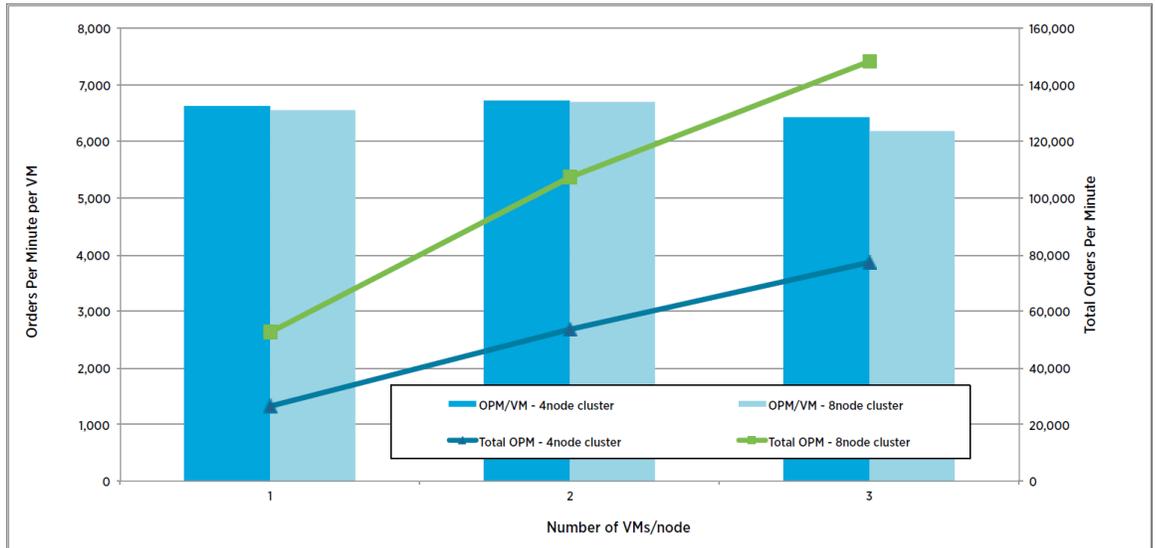


Figure 9. Workload scaling wrt. number of nodes in VSAN cluster

Conclusion

VMware performance labs ran tests to see how well vSphere 5.5 with Virtual SAN handled a workload that simulated a typical transactional database application workload featuring Microsoft SQL Server database. Tests show the DVD Store benchmark with SQL Server performs well when using Virtual SAN. Virtual SAN can scale out linearly, with respect to performance, when adding more virtual machines per node in the Virtual SAN cluster and also by increasing the number of nodes in the Virtual SAN cluster. Experiments show that the Virtual SAN storage sub-system is never the bottleneck when scaling and the workload saturated the host CPU while the I/O latency remained constant.

Appendix A. Hardware Configuration

We used 4 or 8 Dell R720 servers, and each server had the following configuration:

- Dual-socket Intel® Xeon® Processor E5-2650 v2 @ 2.60GHz/16 cores/32 threads. Hyper-Threading is on.
- 128 GB RAM. DDR3 @1866MHz
- 1x Intel S3700 SSD. (INTEL SSDSC2BA40, Firmware: DL04)
- 4x Hitachi 1.1TB, 10K RPM SAS drives (HUCI01212CSS600, Firmware: U5E0)
- 3x Seagate 1.1TB, 10K RPM SAS drives (ST1200MM0007, Firmware: IS04)
- LSI Logic 9207-8i controller (LSI2308_2 chipset)
- 1x Quad-Port Broadcom 1GbE NIC (BCM5720)
- 1x Dual-Port Intel 10GbE NIC (82599EB, fibre optic connector)

One 1GbE port on each machine was connected to an Extreme X350 switch and one 10GbE port of each machine was connected to an Arista 7050 switch. The rest of the NICs were not used in our experiments.

Appendix B. Virtual Machine Configuration

The virtual machines were homogenous and had the following configuration:

- 64-bit Microsoft Windows Server 2008 R2
- VMXNET3 driver version 1.1.30.0, PVSCSI driver version 1.1.1.0
- 50GB disk with operating system on LSI Logic controller
- 200GB database disk and 10GB log disk, both on PVSCSI controller
- Microsoft SQL Server 2008

Appendix C. Workload Configuration

The DVD Store benchmark was configured as follows for all the experiments:

```
target=localhost
n_threads=4
ramp_rate=100
run_time=90
db_size=100GB
warmup_time=2
think_time=0.010
pct_newcustomers=20
n_searches=15
search_batch_size=15
n_line_items=15
virt_dir=ds2
page_type=php
windows_perf_host=
linux_perf_host=
detailed_view=n
```

References

- [1] Dave Jaffe and Todd Muirhead. (2011, December) DVD Store Benchmark.
<http://en.community.dell.com/techcenter/extras/w/wiki/dvd-store.aspx>
- [2] VMware, Inc. (2014, March) Virtual SAN Design and Sizing Guide.
http://www.vmware.com/files/pdf/products/vsan/VSAN_Design_and_Sizing_Guide.pdf

About the Author

Dr. Sankaran Sivathanu is a staff engineer in the VMware Performance Engineering team. His work focuses on the performance aspects of the ESXi storage stack and characterization/modeling of new and emerging I/O workloads. He has a PhD in Computer Science from the Georgia Institute of Technology.

Acknowledgements

The author thanks Julie Brodeur, Lenin Singaravelu, Priti Mishra, and Shilpi Agarwal from the VMware performance team for their valuable suggestions and feedback to improve this paper. The author also thanks Todd Muirhead, one of the authors of the DVD Store benchmark tool, for his help in configuring and running DVD Store.

