White Paper

Microsoft SQL Server 2012 on Cisco UCS with iSCSI-Based Storage Access in VMware ESX Virtualization Environment: Performance Study

Contents

What You Will Learn

The document provides detailed information about performance tests conducted and the results obtained on a single virtual machine running VMware ESX 5.0 hypervisor with the virtual machine running Microsoft SQL Server 2012. The test used the standard data warehouse (TPC-H) test suite to load test the SQL databases. Data warehouse access patterns are typically sequential and read intensive.

VMware ESX 5.0 hypervisor was booted using Small Computer System Interface over IP (iSCSI) on the Cisco UCS® B200 M2 Blade Server with the Cisco UCS M81KR Virtual Interface Card (VIC) providing I/O virtualization on the Cisco Unified Computing System™ (Cisco UCS). The virtual machine was installed with guest OS Microsoft Windows 2008 R2 SP1 using its native iSCSI software-based initiator to access the shared NetApp clustered iSCSI storage system with the native VMware ESX vSwitch and the Cisco® Data Center Virtual Machine Fabric Extender (VM-FEX) distributed virtual switch (DVS) network infrastructure.

Microsoft SQL Server 2012 was installed on the guest Microsoft Windows 2008 R2 virtual machine. Microsoft SQL Server 2012 was configured with a single database consisting of two data files and one log file for storing data, indexes, and log information on iSCSI logical unit numbers (LUNs) exposed over the iSCSI network.

The NetApp clustered iSCSI storage system was provisioned to host Microsoft SQL Server 2012 database data and log LUNs on separate controllers with dedicated disks on aggregate and flex volumes.

The three main areas of interest in the performance study were:

- Capability to saturate the 10-Gbps Cisco VIC with less than 20 ms of disk latency over iSCSI in a VMware virtualization environment with a sequential read-intensive workload on a single-database Microsoft SQL Server 2012
- Capability to achieve higher throughput and lower latency with the Cisco Data Center VM-FEX DVS compared to the VMware ESX native vSwitch using a Microsoft Windows 2008 R2 guest virtual machine with its native iSCSI initiator
- Capability to achieve lower application response time with a Cisco Data Center VM-FEX switch compared to a VMware ESX native vSwitch on a single-database Microsoft SQL Server 2012 hosted on a Microsoft Windows 2008 R2 guest virtual machine over an iSCSI network

Based on the test results, these conclusions were reached:

- Decision support system (DSS) workloads with read-intensive sequential query access saturated the Cisco VIC, achieving more than 10 Gbps and disk latency of less than 17 ms on a single-database Microsoft SQL Server 2012 performing data warehousing in a virtual environment.
- Processing exceeded the VMware ESX vSwitch, with a 30 percent gain in I/O throughput and 43 percent reduction in disk latency with the Cisco Data Center VM-FEX DVS network infrastructure running a DSS work load on a single-database Microsoft SQL Server 2012 in a virtual environment.
- Queries running on a single-database Microsoft SQL Server 2012 showed 12 percent response time reduction with the Cisco Data Center VM-FEX DVS compared to the VMware ESX vSwitch network infrastructure in a virtual environment.

Test Setup

This section provides the configuration details and the methods used to run the tests.

Test Application

Benchmark Factory for Databases is a database performance testing tool that allows you to conduct database workload replay, industry-standard benchmark testing, and scalability testing. This test used the TPC-H profile to run a read-intensive DSS-type workload for throughput and latency analysis. The TPC-H workload was designed to mimic impromptu reports and business analytics running over a large data set with a periodic refresh of the data set. The data set size was 100 GB (with a scale factor of 100 GB), with the database size close to 150 GB. Only the Power stream was run to measure the throughput and response time for read-intensive queries.

Test Environment

Table 1 lists the hardware components and software products used for the performance study.

Table 1. Hardware and Software Details

Test Components	Component Details
Test server	• Model: Cisco UCS B-Series B200 M2 • Processors: Dual-core Intel Xeon processor 5600 series • RAM: 192 GB • Storage: iSCSI boot • Network: Cisco 10-Gbps VIC
Test virtual machine	• CPU: 16 virtual CPUs (vCPU) Memory: 32 GB Guest OS: Windows 2008 R2 SP2 Application: Microsoft SQL 2012 Single Server • Virtual network adapter: VMXNET3 1.2.22.0 or later • iSCSI initiator: Guest-based software iSCSI initiator
iSCSI storage	• Model: NetApp FAS3270 cluster storage system • Disks: Total of 53 SAS 1500-rpm disks • Performance acceleration module (PAM) capacity: 512 GB • iSCSI target: Dual 10 GB network adapter • NetApp Data ONTAP version: 8.0.1
Microsoft SQL Server version	Microsoft SQL Server 2012 Enterprise Edition (64-bit)
Virtualization product version	VMware ESX 5.0

The Cisco UCS B-200 M2 Blade Server was used as the test server, installed with VMware ESX 5.0. The VMware ESX host used a Microsoft Windows 2008 R2 guest virtual machine with Microsoft SQL Server 2012 installed on the virtual machine. Benchmark Factory was installed on the client machine to run the workload. NetApp FAS3270 iSCSI storage was used to meet all the storage requirements, including the LUNs for VMware ESX 5.0 host iSCSI boot and Microsoft SQL Server 2012 database data and log file placement. The test server and the storage were connected over a pair of Cisco Nexus® 5548UP Switches using a virtual PortChannel (vPC) fabric on 10-GB interfaces with the iSCSI protocol. The connectivity between the server and the storage were configured with end-to-end quality of service (QoS) and a maximum transmission unit (MTU) setting to achieve a high throughput and lower latency network.

Figure 1 shows the high-level physical and logical topology with the Cisco UCS, Cisco Nexus 5548UP, and NetApp iSCSI storage system in an iSCSI network design.

Figure 1. Physical and Logical Topology

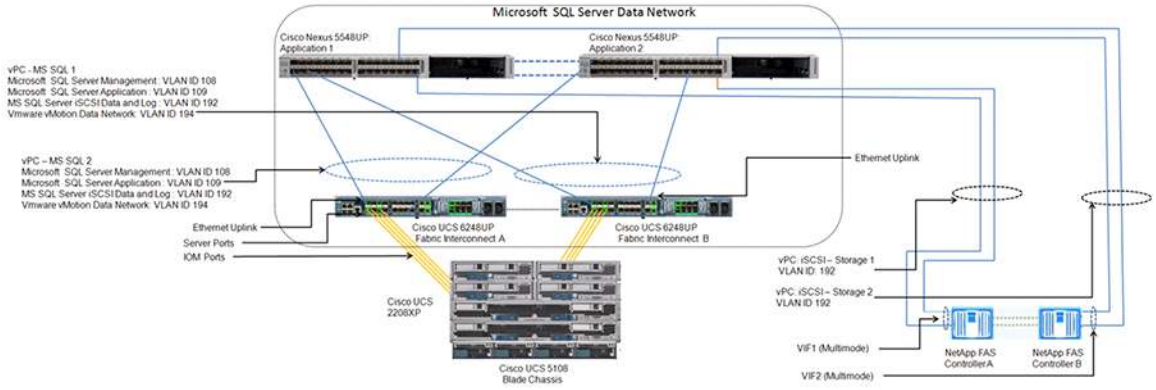

Figure 2 and Figure 3 show the physical and logical views of the test setup used for benchmarking run to compare iSCSI throughput, latency, and response-time results between the VMware ESX vSwitch and Cisco Data Center VM-FEX network infrastructure on Cisco UCS and NetApp storage systems.

Figure 2. Guest-Based iSCSI Initiator on Cisco Data Center VM-FEX DVS

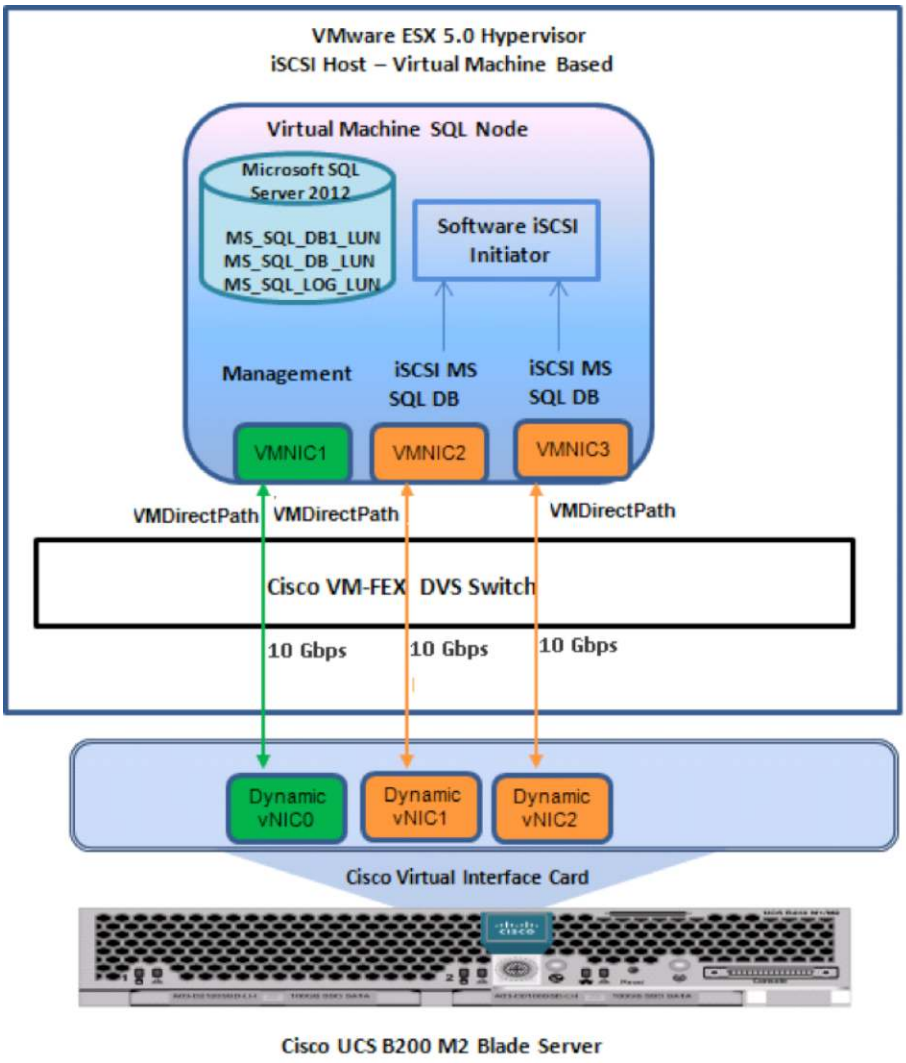

Figure 3. Guest-Based iSCSI Initiator on VMware ESX vSwitch

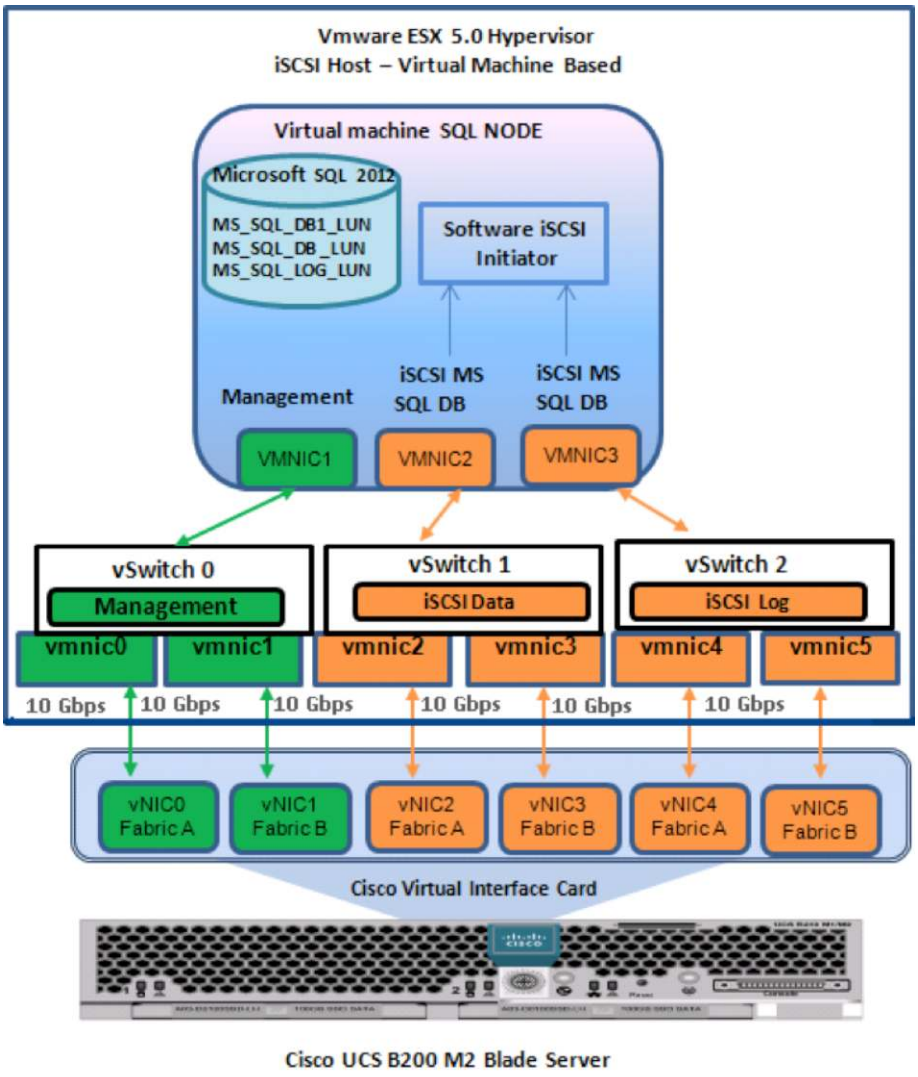

Virtual Machine

Microsoft Windows Server 2008 R2 Enterprise 64-bit was used as the guest OS running Microsoft SQL Server 2012. A guest-based software iSCSI initiator with Microsoft Multipath I/O (MPIO) multipathing was configured to access NetApp iSCSI targets for storing Microsoft SQL Server 2012 database data and log files on LUNs exposed over iSCSI. Management and iSCSI storage network traffic was accessed using dedicated VMXNET3 virtual network adapters as shown in Figure 2 and Figure 3.

The virtual machine had 16 vCPUs and 32 GB of memory reserved for the operating system and the Microsoft SQL Server 2012 single-instance database.

Table 2 lists VMXNET3 driver parameters that were tuned during the performance study for the guest-based software iSCSI initiator with the VMware ESX vSwitch or the Cisco Data Center VM-FEX DVS.

Table 2. VNXNET3 Driver Parameters

Parameter	Value
RSS	Enabled
Jumbo	9000
TCP checksum offload (IPv4)	Enabled

VMware ESX Software Configuration

VMware ESX 5.0 was iSCSI booted on the test server with the default virtual network interface card (vNIC) and adapter configuration setting. For VMware ESX vSwitch–based infrastructure, a separate vSwitch with the appropriate Cisco UCS static vNIC uplinks was created. For the Cisco Data Center VM-FEX DVS, appropriate Cisco Data Center VM-FEX port profiles were applied on the Cisco UCS dynamic vNICs to carry management and Microsoft SQL Server 2012 data and log traffic using iSCSI.

To configure end-to-end jumbo frames with a MTU of 9000 from the test server to the storage path on Cisco UCS, an appropriate QoS class with MTU 9000 policy was defined and applied to the static and dynamic vNICs of the VMware ESX vSwitch and Cisco Data Center VM-FEX network infrastructure. Similar matching QoS class and policy-map values were defined in the Cisco Nexus 5548UP vPC domain, and for NetApp iSCSI target interfaces the Cisco Nexus 5548UP tagged QoS and MTU values by untagging the class-of-service (CoS) values.

Table 3 lists the static and dynamic vNICs adapter policy parameters that were tuned during the performance study for the guest-based software iSCSI initiator with the VMware ESX vSwitch and the Cisco Data Center VM-FEX DVS.

Table 3. Static and Dynamic vNICs Parameters

Parameter	Value
Transmit queues	1
Ring size	256
Receive queues	4
Ring Size	512
Completion queues	28
Interrupts	7

Storage Layout

The Microsoft Windows 2008 R2 guest virtual machine was installed with Microsoft SQL Server 2012. It had two LUNs with 100-GB capacity from each NetApp iSCSI controller for storing Microsoft SQL Server 2012 database data files, and one LUN with 100-GB capacity on a single NetApp iSCSI controller for storing the Microsoft SQL Server 2012 database log file. The LUNs were created on a dedicated aggregation layer configured with RAID-DP striped on SAS disk drives on each NetApp cluster storage system, as shown in Figure 4.

Figure 4. SAS Disks on NetApp Storage Controllers

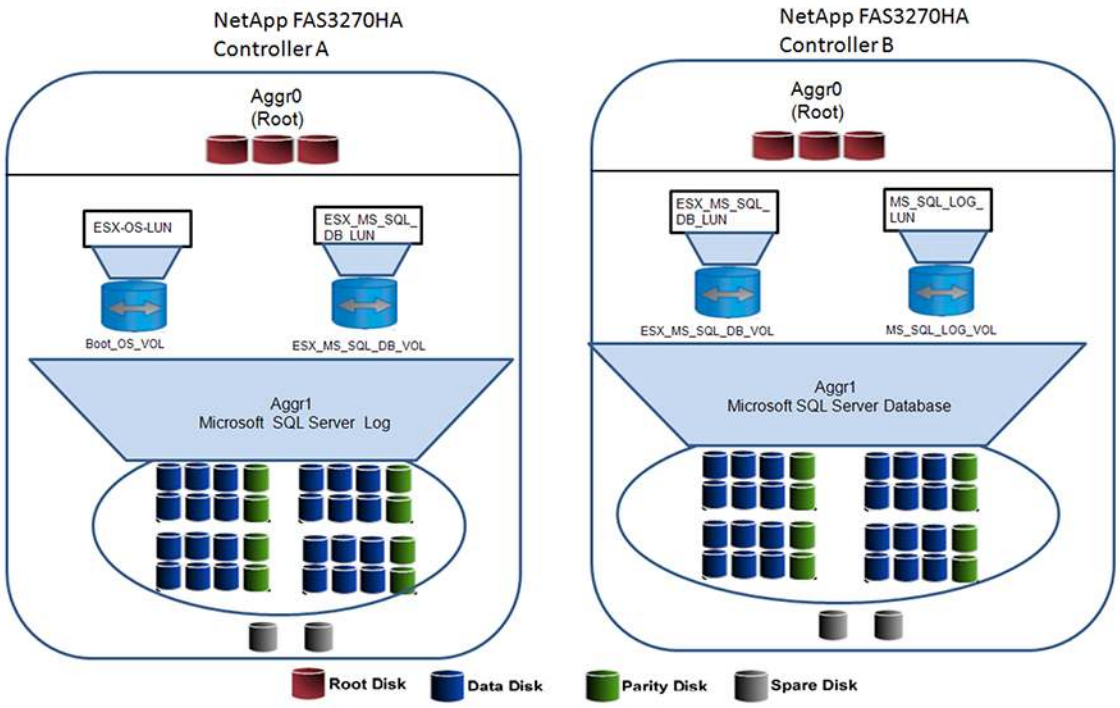

The NetApp iSCSI target was configured with virtual interface–enabled mutimode (Link Aggregation Control Protocol [LACP]) cluster failover on each NetApp controller over dual 10 Gigabit Ethernet interfaces, which were part of the Cisco Nexus 5548UP vPC domain.

In the test, the Microsoft Windows guest-based software iSCSI initiator was used to access the iSCSI storage target in the VMware ESX vSwitch or Cisco Data Center VM-FEX network environment, as shown in Figure 5.

Figure 5. vPC Domain for Guest-Based Software iSCSI Initiator

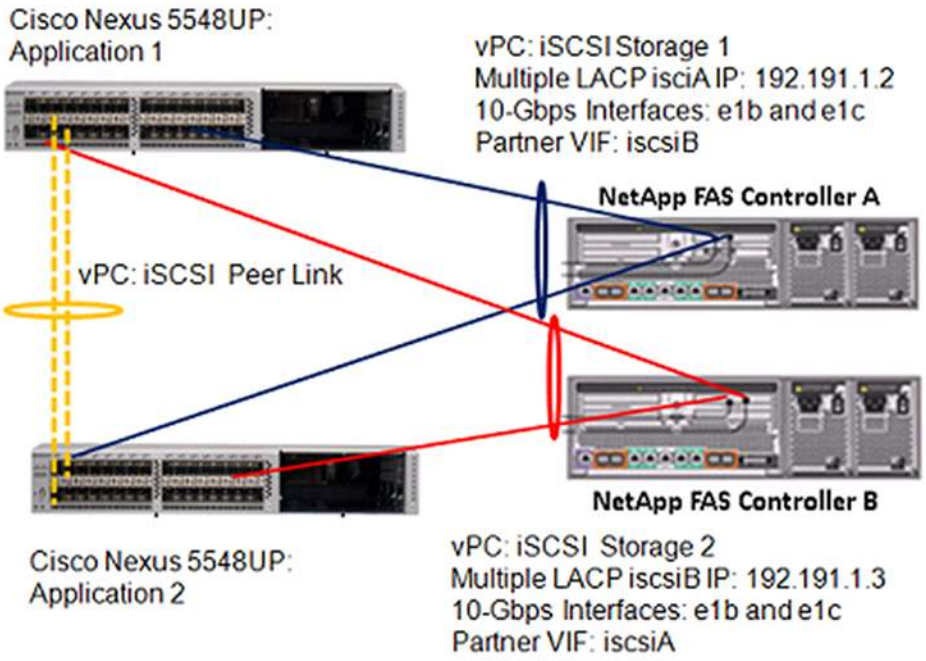

Results

This section discusses the results of the performance tests. Metrics measured achieved storage bandwidth and throughput (MBps), disk latency (ms), and query response time for read-intensive queries.

The iSCSI NetApp storage systems have a 512-GB PAM flash disk that caches the entire database data, eliminating the need for read-intensive queries that depend on disk performance.

Storage Bandwidth Saturation Through iSCSI Network

With a query that has a read-intensive and sequential access pattern, the bandwidth of one 10-Gbps link is saturated on the Cisco VIC with throughput close to 1200 MBps and disk latency of 17 ms in the Cisco Data Center VM-FEX DVS virtual environment on the test server, as shown in Figure 6.

Figure 6. Data Warehousing Operation on Cisco Data Center VM-FEX DVS

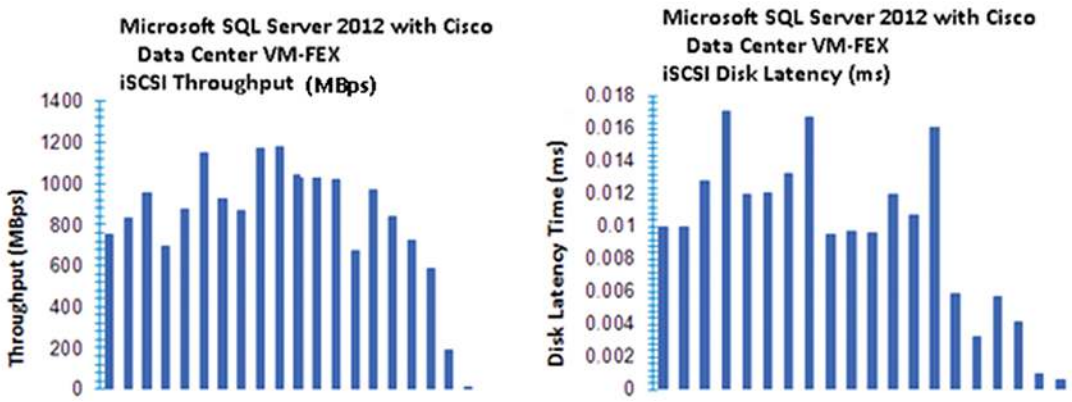

Comparison of iSCSI Network Throughput and Disk Latency with Cisco Data Center VM-FEX and VMware vSwitch

Microsoft SQL Server 2012 derived 30 percent more bandwidth and experienced 43 percent faster response time with the Cisco Data Center VM-FEX DVS with VMDirectPath IO Mode compared to the VMware ESX vSwitch with network emulation infrastructure, as shown in Figure 7.

Figure 7. Database Bandwidth Throughput and Disk Latency

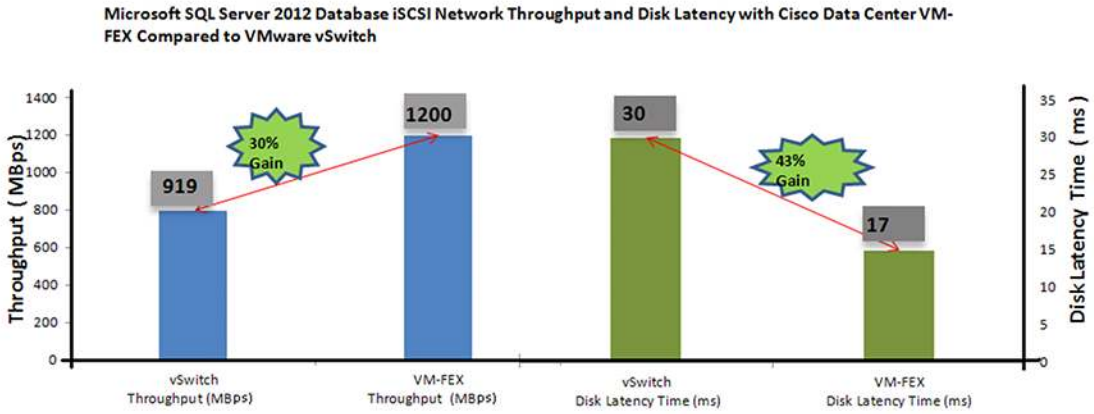

Microsoft SQL Server 2012 Database Query Response Time with iSCSI Storage

In this performance test, a read-intensive TPC-H query response time showed 12 percent gain with the Cisco Data Center VM-FEX DVS compared to the VMware ESX vSwitch network infrastructure, as shown in Figure 8.

Figure 8. Database iSCSI Query Response Time

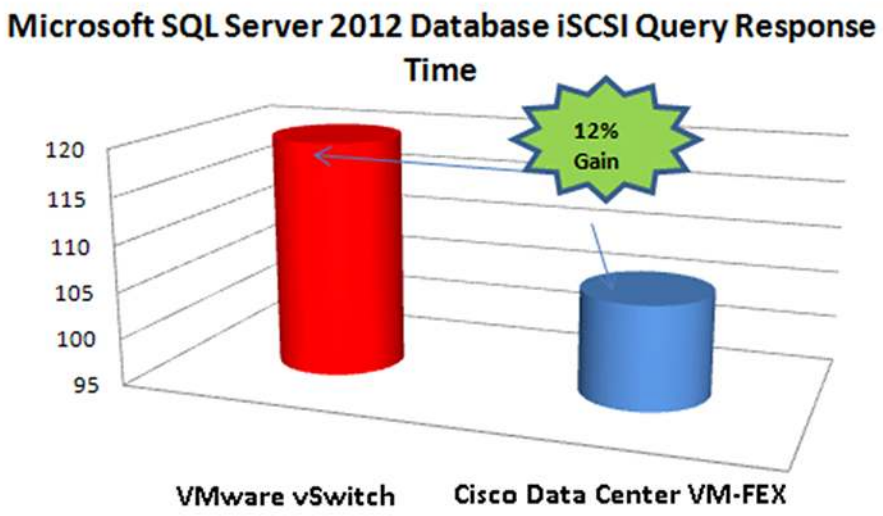

Conclusion

The performance tests show that iSCSI is a scalable protocol for running a data warehouse on Microsoft SQL Server 2012 in a VMware virtual environment. The study shows that with a 10-GB Cisco virtual interface adapter and Cisco Data Center VM-FEX DVS with VMDirectPath IO Mode technology, you can achieve higher throughput and lower latency with fewer resources on virtual machines. The combination of the Cisco Nexus 5548UP vPC fabric, Cisco UCS, Cisco VIC adapters, and NetApp iSCSI storage can be used to provide end-to-end high availability and network-controlled policy over a 10-GB iSCSI network.

For More Information

The documents listed here provide additional information relevant to implementation of Microsoft SQL Server 2012 on the VMware ESX 5.0 hypervisor with the NetApp iSCSI storage system on Cisco UCS B-Series servers.

- Microsoft SQL Server 2012 failover cluster on Cisco UCS with iSCSI-based storage access deployment guide: http://www.cisco.com/en/US/prod/collateral/ps10265/ps10280/guide_c07-707705_1.pdf
- Microsoft SQL Server 2012 installation guide: http://msdn.microsoft.com/en-us/library/bb500469%28v=sql.110%29.aspx
- Cisco Nexus QoS switch configuration guide: http://www.cisco.com/en/US/docs/switches/datacenter/nexus5000/sw/qos/Cisco_Nexus_5000_Series_NX-OS_Quality_of_Service_Configuration_Guide_chapter3.html#con_1150612
- Cisco Data Center VM-FEX configuration guide: http://www.cisco.com/en/US/docs/unified_computing/ucs/sw/vm_fex/vmware/gui/config_guide/b_GUI_VMware_VM-FEX_UCSM_Configuration_Guide.pdf

- Cisco Data Center VM-FEX best practices: http://www.cisco.com/en/US/solutions/collateral/ns340/ns517/ns224/ns944/vm_fex_best_practices_deployment_guide_ps10277_Products_White_Paper.html
- Cisco Data Center VM-FEX troubleshooting guide: http://www.cisco.com/en/US/solutions/collateral/ns340/ns517/ns224/ns944/basic_troubleshooting_vm_fex.html
- Cisco UCS System hardware and software interoperability matrix: http://www.cisco.com/en/US/docs/unified_computing/ucs/interoperability/matrix/r_hcl_B_rel2_0.pdf
- VMware vSphere networking with VMware ESXi 5.0: http://pubs.vmware.com/vsphere-50/topic/com.vmware.ICbase/PDF/vsphere-esxi-vcenter-server-50-networking-guide.pdf

Americas Headquarters
Cisco Systems, Inc.
San Jose, CA

Asia Pacific Headquarters
Cisco Systems (USA) Pte. Ltd.
Singapore

Europe Headquarters
Cisco Systems International BV Amsterdam,
The Netherlands

Cisco has more than 200 offices worldwide. Addresses, phone numbers, and fax numbers are listed on the Cisco Website at **www.cisco.com/go/offices.**

Printed in USA C11-711915-00 08/12