



Technical Report

Microsoft Windows Server 2012 Hyper-V Storage Performance: Measuring SMB 3.0, iSCSI, and FC Protocols

Dan Chilton, NetApp
May 2013 | TR-4175

Abstract

With Windows Server® 2012 Hyper-V™ and the NetApp® clustered Data ONTAP® 8.2 architecture, customers have choices in shared storage design for building virtualized data centers. In addition to traditional block protocols like Fibre Channel (FC) and iSCSI, Server Message Block (SMB) 3.0 brings file protocol storage access for the first time to data centers built with Windows Server 2012 Hyper-V on NetApp. NetApp tested the performance of a virtual server workload over 10Gb SMB 3, 10Gb iSCSI, and 8Gb FC. The results of these tests including throughput and latency are detailed in this report. With this information customers can have confidence in the performance of Windows Server 2012 Hyper-V on NetApp.

TABLE OF CONTENTS

1	Introduction	3
1.1	Windows Server 2012 Hyper-V on NetApp	3
1.2	Windows Server 2012 Hyper-V on NetApp: Testing Performance Across Protocols	4
2	Executive Summary	4
2.1	Test Workload	5
2.2	Performance Summary	5
3	Performance Comparison Details	6
3.1	Relative Throughput Comparison	6
3.2	Relative Latency Comparison	6
4	Test Design and Configuration	7
4.1	Hardware and Software	7
4.2	Network Connections	8
4.3	Storage Provisioning Design FC and iSCSI	9
4.4	Storage Provisioning Design SMB 3.0	10
4.5	Building the Windows Server 2012 Hyper-V Data Center	11
4.6	Workload and Iometer Test Practices	11
4.7	Additional Configuration Details	12
5	Conclusion	12
	References	13
	Acknowledgements	13

LIST OF TABLES

Table 1)	Hardware and software for Windows Server 2012 Hyper-V data center.	7
Table 2)	Hardware and software for clustered Data ONTAP 8.2.	8
Table 3)	Iometer workload generation controller server.	8
Table 4)	FAS6240 aggregate, FlexVol volume, and LUNs for testing block protocols (iSCSI and FC).	10
Table 5)	FAS6240 aggregate, FlexVol, and LUNs for testing SMB 3.0 protocol.	10
Table 6)	Virtual machine components.	11
Table 7)	Outstanding I/Os per VM and the total outstanding I/Os.	12

LIST OF FIGURES

Figure 1)	Relative throughput in IOPs across the three protocols at low, medium, and heavy workload levels.	6
Figure 2)	Relative latency across the three protocols at low, medium, and heavy workload levels.	6
Figure 3)	FC, iSCSI, and SMB 3.0 connections between Windows Server 2012 Hyper-V servers and NetApp shared storage.	9

1 Introduction

As virtualized data centers and private clouds have grown from the margins to the norm during the past few years, companies have embraced this new paradigm. Many companies have reached a point at which virtualized servers are the default and physical servers must be justified. This has brought cost-reduction benefits to corporate IT budgets by reducing data center footprint, power consumption, and server OS licensing. Virtualization at the server, network, and storage layer has also brought about a dramatic increase in application deployment speed while adding flexibility and agility. Microsoft and NetApp are leading this paradigm shift with powerful server virtualization and storage solution technologies that customers can leverage to create private clouds. This report compares the performance of a typical virtual server workload deployed on Microsoft® Windows Server 2012 with Hyper-V and clustered Data ONTAP 8.2 shared storage across various protocols (SMB 3.0, iSCSI, and FC).

1.1 Windows Server 2012 Hyper-V on NetApp

In the past, only block protocols were available for Microsoft virtualization customers. Flash forward to today and the SAN barrier has been broken. With the release of Microsoft Windows Server 2012 and SMB 3.0, customers can deploy Hyper-V on NAS, SAN, or both. When Windows Server 2012 Hyper-V is deployed on clustered Data ONTAP 8.2 storage solutions, the SMB 3.0, iSCSI, FCOE, and FC protocols can all be leveraged on the same unified platform. This allows customers to design the virtualized data centers with standard 10GbE networking for iSCSI, and SMB 3.0 storage connections or 10GbE for FCOE and 8Gb FC for Fibre Channel SAN configurations. The Microsoft Windows and Data ONTAP technologies that customers should consider for new virtualized data centers and private clouds are discussed in the sections that follow.

Microsoft Windows Server 2012 with Hyper-V

With Microsoft Windows Server 2012 with Hyper-V role enabled, Microsoft has built the combination of a well-performing hypervisor with enterprise resiliency for small to large deployments. Performance enhancements include increased size and scale with up to 64 nodes in a single Windows Server 2012 Hyper-V cluster and up to 1,024 virtual machines per Windows Server 2012 Hyper-V host. To increase resiliency, Microsoft dramatically improved the SMB protocol, moving from SMB 2.1 to SMB 3.0. Key features include persistent file handles (continuously available file shares), witness protocol, and clustered client failover. These features, when combined with the right shared storage solution, provide the high availability of a block protocol with the simplified administration and reduced deployment cost of network file share access. SMB 3.0 is ready for prime time as an enterprise storage protocol for Hyper-V and can also be leveraged for Microsoft SQL Server®. However, if your data center is focused on SAN, iSCSI, FCoE, and FC are still available with all of the performance and multipath resiliency that can be expected from a Microsoft Windows Server solution deployed on NetApp clustered Data ONTAP. Virtual Fibre Channel (vFCP) from the guest virtual machine to the storage controllers is new with 2012 and is fully supported for a Windows Server 2012 Hyper-V on NetApp shared solution as well.

NetApp Clustered Data ONTAP 8.2

With clustered Data ONTAP 8.2, NetApp provides a unique, highly available shared storage solution for Hyper-V virtualized data centers. All clustered Data ONTAP solutions are unified, meaning that NAS protocols like SMB 3.0 and NFS as well as SAN protocols like iSCSI and FC are spoken natively by the purpose-built storage operating system. In 2012, the NetApp Data ONTAP operating system was recognized by IDC as the #1 branded storage OS.¹

¹Source: IDC Worldwide Quarterly Disk Storage Systems Tracker Q4 2012, March 2013 (Open Networked Disk Storage Systems revenue)

In clustered Data ONTAP 8.2, NetApp implements SMB 3.0, the protocol's latest version, with features like persistent file handles (continuously available file shares) and fully supports transparent clustered client failover and witness protocol. Data center design is simple with SMB 3.0 file shares since several virtual machine hard drives (VHDs or VHDXs) can be located on a file share. Other features in the SMB 3.0 space that provide added value include scale-out awareness, ODX, and VSS for SMB file services. Offloaded Data Transfer (ODX) is particularly powerful because it can rapidly speed up virtual machine deployment by cloning virtual machines on the storage system with zero network traffic required on the Windows Server 2012 Hyper-V host-side. ODX as implemented in clustered Data ONTAP can even work seamlessly between volumes used for NAS or SAN protocols on the storage, a feature exclusive to NetApp. These features are discussed in greater detail in [TR-4172: Microsoft Hyper-V over SMB 3.0 with Clustered Data ONTAP: Best Practices](#).

NetApp has continued to innovate with SAN protocols in clustered Data ONTAP 8.2. The maximum SAN clustered Data ONTAP storage controller node count was increased from 6 to 8 nodes. The maximum number of volumes was increased to 6,000 and the maximum number of LUNs to 49,152 for the largest clusters (note that these specifications are reduced for smaller platforms). NetApp LUNs that are used as cluster shared volumes (CSVs) can be accessed seamlessly by several Windows Server 2012 Hyper-V nodes while being distributed across several physical NetApp storage nodes. For Microsoft Windows Server 2012 Hyper-V on NetApp, this provides for the first time a truly distributed file system, with everything mounted under a common name space for SAN, NAS, or both. All clustered Data ONTAP solutions continue to provide traditional Data ONTAP features, such as thin provisioning, deduplication, Snapshot™ copies, and FlexClone® technology, to help customers reduce storage space consumption and to quickly back up and recover data.

1.2 Windows Server 2012 Hyper-V on NetApp: Testing Performance Across Protocols

This technical report discusses the results from a recent set of performance tests completed by NetApp to measure the performance of a typical virtual machine workload. Specifically, we deployed Windows Server 2012 with Hyper-V role enabled on a two-node Windows® Failover Cluster with a clustered Data ONTAP 8.2 two-node solution and tested the performance across various protocols. We tested with FC using 8GB FC SAN networks and iSCSI and SMB 3.0 using 10GbE networks. The goal of the tests and this report is to evaluate the performance of each solution and help customers choose how to build their virtual infrastructures.

2 Executive Summary

We set up the test environment according to Microsoft and NetApp best practices. Given the new release of SMB 3.0, it was imperative that we quantify if it could perform up to par with traditional block protocols like iSCSI and FC. These tests were not designed to show the maximum throughput achieved for each protocol but to simulate a real-world environment running real-world loads. We tested with several Windows Server 2012 Hyper-V servers, each hosting several virtual machines all accessing data at the same time. The tests used a realistic I/O pattern, I/O block size, and read/write mix.

For the Windows Server 2012 with Hyper-V servers, we used two Fujitsu RX300S6 servers in a cluster configuration to host the virtual machines. These servers each contained 2-socket 6-core Intel® Xeon® 5645 processors with 48GB of RAM each. We used two NetApp FAS6240 controllers in a high-availability clustered Data ONTAP 8.2 configuration. The two controllers were connected through redundant 10Gb Ethernet cluster interconnect links and switches. Clustered Data ONTAP by default provides storage controller virtualization with storage virtual machines that can span several physical storage controller nodes and disk aggregates. Data ONTAP storage virtual machines provide network port virtualization using logical interfaces (LIFs) that reside on top of storage controller physical interfaces for both Fibre Channel and Ethernet. We used a single storage virtual machine with several volumes located on disk aggregates on both storage controllers. In addition, we leveraged NetApp Flash Cache™ technology by using a 512GB controller-attached PCIe intelligent caching module.

The network connectivity was provided through the best-in-class switching infrastructure. For the Ethernet connectivity, a Cisco[®] Nexus[®] 5020 Ethernet switch was used with various 10GbE network interfaces on the servers and storage controllers. For the Fibre Channel SAN connectivity, we used a Brocade 300 Fibre Channel switch operating at 8Gb speed with numerous ports for each Windows Server 2012 Hyper-V server and storage controller.

2.1 Test Workload

This section describes the access specifications that were used to create a mixed virtual server workload with the publicly available tool Iometer. The goal was to simulate a set of small database servers hosted in a virtualized environment performing an OLTP workload. The workload used was a mix of 70% read and 30% write, and 100% random 8K request size operations. Note that 8K is the default block size of Oracle[®] Database and Microsoft SQL Server servers.

The following statistics were measured for each protocol (FC, iSCSI, and SMB 3.0).

- Throughput in IOPs
- Latency in milliseconds (ms)

For each protocol, we tested performance at three different load levels to simulate, low, medium, and high virtualized server workloads. This can be controlled by adjusting the I/Os outstanding within the Iometer controller software. We tested with two I/Os outstanding, four I/Os outstanding, and six I/Os outstanding. The test results were collected at the point at which the performance of the systems had reached a steady state.

Ten virtual machines running Windows 7 Ultimate SP1 were deployed on each of the two Windows Server 2012 Hyper-V servers for a total of 20 virtual machines. The VMs accessed their respective virtual hard drives (VHDXs) located on cluster shared volumes (CSVs) that were located on NetApp LUNs for SAN protocols tests or SMB 3.0 shares on NetApp volumes for NAS protocol tests. Each VM ran with 2, 4, and 6 outstanding I/Os for a total of 40, 80, and 120 outstanding I/Os to simulate low, medium, and high workloads, respectively. The workload of all VMs running at once created a combined intense workload that stressed the servers, storage, and network connections. We believe this workload effectively simulates mixed database and application servers virtualized on Microsoft Windows Server 2012 Hyper-V deployed on clustered Data ONTAP 8.2 shared storage.

2.2 Performance Summary

NAS storage protocols are becoming a preferred storage protocol for virtualization and cloud platforms for a number of reasons, including ease of management, scalability, object-level granularity, and extended integration into applications and automation workflows. With all of these benefits, many question the performance capability of modern NAS protocols like SMB 3.0 in comparison to more traditional block protocols like iSCSI and FC. The purpose of this report is to validate the performance of all three protocols using a mixed OLTP workload.

As discussed in the introduction, many features were added to SMB 3.0 to make it resilient, but the question remains, does it perform? We are happy to report that the answer is a resounding “Yes!”

- The test results show that SMB 3.0 provides comparable throughput performance with iSCSI and FC protocols at each load level.
- Our test results demonstrate that comparable latency can be achieved with SMB 3.0, iSCSI, and FC at each load level.
- We found that Windows Server 2012’s average total CPU utilization was not significantly different when choosing FC, iSCSI, or SMB 3.0.

With confidence in solution performance, customers can make the optimal choice based upon cost, design, and management needs for their current and future virtualized data centers and private clouds.

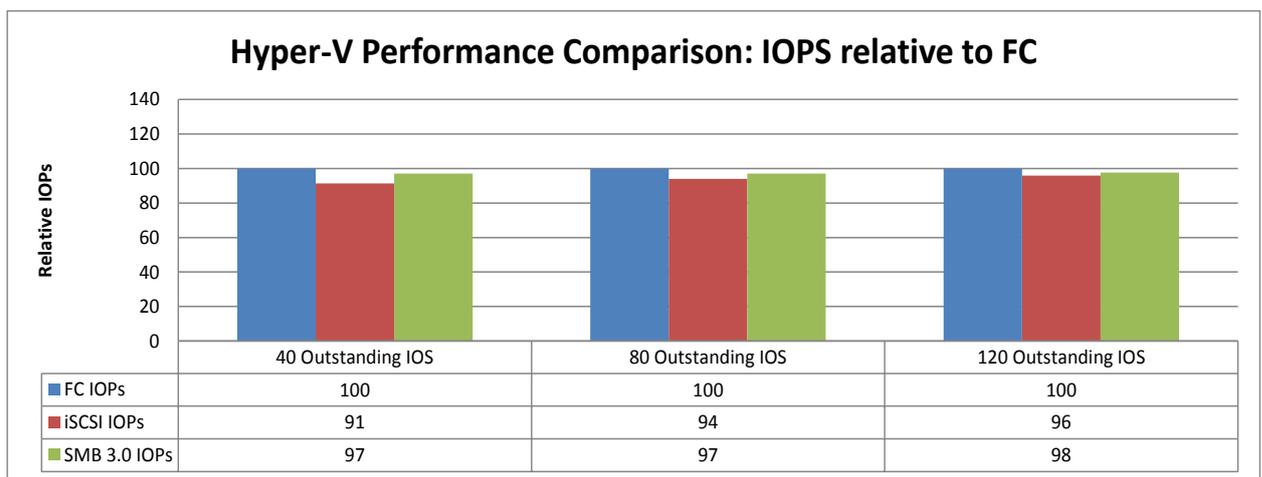
3 Performance Comparison Details

As mentioned above, the workload was composed of 100% random 8K request size I/Os with a mix of 70% reads and 30% writes. To simulate low, medium, and heavy workloads, a constant number of VMs was used, but the number of I/Os outstanding was modified. We used 2, 4, and 6 I/Os outstanding per VM with 20 VMs deployed for a total of 40, 80, and 120 I/Os outstanding to the storage controller.

While reading the numbers, keep in mind that the values for different protocols are presented on a relative scale and are not absolute values. This means that all results are compared to 8Gb FC as the baseline. On the charts, FC is valued as 100 and SMB 3.0 and iSCSI results are calculated relative to that value. For example, when the throughput for SMB 3.0 is shown at 98, this indicates that throughput for SMB 3.0 was 2% less than the throughput for FC.

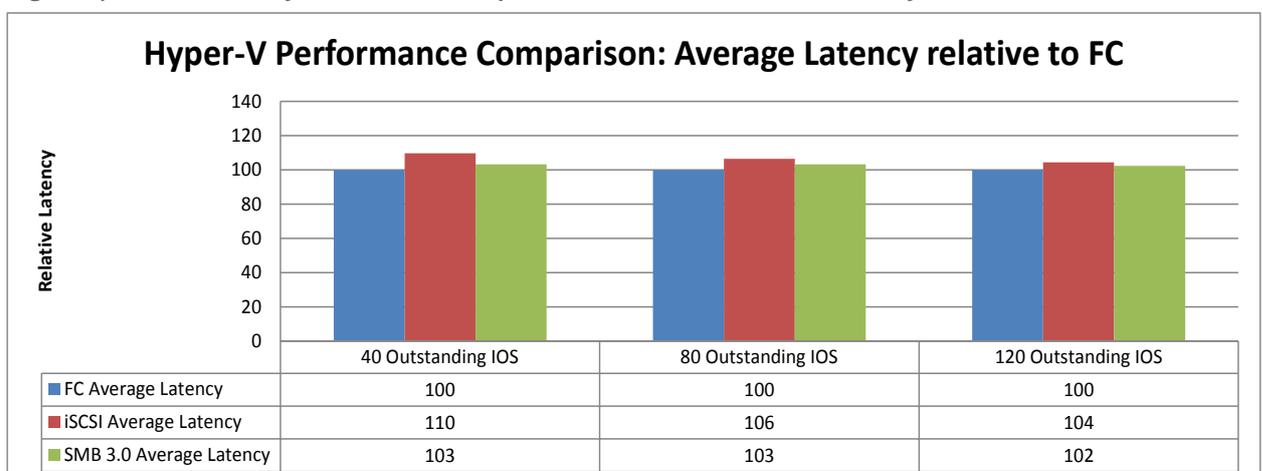
3.1 Relative Throughput Comparison

Figure 1) Relative throughput in IOPs across the three protocols at low, medium, and heavy workload levels.



3.2 Relative Latency Comparison

Figure 2) Relative latency across the three protocols at low, medium, and heavy workload levels.



The graphical representation of the CPU utilization for the Windows Server 2012 Hyper-V servers and the clustered Data ONTAP storage controllers is not provided in this report. We found that Windows Server 2012's average total CPU utilization was not significantly different when choosing FC, iSCSI, or SMB 3.0. We examined the server physical CPU utilization counters gathered through Perfmon and

the utilization was similar across all three protocols. Within Data ONTAP, CPU utilization is not a good representation of overall system utilization. Data ONTAP is a highly parallelized operating system designed to effectively spread workload across all cores within the system; therefore, high CPU utilization indicates effective use of all cores. We examined the CPU utilization of Data ONTAP during the workload of all three protocols using perfstat and found that each protocol acceptably distributed the workload across all storage controller CPU cores. The metrics for gauging Data ONTAP system performance including throughput and latency were examined at the storage level and demonstrated comparable throughput and latency to what was seen at the Windows Server 2012 Hyper-V server level and the Iometer-application level.

4 Test Design and Configuration

This section provides detailed information about the hardware and software deployed for the test environment. Included is network infrastructure, virtual machine provisioning, and storage provisioning. Storage connectivity speeds for SMB 3.0 and iSCSI were 10Gb throughout and FC was 8Gb to provide performance and resiliency that one would expect for an enterprise virtualized data center or private cloud.

4.1 Hardware and Software

Table 1 provides details of the hardware and software components that were used to create the Microsoft Windows Server 2012 Hyper-V data center. The data center includes two Windows Server 2012 servers with Hyper-V role enabled in a single Windows cluster. Microsoft Failover Cluster Manager, installed locally on the Windows Server 2012 host cluster, was used to manage the virtual machine environment.

Table 1) Hardware and software for Windows Server 2012 Hyper-V data center.

Hardware/Software Component	Details
Server	2 x Fujitsu Primergy RX300 S6
Processors	2 Sextet-Core Intel Xeon E5645s @ 2.4GHz
Memory	48GB
OS	Windows Server 2012 Datacenter Edition with Hyper-V role enabled
Fibre Channel network	8Gb FC
Fibre Channel HBA	QLogic QLE2562 dual-port PCIe FC HBA
Dedicated NIC for Hyper-V host management	Intel 82575EB Gigabit Network Connection
Dedicated NIC for Virtual Machine Network Configuration (virtual switch)	Intel 82575EB Gigabit Ethernet Controller
Dedicated NICs for 10GbE iSCSI and SMB3 connections	Intel 82599 10 Gigabit dual-port Ethernet Network Connection
8Gbps Fibre Channel switch	Brocade 300E
Gigabit Ethernet switch	Cisco Catalyst® 4948 Ethernet switch
10 Gigabit Ethernet switch	Cisco Nexus 5020

Table 2 provides details of the clustered Data ONTAP unified storage solution.

Table 2) Hardware and software for clustered Data ONTAP 8.2.

Hardware/Software Component	Details
Storage controller	2x FAS6240 configured as SFO pair
Clustered Data ONTAP version	8.2RC1
Total number of drives	144
Size of drives	450GB
Drive speed	15K RPM
Drive type	SAS
Flash Cache	2 x 512GB
FC target ports	4 x 8Gb (2 per node)
10Gb Ethernet ports	8 x 10Gb (2 per node)
Storage virtual machines	4 x storage virtual machines (1 cluster admin, 1 per node [2 nodes], 1 data)
Data Ethernet LIFs (on 10GbE ports)	4 x data storage virtual machines (2 per node, each connecting to separate private VLAN)
Data FC LIFs (on 8Gb FC ports)	4 x data storage virtual machine (2 per node)

Table 3 provides details of the server used to initiate the Iometer workload.

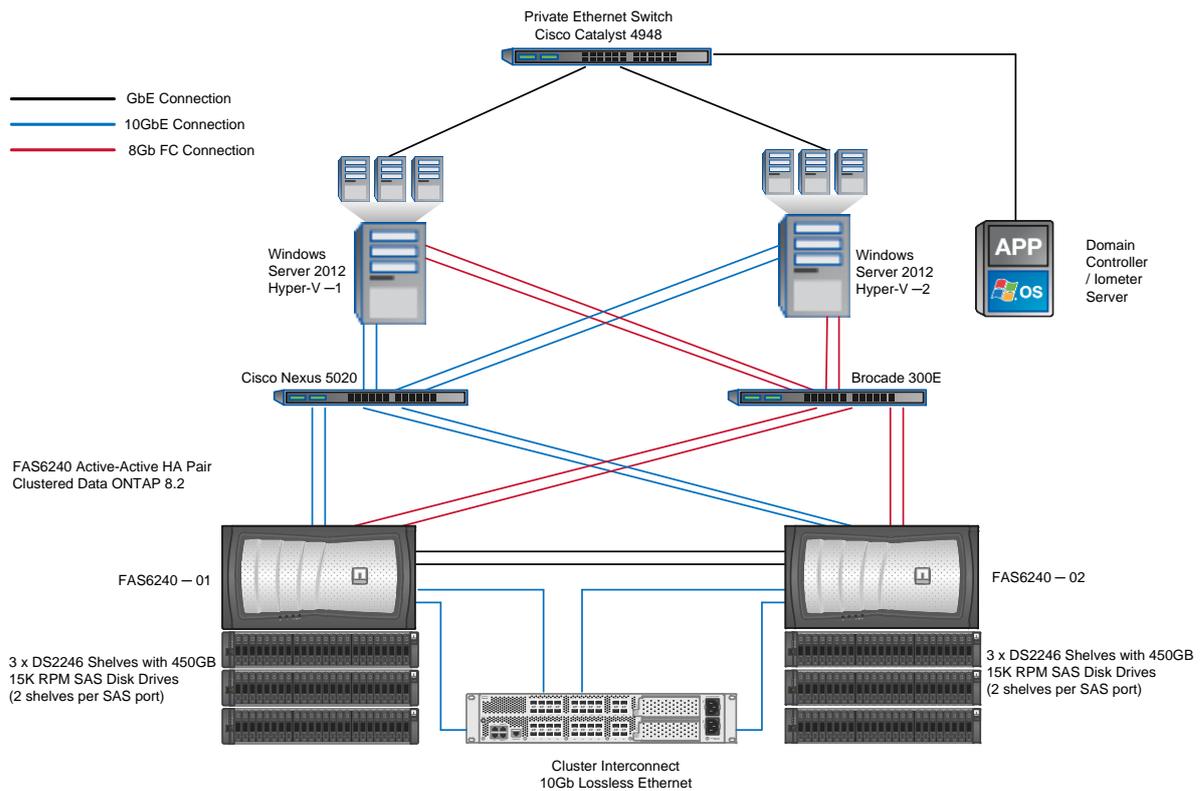
Table 3) Iometer workload generation controller server.

Hardware/Software Component	Details
Server	1 x virtual machine
Processors	1 VCPU Intel Xeon E5645 @ 2.4GHz
Memory	4GB
Operating system	Microsoft Windows Server 2008 R2 SP1 Enterprise Edition

4.2 Network Connections

This section provides the details of the network connectivity between the Microsoft Windows Server 2012 Hyper-V and the clustered Data ONTAP FAS6240 storage system. The network diagram, as illustrated in Figure 3, shows that the SMB 3.0 and iSCSI connections were made through several 10GbE connections on each server and storage virtual machine connected to separate private VLANs configured on the Cisco 5020 switch. The FC SAN was configured using the 8Gb Brocade 300E switch. Several ports were used on servers and storage systems to provide a fully redundant multipath solution. The Microsoft Windows Server 2012 Hyper-V servers and Iometer controller/domain controller were connected through a separate GbE Cisco Catalyst 4948 Ethernet switch.

Figure 3) FC, iSCSI, and SMB 3.0 connections between Windows Server 2012 Hyper-V servers and NetApp shared storage.



4.3 Storage Provisioning Design FC and iSCSI

All VM VHDX volumes were stored on a single large data aggregate on the storage controller. The aggregate contained 68 450GB 15K RPM SAS high-performance disk drives. It included four 17 (15+2) disk RAID-DP[®] groups. The storage controller had one spare disk and a small three-disk aggregate for the operating system and configuration files.

The virtual machine OS hard drives (VHDXs) were deployed on volumes separate from the data drives but on the same large aggregate. Each volume contained a single LUN mapped through FC to the Windows Server 2012 Hyper-V servers, and each LUN contained 20 VHDX drives. The OS VHDX drives for the Windows Server 2012 Hyper-V server 1 were located on storage controller 1 and for the Windows Server 2012 Hyper-V server on storage controller 2. We created two volumes, each containing a single LUN for the VM data drives. These LUNs were used for both SAN protocol tests (iSCSI and FC). It was simple to switch back and forth between iSCSI and FC protocols by simply unmapping the LUNs from their respective igroups for FC and mapping them to igroups for iSCSI. The LUNs were formatted by Disk Management with GPT and the NTFS file system using the default block size. Then we created CSV volumes shared by the Windows Server 2012 Hyper-V servers under a common name space. For each server there were two CSVs.

1. One CSV for the VM C:\ drives with operating system files
2. One CSV for the VM E:\ data drives with test files (10 VHDX data drives per CSV)

The LUNs were mapped to the servers using the following igroups for each respective test:

- Two iSCSI igroups, one for each Windows 2012 Hyper-V server, containing the software iSCSI initiator name
- Two FC igroups, one for each Windows 2012 Hyper-V server, containing the FC initiator WWPNs

Appropriate zoning was completed for the FC SAN on the Brocade switch and the LUNs were connected into the Windows Server 2012 Hyper-V servers as disks; then they were designated as CSVs by using the Windows Cluster Failover Manager. Additionally, the Microsoft Windows native MPIO driver and Microsoft DSM were used to provide multipath access based upon Asymmetric

Logical Unit Assignment (ALUA) communicated by the clustered Data ONTAP storage controllers to the servers. The MPIO policy was retained as the default of Round Robin with Subset. ALUA capability provides optimal path usage from server to storage and is available for iSCSI and FC in clustered Data ONTAP 8.2.

The full storage system design for iSCSI and FC testing is detailed in Table 4.

Table 4) FAS6240 aggregate, FlexVol volume, and LUNs for testing block protocols (iSCSI and FC).

NetApp Storage Controller	Aggregate Name	Volume Name	LUN	Size	Description
Controller 1	Aggr0			3 disks	
		Vol0		348.3GB	Root volume
	C1_Aggr1			68 disks	
		Hv_c_c1vol		1TB	
			vm_cdrive	1TB	VM OS disks
		Hv_e_c1vol		1TB	
			vm_edrive	1TB	Iometer data files
		Hv_e_c1vol2		1TB	
			vm_edrive2	1TB	Iometer data files

4.4 Storage Provisioning Design SMB 3.0

For the SMB 3.0 protocol tests, we maintained the same virtual machine OS hard drives configuration as with iSCSI and FC. The c:\ drives containing the VHDXs with the Windows 7 operating system were located on a separate volume containing a LUN and connected through FC. We created separate volumes for the VM data drives. All volumes were stored on a single large data aggregate on storage controller 2. The aggregate contained 68 450GB 15K RPM SAS high-performance disk drives and was composed of 4 17-disk (15+2) RAID-DP RAID groups. The storage controller also had a single spare disk and a small three-disk aggregate for the Data ONTAP operating system and configuration files.

We created 2 volumes and 2 SMB 3.0 shares as containers for the 20 VM data drives D:\drives. We set the share properties to Continuously Available as required for Windows Server 2012 Hyper-V over SMB 3.0. When deploying VM hard drives on SMB 3.0, it is not necessary to use the Windows Disk Management stack and then apply the Cluster Shared Volume properties on top of the volume. Instead, SMB 3.0 allows simple file sharing in a common name space controller by the storage virtual machine according to the volume and mount point design within clustered Data ONTAP.

The full storage system design for SMB 3.0 testing is detailed in Table 5.

Table 5) FAS6240 aggregate, FlexVol, and LUNs for testing SMB 3.0 protocol.

NetApp Storage Controller	Aggregate Name	Volume Name	LUN	Size	Description
Controller 2	Aggr0			3 disks	
		vol0		348.3GB	Root volume
	C2_Aggr1			68 disks	
		Hv_c_c2vol		1TB	
			vm_cdrive	1TB	VM OS disks
		Hv_d_c2vol_smb		1TB	Iometer data files
		Hv_d_c2vol_smb_2		1TB	Iometer data files

The OS drives were intentionally left running on FC to simplify configuration details for the SMB 3.0 tests. The workload generated by the operating system C:\ drives for the virtual machines is a constant across all three workloads. It has also been established that once virtual machines are booted and running, the storage controller workload for the OS drives is relatively minimal. Instead,

the Windows Server 2012 Hyper-V server provides most of the processing power needed for the VMs to run. Therefore, we chose not to run the VM OS drives over SMB 3.0 during tests. We believe the results would have been equivalent had the VM OS drives been running over SMB 3.0.

4.5 Building the Windows Server 2012 Hyper-V Data Center

A single virtual machine was installed with Windows 7 SP1 Ultimate into a single 15GB VHDX, and configuration changes were made as required. Then we used the Microsoft Sysprep tool to create a golden image of the virtual machine. Sysprep removes the machine name, domain information, and GUID so that it can easily be cloned. We then detached the VHDX from the virtual machine. This operating system VHDX was used as a master from which we cloned 40 virtual machines.

Each virtual machine was provided with a 20GB VHDX data drive that was assigned through disk management as the E:\drive for iSCSI/FCP and D:\drive for SMB 3.0. We used Windows Disk Management to format the drive with the NTFS file system using the default block size. For the test results discussed in this technical report we used a fixed VHDX, since this has traditionally been the most common deployment type used for Hyper-V and NetApp storage solutions.

Microsoft Windows Server 2012 Hyper-V includes dynamic and differencing VHDX drive types. Dynamic disks start as sparse disks and are not populated until the first reads and/or writes occur. Differencing disks start with a fixed disk as the parent. The differencing disks are only 4KB in size and grow as the reads and writes occur. We tested with dynamic and differencing VHDX drives and found the performance to be equivalent to fixed VHDX drives.

Table 6) Virtual machine components.

Component	Details
Operating system	Windows 7 SP1 Ultimate
Number of virtual processors	1
Virtual memory	1GB
OS virtual disk size	15GB
Iometer data virtual disk size	20GB

4.6 Workload and Iometer Test Practices

We used the publicly available workload generator application Iometer to drive the workload. It can be found at www.iometer.org. It is a client-server application that works as both a workload generator and a measurement tool. The server portion is called the "Iometer controller" and was installed on a standalone Windows 2008 R2 SP1 server that was separate from the Microsoft Windows Server 2012 Hyper-V hosts. We used a Microsoft Windows 2008 R2 SP1 server in a dual role as Iometer controller and domain controller/DNS/DHCP provider. The client portion of Iometer works in a distributed manner by installing "Dynamo" on each of the VMs. Twenty Windows 7 SP1 Ultimate VMs split between two Microsoft Windows Server 2012 Hyper-V servers simultaneously executed the Iometer workload.

After the VMs were deployed as described in the section above, Iometer was used to initialize the VM data drives. The steps were as follows.

1. Power on 20 VMs.
2. Execute the Dynamo application on each VM that syncs with the Iometer controller for instructions.
3. Launch the Iometer controller application on the Iometer controller server.
4. From the Iometer controller application, for each of the VMs select the appropriate drive letters for the FC/iSCSI data drives or the SMB 3.0 data drives.
5. Create an Iometer access specification to generate a workload of 100% random reads with an 8K request size with no ramp-up time. Set I/Os to align on a 4K boundary.

6. Execute the specification workload on all VMs for a total of 5 minutes.
7. Allow the test to complete, then discard the results.
8. The test created a 20GB Iometer data file called iobw.tst on each data drive (VHDX) for each VM. These drives were seen as E:\ for iSCSI and FC and D:\ for SMB3 and were used for each workload test thereafter.
9. Power off the virtual machines and create a Snapshot copy of the volumes containing the data drives, as described in Table 4 and Table 5.

After completing the steps above to initialize the dataset and creating a Snapshot copy of the dataset in its initial state, we were ready to test. After completing each test sequence, we leveraged NetApp Snapshot technology to quickly roll back the dataset to its initial state. We powered off the VMs after each test sequence and used the SnapRestore[®] function to restore virtual hard drives to the exact same point in time with the Iometer data files as when first initialized. This helped to enable consistency in our methods and results.

We used the following access specifications to create a mixed simulated virtual machine workload.

- 70% read, 30% write, 100% random 8K request size

This workload type is also commonly used to simulate OLTP workloads. We measured the following statistics for each protocol and workload:

- Throughput in IOPS
- Latency in milliseconds (ms)

Each of the three workloads was run at three different load levels to simulate, low, medium, and high server workloads. This can be controlled by adjusting the I/Os outstanding within the Iometer controller software. The I/O was aligned to the data.

Table 7) Outstanding I/Os per VM and the total outstanding I/Os.

VMs	Outstanding I/Os per VM	Total Outstanding I/Os
20	2	40
20	4	80
20	6	120

We ran the mixed workload for several hours to fully warm up the NetApp Flash Cache. After the throughput achieved a steady state, we tested two times at each level. For each tested level, we ran a 30-second warm-up followed by a 15-minute test and then recorded the results. We collected perfstat data on the NetApp storage controllers and Windows Perfmon data for the Microsoft Windows Server 2012 Hyper-V. Though each virtual machine had a 20GB data drive, we limited the working set size to 10GB per VM in the Iometer configuration files. This created a total working set size of 200GB per Windows Server 2012 Hyper-V server. The working set was large enough to exceed the clustered Data ONTAP storage controller memory cache though a significant amount of the workload was satisfied by the NetApp Flash Cache.

4.7 Additional Configuration Details

On the Windows Server 2012 Hyper-V servers we set the FC HBA execution throttle (queue depth) to 255.

On the NetApp storage controllers we set the following option:

```
options wafl.optimize_write_once off
```

5 Conclusion

In this report, we examine new virtualized data center design technologies for deploying Microsoft Windows Server 2012 Hyper-V on clustered Data ONTAP 8.2 storage solutions. We report the results

from recent tests performed by NetApp to determine the performance of a mixed virtual machine workload with OLTP-type performance characteristics over SMB 3.0, iSCSI, and FC protocols. We find that when measuring throughput, all protocols perform within 9% of one another. In measuring latency the performance of each protocol is within 10% of the other protocols. These results demonstrate that all three storage protocols are production worthy and validate that the performance capabilities of SMB 3.0 are inline with its operational benefits.

References

The following references were used in this technical report:

- Clustered Data ONTAP 8.1.1: Best Practices for NetApp SnapManager for Hyper-V
www.netapp.com/us/system/pdf-reader.aspx?pdfuri=tcm:10-60928-16&m=tr-4004.pdf
- Microsoft Hyper-V R2 Performance Comparison: Fibre Channel and iSCSI
www.netapp.com/us/system/pdf-reader.aspx?pdfuri=tcm:10-59952-16&m=tr-3846.pdf
- VMware vSphere 4.1 Storage Performance: Measuring FCoE, FC, iSCSI, and NFS Protocols
www.netapp.com/us/system/pdf-reader.aspx?pdfuri=tcm:10-59714-16&m=tr-3916.pdf
- TR-4172: Microsoft Hyper-V over SMB 3.0 with Clustered Data ONTAP: Best Practices
<https://fieldportal.netapp.com/?oparams=131268>

Acknowledgements

Special thanks go to the following people for their contributions:

- Keith Griffin, EDCA Lab Support, NetApp
- Brad Kintner, Senior Manager, Workload Engineering, NetApp
- Saad Jafri, Database Performance Engineer, NetApp
- Bhavik Desai, Workload Performance Engineer, NetApp
- Glenn Sizemore, Technical Marketing Engineer, NetApp
- Pavel Lobanov, Technical Marketing Engineer, NetApp

Refer to the [Interoperability Matrix Tool \(IMT\)](#) on the NetApp Support site to validate that the exact product and feature versions described in this document are supported for your specific environment. The NetApp IMT defines the product components and versions that can be used to construct configurations that are supported by NetApp. Specific results depend on each customer's installation in accordance with published specifications.

NetApp provides no representations or warranties regarding the accuracy, reliability, or serviceability of any information or recommendations provided in this publication, or with respect to any results that may be obtained by the use of the information or observance of any recommendations provided herein. The information in this document is distributed AS IS, and the use of this information or the implementation of any recommendations or techniques herein is a customer's responsibility and depends on the customer's ability to evaluate and integrate them into the customer's operational environment. This document and the information contained herein may be used solely in connection with the NetApp products discussed in this document.

Go further, faster®