Storage I/O Performance on VMware vSphere 6.0 U2 over 32 Gigabit Fibre Channel

Performance Study

TECHNICAL WHITE PAPER

vmware

Table of Contents

Introduction

VMware vSphere® 6.0 Update 2 adds support for 32 gigabit per second (32Gb) Fibre Channel (FC) host bus adapters (HBAs). This paper explores vSphere storage performance improvements of 32Gb FC over 16Gb FC. For sequential I/O of 8KB and above, 32Gb FC has higher sequential read throughput and sequential write throughput than 16Gb FC. For random I/O of 8KB and above, 32Gb FC has higher random read IOPS than 16Gb FC. We expect virtualized applications on vSphere that are storage bandwidth limited under 16Gb FC to exhibit better performance with 32Gb FC.

Test Environment

Testbed Hardware and Software

The testbed consisted of the components shown in Table 1, as follows. Numbers before an "x" indicate how many of that item were included.

System	Parameter	Value
vSphere/ESXi Host	System	Dell PowerEdge R720
	CPU	Intel Xeon E5-2643 v2, 3.5GHz, 2x 6-core (Sandy Bridge EP)
	RAM	64GB
	PCIe	Gen3
	FC HBA 32Gb	Emulex LPe32002 with Avago 32Gb SFP
	FC HBA 16Gb	Emulex LPe16002B-M6 with Avago 16Gb SFP
	Build	ESXi 6.0 Update 2
	FC driver	ESXi 6.0 lpfc 11.0.237.0 FC/FCoE Driver LUN queue depth: 250 HBA queue depth: 8192
Storage Array	Model	EMC VNX7500
	Firmware	05.32.000.5.218
	FC Target Ports	8Gb (4x)
	SSDs	Samsung SS160520 200GB SATA Flash (12x)
	RAID Group	RAID 1/0
	LUNs	100MB (32x)
FC Switch	Model	Brocade G620, 32Gb
	Ports To ESXi	32Gb LPe32002 (1x) 16Gb LPe16002B-M6 (1x) Only one port per HBA was used. Only one HBA is used in any experiment.
	Ports to Storage Array	8Gb (4x)
VM	Virtual Hardware	version 11
	vCPU	2
	vMemory	4GB

System	Parameter	Value
	vAdapters	PVSCSI (2x) Adapter queue depth: 254
	vDisks (I/O)	4x 100MB Raw Device Mapping Physical (RDMp)
	Guest OS	64-bit Windows Server 2012 Datacenter

Table 1. Testbed hardware and software

Testbed Architecture

The testbed (see Figure 1) consists of a VMware ESXi™ (vSphere) host, a Brocade G620 FC switch, and an EMC VNX 7500 storage array. The ESXi host is connected to the Brocade G620 32Gb FC switch via Emulex 32Gb and 16Gb FC HBAs. A single port on each HBA is used. The FC switch is in turn connected to the EMC VNX 7500 storage array using four 8Gb FC connections. (The VNX 7500 used does not have 16Gb or 32Gb FC target ports.) The ESXi host is a Dell PowerEdge R720 with two 6-core Intel Xeon E5-2643 v2 processors at 3.5 GHz, running ESXi 6.0 U2. Each virtual machine runs 64-bit Windows Server 2012 Datacenter. Iometer is used to generate the I/O workload. The virtual disks in each VM are accessed with sequential read-only, sequential write-only, and random read-only I/O for a range of I/O sizes that will be described in the following sections. On the VNX 7500 storage array, 32 100GB LUNs are configured using 12 Samsung SS160520 200GB SATA flash drives.

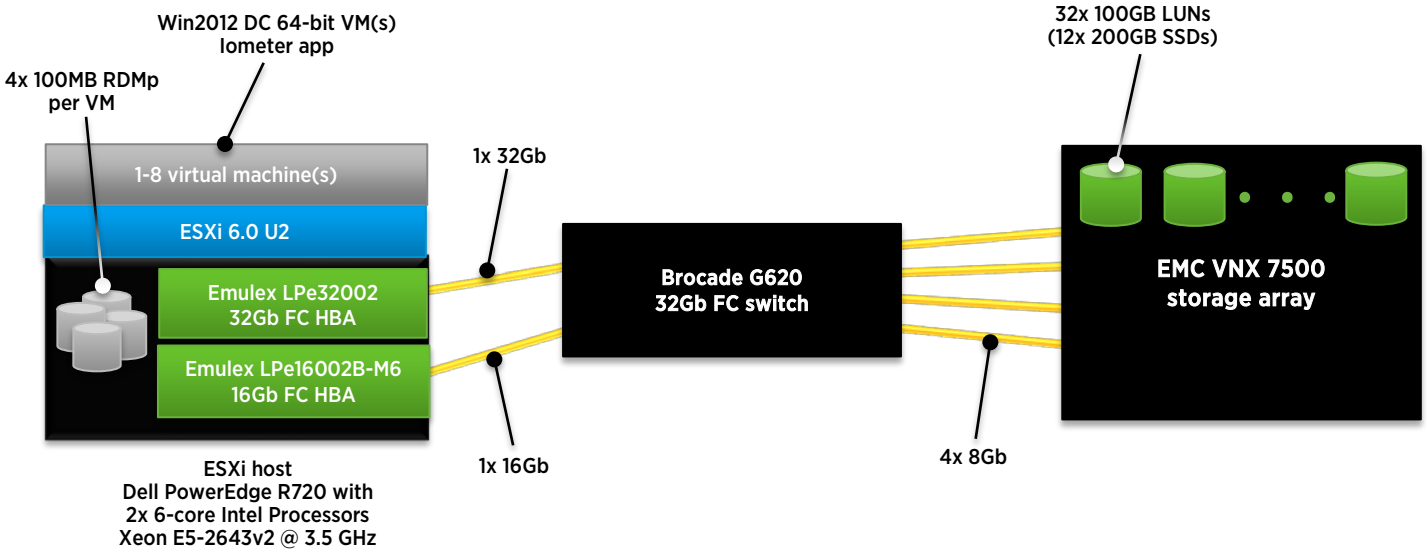

Figure 1. Testbed architecture

I/O Workload

Each virtual machine has four 100MB virtual disks (Raw Mapped LUNs, RDMp). I/O is generated by four Iometer workers in each virtual machine, with one worker per virtual disk. For single virtual machine experiments, 32Gb FC and 16Gb FC are compared using sequential read and write I/Os, with each worker generating 96 outstanding I/Os (OIOs). For multiple virtual machine experiments, 32Gb FC is evaluated using random read I/Os as the number of virtual machines increased from one to eight, with each worker generating 32 outstanding I/Os (OIOs).

Test Results

Sequential Read and Write Throughput for One Virtual Machine

Figure 2 compares single VM sequential read throughput (in megabytes per second, or MB/s) of 32Gb versus 16Gb FC for I/O sizes of 1KB to 256KB. For I/O sizes of 8KB and above, the sequential read throughput of 32Gb FC is higher than that for 16Gb FC. For I/O sizes of 128KB and above, sequential read throughput of 32Gb FC approaches the 32Gb line rate. Sequential read throughput for I/O sizes of 128KB+ for 32Gb FC is higher than that for 16Gb FC by more than 90%.

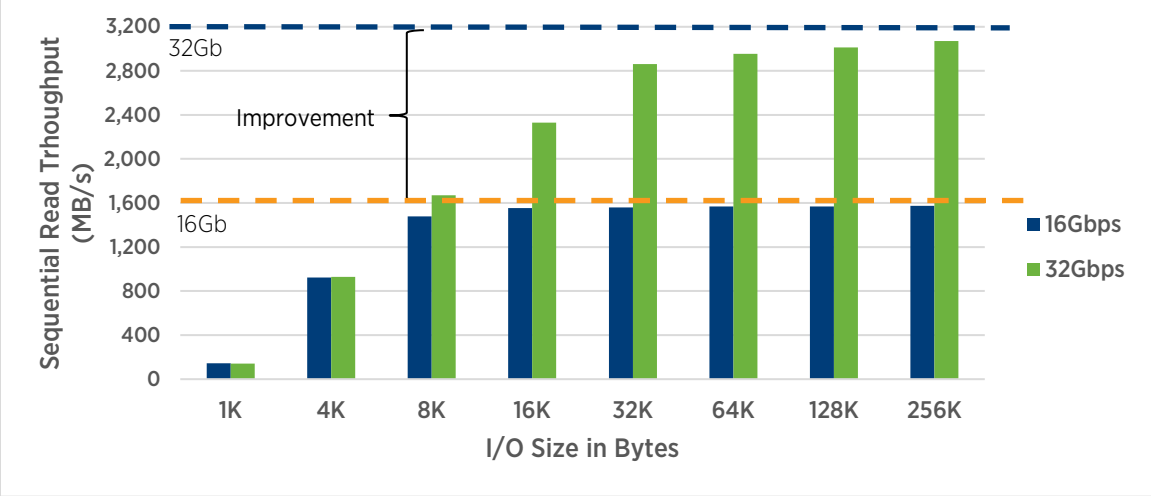

Figure 2. Single-VM: Sequential Read Throughput, comparing 32Gb and 16Gb

Figure 3 compares single VM sequential write throughput of 32Gb versus 16Gb FC for I/O sizes of 1KB to 256KB. For I/O sizes of 8KB and above, sequential write throughput of 32Gb FC is higher than that for 16Gb FC. For I/O sizes of 64KB and above, sequential write throughput of 32Gb FC approaches 32Gb line rate. Sequential write throughput (32KB+) for 32Gb is higher than that for 16Gb by more than 90%.

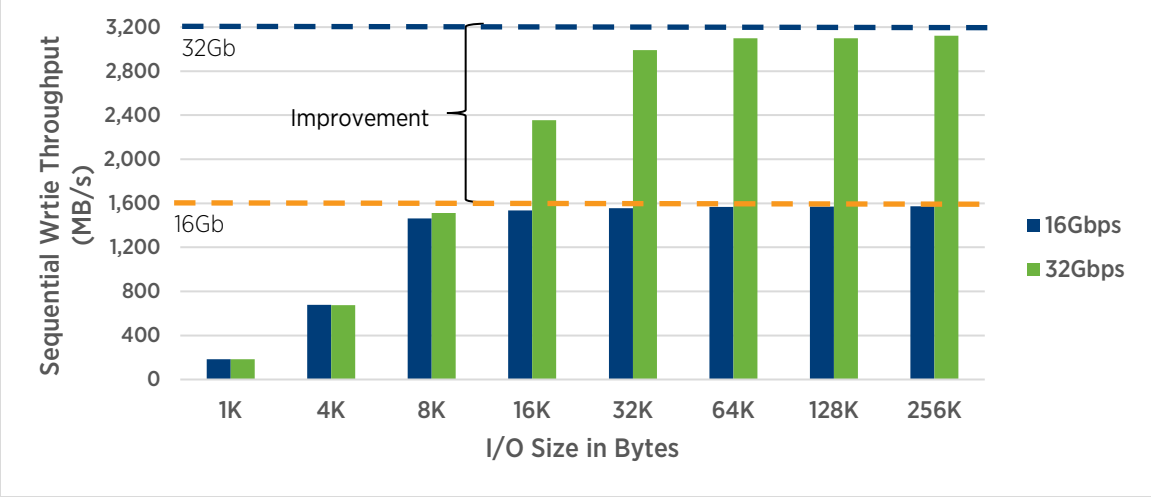

Figure 3. Single-VM: Sequential Write Throughput, comparing 32Gb and 16Gb

Random Read IOPS and Throughput for Multiple VMs – Scaling Up

Figure 4 and Figure 5 show random read IOPS and throughput respectively for 32Gb FC, for 1 to 8 VMs. For I/O sizes of 8KB and above, random read throughput of 32Gb FC exceeds 16Gb line rate. The higher bandwidth of 32Gb FC—higher than the orange horizontal line in Figure 5—supports higher IOPS than that provided by 16Gb FC. For I/O sizes of 64KB and above, random read throughput of 32Gb FC is close to the 32Gb line rate.

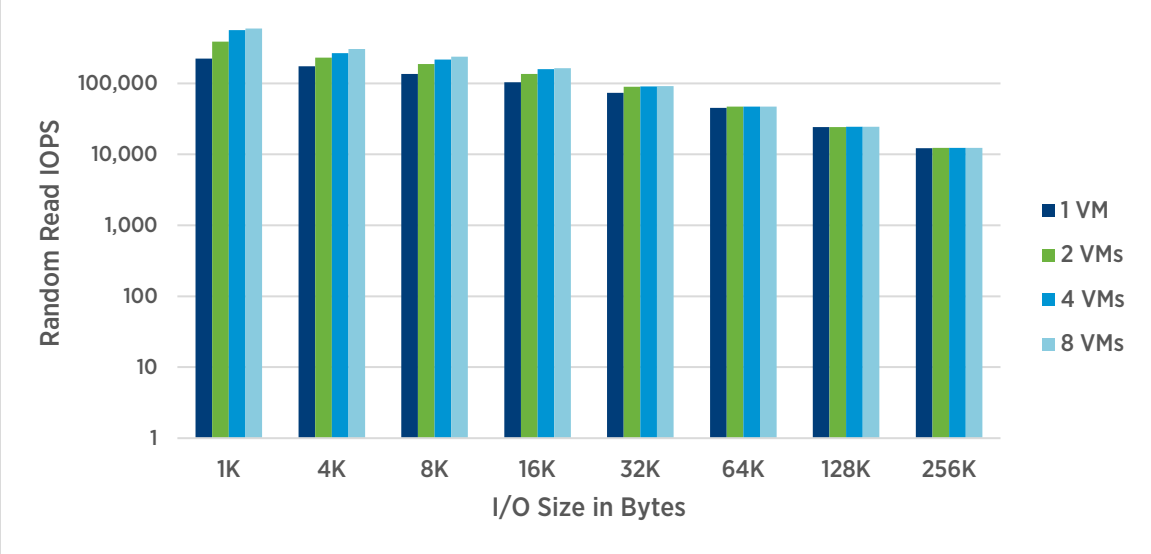

Figure 4. Multi-VM: Random Read IOPS for 1, 2, 4, and 8 VMs, for 32Gb

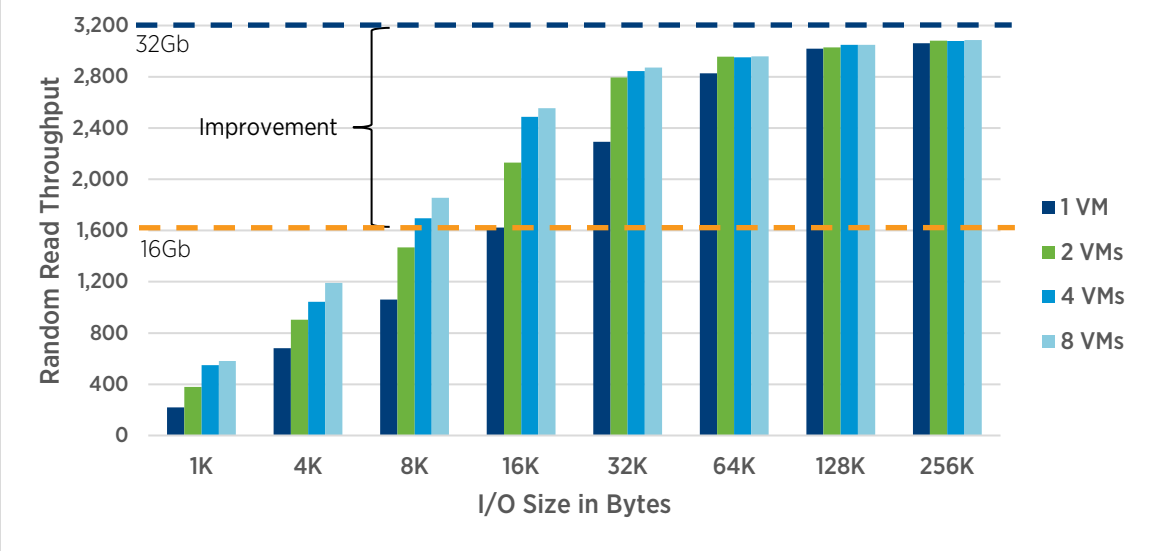

Figure 5. Multi-VM: Random Read Throughput for 1, 2, 4, and 8 VMs, for 32Gb

Conclusion

vSphere 6.0 Update 2 added support for 32Gb FC HBAs. This paper compared the performance of 32Gb FC with 16Gb FC. For sequential I/O, 32Gb FC has higher sequential read throughput and sequential write throughput than 16Gb FC, for I/O sizes of 8KB and above. For random I/O, 32Gb FC has higher random read IOPS than 16Gb FC, for I/O sizes of 8KB and above. Virtualized applications that are storage I/O intensive will benefit from the higher sequential and random I/O performance provided by 32Gb FC HBAs.

References

Details for adjusting the queue depth in the PVSCSI adapter in the VM, as well as the LUN and HBA queue depths in the Emulex FC driver on ESXi may be found in the following articles.

[1] VMware, Inc. (2015, December) Large-scale workloads with intensive I/O patterns might require queue depths significantly greater than Paravirtual SCSI default values. http://kb.vmware.com/kb/2053145

[2] Sudhir Balasubramanian. (2015, July) Queues, Queues and more Queues. http://blogs.vmware.com/apps/2015/07/queues-queues-queues-2.html

About the Author

Kinson Ho is a Staff Engineer in Performance Engineering at VMware with a focus on vSphere storage performance, including FC, NFS and iSCSI performance. His interests include the use of packet capture and analysis for performance troubleshooting on 10G networks.

Acknowledgements

The author thanks Michael Lam for assistance with test environment configuration, Jinpyo Kim for advice on performance tuning, Suraj Kasi, Jinpyo Kim, and Priti Mishra for insightful comments about the paper, and Julie Brodeur for technical editing of the paper.

VMware, Inc. 3401 Hillview Avenue Palo Alto CA 94304 USA Tel 877-486-9273 Fax 650-427-5001 www.vmware.com