PERFORMANCE STUDY – APRIL 2017

WHAT'S NEW IN PERFORMANCE?

VMware vSphere 6.5

vmware®

Table of Contents

Introduction

Underlying each release of VMware vSphere are many performance and scalability improvements. The VMware vSphere 6.5 platform continues to provide industry-leading performance and features to ensure the successful virtualization and management of your entire software-defined datacenter.

Management and Measurement

VMware vCenter Server

VMware vCenter Server continues to evolve in both scale and performance in response to the ever-increasing demands of the software-defined datacenter. To meet these increasing requirements, VMware's vCenter Server appliance has been migrated to VMware® Photon OS™, and supporting components, such as the Inventory Service, have been replaced with more performant and scalable services. These significant changes have created the following improvements:

- 2.3x improvement in throughput (operations/minute)
- 2.3x improvement in latency (VM power on times reduced)
- 2x improvement in churn (VMs per hour)
- 2.5x improvement in CPU efficiency
- 15GB reduction in memory footprint for the largest vCenter
- Embedded vCSA (now the preferred platform) performs better than Windows vCenter with additional features such as native vCenter high availability (HA).

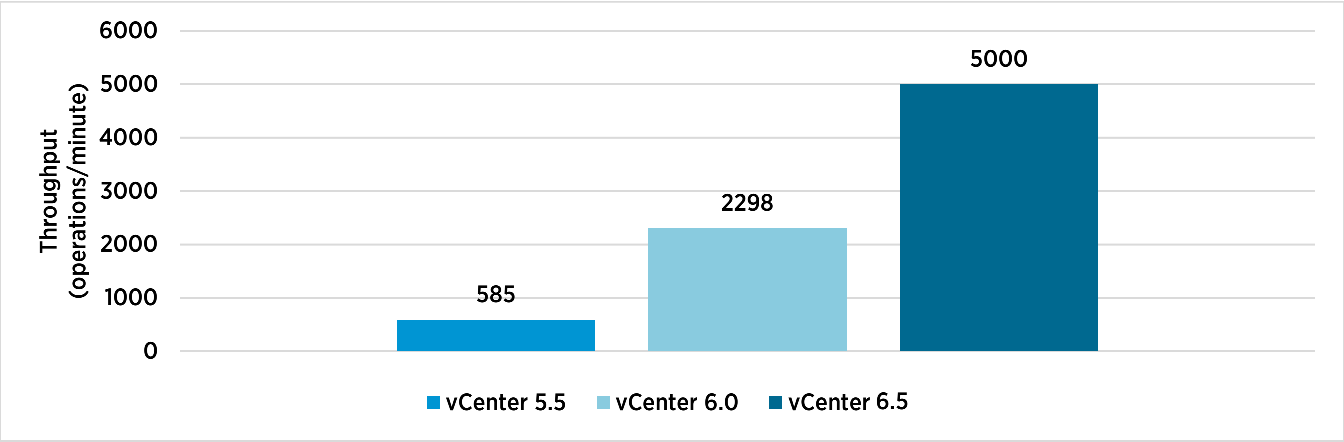

Figure 1. vCenter 6.5 achieves much greater throughput than previous versions (higher is better)

vSphere Web Client (existing vCenter client)

The vSphere Web Client is the primary interface to the vCenter management platform. As such, VMware has placed continued emphasis on improving the experience of using it. The following improvements are highlighted within the vSphere 6.5 release (Table 1).

Action	Improvement in Time to Perform Action
Expand Inventory Tree	93%
Edit VM Settings	56%
View Performance Overview Charts	40%
Refresh VM Power On	50%
VM Migration Wizard: Check Compatibility	67%
VM Migration Wizard: Select Network	33%

Table 1. vSphere Web Client version 6.5 performance improvements over version 6.0

There is also a 30% improvement in the vSphere Web Client's average page response time for single and concurrent browser clients

vSphere Client (new HTML 5 client for vCenter 6.5)

vSphere 6.5 includes a new vSphere Client that is produced in HTML5 instead of Adobe Flex. The HTML5 vSphere Client originated as a VMware Labs fling [1]. The HTML5 vSphere Client outperforms the existing vSphere Web Client, as shown in Figure 1.

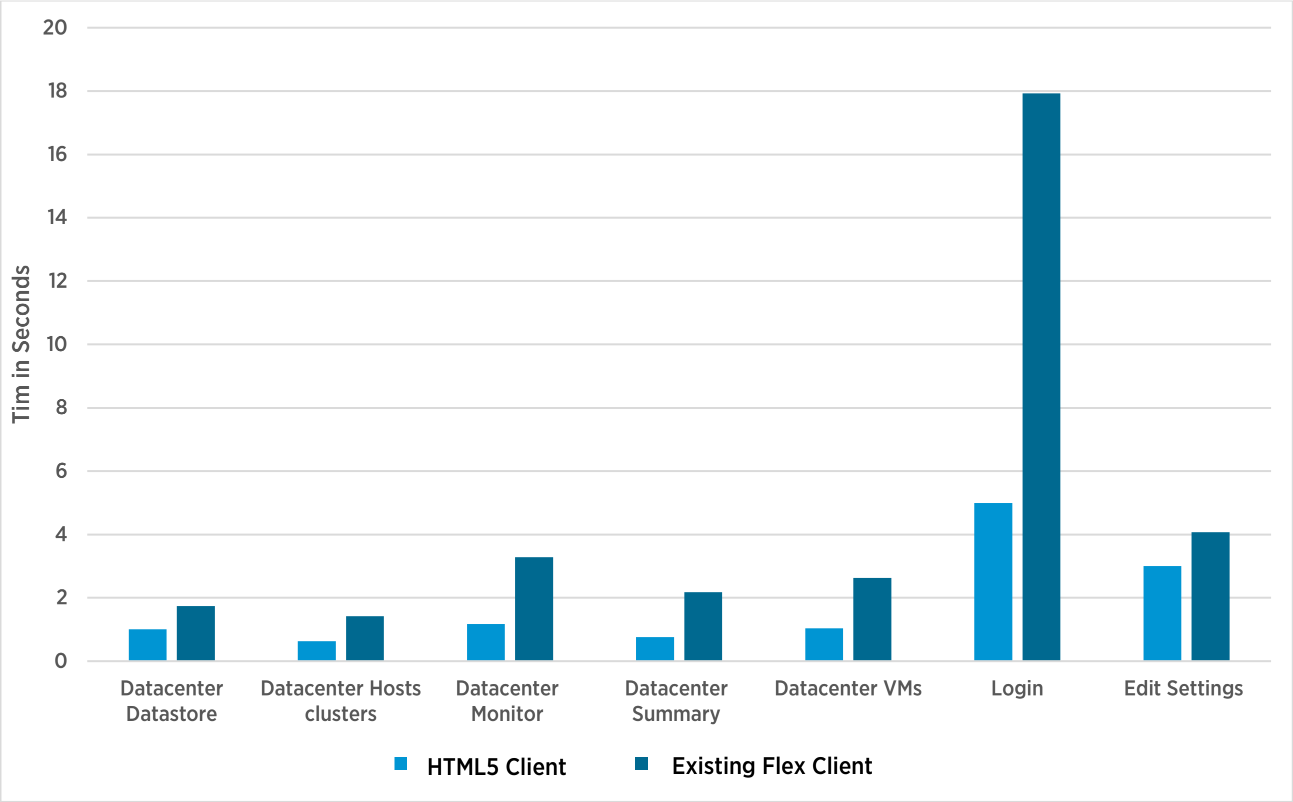

Figure 2. Performance improvements for HTML5 vSphere Client in 6.5 vs. Flex vSphere Client in 6.0 (lower is better)

For an indepth look at performance improvements in vCenter, see vCenter 6.5 Performance: what does 6x mean? [2]

Esxtop

Esxtop in vSphere 6.5 adds a new column **%A/MPERF** in the Power Management screen. The esxtop power screen (obtained by pressing the **p** key) shows power information for the vSphere host. You can add this new column by pressing the **f** key. Use the **%A/MPERF** ratio column to identify at what frequency the processor is currently running.

Aperf and mperf are two hardware registers used to keep track of the actual frequency and nominal frequency of the processor. The ratio tells you how fast the processor is running, which until now was difficult to see or record.

For example, the below screenshot is on vSphere 6.5 running on an Intel Haswell processor, Xeon E5-2699 v3 2.3GHz with up to 3.6GHz Turbo boost. The aperf/mperf ratio of about 150% on CPU 2 and CPU 4 means that the processor is running at around 3.5GHz. For the math, you multiply the percentage by the base rate of your processor. So, 2.3Ghz × 150% = 3.45Ghz, rounded up is approximately 3.5.

Note: We enabled deep C-states like C3 and C6 to get the highest turbo frequency.

```
Power Usage:  238W, Power Cap:  N/A
PSTATE MHZ: 2301 2300 2200 2100 2000 1900 1800 1700 1600 1500 1400 1300 1200

CPU %USED %UTIL %C0 %C1 %C2 %P0 %P1 %P2 %P3 %P4 %P5 %P6 %P7 %P8 %P9 %P10 %P11 %P12 %A/MPERF
  0   0.2   0.3   0   1  98 100   0   0   0   0   0   0   0   0   0    0    0    0    139.4
  1   0.7   0.5   0   3  97  98   0   0   0   0   0   0   0   0   0    0    0    2    140.4
  2 148.9  99.4  99   0   0 100   0   0   0   0   0   0   0   0   0    0    0    0    150.9
  3   0.0   0.0   0   9  91  97   0   0   0   0   0   0   0   0   0    0    0    3    138.9
  4 149.2  99.2  99   1   0 100   0   0   0   0   0   0   0   0   0    0    0    0    150.7
```

Figure 3. Screenshot that includes the new %A/MPERF column (far right)

Core Platform

6TB Virtual RAM

You can now configure virtual machines with up to 6TB of virtual RAM. An ESXi host can now support up to 12TB of physical RAM.

DRS

In vSphere 6.5, VMware has added many new features and several key performance improvements to DRS to ensure more efficient load balancing, faster response times, and simplified cluster management.

- **Higher efficiency:** DRS is now 5x lighter in terms of resource consumption and 3x faster in terms of operational latency.
- **Optimized initial VM placement:** DRS now picks the right host for your VM with very high accuracy, even with several concurrent power-on requests in flight. This means that your VMs are guaranteed to get the best set of resources needed, soon after power-on.
- **Predictive DRS:** DRS now predicts workload changes using vRealize Operations and migrates VMs proactively to avoid any performance impact.
- **Network-aware DRS:** DRS now considers network load along with compute load in the cluster, resulting in more efficient VM placement and load balancing.
- **Pair-wise load balancing:** In addition to its current load balancing, DRS now looks for load imbalance between the least and the most loaded host (pair) in the cluster to try to balance load even more evenly across the cluster.
- **Policies for cluster management:** DRS introduces policies that can help you better manage your cluster load. Some example policies are:
  - Even distribution of VMs: With this policy, DRS tries to distribute VMs evenly across hosts in the cluster, while ensuring that the load balance is not affected.
  - Consumed memory for load balancing - This policy enables DRS to consider consumed memory (instead of active memory) usage for load balancing.

For more information, see vSphere 6.5 DRS Performance [3].

vNUMA

As of vSphere 6.5, changing the `corespersocket` value no longer influences vNUMA and the configuration of the vNUMA topology. The configuration of vSockets and corespersocket only affects the presentation of the virtual processors to the guest OS (generally required for licensing). vNUMA will automatically determine the proper vNUMA topology to present to the guest OS based on the underlying ESXi host

For example, if you create a 4-vSocket virtual machine with 4 corespersocket (total of 16 vCPU) on a dual-socket, 16-core physical ESXi host, prior to vSphere 6.5, vNUMA would have created 4 vNUMA nodes based on the corespersocket setting. As of vSphere 6.5, the guest OS will still see 4 sockets and 4 cores per socket, but vNUMA will now only create 1 vNUMA node for the entire virtual machine since it can be placed in a single physical NUMA node.

This new disconnection of the corespersocket setting with vNUMA allows vSphere to determine the best vNUMA topology automatically in all circumstances.

To revert to the earlier behavior in vSphere 6.0, use and advanced setting set `Numa.FollowCoresPerSocket` to `1`.

For more information, please reference Virtual Machine vCPU and vNUMA Rightsizing – Rules of Thumb [3].

vMotion Enhancements

vSphere 6.5 vMotion introduces Encrypted vMotion that provides the end-to-end encryption of vMotion traffic and protects VM data from eavesdropping occurrences on untrusted networks. Encrypted vMotion provides complete confidentiality, integrity, and authenticity of data transferred over the vMotion network without any additional requirement for specialized hardware.

Encrypted vMotion does not rely on the Secure Sockets Layer (SSL) and Internet Protocol Security (IPsec) technologies for securing vMotion traffic. Instead, it implements a custom encrypted protocol above the TCP layer. To secure VM migration, vSphere 6.5 Encrypted vMotion encrypts all the vMotion traffic, including the TCP payload and vMotion metadata, using the most widely used AES-GCM encryption standard algorithms, provided by the FIPS-certified vmkernel vmkcrypto module.

Encrypted vMotion Duration

Figure 4 shows the vMotion duration in several test scenarios in which we varied vCPU and memory sizes. The figure shows identical performance in all the scenarios with and without encryption enabled on vMotion traffic.

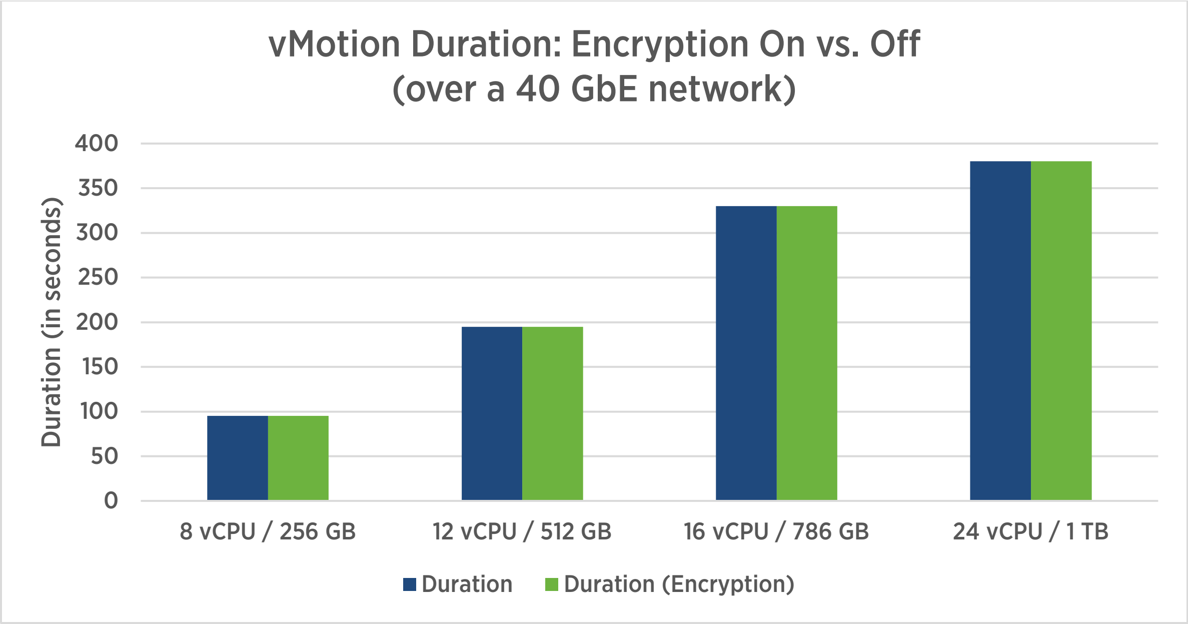

Figure 4. Encryption does not affect vMotion performance over a 40 GbE network

Encrypted vMotion CPU Cost

Figure 5 shows the CPU usage, in terms of the CPU cores, required by vMotion on both source and destinations hosts.

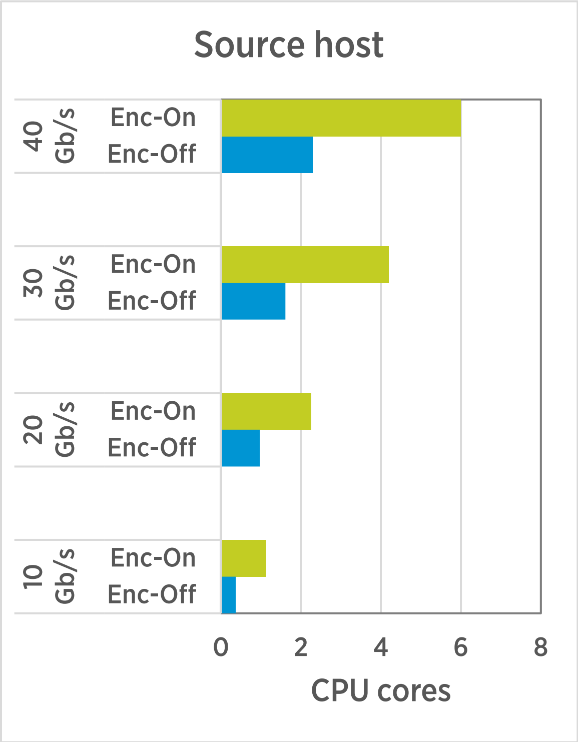

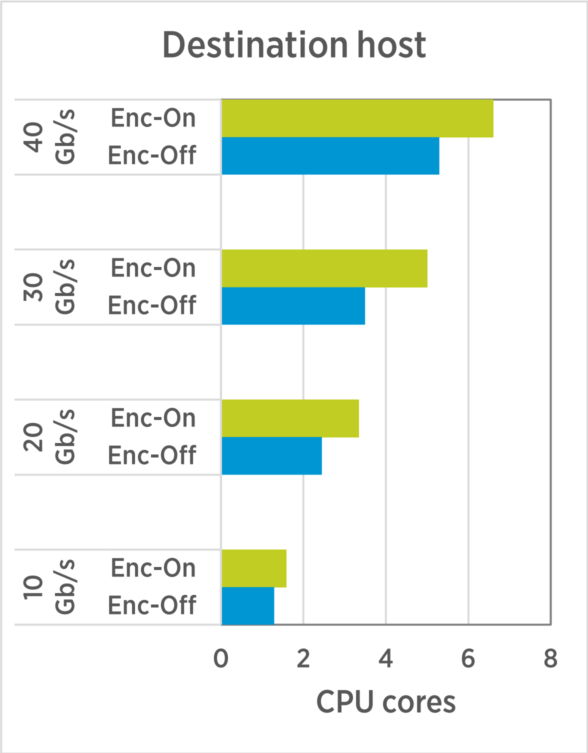

Figure 5. vMotion requires more CPU cores to manage higher amounts of vMotion network traffic

As shown in Figure 5, the CPU requirements of encrypted vMotion are very moderate. For every 10Gb/s of vMotion traffic, vSphere 6.5 Encrypted vMotion requires less than one core on the source host for the encryption-related overheads, and less than half a core on the destination host for the decryption-related overheads.

For the full paper on Encrypted vMotion, see vSphere 6.5 Encrypted vMotion Architecture, Performance, and Best Practices [5].

Storage

Limit Increases

First, ESXi hosts running version 6.5 can now support up to 2,048 paths in total. Second, ESXi hosts running version 6.5 can now support up to 512 devices. This is a 2x increase from previous versions of ESXi where the number of devices supported was limited to 256.

NVMe Adapters

New with vSphere 6.5 is the release of a new NVMe adapter for the virtual machine. vSphere 6.5 supports 4 NVMe adapters per virtual machine with 15 targets each.

Storage I/O Control (SIOC)

VMware engineers designed the Storage I/O Control (SIOC) feature in vSphere 6.5 to differentiate performance between high priority virtual machines and low priority ones. vSphere 6.5 adds two new features for which we tested performance:

- Operation-based policies
- Support for SSD storage

Operation-based policies (in addition to the previously supported share-based policies) are available in SIOC. In addition, now administrators can specify an absolute IOPS number to VMs by assigning the maximum IOPS as a limit and the minimum IOPS as a reservation.

Share-based policies simply indicated which portion of the workload a VM should get. With operation-based policies, administrators can now specify policies based on tiers of service; for example, by silver, gold, and platinum.

IT professionals can also use SIOC in datacenters designed with solid-state drives (SSDs) in their hierarchies. The performance team at VMware has tested SIOC on an SSD system with a variety of workloads.

Here, we take three types of VMs set with operation-based policies of Logging, Analysis, and Production, where each is low, medium, and high priority, respectively.

Type of VM	Priority	Start Time (min)	End Time (min)	# of VMs
Logging	Low	0	40	2
Analysis	Medium	10	40	4
Production	High	20	30	2

Table 2. Setup for three SIOC policies to be tested within the specified time periods

Figure 4 shows the results. From 0-10 minutes, there are only low priority VMs. The SIOC settings always limit the IOPS, so the system workload remains at a low level. In the period of 10–20 minutes and 30–40 minutes, medium priority VMs come into play. This workload fills all the spare IOPS.

During the period of 20–30 minutes, we introduce high priority VMs, and, therefore, the medium priority VMs must give up part of their IOPS to the high performance VMs.

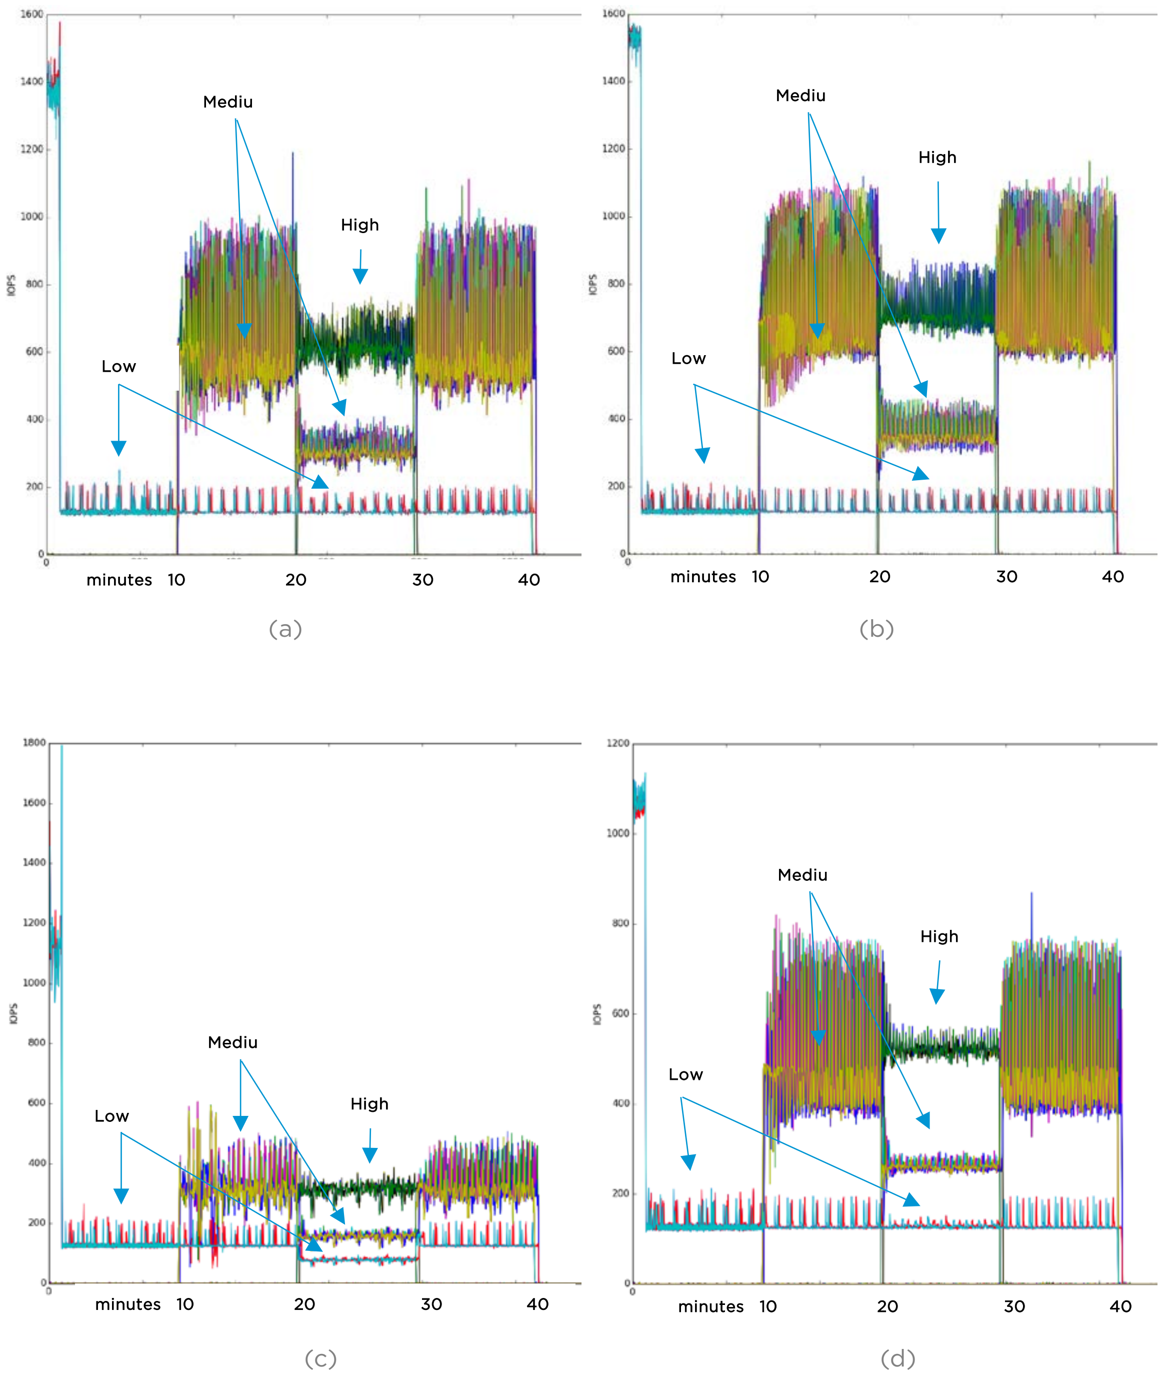

Figure 6. Performance impact of using SIOC in datacenter: (a) sequential write, (b) sequential read, (c) random write, and (d) random read

For more information about SIOC performance, see Performance Implications of Storage I/O-Enabled SSD Datastores in VMware vSphere 6.5 [3].

VMFS6

vSphere 6.5 revises the VMFS file system to VMFS6. Some performance improvements you will notice include faster file creation, device discovery, and device rescanning.

In addition, VMFS6 is now 4K aligned in order to support newer Advanced Format (AF) large-capacity drives in the future, after VMware certifies them. vSphere 6.5 also supports 512 emulation (512e) mode for backward compatibility.

For more information about VMFS6, see "Understanding VMFS Datastores" [4] in the vSphere Storage guide for 6.5 and "VMFS-6" [7] in the core storage whitepaper.

VM Encryption

You can use virtual machine encryption in vSphere 6.5 to create security policies that protect one or more virtual machines and all their associated files. In addition to one or more vSphere 6.5 hosts, vCenter Server and a Key Management Server (KMS) are required. The KMS manages keys with industry-standard technology: the Key Management Interoperability Protocol (KMIP) version 1.1. The encryption actually occurs at the hypervisor level of ESXi and not in the VM. This lets you apply VM encryption policies to any VM, regardless of its guest operating system and datastore type.

VM encryption leverages the latest CPU hardware advances in AES-NI encryption. Advanced Encryption Standard Instruction Set is an extension to the x86 instruction set and provides accelerated encryption and decryption functions on a per-core basis in the CPU.

Performance tests done with Iometer on a Dell PowerEdge R720 with 8 physical cores and 128GB memory show that the process of encryption and decryption does cause some impact on VM performance in the cost of increased CPU cycles. The time spent processing the I/O manifests itself (especially for ultra-low latency storage devices like NVMe) in the form of increased I/O latency, which ultimately also results in a decreased I/O throughput, as shown in Figure 7a. For other types of storage with latencies with a range of a few hundred microseconds (like SSDs), the time spent processing the I/O is a small percentage of the existing I/O path delay. Hence, the overhead incurred due to encryption (and decryption) for such devices is not very noticeable, as shown in Figure 7b.

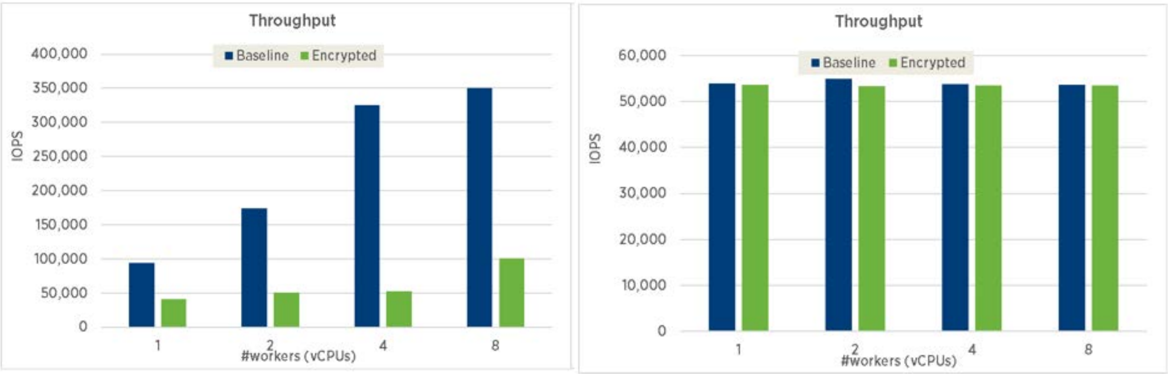

(a) Samsung PM1725 1.5TB NVMe

(b) Intel S3700 400GB SSD

Figure 7. 4KB 100% random write IOPS throughput for NVMe and SSD storage devices (higher is better)

In the case of VM provisioning operations like VM power-on and clone, the overhead of encryption for the Linux VM we tested is less than 20% in the worst case and very minimal for most other cases.

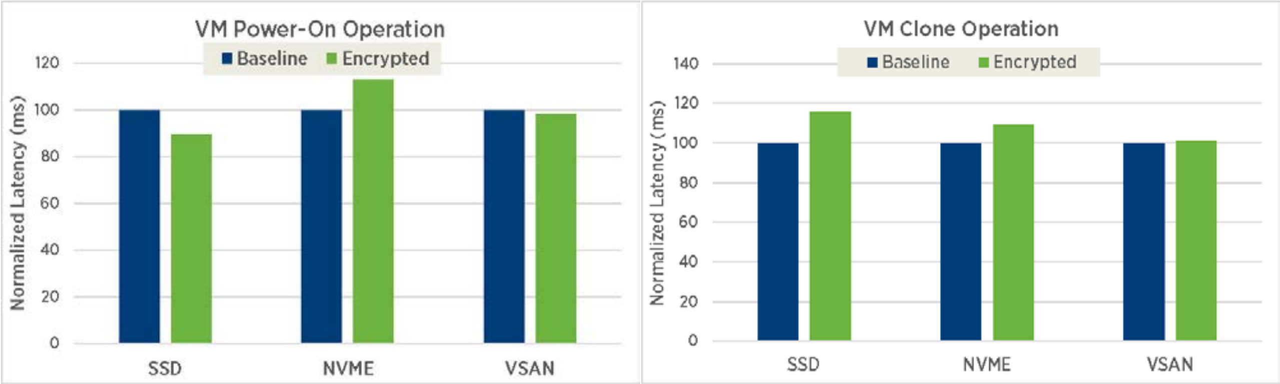

(a) Power-on VM

(b) Clone VM

Figure 8. Performance of powering on VMs and cloning VMs on different storage devices (lower is better)

To learn more about the performance of VM encryption, see vSphere 6.5 Virtual Machine Encryption Performance [8].

Network

100GbE Support

vSphere 6.5 adds support for QLogic FastLinQ QL45604 25/50/100 gigabit Ethernet adapters.

DirectPath I/O

High performance computing (HPC) relies on hardware with lots of CPU and storage. Robust networks are also important to transmit large amounts of data over the wire, typically using technologies like InfiniBand or RCoE (RDMA over converged Ethernet).

vSphere 6.5 achieves parity with bare-metal servers configured with InfiniBand or RoCE in passthrough mode enabled with VMware DirectPath I/O when using the Intel MPI Benchmarks.

This collection of open-source benchmarks performs a set of message passing interface (MPI) performance measurements for point-to-point and global communication operations for a range of message sizes [5]. The generated benchmark data fully characterizes:

- Performance of a cluster system, including node performance, network latency, and throughput
- Efficiency of the MPI implementation used

The PingPong benchmark measures the startup and throughput of a single message sent between two processes.

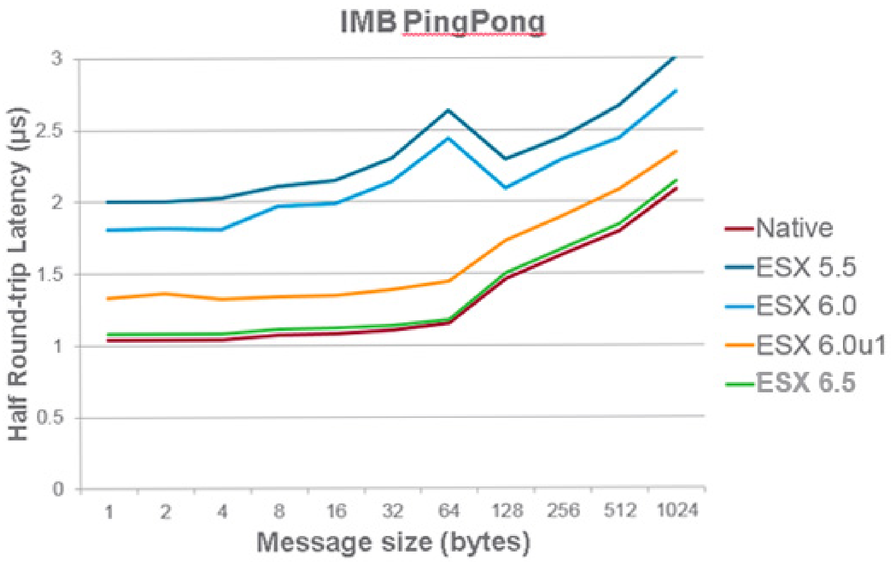

Figure 9. MPI small message InfiniBand latency graph shows the workload run on vSphere 6.5 achieves almost the same round-trip time as that run on a physical machine; latency numbers have also dropped since vSphere 5.5 and 6.0 (lower is better)

Paravirtualzied RDMA (PVRDMA) Adapter

New in vSphere 6.5 is the release of a PVRDMA adapter for virtual machines.

Applications written to leverage RDMA benefit from low latency (<1uS) and high bandwidth (>100GbE), and are now supported under vSphere 6.5 with some limitations [6].

This new adapter exposes a PCIe host to the guest, which helps achieve maximum performance by allowing the hardware to do the data transfer, but maintains enough virtualization to preserve functions like vMotion, snapshots, and memory over-commitment. vSphere 6.5 supports 1 PVRDMA adapter per virtual machine.

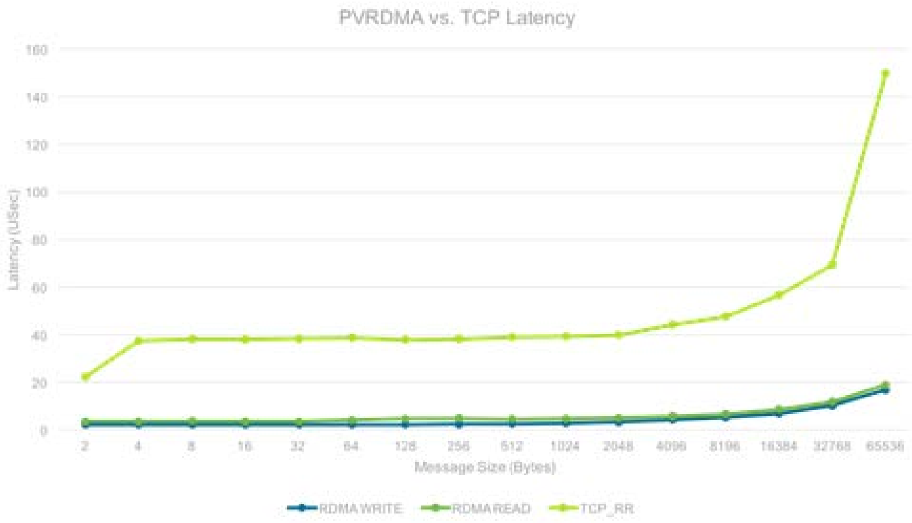

Figure 10. RDMA write and read show lower latencies than that of TCP

Conclusion

Based on these performance, scalability, and feature improvements in vSphere 6.5, VMware continues to demonstrate industry-leading performance.

References

[1] Abhijith Prabhudev , Dennis Lu, and Vishwa Srikaanth. (2017, January) vSphere HTML5 Web Client. https://labs.vmware.com/flings/vsphere-html5-web-client

[2] Ravi Soundararajan. (2017, April) vCenter 6.5 Performance: what does 6x mean? https://blogs.vmware.com/performance/2017/04/vcenter-6-5-performance-what-does-6x-mean.html

[3] Sai Manohar Inabattini, Vikas Madhusudana, and Adarsh Jagadeeshwaran. (2016, November) vSphere 6.5 DRS Performance. http://www.vmware.com/techpapers/2016/drs-vsphere65-perf.html

[4] Mark Achtemichuk. (2017, March) Virtual Machine vCPU and vNUMA Rightsizing – Rules of Thumb. https://blogs.vmware.com/performance/2017/03/virtual-machine-vcpu-and-vnuma-rightsizing-rules-of-thumb.html

[5] Sreekanth Setty. (2017, January) vSphere 6.5 Encrypted vMotion Architecture, Performance, and Best Practices. http://www.vmware.com/techpapers/2017/encrypted-vmotion-vsphere65-perf.html

[6] Yifan Wang. (2017, March) Performance Implications of Storage I/O Control-Enabled SSD Datastores. http://www.vmware.com/techpapers/2017/sioc-ssd-vsphere65-perf.html

[7] VMware, Inc. (2016) Understanding VMFS Datastores. http://pubs.vmware.com/vsphere-65/index.jsp#com.vmware.vsphere.storage.doc/GUID-5EE84941-366D-4D37-8B7B-767D08928888.html

[8] Cormac Hogan. (2016) Core Storage White Paper. https://storagehub.vmware.com/#!/vsphere-core-storage/vsphere-6-5-storage/vmfs-6/1

[9] Aijaz Baig and Sankaran Sivathanu. (2017, March) vSphere 6.5 Virtual Machine Encryption Performance. http://www.vmware.com/techpapers/2016/vm-encryption-vsphere65-perf.html

[10] Intel. (2015) Intel MPI Benchmarks User Guide: Introduction. https://software.intel.com/en-us/node/561887

[11] VMware, Inc. (2016) Remote Direct Memory Access for Virtual Machines. http://pubs.vmware.com/vsphere-65/index.jsp#com.vmware.vsphere.networking.doc/GUID-9AADBB49-876E-4E44-8149-D0523D8ADA6A.html

[12] VMware, Inc. (2016, November) Functionality Updates for the vSphere Client. http://pubs.vmware.com/Release_Notes/en/vsphere/65/vsphere-client-65-html5-functionality-support.html

About the Authors

Mark Achtemichuk currently works as a Staff Engineer within VMware's R&D Operations and Central Services Performance team, focusing on education, benchmarking, collateral, and performance architectures. He has also held various performance focused field, specialist, and technical marketing positions within VMware over the last 6 years. Mark is recognized as an industry expert and holds a VMware Certified Design Expert (VCDX#50) certification, one of less than 250 worldwide. He has worked on engagements with Fortune 50 companies, served as technical editor for many books and publications, and is a sought after speaker at numerous industry events. Mark is a blogger and has been recognized as a VMware vExpert from 2013 to 2016. He is active on Twitter at @vmMarkA where he shares his knowledge of performance with the virtualization community. His experience and expertise from infrastructure to application virtualization helps customers ensure that performance is no longer a barrier, perceived or real, to virtualizing and operating an organization's software-defined assets.

Julie Brodeur is a Senior Technical Writer in the Performance Engineering team at VMware, where she supports the writing activities of a broad organization. She focuses on technical white papers, knowledgebase articles, and blogs about the performance of the vSphere suite of products and related virtualization topics.

Acknowledgments

The authors thank Adarsh Jagadeeshwaran, Aijaz Baig, Josh Simons, Mike Foley, Qasim Ali, Ravi Krishnamurthy, Ravi Soundararajan, Sai Manohar Inabattini, Sankaran Sivathanu, Sreekanth Setty, Vikas Madhusudana, and Yifan Wang for their contributions to and review of this paper.

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

Comments on this document: https://communities.vmware.com/docs/DOC-34263