PERFORMANCE STUDY – MAY 2018

# WHAT'S NEW IN PERFORMANCE?

## VMware vSphere 6.7

vmware®

# Table of Contents

# Introduction

Underlying each release of VMware vSphere® are many performance and scalability improvements. The vSphere 6.7 platform continues to provide industry-leading performance and features to ensure the successful virtualization and management of your entire software-defined datacenter.

# Management

## VMware vCenter Server

vSphere 6.7 delivers an exceptional experience with an enhanced VMware vCenter® Server Appliance™ (vCSA). vSphere 6.7 adds functionality to support not only the typical workflows customers need, but also other key functionality like managing VMware NSX®, VMware vSAN™, VMware vSphere® Update Manager™ (VUM), as well as third-party components.

vSphere 6.7 also delivers phenomenal performance improvements compared to vSphere 6.5:

- **2X faster performance in vCenter operations per second**

  With their benchmark vcbench, VMware performance engineers measured the number of operations per second that vCenter produced. This benchmark stresses the vCenter server by performing typical vCenter operations like power on and off a VM, among several others; for a full list, see Table 1 in VMware vCenter Server 6.0 Performance and Best Practices [1]. vCenter 6.7 performs 16.7 operations per second, which is a twofold increase over the 8.3 operations per second vCenter 6.5 produced.

- **3X reduction in memory usage**

  VMware optimized the core vCenter process (vpxd) to use less memory to complete the same workloads.

- **3X faster DRS-related operations (for example, the latency of powering on a VM)**

  Before vCenter can power on a VM, it first consults several sub-systems, including DRS, to support the initial placement of the VM on a vSphere host. Latency, in this context, is the measure of the duration of this process. VMware made many optimizations in the coordination of these sub-systems to reduce power-on latency from 9.5 seconds to 2.8 seconds.

  For more information about the performance testing and results, see vCenter Performance Improvements from 6.5 to 6.7: What Does 2X Mean? [2]

## vSphere Quick Boot

vSphere Quick Boot is an innovation in conjunction with major server vendors that restarts the VMware ESXi™ hypervisor without rebooting the physical host, skipping time-consuming hardware initialization. If it takes several minutes, or more, for the physical hardware to initialize devices and perform necessary self-tests, then that is the approximate time savings to expect when using Quick Boot. For more information about Quick Boot compatibility, see KB 52477 [3].

# Core Platform

## Binary Translation Deprecation

vSphere 6.5 was the final release that supports binary translation (BT) mode virtualization of operating systems. VMware virtualization technology, introduced in 1998, was founded on a combination of binary translation and direct execution techniques. These software techniques have since been migrated to hardware and are therefore no longer required. In most cases, this will not affect virtual machine performance when you migrate to future vSphere releases, except in some rare instances where old guest operating systems are still being used. For more information, see KB 2147608 [4].

## Processor Deprecation

While ensuring the vSphere 6.7 platform supports the latest generations of processors such as Intel Skylake and Kaby Lake, support for older processors is regularly removed. It's worth noting that vSphere 6.7 no longer supports the following processors:

- AMD Opteron 13xx Series
- AMD Opteron 23xx Series
- AMD Opteron 24xx Series
- AMD Opteron 41xx Series
- AMD Opteron 61xx Series
- AMD Opteron 83xx Series
- AMD Opteron 84xx Series
- Intel Core i7-620LE Processor
- Intel i3/i5 Clarkdale Series
- Intel Xeon 31xx Series
- Intel Xeon 33xx Series
- Intel Xeon 34xx Clarkdale Series
- Intel Xeon 34xx Lynnfield Series
- Intel Xeon 35xx Series
- Intel Xeon 36xx Series
- Intel Xeon 52xx Series
- Intel Xeon 54xx Series
- Intel Xeon 55xx Series
- Intel Xeon 56xx Series
- Intel Xeon 65xx Series
- Intel Xeon 74xx Series
- Intel Xeon 75xx Series

## Host Scalability

There are some minor improvements to vSphere 6.7 ESXi host maximums worth noting.

- Host processor maximums increased from 576 to 768 logical CPUs.
- Host memory maximum increased from 12 TB to 16 TB physical RAM.

## 1 GB Large Memory Pages

Applications with large memory footprints, like SAP HANA, can often stress the hardware memory subsystem (that is, Translation Lookaside Buffer, or TLB) with their access patterns. Modern processors can mitigate this performance impact by creating larger mappings to memory and increasing the memory reach of the application. In prior releases, ESXi allowed guest operating system memory mappings based on 2 MB page sizes. This release introduces memory mappings for 1 GB page sizes.

As shown in the following figure, there is up to 26% improvement in 1 GB memory access performance, compared to the 2 MB page size, through more efficient use of the TLB and processor L1-L3 cache.

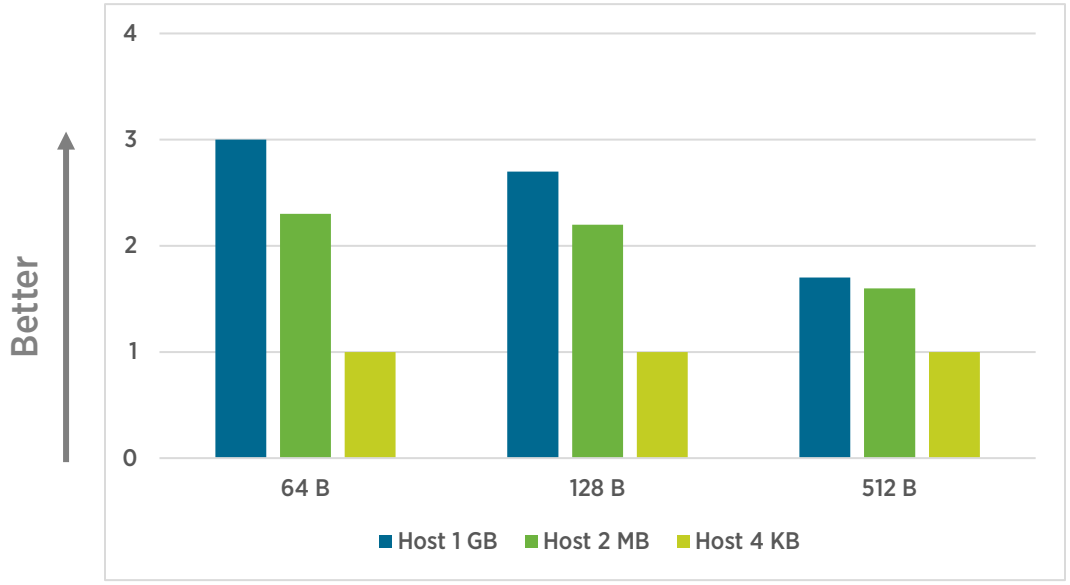


Figure 1. The host with an access size of 1 GB performed 26% better than the host with an access size of 2 MB

## CPU Scheduler Enhancements

Scalability of the vSphere ESXi CPU scheduler is always being improved release-to-release to support current and future requirements. New in vSphere 6.7 is the elimination of the last global lock, which allows the scheduler to support tens of thousands of worlds (various processes running in the VMkernel; for example, each virtual CPU has a world associated with it). This feature ensures vSphere maintains its lead as a platform for containers and microservices.

## Virtual Per VM EVC

vSphere previously implemented Enhanced vMotion Compatibility (EVC) as a cluster-wide attribute because, at the cluster-wide level, you can make certain assumptions about migrating a VM (for example, even if the processor is not the same across all ESXi hosts, EVC still works). However, this policy can cause problems when you try to migrate across vCenter hosts or vSphere clusters. By implementing per-VM EVC, the EVC mode becomes an attribute of the VM rather than the specific processor generation it happens to be booted on in the cluster.

Per-VM EVC allows you the granularity of enabling EVC for particular VMs, rather than for a cluster of hosts. This improves the mobility of a VM beyond a cluster with compatibility for returning. It also improves compatibility for cross-VC migrations without you having to know complicated stretch-cluster concepts.

- Low Mode approximately means more compatible
- High Mode approximately means more performant

## Virtual Hardware 14

Virtual Hardware 14 additionally supports:

- Persistent memory, with a maximum of:
  - 1 NVDIMM controller per VM
  - 64 NVDIMMS per VM
  - 1 TB non-volatile memory per VM
- Virtual Trusted Platform Module (vTPM) – VMware created a new, vTPM 2.0 device to enable Microsoft Virtualization-based Security (VBS). For more information, see Virtualization-based Security (VBS), later in this document.

## Persistent Memory (PMEM)

Persistent memory (PMEM) is a type of non-volatile DRAM (NVDIMM) that has the speed of DRAM but retains contents through power cycles. It is a new layer that sits between NAND flash and DRAM, providing faster performance. It's also non-volatile unlike DRAM.

vSphere 6.7 supports two modes of accessing persistent memory:

- vPMEMDisk – presents NVDIMM capacity as a local host datastore which requires no guest operating system changes to leverage this technology.
- vPMEM – exposes NVDIMM capacity to the virtual machine through a new virtual NVDIMM device. Guest operating systems use it directly as a block device or in DAX mode.

The following chart shows the result of a performance test run using the MySQL benchmark of Sysbench. The benchmark measures the throughput and latency of a MySQL workload. Here, we ran the tests with 3 tables, 9 threads, and an 80-20 read-write ratio. We ran the MySQL server in a VM hosted on vSphere 6.7.

The blue bars show throughput measured in transactions per second. The green line shows latency, measured as the 95$^{th}$ percentile in milliseconds.

We observe that virtual PMEM can improve performance by up to 1.8x better throughput and 2.3x better latency over standard SSD technology.

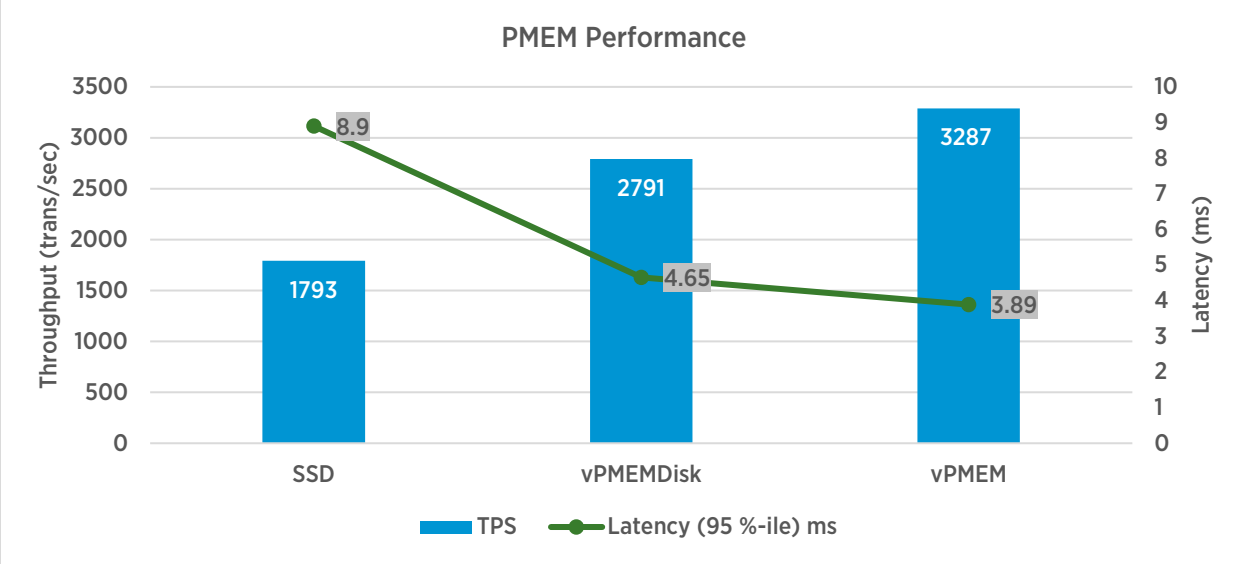


Figure 2. Virtual PMEM can improve performance by up to 1.8x better throughput and 2.3x better latency over standard SSD technology

## Virtualization-based Security (VBS)

Microsoft VBS, a feature of Windows 10 and Windows Server 2016 operating systems, uses hardware and software virtualization to enhance system security by creating an isolated, hypervisor-restricted, specialized subsystem. Starting with vSphere 6.7 and Virtual Hardware 14, you can enable Microsoft virtualization-based security (VBS) on supported Windows guest operating systems.

VMware engineering made a number of vSphere features and enhancements to performantly support VBS-enabled virtual machines.

To measure the performance of a vSphere 6.7 virtual machine running Windows with VBS enabled, we used HammerDB [5], which is a benchmarking application. The test simulated 22 virtual users generating an OLTP workload that wrote to a Microsoft SQL Server 2016 database. This workload was like TPC-C [6].

The testbed included:

- vSphere 6.7 virtual machine:
  - Microsoft Windows 2016 guest operating system
  - VBS enabled
  - 8 vCPUs
  - 12 GB RAM
- vSphere 6.7 host
  - 2x E5-2699 v3 @ 2.3 GHz "Haswell" processors

As shown in the following figure, these engineering efforts resulted in a 33% improvement in transactions per minute.

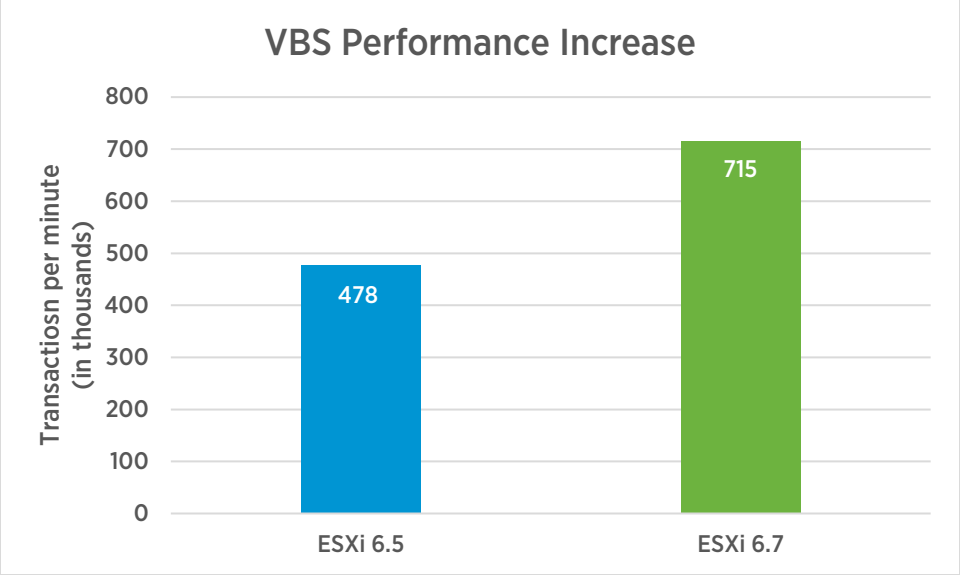


Figure 3. ESXi 6.7 shows a 33% improvement in performance with VBS – higher is better

## Instant Clone (Parentless)

You can use Instant Clone technology to create powered-on virtual machines from the running state of another powered-on virtual machine. The result of an Instant Clone operation is a new virtual machine that is identical to the source virtual machine. With Instant Clone, you can create new virtual machines from a controlled point in time. Instant cloning is very convenient for large-scale application deployments because it ensures memory efficiency and allows for creating numerous virtual machines on a single host.

In the following figure, the Y-axis shows the total completion time of 64 clones in seconds. Linked clones are 128; Instant clones are 44.

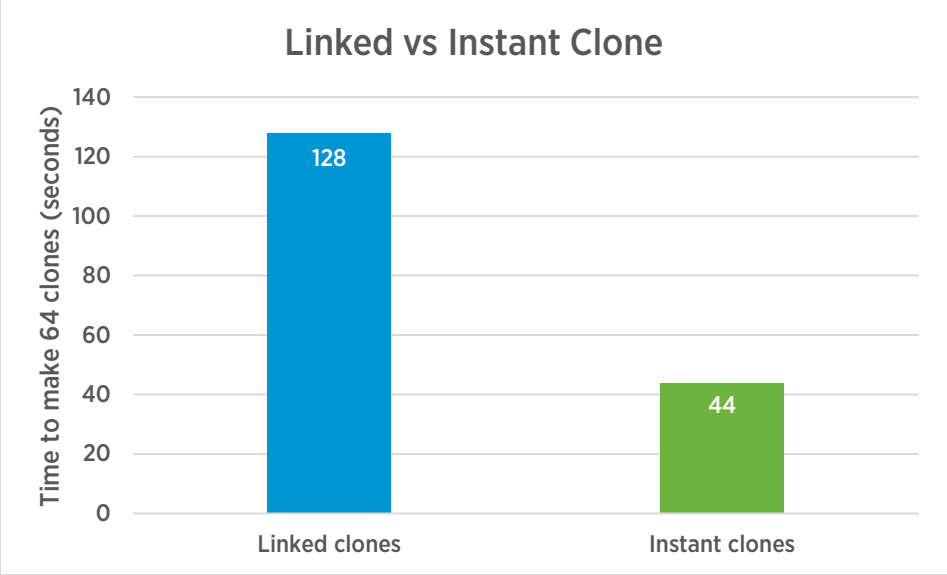


Figure 4. Instant clone performance (fully clone and boot 64 clones) – lower is better

The time to fully deploy and boot 64 clones using vSphere 6.7 Instant Clone showed approximately **2.8x** improvement over the older, Linked Clone architecture.

# Storage

## Maximums

There are some improvements to vSphere 6.7 ESXi host storage maximums worth noting.

- Fibre Channel devices increased from 512 devices per host to 1024.
- Fibre Channel paths increased from 2048 paths per host to 4096.
- The number of virtual disks per PVSCSI controller has been increased from 16 disks to 64, which now allows for 256 disks per virtual machine.

## 4Kn

Similar to firmware in 512e drives, this software emulation layer exposes 512 sectors to the guest operating system while using 4Kn drives as local storage in a server. This approach enables running legacy operating systems, applications, and existing VMs on servers with 4Kn HDD drives. This support also gives you the flexibility of using new 4Kn hardware on your vSphere 6.7 host, but still run older applications in the host's virtual machines.

To learn more about this new architecture in vSphere 6.7, see Support for 4Kn HDDs [7].

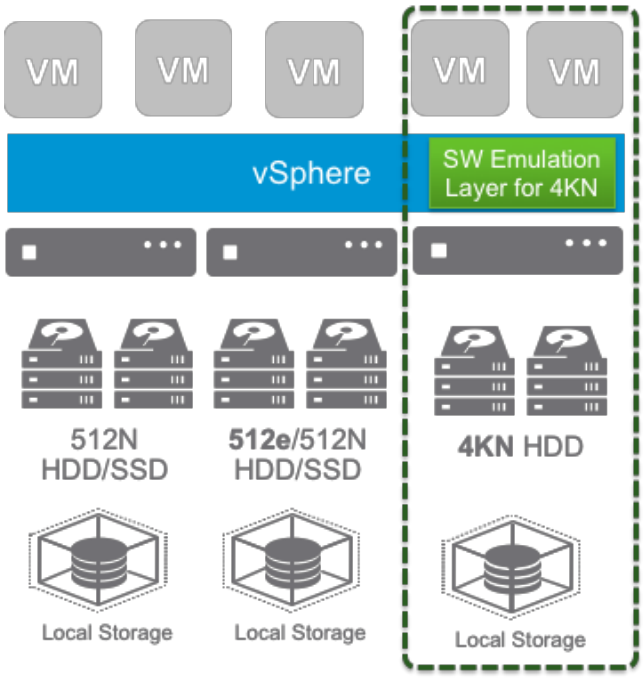


Figure 5. vSphere 6.7 includes software emulation for 4Kn drives

There was no perceivable difference in performance in random I/Os, reads or writes, irrespective of the percentage of 4 KB alignments in the I/O blocks. This test was focused on random I/O because, generally, an ESXi host with many virtual machines requires the storage to be optimized for a random I/O stream.

# Network

## vmxnet

The vSphere 6.7 release includes vmxnet3 version 4, which supports some new features.

- **RSS for UDP** - Receive side scaling (RSS) for the user data protocol (UDP) is now available in the vmxnet3 v4 driver. Performance testing of this feature showed a 28% improvement in receive packets per second. The test used 64-byte packets and four receive queues.
- **RSS for ESP** – RSS for encapsulating security payloads (ESP) is now available in the vmxnet3 v4 driver. Performance testing of this feature showed a 146% improvement in receive packets per second during a test that used IPSec and four receive queues.
- **Offload for Geneve/VXLAN** – Generic network virtualization encapsulation (Geneve) and VXLAN offload is now available in the vmxnet3 v4 driver. Performance testing of this feature showed a 415% improvement in throughput in a test that used a packet size of 64 K with eight flows.

## Native Driver Support for Mellanox Hardware

vSphere 6.7 includes native drivers for the following Mellanox host channel adapters that support 40 and 100 gigabit per second speeds over InfiniBand, iWARP, and RoCE:

- ConnectX-3 EN
- ConnectX-3 Pro with VPI
- ConnectX-4 with VPI
- ConnectX-5 EN

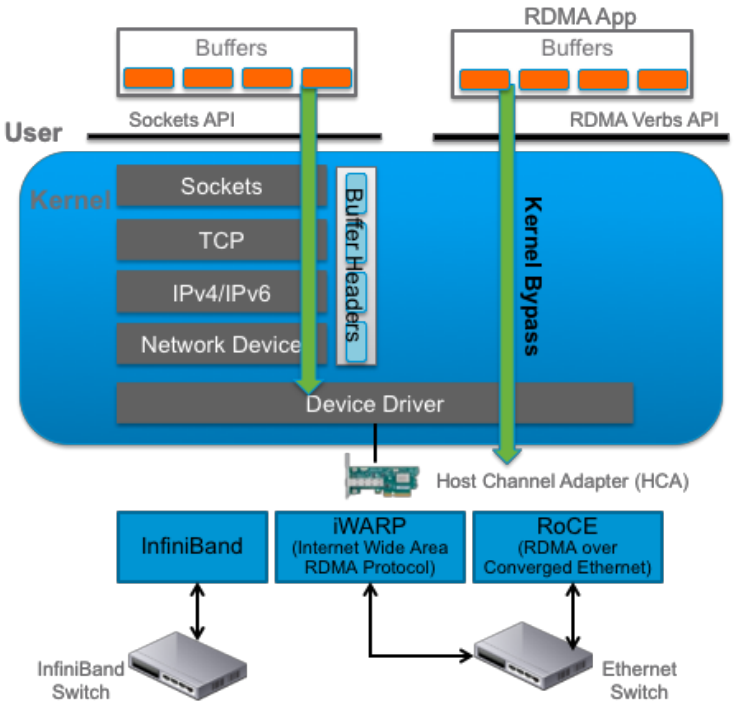


Figure 6. vSphere 6.7 includes a native driver for Mellanox HCAs

# Conclusion

Based on these performance, scalability, and feature improvements in vSphere 6.7, VMware continues to demonstrate industry-leading performance.

# References

[1] Mike Stunes and Ravi Soundararajan. (2015) VMware vCenter Server 6.0 Performance and Best Practices.
https://www.vmware.com/content/dam/digitalmarketing/vmware/en/pdf/techpaper/vmware-vCenter6-perf.pdf

[2] Ravi Soundararajan. (2018, May) vCenter Performance Improvements from 6.5 to 6.7: What Does 2X Mean?
https://blogs.vmware.com/performance/2018/05/vcenter67-perf-improvements.html

[3] VMware. (2018, April) Quick Boot Compatibility.
https://kb.vmware.com/kb/52477

[4] VMware. (2016, November) Deprecation of Hypervisor Binary Translation Mode.
https://kb.vmware.com/s/article/2147608

[5] HammerDB. (2018) HammerDB.
http://www.hammerdb.com/index.html

[6] TPC. (2018) TPC-C.
http://www.tpc.org/tpcc/

[7] VMware, Inc. (2018) Support for 4Kn HDDs.
https://storagehub.vmware.com/t/vsphere-storage/vsphere-6-7-core-storage-1/support-for-4kn-hdds/

## 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 500 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 the product managers and performance engineers who contributed to and reviewed this paper.

**VMware, Inc. 3401 Hillview Avenue Palo Alto CA 94304 USA Tel 877-486-9273 Fax 650-427-5001 www.vmware.com**
Copyright © 2018 VMware, Inc. All rights reserved. This product is protected by U.S. and international copyright and intellectual property laws. VMware products are covered by one or more patents listed at http://www.vmware.com/go/patents. VMware is a registered trademark or trademark of VMware, Inc. in the United States and/or other jurisdictions. All other marks and names mentioned herein may be trademarks of their respective companies.
Comments on this document: https://communities.vmware.com/docs/DOC-34263