What's New in VMware vSphere 6 - Performance

VMware vSphere 6

TECHNICAL MARKETING DOCUMENTATION
V 1.0/JULY 2015

Table of Contents

Introduction

VMware vSphere® 6 continues to enhance the performance features and capabilities of the vSphere platform, making it the most robust and highest-performing cloud platform. vSphere 6 supports larger virtual machines and physical hosts to accommodate even the most demanding workloads. It also introduces several new features that reduce latency and increase throughput for network, storage, and compute.

This paper first looks at the improvements made to VMware vCenter Server™, then to the core platform, storage, and network.

VMware vCenter Server

As the scale and demands of the vSphere management platform continue to grow, manageability has become just as important as the day-to-day performance of applications. The improvements made in the vSphere 6 Web Client and vCenter Server backend provide the flexibility to manage and scale the largest virtualization environments.

vSphere Web Client Performance Improved

Changes to the vSphere Web Client dramatically improve the performance and usability of this tool, with some user actions performing 5 times faster than they did in vSphere 5.5. Major areas of improvement include the login and home pages, summary pages, networking pages, related objects lists, general navigation, performance charts, and action menus (right click). In addition, non-essential data retrieval was reduced, which also serves to lighten the load on vCenter Server [1].

The net result is that the vSphere 6 Web Client is an entirely new experience and easier to use than previous versions of the vSphere Web Client.

vCenter Server Appliance Maximums Raised for vPostgres Database

The vCenter Server Appliance (VCSA) with an embedded vPostgres database now supports up to 1,000 hosts and 3,000 virtual machines. In vSphere 5.5, support was limited to small deployments of 100 hosts and 3,000 virtual machines.

vCenter Server Performance Improvements

vCenter Server 5.5 and 6.0 performance was tested at different inventory sizes (Table 1) for a heavy workload of typical vCenter Server operations against a Microsoft SQL Server 2012 database. Performance in operational throughput (measured in orders per minute) was increased for all inventory sizes: 63% for a small inventory, 73% for a medium inventory, and 67% for a large inventory. This is shown in Figure 1.

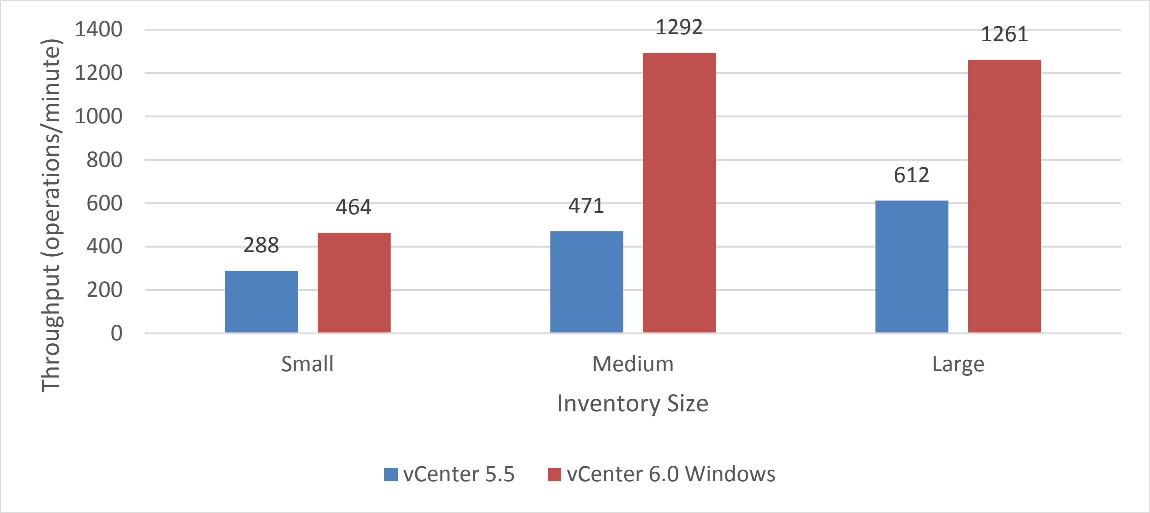

Figure 1. vCenter throughput at several inventory sizes, with heavy workload (higher is better)

In addition, operational latency (measured in seconds taken for the workload to complete) dropped by 60% for a small inventory, 71% for a medium inventory, and 84% for a large inventory. See Figure 2, below.

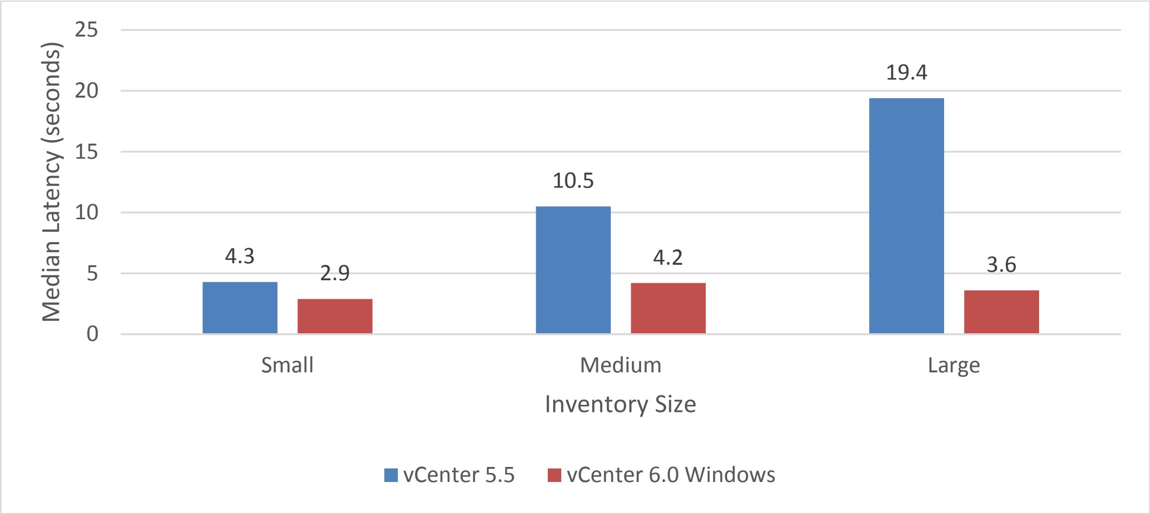

Figure 2. vCenter median latency at several inventory sizes, with heavy workload (lower is better)

Inventory Size	Hosts	VMs
Small	100	1000
Medium	400	4000
Large	1000	10,000

Table 1. Number of hosts and VMs tested in different inventory sizes

vCenter Server Scalability Enhancements

Table 2 shows the increased number of hosts and virtual machines that vCenter Server now supports. Growth in both areas means vCenter Server can support and manage larger clusters and larger VMs running more demanding workloads.

vSphere 5.5	vSphere 6.0
32 Hosts per Cluster	64 Hosts per Cluster
4000 VMs per Cluster	6,000 VMs per Cluster

Table 2. vCenter Server scalability enhancements

Core Platform

vSphere's core platform for the software-defined datacenter is continually being improved to ensure it remains the highest performing platform to host all your enterprise-level applications. Changes for vSphere 6.0 include support for even more vCPUs and hosts, allowing you to consolidate more applications on a host by scaling up (adding vCPUs to a single VM) or scaling out (adding more VMs). Other additions to the core platform pave the way for future features (instant cloning) and underpin new graphics functionality.

128-way vCPU and 4TB Virtual RAM

vSphere 6 introduces the ability to run "monster" virtual machines with up to 128 virtual CPUs (vCPUs) and 4TB of RAM. This doubles the number of vCPUs supported from the previous version and increases the amount of RAM by four times. This new capability allows you to run larger workloads than ever before in a virtual machine [2]. Table 3 shows the differences between vSphere 5.5 and vSphere 6.0.

	vSphere 5.5	vSphere 6.0
VM vCPUs	64	128
VM vRAM	1TB	4TB
Host CPU	320	480
Host RAM	4TB	12TB
VMs per Host	512	1,024

Table 3. Updates for memory and compute for VMs

Performance engineering tested monster virtual machine scalability using the DVD Store 2.1 benchmark workload against both Microsoft SQL [3] and Oracle databases [2].

In one example, testing showed the number of orders per minute (OPM) that could be transacted against an Oracle database running in monster virtual machines of 60 vCPUs and 120 vCPUs. The 60-vCPU virtual machine had affinity set to only two sockets and the 120-vCPU virtual machine had its affinity set to all four sockets. Figure 3 shows that both virtual machines can handle a large number of OPMs, and the scaling is linear from the 60-vCPU virtual machine to the 120-vCPU virtual machine. This indicates good scaling—as the vCPUs are increased in a VM, the VM can handle a greater number of OPMs in an expected growth path.

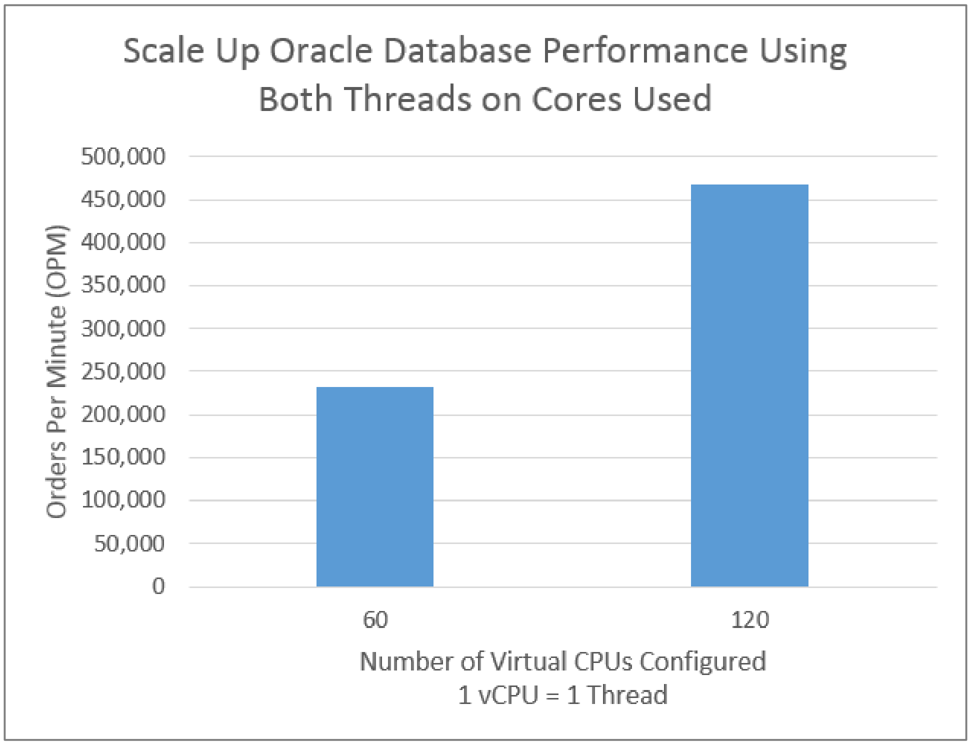

Figure 3. Orders per minute achieved when scaling up monster Oracle VMs

Memory Hot-Add and vNUMA improvements

When a vNUMA virtual machine with the hot-add memory option is enabled and memory is hot-added to it, that memory is now allocated equally across all NUMA regions. In previous releases, all new memory was allocated only to region 0. This enhancement ensures that all regions benefit from the increase in RAM, enabling the virtual machine to scale without requiring any downtime [4].

Instant Clone

Instant cloning, previously known as VMFork, is a part of vSphere 6.0 infrastructure that lays the foundation to rapidly clone and deploy virtual machines within seconds, as much as 10 times faster than cloning a virtual machine in the traditional way.

A new virtual machine is "instantiated" by forking off an existing virtual machine in both memory and disk space, whereby reads come from the source machine but changes or writes are placed into files dedicated to the new virtual machine. This allows for new instantaneous cloning because data is not being copied. Note that the child virtual machine starts where the parent virtual machine left off, so there is no boot process—it is instant on.

Instant cloning is able to provision up to 64 VMs on a single host in under six seconds, making it nearly 13 times faster than prior vSphere releases. While initially being targeted for the VMware Horizon suite, this technology and performance will be valuable to many things including containers within virtual machines.

Note: In vSphere 6.0, instant cloning is achieved through private APIs and is not currently available through the vSphere Web Client.

NVIDIA GRID vGPU Functionality

ESXi 6.0 hosts that are installed with an NVIDIA GRID K1 or K2 physical GPU card can share this card among its virtual machines when NVIDIA GRID vGPU technology is enabled. Virtual machines running Horizon 6 remote desktops can take advantage of this accelerated rendering. Highly graphics-intensive applications benefit from this solution, and the performance level is close to that of a dedicated graphic solution (like vDGA).

Storage

The flexibility of software-defined storage demands that performance of the I/O stack is optimized for the fastest and lowest latency storage in the market.

VMware vSphere Storage Stack Optimizations

Changes were made at the VMkernel level to reduce overhead and increase capabilities to best leverage flash storage.

For NVMe SSDs, the NVMe device driver in vSphere 6.0 was enhanced to create completion threads per socket. This translates into greater I/Os per second (IOPS). For example, Samsung NVMe in vSphere 5.5 drove 240,000 IOPS per second. Now, in vSphere 6.0 that same device drives 710,000 IOPS. Devices of this type are in the process of being certified.

For PCIe SSDs, deferred I/O processing pools were implemented in the VMkernel without modifying the device driver. This will also benefit all flash SAN architectures. For example, an EMC XtremSF device previously drove 200,000 IOPS, and now drives 670,000 IOPS.

In both cases, vSphere 6.0 is driving IOPS at near native "bare-metal" performance, achieving extremely high levels, while deployments today rarely require more than 200,000 IOPS per host.

Virtual Volumes (VVOLs)

VMware vSphere® Virtual Volumes™ is a new integration and management framework that virtualizes SAN/NAS arrays, enabling a more efficient operational model that is optimized for virtualized environments and centers the storage usage model around the application instead of the infrastructure [5].

Unlike LUNs and volumes, which need to be manually changed to meet the growth demands of virtual machines, VVOLs have a more scalable approach because they are "elastic" and can grow or shrink as needed according to support the dynamic disk needs of their virtual machines.

The performance of some storage functions are greatly enhanced while using VVOLs. For example, the performance of a virtual machine with a snapshot is up to 3-4 times better, and the time it takes to create full clones is improved as they are offloaded to the arrays. This varies slightly based on the storage manufacturer and protocol selected. Otherwise, performance is the same as it was previously, but now you have the added flexibility and vSphere integration that VVOLs provides.

Virtual SAN (VSAN) Performance Improved

Among many significant improvements, performance has been dramatically improved for both hybrid and newer all-flash configurations [6]. Table 4 displays the improvements: key areas include an increase in the supported number of hosts in a VSAN cluster, and an increase in IOPS per host.

	VSAN 5.5	VSAN 6.0 Hybrid	VSAN 6.0 All Flash
Hosts/cluster	32	64	64
Number of VMs per host	100	200 2X Improvement over VSAN 5.5	200 2X Improvement over VSAN 5.5
IOPS per host– 70% Read 30% Write	20K	40K 2X Improvement over VSAN 5.5	90K 4.5X Improvement over VSAN 5.5
IOPs per host– 100% Read	60K	120K 2X Improvement over VSAN 5.5	Pending
IOPS on 32 node cluster 100% Read	2 million IOPS	4 Million IOPS 2X Improvement over VSAN 5.5	Pending
IOPS on 64 node cluster 100% Read	N/A	7 Million IOPS 3.5X Improvement over VSAN 5.5	Pending
Latency/response times (DVD Store)	~3.5 ms	< 2.0 ms 40% Improvement over VSAN 5.5	Sub-millisecond with consistently high performance
Max virtual disk size	2TB	62TB 31X greater than VSAN 5.5	62TB 31X greater than VSAN 5.5
Snapshots (DVD Store)	~2 snapshots	Up to 32 snapshots 16X greater than VSAN 5.5 ~20% improvement in performance	Pending

Table 4. Virtual SAN improvements from version 5.5 to 6.0.

Network

vSphere 6 adds a number of powerful new features and enhancements to the networking capabilities of the vSphere platform. These new features enable you to effectively protect and consume bandwidth as well as automatically tune network performance.

Network I/O Control v3 (with reservations)

In vSphere 6.0, VMware has further built on Network I/O Control features to deliver more predictable bandwidth. The goal of introducing these changes has been to allow tighter control on the network resources available to different classes of traffic, irrespective of the traffic originating from other classes of traffic on the host [7]. Here are the key enhancements that Network I/O Control provides in vSphere 6.0:

- Bandwidth reservations for classes of traffic: You can specify the minimum bandwidth that must be reserved for a class of traffic. This guarantees that the bandwidth to the same class of traffic never falls below the specified threshold. As an example, if virtual machine traffic is dedicated 5Gbps of bandwidth, then the combined virtual machine traffic is always guaranteed 5Gbps of bandwidth, irrespective of traffic from other classes of service such as VMware vSphere ® vMotion®, NFS, VMware vSphere® Fault Tolerance (FT), and VMware Virtual SAN™.
- Bandwidth reservations for VMs: Network I/O Control also allows the ability to provide bandwidth reservations to each VM virtual adapter (vNIC). This allows for dedicated bandwidth reservations at a per-VM granularity. Network I/O Control also allows you to create abstract network resource pools that can be attributed to a port group of a distributed virtual switch (DVS). Bandwidth reserved for a resource pool is available only to virtual NICs that are part of the port group associated with the resource pool.
- Load balancing: This feature allows VMware vSphere® Distributed Resource Scheduling™ (DRS) to migrate virtual machines within a cluster of vSphere hosts to accommodate bandwidth reservations assigned to virtual machine ports. This powerful feature allows you to assign bandwidth reservations to virtual machines without worrying about hitting the reservation limit in a single host of the cluster.

Host-Wide Performance Tuning Engine

A new feature in vSphere 6.0, host-wide performance tuning (also called dense mode), allows you to optimize individual ESXi hosts for very high consolidation ratios (many virtual machines per host) where keeping average response times low is more important than the usual action of ESXi, which is keeping all response times low. Dense mode is important in use cases where a cluster of virtual machines on the ESXi host are sharing a job, such as that in a Web farm or other distributed computing environments where sub-millisecond latencies are not required.

When dense mode is enabled, ESXi monitors the number of virtual machines, number of vCPUs, and total CPU utilization; when all three exceed certain thresholds, the dense mode optimizations are implemented.

The dense mode optimizations essentially batch packets more aggressively at various points. This primarily accomplishes two things:

- Reduces hypervisor overhead for packet processing—lower hypervisor overhead leaves more CPU resources for virtual machines. This is beneficial because when CPU utilization is high, average response time is greatly influenced by the availability of spare CPU resources.
- Reduces the number of interruptions to virtual machine execution—reducing the number of interruptions can increase the efficiency of the virtual machines.

This feature has demonstrated up to 10% higher consolidation ratios.

For more information, see *Performance Best Practices for vSphere 6.0* [8].

VMXNET3 LRO Support

Large Receive Offload (LRO) is now supported in VMXNET3 virtual NICs for virtual machines running Microsoft Windows 2012 and Windows 8 as guest operating systems. LRO is a technique to reduce the CPU time for processing TCP packets that arrive from the network at a high rate. LRO reassembles incoming packets into larger ones (but fewer packets) to deliver them to the network stack of the system. LRO processes fewer packets, which reduces its CPU time for networking. Throughput is also correspondingly improved because more CPU is available to deliver additional traffic.

Performance testing shows that LRO considerably increases CPU efficiency by 86% for a message size of 256 bytes, 25% for a message size of 16KB, and 33% for a message size of 64KB (Figure 4).

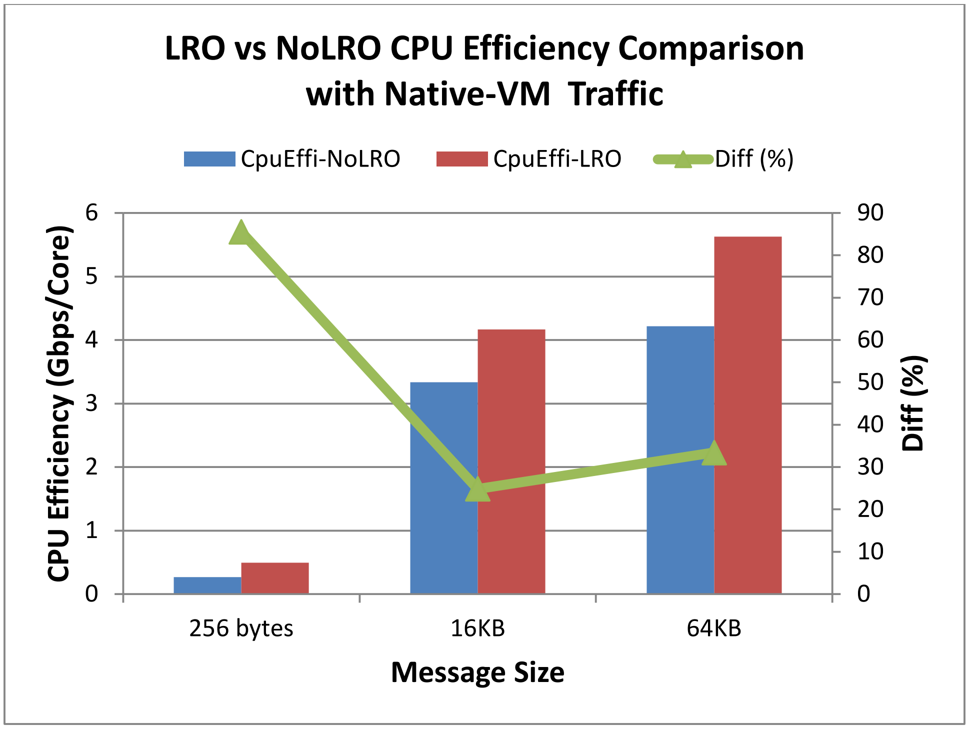

Figure 4. CPU efficiency comparison with and without LRO with Native-VM traffic

LRO is enabled by default for VMXNET3 adapters on vSphere 6.0 hosts, but you must set it to be enabled globally in the guest for Windows 8 virtual machines. For more information about configuring this, see the VMware vSphere 6.0 documentation [9].

vMotion Over Long Distances

vMotion in vSphere 6.0 delivers breakthrough new capabilities that offer you a new level of flexibility and performance in moving virtual machines across their virtual infrastructures. Included with vSphere 6.0 vMotion are features—long-distance migration (supporting vMotion over 100ms latency links), cross-vCenter migration (supporting vMotion between vCenter instances), routed vMotion (supporting vMotion across different IP subnets)—that enable seamless migrations across current management and distance boundaries. For the first time ever, VMs can be migrated across vCenter Servers separated by cross-continental distance with minimal performance impact.

vMotion is fully integrated with all the latest vSphere 6 software-defined data center technologies including Virtual SAN (VSAN) and Virtual Volumes (VVOL). Additionally, the newly re-architected vMotion in vSphere 6.0 now enables extremely fast migrations at speeds exceeding 60 gigabits per second [10].

40GbE Improvements

As the consumption of bandwidth grows and the economics of 40GbE interfaces drop, vSphere must efficiently use these new capabilities [11]. As a result, the following new features have been implemented.

Single vNIC Performance Improvements

vSphere 6.0 includes improvements to the VMXNET3 virtual NIC (vNIC) that allows a single vNIC to achieve line-rate performance with 40GbE physical NICs.

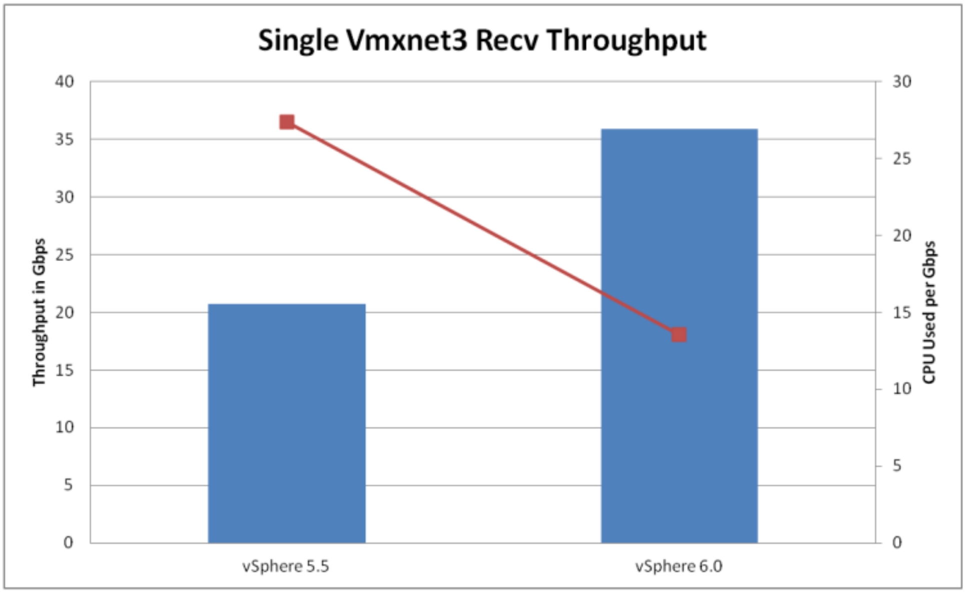

Figure 5. Comparison of throughput and CPU usage for 40GbE with VMXNET3 on vSphere 5.5 and 6.0

vSphere 6.0 can achieve more than 35Gbps of throughput as compared to 20Gbps achieved in vSphere 5.5 (indicated by the blue bars in Figure 5). The CPU used to receive 1Gbps of traffic, on the other hand, is reduced by 50% (indicated by the red line in Figure 5).

CPU Cost Improvements for Mellanox 40GbE NIC

vSphere 6.0 adds a native driver and Dynamic NetQueue for Mellanox 40GbE NICs, and these features significantly improve network performance.

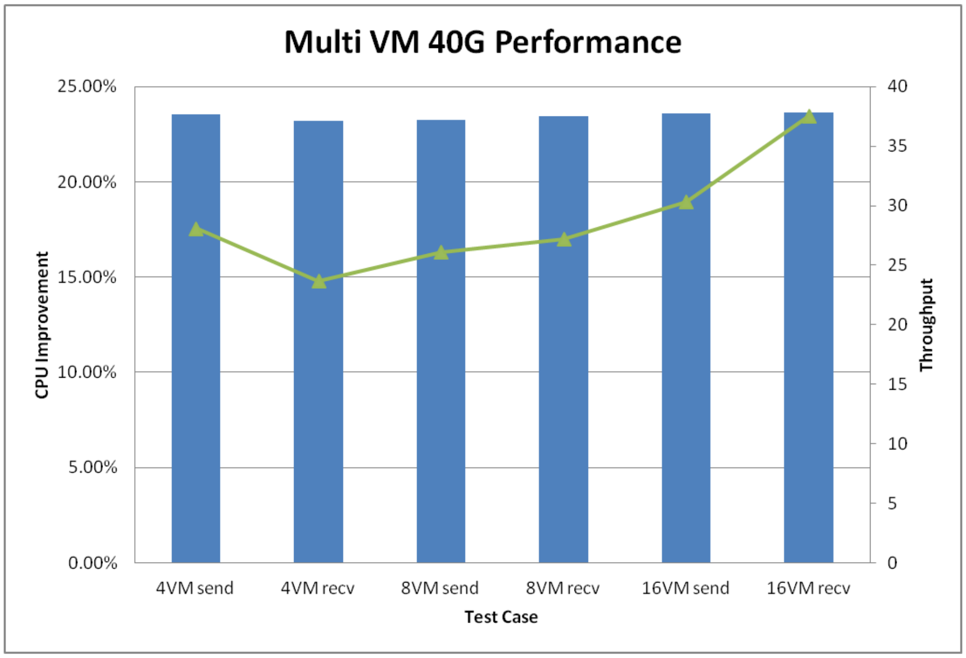

Figure 6. Performance of multiple VMs with 40GbE Melllanox card

As seen Figure 6 above, the new driver can improve CPU efficiency by up to 22%.

NUMA Aware I/O

Another performance feature introduced in 6.0 for high bandwidth NICs is NUMA Aware I/O, which improves performance by collocating highly network-intensive workloads with the device NUMA node.

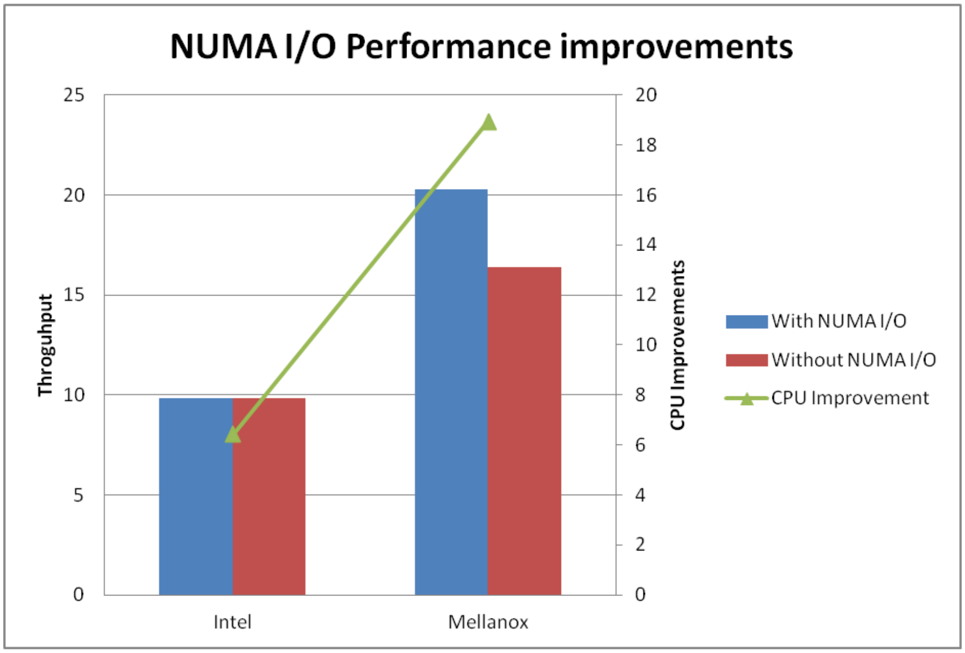

Figure 7. NUMA I/O performance improvements

NUMA I/O can result in about 20% reduced CPU consumption and about 20% higher throughput with a 1-vCPU VM for 40GbE NICs (Figure 7).

Conclusion

vSphere 6.0 contains many performance improvements across the areas of management, compute, core platform, and network. Some of these improvements include:

- vSphere Web Client performance and usability.
- vCenter Server management of very large deployments and operational throughput improved for small, medium, and large inventories.
- Scaling up (more virtual CPU) and out (more VMs per host) higher than ever before.
- Infrastructure improvements like the way nodes are arranged in vNUMA when memory is hot-added
- Graphics improvements leveraged by Horizon.
- Storage stack optimizations, which allow greater IOPS per second. Virtual Volumes are new, adding flexibility and elasticity to storage, and Virtual SAN performance is improved.
- Networking enhancements like Network I/O Control with reservations, dense mode for clusters like Web farms, VMXNET3 support for LRO, and line rate for a vNIC with a 40GbE physical NIC.

References

[1] Dennis Lu. (2015, February) vSphere Web Client.
http://blogs.vmware.com/vsphere/2015/02/vsphere-6-web-client.html

[2] Todd Muirhead. (2015, March) VMware vSphere 6 and Oracle 12c Scalability Study: Scaling Monster Virtual Machines. http://blogs.vmware.com/performance/2015/03/vmware-vsphere-6-oracle-12c-scalability-study-scaling-monster-virtual-machines.html

[3] David Morse. (2015, March) SQL Server VM Performance on VMware vSphere 6.
http://blogs.vmware.com/performance/2015/03/sql-server-vsphere6-perf.html

[4] VMware, Inc. (2015, March) What's New in the VMware vSphere 6.0 Platform.
http://www.vmware.com/files/pdf/vsphere/VMware-vSphere-Platform-Whats-New.pdf

[5] VMware, Inc. (2015) vSphere Virtual Volumes.
http://www.vmware.com/products/vsphere/features/virtual-volumes

[6] Chuck Hollis. (2015, March) How to Double Your VSAN Performance.
http://blogs.vmware.com/storage/2015/03/17/double-vsan-performance/

[7] Amitabha Banerjee. (2015, March) Performance Evaluation of Network I/O Control in VMware vSphere 6.
http://www.vmware.com/files/pdf/techpaper/Network-IOC-vSphere6-Performance-Evaluation.pdf

[8] VMware, Inc. (June, 2015) Performance Best Practices for vSphere 6.0.
http://www.vmware.com/files/pdf/techpaper/VMware-PerfBest-Practices-vSphere6-0.pdf

[9] VMware, Inc. (2015) Enable or Disable LRO on a VMXNET3 Adapter on a Windows Virtual Machine.
http://pubs.vmware.com/vsphere-60/index.jsp?topic=%2Fcom.vmware.vsphere.networking.doc%2FGUID-ECC80415-442C-44E9-BA7A-852DDB174B9F.html

[10] Sreekanth Setty. (2015, Febrary) VMware Pushes the Envelope with vSphere 6.0 vMotion.
http://blogs.vmware.com/performance/2015/02/vmware-pushes-envelope-vsphere-6-0-vmotion.html

[11] Rishi Mehta. (2015, April) Network Improvements in vSphere 6 Boost Performance for 40G NICs.
http://blogs.vmware.com/performance/2015/04/network-improvements-vsphere-6-boost-performance-40g-nics.html

[12] Microsoft. (2012, August) Receive Segment Coalescing (RSC).
https://technet.microsoft.com/en-us/library/jj574168.aspx#bkmk_rsc

About the Authors

Mark Achtemichuk is a Senior Technical Marketing Architect specializing in Performance within the Integrated Systems Business Unit (ISBU) group at VMware. Certified as VCDX #50, @vmMarkA has a strong background in datacenter infrastructures, cloud architectures, experience implementing enterprise application environments and a passion for solving problems. He has driven virtualization adoption and project success by methodically bridging business with technology. His current challenge is ensuring that performance is no longer a barrier, perceived or real, to virtualizing an organization's most critical applications, on their journey to the software-defined datacenter.

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.

Acknowledgements

The authors would like to thank Amitabha Banerjee, Chien-Chia Chen, Dennis Lu, Hari Sivaraman, Jin Heo, Jinpyo Kim, Lan Vu, Lei Chai, Ravi Soundararajan, Rishi Mehta, Sreekanth Setty, Todd Muirhead, and Yifan Wang.

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