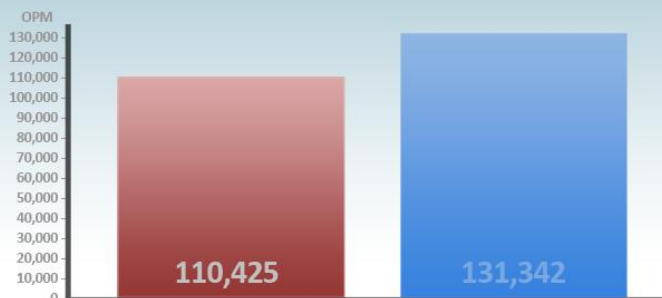


VMware® vSphere™ 5 outperformed Microsoft® Hyper-V™ R2 SP1 by 18.9%



Better scalability and more consistent resource allocation

When you invest in servers to host your virtualized applications, you can maximize the performance these systems yield by fully utilizing available resources. A hypervisor that excels at resource management allows for greater virtual machine (VM) density and better application performance.

We tested two hypervisors—VMware vSphere 5, and Microsoft Hyper-V, part of Windows Server® 2008 R2 SP1—to compare their performance and ability to manage resources at high levels of utilization. When running 30 virtual machines each, VMware vSphere 5 outperformed Microsoft Hyper-V R2 SP1 by 18.9 percent. VMware vSphere 5 also allocated server resources among individual VMs more fairly than Microsoft Hyper-V R2 SP1 (39.2 percent less variation), resulting in more consistent performance across all virtual machines. Finally, VMware vSphere 5 did a better job of scaling from 24 to 30 VMs with our workload: Overall server performance increased by 11.2 percent with VMware vSphere 5, whereas performance *decreased* by 3.3 percent with Microsoft Hyper-V R2 SP1.

These results show that VMware vSphere 5 can deliver superior performance in a densely virtualized environment over Microsoft Hyper-V R2 SP1. When you fully utilize your servers, you need fewer systems to perform the same amount of work, resulting in overall savings for your organization. In addition, more stable performance across VMs allows for consistent performance during high-availability events such as outages or maintenance windows, when VMs are consolidated down to fewer hosts.

ADVANCED RESOURCE MANAGEMENT MAXIMIZES PERFORMANCE

A hypervisor with exceptional resource management lets you optimize virtual machine performance across the entire server. The result? Greater density, scalability, and performance. VMware offers several unique features that enable vSphere 5 to utilize system resources better than competing hypervisors.

Direct driver model. The VMware approach is to install device drivers directly onto the hypervisor, effectively making the hypervisor an intermediary between the physical hardware and VMs that run on the server. The direct driver model improves performance and scalability as the number of VMs on a host increases.

High-performance “gang” scheduler. This feature allows VMware vSphere 5 to handle the challenging CPU and I/O needs of VMs. vSphere 5 is thus able to allocate resources and processor timeslices to the VMs that most need it.

How VMware manages memory

Additional VMware technologies allow vSphere 5 to optimize physical memory allocation, dynamically shifting this critical resource from less active VMs to VMs that are more active. This process is sometimes referred to as *memory oversubscription*, and is accommodated by the following features, working in concert:

Transparent page sharing. Transparent page sharing (TPS) identifies common pages across VMs and stores each in physical memory only once. This is somewhat analogous to deduplication technologies used in storage implementations. All VMs then share only that single copy of the memory page. VMware vSphere 5 determines sharable pages by scanning the content of the virtual machines’ physical memory for sharing opportunities. By eliminating these redundant copies, VMware vSphere 5 frees up memory for use by workloads.

Memory ballooning. When the hypervisor needs to give more memory to VMs that are powering on or experiencing a heavy workload, it “asks” the guest operating systems in other VMs to provide memory to a balloon process that runs in the guest. The hypervisor can then temporarily lend that “ballooned” memory to the busy VMs. When the busy VMs no longer need the extra memory, the hypervisor “deflates” the balloon, returning that memory to the original guest OS.

Memory compression. The innovative memory compression capability in VMware vSphere 5 sets aside a small portion of physical RAM as a cache. Compressing unused memory pages avoids hypervisor swapping and is orders of magnitude faster than disk.

Hypervisor swap. If a system’s memory resources are experiencing intense pressure, hypervisor swap acts as a safety valve, ensuring reliable operation of the host and all workloads. While this may result in a short-term performance hit, it offers the

hypervisor another option to resolve memory issues. Furthermore, a new feature in vSphere 5, called swap to host cache, can use solid-state disks for swap purposes, reducing the impact on performance.

DRS with resource pools. This feature is a safety net of sorts, largely because it ensures that applications receive the resources they need when they need them. It accomplishes this by dynamically load balancing resources throughout a cluster of VMs. This does not apply to standalone hosts such as the one tested for this report, but to vSphere clusters. Using a vSphere-clustered environment with DRS ensures optimization of resources and the ability to accommodate shifting workloads.

How Hyper-V manages memory

In contrast with multiple advanced VMware memory technologies, all designed to maximize memory performance, Microsoft Hyper-V R2 SP1 uses a single feature, Dynamic Memory, which uses hot-add in conjunction with the memory ballooning process. Based on our results, this limited resource management capability offers a less flexible approach to optimizing resources. Hyper-V was thus unable to match the performance of VMware vSphere 5 in our test scenarios

PUTTING THE HYPERVISORS TO THE TEST

To compare the hypervisors we tested, we carried out two scenarios while running VMware vSphere 5 on our server and then repeated the scenarios while running Microsoft Hyper-V R2 SP1. In the first scenario, we used 30 VMs, each allocated 2 virtual CPUs and 4 GB of RAM. Our server had 24 logical CPUs (2 sockets, with a total of 12 physical cores plus hyperthreading) and 96 GB of RAM. In the second scenario, we used a total of 24 VMs.

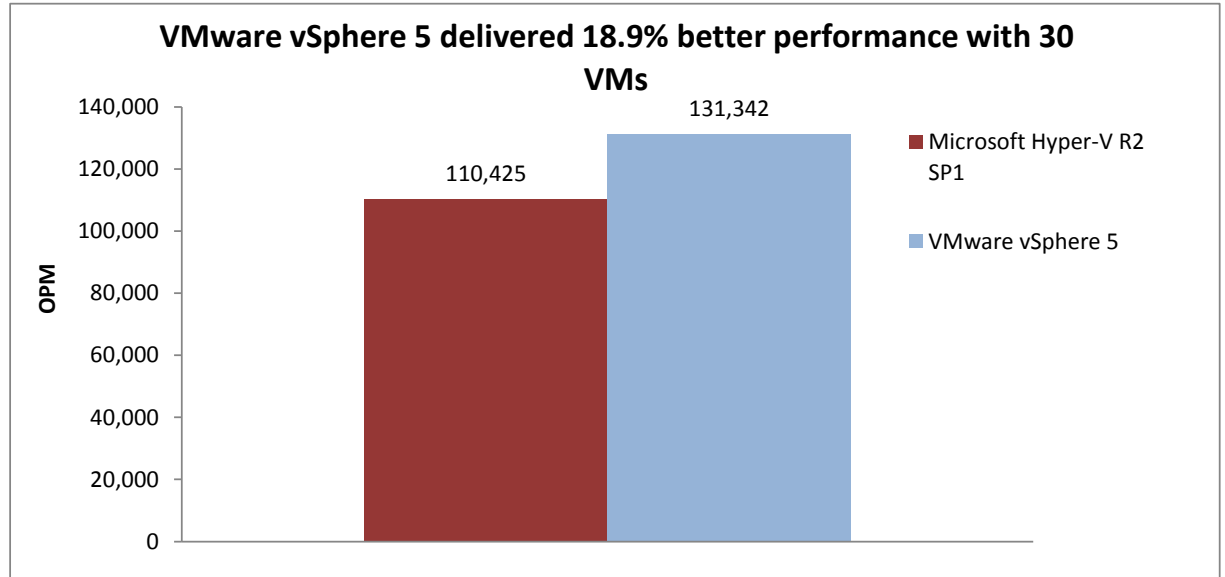
In both scenarios, each VM ran Windows Server 2008 R2 SP1, Microsoft SQL Server® 2008 R2 SP1 and an identical database workload created with the DVD Store Version 2 (DS2) benchmark tool. DS2 simulates an online store, and reports results using the metric of orders per minute (OPM). We chose this database workload for our testing as it heavily utilizes processor, memory, and I/O to create a broad and demanding load on the system. For more details about the DS2 tool, see <http://www.delltechcenter.com/page/DVD+Store>.

For more details on our test bed configuration and for detailed test results, see the full version of this report at <http://www.principledtechnologies.com/clients/reports/VMware/vSphere5density0811.pdf>.

VMWARE VSPHERE 5 DELIVERS

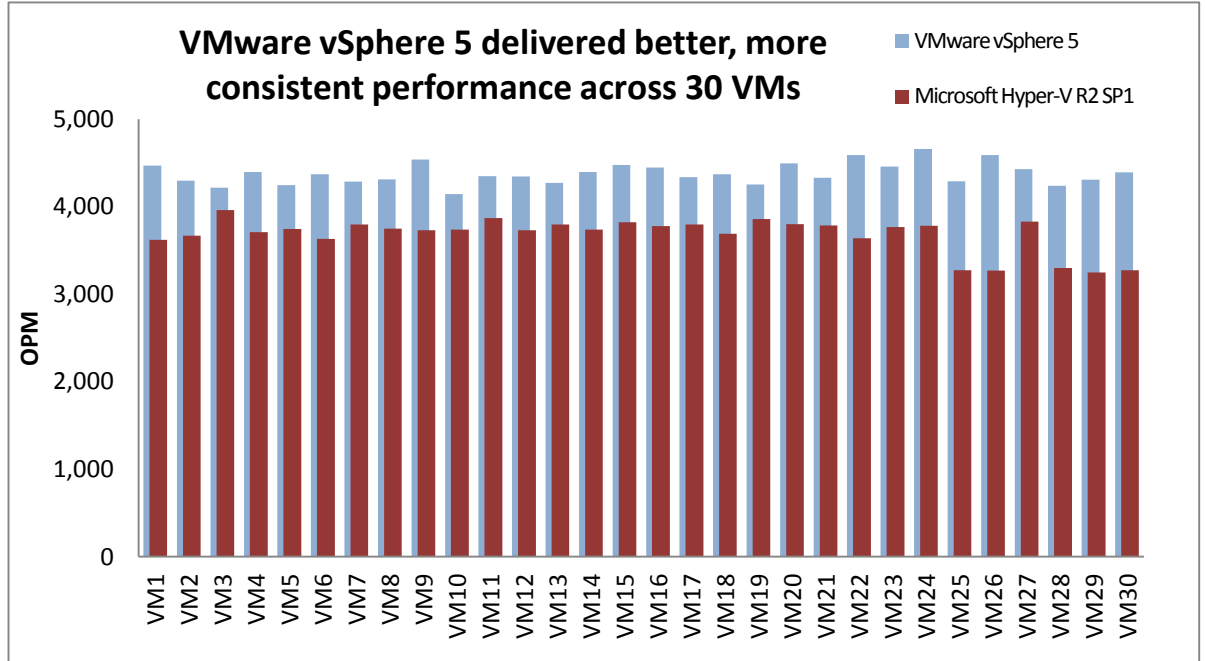
As Figure 1 shows, aggregate performance across the VMs on the server with VMware vSphere 5 was 131,342 OPM; 18.9 percent higher than on the server running Microsoft Hyper-V R2 SP1, where the total performance of all 30 VMs was 110,425 OPM.

Figure 1. Aggregate performance, in orders per minute, for the 30-VM scenario. Higher numbers are better.



As Figure 2 shows, the individual VMs all achieved higher performance levels with VMware vSphere 5 than with Microsoft Hyper-V R2 SP1 at 30 VMs. Figure 2 shows not only that performance levels on the server running VMware vSphere 5 were greater, but also that the levels were much more consistent, showing that VMware vSphere 5 more fairly allocates resources across VMs. When you have multiple VMs running on a server, the goal is to have the hypervisor intelligently manage resource allocation so that your applications deliver predictable levels of service.

Figure 2. Performance per VM, in orders per minute, for the 30-VM scenario. Higher and more consistent numbers are better.



As Figure 3 shows, the standard deviation of the VMware vSphere 5 VMs' performance was 120.7, 39.2 percent less than the standard deviation of 198.7 on Microsoft Hyper-V R2 SP1. A smaller standard deviation means the VMs are running within a narrower, and thus fairer, performance range.

Figure 3. Standard deviation of the OPM counts of the 30 VMs in the oversubscribed scenario. Lower numbers are better.

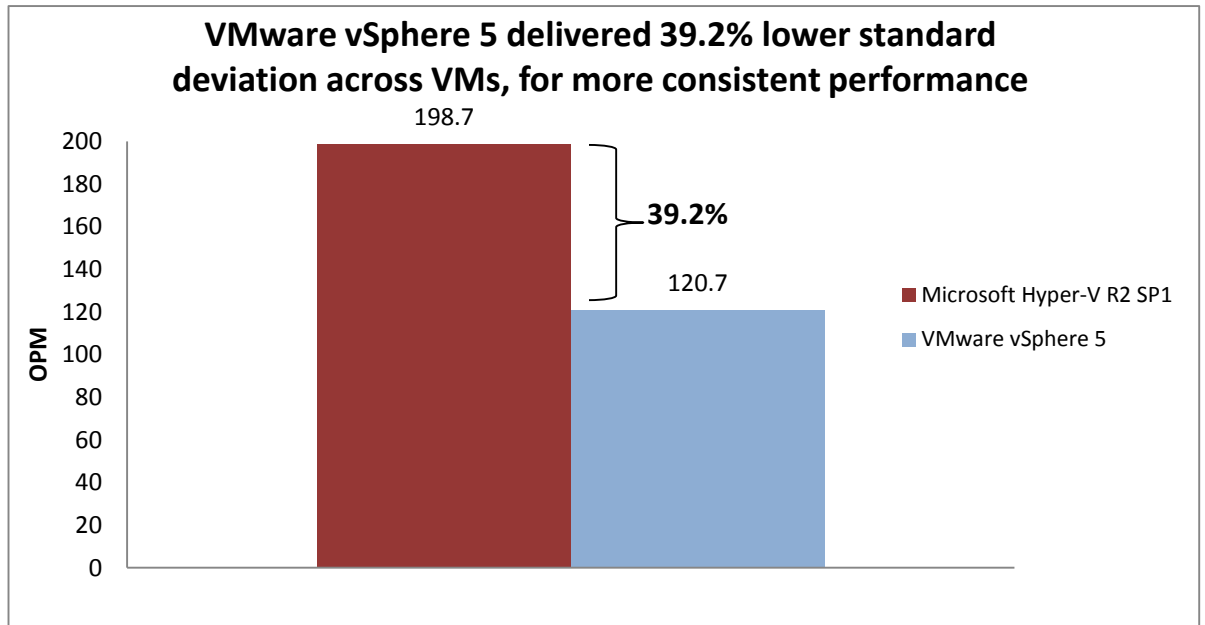
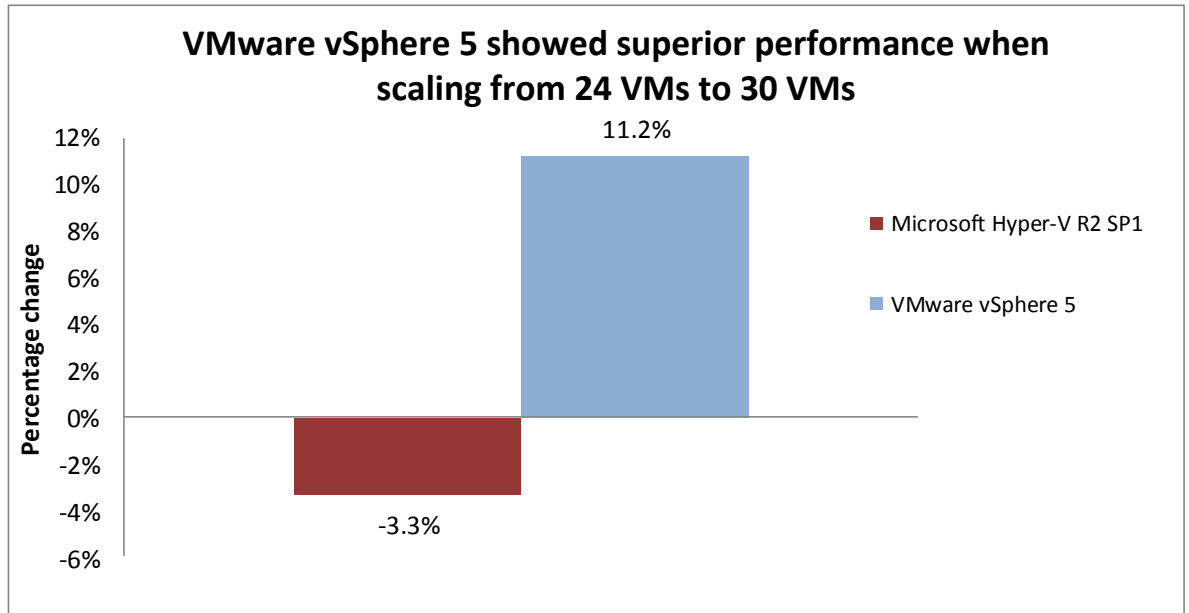


Figure 4 charts the change in the total performance from our second scenario, when we scaled from 24 to 30 VMs (125 percent oversubscribed memory). In this scenario, the total number of OPM decreased by 3.3 percent on the server running Microsoft Hyper-V R2 SP1, meaning the overall application performance of the entire system declined. In contrast, the server running VMware vSphere 5 showed superior VM scaling by actually increasing overall server performance by 11.2 percent in the 30-VM scenario.

Figure 4. Total performance change when scaling from 24 VMs to 30 VMs with the two hypervisors. Higher numbers are better.



CONCLUSION

Using a hypervisor that offers better resource management and scalability can deliver maximum and predictable virtual machine performance on your servers.

In our testing, VMware vSphere 5 allowed virtual machines to outperform those running on Microsoft Hyper-V R2 SP1 by 18.9 percent in total system OPM performance. It also allowed these VMs to operate at a consistent and more fair level, while the VMs running on Microsoft Hyper-V R2 SP1 showed 39.2 percent more deviation than those running on VMware vSphere 5. Furthermore, VMware vSphere 5 scaled better when going from 24 VMs to 30 VMs: Total performance for VMware vSphere 5 increased by 11.2 percent, whereas it decreased by 3.3 percent with Microsoft Hyper-V R2 SP1.

With the consistent resource management and scalability that VMware vSphere 5 offers, you are able to utilize the full capability of your servers with confidence; this translates to fewer servers in the data center, lower costs for your business, and more consistent overall application performance.



Principled Technologies, Inc.
1007 Slater Road, Suite 300
Durham, NC, 27703
www.principledtechnologies.com

Principled Technologies is a registered trademark of Principled Technologies, Inc.
All other product names are the trademarks of their respective owners.

Disclaimer of Warranties; Limitation of Liability:

PRINCIPLED TECHNOLOGIES, INC. HAS MADE REASONABLE EFFORTS TO ENSURE THE ACCURACY AND VALIDITY OF ITS TESTING, HOWEVER, PRINCIPLED TECHNOLOGIES, INC. SPECIFICALLY DISCLAIMS ANY WARRANTY, EXPRESSED OR IMPLIED, RELATING TO THE TEST RESULTS AND ANALYSIS, THEIR ACCURACY, COMPLETENESS OR QUALITY, INCLUDING ANY IMPLIED WARRANTY OF FITNESS FOR ANY PARTICULAR PURPOSE. ALL PERSONS OR ENTITIES RELYING ON THE RESULTS OF ANY TESTING DO SO AT THEIR OWN RISK, AND AGREE THAT PRINCIPLED TECHNOLOGIES, INC., ITS EMPLOYEES AND ITS SUBCONTRACTORS SHALL HAVE NO LIABILITY WHATSOEVER FROM ANY CLAIM OF LOSS OR DAMAGE ON ACCOUNT OF ANY ALLEGED ERROR OR DEFECT IN ANY TESTING PROCEDURE OR RESULT.

IN NO EVENT SHALL PRINCIPLED TECHNOLOGIES, INC. BE LIABLE FOR INDIRECT, SPECIAL, INCIDENTAL, OR CONSEQUENTIAL DAMAGES IN CONNECTION WITH ITS TESTING, EVEN IF ADVISED OF THE POSSIBILITY OF SUCH DAMAGES. IN NO EVENT SHALL PRINCIPLED TECHNOLOGIES, INC.'S LIABILITY, INCLUDING FOR DIRECT DAMAGES, EXCEED THE AMOUNTS PAID IN CONNECTION WITH PRINCIPLED TECHNOLOGIES, INC.'S TESTING. CUSTOMER'S SOLE AND EXCLUSIVE REMEDIES ARE AS SET FORTH HEREIN.
