Power Management and Performance in VMware vSphere® 5.1 and 5.5

Performance Study

TECHNICAL WHITE PAPER

vmware®

Table of Contents

Executive Summary

This paper describes the testing and results of two performance studies: The first shows how power management in VMware vSphere® 5.5 in balanced mode (the default) performs 18% better than the physical host's balanced mode power management setting. The second study compares vSphere 5.1 performance and power savings in two server models that have different generations of processors. Results show the newer servers have 120% greater performance and 24% improved energy efficiency over the previous generation.

Introduction

One of the exciting aspects of virtualization is its ability to cut costs across the datacenter. Cost savings occur when IT professionals consolidate software instances on virtual machines instead of having multiple physical machines each dedicated to a single software instance. IT staff can add multiple virtual machines to a single host, thus allowing the software to most efficiently utilize the compute, memory, and storage available to the host. In addition to consolidation, an important part of the datacenter cost strategy is power consumption. Physical servers frequently offer a power management scheme that puts processors into low power states when not fully utilized. Additionally, virtualization software can have its own power management techniques. How effective are these power-saving options, and how do they compare to each other? Does power efficiency differ between physical servers across different CPU generations? This paper describes two performance studies that were undertaken to discover the answers to these questions.

Benchmark Software

These studies used VMware VMmark 2.5, a multi-host virtualization benchmark that uses a variety of application workloads as well as common datacenter operations to model the demands of the datacenter. Virtual machines running diverse application workloads are grouped into units of load called tiles. VMmark 2.5 scores are based on application and infrastructure workload throughput, while application latency reflects Quality of Service (QoS). For the Mail Server, Olio, and DVD Store 2 workloads, latency is defined as the application's response time. For more details, see the VMmark product page [1].

Power Management in the BIOS versus VMware vSphere 5.1 and 5.5 Host Power Management

Host power management (HPM) in VMware vSphere 5.5 has four settings: High Performance, Balanced, Low Power, and Custom. VMware vSphere power management settings can take effect when the BIOS is in OS control mode. Alternatively, power management can be controlled by the BIOS, in which server manufacturers provide several types of power management profiles. Although they vary by vendor, most include a Performance option, which does not use any power saving techniques, and a Balanced option, which claims to increase energy efficiency with minimal or no impact to performance. This study compares the default Balanced setting in vSphere with the Balanced setting in the physical server's BIOS. Tests were performed to discover if the Balanced setting, either controlled by the BIOS or by VMware ESXi™ (the operating system and hypervisor component of vSphere), reduces performance, and to what degree, relative to the BIOS-controlled Maximum Performance setting.

Three power management profiles were tested: the BIOS-controlled Maximum Performance setting, which uses no power management techniques, the ESXi-controlled Balanced setting, and the BIOS-controlled Balanced setting. The ESXi Balanced setting cuts power by reducing processor frequency and voltage [2] [3]; the BIOS-controlled Balanced setting reduces processor frequency and voltage and puts the processor in deep idle states [4]. Lastly, the ESXi 5.5 Balanced setting was compared against the ESXi 5.1 Balanced setting.

Test Methodology

All tests were conducted on a four-node cluster running VMware vSphere. The cluster's performance and energy efficiency were compared between three power management profiles: the BIOS-controlled Performance setting, the ESXi-controlled Balanced setting, and the BIOS-controlled Balanced setting, also known as "Performance per Watt (Dell Active Power Controller)." The ESXi-controlled Balanced setting was tested on both ESXi 5.1 and ESXi 5.5.

Configuration

Systems Under Test	Four Dell PowerEdge R620 servers
CPUs (per server)	One 8-core Intel® Xeon® E5-2665 @ 2.4 GHz, Hyper-Threading enabled
Memory (per server)	96GB DDR3 ECC @ 1067 MHz
Storage Array	EMC VNX5700 62 Enterprise Flash Drives (SSDs), RAID 0, grouped as three 8-SSD LUNs, seven 5-SSD LUNs, and one 3-SSD LUN
Host Bus Adapter	Two QLogic QLE2562, Dual Port 8Gb Fibre Channel to PCI Express
Network Controller	One Intel Gigabit Quad Port I350 Adapter
Hypervisor	VMware ESXi 5.5 and ESXi 5.1
Virtualization Management	VMware vCenter Server 5.1.0
VMmark version	2.5
Power Meters	Three Yokogawa WT210 measured the power consumption of four servers.

Table 1. Cluster configuration for first performance test

Results

To determine the maximum VMmark load supported for each power management setting, the number of VMmark tiles was increased until the cluster reached saturation, which is defined as the largest number of tiles that still meet Quality of Service (QoS) requirements. Power consumption was also measured when the hosts were idle and no virtual machines were running on the cluster. VMmark scores are normalized to the BIOS Balanced one-tile score.

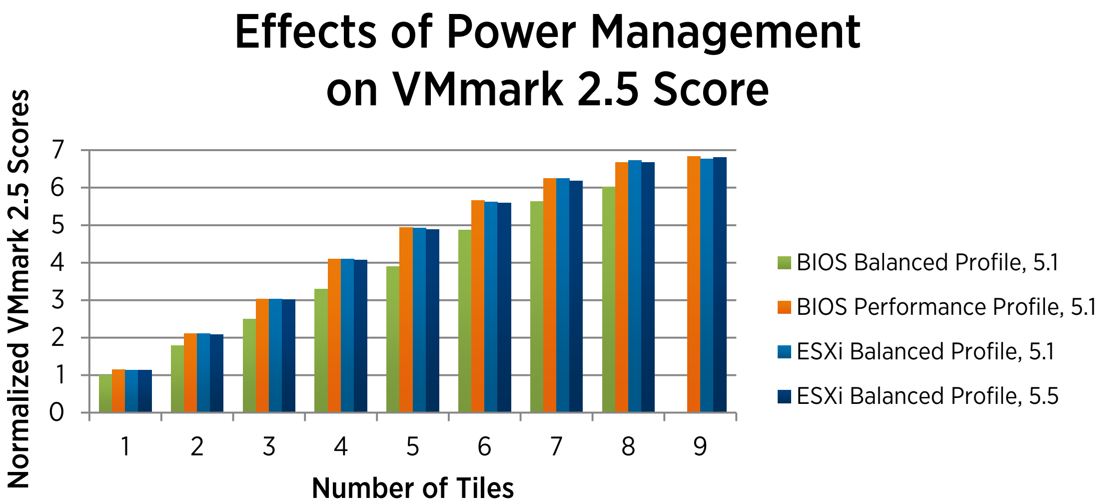

Figure 1. Effects of power management on VMmark 2.5 score (higher is better)

The VMmark score was essentially equivalent between the BIOS Performance setting, the ESXi 5.5 Balanced setting, and the ESXi 5.1 Balanced setting, with less than a 1% difference at all load levels. However, the BIOS Balanced setting showed lower performance and reduced the VMmark score by an average of 15%. Only passing runs are pictured in Figure 1.

On the BIOS Balanced setting, the environment was unable to support nine tiles. As shown in Figure 2, some runs failed QoS requirements, indicating that the applications did not maintain an acceptable Quality of Service, even at low load levels.

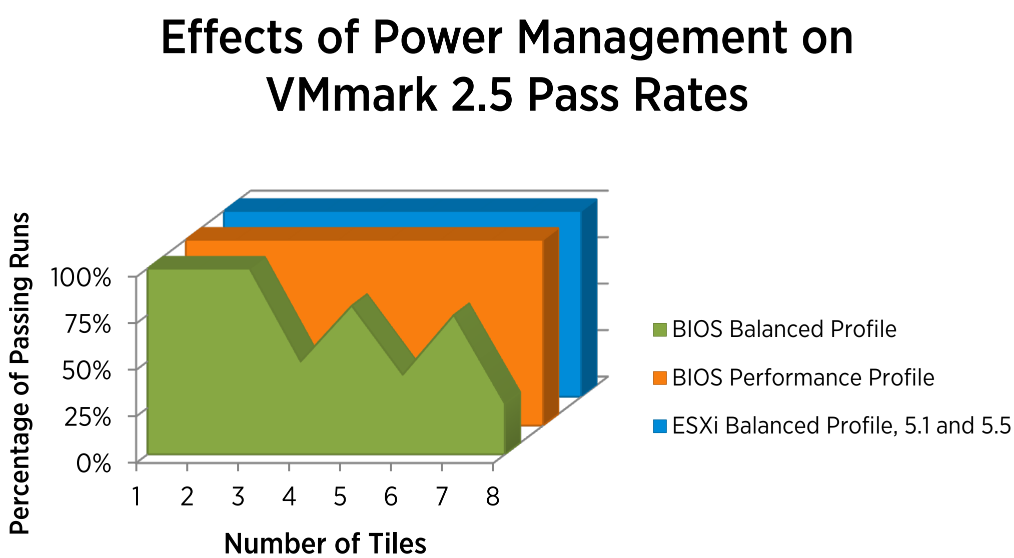

Figure 2. VMmark 2.5 pass rates

The improvements in energy efficiency between the ESXi and BIOS Balanced settings and the BIOS Performance setting were also compared. The Performance Per Kilowatt (PPKW) metric in VMmark represents energy efficiency as the VMmark score per kilowatt of power consumed. More efficient results will have a higher PPKW. Energy efficiency is independent of overall performance.

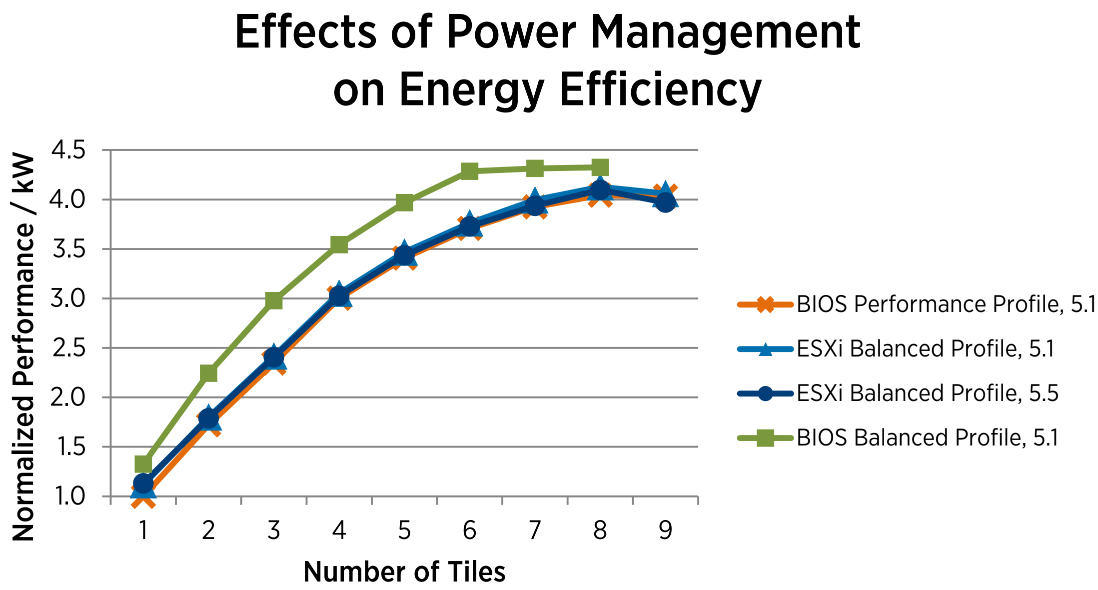

Figure 3. Effects of power management on energy efficiency (higher is better)

Two trends are visible in Figure 3. As expected, the BIOS Performance setting showed the lowest energy efficiency. At every load level, the ESXi 5.5 and ESXi 5.1 Balanced settings performed equivalently, about 3% more energy efficiently than the BIOS Performance setting, despite the fact that the ESXi Balanced setting delivered an equivalent score to the BIOS Performance setting. The BIOS Balanced setting had the greatest energy efficiency at a 20% average improvement over the BIOS Performance setting.

Second, an increase in load is correlated with greater energy efficiency. As the CPUs become busier, throughput increases at a faster rate than the required power. This can be understood by noting that an idle server will still consume power, but with no work to show for it. A highly utilized server is typically the most energy efficient per request completed, which is confirmed by these results. Higher energy efficiency creates cost savings in host energy consumption and in cooling costs.

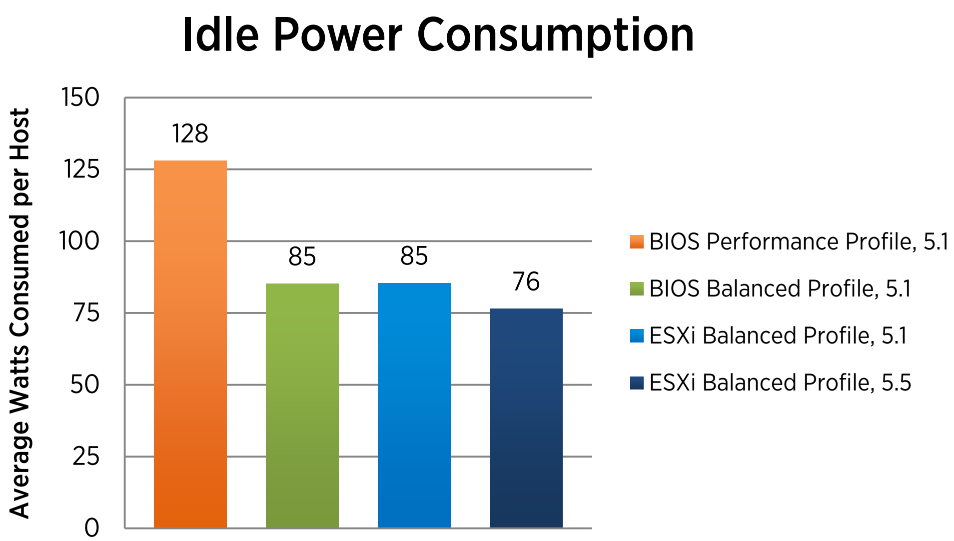

Figure 4. Idle power consumption per host (lower is better)

The bursty nature of most environments sometimes results in idle servers, so each host's idle power consumption was also measured when no virtual machines were running. As shown in Figure 4, the ESXi Balanced settings consumed the least power at idle, with the ESXi 5.5 Balanced setting decreasing power usage by 11% over the ESXi 5.1 Balanced setting. Although the BIOS Performance and ESXi 5.5 Balanced settings performed very similarly under load, hosts using the ESXi 5.5 Balanced setting consumed 40% less power while idle.

The last test in this section looks at how power management settings affect application latency.

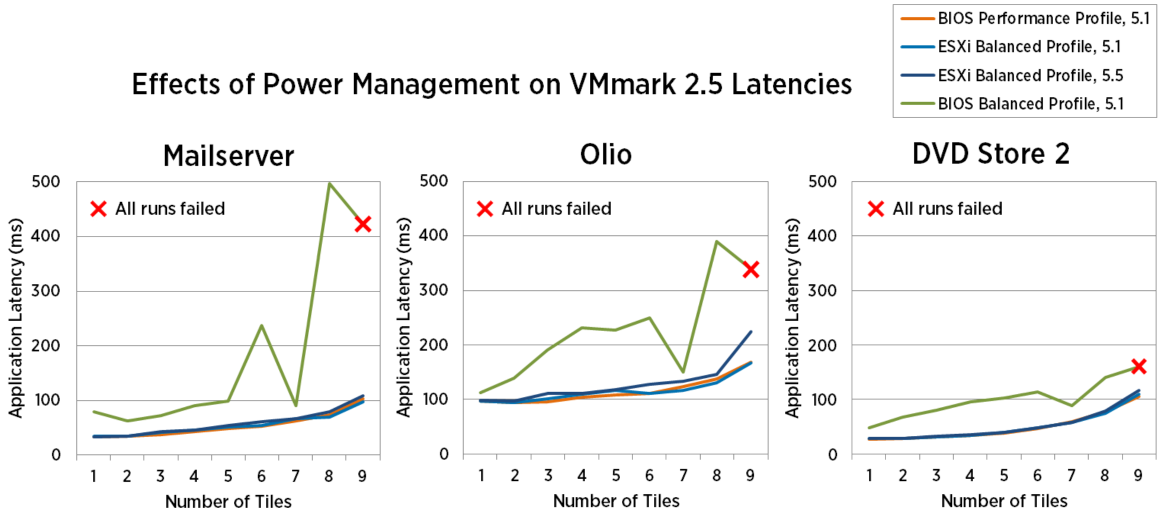

Figure 5. Effects of power management on VMmark 2.5 latencies (lower is better)

While the BIOS Performance and ESXi Balanced latencies tracked closely, BIOS Balanced latencies were significantly higher at all load levels. Furthermore, these latencies were unpredictable even at low load levels, and this unpredictability causes run failures as shown in Figure 2 and Figure 5. Figure 5 includes all runs, whereas Figure 1 and Figure 3 show only passing runs.

Our tests showed that the ESXi 5.5 and ESXi 5.1 Balanced power management settings didn't affect throughput or latency compared to the BIOS-controlled Performance setting, but improved energy efficiency under load by 3%, and the ESXi 5.5 Balanced setting cut power consumption while idle by 40%. The BIOS-controlled Balanced profile did improve power efficiency by an average of 20% over the BIOS-controlled Performance setting, but was so aggressive in cutting power that it often caused VMmark to fail QoS requirements.

Overall, the BIOS-controlled Balanced setting produced substantial efficiency gains but with unpredictable performance, failed runs, and reduced performance at all load levels. This profile may still be suitable for some workloads that can tolerate this unpredictability, but should be used with caution. On the other hand, the ESXi Balanced setting produced efficiency gains while it did an excellent job maintaining high performance across all load levels. These findings make us confident that the ESXi Balanced setting is the better choice for most kinds of virtualized applications.

Performance and Energy Efficiency Across CPU Generations

Each new generation of servers brings advances in hardware components. For IT professionals purchasing or managing new generations of hardware, it is important to understand how these incremental hardware improvements translate into real-world gains in the datacenter. Tests in this section compare the performance and energy efficiency of two different generations of servers in four-node clusters.

Testing Methodology

Tests were conducted on two four-node clusters running VMware vSphere 5.1. Performance and energy efficiency were compared between a cluster of previous generation Dell PowerEdge R310 servers, and a cluster of current generation Dell PowerEdge R620 servers. For simplicity, these clusters are referred to as the *old cluster* and *new cluster*, respectively. Among other hardware differences, the old cluster servers contained four-core Intel Nehalem processors, while the new cluster servers contained eight-core Intel Sandy Bridge EP processors. Memory in the newer servers was appropriately scaled up to accommodate their increased processing power and represents common current server configurations. Software and storage configurations were identical between clusters.

Configuration

OLD CLUSTER	
Systems Under Test	Four Dell PowerEdge R310 servers
CPUs (per server)	One Quad-Core Intel® Xeon® X3460 @ 2.8 GHz, Hyper-Threading enabled
Memory (per server)	32GB DDR3 ECC @ 800 MHz
NEW CLUSTER	
Systems Under Test	Four Dell PowerEdge R620 servers
CPUs (per server)	One 8-Core Intel® Xeon® E5-2665 @ 2.4 GHz, Hyper-Threading enabled
Memory (per server)	96GB DDR3 ECC @ 1067 MHz

SAME FOR EACH CLUSTER	
Storage Array	EMC VNX5700 62 Enterprise Flash Drives (SSDs), RAID 0, grouped as three 8-SSD LUNs, seven 5-SSD LUNs, and one 3-SSD LUN
Hypervisor	VMware vSphere 5.1
Virtualization Management	VMware vCenter Server 5.1.0
VMmark version	2.5
Power Meters	Three Yokogawa WT210

Table 2. Cluster configuration for second performance test

Results

To determine the maximum VMmark load the old cluster could support, the number of VMmark tiles was increased until the cluster reached saturation, which is defined as the largest number of tiles that still meets QoS requirements. The new cluster was then tested at the same number of tiles. All data points represent the mean of four tests in each configuration and VMmark scores are normalized to the old cluster's performance.

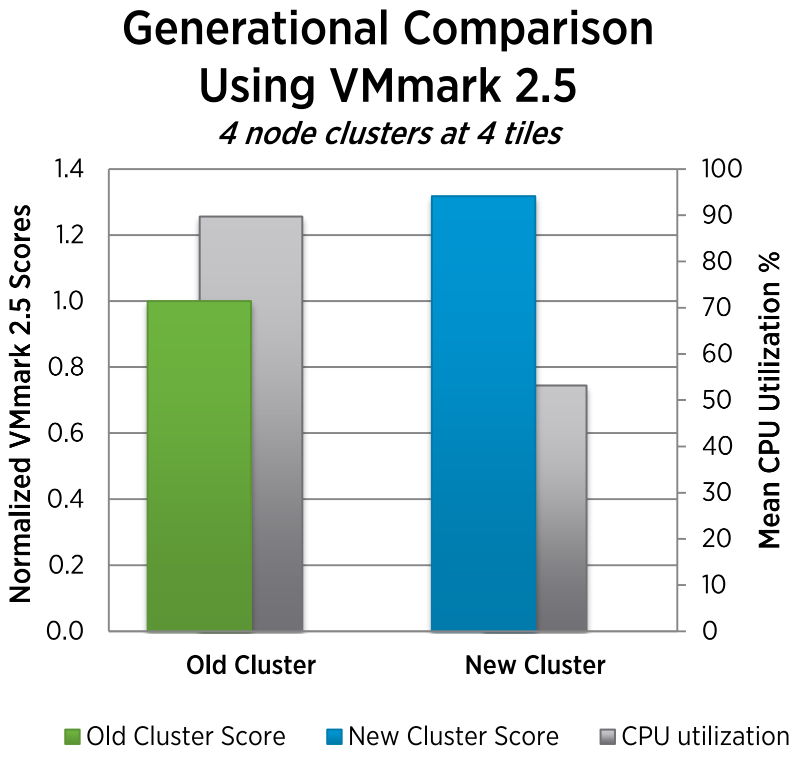

Figure 6. Performance comparison of two clusters with different generations of CPU using VMmark 2.5 at four tiles

As seen in Figure 6, the new cluster had a 32% higher VMmark score in combination with a 41% lower CPU utilization. The new cluster also showed a 24% increase in energy efficiency over the old cluster, which is discussed later. At four tiles, the old cluster was bottlenecked on CPU, resulting in decreased workload throughput, while the new cluster was not. With CPU resources to spare, the new cluster met the requested load at lower latencies, which increased its total throughput and score. Mean I/O latencies remained low for both clusters, at 1.2ms reads and 1.1ms writes for the old cluster, and 1.0ms reads and 0.9ms writes for the new cluster.

The next test (Figure 7) determined the maximum VMmark load the new cluster could support. While the old cluster was saturated at four tiles, the new cluster accommodated more than twice the load at nine tiles and produced a score 120% higher than the old cluster. Mean I/O latencies remained low at 1.0ms.

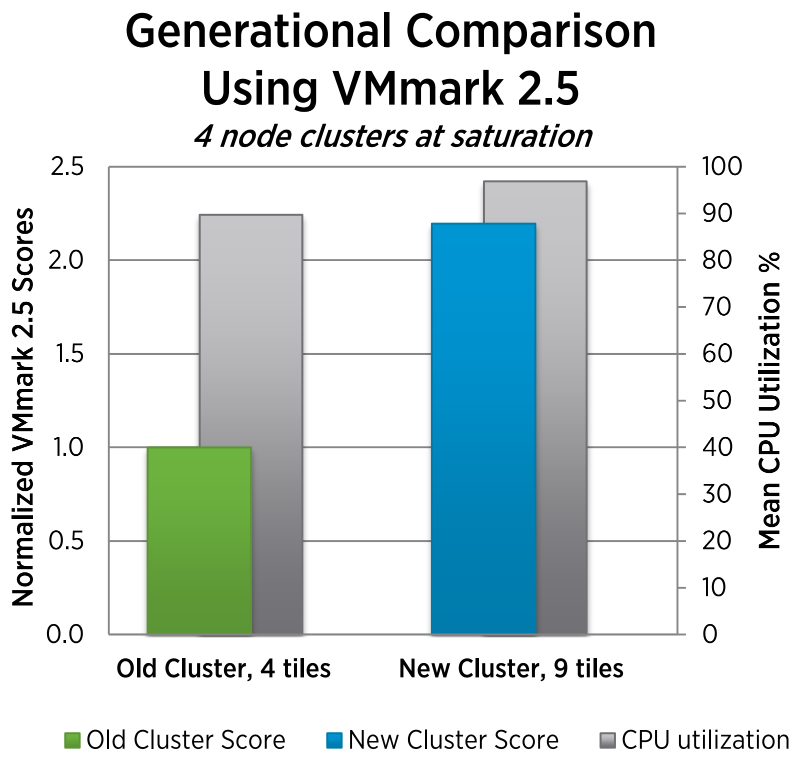

Figure 7. Performance comparison of two clusters, each with a different CPU generation, tested with VMmark 2.5 at saturation

The performance advantages of the PowerEdge R620 over the R310 were largely due to the generational improvements of the R620's eight-core E5-2665 processor versus the R310's quad-core X3460 processor, which includes improved bus speeds and larger L3 cache, and the R620's higher memory capacity.

These performance results suggest that it would be possible to replace four Dell R310 servers with two Dell R620 servers and still achieve higher performance (Figure 8).

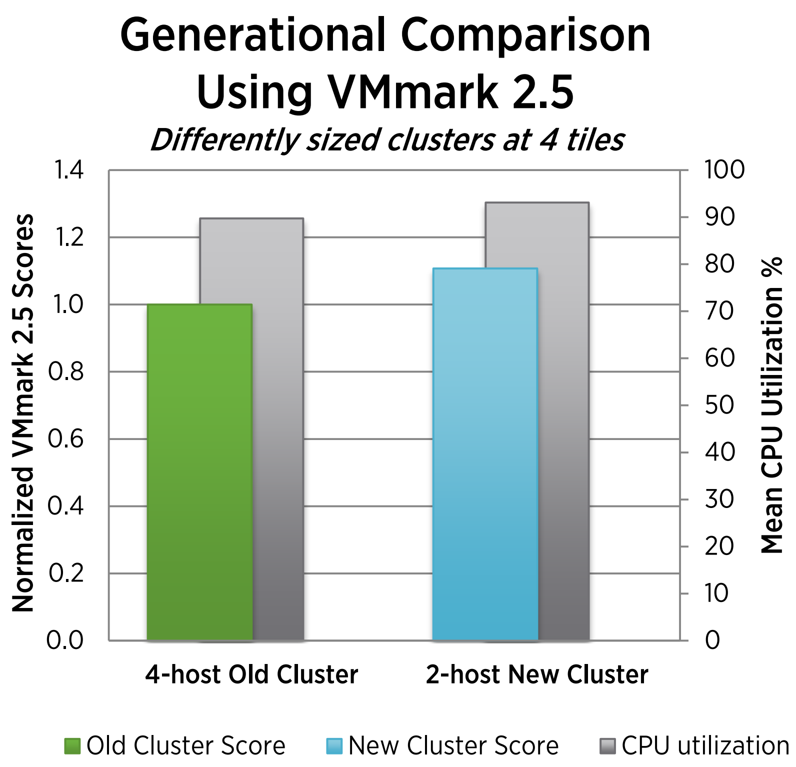

Figure 8. Performance comparison of two differently sized clusters with different generations of CPU using VMmark 2.5 at four tiles

Removing two nodes from the new cluster showed that the two remaining nodes did support four tiles at 93% utilization, with an 11% higher VMmark score and 74% greater energy efficiency than the four-host old cluster.

Lastly, the improvements in energy efficiency were compared. The VMmark Performance Per Kilowatt (PPKW) metric models energy efficiency as a VMmark score per kilowatt of power consumed. Figure 9 plots energy efficiency against the normalized VMmark score. Both clusters were run with their servers' power management set to "maximum performance."

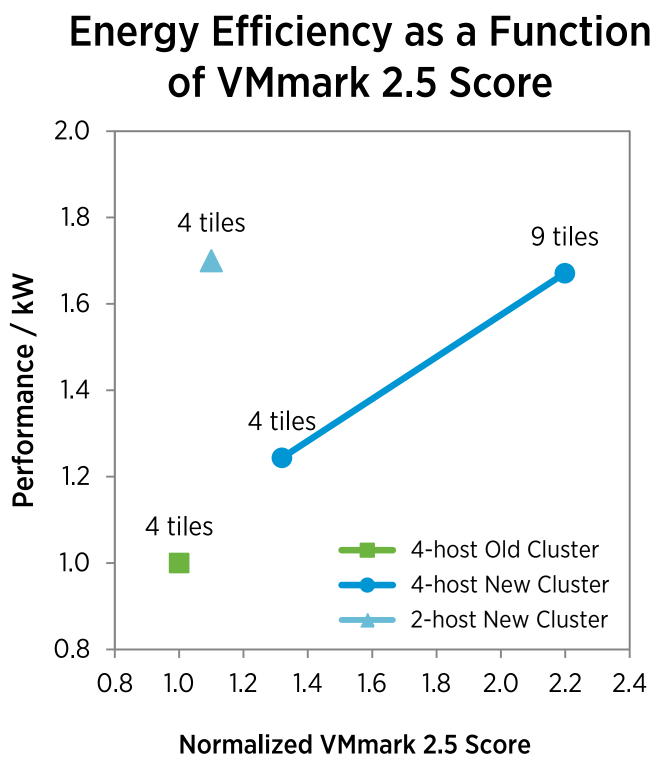

Figure 9. Energy efficiency as a function of VMmark 2.5 score

Two trends emerge from Figure 9. First, at four tiles, the four-host new cluster accomplishes more work at higher energy efficiency than the old cluster. Across the board, the new cluster is more energy efficient than the old cluster. Second, within the four-host new cluster, greater energy efficiency is correlated with an increase in VMmark score. As the CPUs become busier, performance increases at a faster rate than the required power. This can be understood by noting that an idle server will still consume power, but with no work to show for it. A highly utilized server is typically the most energy efficient per request completed, which is confirmed by the two-host new cluster that achieved high efficiency at 93% utilization.

Measurements of idle power consumption showed that the new cluster consumes more energy while idle, at 128 watts per host, than the old cluster, at 88 watts per host. This is likely due to the increased core count and memory of the new cluster hosts. However, the new cluster's increased energy efficiency while utilized more than compensates for its increased power consumption while idle.

Conclusion

In the first study, the ESXi Balanced setting did an excellent job of preserving performance, with no measurable performance impact at all levels of load. Not only was performance on par with expectations, but it did so while producing consistent improvements in energy efficiency, even while idle. In comparison, the BIOS Balanced setting aggressively saved power but created intolerably high latencies and reduced performance and reliability.

Further studies showed that, while running vSphere 5.1, two newer Dell R620 servers are capable of supporting a greater load than four older Dell R310 servers. Because the Dell R620 performance is more than double that of the Dell R310, a four-node Dell R620 cluster reached a 120% higher maximum score than the Dell R310 cluster. In addition to its performance advantages, the Dell R620 cluster performed with greater energy efficiency at each load level, showing that the Dell R620 has both superior performance and greater energy efficiency than the Dell R310.

References

[1] VMware, Inc. (2014) VMmark 2.x.
http://www.vmware.com/products/vmmark/overview.html

[2] Qasim Ali. (2010) Host Power Management in VMware vSphere 5.
https://www.vmware.com/files/pdf/hpm-perf-vsphere5.pdf

[3] Qasim Ali. (2013) Host Power Management in VMware vSphere 5.5.
http://www.vmware.com/files/pdf/techpaper/hpm-perf-vsphere55.pdf

[4] Bruce Wagner. (2014) Dell PowerEdge 12th Generation Server System Profiles.
http://en.community.dell.com/techcenter/extras/m/white_papers/20161975/download.aspx

[5] VMware, Inc. (2013) Performance Best Practices for VMware vSphere 5.5.
http://www.vmware.com/pdf/Perf_Best_Practices_vSphere5.5.pdf

[6] VMware, Inc. (2013) vSphere Resource Management.
http://pubs.vmware.com/vsphere-55/topic/com.vmware.ICbase/PDF/vsphere-esxi-vcenter-server-55-resource-management-guide.pdf

About the Author

Rebecca Grider is a member of the VMware Performance Engineering Group. She uses VMmark to measure hardware and virtualization platform performance, and contributes to VMmark development. She has a Master's degree in Computer Science from the University of Texas at Austin.

Acknowledgements

The author would like to thank colleagues in the Performance Group, other colleagues at VMware, and her manager Bruce Herndon, for their reviews and input on the paper.