

## Evaluating the ESX 4 Hypervisor and VM Density Advantage

September 2009



The summer of 2009 rang in big changes for the server virtualization market. The hypervisor – the engine of every server virtualization offering – underwent significant changes in nearly every vendor’s offering. Hyper-V reached Release 2 status, Citrix released XenServer version 5.5, and VMware announced ESX 4, alongside a number of other virtual infrastructure announcements. Each vendor in this group anticipated significant performance enhancements with their hypervisor update.

Early in 2009, Taneja Group had benchmarked the performance of some of these hypervisors. With these new product versions, we were eager to update the results. This Technology Validation report is an update to our benchmarking of these popular hypervisors.

As with our previous testing, this validation exercise focuses upon *Virtual Machine density*. Virtual Machine (VM) density is the number of virtual machines that can run simultaneously – based on a specific set of application workloads – on a single hypervisor instance without disruptive performance impact. This new study compares the VM density of each of the latest generation of hypervisor products – Microsoft Hyper-V R2 (Release Candidate), XenServer 5.5, and VMware ESX 4. Our comparison is based on a set of tests designed to evaluate hypervisor performance under different types of workloads.

Why is VM density important? Because it has a considerable impact on virtual infrastructure acquisition costs. As we have discussed in earlier reports, and as VMware illustrates through a detailed calculator ([www.vmware.com/go/costperappcalc](http://www.vmware.com/go/costperappcalc)), VM density will determine the cost to purchase and deploy your virtual server infrastructure. That is largely because VM density will determine how many physical machines you must maintain, and how many separate hypervisors you must support with management software and other software licenses. We have before set the stage for what this total cost picture is, when the total suite of necessary management tools is deployed. As we reviewed in that report, a VM density advantage of 1.5:1 could yield a cost of acquisition savings of up to 29%. Depending on the total size of the environment, that savings may range from several thousand dollars to *several hundreds of thousands of dollars*.

**Our Findings:** Our hands-on validation has revealed that VMware ESX 4 achieves a significantly higher VM density than both Microsoft Hyper-V R2 and Citrix XenServer 5.5, allowing users to consolidate more servers on a VMware platform. The magnitude of the difference continues to be quite large: We have concluded that in real world use, VMware ESX 4 demonstrates at least a 80% (1.8 to 1) density advantage over both Hyper-V and XenServer across a variety of application workloads, although in some cases it easily achieved a 100% (2 to 1) advantage. Thus, for a given application workload, if users can safely consolidate 8 servers on Hyper-V or XenServer, they will comfortably be able to consolidate 14 to 16 servers on VMware ESX 4. Given the VM density superiority ESX 4 demonstrated in this study, ESX 4 appears to steadily march forward with increasing its significant cost of acquisition advantage. In this report, we'll take a closer look at what our testing of VM density revealed.

## VM Density: A Key Measure of Hypervisor Efficiency

We've long held to our position that there are marked differences between hypervisor architectures. Some of these differences are based on maturity levels and others on fundamental choices in hypervisor architecture. In the end, the architectural differences between hypervisors can have a major impact on total solution cost. Hypervisor architecture will determine efficiency in processing workloads, and architecture plays an important role in determining whether an infrastructure is flexible and adaptable.

**VM density:** The number of guest VMs that can be effectively consolidated on a hypervisor without disruptive performance impact.

The biggest influencing factor behind total solution cost is hypervisor efficiency, or what we call *VM density* in this Technology Validation report. VM density will determine the total physical equipment,

management infrastructure, and software licensing required for a virtual server infrastructure.

The reason VM density matters is simple. Just take a look around – nearly every core component and/or add-on software tool in the virtual infrastructure is licensed on a per-processor basis. A hypervisor that can run more VMs per processor, and therefore require fewer total hypervisors, can create enormous cost savings in software licensing, and hardware, power and cooling savings only increase that advantage. In an infrastructure with as few as 30 workloads, a hypervisor that can host even 50% more workloads may allow a customer to reduce their number of hypervisor instances by one-third or more. When examined in detail, as is possible using VMware's cost per application calculator ([www.vmware.com/go/costperappcalc](http://www.vmware.com/go/costperappcalc)), it is clear that VM density can have an enormous bottom-line impact.

## Testing Density: Not All Hypervisors – or Virtual Infrastructures – are the Same

While architectural differences between VMware ESX, Microsoft Hyper-V and Citrix XenServer hypervisors are now generally well understood, the impact of hypervisor architecture on VM density is often

**A note on Server Core:** Microsoft advertises Server Core and Hyper-V Server as the ideal platforms for Hyper-V because they are streamlined with smaller memory footprints than Windows Server 2008, which is the product we used in our testing. Hyper-V has suffered in the past from product limitations, many of which are being addressed in an updated R2 version that is based on the Windows Server 2008 R2 code. Unfortunately, Hyper-V Server R2 hadn't yet been released when we conducted this exercise. To meet our timeline, and have access to the appropriate features, we stuck with Windows 2008 Server Data Center R2. Even so, we wanted to fairly consider any potential advantages that Hyper-V Server R2 or Server Core may have had. Previously, we compared the Hyper-V Server / Server Core memory footprint separately and we saw a difference of about 500MB. In our testing, we actually pushed Windows Server 2008 R2 (and separately XenServer) beyond the vendor's recommendations for host OS free memory. By doing so, we were able to squeeze in an extra VM that compensated for the extra memory footprint, with no impact on total performance.

overlooked. We set out to determine the VM density supported by these three hypervisors based on a set of typical application workloads.

## Our Methodology

To evaluate VM density between Hyper-V, XenServer and ESX, we once again turned to two broadly used tests, DBHammer and SPECjbb. The first test – DBHammer – simulates randomized SQL Server workloads within multiple VMs in order to shed light on how total hypervisor differences may show up when virtualizing typical enterprise application workloads. The second test – SPECjbb – uses an atypical, extremely aggressive workload to level the playing field around memory management differences, and examine how hypervisors hold up under worst-case workloads. In performing our updated testing, we looked to SPECjbb for an additional reason as well. Now, with more capacity for virtual guests via next generation hardware, our SPECjbb testing let us peer into how a hypervisor may handle additional loads under duress, and whether in the real world users should be able to expect a density advantage under all circumstances.

To run these tests, we turned to the latest available server hardware – the Intel Nehalem microarchitecture in the form of e5540 2.53ghz processors running on HP BL460c G6 blades. Using a test harness running from an ESX 3.5 infrastructure, we ran tests against Windows 2008 virtual machines on top of the latest hypervisor versions from each of our identified

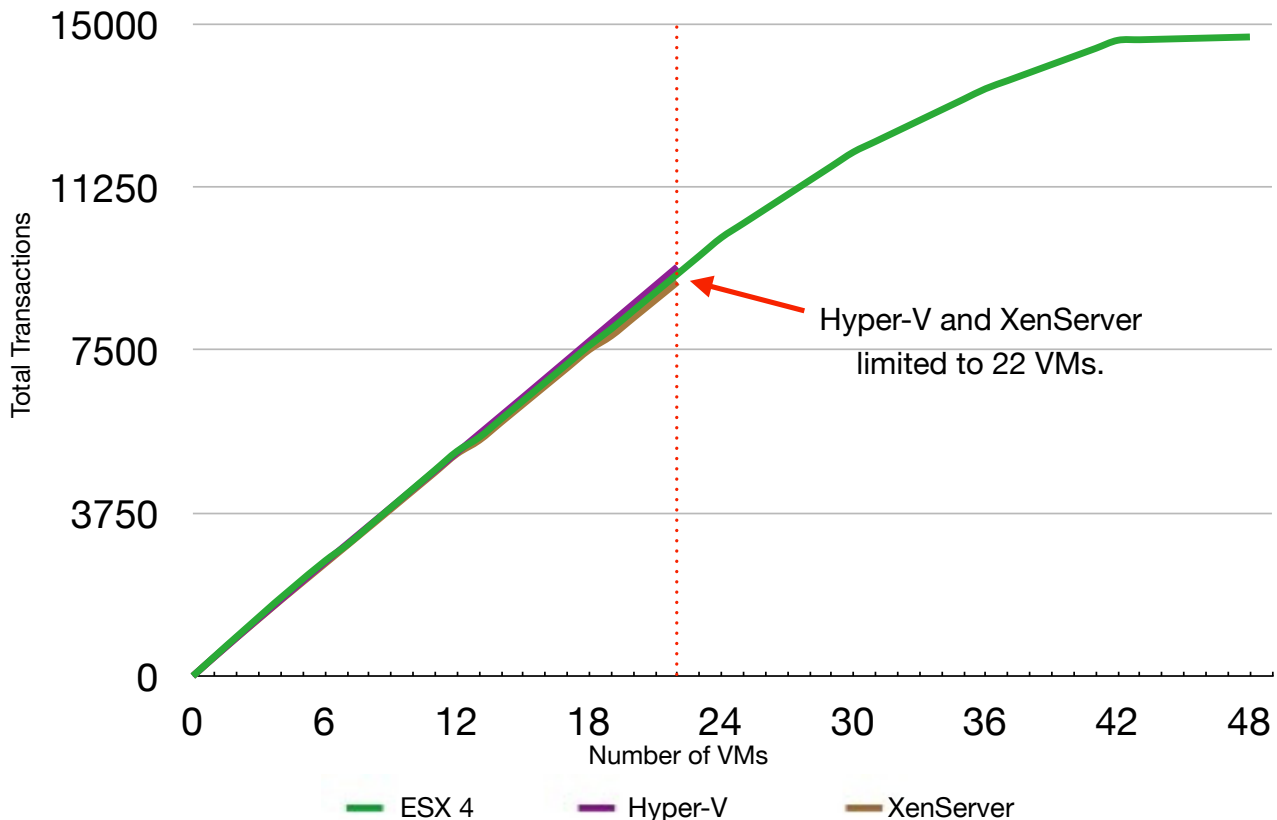
**T E C H N O L O G Y V A L I D A T I O N**

vendors. With DBHammer, our VMs were configured to use 1 vCPU and 1GB of memory, while with SPECjbb our VMs used 2 vCPUs and 2GB of memory – these configurations provided maximum performance on each platform for the respective tests. We configured 1 vNIC in each VM, attached to a gigabit Ethernet network. For storage, we used two separate pools of disk, with a dedicated LUN on each assigned to hosting half of the virtual server images in our deployed infrastructure. One

pool was made up of EqualLogic PS100E storage attached via iSCSI, and the second pool consisted of EMC Celerra NS20 disk accessed via FC. As we tested, we evenly distributed our virtual machines across the two storage pools.

We'll turn next to a discussion of the individual results from these tests, and then with this data in hand, we will discuss how VM density differs between the three leading hypervisors.

### Hypervisor Scaling under DB Hammer



\*Note, testing was performed by groups of 6 virtual guests (1, 6, 12, 18, etc.) until we reached the limit of the hypervisor. Incremental values in between have been averaged from the next highest result.

**Figure 1: ESX 4, Hyper-V, and XenServer scaling under the DBHammer Benchmark**

### The Impact of ASLR

In our previous testing, we suspected a key technology behind VMware ESX's VM density performance was memory oversubscription. In turn, we expected to find that VMware's over provisioning capabilities would be less effective, or more performance constrained, as we moved forward with Windows 2008 VM testing. Why? ASLR, or Address Space Layout Randomization – a memory space security technology that Microsoft incorporates into Windows Vista and later operating systems. Simply put, ASLR scrambles the memory layout of an operating system. This is meant to make the operating system more secure, but it simultaneously makes virtualization more difficult because ASLR makes it more challenging to trim unused memory from virtual guests. VMware memory oversubscription uses just such a trimming approach to increase the number of hosted guests within a given memory configuration. What does ASLR really do to your memory footprint? With Windows 2003, idle VMs take up 300-400 MB of memory each, while with 2008, idle VMs take up 750 to 800MB each.

Despite our anticipation, VMware impressed us. ESX has obviously taken some significant steps in squeezing every bit of performance out of the underlying hardware, as even when their memory footprint advantage was decreased in the face of ASLR, they still demonstrated the same or (in the case of SPECjbb) a significantly greater density advantage over our earlier tests, while maintaining expected performance. While memory did indicate higher usage because of ASLR, we can only conclude that hypervisor architecture improvements managed to compensate for additional memory overhead.

The three hypervisors under consideration, Microsoft Hyper-V, Citrix XenServer and VMware ESX, are quite different in capabilities and features. For example, neither Hyper-V R2 nor XenServer 5.5 supports the memory oversubscription capability included with ESX 4. To ensure as close to an apples-to-apples comparison as possible, each of our test configurations was designed to stay within the limits of all three hypervisors.

### The Tests

Based on a broad collection of data, here are the salient points relevant to our VM density testing.

**DBHammer Testing:** *Under typical workloads, ESX 4 can run twice as many VMs as Hyper-V R2 and XenServer 5.5.*

As illustrated in Figure 1 under typical workloads, VMware ESX 4 scales to twice the number of virtual machines that either Hyper-V R2 or XenServer 5.5 can host using the same amount of physical memory.

Windows 2008 implements a new memory technology called Address Space Layout Randomization (ASLR, see sidebar) that makes efficient hypervisor memory management more challenging. Even with the more challenging-to-virtualize memory footprint of Windows 2008, ESX 4 was able to run significantly more VMs compared to Hyper-V and XenServer.

As we tested Hyper-V and XenServer, we encountered a hard limit to the number of VMs we could start up. Our test system was configured with 24 physical gigabytes of

TECHNOLOGY VALIDATION

memory, and once that amount was used by the host OS or *assigned to a running VM*, no more VMs could be started. But *assigning* memory to a VM does not mean it is being fully utilized. This in turn limited VMs and created underutilization of memory, CPU, and IO. (See sidebar: ***The Usefulness of Memory Oversubscription***)

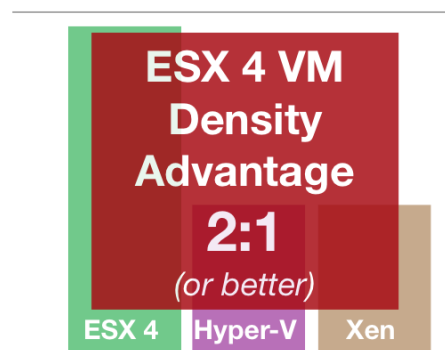
Hyper-V and XenServer both turned in a maximum result of 22 VMs. ESX added additional workloads up to 42 VMs with no performance deterioration. We in fact scaled the workloads past this number and while total system performance began leveling off at around 44 VMs, we only noted user experience deterioration (in administrator remote access sessions) when we virtualized 48 or more workloads.

While we were able to virtualize 22 workloads with XenServer and Hyper-V, we only did so by ignoring the vendors' configuration and management guidelines and performing tricks like starting all of the VMs simultaneously so they would fit within the hypervisor's available memory limit. The astute reader will note that while ESX began to saturate hardware performance at 42 workloads, if we used the more realistic limit of 21 workloads (leaving the vendor's suggested free memory on Xen Server and Hyper-V available) then ESX was running more than twice as many workloads while still delivering peak performance scaling with each incremental workload. This demonstrates how much hardware capability the other hypervisors may leave unused, and is particularly relevant with the latest server hardware.

Connecting the dots between our current testing and previous testing on ESX 3.5 yields an interesting story. Previously, memory oversubscription technologies (like memory ballooning, used to free up the memory footprint of each virtual machine) resulted in a minor hit to total performance. With ESX 4, impact upon performance is negligible until we scale beyond twice as many VMs. While this may be in part attributable to improving server memory buses, we have little doubt that this is at least in part a result of VMware's continued work on hypervisor architecture.

**Our conclusion:** Based on these typical workloads, in some situations ESX can virtualize twice as many VMs as Hyper-V and XenServer.

**DBHammer VM Density**



## TECHNOLOGY VALIDATION

**SPECjbb Testing:** *Under extreme workloads, ESX 4 was able to run at least 50% more VMs than Hyper-V R2 or XenServer 5.5. Moreover, on average, at any given number of VMs ESX 4 delivered 12% (versus XenServer) to 16% (versus Hyper-V) more total performance from the same server hardware.*

We once again turned to SPECjbb to evaluate how a hypervisor performs under extreme load. To conduct these tests, we used VMs configured with 2 vCPUs and 2GB of memory each on our Nehalem based blade infrastructure. While SPECjbb enabled us to evaluate whether ESX could efficiently exceed the hard VM Density limits of the other hypervisors, it also gave

us a feel for how well the inner architectures of the hypervisors handled extreme workload pressure. The results were telling, and are represented in the individual performance graphs in figure 2 (next page).


As is evident from the graphs, competitive hypervisors turned in less total performance, and their performance dropped off earlier in our testing. Compared to both ESX and XenServer, Hyper-V was especially erratic in delivering results. While Hyper-V improvements allowed them to deliver more total performance at their peak compared to our earlier testing, the hypervisor struggled with balancing multiple aggressive workloads. Hyper-V was only able to

### The Usefulness of Memory Oversubscription

Many users take memory oversubscription for granted while using it extensively. Based on our discussions with end users, more than 70% of VMware customers make use of memory oversubscription. Seasoned users of VMware have come to expect that they won't be restricted by physical memory when starting VMs. Many users plan around their ability to temporarily impact performance by moving active machines when they have already assigned 100% or more of a hypervisor's available memory - this is in fact the premise of many virtual server availability and disaster recovery plans today. Yet the basic capabilities of this feature remain sorely missing in competitive offerings.

Moreover, as we noticed during our testing, the story about VMware and memory oversubscription isn't just about squeezing in more servers, but in doing so *effectively*. We expected to see ASLR and SPECjbb both bring out the worst in memory oversubscription - surely, with dynamic optimization and manipulation of memory pages, there would be a big performance hit. With ASLR (see ASLR side bar) we especially expected ESX 4 to turn in poor results, because they would be operating beyond physical memory much earlier in the testing. But our testing pointed out there are two dimensions to memory management - minimizing footprint, *and* optimizing use of memory to make sure the important stuff is there. The latter is likely the reason we could see little impact from ESX VMs exceeding the 24GB of memory in our blades. Without a doubt, when competitors someday deploy memory oversubscription, ESX is setting the bar high for how effectively virtualized memory should work.

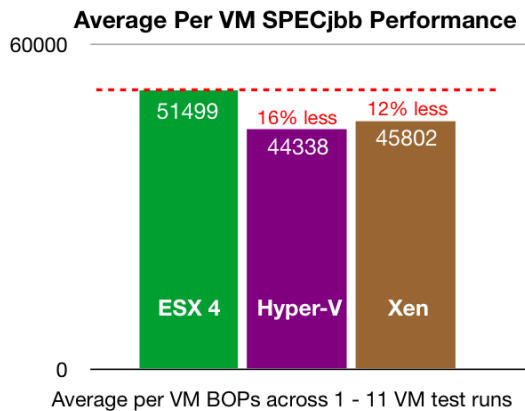
**T E C H N O L O G Y V A L I D A T I O N**



**Figure 2: How each hypervisor performed at different intervals of VMs simultaneously running SPECjbb - up to 11 VMs simultaneously running SPECjbb for each hypervisor**

deliver peak performance when the workloads could be easily balanced across vCPUs and physical cores that closely matched each other in number (e.g., 8 total vCPUS running on our 8 physical cores).

When SPECjbb was run across 11 VMs (XenServer and Hyper-V were limited to 11 VMs due to physical memory), then ESX 4 turned in 16% more performance than Hyper-V and 12% more performance than XenServer when the average performance of every test run is compared.



**Figure 3: Average SPECjbb Business Operations per Second (BOP/s) per VM across 1-11 test runs for each of three major hypervisors**

Moreover, when we virtualized more VMs with ESX, ESX 4 demonstrated with aplomb that it could operate more SPECjbb workloads than the competition, simultaneously running SPECjbb across 17 VMs. This was 1.5 times the memory associated hard limits of Hyper-V and XenServer, while still balancing the demands of all of these aggressive workloads, and demonstrating little drop off in performance.

**Conclusion:** VMware ESX was again able to effectively run more VMs – 17 versus 11 – than the other hypervisors. This translates to a VM Density advantage of slightly better than 1.5 to 1.

**Another twist - how does the hypervisor handle additional loads when saturated?**

There is more to the story that is told by our SPECjbb testing than just number of VMs successfully running the SPECjbb benchmark. More telling is a deeper examination of how each scaled past peak performance. For this examination, we'll turn to figure 4.

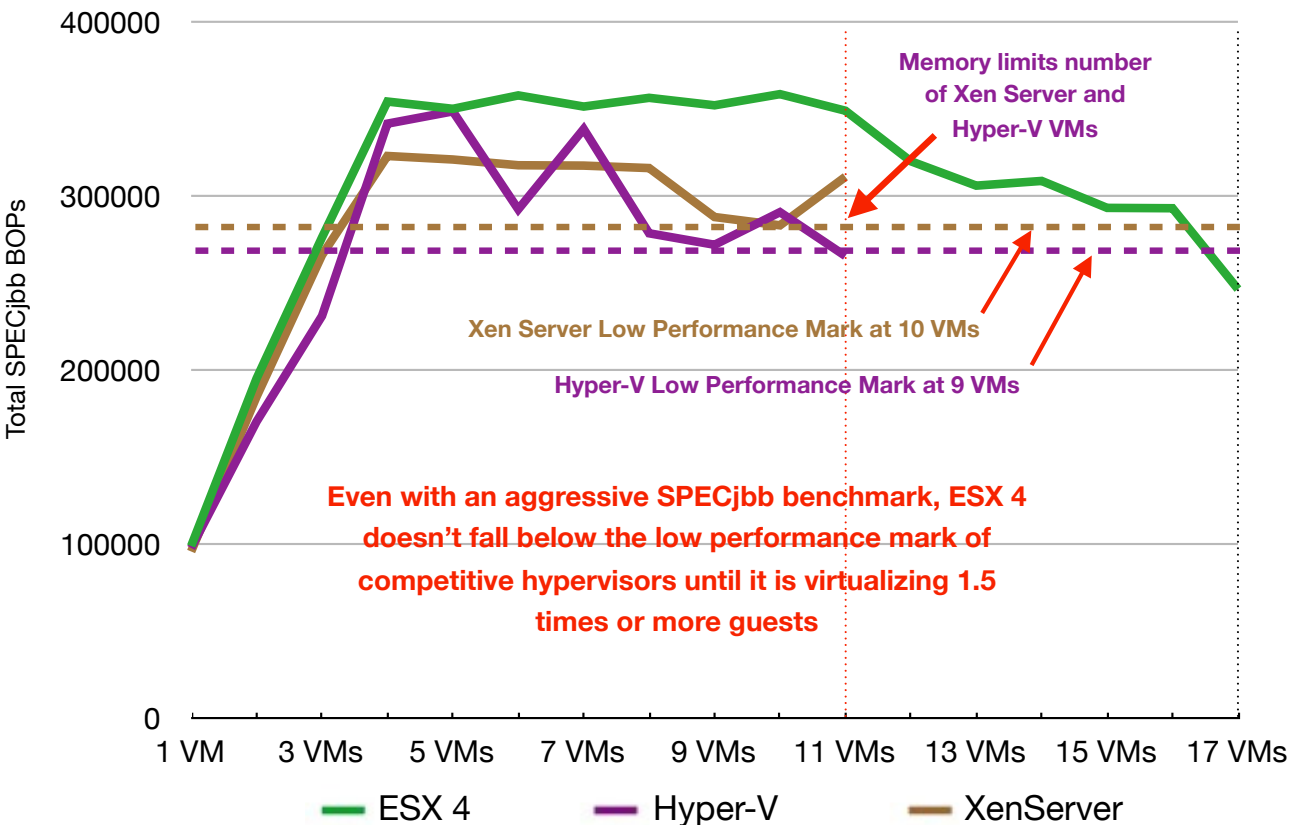
**T E C H N O L O G Y V A L I D A T I O N**

While average SPECjbb performance is telling of how efficiently each hypervisor was able to utilize underlying hardware, the trends that emerged as we added VMs to our test told an additional story. Specifically, the predictability of ESX 4 under a stressful workload shone through. As shown in Figure 4, when looking at performance at different numbers of workloads, ESX 4 performed to the limit of the hardware more consistently. Moreover, under extreme load, performance declined more predictably.

where the other hypervisors turned in their lowest performance (figure 4). VMware ESX 4 scales well *beyond* 1.5 times as many VMs before ESX 4 performance declines to the same level. Specifically, after reaching its peak, Hyper-V declined to fewer than 275,000 BOPS at 9 VMs. ESX did not decline to this performance level until it reached 17 VMs, 1.9 times the density of Hyper-V. After XenServer reached peak performance, it declined to fewer than 285,000 BOPs at 10 VMs. ESX did not decline to this performance level until reaching 17 VMs, or 1.7 times the VM density of XenServer.

In greater detail, if we look at the point

**Testing VM Density with SPECjbb**



**Figure 4: Comparative SPECjbb performance as it was scaled across the hypervisors**

In our view, ESX 4's sustained performance level under SPECjbb duress means that ESX 4 will perform more consistently as an infrastructure changes, and different dynamic workloads ramp up and down. This is especially important at scale, as more workloads might mean more variable demands. By the evidence at hand, it appears that Hyper-V in particular may deliver less performance in the face of more aggressive demands. Compared to XenServer or ESX 4, administrators planning to use Hyper-V within a pool of moving workloads should anticipate that Hyper-V will tolerate a much smaller number of total workloads before demonstrating significant service level impacts. Similarly, XenServer infrastructures will handle fewer VMs than ESX 4 will be able to service. In addition, XenServer will tolerate fewer aggressive workloads than ESX 4. Most likely, this is attributable to hypervisor maturity, and the ESX engineering team's attention to innovation around elegantly time slicing virtual processor contexts, and handling details like locking and preemption – capabilities that are rolled up under the umbrella of ESX 4's CPU Scheduler architecture.

We should point out that some of the issues around sustainable performance may be directly attributable to the way other hypervisors handle memory allocation and management – we saw repeatedly that Xen had some difficulty supporting large memory pages, through Extended Page Tables (EPT), when memory was heavily loaded. Specifically, at around 6 VMs, we were sometimes notified that EPT would be disabled. This behavior was erratic across

multiple test runs, and would not consistently happen in the same place, but always happened before reaching the limit of 11 VMs. The results shown in Figures 2 and 4 reflect the best test runs for each hypervisor.

**Our conclusion:** When digging deeper, ESX's performance gives it a distinct edge in hosting more VMs with aggressive workloads, and ESX easily handles more than 1.5 times the workloads of competitive hypervisors. In our assessment, a practical baseline for these aggressive workloads – based on where the hypervisors deliver the same levels of total performance – gives ESX 4 a VM density advantage of 1.7:1.

### **What's the Real VM Density Factor?**

We tested ESX 4, Hyper-V R2 and XenServer 5.5 to comprehensively evaluate their architectures and determine how many VMs each could effectively host with the same hardware. One of these tests yielded a VM density that was 2:1 in ESX's favor. The other test demonstrated that ESX was up to the task of virtualizing at least 1.5 times the workloads of competitive hypervisors given worst-case workloads. While no infrastructure continuously runs at 100% utilization, the point of this test was that VMware's architecture makes more efficient use of hardware resources, and performs better under duress.

Based on our observations during testing as well as our familiarity with a broad range of virtual server infrastructures, we know there are many realistic workloads under

## TECHNOLOGY VALIDATION

which ESX will demonstrate a 2:1 VM density advantage when compared to Hyper-V and XenServer. Even so, we also realize that exceptional workloads, similar to our SPECjbb testing, will occur in the real world. Consequently, we've tempered our expectations for 2:1 VM density with the 1.7:1 processing advantage seen when looking at the details of workload performance under the aggressive SPECjbb testing, and have selected a comfortable mid-point. Our conclusion is that a 1.8 VM density advantage – or 1.8 times as many VMs as Hyper-V – is a balanced mid-point that fairly compensates for the increasing frequency with which users today are applying server virtualization to more aggressive production workloads.

**Our findings: VMware ESX 4 has at least a 1.8 to 1 VM density advantage over Hyper-V R2 and XenServer**

### Taneja Group Opinion

As we have found in this Technology Validation exercise, the variety of new hypervisors in the marketplace – some of which are offered free of charge – does not mean the hypervisor is becoming commoditized. Reality is in fact far from this – as we've observed, there are distinct differences between hypervisors. Clearly VMware's time in the market has led to a level of hypervisor maturity that pays off in performance and efficiency with hardware.

While Microsoft has achieved some impressive core performance gains with its release of Hyper-V R2, ESX 4 still holds a significant advantage in VM density across the set of application workloads we tested. Our latest round of tests indicates that ESX 4 has a similar VM density advantage over Citrix XenServer 5.5. That advantage varies between 1.5 to 1 and nearly 2 to 1, depending on the specific benchmark. We believe that the architectural design of ESX 4 is largely responsible for the greater VM density, providing it with an inherent advantage over the parent partition architectures employed by Hyper-V and XenServer.

While VM density is an incredibly important factor in determining acquisition costs, in the dynamic data center, where workloads can fluidly move across multiple hardware systems and enhance how users address maintenance windows and unplanned outages, sustainable performance under load is at least as important. Sustainable performance will no doubt determine how effectively virtualization can be used, as well as the ongoing cost of operating a broader pool of hypervisors.

We routinely see users today implementing server virtualization across a pool of resources, where multiple other hypervisors are able to absorb the load of any single hypervisor that may fail or be turned off for maintenance. But this key capability in the virtualized infrastructure relies upon having enough headroom to deal with additional workloads, and to handle their unexpected movement, or unexpected growth over

## TECHNOLOGY VALIDATION

time. If your hypervisor isn't built for this, it may mean you have bigger issues over the long run than just total cost of acquisition. Limited scalability may mean you continue to operate with a hardware centric focus, where you can't with assurance move your virtual servers, and you still must manage within server silos, and to the limits of server hardware. ESX's ability to support a higher VM density lets you safely operate at higher levels of utilization where host failures and VM migrations can be absorbed with less planning and hardware.

In an economic environment that is repeatedly heralding the message "do more with less", the efficiency of hypervisors is an oft-overlooked aspect of virtual infrastructure acquisition that has massive impact on total price. As one of the key measures of hypervisor efficiency, VM density plays a big role in determining total

solution cost. Today, VM density advantages in vSphere 4 can yield total cost savings of up to 29% over equivalent Microsoft installations. Even in small implementations, this is a difference of many thousands of dollars.

We recommend that every organization undertaking a server virtualization initiative sharpen their pencils and review our VM density findings in this report. Then use our findings, along with the associated cost per application assessment approach included in our previous (March 2009) report, or VMware's cost per application calculator ([www.vmware.com/go/costperappcalc](http://www.vmware.com/go/costperappcalc)) to identify the true cost of the virtualization solutions you are considering. Viewing the multitude of virtualization solutions on the market through this lens will help you figure out how to do more with less.

---

***NOTICE:** The information and product recommendations made by the TANEJA GROUP are based upon public information and sources and may also include personal opinions both of the TANEJA GROUP and others, all of which we believe to be accurate and reliable. However, as market conditions change and not within our control, the information and recommendations are made without warranty of any kind. All product names used and mentioned herein are the trademarks of their respective owners. The TANEJA GROUP, Inc. assumes no responsibility or liability for any damages whatsoever (including incidental, consequential or otherwise), caused by your use of, or reliance upon, the information and recommendations presented herein, nor for any inadvertent errors which may appear in this document.*