



How far will you take virtual?

Virtualization Reality: Why Microsoft® Virtualization Solutions Deliver Value When Compared to VMware®

This white paper describes why the VMware cost calculator, based on a white paper by the Taneja Group, does not produce the real-world calculations that it claims.

Copyright © 2009 MICROSOFT CORPORATION

How far will you take virtual?



- Overview** 3
- Why the Taneja Group white paper does not add up** 4
 - Microsoft® System Center provides more functionality than VMware® 4
 - Real-world implementations do not have the ratios claimed 4
 - Tested VMs do not represent real-world VMs or workloads 5
 - The benchmarking tools used do not stress the features tested 6
 - The cost calculations misrepresent the actual cost comparison 7
 - Taneja didn't include the cost of VMware vCenter Server™ 7
 - The cost calculations bought more servers for capacity, rather than more memory..... 7
- Why Memory Overcommit should not be overused** 8
 - VMware does not recommend allocating more than is physically available 8
 - When VM performance suffers because of overcommit, the app performance will also suffer9
 - DRS does not resolve Memory Overcommit issues 9
- Conclusion** 10
 - More Information 10

How far will you take virtual?



Overview

In March 2009, VMware® released a new cost-per-application calculator on its company web site, based on a white paper written by the Taneja Group (which can be downloaded at www.tanejagroup.com), which produces a direct capital cost comparison to Microsoft® virtualization solutions. However, the actual calculations seem contrary to both the cost comparisons done by Microsoft and the real-world recommendations and results that many customers have been experiencing throughout the industry. A review of that white paper was necessary to see how VMware came to those conclusions.

One of the key issues with the white paper is that the Taneja Group and VMware do not factor in the additional benefits and capabilities provided by the Microsoft server virtualization solution, including Microsoft Hyper-V™ and the Microsoft System Center suite. System Center provides detailed in-guest knowledge of the virtual machines (VMs)—something VMware does not have—and System Center allows optimizations to leverage that knowledge. Finally, the Microsoft solution has a cost advantage when compared to the VMware solutions, which would require purchasing additional third-party software.

Another issue with the white paper, and thus with the calculator itself, is the reliance on a function in VMware ESX called Memory Overcommit, which is a group of features that allows the host's VMs to be allocated more memory than is physically available on the host. While this is certainly not a new function of ESX, the ratios of Memory Overcommit and the testing methodology used in the white paper produced results well beyond what is reasonable or recommended even by VMware.

Why would higher memory overcommit ratios be important to VMware? As more and more customers have implemented Microsoft virtualizations solutions for their performance and value, VMware is under increasing pressure to justify the cost and average sale price of VMware software. By changing the cost discussion from the cost of software to a calculation based on the number of VMs per server ("per app" or "per VM"), VMware shifts the focus away from the actual cost of buying their software. This shift is especially important since the VMware server virtualization solution is not as cost-effective as Microsoft server virtualization solutions. VMware then leverages overcommit and other variables to reach consolidation ratios not seen in the real world, to make the cost comparisons work.

This white paper will present why both the calculations and the testing methodology in the Taneja Group whitepaper are not representative of real-world configurations, why the overuse of Memory Overcommit can cause problems in real-world deployments, and why the Microsoft server virtualization solution with Hyper-V and System Center is still the most cost-effective.

How far will you take virtual?



Why the Taneja Group white paper does not add up

As we take a closer look at the cost comparison and performance scaling white paper, there are several key points we must review in order to understand why the conclusions do not add up.

Microsoft System Center provides more functionality than VMware

One of the most important points is that the System Center suite, with Microsoft Configuration Manager 2007, Microsoft Operations Manager 2007, Microsoft Data Protection Manager 2007, and Microsoft Virtual Machine Manager 2008, provides significantly greater functionality than VMware vSphere™ 4. The System Center suite provides both host management and in-guest management of all VMs. This allows Microsoft customers to have an end-to-end solution, with a single set of applications to manage both the physical and virtual machines.

Also, the level of functionality provided by the System Center products far exceeds VMware, including integrated software distribution and patch management (including virtualized applications with Microsoft App-V), compliance, system, application, and service level monitoring, full backup and recovery of hosts, the VMs, and the applications inside the VMs, as well as integrated VM management throughout the suite.

In addition to the application and service level knowledge that System Center brings, the Microsoft solution also manages VMware systems through Virtual Machine Manager 2008 (VMM). By using VMM, customers can integrate the management of their entire virtual infrastructure with a single set of tools, providing not only the end-to-end management described above but also multi-hypervisor management.

Finally, without the use of the System Center suite, VMware customers would be required to add all the additional third-party applications for in-guest management, such as backup and software distribution. Many of these third-party applications are charged on a per-VM basis. By contrast, the System Center suite is licensed at the host level, not the VM level. When using the VMware solution, you must purchase additional software, which is charged at a per VM cost, while the Microsoft solution includes all of those functions and is a fixed cost at the host level. Thus, as you scale more VMs per server, the capital cost benefits with Microsoft increase and the benefits and TCO of VMware decreases.

Real-world implementations do not have the ratios claimed

In its white paper, the Taneja Group claimed VMware ESX achieved VM density of up to 100 percent more than Hyper-V, implementing twice as many VMs (with twice the memory allocated) than could otherwise be implemented based on the actual memory on the physical host. This ratio is achieved through the use of the Memory Overcommit (or Memory Oversubscription) feature set found in VMware ESX. As mentioned earlier, Memory Overcommit is a group of features that allows more memory to be allocated to the host's VMs than is physically available on the host. While there are certain scenarios where Memory Overcommit is useful, the

How far will you take virtual?



increased ratios described in the white paper cannot be unconditionally applied to all scenarios, especially in production.

According to our research, customers just don't see the ratios VMware use in their cost calculators. As Tom Bittman of Gartner stated in ITworld.ca:

"On average, what we've seen is that VMware users have about 120 percent of virtual memory configured to physical memory, so a 20 percent higher density than they would have if they didn't have Memory Overcommit... So when [the Taneja Group] says eight to 12 or eight to 16, they are giving ESX too much credit. I think it's real, but it's not that wide of a gap."

It also goes beyond the Memory Overcommit ratios. Most companies are not reaching the consolidation ratios for which the paper leverages Memory Overcommit. Industry wide, the actual consolidation ratios (the numbers of VMs per physical host) is still around 10–12 VMs per system, not 24. As ESG Research commented on CNET:

"The average number of virtual machines per physical server is between 5 and 10. This data was gathered from a survey of over 200 enterprise organizations (i.e. organizations with 1,000 employees or more)."

Finally, Microsoft customers are easily running 11:1 and 15:1 consolidation ratios with Windows Server® 2008 Hyper-V. As examples, read the following case studies, where you'll see that current uses of server virtualization do not reach the levels of consolidation that require the use of Memory Overcommit features.

[SABMiller RUS](#)

[Caja de Compensacion Compensar](#)

[Ingersoll Rand](#)

Tested VMs do not represent real-world VMs or workloads

In the Taneja Group white paper, the VMs were configured to use only half the memory they were allocated. Thus, inherently, the VMs were configured to reach the ratio desired (configuring VMs to only use half the allocated memory leads to the 2:1 ratio). It would be difficult to assume that all real-world VMs would only use half the memory that is allocated to them. Also, while a SQL Server®-based benchmarking tool was used, the Taneja Group used a static workload for the testing, ensuring that the memory usage would never spike and cause the significant performance drops that occur when Memory Overcommit actually uses more memory than is available. Very few SQL administrators would create the VMs in the configurations that were used in the test (VMs with 1GB allocated with only 512MB of RAM in use by SQL).

How far will you take virtual?



Real-world applications have base static loads with peaks of memory utilization based on the application. Over-allocating memory to a VM can result in under-utilization of the resources on the host system, which goes against one of the driving factors for using virtualization for server consolidation in the first place. A solution to prevent this is to allocate the right amount of memory for the VM, use in-guest application aware monitoring to watch the application (a capability VMware does not have), and utilize that in-guest knowledge to optimize the virtualization hosts and clusters.

Microsoft is uniquely positioned to provide both application-level knowledge and support, allowing a greater level of integration with virtualization. With the Microsoft System Center suite, Microsoft enables a feature called PRO, Performance Resource Optimization, which allows optimization at both a host and application level. This allows administrators to optimize and manage VMs.

The benchmarking tools used do not stress the features tested

As noted above, the Taneja Group used a SQL benchmarking tool, DBHammer, to simulate server load. While this does simulate SQL server usages and loads, the test was used with a fixed load that did not vary the memory usage, which is not characteristic of a real-world SQL database. If the memory load were varied, the actual amount of memory used would change and the scaling and performance of the VMs would have changed dramatically.

The white paper also used another test, SPECjbb, for the actual testing of CPU performance of the hypervisors. However, SPECjbb is not an effective way to benchmark a hypervisor due to several reasons:

- **No hypervisor involvement:** SPECjbb2005 from Standard Performance Evaluation Corporation (SPEC) evaluates the performance of server-side Java by emulating a three-tier client/server system with emphasis on the middle tier, and not focused on the kernel level. This means the hypervisor is not used extensively to get comparable results.
- **No I/O measurements:** Instead of testing with a possible disk-intensive database system, SPECjbb uses tables of objects, implemented by Java Collections, rather than a separate database.

How far will you take virtual?



The cost calculations misrepresent the actual cost comparison

Independent of the technical considerations discussed above and below, the actual cost comparisons are not applicable for what most customers would do.

Taneja didn't include the cost of VMware vCenter Server

In its cost comparison, the Taneja Group compared VMware VI3 versus Windows Server 2008 Hyper-V with the System Center suite, which includes Configuration Manager 2007, Operations Manager 2007, Data Protection Manager 2007, and Virtual Machine Manager 2008. The actual cost comparison didn't include the cost for vCenter Server, which is required to manage the ESX servers. This is an additional cost.

In addition to omitting the management server costs, the Taneja Group also omitted the VMs needed to manage the server. While the Microsoft comparison has additional VMs for System Center servers, the additional systems required for VMware vCenter Server and other VMware tools (such as backup and patching) are not included. This would result in either additional VMs or physical machines just for VMware management.

The cost calculations bought more servers for capacity, rather than more memory

When performing the final calculations, Taneja and VMware are assuming customers will buy more servers, rather than simply adding memory to the system. Why would they do that? They have made that assumption because the cost calculator attempts to add power, space, and OS licensing costs, which would only increase by having more servers. In fact, the only way that VMware comes close to being as cost-effective as Microsoft is by adding more servers. Most enterprise customers will increase existing server capacity by adding memory, CPU, or storage, with the smaller incremental cost, rather than simply ordering more servers.

How far will you take virtual?



Why Memory Overcommit should not be overused

One of the fundamental questions asked about using Memory Overcommit is: how does the host allow VMs to use more memory than the host actually has? The answer is that it really doesn't. Just like other technologies that expand memory beyond what is physically installed, the Memory Overcommit features give the VMs the impression that they have more memory than has been allocated. If the total amount of memory used actually exceeds the host's physical memory, ESX must use disk paging to support the additional requests.

What this means is while Memory Overcommit has its uses in certain configurations, it's not a panacea for virtualization administrators. Memory Overcommit should be used judiciously and only when absolutely necessary. Even when choosing to use Memory Overcommit, the following considerations should be kept in mind.

VMware does not recommend allocating more than is physically available

Until recently, VMware recommended against using Memory Overcommit in production environments. Memory Overcommit was originally implemented to work around scaling issues in VMware ESX 2.X, where overcommitting memory was necessary. Even today, there are specific warnings about using Memory Overcommit in VMware's own performance-tuning document, **Performance Tuning Best Practices for ESX Server 3**, including the following:

Make sure the host has more physical memory than the total amount of memory that will be used by ESX plus the sum of the working set sizes that will be used by all the virtual machines running at any one time.

and

If the working set (active memory) of the virtual machine resides in physical memory, using the swapping mechanism and having inactive pages swapped out does not affect performance. However, if the working set is so large that active pages are continuously being swapped in and out (that is, the swap I/O rate is high), then performance may degrade significantly.

These recommendations point to what has been stated above. Memory Overcommit should only be used if you don't actually use the memory you are committing. If the VMs use all physical the memory, the performance will suffer. Thus, the same result can be achieved with properly sized VMs (as VMware recommends before Memory Overcommit).

How far will you take virtual?



When VM performance suffers because of Memory Overcommit, the app performance will also suffer

One unintended result of poor use of Memory Overcommit is that when disk swapping occurs, the VMs will suffer significantly and the applications will bear the brunt of the problems. Thus, the initiation point of issues moves from the VM level to the application level, and users/customers will be the first to notice there are issues with VM.

This problem is further exacerbated by the lack of in-guest application and service monitoring within VMware vCenter Server. Unless the application monitoring group is also the VM management group, or the VM management group is using the System Center family of products, the VM administrator will not be aware of problems within the VM. While the administrator will see some metrics indications of memory and disk usage, the understanding of what the impact on the actual application is missing. By contrast, with the Microsoft solution, the System Center suite provides administrators with management of both the application/service level view and the host/VM view.

DRS does not resolve Memory Overcommit issues

When discussing issues with performance problems related to Memory Overcommit, the VMware Distributed Resource Scheduler (DRS) feature is often cited as the solution to the issue. The problem with relying on DRS to solve issues is that DRS is not Memory Overcommit-aware. When using DRS to place VMs, actual memory usage is not factored into the calculation. Only the requested memory is factored into the calculation. Thus, DRS, in placement, may exacerbate the memory commit problem, not prevent it.

When using DRS to perform post placement optimization, even greater care needs to be taken. Unlike the PRO feature of System Center Virtual Machine Manager, DRS can only monitor CPU utilization and thus, cannot react automatically to application issues. In addition, since DRS does not distinguish between overcommitted and non-overcommitted resources, the VM could be moved to another host, which would then be overcommitted, triggering another DRS action. If the server resource pool is fully maximized and overcommitted (which VMware requires to reach the high ratios described in the Taneja Group white paper), the system could DRS VMs all over the pool, causing what many administrators call "vMotion sickness." This results in reduced performance of not only the pool but of vCenter Server itself.

How far will you take virtual?



Conclusion

It is clear that the new cost calculation and technical assertions made by both VMware and the Taneja Group are not as simple and significant as they make them out to be. As the analysts and Microsoft case studies show, enterprise customers are not seeing the ratios of overcommitted VMs that were claimed in the white paper. Combined with the unrealistic comparison scenarios in the calculations in both the white paper and the calculator, it is easy to see how Microsoft still provides the cost value that customers are looking for from virtualization.

The best way to determine the true cost of Microsoft Virtualization is to try it out for your own company. Enabling Hyper-V in Windows Server 2008 is as simple as enabling the role in Server Manager. The free, standalone hypervisor, Microsoft Hyper-V Server 2008, is available for download from www.microsoft.com.

More Information

[Microsoft Virtualization Web Site](#)

[Microsoft Windows Server 2008 Hyper-V Web Site](#)

[Microsoft System Center Web Site](#)

Costs are based on a comparison of Windows Server® 2008 Enterprise with Hyper-V™ and Microsoft® System Center Server Management Suite Enterprise with VMware's VMware Infrastructure Enterprise with VMware vCenter Server. Includes 2 years support costs for both. Based on Microsoft estimated retail prices and published VMware prices available at www.vmware.com/vmwarestore as of 02/04/2009. Actual reseller prices may vary.

The information contained in this document represents the current view of Microsoft Corporation on the issues discussed as of the date of publication. Because Microsoft must respond to changing market conditions, it should not be interpreted to be a commitment on the part of Microsoft, and Microsoft cannot guarantee the accuracy of any information presented after the date of publication.

This white paper is for informational purposes only. MICROSOFT MAKES NO WARRANTIES, EXPRESS OR IMPLIED, IN THIS DOCUMENT.

Complying with all applicable copyright laws is the responsibility of the user. Without limiting the rights under copyright, no part of this document may be reproduced, stored in, or introduced into a retrieval system, or transmitted in any form or by any means (electronic, mechanical, photocopying, recording, or otherwise), or for any purpose, without the express written permission of Microsoft Corporation.

Microsoft may have patents, patent applications, trademarks, copyrights, or other intellectual property rights covering subject matter in this document. Except as expressly provided in any written license agreement from Microsoft, the furnishing of this document does not give you any license to these patents, trademarks, copyrights, or other intellectual property.

© 2009 Microsoft Corporation. All rights reserved.

Microsoft is a registered of Microsoft Corporation in the United States and/or other countries.

The names of actual companies and products mentioned herein may be the trademarks of their respective owners.