

Microsoft Office SharePoint Server 2007 Performance on VMware vSphere 4.1

VMware vSphere 4.1

One of the key benefits of virtualization is the ability to consolidate multiple applications onto a single physical server. The importance of this can be seen clearly with multi-tier applications, where the load placed on the servers by the individual tiers often varies widely, leaving many servers under-utilized. Virtualization allows multiple tiers to be consolidated on a single server, thus lowering the cost of deployment. VMware vSphere™ 4.1 supports this consolidation with high performance and advanced features that can actually improve the overall user experience.

Microsoft Office SharePoint Server is a prime example of a multi-tier application that will benefit from virtualization on VMware vSphere 4.1. In this paper, we present the results of experiments which demonstrate the high performance and scalability that can be achieved when deploying SharePoint on vSphere. These experiments, in which up to 171,600 heavy SharePoint users were supported on a single physical server, also highlight the benefits gained by the ability to easily deploy additional SharePoint virtual machines (VMs) as needed to satisfy changing demands. This paper also discusses some of the advanced features of VMware vSphere 4.1 that allow it to efficiently virtualize resource-intensive and latency-sensitive applications such as SharePoint. The paper concludes with a set of recommended best-practices for achieving optimal SharePoint performance on VMware vSphere 4.1.

Experiment Configuration and Methodology

In order to understand the performance of SharePoint in a virtualized and consolidated environment, we performed a series of experiments using SharePoint on vSphere 4.1. In these experiments, we ran the Web server, application server, and database tiers in VMs consolidated on a single physical server (the system-under-test or SUT). In the following sections, we give a detailed description of the testbed configuration and describe the load drivers and workload used to perform these experiments.

Testbed Configuration

In our testbed configuration, the SUT was a Dell PowerEdge R710 server, which was configured with two quad-core Intel Xeon X5570 processors and 96GB of physical memory. SharePoint Server 2007 SP2 was installed on Windows Server 2008 R2. VMware vSphere 4.1 was used to virtualize the Microsoft Office SharePoint Server 2007 environment. Table 1 details the testbed configuration.

Table 1. Hardware and software configuration

Component	Details
System Under Test (SUT)	
Server	Dell PowerEdge R710
Processors	2 quad-core Intel Xeon X5570 (Nehalem) processors @ 2.93GHz, with hyper-threading
Memory	96GB DDR-3 DIMM
HBA	QLogic 8Gb, dual-port QLE2562 adapter for Fibre Channel
Virtualization Software	VMware ESX 4.1
Guest Operating System	Microsoft Windows Server 2008 Datacenter Edition R2 (64-bit)
Application	Microsoft SharePoint Server 2007 SP2 (64-bit) and SQL server 2008 SP1(64-bit)
Storage	
Storage Enclosure	EMC CX3-40 with 2 storage processors
Hard Drives	60 146GB 15K RPM drives
Clients	
Used 2 machines to drive the workload	
Client 1 and 2	Dell PowerEdge 1950, 1 quad-core Intel Xeon 2.66 GHz X5355 with 16GB memory
Operating System	Microsoft Windows Server 2003 Enterprise Edition R2 (64-bit)
Application	Microsoft Visual Studio Test Agent 2008
Server hosting Active Directory and DNS services	
Server	Dell PowerEdge 2950
Processors	2 quad-core Intel Xeon 2.66 GHz X5355
Memory	32GB DDR-2 DIMM
Operating System	Microsoft Windows Server 2008 Enterprise Edition (64-bit)

Figure 1 shows the physical layout of the testbed. We used two Dell PowerEdge 1950 machines to serve as the client machines. The SharePoint server was a Dell PowerEdge R710. Active Directory and DNS were run from a Dell PowerEdge 2950.

Figure 1. Physical layout of the testbed





Figure 2 shows how the virtualized SharePoint farm is deployed as multiple virtual machines on a single host.

Figure 2. Virtualized testbed layout; servers comprising SharePoint are consolidated onto a single host



Test and Measurement Tools

We ran Microsoft Visual Studio Team System 2008 Test Load Agent (VSTT) on two client systems to simulate the Web requests from SharePoint users.

VSTT provides tools to set a distribution percentage for the transaction mix and reports the number of requests-per-second passed along with the average response-time during load testing. We set the think-time to zero to allow VSTT to create a continuous workload on the system for load testing.

We used a workload mix of 80% read, 10% search, and 10% modify transactions. The average response-time for each transaction type was required to be less than 3 seconds over a 1 hour period. The reported requests-per-second was used to compute the maximum number of users that could be supported for the configuration. More details about the computation of the maximum users will be given in the results section.

In order to generate sufficient load on the system, we used two client machines. One client was installed with the VSTT controller and agent, and the other was dedicated to the VSTT agent.

For VMware vSphere 4.1, we used esxtop to record both ESX server and virtual machine related performance counters. We configured esxtop to log CPU, memory, disk, and network and system counters during the runs.

Table 2. Test and measurement tools

Test Tool	Measurement
Microsoft Visual Studio Team System 2008 Test Agent	Requests per Second and Average Response Time
VMware esxtop	VMware ESX 4.1 performance statistics for server and VMs

Experiment Results and Performance Analysis

In order to maximize the performance of a multi-tier application, it is necessary to provide sufficient resources to the tier which places the highest demands on the underlying infrastructure. In our experiments with SharePoint on vSphere, we initially configured SharePoint with 3 VMs, using a single Web, application, and database server. Using performance monitoring tools provided by vSphere, we identified the Web server as the performance bottleneck in our initial configuration. We then used the facilities provided by vSphere to quickly deploy additional instances of the Web server. As a result, we were able to maximize the performance and utilization of the physical server.

In the following sections, we discuss our test cases, show how we were able to maximize the utilization of the physical server by scaling out the Web server tier, and discuss our results for the maximum number of SharePoint users supported.

Test Cases and Test Method

The VM configurations that we used in our experiments are shown in Table 3. The initial configuration used one VM for each of the Web, application, and database tiers. Based on the initial performance data, we then scaled out the Web tier to use two and then three VMs. As will be shown in the next sections, this enabled us to maximize the number of SharePoint users supported on a single physical server.

Table 3. Number of 2 vCPU VMs used for scale-out and corresponding number of physical servers

Number of VMs in Testbed	Configured as Number of Servers
3 VMs	1 Web Server, 1 Application Server, 1 SQL Server
4 VMs	2 Web Servers, 1 Application Server, 1 SQL Server
5 VMs	3 Web Servers, 1 Application Server, 1 SQL Server

The Web server VMs were configured with the Web Server role, as well as a Query Server role to support searching. Microsoft Network Load Balancing (NLB) was installed on each Web server VM with equal load and multicast-mode settings. An Application server VM was dedicated as an index server and had completed full indexing before our experiments. The database tier used SQL server, and had 200GB of SharePoint user data. All VMs had large pages enabled in the guest OS, and used the VMXNET3 virtual network driver. Table 4 lists the SharePoint VM server roles and configuration.

Table 4. VM configuration and SharePoint server roles

Type of Server	SharePoint Role	CPU and Memory Configuration
Web	Web Server and Query Server	2 vCPUs with 4 GB memory
Application	Index Server	2 vCPUs with 4 GB memory
Database	Database Server	2 vCPUs with 16 GB memory

Maximizing SharePoint Throughput


As will be discussed in detail in the next section, we determined the peak number of SharePoint users supported by a given configuration based on the average operation response-time. In our initial configuration, the maximum number of supported users was reached when the overall system CPU

utilization was still relatively low. The results in this section show how we were able to use CPU utilization data from the individual VMs to identify the bottlenecked service in the SharePoint farm. We then show how we were able to maximize the utilization of the underlying server by replicating the bottlenecked service in multiple VMs.

Virtual Machine CPU Utilization

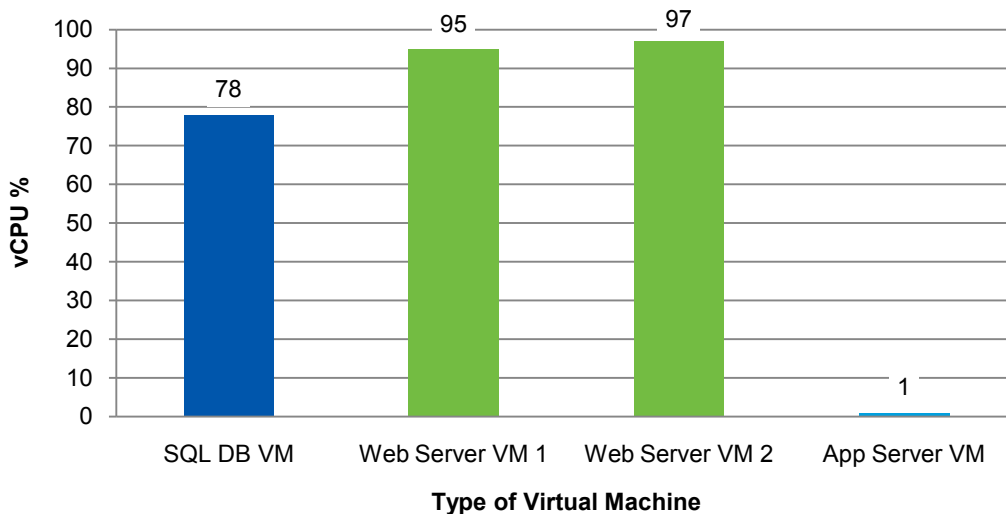
Figure 3 shows that in the 3 VM configuration (1 SQL Server VM + 1 Web server VM + 1 Application server VM), the single Web server is the bottleneck with a 100% vCPU utilization. The SQL Server has only 43% CPU usage, and the utilization of the Application server is minimal.

Figure 3. Percentage of guest vCPUs utilized in the 3 VM configuration (1 Web server)




After adding a second Web server VM (the 4 VM case), the Web servers have 95% and 97% CPU utilization, while SQL server CPU usage increased to 78%. This indicated that the SQL Server VM could support more load while the Web servers remained a bottleneck.

Figure 4. Percentage of guest vCPUs utilized in the 4 VM configuration (2 Web servers)



After adding a third Web server (the 5 VMs case), the average CPU utilization of the Web servers is 88%, and the utilization of SQL Server is 96%. The SQL Server VM is close to saturation, and the Web servers are no longer the bottleneck.


Figure 5. Percentage of guest vCPUs utilized in the 5 VM configuration (3 Web servers)



Host CPU Utilization


The following chart shows the average physical host CPU utilization when supporting the maximum number of users in the 3, 4, and 5 VM cases. The increased CPU utilization on the host means that the host could process more requests without the response time of users going over 3 seconds. As a result, more SharePoint users were supported on the host without requiring an additional physical machine.

Figure 6. Average of host CPU utilized for 3, 4, and 5 VMs during test run



Maximum Number of SharePoint Users Supported

The maximum number of SharePoint users support by each of the 3, 4, and 5 VM configurations is shown in Figure 7. As discussed below, these are “heavy” users as defined by Microsoft. The maximum number of users supported in the 3 VM case is 72,600 heavy users. The 4 VM case can support 135,600 heavy users, an 87% increase. The 5 VM case can support up to 171,600 heavy users, a 136% increase over the 3 VM case. These scale-out experiments demonstrate that a virtualized SharePoint environment can support a large number of users by increasing the utilization of its physical server.

Figure 7. Number of heavy users at 1% concurrency supported at each stage of scale-out

We used Microsoft's "Estimating throughput targets" guideline (<http://technet.microsoft.com/en-us/library/cc261795.aspx>) to determine the maximum number of heavy users that could be supported while ensuring average response-times remained within acceptable limits. The acceptable user response-time was set to under 3 seconds for each type of transaction. The transaction mix was 80 (read) / 10 (write) / 10 (search). As per Microsoft standard, a heavy user performs 60 requests per hour (RPH), that is, one request every 60 seconds. If you have 1% of users using the system simultaneously, this would be 6,000 total users. The number of users shown is at 1% concurrency, which means that for every 100 users on the system, 1 user is involved in a transaction. The maximum user capacity is derived from the following formula:

seconds per hour / RPH / Concurrency% * RPS = # of users

Example for RPS of 28.6:

$3600 / 60 / 1\% * 28.6 = 171,600$ (supported user capacity for 1% concurrency)

Benefit of Virtualization Features in a Consolidation Environment

VMware vSphere 4.1 has a number of advanced features that can help to maximize performance and overall resources utilization. In this section, we examine the performance benefit of some of those features for a SharePoint deployment.

Memory compression—Memory over-commitment occurs when more memory is allocated to virtual machines than is physically present in a VMware ESX host. Using sophisticated techniques, such as ballooning and transparent page sharing, ESX is typically able to handle memory over-commitment without any performance degradation. However, if more memory is being actively used than is present on the server, ESX may have to resort to swapping out portions of a VM's memory.

Memory compression is a feature that reduces the performance penalty of ESX-host-level swapping in memory over-commitment scenarios. vSphere 4.1 implements a compression cache to store a compressed swap candidate page. When a guest page needs to be swapped, ESX will first try to compress the page. If the page can be compressed to 2KB or less, the page will be stored in the compression cache. Otherwise, the page will be swapped out to disk, which could cause a significant performance impact when such pages are referenced later. The maximum compression cache size by default is 10% of configured VM memory size.


Figure 8 presents how memory compression helps improve SharePoint performance when host memory size is reduced. The total memory size of the 5 SharePoint VMs was 32GB. The SQL server VM was given a full 16GB memory reservation because any memory reclamation through ballooning or swapping from this VM significantly degraded the SharePoint performance. The Application and Web server VMs were not

given any reservation. For comparison, we also present the results of different memory over-commit levels without using the memory compression feature.

The results:

- 36GB to 28GB cases: There is only 6% performance degradation. This is because paging sharing and ballooning techniques efficiently reclaim the memory with very little performance penalty.
- 26GB case: In the case without memory compression, performance is severely reduced to only 5% of the original throughput. With memory compression, the system can still achieve 93% of the original throughput. The large improvement in the memory compression case is mainly a result of the fact that most of the missing pages can be referenced in the compression caches, thus eliminating a large amount of swapping from disk.
- 24GB case: Even with memory compression in this really low host memory situation, the throughput drops to 25% of the throughput in the 36GB case. This is because the compression caches of the VMs reach their maximum size in this case.

Figure 8. Throughput compression with and without Memory Compression



- **NUMA-aware resource management**— vSphere CPU and memory resource scheduler has built-in optimizations for applications running on servers with a NUMA architecture. The intent of the NUMA load-balancer in vSphere is to balance CPU and memory load between NUMA nodes while maximizing local memory accesses, which are faster than remote accesses. For this purpose, a virtual machine is associated with a home node where most of its CPU and memory demands are satisfied. If a virtual machine's home node is more heavily loaded than other nodes, the load balancer will migrate the virtual machine to a less loaded node to improve overall performance.

This NUMA-aware resource management improves throughput by 8% in the 5 SharePoint VMs case when compared with the case where the NUMA-aware optimization is disabled. The NUMA load-balancer also provides fairness in resource distribution for multiple VMs. The CPU utilization for the 3

Web server VMs was within 1% of each other. This resource distribution fairness allows SharePoint administrators to assign Network Load Balancing distributions for SharePoint user requests in a predictable manner.

- **Inter-VM communication**— Network communication among VMs that are connected to the same virtual switch (vSwitch) on a host is called inter-VM communication. In this case, the ESX networking subsystem can pass network traffic between the virtual machines directly through the vSwitch layer. This eliminates the overhead of networking traffic passing through the physical network card.

Inter-VM communication improves throughput by 1% in the 3 SharePoint VMs case when compared with the case where the Web server VM is deployed on its own physical machine. Thus further networking performance benefits are gained when you consolidate more physical machines of your SharePoint infrastructure into multiple VMs on a single host.

Best Practices for vSphere 4.1

This section describes some best-practice recommendations for improving virtualized SharePoint performance:

- **Enable NUMA for your NUMA processor.** NUMA-aware resource management in vSphere achieves fast memory accesses and balanced load-distribution among NUMA nodes. This provides optimal performance in the NUMA system.
- **Make sure the system has enough memory to avoid ESX host swapping; otherwise, performance in the virtual machines is significantly reduced.** The memory compression feature is designed to lessen the performance penalty of host-level swapping by replacing host-level swap-out with compressed pages on memory. Pay attention to the balloon driver's inflation, which may induce guest swapping. Detailed information is available at http://www.vmware.com/files/pdf/perf-vsphere-memory_management.pdf.
- **Use the VMXNET3 virtual network device.** VMXNET3 is the latest virtual network driver. It provides lower CPU cost and better latency. For details, see "Performance Evaluation of VMXNET3 Virtual Network Device" at http://www.vmware.com/pdf/vsp_4_vmxnet3_perf.pdf.
- **Enable large pages for SQL Server VM.** By enabling large pages for the SQL Server VM, costly translation lookaside buffer (TLB) misses are reduced. Also see "Performance and Scalability of Microsoft SQL Server on VMware vSphere 4" at http://www.vmware.com/files/pdf/perf_vsphere_sql_scalability.pdf.
- **Use multicast mode for Microsoft Network Load Balancing software in VMs.** Multicast mode is recommended because unicast mode forces the physical switches on the LAN to broadcast all Network Load Balancing traffic to every machine on the LAN. Also see "Implementing Microsoft Network Load Balancing in a Virtualized Environment" at http://www.vmware.com/files/pdf/implmenting_ms_network_load_balancing.pdf.

Conclusion

This paper shows that VMware vSphere™ 4.1 is an excellent platform on which to virtualize and consolidate a SharePoint implementation that supports a large number of users.

We determined the configurations that result in the best performance, based on a large number of users at 1% concurrency, as:

- 3 VMs (1 Web server, 1 application server, 1 SQL server) for 72,600 heavy users
- 4 VMs (add a 2nd Web server) for 135,600 heavy users
- 5 VMs (add a 3rd Web server) for 171,600 heavy users

Our scalability result shows that you can leverage the Infrastructure as a Service model for your SharePoint farm in the consolidation environment while ensuring average response times under 3 seconds to maintain a positive user experience. Using vSphere 4.1, you can simply deploy additional VMs for a SharePoint server role that needs more resources to meet increased demand. Adding more VMs in this way also increases the physical server utilization.

About the Author

Vincent Lin is a Performance Engineer at VMware. In this role, his primary focus is to evaluate and help improve the performance of VMware products in better supporting key Enterprise applications. Prior to VMware, Vincent was a Principal Performance Engineer at Oracle. Vincent received a Master of Science degree from the University of West Florida.

Acknowledgements

The author would like to thank the VMware development teams for their efforts in developing and optimizing vSphere. He would also like to thank Shandar Varadarajan for the Microsoft SharePoint production deployment insights. Finally, he would like to thank members of the VMware Performance team who have helped with the performance analysis and creation of this document: Julie Brodeur, Scott Drummonds, Fei Guo, Seongbeom Kim, Todd Muirhead, Aravind Pavuluri, Harold Rosenberg, Priya Sethuraman, and Lenin Singaravelu.

If you have comments about this documentation, submit your feedback to: docfeedback@vmware.com

VMware, Inc. 3401 Hillview Ave., Palo Alto, CA 94304 www.vmware.com

Copyright © 2010 VMware, Inc. All rights reserved. This product is protected by U.S. and international copyright and intellectual property laws. VMware products are covered by one or more patents listed at <http://www.vmware.com/go/patents>. VMware is a registered trademark or trademark of VMware, Inc. in the United States and/or other jurisdictions. All other marks and names mentioned herein may be trademarks of their respective companies.

Item: EN-000348-00
