White Paper

Intel® Xeon® Processor
VMware vSphere* 4.0
Web Server Consolidation

Achieving Twenty-Four-to-One Web Server Consolidation

Using VMware vSphere* 4.0 on Intel® Xeon® Processor-Based Servers

Quantifying the Value of Web Server Consolidation

Server performance and energy-efficiency have improved many-fold in recent years, yet a high percentage of organizations around the world continue to rely on Web servers running on hardware that is five years old or older. Consolidating these aging Web servers into virtual machines running on the latest Intel® Xeon® processor-based systems can deliver substantial business and IT benefits, not only through improved performance, but also by dramatically reducing data center space, power, cooling and maintenance requirements.

To quantify the benefits, Intel conducted tests to ascertain how many Web servers could be consolidated onto a single four-socket Intel® Xeon® processor-based server running VMware ESX* virtualization software. A consolidation ratio of 24:1 was easily achieved, along with an 87 percent reduction in energy consumption.

This paper describes the tests and the results in detail. It will be useful for anyone interested in delivering better performance to more Internet users, while reducing operating costs and freeing up data center resources for other needs.

Author: Aamir Yunus,
Senior Software Engineer, Intel Corporation
aamir.b.yunus@intel.com

The Performance Benchmark

Web workload is a useful benchmark for gauging the performance of Web servers. It provides a workload that stresses all aspects of the system under test, including CPU, disk I/O and network resources. (To help exercise the disk and network subsystems, Web workload uses long files, which are typical in online support requests.)

Essential components of Web workload include:

- A Web server
- A back-end application server/database server (BeSim)
- Client systems to generate the load. PHP or JavaServer Pages* (JSP) can be used to generate dynamic Web content

Web workload performance is measured as the maximum number of simultaneous user sessions a Web server can support while meeting the following quality of service requirements:

- 95 percent of page requests must come back within 3 seconds
- 99 percent of page requests must come back within 5 seconds

Test Configurations

Two test environments were set up. The first was used to measure Web workload performance on a dedicated physical server running a single instance of the Web server software. The second was used to measure performance on a virtualized physical server running multiple instances of the Web server software.

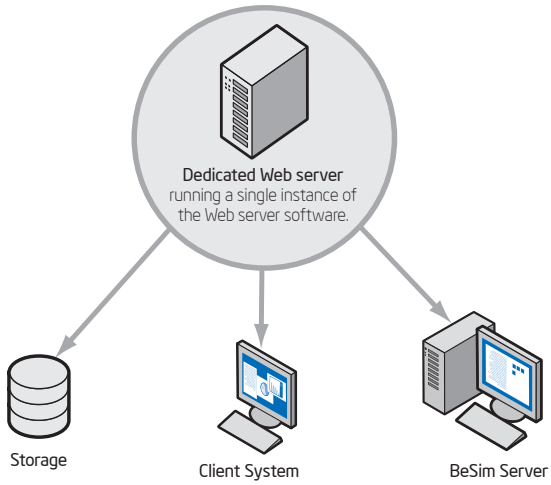

Figure 1. Test Configuration for the Dedicated Web Server. Performance was measured for a two-socket Intel® Xeon® processor 3.60 GHz-based server running a single instance of Microsoft IIS.*

Dedicated Web Server

The setup for the dedicated Web server test is shown in Figure 1. The Web server hardware and software were configured to provide a server environment similar to the solutions that are widely deployed in today's production environments.

- **The Web server hardware** consisted of a two-socket Intel® Xeon® processor 3.60 GHz-based server with 4 GB of memory. It was connected to 256 GB of high-speed storage and had 1 Gbps of network bandwidth to support traffic among the Web server, the BeSim server and the client system.
- **The Web server software** consisted of Microsoft Internet Information Services* (IIS) running on the Microsoft Server 2003* operating system. This software was used because it represents today's most widely deployed Web server software environment.

Detailed configurations for all test configurations are shown in Table 1, on page 4.

Consolidated Web Server

The setup for the consolidated Web server test is shown in Figure 2.

- **The Web server hardware** consisted of a four-socket system configured with the Intel Xeon processor X7560.[Δ] Memory was increased to 256 GB and storage capacity to 1.5 TB to handle the increased requirements of the consolidated environment. Since total client traffic was found to exceed 2.5 Gbps, a 10 Gigabit Ethernet network adapter was used. A separate 1 Gigabit Ethernet adapter was used to support communications between the Web server and the BeSim server.

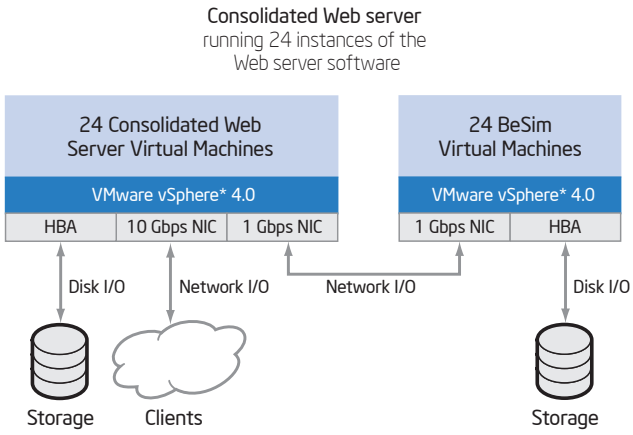

Figure 2. Test Configuration for the Consolidated Web Server. Performance was measured for a single Intel® Xeon® processor X7560 series-based server running VMware ESX* Server 4.0 and twenty-four instances of Microsoft IIS.*

- **The virtualization software** in the primary test was VMware ESX Server 4.0. Twenty-four virtual machines (VMs) were created on the physical Web server to run the Web server software. Each VM was allocated two virtual CPUs and 4 GB of memory.
- **The Web server software** consisted of twenty-four instances of Microsoft Internet Information Services (IIS) and the Microsoft Windows Server 2008 operating system running in the twenty-four VMs. Other software configurations within each VM were the same as for the dedicated server, equating to a simple consolidation of existing production Web servers onto the new physical server.

VMware ESX Server was also used to create twenty-four VMs on the back-end application and database server, one VM for each instance of BeSim. Resource utilization for BeSim virtual machines is so low that all twenty-four instances could have been run on the system under test, which would have eliminated the need for a separate, back-end server. This would have provided another level of consolidation and additional power savings. However, the run rules for Web workload require that BeSim be hosted on a separate system.

Test Results

Twenty-Four-to-One Consolidation

The consolidated Intel Xeon processor X7560-based server running VMware ESX Server 4.0 increased performance by a factor of twenty-four compared with the dedicated Intel Xeon processor 3.60 GHz-based server (Figure 3). The dedicated server was able to support 750 simultaneous Web workload users, while satisfying the stipulated quality of service requirements. The Intel Xeon processor X7560-based server was able to support 18,000 simultaneous users, while delivering the same quality of service.

Almost $5,800 Annual Cost Savings

Power consumption per 18,000 users was dramatically reduced for the consolidated Web server.

- **The dedicated Web server** consumed an average of 319 watts to support 750 Web workload users. It would take twenty four of these dedicated servers to support 18,000 users, and those servers would consume 24x319 watts = 7,656 watts.
- **The consolidated Web server** consumed 1,011 watts to support 18,000 users, reducing power consumption by 6,645 watts.

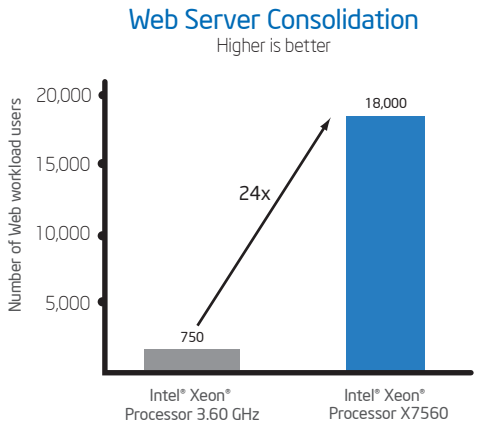

Figure 3. Web Workload Performance Comparison. The Intel® Xeon® processor X7560-based server running VMware ESX* Server 4.0 was able to support 24 times as many Web Workload users as the dedicated Intel Xeon processor 3.60 GHz-based server.

The potential cost savings are substantial. Assuming the Web servers operate 24/7, the twenty-four dedicated servers would consume a total of 67,067 kilowatt hours per year. At 10 cents per kilowatt hour, that would create an annual energy bill of $6,707. The same calculation for the consolidated Web server would yield an annual energy bill of just $886, delivering yearly savings of $5,821.

Of course, those savings don't include cooling costs which could be reduced by a comparable amount, depending on the cooling infrastructure and data center policies. At the very least, the consolidated Web server environment would reduce the load on the cooling system, which would free up cooling resources for other uses.

Conclusion

Web servers can be consolidated very efficiently on today's Intel Xeon processor-based servers using VMware ESX 4.0. Based on our tests with Web workload, we were able to consolidate twenty-four Web servers on a standard four-socket system and reduce power consumption by 87 percent. This level of consolidation would not only free up data center space, but could also save more than $5,821 annually in energy costs alone.

Additional savings could be realized through reduced cooling requirements and system maintenance costs. IT organizations could also take advantage of VMware management tools to balance workloads, simplify management and improve availability through automatic failover.

Table 1. Detailed System Configurations

	Dedicated Web Server Test	Consolidated Web Server Test
Web Server		
Platform	Intel® White Box (Coyote)	Intel® Pre-Production Four-socket Platform
Processor	2x Intel® Xeon® processor 3.60 GHz (2 MB Cache, 800 MHz FSB, Intel® Hyper-Threading Technology enabled)	4x Intel® Xeon® processor X7560 (24 MB Cache, 2.27 GHz, 6.40 GT/s Intel® QPI)
Memory	4 GB PC-3200 DDR	256 GB DDR3
Storage	8x 32 GB Intel® SSDs, RAID 0	1.5 TB Intel® SSDs, RAID 0
Network Adapter	1 Gbps Ethernet	Intel® 10 Gigabit XF SR Server Adapter
Operating System	Microsoft Windows Server 2003* (32-bit) running natively	Microsoft Windows Server 2008* R2 (64-bit) (1 instance per VM)
Web Server Software	Microsoft Internet Information Services* 6 running natively	Microsoft Internet Information Services* 7 (1 instance per VM)
Virtualization Software	N/A	VMware ESX* 4.0 UI
Virtual Machine Configurations	N/A	24 VMs, each with 2 virtual CPUs and 4 GB memory
Application and Database Server (BeSim)		
Processor	2 x Dual-Core Intel® Pentium® 4 processor 3.73 GHz	2x Intel® Xeon® processor X5460 3.16 GHz
Network Adapter	1 Gbps Ethernet	1 Gbps Ethernet
Virtualization Software	N/A	VMware ESX 4.0 UI
Virtual Machine Configurations	N/A	24 VMs, each with 1 virtual CPU and 512 MB memory
Client		
Processor	Dual-Core Intel® Pentium® 4 processor-based system	

For more information about Intel Xeon processors, visit www.intel.com/p/en_US/products/server/processor

For more information about VMware products and services, visit www.vmware.com

[A]Intel processor numbers are not a measure of performance. Processor numbers differentiate features within each processor family, not across different processor families. See www.intel.com/products/processor_number for details.

Printed in USA 0310/RR/OCG/XX/PDF Please Recycle 323585-001US