

Sun Microsystems Sun Fire X4450 server with Intel Xeon 7400 Series processors: Performance and Scalability for Virtual Desktop Infrastructure

Independent Test Report Prepared for Sun Microsystems

Executive Summary

Sun Microsystems, Inc. (hereinafter “Sun”) commissioned VeriTest, the testing service of Lionbridge Technologies, Inc., to conduct a study aimed at demonstrating a theoretical maximum threshold for the number of “Desktop” type virtual machines (VMs) commonly used in a Virtual Desktop Infrastructure (VDI) that can be supported on a single Sun Fire X4450 server running the Intel Xeon 7400 Series processor processors running on VMware Infrastructure 3 (ESX 3.5.0 Build 110181 with VirtualCenter Server 2.5). Testing was performed at the Lionbridge test facility in Oakdale, Minnesota.

Virtual Desktop Infrastructure, commonly known as VDI, is a computing model that is designed around the architecture of centralizing the desktop computing resources into the data center and replacing each desktop with a more secure and cost-efficient device. The basic concept is to shift physical desktop workstations and their corresponding workloads at each user desk space, to independent and dedicated virtual machines that are residing on secure, robust servers in a centralized data center. Each specific end-users desktop operating environment is then remotely accessed from a more secure and more easily manageable device, such as a thin client. This model is intended to ensure data security, increase the administrators ability to better manage and service each individual end-user, all while providing secure access to the end-users identical environment they are already familiar with, such as Windows XP or Vista.

Testing was performed at different levels of total load with a goal of achieving as many VMs on a single Sun Fire X4450 server as possible without exceeding an average CPU utilization of over 95% for the duration of the test simulation. In our tests, this resulted in a simulation of 120 VMs. In an effort to appropriately simulate real-world usage, testing was focused on workloads that are most commonly seen in an average VDI desktop user scenario.

PCWorld’s WorldBench 5 tool was used to generate typical Windows user load on the VMs. WorldBench 5 is the fifth generation of PC World’s industry-standard benchmarking application designed to measure the performance of a wide range of personal computers. WorldBench 5 uses real applications running real-world tasks to assess a PC’s overall performance.

Key Findings

- ❑ Testing showed that the Sun Fire X4450 server with the Intel Xeon 7400 Series processor platform was capable of supporting as many as 120 Windows XP Desktop Virtual Machines.
- ❑ When using the Intel Xeon 7400 Series processor architecture, the X4450 CPU load is 12% lower with 100 VMs than with the Harpertown architecture.
- ❑ Due to the nature of the WorldBench benchmark being geared to run its specified workload at a higher rate than a real-world user, it is believed that more than 120 VMs could be supported.

Testing Methodology

VeriTest developed a test methodology in which Windows XP x86 (SP3) Virtual Machines were configured with a workload that simulated an average desktop user experience. To simulate this workload, VeriTest used the PC World WorldBench v5 application. The WorldBench benchmark is able to simulate various desktop workloads such as office applications, multimedia and internet activity. Since testing was focused on typical enterprise desktop environments, the Office XP workload simulation was selected as most appropriate. This workload simulates the usage of the core Office suite applications Word, Excel, PowerPoint, Outlook and Access. These applications are each launched and then perform common desktop tasks such as opening, closing, creating, and modifying .doc, .xls and .ppt files as well as composing e-mails. During the workload, all Office applications are opened simultaneously and a variety of tasks are executed. These include such things as:

- Reading, composing, deleting, and sending mail, creating and viewing calendar appointments, entering, modifying, and deleting contacts, creating and editing tasks in Microsoft Outlook.
- Typing and formatting text, scrolling from page to page, running spell check, conducting mail merges, print previewing documents in Microsoft Word.
- Performing calculations, creating charts, sorting data, formatting the screen, and previewing spreadsheets in Microsoft Excel
- Creating slides, formatting slides, modifying the master template, adding, editing, and deleting test, and viewing slideshows in Microsoft PowerPoint
- Entering data, executing queries, and generating reports in Microsoft Access

It should be noted that WorldBench is designed to run the various simulations simultaneously at the fastest rate the target system can execute the program. This rate of execution is indisputably a faster rate than a “real” user would execute the same various applications and tasks.

Each WorldBench instance was configured to execute the Office-XP profile for a total of 15 cycles. The number of cycles was used to ensure that 3 hours of test execution time could be captured to best determine the maximum number of VMs that could be supported on the Sun Fire X4450 server with the simulated workload.

Each World Bench cycle contains numerous system and application setup stages prior to, and after, the execution of the benchmark suite. These tasks are performed on each VM running the WorldBench benchmark suite for Office XP. Each of these stages are executed by the WorldBench application as part of the test cycle. These WorldBench stages include the following:

1. Backup the registry.
2. System Modification by WorldBench
3. Reboot.
4. Setup system files.
5. Setup alluser files.
6. Setup user files.
7. Reboot.
8. Registry setup for WB.
9. Setup application files -> msoffice
10. Setup registry for msoffice.
11. Clear system buffers.
12. Clear background tasks.
13. Wait for system to become idle.
14. Execute benchmark.
15. Remove application files.
16. Setup alluser files.
17. Setup user files.
18. Setup the registry for WB.
19. Setup application files -> msoffice
20. Setup registry for msoffice.

21. Clear system buffers.
22. Clear background tasks
23. Wait for system to become idle.
24. Execute benchmark.
Steps 15-24 repeated for each test cycle.
25. Remove application files.
26. Restore alluser files.
27. Restore user files.
28. Restore system files.
29. Clean up the registry.
20. Reboot.
21. Test complete.

WorldBench 5 combines the results of scripted application tests and then compares them to the scores of a reference system – published as a system with a 2.2-GHz Athlon 64 FX-51 CPU with 1MB of Level 2 cache and 1GB of RAM, as well as an NVidia GeForce FX 5950 Ultra graphics card with 256MB of RAM. For the purposes of this study, WorldBench was used as a load generation tool to simulate workload that is most common in an enterprise desktop environment. Due to the rapid nature of the application tasking WorldBench performs, the load placed on the ESX host by the VM simulation was expected to be greater than what a typical user would achieve using the same applications. Such a situation allows for some conclusions to be made with regard to how many VMs can be supported.

WorldBench is a tool designed for the physical environment and cannot take into account the nuances of shared resources in a virtual environment. As a result, the overall WorldBench score was not relevant to the testing. The focus of this study was on the performance of the server under test. However, during all test runs, the individual VMs were sampled to ensure their application performance fell within a reasonable range.

In addition, since WorldBench runs locally on the target system, the effects associated with accessing the Windows XP VMs over a network connection (such as from a Sun Ray virtual display client) was not considered as part of the performance associated with the Sun Fire X4450 platform under test.

Test Bed

Figure 1 below shows the test environment configured in support of this testing.
SUN VDI Benchmarking – Network Topology

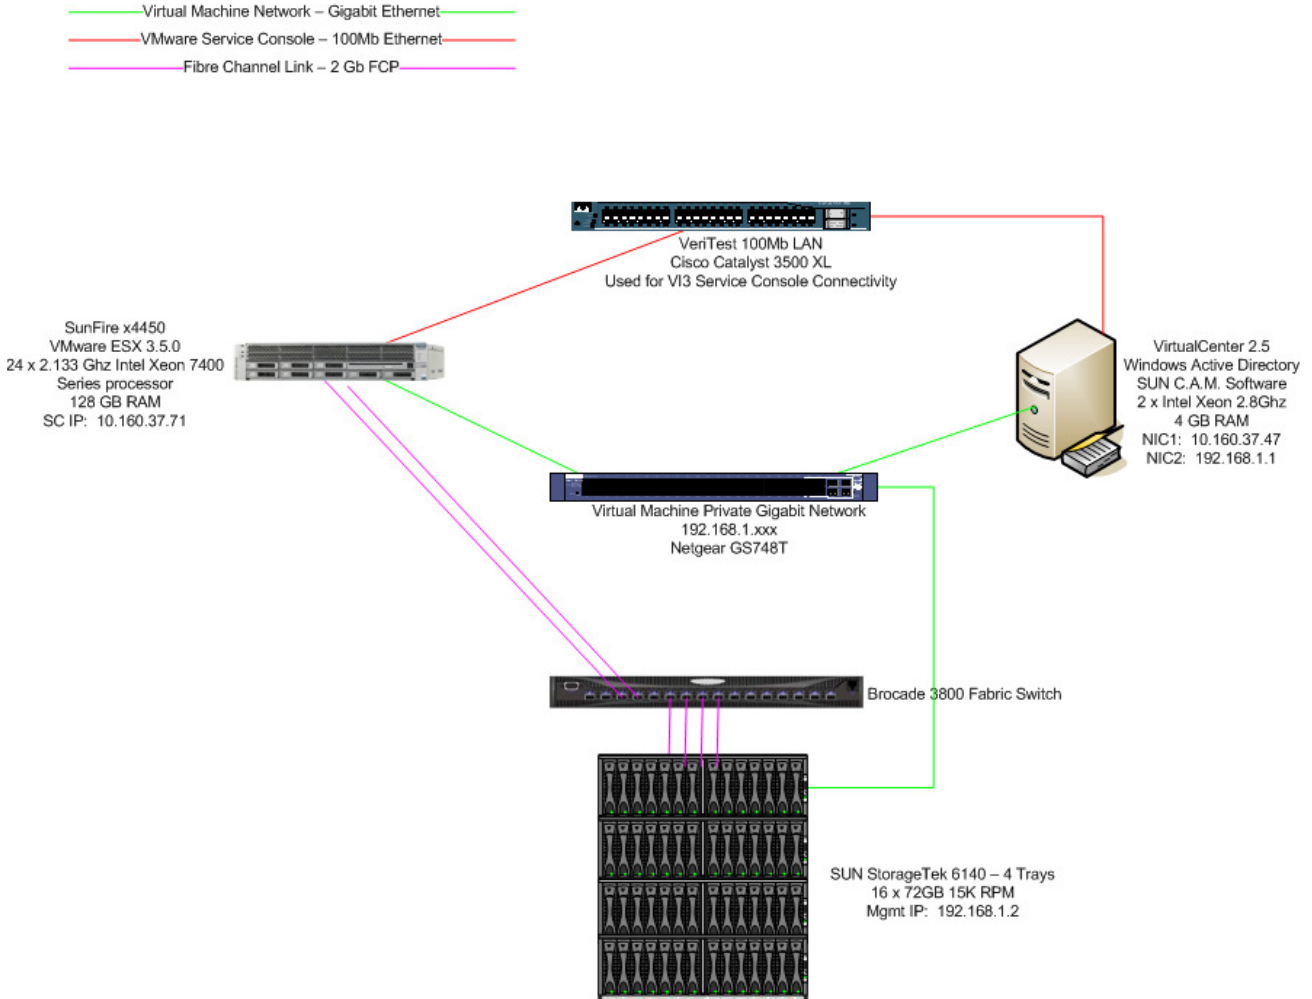


Figure 1: Test Bed Diagram

The server under test was a Sun Fire X4450 server. This system contained four (4) Six-core Intel Xeon L7455 processors running at 2.13 GHz, 128 GB RAM, 4 GbE Ethernet ports, and 4 x 146 GB 10K RPM SAS drives. VMware ESX Server 3.5.0 Update 2 was installed onto one of the internal disks and a default Datastore configured. This local Datastore was configured to hold the template Windows XP x86 w/SP3 Virtual Machine image.

To provide storage for the Virtual Machines, a Sun StorageTek 6140 array was configured with 8 volumes using the Microsoft_NTFS profile. This configuration allocated 4 disks in a 3+1 configuration using RAID-5 for each of the 8 volumes. A total of 20 Virtual Machines could be allocated to each volume if needed. The STK6140 was in a dual controller configuration. Each controller had two Fibre Channel ports connected to the Fibre Channel switch. In addition, the Sun Fire X4450 server had a single dual-port QLE2462 HBA installed. Each port of the HBA was connected to the Fibre Channel switch. The Fibre Channel switch was configured with two zones that dedicated port1 from the server's FC HBA to port1 on the STK6140 Controller

A and port2 of the server's FC HBA to Controller B port1 on the STK6140. This resulted in a configuration in which the 8 storage volumes were evenly balanced across the two FC HBA ports on the Sun Fire X4450 as well as evenly balanced across the two controllers of the STK6140.

To represent the virtualized desktops, VeriTest first installed a single source Virtual Machine (VM) which contained the Windows XP x86 with SP3 Operating System along with the WorldBench 5 application. Each VM was configured with a single vCPU, 512 MB of RAM, a single 8 GB HDD drive image, and a single NIC using the Flexible adapter. This image was then replicated using the VMware VI3 cloning capabilities. Each cloned Virtual Machine image was then made unique by modifying the system hostname and then joining the VM to the Windows Active Directory domain.

A separate physical server was used to run Windows Active Directory services. This was equipped with dual 2.8 Ghz Intel Xeon CPUs, 4 GB of RAM, dual Gigabit Ethernet NICs, and 2 x 36 GB 15K RPM U320 SCSI HDDs in a RAID-1 configuration running Windows Server 2003 R2 Enterprise Edition with SP2 x86. This system was used to provide Active Directory services to the Windows XP Virtual Machines. This system also had VirtualCenter Server 2.5 Update 2 installed.

Finally, to measure the CPU and RAM performance of the Sun Fire X4450 server the VMware VI3 utility esxtop was utilized. The command was run in batch mode to capture all performance metrics every 10 seconds. In addition to the esxtop command running, the Sun Fire X4450 server was monitored with VirtualCenter to ensure that no warnings or alerts were reported by ESX with regard to other monitored resources such as RAM usage.

Test Execution

To determine the maximum number of Virtual Machines that could be successfully run on the Sun Fire X4450 server without exceeding a maximum CPU load level that would result in degraded performance on the VMs it was necessary to gradually increase the number of VMs running the simulated workload and monitor the CPU performance with esxtop. It was determined that the maximum CPU load to achieve would be an average of 95%. This allows for some overhead room to exist in the event a small number of the VMs experience a short spike in usage.

Using the results from the Sun Fire X4450 Intel Xeon 7300 Series processor as a guide, VeriTest started with 100 VMs running the workload described above as the starting point for determining the maximum number of VMs that could be supported on the Sun Fire X4450 with the Intel Xeon 7400 Series processor architecture. All tests were executed twice to ensure repeatability of the results. The maximum VM load was allowed to run for a minimum of three hours for each test run to ensure the average CPU load was sustained at the 95% range for this extended period of time.

In addition to CPU load, the Sun Fire X4450 server was monitored to ensure there weren't any other performance related alerts.

Test Results

As described in the Test Methodology section above and illustrated in Figure 2 below, the first test performed used 100 VMs running the WorldBench Office-XP workload simulation. The CPU utilization was collected from VMware's esxtop and was sampled approximately every 10 seconds. The Y axis on the chart shows the execution time in "hours:minutes:seconds" beginning at zero as the start of testing and continuing to the final sampling time roughly three hours later.

Testing showed that with 100 VMs running the sustained average CPU load was 84.10% over the three hour sample period. Thus, additional VMs could be added in order to increase the CPU load towards the 95% target threshold.

In order to produce a test cycle that would run over a three hour period, WorldBench was configured to run the OfficeXP workload fifteen times. At the beginning of each test cycle, WorldBench performs a variety of configuration changes and updates and then reboots each VM. These changes are subsequently reset to the original settings at the end of the test cycle and each VM is again rebooted. First-hand observations by VeriTest engineers indicate the "valleys" in Figure below appear to be directly correlated to the WorldBench application preparation stages, steps 1 thru 13, as described in the Test Methodology section above. The peaks in the below figure are the periods in which WorldBench was in the process of executing the Office XP workload simulation.

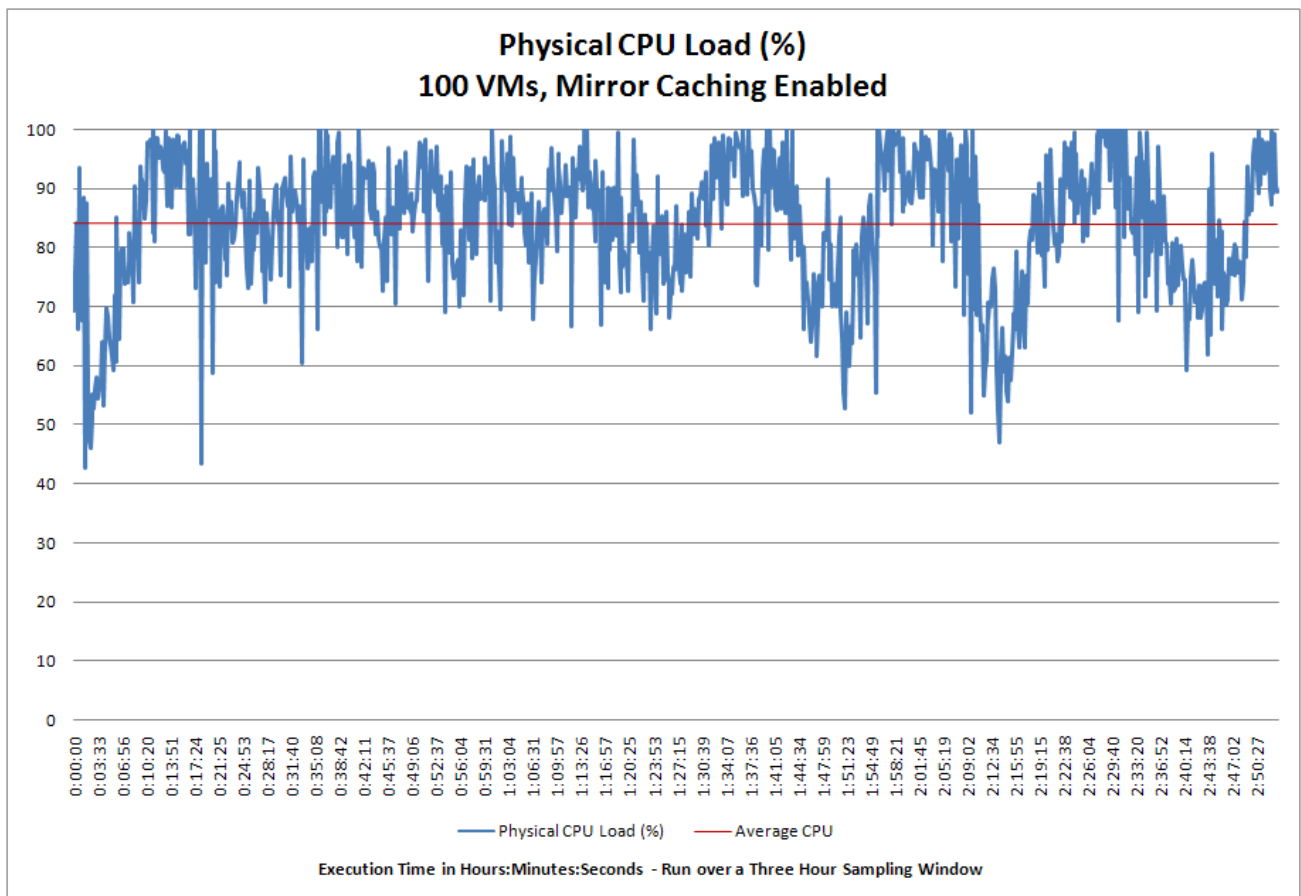


Figure 2: CPU Load – 100 Virtual Machines – Mirror Cache Enabled

For the next simulation, VeriTest increased the number of VMs running the WorldBench OfficeXP simulation to 120. The CPU utilization was collected from VMware’s esxtop and was sampled approximately every 10 seconds. The Y axis on the chart shows the execution time in “hours:minutes:seconds” beginning at zero as the start of testing and continuing to the final sampling time roughly three hours later.

Testing showed that with 120 VMs, the average sustained CPU load over the three hour sample period was 88.04%.

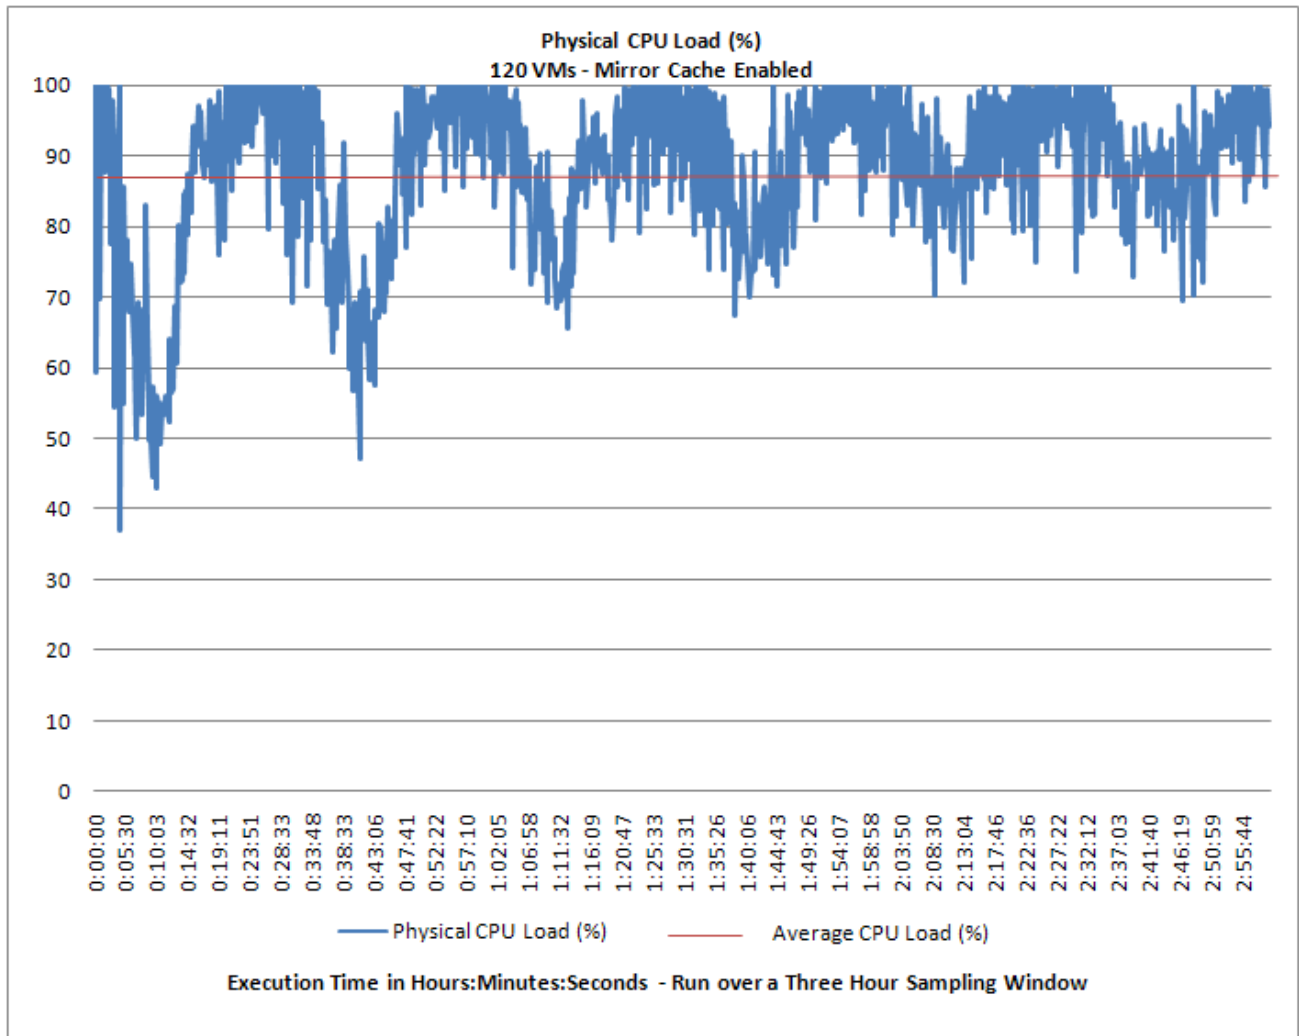


Figure 3: CPU Load – 120 Virtual Machines – Mirror Cache Enabled

Upon completion of this test cycle, it was determined that a bottleneck existed elsewhere in the environment. Upon investigation, it appeared that the particular storage array configuration in place for this test was at the root of the performance. Through conversations with SUN technical representatives, it was suggested that in this benchmark configuration performance may benefit from the disabling of the storage controller Mirror Caching. Mirror Caching provides an additional layer of data protection in the environment that is generally required for production environments, but is not necessarily required in benchmark situations where the storage array is not the product being benchmarked.

After Mirror Cache was disabled on the STK6140 array, VeriTest again ran the WorldBench OfficeXP simulation with 120 VMs. The CPU utilization was collected from VMware's esxtop and was sampled approximately every 10 seconds. The Y axis on the chart shows the execution time in "hours:minutes:seconds" beginning at zero as the start of testing and continuing to the final sampling time roughly three hours later.

Testing showed that with 120 VMs running, and Mirror Cached disabled on the STK6140, the sustained average CPU load was 94.73% over the three hour sample period. This result met the goal of finding the number of VMs that could be supported on the X4450 with 95% being the target threshold.

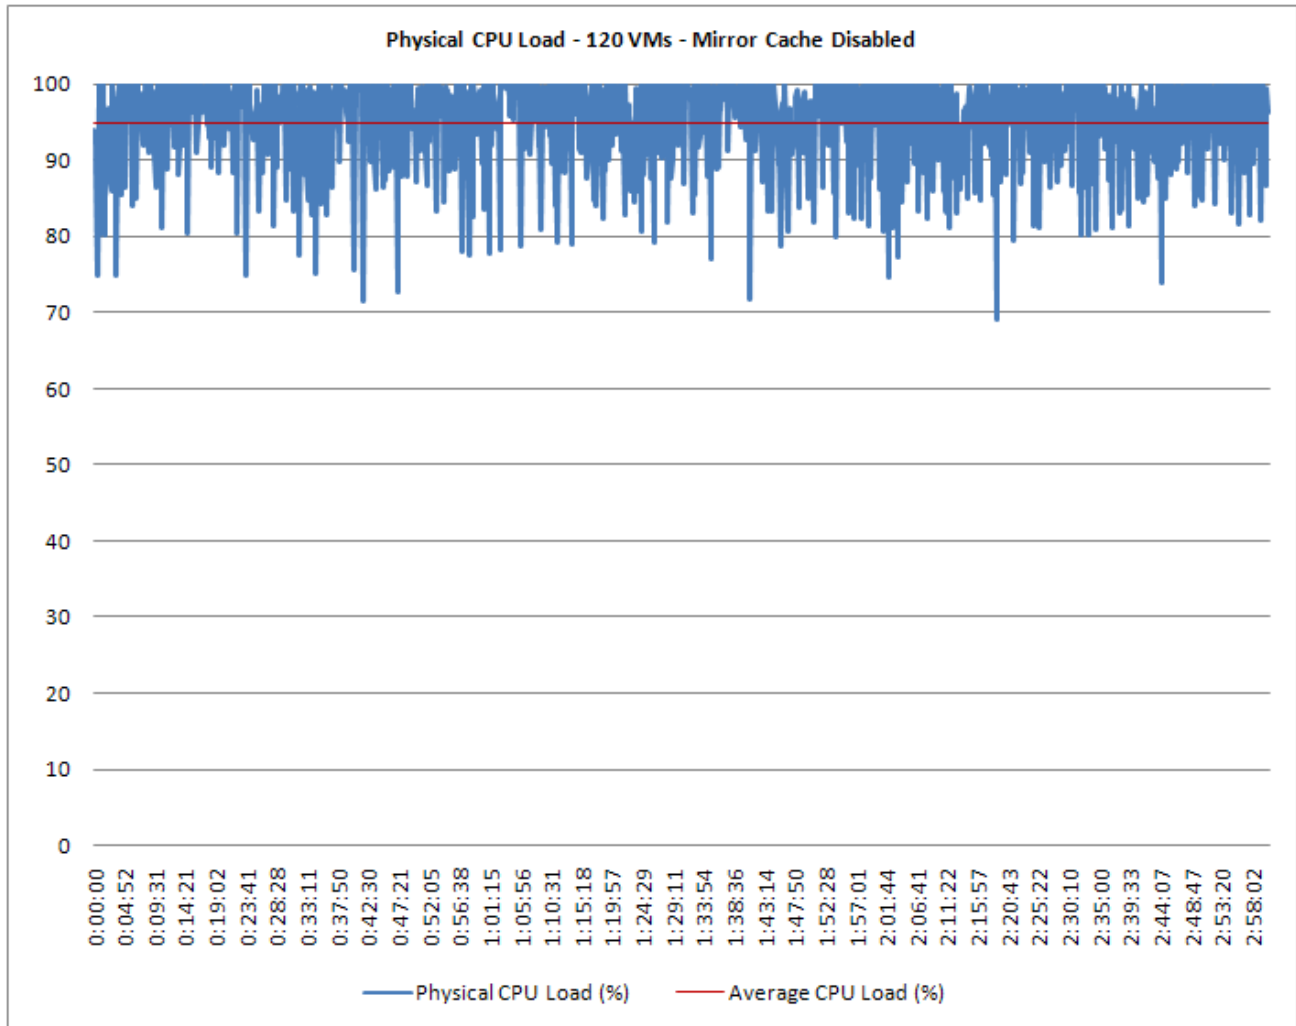


Figure 4: CPU Load – 120 Virtual Machines – Mirror Cache Disabled

ESX was configured to issue an alert whenever the ESX host would experience CPU load that reach 95% or above. It should be noted that during all test simulations, no ESX alerts were recorded pertaining to resource contentions with the Sun Fire X4450 server. ESX was configured to issue an alert whenever the ESX host would experience CPU load that reach 95% or above. The only alerts that appeared were with regard to CPU utilization, which was expected when the load began to reach the 95% plus mark during the second running of 120 VMs.

Conclusion

Testing showed that the Sun Fire X4450 server, with the Intel Xeon 7400 Series processor platform, was capable of supporting as many as 120 Virtual Machines running the WorldBench Office XP workload simulation at an average CPU utilization of 95%. The only system resource that reported alerts indicating a resource constraint was the CPU utilization reaching our intentionally pre-set mark of 95%. This occurred during the second 120 VMs test, indicating this was the maximum number of VMs that could be supported under the test environment conditions.

Further, since the WorldBench simulation generates a more consistent and intensive workload than might typically be found in the enterprise, it is reasonable to expect customers could achieve an even greater number of desktop virtual machines on a single Sun Fire X4450 server in their environment.

Appendix A. Hardware Disclosures

Server Under Test

Model	Sun Fire X4450 server
BIOS	S93_3B53
CPU	24 x Intel Xeon L7455 2.13 Ghz
RAM	128 GB, DDR2 667 Mhz
HDD	4 x 144 GB 15K RPM SAS
NIC	4 x 10/100/1000 Onboard
Fibre Channel HBA	QLE2462
BIOS	1.04
ESX Version	3.5.0 Update 2 Build 110181

Storage Array

Model:	Sun STK6140
Firmware:	6.60.11.10
Controllers:	2
Trays:	4
Disks:	64 x 72GB 15K RPM FC-SCSI
Volumes:	8
RAID Type:	5
RAID Profile:	Microsoft NTFS
Number of Disks	4 (3+1)

Fabric Switch

Model	Brocade 3800 Silkworm
Firmware	3.2.0a

Network Switch

Model	Netgear GS748T
Firmware	V2.0.2_02

Windows Domain Controller

CPU	2 x 3.0Ghz Intel Xeon
RAM	4 GB
HDD	2 x 72GB 15K RPM U320 SCSI (RAID-1)
NIC	2 x 10/100/1000 Broadcom (onboard)
NIC1	Remote Admin Access
NIC2	VM Network Access
OS	Windows Server 2003 EE SP2 x86

Virtual Machine Configuration

OS	Windows XP Pro SP3 x86
vCPU	1
RAM	512 MB
HDD	8 GB (Bus Logic)
NIC	1 (Flexible Adapter)
Benchmark Software	PC World WorldBench v5.0

VeriTest (www.veritest.com), the testing service of Lionbridge Technologies, Inc., provides outsourced testing solutions that maximize revenue and reduce costs for our clients. For companies who use high-tech products as well as those who produce them, smoothly functioning technology is essential to business success. VeriTest helps our clients identify and correct technology problems in their products and in their line of business applications by providing the widest range of testing services available.

VeriTest created the suite of industry-standard benchmark software that includes WebBench, NetBench, Winstone, and WinBench. We've distributed over 20 million copies of these tools, which are in use at every one of the 2001 Fortune 100 companies. Our Internet BenchMark service provides the definitive ratings for Internet Service Providers in the US, Canada, and the UK.

Under our former names of ZD Labs and eTesting Labs, and as part of VeriTest since July of 2002, we have delivered rigorous, objective, independent testing and analysis for over a decade. With the most knowledgeable staff in the business, testing facilities around the world, and almost 1,600 dedicated network PCs, VeriTest offers our clients the expertise and equipment necessary to meet all their testing needs.

For more information email us at info@veritest.com or call us at 919-380-2800.

Disclaimer of Warranties; Limitation of Liability:

VERITEST HAS MADE REASONABLE EFFORTS TO ENSURE THE ACCURACY AND VALIDITY OF ITS TESTING, HOWEVER, VERITEST SPECIFICALLY DISCLAIMS ANY WARRANTY, EXPRESSED OR IMPLIED, RELATING TO THE TEST RESULTS AND ANALYSIS, THEIR ACCURACY, COMPLETENESS OR QUALITY, INCLUDING ANY IMPLIED WARRANTY OF FITNESS FOR ANY PARTICULAR PURPOSE. ALL PERSONS OR ENTITIES RELYING ON THE RESULTS OF ANY TESTING DO SO AT THEIR OWN RISK, AND AGREE THAT VERITEST, ITS EMPLOYEES AND ITS SUBCONTRACTORS SHALL HAVE NO LIABILITY WHATSOEVER FROM ANY CLAIM OF LOSS OR DAMAGE ON ACCOUNT OF ANY ALLEGED ERROR OR DEFECT IN ANY TESTING PROCEDURE OR RESULT.

IN NO EVENT SHALL VERITEST BE LIABLE FOR INDIRECT, SPECIAL, INCIDENTAL, OR CONSEQUENTIAL DAMAGES IN CONNECTION WITH ITS TESTING, EVEN IF ADVISED OF THE POSSIBILITY OF SUCH DAMAGES. IN NO EVENT SHALL VERITEST'S LIABILITY, INCLUDING FOR DIRECT DAMAGES, EXCEED THE AMOUNTS PAID IN CONNECTION WITH VERITEST'S TESTING. CUSTOMER'S SOLE AND EXCLUSIVE