Scalability of XenDesktop 4 on Microsoft Windows Server 2008 R2 Hyper-V

Completed at HP's Solution Center Labs

August 2010

Contents

Introduction

Citrix partners and customers alike are looking to build scalable XenDesktop environments on Microsoft Windows Server 2008 R2 Hyper-V and have requested guidance around what to expect when scaling up those environments. The Citrix team, with help from HP and Microsoft, collaborated on creating a large-scale virtual desktop environment at HP's Houston Solution Center. This document provides the information gathered during that seven-week project.

The virtualization environment was entirely Microsoft Windows Server 2008 R2 Hyper-V. The desktop hosting infrastructure consisted of 3500 desktop virtual machines each receiving a streamed Windows 7 image from Citrix Provisioning Server. The server virtualization infrastructure consisted of 15 virtual machines that hosted the support services for delivering and testing the desktops.

Login VSI 2.1 from Login Consultants was used to simulate the workload inside the virtual desktops. The Login VSI medium level workload consists of Microsoft Office 2007, Microsoft Internet Explorer, and Adobe tasks executed in a repeating 12-minute loop. The Login VSI script recorded response times so the user experience could be quantified for comparison.

Executive Summary

Results

A Windows 7 XenDesktop farm was built with 80 HP BL460c hosts servers running Hyper-V over the course of the seven-week project. On average, each of the 80 Hyper-V hosts were running 44 virtual desktops concurrently which provided a total of around 3520 active desktops in the pool. Upon completion of this project, the desktop farm supported over 3500 desktops with all desktops reporting acceptable response times under 1800 milliseconds for the Login VSI medium workload.

Key Findings

The analysis shows that one of the primary factors for the project's success was that the processors on the blades were not fully utilized. This behavior differs from typical session-based tests where the processors are normally the bottleneck. The single-factor that limited the capacity of the farm was the amount of RAM available in the BL460c servers. The single-server scalability tests showed the blade servers capable of 75-90 desktops when the available RAM was increased to 96GB. However, since the hardware had been ordered with only 48GB of RAM, the number of virtual machines was capped at 44 per host since each desktop required 1GB of RAM and approximately 3-4GB was used by the hypervisor.

The HP P4000 storage unit provided all the needed performance for the environment and linear progression analysis shows that the two clusters could have probably handled almost twice as many desktops. The addition of the HP Multipath I/O drivers significantly improved performance for the host LUNs.

Although this paper is concerned primarily with the actual test results, conclusions can be drawn from the observed performance metrics during the test. After careful analysis, the solution architects believe that the hardware and storage would have supported approximately 6000 Windows 7 desktops provided each Hyper-V host had 96GB of RAM and the same workload was used. Furthermore, the team believes that the gating factor would most likely be the storage platform since according to the single-server scalability tests that RAM and processor loads would have been well within the capabilities of the BL460c at that level.

Testing Methodology

The scalability testing completed during this project focused on two areas: single-server scalability and large-farm scalability. The first objective was to be able to identify how many desktops could run on a single BL460c blade server. The second objective was to determine how many desktops could run successfully with the given infrastructure, particularly the storage.

Desktop Stages

With those goals in mind, performance data was gathered during three primary desktop stages: Bootup, Logon, and Workload. Each stage and their characteristics are discussed below.

Bootup: This stage represents the time between when a desktop is powered on and when the desktop is registered ("Idle") and ready to accept a logon. During the testing the XenDesktop delivery controller controls the "spin-up" process and starts the desktops at a consistent, but throttled, rate. For the single-server tests, the spin-up rate was 20 desktops concurrently, with new desktops booted as existing ones registered. For large-farm tests, the spin-up rate was 10% of the total desktop pool, or approximately 350 desktops concurrently across all the hosts, with the same behavior of a successful registration initiating another desktop start. The desktop selection algorithm is random so the startups occur evenly across the farm.

The desktop bootup sequence is the most resource-intensive stage. As the desktop boots, the storage infrastructure takes the largest hit as it must create the page file and initialize the environment. In an environment where the operating system is streamed (or provisioned), such as tested here, the network infrastructure is also stressed significantly as the number of simultaneously booted desktops grows and the boot files are read across the network. Typically this stage has the highest storage I/O requirement and a moderate to high CPU requirement.

Logon/Logoff: The logon stage is the time between when a user requests a desktop and the workload is initiated. The logoff stage is the time between when the user completes the workload and the desktop returns to the "Idle" state. The logoff stage is slightly less resource intensive than the logon stage, so for the purposes of tests during this project, the results focus on the logon stage.

During the logon stage, the operating system must load the user's roaming profile and create a local copy of the profile. During logoff, the system must update the user's roaming profile. Typically these stages last 30-90 seconds each and have the second highest storage I/O requirement and a moderately high CPU requirement.

Workload: This stage represents the user's daily activities and is the period between when logon completes and logoff starts. Generally, this state has the lowest storage I/O requirement, but usually higher CPU utilization. For this project, a third-party tool from Login Consultants simulated the user workload.

Scalability Tests

During the single-server scalability tests, the goal was to identify the maximum number of virtual desktops a single server supports by evaluating the server during the three primary stages. To evaluate the different stages, the host servers had an agent from the test harness installed that gathered key performance counters at 15-second intervals. To measure the user experience, the simulated workload script recorded response times within the session. For the single-server tests, the desktops launched at a fixed rate of three per minute.

During the large-farm scalability tests, the goal was to identify the performance of the supporting infrastructure as the desktop load increased. As such, performance counters on each of the infrastructure components were monitored during five different load tests. For this environment, each C7000 chassis had approximately 720 desktops residing on the 16 BL460c hosts. Since running a fully-loaded chassis is not recommended, the defined testing intervals became 700, 1400, 2100, 2800, and 3505 desktops. These levels allowed the infrastructure load to increase linearly each time so the infrastructure could be assessed. The launch rate for desktops during these tests was configured so that all desktops would be logged in within a 30-minute window.

Test Tools

This section discusses the two primary test tools used in the environment: Login VSI and STAT. Login VSI was the workload simulation tool from Login Consultants (www.loginconsultants.com) and is publicly available. The test harness, referred to as STAT is an internal Citrix tool used to gather performance metrics and coordinate the ICA session launches into the XenDesktop farm.

Login VSI

The key to providing scalability guidance is to have a "random" workload that produces repeatable results when executed in different configurations. In the case of this project, the Login VSI 2.1 virtual session indexer was selected. The Login VSI workload is gaining popularity in the community and as such will provide scores which can be compared across hypervisor and hardware platforms.

Comparing results from earlier versions of Login VSI cannot be done because the test itself was changed between version 1.x and 2.x. The Login VSI documentation explains the differences but in general, the medium workload was modified, the loop length shortened, and the formula for creating the index updated.

The VSI Express edition comes with the medium workload which does the following:

- This workload emulated a medium knowledge worker using Office, IE and PDF
- Once a session has been started the medium workload will repeat every 12 minutes
- During each loop the response time is measured every 2 minutes
- The medium workload opens up to 5 apps simultaneously
- The type rate is 160 ms for each character
- The medium workload in VSI 2.0 is approximately 35% more resource intensive than VSI 1.0
- Approximately 2 minutes of idle time is included to simulate real-world users

Each loop will open and use:

- Microsoft Outlook 2007, browse 10 messages
- Microsoft Internet Explorer, one instance is left open (BBC.co.uk), one instance is browsed to Wired.com, Lonelyplanet.com and gettheglass.com (heavy flash app)
- Microsoft Word 2007, one instance to measure response time, one instance to review and edit document
- Bullzip PDF Printer & Acrobat Reader, the word document is printed and reviewed to PDF
- Microsoft Excel 2007, a very large randomized sheet is opened
- Microsoft PowerPoint 2007, a presentation is reviewed and edited
- 7-zip: using the command line version the output of the session is zipped

More information about the workload timing can be found in the Appendix.

Test Automation Harness

The testing leveraged a test harness called STAT to manage test runs and gather relative data. The tool uses SQL database queues for launch management and provides agents to gather performance counters on all Windows-based servers. Installing the STAT software required that both the .NET Framework 3.5 SP1 and SQL 2008 native driver be present on all servers. In this environment all servers had the STAT software installed.

The STAT tool provides a way to control the launch rate of desktop sessions into the XenDesktop farm as well as to correlate events, such as login, with performance counters recorded on the host servers. The tool was not used to monitor any response times inside the virtual desktops, but it was configured to monitor key performance counters on the following infrastructure components:

- Hyper-V Hosts
- Desktop Delivery Controllers
- Provisioning Services Servers
- Active Directory Domain Controllers
- Roaming Profile Server
- SQL Servers
- System Center Virtual Machine Manager Servers

The STAT tool also automatically recorded two key session-level statistics used in this environment:

XML Resolutions: The time elapsed from request to receipt of XML tickets from the Desktop Delivery Controller (DDC). This time should normally be under 2500ms and is an early indicator of IMA service or DDC issues.

Login Time: The time elapsed from launch of the ICA session to the time the STAT agent launched the Login VSI workload. This time should normally be under 30-seconds and higher values can indicate issues with storage, profiles, or desktop licensing.

Test Environment

The test environment was built in the HP Solution Center labs using the HP blade technologies. Storage was provided by an HP StorageWorks P4000 iSCSI SAN(LeftHand) using P4500 storage nodes. The SAN storage was used exclusively to store the virtual machine configuration files and Provisioning Services write-cache drives. Technical specifications for the hardware can be found in the Appendix. Both the physical and logical architectures are discussed in this section.

Physical Architecture

The physical architecture of the lab is detailed in Figure 1 below.

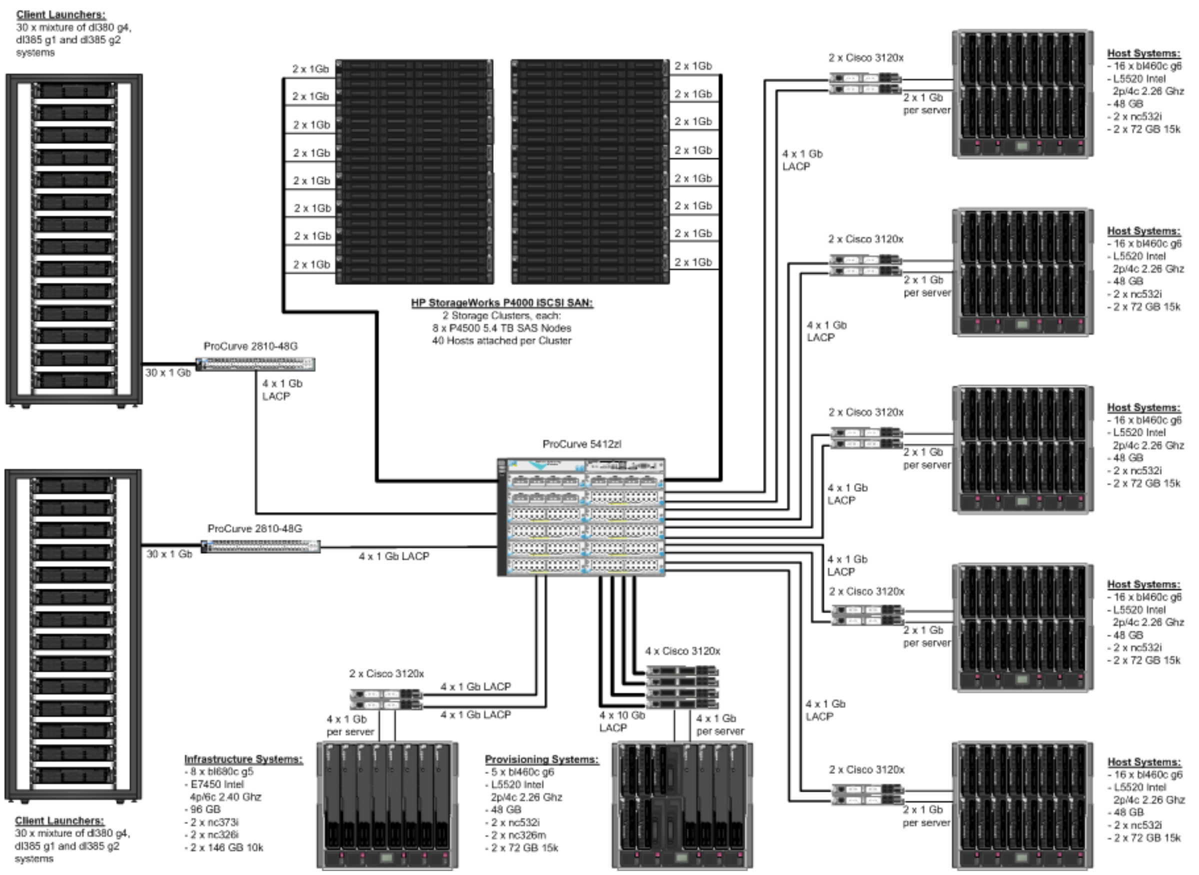

Figure 1: Physical layout of the test environment

Servers

As seen in Figure 1, physical servers in the XenDesktop environment were either BL680 G5 (full height) or BL460c G6 hosts. Additional servers, such as DL380 and DL385, were used outside the environment as client launchers to generate the ICA sessions into the farm. The ten BL680c blades were used primarily as Hyper-V hosts for virtual machines running the infrastructure. The 80 BL460c blades were used as the Hyper-V hosts for the virtual desktops. The desktop host blade servers were housed in five C7000 blade chassis. Two chassis were used for infrastructure and contained five BL460c (one of the BL460c's was unused) and ten BL680c servers. The configuration and role information for the infrastructure server hosts is provided in Table 1 below.

Hostname	Configuration	Purpose
INFRA1	Windows Server 2008 R2 x64 BL680c-G5 96GB, 12-cores	Hyper-V Host Guests: DDC2, DDC3
INFRA2	Windows Server 2008 R2 x64 BL680c-G5 96GB, 12-cores	Hyper-V Host Guests: DDC1, STAT, STAT2
INFRA3	Windows Server 2008 R2 x64 BL680c-G5 96GB, 12-cores	Hyper-V Host Guests: SCVMM5, SQL5
INFRA4	Windows Server 2008 R2 x64 BL680c-G5 96GB, 12-cores	Hyper-V Host Guests: SCVMM2, SQL2
INFRA5	Windows Server 2008 R2 x64 BL680c-G5 96GB, 12-cores	Hyper-V Host Guests: SCVMM3, SQL3
INFRA6	Windows Server 2008 R2 x64 BL680c-G5 96GB, 12-cores	Hyper-V Host Guests: SCVMM4, SQL4
INFRA7	Windows Server 2008 R2 x64 BL680c-G5 96GB, 12-cores	Hyper-V Host Guests: SCVMM6, SQL6

Table 1: Infrastructure Servers

Networking

At the core of the networking infrastructure was an HP ProCurve 5412 gigabyte switch. The C7000 chassis were connected to the core 5412 using Cisco 3120x switches with 4x1Gb aggregated links. The client launcher servers were connected to the core 5412 with ProCurve 2810 gigabit switches and a 4x1Gb aggregated link. If the Cisco 3120x chassis switches had been replaced by HP Flex10 technology, each Hyper-V network port could have supported four physical NIC devices (FlexNICs) enabling total bandwidth of 10Gb[1]. With the Cisco switch configuration each host had only two 1 Gb NICs per host.

[1] For more information on the Flex-10 technology refer to the HP website.

Two VLANs were configured for the environment. VLAN61 represented the normal networking traffic and included a default gateway. VLAN62 was dedicated for iSCSI storage traffic and had no default gateway associated with it.

All 80 Hyper-V virtual desktop hosts had two network adapters. One network adapter was associated with VLAN61 and the other was associated with VLAN62, each with 1Gb of bandwidth.

Storage

The storage provided was an HP StorageWorks P4000 G2 iSCSI SAN. The P4000 G2 SANs, representing a virtualized pool of storage resources, deliver enterprise SAN functionality that enhances virtual environments, simplifies management, and reduces cost. The P4000 is a scale-out, clustered storage system which is comprised of individual "storage nodes" which are clustered together to provide a pool of storage from which iSCSI LUNs are created. These pools are called "clusters". The clusters are self-protecting in that all data is replicated, and I/O is automatically load-balanced, across all the storage nodes in the cluster. Clusters can be expanded at any time by adding additional storage nodes, even while the SAN and volumes are active.

Another consideration is determining the availability settings for the SAN. The P4000 has two levels of protection. First, each storage node in the cluster can be configured to use disk (or hardware) RAID to protect against disk failures. Typical RAID levels (0,5,6,10) are available. In addition, the P4000 offers an additional protection mechanism called Network RAID. Network RAID replicates data blocks across nodes such that any node failure or network connection problems will not affect the availability of the LUNs. Network RAID is set on a per LUN basis. Since both RAID mechanisms impact the usable capacity of the SAN, there is a tradeoff to be made here between availability and capacity. This innovative approach to storage provides a unique data protection level across the entire SAN, reducing vulnerability without driving up costs.

The SAN in the test configuration consisted of 16 P4500 G2 storage shelves or nodes, each with 12 450GB/300K SAS drives providing 5.4TB of raw storage capacity and 2x1Gb iSCSI links. The storage was configured in two clusters each with eight shelves configured with hardware RAID 10 providing about 19TB of storage capacity per cluster. Each storage node had two 1GbE NICs, bonded together for an aggregate of 16Gb/sec of network throughput to the SAN.

The Hyper-V hosts used Windows Server 2008 R2 iSCSI initiator with the Multi-Path I/O (MPIO) drivers to improve response times. The HP P4000 DSM for MPIO enables fault-tolerant paths to the SAN storage nodes while increasing available bandwidth to the SAN[2]. The HP P4000 DSM for MPIO enables a Hyper-V server to continually communicate with the storage node that is hosting a copy of the data that the server is requesting. In other words, the MPIO drivers allow each host to keep a copy of the data mapping so that data lookup services and I/O are distributed properly across the storage nodes. When MPIO[3] was enabled in the environment, a significant boost in SAN response time was noted and the CPU utilization was no longer biased to a single storage node.

---

[2] For more information on the HP P4000 DSM for MPIO refer to the HP website at http://www.hp.com/go/P4000.

[3] The multi-path driver version used for the P4500 G2 was MPIO 8.1 with patch 8.1.0.85

Each Hyper-V host had a 200GB LUN configured for maximum availability. Each LUN was assigned to a host via the Windows iSCSI initiator. This LUN hosted the virtual machine configurations, write-cache drives, and Hyper-V saved-state file.

Logical

This section discusses the logical architecture for the environment. The architecture includes support services, hosting infrastructure and desktop delivery components. Each of the components is discussed below.

Support Services

The support services for this environment contain all the network and domain services required to support a virtual desktop environment. Table 2: Infrastructure Support Servers contains the information about each of the servers running in the environment.

Hostname	Configuration	Purpose
DC1	Windows Server 2008 R2 x64 Virtual - 4GB RAM, 4vCPU Hosted on INFRA2	Active Directory Domain Controller in 2003 Forest/Domain Functional Level DHCP Services
DC2	Windows Server 2008 R2 x64 BL460c – G6 48GB, 8-cores	Active Directory Domain Controller[4] DNS Services (Forward/Reverse lookup zones both)
FS1	Windows Server 2008 R2 x64 BL680c-G5 96GB, 12-cores	File server hosting the lab files and the Login VSI file share.
PROFILE1	Windows Server 2008 R2 x64	File server hosting roaming

[4] A physical server was used because early there were issues with DC1 and the networking on INFRA1 where it was hosted. As a stop-gap measure, a second DC was created on a physical server. Later the issue with INFRA1 was discovered to be the order of installation for the HP NIC Teaming drivers.

Hostname	Configuration	Purpose
	BL680c-G5 96GB, 12-cores	profiles for USER1-USER4000.
SQL1	Windows Server 2008 R2 x64 SQL Server 2008 R2 SP1 BL680c-G5 96GB, 12-cores	SQL Server for farm databases Hosted XenDesktop, Provisioning Services, and STAT databases.[5]

Table 2: Infrastructure Support Servers

Hosting Infrastructure

Eighty Hyper-V hosts were used to host the virtual machines running XenDesktop. Each host was on a BL460c with 48GB of RAM and two network adapters. As mentioned earlier, one network adapter was configured for the iSCSI storage VLAN62 and the other was assigned to VLAN61 for normal traffic.

Hostname	Configuration	Purpose
HOST01 – HOST80	Windows Server 2008 R2 x64 BL460c-G6, 48GB, 8-cores	Hyper-V Host for 46 guest VMs

Table 3: Host Servers

All Hyper-V hosts were running a full version of Windows Server 2008 R2. The core installation mode was not used because the STAT agent was installed and configured to gather performance data and the STAT agent required a normal install mode. The VLAN61 network adapter was configured in Hyper-V as the host network. Forty-six virtual machines were created on each host server providing a total of 3680 VMs available to the farm; however due to memory constraints, no more than 44 VMs were ever running simultaneously on a single host.

After operating system installation, all critical and recommended Windows Updates were installed and the HP NIC drivers installed. Later the servers were updated with the latest MPIO iSCSI drivers.[6]

---

[5] A physical SQL Server was used because the STAT database required the additional horsepower as 160+ agents were reporting performance monitor counters at 15-second intervals and using the SQL queues for launch processing.

[6] The multi-path driver version used P4500 G2 MPIO 8.1 with patch 8.1.0.85

Each server contained an HP NC532i Dual Port Adapter and one network port was configured to for VLAN62 which was exclusively used for the iSCSI traffic. All servers were then configured with an individual 200GB iSCSI LUN that resided on the P4000 SAN.

XenDesktop Infrastructure

The XenDesktop 4 delivery infrastructure included the XenDesktop Desktop Delivery Controllers (DDC), Provisioning Servers (PVS), and License Server. In addition, Microsoft System Center Virtual Machine Manager (VMM) servers were configured at the ratio of one SCVMM server per 16 hosts. Each VMM server was virtual and had a dedicated virtual SQL Server on the same physical host. Table 4: XenDesktop Infrastructure Servers provides information about each of the hosts providing XenDesktop related services within the environment.

Hostname	Configuration	Purpose
DDC1	Windows Server 2003 R2 x64 Virtual – 8GB, 4 vCPUs[7] Hosted on INFRA2	XenDesktop Delivery Controller XD Pool Master Citrix License Server
DDC2	Windows Server 2003 R2 x64 Virtual – 8GB, 4 vCPUs[7] Hosted on INFRA1	XenDesktop Delivery Controller XML Resolutions VDA Registrations
DDC3	Windows Server 2003 R2 x64 Virtual – 8GB, 4 vCPUs[7] Hosted on INFRA2	XenDesktop Delivery Controller XML Resolutions VDA Registrations
PROV1	Windows Server 2008 R2 x64 BL460c-G6 48GB, 8-cores 4x1Gb Teamed NICs	Provisioning Services Windows 7 vDisk
PROV2	Windows Server 2008 R2 x64 BL460c-G6 48GB, 8-cores 4x1Gb Teamed NICs	Provisioning Services Windows 7 vDisk
PROV3	Windows Server 2008 R2 x64 BL460c-G6 48GB, 8-cores 4x1Gb Teamed NICs	Provisioning Services Windows 7 vDisk

[7] This configuration of Windows Server 2003 R2 x64 is not explicitly supported by Microsoft on Hyper-V with 4 vCPUs.

Hostname	Configuration	Purpose
SCVMM2	Windows Server 2008 R2 x64 SCVMM 2008 R2 Virtual 4GB, 4 vCPUs Hosted on INFRA4	VMM Server assigned to HOST65-HOST80
SCVMM3	Windows Server 2008 R2 x64 SCVMM 2008 R2 Virtual 4GB, 4 vCPUs Hosted on INFRA5	VMM Server assigned to HOST01-HOST16
SCVMM4	Windows Server 2008 R2 x64 SCVMM 2008 R2 Virtual 4GB, 4 vCPUs Hosted on INFRA6	VMM Server assigned to HOST17-HOST32
SCVMM5	Windows Server 2008 R2 x64 SCVMM 2008 R2 Virtual 4GB, 4 vCPUs Hosted on INFRA3	VMM Server assigned to HOST33-HOST48
SCVMM6	Windows Server 2008 R2 x64 SCVMM 2008 R2 Virtual 4GB, 4 vCPUs Hosted on INFRA7	VMM Server assigned to HOST49-HOST64
SQL2	Windows Server 2008 R2 x64 SQL Server 2008 R2 SP1 Virtual 16GB, 4 vCPUs Hosted on INFRA4	SQL Server assigned to support SCVMM2
SQL3	Windows Server 2008 R2 x64 SQL Server 2008 R2 SP1 Virtual 16GB, 4 vCPUs Hosted on INFRA5	SQL Server assigned to support SCVMM3
SQL4	Windows Server 2008 R2 x64 SQL Server 2008 R2 SP1 Virtual 16GB, 4 vCPUs	SQL Server assigned to support SCVMM4

Hostname	Configuration	Purpose
	Hosted on INFRA6	
SQL5	Windows Server 2008 R2 x64 SQL Server 2008 R2 SP1 Virtual 16GB, 4 vCPUs Hosted on INFRA3	SQL Server assigned to support SCVMM5
SQL6	Windows Server 2008 R2 x64 SQL Server 2008 R2 SP1 Virtual 16GB, 4 vCPUs Hosted on INFRA7	SQL Server assigned to support SCVMM6

Table 4: XenDesktop Infrastructure Servers

After installation of the operating system, Windows Update was initiated and all critical and recommended updates were installed. In the case of the System Center Virtual Machine Manager servers, Windows Update was run an additional time to pick up some critical fixes for VMM Server.

Farm Configuration

Following a standard XenDesktop farm installation, the following additional items were configured or installed:

1. Installed Microsoft Hotfix for unsigned PCL printer drivers KB961118 (http://support.microsoft.com/kb/961118)
2. Installed Citrix Pool Management hotfix 4.0.4530. The final 3500-desktop test used a privately issued CdsPoolMgr binary that was updated with specific fixes for this environment. These fixes should now be available in Hotfix 3 or through Citrix Support.
3. Installed Citrix Desktop Delivery Controller hotfix 4.0.4529
4. Disabled SpeedScreen Multimedia
5. Created XenDesktop policy to disable client printer mappings
6. Configured DDC1 as Farm master and disabled accepting registrations per CTX117477 (http://support.citrix.com/article/ctx117477)
7. Configured DDC2 as Backup Farm master per CTX117477 (http://support.citrix.com/article/ctx117477)
8. Configured DDC3 as Member server per CTX117477 (http://support.citrix.com/article/ctx117477)
9. Created five Desktop groups, one for each chassis and each pointing to a single SCVMM server dedicated to that chassis.

The Provisioning Services farm was also created. The following items represent additional changes to the environment after the initial default installation.

1. Installed NIC teaming drivers and configured a single network that combined all 4 1Gb network cards into a single 4Gb link
2. Re-ran the Provisioning Services Configuration wizard to reconfigure the services for the new network
3. Configured the bootstrap file to contain the static IP address assigned to each of the provisioning servers
4. Created a local vDisk store named "LocalStorage" for each of the provisioning servers and configured it for the D: drive
5. Copied the Windows7 vDisk to each servers LocalStorage location

Testing Results

This section covers the high-level results for the testing completed at the HP Solution Center. The first section reports on the results achieved from testing a single HP BL460c – G6 server in four configurations. The second section discusses the results from the large farm testing where 3505 desktops received a perfect Login VSI Max score.

Single-Server Scalability

The lab consisted of 80 BL460c G6 servers each with 48GB of RAM installed. The primary driver of the single-server scalability (SSS) testing was to determine a feasible number of guests for each of the 80 servers. Since the opportunity presented itself, three additional configurations were also tested. Table 5: SSS Tested Configurations outlines the four tested configurations.

Test Run	Configuration	Notes
48GB – SAN	BL460c – G6 48GB RAM, 8-cores 2 x 2.26 GHz L5520 Nehalem HP NC532i Dual Port Adapter	Matched host configuration that would be used in the large-farm tests.
96GB – SAN	BL460c – G6 96GB RAM, 8-cores 2 x 2.26 GHz L5520 Nehalem HP NC532i Dual Port Adapter	RAM doubled to a more economical configuration.
48GB – SSD	BL460c – G6 48GB RAM, 8-cores 2 x 2.26 GHz L5520 Nehalem 2 x 120GB SSD	Local solid-state drives were configured in a RAID0 array and included the operating system.
96GB – FIO	BL460c – G6 96GB RAM, 8-cores 2 x 2.26 GHz L5520 Nehalem 1 Fusion I/O 320GB	The Fusion I/O mezzanine card was configured with maximum storage space.

Table 5: SSS Tested Configurations

When testing with the SAN configuration, no other load was placed on the SAN or the blade chassis, so optimum throughput would be achieved. In all instances of the single-server testing, the Provisioning Server write-cache drives for the guests were placed entirely on the storage. In addition, the VM configuration files for Hyper-V and the save-state BIN file were also placed on the storage with the write-cache drives. With the exception of the 48GB – SSD test run, the operating system was always on a different partition than the write-cache and save-state files.

Procedure

When performing the single-server scalability test the following procedure was executed for each test run:

1. Reboot all of the guest virtual machines. This step resets the operating system back to a new image with no user profiles and causes the system page file to be recreated.
2. Wait for all the desktops to report an "Idle" status in the Delivery Services Console.
3. Launch user sessions at the rate of one session every 20 seconds or three per minute.
4. Automatically start the Login VSI medium workload upon desktop login.
5. When the last user session has launched, wait 15 minutes minimum to allow the Login VSI workload (12-minute cycle) to complete.
6. Add the Logoff.txt file to the Login VSI share to signal all desktops to logoff at the end of their cycle.
7. Analyze the Login VSI results and record the data.

The results from the single-server scalability tests are shown in Table 6: SSS Test Results.

Test Run	VSIMax Score	VM Density per core	VSI Response Time Avg./Max (ms)	Login Time Avg.	Minimum Idle CPU%
48GB – SAN	45	5.625	634/2561	93.31[8]	83.57%
96GB – SAN	90	11.25	1194/3784	38.33	4.64%
48GB – SSD	45	5.625	633/1212	20.39	87.34%
96GB – FIO	88	11	1059/3365	24.46	37.66%

Table 6: SSS Test Results

The testing with the Fusion I/O card was repeated multiple times in an effort to determine why the performance was lower given that the Fusion I/O card had specifications that were significantly higher than those for the solid-state drives. Additional configuration options, including maximize for write operations (at a cost of 50% of the card capacity), smaller block size (4K instead of 64K), and

[8] Approximately 32 workstations took longer than 120 seconds to login, which was later determined to be related to a Microsoft licensing issue and resolved. The adjusted number would be approximately 29.2 seconds.

different card (160GB) were also tested, but the performance and density were not improved significantly.

Conclusions from SSS Testing

The SSS testing portion is undertaken for the sole purpose of identifying the number of VMs that a single server can ideally manage in the environment. The results showed that the number of VMs per server for the 80 hosts in the farm should be no more than 45. Given that the server would be collecting performance statistics, 44 guests per host was selected as the per host maximum. With 44 guests per host, the total number of virtual desktops available would be 3520. In addition, the following conclusions were drawn based on the results of the testing performed:

- The increase in response time between the 48GB – SAN and the 96GB – SAN is attributed to the high CPU utilization which slowed user response times and possibly throttled SAN communication. A later unpublished run with 75 guests shows a minimum idle CPU performance around 75% and a more respectable average/max response time of 748ms/1482ms, suggesting that 75 guests is a better fit for the hardware.
- From the data gathered, it appears that VM density per core is at least somewhat dependent on the speed and responsiveness of the storage layer.
- The SSDs provided the best response times for the login test and the fastest boot times for the images. In fact, if the SSDs are used to host the write-cache drives the boot time for the VM is under a 30 seconds suggesting that an idle pool of desktops may not be necessary in situations where users are willing to wait for their desktop to be started.
- SSDs are a viable alternative for hosting the PVS write-cache drives and the Hyper-V memory save-state files. The SSS score was the same as a SAN and the cost to maintain and configure SSDs may provide a considerable savings over that of a SAN. The only unknown with this configuration is the endurance of SSDs in this type of VDI environment.
- The prevailing theory around the reason for the lower score on the Fusion I/O card is that the data management requires host memory and CPU resources, effectively decreasing the total number of workstations that the host can manage at full capacity. In addition, the write-cache drives, which are almost exclusively write operations, require a considerable amount of resources since existing data must be erased before being written.

Large-Farm Testing

The large-farm testing takes the single-server scalability results to the next level with a goal of 3500 Windows 7 desktops across 80 host servers. As mentioned earlier, the test runs were configured on 700-desktop increments. Although more than five test runs were completed in the lab, Table 7: Large-Farm Tested Configurations outlines the five primary test runs executed in the environment.

Test Run	Configuration	Notes
700	16 Hosts Single C7000 Chassis Single P4500 RAID5 cluster	Only test run to use RAID5 SAN configuration
1400	80 Hosts Five C7000 Chassis Dual P4500 RAID10 clusters	First run with Multipath I/O enabled
2100	80 Hosts Five C7000 Chassis Dual P4500 RAID10 clusters	High number of stuck sessions caused by Login VSI scripts hanging
2800	80 Hosts Five C7000 Chassis Dual P4500 RAID10 clusters	Fixed hung session issues with a workaround provided by Login Consultants
3500	80 Hosts Five C7000 Chassis Dual P4500 RAID10 clusters	First run to execute with the power management settings configured for Static High Performance Mode

Table 7: Large-Farm Tested Configurations

Since the ultimate goal was in fact 3500 desktops, the earlier test levels are used to refine the configuration and identify possible barriers to success at the 3500 desktop level. Since all the tests were not executed with the exact same configuration, the significant differences between the test runs are outlined below.

Test run at 700: The 700-desktop test run occurred early on while the farm was still being built. All 700 desktops were confined to a single C7000 chassis to verify that SAN traffic would not bottleneck at the chassis level. Also the storage traffic was confined to single cluster of seven shelves on the P4500 chassis that was configured using hardware RAID5 (for disk striping) while the remaining shelves were being shipped.

Test runs at 1400, 2100, 2800: These test runs occurred with the desktops being evenly distributed across all five C7000 chassis. In addition, the hosts were evenly distributed across the SAN with HOST01-HOST40 using cluster 1 and HOST41-HOST80 using cluster 2.

Test run at 3500: This test run had the same configuration as runs 1400-2800, but one additional change was made to disable the BIOS dynamic power feature to allow all the cores to run at full capacity. This change was made on the final run because the hosts were reporting event log entries indicating the BIOS was limiting the available processor.

Procedure

The following test procedure was executed for each test run:

1. Set the idle pool for each of the desktop groups to 0.
2. Power off all the guest virtual machines.
3. Set the idle pool for each of the five desktop groups equal to 20% of the total number of desktops expected. For instance with 1400 desktops, each of the five desktop groups would be set to provide a total of 280 desktops at idle. This approach allowed the system to randomly select desktops for the test.[9]
4. Once the desktops are registered as "Idle", set all other desktops to maintenance mode to prevent startup during the test.
5. Set the launch rate to have all desktops logged in within 30 minutes. The launch rates were as follows:

700 Desktops	1 desktop every 2.5 seconds
1400 Desktops	1 desktop every 1.25 seconds
2100 Desktops	1 desktop every .9 seconds
2800 Desktops	1 desktop every .6 seconds
3500 Desktops	1 desktop every .5 seconds

6. Launch the user sessions.
7. Automatically start the Login VSI medium workload upon desktop login.
8. When the last user session has launched, wait 15 minutes minimum to allow a complete Login VSI workload (12-minute cycle) to complete.
9. Add the Logoff.txt file to the Login VSI share to signal all desktops to logoff at the end of their cycle.
10. Analyze the Login VSI results and record the data.

Login VSI Test Results

This section provides detail around the scores and response times measured during the testing. The Table 8: Login VSI COPI scores for Primary Test Runs shows the scoring model values from each of the primary test runs. In addition, Table 9: Large-Farm Performance Results displays the key performance indicators for each of the primary test runs.

[9] Since the testing for the 700 desktops was confined to a single chassis, only a single desktop group was used.

Test Run	700	1400	2100	2800	3500
Total Sessions Launched	700	1400	2100	2800	3505
Uncorrected Optimal Performance Index (UOPI)	700	1400	2100	2800	3505
Stuck Session Count before UOPI (SSC)	0	31	189	85	0
Lost Session Count before UOPI (LSC)	0	0	0	4	0
Corrected Optimal Performance Index (COPI = UOPI - (SSC*50%) - LSC)	700	1369	1911	2711	3505

Table 8: Login VSI COPI Scores for Primary Test Runs

Test Run	VSI COPI Score	VSI Response Time Avg./Max (ms)	Average Login Time
700	700	612/1791	34.34
1400	1369	770/960	20.53
2100	1911	619/1875	22.71
2800	2711	645/1551	23.03
3500	3505	575/1761	59.19[10]

Table 9: Large-Farm Performance Results

[10] Due to time constraints with the environment, the actual cause of this increased login time was not determined.

Figure 2: Login VSI Response Time Graph shows the Login VSI Max response times as measured within the guest operating system during the 3500-desktop test run.

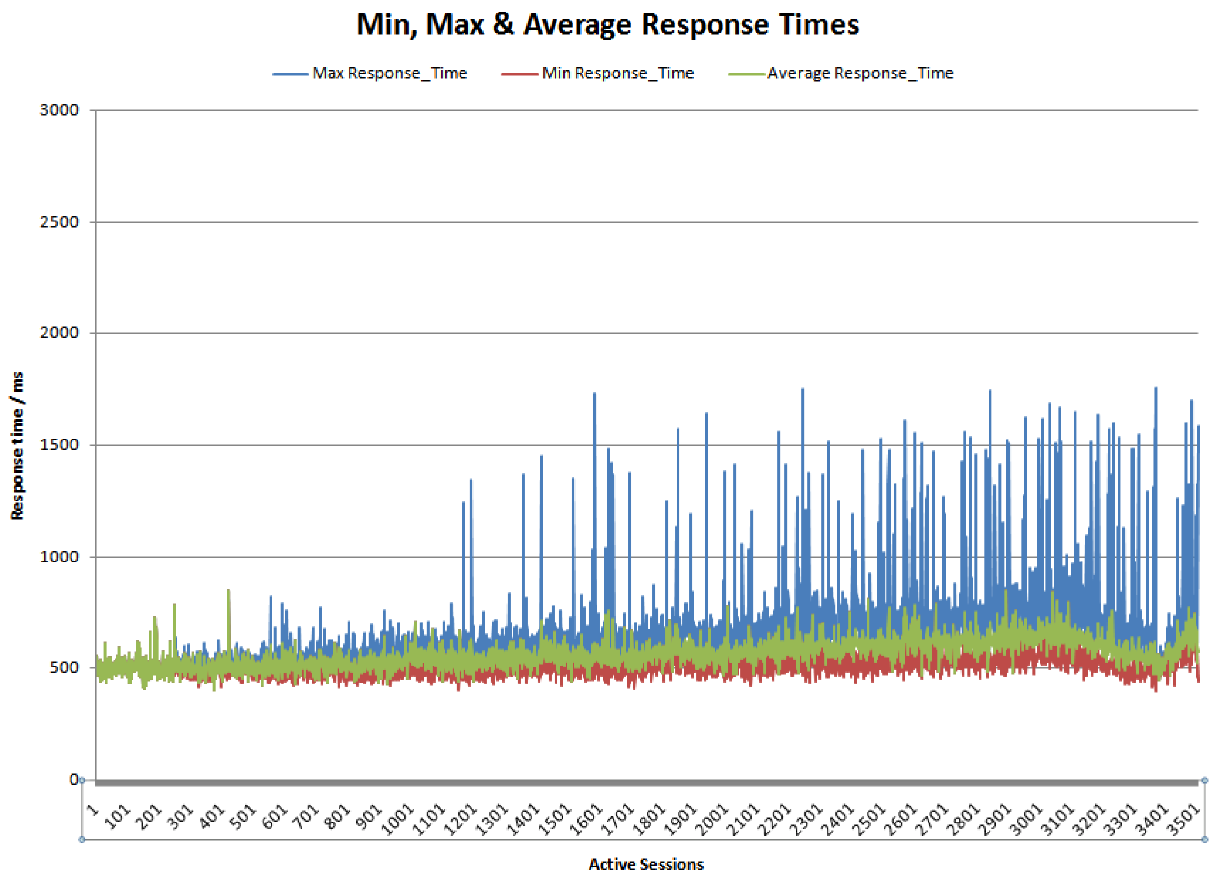

Figure 2: Login VSI Response Time Graph

As indicated by the graph, the maximum response time does grow over the course of the test but is entirely contained within the 2-second "optimum" response time window. The response time is generally an indicator of host CPU resources and since the VM density (5.6 per core) is just over half of the maximum results (11 per core) obtained during the SSS testing, these results are to be

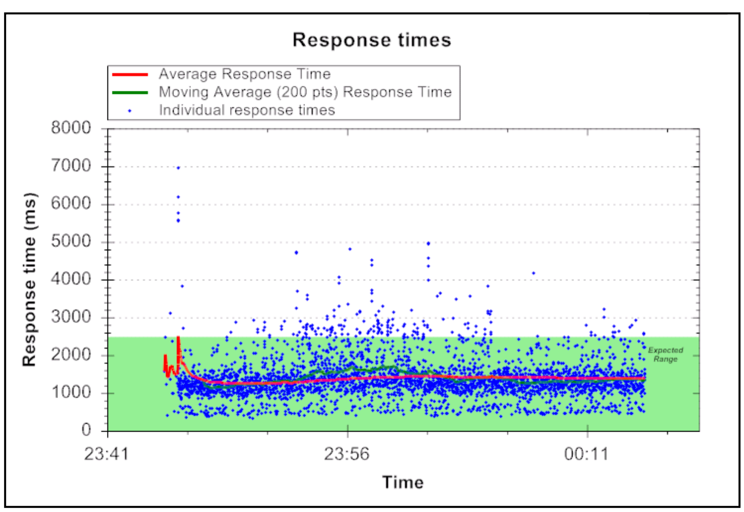

Figure 3: XML Response Times

expected. The average response time provides a great indicator of user experience, and in the 3500-desktop test it stayed around 600 milliseconds indicating an excellent user experience.

The STAT tool also records various metrics during the course of the test. One such metric is the XML response time. This metric is generally a good indicator of DDC health during the test. When the response times grow beyond 2500ms on average, the infrastructure is getting bogged down. Figure 3: XML Response Times shows the XML response times recorded during the 3500-desktop run.

Figure 3 indicates that both the Average Response Time and Moving Average trend line are within the 2500 ms acceptable response time range. The longest was 7000 ms which occurred at the beginning of the test; presumably because both DDCs had recently been rebooted and the .NET Framework had not yet been loaded so the first requests took longer to be serviced.

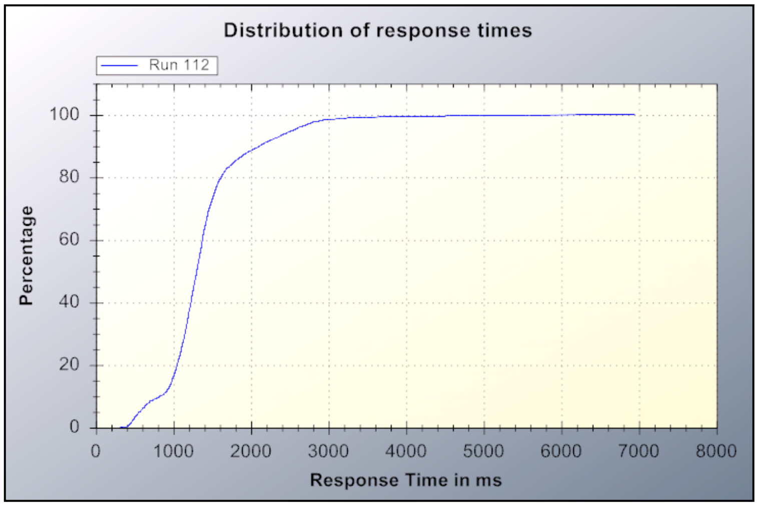

Figure 4: Distribution of XML Response Times

Figure 4: Distribution of XML Response Times provides a view of the same information by distribution. This graph shows the percentage of responses that completed within a particular time frame. In this case, 95% of the response times completed within the expected range of 2500ms. The ideal graph would have 100% of the responses within the 2500ms range.

Conclusions from Large-Farm Testing

Receiving a perfect Login VSI score for a large-farm test is unusual and in this case it is an indication that the hardware was fully utilized. The SSS testing indicated that the BL460c server is capable of hosting up to 90 desktops, yet in the large-farm testing, only 44 desktops were running on each blade due to limitations on available RAM. Several conclusions can be drawn from the large-farm test results:

- The farm could be grown easily by adding RAM to the existing servers assuming the user workload remained close to the tested Login VSI Medium workload. Conversely, the existing farm could support a workload that was more CPU-intense than the one tested.
- The five percent of XML response times that fall outside of the recommended response period seem to indicate the DDCs were experiencing a resource-constrained.

Detailed Performance

This section provides the key performance counters on the core infrastructure items as gathered during the testing of the 3500-desktop run. One section is provided for each infrastructure component: Desktop Delivery Controller, Provisioning Server, System Center Virtual Machine Manager, and the SQL Servers associated with the SCVMM databases. Keep in mind that not every metric collected is shown here.

The format of the metrics is as follows:

- ***Analysis***: Relevant comments regarding the performance monitoring data. When trends or issues can be identified in the chart, they will be noted in the Analysis field.
- ***Chart***: Excel chart of the perfmon counter data. The chart title indicates the XenDesktop component, performance counter, and test cycle (Bootup or Login) and uses the following format: **[Component role]-[Perfmon Counter]-[Bootup|Login]**

Thus a chart title of **DDC - CPU Idle Time - Bootup** would include the CPU Idle Time data for all DDCs during the Bootup cycle of the desktops. In order to get the data in a readable format the charts are set to generate the x- and y-axis labels automatically, so the axis labels vary and should be noted for each chart.

Desktop Delivery Controller Performance

During all the test runs DDC1 was functioning as the Pool Master and Farm Master. Likewise DDC2 and DDC3 were functioning solely in a supportive role handling only XML resolutions and VDA registrations.

Analysis	Chart
During the Bootup phase of the tests all three DDCs were significantly busy.	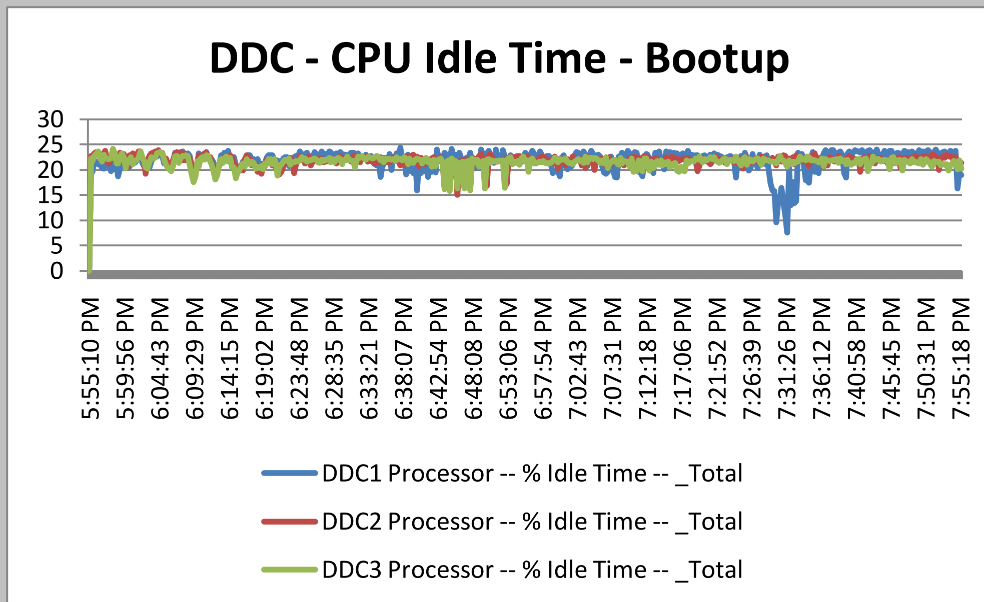
The available CPU cycles during the Login cycle are not much better than during the Bootup cycle. This chart suggests that the loads were pretty evenly distributed across all the DDCs.	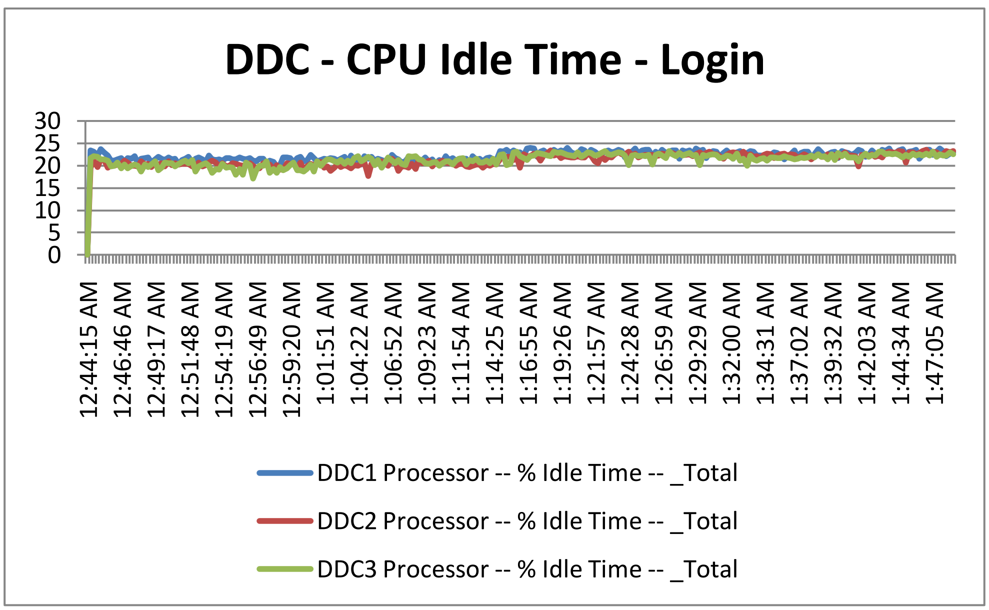

Analysis **Chart**

DDC1 was the pool manager and that process is consuming the majority of the CPU cycles on DDC1, with 20-60% of a single vCPU being used by the Pool Manager service as it managed the startup requests for all the desktops in the farm.

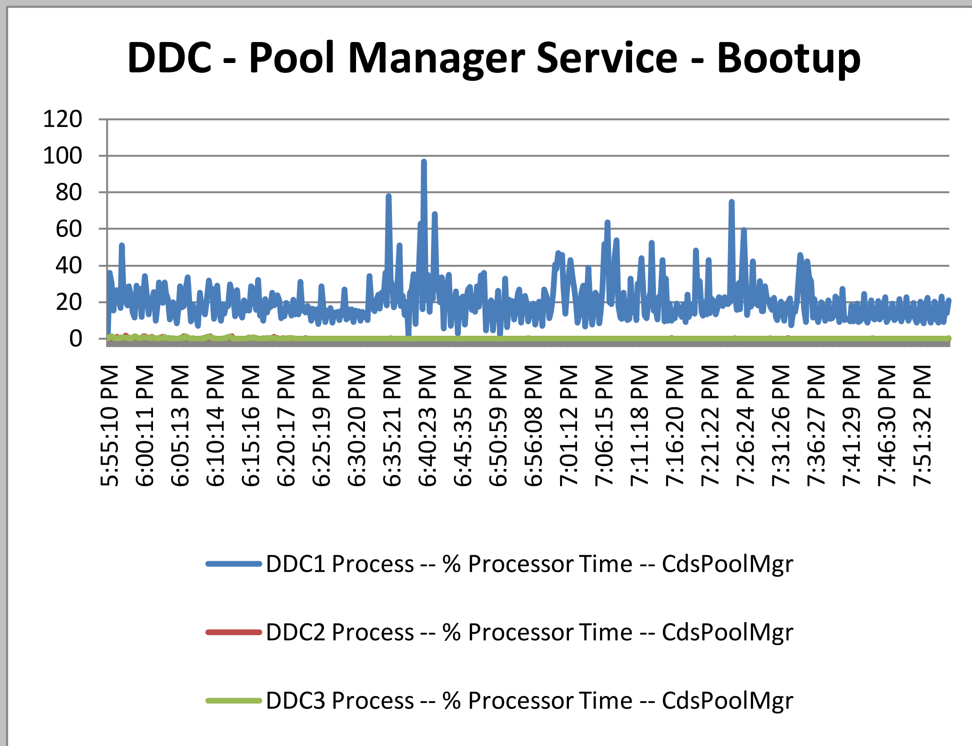

DDC1 performing the Pool Master role during Login phase as well. This chart suggests that even during the Login phase the Pool Manager service is busy communicating with SCVMM for desktop status before routing users to the desktops.

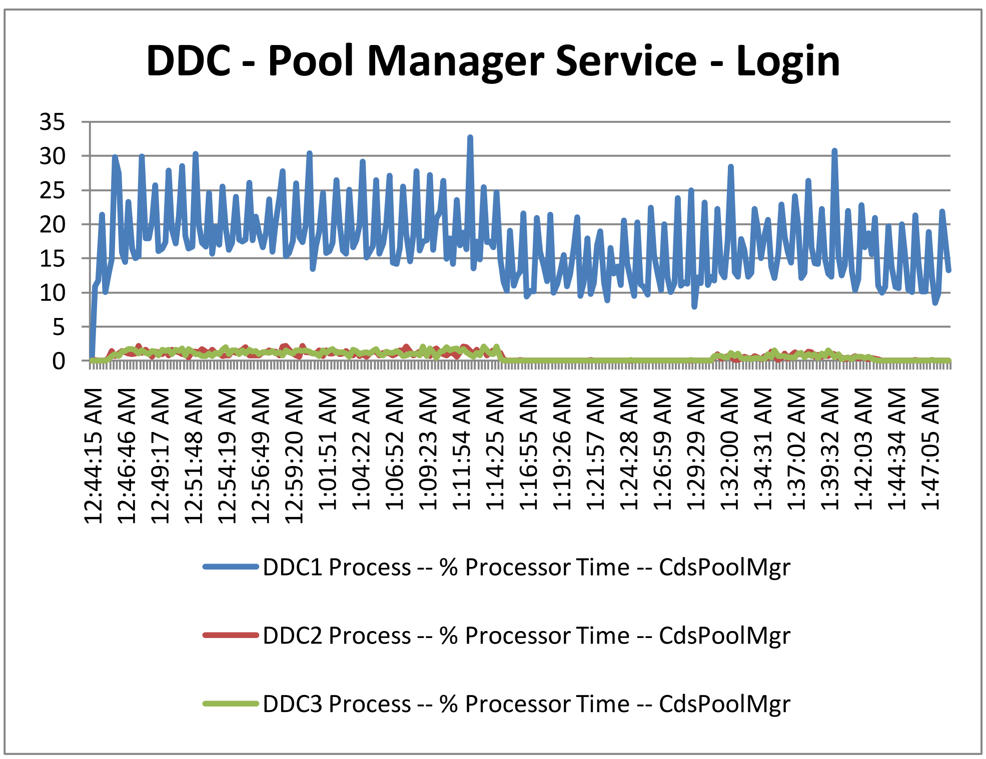

Analysis	Chart
On a 4 vCPU guest, the maximum %Processor Time is 400%. This graph indicates that approximately one and a half vCPUs are required to meet the IMA service load generated during the Bootup phase.	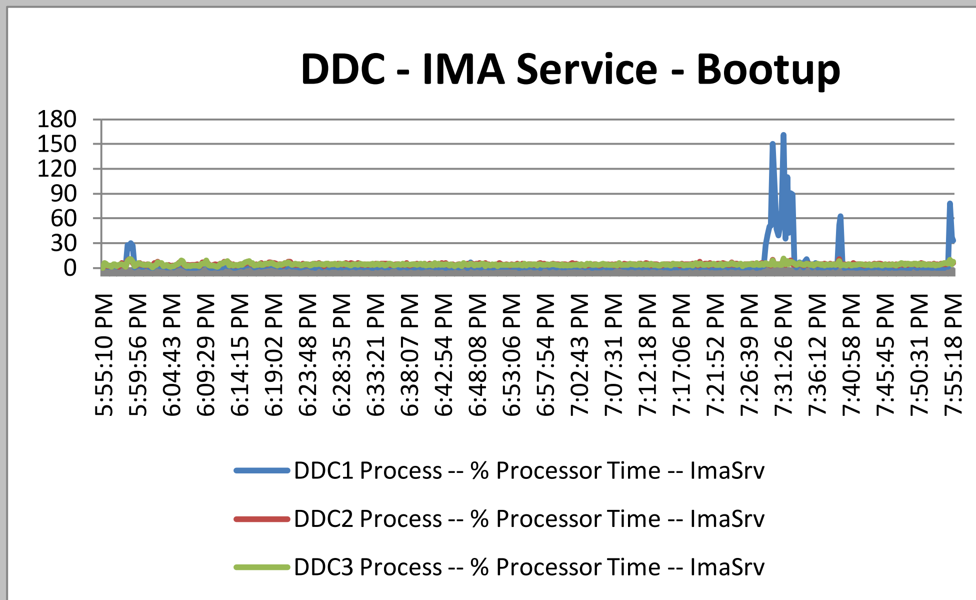
IMA Service is responsible for XML resolutions and also brokering the desktops, which explains the utilization levels during the first 30-minutes of the test when the desktops are launching.	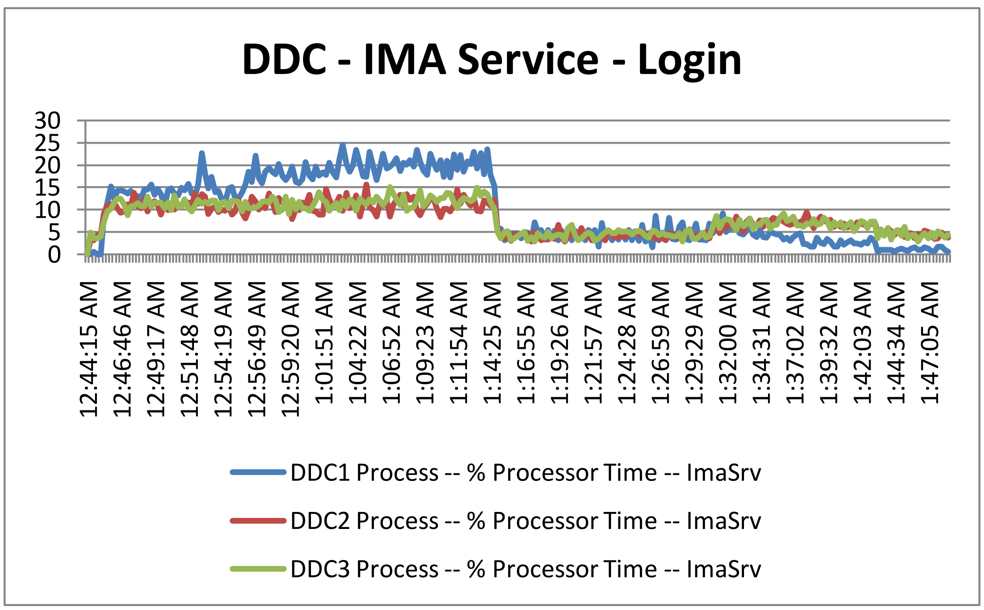

Analysis	Chart
Among other things, the desktop service is responsible for all VDA registrations which would be the primary consumer of this service during the Bootup phase as the desktops randomly select a DDC to register with. DDC1 was configured to reject all registrations.	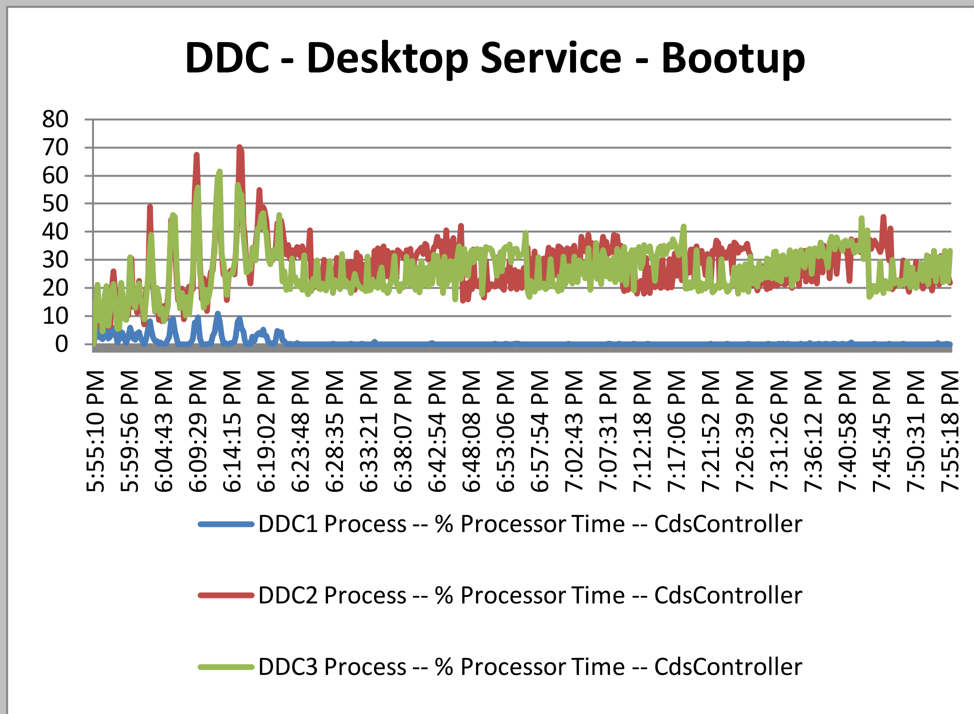
The desktop service is also responsible for managing the connection with VDA and pushing the XenDesktop policies to the desktops as the logins occur.	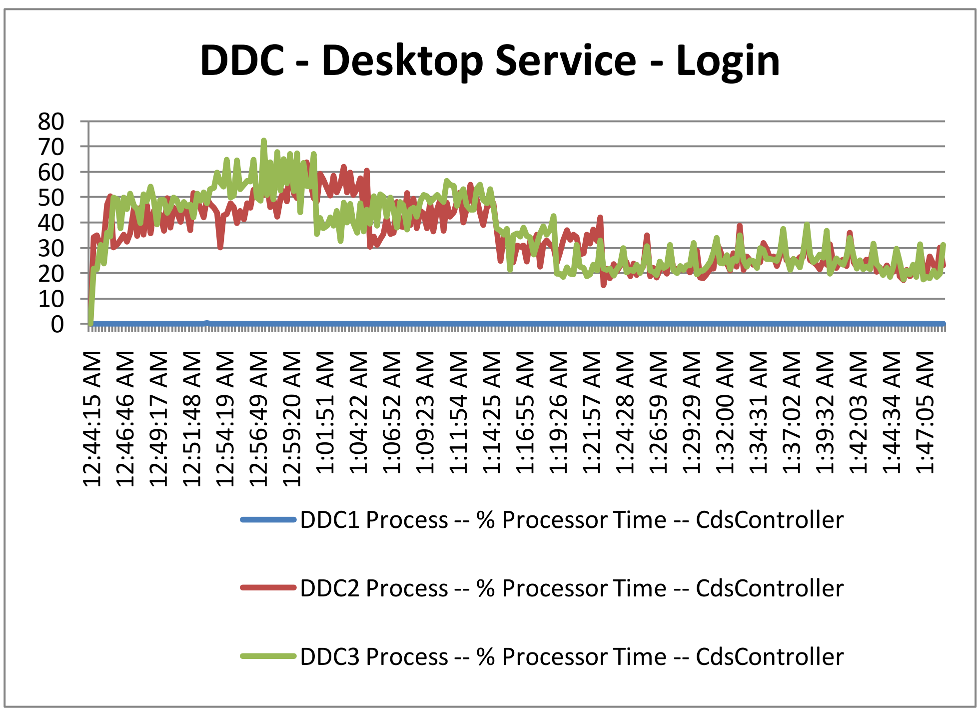

<table>
<tr><th>Analysis</th><th>Chart</th></tr>
<tr><td>Not a significant amount of disk utilization on the DDC, which would be expected, since the DDC is typically RAM and CPU-intensive.</td><td>
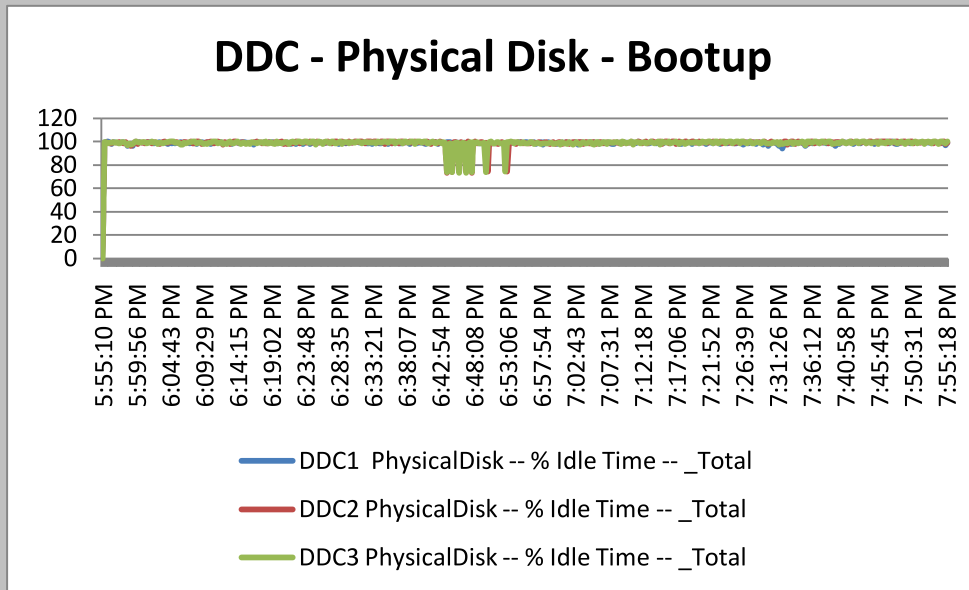

</td></tr>
<tr><td>Even less disk activity is occurring on the physical disk of the DDC during the login cycle.</td><td>
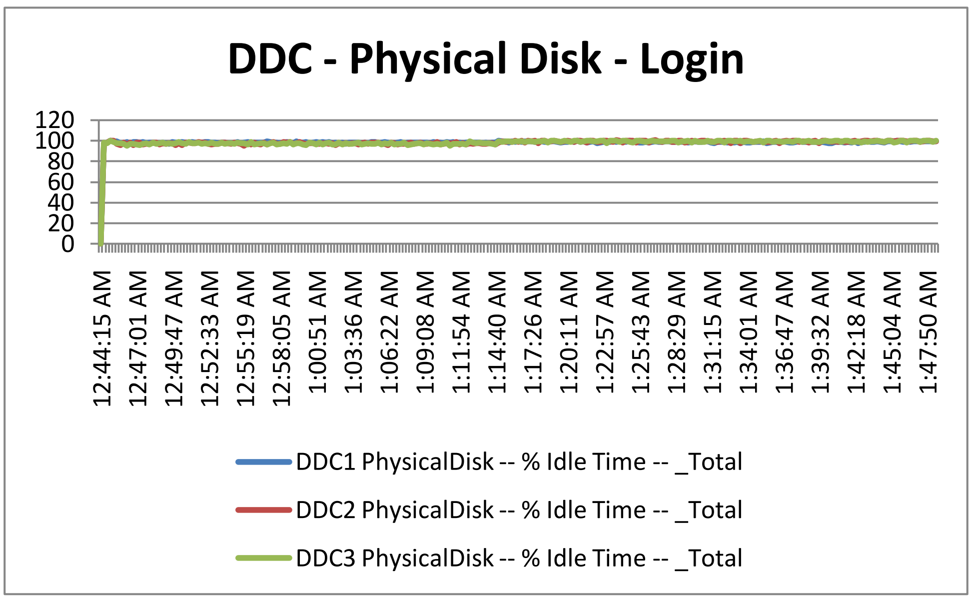

</td></tr>
</table>

Analysis	Chart
The context switching is well within normal operating parameters.	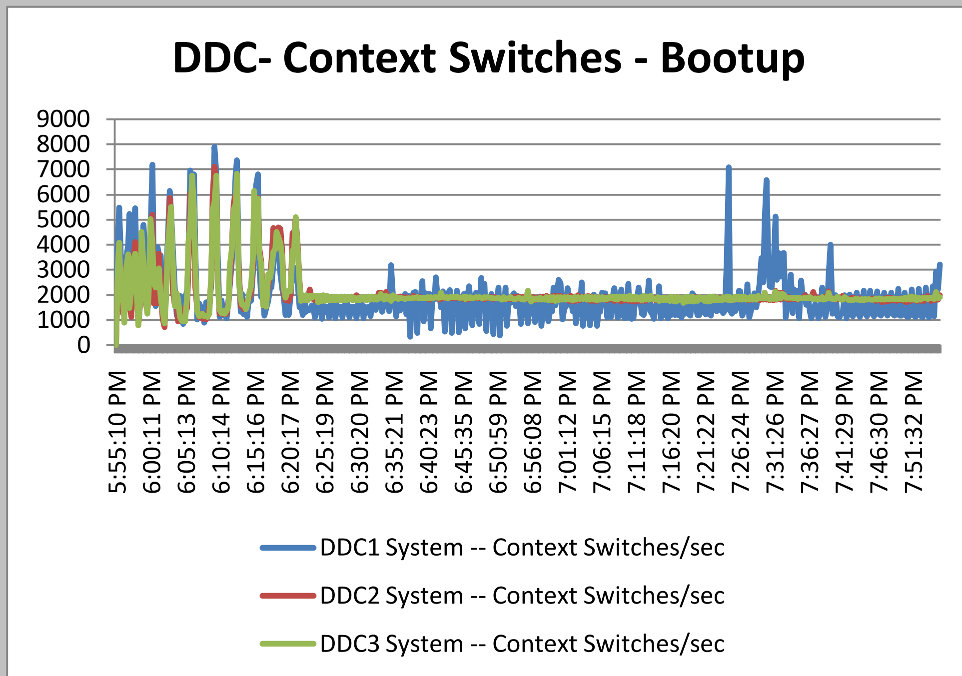
The context switching is well within normal operating parameters.	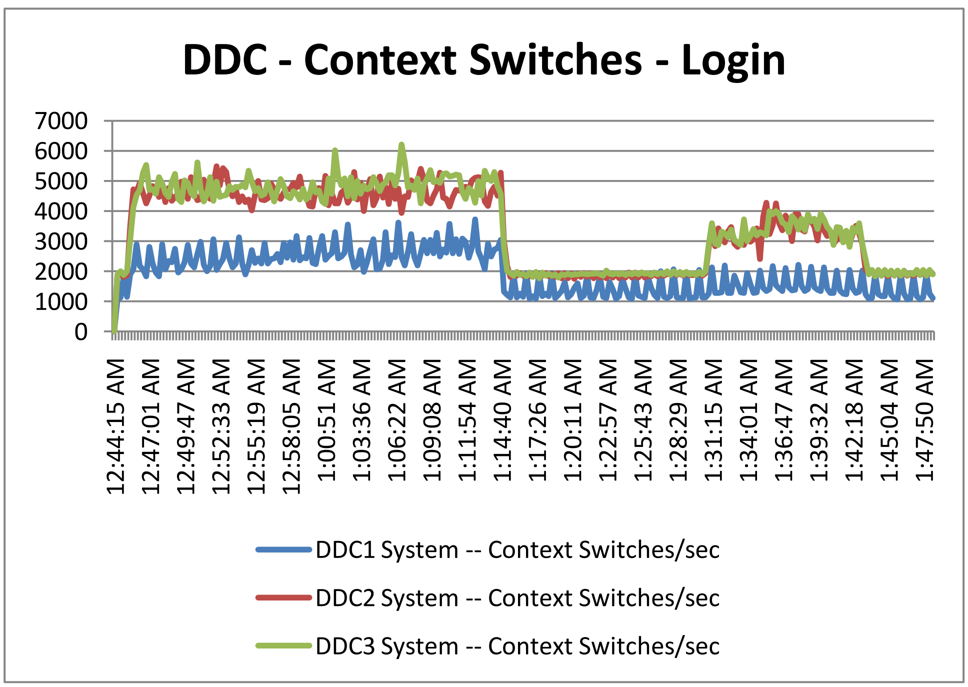

Analysis	Chart
Though this looks like a lot of data is getting sent out, 1 million bytes is about 7.62MB of traffic. The network traffic is well within normal operating parameters.	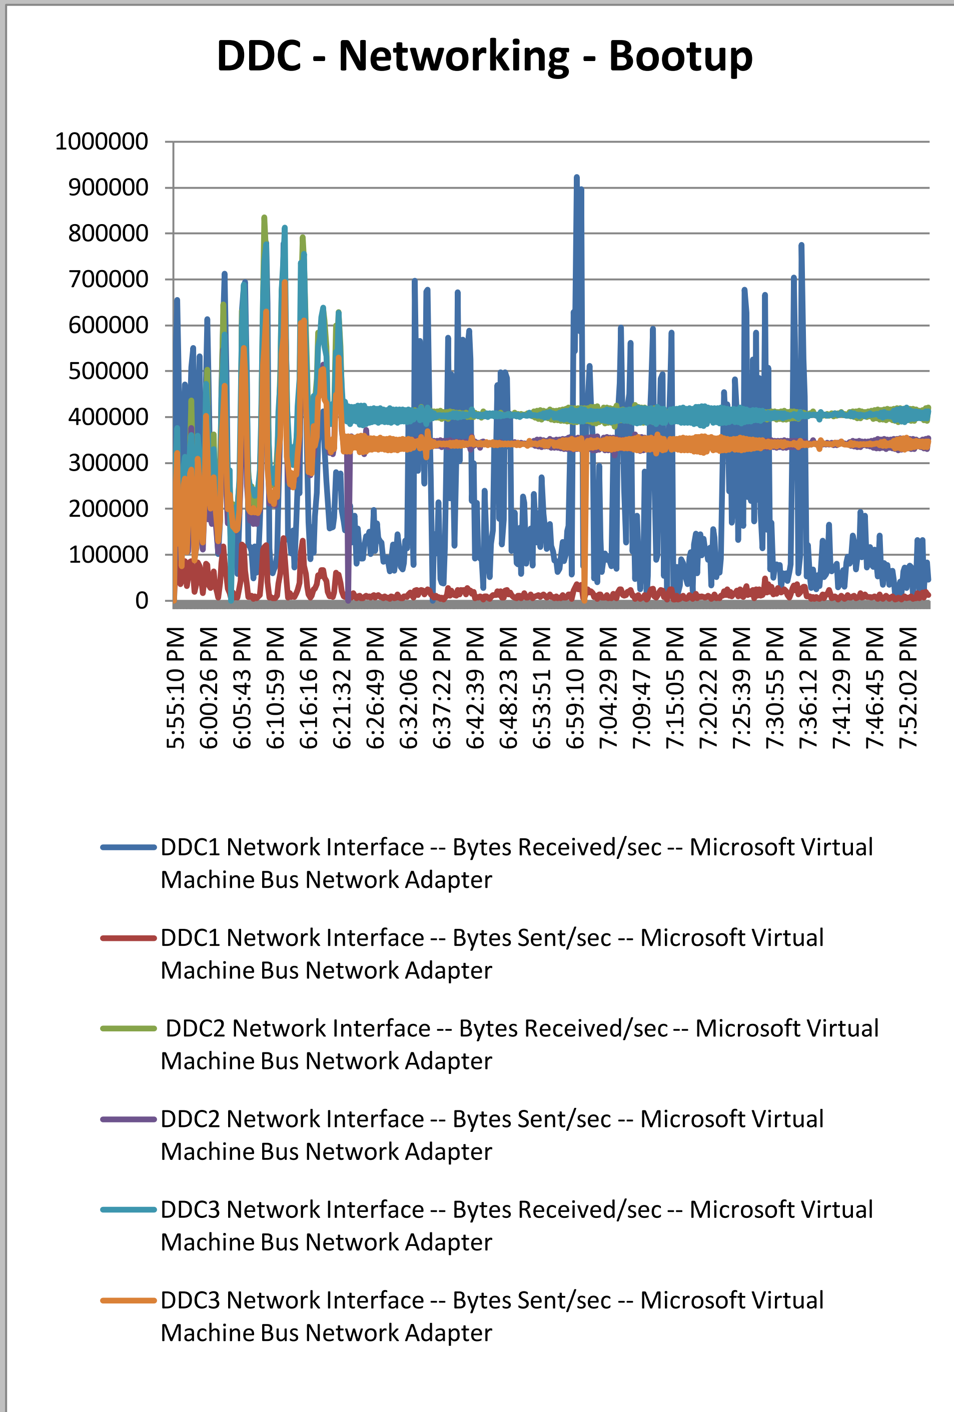

Analysis	Chart
Very similar to the earlier graph. 800000 bytes is about 6.10MB worth of traffic. This shows that the majority of networking traffic is around the communication with the VDA and XML resolutions. DDC1 (Farm Master) does not appear to be too busy on the network during the Login phase.	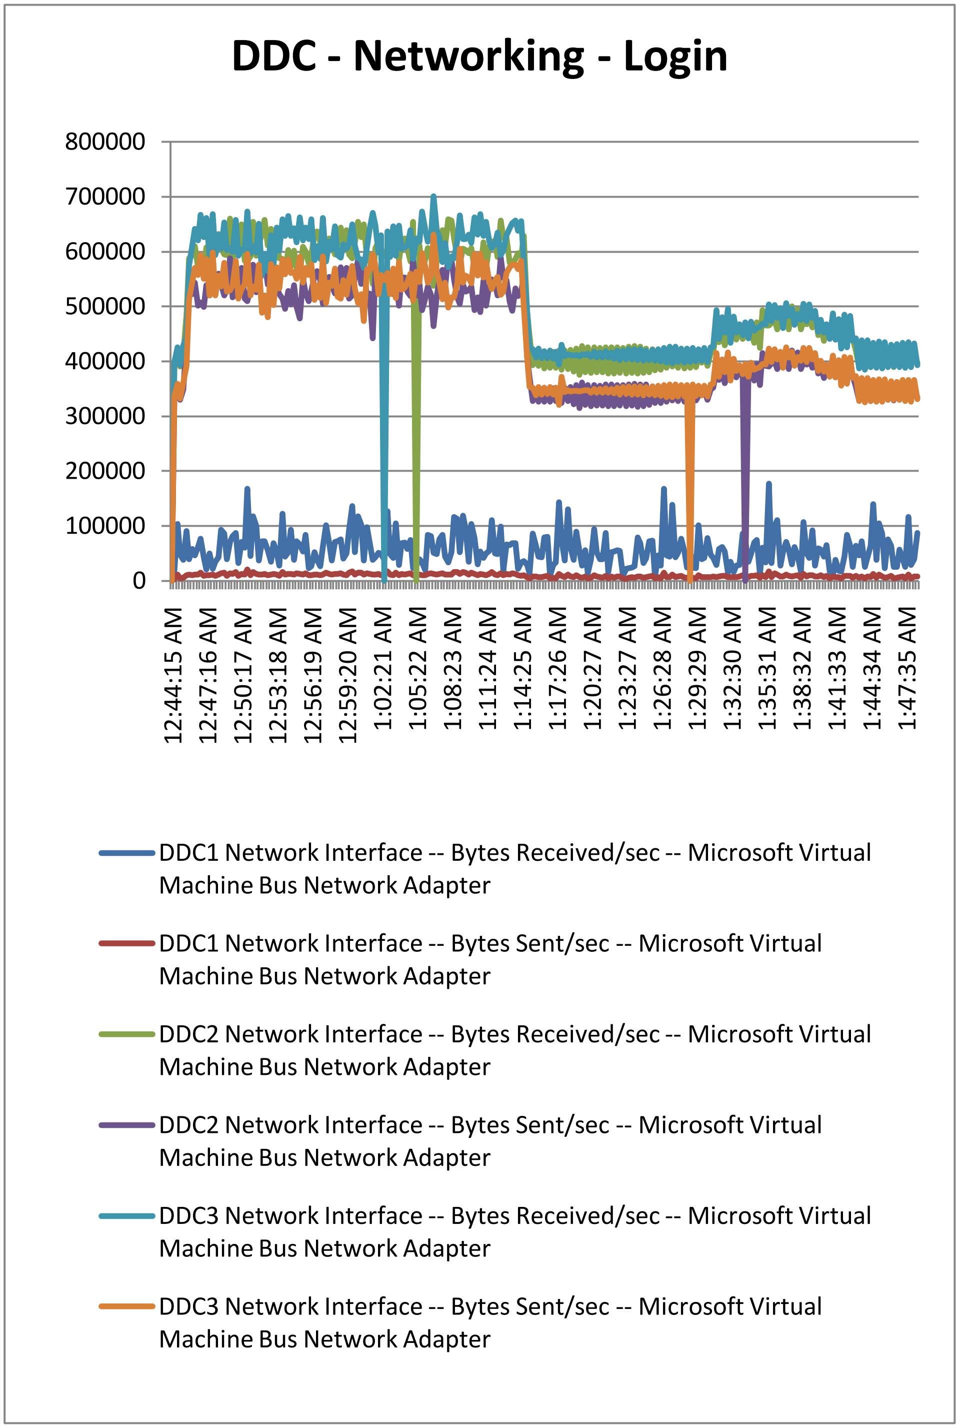

Analysis	Chart
All the DDCs had 8GB of RAM assigned. The graph indicates they could have been successful with only 4GB assigned to the VM.	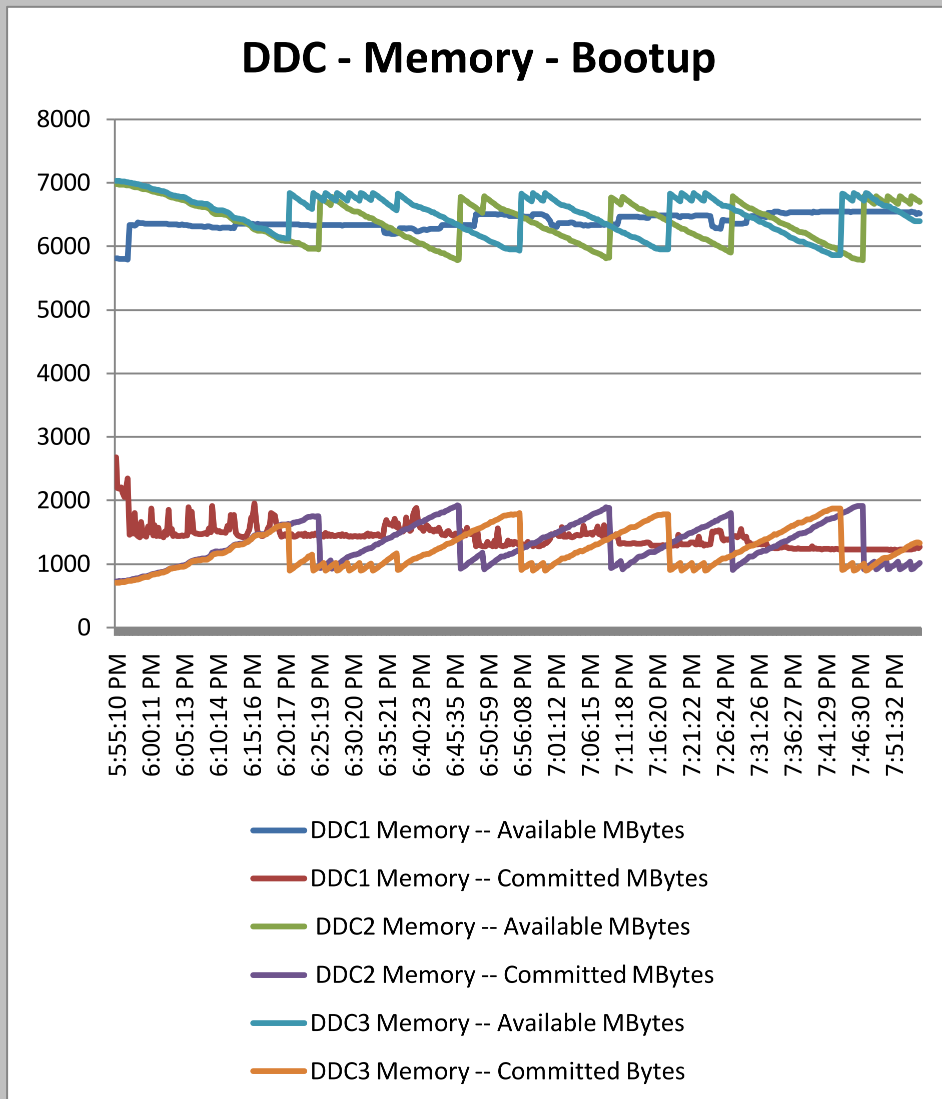

Analysis	Chart
All the DDCs had 8GB of RAM assigned. The graph indicates they could have been successful with only 4GB assigned to the VM.	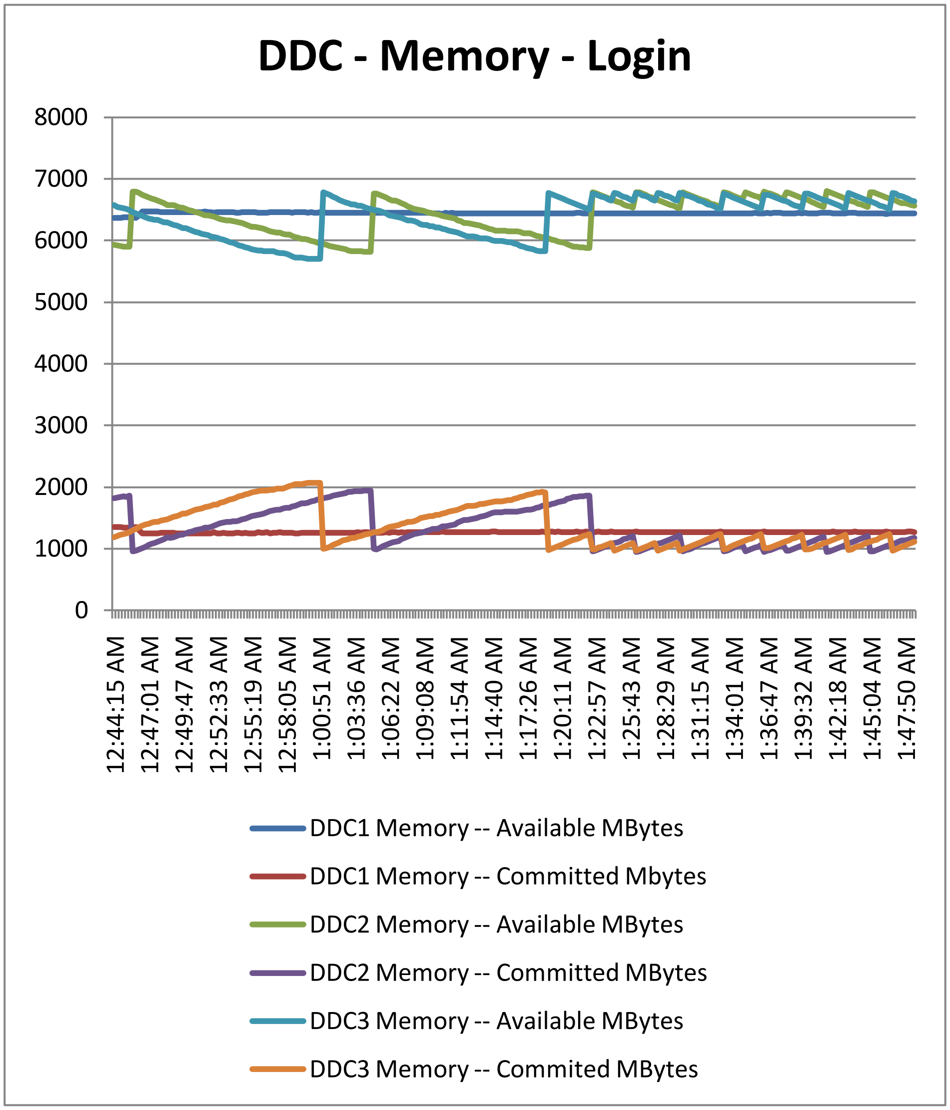

DDC Performance Final Analysis

From the DDC charts it looks like the DDCs were well positioned for the environment. The VMs could have probably run successfully on only 4GB of RAM. In general the DDC role is not very network or disk-intensive. The three DDCs configured with one farm/pool master and two member servers is the recommended design for a 3500-desktop deployment. Adding more desktops will most likely require a fourth DDC configured as a member server.

Provisioning Services Performance

The provisioning servers PROV1 – PROV3 were all configured as farm servers and were all included on the bootstrap server list. This configuration allowed all three to respond to any network boot requests.

Analysis	Chart
The Bootup phase is the most resource intensive phase for Provisioning Services. The chart shows that the servers had 50% of CPU cycles always available. What is remarkable is how evenly distributed the workload is across the three servers.	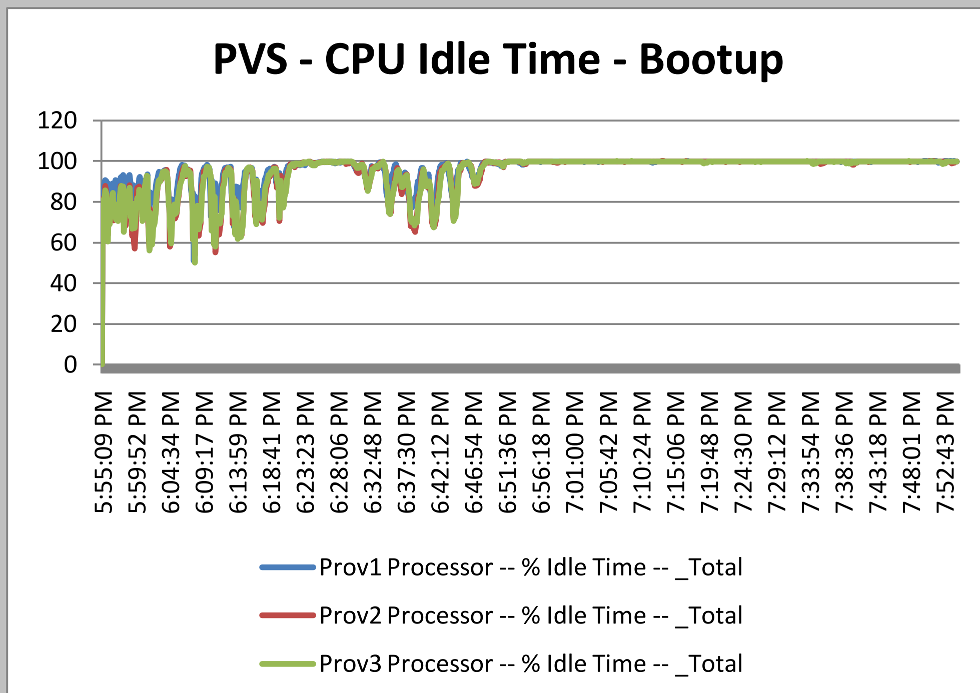

Analysis	Chart
The Login phase does little to affect the CPU usage on the PVS hosts, primarily because few additional files need to be streamed from the vDisk during login.	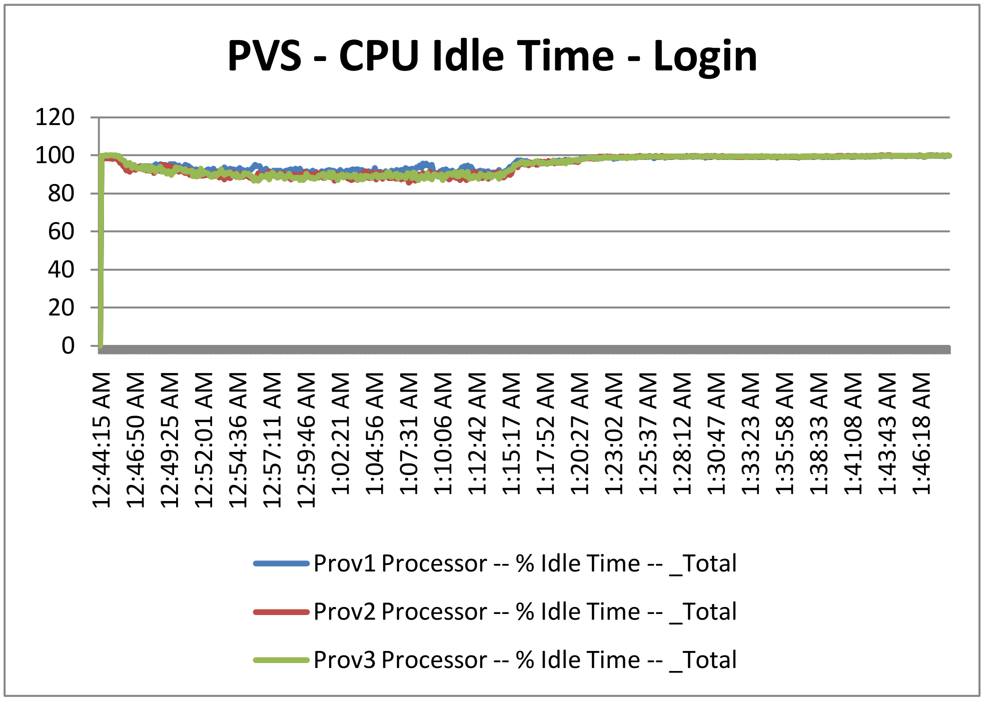
The context switches per second peak above the recommended limits which would indicate the processor was busy. The current theory is that the high context switches are related to the processor waiting on the slower network I/O traffic.	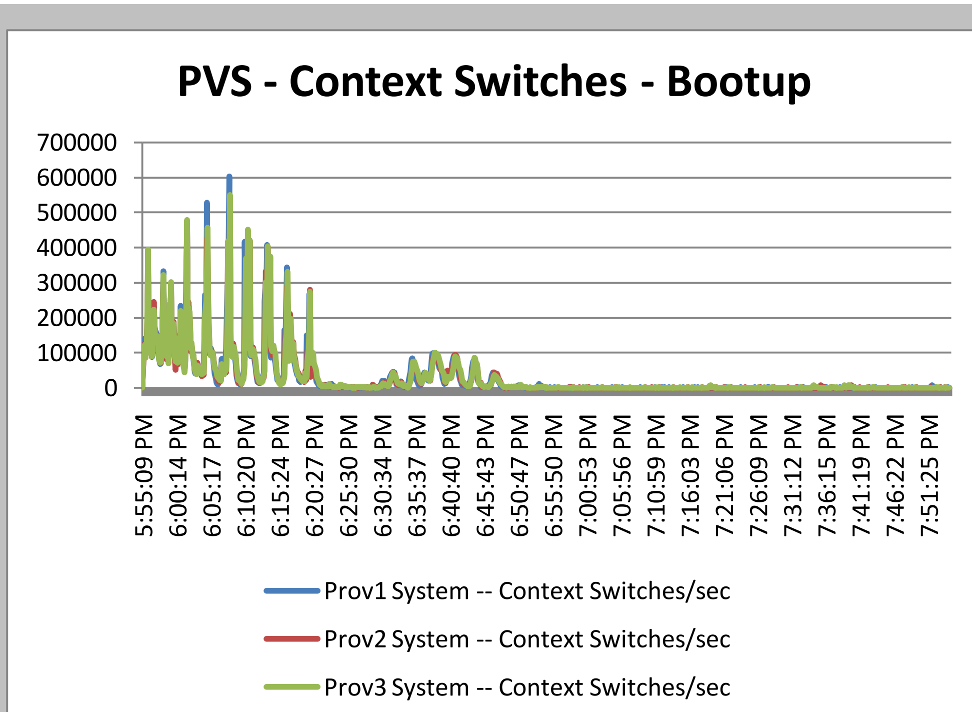

Analysis	Chart
The context switches per second are much lower with the Login phase and within normal operating parameters. This behavior is probably because the I/O streams are smaller and less intense during the Login phase.	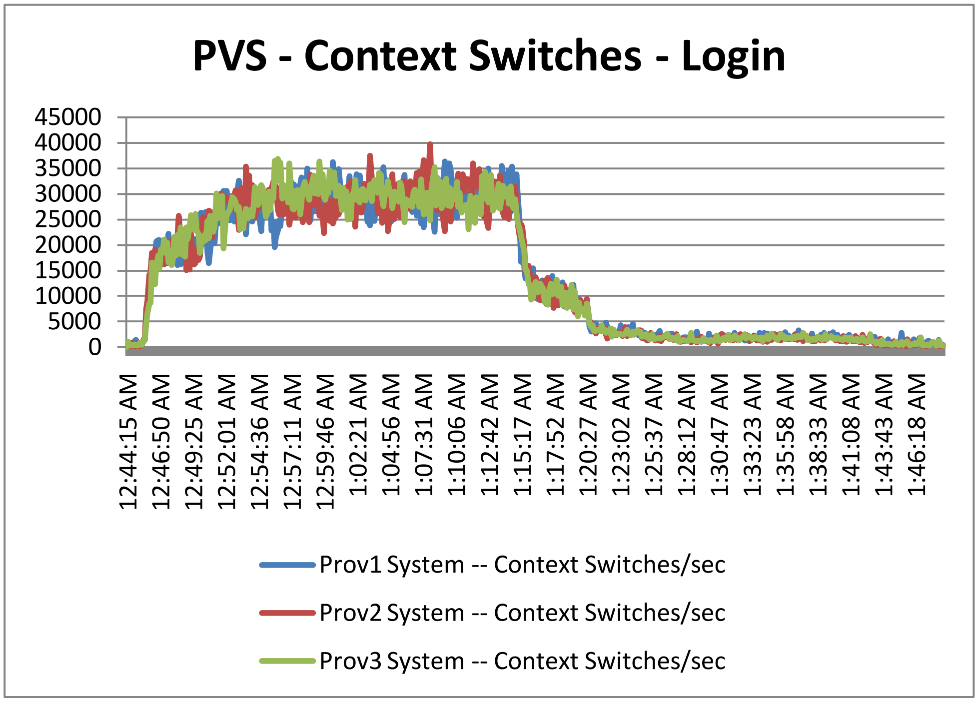
Surprisingly enough, even though the servers had 48GB RAM available the system only uses around 3GB of RAM. This amount of RAM is in line with earlier results which indicate PVS uses approximately 1GB of RAM for the OS and about 2GB of RAM per vDisk.	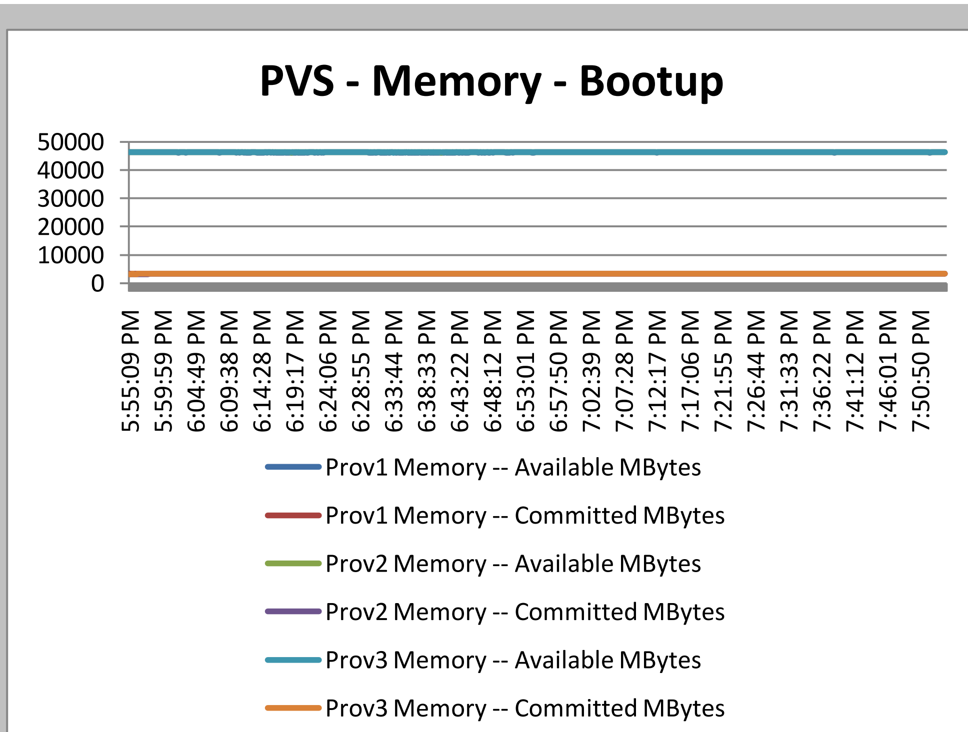

<table>
<tr><th>Analysis</th><th>Chart</th></tr>
<tr><td>This amount of RAM is in line with earlier results which indicate PVS uses approximately 1GB of RAM for the OS and about 2GB of RAM per vDisk.[11]</td><td>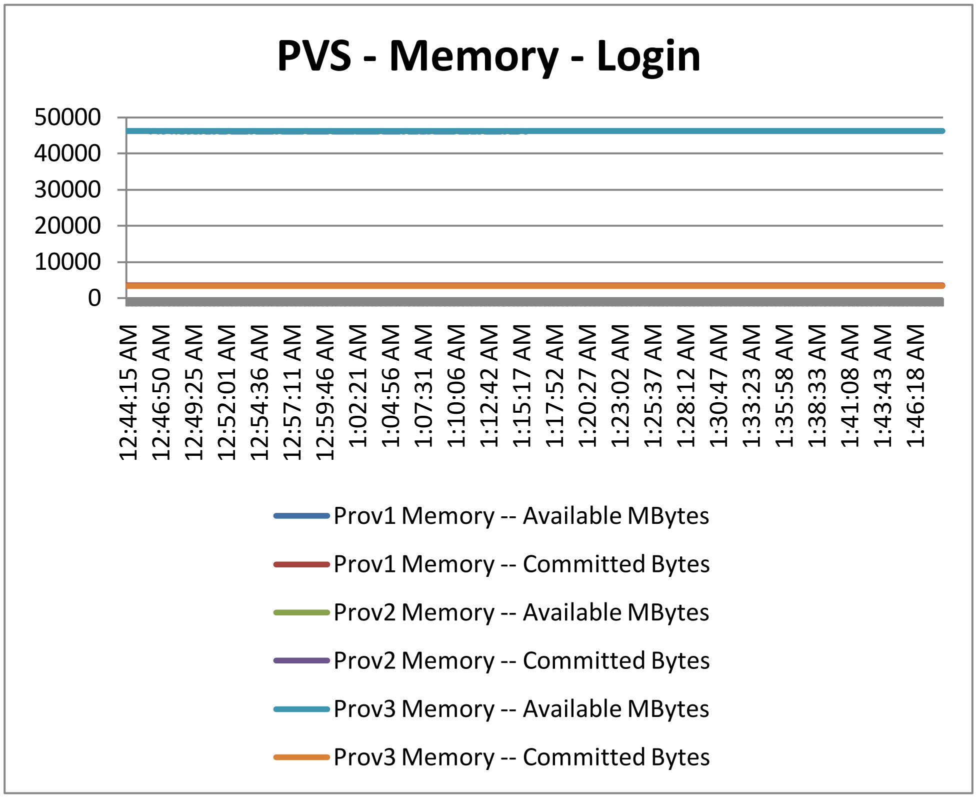
</td></tr>
</table>

11 See the Citrix KB article "Advanced Memory and Storage Considerations for Provisioning Services" (CTX125126) for more information on calculating RAM requirements for Provisioning Services.

Analysis | **Chart**

The Bootup phase will have the highest network activity as the vDisk image boot files are sent across the network. For reference, 400000000 represents about 3000 MBits of traffic. The peak usage represents about 2670MBits of the available 4000Mbits of teamed bandwidth.

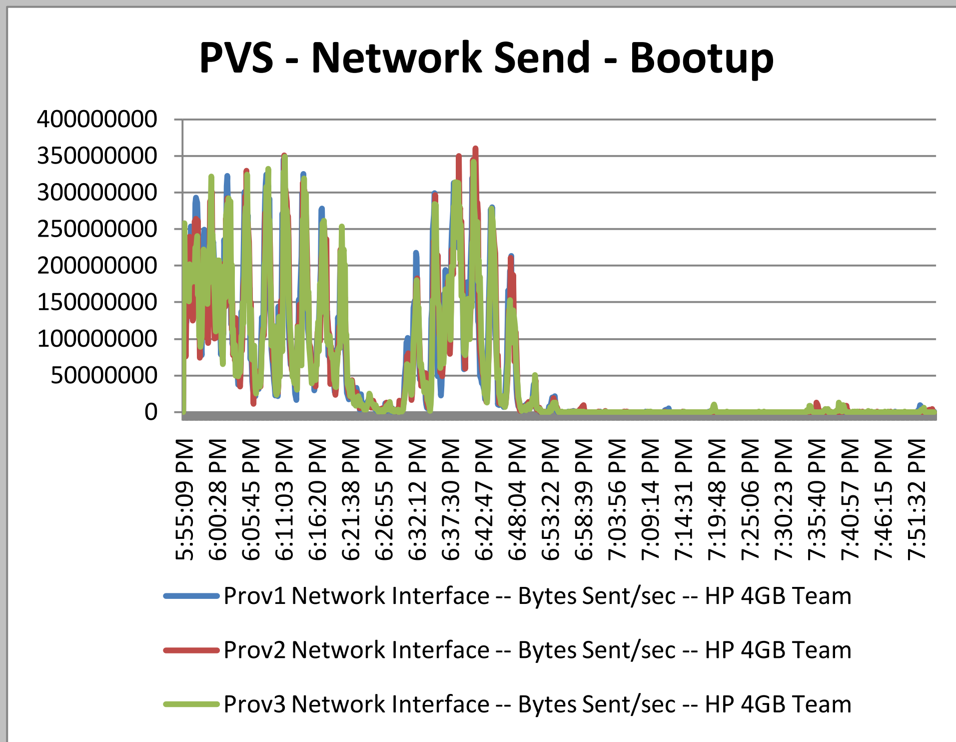

The top of this chart 250000000 represents about 1910MBits of traffic. The Login phase is clearly less demanding on the PVS servers from a network perspective. The peak in the graph is about 1600MBits.

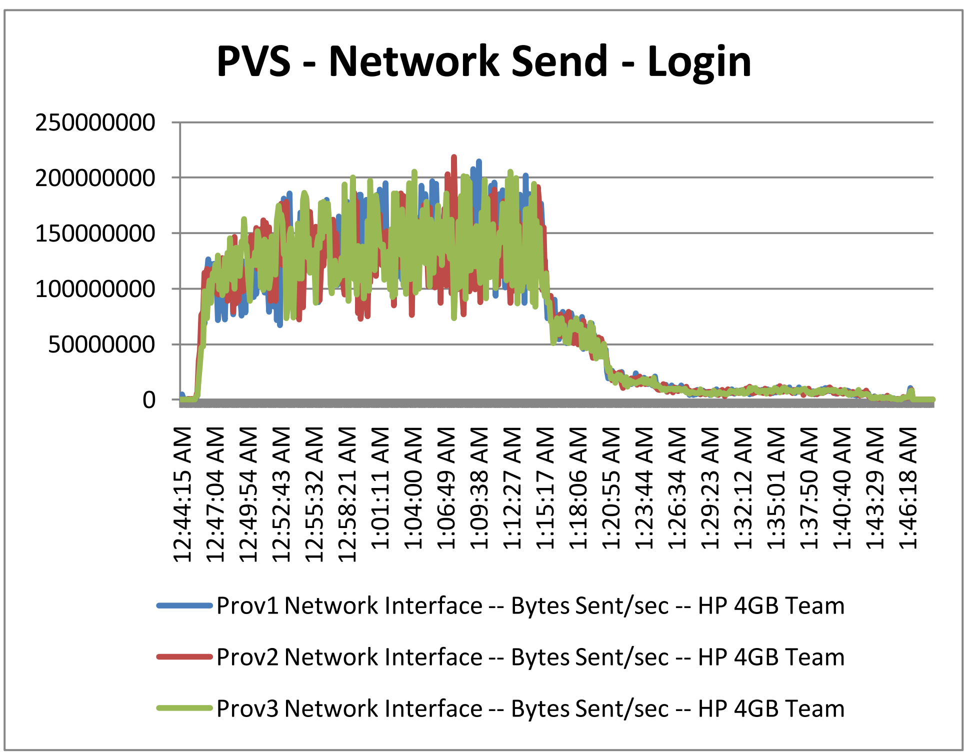

Analysis	Chart
The top of this chart represents about 91MBits of traffic.	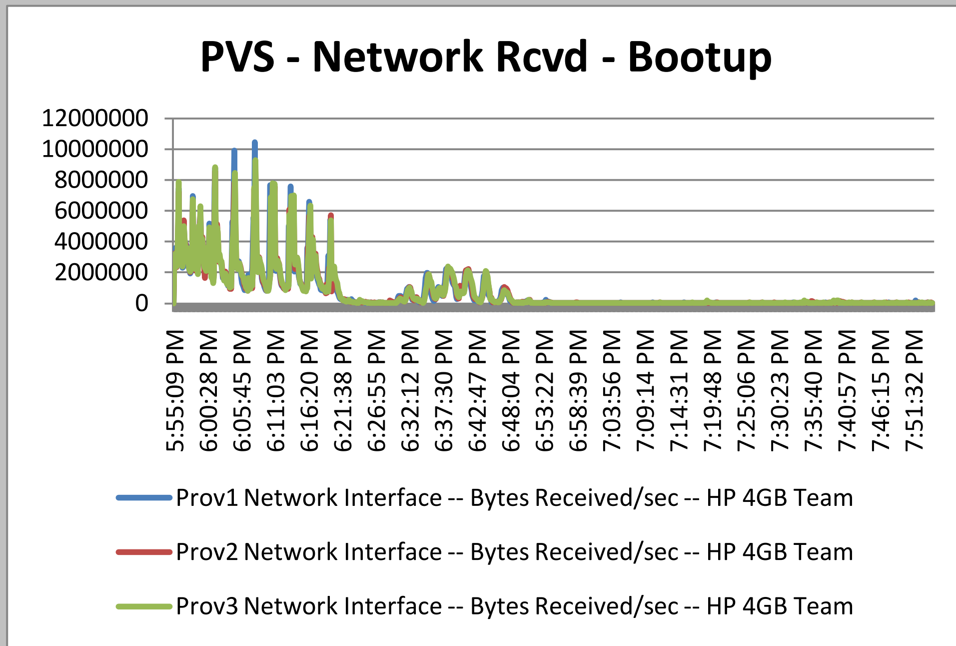
The top of this chart represents about 91MBits of traffic.	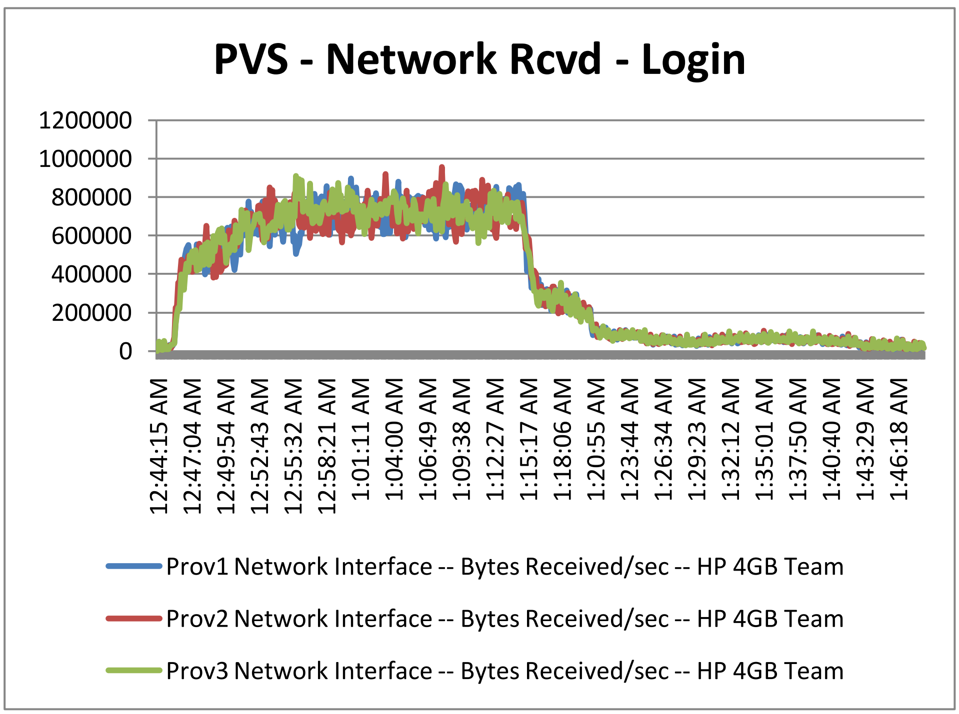

Provisioning Services Performance Final Analysis

As mentioned at the beginning of the section the three PVS servers were configured identically. And from the performance metrics recorded across all three servers it appears they perform at the same level. This trait indicates that the PVS farm (at least at this size) scales linearly and some additional observations can be made based on this assumption of equal processing/load across all the farm servers.

Some interesting facts regarding the data gathered during the 3500-desktop test, which actually had 3505 active desktops:

- Each PVS server managed about 1168 guests during the test.
- Each PVS server averaged 146 (3505/24) guests per physical core or 292 (3505/12) per core in use assuming the unused CPU capacity is adjusted out.
- Each 1Gb of network bandwidth supported approximately 450 (3505/(2670*3/1024)) guests for bootup and about 745 (3505/(1600*3/1024)) guests for login.

One final observation is that the CPU utilization indicates that the processors are likely to be the next constrained resource once the networking limits are removed. Using a TCP-Offload enabled network adapter or moving to a higher bandwidth adapter such as 10GbE may provide a significant increase in throughput and could possibly reduce the CPU utilization observed allowing for higher density.

System Center Virtual Machine Manager Performance

As related earlier five SCVMM servers were configured for the environment. Each was virtualized with 4vCPU and 4GB RAM per Microsoft's recommendation. This is a standard recommendation for the field when working with server virtualization, so this was the starting configuration for a desktop virtualization deployment.

Each chart in this section will contain the metrics from all five SCVMM servers. The metrics from the corresponding SQL servers are contained in the next section.

Analysis	Chart
This graph indicates the five servers were quite busy during the Bootup phase, with most servers having less than 30% available capacity on at least one occasion. Presumably this is the result of the large job queue that occurs during the Bootup phase as the DDC sends requests to boot up to 350 guests at once.	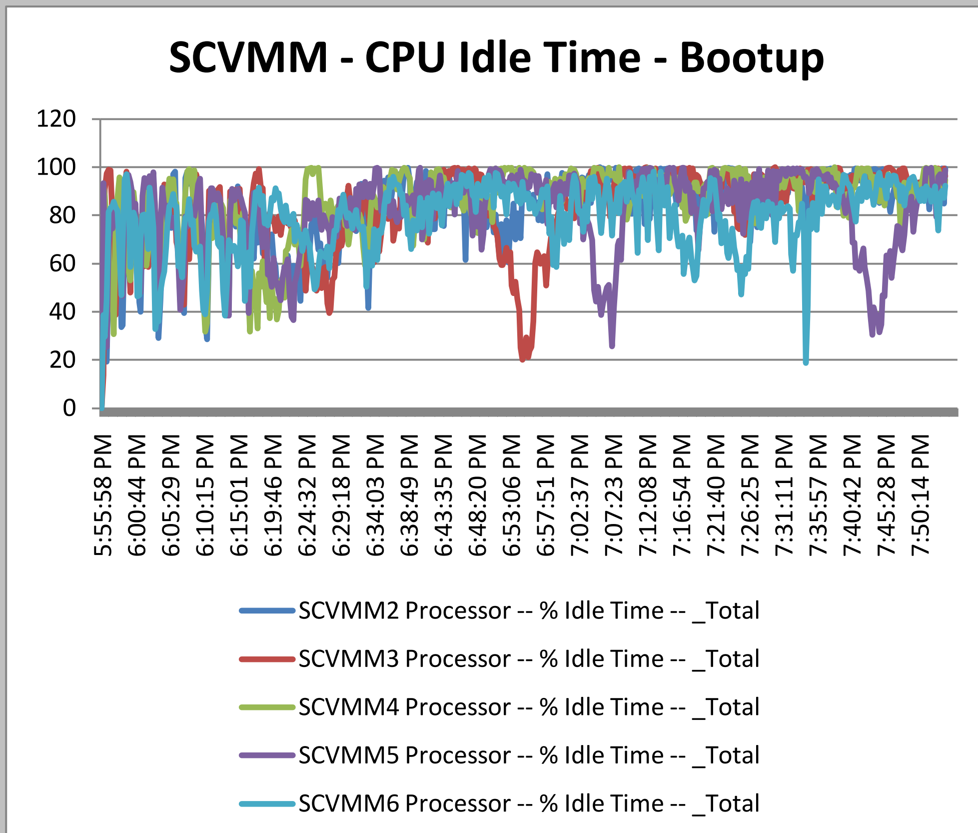

Analysis	Chart
During the Login phase the SCVMM servers have around 75% available capacity. Presumably this traffic is the normal status job checks.	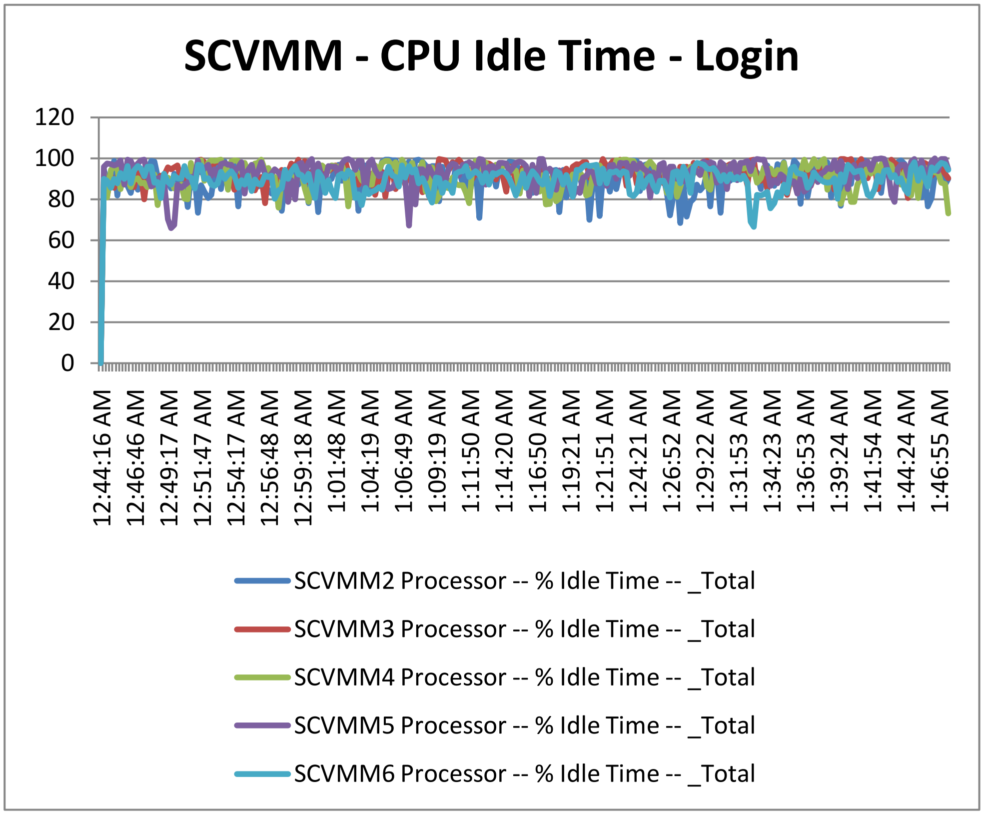

Analysis	Chart
This chart shows the majority of the processor cycles were being consumed by the VMM Service during the Bootup phase.	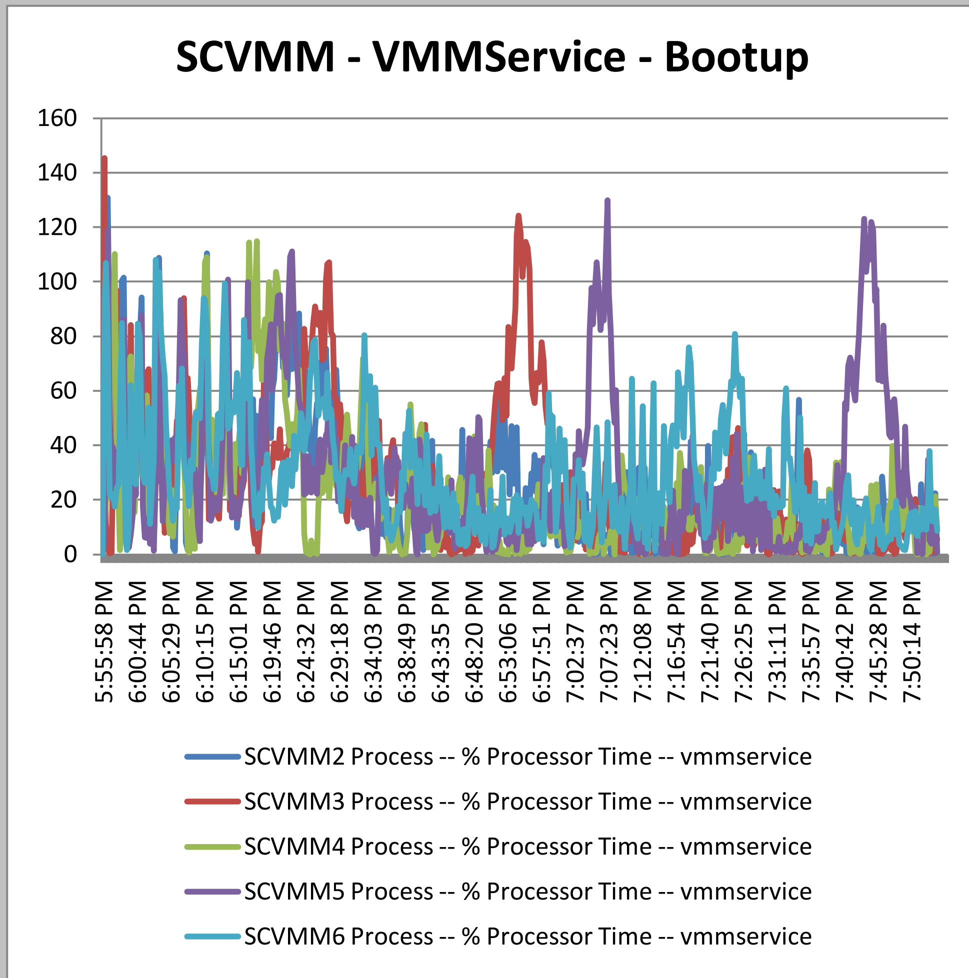

Analysis	Chart
The VMM Service is still quite active during the Login phase.	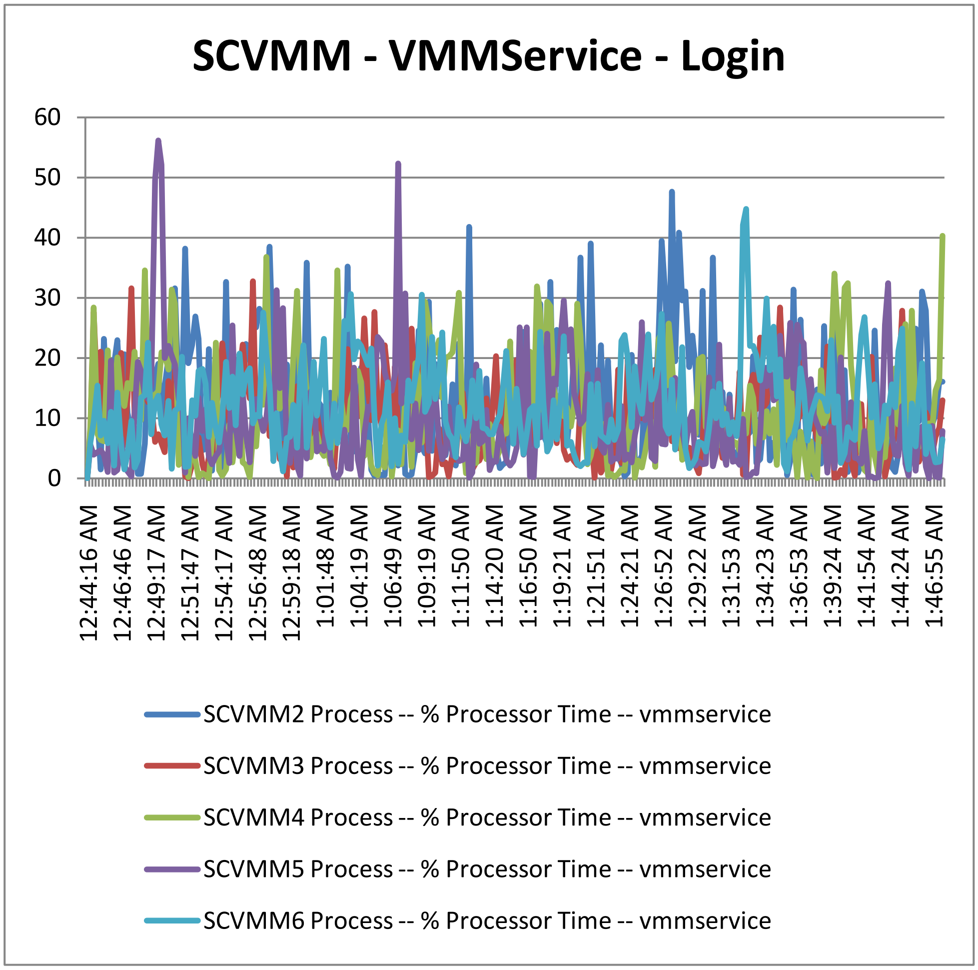

Analysis	Chart
The context switches per second are within the recommended range.	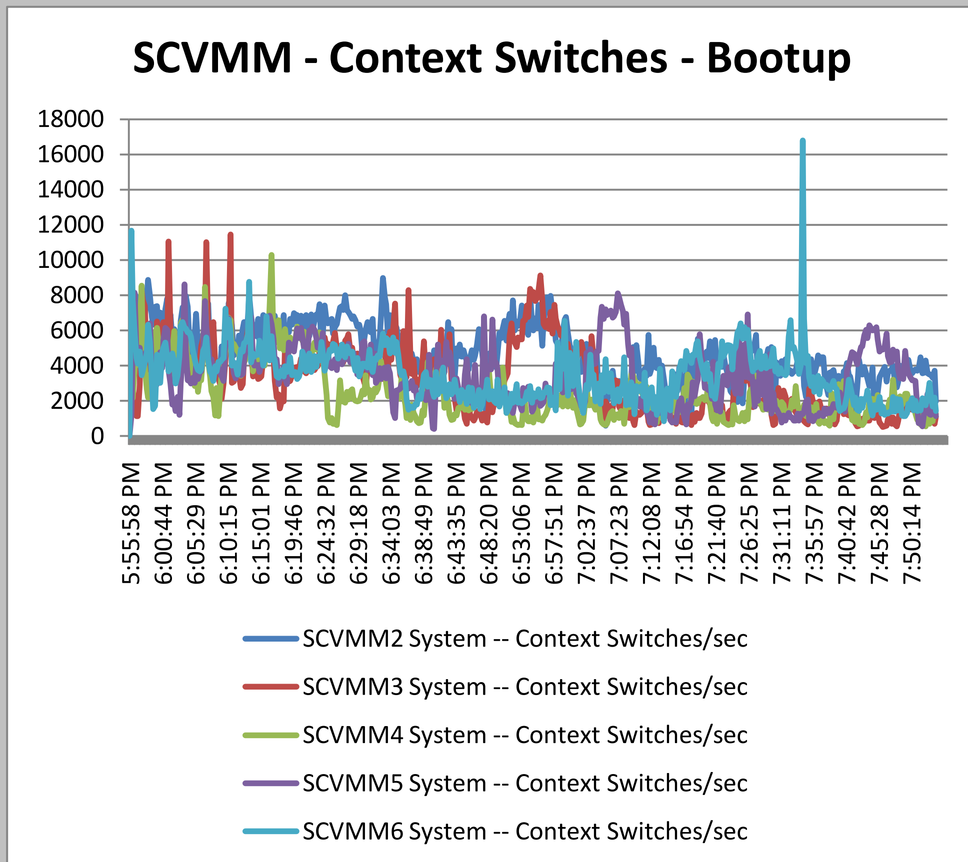

Analysis	Chart
The Login phase is less CPU intensive than the Bootup phase as expected.	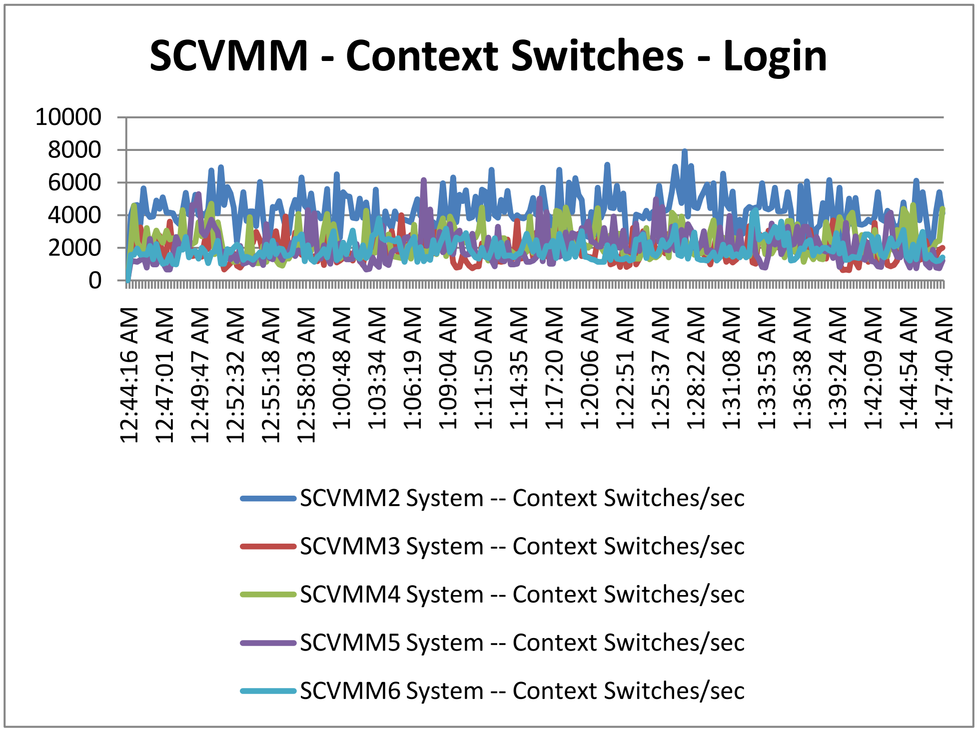
This chart indicates the SCVMM server is not disk intensive.	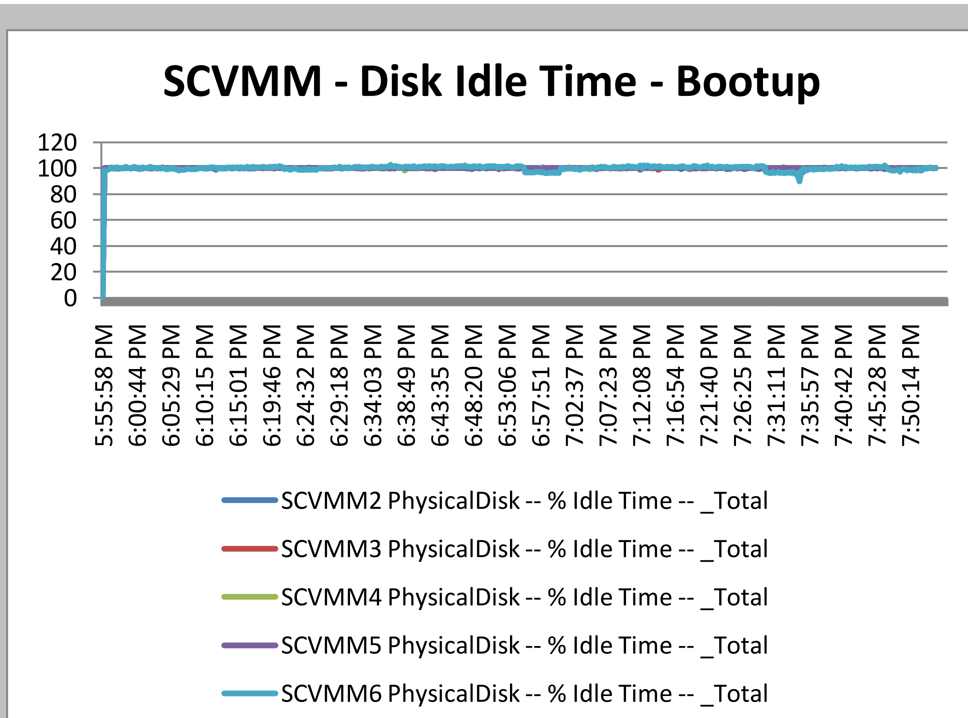

Analysis	Chart
This chart indicates the SCVMM server is not disk intensive.	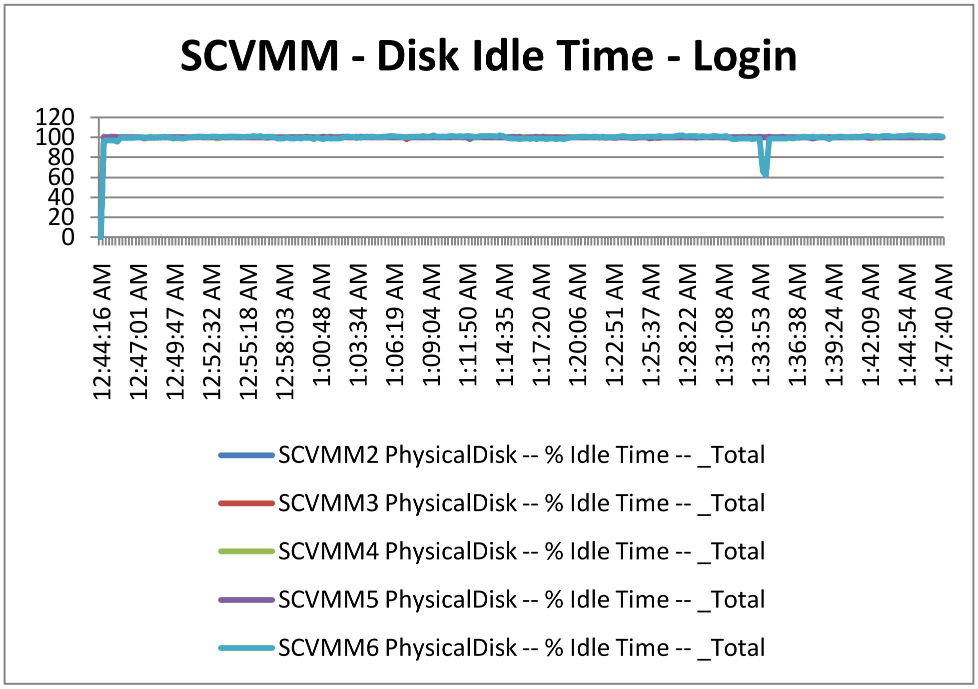

Analysis	Chart
The chart maximum of 8000000 represents about 61MBits of throughput.	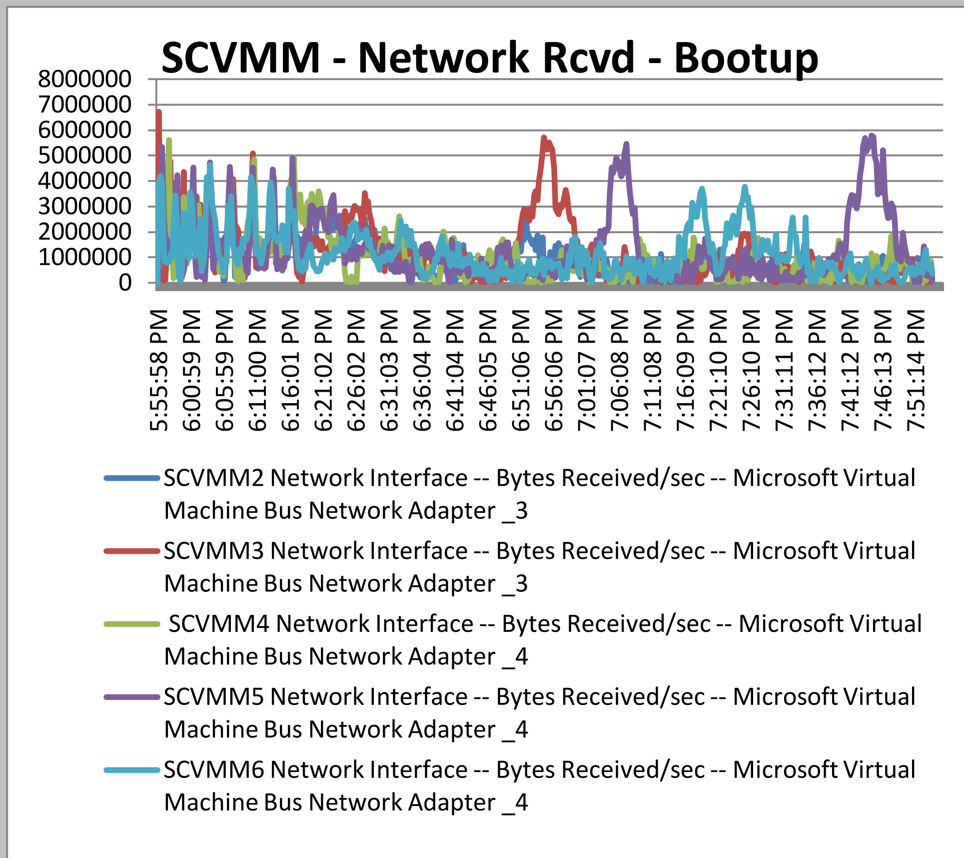

Analysis	Chart
The top of this chart 2500000 represents about 20MBits of traffic.	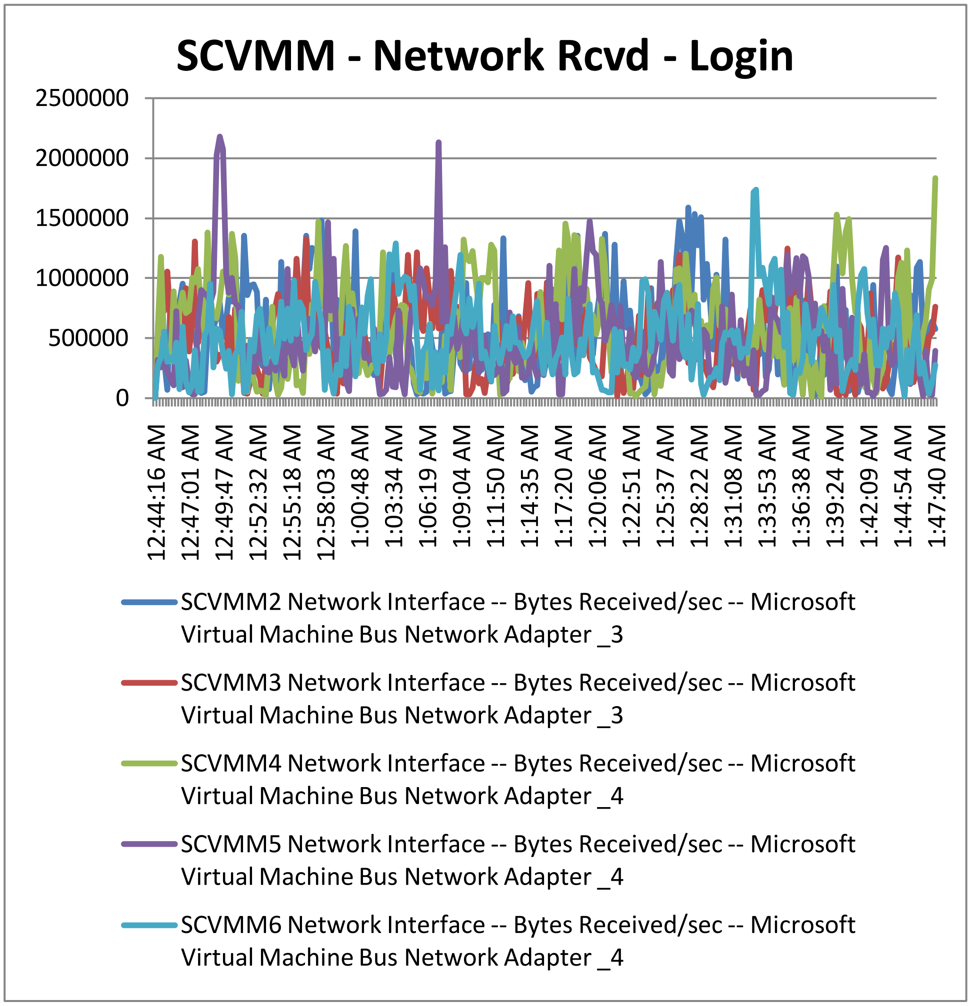

Analysis	Chart
Top of the chart 1400000 represents about 11MBits.	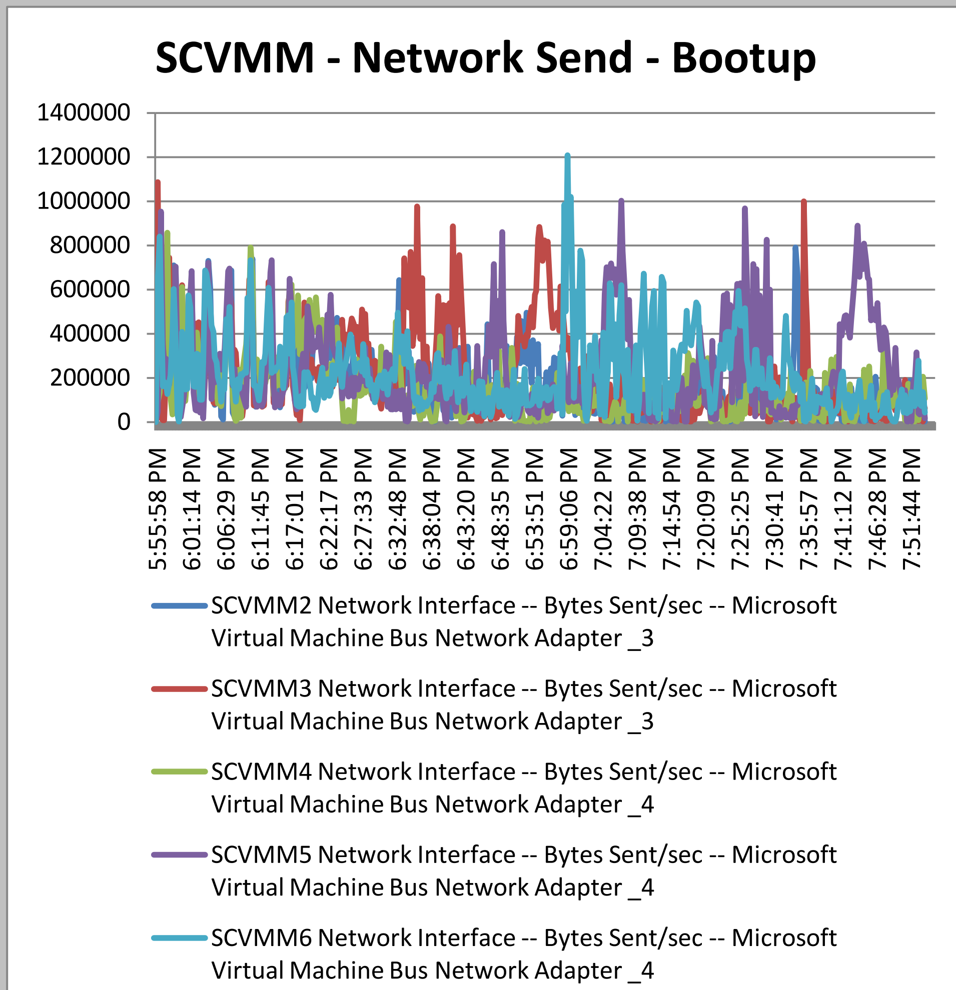

Analysis	Chart
Top of chart 500000 represents about 3.8MBits of traffic.	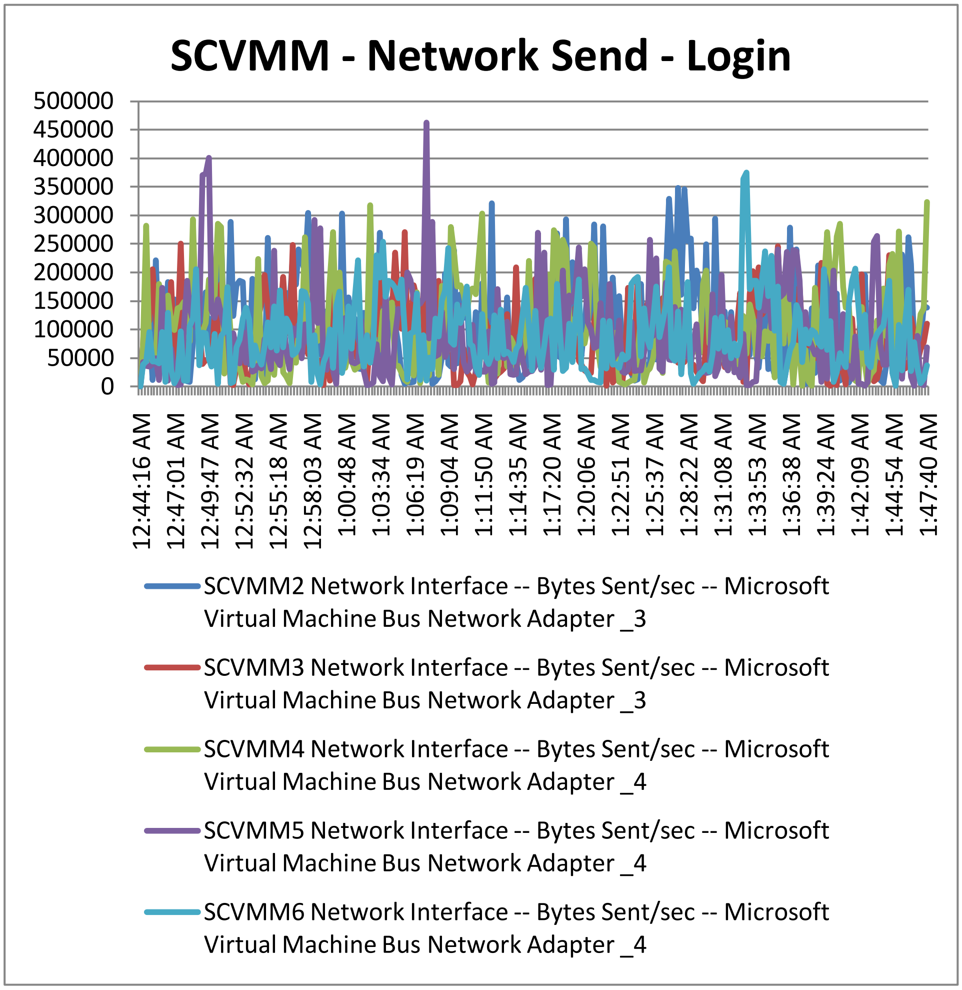

Analysis	Chart
This chart shows that most of the 4GB of RAM is available. Interestingly enough, it looks like near the end of the test the RAM is reclaimed and the available amount on three of the servers increases.	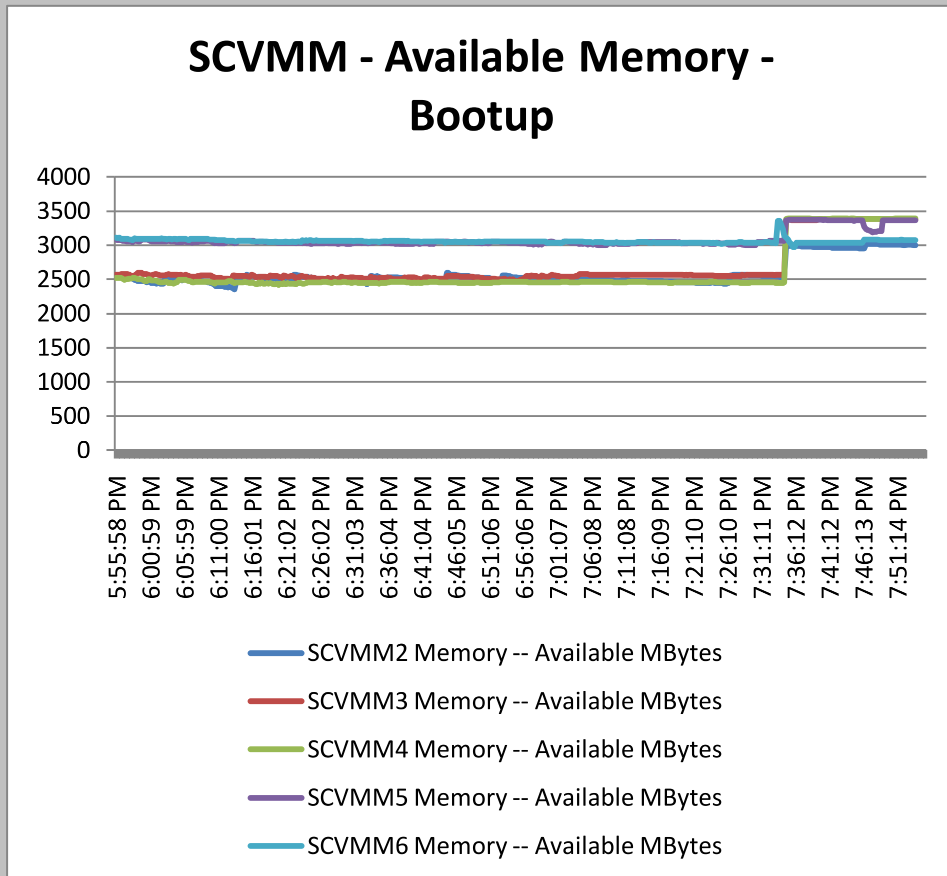

Analysis	Chart
Plenty of memory is available to the system during the Login phase.	

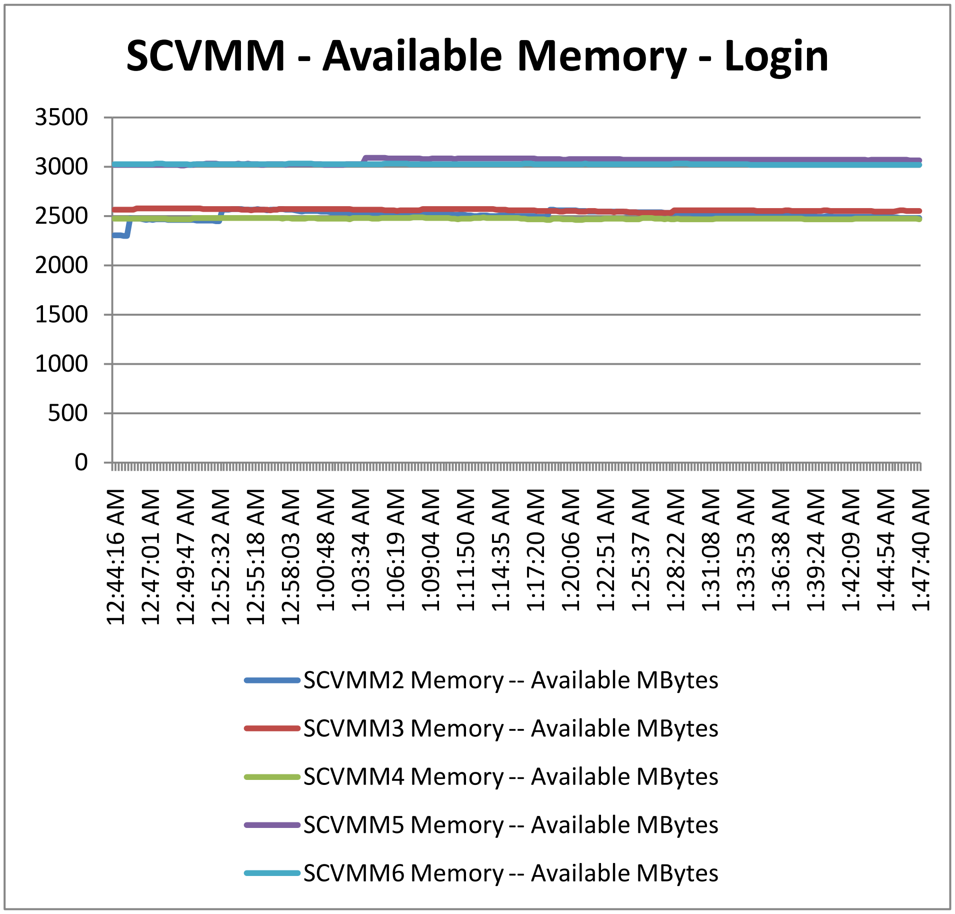

Analysis	Chart

This chart shows that for most of the Bootup phase, less than half of the system's 4GB of memory is in use.

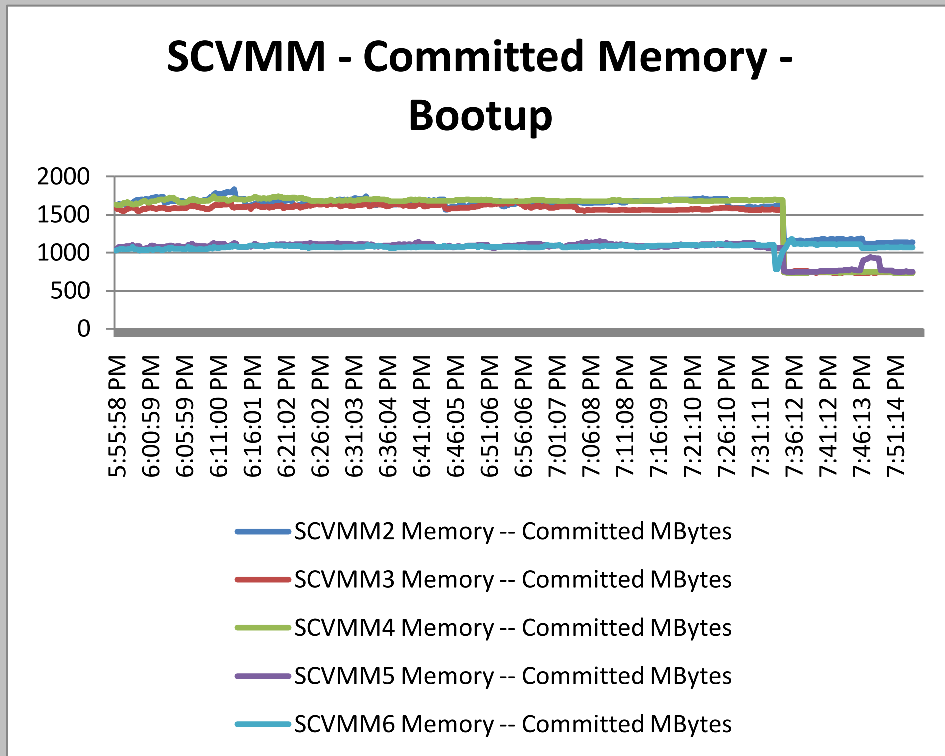

The memory requirements seem to be about the same for the Login and the Bootup phases.

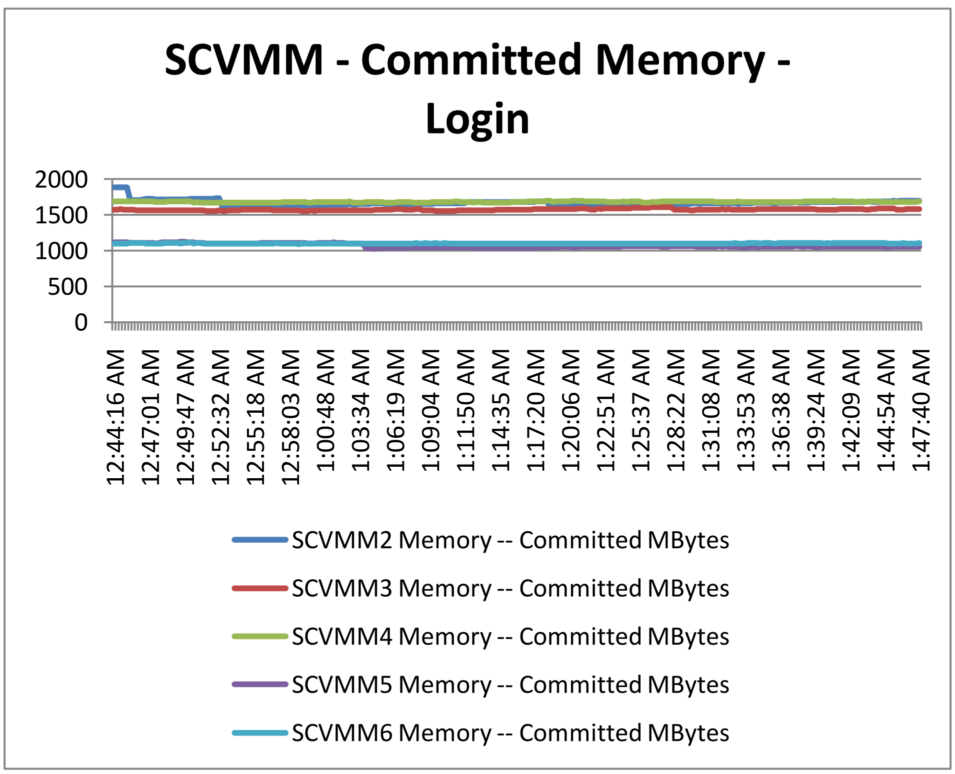

System Center Virtual Machine Manager Performance Final Analysis

The measured performance data indicates that the VMM servers are not network or disk intensive. Most of the resources needed by the SCVMM servers appear to be related to the processor with moderate memory usage. The 4GB of RAM assigned to the virtual machine is more than adequate to meet the needs at this level of use.

SQL Server Performance

As mentioned earlier, each SCVMM server had a dedicated virtualized SQL Server residing on the same physical host. The SQL Server had the same number as its paired SCVMM server, for example SCVMM2 was paired with SQL2. SQL1 was used to host the XenDesktop and Provisioning Services databases. For comparative purposes, the performance metrics from SQL2 – SQL6 are on a single chart since they are all performing the same role. SQL1 statistics are limited to the XenDesktop (XD4) and Provisioning Services (ProvisioningServices) databases since the majority of the processing and memory were used by the STAT test harness.

Analysis	Chart
Relatively speaking the XenDesktop and Provisioning Services database were not busy during the Bootup phase. The Provisioning Services database was the busiest since it was constantly updating the records in the database as each guest VM came online.	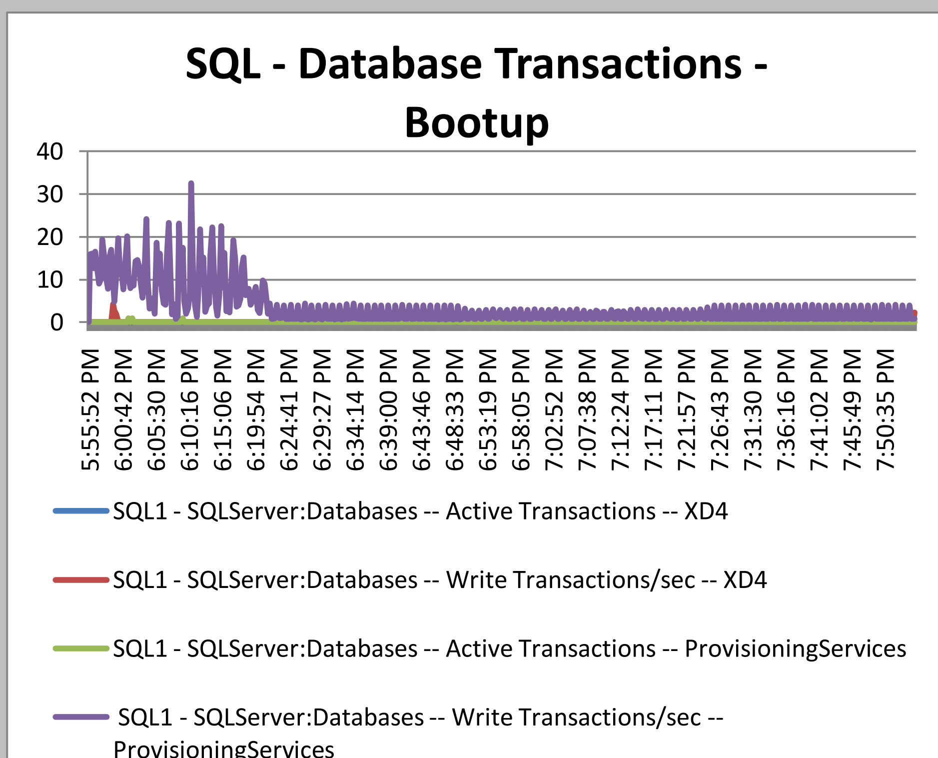

Analysis	Chart
The Login phase is not near as transaction intensive as the Bootup phase. Provisioning Services is still the most active database.	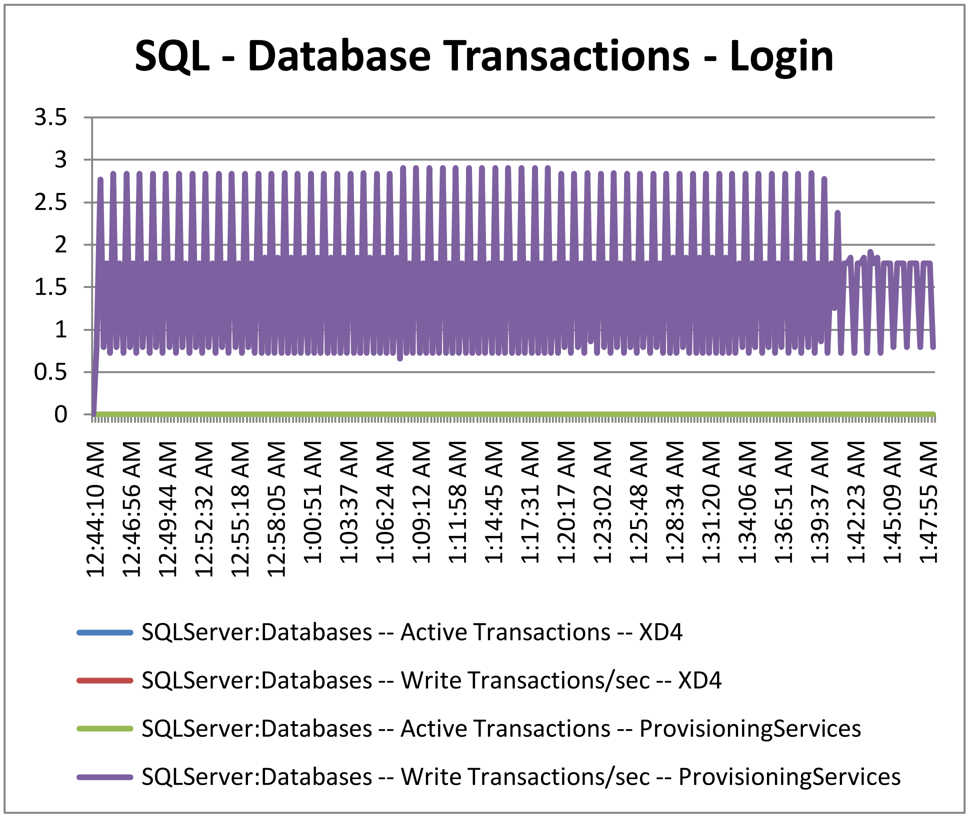

Analysis **Chart**

The SQL Servers had plenty of headroom for CPU utilization even though only 4 vCPUs were assigned to each database server.

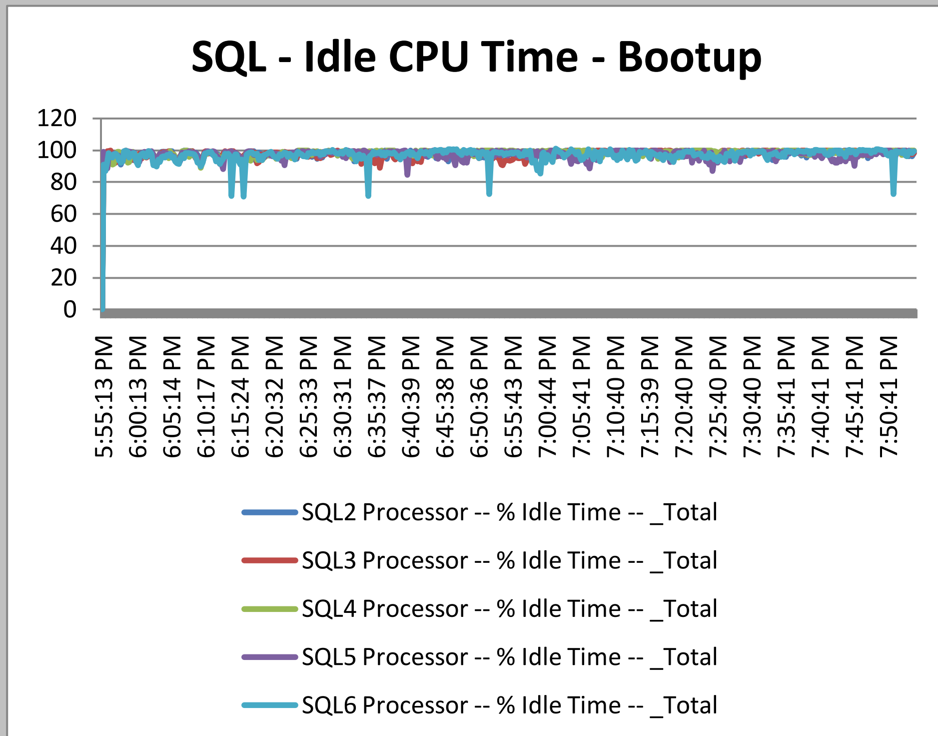

Still plenty of headroom during the Login phase. This would be expected since the server is only servicing a single SCVMM server and it is less busy during the Login phase.

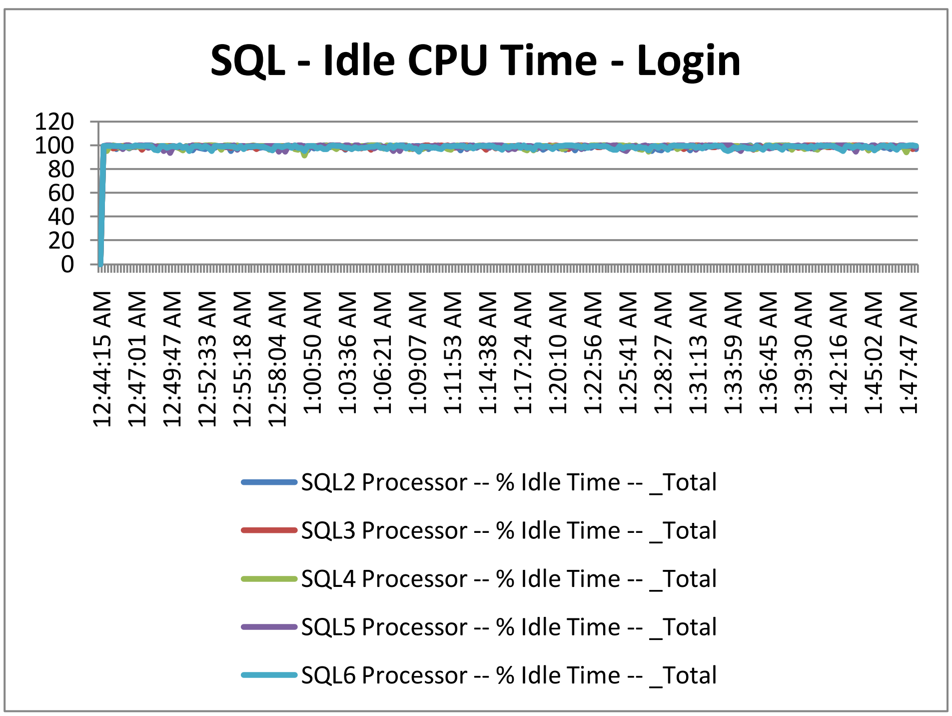

Analysis	Chart
One of the main concerns early on was the capacity of the disk for the SQL Server database since the server had only two local disks in a RAID1 configuration. This chart shows that for the majority of the time the disk was idle. Though not graphed, the disk queue length was normal.	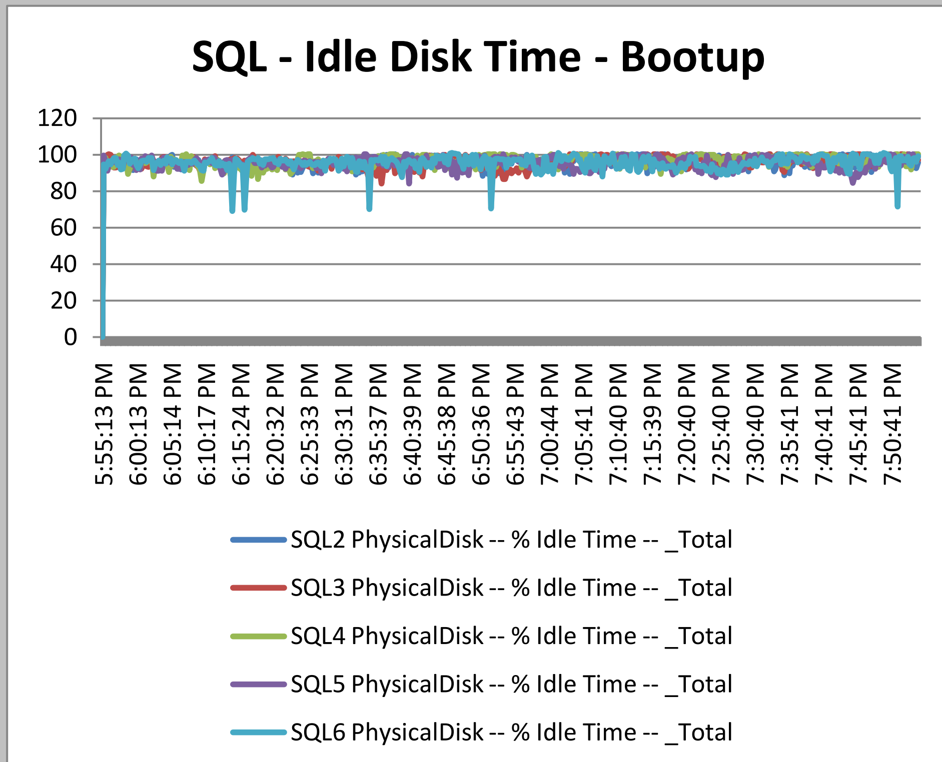

Analysis	Chart
Disks were mostly idle during the Login phase. Though not shown the disk queue length was normal.	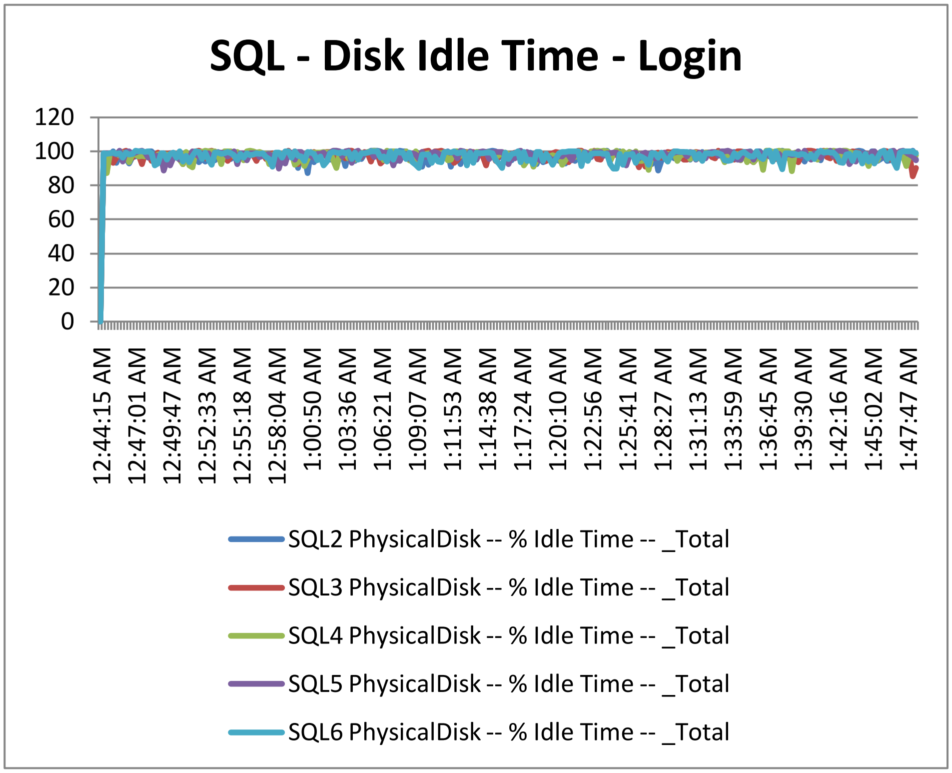

Analysis	Chart
Memory in use was always less than 3GB. Each server had 16GB installed.	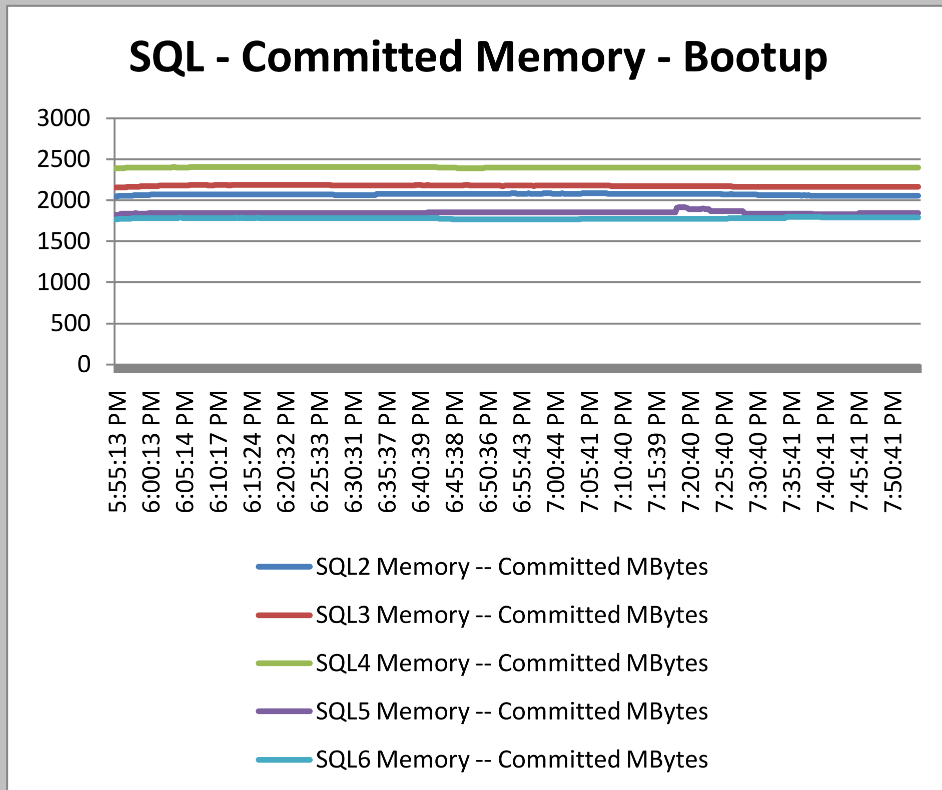

Analysis	Chart
Memory in use was always less than 3GB. Each server had 16GB installed.	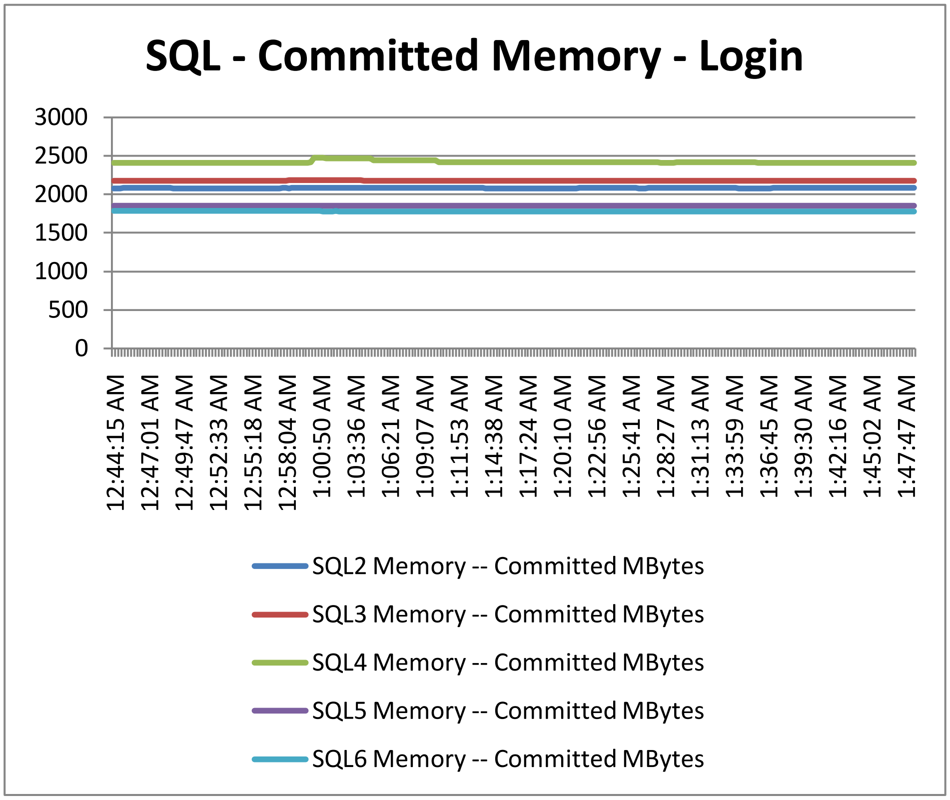

Analysis	Chart
The VMM Database reached slightly over 200 active transactions per second during the Bootup phase.	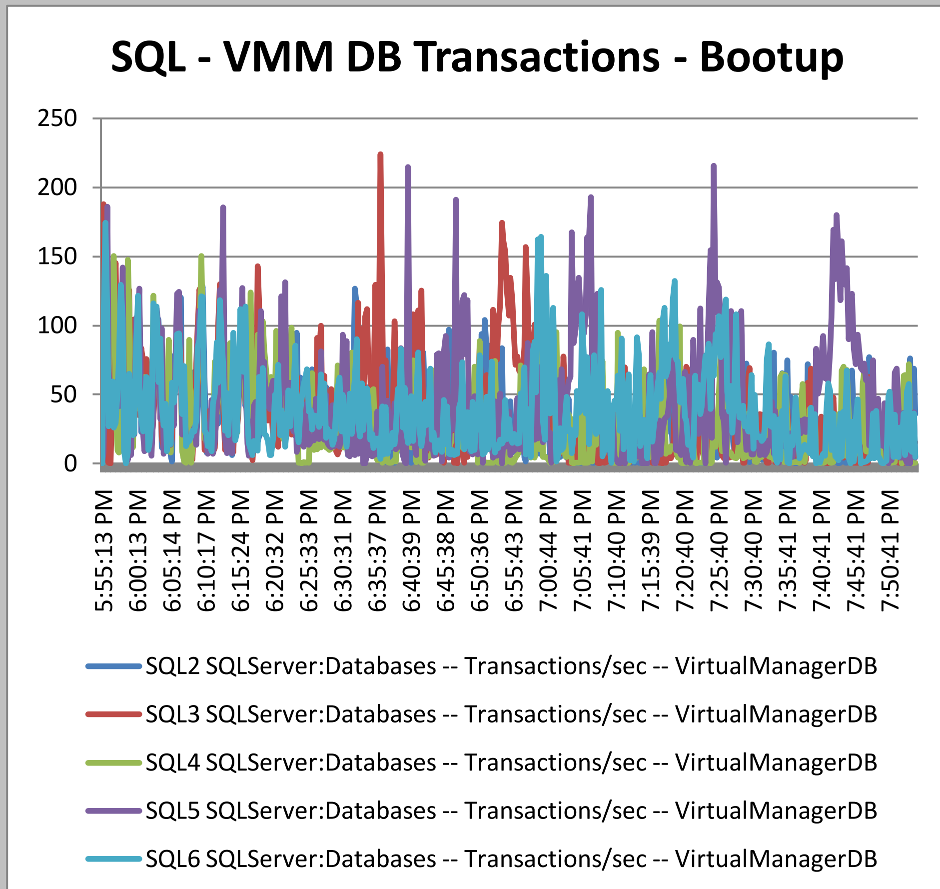

Analysis	Chart
Even during the less resource-intensive Login phase, the VMM database was reaching almost 120 transactions per second.	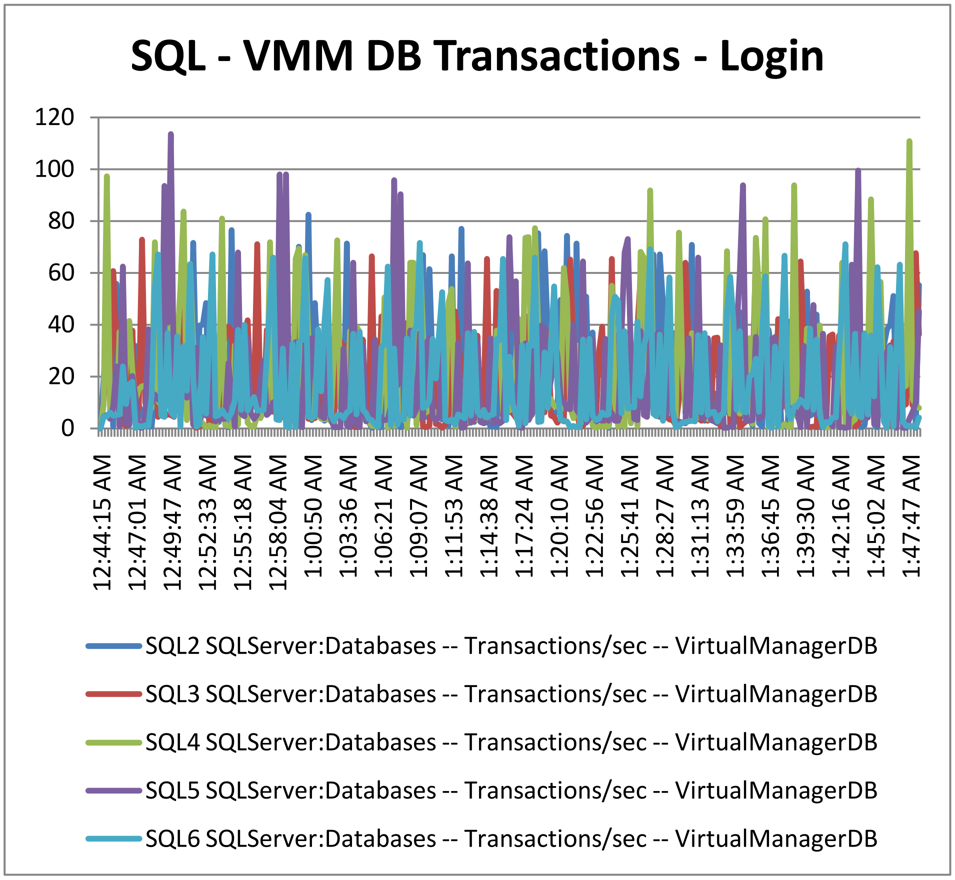

SQL Server Performance Final Analysis

The SQL Servers supporting the SCVMM environment were not bottlenecked on any of the primary resources. The SQL servers could have probably been assigned less RAM. With 16GB available, SQL server was using less than 3GB of it. The database servers did not appear to be stressed since SQL can easily handle the 200+ transactions per second that the VMM database required.

A single SQL Server of sufficient capacity could have handled the load generated by the five SCVMM servers.

HP P4500 (LeftHand) Storage Performance

This section covers the detailed performance metrics gathered during the different test runs. Due to a small clerical error, the performance statistics for the 700-desktop run were accidentally discontinued during the run. However, data for the other runs has been included here for analysis.

Analysis	Chart
Total IOPS between the two clusters was peaking around 28000 for the Login phase. The graph shows an increment of about 7000 IOPS for each 700 users above 1400 during the first 30 minutes of the run. This behavior shows the linear scaling provided by the SAN.	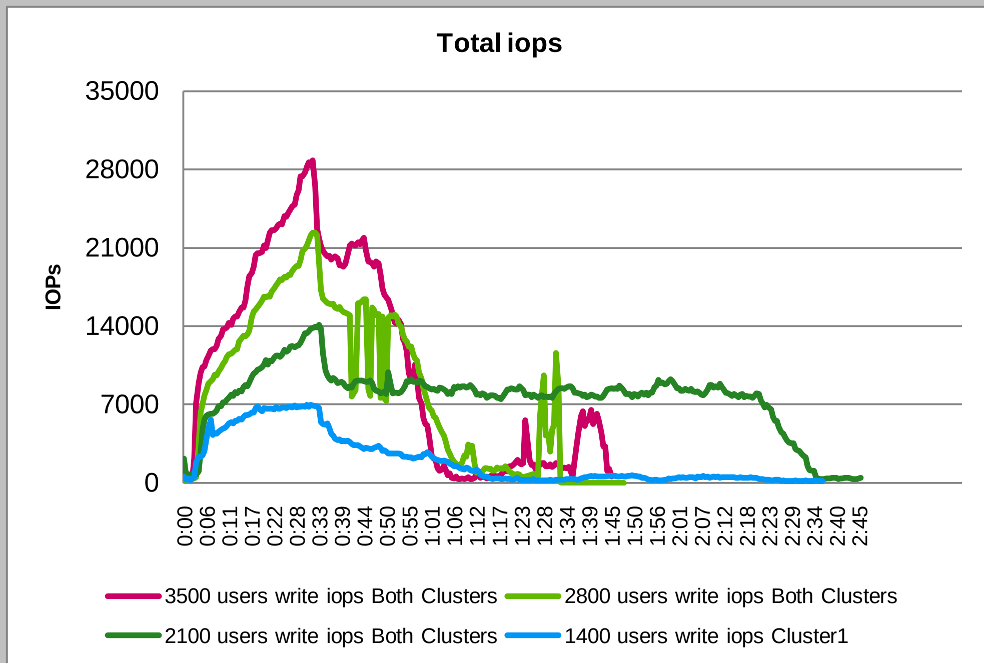
The cluster nodes were not exceeding 30% utilization. The implementation of MPIO drivers also prevented one node from being constantly bombarded with data lookup requests.	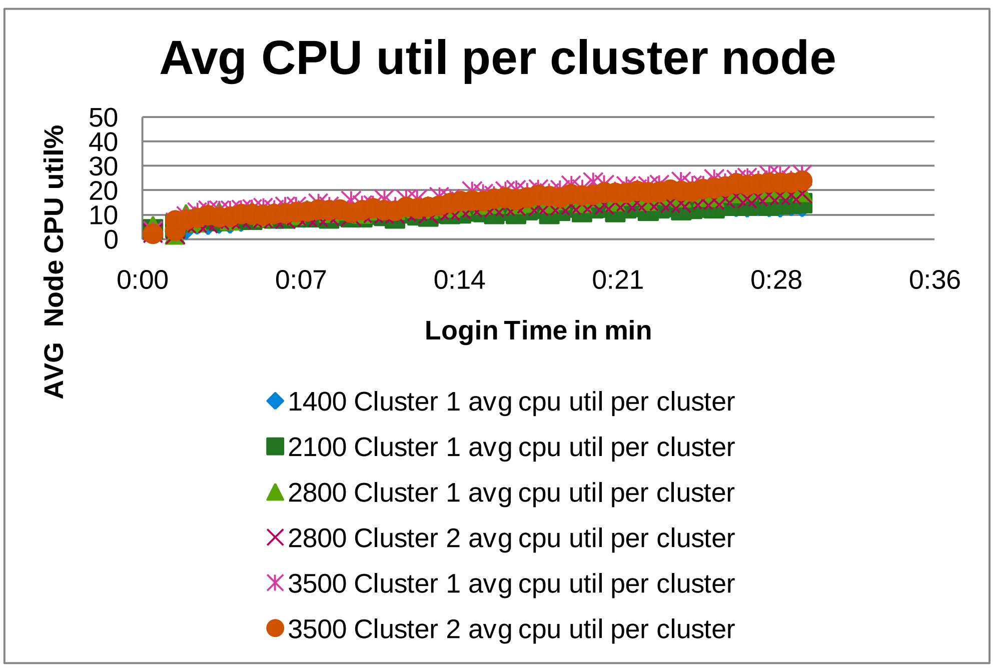

This chart shows the write latency in milliseconds. The cluster was handling the latency respectably with a peak around 5ms.

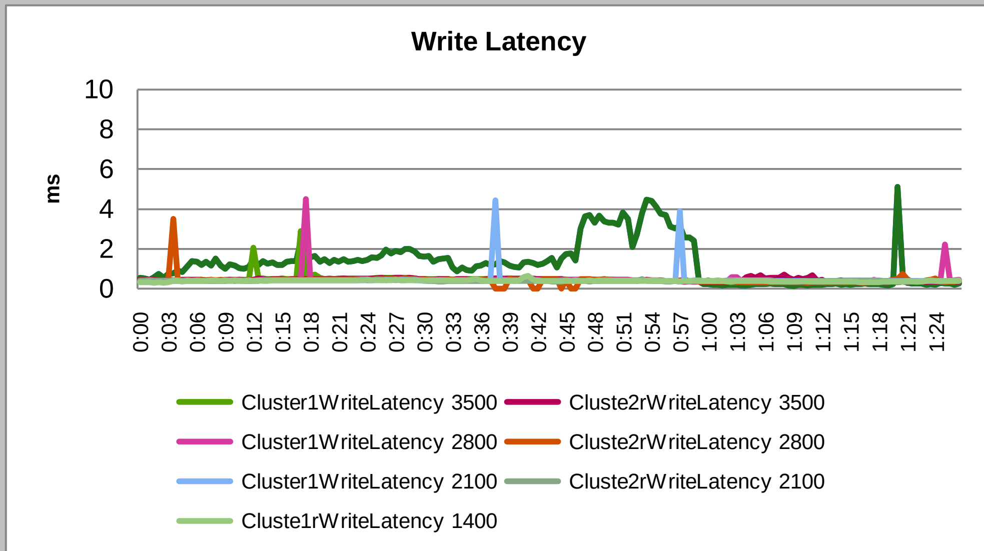

The read latency was also well within expected response times, although surprisingly higher than expected.

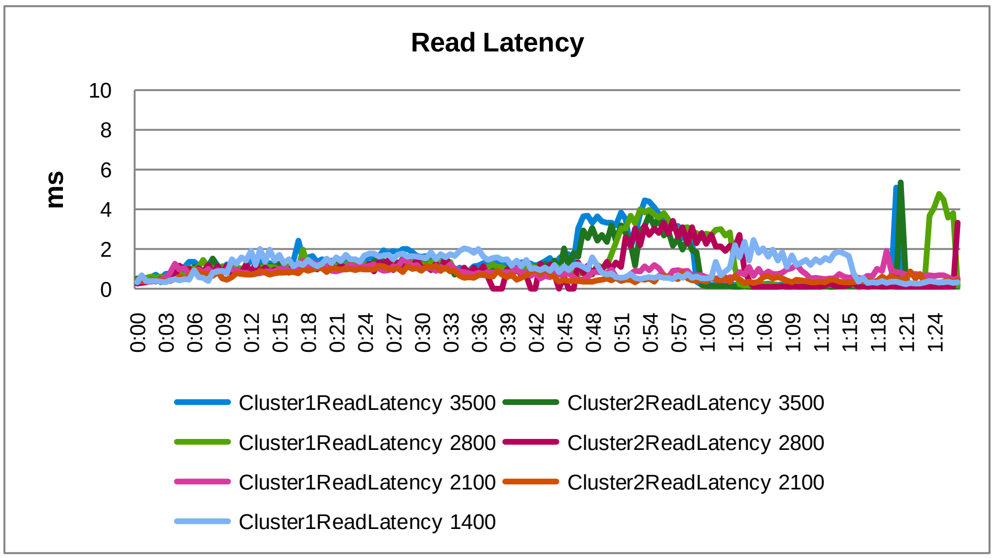

HP P4500 Performance Final Analysis

As mentioned earlier, the P4500 G2 SAN was configured with both Network RAID 10 and Hardware RAID 10 for optimum performance. The storage performed well during the 3500 test. In fact, projections based on the trending show that this storage unit as configured is most likely capable of supporting a higher number of desktops.

Figure 5: P4500 CPU Utilization Projections provides a linear approximation of the CPU utilization for a single cluster based on the data gathered during the testing. The graph depicts the logging time ramp up and the CPU utilization observed during that period for the 2100, 2800, and 3500 desktop test runs. Also shown on the graph is the expected CPU utilization rates at 4200, 4900, and 5600 desktops based on a linear projection.

The projection confirms that P4500 as configured could confidently support more than 5600 concurrent users and still maintain the response times under 5ms.

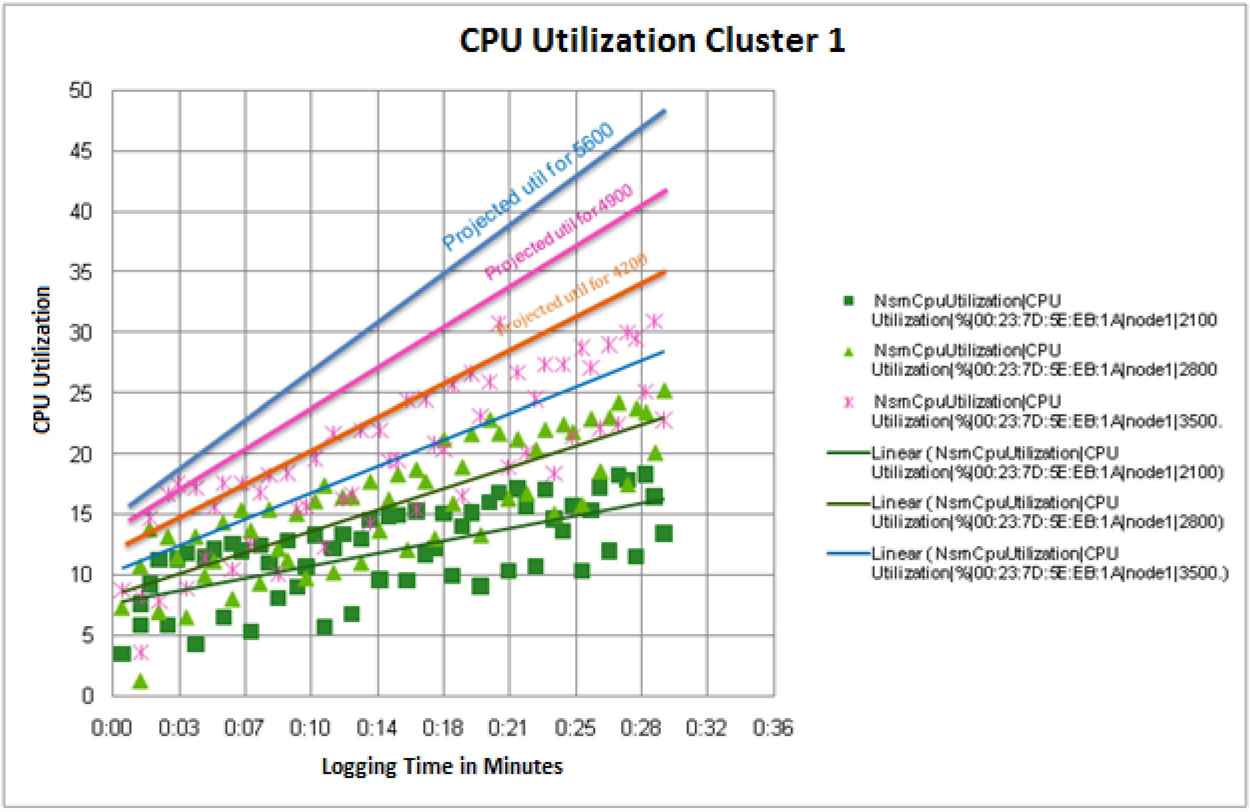

Figure 5: P4500 CPU Utilization Projections

Conclusion

The analysis shows that one of the primary factors for the project's success was that the processors on the blades were not fully utilized. This behavior differs from typical session-based tests where the processors are normally the bottleneck. In these tests the system memory was the limiting factor to the number of guests that could successfully run. A realistic number of guests per server without reducing performance is expected to be around 75, bringing the farm capacity to 6000 desktops for this workload level.

As mentioned above, the single-factor that limited the capacity of the farm was the amount of RAM available in the BL460c servers. The single-server scalability tests showed the blade servers capable of 75-90 desktops when the available RAM was increased to 96GB. However, since the hardware had been ordered with only 48GB of RAM the number of guests per host was capped at around 44, leaving unused processor capacity. Increasing the RAM to 96GB would allow for 75 guests per host.

Growing this farm would also require an additional DDC member server be added to the farm to provide XML and VDA registrations. The PVS analysis section indicated that the servers could probably handle the additional desktops, but the bootup times may be drawn out since the network cards were at 66% of maximum capacity so an additional PVS server would be recommended for both performance and high availability concerns.

To support 6000 desktops, an additional SCVMM server will be necessary so that each SCVMM server could manage 1000 desktops. The five SQL servers used for supporting the SCVMM servers could be consolidated into a single physical SQL server and easily support the load from all six SCVMM servers.

Based on the linear approximations, the storage analysis projections indicate that the two clusters could handle around 6000 desktops with just over 50% CPU utilization.

Appendix

Product Specifications

HP ProLiant BL680c G5 Blade Server

Model	HP BL680c G5 (492335-B21)
Processors	2 x 2.40 GHz E7450 Hex-core Xeon
RAM (Maximum)	96GB (128GB)
Memory slots	16 DIMM slots
Memory protection	Advanced ECC Online Spare Mirrored Memory Memory interleaving (2x1)
Expansion slots	3
Network Controller	2x1GbE NC373i Multifunction 2 ports 1x1GbE NC326i 2 ports
Disks	2 x 146GB 10K
Storage Controller	1-Smart Array P400i/256MB
Power Supply	Enclosure-based power
Form Factor (fully configured)	4 (c3000), 8 (c7000)
Rack height	Full-height

HP ProLiant BL460c G6 Blade Server

Model	HP BL460c G6 (532020-B21)
Processors	2 x 2.26 GHz L5520 Quad-core Nehalem
RAM (Maximum)	48GB 6x8GB PC3-8500 Dual Rank x 4 (192GB)
Memory Slots	12 DIMM slots
Memory protection	Advanced ECC Mirrored Memory
Expansion Slots	2
Network Controller	1x10GbE NC532i Flex-10 Multifunction 2 ports
Storage Controller	1-Smart Array P410i/Zero Memory
Disks	2 x 72GB 15K
Power Supply	Enclosure-based power
Form Factor (fully configured)	8 (c3000), 16 (c7000)
Rack height	Half-height

HP StorageWorks P4500 G2 SAS SAN Solution

Model	HP P4500 G2 (AX697A)
Shelves	16
Drives	192 x 450GB 3GB 15K LFF Dual-port SAS

Capacity	86.4TB (36.8TB RAID 10)
Clusters	2
Controllers	2 redundant controllers with 2,048 MB battery-backed cache.
Host Interface	1 Gb/sec iSCSI (32) ports

Solid State Disks (used for Single-Server Scalability Testing)

Model	HP 120GB 3G SATA 2.5in MDL	
Capacity	120GB	
Height	15mm	
Length x Width	Standard SFF	
Interface	SATA	
Transfer Rate Synchronous (Maximum)	3 Gb / sec	
Performance	Rotational Speed: Random Reads: Random Writes: Sequential Reads: Sequential Writes:	N/A >20,000 IOPS >5,000 IOPS 230MB/s 180MB/s
Physical Configuration	Logical blocks	512 bytes
Power	2W	
Operating Temperature	0-60° C	

HP StorageWorks IO Accelerator for BladeSystem c-Class

Model	AJ878A	
Capacity	320GB	
Height	1.38cm / 0.543”	
Length x Width	10.08cm x 11.33cm / 3.970” x 4.460”	
Interface	PCI Express Gen-1 x4	
Form factor	Type 1 c-Class Mezzanine	
Performance	Rotational Speed: Random Reads: Random Writes: Sequential Reads: Sequential Writes:	N/A ~100,000 IOPS ~100,000 IOPS 700MB/s 600MB/s
Physical Configuration (smaller block sizes require more physical RAM)	Logical blocks	512 – 8K
Power	7.5 watts (70% read / 30% write ratio)	
Operating Temperature	0-60° C	

HP ProCurve Switch 5412zl Intelligent Edge

Model	J8698A
Ports	12 open module slots 1 RS-232C DB-9 console port Maximum 288 10/100/1000 ports or 48 10-GbE ports or 288 mini-GBICs, or combination
Memory and Processor Gigabit Module	ARM9 @ 200 MHz; packet buffer size 144Mb QDR SDRAM
Memory and Processor 10G Module	ARM9 @ 200 MHz; packet buffer size 36Mb QDR SDRAM
Memory and Processor Management Module	Freescale PowerPC 8540 @ 666 MHz, 4 MB flash, 128 MB compact flash, 256 MB DDR SDRAM
Dimensions	17.75(d) x 17.5(w) x 12.1(h) in 45.09 x 44.45 x 30.73 cm (7U)
Latency	1000 Mb: < 3.7 µs (FIFO 64-byte packets) 10 Gbps: < 2.1 µs (FIFO 64-byte packets)
Throughput	Up to 480.3 million pps
Routing/Switching capacity	645.6 Gbps
Switch fabric speed	691.2 Gbps
Power	100-127 / 200-240 VAC 50/60 Hz 4 power supply slots (875W / 1500W)
Operating Temperature	0-55° C (0-40° C with J8706A or J8707A)

HP ProCurve 2810-48G Switch

Model	J9022A
Ports	44 10/100/1000 ports 1 RJ-45 serial console port 4 dual-personality ports (RJ-45 10/100/1000) or mini-GBIC
Dimensions (d x w x h)	12.7x17.4x1.7in (32.2 x 44.2 x 4.32 cm) 1U
Latency	< 5.4 µs (FIFO 64-byte packets)
Throughput	Up to 71.4 million pps
Routing/Switching capacity	96 Gbps
Power	100-127 / 200-240 VAC 50/60 Hz 100 Watts
Operating Temperature	0°C to 45°C

Cisco Catalyst Blade Switch 3120 for HP Hardware

Model	WS-CBS3120X-S
Ports	Up to 8 GbE uplink ports: 4 10/100/1000BASE-T 4 Small Form-Factor Pluggable (SFP) GbE
Dimensions (H x D x W)	1.1 x 7.6 x 10.76 in. [2.8 x 19.3 x 27.3 cm]
Power	Supplied by C7000 chassis
Operating Temperature	0 to 43°C

Acronyms

Abbreviation	Description
DC	Microsoft Active Directory Domain Controller
DDC	Citrix Desktop Delivery Controller
DSM	Device Specific Module
IOPS	Input / output operations per second
LS	Citrix License Server
MPIO	Multi-path input/output
PVS	Provisioning Server or Provisioning Services
SCVMM	Microsoft System Center Virtual Machine Manager
SQL	Microsoft SQL Server
SSS	Single-Server Scalability, usually refers to a test
VMM	Microsoft System Center Virtual Machine Manager
VDA	Virtual Desktop Agent
VSI	Login-Consultant’s Virtual Session Indexer
XD	Citrix XenDesktop
XML	Extensible Markup Language

Login VSI 2.1 Workload Timing Index

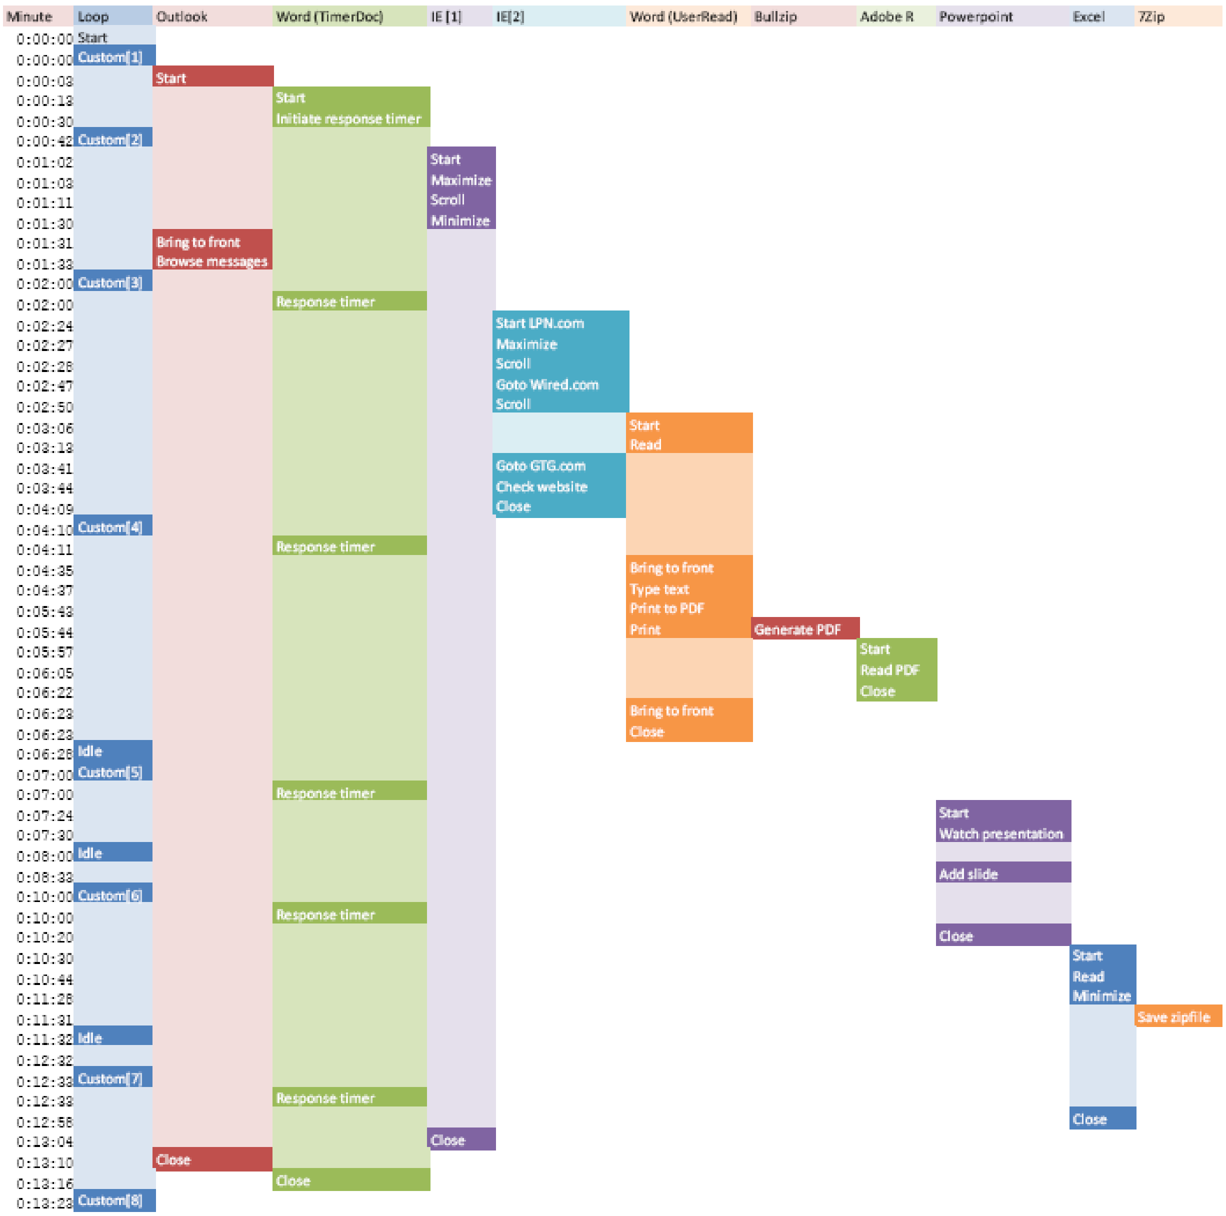

Revision History

Revision	Change Description	Updated By	Date
1.0	Original Draft	Citrix Consulting	August 2010

Citrix Consulting - Americas
Citrix Systems, Inc.
851 W Cypress Creek Road
Fort Lauderdale, FL 33309 USA
Phone: 1-800-393-1888 Option 55
Phone: 1-954-229-6170

Citrix Consulting - Pacific
Citrix Systems Singapore Pte Ltd

8, Temasek Boulevard
#34-02, Suntec Tower 3
Singapore, 038988
Singapore

Citrix Consulting - Europe
Citrix Systems UK, Ltd.
Chalfont House
Chalfont Park, Gerrards Cross
Bucks SL9 0BG
Phone: +44 1753-276-200

Citrix Systems Japan KK

Kasumigaseki Common Gate West Tower 24F
3-2-1, Kasumigaseki, Chiyoda-ku
Tokyo, 100-0013 Japan

About Citrix

Citrix Systems, Inc. (NASDAQ:CTXS) is a leading provider of virtual computing solutions that help companies deliver IT as an on-demand service. Founded in 1989, Citrix combines virtualization, networking, and cloud computing technologies into a full portfolio of products that enable virtual workstyles for users and virtual datacenters for IT. More than 230,000 organizations worldwide rely on Citrix to help them build simpler and more cost-effective IT environments. Citrix partners with over 10,000 companies in more than 100 countries. Annual revenue in 2009 was $1.61 billion.