Citrix XenDesktop on FlexPod with Microsoft Private Cloud

A Cisco Validated Design for 2000 Virtual Desktops featuring
Cisco Unified Computing System Blade Servers, NetApp FAS Storage,
Citrix XenDesktop 5.6, and Microsoft Hyper-V Server 2008 R2 SP1

Last Updated: May 24, 2012

Cisco Validated Design

Building Architectures to Solve Business Problems

About the Authors

Mike Brennan, Senior Technical Marketing Engineer, Cisco Systems

Mike Brennan is a Cisco Unified Computing System architect, focusing on Virtual Desktop Infrastructure solutions with extensive experience with Microsoft Hyper-V, Citrix XenDesktop and Provisioning Services. He has expert product knowledge in application and desktop virtualization across all three major hypervisor platforms, both major desktop brokers, Microsoft Windows Active Directory, Profiles, DNS, DHCP and Cisco networking technologies.

Rob Briggs, VDI Reference Architect, Microsoft Group, NetApp

Rob Briggs is a VDI Reference Architect in the NetApp Microsoft Group. Rob is focused on developing, validating and supporting VDI initiatives and solutions that include NetApp products. Rob has extensive experience with VDI on Microsoft and Citrix Virtualization.

Loay Shbeilat, Program Manager, Microsoft

Loay Shbeilat is a Program Manager at the Enterprise Engineering Center, focusing on real world deployments of Virtual Desktop Infrastructure solutions. Loay has expert product knowledge in the Microsoft Windows Server, with emphasis on Hyper-V and high availability.

Frank Anderson, Principle Solutions Architect, Strategic Alliance, Citrix Systems

Frank Anderson is a Principal Solutions Architect at Citrix, focusing on Desktop and Application Virtualization. Responsibilities include solutions validation, strategic alliances, technical content creation, testing, and benchmarking.

About Cisco Validated Design (CVD) Program

The CVD program consists of systems and solutions designed, tested, and documented to facilitate faster, more reliable, and more predictable customer deployments. For more information visit http://www.cisco.com/go/designzone.

Citrix XenDesktop on FlexPod with Microsoft Private Cloud

Citrix XenDesktop on FlexPod with Microsoft Private Cloud

Overview

Industry trends indicate a vast data center transformation toward shared infrastructures. Enterprise customers are moving away from silos of information and toward shared infrastructures, to virtualized environments, and eventually to the cloud to increase agility and reduce costs.

FlexPod™ is a predesigned configuration that is built on the Cisco® Unified Computing System® (Cisco UCS™), the Cisco Nexus® family of data center switches, NetApp® FAS storage components, and Microsoft® Windows Server® and System Center software. FlexPod is a base configuration, but can scale up for greater performance and capacity, or it can scale out for environments that require consistent, multiple deployments. It has the flexibility to be sized and optimized to accommodate many different use cases.

FlexPod is a platform that can address current virtualization needs and simplify the evolution to IT-as-a-service (ITaaS) infrastructure. FlexPod for Microsoft Private Cloud can help improve agility and responsiveness, reduce total cost of ownership (TCO), and increase business alignment and focus.

This document focuses on deploying Virtual Desktop Infrastructure (VDI) on a Microsoft Private Cloud built on FlexPod with Microsoft System Center 2012. It leverages the core infrastructure detailed in the Cisco white paper document titled, FlexPod with Microsoft Private Cloud: Architecture Overview released in April 2012. Readers should refer to that document for a detailed description of the core components and how they should be configured.

Corporate Headquarters:
Cisco Systems, Inc., 170 West Tasman Drive, San Jose, CA 95134-1706 USA

Audience

This document describes the architecture and deployment procedures of an infrastructure comprised of Cisco, NetApp, Microsoft, and Citrix virtualization technologies. The intended audience includes, but is not limited to, sales engineers, field consultants, professional services personnel, IT managers, partner engineering staff, and customers who want to deploy the core FlexPod architecture.

Solution Component Benefits

Each of the components of the overall solution materially contribute to the value of functional design contained in this document.

Benefits of the Cisco Unified Computing System

Cisco Unified Computing System™ is the first converged data center platform that combines industry-standard, x86-architecture servers with networking and storage access into a single converged system. The system is entirely programmable using unified, model-based management to simplify and speed deployment of enterprise-class applications and services running in bare-metal, virtualized, and cloud computing environments.

Benefits of the Cisco Unified Computing System include:

Architectural Flexibility

- Cisco UCS B-Series blade servers for infrastructure and virtual workload hosting
- Cisco UCS C-Series rack-mount servers for infrastructure and virtual workload hosting
- Cisco UCS 6200 Series second generation Fabric Interconnects provide unified blade, network and storage connectivity
- Cisco UCS 5108 Blade Chassis provide the perfect environment for multi-server type, multi-purpose workloads in a single containment

Infrastructure Simplicity

- Converged, simplified architecture drives increased IT productivity
- Cisco Unified Computing System management results in flexible, agile, high performance, self-integrating information technology with faster ROI
- Cisco Fabric Extender (FEX) technology reduces the number of system components to purchase, configure and maintain
- Standards-based, high bandwidth, low latency virtualization-aware unified fabric delivers high density, excellent virtual desktop user experience

Business Agility

- Model-based management means faster deployment of new capacity for rapid and accurate scalability
- Scale up to 16 chassis and up to 128 blades in a single Cisco UCS management domain
- Leverage Cisco UCS Management Packs for System Center 2012 for integrated management

Benefits of Cisco Nexus 5548UP Switches

The Cisco Nexus 5548UP Switch delivers innovative architectural flexibility, infrastructure simplicity, and business agility, with support for networking standards. For traditional, virtualized, unified, and high-performance computing (HPC) environments, it offers a long list of IT and business advantages, including:

Architectural Flexibility

- Unified ports that support traditional Ethernet, Fibre Channel (FC), and Fibre Channel over Ethernet (FCoE)
- Synchronizes system clocks with accuracy of less than one microsecond, based on IEEE 1588
- Offers converged fabric extensibility, based on the emerging IEEE 802.1BR standard, with FEX technology portfolio, including:
  - Cisco Nexus 2000 FEX
  - Adapter FEX
  - VM-FEX

Infrastructure Simplicity

- Common high-density, high-performance, data-center-class, fixed-form-factor platform
- Consolidates LAN and storage
- Supports any transport over an Ethernet-based fabric, including Layer 2 and Layer 3 traffic
- Supports storage traffic, including iSCSI, NAS, FC, RoE, and IBoE
- Reduces management points with FEX Technology

Business Agility

- Meets diverse data center deployments on one platform
- Provides rapid migration and transition for traditional and evolving technologies
- Offers performance and scalability to meet growing business needs

Specifications at-a Glance

- A 1 -rack-unit, 1/10 Gigabit Ethernet switch
- 32 fixed unified ports on base chassis and one expansion slot totaling 48 ports
- The slot can support any of the three modules: Unified ports, 1/2/4/8 native FC, and Ethernet or FCoE
- Throughput of up to 960 Gbps

Benefits of the NetApp FAS Family of Storage Controllers

Planning your storage implementation should take into account that VDI environments are extremely I/O intensive. IOPS range from majority reads to majority writes depending on the system state. When in a boot storm, the storage back end will see a steady increase in read IOPS. When in production, heavy-write IOPS might be noticed, especially during high end user workloads. NetApp recommends sizing the storage for high IOPS with small I/O sizing.

NetApp provides a scalable, unified storage and data management solution for VDI. The benefits of the NetApp solution are:

Storage efficiency: Significant cost savings with multiple levels of storage efficiency for all the virtual machine data components. These storage efficiencies include:

- NetApp Thin Provisioning, a way of logically presenting more storage to hosts than physically available.
- NetApp Deduplication, which saves space on primary storage by removing redundant copies of blocks within a volume.
- NetApp FlexClones, which provides hardware-assisted rapid creation of space-efficient, writeable, point-in-time images of individual files, LUNs, or flexible volumes.

Performance: Enhanced user experience with transparent read and write I/O optimization that strongly complements NetApp's storage efficiency capabilities. NetApp provides performance enhancements with:

- NetApp Transparent Storage Cache Sharing, which allows customers to benefit from storage efficiency and at the same time significantly increase I/O performance.
- NetApp Flash Cache, which increases the amount of available cache to help reduce virtual desktop storm activities and drastically improves read I/O.
- NetApp Write Optimization to optimize write operations in RAID- Double Parity (RAID-DP™).
- NetApp Flexible Volumes and Aggregates to allow the performance and capacity to be shared by all desktops in the volume or aggregate.

Data protection: Enhanced protection of both the virtual desktop operating system data and the user data with very low overhead for both cost and operations. Superior NetApp data protection is achieved with RAID-DP. NetApp RAID-DP is an advanced RAID technology that provides the default RAID level on all storage systems. RAID-DP protects against the simultaneous loss of two drives in a single RAID group. It is very economical to deploy and the overhead with default RAID groups is a mere 12.5 percent. This level of resiliency and storage efficiency make data residing on RAID-DP safer than data residing on RAID 5, and more cost effective than RAID 10.

Benefits of Microsoft Hyper-V 2008 R2 SP1

Microsoft Windows Server 2008 Release 2 Service Pack 1 (R2 SP1) Hyper-V builds on the architecture and functionality of Windows Server 2008 Hyper-V by adding multiple new features that enhance product flexibility, including dynamic memory, live migration, RemoteFx, and Second Level Address Translation (SLAT).

The adoption of virtualization in the enterprise has increased flexibility in the deployment and life cycle management of applications. IT professionals deploy and use virtualization to consolidate workloads and reduce server sprawl. Additionally, they deploy virtualization with clustering technologies to provide a robust IT infrastructure with high availability and fast disaster recovery.

Hyper-V provides a dynamic, reliable, and scalable virtualization platform combined with a single set of integrated management tools to manage both physical and virtual resources, enabling creation of an agile and dynamic data center.

Benefits of Citrix XenDesktop and Provisioning Server

If you are licensed to use the Advanced, Enterprise, or Platinum editions of Citrix XenDesktop, you can install Provisioning Server and use it to create a single desktop operating system image (vDisk) that you can stream to multiple desktops hosted in the virtual machine infrastructure.

Summary of Main Findings

The combination of technologies from Cisco Systems, Citrix Systems, Microsoft, and NetApp produced a highly efficient, robust and scalable Virtual Desktop Infrastructure (VDI) for a hosted virtual desktop deployment. Key components of the solution included:

- The combination of Cisco UCS compute, Nexus switching, and NetApp storage hardware with Microsoft Hyper-V 2008 R2, Citrix Provisioning Server 6.1, Citrix XenDesktop 5.6, and Microsoft System Center Virtual Machine Manager 2012 software produces a high density per blade and chassis virtual desktop delivery system.
- The Cisco UCS B230 M2 half-width blade with dual 10-core processors and 256GB of memory supports up to 145 users per blade and 22.7% more virtual desktop workloads than the previously studied full width blade using a medium workload with flash.
- A single FlexPod VDI design based on two Cisco UCS chassis, each with one Cisco UCS B200 M2 blade with dual six-core processors and 96 GB of memory and seven Cisco UCS B230 M2 blades with dual 10-core processors, 256GB of memory and a M81KR (Palo) converged network adapter supports 2000 virtual desktop workloads running a medium workload with Flash. This is more than 2.25 times the density of previously studied chassis with full width blades running infrastructure services in the same chassis.
- We were able to ramp up (log in and start workloads) to a steady state faster without pegging the processor, exhausting memory or storage subsystems.
- Compared to previous studies with full width blades, from a Cisco UCS blade and chassis basis, the rack space required to support 2000 users was reduced from 30 Rack Units to 12 Rack Units (RUs).
- Pure Virtualization: We continue to present a validated design that is 100 percent virtualized on Server 2008 R2 SP1 with Hyper-V role enabled and Hyper-V 2008 R2 SP1. All of the Windows 7 SP1 virtual desktops and supporting infrastructure components (including Active Directory, Profile Servers, Provisioning Servers, SQL Servers, and XenDesktop delivery controllers) were hosted as virtual servers.
- Cisco maintains industry leadership with the new Cisco UCS Manager 2.0 software that makes scaling simple, consistency guaranteed, and maintenance simple.
- The Cisco 10G unified fabric story gets additional validation on second generation Cisco UCS 6200 Series Fabric Interconnects and second generation Nexus 5500 Series access switches as we run more challenging workload testing, maintaining unsurpassed user response times.
- NetApp's FAS3240 system provides storage consolidation and efficiency. Both block and iSCSI storage resources were provided by a single system, utilizing NetApp Flash Cache technology.
- Validated that the entire environment including Boot from SAN Host OS's, Infrastructure Virtual Machines and 2000 Virtual Desktops could be run from one NetApp FAS3240 High Availability Storage Controller.
- Provisioned a 2000 seat VDI environment more than 25 times faster than VDI designs built with conventional storage by using Microsoft, Citrix and NetApp PowerShell scripts.

- Observed an 80 percent storage savings in comparison to a similar VDI environment without NetApp storage efficiencies.
- Citrix HDX technology, extended in XenDesktop 5.6 software, provides excellent performance with host-rendered Flash video and other demanding applications.

Architecture

The FlexPod architecture is highly modular or "pod" like. While each customer's FlexPod unit might vary in its exact configuration, once a FlexPod unit is built, it can easily be scaled as requirements and demands change. This includes scaling both up (adding additional resources within a FlexPod unit) and out (adding additional FlexPod units).

FlexPod is a defined set of hardware and software that serves as an integrated foundation for all virtualization solutions. Microsoft Private Cloud Solution validated with FlexPod includes NetApp storage, Cisco networking, the Cisco Unified Computing System, and Microsoft virtualization software in a single package in which the computing and storage can fit in one data center rack with the networking residing in a separate rack or deployed according to a customer's data center design. Due to port density, the networking components can accommodate multiple such configurations.

This document details the deployment of Citrix XenDesktop 5.6 with Provisioning Server 6.0 on a standard FlexPod for Microsoft Private Cloud. In addition to the Citrix software, Microsoft System Center VM Manager 2012 was deployed to manage the virtual machine infrastructure.

Figure 1 FlexPod for Microsoft Private Cloud Components

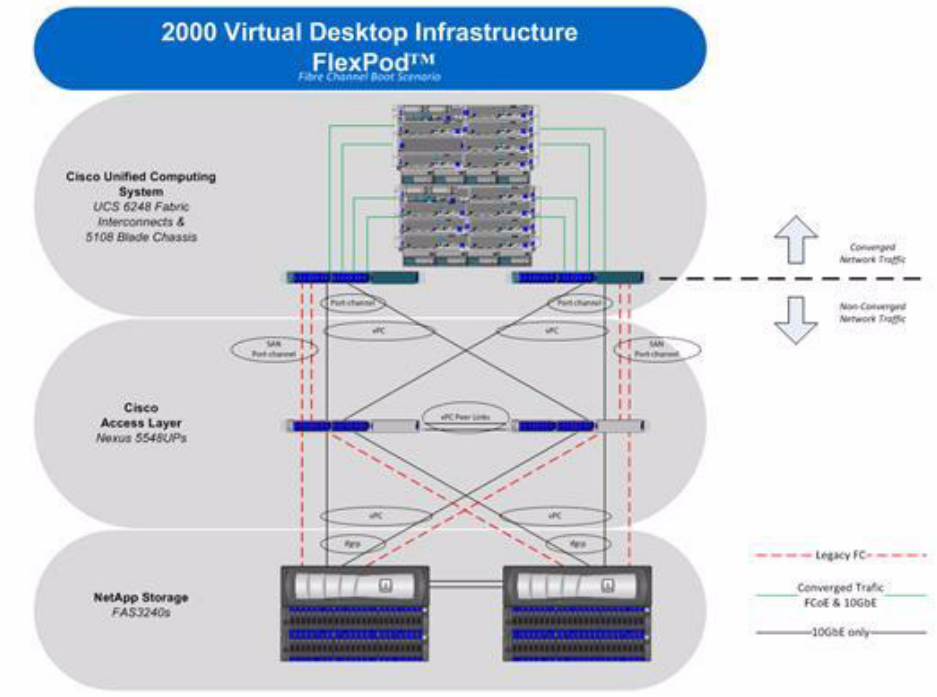

The reference configuration includes:

- Two Cisco Nexus 5548 switches
- Two Cisco UCS 6248 Series Fabric Interconnects

- Two Cisco UCS 5108 Blade Server Chassis with two 2104XP fabric extenders per chassis
- Two Cisco UCS B200 M2 Blade servers for infrastructure services
- Fourteen Cisco UCS B230 M2 Blade servers for VDI workloads
- One NetApp FAS3240A dual controller for HA

For low-level details of deploying Microsoft Hyper-V on top of a FlexPod for Microsoft Private Cloud, please refer to the Cisco Validated Design document titled, FlexPod with Microsoft Private Cloud: Architecture Overview released in April 2012.

Storage is provided by a NetApp FAS 3240A (HA configuration within a single chassis) with accompanying disk shelves. All systems and fabric links feature redundancy, providing for end-to-end HA. For desktop virtualization, the deployment includes MS Hyper-V and System Center VM Manager (VMM) 2012.

Software Revisions

Table 1 Software Revisions

Layer	Compute	Version or Release	Details
Compute	Cisco UCS Fabric Interconnect Cisco UCS B200 M2 Cisco UCS B230 M2	2.0 (1w) 2.0 (1w) 2.0 (1w)	Embedded Management Hardware BIOS Hardware BIOS
Network	Nexus Fabric Switch	5.0 (3) N2 (2a)	Operating System Version
Storage	NetApp FAS3240 HA	ONTAP 8.1.0 RC2	Operating System Version
Software	Cisco UCS Blade Hosts	B200: Microsoft Windows Server 2008 R2 SP1 Data Center Edition + MS Hyper-V Role. B230: Microsoft Hyper-V Server 2008 R2 SP1	Operating System Version
	.NET Framework	Various	Windows Feature
	Microsoft Hotfixes	6.4 64-bit	Microsoft Updates for all Microsoft Components
	NetApp SnapDrive for Windows	3.5	NetApp Integration within Windows
	Data ONTAP DSM	1.0	Windows MPIO Software
	SnapManager for Hyper-V	2008 SP2	Backup/Reture of Hyper-V Virtual Machines
	MS SQL Server	2012+Cumulative Update 1	VMM, XenDesktop, PVS, DBs
	SC VM Manager	3.0	VM Control
	OnCommand Plug-In	2.1.0	NetApp System Center Integration
	Cisco UCS Management Pk R2 Cisco Integration with SCOM Cisco UCS Power Tools Cisco UCS Power Shell Management	0.9.3.1	Commandlets

Configuration Guidelines

The FlexPod solution validated with Microsoft Private Cloud released in February 2012 provides details for configuring a fully redundant, highly-available configuration. As it applies to this document for building a VDI on a Microsoft Private Cloud built on FlexPod, references refer to which redundant component is being configured with each step, whether that be A or B. For example, Controller A and

Controller B are used to identify the two NetApp storage controllers that are provisioned with this document while Nexus A and Nexus B identify the pair of Cisco Nexus switches that are configured. The Cisco UCS Fabric Interconnects are configured likewise.

For this document, only the configurations that are specific for the Citrix XenDesktop 5.6 with Provisioning Server 6.0 FlexPod will be covered.

This document is intended to allow the reader to configure the Citrix XenDesktop 5.6 with Provisioning Server 6.0 FlexPod customer environment as an addition to an existing Microsoft Private Cloud Flexpod. Alternatively, the document will provide instructions for the reader to configure a stand-alone VDI pod utilizing Microsoft Hyper-V and System Center VM Manager 2012.

For the Citrix XenDesktop 5.6 with Provisioning Server 6.0 FlexPod, we utilized VLANs to isolate and apply access strategies to various types of network traffic. Table 2 details the VLANs used in this study.

Table 2 VLANs

VLAN Name	VLAN ID	Native
App/Cluster	2222	No
Public	2233	No
VM-PVS	2244	No
Management	2255	Yes
Clustered Shared Volumes (CSV)	2266	Yes
Live Migration	2277	Yes
iSCSI-A	2288	No
iSCSI-B	2299	No

Infrastructure Components

This section describes all of the Cisco Unified Computing System infrastructure components used in the configuration.

Cisco Unified Computing System

Cisco Unified Computing System is a set of pre-integrated data center components that include blade servers, adapters, fabric interconnects, and extenders that are integrated under a common embedded management system (Figure 2). This approach results in far fewer system components and much better manageability, operational efficiency, and flexibility than comparable data center platforms.

Figure 2 Cisco Unified Computing System Components

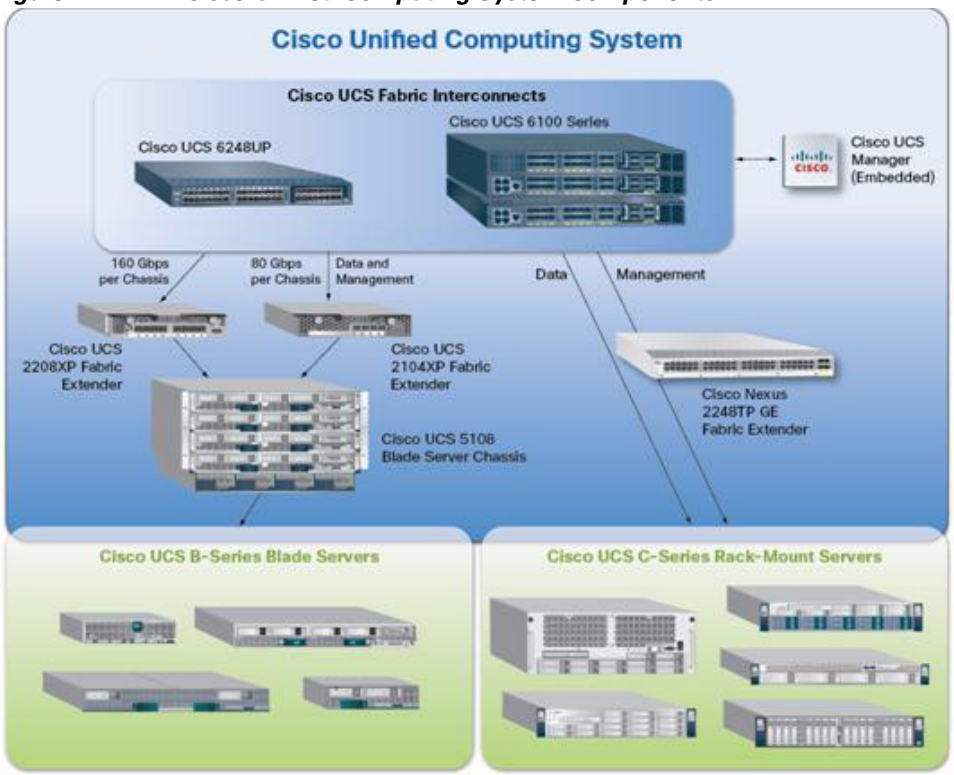

The Cisco Unified Computing System is designed from the ground up to be programmable and self integrating. A server's entire hardware stack, ranging from server firmware and settings to network profiles, is configured through model-based management. With Cisco virtual interface cards (VICs), even the number and type of I/O interfaces are programmed dynamically, making every server ready to power any workload at any time.

With model-based management, administrators manipulate a model of a desired system configuration, associate a model's service profile with hardware resources, and the system configures itself to match the model. This automation speeds provisioning and workload migration with accurate and rapid scalability. The result is increased IT staff productivity, improved compliance, and reduced risk of failures due to inconsistent configurations.

Cisco Fabric Extender technology reduces the number of system components to purchase, configure, manage, and maintain by condensing three network layers into one. It eliminates both blade server and hypervisor-based switches by connecting fabric interconnect ports directly to individual blade servers and VMs. Virtual networks are now managed exactly as physical networks are, but with massive scalability. This represents a radical simplification over traditional systems, reducing capital and operating costs while increasing business agility, simplifying and speeding deployment, and improving performance.

Fabric Interconnect

Cisco UCS Fabric Interconnects create a unified network fabric throughout the Cisco UCS. They provide uniform access to both networks and storage, eliminating the barriers to deploying a fully virtualized environment based on a flexible, programmable pool of resources.

Cisco Fabric Interconnects are a family of line-rate, low-latency, lossless 10-GE, Cisco Data Center Ethernet, and FCoE interconnect switches. Based on the same switching technology as the Cisco Nexus 5000 Series, Cisco UCS 6000 Series Fabric Interconnects provide the additional features and management capabilities that make them the central nervous system of the Cisco UCS.

The Cisco UCS Manager software runs inside the Cisco UCS Fabric Interconnects. The Cisco UCS 6000 Series Fabric Interconnects expand the Cisco UCS networking portfolio and offer higher capacity, higher port density, and lower power consumption. These interconnects provide the management and communication backbone for the Cisco UCS B-Series Blade servers and Cisco UCS Blade Server Chassis.

All chassis and all blades that are attached to the Cisco Fabric Interconnects are part of a single, highly available management domain. By supporting unified fabric, the Cisco UCS 6000 Series Fabric Interconnects provide the flexibility to support LAN and SAN connectivity for all blades within their domain right at configuration time. Typically deployed in redundant pairs, the Cisco UCS Fabric Interconnect provides uniform access to both networks and storage, facilitating a fully virtualized environment.

The Cisco UCS Fabric Interconnect family is currently comprised of the Cisco 6100 Series and Cisco 6200 Series Fabric Interconnects.

Cisco UCS 6248UP 48-Port Fabric Interconnect

The Cisco UCS 6248UP 48-Port Fabric Interconnect is a 1 RU, 10-GE, Cisco Data Center Ethernet, FCoE interconnect providing more than 1Tbps throughput with low latency. It has 32 fixed ports of FC, 10-GE, Cisco Data Center Ethernet, and FCoE SFP+ ports.

One expansion module slot can be up to 16 additional ports of FC, 10-GE, Cisco Data Center Ethernet, and FCoE SFP+.

Cisco UCS U6120XP 20-Port Fabric Interconnect

The Cisco UCS U6120XP 20-Port Fabric Interconnect is a 1 RU, 10-GE, Cisco Data Center Ethernet, FCoE interconnect providing more than 500-Gbps throughput with very low latency. It has 20 fixed 10-GE, Cisco Data Center Ethernet, and FCoE SFP+ ports.

One expansion module slot can be configured to support up to six additional 10-GE, Cisco Data Center Ethernet, and FCoE SFP+ ports.

Cisco UCS U6140XP 40-Port Fabric Interconnect

The Cisco UCS U6140XP 40-Port Fabric Interconnect is a 2 RU, 10-GE, Cisco Data Center Ethernet, and FCoE interconnect built to provide 1.04 Tbps throughput with very low latency. It has 40 fixed 10-GE, Cisco Data Center Ethernet, and FCoE SFP+ ports.

Two expansion module slots can be configured to support up to 12 additional 10-GE, Cisco Data Center Ethernet, and FCoE SFP+ ports.

Cisco UCS 2100 and 2200 Series IO Module

The Cisco UCS 2100/2200 Series FEX multiplexes and forwards all traffic from blade servers in a chassis to a parent Cisco UCS Fabric Interconnect from 10-Gbps unified fabric links. All traffic, even traffic between blades on the same chassis or virtual machines on the same blade, is forwarded to the parent interconnect, where network profiles are managed efficiently and effectively by the Fabric Interconnect. At the core of the Cisco UCS Fabric Extender are ASIC processors developed by Cisco that multiplex all traffic. Up to two Fabric Extenders can be placed in a blade chassis.

Cisco UCS 2104 has eight 10GBASE-KR connections to the blade chassis midplane, with one connection per Fabric Extender for each of the chassis' eight half slots. This gives each half-slot blade server access to each of two 10-Gbps unified fabric-based networks via SFP+ sockets for both throughput and redundancy. It has four ports connecting the Fabric Interconnect.

The Cisco UCS 2208 I/O module has 32 10GBASE-KR connections to the blade chassis midplane, with one connection per Fabric Extender for each of the chassis' eight half slots. This gives each half-slot blade server access to each of two 4x10-Gbps unified fabric-based networks via SFP+ sockets for both throughput and redundancy. It has eight ports connecting the Fabric Interconnect.

Cisco UCS Chassis

The Cisco UCS 5108 Series Blade Server Chassis is a 6RU blade chassis that will accept up to eight half-width Cisco UCS B-Series Blade Servers, up to four full-width Cisco UCS B-Series Blade Servers, or a combination of the two. The UCS 5108 Series Blade Server Chassis can accept four redundant power supplies with automatic load sharing and failover and two Cisco UCS Fabric Extenders, either 2100 or 2200 series. The chassis is managed by Cisco UCS Chassis Management Controllers, which are mounted in the Cisco UCS Fabric Extenders and work in conjunction with the Cisco UCS Manager to control the chassis and its components.

A single Cisco Unified Computing System managed domain can theoretically scale to up to 40 individual chassis and 320 blade servers. At this time Cisco supports up to 20 individual chassis and 160 blade servers.

Basing the I/O infrastructure on a 10-Gbps unified network fabric allows the Cisco Unified Computing System to have a streamlined chassis with a simple yet comprehensive set of I/O options. The result is a chassis that has only five basic components:

- The physical chassis with passive midplane and active environmental monitoring circuitry
- Four power supply bays with power entry in the rear and hot-swappable power supply units accessible from the front panel
- Eight hot-swappable fan trays, each with two fans
- Two Fabric Extender slots accessible from the back panel
- Eight blade server slots accessible from the front panel

Cisco UCS B200 M2 Blade Server

The Cisco UCS B200 M2 Blade Server balances simplicity, performance, and density for production-level virtualization and other mainstream data -center workloads. The server is a half-width, 2-socket blade server with substantial throughput and scalability. The UCS B200 M2 server extends the capabilities of the Cisco Unified Computing System. It uses Intel® Xeon® 5600 series multicore processors to deliver even better performance and efficiency. Two UCS B200 M2 blades were used in this study to host infrastructure virtual machines.

Cisco UCS B230 M2 Blade Server

The Cisco UCS B230 M2 Blade Server is a full-slot, two-socket blade server featuring the performance and reliability of Intel Xeon Processor E7-2800 product family and up to 32 DIMM slots which support up to 512 GB of memory. The Cisco UCS B230 M2 Blade Server supports two SSD drives and one CNA mezzanine slot for up to 20 Gbps of I/O throughput. The server delivers outstanding performance,

memory, and I/O capacity to meet the diverse needs of a virtualized environment with advanced reliability and exceptional scalability for the most demanding applications. Fourteen B230 M2 blades were used in this study to host Windows 7 SP1 virtual desktops.

Intel Xeon 5600 Series Processor

Cisco UCS B200 M2 servers utilize x86-based Intel Xeon 5600 Series processors. The advanced reliability, availability, and serviceability (RAS) features of Intel Xeon, combined with the highly available Cisco UCS architecture with redundant, hot-swappable components is designed to protect data integrity and reduce downtime. Intel Xeon processors also automatically and intelligently adjust server performance according to application requirements, increasing performance as needed and achieving substantial energy savings when performance requirements are low.

The Intel Xeon processor has become ubiquitous throughout the most demanding IT environments, supporting multiple OSs, including Sun Solaris and varieties of Linux and Microsoft Windows systems, enabling a consistent deployment platform across organizations. With Intel's continued support of virtual server environments, the x86 architecture has also become the standard platform used for virtualization.

Intel Xeon E7 Series Processor

The Intel Xeon processor E7-8800/4800/2800 product families deliver the highest performance of Intel® Xeon® processors for enterprise, mission-critical, and high-performance computing (HPC) deployments. With up to 10 cores and 20 threads, up to 30 MB of last-level cache, and supporting up to 32 GB DIMMs, the Intel Xeon processor E7 families are ideal for large data centers, ERP, CRM, and SCM applications, and the most demanding scientific and financial workloads. The E7-2870 processor was used in the B230 M2 blades in this study.

Cisco UCS Converged Network Adapter

A Cisco innovation, the Cisco UCS M81KR (Palo) VIC is a virtualization-optimized FCoE mezzanine card designed for use with Cisco UCS B-Series Blade Servers (Figure 3). The VIC is a dual-port 10 Gigabit Ethernet mezzanine card that supports up to 128 Peripheral Component Interconnect Express (PCIe) standards-compliant virtual interfaces that can be dynamically configured so that both their interface type (network interface card [NIC] or HBA) and identity (MAC address and worldwide name [WWN]) are established using just-in-time provisioning. In addition, the Cisco UCS M81KR supports Cisco VN-Link technology, which adds server-virtualization intelligence to the network.

Figure 3 Cisco UCS VIC M81KR Converged Network Adapter

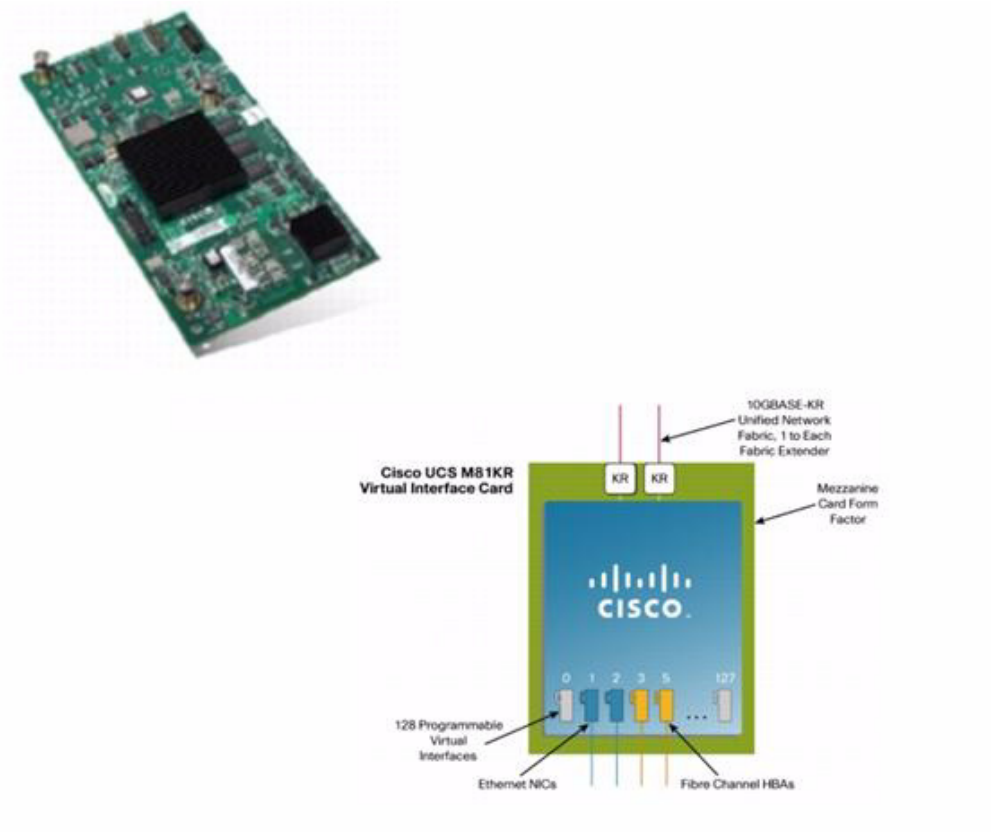

Cisco Extended Memory Architecture

A crucial innovation of the Cisco UCS, Cisco patented Extended Memory Technology provides more than twice as much memory (384 GB) as traditional two-socket servers, increasing performance and capacity for demanding virtualization and large data-set workloads. Alternatively, this technology offers a more cost-effective memory footprint for less-demanding workloads.

Building on the power of the Intel Xeon 5500 Series processors in the Cisco UCS, Cisco's Extended Memory Technology enables up to 384 GB of memory on a single server. It is available on the Cisco UCS B250 M1 Blade Server and the Cisco UCS C250 M1 Rack-Mount Server. This technology provides more than double the industry-standard memory footprint when compared even to other Xeon 5500 Series processor-based systems. Cisco Extended Memory Technology enables memory scalability decoupled from the traditional cost. With reduced costs and larger-than-ever memory footprints, IT departments can now consolidate more applications and virtual machines more economically.

Cisco UCS C-Series Rack-Mount Servers

Cisco UCS C-Series Rack-Mount Servers (Figure 4) extend Cisco UCS innovations to a rack-mount form factor, including a standards-based unified network fabric, Cisco VN-Link virtualization support, and Cisco Extended Memory Technology. Designed to operate both in standalone environments and as part of the Cisco Unified Computing System, these servers enable organizations to deploy systems incrementally-using as many or as few servers as needed-on a schedule that best meets the organization's

timing and budget. Cisco UCS C-Series servers offer investment protection through the capability to deploy them either as standalone servers in heterogeneous data centers or as part of the Cisco Unified Computing System.

Although this deployment used the Cisco UCS B-Series blade servers, the Cisco UCS C-Series Rack-Mount servers extend the same benefits to customers. Future desktop virtualization deployments are planned on this server platform.

Each server platform addresses varying workload challenges through a balance of processing, memory, I/O, and internal storage resources.

- The Cisco UCS C460 M2 High-Performance Rack-Mount Server is a 4-socket, 4RU enterprise-critical server for data-demanding applications and bare-metal and virtualized workloads.
- The Cisco UCS C260 M2 Rack-Mount Server is one of the industry's highest-density and most expandable 2-socket, 2RU servers for enterprise-critical workloads ranging from storage serving to online transaction processing (OLTP) and data warehousing.
- The Cisco UCS C250 M2 Extended Memory Rack-Mount Server is a 2-socket, 2RU server that features Cisco Extended Memory Technology, which increases performance and capacity for a wide range of memory-intensive enterprise workloads.
- The Cisco UCS C240 M3 Rack Server is a third generation 2RU server designed for both performance and expandability over a wide range of storage-intensive infrastructure workloads from Big Data to collaboration.
- The Cisco UCS C220 M3 Rack Server is a third generation 1RU server designed for performance and density over a wide range of business workloads from Web serving to distributed database.
- The Cisco UCS C210 M2 High-Density Rack-Mount Server is a 2-socket, 2RU server that offers up to 16 internal disk drives for up to 8 terabytes (TB) of disk space for storage-intensive workloads.
- The Cisco UCS C200 M2 High-Density Rack-Mount Server is a 2-socket, 1RU server that balances simplicity, performance, and density for production-level virtualization, IT and Web infrastructure, standalone applications, and other mainstream data center workloads.

Figure 4 Cisco UCS C-Series Rack-Mount Servers

Citrix XenDesktop

Citrix XenDesktop is a desktop virtualization solution that delivers Windows desktops as an on-demand service to any user, anywhere. With FlexCast™ delivery technology, XenDesktop can quickly and securely deliver individual applications or complete desktops to the entire enterprise, whether users are task workers, knowledge workers or mobile workers. Users now have the flexibility to access their desktop on any device, anytime, with a high definition user experience. With XenDesktop, IT can

manage single instances of each OS, application, and user profile and dynamically assemble them to increase business agility and greatly simplify desktop management. XenDesktop's open architecture enables customers to easily adopt desktop virtualization using any hypervisor, storage, or management infrastructure.

Enhancements in Citrix XenDesktop 5.6

Citrix XenDesktop 5.6 makes virtual desktops personal and cost-effective. Enterprises around the world are transforming their desktop environments from device-centric management to user-centric private clouds where desktops and applications are delivered as a service, on-demand. XenDesktop 5.6 accelerates desktop transformation by delivering high-performance personal desktops and applications with all the flexibility, performance, and user experience of a PC but optimized for network, server, and storage resources.

Citrix XenDesktop 5.6 features include:

- **Citrix Personal vDisk technology**—Drives down the cost of implementing desktop virtualization by allowing IT to supply even the most demanding users with flexible, personalized, and persistent virtual desktops while benefitting from cost effective and easy-to-maintain pooled virtual desktops. Personal vDisk technology is now fully integrated into Desktop Studio, Desktop Director, and Citrix Provisioning Services.
- **Mobile application access**—XenApp dynamically transforms an application's user interface to look and feel like the native user interface of smartphones and tablet devices. Now existing Windows applications adapt to the way users interact with applications on smaller devices.
- **CloudGateway Express**—Aggregates and centrally delivers virtual applications and desktops to provide users with an intuitive single point of access and self-service to all their business applications on any device, anywhere.
- **Microsoft System Center 2012 ready**—Updates integration support for Microsoft System Center 2012 Configuration Manager to make pooled VDI virtual desktops look like a standard desktop from a System Center Configuration Manager perspective and leverages its policy enforcement and reporting tools.
- **XenClient 2.1**—Windows Dynamic Layering, a single base image management technology, provides an easier and more reliable way of managing updates on XenClient devices and includes new multi-lingual support in German, French, Spanish, Japanese, and Simplified Chinese languages.

FlexCast Technology

Citrix XenDesktop with FlexCast is an intelligent delivery technology that recognizes the user, device, and network, and delivers the correct virtual desktop and applications specifically tailored to meet the performance, security, and flexibility requirements of the user scenario. FlexCast technology delivers any type of virtual desktop to any device and can change this mix at any time. FlexCast also includes on-demand applications to deliver any type of virtual applications to any desktop, physical or virtual.

The FlexCast delivery technologies can be broken down into the following categories:

- **Hosted Shared Desktops** provide a locked down, streamlined and standardized environment with a core set of applications, ideally suited for task workers where personalization is not needed or allowed.
- **Hosted VDI Desktops** offer a personalized Windows desktop experience, typically needed by office workers, which can be securely delivered over any network to any device.

- **Streamed Virtual Hard Disk (VHD) Desktops** use the local processing power of rich clients while providing centralized single image management of the desktop. These types of desktops are often used in computer labs and training facilities and when users require local processing for certain applications or peripherals.
- **Local VM Desktops** utilize XenClient to extend the benefits of centralized, single-instance management to mobile workers that need to use their laptops offline. When they are able to connect to a suitable network, changes to the OS, applications, and user data are automatically synchronized with the data center.
- **On-demand Applications** allows any Windows® application to be centralized and managed in the data center, hosted either on multi-user terminal servers or VMs and instantly delivered as a service to physical and virtual desktops. Optimized for each user device, network, and location, applications are delivered through a high speed protocol for use while connected or streamed through Citrix application virtualization or Microsoft App-V directly to the endpoint for use when offline.

High-Definition User Experience Technology

Citrix High-Definition User Experience (HDX) technology is a set of capabilities that delivers a high definition desktop virtualization user experience to end users for any application, device, or network. These user experience enhancements balance performance with low bandwidth. Anything else becomes impractical to use and scale. Citrix HDX technology provides network and performance optimizations to deliver the best user experience over any network, including low bandwidth and high latency WAN connections.

Citrix XenDesktop Hosted VDI Overview

Hosted VDI uses a hypervisor to host all the desktops in the data center. Hosted VDI desktops can either be pooled or assigned. Pooled virtual desktops use Citrix Provisioning Services to stream a standard desktop image to each desktop instance upon boot-up. Therefore the desktop is always returned to its clean, original state. Citrix Provisioning Services enables the streaming of a single desktop image to create multiple virtual desktops on one or more hypervisors in a data center. This feature greatly reduces the amount of storage required compared to other methods of creating virtual desktops. The high-level components of a Citrix XenDesktop architecture utilizing the Hosted VDI model for desktop delivery are shown in Figure 5.

Figure 5 Citrix XenDesktop on Microsoft Hyper-V Architecture

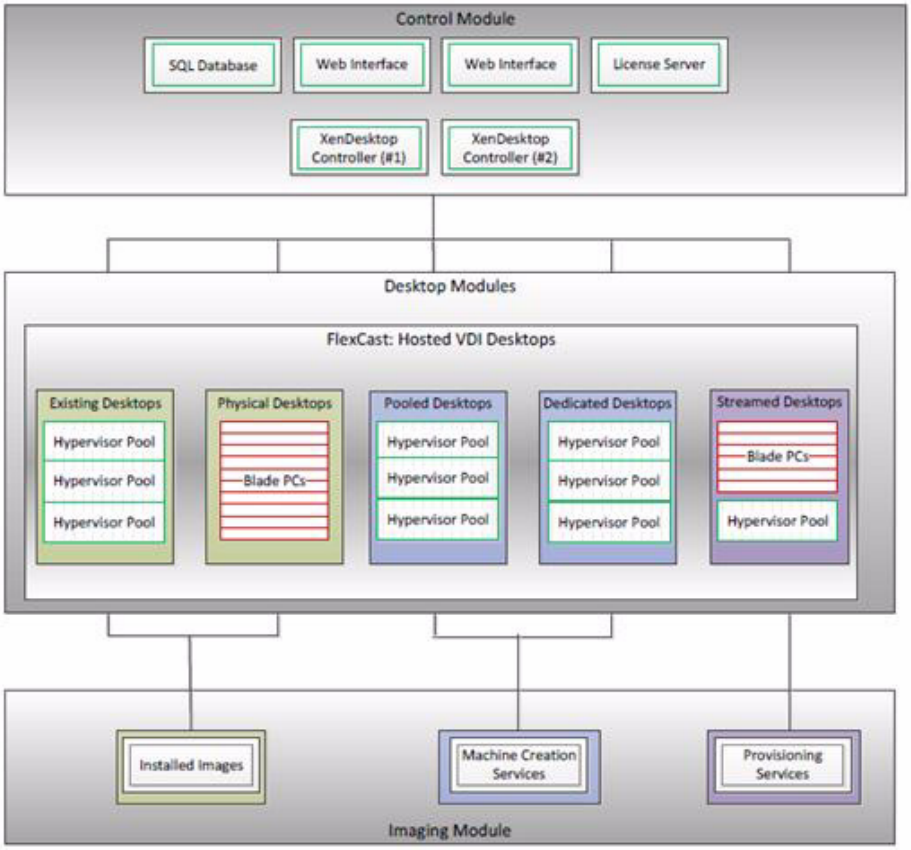

Components of a Citrix XenDesktop architecture using Hosted VDI include:

- **Web Interface**: The Web Interface provides the user interface to the XenDesktop environment. Web Interface brokers user authentication, enumerates the available desktops and, upon launch, delivers an .ica file to the Citrix Receiver on the user's local device to initiate a connection. Because Web Interface is a critical component, redundant servers must be available to provide fault tolerance.
- **License Server**: The Citrix License Server is responsible for managing the licenses for all of the components of XenDesktop 5.5. XenDesktop has a 90 day grace period which allows the system to function normally for 90 days if the license server becomes unavailable. This grace period offsets the complexity involved with building redundancy into the license server.
- **Domain Controller**: The Domain Controller hosts Active Directory, Dynamic Host Configuration Protocol (DHCP), and Domain Name System (DNS). Active Directory provides a common namespace and secure method of communication between all the servers and desktops in the environment. DNS provides IP Host name resolution for the core XenDesktop infrastructure components. DHCP is used by the virtual desktop to request and obtain an IP address from the DHCP service. DHCP uses Option 66 and 67 to specify the bootstrap file location and file name to a virtual desktop. The DHCP service receives requests on UDP port 67 and sends data to UDP port 68 on a virtual desktop. The virtual desktops then have the operating system streamed over the network utilizing Citrix Provisioning Services (PVS).
- **Provisioning Services**: PVS creates and provisions virtual desktops from a single desktop image (vDisk) on demand, optimizing storage utilization and providing a pristine virtual desktop to each user every time they log on. Desktop provisioning also simplifies desktop images, provides the best flexibility, and offers fewer points of desktop management for both applications and desktops. The

Trivial File Transfer Protocol (TFTP) and Pre-boot eXecution Environment (PXE) services are required for the virtual desktop to boot off the network and download the bootstrap file which instructs the virtual desktop to connect to the PVS server for registration and vDisk access instructions.

- **Desktop Delivery Controller**: The XenDesktop controllers are responsible for maintaining the proper level of idle desktops to allow for instantaneous connections, monitoring the state of online and connected virtual desktops and shutting down virtual desktops as needed. The primary XD controller is configured as the farm master server. The farm master is able to focus on its role of managing the farm when an additional XenDesktop Controller acts as a dedicated XML server. The XML server is responsible for user authentication, resource enumeration, and desktop launching process. A failure in the XML broker service will result in users being unable to start their desktops. This is why multiple controllers per farm are recommended.
- **Data Store**: Each XenDesktop farm requires a database called the data store. Citrix XenDesktops use the data store to centralize configuration information for a farm in one location. The data store maintains all the static information about the XenDesktop environment.
- **Virtual Desktop Agent**: The Virtual Desktop Agent (VDA) is installed on the virtual desktops and enables direct Independent Computing Architecture (ICA) connections between the virtual desktop and user devices with the Citrix online plug-in.
- **Citrix Receiver**: Installed on user devices, the Citrix Receiver enables direct HDX connections from user devices to virtual desktops. The Receiver is available for a range of different devices so users can connect to on-demand applications from various platforms.
- **Hypervisor**: XenDesktop has an open architecture that supports the use of XenServer, Microsoft Hyper-V, VMware ESX, or vSphere. For the purposes of the testing documented in this paper, Microsoft Hyper-V was the hypervisor of choice.
- **Citrix XenApp**: Citrix XenApp is an on-demand application delivery solution that enables any Windows application to be virtualized, centralized, managed in the data center, and instantly delivered as a service to users anywhere on any device. XenApp can be used to deliver both virtualized applications and virtualized desktops. In the Hosted VDI model, XenApp is typically used for on-demand access to streamed and hosted applications.

All of the aforementioned components interact to provide a virtual desktop to an end user based on the FlexCast Hosted VDI desktop delivery model leveraging the Provisioning Services feature of XenDesktop. This architecture provides the end user with a pristine desktop at each logon based on a centralized desktop image that is owned and managed by IT. The following steps in Figure 6 outline the sequence of operations executed by XenDesktop to deliver a Hosted VDI virtual desktop to the end user.

Figure 6 Operational Sequence

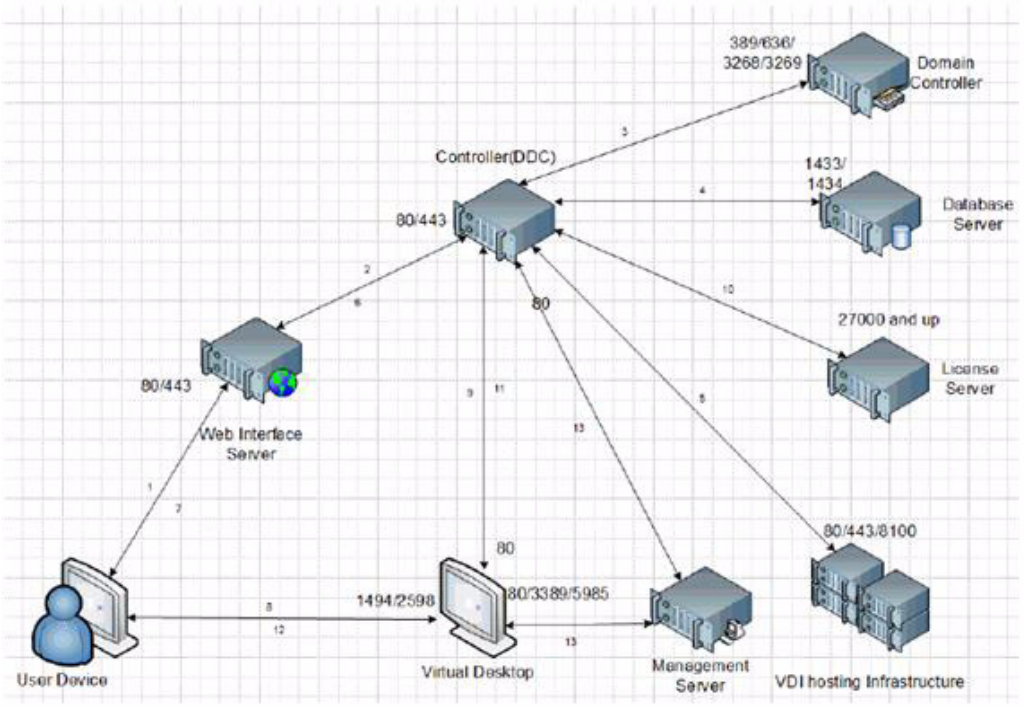

1. The end user launches an Internet browser to access Web Interface and submits their credentials.
2. Web interface passes the user credentials to the XenDesktop Controller.
3. Controller verifies user authorization by performing a Microsoft Active Directory query with the end user's credentials.
4. Controller queries the site database for the end user's assigned desktop groups, named instance uses ports 1434 and 1433.
5. Using the desktop group obtained from the database, controller queries the hypervisor about the status of desktops within that group.
6. Controller identifies to Web Interface the desktop it assigned for this particular session.
7. Web Interface sends an ICA file to the online plug-in, which points to the virtual desktop identified by the hypervisor.
8. Online plug-in establishes an ICA connection to the specific virtual desktop that was allocated by the controller for this session.
9. Virtual Desktop Agent verifies the license file with the controller.
10. Controller queries Citrix license server to verify that the end user has a valid ticket.
11. Controller passes session policies to the Virtual Desktop Agent (VDA), which then applies those policies to the virtual desktop.
12. Online plug-in displays the virtual desktop to the end user.
13. Administrator and helpdesk personnel use Desktop Director and Studio tools to manage the desktops from the management server.

Citrix Provisioning Services

Citrix Provisioning Server provides images to physical and virtual desktops. Desktops utilize network booting to obtain the image and only portions of the desktop images are streamed across the network as needed. Provisioning Server does require additional server resources but these can be either physical or virtual servers depending on the capacity requirements and hardware configuration. Also, Provisioning Server does not require the desktop to be virtualized as Provisioning Server can deliver desktop images to physical desktops.

Citrix XenDesktop 5.6 Advantages and Value Proposition

Personal VDI desktops for any user expand the use cases for virtual desktops to demanding knowledge workers by permitting all of the customization, application flexibility, and persistence they expect from their desktop while achieving the ease of management and storage optimizations needed for large-scale deployments.

- Virtual user drive technology called Personal vDisk in Desktop Studio permits administrators or even users to install applications without impacting the master image.
- New user profile manager unifies and manages user settings across Windows platforms.
- Windows folder redirection keeps user data out of the VM for efficiency and high availability.

Support for More than One Billion Devices

Citrix XenDesktop 5.6 provides secure access to any desktop or application over any device from any location. It enables delivery of Windows applications to mobile devices that adapt to the way users interact with applications on smaller devices such as smartphones and tablets without source code changes.

- New Citrix Receiver clients support over one billion devices including iOS, Android, WebOS, Chrome OS, Mac OS, Linux and Windows
- New Receiver clients deliver over three times faster Windows performance
- New Receiver client for Linux powers a new generation of low-cost thin clients

Scalable, High-Performance Multimedia

Optimize your users' experience by delivering high-quality voice, video and multimedia while reducing bandwidth and server-side rendering to drive down the cost of deployment.

- Real-time audio stream reduces the impact of network latency for voice traffic
- Webcam compression enables enterprise-scale video conferencing
- Graphics command redirection decreases bandwidth an additional 33% while delivering a "local" experience
- Intelligent local rendering takes the load off the server, increasing density by 10 times

Breakthrough WAN Performance for Remote, Branch and Cloud Deployments

Expand virtual desktop deployments across the enterprise from remote branches to mobile users with high quality of service (QoS) to meet the most demanding service level agreements.

- New multi-stream ICA splits virtual desktop traffic into five separate streams, permitting granular QoS controls

- HDX WAN Optimization powered by Branch Repeater 6.0 doubles the number of users that can be supported on a given network
- HDX WAN Optimization improves performance for both mobile users and branch employees

NetApp FAS Systems

Challenges

There is an ongoing challenge to efficiently meet the growing storage needs of business applications such as VDI. What is needed is storage that is both efficient and flexible with high-end availability and performance to effectively meet the challenges of virtualized and traditional IT environments.

The Solution

The NetApp FAS3200 series enables businesses to cost-effectively meet the storage needs of business applications in both virtual and traditional environments. It handles today's workloads with Data ONTAP industry-leading storage efficiency through the NetApp Unified Storage Architecture running NetApp's storage operating system. Users can consolidate diverse data sets and be ready to respond to changes more easily and non-disruptively with the extra PCIe slots, high performance, and enterprise-class availability of the FAS3200 series.

Midsize environments can gain high-end benefits without the budget or space required to support frame array-class systems. Regional data centers, replicated sites, and departmental systems that need full-featured yet efficient storage with advanced availability and performance capabilities can benefit too.

The FAS3200 series includes two models to tailor-fit a solution that is right for all environments: FAS3240, and FAS3270. NetApp systems are easy to install, provision, and upgrade to meet customer needs. Maximize productivity with NetApp's common suite of application-aware management software. The FAS3200 systems, with the best value for mixed workloads, can give customers an edge that other midrange storage platforms simply cannot match.

Lower Costs with Highly Efficient Systems

NetApp delivers a truly unified storage architecture that uses a single platform with common software and processes across all tiers of storage. Users can consolidate their diverse workloads with multi-protocol support and also benefit from integrated data protection and one operating system across the entire family. This helps customers maximize the efficiency of their virtual servers by delivering storage when and where it is needed.

Additionally, common management across the unified storage architecture assists in consolidation of diverse data sets and data-in-place controller upgrades to more powerful FAS systems. This lowers administrative costs and makes it easier for customers to deploy new capabilities across the enterprise. Deploying a FAS3200 series system with a NetApp DS2246 disk shelf results in dramatically decreased space, power, and cooling consumption. This leading-edge disk shelf features the latest SAS technology by using small form factor SAS 2.5" disk drives that are capable of doubling capacity per rack unit, conserving valuable data center resources.

Increase Flexibility

The FAS3200 series scales to nearly 3PB of versatile storage that adapts readily to growing storage demands. If the environment needs extra connectivity, the expanded I/O configurations of the FAS3240 and FAS3270 models significantly increase the number of PCIe expansion slots available. Moreover, all FAS3200 systems support Data ONTAP in both standard and cluster modes, providing the flexibility of up to 24 nodes.

Proven Availability and Performance

The FAS3200 series is built on the proven enterprise-class availability of the NetApp storage infrastructure. The FAS3200 models leverage high-end systems by introducing features such as Alternate Control Path (ACP) and service processor. These enhance NetApp's already highly available architecture by enabling additional diagnostics and non-disruptive recovery.

Data availability can be further boosted with zero planned and unplanned downtime by combining the FAS3200 series with the NetApp MetroCluster™ solution, which promotes continuous access to data and prevents data loss. MetroCluster delivers distance array-based clustering to protect against outages in the data center, across campus, or citywide due to hardware, power, and network failures and environmental faults.

The FAS3200 series features leading-edge technology for high-performance storage. Use the NetApp DS4243 with solid state drives (SSDs) when every I/O read must be fast. Or boost system performance by adding up to 2TB of Flash Cache. These intelligent caching modules automatically increase read rates and reduce average latency for frequently accessed data, without adding more disk drives. NetApp Flash Cache combined with hard disk drives is an effective and typically more affordable alternative to SSDs.

NetApp Technology Differentiators

Single Scalable Unified Architecture

The NetApp Unified Storage Architecture provides customers with an agile and scalable storage platform. NetApp's innovative storage solutions provide customers with new alternatives and expanded possibilities compared to solutions from traditional storage vendors. All NetApp storage systems utilize the Data ONTAP® operating system to provide SAN (FCoE, FC, and iSCSI), NAS (CIFS, NFS), primary storage, and secondary storage within a single unified platform so that all virtual desktop data components can be hosted on the same storage array. A single process for activities such as installation, provisioning, mirroring, backup, and upgrading is used throughout the entire product line from the entry level to enterprise-class controllers. Having a single set of software and processes brings welcome simplicity to even the most complex enterprise data management challenges.

Unifying storage and data management software and processes reduces the complexity of data ownership, enables companies to adapt to their changing business needs without interruption, and results in a dramatic reduction in TCO.

For large, scalable VDI environments, the NetApp solution provides the following benefits:

- At least 50 percent savings in storage, power, and cooling requirements
- Agile and operationally efficient storage solutions
- Best-in-class data protection and business continuance solutions to address any level of data availability demands

Storage Efficiency

One of the critical barriers to VDI adoption is the increased cost of using shared storage to obtain a highly available enterprise quality infrastructure. Virtual desktop deployment creates a high level of data redundancy, especially for the virtual machine OS data. Using traditional storage, this means you need storage equal to the sum of the storage required by each virtual machine. For example, if each virtual machine is 20 GB in size and there are supposed to be 1000 virtual machines in the solution, at least 2 GB of usable data would be required on the shared storage.

Thin provisioning, data deduplication, and FlexClone® thin-cloning technology are the critical components of the NetApp solution and offer multiple levels of storage efficiency across the virtual desktop OS data, installed applications, and user data. This helps customers save 50 percent to 90 percent

on the cost associated with shared storage (based on existing customer deployments and NetApp solutions lab validation). NetApp is the only storage vendor that offers block-level data deduplication for live virtual machines, without any negative trade-offs.

Thin Provisioning

Thin provisioning is a way of logically presenting more storage to hosts than is physically available. With thin provisioning, the storage administrator is able to utilize a pool of physical disks (known as an aggregate) and to create logical volumes for different applications to use, while not allocating space to those volumes. The space gets allocated only when the host needs it. The unused aggregate space is available for the existing thinly provisioned volumes to expand or for use in creation of new volumes. For details about thin provisioning, refer to NetApp TR 3563: NetApp Thin Provisioning.

NetApp recommends using thinly provisioned logical unit numbers (LUNs) where possible in the Hyper-V environment for maximum storage efficiency. Note that when using thin provisioning it is important to monitor capacity utilization. Administrators should also configure storage management policies on the volumes that contain the thin-provisioned LUNs. The use of these policies aids in providing the thin-provisioned LUNs with storage capacity as they require it. The policies include automatic sizing of a volume, automatic snapshot deletion, and LUN fractional reserve.

NetApp Deduplication

NetApp deduplication saves space on primary storage by removing redundant copies of blocks within a volume hosting hundreds of virtual desktops. This process is transparent to the application and user and can be enabled and disabled on the fly. In a VDI environment, deduplication provides significant space savings, given that each virtual machine is an identical copy of the OS, applications, and patches. The savings are also achieved for the user data hosted on CIFS home directories. For more information on NetApp deduplication, refer to NetApp TR-3505: NetApp Deduplication for FAS, Deployment and Implementation Guide.

Using NetApp deduplication and file FlexClone not only can reduce the overall storage footprint of virtual desktops but also can improve performance by using transparent storage cache sharing. Data that is deduplicated or nonduplicated, in the case of file FlexClone data, on disk exists in the storage array cache only once per volume. All subsequent reads from any of the virtual machine disks of a block that is already in cache are read from cache and not from disk, improving performance by 10x.

Any non-deduplicated data that is not in cache must be read from disk. Data that is deduplicated but does not have as many block references as a heavily deduplicated data appear in cache only once. But based on the frequency of access might be evicted earlier than data that has many references or is heavily used. For more information on deduplication, refer to NetApp TR-3505: NetApp Deduplication for FAS Deployment and Implementation Guide.

FlexClone

NetApp FlexClone technology is hardware-assisted rapid creation of space-efficient, writable, point-in-time images of individual files, LUNs, or flexible volumes. The use of FlexClone technology in VDI deployments provides the flexibility to provision and redeploy thousands of virtual machines rapidly.

FlexClone adds a new level of agility and efficiency to storage operations. FlexClone volumes take only seconds to create and are non-disruptive to the parent FlexVol® volume or virtual machine. FlexClone copies share the same physical data space as the source and occupy negligible space (metadata) on the storage system. FlexClone file-level or volume-level clones use space very efficiently, leveraging the Data ONTAP architecture to store only data that changes between the source and clone. In addition to all these benefits, file-level or volume-level FlexClone volumes have the same high performance as other FlexVol volumes or files hosted on the volumes.

Additionally, FlexClone technology provides significant benefits with disaster recovery (DR) testing. DR testing with FlexClone is safe, risk free, and can be done during operational hours at any time. For more information on FlexClone technology concepts, refer to NetApp TR-3347: FlexClone Volumes: A Thorough Introduction.

Performance

Virtual desktops can be both read-intensive and write-intensive at different times during the lifecycle of the desktop, depending on the user activity and the desktop maintenance cycle. The performance-intensive activities are experienced by most large-scale deployments and are referred to as storm activities such as:

- Boot storms
- Login storms
- Virus scan and/or definition update storms

With physical desktops, this was not a problem as each machine had its own disks and I/O was contained within a single desktop. With VDI using a shared storage infrastructure, significant performance issues might arise during these critical operations. This essentially means the solution would require a large number of additional spindles to meet the performance requirements, resulting in an increased overall solution cost.

To solve this problem, the NetApp solution contains transparent storage cache sharing (TSCS). TSCS is a core component of Data ONTAP and is extended with Flash Cache (formerly Performance Acceleration Module II or PAM II). These solution components save customers money by:

- Requiring fewer disks and less cache
- Serving read data from cache freeing up disk I/O to perform writes
- Providing better throughput and system utilization
- Providing faster response times and a better overall end user experience

Transparent Storage Cache Sharing

Transparent storage cache sharing (TSCS) allows customers to benefit from NetApp's storage efficiency and at the same time significantly increase I/O performance. TSCS is natively built into the Data ONTAP operating system and works by using block-sharing technologies such as NetApp primary storage deduplication and file/volume FlexClone to reduce the amount of cache required and eliminate duplicate disk reads. Only one instance of any duplicate block is read into cache, thus requiring less cache than traditional storage solutions. Since VDI implementations can see as much as 99% initial space savings (validated in the NetApp solutions lab) using NetApp space-efficient cloning technologies, this translates into higher cache deduplication and high cache hit rates. TSCS is especially effective in addressing the simultaneous system boot or "boot storm" of hundreds to thousands of virtual desktop systems that can overload a traditional legacy storage system.

The following are the main benefits of transparent storage cache sharing:

- **Increased performance**: With transparent storage cache sharing in combination with FlexClone and deduplication, latencies decrease significantly by a factor of 10 versus serving data from the fastest spinning disks available, giving sub millisecond data access. Decreasing the latency results in higher throughput and lower disk utilization, which directly translate into fewer disks reads.
- **Lowering TCO**: Requiring fewer disks and getting better performance allow customers to increase the number of virtual machines on a given storage platform, resulting in a lower TCO.
-

- **Green benefits**: Power and cooling costs are reduced because the overall energy needed to run and cool the Flash Cache module is significantly less than even a single shelf of FC disks. A standard disk shelf of 300GB 15K RPM disks can consume as much as 340 watts (W)/hr and generate heat up to 1394BTU/hr. By contrast, the Flash Cache module consumes a mere 18W/hr and generates 90BTU/hr. By not deploying a single shelf, the power savings alone can be as much as 3000kWh/year per shelf. In addition to the environmental benefits of heating and cooling, you can save 3U of rack space per shelf. A real-world deployment, a NetApp solution with Flash Cache as a primary component would typically replace several such storage shelves. Therefore, the savings could be considerably higher.

NetApp Flash Cache

NetApp Flash Cache is a hardware device that extends the native Data ONTAP TSCS capabilities. Flash Cache increases the amount of available cache which helps reduce virtual desktop storm activities.

NetApp Write Optimization

Virtual desktop I/O patterns are often random in nature. Random writes are the most expensive operation for almost all RAID types because each write operation requires more than one disk operation. The ratio of VDI client operation to disk operation also depends on the RAID type for the back-end storage array. In a RAID 5 configuration on a traditional storage array, each client write operation requires up to four disk operations. A large write cache might help, but traditional storage arrays still require at least two disk operations. (Some coalescing of requests will happen if you have a big enough write cache. Also, there is a chance that one of the reads might come from the read cache.)

In a RAID 10 configuration, each client write operation requires two disk operations. The cost of RAID 10 is very high compared to RAID 5. However, RAID 5 offers lower resiliency or protection against single disk failure. (Imagine dual disk failure in the middle of the day, making hundreds to thousands of users unproductive.)

With NetApp, write operations have been optimized for RAID-DP by the core operating system Data ONTAP® and Write Anywhere File Layout (WAFL®) since their invention. NetApp arrays combine multiple client write operations and send them to disk as a single IOP. Therefore, the ratio of client operations to disk operations is always less than 1, as compared to traditional storage arrays with RAID 5 or RAID 10 which require at least 2x disk operations per client operation. Also, NetApp RAID-DP® provides the desired resiliency-or protection against dual disk failure-and performance, comparable to RAID 10 but at the cost of RAID 5.

Flexible Volumes and Aggregates

Flexible volumes (FlexVol) and aggregates provide pools of storage. This storage virtualization allows the performance and capacity to be shared by all desktops in the volume or aggregate. In much the same way that Citrix virtualizes computing resources, NetApp virtualizes storage resources.

Data Protection

The availability of thousands of virtual desktops depends on the availability of the shared storage on which the virtual desktops are hosted. Thus, using the proper RAID technology is critical. Also, being able to protect the virtual desktop images and/or user data is very important. RAID-DP®, the Citrix StorageLink virtual machine Backup and Recovery function, NetApp SnapMirror® replication technology, and NetApp Snapshot™ copies are critical components of the NetApp solution that help address storage availability.

RAID-DP

With any Citrix XenDesktop deployment, data protection is critical because any RAID failure could result in hundreds to thousands of end users being disconnected from their desktops, resulting in lost productivity. NetApp RAID DP provides performance that is comparable to that of RAID 10 yet requires

fewer disks to achieve equivalent protection. RAID DP protects against double disk failure as compared to RAID 5, which can protect against only one disk failure per RAID group. For more information about RAID DP, refer to NetApp TR-3298: RAID-DP: NetApp Implementation of RAID Double Parity for Data Protection.

NetApp FAS3240 Series Controller and DS2246 Disk Shelves Used in Test

NetApp FAS3240 series controllers and DS2246 disk shelves are ideally suited for SAN and NAS storage infrastructures that support enterprise applications such as VDI. With FAS3420 series controllers, you can take advantage of NetApp unified storage by simultaneously running SAN-based and NAS-based applications on the same storage system.

The FAS3240 chassis is 3U high and is designed for flexibility. It supports up to two controllers in a HA failover configuration. It also supports single-controller and single-controller plus I/O expansion module (IOXM) configurations. In addition to the wide array of connectivity, with PCIe expansion capabilities the FAS3240 can have additional ports added for FCoE or iSCSI. This expansion ability also enables the use of on-board memory for expanded read capabilities. This is referred to as Flash Cache. Flash Cache is offered for the FAS3240 in up to 1TB configurations.

Each DS2246 disk shelf is 2U high and supports up to 24 x 2.5 inch hard drives. It can be configured with either 450GB or 600GB 10K RPM SAS drives. This offers ideal capacity and increased IOPS per rack unit. For this test, 600GB 10K RPM SAS drives were used. The DS2246 has excellent SAS bandwidth, achieving 24GB SAS rates at 6GB x 4 per wide port.

For testing purposes, one NetApp FAS3240 single chassis HA solution was deployed. Connected to each of its controllers were four NetApp DS2246 disk shelves. The disk shelves were connected to each controller using dual-path connectivity. This allowed for higher performance in addition to failover capabilities. In addition, each FAS3240 controller was configured with 256GB of Flash cache for read optimization.

Figure 7 shows the connections between the FAS3240 Controllers and the DS2246 shelves.

Figure 7 FAS3240 and DS2246 Connections

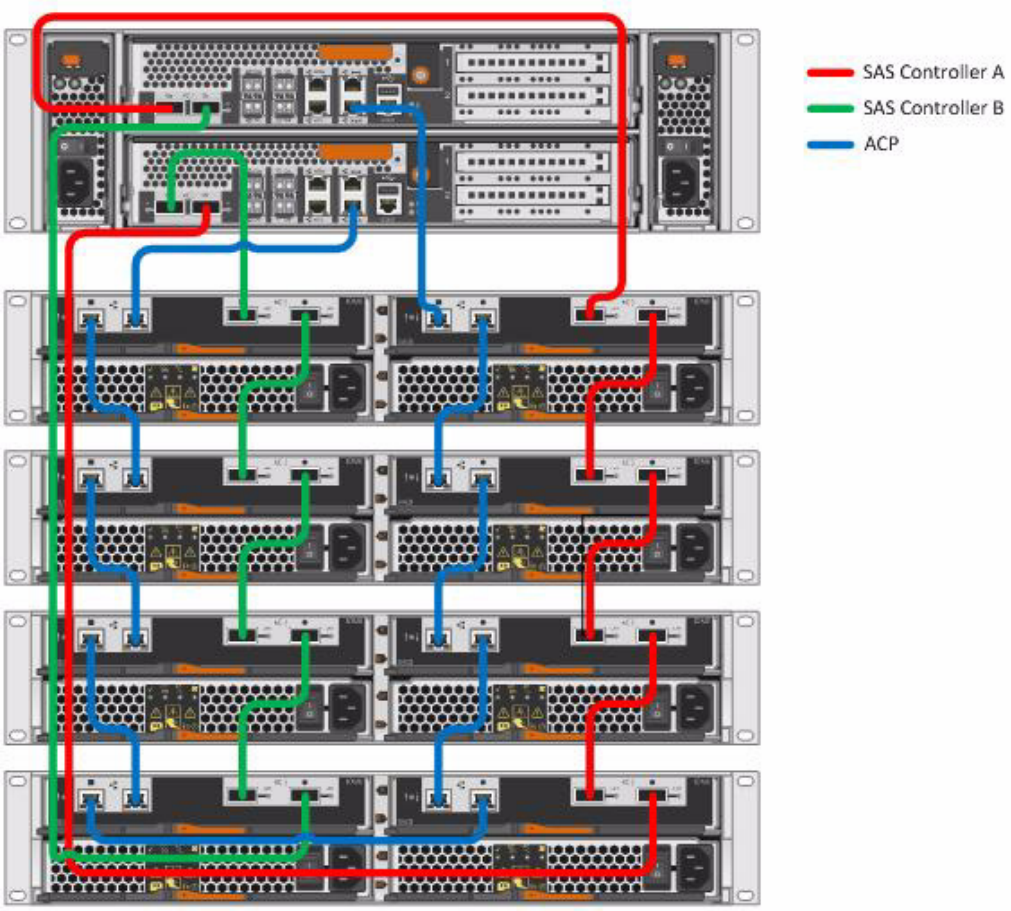

Figures 8 and 9 show the back of a NetApp FAS3240 dual-controller configuration with further details on port/slot configuration for each.

Figure 8 FAS3240 Dual-Controller Rear

Figure 9 FAS3240 Port/Slot Configuration

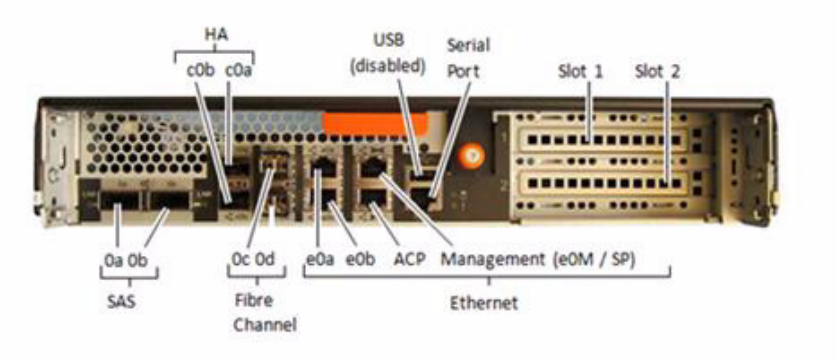

Microsoft Technologies

The VDI software solution is built on top of Microsoft Windows Server 2008 R2, Microsoft Hyper-V Server 2008 R2, System Center 2012 technologies and SQL Server 2008 R2. The Desktop Virtual Desktops deployed are Microsoft Windows 7 Enterprise.

Windows Server 2008 R2

Windows Server 2008 R2 is a multi-purpose server designed to increase the reliability and flexibility of your server or private cloud infrastructure, helping customers to save time and reduce costs. It provides powerful tools to react to business needs faster than ever before with greater control and confidence. Multiple components of Windows Server have been implemented in the VDI solution.

Active Directory

Active Directory helps customers manage corporate identities, credentials, information protection, and system and application settings. Active Directory Domain Services (AD DS) is the central location for configuration information, authentication requests, and information about all of the objects that are stored within the AD forest. Using AD, users can efficiently manage company users, computers, groups, printers, applications, and other directory-enabled objects from one secure, centralized location.

File Services

Windows Server 2008 R2 offers a cost-effective, enterprise-ready file serving platform for Windows and mixed environments. It offers an abundance of capabilities that have been requested over the years by IT organizations. Organizations can leverage the extensive file serving, data management, and data protection capabilities found in Windows Server 2008 R2.

Failover Clustering

Server availability is a higher priority than ever. The demands of a 24x7 global marketplace mean downtime can equate to lost customers, revenue, and productivity.

Windows Server 2008 brought many new or enhanced configuration, management, and diagnostic features to failover clustering that made setting up and managing the clusters easier for IT staff. Windows Server 2008 R2 builds on that work with improvements aimed at enhancing the validation

process for new or existing clusters, simplifying the management of clustered VMs which run with Hyper-V, providing a Windows PowerShell interface, and providing more options for migrating settings from one cluster to another. These enhancements provide customers with a near turnkey solution for making applications and services highly available.

Hyper-V

Hyper-V is an integral part of Windows Server and provides a foundational virtualization platform that enables customers to transition to the cloud. With Windows Server 2008 R2, customers get a compelling solution for core virtualization scenarios, including production server consolidation, dynamic datacenter, business continuity, VDI, and test and development. Hyper-V provides better flexibility with features like live migration and cluster shared volumes for storage flexibility.

Hyper-V Server

Microsoft Hyper-V Server 2008 R2 is the hypervisor-based server virtualization product that allows users to consolidate workloads onto a single physical server. It is a stand-alone product that provides a reliable and optimized virtualization solution enabling organizations to improve server utilization and reduce costs. Since Hyper-V Server is a dedicated stand-alone product containing only the Windows Hypervisor, Windows Server driver model, and virtualization components, it provides a small footprint and minimal overhead.

Deeper Dive to HyperV Features

The release of SP1 for Windows Server 2008 R2 provides new virtualization technology in Hyper-V, enabling users to deliver more advanced capabilities to the business for increased IT efficiency and agility.

Dynamic Memory

New in Windows Server 2008 R2 with SP1, Dynamic Memory enables customers to better utilize the memory resources of Hyper-V hosts by balancing how memory is distributed between running VMs. Memory can be dynamically reallocated between different VMs in response to the changing workloads of these machines. It enables more efficient use of memory while maintaining consistent workload performance and scalability. Implementing Dynamic Memory enables the achievement of higher levels of server consolidation with minimal impact on performance. Dynamic Memory also means larger numbers of virtual desktops per Hyper-V host for VDI scenarios. The net result for both scenarios is more efficient use of expensive server hardware resources, which can translate into easier management and lower costs.

Live Migration

Windows Server 2008 R2 with Hyper-V includes the much-anticipated live migration feature. Data centers with multiple Hyper-V physical hosts can move running VMs to the best physical computer for performance, scaling, or optimal consolidation without affecting users, thereby reducing costs and increasing productivity. Service and maintenance can now be done in a controlled fashion during business hours, increasing productivity for users and server administrators. Data centers can now also reduce power consumption by dynamically increasing consolidation ratios and powering off un-used physical hosts during lower demand times.

Hardware Support for Hyper-V VMs

Windows Server 2008 R2 now supports up to 64 logical processors in the host processor pool. This is a significant upgrade from previous versions and allows not only greater VM density per host but also gives IT administrators more flexibility in assigning CPU resources to VMs. Also new, Hyper-V processor compatibility mode for live migration allows for migration across different CPU versions within the same processor family (for example, "Intel Core 2-to-Intel Pentium 4" or "AMD Opteron-to-AMD Athlon"), enabling migration across a broader range of server host hardware.

Cluster Shared Volumes

With Windows Server 2008 R2, Hyper-V uses Cluster Shared Volumes (CSV) storage to simplify and enhance shared storage usage. CSV enables multiple Windows Servers to access SAN storage using a single consistent namespace for all volumes on all hosts. Multiple hosts can access the same LUN on SAN storage. CSV enables faster live migration and easier storage management for Hyper-V when used in a cluster configuration. Cluster Shared Volumes are available as part of the Windows Failover Clustering feature of Windows Server 2008 R2

Performance and Power Consumption

Hyper-V in Windows Server 2008 R2 adds enhancements that reduce VM power consumption. Hyper-V now supports Second Level Address Translation (SLAT), which uses new features on today's CPUs to improve VM performance while reducing processing load on the Windows Hypervisor. New Hyper-V VMs also consume less power due to the new Core Parking feature implemented in Windows Server 2008 R2.

Networking Support

In Windows Server 2008 R2 there are three new networking features that improve the performance of virtual networks. Support for Jumbo frames, previously available in non-virtual environments, has been extended to work with VMs. This feature enables VMs to use Jumbo Frames up to 9014 bytes if the underlying physical network supports it. Supporting Jumbo frames reduces the network stack overhead incurred per byte and increases throughput. In addition, there is a significant reduction of CPU utilization due to the lower number of calls from the network stack to the network driver. The VM Queue (VMQ) feature allows physical computer NICs to use direct memory access (DMA) to place the contents of packets directly into VM memory, increasing I/O performance.

Windows 7 Enterprise

Windows 7 is the most advanced Windows operating system to date, designed to meet the evolving needs of end users and IT professionals both in and out of the office. With Windows 7 Enterprise, users can take advantage of the following features that are not available in Windows 7 Professional:

- **DirectAccess**: Gives mobile users seamless access to corporate networks without a need to VPN.
- **BranchCache**: Decreases the time branch office users spend waiting to download files across the network.
- **Federated Search**: Finds information in remote repositories, including SharePoint sites, with a simple user interface.

- **BitLocker and BitLocker To Go**: Helps protect data on PCs and removable drives, with manageability to enforce encryption and backup of recovery keys.
- **AppLocker**: Specifies what software is allowed to run on a user's PCs through centrally managed but flexible Group Policies.
- **VDI Optimizations**: Improves user experience for VDI with multimon and microphone support, which have the ability to reuse virtual hard drive (VHD) images to boot a physical PC.

SQL Server

Microsoft SQL Server 2008 R2 includes a number of new services, including PowerPivot for Excel and SharePoint, Master Data Services, and StreamInsight. It also includes Report Builder 3.0, Reporting Services Add-in for SharePoint, a data-tier function in Visual Studio that enables packaging of tiered databases as part of an application, and Utility Control Point, which is part of Application and Multi-Server Management (AMSM) used to manage multiple SQL servers.

System Center VM Manager 2012

System Center VM Manager (VMM) 2012 helps to enable centralized management of physical and virtual IT infrastructure, increased server utilization, and dynamic resource optimization across multiple virtualization platforms. It includes end-to-end capabilities such as planning, deploying, managing, and optimizing the virtual infrastructure. VMM centrally creates and manages VMs across data centers, easily consolidates multiple physical servers onto virtual hosts, rapidly provisions and optimizes VMs, and dynamically manages virtual resources through management packs. The following is an overview of the new features offered in VMM 2012.

Fabric Management

VMM 2012 offers a robust set of new features to enable better management of the fabric supporting the virtual infrastructure.

- **Storage management**: storage arrays using the Storage Management Initiative Specification protocol (SMI-S) can be added to VMM 2012 for management. VMM then can discover, classify, and provision storage to Hyper-V hosts or clusters. The current arrays supported in VMM 2012 are:
  - NetApp FAS
  - EMC Symmetrix
  - EMC CLARiiON CX
  - HP StorageWorks Enterprise Virtual Array (EVA)
- **Bare metal provisioning**: VMM 2012 integrates with new or existing deployments of Windows Deployment Services (WDS) to deploy Windows Server 2008 R2 on bare metal machines using boot from VHD. VMM communicates with the bare metal machines via BMC. In order for VMM to communicate with the host, the host BMC must support one of the following protocols:
  - Data Center Management Interface (DCMI)
  - Systems Management Architecture for Server Hardware (SMASH)
  - Intelligent Platform Management Interface (IPMI)

- **Hyper-V cluster provisioning**: Leveraging VMM 2012's storage management features, shared storage can be provisioned to Hyper-V hosts and used to create host clusters from within VMM. With this feature and bare metal provisioning users can go from having bare metal machines to having Hyper-V clusters all done from VMM 2012.
- **Fabric updates**: VMM 2012 integrates with Windows Server Update Management server (WSUS) to manage updates to the VMM fabric servers. VMM fabric servers including Hyper-V hosts, Hyper-V clusters, library servers, Pre-Boot Execution Environment (PXE) servers, WSUS server, and the VMM management server. With host Hyper-V clusters, VMM can orchestrate the update process. If the cluster supports live migrations, running VMs will be automatically evacuated from the host being patched. If live migration is not supported, running VMs are put on saved state and brought back online once the host has been updated.
- **Multiple hypervisor management**: VMM 2012 can now manage XenServer hosts and pools with support for XenMotion. In addition to the support of XenServer, there are also improvements to VMWare ESX management. In VMM 2012 ESX hosts can now be added individually instead of importing the complete tree structure from vCenter. When importing VMWare templates to VMM, VMM no longer copies the .vmdk to the VMM library, the files remain on the ESX data store.

Private Cloud Management and Self-Service

Administrators can now create private clouds in VMM 2012. In VMM, private clouds are created from resources on premise. When creating a private cloud, administrators can choose a set of resources, such as storage and networking resources, that will support the private cloud. Administrators can then delegate those resources to self-service users to be used for creation of VMs but the details of the underlying fabric are hidden from the users. An added benefit of using private clouds is the elasticity it provides. Administrators can add resources to the cloud to increase the cloud capacity at the same time resource usage on the cloud is limited by the cloud capacity and by user role quotas.

The self-service user experience in VMM 2012 has been enhanced. Thanks to the opacity provided by the use of private clouds, self-service users can now use the VMM administrator console to create and manage their VMs on their assigned clouds.

Service Lifecycle Management

New to VMM 2012 is the concept of services. In VMM, services are a group of VMs that are configured, deployed, and managed from a single entity. Services are created from service templates and once they are deployed the service instances retain their relationship to the original service template. Administrators then can choose to make updates to the service template and those changes can be rolled out to the deployed services. In addition to deploying the VMs, applications can be installed to VMs deployed in services. VMM supports installation of Microsoft Server Application Virtualization (Server App-V) applications, Microsoft Web Deploy applications, and Microsoft SQL Server data-tier applications (DACs).

Modular Virtual Desktop Infrastructure Technical Overview

Audience

The key challenges all businesses face when considering VDIs include understanding the business drivers around VDI, understanding VDI desktop user candidate groups, understanding the virtual desktop workloads used within each group, and, perhaps most importantly, understanding where the data resides for a user group and workload pair.

Understanding these key drivers and uses allow the enterprise to analyze its VDI use cases, to select those that have the best potential for ROI, and to plan for successful VDI deployment. The analysis must identify the acceptable end user experience that will define a successful project.

Modular Architecture

Today's IT departments are facing a rapidly-evolving workplace environment. The workforce is becoming increasingly diverse and geographically distributed and includes offshore contractors, distributed call center operations, knowledge and task workers, partners, consultants, and executives connecting from locations around the globe at all times.

An increasingly mobile workforce wants to use a growing array of client computing and mobile devices that they can choose based on personal preference. These trends are increasing pressure on IT to ensure protection of corporate data and to prevent data leakage or loss through any combination of user, endpoint device, and desktop access scenarios (Figure 10). These challenges are compounded by desktop refresh cycles to accommodate aging PCs and bounded local storage and migration to new operating systems, specifically Microsoft Windows 7.

Figure 10 The Evolving Workplace Landscape

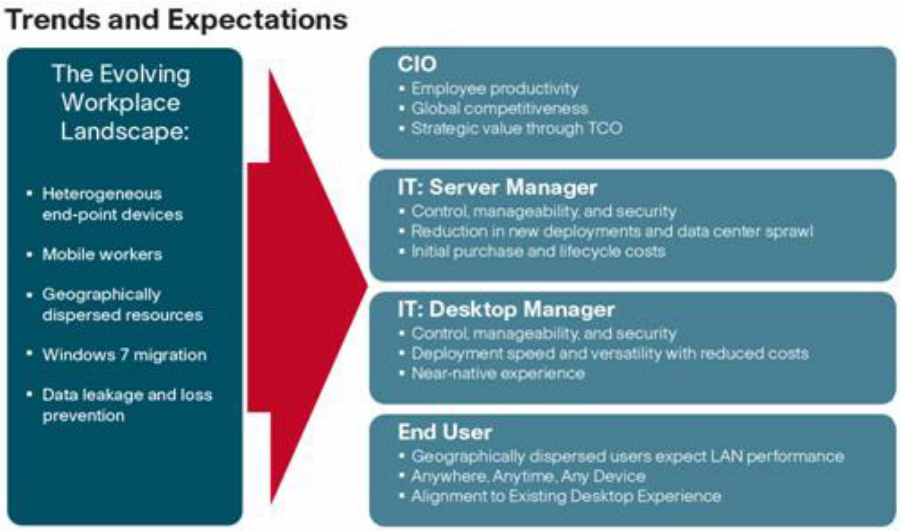

Some of the key drivers for desktop virtualization are increased data security and reduced TCO through increased control and reduced management costs.

Cisco Data Center Infrastructure for Desktop Virtualization

Cisco focuses on three key elements to deliver the best desktop virtualization data center infrastructure: simplification, security, and scalability. The software combined with platform modularity provides a simplified, secure, and scalable desktop virtualization platform (Figure 11).

Figure 11 Citrix XenDesktop on Cisco Unified Computing System

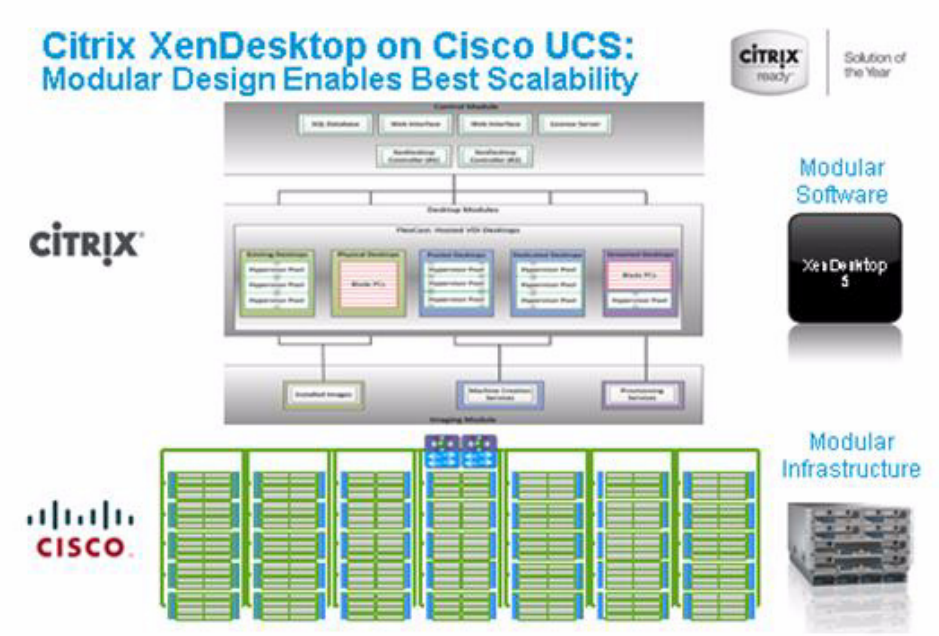

Simplified

Cisco Unified Computing System provides a radical new approach to industry standard computing and provides the heart of the data center infrastructure for desktop virtualization and the Cisco Virtualization Experience (VXI). Among the many features and benefits of Cisco Unified Computing System are the drastic reductions in the number of servers needed and number of cables per server and the ability to very quickly deploy or re-provision servers through Cisco UCS Service Profiles. With fewer servers and cables to manage and with streamlined server and virtual desktop provisioning, operations are significantly simplified .Thousands of desktops can be provisioned in minutes with Cisco Service Profiles and Cisco storage partners' storage-based cloning. This speeds time to productivity for end users, improves business agility, and allows IT resources to be allocated to other tasks.

IT tasks are further simplified through reduced management complexity, provided by the highly integrated Cisco UCS Manager, along with fewer servers, interfaces, and cables to manage and maintain. This is possible due to the industry-leading, highest virtual desktop density per blade of Cisco Unified Computing System along with the reduced cabling and port count due to the unified fabric and unified ports of Cisco Unified Computing System and desktop virtualization data center infrastructure.

Simplification also leads to improved and more rapid success of a desktop virtualization implementation. Cisco and its partners -Citrix (XenDesktop and Provisioning Server) and NetApp - have developed integrated, validated architectures, including available pre-defined, validated infrastructure packages, known as FlexPod.

Secure

While virtual desktops are inherently more secure than their physical world predecessors, they introduce new security considerations. Desktop virtualization significantly increases the need for virtual machine-level awareness of policy and security, especially given the dynamic and fluid nature of virtual machine mobility across an extended computing infrastructure. The ease with which new virtual desktops can proliferate magnifies the importance of a virtualization-aware network and security infrastructure. Cisco UCS and Nexus data center infrastructure for desktop virtualization provides

stronger data center, network, and desktop security with comprehensive security from the desktop to the hypervisor. Security is enhanced with segmentation of virtual desktops, virtual machine-aware policies and administration, and network security across the LAN and WAN infrastructure.

Scalable

Growth of a desktop virtualization solution is all but inevitable and it is critical to have a solution that can scale predictably with that growth. The Cisco solution supports more virtual desktops per server and additional servers scale with near linear performance. Cisco data center infrastructure provides a flexible platform for growth and improves business agility. Cisco UCS Service Profiles allow for on-demand desktop provisioning, making it easy to deploy dozens or thousands of additional desktops.

Each additional Cisco Unified Computing System server provides near linear performance and utilizes Cisco's dense memory servers and unified fabric to avoid desktop virtualization bottlenecks. The high performance, low latency network supports high volumes of virtual desktop traffic, including high resolution video and communications.

Cisco Unified Computing System and Nexus data center infrastructure is an ideal platform for growth, with transparent scaling of server, network, and storage resources to support desktop virtualization.

Savings and Success

As demonstrated above, the simplified, secure, scalable Cisco data center infrastructure solution for desktop virtualization will save time and cost. There will be faster payback, better ROI, and lower TCO with the industry's highest virtual desktop density per server, meaning there will be fewer servers needed, reducing both capital expenditures (CapEx) and operating expenditures (OpEx). There will also be much lower network infrastructure costs, with fewer cables per server and fewer ports required, via the Cisco UCS architecture and unified fabric.

The simplified deployment of Cisco Unified Computing System for desktop virtualization speeds up time to productivity and enhances business agility. IT staff and end users are more productive more quickly and the business can react to new opportunities by simply deploying virtual desktops whenever and wherever they are needed. The high performance Cisco systems and network deliver a near-native end-user experience, allowing users to be productive anytime, anywhere.

Solution Components: Cisco, Citrix, and NetApp Reference Architecture

Cisco's desktop virtualization solution is the first in the industry to bind together the three critical elements of an end-to-end deployment: the end-user, the network, and the data center. It draws on Cisco's architectural advantage to provide a solution that supports a diversity of endpoint devices, extends pervasive security and policy management to each virtual desktop, and uses Cisco Unified Computing System—a new and innovative virtualization-optimized stateless server computing model.

Base Components

The Cisco UCS computing platform includes:

- Cisco UCS 6200 Series Fabric Interconnects
- Cisco UCS 2100 or 2200 Series IO Modules
- Cisco UCS 5108 Blade Chassis
- Cisco UCS B230 M2 Blade Servers for virtual desktop hosting
- Cisco UCS B200 M2 Blade Servers for infrastructure

Other key components include:

- Access Layer Switch: Cisco Nexus 5500 Series switches
- Storage System: NetApp 3240A Storage System with block and iSCSI storage and Flash Cache
- Hypervisor: Microsoft Hyper-V 2008 R2
- Virtual Desktop Connection Broker: Citrix XenDesktop 5.5 with Provisioning Server 5.6 SP1

Understanding Desktop User Groups

There must be a considerable effort within the enterprise to identify desktop user groups and their memberships. The most broadly recognized, high level user groups are:

- **Task Workers**—Groups of users working in highly specialized environments where the number of tasks performed by each worker is essentially identical. These users are typically located at a corporate facility (for example, call center employees).
- **Knowledge/Office Workers**—Groups of users who use a relatively diverse set of applications that are Web-based and installed and whose data is regularly accessed. They typically have several applications running simultaneously throughout their workday and a requirement to utilize Flash video for business purposes. This is not a singular group within an organization. These workers are typically located at a corporate office (for example, workers in accounting groups).
- **Power Users**—Groups of users who run high-end, memory, processor, disk IO, and/or graphic-intensive applications, often simultaneously. These users have high requirements for reliability, speed, and real-time data access (for example, design engineers).
- **Mobile Workers**—Groups of users who may share common traits with Knowledge/Office Workers, with the added complexity of needing to access applications and data from wherever they are whether at a remote corporate facility, customer location, at the airport, at a coffee shop, or at home all in the same day (for example, a company's outbound sales force).
- **Remote Workers**—Groups of users who could fall into the Task Worker or Knowledge/Office Worker groups but whose experience is from a remote site that is not corporate owned, most often from the user's home. This scenario introduces several challenges in terms of type, available bandwidth, and latency and reliability of the user's connectivity to the data center (for example, a work-from-home accounts payable representative).
- **Guest/Contract Workers**—Groups of users who need access to a limited number of carefully controlled enterprise applications and data and resources for short periods of time. These workers may need access from the corporate LAN or remote access (for example, a medical data transcriptionist).

There is good reason to search for and identify multiple sub-groups of the major groups listed above in the enterprise. Typically, each sub-group has different application and data requirements.

Understanding Applications and Data

When the desktop user groups and sub-groups have been identified, the next task is to catalog group application and data requirements. This can be one of the most time-consuming processes in the VDI planning exercise, but is essential for the VDI project's success. If the applications and data are not identified and co-located, performance will be negatively affected.

The process of analyzing the variety of application and data pairs for an organization will likely be complicated by the inclusion cloud applications, like SalesForce.com. This application and data analysis is beyond the scope of this Cisco Validated Design, but should not be omitted from the planning process. There are a variety of third party tools available to assist organizations with this crucial exercise.

Project Planning and Solution Sizing Sample Questions

Now that user groups, their applications and their data requirements are understood, some key project and solution sizing questions may be considered.

General project questions should be addressed at the outset, including:

- Has a VDI pilot plan been created based on the business analysis of the desktop groups, applications and data?
- Is there infrastructure and budget in place to run the pilot program?
- Are the required skill sets to execute the VDI project available? Can we hire or contract for them?
- Do we have end user experience performance metrics identified for each desktop sub-group?
- How will we measure success or failure?
- What is the future implication of success or failure?

Provided below is a short, non-exhaustive list of sizing questions that should be addressed for each user sub-group:

- What is the desktop OS planned? Windows 7 or Windows XP?
- 32 bit or 64 bit desktop OS?
- How many virtual desktops will be deployed in the pilot? In production? All Windows 7?
- How much memory per target desktop group desktop?
- Are there any rich media, Flash, or graphics-intensive workloads?
- What is the end point graphics processing capability?
- Will XenApp be used for Hosted Shared Server Desktops or exclusively XenDesktop?
- Are there XenApp hosted applications planned? Are they packaged or installed?
- Will Provisioning Server or Machine Creation Services be used for virtual desktop deployment?
- What is the hypervisor for the solution?
- What is the storage configuration in the existing environment?
- Are there sufficient IOPS available for the write-intensive VDI workload?
- Will there be storage dedicated and tuned for VDI service?
- Is there a voice component to the desktop?
- Is anti-virus a part of the image?
- Is user profile management (e.g., non-roaming profile based) part of the solution?
- What is the fault tolerance, failover, disaster recovery plan?
- Are there additional desktop sub-group specific questions?

Cisco Services

Cisco offers assistance for customers in the analysis, planning, implementation, and support phases of the VDI lifecycle. These services are provided by the Cisco Advanced Services group. Some examples of Cisco services include:

- Cisco VXI Unified Solution Support
- Cisco VXI Desktop Virtualization Strategy Service

- Cisco VXI Desktop Virtualization Planning and Design Service

The Solution: A Unified, Pre-Tested and Validated Infrastructure

To meet the challenges of designing and implementing a modular desktop infrastructure, Cisco, Citrix, NetApp and Microsoft have collaborated to create the data center solution for virtual desktops outlined in this document.

Key elements of the solution include:

- A shared infrastructure that can scale easily
- A shared infrastructure that can accommodate a variety of virtual desktop workloads

Cisco Networking Infrastructure

This section describes the Cisco networking infrastructure components used in the configuration.

Cisco Nexus 5548 Switch

The Cisco Nexus 5548 Switch is a 1RU, 10 Gigabit Ethernet, FCoE access-layer switch built to provide more than 500 Gbps throughput with very low latency. It has 20 fixed 10 Gigabit Ethernet and FCoE ports that accept modules and cables meeting the Small Form-Factor Pluggable Plus (SFP+) form factor. One expansion module slot can be configured to support up to six additional 10 Gigabit Ethernet and FCoE ports, up to eight FC ports, or a combination of both. The switch has a single serial console port and a single out-of-band 10/100/1000-Mbps Ethernet management port. Two N+1 redundant, hot-pluggable power supplies and five N+1 redundant, hot-pluggable fan modules provide highly reliable front-to-back cooling.

Cisco Nexus 5500 Series Feature Highlights

The switch family's rich feature set makes the series ideal for rack-level, access-layer applications. It protects investments in data center racks with standards-based Ethernet and FCoE features that allow IT departments to consolidate networks based on their own requirements and timing.

- The combination of high port density, wire-speed performance, and extremely low latency makes the switch an ideal product to meet the growing demand for 10 Gigabit Ethernet at the rack level. The switch family has sufficient port density to support single or multiple racks fully populated with blade and rack-mount servers.
- Built for today's data centers, the switches are designed just like the servers they support. Ports and power connections are at the rear, closer to server ports, helping keep cable lengths as short and efficient as possible. Hot-swappable power and cooling modules can be accessed from the front panel, where status lights offer an at-a-glance view of switch operation. Front-to-back cooling is consistent with server designs, supporting efficient data center hot-aisle and cold-aisle designs. Serviceability is enhanced with all customer replaceable units accessible from the front panel. The use of SFP+ ports offers increased flexibility to use a range of interconnect solutions, including copper for short runs and fibre for long runs.
- FCoE and IEEE data center bridging features support I/O consolidation, ease management of multiple traffic flows, and optimize performance. Although implementing SAN consolidation requires only the lossless fabric provided by the Ethernet pause mechanism, the Cisco Nexus 5500 Series switches provide additional features that create an even more easily managed, high-performance, unified network fabric.

Features and Benefits

This section details the specific features and benefits provided by the Cisco Nexus 5500 Series.

10GB Ethernet, FCoE, and Unified Fabric Features

The Cisco Nexus 5500 Series is first and foremost a family of outstanding access switches for 10 Gigabit Ethernet connectivity. Most of the features on the switches are designed for high performance with 10 Gigabit Ethernet. The Cisco Nexus 5500 Series also supports FCoE on each 10 Gigabit Ethernet port that can be used to implement a unified data center fabric, consolidating LAN, SAN, and server clustering traffic.

Low Latency

The cut-through switching technology used in the Cisco Nexus 5500 Series ASICs enables the product to offer a low latency of 3.2 microseconds, which remains constant regardless of the size of the packet being switched. This latency was measured on fully configured interfaces, with access control lists (ACLs), QoS, and all other data path features turned on. The low latency on the Cisco Nexus 5500 Series enables application-to-application latency on the order of 10 microseconds (depending on the NIC). These numbers, together with the congestion management features described in the next section, make the Cisco Nexus 5500 Series a great choice for latency-sensitive environments.

Other features include: Nonblocking Line-Rate Performance, Single-Stage Fabric, Congestion Management, Virtual Output Queues, Lossless Ethernet (Priority Flow Control), Delayed Drop FC over Ethernet, Hardware-Level I/O Consolidation, and End-Port Virtualization.

Architecture and Design of XenDesktop 5.6 on Cisco Unified Computing System and NetApp FAS Storage

Design Fundamentals

There are many reasons to consider a virtual desktop solution such as an ever growing and diverse base of user devices, complexity in management of traditional desktops, security, and even Bring Your Own Computer (BYOC) to work programs. The first step in designing a virtual desktop solution is to understand the user community and the type of tasks that are required to successfully execute their role. The following user classifications is provided:

- **Knowledge Workers** today do not just work in their offices all day—they attend meetings, visit branch offices, work from home, and even coffee shops. These anywhere workers expect access to all of their same applications and data wherever they are.
- **External Contractors** are increasingly part of your everyday business. They need access to certain portions of your applications and data, yet administrators still have little control over the devices they use and the locations they work from. Consequently, IT is stuck making trade-offs on the cost of providing these workers a device vs. the security risk of allowing them access from their own devices.
- **Task Workers** perform a set of well-defined tasks. These workers access a small set of applications and have limited requirements from their PCs. However, since these workers are interacting with your customers, partners, and employees, they have access to your most critical data.

- **Mobile Workers** need access to their virtual desktop from everywhere, regardless of their ability to connect to a network. In addition, these workers expect the ability to personalize their PCs, by installing their own applications and storing their own data, such as photos and music, on these devices.
- **Shared Workstation** users are often found in state-of-the-art university and business computer labs, conference rooms or training centers. Shared workstation environments have the constant requirement to re-provision desktops with the latest operating systems and applications as the needs of the organization change, tops the list.

After the user classifications have been identified and the business requirements for each user classification have been defined, it becomes essential to evaluate the types of virtual desktops that are needed based on user requirements. There are essentially five potential desktops environments for each user:

- **Traditional PC**: A traditional PC is what typically constituted a desktop environment, a physical device with a locally installed operating system.
- **Hosted Shared Desktop**: A hosted, server-based desktop is a desktop where the user interacts through a delivery protocol. With hosted, server-based desktops, a single installed instance of a server operating system, such as Microsoft Windows Server 2008 R2, is shared by multiple users simultaneously. Each user receives a desktop "session" and works in an isolated memory space. Changes made by one user could impact the other users.
- **Hosted Virtual Desktop**: A hosted virtual desktop is a virtual desktop running either on virtualization layer (XenServer, Hyper-V or ESX) or on bare metal hardware. The user does not work with and sit in front of the desktop, but instead the user interacts through a delivery protocol.
- **Streamed Applications**: Streamed desktops and applications run entirely on the user's local client device and are sent from a server on demand. The user interacts with the application or desktop directly but the resources may only available while they are connected to the network.
- **Local Virtual Desktop**: A local virtual desktop is a desktop running entirely on the user's local device and continues to operate when disconnected from the network. In this case, the user's local device is used as a type 1 hypervisor and is synced with the data center when the device is connected to the network.

For the purposes of the validation represented in this document only hosted virtual desktops were validated. Each of the sections provides some fundamental design decisions for this environment.

Hosted VDI Design Fundamentals

Citrix XenDesktop 5.6 can be used to deliver a variety of virtual desktop configurations. When evaluating a Hosted VDI deployment, consider the following:

- Hypervisor Selection
- Provisioning Services
- System Center Virtual Machine Manager 2012

Hypervisor Selection

Citrix XenDesktop is hypervisor agnostic, so any of the following three hypervisors can be used to hosted VDI-based desktops:

- **Hyper-V**: Microsoft Windows Server 2008 R2 Hyper-V builds on the architecture and functions of Windows Server 2008 Hyper-V by adding multiple new features that enhance product flexibility. Hyper-V is available in a Standard, Server Core and free Hyper-V Server 2008 R2 versions. More information on Hyper-V can be obtained at the company Web site.
- **vShpere**: VMware vSphere consists of the management infrastructure or virtual center server software and the hypervisor software that virtualizes the hardware resources on the servers. It offers features like Distributed Resource Scheduler, vMotion, high availability, Storage vMotion, VMFS, and a mutlipathing storage layer. More information on vSphere can be obtained at the company Web site.
- **XenServer**: Citrix® XenServer® is a complete, managed server virtualization platform built on the powerful Xen® hypervisor. Xen technology is widely acknowledged as the fastest and most secure virtualization software in the industry. XenServer is designed for efficient management of Windows® and Linux® virtual servers and delivers cost-effective server consolidation and business continuity. More information on Hyper-V can be obtained at the company Web site.

For this study, we utilized Microsoft Windows Server 2008 R2 SP1 with Hyper-V, Hyper-V Server 2008 R2 SP1 and System Center Virtual Machine Manager 2012 as the hypervisor.

Provisioning Services

Hosted-VDI desktops can be deployed with and without Citrix Provisioning Services, but Citrix Provisioning Services enables you to stream a single desktop image to create multiple virtual desktops on one or more servers in a data center. This facility greatly reduces the amount of storage required compared to other methods of creating virtual desktops. Citrix Provisioning Services desktops can be deployed as Pooled or Private:

- **Private Desktop**: A private desktop is a single private desktop for assigned to one distinct user.
- **Pooled Desktop**: A pooled virtual desktop uses Citrix Provisioning Services to stream a standard desktop image to multiple desktop instances upon boot-up.

When considering a Provisioning Services deployment, there are some design decisions that need to be made regarding the write-cache for the virtual desktop device leveraging provisioning. The write-cache is a cache of all data that the target device has written. If data is written to the Provisioning Server vDisk in a caching mode, the data is not written back to the base vDisk. Instead it is written to a write-cache file in one of the locations specified below. The following options exist for the Provisioning Services write cache:

- **Cache on device HD**: Cache on local HD is stored in a file on a secondary local hard drive of the device. It gets created as an invisible file in the root folder of the local HD. The Cache file size grows as needed, but never gets larger than the original vDisk, and frequently not larger than the free space on the original vDisk.
- **Cache in device RAM**: Cache is stored in client RAM (Memory), The Cache maximum size is fixed by a setting in vDisk properties. All written data can be read from local RAM instead of going back to server.RAM Cache is faster than server cache and works in a high availability environment.
- **Cache on server**: Server Cache is stored in a file on the server, or on a share, SAN, or other. The file size grows as needed, but never gets larger than the original vDisk, and frequently not larger than the free space on the original vDisk. It is slower than RAM cache because all reads/writes have to go to the server and be read from a file. Cache gets deleted when the device reboots, in other words, on every boot the device reverts to the base image. Changes remain only during a single boot session.

- **Cache on device hard drive persisted**: (Experimental Phase) The same as Cache on device hard drive, except cache persists. At this time, this write cache method is an experimental feature only, and is only supported for NT6.1 or later (Windows 7 and Windows 2008 R2 and later). This method also requires a different bootstrap.
- **Cache on server persisted**: This cache option allows for the saving of changes between reboots. Using this option, after rebooting, a target device is able to retrieve changes made from previous sessions that differ from the read only vDisk image. If a vDisk is set to Cache on server persistent, each target device that accesses the vDisk automatically has a device-specific, writable disk file created. Any changes made to the vDisk image are written to that file, which is not automatically deleted upon shutdown.

The alternative to Citrix Provisioning Services for pooled desktop deployments is Citrix Machine Creation Services, which is integrated directly with the XenDesktop Studio console.

For this study, we used Provisioning Server 6.1 for management of Pooled Desktops with Cache on device HD of each virtual machine. Provisioning Server 6.1 was used for Active Directory machine account creation and management as well as for streaming the shared disk to the hypervisor hosts.

System Center 2012 Virtual Machine Manager

System Center 2012 Virtual Machine Manager (SCVMM) is a management solution for the virtualized datacenter, enabling you to configure and manage your virtualization host, networking, and storage resources in order to create and deploy virtual machines and services to private clouds that you have created.

Microsoft System Center 2012 cloud and datacenter management solutions empower you with a common management toolset for your private and public cloud applications and services.

SCVMM is an integral part of the System Center 2012 Application Management component.

Figure 12 Microsoft System Center 2012 Management Suite

We utilized System Center 2012 VMM for the VDI deployment outlined in this study.

Designing a Citrix XenDesktop 5.6 Deployment

To implement our pooled desktop delivery model for this study, known as Hosted VDI Pooled Desktops, we followed the Citrix Reference Architecture for local desktop delivery.

Figure 13 Pooled Desktop Infrastructure

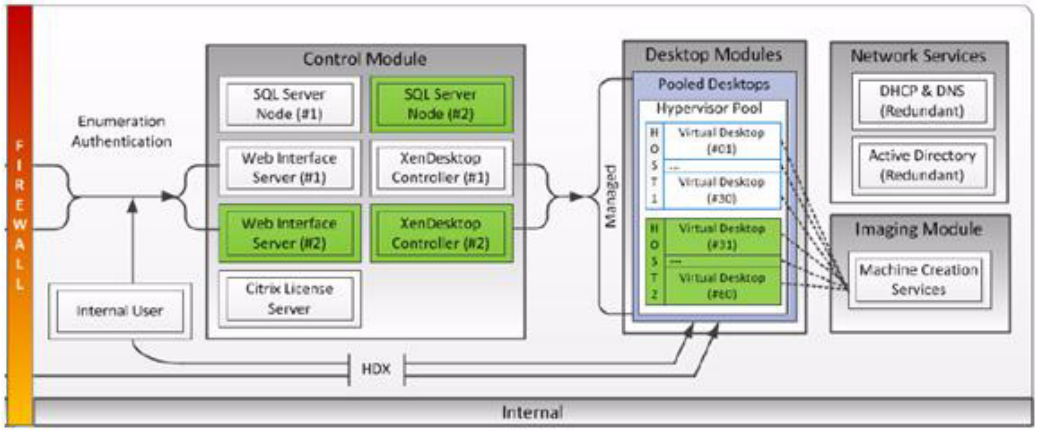

We elected to use Provisioning Services 6.1 in place of Machine Creation Services for this study so that our design would scale to many thousands of desktops.

Learn more about Citrix's XenDesktop Reference Architecture - Pooled Desktops (Local and Remote) here. Learn more about XenDesktop 5.6 Planning and Design here.

Storage Architecture Design

In a VDI environment, the availability and performance of the storage infrastructure are very critical because thousands of users will be affected by storage outages or performance issues. Thus the storage architecture must provide the level of availability and performance typical for business-critical applications. NetApp has all the software and hardware solutions to address availability and performance for large, scalable VDI environments.

Storage System Configuration Best Practices

This section is a high-level overview of the components and features to consider when deploying a Citrix XenDesktop infrastructure on NetApp. For detailed information on storage resiliency, refer to the following:

- NetApp TR-3437: Storage Best Practices and Resiliency Guide
- NetApp TR-3450: Active-Active Controller Overview and Best Practices Guidelines

Building a Resilient Storage Architecture

Active NetApp Controllers: The controller in a storage system can be a single point of failure if not designed correctly. Active-active controllers provide controller redundancy and simple automatic transparent failover in the event of a controller failure to deliver enterprise-class availability. Providing transparent recovery from component failure is critical because all desktops rely on the shared storage.

Multipath High Availability (HA): Multipath HA storage configuration further enhances the resiliency and performance of active-active controller configurations. Multipath HA-configured storage enhances storage resiliency by reducing unnecessary takeover by a partner node due to a storage fault, improving overall system availability and promoting higher performance consistency. Multipath HA provides added protection against various storage faults, including HBA or port failure, controller-to-shelf cable failure, shelf module failure, dual intershelf cable failure, and secondary path failure. Multipath HA helps provide consistent performance in active-active configurations by providing larger aggregate storage loop bandwidth.

RAID Data Protection: Data protection against disk drive failure using RAID is a standard feature of most shared storage devices, but with the capacity and subsequent rebuild times of current hard drives where exposure to another drive failure can be catastrophic, protection against double disk failure, is now essential. NetApp RAID-DP is an advanced RAID technology that is provided as the default RAID level on all FAS systems. RAID-DP offers performance that is comparable to that of RAID 10, with much higher resiliency. It provides protection against double disk failure as compared to RAID 5, which can protect against only one disk failure. NetApp strongly recommends using RAID-DP on all RAID groups that store Citrix XenDesktop data.

Remote LAN Management (RLM) Card: The RLM card improves storage system monitoring by providing secure out-of-band access to the storage controllers, which can be used regardless of the state of the controllers. The RLM offers a number of remote management capabilities for NetApp controllers, including remote access, monitoring, troubleshooting, logging, and alerting features. The RLM also extends AutoSupport™ capabilities of the NetApp controllers by sending alerts or "down storage system" notification with an AutoSupport message when the controller goes down, regardless of whether the controller can send AutoSupport messages. These AutoSupport messages also send proactive alerts to NetApp to help provide faster service.

Networking infrastructure design (FCoE, Fibre Channel, or IP): A network infrastructure (FCoE, Fibre Channel, or IP) should have no single point of failure. A highly available solution includes having two or more Fibre Channel, FCoE or IP network switches; two or more CNAs, HBAs, or NICs per host; and two or more target ports or NICs per storage controller. In addition, if using Fibre Channel, two independent fabrics are required to have a truly redundant architecture.

Top Resiliency Practices

- Use RAID-DP, the NetApp high-performance implementation of RAID 6, for better data protection.
- Use multipath HA with active-active storage configurations to improve overall system availability as well as promote higher performance consistency.
- Use the default RAID group size (16) when creating aggregates.
- Allow Data ONTAP to select disks automatically when creating aggregates or volumes.
- Use the latest Data ONTAP general availability release available on the NetApp Support site (formerly NOW®).
- Use the latest storage controller, shelf, and disk firmware available on the NetApp Support site.
- Disk drive differences are Fibre Channel, SAS, SATA disk drive types, disk size, and rotational speed (RPM).
- Maintain two hot spares for each type of disk drive in the storage system to take advantage of Maintenance Center.
- Do not put user data into the root volume (vol0) due to lack of disk spindles.
- Replicate data with NetApp SnapMirror replication technology or NetApp SnapVault® backup software for disaster recovery (DR) protection.
- Replicate to remote locations to increase data protection levels.

- Use an active-active storage controller configuration (clustered failover) to eliminate single points of failure (SPOFs).
- Deploy NetApp SyncMirror® synchronous mirroring software and RAID-DP for the highest level of storage resiliency.

For more information, refer to NetApp TR-3437: Storage Best Practices and Resiliency Guide.

Building a High-Performance Storage Architecture

A VDI workload can be very I/O intensive, especially during the simultaneous boot up, login, and virus scan within the virtual desktops. A boot storm, depending on how many servers and guests are attached to the storage, can create a significant performance effect if the storage is not sized properly. A boot storm can affect both the speed in which the desktops are available to the customer and overall customer experience. A "virus scan storm" is similar to a boot storm in I/O but might last longer and can significantly affect customer experience.

Due to these factors, it is important to make sure that the storage is architected in such a way as to eliminate or decrease the effect of these events.

- **Aggregate sizing**. An aggregate is NetApp's virtualization layer, which abstracts physical disks from logical datasets, which are referred to as flexible volumes. Aggregates are the means by which the total IOPS available to all of the physical disks are pooled as a resource. This design is well suited to meet the needs of an unpredictable and mixed workload. NetApp recommends that whenever possible a small aggregate should be used as the root aggregate. This root aggregate stores the files required for running and providing GUI management tools for the storage system. The remaining storage should be placed into a small number of large aggregates. The overall disk I/O from virtualization environments is traditionally random by nature, so this storage design gives optimal performance because a large number of physical spindles are available to service I/O requests. On smaller storage systems, it might not be practical to have more than a single aggregate, due to the restricted number of disk drives on the system. In these cases, it is acceptable to have only a single aggregate.
- **Disk configuration summary**. When sizing your disk solution, consider the number of desktops being served by the storage controller/disk system and the number of IOPS per desktop. This way you can make a calculation to arrive at the number and size of the disks needed to serve the given workload. Remember, keep the aggregates large, spindle count high, and rotational speed fast. When one factor needs to be adjusted, NetApp Flash Cache can help eliminate potential bottlenecks to the disk.
- **Flexible Volumes**. Flexible volumes contain either LUNs or virtual disk files that are accessed by Citrix XenDesktop servers. NetApp recommends a one-to-one alignment of Citrix XenDesktop datastores to flexible volumes. This design offers an easy means to understand the Citrix XenDesktop data layout when viewing the storage configuration from the storage system. This mapping model also makes it easy to implement Snapshot backups and SnapMirror replication policies at the datastore level, because NetApp implements these storage side features at the flexible volume level.
- **Flash Cache**. Flash Cache enables transparent storage cache sharing and improves read performance and in turn increases throughput and decreases latency. It provides greater system scalability by removing IOPS limitations due to disk bottlenecks and lowers cost by providing the equivalent performance with fewer disks. Using Flash Cache in a dense (deduplicated) volume allows all the shared blocks to be accessed directly from the intelligent, faster Flash Cache versus disk. Flash Cache provides great benefits in Citrix XenDesktop environments, especially during a

boot storm, login storm, or virus storm, because only one copy of deduplicated data needs to be read from the disk (per volume). Each subsequent access of a shared block is read from Flash Cache and not from disk, increasing performance and decreasing latency and overall disk utilization.

Solution Validation

This section details the configuration and tuning that was performed on the individual components to produce a complete, validated solution.

Configuration Topology for Scalable Citrix XenDesktop 5.6 Virtual Desktop Infrastructure on Cisco Unified Computing System and NetApp FAS Storage System

Figure 14 Architecture Block Diagram

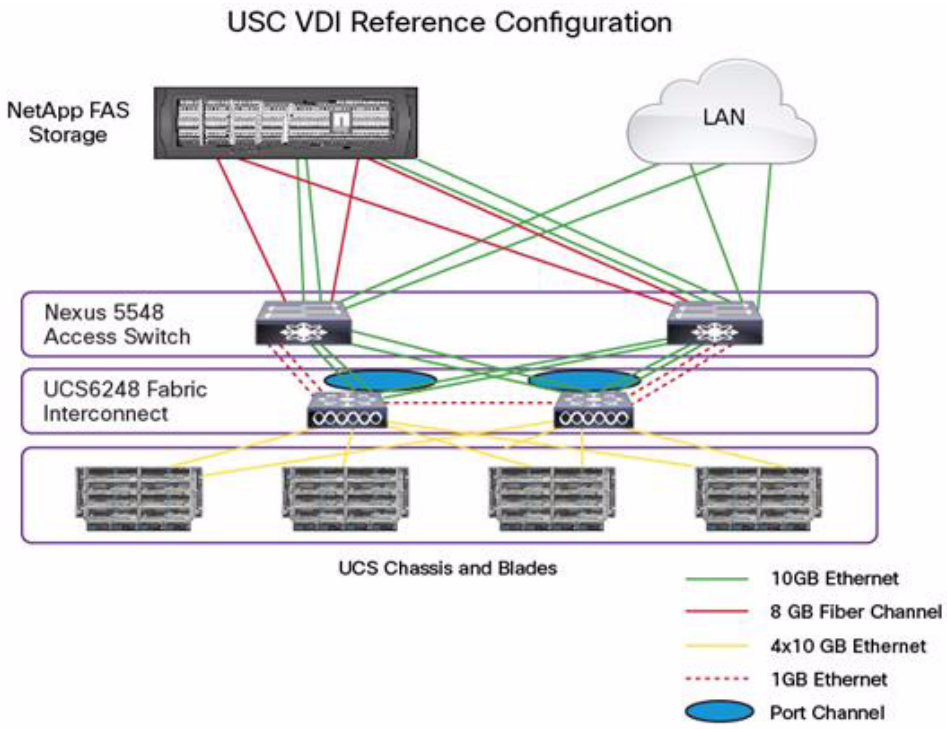

The figure above captures the architectural diagram for the purpose of this study. The architecture is divided into four distinct layers:

- Cisco UCS Compute Platform
- The VDI that runs on Cisco UCS blade hypervisor hosts
- Network access layer and LAN
- Storage Area Network (SAN) and NetApp FAS System storage

Figure 15 Detailed Architecture of the Configuration

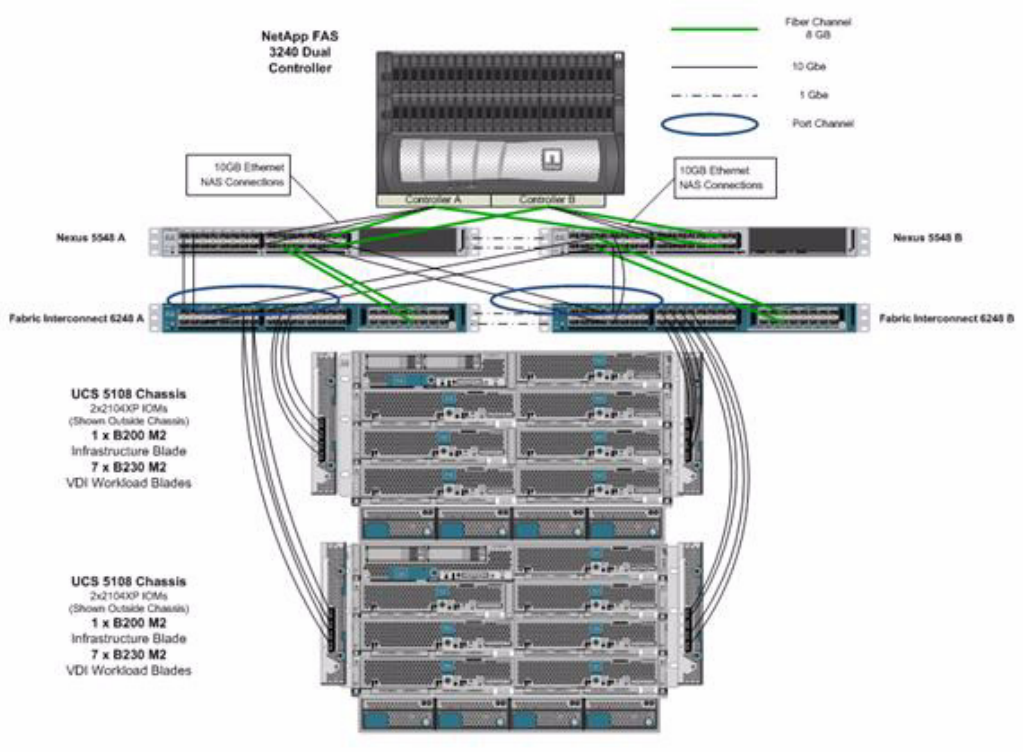

Cisco UCS Configuration

This section talks about the Cisco UCS configuration that was done as part of the infrastructure build out. The racking, power and installation of the chassis are described in the Installation Guide and it is beyond the scope of this document. More details on each step can be found in the following documents:

- Cisco UCS CLI Configuration Guide
- Cisco UCS-M GUI Configuration Guide

Base Cisco UCS System Configuration

To configure the Cisco Unified Computing System, perform the following steps.

1. Bring up the Fabric interconnect and from a Serial Console connection set the IP address, gateway, and the hostname of the primary fabric interconnect. Now bring up the second fabric interconnect after connecting the dual cables between them. The second fabric interconnect automatically recognizes the primary and ask if you want to be part of the cluster, answer yes and set the IP address, gateway and the hostname. When this is done all access to the FI can be done remotely. You will also configure the virtual IP address to connect to the FI, you need a total of three IP address to bring it online. You can also wire up the chassis to the FI, using either 1, 2 or 4 links per IO Module, depending on your application bandwidth requirement. We connected all the four links to each module.
2. Now connect using a browser to the Virtual IP and launch the Cisco UCS Manager. The Java based Cisco UCS Manager will let you do everything that you could do from the CLI and we will highlight the GUI methodology here.

3. Check the firmware on the system and see if it is current. The firmware release used in this document is 2.0(1w).

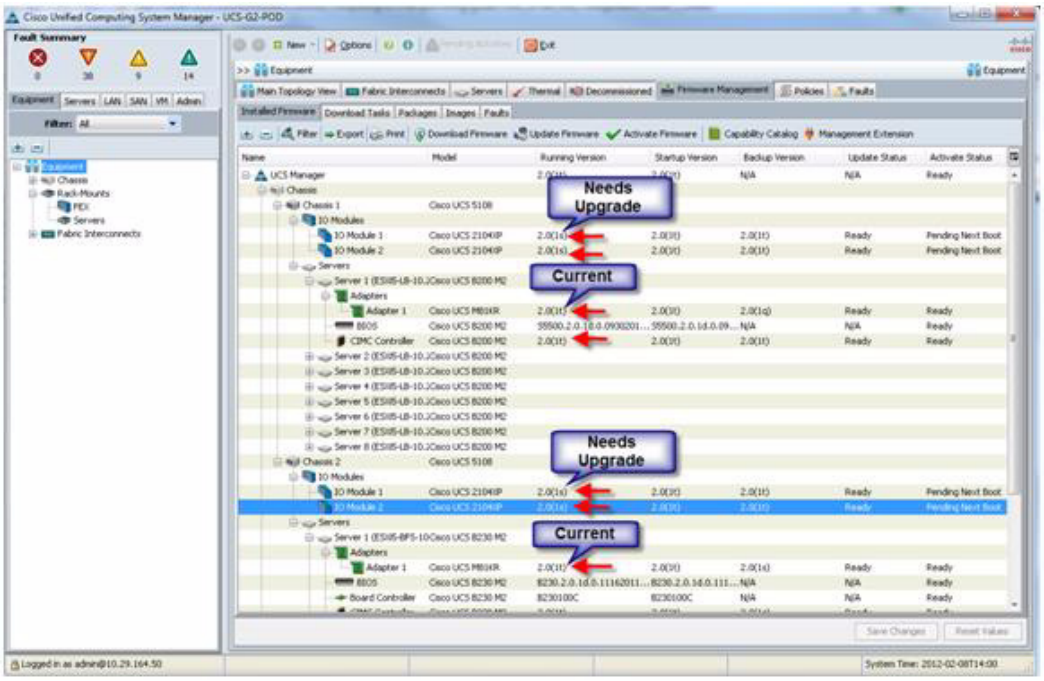

If the firmware is not current, follow the installation and upgrade guide to upgrade the Cisco UCS firmware. Also do not forget to upgrade the BIOS to the latest level and associate it with all the blades. Note: The Bios and Board Controller version numbers do not track the IO Module, Adapter, nor CIMC controller version number.

4. Configure and enable the server port on the FI. To bring the chassis online acknowledge the chassis.

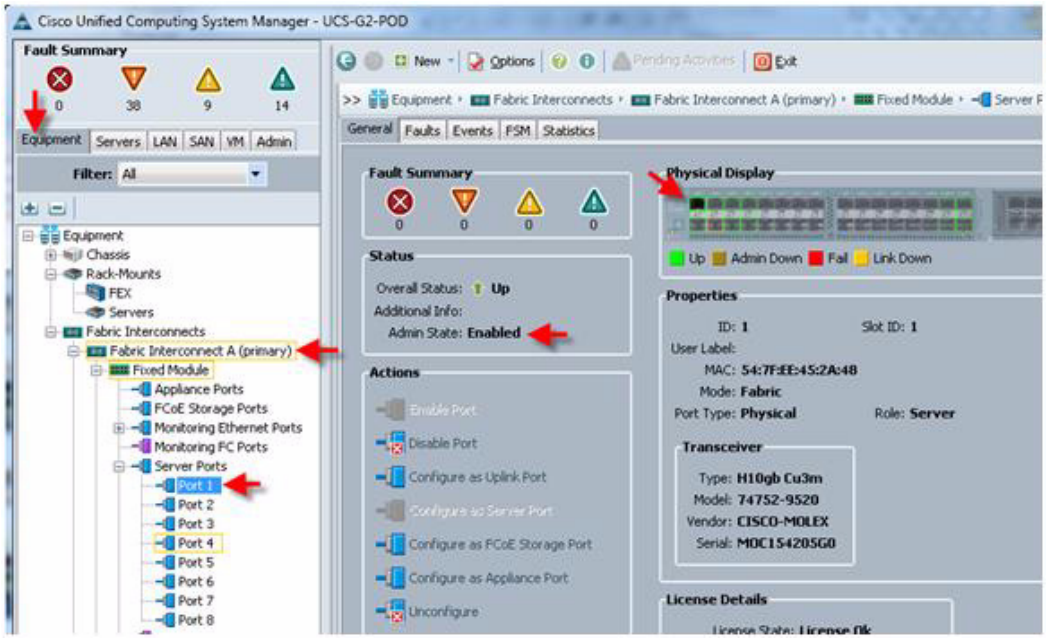

5. Configure and enable upstream Ethernet links and FC links.

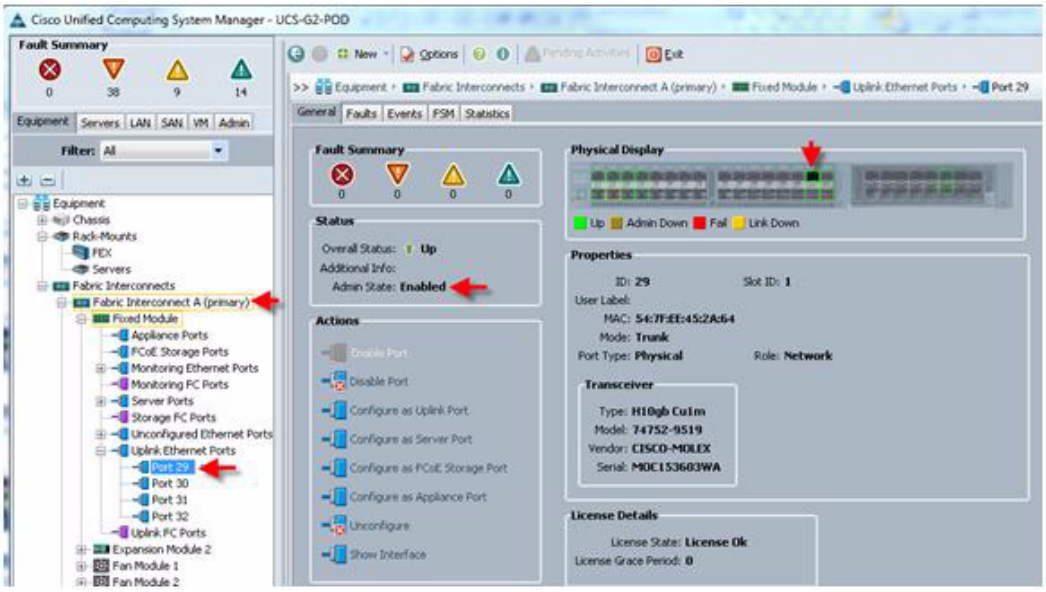

6. When the blades are discovered, it is time to set the KVM IP addresses for each of the blades. This is done through the admin tab>communication management>Management IP address pool. Note: Make sure you have a sufficient number of IP address for all the blades and make sure the gateway and subnet mask is set correctly.

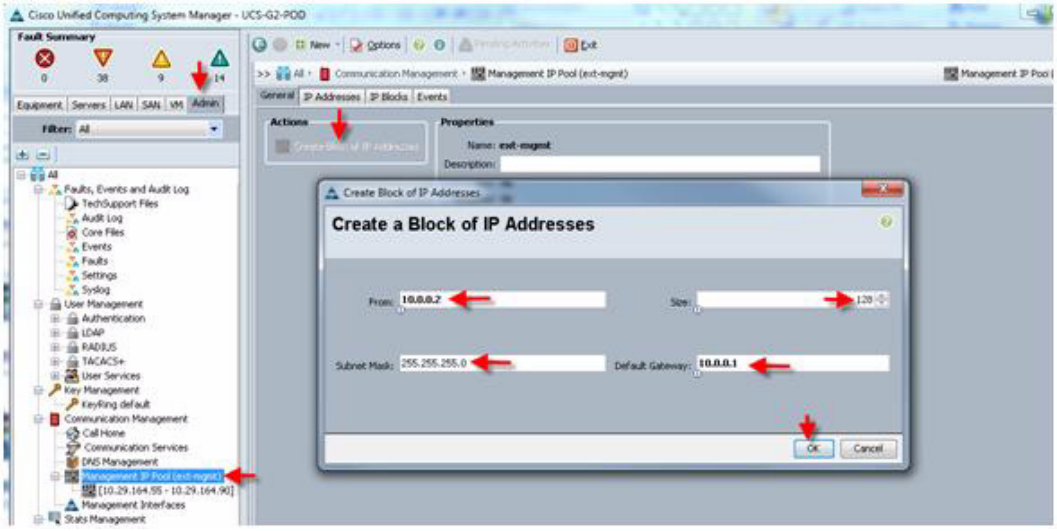

7. Create all the pools: MAC pool, WWPN pool, WWNN pool, UUID pool, Server pool.

a. MAC Pool.

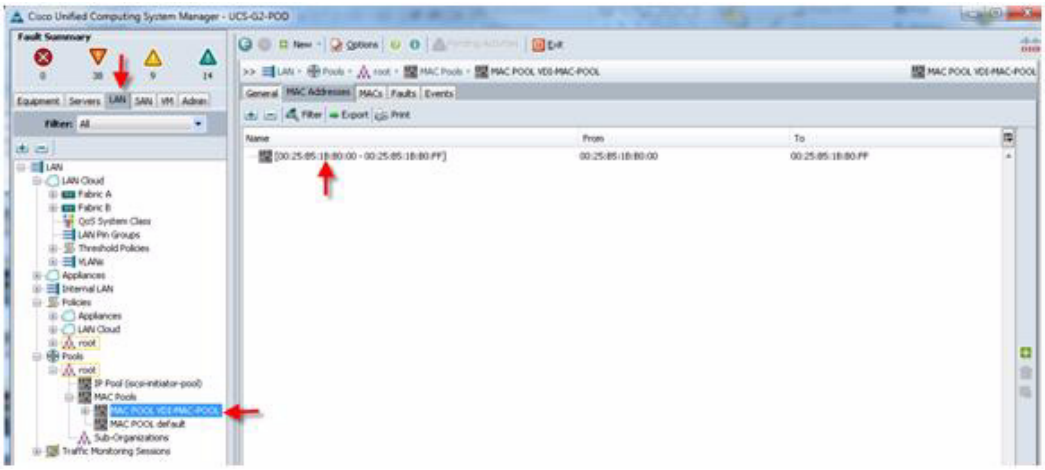

b. WWPN Pool.

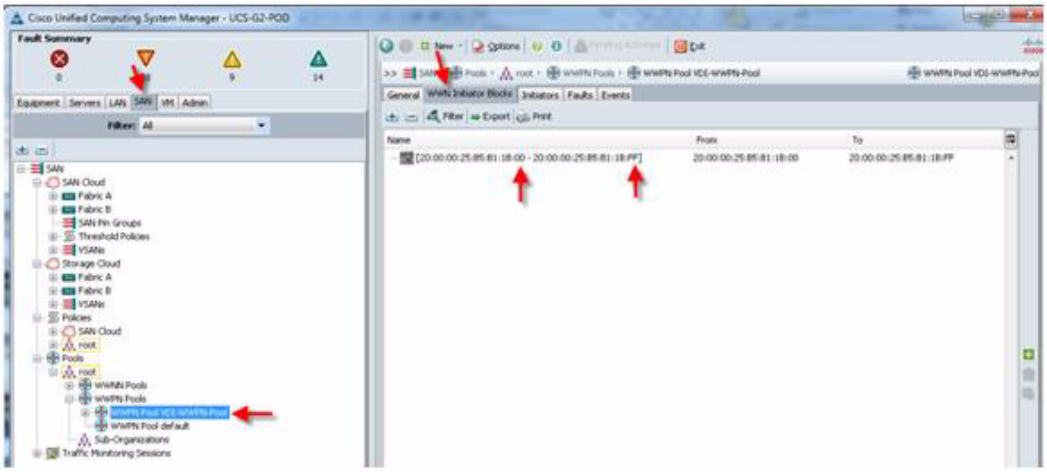

c. WWNN Pool.

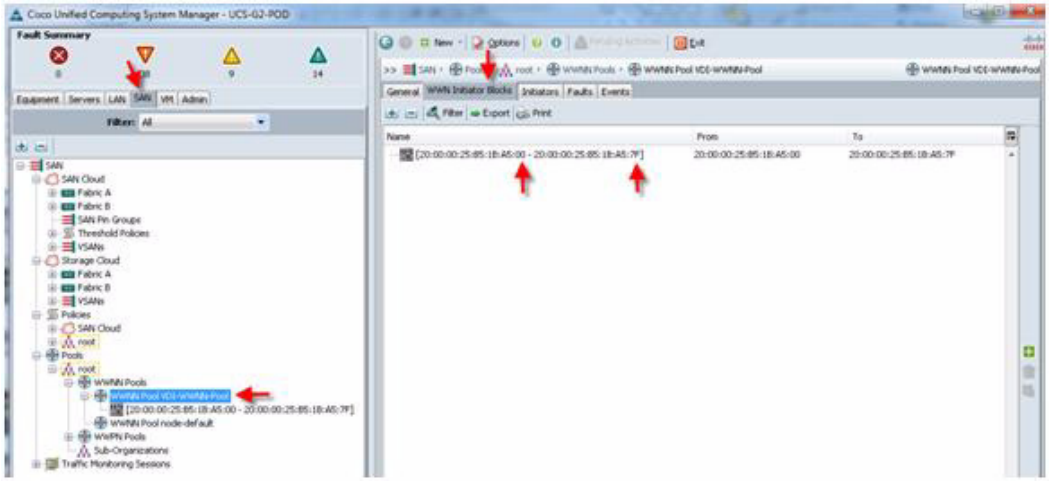

d. UUID Pool.

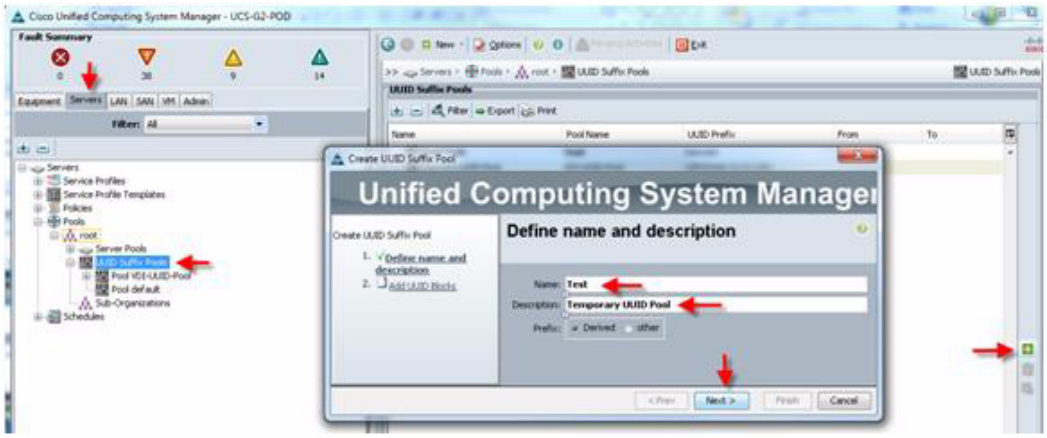
Cisco Unified Computing System Manager - UCS-G2-POD
Fault Summary
Create UUID Suffix Pool
Unified Computing System Manager
Define name and description
Create UUID Suffix Pool
1. Define name and description
2. Add UUID Blocks
Name: Test
Description: Temporary UUID Pool
Prefix: Derived other
< Prev
Next >
Finish
Cancel

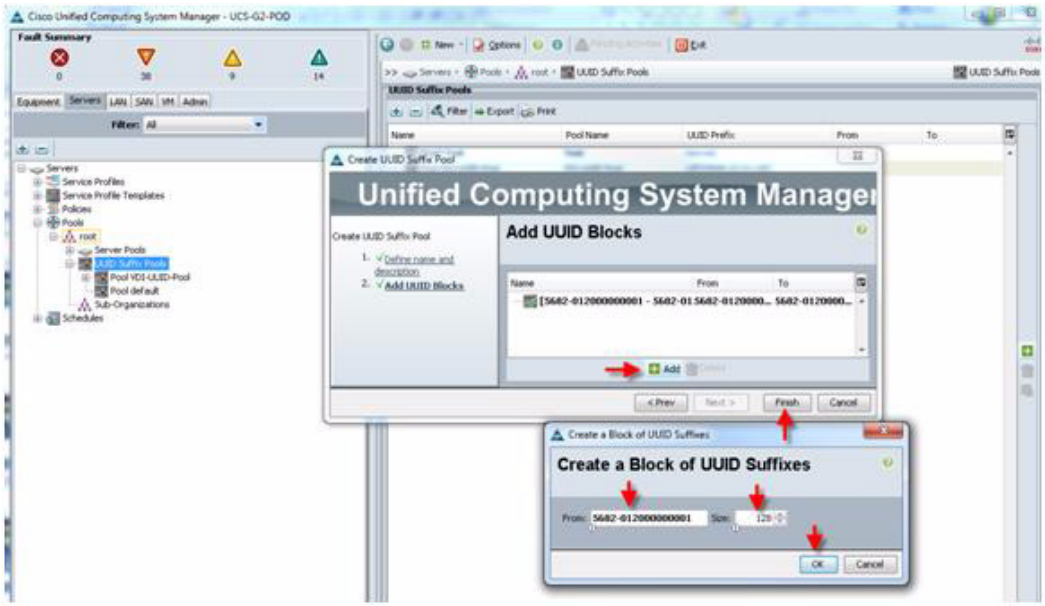
Cisco Unified Computing System Manager - UCS-G2-POD
Fault Summary
Create UUID Suffix Pool
Unified Computing System Manager
Add UUID Blocks
Create UUID Suffix Pool
1. Define name and description
2. Add UUID Blocks
Name
From
To
Add
< Prev
Next >
Finish
Cancel
Create a Block of UUID Suffixes
Create a Block of UUID Suffixes
From: 5682-012000000001
Size:
OK
Cancel

e. Server Pool.

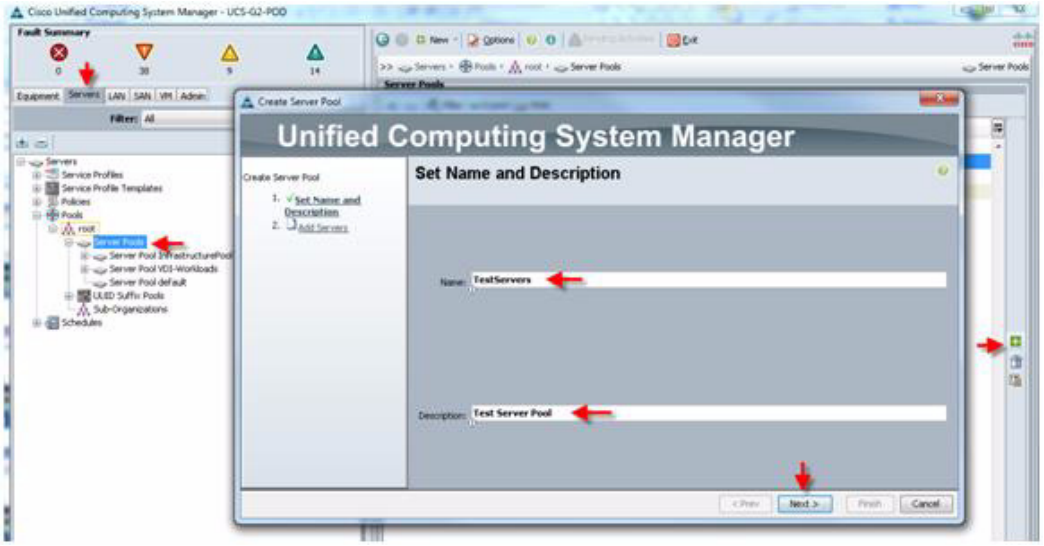

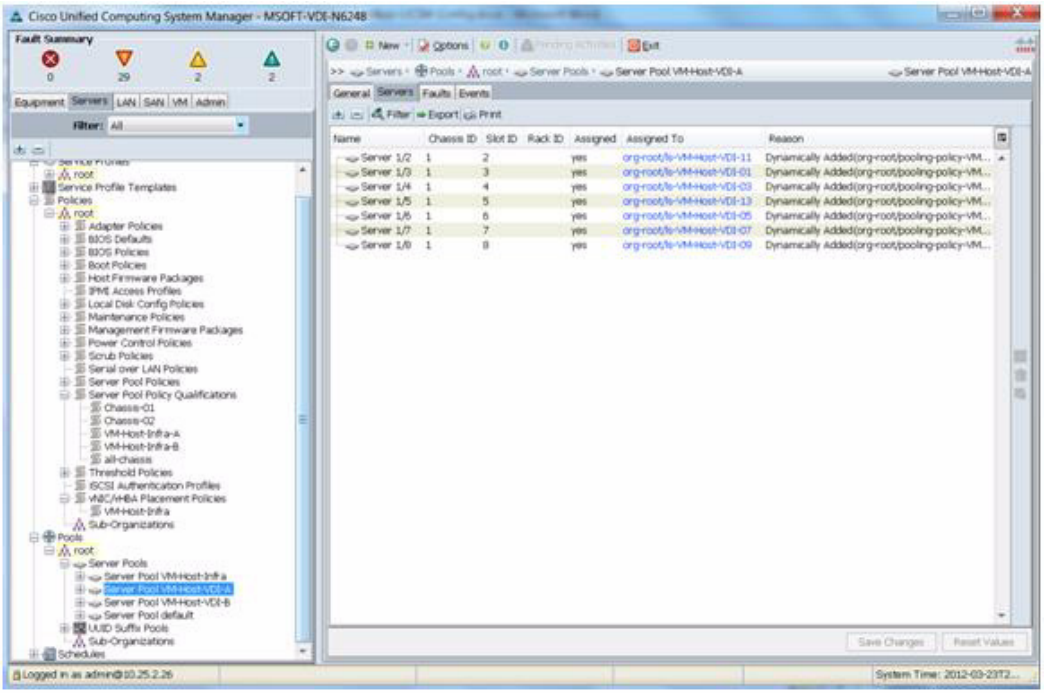

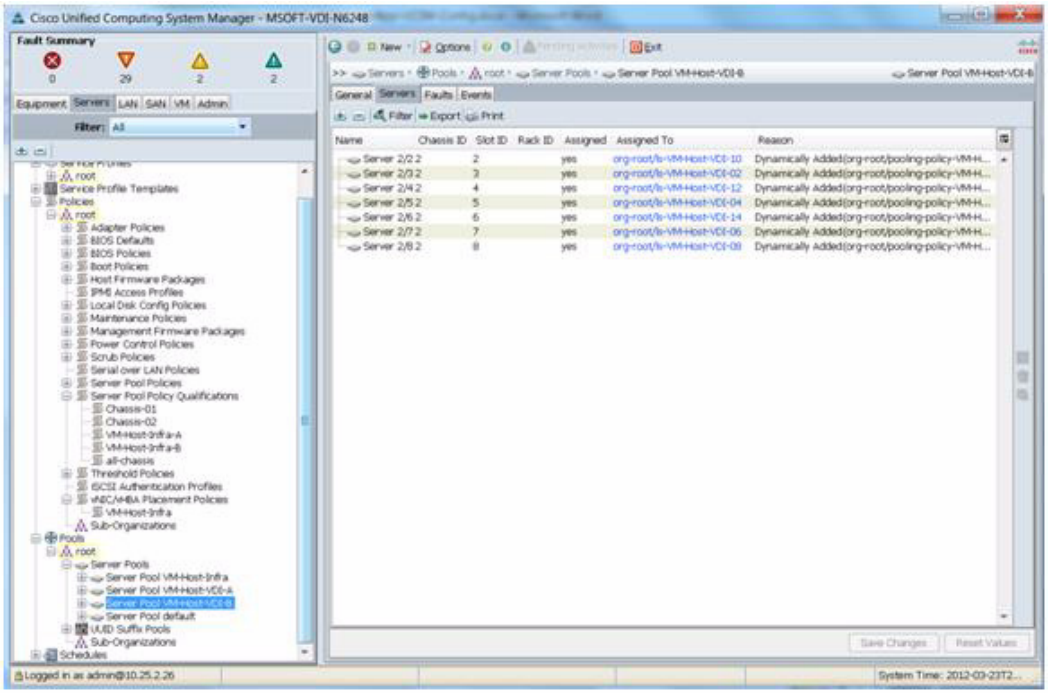

Two server pools were created to balance load across the fabric and storage infrastructure.

8. Create a vHBA template. Create at least one HBA template for each fabric interconnect if block storage will be used.

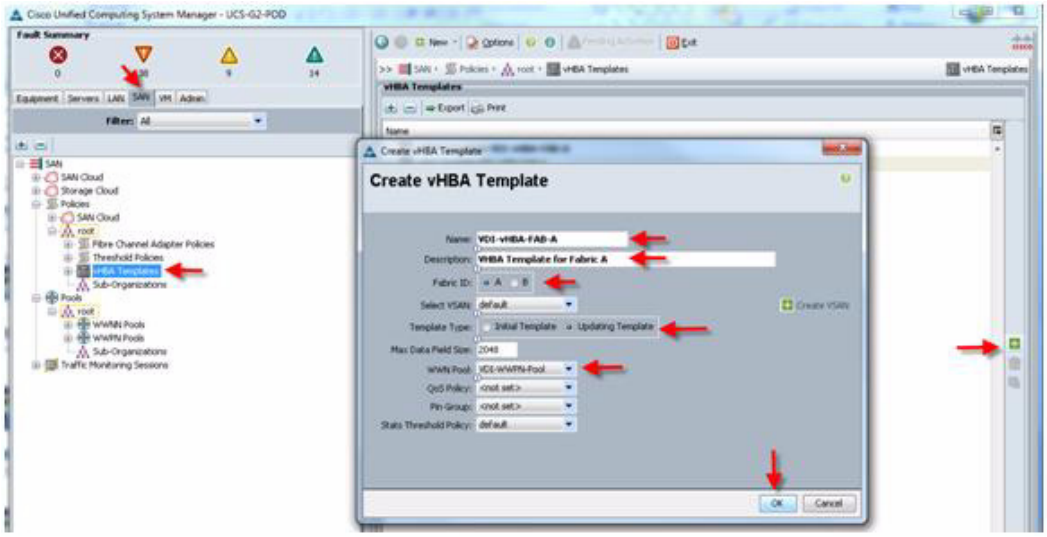

9. Create VLANs (Optional).

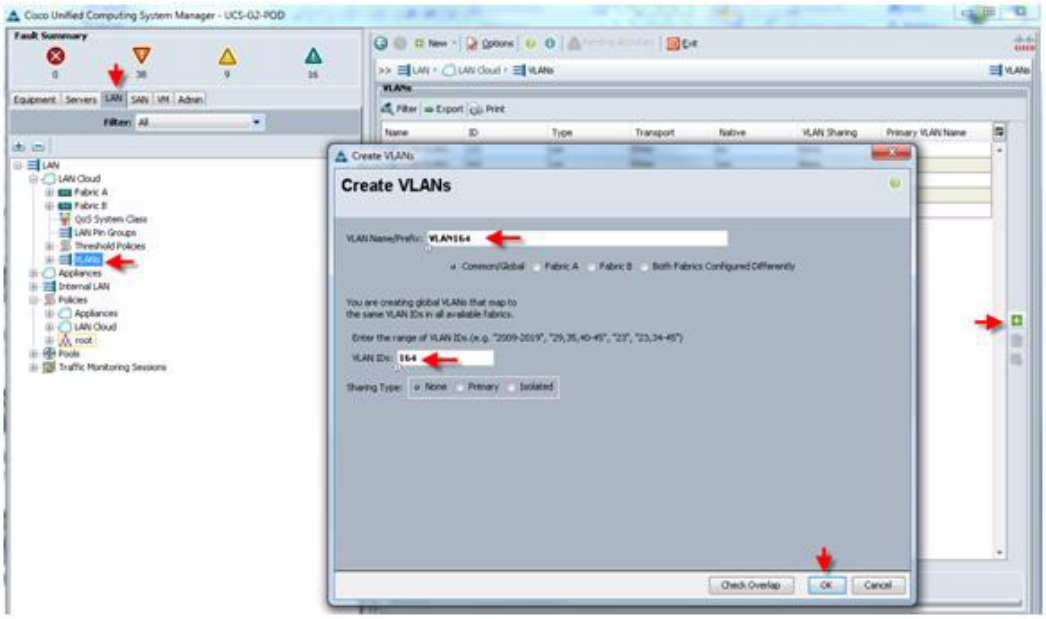

10. Create vNIC templates for both fabrics, check Enable Failover, select VLANs supported on adapter (optional,) set the MTU size, select the MAC Pool, then click OK.

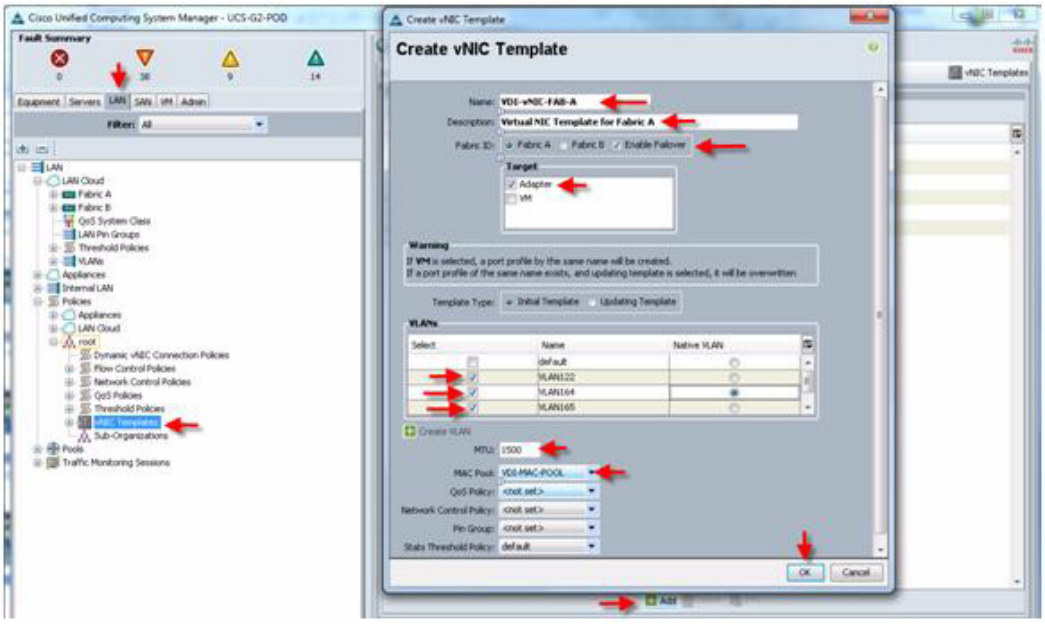

11. Create boot from SAN policies, adapter policies, Server Pool policies, Server Pool Policy Qualifications, and Local boot policy if desired.

12. Create Performance BIOS policies for each blade type to insure that your Low Voltage 8GB-1333 MHz DIMMS will operate at top speed.

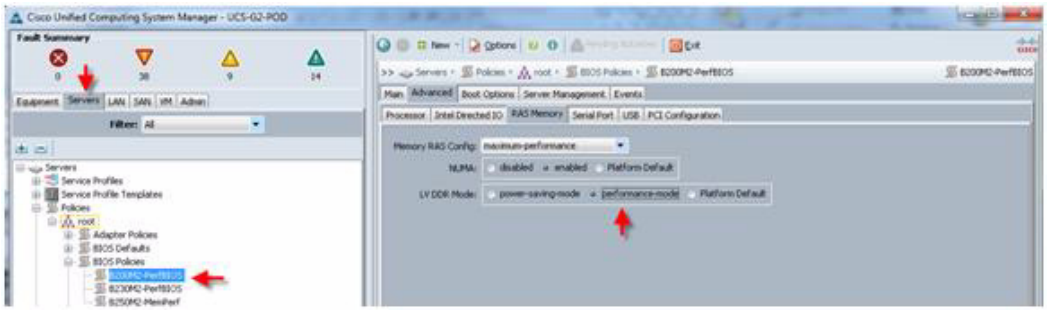

Note Be sure to Save Changes at the bottom of the page to preserve this setting. Be sure to add this policy to your blade service profile template.

13. Create a service profile template using the pools, templates, and policies configured above.

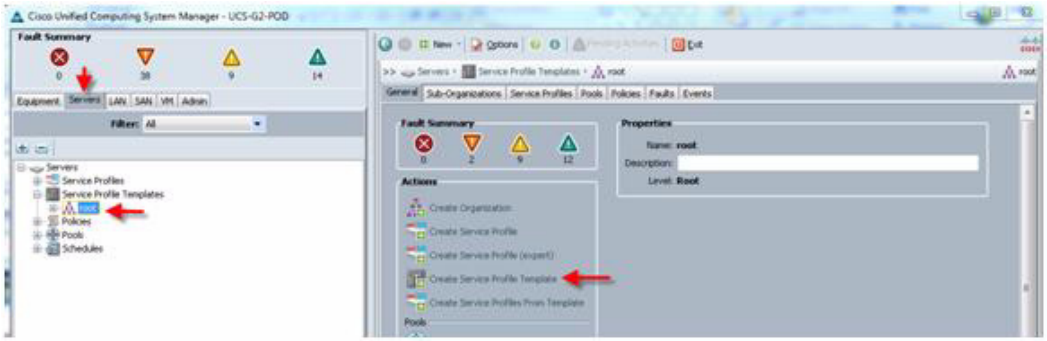

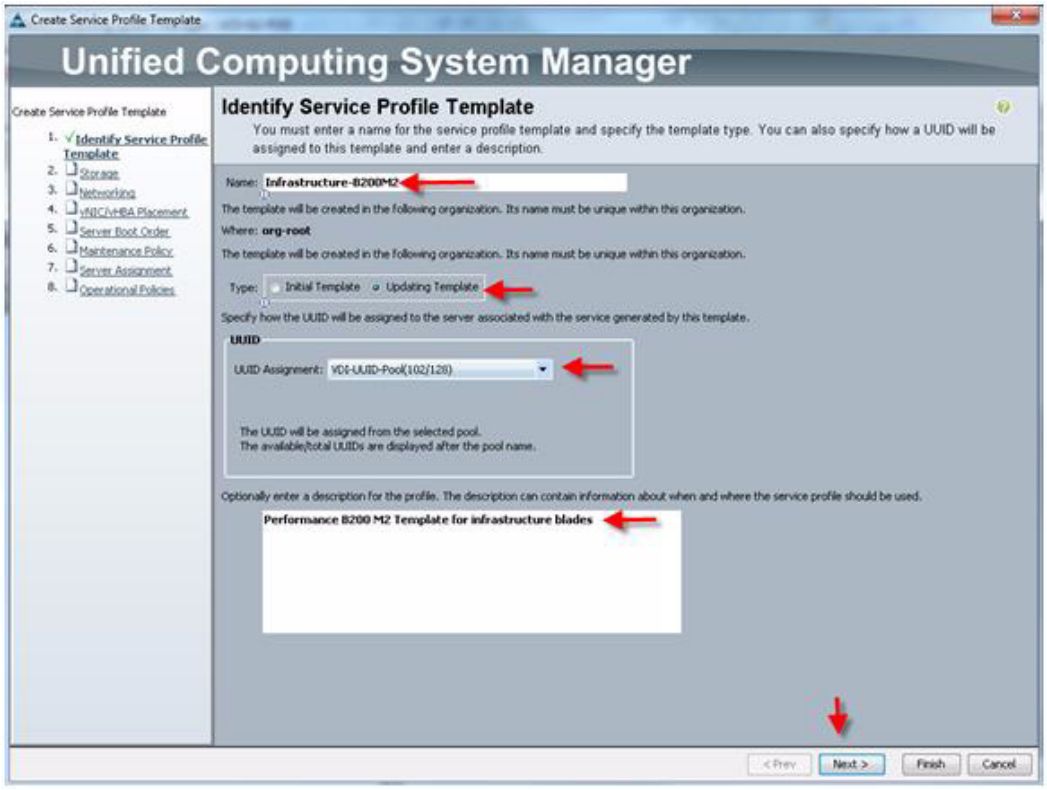

14. Follow through each section, utilizing the policies and objects you created earlier, then click Finish.

Note On the Operational Policies screen, select the appropriate performance BIOS policy you created earlier to insure maximum LV DIMM performance. For automatic deployment of service profiles from your template(s), you must associate a server pool that contains blades with the template.

15. Right-click a Service Profile Template to deploy as many service profiles as you need and the Cisco UCS Manager will automatically start configuring these new service profile templates on the selected blade servers.
16. At this point, the servers are ready for OS provisioning, we would recommend setting up a PXE server to fasten the OS install. Virtual media CD based OS installation is also possible.

QoS and CoS in Cisco Unified Computing System

Cisco Unified Computing System provides different system class of service to implement quality of service including:

- System classes that specify the global configuration for certain types of traffic across the entire system
- QoS policies that assign system classes for individual vNICs
- Flow control policies that determine how uplink Ethernet ports handle pause frames.

Applications like the Cisco Unified Computing System and other time sensitive applications have to adhere to a strict QOS for optimal performance.

System Class Configuration

Systems Class is the global operation where entire system interfaces are with defined QoS rules.

- By default system has Best Effort Class and FCoE Class.
  - Best effort is equivalent in MQC terminology as "match any"
  - FCoE is special Class define for FCoE traffic. In MQC terminology “match cos 3”
- System class allowed with 4 more users define class with following configurable rules.
  - CoS to Class Map
  - Weight: Bandwidth
  - Per class MTU
  - Property of Class (Drop v/s no drop)
- Max MTU per Class allowed is 9216.
- Via UCS we can map one CoS value to particular class.
- Apart from FCoE class there can be only one more class can be configured as no-drop property.
- Weight can be configured based on 0 to 10 numbers. Internally system will calculate the bandwidth based on following equation (there will be rounding off the number)

$$\%\text{b/w shared of given Class} = \frac{\text{(Weight of the given priority * 100)}}{\text{Sum of weights of all priority}}$$

System Class Defined and Mapping

Cisco Unified Computing System defines system class names as follows:

- Platinum
- Gold
- Silver
- Bronze
- Best Effort
- FC (Fibre Channel)

The following tables show the relationships between Class Names and Default Class Values and Weights.

Table 3 Class to CoS Map by default in Cisco Unified Computing System

Cisco UCS Class Names	Cisco UCS Default Class Value
Best effort	Match any
FC	3
Platinum	5
Gold	4
Silver	2
Bronze	1

Table 4 Default Weight in Cisco Unified Computing System

Cisco UCS Class Names	Weight
Best effort	5
FC	5

The Nexus OS uses a similar system, but the names are slightly different. By using corresponding classes of service to set QoS on the Cisco Unified Computing System and Nexus switches, end-to-end QoS can be implemented. Table 5 shows the mapping relationship between the Cisco Unified Computing System and Nexus.

Table 5 Name Map between Cisco Unified Computing System and the NXOS

Cisco UCS Names	NXOS Names
Best effort	Class-default
FC	Class-fc
Platinum	Class-Platinum
Gold	Class-Gold
Silver	Class-Silver
Bronze	Class-Bronze

Enable QoS on the Cisco Unified Computing System

To enable QoS, do the following:

1. Configure the Platinum policy by checking the Platinum policy box and if you want jumbo frames enabled change MTU from normal to 9000. Notice the option to set no packet drop policy during this configuration.

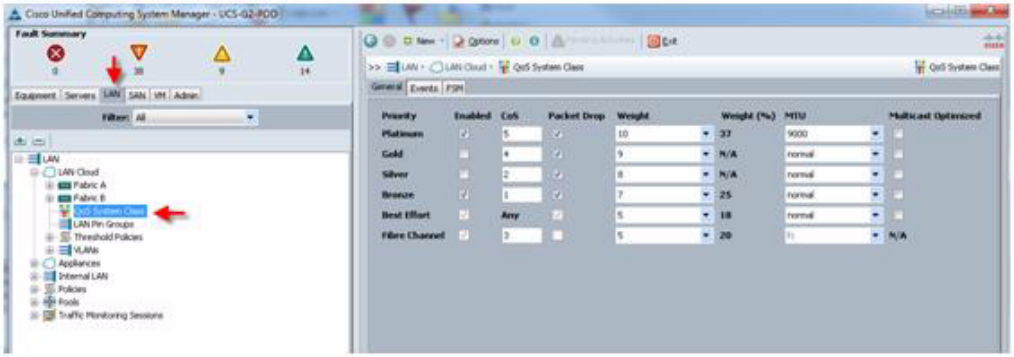

2. In the LAN tab under Policies, Root, QoS Polices, verify a QoS Policy Platinum exists and that Platinum is set as the Priority.

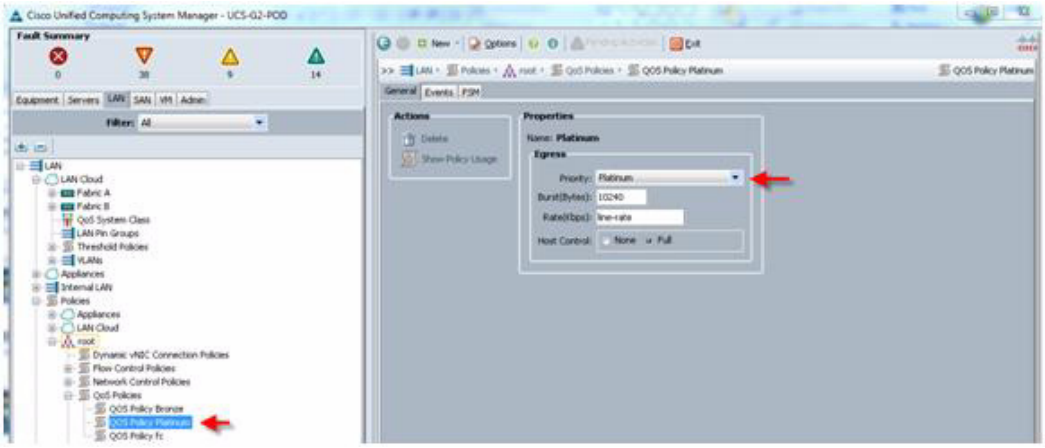

3. Include the QoS Policy Platinum policy into the vNIC template under the QoS policy.

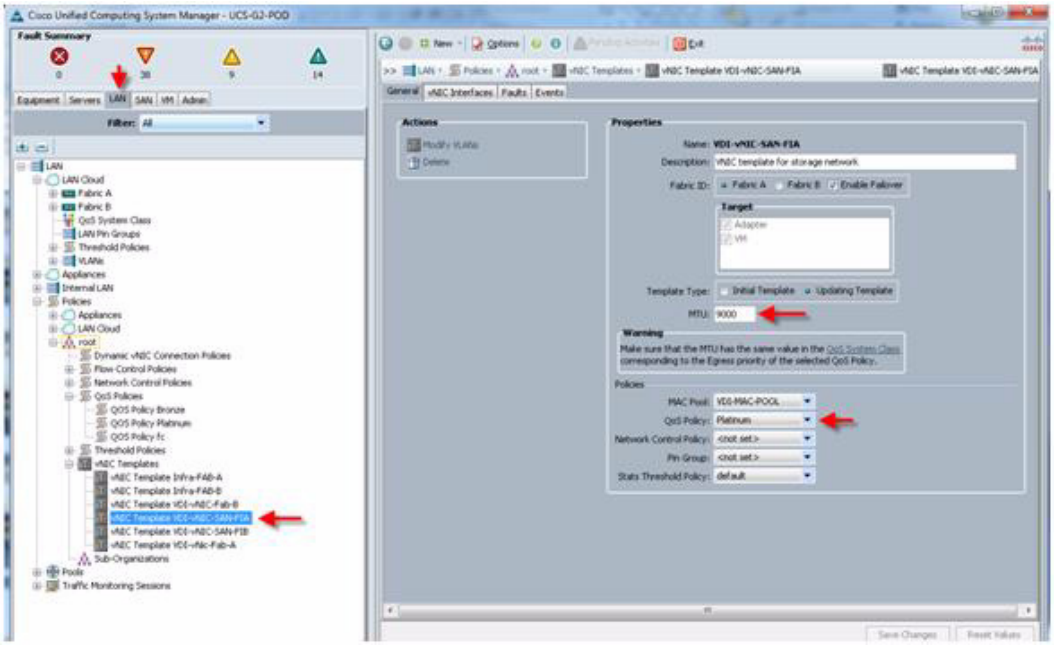

This is a unique value proposition for the Cisco Unified Computing System with respect to end-to-end QOS. For example, we have a VLAN for the NetApp filer iSCSI LUNs, configure Platinum policy with Jumbo frames and get an end-to-end QOS and performance guarantees.

LAN Configuration

The access layer LAN configuration consists of a pair of Cisco Nexus 5548s, a family member of our low-latency, line-rate, 10 Gigabit Ethernet and FCoE switches for our VDI deployment.

Cisco UCS Connectivity

Four 10 Gigabit Ethernet uplink ports are configured on each of the Cisco UCS 6248 fabric interconnects, and they are connected to the Cisco Nexus 5548 pair in a bow tie manner as shown below in a port channel.

The 6248 Fabric Interconnect is in End host mode, as we are doing both Fibre Channel as well as Ethernet (NAS) data access as per the recommended best practice of the Cisco Unified Computing System. We built this out for scale and have provisioned more than 40 G per Fabric interconnect.

Note The upstream configuration is beyond the scope of this document; there are some good reference document [4] that talks about best practices of using the Cisco Nexus 5000 and 7000 Series Switches. New with the Nexus 5500 series is an available Layer 3 module that was not used in these tests and that will not be covered in this document.

Figure 16 Ethernet Network Configuration with Upstream Cisco Nexus 5000 Series from the Cisco Unified Computing System Fabric Interconnects

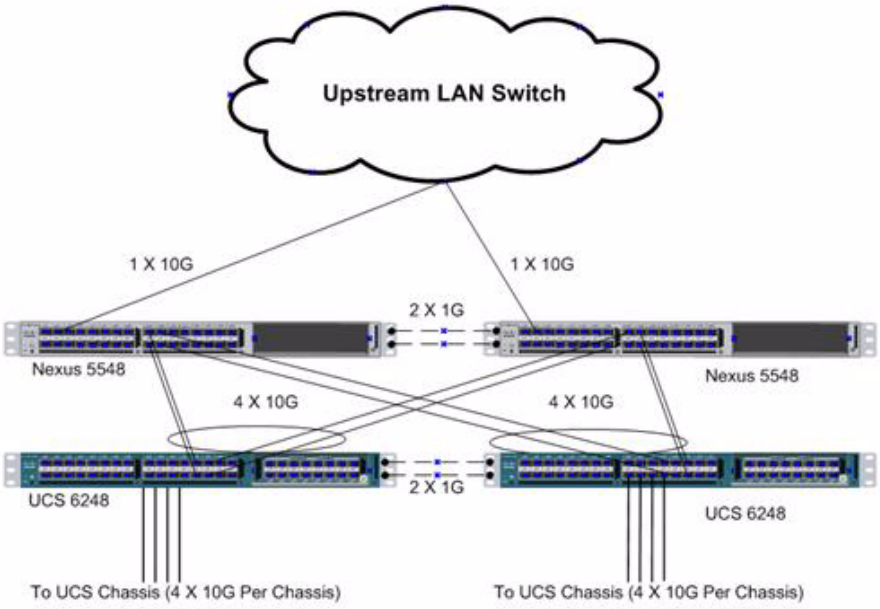

FAS 3240 LAN Connectivity

The Cisco Nexus 5548 is used to connect to the NetApp FAS3240 storage system for Fibre Channel, iSCSI and file-based access.

The FAS3240 is equipped with dual-port 8GB FC modules on each controller. These are connected to the pair of Nexus 5548s to provide block storage access to the environment.

The FAS3240 supports dual-port 10G Ethernet modules on each controller which are configured in a port channel and connected to the pair of N5Ks downstream. This allows end-to-end 10G access for iSCSI and file-based storage traffic. We have implemented jumbo frames on the ports and have priority flow control on, with platinum CoS and QoS assigned to the vNICs carrying the storage data access on the Fabric Interconnects.

The NetApp ethernet connectivity diagram is shown below. There is a total of 40G bandwidth available for the servers.

Figure 17 NetApp FAS 3240 Ethernet Connectivity

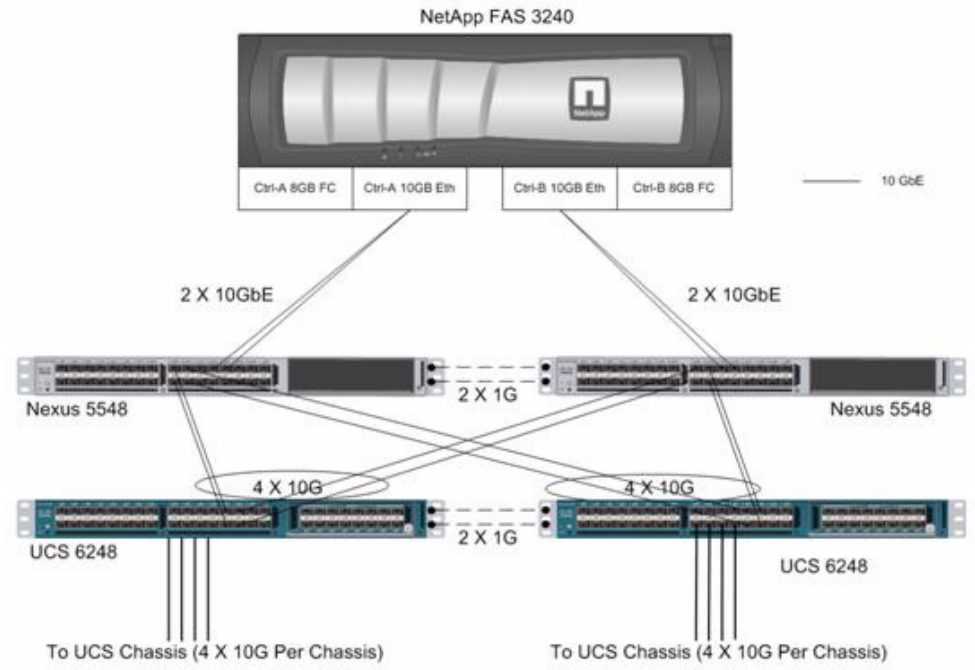

For information on configuring ethernet connectivity on a NetApp FAS3240, refer to FlexPod Validated with Microsoft Private Cloud.

SAN Configuration

The same pair of 5548 switches were used in the configuration to connect between the FC ports on the EMC VNX5500 and the FC ports of the UCS 6248 Fabric Interconnect expansion module.

Boot from SAN Benefits

Booting from SAN is another key feature which helps in moving towards stateless computing in which there is no static binding between a physical server and the OS / applications it is tasked to run. The OS is installed on a SAN LUN and boot from SAN policy is applied to the service profile template or the

service profile. If the service profile were to be moved to another server, the pwwn of the HBAs and the Boot from SAN (BFS) policy also moves along with it. The new server now takes the same exact view of the old server, the true stateless nature of the blade server.

The key benefits of booting from the network:

- Reduce Server Footprints: Boot from SAN alleviates the necessity for each server to have its own direct-attached disk, eliminating internal disks as a potential point of failure. Thin diskless servers also take up less facility space, require less power, and are generally less expensive because they have fewer hardware components.
- Disaster and Server Failure Recovery: All the boot information and production data stored on a local SAN can be replicated to a SAN at a remote disaster recovery site. If a disaster destroys functionality of the servers at the primary site, the remote site can take over with minimal downtime.
- Recovery from server failures is simplified in a SAN environment. With the help of snapshots, mirrors of a failed server can be recovered quickly by booting from the original copy of its image. As a result, boot from SAN can greatly reduce the time required for server recovery.
- High Availability: A typical data center is highly redundant in nature - redundant paths, redundant disks and redundant storage controllers. When operating system images are stored on disks in the SAN, it supports high availability and eliminates the potential for mechanical failure of a local disk.
- Rapid Redeployment: Businesses that experience temporary high production workloads can take advantage of SAN technologies to clone the boot image and distribute the image to multiple servers for rapid deployment. Such servers may only need to be in production for hours or days and can be readily removed when the production need has been met. Highly efficient deployment of boot images makes temporary server usage a cost effective endeavor.
- Centralized Image Management: When operating system images are stored on networked disks, all upgrades and fixes can be managed at a centralized location. Changes made to disks in a storage array are readily accessible by each server.

With Boot from SAN, the image resides on a SAN LUN and the server communicates with the SAN through a host bus adapter (HBA). The HBAs BIOS contain the instructions that enable the server to find the boot disk. All FC capable Converged Network Adapter (CNA) cards supported on Cisco UCS B-series blade servers support Boot from SAN.

After power on self test (POST), the server hardware component fetches the boot device that is designated as the boot device in the hardware BOIS settings. Once the hardware detects the boot device, it follows the regular boot process.

Configuring Boot from SAN Overview

There are three distinct phases during the configuration of Boot From SAN. The high level procedures include:

- SAN zone configuration on the Nexus 5548s
- Storage array host initiator configuration
- Cisco UCS configuration of service profile

In the following sections, each high-level phase will be discussed.

SAN Configuration on Nexus 5548

The FCoE and NPIV feature has to be turned on in the Nexus 5500 series switch. Make sure you have 8 GB SPF+ modules connected to the UCS 6200 series fabric interconnect expansion ports. The port mode is set to AUTO as well as the speed is set to AUTO. Rate mode is "dedicated" and when everything is configured correctly you should see something like the output below on a Nexus 5500 series switch for a given port (for example, Fc1/17).

Note A Nexus 5500 series switch supports a single VSAN configuration. If more VSANs are required, the use of a SAN switch like the MDS 9100 series should be deployed.

Cisco Fabric Manager can also be used to get a overall picture of the SAN configuration and zoning information. As discussed earlier, the SAN zoning is done up front for all the pwwn of the initiators with the NetApp FAS 3240 target pwwn.

```
VDI-N5548-A# show feature | grep npiv
npiv                  1         enabled


 VDI-N5548-A# show interface brief
--------------------------------------------------------------------------------
Interface  Vsan   Admin  Admin   Status         SFP  Oper  Oper    Port
                   Mode   Trunk                      Mode Speed  Channel
                          Mode                            (Gbps)
--------------------------------------------------------------------------------
fc1/17     1      auto   on      up             swl   F     8    --
fc1/18     1      auto   on      up             swl   F     8    --
```

The FC connection was used for configuring boot from SAN for all of server blades. In addition, a general purpose 1TB infrastructure LUN for infrastructure virtual machine storage and 14 write-cache LUNs for each VDI host were provisioned.

Single vSAN zoning was set up on the Nexus 5548's to make those FAS3240 LUNs visible to the infrastructure and test servers.

An example SAN zone configuration is shown below on the Fabric A side:

```
VDI-N5548-A# sh zoneset active vsan 1  zoneset name MS-VDI vsan 1
  zone name vm-host-infra-01-fc0 vsan 1
  * fcid 0xe30018 [pwwn 20:00:00:25:b5:02:02:3f] [vm-host-infra-01-fc0]
  * fcid 0xe300ef [pwwn 50:0a:09:81:8d:42:bd:e8] [FAS3240-B-0c]
  * fcid 0xe301ef [pwwn 50:0a:09:81:9d:42:bd:e8] [FAS3240-A-0c]


  zone name vm-host-infra-02-fc0 vsan 1
  * fcid 0xe3000d [pwwn 20:00:00:25:b5:02:02:1f] [vm-host-infra-02-fc0]
  * fcid 0xe301ef [pwwn 50:0a:09:81:8d:42:bd:e8] [FAS3240-B-0c]
  * fcid 0xe300ef [pwwn 50:0a:09:81:9d:42:bd:e8] [FAS3240-A-0c]
```

Where 20:00:00:25:b5:02:02:3f /20:00:00:25:b5:02:02:1f are blade servers pwwn's of their respective Converged Network Adapters (CNAs) that are part of the Fabric A side.

The NetApp FC target ports are 50:0a:09:81:9d:42:bd:e8/50:0a:09:81:9d:42:bd:e8 and belong to FC modules on Controller-A.

Similar zoning is done on the second Nexus 5548 in the pair to take care of the Fabric B side as shown below.

VDI-N5548-B# sh zoneset active vsan 1 zone name vm-host-infra-01-fc1 vsan 1

* fcid 0xd501ef [pwwn 50:0a:09:82:9d:42:bd:e8] [FAS3240-A-0d]

* fcid 0xd500ef [pwwn 50:0a:09:82:8d:42:bd:e8] [FAS3240-B-0d]

* fcid 0xd50010 [pwwn 20:00:00:25:b5:02:02:38] [vm-host-infra-01-fc1]

zone name vm-host-infra-02 vsan 1

* fcid 0xd5000f [pwwn 20:00:00:25:b5:02:02:0f] [vm-host-infra-02-fc1]

* fcid 0xd500ef [pwwn 50:0a:09:82:8d:42:bd:e8] [FAS3240-B-0d]

* fcid 0xd501ef [pwwn 50:0a:09:82:9d:42:bd:e8] [FAS3240-A-0d]

Where 20:00:00:25:b5:02:02:3f /20:00:00:25:b5:02:02:0f are blade servers pwwn's of their respective Converged Network Adapters (CNAs) that are part of the Fabric B side.

The NetApp FC target ports are 50:0a:09:82:9d:42:bd:e8/50:0a:09:82:8d:42:bd:e8 and belong to FC modules on Controller. They were spread across the two controllers for redundancy as shown in Figure 18.

Figure 18 FAS3240 FC Target Ports

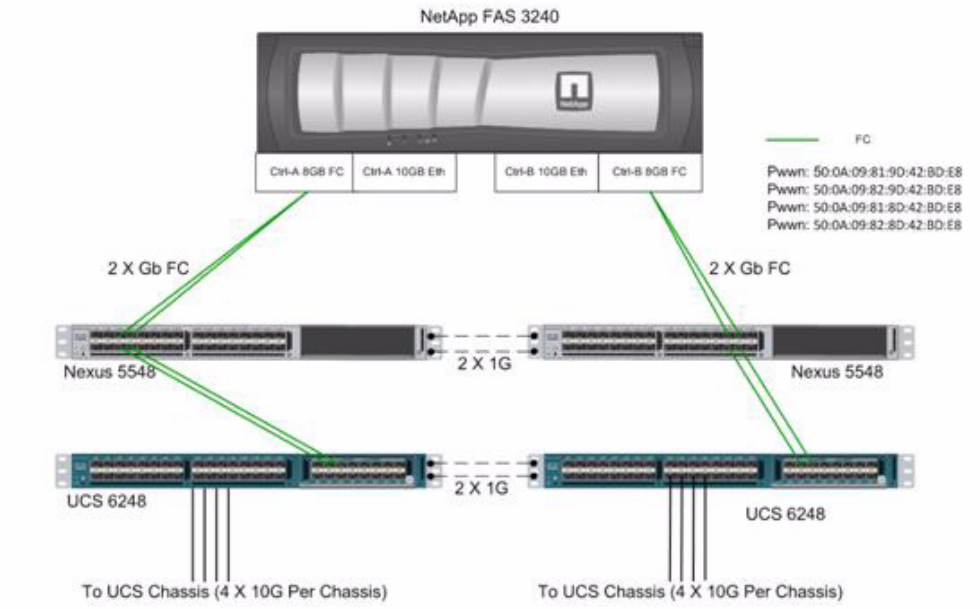

For detailed Nexus 5500 series switch configuration, refer to Cisco Nexus 5500 Series NX-OS SAN Switching Configuration Guide. (See the Reference Section of this document for a link.)

Storage Array Host Initiator Configuration

The following section provides detailed procedures for configuring fibre channel initiator groups (or igroups) for mapping to a NetApp LUN for Boot from SAN.

Create Initiator Groups

Figure 22 shows a PowerShell Window capture of the commands to create Initiator Groups.

Figure 19 PowerShell Window - Create Initiator Groups

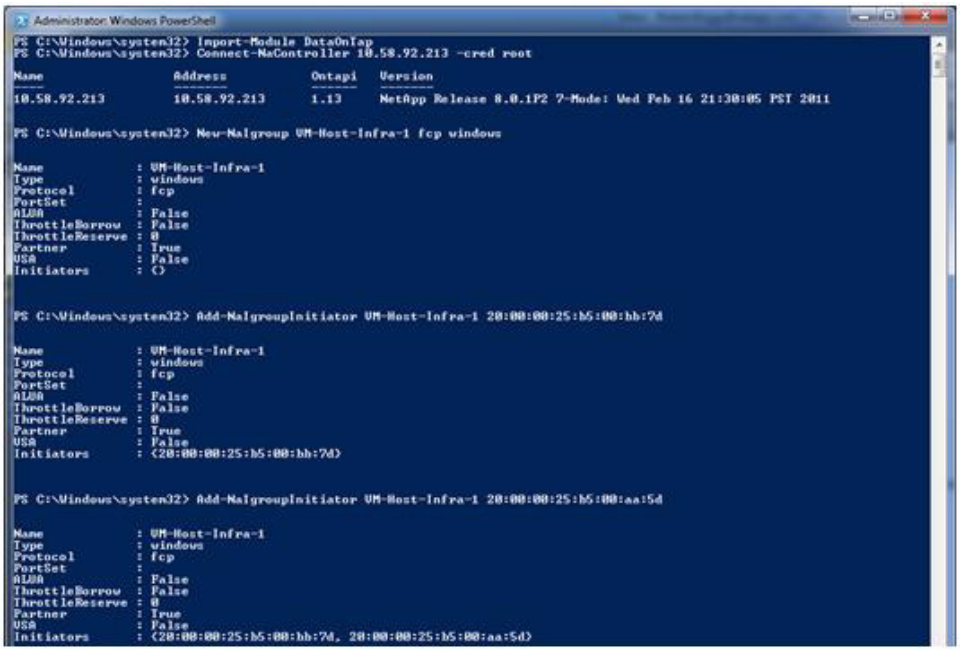

1. Open an elevated Windows PowerShell Window as a Local or Domain Administrator.
2. Type **import-module dataontap** at the prompt and hit enter. Ensure that you have the most recent version of the NetApp PowerShell Toolkit installed in the **C:\Windows\System32\WindowsPowerShell\v1.0\Modules** directory.
3. Type **Connect-NaController <IP Address of Controller> -cred root** and hit enter. You will be prompted to enter the password for root. Enter password and hit Enter.
4. When connected to the controller, at the prompt type **New-NaIgroup <igroup name> fcp windows** and click Enter. This will create a new initiator group on your controller called <igroup name>.
5. At the prompt, type **Add-NaIgroupInitiator <igroup name> <Fabric A WWPN>** and click Enter.
6. Type **Add-NaIgroupInitiator <igroup name> <Fabric B WWPN>** and click Enter. Now the <igroup name> igroup will have both Fabric A and Fabric B WWPNs assigned to it.

SAN Configuration on Cisco UCS Manager

To enable Boot from SAN on the Cisco UCS Manager 2.0 series, do the following:

1. Create a boot policy in the Servers tab. To do this, Select the policies and on the right plane select boot policies and select Add button. Enter name, select reboot on change, and do not select enforce vHBA name.

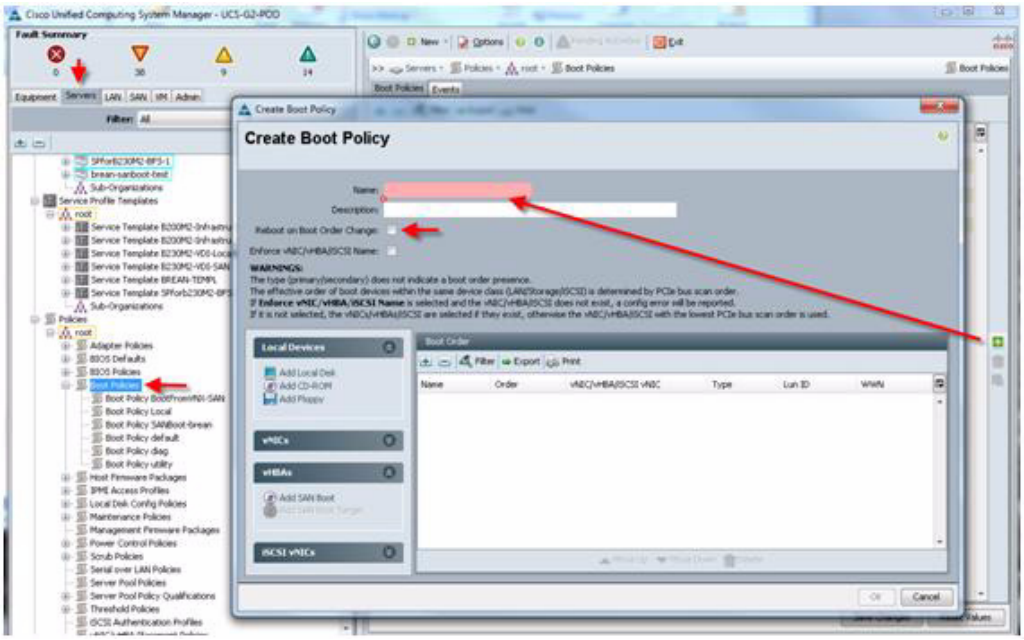

2. Add SAN Boot for primary to the new policy. The vHBA name is optional and can be left blank and you do not have to enforce the vHBA name. Click OK.

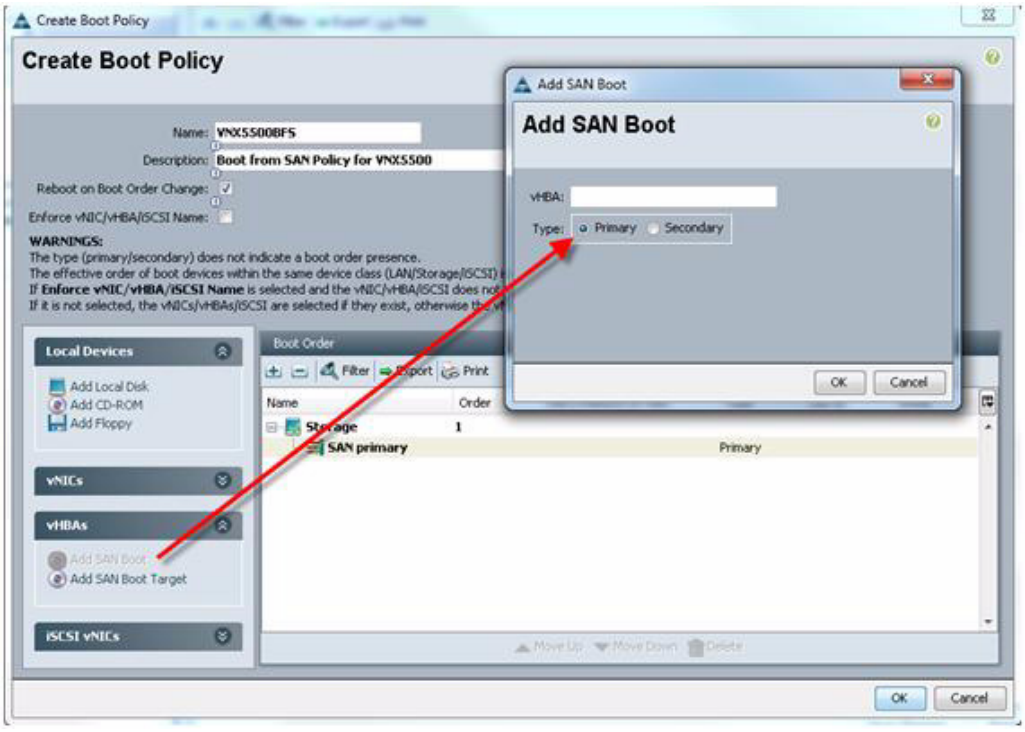

3. Add SAN boot for SAN Secondary. Click OK. Leave the optional vHBA name blank.

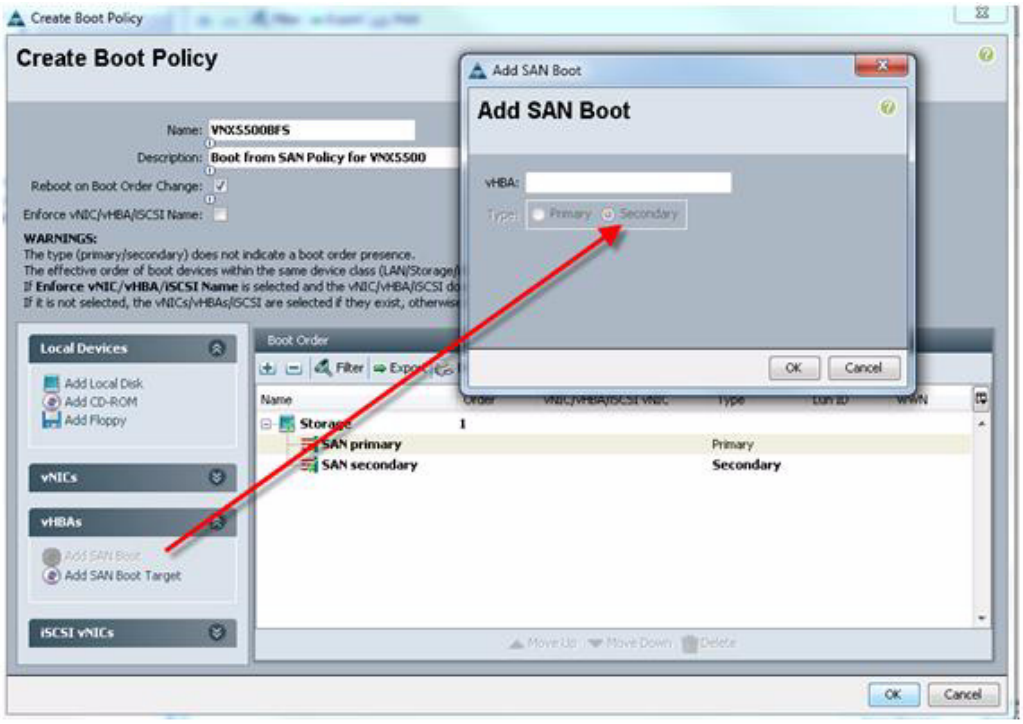

4. Add Boot target WWPN to the SAN Primary, make sure this is exactly matches the NetApp FAS pwwn. To avoid any typos, copy and paste from Nexus 5500 Series command as follows from each switch:

VDI-N5548-A# show fcns database vsan 1

0x470300 N 50:0a:09:81:9d:42:bd:e8 (NetApp) scsi-fcp

0x470200 N 50:0a:09:81:8d:42:bd:e8 (NetApp) scsi-fcp

VDI-N5548-B # show fcns database vsan 1

0x470400 N 50:0a:09:82:9d:42:bd:e8 (NetApp) scsi-fcp

0x470500 N 50:0a:09:82:8d:42:bd:e8 (NetApp) scsi-fcp

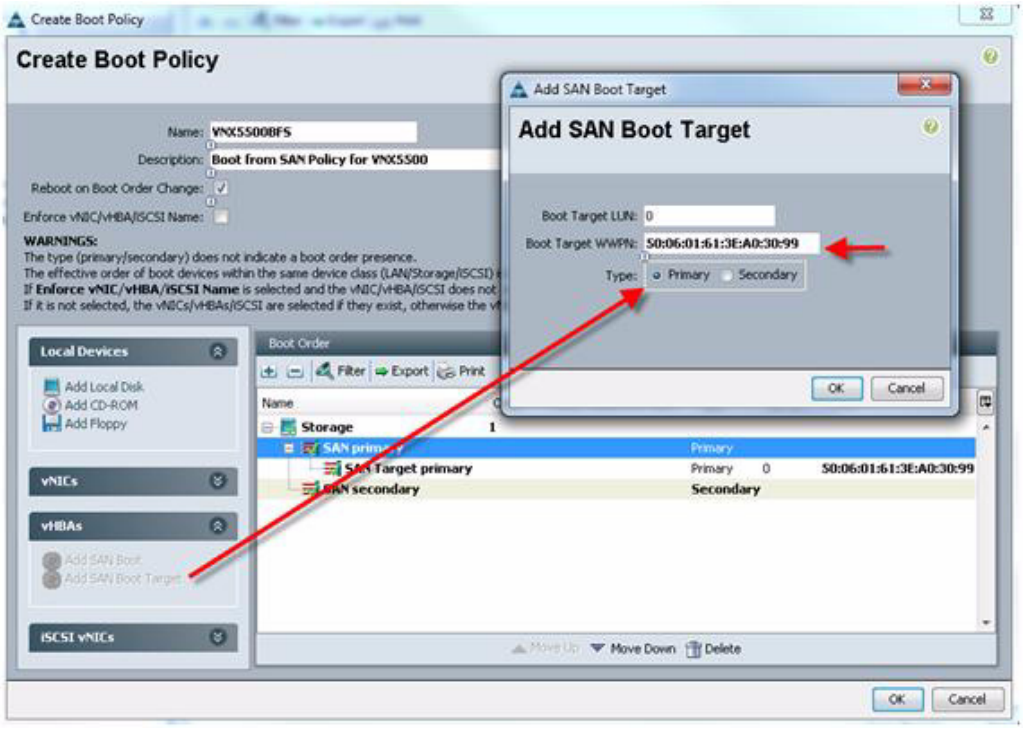

5. Repeat step 4 for SAN primary's SAN Target Secondary
6. Repeat step 4 for SAN Secondary's—AN Target Primary
7. Repeat step 4 for SAN Secondary's—SAN Target Secondary
8. At the end your Boot from SAN policy should look like:

   Boot Fabric A Policy

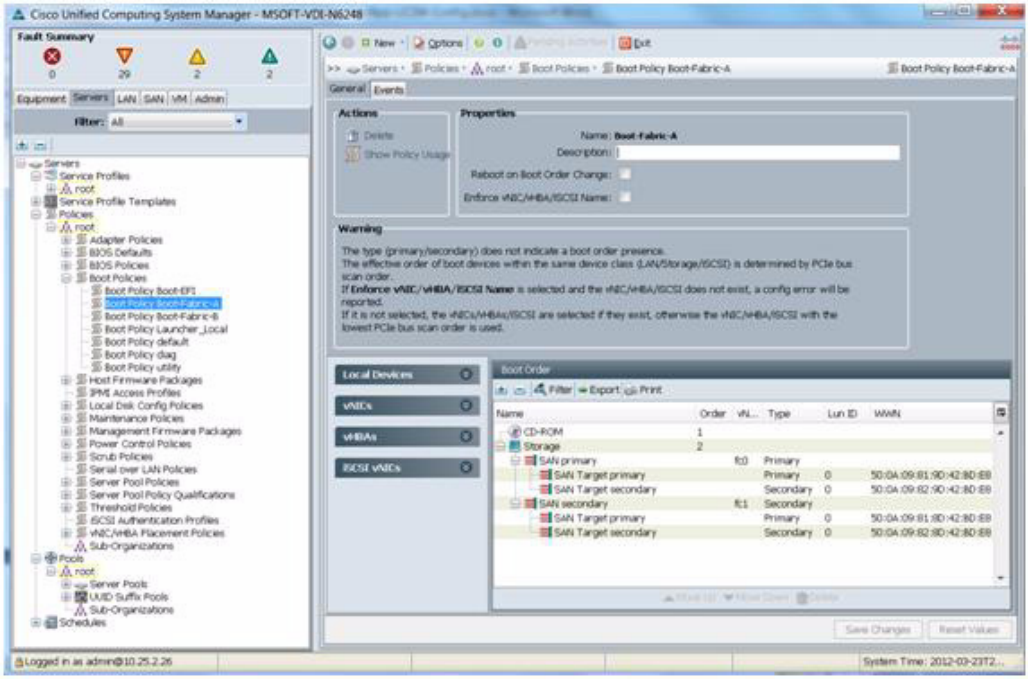

Boot Fabric B Policy

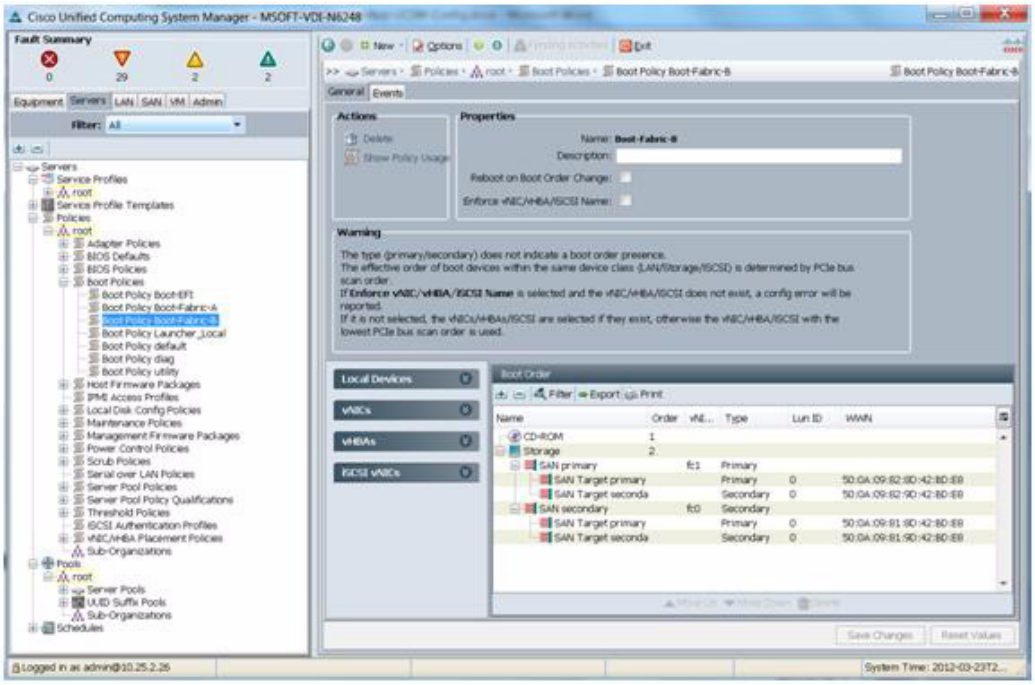

9. Associate the service profile template to the Boot from SAN policy during the service profile template configuration.

You can also modify the Boot policy in use in a Service Profile Template as shown below:

For Fabric A

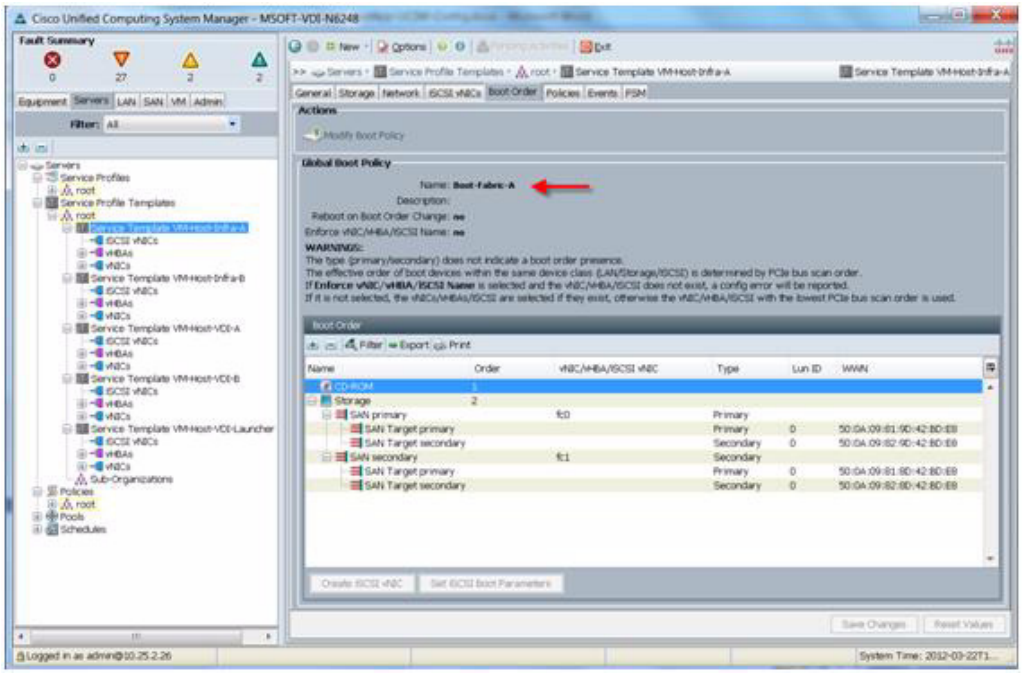

For Fabric B

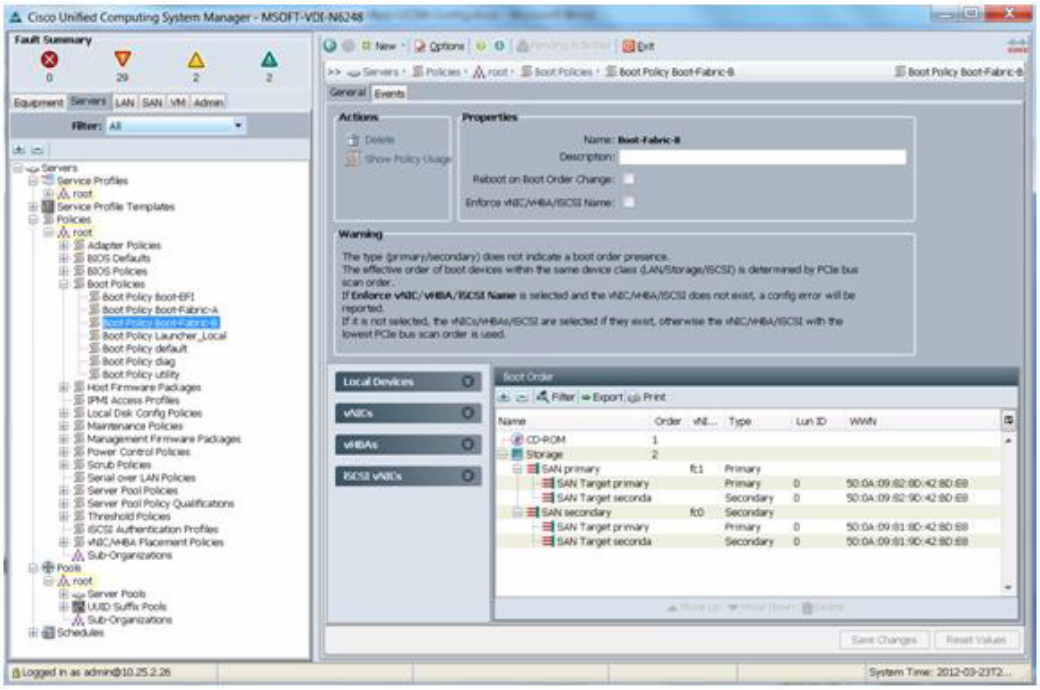

10. This completes the BFS configuration on the Cisco UCS Manager. When the service profile is created out of the template, each server will be ready to boot from SAN provided the appropriate OS installation steps has taken place.

NetApp FAS3240 Storage Configuration

The following sections provide a detailed overview and configuration instructions for allocating storage on the NetApp FAS3240. This overview directly pertains to the VDI workload which is built on the FlexPod 2.0 specification. For details about configuring NetApp storage, refer to the Cisco CVD FlexPod Validated with Microsoft Private Cloud - February 2012.

Storage Overview

Each FAS3240 Controller is connected to two NetApp DS2246 Disk Shelves. Each DS2246 shelf contains 24 600GB 10,000rpm SAS drives. This gives each FAS3240 Controller direct control of and access to a total of 48 disks.

Aggregate Overview

For this Reference Architecture it was decided that two Raid-DP Groups of 16 disks would be used. This totaled to 32 disks for each controller. Two physical disks per Raid Group were allocated for parity, so for sizing calculations, 28 physical disks per controller were used.

One aggregate per controller, containing these Raid-DP Groups was created. This aggregate was named aggr1.

Volume Overview

To help ensure that each controller maintained peak performance throughout all of the lifecycles of a VDI environment, volumes were created symmetrically across both controllers. The following is a brief list of the volumes and their purpose:

- **Cisco UCS Boot Volume** contains all Boot from SAN OS LUNs for each physical UCS Server.
- **VDI Host Write Cache Volume** contains write cache LUNs for each VDI Hyper-V Server.
- **Infrastructure VM Non-Clustered Volume** contains a LUN for non-clustered infrastructure VMs (SQL and Active Directory).
- **Infrastructure VM Clustered Volume** contains a Cluster Shared Volume (CSV) for clustered infrastructure VMs (PVS, SCVMM, XenDesktop Desktop Controller, File Services).
- **SQL Clustered Volume** contains CSVs for clustering SQL at an application level. Included are CSVs for SQL instances, LOG files and data.
- **User Profile Clustered Volume** contains a CSV for clustering file services for user profile management.
- **Quorum Volume** contains quorum disks for SQL, Infrastructure VMs and File Services.

LUN Overview

All LUNs were created within their respective volume containers. They were all mapped to FC initiator groups for presentation to respective physical hosts. One exception to this was the SQL CSVs. These were mapped to iSCSI initiators for presentation to VMs with SQL installed.

Figure 20 shows two NetApp FAS3240 Controllers in Active-Active HA configuration with a detailed breakdown of Aggregate, Volume, and LUN allocation and placement.

Figure 20 NetApp Storage Overview

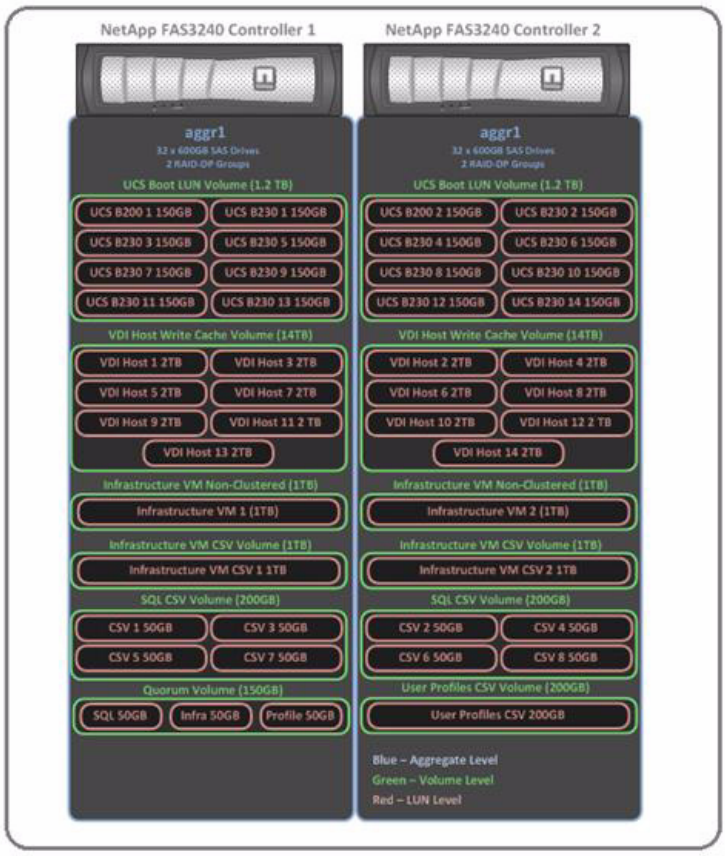

6.5.1 Example NetApp FAS Boot Volume and LUN Configuration

The following section provides detailed procedures for creating a Cisco UCS boot volume and LUN on which an OS image may be installed. In addition it describes the steps for mapping the LUN to an igroup to be accessible from a Cisco UCS server.

Create Thin Provisioned Volume and Boot from SAN LUN

Figure 21 shows a PowerShell Window capture of the commands to create a thin provisioned volume, LUN, and an igroup attached.

Figure 21 PowerShell Window - Boot Volume and LUN Creation

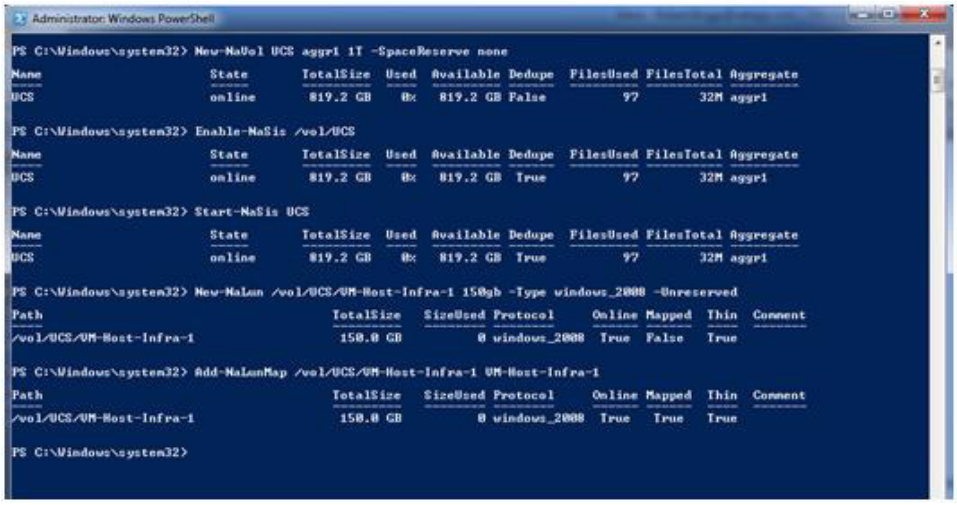

1. Open an elevated Windows PowerShell Window as a Local or Domain Administrator.
2. Type **import-module dataontap** at the prompt and hit enter. Please ensure that you have the most recent version of the NetApp PowerShell Toolkit installed in the **C:\Windows\System32\WindowsPowerShell\v1.0\Modules** directory.
3. Type **Connect-NaController <IP Address of Controller> -cred root** and click Enter. You will be prompted to enter the password for root. Enter the password and click Enter.
4. Type **New-NaVol <volume name> aggr1 1T -SpaceReserve none** and click Enter. This will create a 1 TB Volume called <volume name>. It will be thin provisioned.
5. Type **Enable-NaSis /vol/<volume name>** and click Enter.
6. Type **Start-NaSis <volume name>** and click Enter. This will enable and activate NetApp storage efficiencies (such as deduplication) on this new volume.
7. Type **New-NaLun /vol/<volume name>/<LUN name> 150gb -Type windows_2008 -Unreserved** and click Enter. This will create a 150GB LUN called <LUN name> in volume <volume name>. It will be thin provisioned.
8. Type **Add-NaLunMap /vol/<volume name>/<LUN name> <igroup name>** and click Enter. This will add the igroup <igroup name> to the LUN <LUN name>.
9. The LUN will now be presented to the physical machine with the FC WWPN associated with the igroup you mapped to this LUN. Return to your server and refresh disk manager.

Example NetApp FAS Volume Configuration for Write Cache

The following section provides detailed procedures for creating a volume and LUN for placing individual VM Write Cache VHDs. In addition it describes the steps for mapping the LUN to an igroup to be accessible from a Cisco UCS server.

Create Thin Provisioned Volume and LUN for VM Write Cache

Figure 22 shows a PowerShell Window capture of the commands to create a thin provisioned volume, LUN, and then to attach an igroup to it.

Figure 22 PowerShell Window - Write Cache Volume and LUN Creation

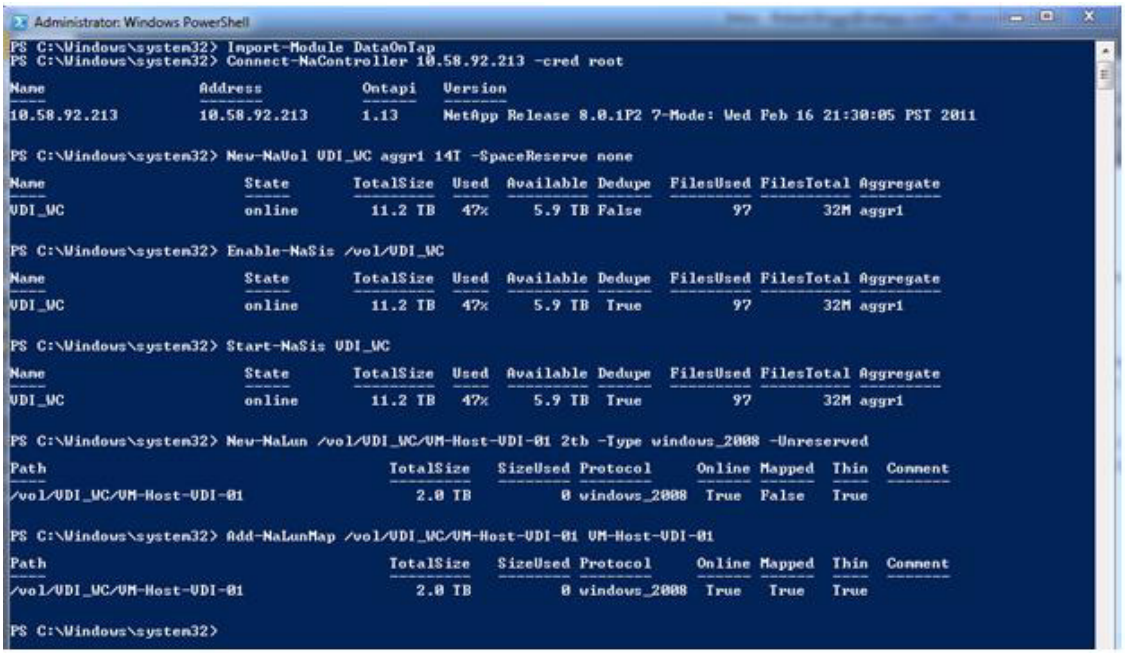

1. Open an elevated Windows PowerShell Window as a Local or Domain Administrator.
2. Type **import-module dataontap** at the prompt and click Enter. Make sure that you have the most recent version of the NetApp PowerShell Toolkit installed in the **C:\Windows\System32\WindowsPowerShell\v1.0\Modules** directory.
3. Type **Connect-NaController <IP Address of Controller> -cred root** and click Enter. You will be prompted to enter the password for root. Enter the password and click Enter.
4. Type **New-NaVol <volume name> aggr1 14T -SpaceReserve none** and click Enter. This will create a 14 TB Volume called <volume name>. It will be thin provisioned.
5. Type **Enable-NaSis /vol/<volume name>** and click Enter.
6. Type **Start-NaSis <volume name>** and click Enter. This will enable and activate NetApp storage efficiencies (such as deduplication) on this new volume.
7. Type **New-NaLun /vol/<volume name>/<LUN name> 2tb -Type windows_2008 -Unreserved** and click Enter. This will create a 2TB LUN called <LUN name> in volume <volume name>. It will be thin provisioned.
8. Type **Add-NaLunMap /vol/<volume name>/<LUN name> <igroup name>** and click Enter. This will add the igroup <igroup name> to the LUN <LUN name>.
9. The LUN will now be presented to the physical machine with the FC WWPN associated with the igroup you mapped to this LUN. Return to your server and refresh the disk manager.

Example NetApp iSCSI LUN Configuration for File and SQL Clustering

This section provides detailed procedures for creating a volume and LUN for providing CSVs for a virtualized instance of SQL or a file server. Since the initiator is a VM, iSCSI is required to attach the CSV LUN. Also detailed will be the process of creating an iSCSI igroup to map to the created LUN.

Create a Thin Provisioned Volume and LUN for a Virtualized SQL or File Server Instance

Figure 23 shows a PowerShell Window capture of the commands to create a thin provisioned volume, LUN, and then attach an iSCSI igroup to it.

Figure 23 PowerShell Window - Virtualized SQL/File Services Cluster Volume Creation

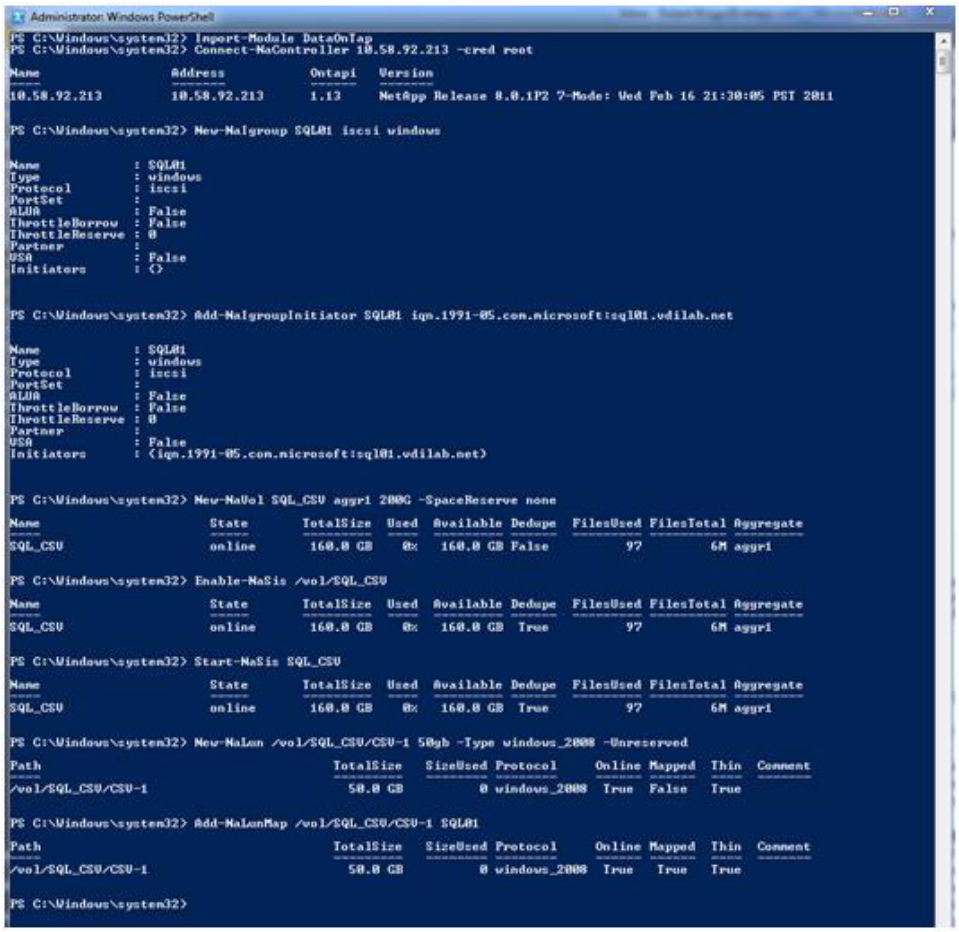

1. Open an elevated Windows PowerShell Window as a Local or Domain Administrator.
2. Type **import-module dataontap** at the prompt and hit Enter. Make sure that you have the most recent version of the NetApp PowerShell Toolkit installed in the **C:\Windows\System32\WindowsPowerShell\v1.0\Modules** directory.
3. Type **Connect-NaController <IP Address of Controller> -cred root** and click Enter. You will be prompted to enter the password for root. Enter the password and click Enter.
4. When connected to the controller, at the prompt, type **New-NaIgroup <igroup name> iscsi windows** click Enter. This will create a new iSCSI initiator group on your controller called <igroup name>.
5. At the prompt, type **Add-NaIgroupInitiator <igroup name> <iSCSI IQN>** and click Enter. Now the <igroup name> igroup will have the iSCSI IQN assigned to it.
6. Type **New-NaVol <volume name> aggr1 200G -SpaceReserve none** and click Enter. This will create a 200GB Volume called <volume name>. It will be thin provisioned.
7. Type **Enable-NaSis /vol/<volume name>** and click Enter.
8. Type **Start-NaSis <volume name>** and click Enter. This will enable and activate NetApp storage efficiencies (such as deduplication) on this new volume.

9. Type **New-NaLun /vol/<volume name>/<LUN name> 50gb -Type windows_2008 -Unreserved** and click Enter. This will create a 50GB LUN called <LUN name> in volume <volume name>. It will be thin provisioned.
10. Type **Add-NaLunMap /vol/<volume name>/<LUN name> <igroup name>** and click Enter. This will add the igroup <igroup name> to the LUN <LUN name>.
11. The LUN will now be presented to the physical or VM with the iSCSI IQN you entered. Return to your server and refresh the disk manager.

Cisco UCS Manager Configuration for Microsoft Hyper-V Server 2008 R2 SP1 and Server 2008 R2 SP1

This section addresses the creation of the service profiles and VLANs to support the project.

There are two types of service profiles required to support two different blade server roles, Table 6.

Table 6 Role, Server and OS Deployment

Role	Blade Server Used	OS Deployed
Infrastructure	Cisco UCS B200 M2	Server 2008 R2 SP1 Enterprise Edition
VDI Hosts	Cisco UCS B230 M2	Hyper-V Server 2008 R2 SP1

To support those roles and hardware platforms, service profile templates were created, utilizing various policies created earlier. The templates were created in such a way as to split the load between the two fabrics.

The service profile templates were then used to quickly deploy service profiles for each blade server in the Cisco Unified Computing System. When each blade server booted for the first time, the service profile was deployed automatically, providing the perfect configuration for Server 2008 R2 SP1 and Hyper-V Server 2008 R2 SP1 installation.

In addition, to control network traffic in the infrastructure, virtual LANs (VLANs) were created on the Nexus 5548s, on the Cisco UCS Manager (Fabric Interconnects,) and on the Virtual Machine Manager 2012. The virtual machines in the environment used the VLANs depending on their role in the system.

The following sections describe the Service Profile Templates that were created and used and the VLAN configuration on Cisco UCS Manager.

Service Profile Templates For Server 2008 R2 SP1 and Hyper-V Server 2008 R2 SP1

There were four service profile templates created in this configuration:

- VM-Host-Infra-A: For B200 M2 Infrastructure Server on Fabric A Running Server 2008 R2 SP1 with Hyper-V Role
- VM-Host-Infra-B: For B200 M2 Infrastructure Server on Fabric B Running Server 2008 R2 SP1 with Hyper-V Role
- VM-Host-VDI-A: For B230 M2 VDI Hosts on Fabric A Running Hyper-V Server 2008 R2 SP1
- VM-Host-VDI-B: For B230 M2 VDI Hosts on Fabric B Running Hyper-V Server 2008 R2 SP1

Figure 24 Service Profile Templates—General Configuration

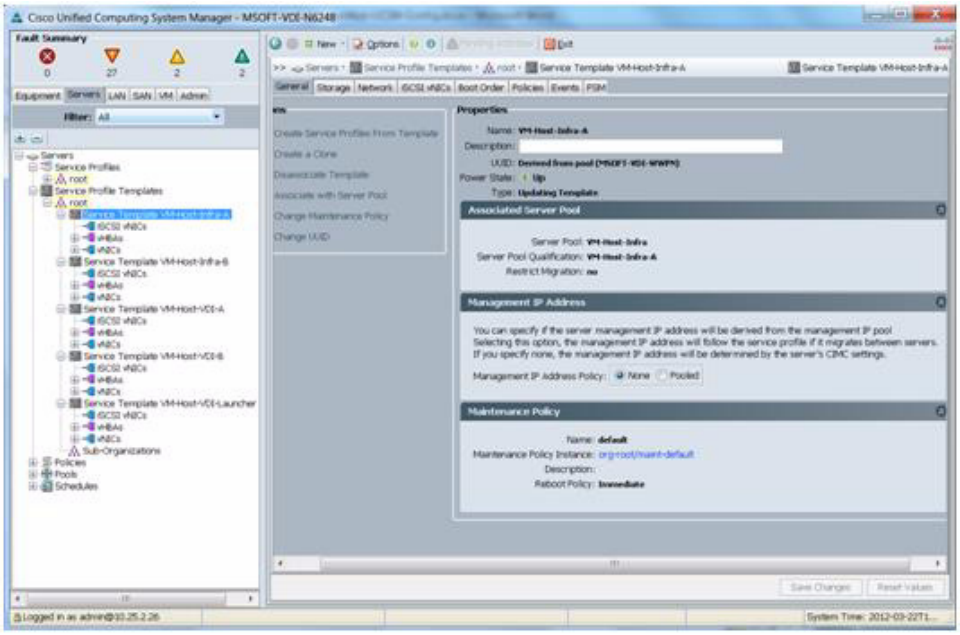

The Server Pool and Server Pool Qualification policy are applied to select the servers that should receive this template. This template is an Updating Template so that changes made to the template will be applied to each associated Service Profile automatically when changes are saved.

Figure 25 Service Profile Templates - Storage Configuration

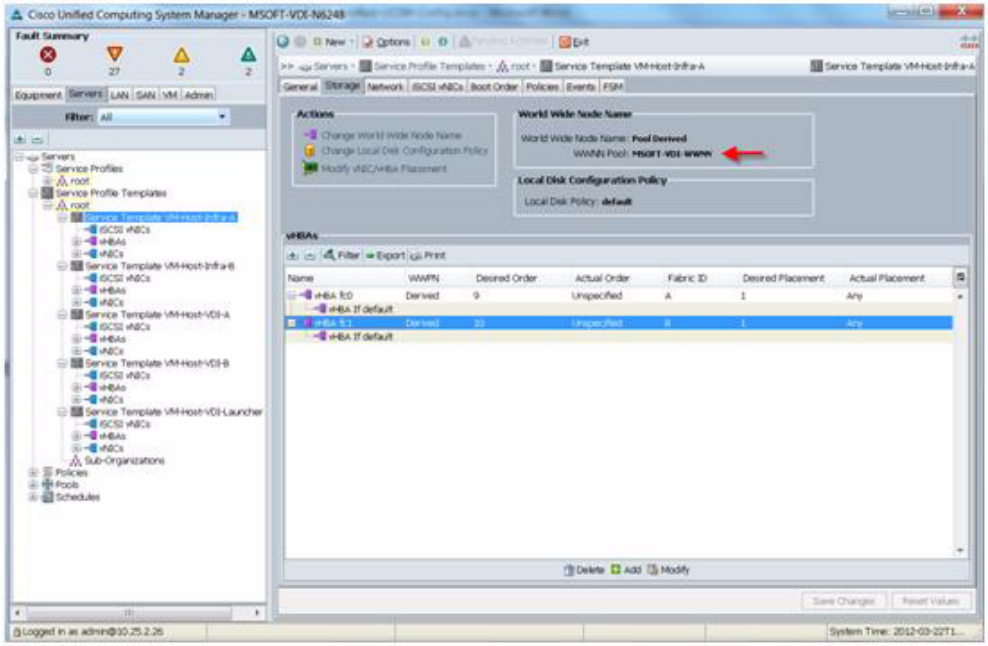

On the Storage tab of the Service Profile Template, the World Wide Node Name Pool that was configured earlier is assigned to the virtual host bus adapters, vHBAs, for both fabrics. Note that vHBA fc0 is assigned to Fabric A while vHBA fc1 is assigned to Fabric B.

Figure 26 Service Profile Templates - Network Configuration

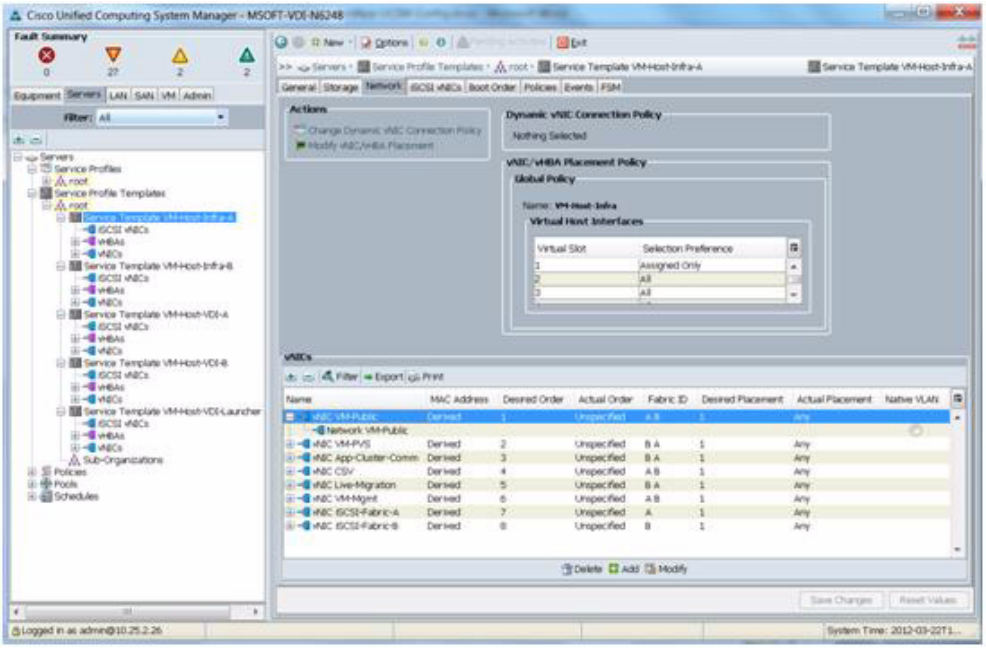

On the Network tab, we create the virtual NICs, vNICS, that the blade will have when a Service Profile derived from this template is applied. The vNICs created here are in support of the VLANs used in the project, which are detailed in the next section. Note the Fabric ID column, where the primary fabric and the failover fabric, following the primary if available are specified.

There is no hardware iSCSI vNIC configuration because Server 2008 R2 SP1 and Hyper-V Server 2008 R2 support software iSCSI initiators only.

Figure 27 Service Profile Templates - Boot Order

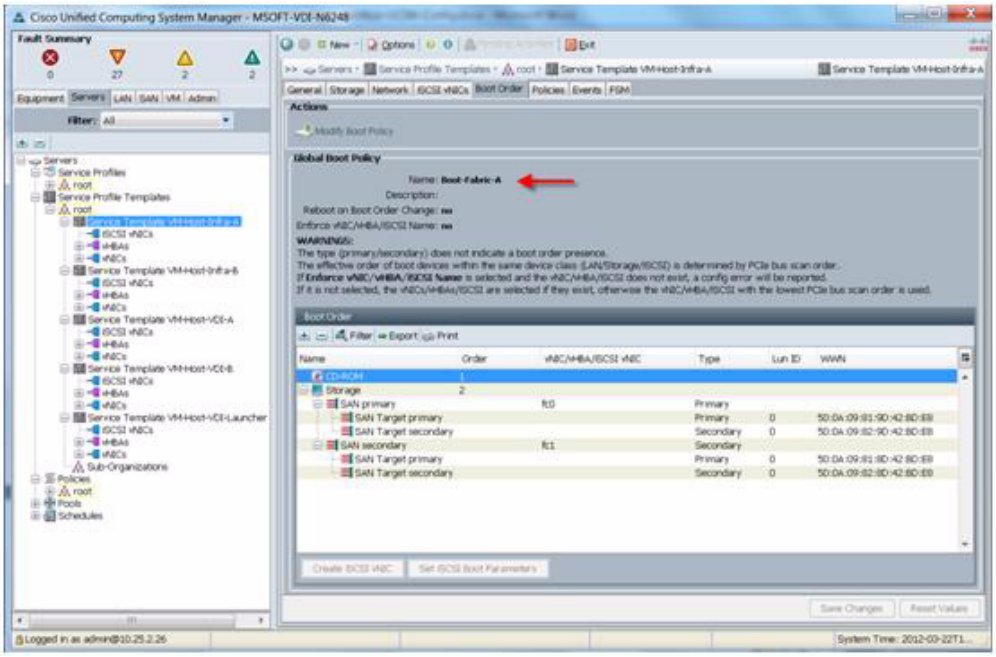

The Boot Order tab on the Service Profile Template utilizes the Boot policies configured earlier in the document.

Figure 28 Service Profile Templates - Policies

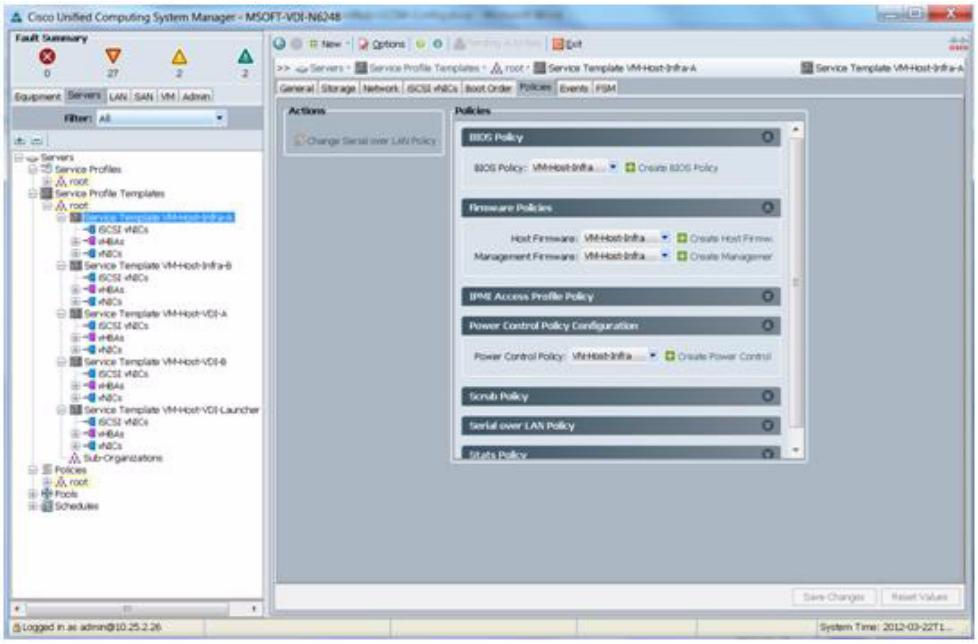

On the Policies tab of the Service Profile Template, we apply the BIOS, Host Firmware, Management Firmware, and Power Control Policies created in earlier sections of the document. Policies not expanded in the figure above were set at defaults.

VLAN Configuration for Hyper-V Server 2008 R2 SP1 and Server 2008 R2 SP1 with Hyper-V Installation

A total of nine Virtual LANs, VLANs, were utilized for the project. Table 7 identifies them and describes their use.

Table 7 VLAN Naming and Use

VLAN Name	VLAN ID	Use
App/Cluster	2222	Server and Application Cluster Communications
Public	2233	VDI Virtual Machine Communications
VM-PVS	2244	Legacy NIC Provisioning Server Boot Communications
Management	2255	VMM 2012, Hyper-V Server 2008 R2 SP1 Communications
CSV	2266	Clustered Shared Volume Communications for SQL and Profiles
Live Migration	2277	Virtual Machine Migration Communications
iSCSI-A	2288	iSCSI Fabric A Communications
iSCSI-B	2299	iSCSI Fabric B Communications

VLANs are configured in Cisco UCS Manager on the LAN tab, LAN\VLANs node in the left pane of Cisco UCS Manager. Figure 29 details the VLANs created for the project.

Figure 29 Cisco UCS Manager VLAN Configuration

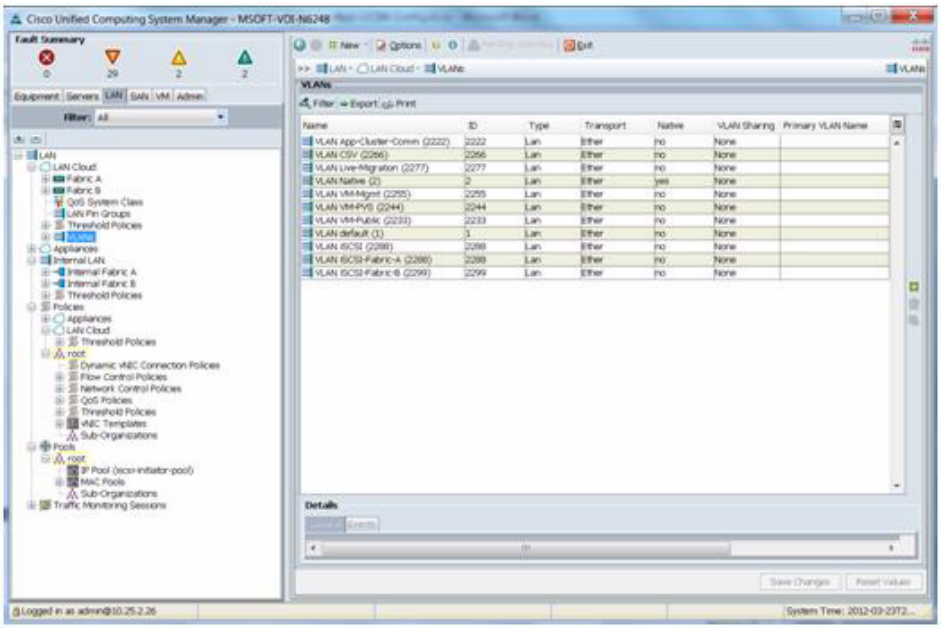

Note VLAN Default and VLAN Native were not used in the project. VLANs can be added by using the green + symbol at the right edge of the VLANs right pane. Similarly, VLANs can be removed by highlighting a VLAN and clicking the trash can icon at the right edge of the VLANs right pane.

Installing Microsoft Server 2008 R2 SP1 and Hyper-V Server 2008 R2 SP1 Using Cisco UCS Manager

OS Base Installation

The following are the installation steps for Microsoft Server 2008 R2 and Hyper-V Server 2008 R2 SP1.

1. In Cisco UCS Manager launch the KVM view for the server where the OS will be installed.
2. Under Virtual Media tab, click Add Image:
   a. Select the ISO for the Microsoft Hyper-V Server 2008R2 SP1.
   b. Select the ISO for the Cisco UCS drivers: 2.0.1e.
   c. Check the Mapped box for the Hyper-V server ISO.

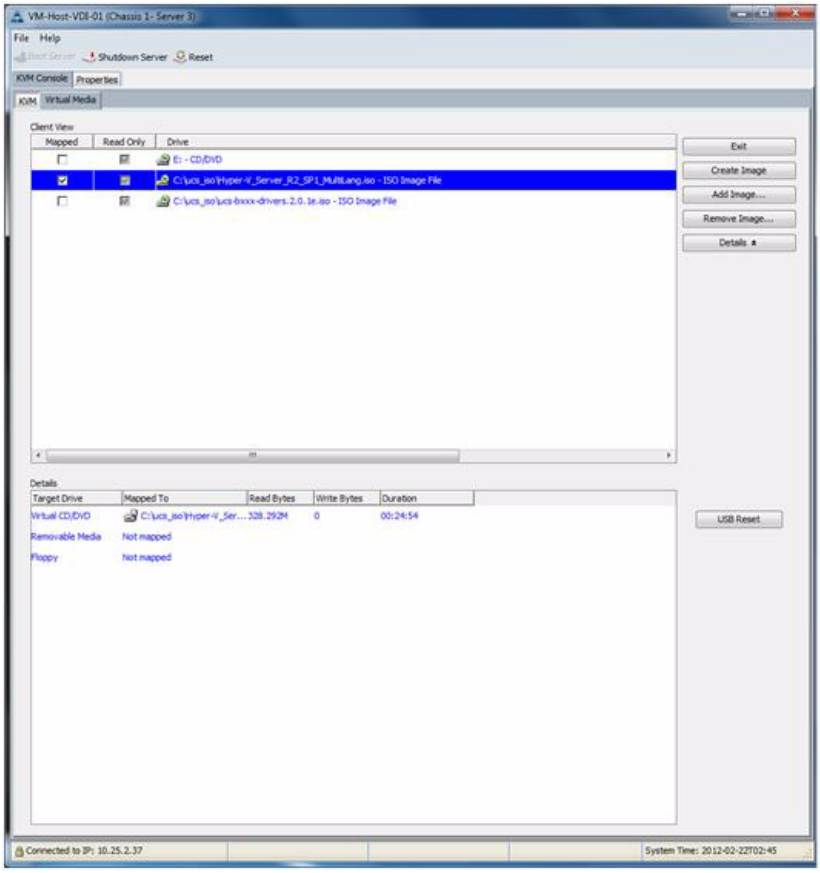

3. Go back to KVM tab.
4. Boot up the OS.
5. Select the language / local.
6. Click Install now.
7. Read and accept the Microsoft Software License Terms.
8. Select Custom (advanced) for type of installation.
9. The OS will not be able to locate the target drive to install the OS.
10. This is the point where you will need to install the required drivers.
11. Switch to the Virtual Media tab.
12. Unchecked the box for Hyper-V OS and check the box for the Cisco UCS drivers.

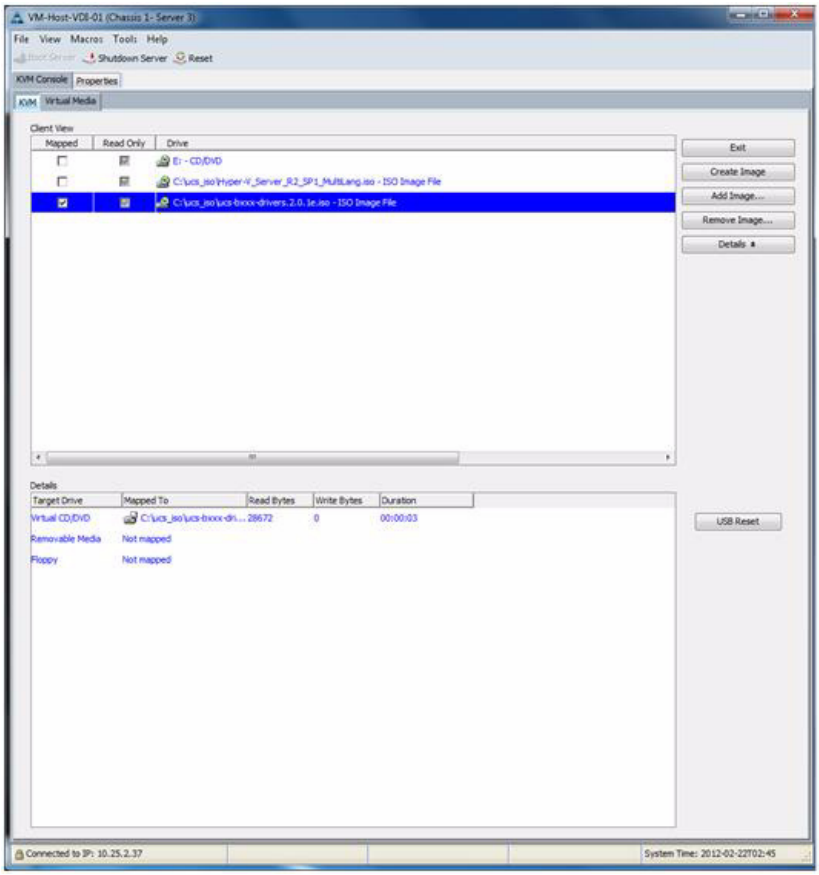

13. Switch back to the KVM tab.
14. Click Load Driver.
15. Under the CD-Rom, browse to Storage>Cisco>M81KR>W2K8R2>x64.
16. Click OK.

    The driver should be identified as shown below.

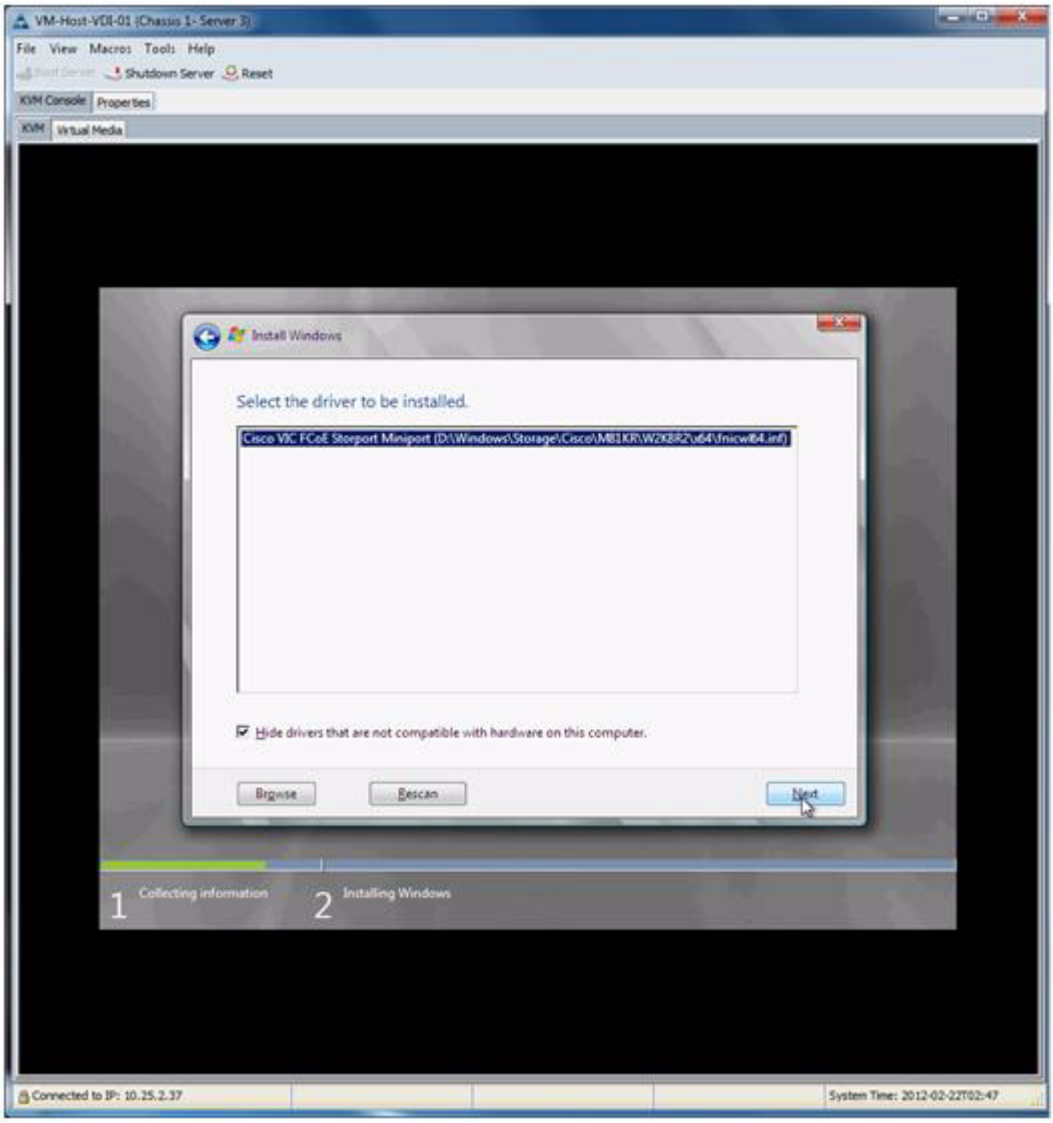

17. Click Next.
18. Repeat the process to install the Network driver.
19. Click Load Driver.
20. Under the CD-Rom, browse to Network>Cisco>M81KR>W2K8R2>x64.
21. Click OK.
22. The driver should be identified as shown below.

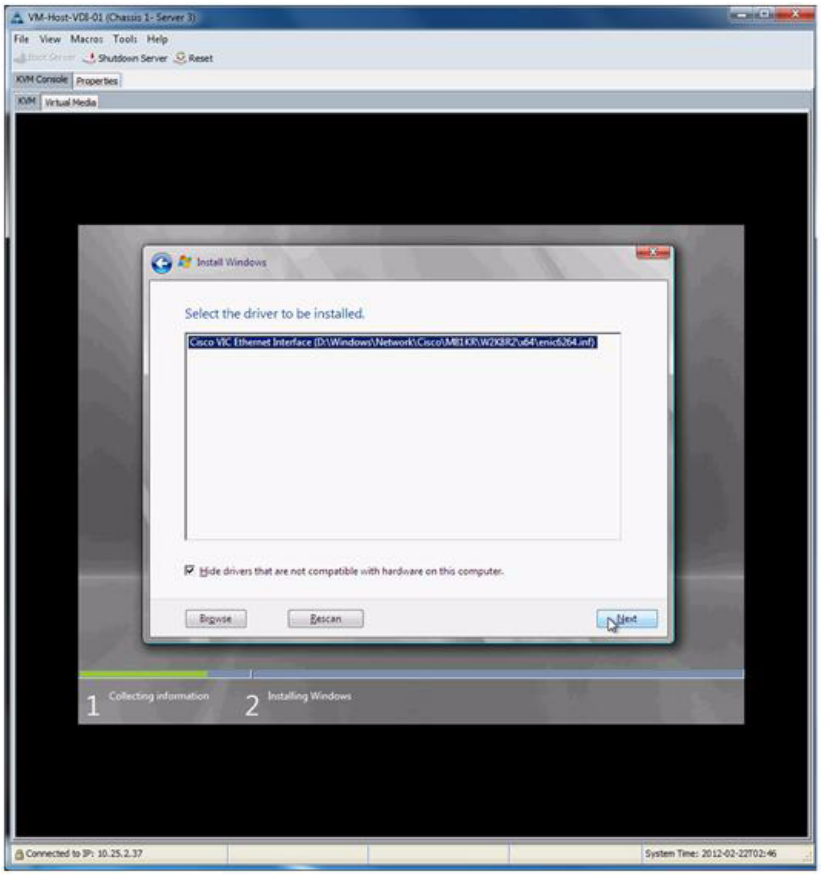

23. The windows installer should be able to locate the disk driver for the target OS. If that is not the case, click Refresh and that should locate the disk drive.
24. Before proceeding with the OS install, make sure you switch the ISO back to the hyper-V Server.
25. Switch to the Virtual Media tab.
26. Uncheck the ucs drivers ISO.
27. Check the Hyper-v Server ISO.
28. Switch back to KVM tab.
29. Highlight the drive.
30. Click Next.
31. When the OS installation completes, you will be prompted to change the Administrator password.
32. When completed, you are ready to proceed to configuration.

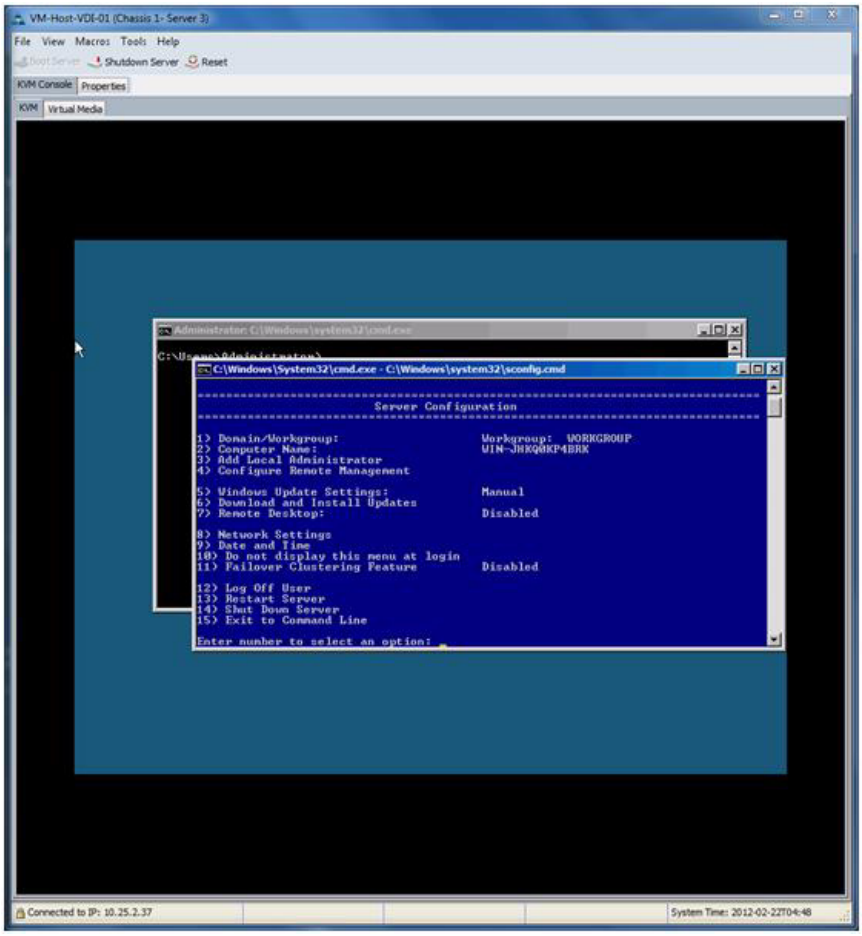

Note By default, blue sconfig.cmd shell is launched inside Microsoft Hyper-V server. This tool is very handy for configuration changes.

Configuration for the Pre-sysprepped Image

For configuration there are multiple ways to perform tasks, some of which are: using the command line, sconfig, or Core Configurator. The Core configurator is a very handy mini GUI built on top of powershell that can be downloaded here.

To enable RDP using sconfig:

1. Select 7 (remote desktop).
2. Select E for Enable.
3. Select 2 to allow any client version to connect.
4. Click OK.

To enable remote management using sconfig:

1. Select 4 to configure remote management.
2. Select 1 to allow MMC remote management.

3. Click OK.
4. Select 2 to Enable Windows PowerShell.
5. Click yes to restart server after installation.
6. When the server is restarted, select 4.
7. Select 3 to allow Server Manager Remote Management.
8. Click OK.
9. Select 5 to go back to the main menu.

To install the MPIO feature; sconfig does not facilitate this task. This task can be accomplished through the command line or Core configurator, if that was installed on the server.

1. Using command line: ocsetup MultipathIo.
2. Using the Core configurator:
   a. Browse to Core configurator directory.
   b. Start core configurator by typing Start_Coreconfig.wsf.
   c. Click computer settings.
   d. Click n Add or Remote Roles.
   e. Select MultipathIo and click Apply.

Sysprep generalize; this task is not available through sconfig or core configurator.

1. Browse to %homedrive%\Windows\System32\sysprep.
2. Type: sysprep /generalize /oobe /shutdown.
3. After the server is shutdown, clone the LUN and present to all 14 VDI hosts.

Configuration for the Post-sysprepped Image

To configure the post-sysprepped image, do the following:

1. Boot up all the VDI hosts, they should go through mini setup. After mini setup is complete, logon to the server to complete the configuration. With each of the server having three adapters—Management, Public and PVS—it is important to start by identifying those for a successful and smooth configuration.

Note The following tasks have to be completed on all 14 VDI hosts

2. Using your preferred editor, create a table of the following format and update that as you build the environment, this will be very handy in the process below. This is the most labor intensive part of the process, once this is identified, the rest is smoother. Customer scripting would always help in the process. We will use host 14 as an example.

Host Name	Adapter Name	Adapter Description	MAC address	VLAN
Vm-host-vdi-14				

3. Start by identifying the NICs to VLANs mapping.
4. Run Ipconfig /all on the host and update the table with the relevant information.

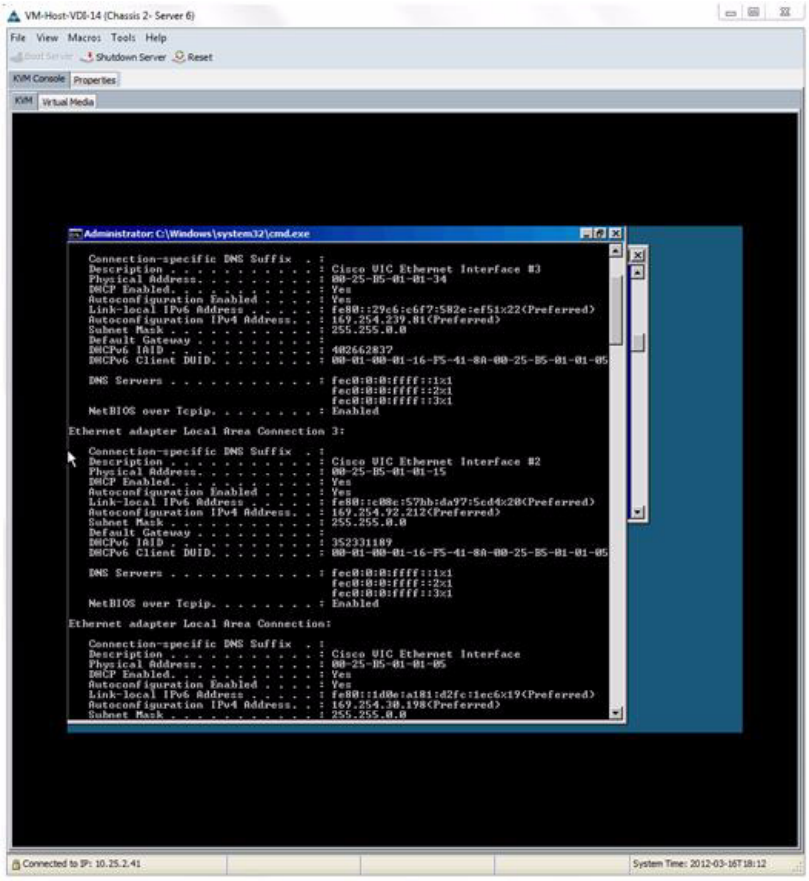

5. Fill in the sections as identified in Table 8.

Table 8 vNIC-MAC Address Mapping

Host Name	Adapter Name	Adapter Description	MAC address	VLAN
Vm-host-vdi-14	Local Area Connection 2	Cisco VIC Ethernet Interface 3	00-25-B5-01-34	To be determined
Vm-host-vdi-14	Local Area Connection 3	Cisco VIC Ethernet Interface 2	00-25-B5-01-15	To be determined
Vm-host-vdi-14	Local Area Connection	Cisco VIC Ethernet Interface	00-25-B5-01-05	To be determined

Note You might notice that the Public/PVS networks will not have IP addresses even though there is a DHCP server on those networks. This is because the two virtual NICs were not set as native networks in the Cisco UCS Manager. You will need to enable VLAN tagging before they can successfully communicate on the network.

6. Identify NIC to VLAN mapping:

a. From the Cisco UCS Manager, go to the Properties tab.

b. Click the Network tab.

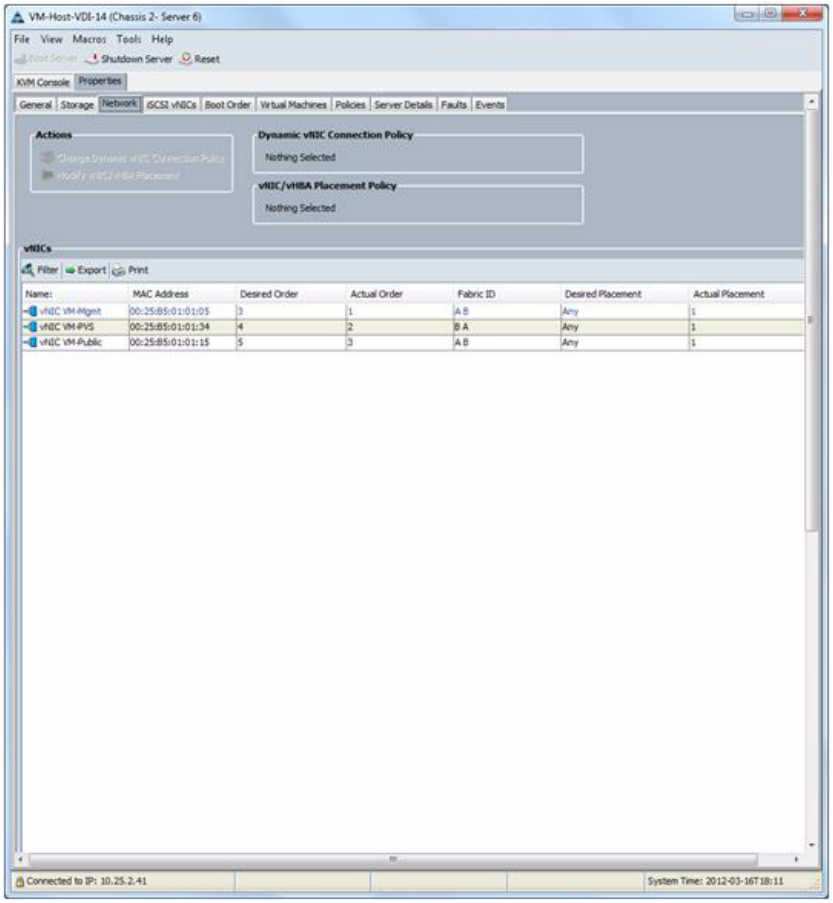

7. Using the above mapping between MAC and VLANs, update the table with the correct VLANs.

Host Name	Adapter Name	Adapter Description	MAC address	VLAN
Vm-host-vdi-14	Local Area Connection 2	Cisco VIC Ethernet Interface 3	00-25-B5-01-34	PVS
Vm-host-vdi-14	Local Area Connection 3	Cisco VIC Ethernet Interface 2	00-25-B5-01-15	Public
Vm-host-vdi-14	Local Area Connection	Cisco VIC Ethernet Interface	00-25-B5-01-05	Mgmt

8. Rename NICs to SubnetNameNIC. We added a NIC suffix to the name to differentiate that from the virtual switch that will be created later on (use the subnet name).

9. Using the command line run the following commands:

a. Netsh int set int “local area connection 2” newname=”PVSNIC”

b. Netsh int set int “local area connection 3” newname=”PUBLICNIC”

c. Netsh int set int “local area connection” newname=”MGMTNIC”

Note You can also accomplish this task using core configurator.

10. After renaming the NICs update the table to map the changes.

Host Name	Adapter Name	Adapter Description	MAC address	VLAN
Vm-host-vdi-14	PVSNIC	Cisco VIC Ethernet Interface 3	00-25-B5-01-34	PVS
Vm-host-vdi-14	PUBLICNIC	Cisco VIC Ethernet Interface 2	00-25-B5-01-15	Public
Vm-host-vdi-14	MGMTNIC	Cisco VIC Ethernet Interface	00-25-B5-01-05	Mgmt

11. Set the IP for the management NIC. Use the following command line to complete this task:

```
NETSH INT IPV4 SET ADDR MGMTNIC STATIC IPaddr ubnetmask
NETSH INT IPV4 SET ADDR MGMTNIC STATIC 10.13.0.113 255.255.0.0
```

Note You can accomplish this task via sconfig and Core configurator as well. Also, you do not need to setup the PVS and Public NIC at this point of time, since they will not communicate on the network prior to VLAN tagging those on the Virtual Switch.

12. Rename the server using sconfig:
    a. Selection option 2 (Computer Name.).
    b. Input server name (vm-host-vdi-14)
    c. When it prompts you to restart, click Yes.

Note You can accomplish this task via command line (netdom) or via Core Configurator.

When the server restarts, you can proceed with the rest of the configuration.

13. Outside of custom scripting, the only option to configure the virtual switch is from a remote server that has the hyper-V console installed. Execute the following steps from a remote server with Hyper-V console installed:
    a. Open the Hyper-V Manager console.
    b. Right-click the hyper-V Manager node.
    c. Click Connect to Server….
    d. Click the Another computer radio button.
    e. Type the IP of the server to connect to. In this example 10.13.0.114.

Note If you get an RPC error, make sure that "remote management" performed on pre-sysprepped Image is enabled.

14. After successfully connecting remotely to the target VDI host, proceed with creating the virtual switch for the PUBLIC network as follows:
    a. Right-click Virtual Network Manager…under the Actions pane.
    b. Make sure External network type is highlighted.
    c. Click Add .
    d. Under Name filed, type PUBLIC.
    e. Under the external radio button selection, use the table created previously to identify which NIC to use.

Host Name	Adapter Name	Adapter Description	MAC address	VLAN
Vm-host-vdi-14	PVSNIC	Cisco VIC Ethernet Interface 3	00-25-B5-01-34	PVS
Vm-host-vdi-14	PUBLICNIC	Cisco VIC Ethernet Interface 2	00-25-B5-01-15	Public
Vm-host-vdi-14	MGMTNIC	Cisco VIC Ethernet Interface	00-25-B5-01-05	Mgmt

    f. In this example PUBLIC should map to Cisco VIC Ethernet Interface 2.
    g. Make sure "Allow management operating system to share his network adapter" is checked.
    h. Check "Enable virtual LAN identification for management operating system."
    i. Fill in the VLAN ID for the public network. In this environment it will be 2233.
15. Repeat the same task for the PVS network, with the following minor differences:
    a. Right-click the "Virtual Network Manager…" under the Actions pane.
    b. Make sure External network type is highlighted.
    c. Click Add.
    d. Under Name filed, type PVS.
    e. Under the external radio button selection, use the table created previously to identify what NIC will be used.
    f. In this example PVS maps to "Cisco VIC Ethernet Interface 3."
    g. Uncheck "Allow management operating system to share his network adapter"."
    h. Check "Enable virtual LAN identification for management operating system."
    i. Fill in the VLAN ID for the public network. In this environment it will be 2244.

Note If you mistakenly bound the virtual switch to the management adapter. Now you can connect remotely to fix the issue but cannot locally fix the problem. What should you do? Two very helpful tools can be used. NVSPBIND can be downloaded here and NVSPCRUB can be downloaded here.

16. When the virtual switches for the PVS and the Public have been created, you might notice that the IPV4 protocol was not bound from the NICs. This occurs on Hyper-V server when configured remotely.
17. To overcome that, reboot the server after configuring the virtual switches.
18. After the server restarts, you can proceed with the rest of the configuration.
19. Rename the Virtual switch adapter name as follows:
    a. Usually the adapter name would default to “Local Area Connection.”
    b. Netsh int set int “local area connection” newname=”PUBLIC”
20. Set IP address and DNS server for PUBLIC adapter
    c. NETSH INT IPV4 SET ADDR PUBLIC STATIC 10.11.0.114 255.255.0.0 10.11.0.1
    d. NETSH INT IPV4 ADD DNSS PUBLIC ADDRESS=10.11.0.6 INDEX=1
    e. NETSH INT IPV4 ADD DNSS PUBLIC ADDRESS=10.11.0.5

    This is not needed for PVS (since the host doesn't use that network).
21. Join the domain using sconfig as follows:
    a. Choose Option 1 (Domain/Workgroup).
    b. Choose D (for domain).
    c. Enter the domain name. In this case it is VDILAB.NET.
    d. Enter the authorized user to join the domain.
    e. When prompted, enter the password.
    f. Click No, when prompted to change the computer name.
    g. Click No, when prompted to restart the server.

Note This can also be completed using command line (netdom) as well as core configurator

22. Configure MPIO as follows:”
    a. On the command line run MPIOCPL
    b. Go to "Discover Multi-Paths” tab.
    c. Under Device Hardware ID select NETAPP LUN.
    d. Click Add.
    e. Do not restart when prompted

Installing NetApp MPIO DSM

Run the installation program to install the DSM code and to set required parameters for HBAs in Windows registry.

Before you Begin

This process is for new DSM installations.You must have already completed the following tasks:

- Close all applications and stop I/O.
- For hosts running Microsoft Clustering, stop the Cluster Service.
- Obtain an MPIO license key from NetApp.
- Back up any critical data on your host.
- Enable the Windows Multipath I/O feature.

A reboot of the Windows host is required to complete the installation.

The installation program displays the current and new versions of DSM and Windows MPIO Components. However, the installation program never installs new Windows MPIO components for Windows Server 2008 R2 SP1.

Steps

1. Change to the directory to which you downloaded the executable file.
2. Launch the installation program and follow the instructions on the screen.
3. Enter the MPIO license key when prompted.
4. Select the Use the default system account check box. Or optionally enter the user name and password of the account on the Windows host under which the DSM management service will be logged on. This account must be in the Windows Administrators group.
5. The DSM service requires an Administrator-level account to allow it to manage disks and paths on the Windows host.
6. Choose whether to install the Hyper-V Guest Utilities.
7. When prompted, click Yes to reboot the Windows host and complete the installation.

After you Finish

If the installer reports a problem, such as a required hotfix not found, correct the problem and run the installation again.

The installation program might also display a message instructing you to install Windows hotfixes after installing the DSM. If so, download the specified hotfixes from the Microsoft support site and install them.

Use Windows Disk Management to verify that all existing disks are online. If any disks are offline, set them online.

Finish the Installation

To finish the installation, do the following:

1. Restart server using sconfig.
2. Option 13.
3. Click OK when prompted to restart.

Installing and Configuring Virtual Machine Manager 2012

This section provides a high-level walkthrough on how to prepare Virtual Machine Manager Pre-requisites, install Virtual Machine Manager and configure Virtual Machine Manager for use in the VDI setup.

Setting up VMM Pre-Requisites

VMM requires the .NET Framework 3.5.1 feature, Windows Automated Installation Kit, and SQL Server 2008 R2 Command Line Utilities be installed on the server before proceeding.

Install .NET Framework 3.5.1

To install .NET Framework 3.5.1, do the following:

1. In server Manager, browse to feature.
2. Click Add Feature.
3. Check.NET Framework 3.5.1.]

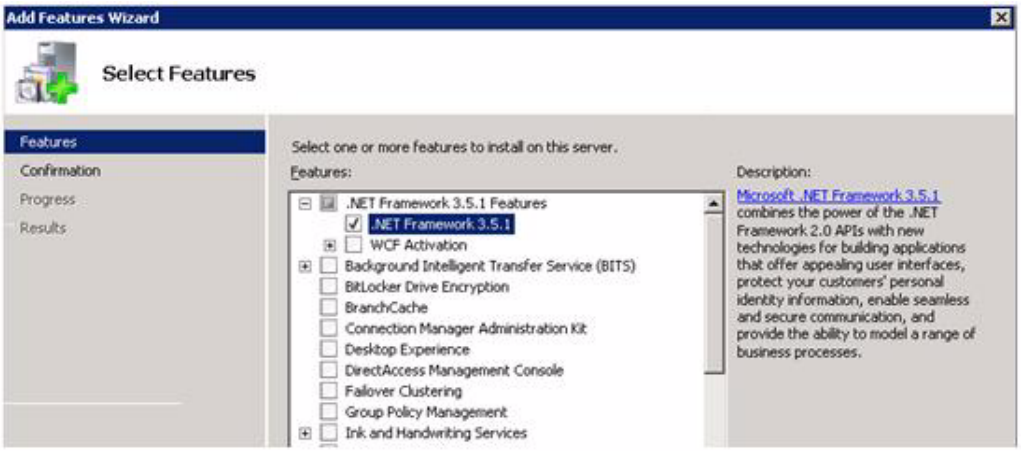

4. Click Next.
5. At the confirmation page, click Install.

Windows Automated Installation Kit

To install the Microsoft Window Automated Installation kit, do the following:

1. Insert CD for WAIK for Windows 7.
2. Browse into CD/ISO and double click on Start CD
3. Click Windows AIK Setup.

4. On the Welcome Screen click Next.
5. Read the License Terms; you will need to select "I Agree" to proceed.
6. Click Next.
7. Select the Installation folder and click Next.
8. On the confirmation page click Next.
9. When setup is complete click Next.

SQL Server 2008 R2 Command Line Utilities

To instal the SQL Server 2008 Native Client, do the following:

1. Command line utilities require SQL Server 2008 R2 Native Client be installed prior to proceeding.

2. Follow the on screen instructions to install SQL Server 2008 R2 Native Client. Under the feature selection make sure you pick SQL Server Native client SDK

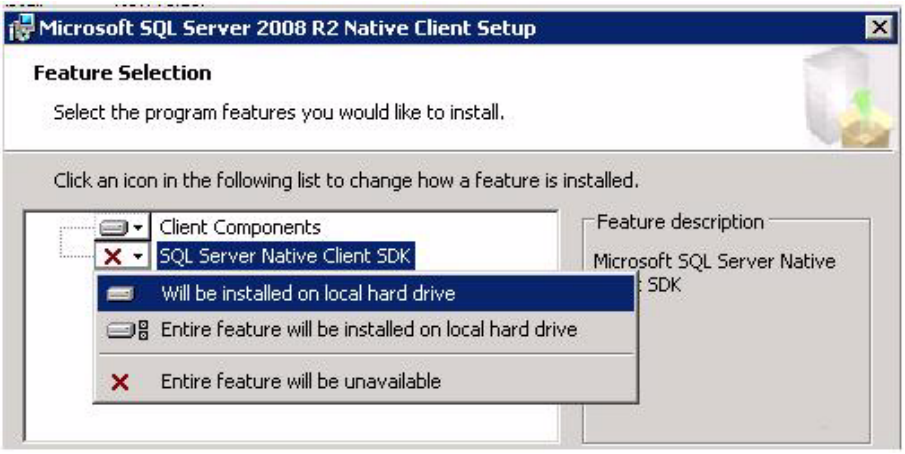

3. After completion proceed to installing SQL Server 2008 R2 Command Line Utilities.
4. After completion you are ready to start with VMM installation.

Installing Virtual Machine Manager

1. Browse to the VMM 2012 CD and start setup.exe.
2. On the initial screen click Install.

3. On the feature selection page pick VMM management server and VMM console; then click Next.
4. The VMM console is processor and memory intensive; we strongly advise that you also install the VMM Console on a separate workstation.
5. On Product registration information page, enter the details required; then click Next.
6. Read the license agreement. You will need to agree with the terms before you can proceed. Click Next.
7. Choose if you would like to join the customer experience improvement program and click next.
8. It is advised that you turn on Microsoft Update feature. Click Next

9. Choose the program file path. Click Next.
10. The prerequisites should be validated successfully.
11. In the Database configuration:
    a. Point it to the already created SQL server (in this case it is SCSQL).
    b. Select the instance Name (in this case it is SCSQL).
    c. Select New database and enter the name.

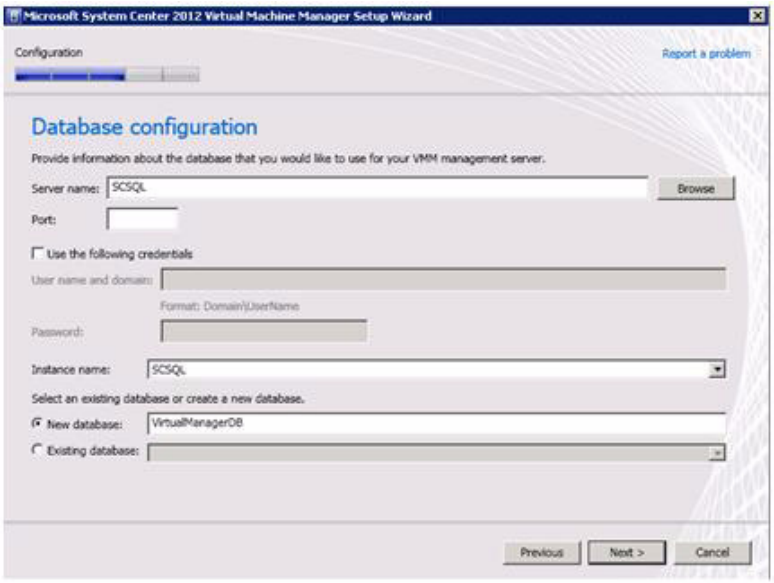

12. For the service account configuration, enter the pre-created domain user with local admin rights on the VMM server. Enter the password and click Next.
13. For the port configuration, use the default port configuration and click Next.
14. For the library configuration, create a new library share, use the defaults and click Next.
15. At the Installation summary, validate all information and click Install.
16. After VMM installation is complete, make sure you install the cumulative update 1. It is critical for VDI deployments.

Configuring Virtual Machine Manager 2012

This section tackles the configuration required after completing the installation. On a high-level, three actions are needed:

- Add the VDI hyperV host to the VMM server
- Configure virtual machine placement location
- Configure VLANs in the VMM server.

Add Hyper-V Hosts

Now that VMM is installed and the 14 Hyper-V hosts have been fully configured, it is time to add those hosts to VMM.

1. Right-click on “All Hosts” and click “Add Hyper-V Hosts and Clusters.”

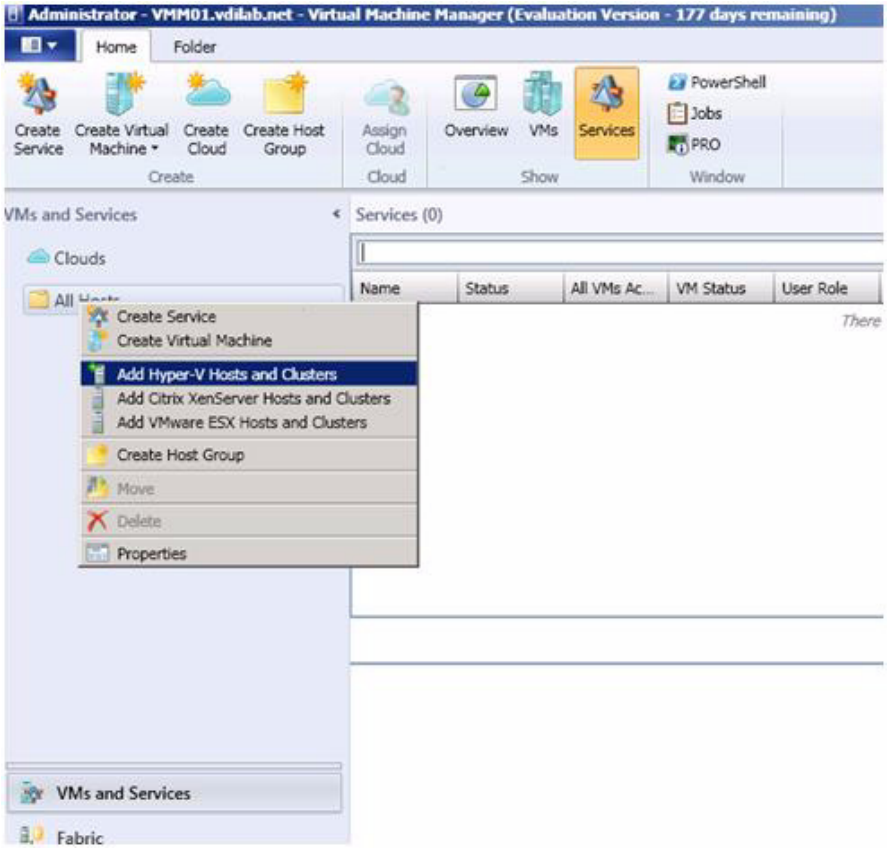

2. For the resource location, leave the default and click Next .
3. For the credentials, you will be required to use a "Run as" account that is different than the VMM service account. Click Browse.
4. Click Create Run As Account.

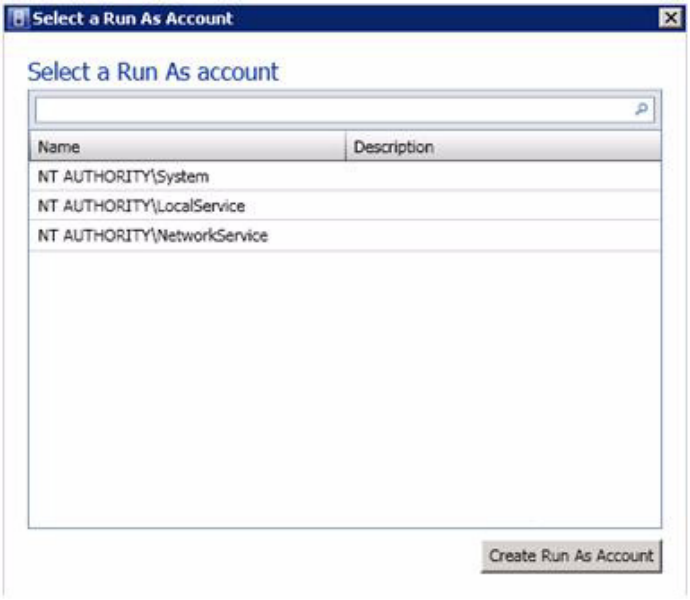

5. Enter the details required for Run As. The account needs to be a domain user. After completion, click OK.
6. Select the newly created account and click Next.
7. For the discovery scope, select Specify Windows Server computers by names, enter computer names (or host suffix name) and click Next.

8. For the target resources, check all 14 servers to be managed by VMM (Select all button). Click Next.
9. For host settings, accept defaults and click Next.
10. Review summary page and click Finish.

Configure Virtual Machine Placement Location

This section will explain how to setup the virtual machine placement. This is the path used by VMM to place the Virtual machine metadata (xml) on the target hosts. By default this path is: C:\ProgramData\Microsoft\Windows\Hyper-V

In our case, all the virtual machines should be located on drive V: on the VDI hosts.

Note The process below needs to be repeated for each of the 14 VDI hosts managed by the VMM server.

1. Right-click on Host 1 (in this case it is Vm-host-vdi-01) and click Properties.
2. On the right side, click on Placement.
3. Click the Add button.
4. If the hyper-V is correctly configured, you should be able to see the V: drive (LUN to be used for VM meta-data and VHDs). Highlight V and click OK.
5. Highlight C:\ProgramData\Microsoft\Windows\Hyper-V and click Remove.
6. Click OK.

Configure Logical Networks Sites and VLANs

This section explains how to setup the Network Sites and VLAN tags to be used by VMM. This is needed by VMM to identify any VLAN tagging required on a virtual machine's virtual adapter.

1. Click on the Fabric category on the left lower corner.
2. Browse to Networks and Logical Networks.

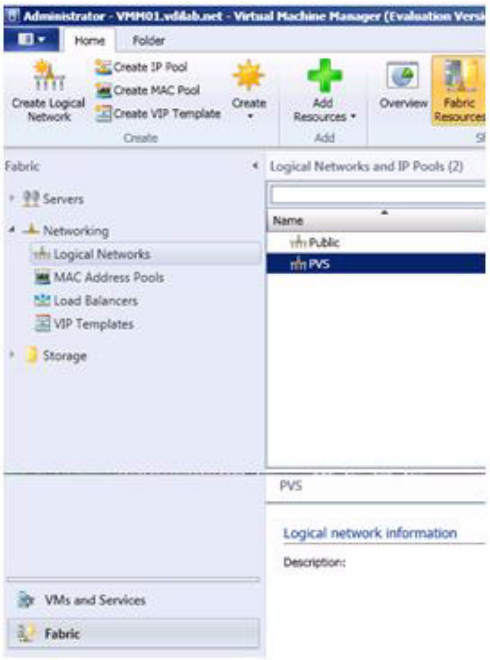

3. Right-click on the Public Network and click properties.
4. Switch to Network Site.
5. Click Add.

a. Name it Public.

b. Select All Hosts.

c. Insert a new row describing the VLAN and the associated IP subnet.

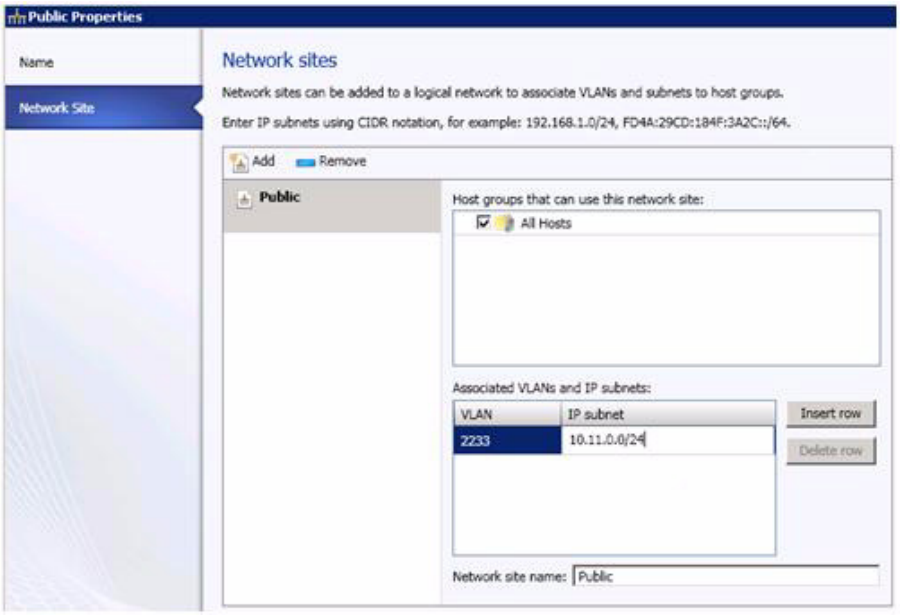

6. Click OK.

7. Repeat the same process for the PVS network and enter the associated VLAN tag/Subnet.

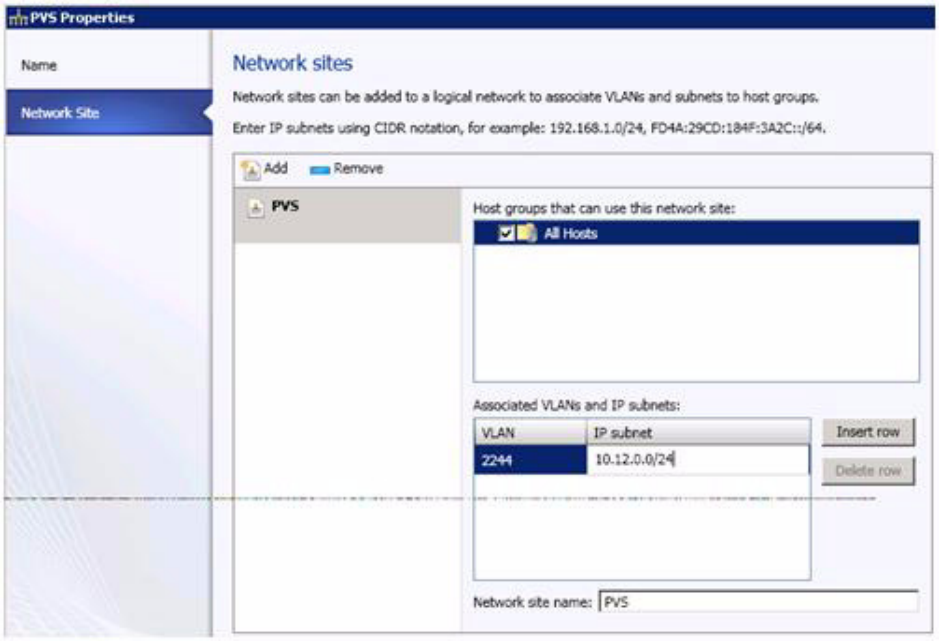

Create Virtual Desktop Virtual Machines, Write-Cache Drives and Export VM Properties

To create the virtual desktop shell machines, the user can use the Provisioning Services 6.1 wizard. The current wizard does the job but takes an extended period of time to provision 2000 virtual desktops. To showcase System Center 2012 Virtual Machine Manager and NetApp capabilities of thin provisioning and cloning, we opted to use a set of PowerShell scripts that are efficient and make use of the NetApp full potential. This section describes the process.

The following figure outlines the process flow for this section.

Figure 30 Machine Creation Process

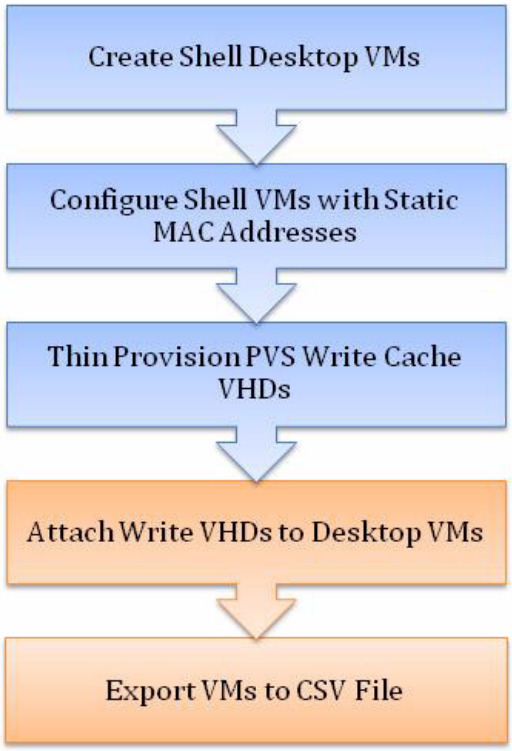

Create Virtual Machine Shell Devices

Copy the scripts identified in Appendix B to the VMM server intended to manage the virtual desktop environment.

1. Launch the powershell console from the VMM console by clicking Power Shell; this will automatically import the required VMM module.

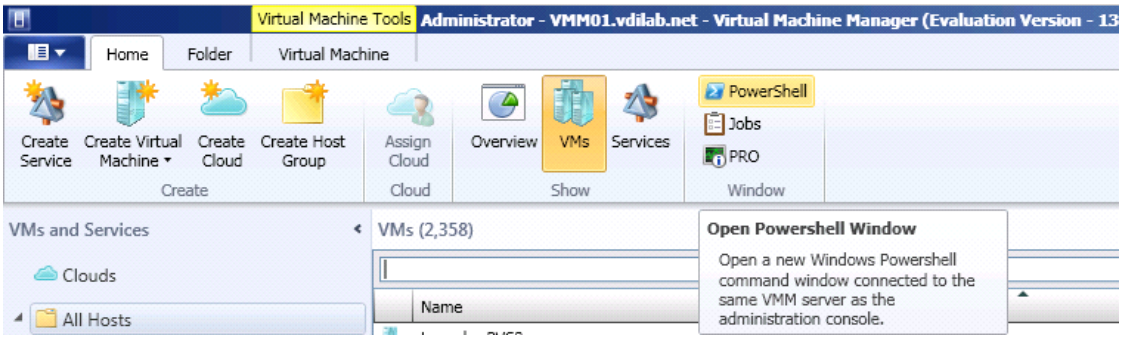

2. If you opt to start the console out of band, make sure you import the required module: import-module VirtualMachineManager.psd1.
3. Run the following script from the powershell console: GenVMs.ps1.
4. This is the initial script that will create the shell VMs in VMM server and on the VDI Hyper-V hosts.
5. The script will prompt you to enter the parameters manually, type y. Enter the appropriate parameters. The figure below shows the variables that were applicable to our environment.

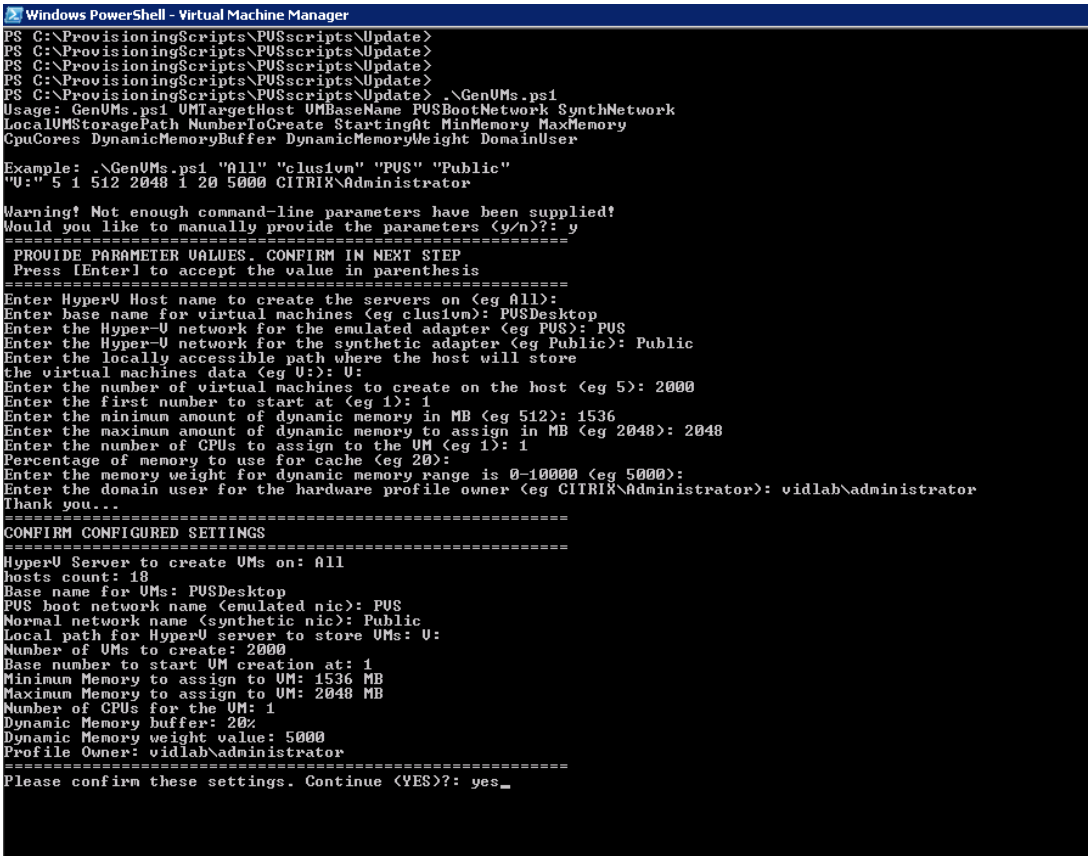

The values in parenthesis are the defaults. If you don't enter a value, the default will be used. This step should take approximately 1.5 hours. That is dependent on the overall system performance. When this step is completed, you can proceed to following step.

Create Fixed Write-Cache VHD using NetApp Storage

Microsoft Hyper-V allows us to create three different types of Virtual Disks (VHD). These are:

- **Fixed Size VHD**—this type of virtual disk provides better performance and is recommended for servers running applications with high levels of disk activity. Once created, the virtual disk consumes the full amount of disk space and does not expand or shrink.
- **Dynamic VHD**—this type of virtual disk provides better use of physical storage and is not recommended for servers running applications with high levels of disk activity. Once created, the virtual disk is small and expands as data is written to it.
- **Differencing VHD**—this type of virtual disk is associated in a parent-child relationship with another disk that is meant to be left intact. Changes can be made to the data or operating system without affecting the parent disk. This allows for easy changes.

Excluding Differencing VHDs, NetApp Storage gives the ability to thin provision a fixed size virtual disk. That is to say, a virtual disk can be created as a fixed size; but it has the ability to expand as data is written to it. This gives the benefit of fixed size performance with the space saving attributes of a dynamically expanding virtual disk. Fixed size thin provisioned disks were used for Provisioning Services write-cache disks for this study.

Creating a Fixed Size Thin Provisioned Virtual Disk

The following figure shows a PowerShell Window capture of the commands to create a thin provisioned virtual disk.

Figure 31 PowerShell Window—Create Fixed Size Thin Provisioned VHD

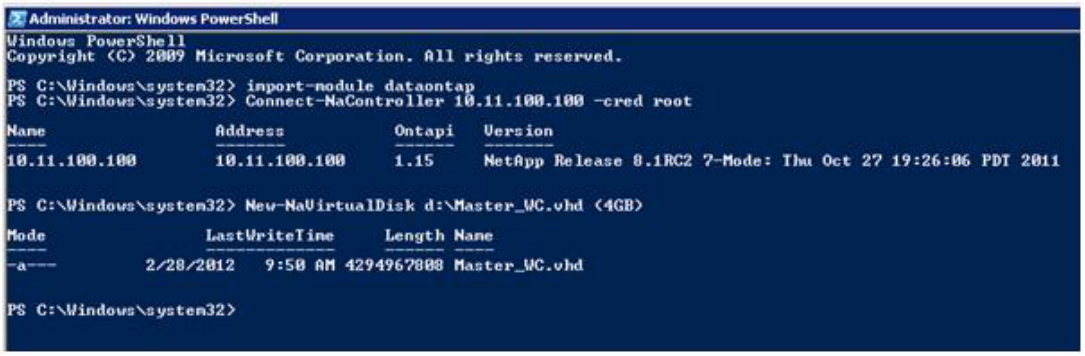

1. Open an elevated Windows PowerShell Window as a Local or Domain Administrator.
2. Type import-module dataontap at the prompt and hit enter. Please ensure that you have the most recent version of the NetApp PowerShell Toolkit installed in the C:\Windows\System32\WindowsPowerShell\v1.0\Modules directory.
3. Type Connect-NaController <IP Address of Controller> -cred root and hit enter. You will be prompted to enter the password for root. Please enter it and hit enter.
4. When connected to the controller, at the prompt, type New-NaVirtualDisk d:\Master_WC.vhd (4GB) hit enter. This will create a new, thin provisioned, fixed size, 4GB virtual disk at d:\Master_WC.vhd.

Cloning Master Virtual Disk to Virtual Machine Shell Subfolders

When the PowerShell Scripts for creating and modifying the Virtual Machines have completed (see section 6.9.1 above), it will be time to assign each virtual machine "shell" its own write cache virtual disk. This is completed using a quick PowerShell script that utilizes the NetApp DataOnTap PowerShell Toolkit. The script will use NetApp FlexClone technology to create identical virtual disks in each Virtual Machine's subfolder.

The following figure shows a PowerShell Window capture of the commands to FlexClone a Master Write Cache Virtual Disk.

Figure 32 PowerShell Window—Cloning Virtual Disks

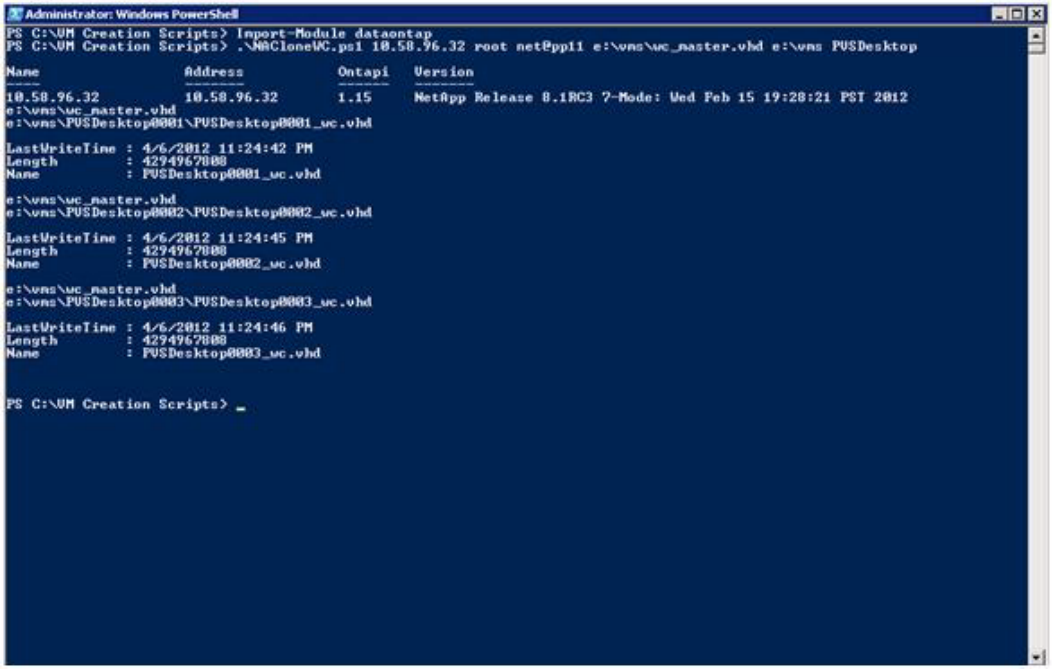

1. Open an elevated Windows PowerShell Window as a Local or Domain Administrator.
2. Type import-module dataontap at the prompt and hit enter. Please ensure that you have the most recent version of the NetApp PowerShell Toolkit installed in the C:\Windows\System32\WindowsPowerShell\v1.0\Modules directory.
3. Referencing the PowerShell Script NACloneWC.ps1 . (See Appendix C)
4. At the prompt, type NACloneWC.ps1 <IP Address of Controller> <User Name> <Password> <Source VHD File> <Destination Base> <VM Base Name> then hit enter. This will FlexClone the Master Write Cache Virtual Disk to each Virtual Machine's sub-folder.

Update the MAC Address on the Virtual Desktops

In this step, run the following powershell script: ModifyVMs.ps1

This script will switch the Virtual Desktops MAC addresses from Dynamic to Static (on the PVS legacy network). This is a requirement to stream the virtual desktop using the PVS server. The script will prompt you to enter parameters manually. Type y.

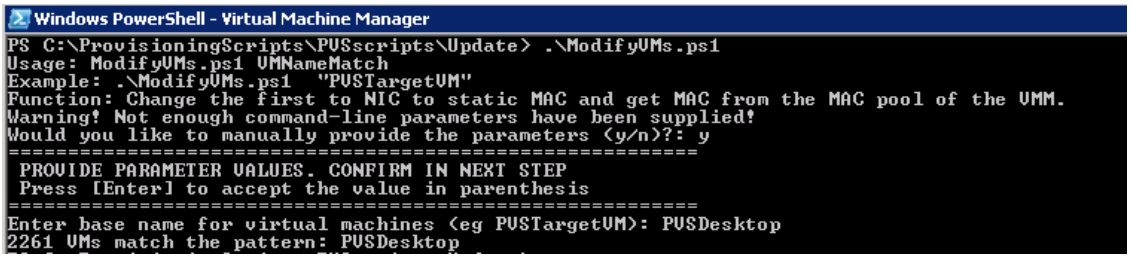

This script is fairly quick and should complete in about 20 minutes while cloning the VHD's is in progress.

Attach Write-Cache Disks to Virtual Desktops

In this step, run the following powershell script: AttachVHD.ps1

This script will attach the VHD's created in section Cloning Master Virtual Disk to Virtual Machine Shell Subfolders using NetApp cloning technology to the shell virtual desktops created in section Create Virtual Machine Shell Devices.

In this case the script requires that you enter the parameter directly (no manual parameter option).

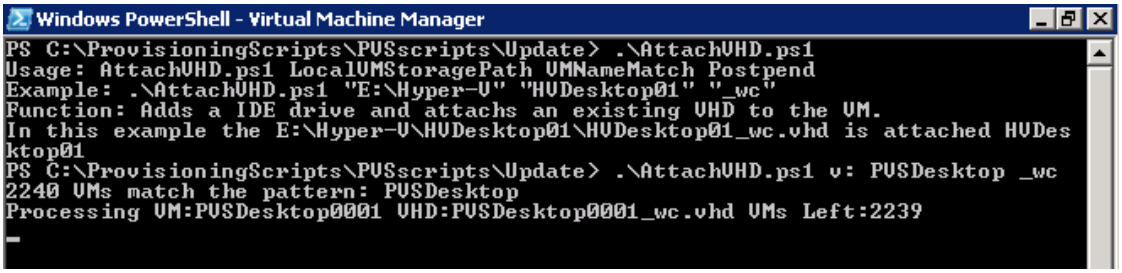

While this step is running, you can proceed to next step and execute in parallel. This step is expected to complete in about 20 to 30 minutes.

Export Virtual Machines for Import into Provisioning Server and Desktop Controllers

In this step, run the following powershell script: GenPVSFile.ps1

This script exports the virtual desktop properties from VMM into a CSV file format that can be imported by the PVS server.

The script requires that you enter the parameter directly (no manual parameter option).

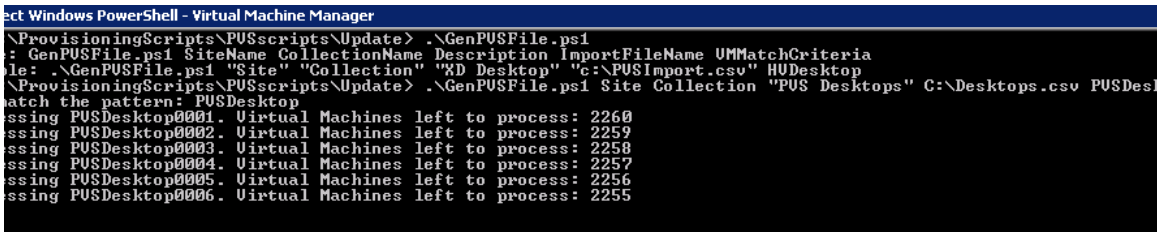

This script is expected to complete in about 3 minutes for the 2000 Desktops. The file produced by the script, Desktops.csv in the example above, will be used to import the machines into the Provisioning Server 6.1 Device Collection.

Desktop Delivery Infrastructure and Golden Image Creation

Many components are required to support the Virtual Desktop Infrastructure used in this study. This section details the key servers and their roles in the solution. Ultimately, Provisioning Server 6.1 in combination with XenDesktop 5.6 managed the VDI environment in this project.

The next section includes:

- An overview of the component servers in the solution
- Virtual Microsoft Server 2008 R2 SP1 Machines for Roaming Profiles
- Citrix Provisioning Server 6.1 Configuration
- Importing VMs created in Section 6.9 into Provisioning Server 6.1 and creating AD machine accounts
- Creating the Windows 7 SP1 Golden Image and converting it to a Provisioning Services vDisk
- Citrix XenDesktop 5.6 Desktop Delivery Controller configuration with the Desktop Studio

Overview of Solution Components

Figure 33 provides a logical overview of the solution components in the environment.

Figure 33 Citrix XenDesktop 5.6 and Provisioning Server 6.1 Logical Diagram

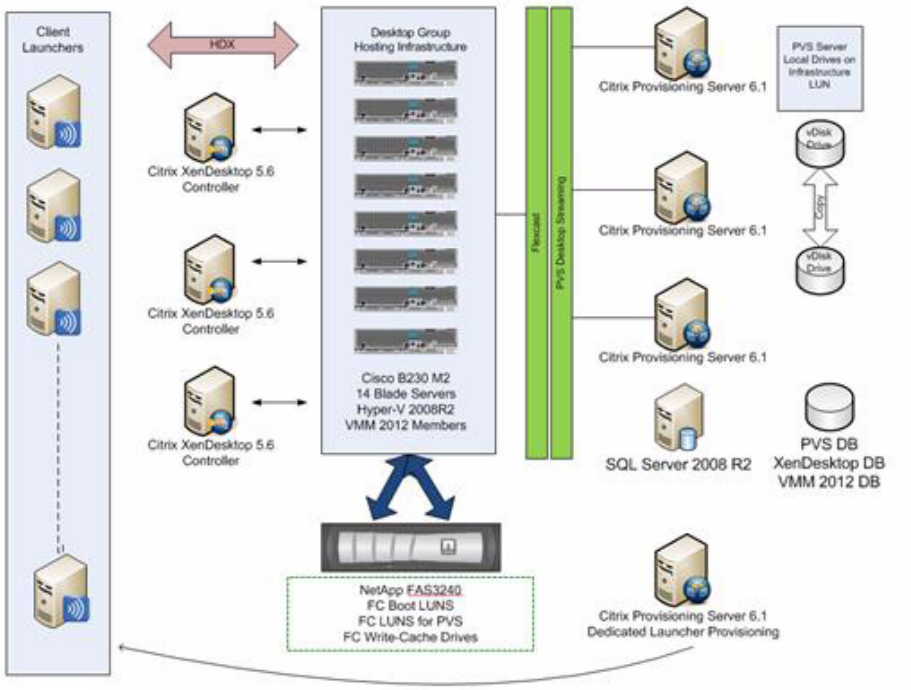

Summary of the Environment:

- 2 XenDesktop 5.6 Delivery Controllers
- 3 Provisioning Server 6.1 Server for Virtual Desktops
- 1 Provisioning Server 6.1 Server for Launchers
- 2 System Center 2012 Virtual Machine Managers for Windows 7 VMs
- 4 Launcher blades in a separate chassis for Login VSI client launchers.
- 2000 Virtual Desktops on 14 VDI blades
- 1 Citrix Licensing Server on XenDesktop Controller 1
- 1 Web Interface Server on XenDesktop Controller 1
- 2 Profile Servers for Microsoft Roaming Profiles, using a Clustered Shared volume
- 2 Microsoft SQL Server 2008 R2 for Provisioning Services, XenDesktop, and VMM using SQL Clustering and a Clustered Shared Volume
- 1 NetApp FAS 3240 dual-controller system with Flash Cache, 14 FC Write Cache LUNs, 16 Fibre Channel Boot LUNS, 1 shared infrastructure Fibre Channel LUN and two iSCSI LUNS.
- 95 Login VSI Launchers

The following tables provide details on the configuration of the solution components.

Hyper-V Server 2008 R2 SP1 VDI Hosts (14)			
Hardware:	Cisco B-Series blade servers	**Model:**	B230 –M2
OS:	Hyper-V Server 2008 R2	**Service Pack:**	SP1
CPU:	2 x 10 Core Intel Westmere EX E7-2870, 2.4 GHz (40 Logical Cores Total)	**RAM:**	256 GB
Disk:	Boot From SAN	**Network:**	4x 10GbE
. Updated M81KR (Palo) NIC driver to enic_driver_2.1.2.22-564611 . Updated M81KR (Palo) HBA driver to fnic_driver_1.5.0.7-563432			

Server 2008 R2 SP1 Enterprise Infrastructure Hosts (2)			
Hardware:	Cisco B-Series blade servers	**Model:**	B200-M2
OS:	Server 2008 R2 Enterprise Edition	**Service Pack:**	SP1
CPU:	2 x 6 Core Intel E5649, 2.53 GHz (24 Logical Cores Total)	**RAM:**	96 GB
Disk:	Boot From SAN	**Network:**	1 x 10 GbE
. Updated. **. Updated**	M81KR (Palo) NIC driver to enic_driver_2.1 M81KR (Palo) HBA driver to fnic_driver_1.5	.2.22-564611 .0.7-563432	

Citrix Provisioning Server 6.1 (3)			
OS:	Windows 2008 R2 Enterprise 64bit	**Service Pack:**	SP1
CPU:	4 x vCPU	**RAM:**	8192MB
Disk:	1 x100GB Virtual Disk (hosted on Infrastructure FC LUN)	**Network:**	1 x 10 GbE
• **Database for PVS hosted on separate Microsoft SQL Server 2008 R2 SP1 64bit**			

Citrix XenDesktop 5.6 DDCs (2)			
OS:	Windows 2008 R2 Enterprise 64bit	**Service Pack:**	SP1
CPU:	4 x vCPU Xeon 5690 3.46 GHz	**RAM:**	8192MB
Disk:	1 x100GB Virtual Disk (hosted on Infrastructure FC LUN)	**Network:**	1 x 10 GbE
• **Database for DDC hosted on separate Microsoft SQL Server 2008 R2 SP1 64bit** • **XDC01 Hosts the Web Interface and License Server for XenDesktop and PVS**			

Microsoft Virtual Machine Manager 2012 (2)			
OS:	Windows 2008 R2 Enterprise 64bit	**Service Pack:**	SP1
CPU:	2 x vCPU	**RAM:**	8192MB
Disk:	1 x100GB Virtual Disk (hosted on Infrastructure FC LUN)	**Network:**	1 x 10 GbE
• **Database for VMM hosted on separate Microsoft SQL Server 2008 R2 SP1 64bit**			

Microsoft SQL Server 2008 R2 (2)			
OS:	Windows 2008 R2 Enterprise 64bit	**Service Pack:**	SP1
CPU:	4 x vCPU	**RAM:**	8192MB
Disk:	1 x100GB Virtual Disk (hosted on Infrastructure FC LUN) 1x1000GB Clustered Shared Volume (SQL Databases)	**Network:**	1 x 10 GbE

The other dedicated Infrastructure Virtual Machines, all running Server 2008 R2 SP1:

- Two Active Directory Servers (Directory Services, DNS, and DHCP)

- Stand-alone File Server (collects Login VSI data on local drive)

Microsoft Server 2008 R2 SP1 Profile Servers

This section explains the installation and configuration of the profile cluster. The installation did the following:

- Used the Windows iSCSI initiator to connect to the iSCSI targets created on the NetApp
- Clustered the two virtual machines (profile 01/02)
- Create a highly available file share

Setting up iSCSI LUNs

This section details how to set the iSCSI LUNs using the windows iSCSI initiator. The steps below need to be completed on both Profile servers: profile01 and profile02. Before proceeding with this step, make sure that the target LUNs have been created on the NetApp storage.

1. Considering the LUNs are dual path, it is important to start by installing MPIO feature within the server:
   a. Open the Server Manager.
   b. Browse to the Feature node.
   c. Click Add Feature.
   d. Check Multipath I/O.
   e. Click Install.
2. When the Multipath I/O is installed, proceed to the Start Menu >Administrative Tools>iSCSI Initiator.
3. Under Target type the IP addresses for the NetApp Target. In our case that would be: 10.16.0.14 and 10.17.0.14

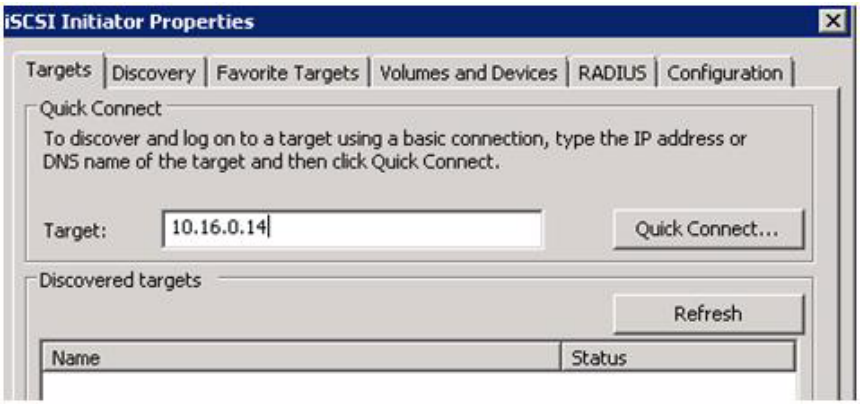

4. Click on Quick Connect…
5. When this is completed use the following Names to complete the setup on the NetApp storage.

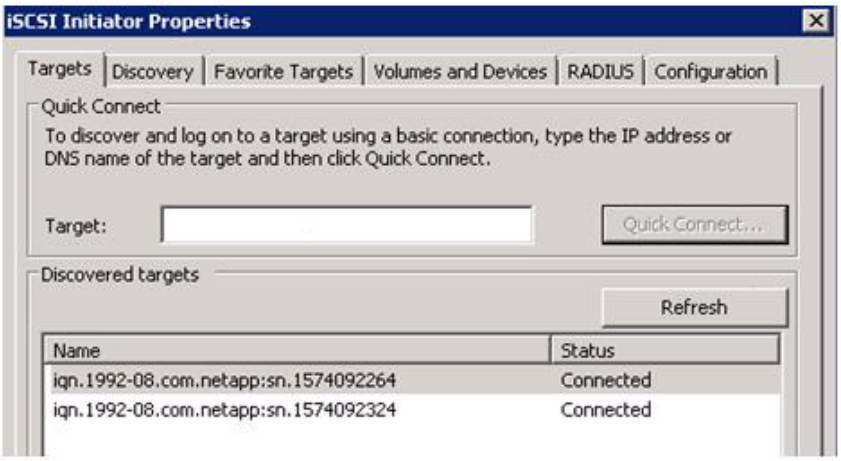

The names will be used to connect the LUNs to the hosts. When that is completed, rescan the disks under disk management.

6. Online, Initialize and format the two disks.

Setting up a Two-Node Profile Server Cluster

The two LUNs created in the section above, will be used for Quorum and file server clustered disk. Please follow the steps below to complete the clustering setup.

1. Install Failover Cluster feature with both servers:
   a. Open the Server manager.
   b. Browse to the feature node.
   c. Click Add Feature.
   d. Check Failover Clustering.
   e. Click Install.
2. When this is completed, proceed to the Start Menu>Administrative Tools>Failover Cluster Manager
3. Click Validate a Configuration and follow the on screen instructions to validate the two nodes: Profile01 and Profile02. When that succeeds, proceed to the next step.
4. Click Create a Cluster and follow the on screen instructions to create a two node cluster.

Note When the setup completes, one LUN should be used as a Quorum Disk and the other should be available for the cluster.

Setting up a Highly Available Profile Share

To set up a highly available profile share, do the following:

1. Install File Services on both servers.
   a. Open the Server Manager.
   b. Browse to the Role node.
   c. Click Add Roles.

d. Check File Services.

e. Click Next.

f. Check File Server.

g. Click Next then Install.

When that is completed on both Nodes, proceed to Failover Cluster Manager.

2. Right-click Services and applications node and click Configure a Service or Application…
3. Select File Server and click Next.
4. Input a File Server name and IP address; then click Next.
5. Select the cluster disk to be used then click Next.
6. When that is complete, click Add a shared folder in the Actions pane.
7. Click Browse and set the profile folder intended for the profiles; then click Next.
8. Leave the NTFS permission and click Next.
9. Validate SMB is checked and input Share name and then click Next.
10. Accept defaults and click Next.
11. Check Users and groups have custom share permission.
12. Click Permissions and set permissions to Everyone at Full Control; then click OK.
13. Click Next.
14. Accept Defaults then click Next.
15. Review summary and click Create.

Now you can use the Share name created in step 10 above to create highly available roaming user profiles.

Citrix Provisioning Services

Citrix Provisioning Server (PVS) is part of the XenDesktop Enterprise and Platinum suites and was used in all tested scenarios. Provisioning provides the ability to create and manage 1000's of virtual machines hosted on hypervisor servers with identical virtual machine configurations. It allows those virtual machines to PXE boot from a single golden Windows 7 Image.

Storage Configuration for Hyper-V Server Hosting PVS and Virtual Desktop vDisks

The test environment utilized a single NetApp FAS3240 storage system to provide both boot-from-SAN LUNS for the B230 M2 blades hosting the Windows 7 SP1 virtual machines and LUNs for volumes that were used for:

- Infrastructure virtual machines hard drives (One common Infrastructure LUN)
- Write-cache drives provisioned with each Windows 7 SP1 virtual machine (One LUN per VDI host)
- Windows 7 SP1 vDisks storage accessed by the Provisioning Servers (One common Infrastructure LUN)

The Launcher vDisks were stored on the Launcher blades' local storage devices.

Citrix Provisioning Server (PVS) 6.1 for use with Standard Mode Desktops

The Windows 7 SP1 desktop image is converted into a vDisk (.vhd) image. The vDisk is then locked in a Shared (Read-only) mode and hosted on the PVS server's local disk or on a shared file location.

- PVS can be used to create the desired number of virtual machines based on parameters specified using built-in setup wizard. In this study, those wizards were not employed.
- PVS was used to create machine accounts in Active Directory after the VMs were created by Virtual Machine Manager 2012.
- Virtual desktops are then configured to PXE boot on Hypervisor server.
- PVS streams the vDisk image on start up to the Hypervisor and is loaded into RAM.
- PVS injects a Security Identifier (SID) and host name associated with the virtual machine as each desktop boots to maintain uniqueness in AD. These object mappings are maintained and managed within the PVS server and are visible in the PVS Console under "Collections" view.
- Each virtual desktop is assigned a "Write Cache" (temporary file) where any delta changes (writes) to the default image are recorded and is used by the virtual windows operating system throughout its working life cycle. The Write Cache is written to a dedicated 4GB hard drive created by NetApp thin provisioning and attached to each new virtual desktop using PowerShell scripts.

For best performance, a copy of the vDisk was hosted and maintained on each PVS server's local drive to provide high availability and load balancing by all servers within the farm. The drive assigned by the hypervisor to each PVS server for vDisk storage was on a dedicated NFS mount. We assigned PVS servers 8GB RAM to facilitate the image(s) remaining persistent and be serviced by RAM after it is initially read for the first time by each server.

Two PVS servers were configured in a farm to provide high availability and resilience for virtual desktop provisioning. Connections are automatically failed over to a working server/s within the farm in the event of a failure without interruption to the desktop.

A separate PVS server with local storage was used to provision Login VSI Launcher machines for workload testing. We used 14 FC LUNS on the dual-controller NetApp FAS3240 to create and store each virtual machine's Write-Cache drive. Each controller hosted seven FC LUNs. Each LUN was zoned to an individual Hyper-V host.

It is important to consider where the Write Cache is placed when scaling virtual desktops using PVS server. There are several options as to where the Write Cache can be placed:

- PVS Server
- Hypervisor RAM
- Device Local Disk (an additional Virtual Disk for VDI instances)

For optimal performance and scalability the Cache on device HD option is used. A 4GB virtual disk is assigned to the virtual machine templates used in the clone creation process.

The PVS Target device agent installed in the Windows 7 gold image automatically places the Windows swap file on the same drive used by the PVS Write Cache when this mode is enabled.

The figure below illustrates multiple virtual machine instances hosted on a hypervisor server booting from a PVS single master image, each one has a virtual disk hosted on different NetApp FAS3240-provided LUNs associated with the VM's hypervisor host where the PVS cache is placed. This ensures that all write IO takes place on the NetApp FAS System filer storage over FC using high performance storage.

Figure 34 vDisk hosting on FC volumes

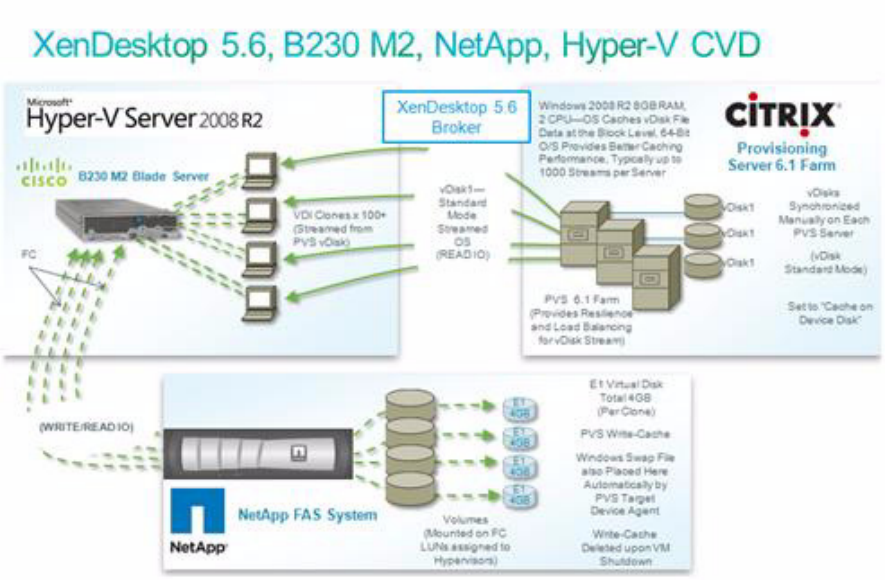

Import Virtual Machines into PVS Collection and Create AD Machine Accounts

In a previous section we created a comma separated file, named desktops.csv which contained the properties of the 2000 virtual desktop machine shells with 4GB write-cache drives attached. In this section, we will use that file to import the machines into our PVS device collection.

Import Virtual Machines into PVS Collection

To import virtual machines into a PVS collection, do the following:

1. Right-click the Site container and choose Import Devices.

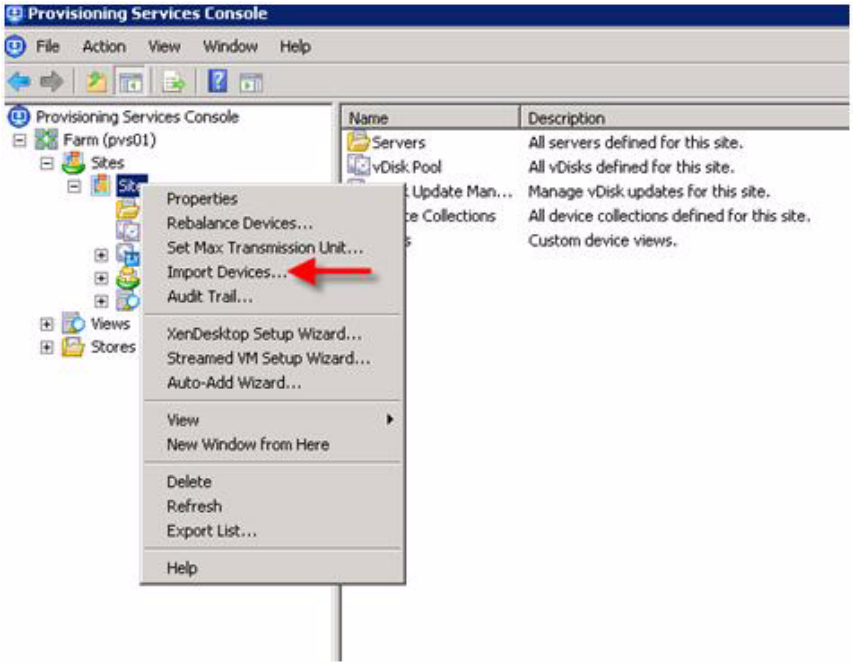

2. Click Next on the welcome page, browse to the location of the desktops.csv, then Click Next.

3. On the Import Target Device Options, Click Next.
4. On the New Target Device Selections, with all machines selected, Click Finish.

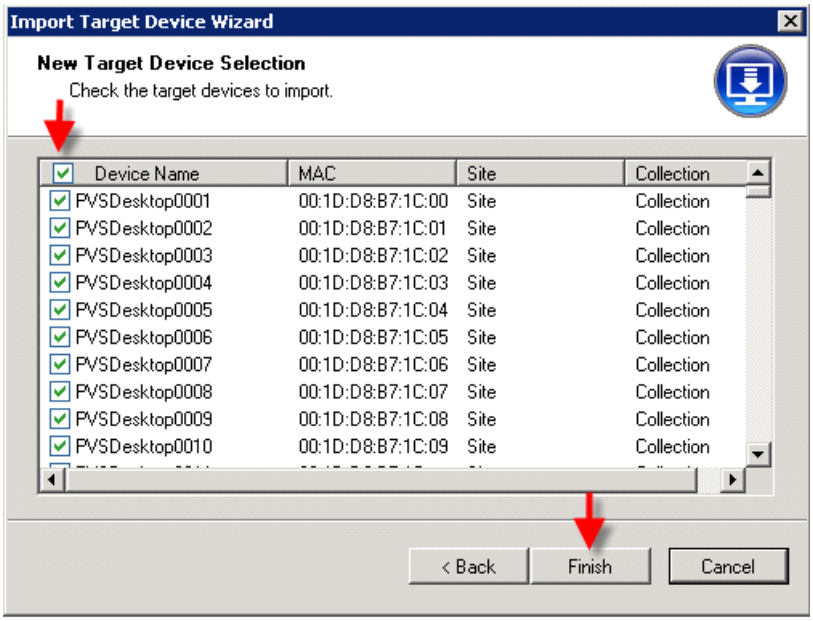

5. Expand your Collection and verify that all of the machines appear.

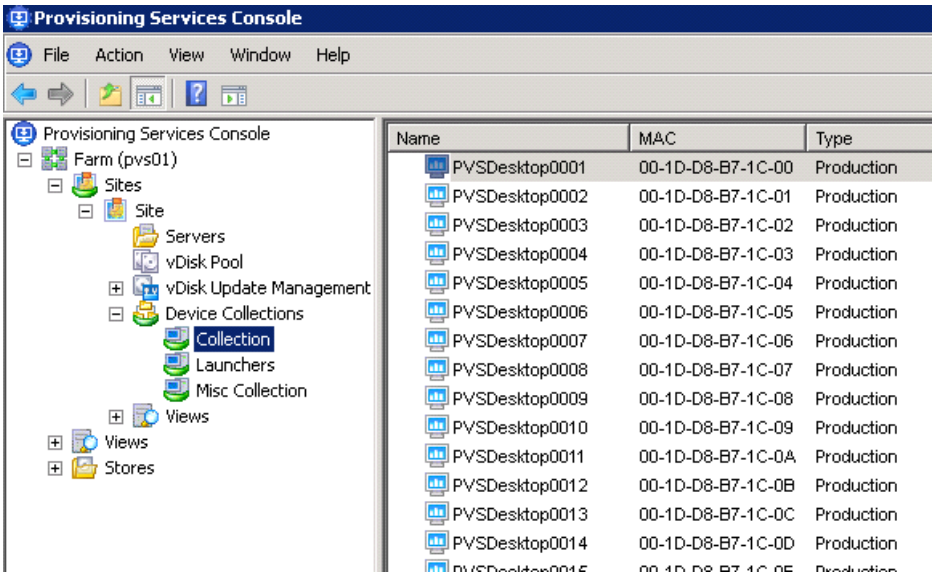

Create Active Directory Machine Accounts

When the devices are successfully added to your Collection in the PVS console, you must create Active Directory machine accounts in your virtual machine organizational unit. This section outlines how to accomplish that using the Provisioning Services Console on PVS01.

1. From the Collection container, select all of your virtual machines, then right-click and select Active Directory.

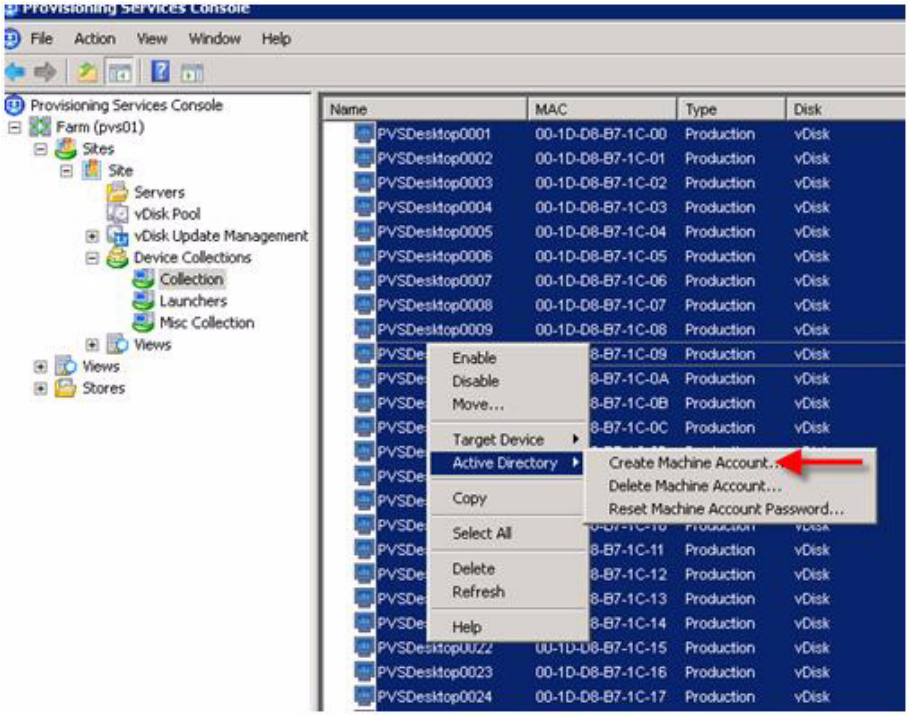

2. Select the Domain, Organization unit, then click Create Account.

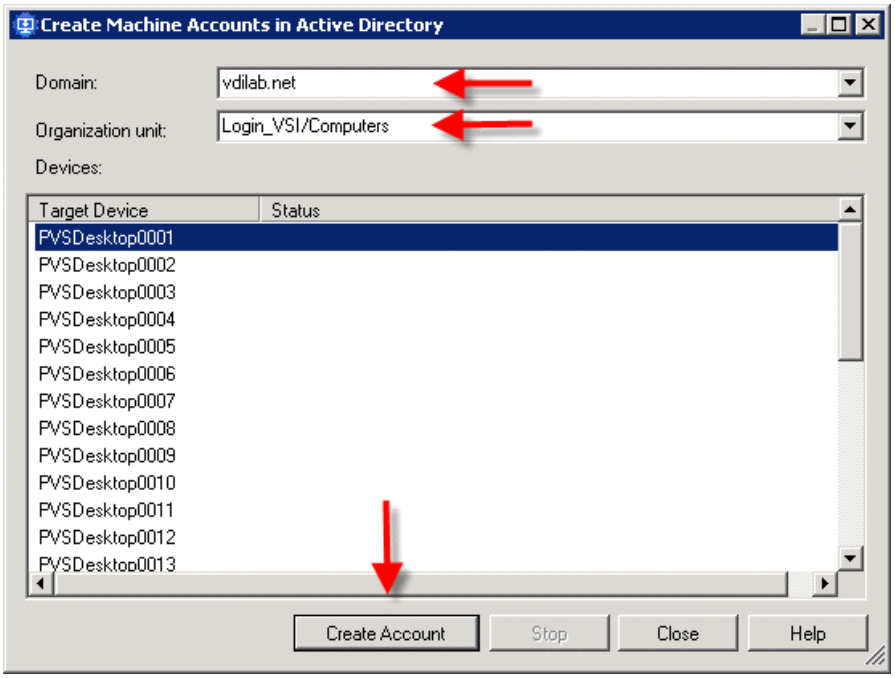

3. The Create Machine Account wizard should list Success next to each machine on the list.

Microsoft Windows 7 SP1 Image Creation

The Microsoft Windows 7 SP1 master or golden image with additional software was initially installed and prepared as a standard virtual machine on Microsoft Hyper-V 2008 R2 prior to being converted into a separate Citrix Provisioning Server vDisk file. The vDisk is used in conjunction with Provisioning Server 6.1 and the XenDesktop 5.6 controller to create more than 2000 new desktop virtual machines on the Hyper-V hosts.

With XenDesktop 5.6 and Provisioning Server 6.1, the XenDesktop Setup Wizard was not utilized. System Center Virtual Machine Manager 2012 power shell scripting was used to create the virtual desktop virtual machines. The properties of those machines were then exported and imported into the appropriate Provisioning Server collection.

Each virtual desktop virtual machine was created with a 3.0 GB write cache disk. More information as to why the additional virtual disks are needed can be found in section 5.1 Configuration Topology for Scalability of Citrix XenDesktops on Cisco UCS and NetApp FAS Storage.

The following section describes the process used to create the master or golden image and centralized Windows 7 vDisk used by Provisioning Services.

Citrix XenDesktop 5.6 Desktop Delivery Controller

Two XenDesktop 5.6 Desktop Delivery Controllers (DDCs) were virtualized on Server 2008 R2 SP1 with the Hyper-V role installed running on Cisco B200 M2 infrastructure blades.

Beginning with XenDesktop 5, Citrix replaced the proprietary IMA protocol encrypted data store in favor of Microsoft SQL Server databases. Concurrently the concept of XenDesktop Farms (used in XenDesktop 4 and earlier) was eliminated in favor of the concept of Sites.

There is no master DDC. All DDCs communicate directly with the SQL database continuously to update status and usage, providing seamless fault tolerance. Citrix recommends at least two DDCs per site.

For this study, we created and utilized two DDCs. There were no public hot fixes or service packs for XenDesktop 5.6 posted during the study.

From a management standpoint, Citrix introduced two new management consoles beginning with XenDesktop 5.

- Desktop Studio
- Desktop Director

The Desktop Studio is the main administration console where hosts, machine catalogs, desktop groups and applications are created and managed. The Desktop Studio is where HDX policy is configured and applied to the site. The Desktop Studio is a Microsoft Management Console snapin and fully supports PowerShell.

Figure 35 XenDesktop 5.6 Desktop Studio

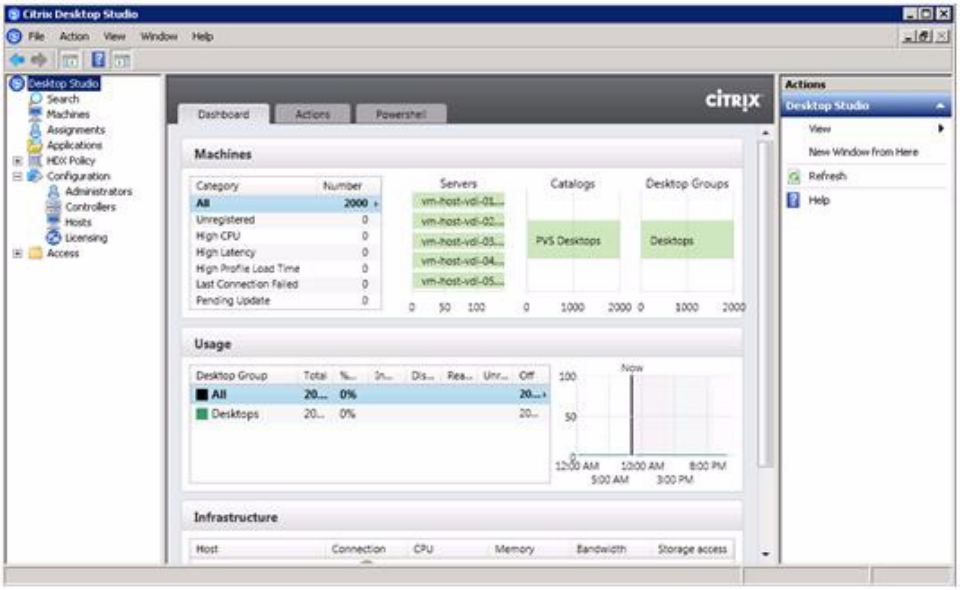

The Desktop Director is designed for use by level 1 and level 2 help desk personnel. It provides real-time status and limited management capability for the running XenDesktop infrastructure. Help Desk personnel can provide real-time status to end users, shadow their sessions, and restart their desktops. Desktop Director utilizes Internet Explorer and flash to present data.

Figure 36 XenDesktop 5.6 Desktop Director

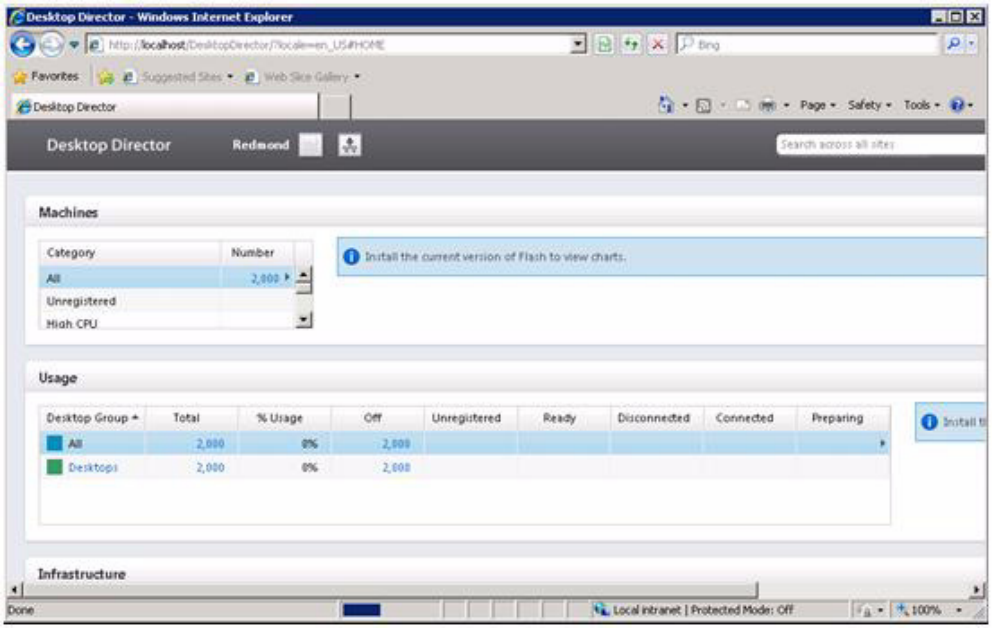

Site Configuration

In addition to the standard XenDesktop site installation, the following additional items were configured one Machine Catalog, named PVS Desktops, to hold all virtual desktops in the test environment.

Figure 37 XenDesktop Catalog

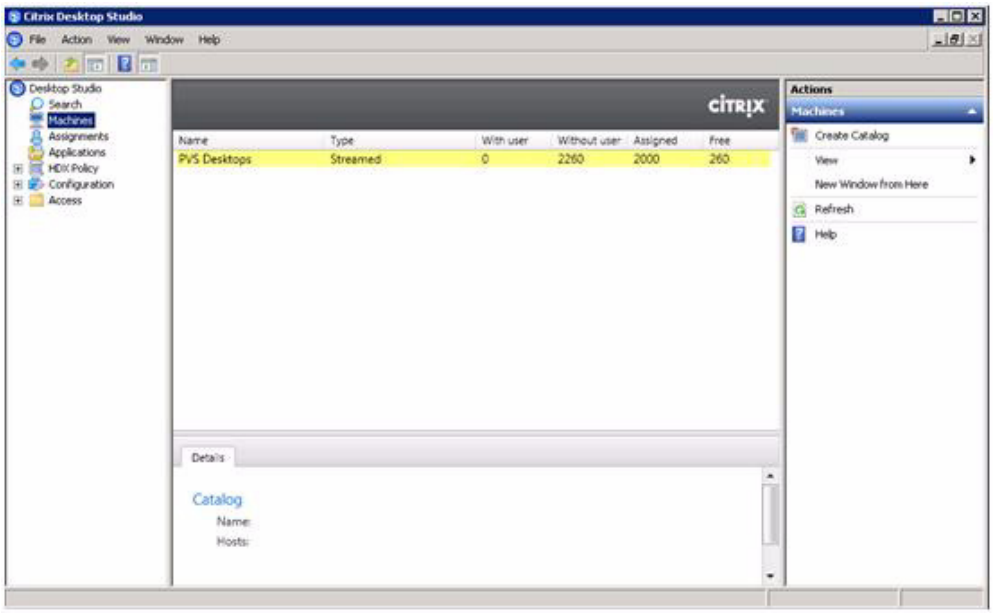

One Desktop Group, named Desktops containing 2000 of the desktops from the catalog.

Figure 38 Desktop Group

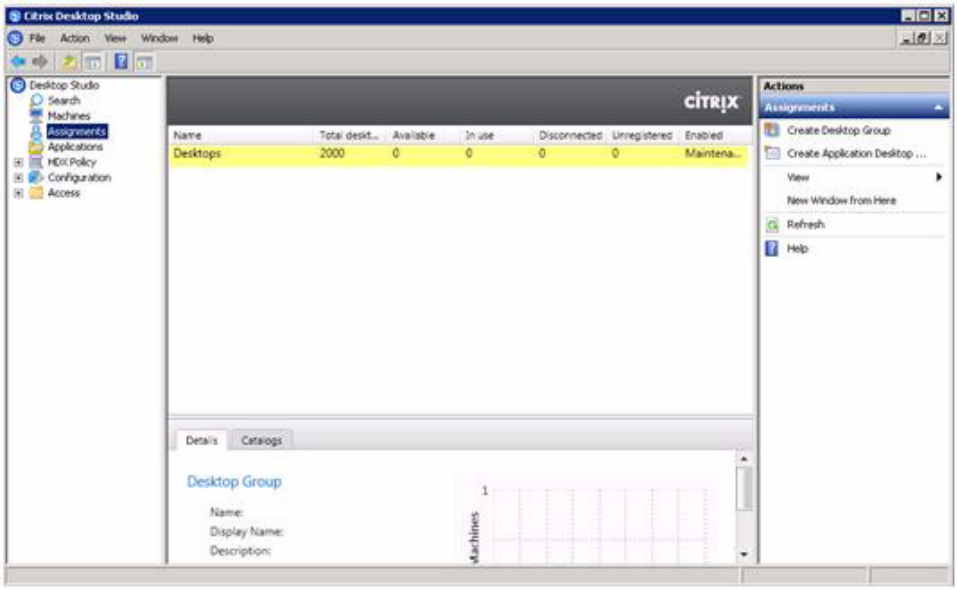

Created XenDesktop HDX user policies to disable auto connect client drives and client printer redirection.

Figure 39 XenDesktop HDX User Policy

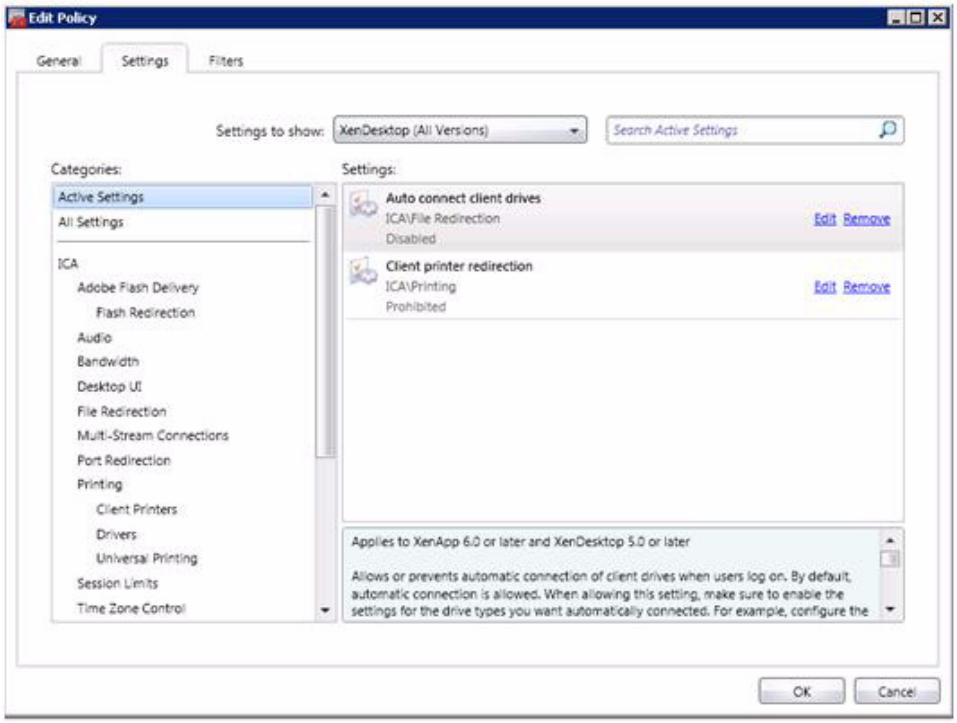

Configured the Power Settings for the Desktop Group to start the desired number of virtual desktops when the group was removed from Maintenance Mode in Assignments node of the Desktop Studio. Set action on disconnect to No action.

Figure 40 XenDesktop Power Settings

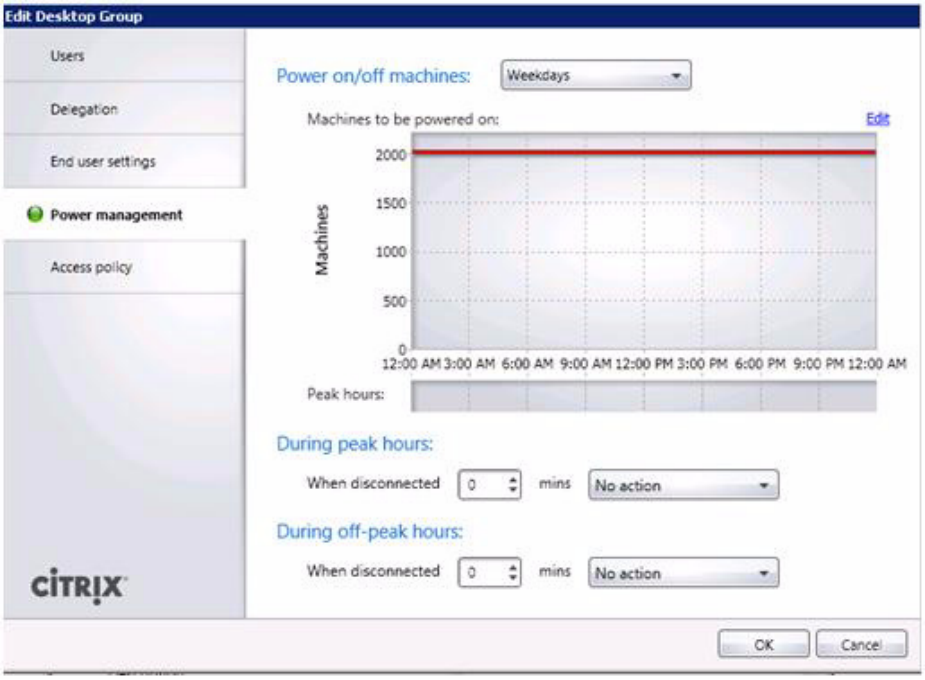

Configured the Max active actions and Max new actions per minute from the default 10 value to 75 for the multi-blade test in the Advanced Host Details section of the Configuration/Hosts node of the Desktop Studio.

Figure 41 XenDesktop Maximum Active and New Actions

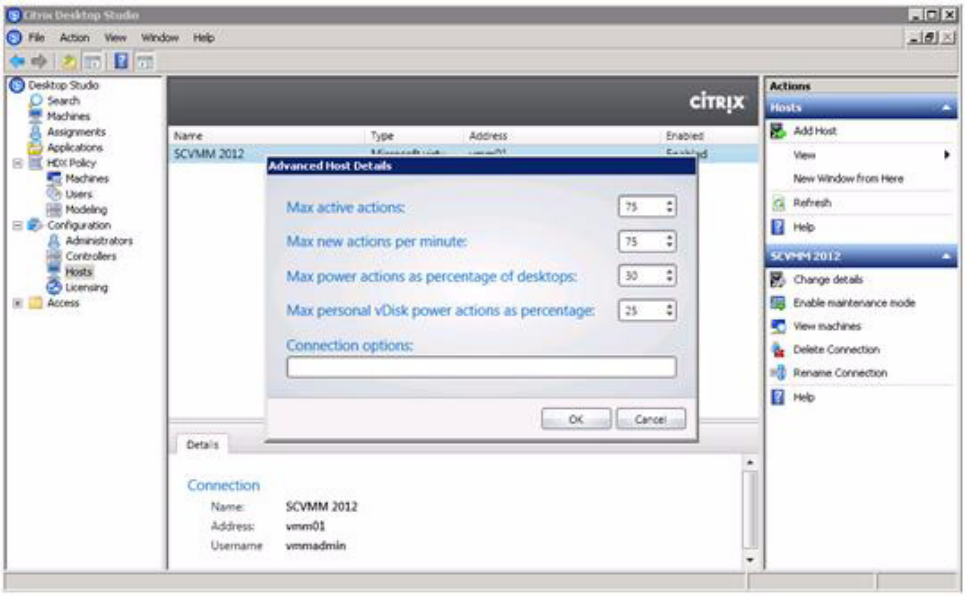

When the environment was fully configured, the Desktop Controller was used during testing to boot, monitor desktop availability/usage status, and shut down the virtual desktops.

Test Setup and Configurations

In this project, we tested a single Cisco UCS B230 M2 blade in a single chassis and fourteen B230 M2 blades in two chassis to illustrate linear scalability.

Cisco UCS Test Configuration for Single Blade Scalability

Figure 42 Cisco UCS B230 M2 Blade Server for Single Server Scalability

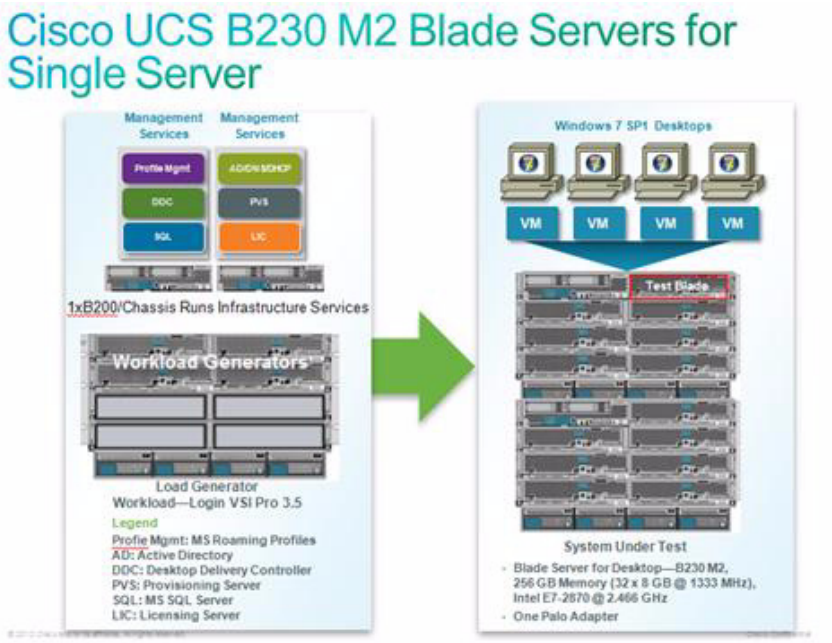

Hardware Components

- 1 X Cisco UCS B230-M2 (E7-2870 @ 2.4 GHz) blade server with 256GB of memory (8 GB X 32 DIMMS @ 1333 MHz) Windows 7 SP1 Virtual Desktops
- 2 X Cisco UCS B200-M2 (5690 @ 3.466 GHz) blade servers with 96 GB of memory (8 GB X 12 DIMMS @ 1333 MHz) Infrastructure Servers
- 4 X Cisco UCS B230-M2 (5690 @ 3.466 GHz) blade servers with 256 GB of memory (8 GB X 32 DIMMS @ 1333 MHz) Load Generators
- 1 X M81KR (Palo) Converged Network Adapter/Blade
- 2 X Cisco Fabric Interconnect 6248s
- 2 X Cisco Nexus 5548 Access Switches
- 1 X NetApp FAS3240 System storage array, two controllers, 2 x dual port 8GB FC cards, 2 x dual port 10 GbE cards with four DS2246 shelves, each containing 24 600GB 10,000rpm SAS drives

Software Components

- Cisco UCS firmware 2.0(1w)
- Hyper-V Server 2008 R2 SP1 for VDI Hosts

- Server 2008 R2 SP1 with Hyper-V Role for Infrastructure Hosts
- System Center 2012 Virtual Machine Manager
- XenDesktop 5.6
- Provisioning Server 6.1
- Windows 7 SP1 32 bit, 1vCPU, 1.5 GB of memory, 20 GB/VM, Office 2010

Cisco UCS Configuration for Two Chassis—Fourteen Blade Test

Figure 43 Two chassis test configuration-14 x B250 Blade Server

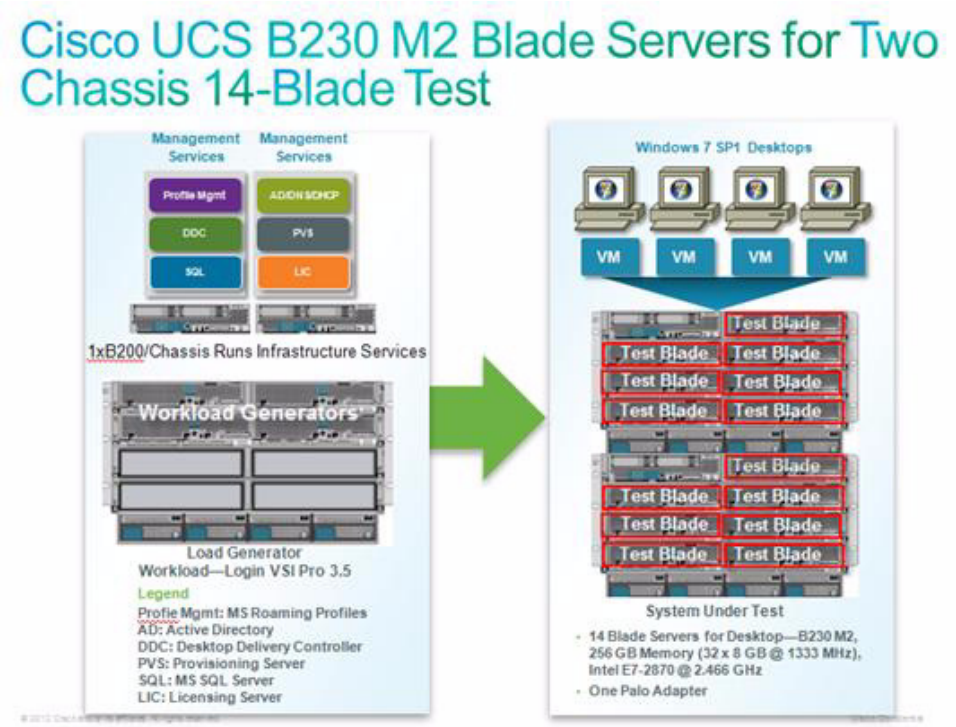

Hardware Components

- 1 X Cisco UCS B230-M2 (E7-2870 @ 2.4 GHz) blade server with 256GB of memory (8 GB X 32 DIMMS @ 1333 MHz) Windows 7 SP1 Virtual Desktops
- 2 X Cisco UCS B200-M2 (5690 @ 3.466 GHz) blade servers with 96 GB of memory (8 GB X 12 DIMMS @ 1333 MHz) Infrastructure Servers
- 4 X Cisco UCS B230-M2 (5690 @ 3.466 GHz) blade servers with 256 GB of memory (8 GB X 32 DIMMS @ 1333 MHz) Load Generators
- 1 X M81KR (Palo) Converged Network Adapter/Blade
- 2 X Cisco Fabric Interconnect 6248s
- 2 X Cisco Nexus 5548 Access Switches
- 1 X NetApp FAS3240 System storage array, two controllers, 2 x dual port 8GB FC cards, 2 x dual port 10 GbE cards with four DS2246 shelves, each containing 24 600GB 10,000rpm SAS drives

Software Components

- Cisco UCS firmware 2.0(1w)
- Hyper-V Server 2008 R2 SP1 for VDI Hosts
- Server 2008 R2 SP1 with Hyper-V Role for Infrastructure Hosts
- System Center 2012 Virtual Machine Manager
- XenDesktop 5.6
- Provisioning Server 6.1
- Windows 7 SP1 32 bit, 1vCPU, 1.5 GB of memory, 20 GB/VM, Office 2010

Testing Methodology and Success Criteria

All validation testing was conducted on-site within the Microsoft Engineering Center with joint support from both Citrix and NetApp resources.

The testing results focused on the entire process of the virtual desktop lifecycle by capturing metrics during the desktop boot-up, user logon and virtual desktop acquisition, also referred to as ramp-up, user workload execution, also referred to as steady state, and user logoff for the Hosted VDI model under test.

Test metrics were gathered from the hypervisor, virtual desktop, storage, and load generation software to assess the overall success of an individual test cycle. Each test cycle was not considered passing unless all of the planned test users completed the ramp-up and steady state phases described below and unless all metrics were within the permissible thresholds as noted as success criteria.

Three successfully completed test cycles were conducted for each hardware configuration and results were found to be relatively consistent from one test to the next.

Load Generation

Within each test environment, load generators were utilized to put demand on the system to simulate multiple users accessing the XenDesktop 5.6 environment and executing a typical end-user workflow. To generate load within the environment, an auxiliary software application was required to generate the end user connection to the XenDesktop environment, to provide unique user credentials, to initiate the workload, and to evaluate the end user experience.

In the hosted VDI test environment, sessions launchers were used simulate multiple users making a direct connection to XenDesktop 5.6 via a Citrix HDX protocol connection.

User Workload Simulation—LoginVSI From Login Consultants

One of the most critical factors of validating a XenDesktop deployment is identifying a real-world user workload that is easy for customers to replicate and standardized across platforms to allow customers to realistically test the impact of a variety of worker tasks. To accurately represent a real-world user workload, a third-party tool from Login Consultants was used throughout the Hosted VDI testing.

The tool has the benefit of taking measurements of the in-session response time, providing an objective way to measure the expected user experience for individual desktop throughout large scale testing, including login storms.

The Virtual Session Indexer (Login Consultants' Login VSI 3.5) methodology, designed for benchmarking Server Based Computing (SBC) and Virtual Desktop Infrastructure (VDI) environments is completely platform and protocol independent and hence allows customers to easily replicate the testing results in their environment. NOTE: In this testing, we utilized the tool to benchmark our VDI environment only.

Login VSI calculates an index based on the amount of simultaneous sessions that can be run on a single machine.

Login VSI simulates a medium workload user (also known as knowledge worker) running generic applications such as: Microsoft Office 2007 or 2010, Internet Explorer 8 including a Flash video applet and Adobe Acrobat Reader. (Note: For the purposes of this test, applications were installed locally, not streamed nor hosted on XenApp).

Like real users, the scripted Login VSI session will leave multiple applications open at the same time. The medium workload is the default workload in Login VSI and was used for this testing. This workload emulated a medium knowledge working using Office, IE, printing and PDF.

- When 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 160ms for each character.
- Approximately 2 minutes of idle time is included to simulate real-world users.

Each loop will open and use:

- Outlook 2007/2010, browse 10 messages.
- Internet Explorer, one instance is left open (BBC.co.uk), one instance is browsed to Wired.com, Lonelyplanet.com and heavy
- 480 p Flash application gettheglass.com.
- Word 2007/2010, 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.
- Excel 2007/2010, a very large randomized sheet is opened.
- PowerPoint 2007/2010, a presentation is reviewed and edited.
- 7-zip: using the command line version the output of the session is zipped.

Figure 44 Graphical Representation of Medium Workload

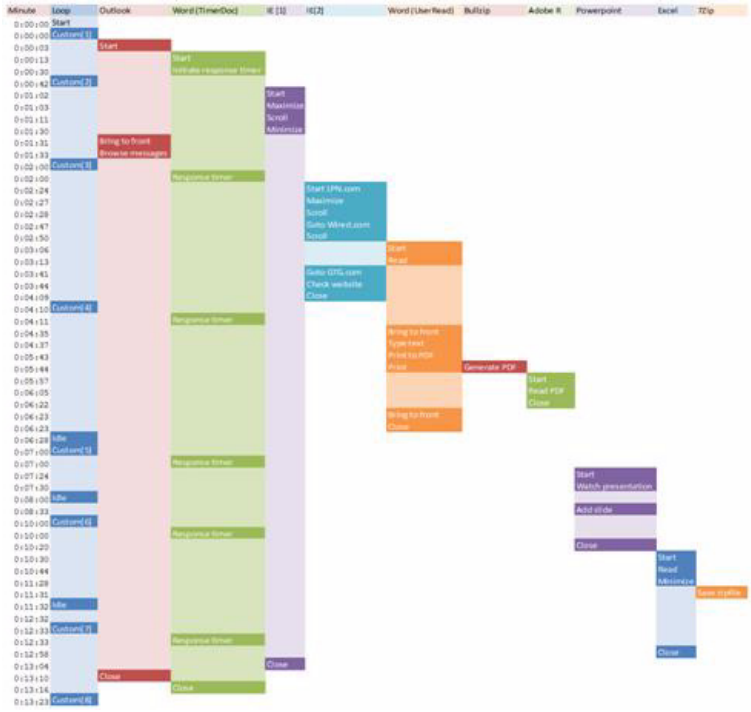

You can obtain additional information on Login VSI from http://www.loginvsi.com.

Testing Procedure

The following protocol was used for each test cycle in this study to insure consistent results.

Pre-Test Setup For Single and Multi-Blade Testing

All virtual machines were shut down utilizing the Citrix XenDesktop 5.6 Desktop Studio. All Launchers for the test were shut down. They were then restarted in groups of 10 each minute until the required number of launchers was running with the Login VSI Agent at a waiting for test to start state.

All Hyper-V Server 2008 R2 SP1 VDI host blades to be tested were restarted prior to each test cycle.

Test Run Protocol

For each of the three consecutive runs on single blade (145 User) and 14-blade (2000 User) tests, the following process was followed:

1. Time 0:00:00 Started PerfMon Logging on the following systems:

- VDI Host Blades used in test run
- VMM 2012 Server used in test run
- PVS Server(s) used in test run
- DDCs used in test run
- Profile Servers used in test run
- SQL Servers used in test run
- 3 Launcher virtual machines

2. Time 0:00:10 Started NetApp Performance Logging on Controller A
3. Time 0:00:15 Started NetApp Performance Logging on Controller B
4. Time 0:05 Take 145 or 2000 desktops out of maintenance mode on XenDesktop Studio
5. Time 0:06 First machines boot
6. Time 0:33 145 or 2000 desktops booted on 1 or 14 blades
7. Time 0:35 145 or 2000 desktops available on 1 or 14 blades
8. Time 0:50 Start Login VSI 3.5 Test with 145 or 2000 desktops utilizing 6 or 95 Launchers
9. Time 1:20 145 or 2000 desktops launched
10. Time 1:22 145 or 2000 desktops active
11. Time 1:35 Login VSI Test Ends
12. Time 1:50 145 or 2000 desktops logged off
13. Time 2:00 All logging terminated

Success Criteria

There were multiple metrics that were captured during each test run, but the success criteria for considering a single test run as pass or fail was based on the key metric, VSI Max. The Login VSI Max evaluates the user response time during increasing user load and assesses the successful start-to-finish execution of all the initiated virtual desktop sessions.

Login VSI Max

VSI Max represents the maximum number of users the environment can handle before serious performance degradation occurs. VSI Max is calculated based on the response times of individual users as indicated during the workload execution. The user response time has a threshold of 4000ms and all users response times are expected to be less than 4000ms in order to assume that the user interaction with the virtual desktop is at a functional level. VSI Max is reached when the response times reaches or exceeds 4000ms for 6 consecutive occurrences.

If VSI Max is reached, that indicates the point at which the user experience has significantly degraded. The response time is generally an indicator of the host CPU resources, but this specific method of analyzing the user experience provides an objective method of comparison that can be aligned to host CPU performance.

Note In the prior version of Login VSI, the threshold for response time was 2000ms. The workloads and the analysis have been upgraded in Login VSI 3 to make the testing more aligned to real-world use. In the medium workload in Login VSI 3.0 a CPU intensive 480p flash movie is incorporated in each test loop. In general, the redesigned workload would result in an approximate 20% decrease in the number of users passing the test versus Login VSI 2.0 on the same server and storage hardware.

Calculating VSIMax

Typically the desktop workload is scripted in a 12-14 minute loop when a simulated Login VSI user is logged on. After the loop is finished it will restart automatically. Within each loop the response times of seven specific operations is measured in a regular interval: six times in within each loop. The response times if these seven operations are used to establish VSImax.

The seven operations from which the response times are measured are:

- Copy new document from the document pool in the home drive
  - This operation will refresh a new document to be used for measuring the response time. This activity is mostly a file-system operation.
- Starting Microsoft Word with a document
  - This operation will measure the responsiveness of the Operating System and the file system. Microsoft Word is started and loaded into memory, also the new document is automatically loaded into Microsoft Word. When the disk I/O is extensive or even saturated, this will impact the file open dialogue considerably.
- Starting the File Open dialogue
  - This operation is handled for small part by Word and a large part by the operating system. The file open dialogue uses generic subsystems and interface components of the OS. The OS provides the contents of this dialogue.
- Starting Notepad
  - This operation is handled by the OS (loading and initiating notepad.exe) and by the Notepad.exe itself through execution. This operation seems instant from an end-user's point of view.
- Starting the Print dialogue
  - This operation is handled for a large part by the OS subsystems, as the print dialogue is provided by the OS. This dialogue loads the print-subsystem and the drivers of the selected printer. As a result, this dialogue is also dependent on disk performance.
- Starting the Search and Replace dialogue
  - This operation is handled within the application completely; the presentation of the dialogue is almost instant. Serious bottlenecks on application level will impact the speed of this dialogue.
- Compress the document into a zip file with 7-zip command line
  - This operation is handled by the command line version of 7-zip. The compression will very briefly spike CPU and disk I/O.

These measured operations with Login VSI do hit considerably different subsystems such as CPU (user and kernel), Memory, Disk, the OS in general, the application itself, print, GDI, etc. These operations are specifically short by nature. When such operations are consistently long: the system is saturated because of excessive queuing on any kind of resource. As a result, the average response times will then escalate. This effect is clearly visible to end-users. When such operations consistently consume multiple seconds the user will regard the system as slow and unresponsive.

With Login VSI 3.0 it is now possible to choose between VSImax Classic and VSImax Dynamic results analysis. For these tests, we utilized VSImax Classic analysis.

VSIMax Classic

VSImax Classic is based on the previous versions of Login VSI, and is achieved when the average Login VSI response time is higher than a fixed threshold of 4000ms. This method proves to be reliable when no anti-virus or application virtualization is used.

To calculate the response times the seven activities listed in the previous section are totaled. To balance these measurements are weighted before they are summed. Without weighting individual response times before they are totaled, one specific measurement (out of seven) could dominate the results.

Within VSImax Classic two measurements are weighted before they are added to the total VSImax response time:

1. Starting Microsoft Word with a document is divided by two (50%)

2. Starting the Search and Replace dialogue is multiplied by five (500%)

A sample of the VSImax Classic response time calculation is displayed in Table 9.

Table 9 VSImax Classic Response Time Calculation

Activity (RowName)	Result (ms)	Weight (%)	Weighted Result (ms)
Refresh document (RFS)	160	100%	160
Start Word with new doc (LOAD)	1400	50%	700
File Open Dialogue (OPEN)	350	100%	350
Start Notepad (NOTEPAD)	50	100%	50
Print Dialogue (PRINT)	220	100%	220
Replace Dialogue (FIND)	10	500%	50
Zip documents (ZIP)	130	100%	130
		VSImax Classic Response Time	**1660**

Then the average VSImax response time is calculated based on the amount of active Login VSI users logged on to the system. When the average VSImax response times are consistently higher than the default threshold of 4000ms, VSImax is achieved.

In practice however, tests have shown a substantial increase of application response time when antivirus and/or application virtualization is used. The baseline response time is typically around 1400 - 1800 ms without application virtualization or antivirus. However, when anti-virus or application virtualization is used, the baseline response time varies between 2500 - 3500 ms.

When the baseline response time is already so high the VSImax threshold of 4000ms is too easily reached. 'VSImax Classic' will report a maximum long before system resources like CPU, mem or disk are actually saturated.

It was decided further optimize VSImax calculation. Now in Login VSI 3.0 VSImax Dynamic is introduced to be able to support wildly varying baseline response times when anti-virus and/or application virtualization is used.

VSIMax Classic was not used in this study.

VSIMax Dynamic

Similar to 'VSImax Classic', VSImax Dynamic is calculated when the response times are consistently above a certain threshold. However, this threshold is now dynamically calculated on the baseline response time of the test.

- Individual measurements are weighted to better support this approach:
- Copy new doc from the document pool in the home drive: 100%
- Microsoft Word with a document: 33.3%
- Starting the "File Open" dialogue: 100%
- Starting "Notepad": 300%
- Starting the "Print" dialogue: 200%
- Starting the "Search and Replace" dialogue: 400%
- Compress the document into a zip file with 7-zip command line 200%

A sample of the VSImax Dynamic response time calculation is displayed in Table 10.

Table 10 VSImax Dynamic Response Time Calculation

Activity (RowName)	Result (ms)	Weight (%)	Weighted Result (ms)
Refresh document (RFS)	160	100%	160
Start Word with new doc (LOAD)	1400	33.3%	467
File Open Dialogue (OPEN)	350	100%	350
Start Notepad (NOTEPAD)	50	300%	150
Print Dialogue (PRINT)	220	200%	440
Replace Dialogue (FIND)	10	400%	40
Zip documents (ZIP)	130	200%	230
VSImax Dynamic Response Time			**1837**

Then the average VSImax response time is calculated based on the amount of active Login VSI users logged on to the system. For this the average VSImax response times need to consistently higher than a dynamically calculated threshold.

To determine this dynamic threshold, first the average baseline response time is calculated. This is done by averaging the baseline response time of the first 15 Login VSI users on the system.

The formula for the dynamic threshold is: Average Baseline Response Time x 125% + 3000. As a result, when the baseline response time is 1800, the VSImax threshold will now be 1800 x 125% + 3000 = 5250ms.

When the application virtualization is used, the baseline response time can wildly vary per vendor and streaming strategy. Therefore it is recommend to use VSImax Dynamic when comparisons are made with application virtualization or anti-virus agents. The result VSImax Dynamic scores are aligned again with saturation on a CPU, Memory or Disk level, also when the baseline response time are relatively high.

Determining VSIMax

The Login VSI analyzer will automatically identify the VSImax. In the example below the VSImax is 98. The analyzer will automatically determine stuck sessions and correct the final VSImax score.

- Vertical axis: Response Time in milliseconds
- Horizontal axis: Total Active Sessions

Figure 45 Sample Login VSI Analyzer Graphic Output

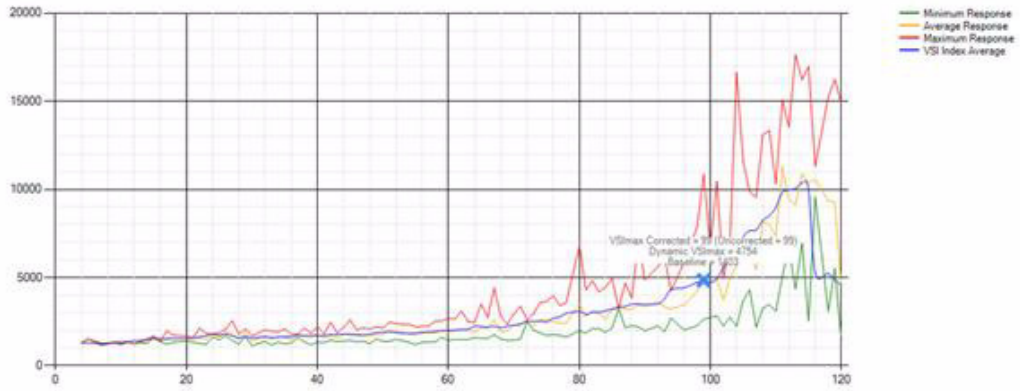

- Red line: Maximum Response (worst response time of an individual measurement within a single session)

- Orange line: Average Response Time within for each level of active sessions
- Blue line: the VSImax average.
- Green line: Minimum Response (best response time of an individual measurement within a single session)

In our tests, the total number of users in the test run had to login, become active and run at least one test loop and log out automatically without reaching the VSI Max to be considered a success.

Note We discovered a technical issue with the VSIMax dynamic calculation in our testing on Cisco B230 M2 blades where the VSIMax Dynamic was not reached during extreme conditions. Working with Login Consultants, we devised a methodology to validate the testing without reaching VSIMax Dynamic until such time as a new calculation is available.

Our Login VSI pass criteria, accepted by Login Consultants for this testing is as follows:

Cisco will run tests at a session count level that effectively utilizes the blade capacity measured by CPU utilization, memory utilization, storage utilization, and network utilization. We will use Login VSI to launch version 3.5 medium workloads. The number of launched sessions must equal active sessions within two minutes of the last session launched in a test.

The Citrix Desktop Studio will be monitored throughout the steady state to insure that:

- All running sessions report In Use throughout the steady state
- No sessions move to unregistered or available state at any time during steady state

Within 20 minutes of the end of the test, all sessions on all launchers must have logged out automatically and the Login VSI Agent must have shut down. We will publish our CVD with our recommendation following the process above and will note that we did not reach a VSIMax dynamic in our testing due to a technical issue with the analyzer formula that calculates VSIMax.

Test Results

The purpose of this testing is to provide the data needed to validate Citrix XenDesktop 5.6 Hosted VDI FlexCast model and Citrix Provisioning Services 6.1 using Microsoft Hyper-V Server 2008 R2 SP1 and System Center 2012 Virtual Machine Manager to virtualize Microsoft Windows 7 SP1 desktops on Cisco UCS B230 M2 blade servers using a dual-controller NetApp FAS3240 storage system.

The information contained in this section provides data points that a customer may reference in designing their own implementations. These validation results are an example of what is possible under the specific environment conditions outlined in this paper, and do not represent the full characterization of XenDesktop with Microsoft Hyper-V.

Two test sequences were performed to establish single blade performance and multi-blade, linear scalability across the FlexPod.

Single Cisco UCS B230-M2 Blade Server Validation - 145 Users

This section details the results from the XenDesktop Hosted VDI single blade server validation testing.

The primary success criteria metrics are provided to validate the overall success of the test cycle as an output chart from Login Consultants VSI Analyzer Professional Edition, VSIMax Dynamic for the Medium workload (with Flash.)

Additional graphs detailing the CPU and Memory utilization during peak session load are also presented. Given adequate storage capability, the limiting factor in the testing was CPU utilization.

The single server graphs shown in this section are representative of a single Hyper-V host in the larger environment for validation purposes, but it should be noted that these graphs are representative of the behavior for all servers in the respective environment.

We present performance information on key infrastructure virtual machines with the tested blade.

Figure 46 145 Desktop Sessions on Microsoft Hyper-V Server 2008 R2 SP1 Below 4000 ms

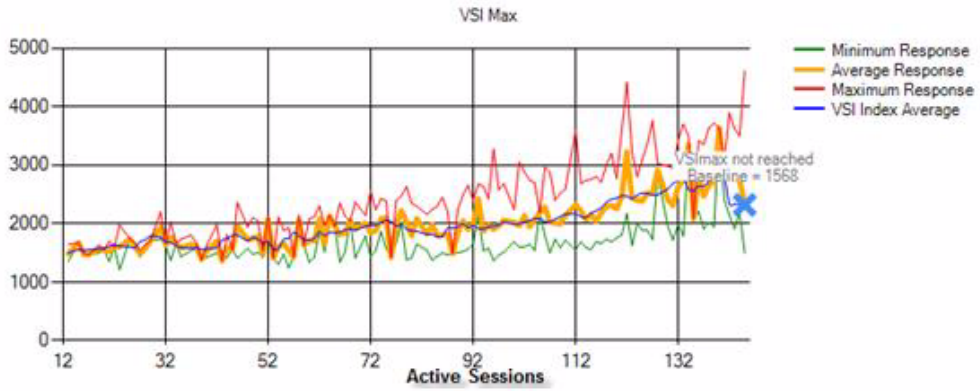

The following graphs detail CPU, Memory, Disk and Network performance on the Single Cisco UCS B230-M2 Blades.

Figure 47 User Single B230 M2 CPU Utilization All Phases

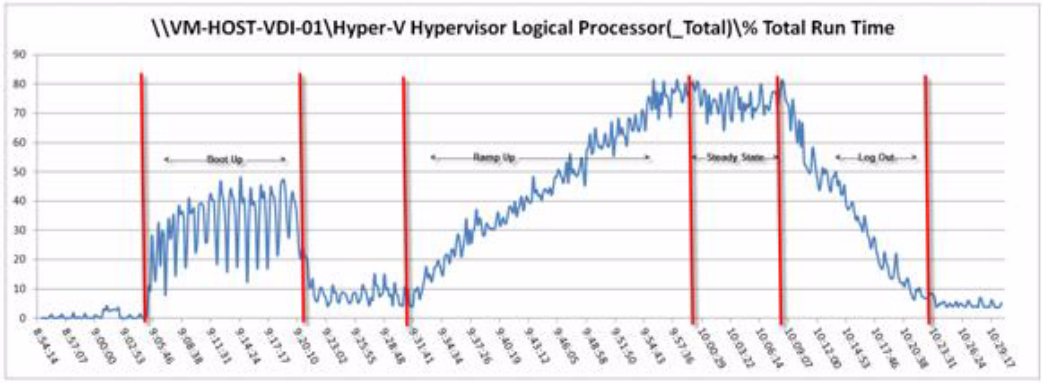

Figure 48 145 User Single B230 M2 Available Memory

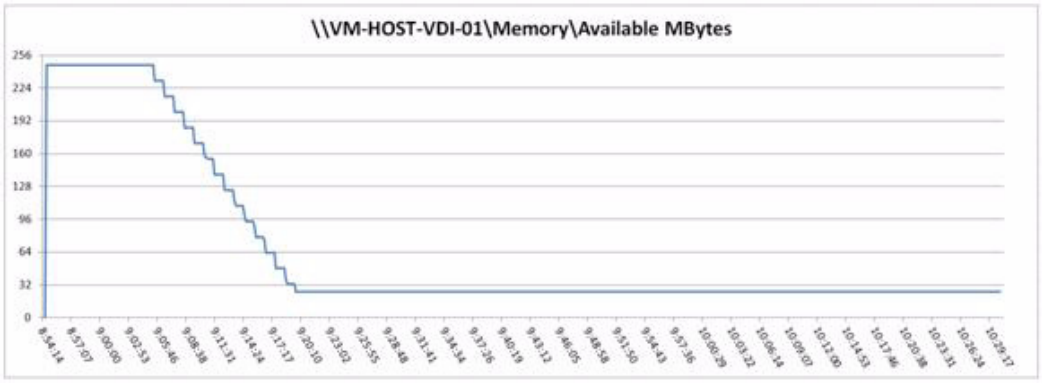

Figure 49 Single B230 M2 Cisco M81KR VIC Mbits Received Per Second

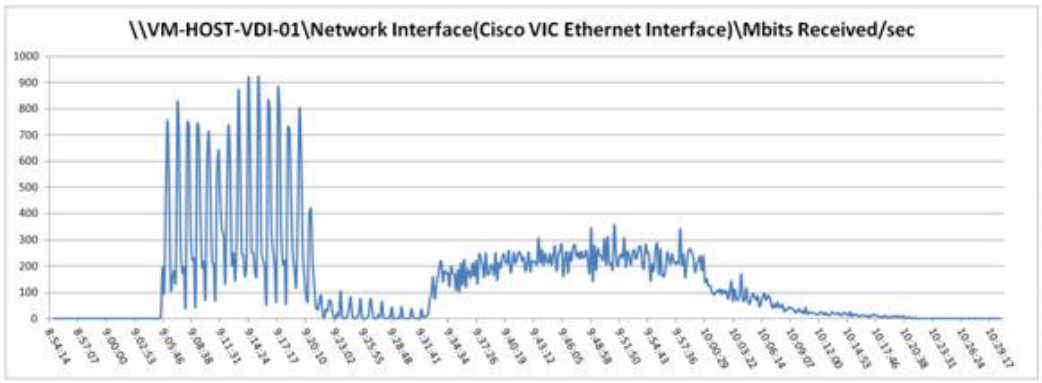

Figure 50 Single B230 M2 Cisco M81KR VIC Mbits Sent Per Second

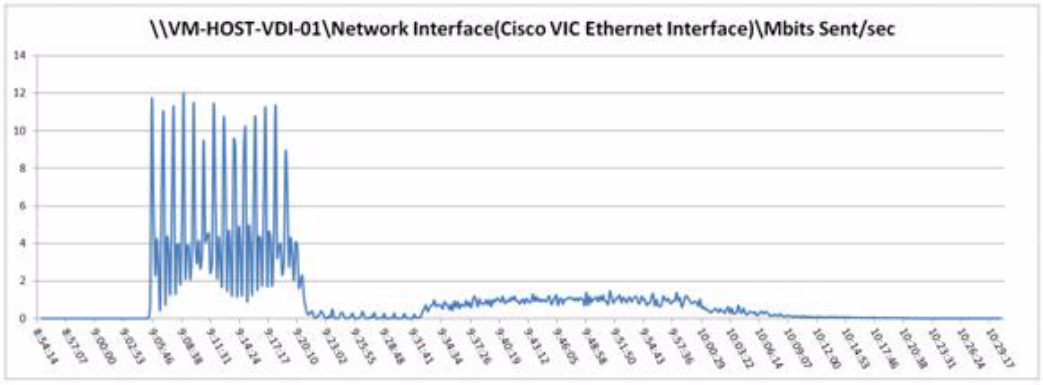

The following graphs detail performance of the NetApp FAS3240 System during the single blade, 145 user test.

Figure 51 145 User NetApp Read IOPS

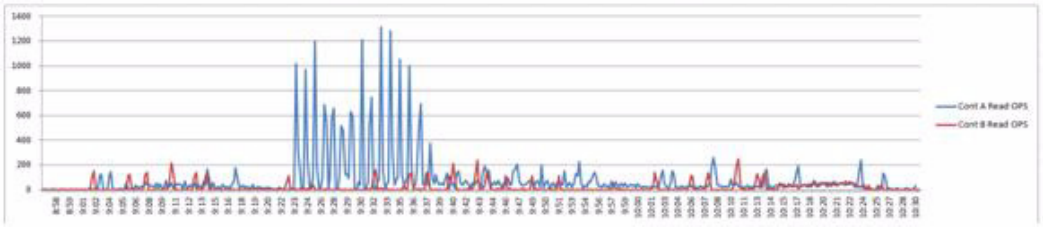

Figure 52 145 User NetApp Write IOPS

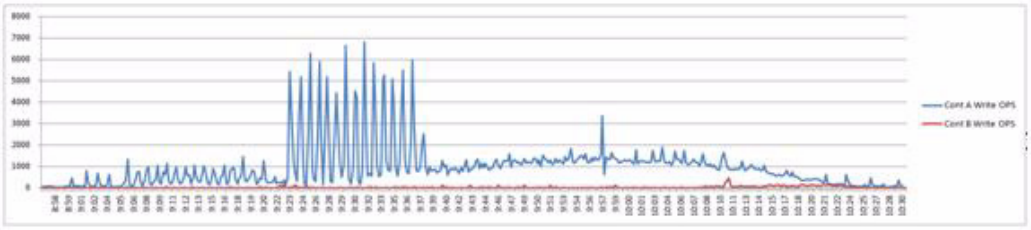

Figure 53 145 User NetApp Read Latency

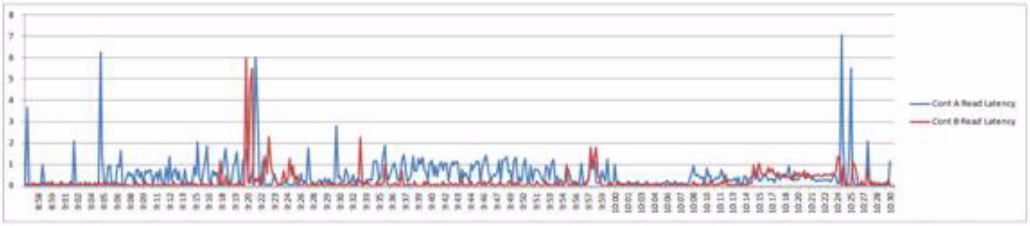

Figure 54 145 User NetApp Write Latency

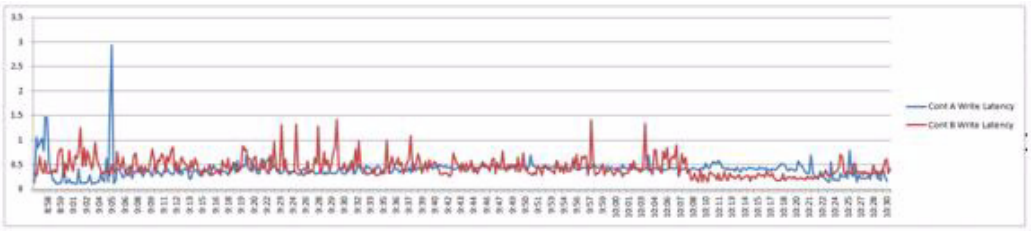

Figure 55 145 User NetApp Processor Average Utilization

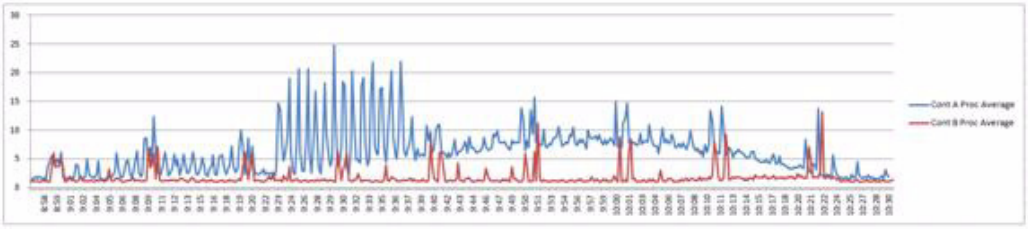

Figure 56 145 User NetApp Processor Total Utilization

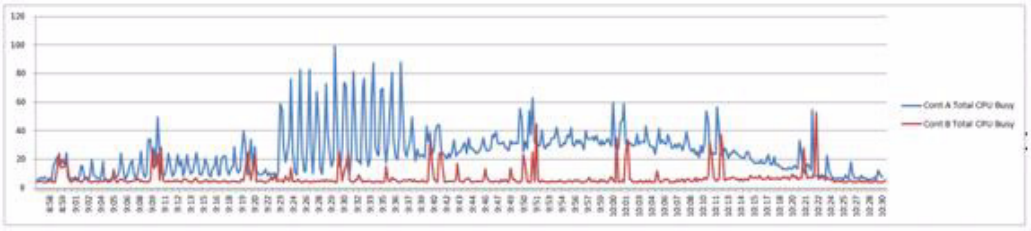

The following charts detail infrastructure server performance during the single blade, 145 User test:

Figure 57 145 User Virtual Machine Manager 2012 CPU Utilization

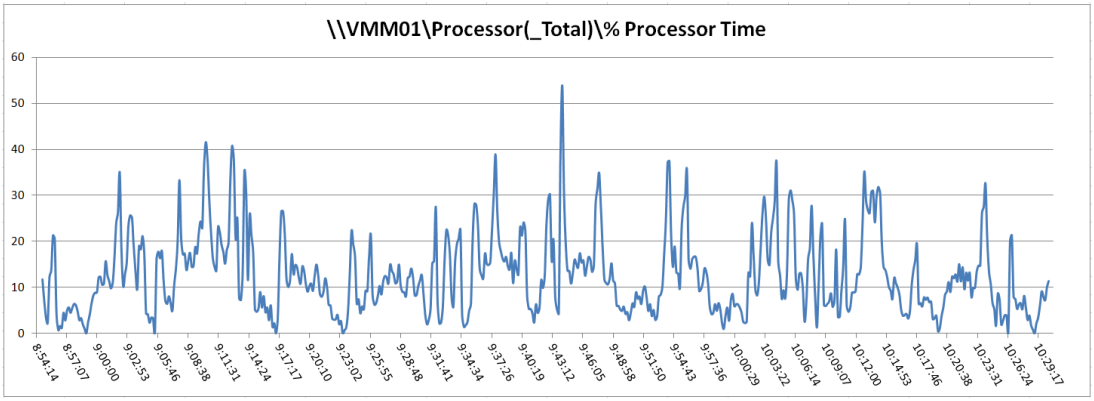

Figure 58 145 User Provisioning Server PVS01 CPU utilization

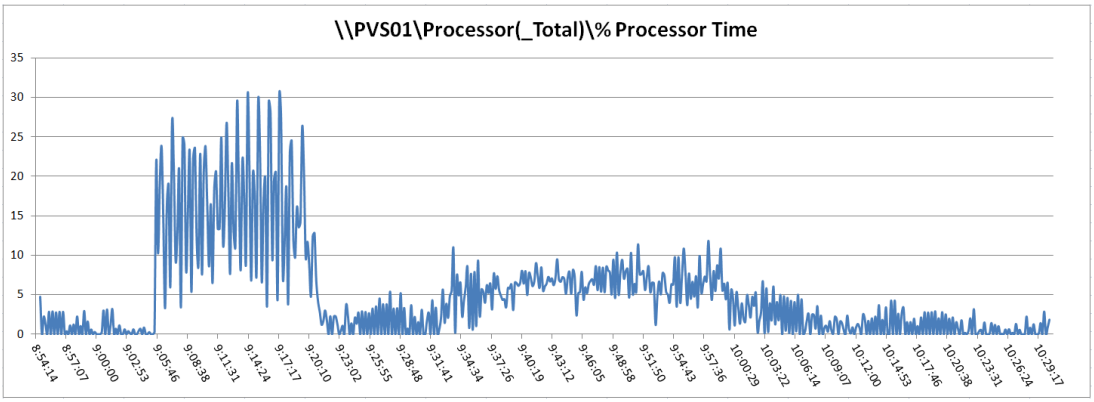

Figure 59 145 User Provisioning Server PVS01 Available Memory

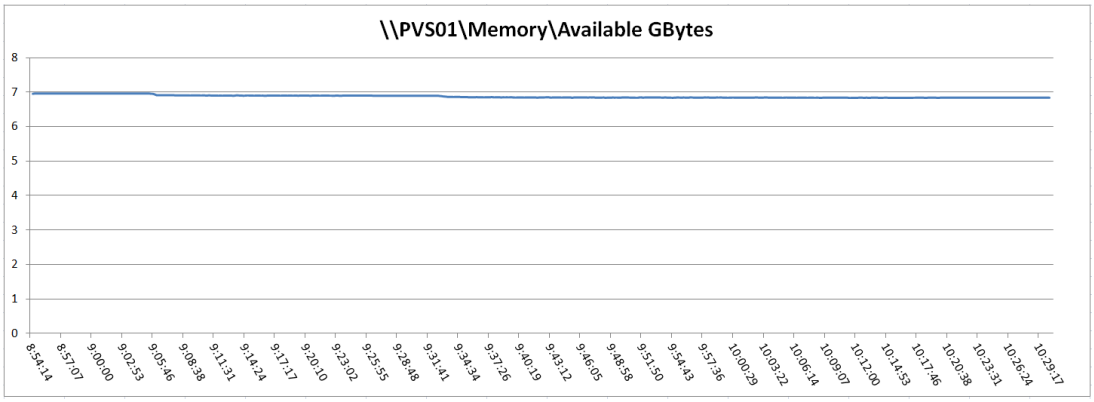

Figure 60 145 User Provisioning Server PVS01 Mbits Sent/Second

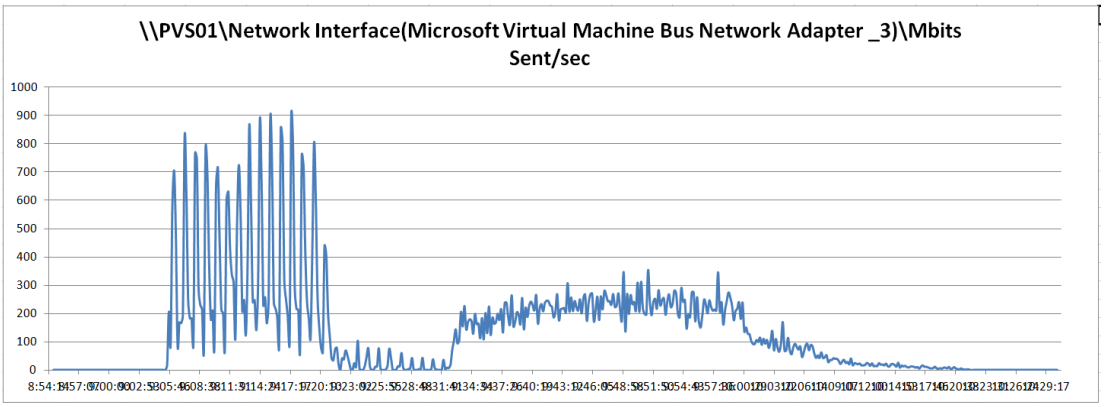

Figure 61 145 User XenDesktop 5.6 Controller XDC01 CPU Utilization

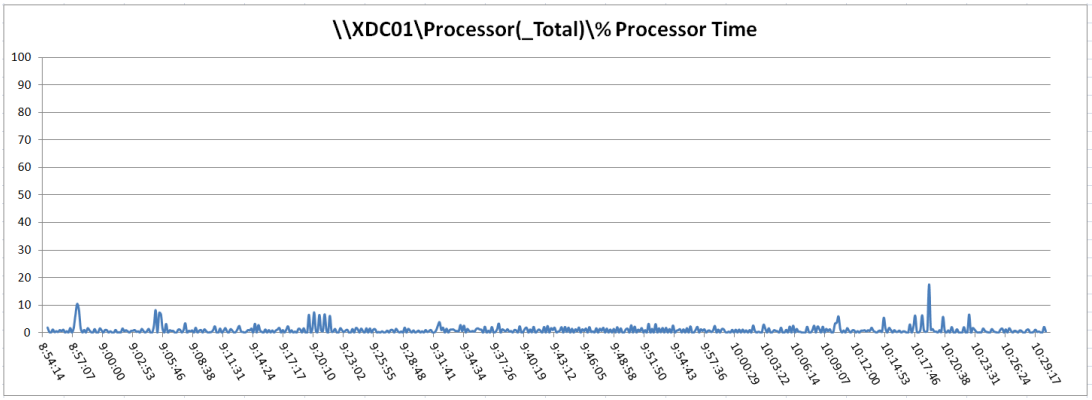

Fourteen Cisco UCS B230-M2 Blade Server Validation - 2000 Users

This section details the results from the XenDesktop Hosted VDI fourteen blade server validation testing. It illustrates linear scalability from one blade at 145 Users to fourteen blades at 2000 users.

Note 2000 users on fourteen blades represents 143 users running on each of twelve blades and 142 users running on two additional blades. The 2000 user stopping point was based on Microsoft's currently supported limit on the number of virtual desktops managed by a System Center 2012 Virtual Machine Manager.

The primary success criteria metrics are provided to validate the overall success of the test cycle as an output chart from Login Consultants' VSI Analyzer Professional Edition, VSIMax Dynamic for the Medium workload (with Flash.)

Additional graphs detailing the CPU and Memory utilization during peak session load are also presented. Given adequate storage capability, the limiting factor in the testing was CPU utilization. Each of the fourteen B230 M2 blade's performance charts are essentially identical for the multi-blade runs. We are including data on one randomly chosen blade in this document to represent all of the blades in this portion of the study.

The single server graphs shown in this section are representative of a single Hyper-V host in the larger environment for validation purposes, but it should be noted that these graphs are representative of the behavior for all servers in the respective environment.

We present performance information on key infrastructure virtual machines with the randomly chosen tested blade information.

Figure 62 2000 Desktop Sessions on Microsoft Hyper-V Server 2008 R2 SP1 Below 4000 ms

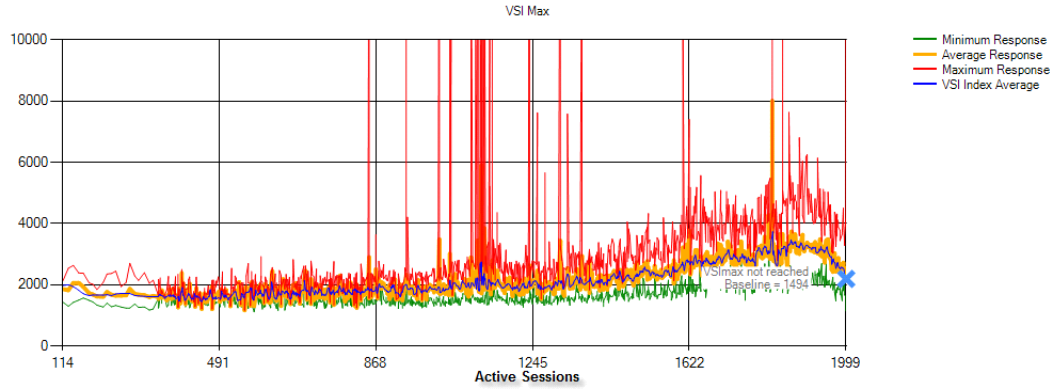

Figure 63 2000 User Fourteen B230 M2 Available Memory All Phases - Single Server Example

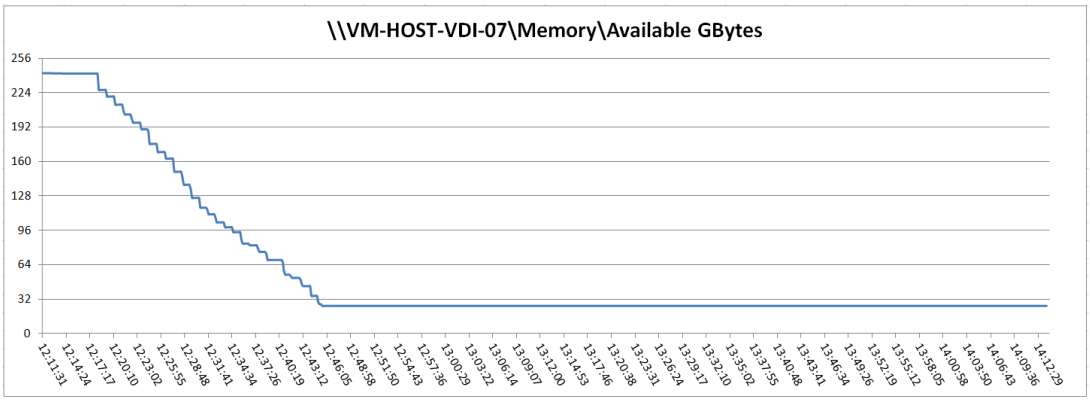

Figure 64 2000 User NetApp Read IOPS All Phases

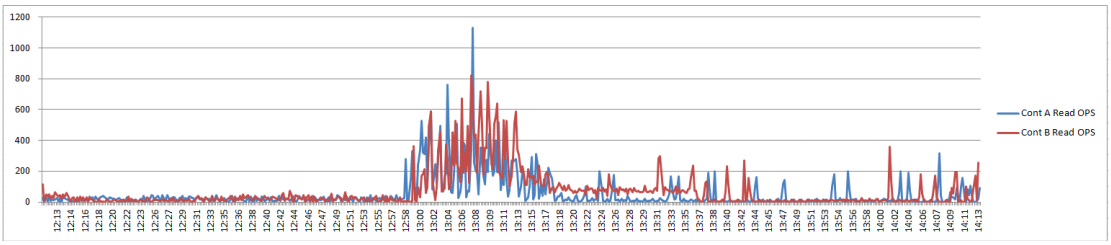

Figure 65 2000 User NetApp Write IOPS All Phases

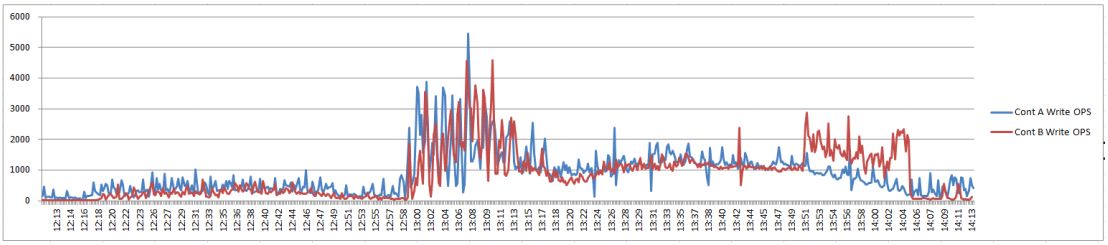

Figure 66 2000 User NetApp Read Latency

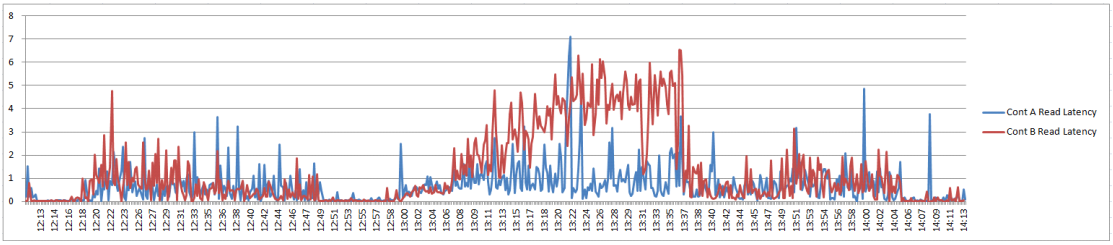

Figure 67 2000 User NetApp Write Latency

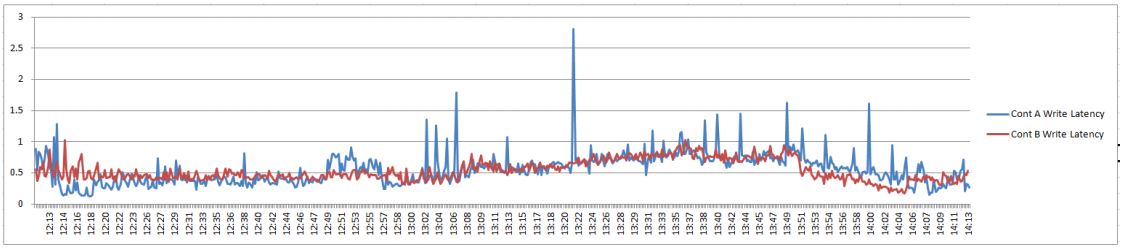

Figure 68 2000 User NetApp Processor Average Utilization

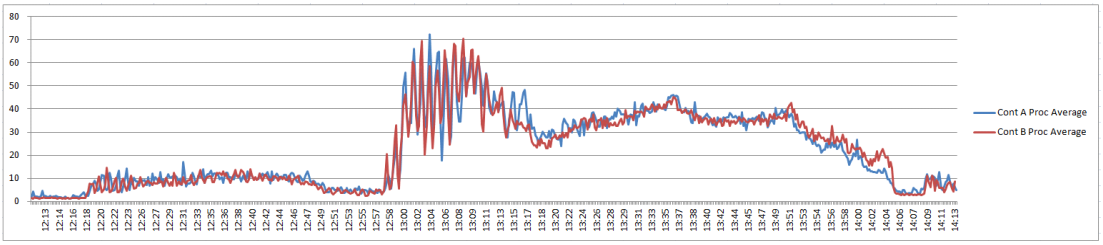

Figure 69 2000 User NetApp Processor Total CPU Busy

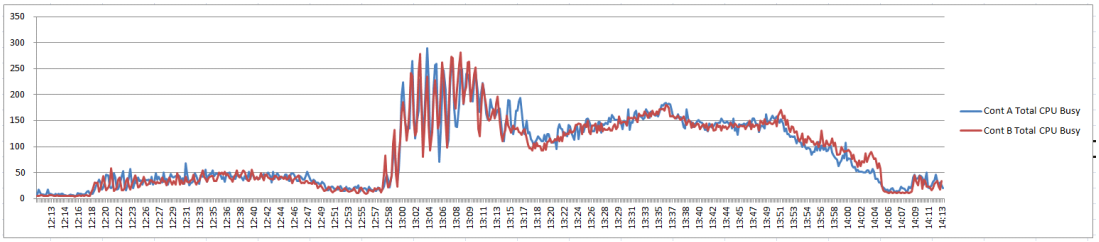

Figure 70 2000 User NetApp FC OPS

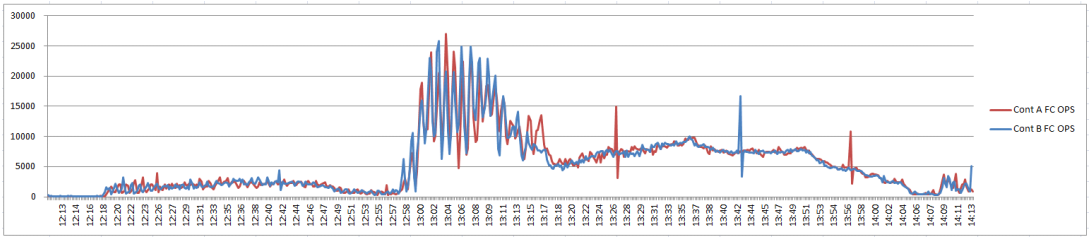

Figure 71 2000 User NetApp iSCSI OPS

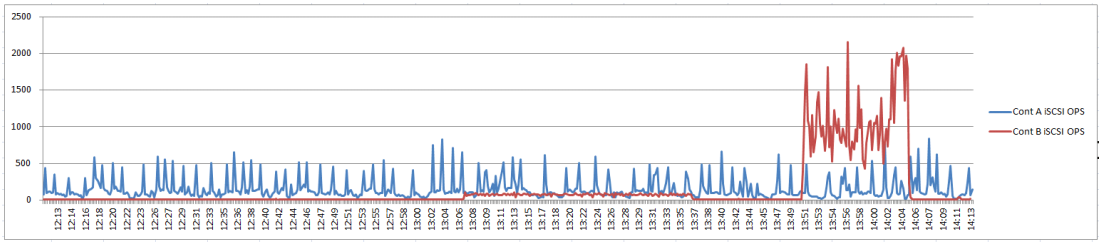

The following charts detail infrastructure server performance during the fourteen blade, 2000 User test.

Figure 72 2000 User Provisioning Services 6.1 PVS01 CPU Utilization

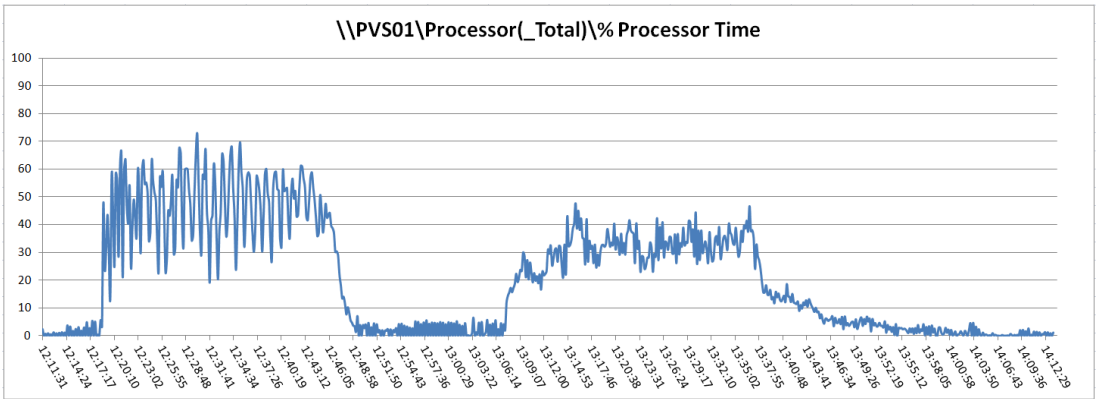

Figure 73 2000 User Provisioning Services 6.1 PVS01 Available Memory

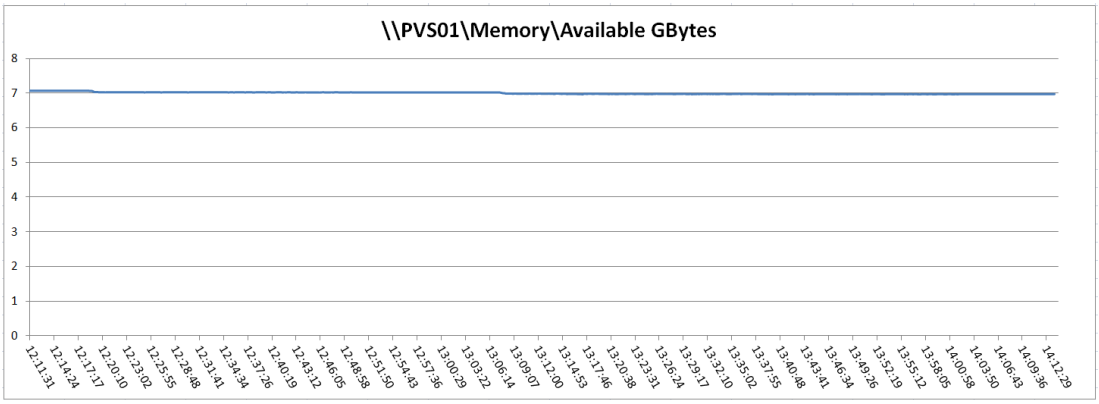

Figure 74 2000 User Provisioning Services 6.1 PVS01 Mbits/Second Sent

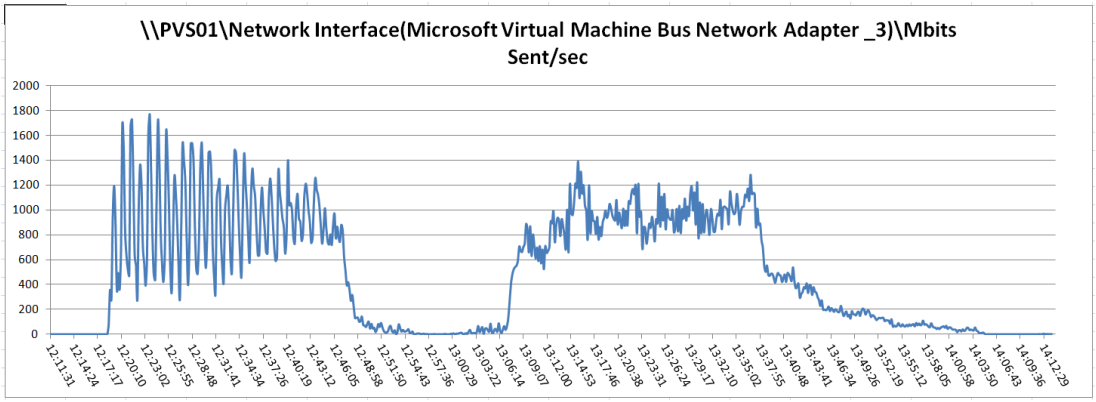

Figure 75 2000 User Provisioning Services 6.1 PVS02 CPU Utilization

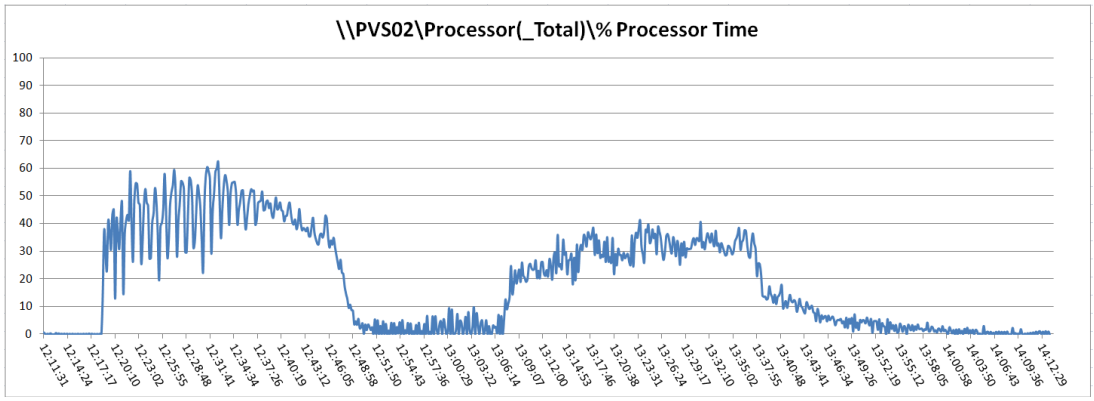

Figure 76 2000 User Provisioning Services 6.1 PVS02 Available Memory

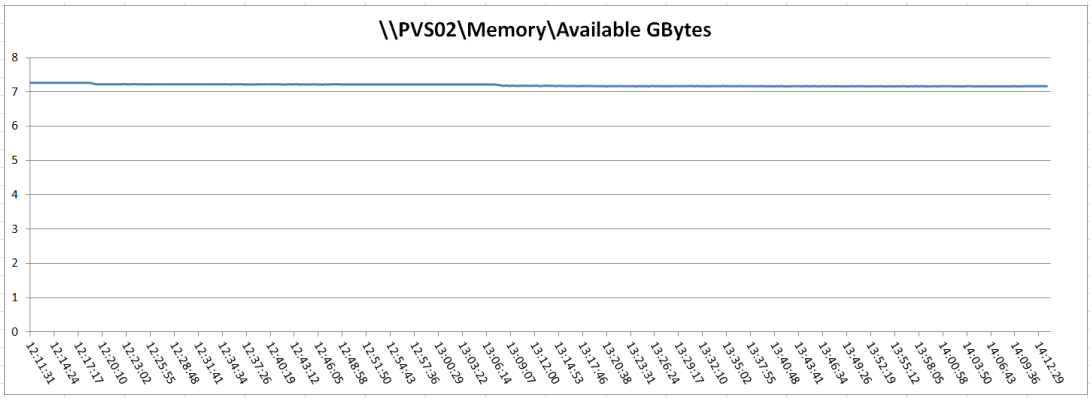

Figure 77 2000 User Provisioning Services 6.1 PVS02 Mbits/Second Sent

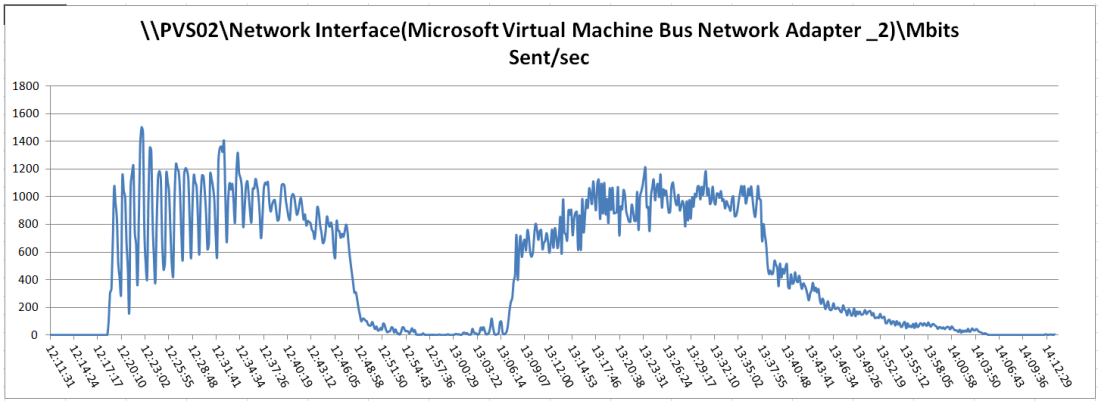

Figure 78 2000 User Provisioning Services 6.1 PVS03 CPU Utilization

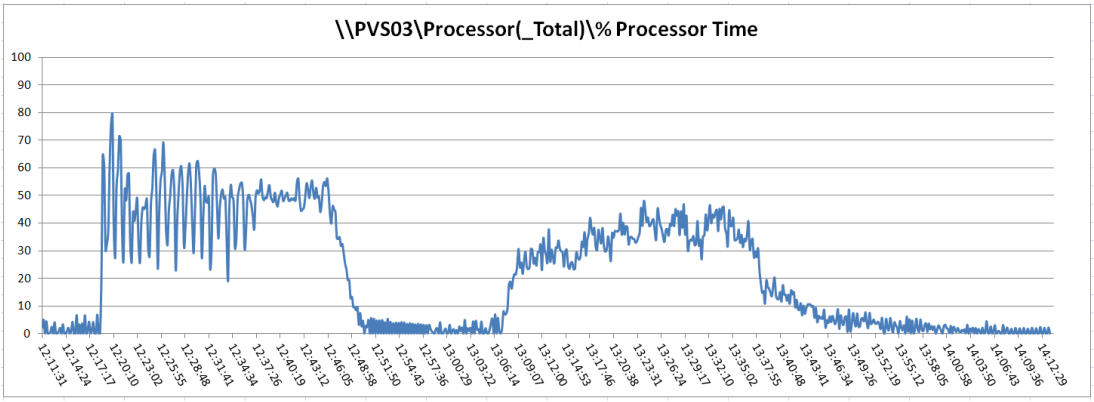

Figure 79 2000 User Provisioning Services 6.1 PVS03 Available Memory

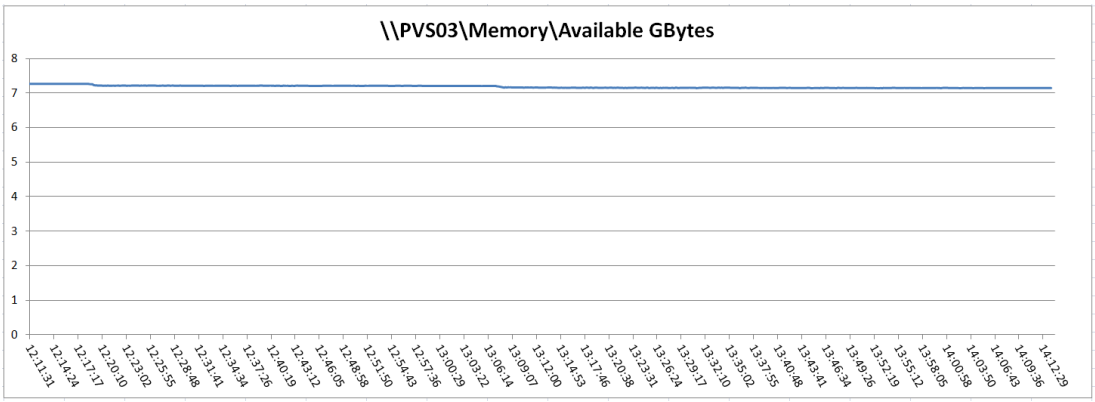

Figure 80 2000 User Provisioning Services 6.1 PVS03 Mbits/Second Sent

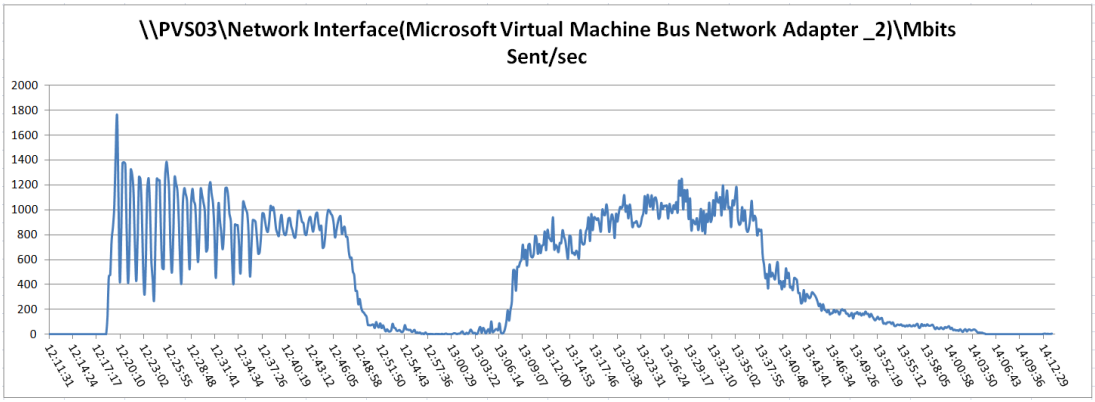

Figure 81 2000 User XenDesktop 5.6 Controller XDC01 CPU Utilization

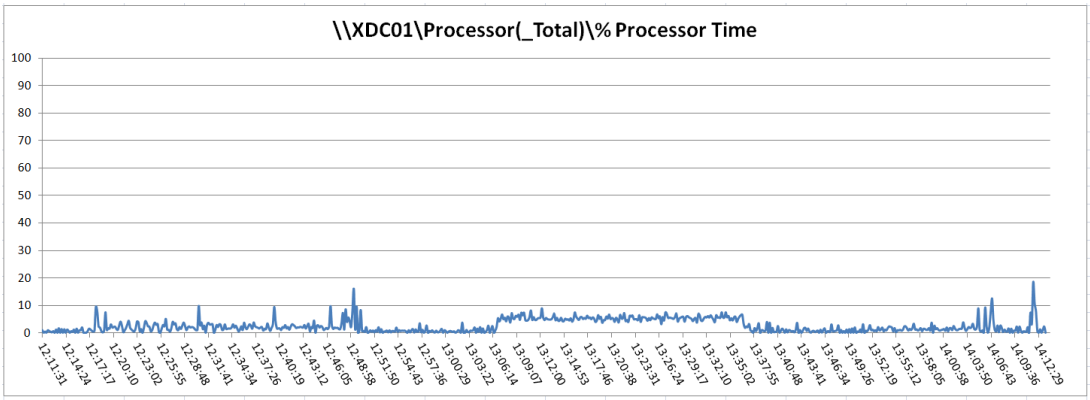

Figure 82 2000 User XenDesktop 5.6 Controller XDC02 CPU Utilization

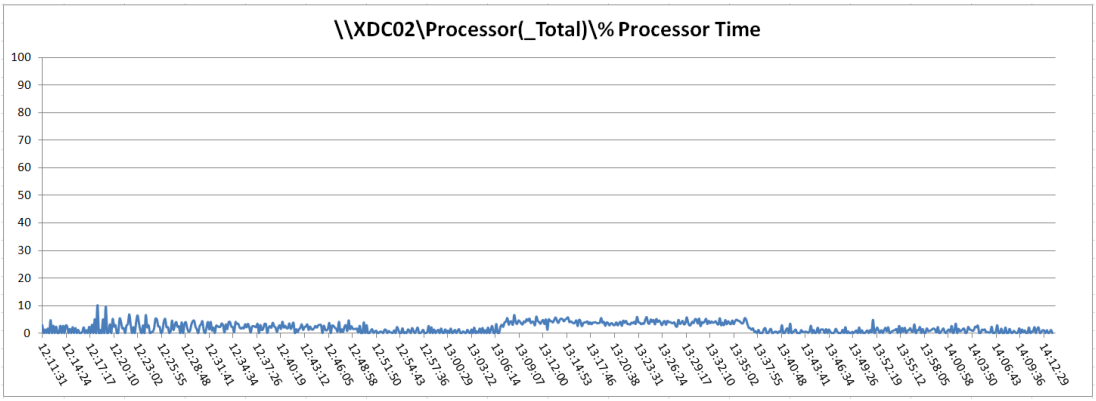

Scalability Considerations and Guidelines

There are many factors to consider when you begin to scale beyond the FlexPod two chassis, 14 VDI host server configuration, which this reference architecture has successfully tested. In this section we give guidance to scale beyond two Cisco UCS chassis.

Cisco UCS System Configuration

As our results indicate, we have proven linear scalability in the Cisco UCS Reference Architecture as tested.

Table 11 Hyper-V Server 2008 R2 SP1 and System Center 2012 Virtual Machine Manager hosting XenDesktop 5.6 with Provisioning Services 6.1

Number of Chassis Tested	Number of Cisco UCS B230 M2 Servers Tested	Number of VMs Hosted	VM/Physical Core Ratio
1	1	145	7.25
2	14	2000*	

*Scale test was limited to 2000 virtual desktops per Microsoft supported number of virtual desktops for Virtual Machine Manager 2012 installation.

- Cisco UCS 2.0 management software supports up to 20 chassis within a single Cisco UCS domain on our second generation Cisco UCS Fabric Interconnect 6249 and 6296 models.
- Provided we had adequate storage capability to support 10 Flexpod for VDI two-chassis units, we could easily scale to 20,000 users in a single UCS domain. We would need to add additional hypervisor virtual machine management virtual servers on each two-chassis FlexPod for VDI deployment to manage the VDI environment.
- To accommodate the Cisco Nexus 5500 upstream connectivity in the way we describe in the LAN Configuration section, we need four Ethernet uplinks to be configured on the Cisco UCS Fabric interconnect. And based on the number of uplinks from each chassis, we could calculate how many desktops can be hosted in a single Cisco UCS domain. Assuming eight links per chassis, four to each 6248, scaling beyond 10 chassis would require a pair of Cisco UCS 6296 fabric interconnects. A 20,000 virtual desktop building block can be built out of the RA described in this study with eight links per chassis and 20 Cisco UCS chassis comprised of seven B230 M2 and one B200 M2 blades servers in each chassis.

The backend storage has to be scaled accordingly, based on the IOP considerations as described in the NetApp scaling section. Please refer the NetApp section that follows this one for scalability guidelines.

Storage Sizing Best Practices

Storage estimation for deploying VDI solutions on NetApp includes the following:

- Gather essential solution requirements
- Perform performance-based and capacity-based storage estimation
- Get recommendations on storage system physical and logical configuration

Gather Essential Solution Requirements

The first step of the storage sizing process is to gather the solution requirements. This is essential to size the storage system correctly in terms of the model and the number of required NetApp storage controllers, type and quantity of disk spindles, software features, and general configuration recommendations.

The main storage sizing elements are:

- Total number of virtual machines for which the system has to be designed (for example, 2000 virtual machines).
- The types and percentage of different types of desktops being deployed. For example, if Citrix XenDesktop is used, different desktop delivery models might require special storage considerations.
- Size per virtual machine (for example, 20GB C: drive, 2GB data disk).
- Virtual machine OS for example, Windows XP, Windows 7, and so on).
- Worker workload profile (type of applications on the virtual machine, IOPS requirement, read-write ratio, if known).
- Number of years for which the storage growth has to be considered.
- Disaster recovery and business continuance requirements.

- Size of NAS (CIFS) home directories.
- NetApp strongly recommends storing user data on NAS (CIFS) home drives. By using NAS home drives, companies can more efficiently manage and protect the user data and eliminate the need to back up the virtual desktops.
- For most of the Citrix XenDesktop deployments, companies might also plan to implement roaming profiles and/or folder redirection.

For detailed information on implementing these technologies, consult the following documentation:

- Microsoft Configuring Roaming User Profiles
- NetApp TR-3367: NetApp Systems in a Microsoft Windows Environment
- Microsoft Configuring Folder Redirection

Citrix XenDesktop Considerations

When implementing Citrix XenDesktop, decide on the following:

- Types of desktops that will be deployed for different user profiles
- Data protection requirements for different data components (OS disk, user data disk, CIFS home directories) for each desktop type being implemented
- For Citrix Provisioning Server pooled desktops, write cache size needs to be calculated based on how often the user reboots the desktop and what applications the user uses. We recommend using a write cache 2x the RAM allocated to each individual Virtual Machine. For example, if a Virtual Machine is allocated with 2GB RAM, use a 4GB write cache VHD for each VM.
- NetApp thin provisioning, deduplication, and NetApp snapshot can be used to achieve the desired storage efficiency and data protection for all virtual disks and storage.

Performance-Based and Capacity-Based Storage Estimation Processes

There are two important considerations for sizing storage for Citrix XenDesktop. The storage system should be able to meet both the performance and capacity requirements of the project and be scalable to account for future growth.

The steps for calculating these storage requirements are:

- Determine storage sizing building block
- Perform detailed performance estimation
- Perform detailed capacity estimation
- Obtain recommendations on the storage system physical and logical configuration

Getting Recommendations on Storage System Physical and Logical Configuration

After determining the total capacity and performance requirements, contact your local NetApp technical resource to determine the appropriate storage system configuration. Provide the total capacity and performance requirements to the NetApp SE and obtain appropriate storage system configuration. If required, NetApp can help you in each phase of the process discussed above. NetApp has detailed sizing tools specific to VDI that can help architect deployments of any scale. The tools are designed to factor in all the NetApp storage efficiency and performance acceleration components discussed earlier.

This step also involves planning the logical architecture (the total number of template and the associated FlexClone volumes that should be provisioned per aggregate). The recommendation is to provision fewer large aggregates over more, smaller aggregates. The advantages to larger aggregates are that the I/O has more disks to write across, therefore increasing the performance of all volumes contained within the aggregate.

Based on the estimated volume size from the capacity calculations section earlier, determine the number of template and associated FlexClone volumes that can be hosted in the largest possible aggregate. It is also a good idea to leave some room to grow the aggregates to handle situations when unexpected growth occurs. Also, disable scheduled aggregate Snapshot copies and set the aggregate snap reserve to zero. Make sure the data disk in the aggregate satisfies the performance requirements for the proposed number of virtual machines for volumes to be hosted in the aggregate.

References

This section provides links to additional information for each partner's solution component of this document.

Cisco Reference Documents

FlexPod with Microsoft Private Cloud: Architecture Overview

http://www.cisco.com/en/US/solutions/collateral/ns340/ns517/ns224/ns1150/ns1154/ucs_flexpod_ms_netapp.pdf

Cisco Unified Computing System Manager Home Page

http://www.cisco.com/en/US/products/ps10281/index.html

Cisco UCS B230 M2 Blade Server Resources

http://www.cisco.com/en/US/products/ps11583/index.html

Cisco UCS B200 M2 Blade Server Resources

http://www.cisco.com/en/US/products/ps10915/index.html

Cisco UCS B200 M3 Blade Server Resources

http://www.cisco.com/en/US/products/ps12288/index.html

Download Cisco UCS Manager and Blade Software Version 2.0(1w)

http://tools.cisco.com/search/JSP/search-results.get?strQueryText=ucs+2.0%281w%29&strDocsPerPage=10&strStartDoc=1&strSortBy=cdcdevfour&strqueryid=2&websessionid=6VZ_ymKQsRIgLlEEiyDynfL&country=US&language=en&profile=enushomesppublished

Citrix Reference Documents

Citrix Consulting White Papers

http://support.citrix.com/product/xd/v5.5/consulting/

Microsoft Reference Documents

Windows Server 2008 R2

http://www.microsoft.com/en-us/server-cloud/windows-server/default.aspx

Active Directory

http://www.microsoft.com/en-us/server-cloud/windows-server/active-directory.aspx

File Services

http://www.microsoft.com/en-us/server-cloud/windows-server/file-and-print-services.aspx

Failover Clustering

http://www.microsoft.com/en-us/server-cloud/windows-server/failover-clustering-network-load-balancing.aspx

Hyper-V Role on Server 2008 R2

http://www.microsoft.com/en-us/server-cloud/windows-server/hyper-v.aspx

Hyper-V server

http://www.microsoft.com/en-us/server-cloud/hyper-v-server/default.aspx

Deeper Dive to HyperV features

http://www.microsoft.com/en-us/server-cloud/windows-server/hyper-v-features.aspx

Windows 7 Enterprise

http://www.microsoft.com/en-us/windows/enterprise/products-and-technologies/windows-7/features.aspx

NetApp Reference Documents

NetApp TR-3563: NetApp Thin Provisioning http://media.netapp.com/documents/tr-3563.pdf

NetApp TR-3505: NetApp Deduplication for FAS, Deployment and Implementation Guide http://media.netapp.com/documents/tr-3505.pdf

NetApp TR-3347: FlexClone Volumes: A Thorough Introduction http://media.netapp.com/documents/tr-3347.pdf

NetApp TR-3437: Storage Best Practices and Resiliency Guide http://media.netapp.com/documents/wp-3437.pdf

Appendix A—Nexus Configurations

Nexus 5548-A

version 5.1(3)N1(1a)

feature fcoe

hostname msoft-vdi-n5548-a

```
feature npiv
feature fport-channel-trunk
feature telnet
no feature http-server
cfs eth distribute
feature interface-vlan
feature lacp
feature vpc
feature lldp
feature fex
username admin password 5 $1$nLc3JOYW$GzDHEH0ksSioYa/UjdJr.0  role network-admin
banner motd #Nexus 5000 Switch
#
ip domain-lookup
logging event link-status default
ip access-list classiffy_silver
  10 permit ip 10.16.0.0/16 any
  20 permit ip any 10.16.0.0/16
  30 permit ip 10.17.0.0/16 any
  40 permit ip any 10.17.0.0/16
ip access-list classify_silver
class-map type qos class-fcoe
class-map type qos match-all class-gold
  match cos 4
class-map type qos match-all class-silver
  match access-group name classiffy_silver
class-map type queuing class-fcoe
  match qos-group 1
class-map type queuing class-gold
  match qos-group 3
class-map type queuing class-silver
  match qos-group 4
class-map type queuing class-all-flood
  match qos-group 2
class-map type queuing class-ip-multicast
  match qos-group 2
policy-map type qos system_qos_policy
  class class-gold
```

```
    set qos-group 3
  class class-silver
    set qos-group 4
  class class-fcoe
    set qos-group 1
  class class-default
    set qos-group 0
policy-map type queuing system_q_in_policy
  class type queuing class-fcoe
    bandwidth percent 20
  class type queuing class-gold
    bandwidth percent 33
  class type queuing class-silver
    bandwidth percent 29
  class type queuing class-default
    bandwidth percent 18
policy-map type queuing system_q_out_policy
  class type queuing class-fcoe
    bandwidth percent 20
  class type queuing class-gold
    bandwidth percent 33
  class type queuing class-silver
    bandwidth percent 29
  class type queuing class-default
    bandwidth percent 18
class-map type network-qos class-fcoe
  match qos-group 1
class-map type network-qos class-gold
  match qos-group 3
class-map type network-qos class-silver
  match qos-group 4
class-map type network-qos class-all-flood
  match qos-group 2
class-map type network-qos class-ip-multicast
  match qos-group 2
policy-map type network-qos system_nq_policy
  class type network-qos class-gold
    set cos 4
```

```
    mtu 9000
  class type network-qos class-fcoe
    pause no-drop
    mtu 2158
  class type network-qos class-silver
    set cos 2
    mtu 9000
  class type network-qos class-default
    mtu 9000
    multicast-optimize
system qos
  service-policy type qos input system_qos_policy
  service-policy type queuing input system_q_in_policy
  service-policy type queuing output system_q_out_policy
  service-policy type network-qos system_nq_policy
slot 1
  port 21-32 type fc
snmp-server user admin network-admin auth md5 0x43c8181d058c52019b4961b64ee16164 priv
0x43c8181d058c52019b4961b64ee16164 localizedkey
vrf context management
vlan 1,8
vlan 2222
  name App-Cluster-Comm
vlan 2233
  name public
vlan 2244
  name vm-pvs
vlan 2255
  name mgmt
vlan 2266
  name CSV
vlan 2277
  name Live-Migration
vlan 2288
  name iSCSI-A
vlan 2299
  name iSCSI-B
spanning-tree port type edge bpduguard default
```

```
spanning-tree port type edge bpdufilter default
spanning-tree vlan 1,2233,2244 priority 24576
vpc domain 10
  peer-keepalive destination 10.25.2.15
device-alias database
  device-alias name citrixvdi-a-0c pwwn 50:0a:09:81:9d:42:bd:e8
  device-alias name citrixvdi-b-0c pwwn 50:0a:09:81:8d:42:bd:e8
  device-alias name c1-s1-b200m2-fc0 pwwn 20:00:00:25:b5:c1:b1:01
  device-alias name hyperV-ucs-2-fc0 pwwn 20:00:00:25:b5:c1:b2:01
  device-alias name hyperV-ucs-3-fc0 pwwn 20:00:00:25:b5:c1:b3:01
  device-alias name vm-host-vdi-01-fc0 pwwn 20:00:00:25:b5:02:02:3f
  device-alias name vm-host-vdi-02-fc0 pwwn 20:00:00:25:b5:02:02:1f
  device-alias name vm-host-vdi-03-fc0 pwwn 20:00:00:25:b5:02:02:3e
  device-alias name vm-host-vdi-04-fc0 pwwn 20:00:00:25:b5:02:02:3b
  device-alias name vm-host-vdi-05-fc0 pwwn 20:00:00:25:b5:02:02:1e
  device-alias name vm-host-vdi-06-fc0 pwwn 20:00:00:25:b5:02:02:1b
  device-alias name vm-host-vdi-07-fc0 pwwn 20:00:00:25:b5:02:02:3d
  device-alias name vm-host-vdi-08-fc0 pwwn 20:00:00:25:b5:02:02:3a
  device-alias name vm-host-vdi-09-fc0 pwwn 20:00:00:25:b5:02:02:1d
  device-alias name vm-host-vdi-10-fc0 pwwn 20:00:00:25:b5:02:02:1a
  device-alias name vm-host-vdi-11-fc0 pwwn 20:00:00:25:b5:02:02:3c
  device-alias name vm-host-vdi-12-fc0 pwwn 20:00:00:25:b5:02:02:29
  device-alias name vm-host-vdi-13-fc0 pwwn 20:00:00:25:b5:02:02:1c
  device-alias name vm-host-vdi-14-fc0 pwwn 20:00:00:25:b5:02:02:09
  device-alias name vm-host-infra-01-fc0 pwwn 20:00:00:25:b5:02:02:28
  device-alias name vm-host-infra-02-fc0 pwwn 20:00:00:25:b5:02:02:08

device-alias commit

fcdomain fcid database
  vsan 1 wwn 20:1f:54:7f:ee:1b:fe:00 fcid 0x410000 dynamic
  vsan 1 wwn 50:0a:09:81:8d:42:bd:e8 fcid 0x410001 dynamic
!           [citrixvdi-b-0c]
  vsan 1 wwn 50:0a:09:81:9d:42:bd:e8 fcid 0x410002 dynamic
!           [citrixvdi-a-0c]
  vsan 1 wwn 20:00:00:25:b5:02:02:3a fcid 0x410003 dynamic
!           [vm-host-vdi-08-fc0]
  vsan 1 wwn 20:00:00:25:b5:02:02:1b fcid 0x410004 dynamic
```

```
!          [vm-host-vdi-06-fc0]
  vsan 1 wwn 20:00:00:25:b5:02:02:09 fcid 0x410005 dynamic
!          [vm-host-vdi-14-fc0]
  vsan 1 wwn 20:00:00:25:b5:02:02:1d fcid 0x410006 dynamic
!          [vm-host-vdi-09-fc0]
  vsan 1 wwn 20:00:00:25:b5:02:02:3b fcid 0x410007 dynamic
!          [vm-host-vdi-04-fc0]
  vsan 1 wwn 20:00:00:25:b5:02:02:3d fcid 0x410008 dynamic
!          [vm-host-vdi-07-fc0]
  vsan 1 wwn 20:00:00:25:b5:02:02:29 fcid 0x410009 dynamic
!          [vm-host-vdi-12-fc0]
  vsan 1 wwn 20:00:00:25:b5:02:02:1e fcid 0x41000a dynamic
!          [vm-host-vdi-05-fc0]
  vsan 1 wwn 20:00:00:25:b5:02:02:1f fcid 0x41000b dynamic
!          [vm-host-vdi-02-fc0]
  vsan 1 wwn 20:00:00:25:b5:02:02:1a fcid 0x41000c dynamic
!          [vm-host-vdi-10-fc0]
  vsan 1 wwn 20:00:00:25:b5:02:02:08 fcid 0x41000d dynamic
!          [vm-host-infra-02-fc0]
  vsan 1 wwn 20:00:00:25:b5:02:02:1c fcid 0x41000e dynamic
!          [vm-host-vdi-13-fc0]
  vsan 1 wwn 20:00:00:25:b5:02:02:3e fcid 0x41000f dynamic
!          [vm-host-vdi-03-fc0]
  vsan 1 wwn 20:00:00:25:b5:02:02:3c fcid 0x410010 dynamic
!          [vm-host-vdi-11-fc0]
  vsan 1 wwn 20:00:00:25:b5:02:02:3f fcid 0x410011 dynamic
!          [vm-host-vdi-01-fc0]
  vsan 1 wwn 20:00:00:25:b5:02:02:28 fcid 0x410012 dynamic
!          [vm-host-infra-01-fc0]
  vsan 1 wwn 20:20:54:7f:ee:1b:fe:00 fcid 0x410013 dynamic
  vsan 1 wwn 20:00:00:25:b5:02:02:07 fcid 0x410014 dynamic
  vsan 1 wwn 20:00:00:25:b5:02:02:26 fcid 0x410015 dynamic
  vsan 1 wwn 20:00:00:25:b5:02:02:25 fcid 0x410016 dynamic
  vsan 1 wwn 20:00:00:25:b5:02:02:06 fcid 0x410017 dynamic


interface Vlan1
```

```
interface Vlan2233
  no shutdown
  ip address 10.11.240.2/16

interface san-port-channel 1
  channel mode active

interface port-channel10
  switchport mode trunk
  switchport trunk native vlan 2
  switchport trunk allowed vlan 2222,2233,2244,2255,2266,2277,2288,2299
  spanning-tree port type network
  vpc peer-link

interface port-channel11

interface port-channel17
  switchport access vlan 2299
  vpc 17

interface port-channel18
  switchport access vlan 2288
  vpc 18

interface port-channel19
  switchport mode trunk
  spanning-tree port type edge trunk
  vpc 19

interface port-channel20
  switchport mode trunk
  spanning-tree port type edge trunk
  vpc 20

interface fc1/21
  switchport trunk mode off
  no shutdown
```

```
interface fc1/22
  switchport trunk mode off
  no shutdown

interface fc1/23
  switchport trunk mode off
  no shutdown

interface fc1/24
  switchport trunk mode off
  no shutdown

interface fc1/25

interface fc1/26

interface fc1/27

interface fc1/28

interface fc1/29

interface fc1/30

interface fc1/31

interface fc1/32

interface Ethernet1/1
  switchport access vlan 2233

interface Ethernet1/2

interface Ethernet1/3

interface Ethernet1/4

interface Ethernet1/5
```

```
interface Ethernet1/6

interface Ethernet1/7
  switchport access vlan 2233
  speed 1000

interface Ethernet1/8
  switchport access vlan 2233
  speed 1000

interface Ethernet1/9

interface Ethernet1/10

interface Ethernet1/11

interface Ethernet1/12

interface Ethernet1/13
  switchport mode trunk
  channel-group 19 mode active

interface Ethernet1/14
  switchport mode trunk
  channel-group 19 mode active

interface Ethernet1/15
  switchport mode trunk
  channel-group 20 mode active

interface Ethernet1/16
  switchport mode trunk
  channel-group 20 mode active

interface Ethernet1/17
  description netapp-fas3240-ctrl-b-10gb
  switchport access vlan 2299
```

```
  channel-group 17 mode active

interface Ethernet1/18
  description netapp-fas3240-ctrl-a-10gb
  switchport access vlan 2288
  channel-group 18 mode active

interface Ethernet1/19
  switchport mode trunk
  switchport trunk native vlan 2
  switchport trunk allowed vlan 2222,2233,2244,2255,2266,2277,2288,2299
  channel-group 10

interface Ethernet1/20
  switchport mode trunk
  switchport trunk native vlan 2
  switchport trunk allowed vlan 2222,2233,2244,2255,2266,2277,2288,2299
  channel-group 10

interface mgmt0
  ip address 10.25.2.16/22
line console
line vty
boot kickstart bootflash:/n5000-uk9-kickstart.5.1.3.N1.1a.bin
boot system bootflash:/n5000-uk9.5.1.3.N1.1a.bin
interface fc1/21
interface fc1/22
interface fc1/23
interface fc1/24
interface fc1/25
interface fc1/26
interface fc1/27
interface fc1/28
interface fc1/29
interface fc1/30
interface fc1/31
interface fc1/32
!Full Zone Database Section for vsan 1
```

```
zone name c1-s1-b200m2-fc0 vsan 1
   member pwwn 50:0a:09:81:9d:42:bd:e8
!            [citrixvdi-a-0c]
   member pwwn 20:00:00:25:b5:c1:b1:01
!            [c1-s1-b200m2-fc0]

zone name hyperV-ucs-2-fc0 vsan 1
   member pwwn 20:00:00:25:b5:c1:b2:01
!            [hyperV-ucs-2-fc0]
   member pwwn 50:0a:09:81:9d:42:bd:e8
!            [citrixvdi-a-0c]

zone name hyperV-ucs-3-fc0 vsan 1
   member pwwn 20:00:00:25:b5:c1:b3:01
!            [hyperV-ucs-3-fc0]
   member pwwn 50:0a:09:81:9d:42:bd:e8
!            [citrixvdi-a-0c]

zone name vm-host-vdi-01-fc0 vsan 1
   member pwwn 20:00:00:25:b5:02:02:3f
!            [vm-host-vdi-01-fc0]
   member pwwn 50:0a:09:81:9d:42:bd:e8
!            [citrixvdi-a-0c]
   member pwwn 50:0a:09:81:8d:42:bd:e8
!            [citrixvdi-b-0c]

zone name vm-host-vdi-02-fc0 vsan 1
   member pwwn 20:00:00:25:b5:02:02:1f
!            [vm-host-vdi-02-fc0]
   member pwwn 50:0a:09:81:9d:42:bd:e8
!            [citrixvdi-a-0c]
   member pwwn 50:0a:09:81:8d:42:bd:e8
!            [citrixvdi-b-0c]

zone name vm-host-vdi-03-fc0 vsan 1
   member pwwn 20:00:00:25:b5:02:02:3e
!            [vm-host-vdi-03-fc0]
   member pwwn 50:0a:09:81:9d:42:bd:e8
```

```
!           [citrixvdi-a-0c]
  member pwwn 50:0a:09:81:8d:42:bd:e8
!           [citrixvdi-b-0c]

zone name vm-host-vdi-04-fc0 vsan 1
  member pwwn 20:00:00:25:b5:02:02:3b
!           [vm-host-vdi-04-fc0]
  member pwwn 50:0a:09:81:9d:42:bd:e8
!           [citrixvdi-a-0c]
  member pwwn 50:0a:09:81:8d:42:bd:e8
!           [citrixvdi-b-0c]

zone name vm-host-vdi-05-fc0 vsan 1
  member pwwn 20:00:00:25:b5:02:02:1e
!           [vm-host-vdi-05-fc0]
  member pwwn 50:0a:09:81:9d:42:bd:e8
!           [citrixvdi-a-0c]
  member pwwn 50:0a:09:81:8d:42:bd:e8
!           [citrixvdi-b-0c]

zone name vm-host-vdi-06-fc0 vsan 1
  member pwwn 20:00:00:25:b5:02:02:1b
!           [vm-host-vdi-06-fc0]
  member pwwn 50:0a:09:81:9d:42:bd:e8
!           [citrixvdi-a-0c]
  member pwwn 50:0a:09:81:8d:42:bd:e8
!           [citrixvdi-b-0c]

zone name vm-host-vdi-07-fc0 vsan 1
  member pwwn 20:00:00:25:b5:02:02:3d
!           [vm-host-vdi-07-fc0]
  member pwwn 50:0a:09:81:9d:42:bd:e8
!           [citrixvdi-a-0c]
  member pwwn 50:0a:09:81:8d:42:bd:e8
!           [citrixvdi-b-0c]

zone name vm-host-vdi-08-fc0 vsan 1
  member pwwn 20:00:00:25:b5:02:02:3a
```

```
!            [vm-host-vdi-08-fc0]
   member pwwn 50:0a:09:81:9d:42:bd:e8
!            [citrixvdi-a-0c]
   member pwwn 50:0a:09:81:8d:42:bd:e8
!            [citrixvdi-b-0c]

zone name vm-host-vdi-09-fc0 vsan 1
   member pwwn 20:00:00:25:b5:02:02:1d
!            [vm-host-vdi-09-fc0]
   member pwwn 50:0a:09:81:9d:42:bd:e8
!            [citrixvdi-a-0c]
   member pwwn 50:0a:09:81:8d:42:bd:e8
!            [citrixvdi-b-0c]

zone name vm-host-vdi-10-fc0 vsan 1
   member pwwn 50:0a:09:81:9d:42:bd:e8
!            [citrixvdi-a-0c]
   member pwwn 50:0a:09:81:8d:42:bd:e8
!            [citrixvdi-b-0c]
   member pwwn 20:00:00:25:b5:02:02:1a
!            [vm-host-vdi-10-fc0]

zone name vm-host-vdi-11-fc0 vsan 1
   member pwwn 50:0a:09:81:9d:42:bd:e8
!            [citrixvdi-a-0c]
   member pwwn 20:00:00:25:b5:02:02:3c
!            [vm-host-vdi-11-fc0]
   member pwwn 50:0a:09:81:8d:42:bd:e8
!            [citrixvdi-b-0c]

zone name vm-host-vdi-12-fc0 vsan 1
   member pwwn 20:00:00:25:b5:02:02:29
!            [vm-host-vdi-12-fc0]
   member pwwn 50:0a:09:81:9d:42:bd:e8
!            [citrixvdi-a-0c]
   member pwwn 50:0a:09:81:8d:42:bd:e8
!            [citrixvdi-b-0c]
```

```
zone name vm-host-vdi-13-fc0 vsan 1
   member pwwn 20:00:00:25:b5:02:02:1c
!            [vm-host-vdi-13-fc0]
   member pwwn 50:0a:09:81:9d:42:bd:e8
!            [citrixvdi-a-0c]
   member pwwn 50:0a:09:81:8d:42:bd:e8
!            [citrixvdi-b-0c]

zone name vm-host-vdi-14-fc0 vsan 1
   member pwwn 20:00:00:25:b5:02:02:09
!            [vm-host-vdi-14-fc0]
   member pwwn 50:0a:09:81:9d:42:bd:e8
!            [citrixvdi-a-0c]
   member pwwn 50:0a:09:81:8d:42:bd:e8
!            [citrixvdi-b-0c]

zone name vm-host-infra-01-fc0 vsan 1
   member pwwn 20:00:00:25:b5:02:02:28
!            [vm-host-infra-01-fc0]
   member pwwn 50:0a:09:81:9d:42:bd:e8
!            [citrixvdi-a-0c]
   member pwwn 50:0a:09:81:8d:42:bd:e8
!            [citrixvdi-b-0c]

zone name vm-host-infra-02-fc0 vsan 1
   member pwwn 20:00:00:25:b5:02:02:08
!            [vm-host-infra-02-fc0]
   member pwwn 50:0a:09:81:9d:42:bd:e8
!            [citrixvdi-a-0c]
   member pwwn 50:0a:09:81:8d:42:bd:e8
!            [citrixvdi-b-0c]

zoneset name ms-vdi vsan 1
   member c1-s1-b200m2-fc0
   member hyperV-ucs-2-fc0
   member vm-host-vdi-01-fc0
   member vm-host-vdi-02-fc0
   member vm-host-vdi-03-fc0
```

```
  member vm-host-vdi-04-fc0
  member vm-host-vdi-05-fc0
  member vm-host-vdi-06-fc0
  member vm-host-vdi-07-fc0
  member vm-host-vdi-08-fc0
  member vm-host-vdi-09-fc0
  member vm-host-vdi-10-fc0
  member vm-host-vdi-11-fc0
  member vm-host-vdi-12-fc0
  member vm-host-vdi-13-fc0
  member vm-host-vdi-14-fc0
  member vm-host-infra-01-fc0
  member vm-host-infra-02-fc0

zoneset activate name ms-vdi vsan 1
```

Nexus 5548-B

```
version 5.1(3)N1(1a)
feature fcoe
hostname msoft-vdi-n5548-b
feature npiv
feature fport-channel-trunk
feature telnet
no feature http-server
cfs eth distribute
feature interface-vlan
feature lacp
feature vpc
feature lldp
feature fex
username admin password 5 $1$kjpHgfqH$qWFuj6a1QIYVVpBnuopgn0  role network-admin

banner motd #Nexus 5000 Switch
#

ip domain-lookup
logging event link-status default
```

```
ip access-list classiffy_silver
  10 permit ip 10.16.0.0/16 any
  20 permit ip any 10.16.0.0/16
  30 permit ip 10.17.0.0/16 any
  40 permit ip any 10.17.0.0/16
class-map type qos class-fcoe
class-map type qos match-all class-gold
  match cos 4
class-map type qos match-all class-silver
  match access-group name classiffy_silver
class-map type queuing class-fcoe
  match qos-group 1
class-map type queuing class-gold
  match qos-group 3
class-map type queuing class-silver
  match qos-group 4
class-map type queuing class-all-flood
  match qos-group 2
class-map type queuing class-ip-multicast
  match qos-group 2
policy-map type qos system_qos_policy
  class class-gold
    set qos-group 3
  class class-silver
    set qos-group 4
  class class-fcoe
    set qos-group 1
  class class-default
    set qos-group 0
policy-map type queuing system_q-in_policy
  class type queuing class-fcoe
    bandwidth percent 20
  class type queuing class-gold
    bandwidth percent 33
  class type queuing class-silver
    bandwidth percent 29
  class type queuing class-default
    bandwidth percent 18
```

```
policy-map type queuing system_q_in_policy
  class type queuing class-fcoe
    bandwidth percent 20
  class type queuing class-gold
    bandwidth percent 33
  class type queuing class-silver
    bandwidth percent 29
  class type queuing class-default
    bandwidth percent 18
policy-map type queuing system_q_out_policy
  class type queuing class-fcoe
    bandwidth percent 20
  class type queuing class-gold
    bandwidth percent 33
  class type queuing class-silver
    bandwidth percent 29
  class type queuing class-default
    bandwidth percent 18
class-map type network-qos class-fcoe
  match qos-group 1
class-map type network-qos class-gold
  match qos-group 3
class-map type network-qos class-silver
  match qos-group 4
class-map type network-qos class-all-flood
  match qos-group 2
class-map type network-qos class-ip-multicast
  match qos-group 2
policy-map type network-qos system_nq_policy
  class type network-qos class-gold
    set cos 4
    mtu 9000
  class type network-qos class-fcoe
    pause no-drop
    mtu 2158
  class type network-qos class-silver
    set cos 2
    mtu 9000
```

```
  class type network-qos class-default
    mtu 9000
    multicast-optimize
system qos
  service-policy type qos input system_qos_policy
  service-policy type queuing output system_q_out_policy
  service-policy type network-qos system_nq_policy
  service-policy type queuing input system_q_in_policy
slot 1
  port 21-32 type fc
snmp-server user admin network-admin auth md5 0x60b6015e3c29399aa87578dd5a0300db
 priv 0x60b6015e3c29399aa87578dd5a0300db localizedkey
vrf context management
vlan 1
vlan 2
  name Native
vlan 2222
  name App-Cluster-Comm
vlan 2233
  name public
vlan 2244
  name vm-pvs
vlan 2255
  name mgmt
vlan 2266
  name CSV
vlan 2277
  name Live-Migration
vlan 2288
  name iSCSI-A
vlan 2299
  name iSCSI-B
spanning-tree port type edge bpduguard default
spanning-tree port type edge bpdufilter default
spanning-tree vlan 1,2233,2244 priority 57344
vpc domain 10
  peer-keepalive destination 10.25.2.16
device-alias database
```

```
  device-alias name citrixvdi-a-0d pwwn 50:0a:09:82:9d:42:bd:e8
  device-alias name citrixvdi-b-0d pwwn 50:0a:09:82:8d:42:bd:e8
  device-alias name c1-s1-b200m2-fc1 pwwn 20:00:00:25:b5:c1:b1:02
  device-alias name hyperV-ucs-2-fc1 pwwn 20:00:00:25:b5:c1:b2:02
  device-alias name vm-host-vdi-01-fc1 pwwn 20:00:00:25:b5:02:02:2f
  device-alias name vm-host-vdi-02-fc1 pwwn 20:00:00:25:b5:02:02:0f
  device-alias name vm-host-vdi-03-fc1 pwwn 20:00:00:25:b5:02:02:2e
  device-alias name vm-host-vdi-04-fc1 pwwn 20:00:00:25:b5:02:02:2b
  device-alias name vm-host-vdi-05-fc1 pwwn 20:00:00:25:b5:02:02:0e
  device-alias name vm-host-vdi-06-fc1 pwwn 20:00:00:25:b5:02:02:0b
  device-alias name vm-host-vdi-07-fc1 pwwn 20:00:00:25:b5:02:02:2d
  device-alias name vm-host-vdi-08-fc1 pwwn 20:00:00:25:b5:02:02:2a
  device-alias name vm-host-vdi-09-fc1 pwwn 20:00:00:25:b5:02:02:0d
  device-alias name vm-host-vdi-10-fc1 pwwn 20:00:00:25:b5:02:02:0a
  device-alias name vm-host-vdi-11-fc1 pwwn 20:00:00:25:b5:02:02:2c
  device-alias name vm-host-vdi-12-fc1 pwwn 20:00:00:25:b5:02:02:39
  device-alias name vm-host-vdi-13-fc1 pwwn 20:00:00:25:b5:02:02:0c
  device-alias name vm-host-vdi-14-fc1 pwwn 20:00:00:25:b5:02:02:19
  device-alias name vm-host-infra-01-fc1 pwwn 20:00:00:25:b5:02:02:38
  device-alias name vm-host-infra-02-fc1 pwwn 20:00:00:25:b5:02:02:18

device-alias commit

fcdomain fcid database
  vsan 1 wwn 20:1f:54:7f:ee:1b:ff:00 fcid 0x850000 dynamic
  vsan 1 wwn 20:20:54:7f:ee:1b:ff:00 fcid 0x850001 dynamic
  vsan 1 wwn 20:00:00:25:b5:02:02:0d fcid 0x850002 dynamic
!           [vm-host-vdi-09-fc1]
  vsan 1 wwn 20:00:00:25:b5:02:02:2d fcid 0x850003 dynamic
!           [vm-host-vdi-07-fc1]
  vsan 1 wwn 20:00:00:25:b5:02:02:0e fcid 0x850004 dynamic
!           [vm-host-vdi-05-fc1]
  vsan 1 wwn 20:00:00:25:b5:02:02:0c fcid 0x850005 dynamic
!           [vm-host-vdi-13-fc1]
  vsan 1 wwn 20:00:00:25:b5:02:02:2e fcid 0x850006 dynamic
!           [vm-host-vdi-03-fc1]
  vsan 1 wwn 20:00:00:25:b5:02:02:2c fcid 0x850007 dynamic
!           [vm-host-vdi-11-fc1]
```

```
  vsan 1 wwn 20:00:00:25:b5:02:02:38 fcid 0x850008 dynamic
!           [vm-host-infra-01-fc1]
  vsan 1 wwn 20:00:00:25:b5:02:02:2f fcid 0x850009 dynamic
!           [vm-host-vdi-01-fc1]
  vsan 1 wwn 20:00:00:25:b5:02:02:2a fcid 0x85000a dynamic
!           [vm-host-vdi-08-fc1]
  vsan 1 wwn 20:00:00:25:b5:02:02:0b fcid 0x85000b dynamic
!           [vm-host-vdi-06-fc1]
  vsan 1 wwn 20:00:00:25:b5:02:02:19 fcid 0x85000c dynamic
!           [vm-host-vdi-14-fc1]
  vsan 1 wwn 20:00:00:25:b5:02:02:2b fcid 0x85000d dynamic
!           [vm-host-vdi-04-fc1]
  vsan 1 wwn 20:00:00:25:b5:02:02:39 fcid 0x85000e dynamic
!           [vm-host-vdi-12-fc1]
  vsan 1 wwn 20:00:00:25:b5:02:02:0f fcid 0x85000f dynamic
!           [vm-host-vdi-02-fc1]
  vsan 1 wwn 20:00:00:25:b5:02:02:0a fcid 0x850010 dynamic
!           [vm-host-vdi-10-fc1]
  vsan 1 wwn 20:00:00:25:b5:02:02:18 fcid 0x850011 dynamic
!           [vm-host-infra-02-fc1]
  vsan 1 wwn 50:0a:09:82:8d:42:bd:e8 fcid 0x850012 dynamic
!           [citrixvdi-b-0d]
  vsan 1 wwn 50:0a:09:82:9d:42:bd:e8 fcid 0x850013 dynamic
!           [citrixvdi-a-0d]
  vsan 1 wwn 20:00:00:25:b5:02:02:17 fcid 0x850014 dynamic
  vsan 1 wwn 20:00:00:25:b5:02:02:36 fcid 0x850015 dynamic
  vsan 1 wwn 20:00:00:25:b5:02:02:35 fcid 0x850016 dynamic
  vsan 1 wwn 20:00:00:25:b5:02:02:16 fcid 0x850017 dynamic


interface Vlan1

interface Vlan2233
  no shutdown
  ip address 10.11.240.3/16

interface san-port-channel 2
  channel mode active
```

```
interface port-channel10
  switchport mode trunk
  spanning-tree port type network
  vpc peer-link

interface port-channel17
  switchport access vlan 2299
  vpc 17

interface port-channel18
  switchport access vlan 2288
  vpc 18

interface port-channel19
  switchport mode trunk
  spanning-tree port type edge trunk  vpc 19

interface port-channel20
  switchport mode trunk
  spanning-tree port type edge trunk
  vpc 20

interface fc1/21
  switchport trunk mode off
  no shutdown

interface fc1/22
  switchport trunk mode off
  no shutdown

interface fc1/23
  switchport trunk mode off
  no shutdown

interface fc1/24
  switchport trunk mode off
  no shutdown
```

```
interface fc1/25

interface fc1/26

interface fc1/27

interface fc1/28

interface fc1/29

interface fc1/30

interface fc1/31

interface fc1/32

interface Ethernet1/1
  switchport access vlan 2233

interface Ethernet1/2

interface Ethernet1/3

interface Ethernet1/4

interface Ethernet1/5

interface Ethernet1/6

interface Ethernet1/7
  switchport access vlan 2233
  speed 1000

interface Ethernet1/8
  switchport access vlan 2233
  speed 1000
```

```
interface Ethernet1/9

interface Ethernet1/10

interface Ethernet1/11

interface Ethernet1/12

interface Ethernet1/13
  switchport mode trunk
  channel-group 19 mode active

interface Ethernet1/14
  switchport mode trunk
  channel-group 19 mode active

interface Ethernet1/15
  switchport mode trunk
  channel-group 20 mode active

interface Ethernet1/16
  switchport mode trunk
  channel-group 20 mode active

interface Ethernet1/17
  description netapp-fas3240-ctrl-b-10gb
  switchport access vlan 2299
  channel-group 17 mode active

interface Ethernet1/18
  description netapp-fas3240-ctrl-a-10gb
  switchport access vlan 2288
  channel-group 18 mode active

interface Ethernet1/19
  switchport mode trunk
  channel-group 10
```

```
interface Ethernet1/20
  switchport mode trunk
  channel-group 10

interface mgmt0
  ip address 10.25.2.15/22
line console
line vty
boot kickstart bootflash:/n5000-uk9-kickstart.5.1.3.N1.1a.bin
boot system bootflash:/n5000-uk9.5.1.3.N1.1a.bin
interface fc1/21
interface fc1/22
interface fc1/23
interface fc1/24
interface fc1/25
interface fc1/26
interface fc1/27
interface fc1/28
interface fc1/29
interface fc1/30
interface fc1/31
interface fc1/32
!Full Zone Database Section for vsan 1
zone name c1-s1-b200m2-fc1 vsan 1
    member pwwn 50:0a:09:82:9d:42:bd:e8
!           [citrixvdi-a-0d]
    member pwwn 20:00:00:25:b5:c1:b1:02
!           [c1-s1-b200m2-fc1]

zone name hyperV-ucs-2-fc1 vsan 1
    member pwwn 50:0a:09:82:9d:42:bd:e8
!           [citrixvdi-a-0d]
    member pwwn 20:00:00:25:b5:c1:b2:02
!           [hyperV-ucs-2-fc1]

zone name vm-host-infra-01-fc1 vsan 1
    member pwwn 50:0a:09:82:9d:42:bd:e8
!           [citrixvdi-a-0d]
```

```
   member pwwn 50:0a:09:82:8d:42:bd:e8
!            [citrixvdi-b-0d]
   member pwwn 20:00:00:25:b5:02:02:38
!            [vm-host-infra-01-fc1]

zone name vm-host-vdi-02-fc1 vsan 1
   member pwwn 20:00:00:25:b5:02:02:0f
!            [vm-host-vdi-02-fc1]
   member pwwn 50:0a:09:82:9d:42:bd:e8
!            [citrixvdi-a-0d]
   member pwwn 50:0a:09:82:8d:42:bd:e8
!            [citrixvdi-b-0d]

zone name vm-host-vdi-03-fc1 vsan 1
   member pwwn 20:00:00:25:b5:02:02:2e
!            [vm-host-vdi-03-fc1]
   member pwwn 50:0a:09:82:9d:42:bd:e8
!            [citrixvdi-a-0d]
   member pwwn 50:0a:09:82:8d:42:bd:e8
!            [citrixvdi-b-0d]

zone name vm-host-vdi-04-fc1 vsan 1
   member pwwn 20:00:00:25:b5:02:02:2b
!            [vm-host-vdi-04-fc1]
   member pwwn 50:0a:09:82:9d:42:bd:e8
!            [citrixvdi-a-0d]
   member pwwn 50:0a:09:82:8d:42:bd:e8
!            [citrixvdi-b-0d]

zone name vm-host-vdi-05-fc1 vsan 1
   member pwwn 20:00:00:25:b5:02:02:0e
!            [vm-host-vdi-05-fc1]
   member pwwn 50:0a:09:82:9d:42:bd:e8
!            [citrixvdi-a-0d]
   member pwwn 50:0a:09:82:8d:42:bd:e8
!            [citrixvdi-b-0d]

zone name vm-host-vdi-06-fc1 vsan 1
```

```
   member pwwn 20:00:00:25:b5:02:02:0b
!            [vm-host-vdi-06-fc1]
   member pwwn 50:0a:09:82:9d:42:bd:e8
!            [citrixvdi-a-0d]
   member pwwn 50:0a:09:82:8d:42:bd:e8
!            [citrixvdi-b-0d]

zone name vm-host-vdi-07-fc1 vsan 1
   member pwwn 20:00:00:25:b5:02:02:2d
!            [vm-host-vdi-07-fc1]
   member pwwn 50:0a:09:82:9d:42:bd:e8
!            [citrixvdi-a-0d]
   member pwwn 50:0a:09:82:8d:42:bd:e8
!            [citrixvdi-b-0d]

zone name vm-host-vdi-08-fc1 vsan 1
   member pwwn 50:0a:09:82:9d:42:bd:e8
!            [citrixvdi-a-0d]
   member pwwn 50:0a:09:82:8d:42:bd:e8
!            [citrixvdi-b-0d]
   member pwwn 20:00:00:25:b5:02:02:2a
!            [vm-host-vdi-08-fc1]

zone name vm-host-vdi-09-fc1 vsan 1
   member pwwn 50:0a:09:82:9d:42:bd:e8
!            [citrixvdi-a-0d]
   member pwwn 50:0a:09:82:8d:42:bd:e8
!            [citrixvdi-b-0d]
   member pwwn 20:00:00:25:b5:02:02:0d
!            [vm-host-vdi-09-fc1]

zone name vm-host-vdi-10-fc1 vsan 1
   member pwwn 50:0a:09:82:9d:42:bd:e8
!            [citrixvdi-a-0d]
   member pwwn 50:0a:09:82:8d:42:bd:e8
!            [citrixvdi-b-0d]
   member pwwn 20:00:00:25:b5:02:02:0a
!            [vm-host-vdi-10-fc1]
```

```
zone name vm-host-vdi-11-fc1 vsan 1
   member pwwn 50:0a:09:82:9d:42:bd:e8
!            [citrixvdi-a-0d]
   member pwwn 50:0a:09:82:8d:42:bd:e8
!            [citrixvdi-b-0d]
   member pwwn 20:00:00:25:b5:02:02:2c
!            [vm-host-vdi-11-fc1]

zone name vm-host-vdi-12-fc1 vsan 1
   member pwwn 20:00:00:25:b5:02:02:39
!            [vm-host-vdi-12-fc1]
   member pwwn 50:0a:09:82:9d:42:bd:e8
!            [citrixvdi-a-0d]
   member pwwn 50:0a:09:82:8d:42:bd:e8
!            [citrixvdi-b-0d]

zone name vm-host-vdi-13-fc1 vsan 1
   member pwwn 50:0a:09:82:8d:42:bd:e8
!            [citrixvdi-b-0d]
   member pwwn 50:0a:09:82:9d:42:bd:e8
!            [citrixvdi-a-0d]
   member pwwn 20:00:00:25:b5:02:02:0c
!            [vm-host-vdi-13-fc1]

zone name vm-host-vdi-14-fc1 vsan 1
   member pwwn 20:00:00:25:b5:02:02:19
!            [vm-host-vdi-14-fc1]
   member pwwn 50:0a:09:82:9d:42:bd:e8
!            [citrixvdi-a-0d]
   member pwwn 50:0a:09:82:8d:42:bd:e8
!            [citrixvdi-b-0d]

zone name vm-host-vdi-01-fc1 vsan 1
   member pwwn 20:00:00:25:b5:02:02:2f
!            [vm-host-vdi-01-fc1]
   member pwwn 50:0a:09:82:9d:42:bd:e8
!            [citrixvdi-a-0d]
```

```
    member pwwn 50:0a:09:82:8d:42:bd:e8
!           [citrixvdi-b-0d]

zone name vm-host-infra-02-fc0 vsan 1
    member pwwn 20:00:00:25:b5:02:02:18
!           [vm-host-infra-02-fc1]
    member pwwn 50:0a:09:82:9d:42:bd:e8
!           [citrixvdi-a-0d]
    member pwwn 50:0a:09:82:8d:42:bd:e8
!           [citrixvdi-b-0d]

zoneset name ms-vdi vsan 1
    member c1-s1-b200m2-fc1
    member hyperV-ucs-2-fc1
    member vm-host-infra-01-fc1
    member vm-host-vdi-02-fc1
    member vm-host-vdi-03-fc1
    member vm-host-vdi-04-fc1
    member vm-host-vdi-05-fc1
    member vm-host-vdi-06-fc1
    member vm-host-vdi-07-fc1
    member vm-host-vdi-08-fc1
    member vm-host-vdi-09-fc1
    member vm-host-vdi-10-fc1
    member vm-host-vdi-11-fc1
    member vm-host-vdi-12-fc1
    member vm-host-vdi-13-fc1
    member vm-host-vdi-14-fc1
    member vm-host-vdi-01-fc1
    member vm-host-infra-02-fc0

zoneset activate name ms-vdi vsan 1
```

Appendix B—Virtual Machine Manager PowerShell Scripts

Generate Virtual Machines

```
# Purpose:    Generate up to virtual machines using the command-line parameters supplied
#             for customization of the new virtual machine.
# Date:       29 August 2011
# Version:    7
# Author:     Paul Wilson
# Modify by:  Duy Le
# Notes:      The script only creates VMs on a single host. To create VMs on multiple hosts
#             run multiple instances of the script from a batch file or create an outer loop.
#             The LocalVMStorage path and networks must exist or the script fails. I have
#             not added any data validation checks to the script.  This has been modified to work
#          with SCVMM 2012.Must be run from a System Center VMM 2012 Server.

# Add the Virtual Machine Manager snapin so the commands work.
# import-module VirtualMachineManager.psd1

# Set Default Values:
#LOAY - defaults updated

   $DefVMHost = "All"
   $DefVMBaseName = "clus1vm"
   $DefNetworkName = "PVS"
   $DefNetworkName2 = "Public"
   $DefVMPath = "V:"
   $DefVMCount = 5
   $DefStartCount = 1
   $DefMemMin = 512
   $DefMemMax = 2048
   $DefNumCPUs = 1
   $DefMemBuffer = 20
   $DefMemWeight = 5000
   $DefDomainUser = "CITRIX\Administrator"



# Parse the command-line and verify the 13 required parameters are present, if not display usage info
```

```
if ($args -eq $null -or $args.Count -lt 13)
{
   write-host "Usage: GenVMs.ps1 VMTargetHost VMBaseName PVSBootNetwork SynthNetwork "
   write-host "LocalVMStoragePath NumberToCreate StartingAt MinMemory MaxMemory"
   write-host "CpuCores DynamicMemoryBuffer DynamicMemoryWeight DomainUser"
   write-host " "
   write-host "Example: .\GenVMs.ps1 ""$DefVMHost"" ""$DefVMBaseName""
""$DefNetworkName"" ""$DefNetworkName2"" "
   write-host """$DefVMPath"" $DefVMCount $DefStartCount $DefMemMin $DefMemMax
$DefNumCPUs $DefMemBuffer $DefMemWeight $DefDomainUser"
   write-host " "
   write-host "Warning! Not enough command-line parameters have been supplied!"
   $strAnswer = read-host "Would you like to manually provide the parameters (y/n)?"
   switch ($strAnswer)
   {
    N {exit 1}
    Y {
      write-host "============================================================"
      write-host " PROVIDE PARAMETER VALUES. CONFIRM IN NEXT STEP"
      write-host " Press [Enter] to accept the value in parenthesis"
      write-host "============================================================"

      $VMHost = read-host "Enter HyperV Host name to create the servers on (eg $DefVMHost)"
      if($VMHost.length -eq 0){ $VMHost = $DefVMHost}

      $VMBaseName = read-host "Enter base name for virtual machines (eg $DefVMBaseName)"
      if($VMBaseName.length -eq 0){ $VMBaseName = $DefVMBaseName}

      $NetworkName = read-host "Enter the Hyper-V network for the emulated adapter (eg
$DefNetworkName)"
      if($NetworkName.length -eq 0){ $NetworkName = $DefNetworkName}

      $NetworkName2 = read-host "Enter the Hyper-V network for the synthetic adapter (eg
$DefNetworkName2)"
      if($NetworkName2.length -eq 0){ $NetworkName2 = $DefNetworkName2}

      $VMPath = read-host "Enter the locally accessible path where the host will store`r`nthe virtual
machines data (eg $DefVMPath)"
      if($VMPath.length -eq 0){ $VMPath = $DefVMPath}
```

```
        [int]$VMCount = read-host "Enter the number of virtual machines to create on the host (eg
$DefVMCount)"
        if($VMCount -eq ""){ [int]$VMCount = $DefVMCount}

      [int]$StartCount = read-host "Enter the first number to start at (eg $DefStartCount)"
        if($StartCount -eq ""){ [int]$StartCount = [int]$DefStartCount}

       $MemMin = read-host "Enter the minimum amount of dynamic memory in MB (eg $DefMemMin)"
        if($MemMin.length -eq 0){ $MemMin = $DefMemMin}

        $MemMax = read-host "Enter the maximum amount of dynamic memory to assign in MB (eg
$DefMemMax)"
        if($MemMax.length -eq 0){ $MemMax = $DefMemMax}

       [int]$NumCPUs = read-host "Enter the number of CPUs to assign to the VM (eg $DefNumCPUs)"
        if($NumCPUs -eq ""){ [int]$NumCPUs = [int]$DefNumCPUs}

        $MemBuffer = read-host "Percentage of memory to use for cache (eg $DefMemBuffer)"
        if($MemBuffer.length -eq 0){ $MemBuffer = $DefMemBuffer}

        $MemWeight = read-host "Enter the memory weight for dynamic memory range is 0-10000 (eg
$DefMemWeight)"
        if($MemWeight.length -eq 0){ $MemWeight = $DefMemWeight}

        $DomainUser = read-host "Enter the domain user for the hardware profile owner (eg
$DefDomainUser)"
        if($DomainUser.length -eq 0){ $DomainUser = $DefDomainUser}

        write-host "Thank you..."
        }
     Default {exit 1}
   }
}
else
{
   # Place the command-line parameters into named variables for later use.
   $VMHost = $args[0]
   $VMBaseName = $args[1]
```

```
    $NetworkName = $args[2]
    $NetworkName2 = $args[3]
    $VMPath = $args[4]
    [int]$VMCount = $args[5]
    [int]$StartCount = $args[6]
    $MemMin = $args[7]
    $MemMax = $args[8]
    [int]$NumCPUs = $args[9]
    $MemBuffer = $args[10]
    $MemWeight = $args[11]
    $DomainUser = $args[12]
}

#LOAY get hosts and sort by name
$VMhosts = Get-SCVMHost | sort Name
$VMhostsCount = $VMhosts.count

# Post back the settings to the user for confirmation
write-host "============================================================="
write-host "CONFIRM CONFIGURED SETTINGS"
write-host "============================================================="
write-host "HyperV Server to create VMs on: $VMHost"
write-host "hosts count: $VMhostsCount"
write-host "Base name for VMs: $VMBaseName"
write-host "PVS boot network name (emulated nic): $NetworkName"
write-host "Normal network name (synthetic nic): $NetworkName2"
write-host "Local path for HyperV server to store VMs: $VMPath"
write-host "Number of VMs to create: $VMCount"
write-host "Base number to start VM creation at: $StartCount"
write-host "Minimum Memory to assign to VM: $MemMin MB"
write-host "Maximum Memory to assign to VM: $MemMax MB"
write-host "Number of CPUs for the VM: $NumCPUs"
write-host "Dynamic Memory buffer: $MemBuffer%"
write-host "Dynamic Memory weight value: $MemWeight"
write-host "Profile Owner: $DomainUser"

write-host "============================================================="
```

```
$strConfirm = read-host "Please confirm these settings. Continue (YES)?"
if ($strConfirm -ne "YES")
 {
  write-host "You did not type out the word YES. Aborting!"
  exit 1
 }

# Get the name of the SCVMM server we are running this on. The VMM server could be passed as a parameter as well.

$VMMServer = Get-SCVMMServer -Computername "localhost"

# Create a new Hardware Profile for a XenDesktop and set the default values or use the existing profile. Updating an existing profile is not supported.
# If the profile already exists and you want to make changes, you will need to change it through PowerShell, SCVMM, or delete and recreate the profile with the script.

   $HWProfile = Get-SCHardwareProfile | where {$_.Name -eq "XD5Profile"}

   if ($HWProfile -eq $null)
   {
      write-output "Hardware profile not found. Creating a default profile."
      $HWProfile = New-SCHardwareProfile -Owner "$DomainUser" -Description "Hosted XenDesktop" -Name "XD5Profile" -CPUCount $NumCPUs -BootOrder PXEBoot,IDEHardDrive,CD,Floppy -DynamicMemoryEnabled $True -DynamicMemoryMaximumMB $MemMax -DynamicMemoryBufferPercentage $MemBuffer -MemoryWeight $MemWeight -MemoryMB $MemMin
   }

# Calculate the ending value for the VM generation loop
$EndCount = $StartCount + $VMCount - 1
$VMsperHost = $VMCount/$VMhostsCount

#LOAY = change loop to double loop
# Create VMs in a loop
for ($i=$StartCount; $i -le $VMsperHost; $i++)
{
for ($j = 0; $j -le ($VMhostsCount-1); $j++)
{
```

```
  # Create the Virtual Machine and assign the VM Name. Use the number after the format type to control the number of leading 0's.
  # Format types: D=Decimal, X=Hexadecimal. ie. D3=(001,002,...,099,100,...999,1000,...n) D2=(01,02...,99,100,...n) X2=(...,0A,OB,...,FF,100,...n)
  # LOAY - code had beend upgraded to handle 4 digit numbers (instead of 3)
  $VMName = "{1}{0:D4}" -f ($j*$VMsperHost + $i), $VMBaseName
  write-host "Creating $VMName..."

  # Create a job group id to link the items together and create them as a group with the New-VM command

  $JobGroupID = [System.Guid]::NewGuid().ToString()

  # Get a network objects for creating the network adapters. If a second network adapter (synthetic usually) comment out the $VNetwork2= line

  $LNetwork = Get-SCLogicalNetwork -Name $NetworkName -VMMServer localhost
  $LNetwork2 = Get-SCLogicalNetwork -Name $NetworkName2 -VMMServer localhost

  New-SCVirtualNetworkAdapter -JobGroup $JobGroupID -LogicalNetwork $LNetwork



  # In case a second synthetic adapter is not necessary comment out the line below

  New-SCVirtualNetworkAdapter -JobGroup $JobGroupID -synthetic -LogicalNetwork $LNetwork2

  # Create a virtual DVD

  New-SCVirtualDVDDrive -JobGroup $JobGroupID -Bus 1 -LUN 0

  # Build Virtual Machine using SCVMM Powershell API.
  # LOAY - code has been upgraded to loop thru the nodes

  New-SCVirtualMachine -VMMServer $VMMServer -Name $VMName -VMHost $VMhosts[$j].computername -Path $VMPath -HardwareProfile $HWProfile -JobGroup $JobGroupID -RunAsynchronously -SkipInstallVirtualizationGuestServices -StartAction NeverAutoTurnOnVM -StopAction TurnOffVM
  }
  write-host "Sleeping 30 secs..."
  sleep 30
```

```
}
```

Modify Virtual Machines

```
# Purpose:    Modify the legacy NIC to use static MAC and assign it one from MAC Address pool
# Date:       08 Nov. 2011
# Version:    1
# Notes:      This script is created for SCVMM 2012.  It sets the first NIC, the Legacy NIC, to use
#             static MAC address and then assign a MAC from the MAC pool.
#           Must be run from a System Center VMM 2012 Server.


# Set Default Values:

   $DefVMBaseName = "PVSTargetVM"


# Parse the command-line and verify the required parameter is present, if not display usage info
if ($args -eq $null -or $args.count -lt 1)
{
   write-output "Usage: ModifyVMs.ps1 VMNameMatch"
   write-output "Example: .\ModifyVMs.ps1  ""PVSTargetVM"" "
   write-output "Function: Change the first to NIC to static MAC and get MAC from the MAC pool of
the VMM."
   write-host "Warning! Not enough command-line parameters have been supplied!"
   $strAnswer = read-host "Would you like to manually provide the parameters (y/n)?"
   switch ($strAnswer)
   {
    N {exit 1}
    Y {
      write-host "============================================================="
      write-host " PROVIDE PARAMETER VALUES. CONFIRM IN NEXT STEP"
      write-host " Press [Enter] to accept the value in parenthesis"
      write-host "============================================================="


      $VMBaseName = read-host "Enter base name for virtual machines (eg $DefVMBaseName)"
      if($VMBaseName.length -eq 0){ $VMBaseName = $DefVMBaseName}
```

```
        }
       Default {exit 1}
     }
}
else
{
# Place the command-line parameters into named variables for later use.

$VMBaseName = $args[0]
}


# Get the list of VMs that match the VMNameMatch provided on the command-line

$AllVMs = Get-SCVirtualMachine | where { $_.Name -match "$VMBaseName" } | sort Name

# Determine how many VM's meet the VMNameMatch criteria. Save the count for later output.

if ($AllVMs -eq $null)
{
write-output "No VMs match the pattern: $VMBaseName"
exit 1
}
else
{
   $LeftToGo = $AllVMs.Count
   if ($LeftToGo -eq $null)
   {
      $matchString = "Only one VM matched the pattern: {0}" -f $VMBaseName
      $LeftToGo = 1
   }
   else
   {
  $matchString = "{0} VMs match the pattern: {1}" -f $AllVMs.Count, $VMBaseName
}
   write-output $matchString
}
```

```
# Get MAC Address Pool name
$MACPool = Get-SCMACAddressPool



# Process each VM and attempt to assign a static MAC Address.
foreach ($myVM in $AllVMs)
{

   try
   {

   # Get Virtual Network adapters from VM
   $NIC = Get-SCVirtualNetworkAdapter -VM $myVM


      #LOAY
   #Grant MAC address
$PooledMACAddress = Grant-SCMACAddress -MacAddressPool $MACPool[1]
-VirtualNetworkAdapter $NIC[0]


#LOAY
   #Set virtual network adapter
   Set-SCVirtualNetworkAdapter -PhysicalAddressType Static -PhysicalAddress $PooledMACAddress
-VirtualNetworkAdapter $NIC[0] -RunAsynchronously
   }
   catch { }


}
```

Attach VHDs to Virtual Machines

```
# Purpose:    This script attaches an existing VHD to a virtual machine. Designed for deploying
XenDesktop and attaching write
#             cache drives to existing VMs.
# Date:       28 Oct 2010
# Authors:    Loay Shbeilat and Paul Wilson (no implied or expressed warranties) with content taken
from Taylor Brown's blog:
```

```
#
http://blogs.msdn.com/b/taylorb/archive/2008/10/13/pdc-teaser-attaching-a-vhd-to-a-virtual-machine.aspx
# Modify by:  Duy Le
# Notes:      This script will add a New IDE Virtual disk drive to attach the VHD and attach a boot ISO if specified.  Must be run from a System Center VMM 2012 Server.

# Add the Virtual Machine Manager snapin so the commands work.
# import-module VirtualMachineManager.psd1

# Function ProcessWMIJob used to add the new Virtual Disk and VHD.

filter ProcessWMIJob
{
   param
   (
      [string]$WmiClassPath = $null,

      [string]$MethodName = $null
   )

   $errorCode = 0

   if ($_.ReturnValue -eq 4096)
   {
      $Job = [WMI]$_.Job

      while ($Job.JobState -eq 4)
      {
         Write-Progress $Job.Caption "% Complete" -PercentComplete $Job.PercentComplete
         Start-Sleep -seconds 1
         $Job.PSBase.Get()
      }
      if ($Job.JobState -ne 7)
      {
         if ($Job.ErrorDescription -ne "")
         {
            Write-Error $Job.ErrorDescription
```

```
            Throw $Job.ErrorDescription
        }
        else
        {
            $errorCode = $Job.ErrorCode
        }
    }
    Write-Progress $Job.Caption "Completed" -Completed $TRUE
}
elseif($_.ReturnValue -ne 0)
{
    $errorCode = $_.ReturnValue
}

if ($errorCode -ne 0)
{
    Write-Error "Hyper-V WMI Job Failed!"
    if ($WmiClassPath -and $MethodName)
    {
        $psWmiClass = [WmiClass]$WmiClassPath
        $psWmiClass.PSBase.Options.UseAmendedQualifiers = $TRUE
        $MethodQualifiers = $psWmiClass.PSBase.Methods[$MethodName].Qualifiers
        $indexOfError = [System.Array]::IndexOf($MethodQualifiers["ValueMap"].Value,
[string]$errorCode)
        if ($indexOfError -ne "-1")
        {
            Throw "ReturnCode: ", $errorCode, " ErrorMessage: '",
$MethodQualifiers["Values"].Value[$indexOfError], "' - when calling $MethodName"
        }
        else
        {
            Throw "ReturnCode: ", $errorCode, " ErrorMessage: 'MessageNotFound' - when calling
$MethodName"
        }
    }
    else
    {
        Throw "ReturnCode: ", $errorCode, "When calling $MethodName - for rich error messages
provide classpath and method name."
```

```
        }
    }
    return $_
}

# Parse the command-line and verify the 3 required parameters are present, if not display usage info
if ($args -eq $null -or $args.Count -lt 3)
{
    write-output "Usage: AttachVHD.ps1 LocalVMStoragePath VMNameMatch Postpend"
    write-output "Example: .\AttachVHD.ps1 ""E:\Hyper-V"" ""HVDesktop01"" ""_wc"" "
    write-output "Function: Adds a IDE drive and attachs an existing VHD to the VM."
    write-output "In this example the E:\Hyper-V\HVDesktop01\HVDesktop01_wc.vhd is attached
HVDesktop01"
    exit 1
}

# Place the command-line parameters into named variables for later use.

$VHDPath = $args[0]
$VMNameMatches = $args[1]
$PostPend = $args[2]

# Get the VMM server name

$VMMServer = Get-SCVMMServer -Computername "localhost"

# Get the list of VMs that match the VMNameMatch provided on the command-line
# LOAY - took away the sorting (this causes the script to go slower)
$AllVMs = Get-SCVirtualMachine | where { $_.Name -match "$VMNameMatches" }

# Determine how many VM's meet the VMNameMatch criteria. Save the count for later output.

if ($AllVMs -eq $null)
{
write-output "No VMs match the pattern: $VMNameMatches"
exit 1
}
else
```

```
{
   $LeftToGo = $AllVMs.Count
   if ($LeftToGo -eq $null)
   {
      $matchString = "Only one VM matched the pattern: {0}" -f $VMNameMatches
      $LeftToGo = 1
   }
   else
   {
  $matchString = "{0} VMs match the pattern: {1}" -f $AllVMs.Count, $VMNameMatches
}
   write-output $matchString
}

# Process each VM and attempt to mount the VHD. The VHD needs to exist first.
#LOAY - added counter for throttling.
$counter = 0
foreach ($myVM in $AllVMs)
{
   $counter++
   if ($counter -eq 20)
   {
      $counter = 0
      Sleep 5
   }
   $LeftToGo = $LeftToGo - 1
   $HyperVGuest = $myVM.Name
   $server = $myVM.hostname

   # Modify $vhdToMount variable to match the path to the VHD if yours is not in the VM directory.

   $vhdToMount = "{1}{2}.vhd" -f $VHDPath, $myVM.Name, $PostPend
   $vhdloc = "{0}\{1}" -f $VHDPath, $myVM.Name, $PostPend

  $Status = "Processing VM:{0} VHD:{1} VMs Left:{2}" -f $myVM.Name, $vhdToMount, $LeftToGo
   Write-output $Status

   try
```

```
   {
##LOAY - updated the attach to be Async... Also commented the DVD mounting.
 #NOT NEEDED:  $vm = Get-SCVirtualMachine -name $myVM.Name
$jbgrp = [GUID]::NewGUID().ToString()

New-SCVirtualDiskDrive -IDE -Bus 0 -LUN 0 -UseLocalVirtualHardDisk -FileName $VhdToMount
-Path $vhdloc -jobgroup $jbgrp
set-scvirtualmachine -vm $myVM -jobgroup $jbgrp -RunAsynchronously
 # Get DVD Drive attached to the VM and then mount the ISO file specified on the command line

 #$DVDDrive = Get-SCVirtualDVDDrive -VM $vm | where { $_.Bus -eq 1 -and $_.LUN -eq 0 }
 #Set-SCVirtualDVDDrive -VirtualDVDDrive $DVDDrive -ISO $ISO[0] -runasychronously
   }
   catch {
write-output "Exception"
 }

}
```

Generate PVS Machine File

```
# Purpose:      Create a CSV file that can be imported by Provisioning services
# Date:         29 August 2011
# Version:      1.03
# Author:       Paul Wilson
# Notes:        The CSV file may need to opened and saved in the ANSI format on some computers.
#               This script automatically appends information to any existing file in case you need to
#               run the command multiple times with different match criteria.


# Parse the command-line and verify the five required parameters are present, if not display usage info
if ($args -eq $null -or $args.Count -lt 5)
{
   write-output "Usage: GenPVSFile.ps1 SiteName CollectionName Description ImportFileName
VMMatchCriteria"
   write-output "Example: .\GenPVSFile.ps1 ""Site"" ""Collection"" ""XD Desktop""
""c:\PVSImport.csv"" HVDesktop "
   exit 1
}
```

```
# Pulls the VM Name Match criteria off the command-line
$VMNameMatches = $args[4]

# Connects to the local SCVMM Server
$VMMServer = Get-SCVMMServer -Computername "localhost"

# Finds all matching VMs and sorts by their machine name
$AllVMs = Get-SCVirtualMachine | where { $_.Name -match "$VMNameMatches" } | sort Name

if ($AllVMs -eq $null)
{
write-output "No VMs match the pattern: $VMNameMatches"
exit 1
}
else
{
   $LeftToGo = $AllVMs.Count
   if ($LeftToGo -eq $null)
   {
      $matchString = "Only one VM matched the pattern: {0}" -f $VMNameMatches
      $LeftToGo = 1
   }
   else
   {
  $matchString = "{0}VMs match the pattern: {1}" -f $VMs.Count, $VMNameMatches
}
   write-output $matchString
}



# The following loop gets the MAC address of the primary NIC then writes
# that output to the CSV file along with the other fields required for the PVS import
# most of which were supplied as parameters on the command-line.
# This code assumes the first NIC is the PVS boot NIC. If not, change the $nicDetails[0] to
$nicDetails[1]

foreach ($vm in $AllVms)
{
```

```
    $LeftToGo = $LeftToGo -1
$nicDetails = Get-SCVirtualNetworkAdapter -VM $vm


    # Look to see if more than one network adapter is defined, if so, grab the first one in the array
    if ($nicDetails.count -lt 1)
      # If only one network adapter is defined, this evaluation will return True because no array is created
    {
     $csvString = "{0},{1},{2},{3},{4}" -f $vm.Name, $nicDetails.PhysicalAddress, $args[0], $args[1],
$args[2]
    add-content $args[3] $csvString
       $StatusString = "Processing {0}. Virtual Machines left to process: {1}" -f $vm.Name, $LeftToGo
       write-output $StatusString
    }
    else
     # If that evalution was False, more than one NIC exists and we grab the MAC of the first one in the
array
    {
       $csvString = "{0},{1},{2},{3},{4}" -f $vm.Name, $nicDetails[0].PhysicalAddress, $args[0],
$args[1], $args[2]
    add-content $args[3] $csvString
       $StatusString = "Processing {0}. Virtual Machines left to process: {1}" -f $vm.Name, $LeftToGo
       write-output $StatusString


    }
}
```

Appendix C—NetApp PowerShell Script

Clone Write-Cache Disks for 2000 Virtual desktops.

```
# This script will clone a VHD file to a sub-directory structure already present on a NetApp LUN.
# It will connect to a NetApp controller and then clone the "Source" file to a destination root with
sub-folders
# i.e. source file is c:\clusteredstorage\volume1\wc.vhd and destination is
c:\clusteredstorage\volume1\VM001 etc.
#
# powershell import of NetApp Powershell Toolkit 1.6
```

```
import-module dataontap

# parse the command line and verify the necessary parameters are present.  If not display usage info.

if ($args -eq $null -or $args.Count -lt 6)
{
  write-output "Usage: NACloneWC.ps1 NAController NAUser NAPasswd SourceFile DestinationBase
VMMBase"
   write-output "Example: .\NACloneWC.ps1 10.2.2.2 root password c:\File.vhd d:\ VMMName"
   exit 1
}

$nacontroller = $args[0]
$nauser = $args[1]
$napasswd = $args[2]
$source = $args[3]
$baseDestination = $args[4]
$baseVMName = $args[5]



# provide credential passthrough and then log into NetApp controller

$password = ConvertTo-SecureString $napasswd -AsPlainText -Force
$cred = New-Object -TypeName System.Management.Automation.PSCredential -ArgumentList
$nauser,$password
Connect-NaController $nacontroller -Credential $cred



$Matchingsubfolders = get-childitem $baseDestination -name -include ($baseVMName + "*")

foreach ($subfolder in $Matchingsubfolders)
{
$destination = $baseDestination + "\" + $subfolder  + "\" + $subfolder + "_wc.vhd"
write-output $source
Write-output $destination
copy-nahostfile $source $destination
}
```

Appendix D—VDI Blade Performance Charts: 2000 Users

2000 User, Fourteen Cisco UCS B230 M2 All Phases- M-Host-VDI-01

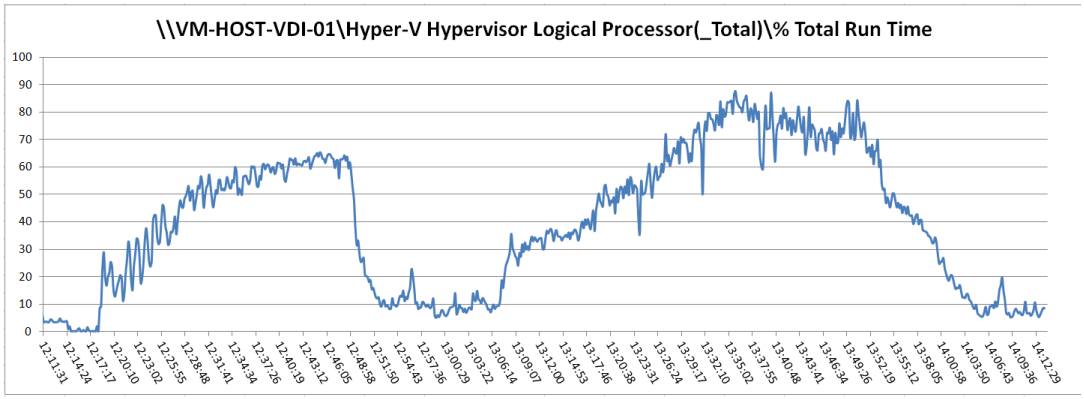

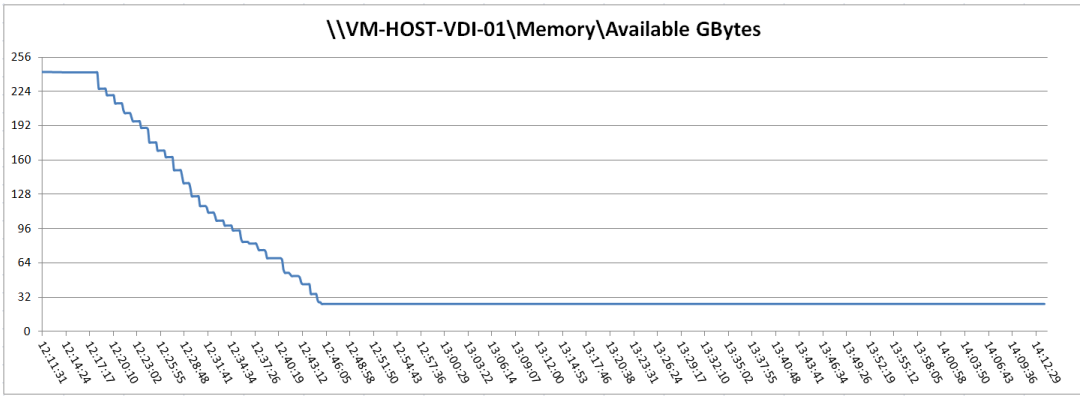

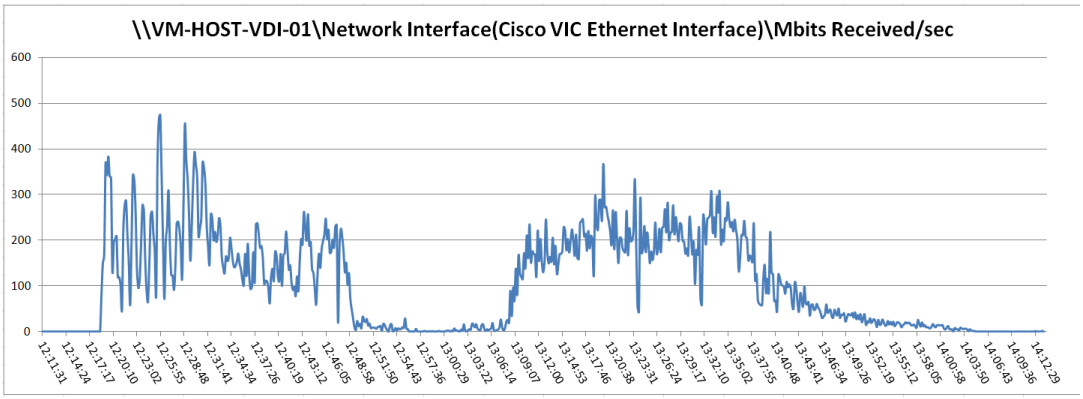

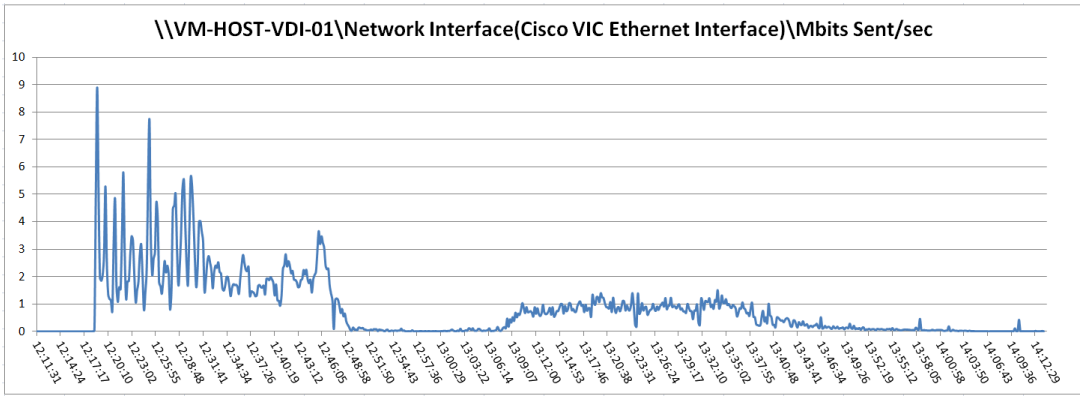

2000 User, Fourteen Cisco UCS B230 M2 All Phases-VM-Host-VDI-02

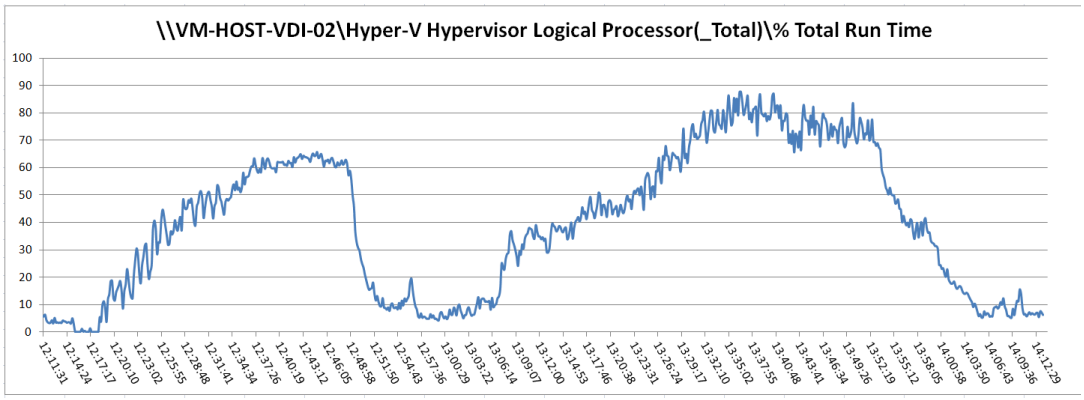

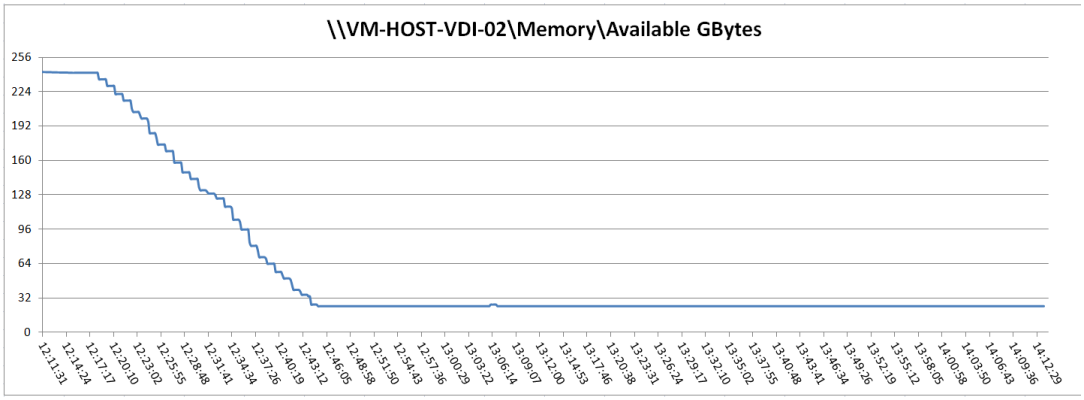

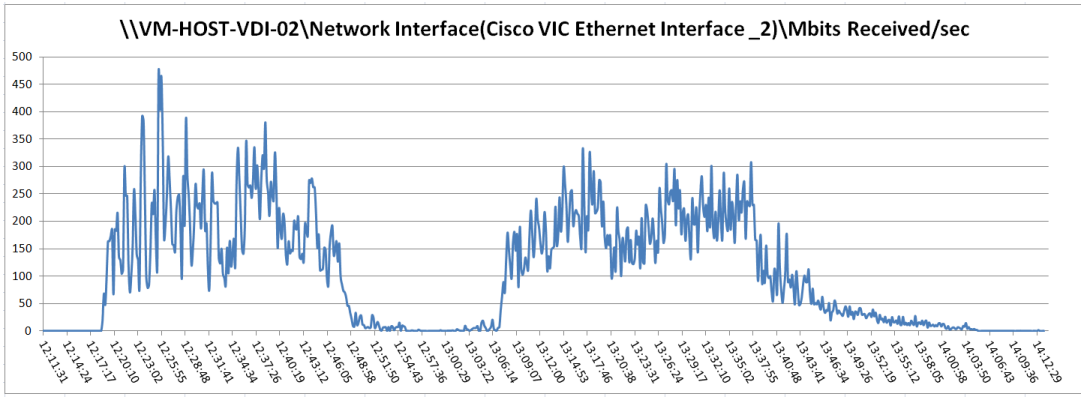
\\VM-HOST-VDI-02\Network Interface(Cisco VIC Ethernet Interface _2)\Mbits Received/sec
500
450
400
350
300
250
200
150
100
50
0
12:11:31
12:14:24
12:17:17
12:20:10
12:23:02
12:25:55
12:28:48
12:31:41
12:34:34
12:37:26
12:40:19
12:43:12
12:46:05
12:48:58
12:51:50
12:54:43
12:57:36
13:00:29
13:03:22
13:06:14
13:09:07
13:12:00
13:14:53
13:17:46
13:20:38
13:23:31
13:26:24
13:29:17
13:32:10
13:35:02
13:37:55
13:40:48
13:43:41
13:46:34
13:49:26
13:52:19
13:55:12
13:58:05
14:00:58
14:03:50
14:06:43
14:09:36
14:12:29

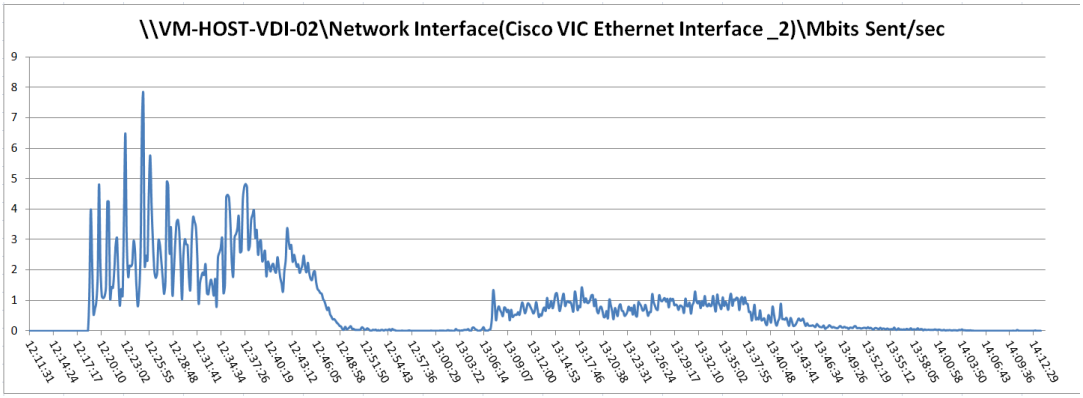
\\VM-HOST-VDI-02\Network Interface(Cisco VIC Ethernet Interface _2)\Mbits Sent/sec
9
8
7
6
5
4
3
2
1
0
12:11:31
12:14:24
12:17:17
12:20:10
12:23:02
12:25:55
12:28:48
12:31:41
12:34:34
12:37:26
12:40:19
12:43:12
12:46:05
12:48:58
12:51:50
12:54:43
12:57:36
13:00:29
13:03:22
13:06:14
13:09:07
13:12:00
13:14:53
13:17:46
13:20:38
13:23:31
13:26:24
13:29:17
13:32:10
13:35:02
13:37:55
13:40:48
13:43:41
13:46:34
13:49:26
13:52:19
13:55:12
13:58:05
14:00:58
14:03:50
14:06:43
14:09:36
14:12:29

2000 User, Fourteen Cisco UCS B230 M2 All Phases-VM-Host-VDI-03

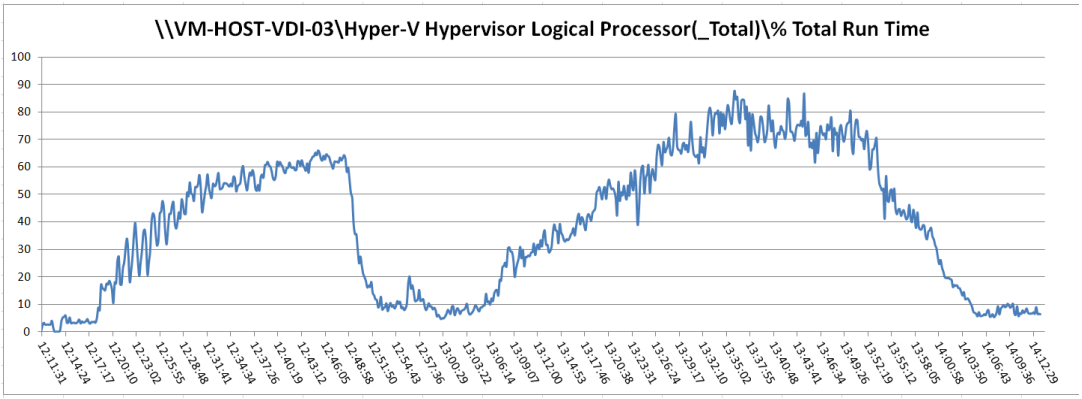

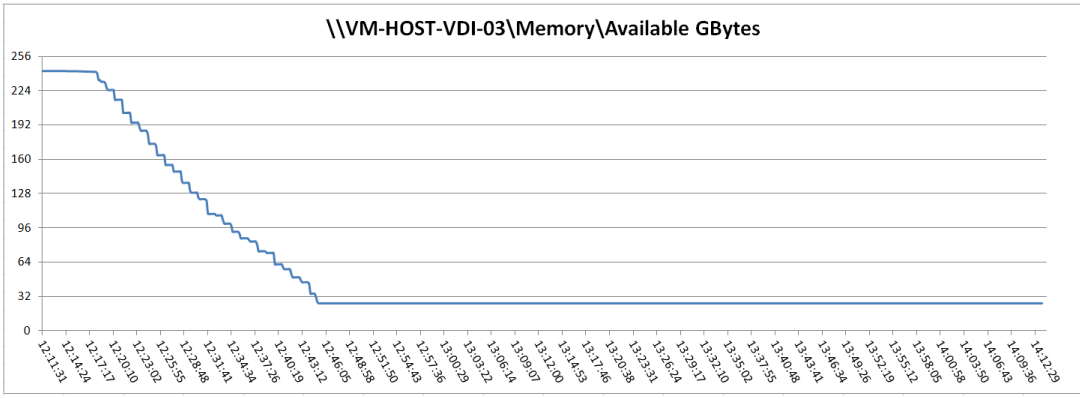

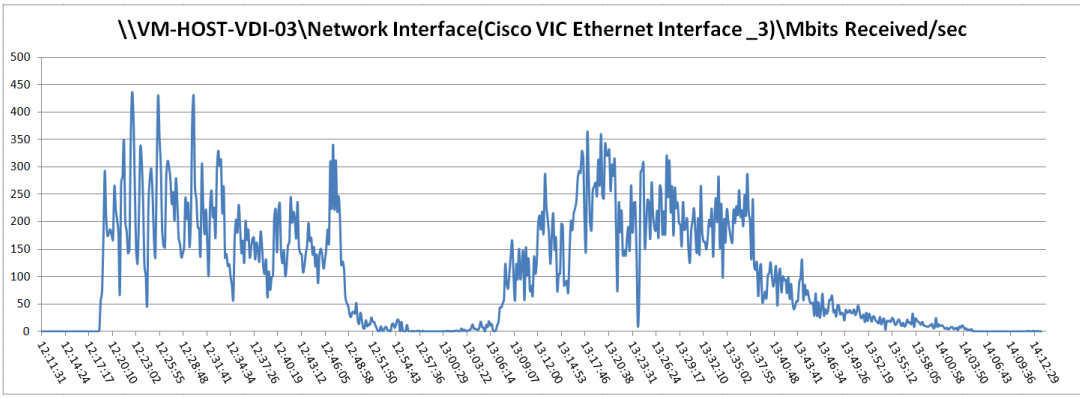

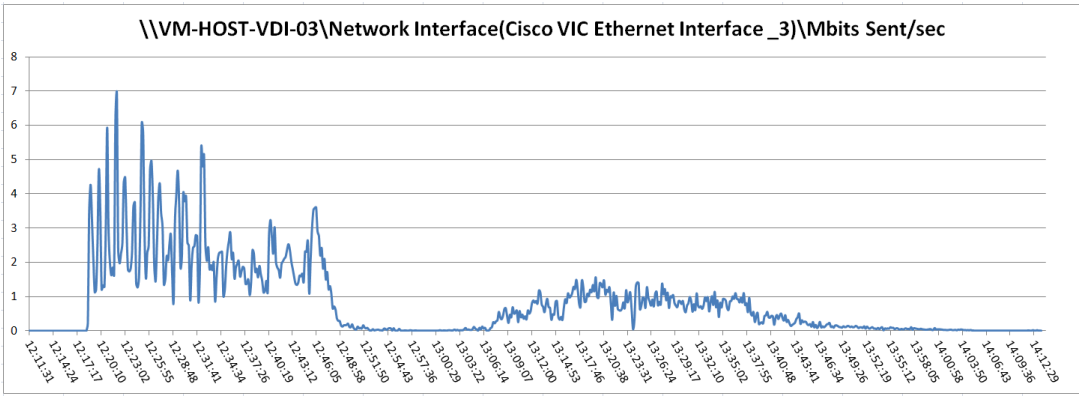

2000 User, Fourteen Cisco UCS B230 M2 All Phases-VM-Host-VDI-04

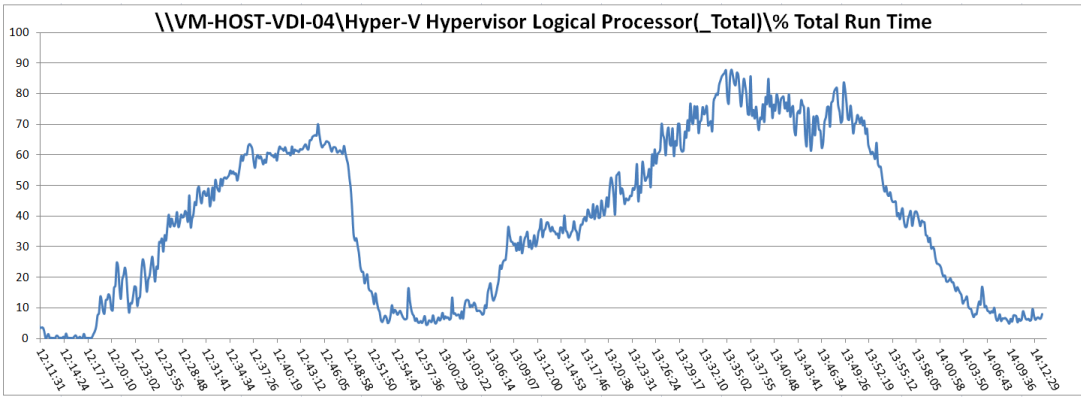

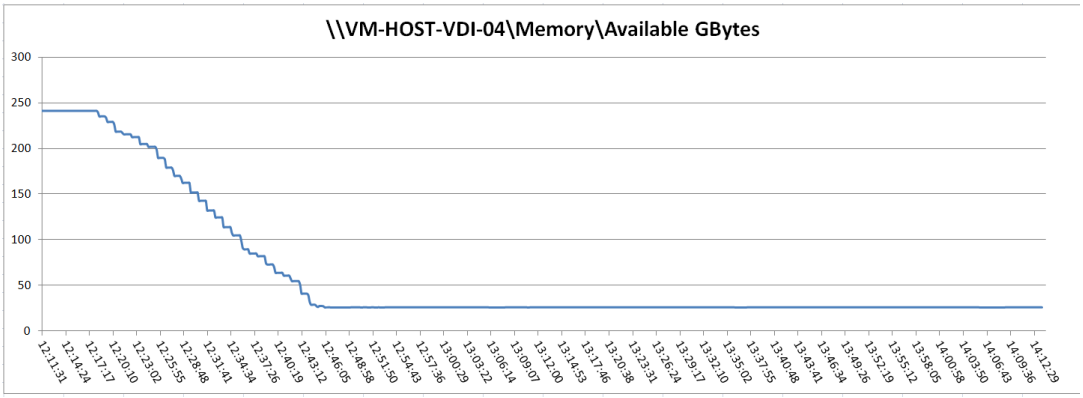

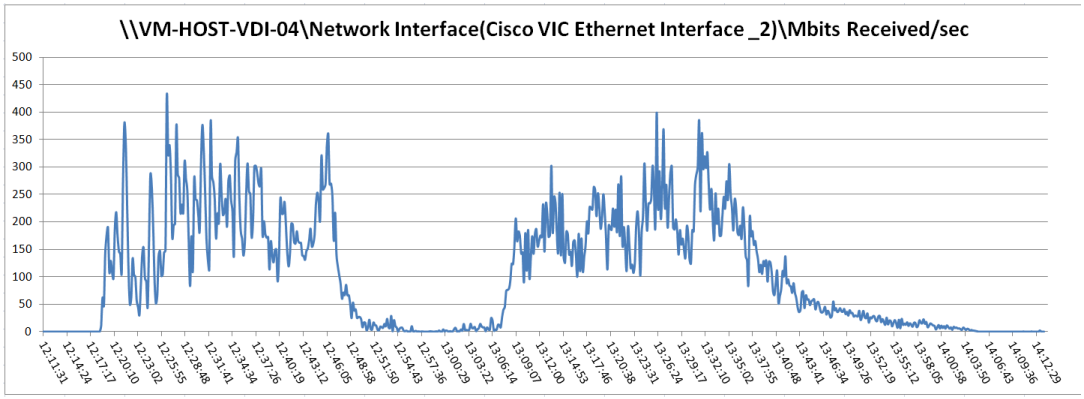
\\VM-HOST-VDI-04\Network Interface(Cisco VIC Ethernet Interface _2)\Mbits Received/sec
500
450
400
350
300
250
200
150
100
50
0
12:11:31
12:14:24
12:17:17
12:20:10
12:23:02
12:25:55
12:28:48
12:31:41
12:34:34
12:37:26
12:40:19
12:43:12
12:46:05
12:48:58
12:51:50
12:54:43
12:57:36
13:00:29
13:03:22
13:06:14
13:09:07
13:12:00
13:14:53
13:17:46
13:20:38
13:23:31
13:26:24
13:29:17
13:32:10
13:35:02
13:37:55
13:40:48
13:43:41
13:46:34
13:49:26
13:52:19
13:55:12
13:58:05
14:00:58
14:03:50
14:06:43
14:09:36
14:12:29

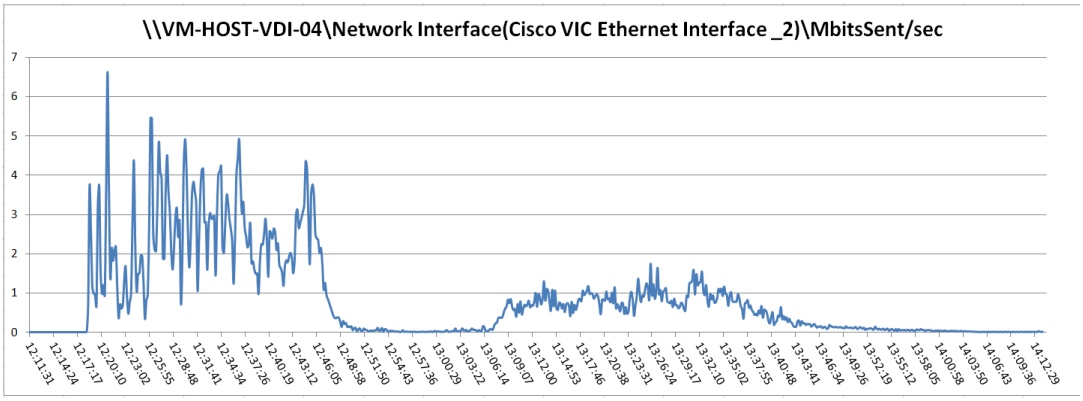
\\VM-HOST-VDI-04\Network Interface(Cisco VIC Ethernet Interface _2)\MbitsSent/sec
7
6
5
4
3
2
1
0
12:11:31
12:14:24
12:17:17
12:20:10
12:23:02
12:25:55
12:28:48
12:31:41
12:34:34
12:37:26
12:40:19
12:43:12
12:46:05
12:48:58
12:51:50
12:54:43
12:57:36
13:00:29
13:03:22
13:06:14
13:09:07
13:12:00
13:14:53
13:17:46
13:20:38
13:23:31
13:26:24
13:29:17
13:32:10
13:35:02
13:37:55
13:40:48
13:43:41
13:46:34
13:49:26
13:52:19
13:55:12
13:58:05
14:00:58
14:03:50
14:06:43
14:09:36
14:12:29

2000 User, Fourteen Cisco UCS B230 M2 All Phases-VM-Host-VDI-05

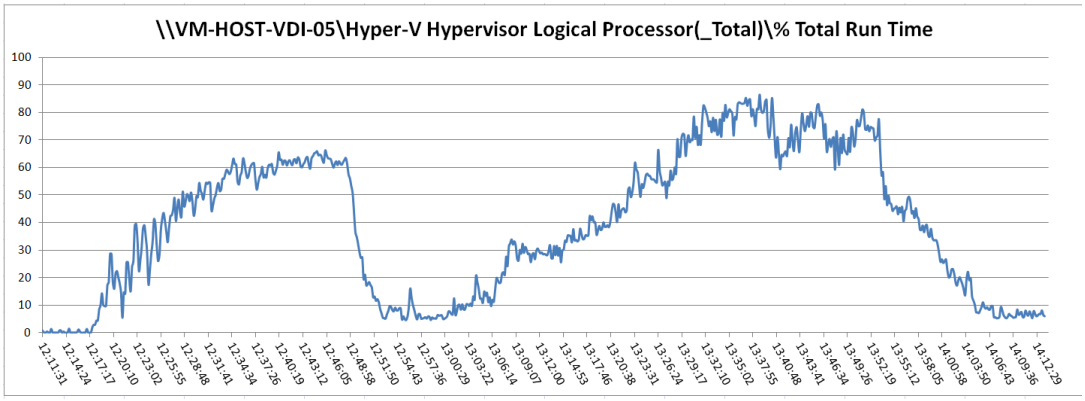

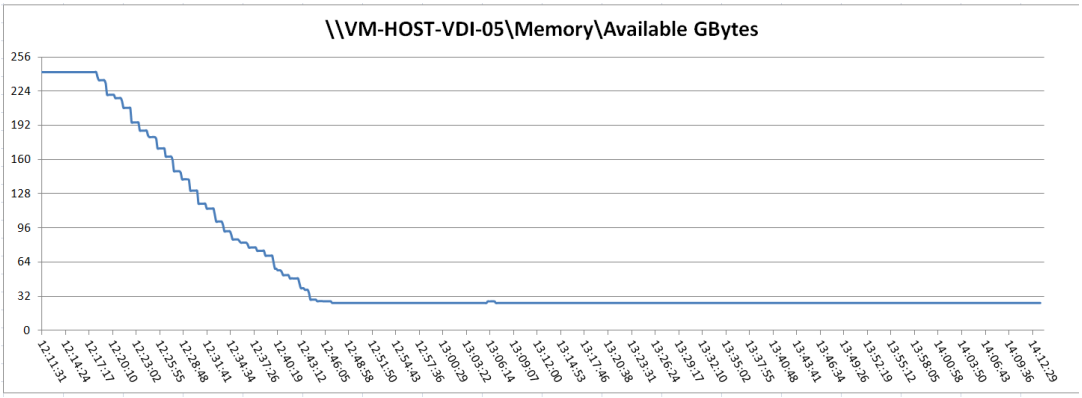

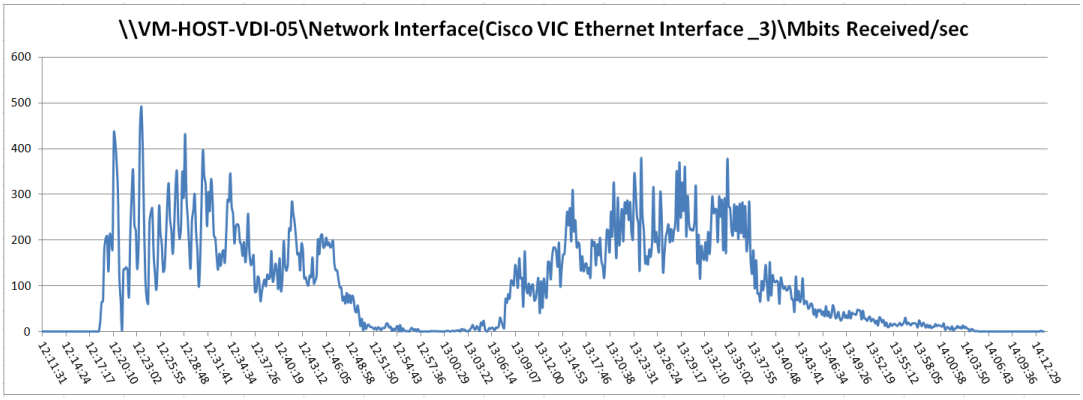

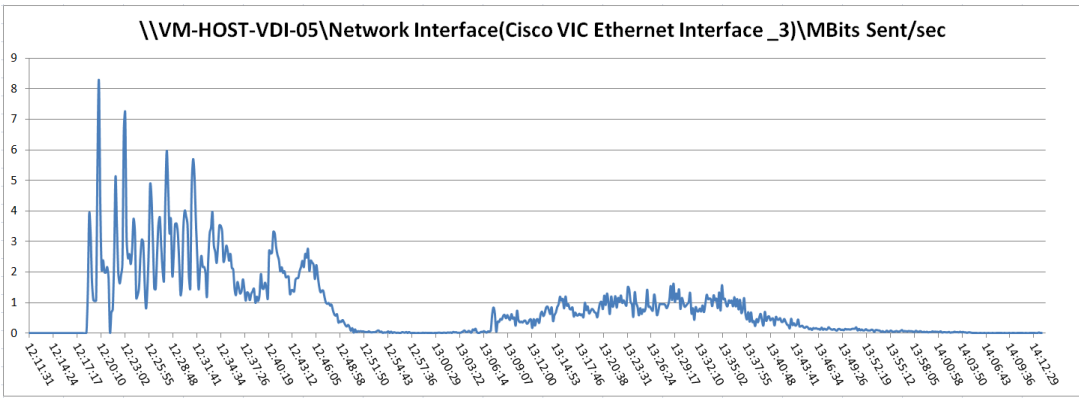

2000 User, Fourteen Cisco UCS B230 M2 All Phases-VM-Host-VDI-06

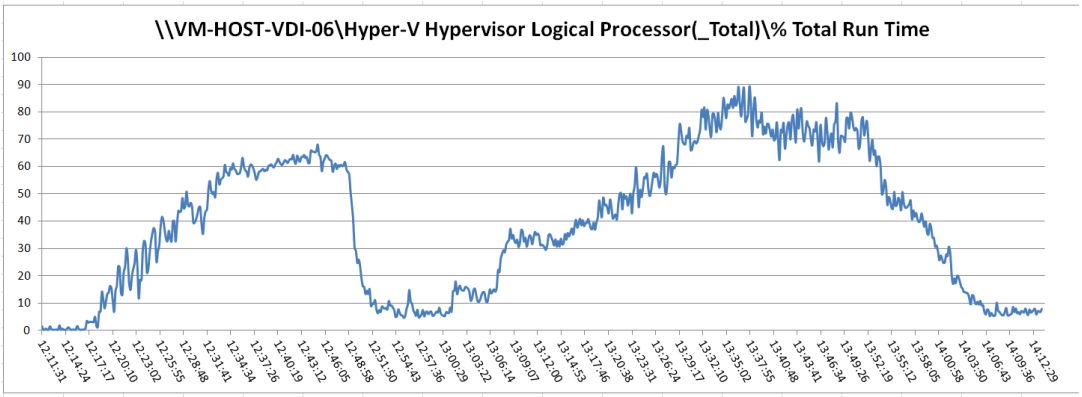

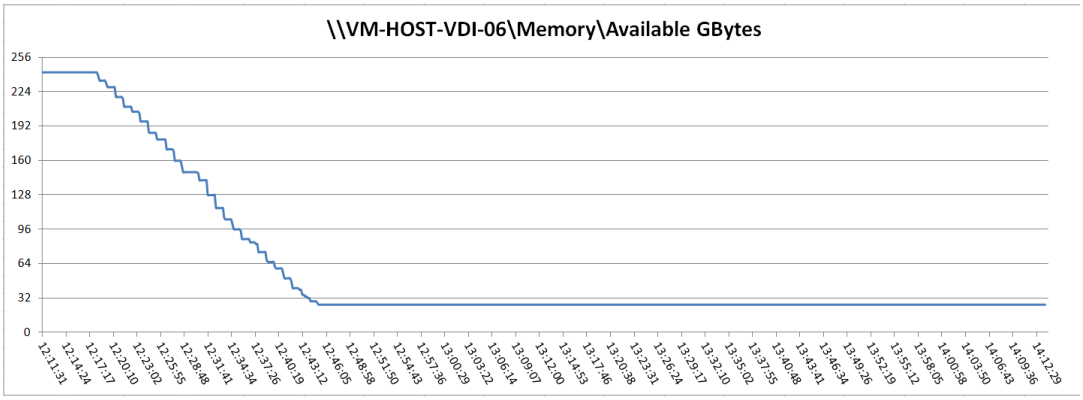

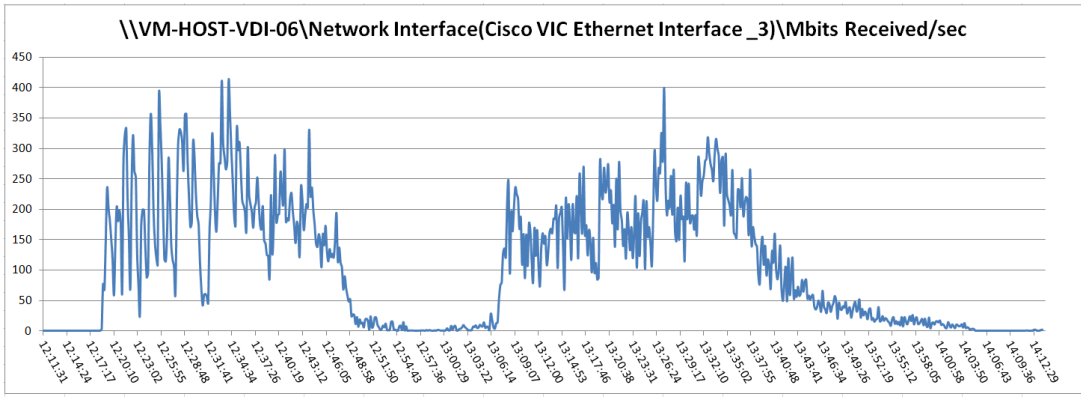
\\VM-HOST-VDI-06\Network Interface(Cisco VIC Ethernet Interface _3)\Mbits Received/sec
450
400
350
300
250
200
150
100
50
0
12:11:31
12:14:24
12:17:17
12:20:10
12:23:02
12:25:55
12:28:48
12:31:41
12:34:34
12:37:26
12:40:19
12:43:12
12:46:05
12:48:58
12:51:50
12:54:43
12:57:36
13:00:29
13:03:22
13:06:14
13:09:07
13:12:00
13:14:53
13:17:46
13:20:38
13:23:31
13:26:24
13:29:17
13:32:10
13:35:02
13:37:55
13:40:48
13:43:41
13:46:34
13:49:26
13:52:19
13:55:12
13:58:05
14:00:58
14:03:50
14:06:43
14:09:36
14:12:29

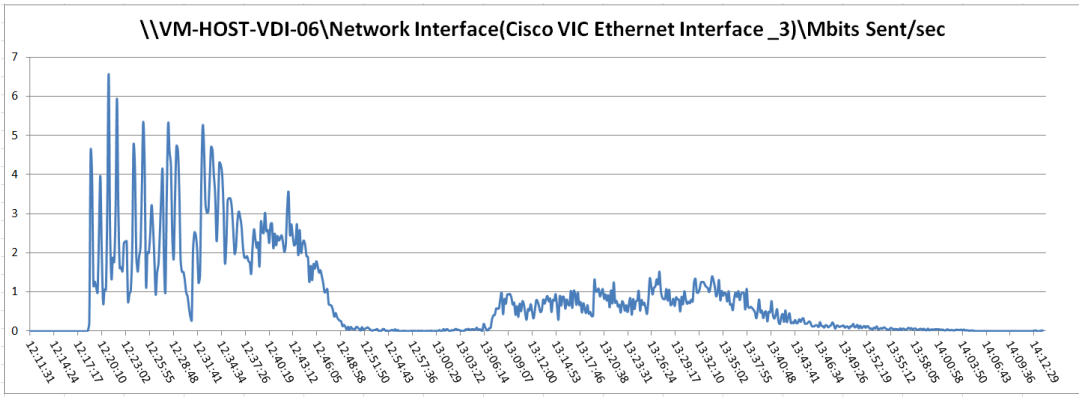
\\VM-HOST-VDI-06\Network Interface(Cisco VIC Ethernet Interface _3)\Mbits Sent/sec
7
6
5
4
3
2
1
0
12:11:31
12:14:24
12:17:17
12:20:10
12:23:02
12:25:55
12:28:48
12:31:41
12:34:34
12:37:26
12:40:19
12:43:12
12:46:05
12:48:58
12:51:50
12:54:43
12:57:36
13:00:29
13:03:22
13:06:14
13:09:07
13:12:00
13:14:53
13:17:46
13:20:38
13:23:31
13:26:24
13:29:17
13:32:10
13:35:02
13:37:55
13:40:48
13:43:41
13:46:34
13:49:26
13:52:19
13:55:12
13:58:05
14:00:58
14:03:50
14:06:43
14:09:36
14:12:29

2000 User, Fourteen Cisco UCS B230 M2 All Phases-VM-Host-VDI-07

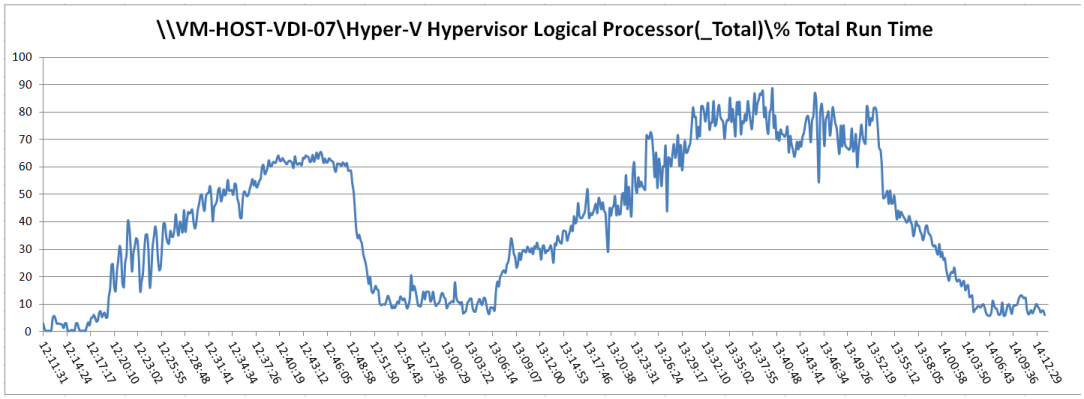

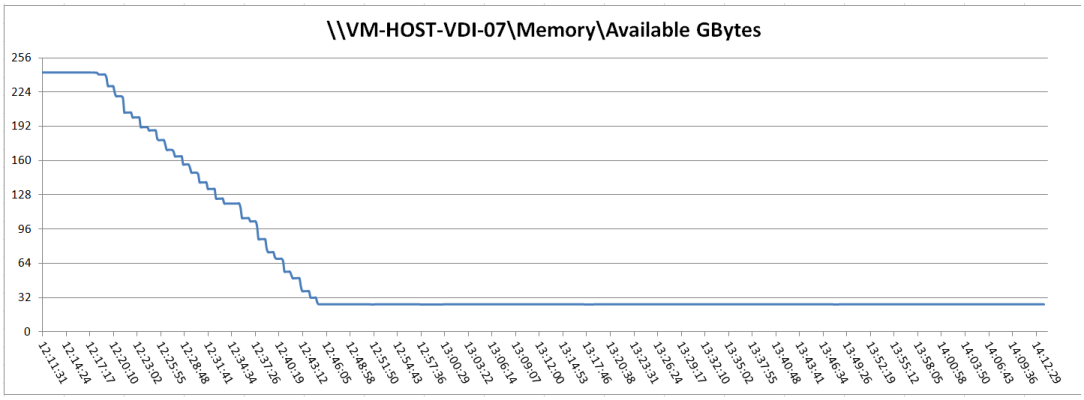

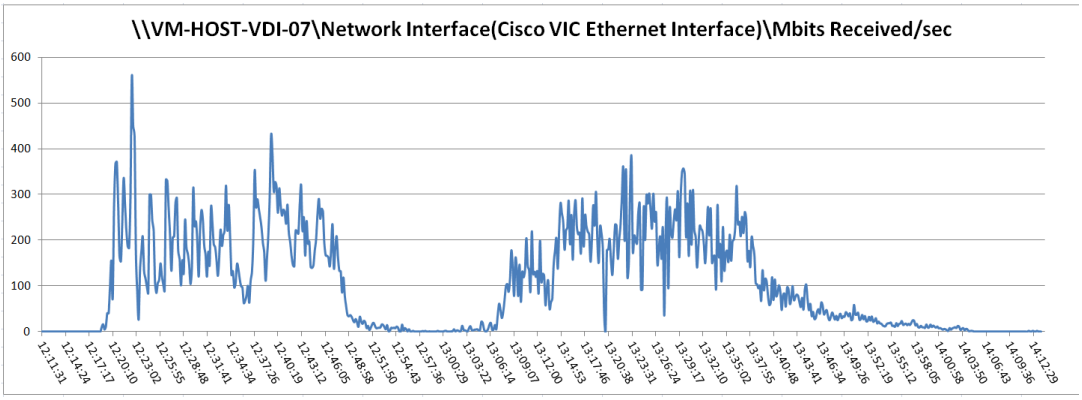

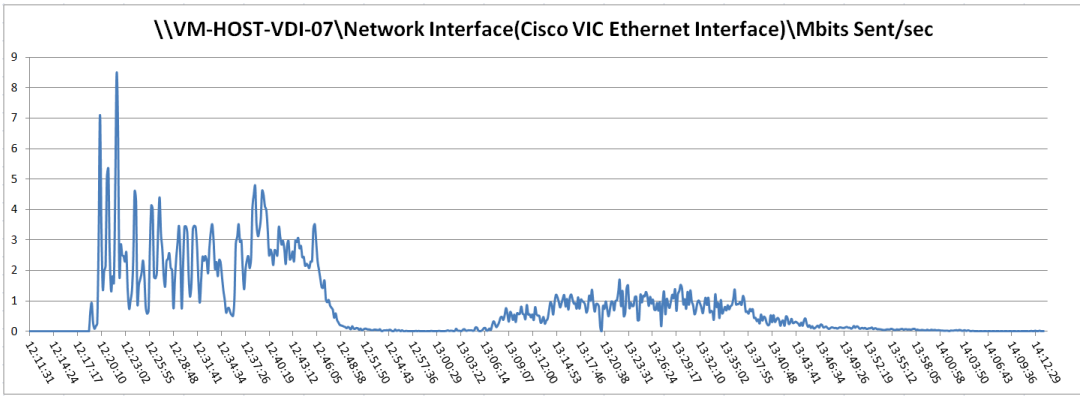

2000 User, Fourteen Cisco UCS B230 M2 All Phases-VM-Host-VDI-08

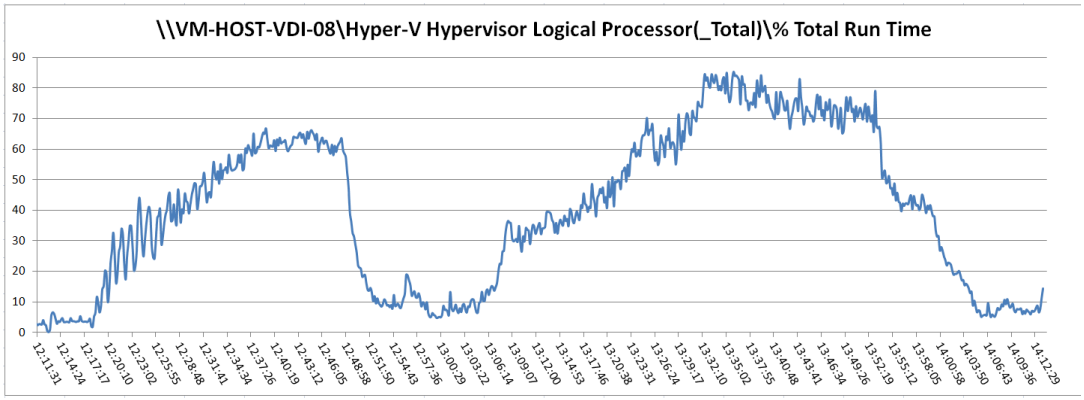

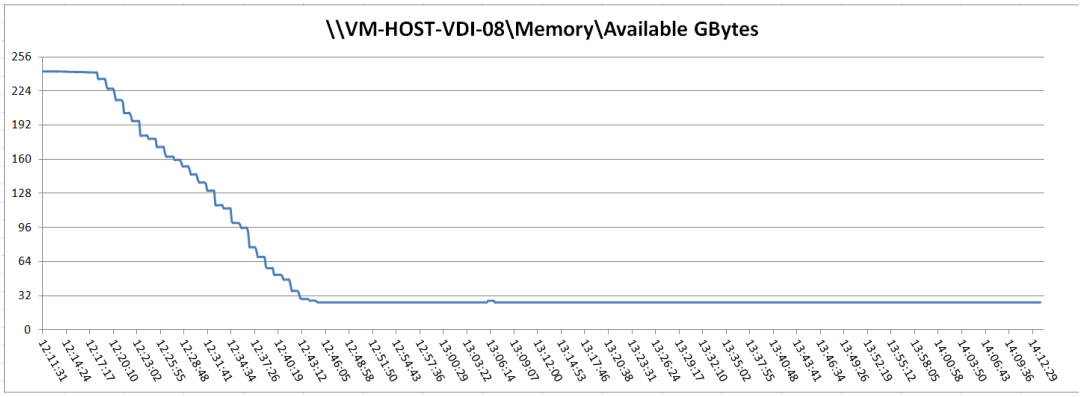

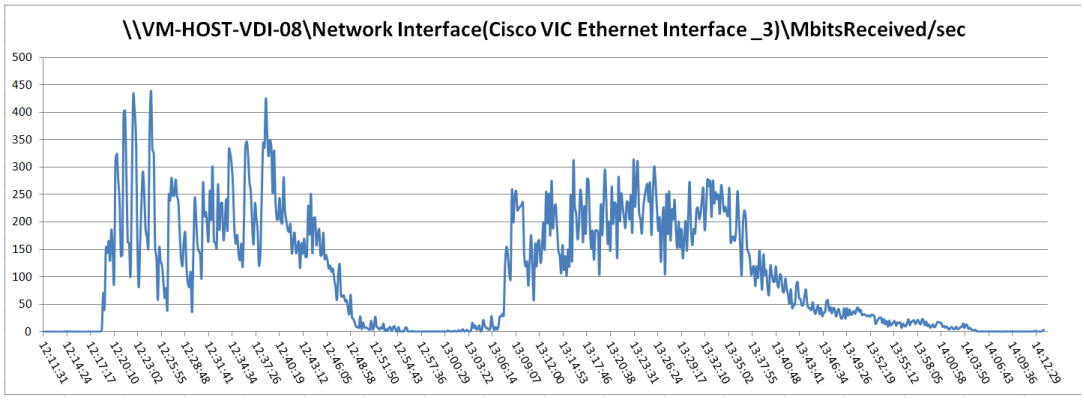
\\VM-HOST-VDI-08\Network Interface(Cisco VIC Ethernet Interface _3)\MbitsReceived/sec
500
450
400
350
300
250
200
150
100
50
0
12:11:31
12:14:24
12:17:17
12:20:10
12:23:02
12:25:55
12:28:48
12:31:41
12:34:34
12:37:26
12:40:19
12:43:12
12:46:05
12:48:58
12:51:50
12:54:43
12:57:36
13:00:29
13:03:22
13:06:14
13:09:07
13:12:00
13:14:53
13:17:46
13:20:38
13:23:31
13:26:24
13:29:17
13:32:10
13:35:02
13:37:55
13:40:48
13:43:41
13:46:34
13:49:26
13:52:19
13:55:12
13:58:05
14:00:58
14:03:50
14:06:43
14:09:36
14:12:29

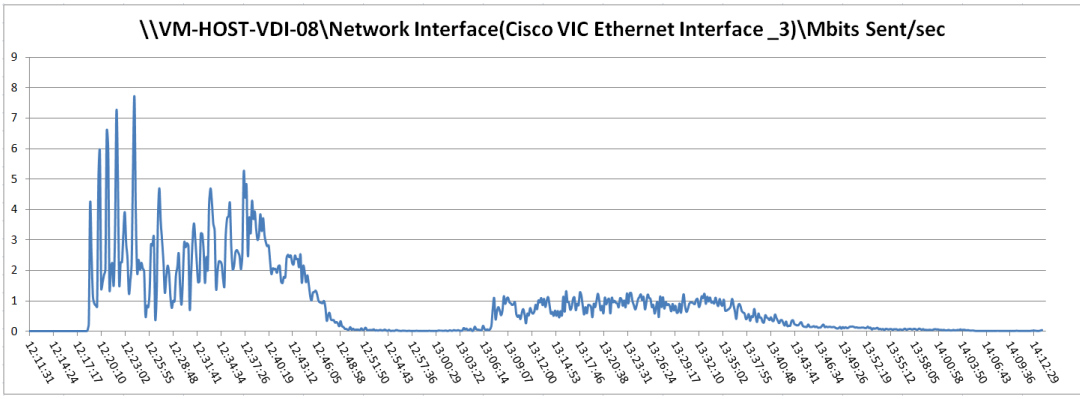
\\VM-HOST-VDI-08\Network Interface(Cisco VIC Ethernet Interface _3)\Mbits Sent/sec
9
8
7
6
5
4
3
2
1
0
12:11:31
12:14:24
12:17:17
12:20:10
12:23:02
12:25:55
12:28:48
12:31:41
12:34:34
12:37:26
12:40:19
12:43:12
12:46:05
12:48:58
12:51:50
12:54:43
12:57:36
13:00:29
13:03:22
13:06:14
13:09:07
13:12:00
13:14:53
13:17:46
13:20:38
13:23:31
13:26:24
13:29:17
13:32:10
13:35:02
13:37:55
13:40:48
13:43:41
13:46:34
13:49:26
13:52:19
13:55:12
13:58:05
14:00:58
14:03:50
14:06:43
14:09:36
14:12:29

2000 User, Fourteen Cisco UCS B230 M2 All Phases-VM-Host-VDI-09

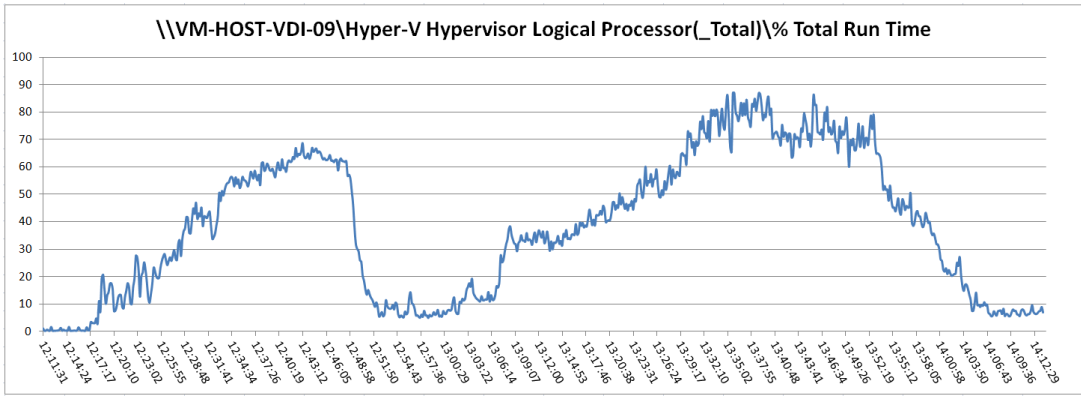

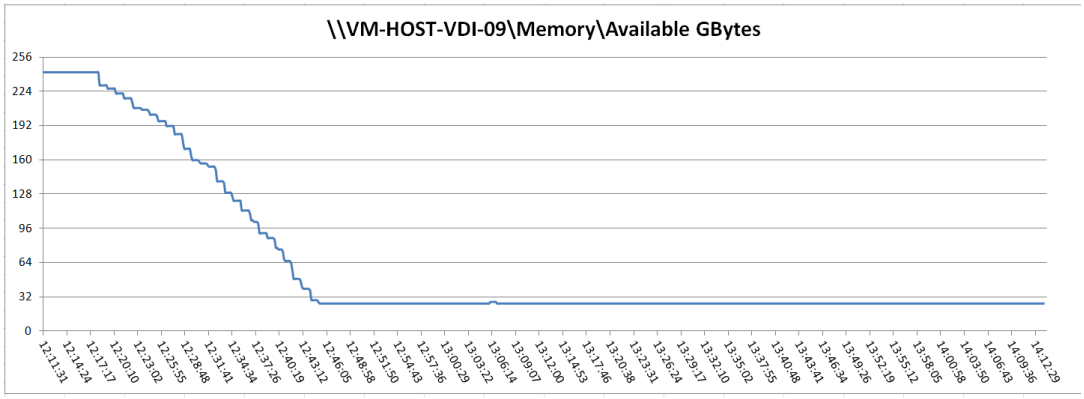

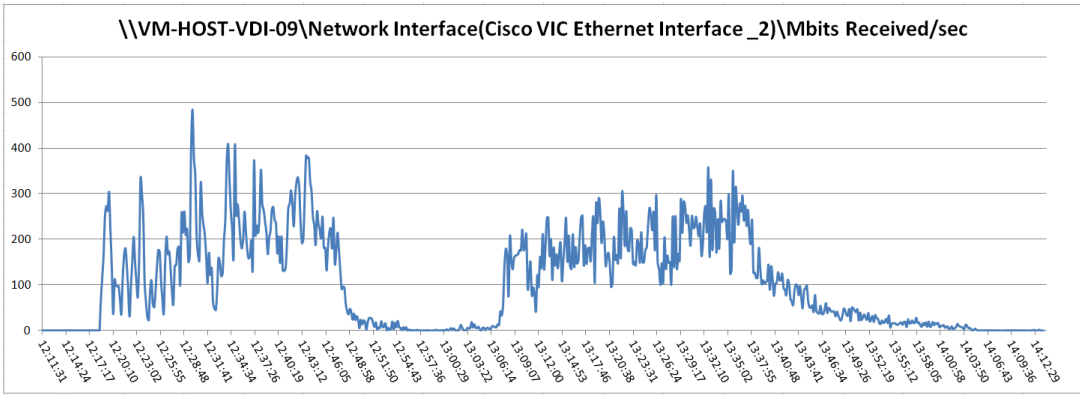

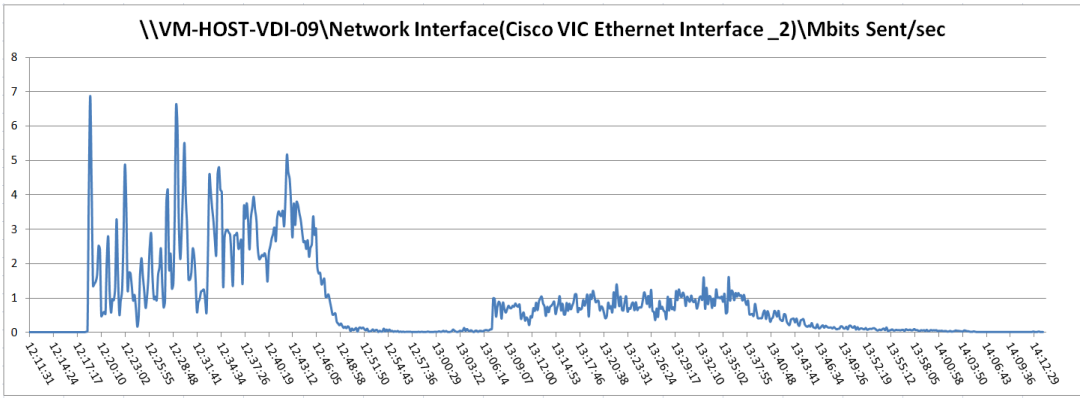

2000 User, Fourteen Cisco UCS B230 M2 All Phases-VM-Host-VDI-10

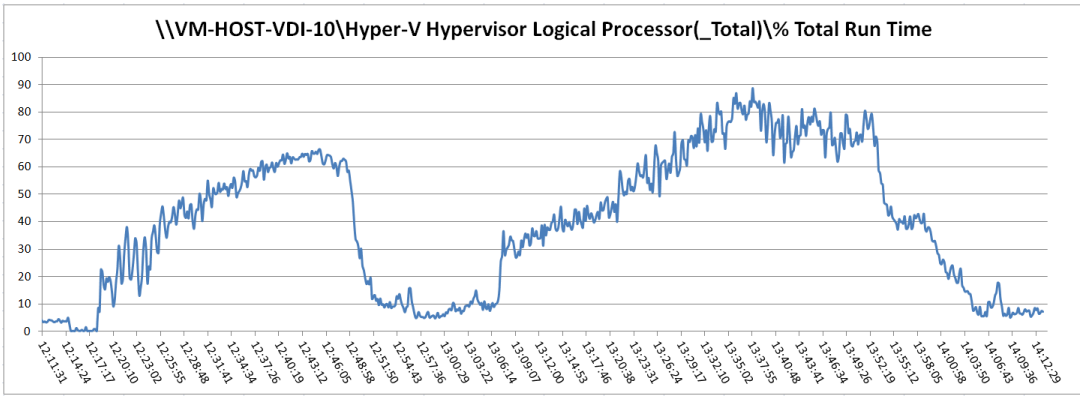

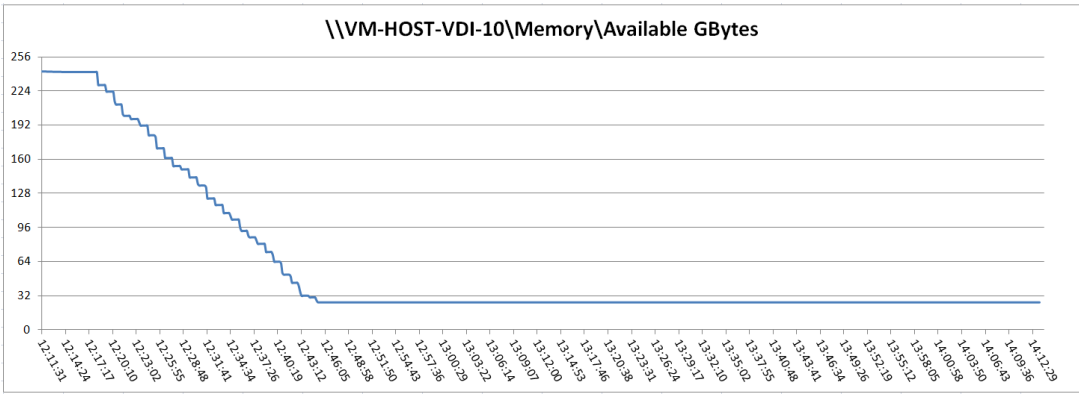

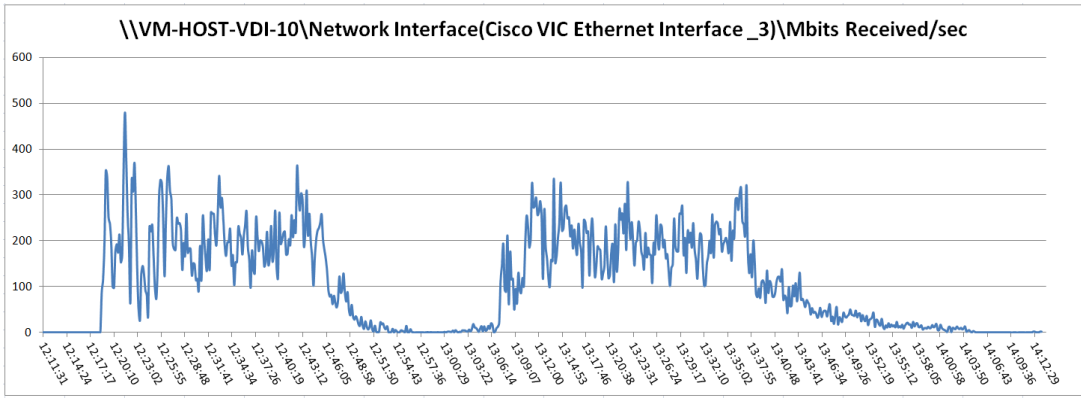
\\VM-HOST-VDI-10\Network Interface(Cisco VIC Ethernet Interface _3)\Mbits Received/sec
600
500
400
300
200
100
0
12:11:31
12:14:24
12:17:17
12:20:10
12:23:02
12:25:55
12:28:48
12:31:41
12:34:34
12:37:26
12:40:19
12:43:12
12:46:05
12:48:58
12:51:50
12:54:43
12:57:36
13:00:29
13:03:22
13:06:14
13:09:07
13:12:00
13:14:53
13:17:46
13:20:38
13:23:31
13:26:24
13:29:17
13:32:10
13:35:02
13:37:55
13:40:48
13:43:41
13:46:34
13:49:26
13:52:19
13:55:12
13:58:05
14:00:58
14:03:50
14:06:43
14:09:36
14:12:29

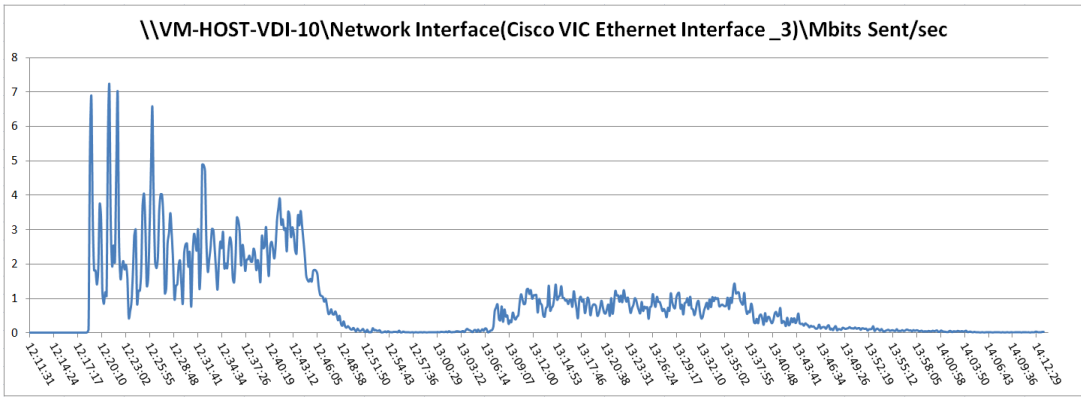
\\VM-HOST-VDI-10\Network Interface(Cisco VIC Ethernet Interface _3)\Mbits Sent/sec
8
7
6
5
4
3
2
1
0
12:11:31
12:14:24
12:17:17
12:20:10
12:23:02
12:25:55
12:28:48
12:31:41
12:34:34
12:37:26
12:40:19
12:43:12
12:46:05
12:48:58
12:51:50
12:54:43
12:57:36
13:00:29
13:03:22
13:06:14
13:09:07
13:12:00
13:14:53
13:17:46
13:20:38
13:23:31
13:26:24
13:29:17
13:32:10
13:35:02
13:37:55
13:40:48
13:43:41
13:46:34
13:49:26
13:52:19
13:55:12
13:58:05
14:00:58
14:03:50
14:06:43
14:09:36
14:12:29

2000 User, Fourteen Cisco UCS B230 M2 All Phases-VM-Host-VDI-11

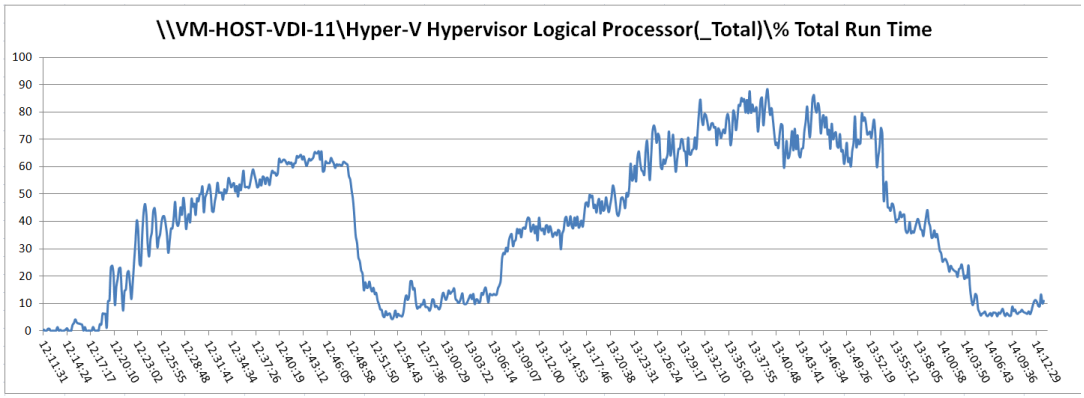

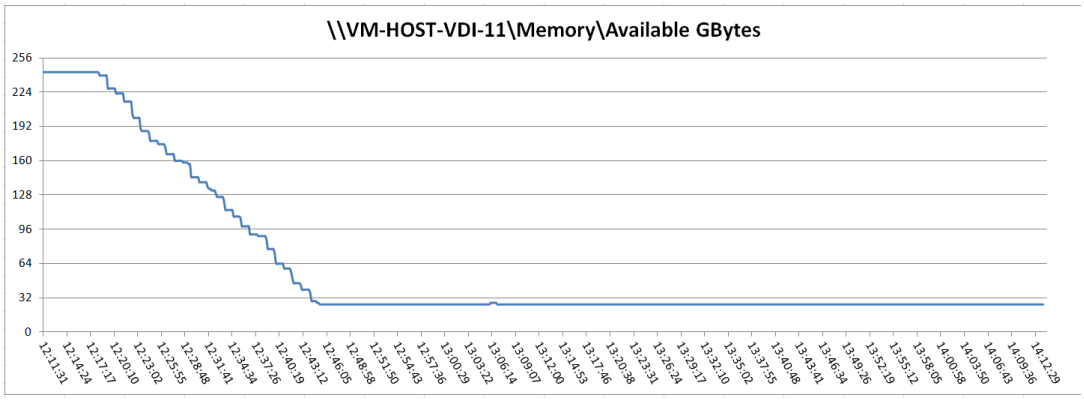

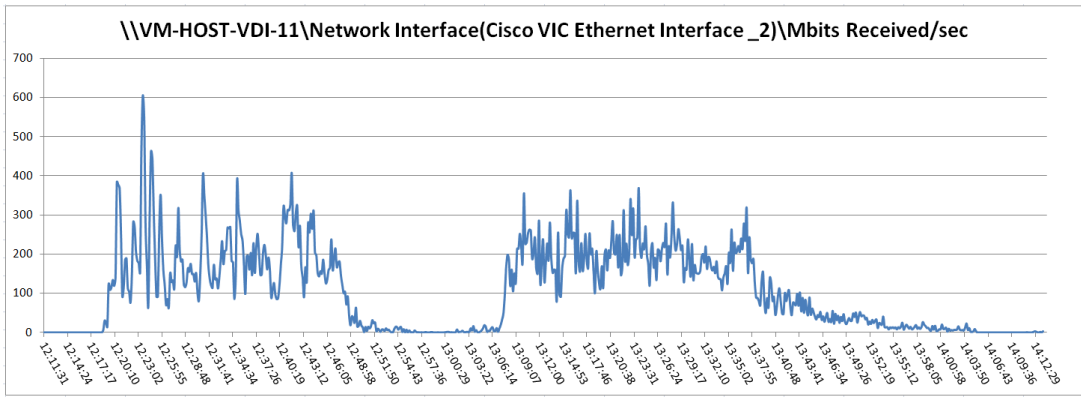

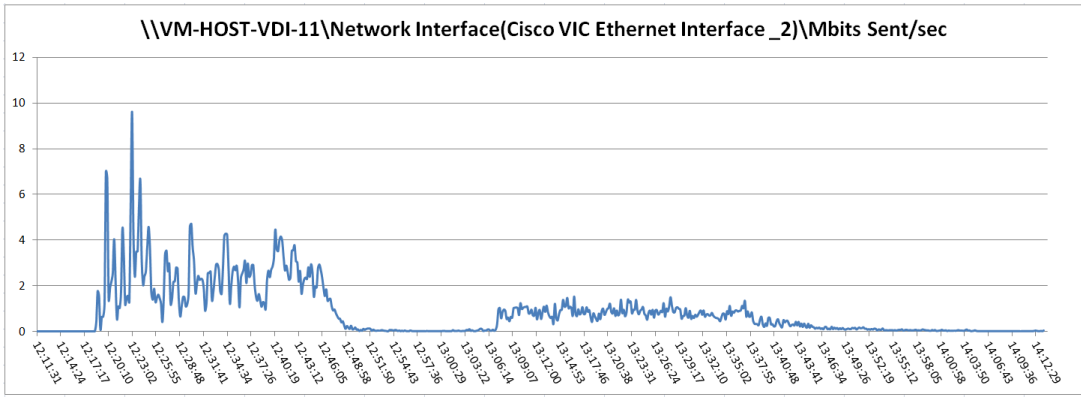

2000 User, Fourteen Cisco UCS B230 M2 All Phases-VM-Host-VDI-12

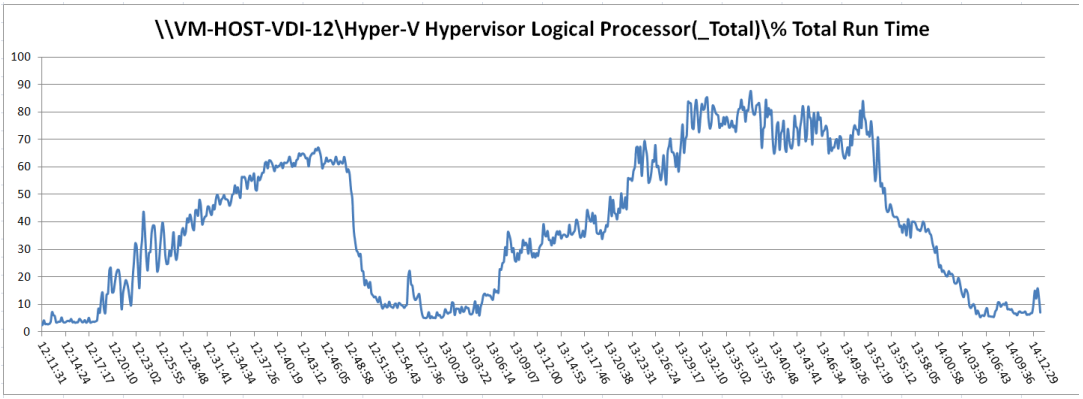

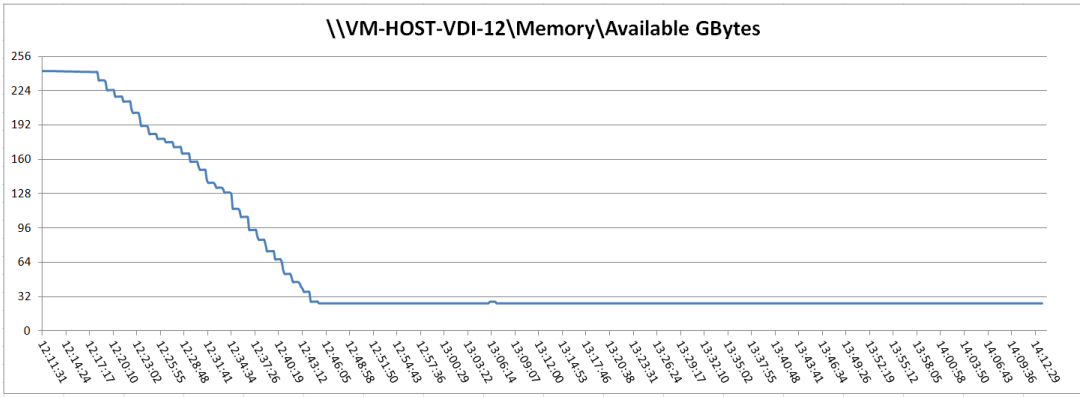

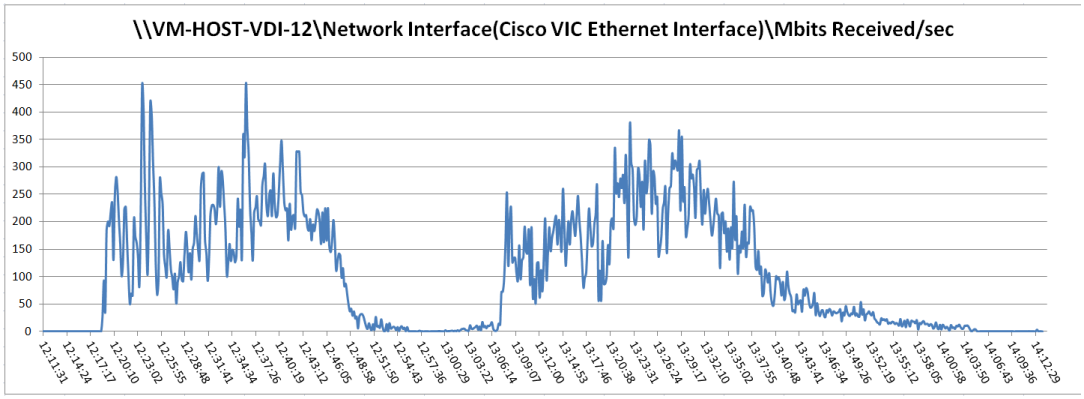

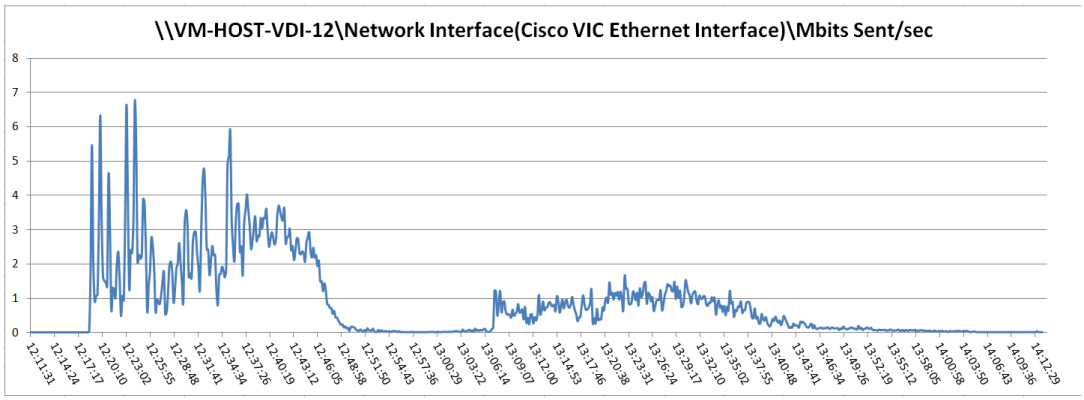

2000 User, Fourteen B230 M2 All Phases- VM-Host-VDI-13

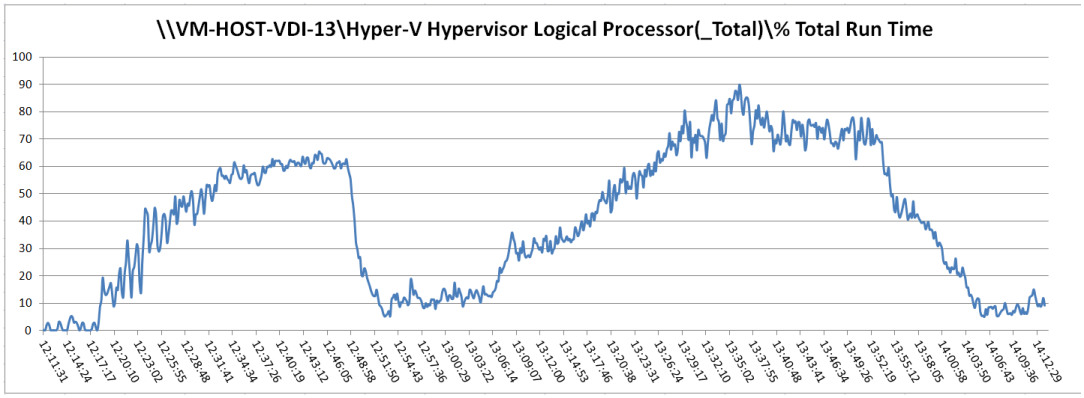

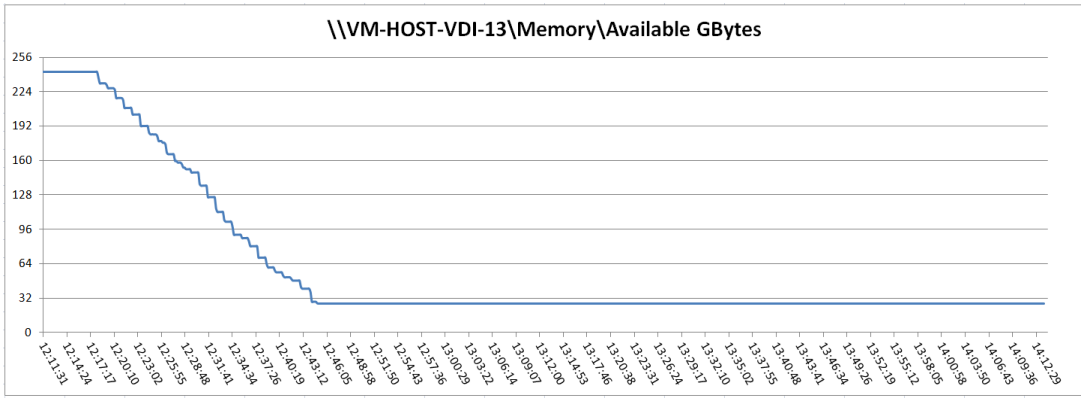
\\VM-HOST-VDI-13\Memory\Available GBytes
256
224
192
160
128
96
64
32
0
12:11:31
12:14:24
12:17:17
12:20:10
12:23:02
12:25:55
12:28:48
12:31:41
12:34:34
12:37:26
12:40:19
12:43:12
12:46:05
12:48:58
12:51:50
12:54:43
12:57:36
13:00:29
13:03:22
13:06:14
13:09:07
13:12:00
13:14:53
13:17:46
13:20:38
13:23:31
13:26:24
13:29:17
13:32:10
13:35:02
13:37:55
13:40:48
13:43:41
13:46:34
13:49:26
13:52:19
13:55:12
13:58:05
14:00:58
14:03:50
14:06:43
14:09:36
14:12:29

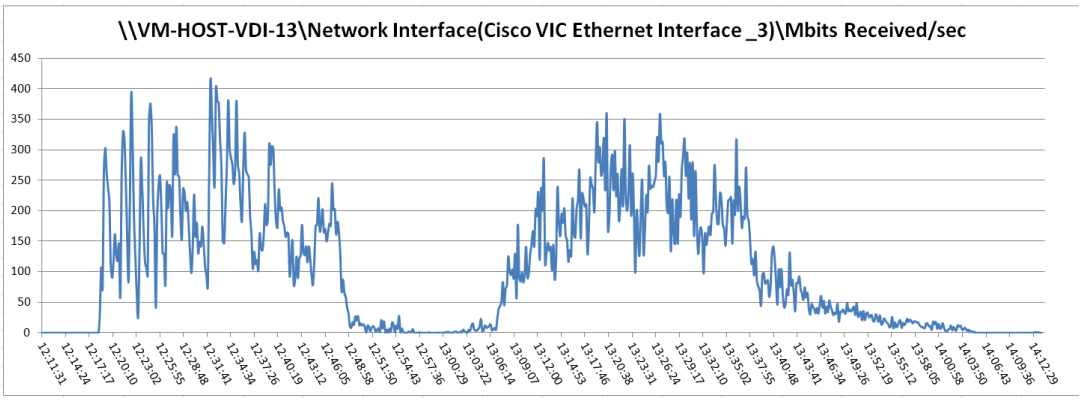
\\VM-HOST-VDI-13\Network Interface(Cisco VIC Ethernet Interface _3)\Mbits Received/sec
450
400
350
300
250
200
150
100
50
0
12:11:31
12:14:24
12:17:17
12:20:10
12:23:02
12:25:55
12:28:48
12:31:41
12:34:34
12:37:26
12:40:19
12:43:12
12:46:05
12:48:58
12:51:50
12:54:43
12:57:36
13:00:29
13:03:22
13:06:14
13:09:07
13:12:00
13:14:53
13:17:46
13:20:38
13:23:31
13:26:24
13:29:17
13:32:10
13:35:02
13:37:55
13:40:48
13:43:41
13:46:34
13:49:26
13:52:19
13:55:12
13:58:05
14:00:58
14:03:50
14:06:43
14:09:36
14:12:29

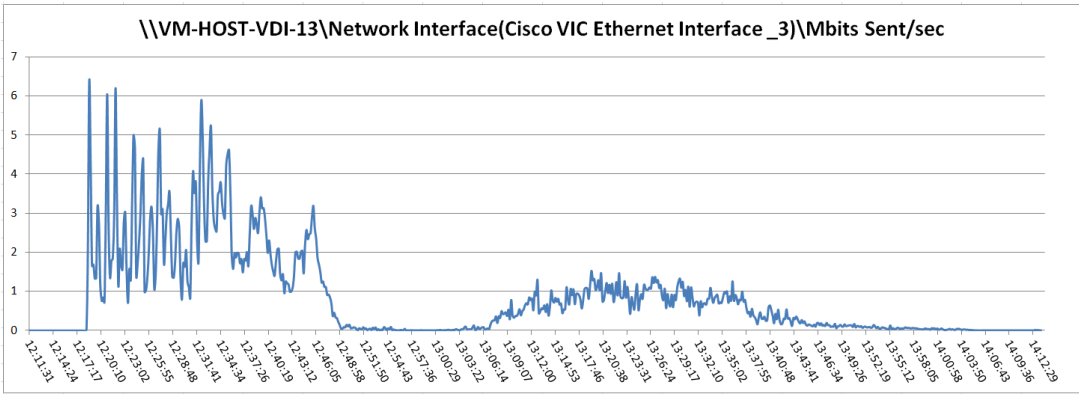
\\VM-HOST-VDI-13\Network Interface(Cisco VIC Ethernet Interface _3)\Mbits Sent/sec
7
6
5
4
3
2
1
0
12:11:31
12:14:24
12:17:17
12:20:10
12:23:02
12:25:55
12:28:48
12:31:41
12:34:34
12:37:26
12:40:19
12:43:12
12:46:05
12:48:58
12:51:50
12:54:43
12:57:36
13:00:29
13:03:22
13:06:14
13:09:07
13:12:00
13:14:53
13:17:46
13:20:38
13:23:31
13:26:24
13:29:17
13:32:10
13:35:02
13:37:55
13:40:48
13:43:41
13:46:34
13:49:26
13:52:19
13:55:12
13:58:05
14:00:58
14:03:50
14:06:43
14:09:36
14:12:29

2000 User, Fourteen Cisco UCS B230 M2 All Phases-VM-Host-VDI-14

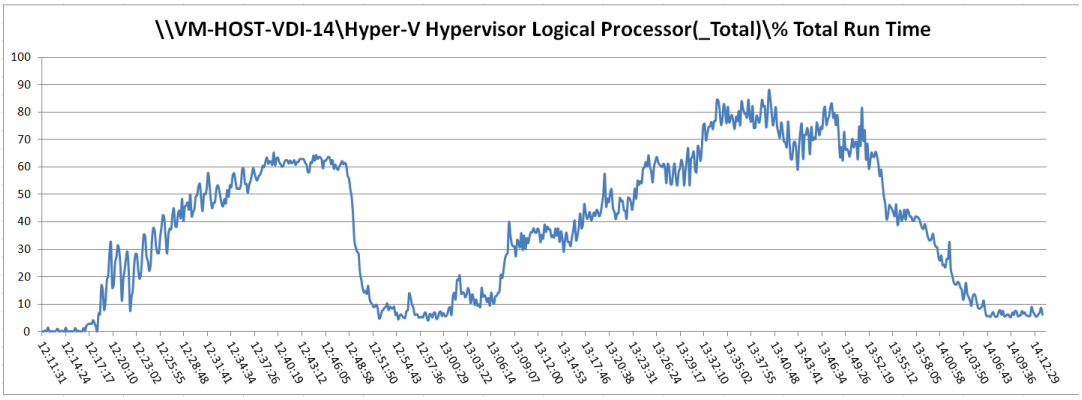

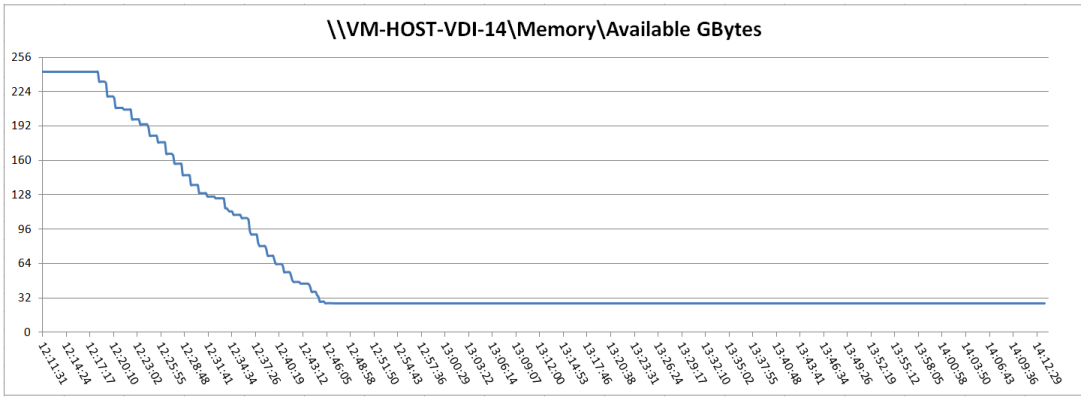

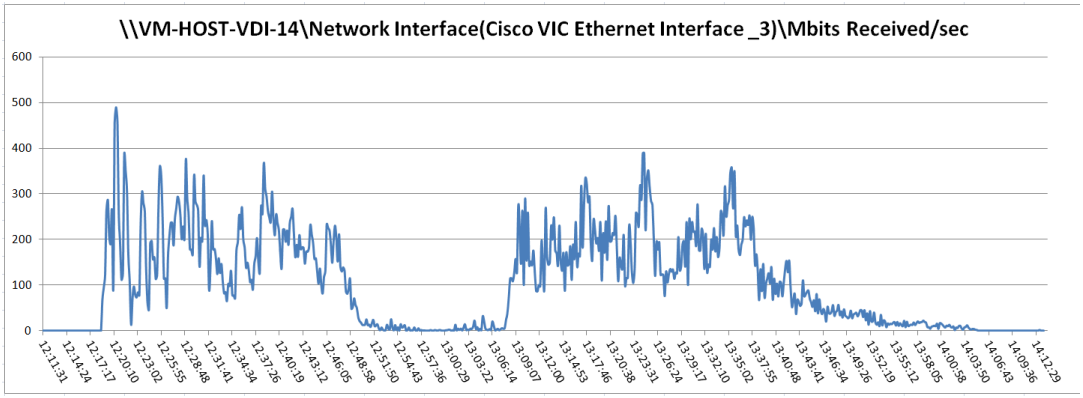

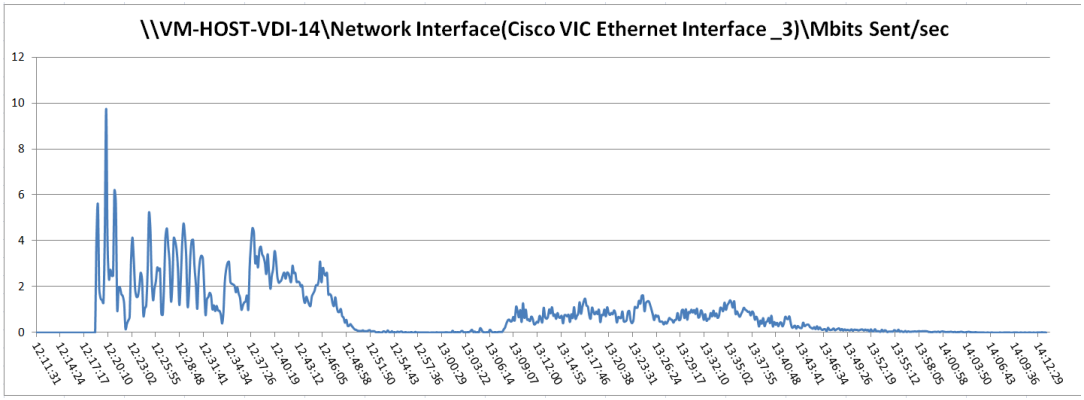
\\VM-HOST-VDI-14\Network Interface(Cisco VIC Ethernet Interface _3)\Mbits Sent/sec
12
10
8
6
4
2
0
12:11:31
12:14:24
12:17:17
12:20:10
12:23:02
12:25:55
12:28:48
12:31:41
12:34:34
12:37:26
12:40:19
12:43:12
12:46:05
12:48:58
12:51:50
12:54:43
12:57:36
13:00:29
13:03:22
13:06:14
13:09:07
13:12:00
13:14:53
13:17:46
13:20:38
13:23:31
13:26:24
13:29:17
13:32:10
13:35:02
13:37:55
13:40:48
13:43:41
13:46:34
13:49:26
13:52:19
13:55:12
13:58:05
14:00:58
14:03:50
14:06:43
14:09:36
14:12:29