# Reference Architecture-Based Design Citrix XenDesktop Built on FlexPod Using Cisco Unified Computing System, Microsoft Hyper-V, and NetApp Storage

Last Updated: August 7, 2011

Building Architectures to Solve Business Problems

# Reference Architecture-Based Design Citrix XenDesktop Built on FlexPod Using Cisco Unified Computing System, Microsoft Hyper-V, and NetApp Storage

# Introduction

This Cisco® Validated Design reports the results of a study evaluating the scalability of the Citrix XenDesktop environment on Cisco® UCS B-Series Blade Servers running Microsoft Hyper-V connected to the NetApp Storage array. We also provide the best practice recommendations and sizing guidelines for large-scale customer deployments of XenDesktop on the Cisco Unified Computing System™.

Reference Architecture-Based Design

Citrix XenDesktop Built on FlexPod

Using Cisco Unified Computing System, Microsoft Hyper-V, and NetApp Storage

## Audience

This document is intended to assist solution architects, sales engineers, field engineers and consultants in planning, design, and deployment of Citrix XenDesktop hosted desktop virtualization solutions on the Cisco Unified Computing System. This document assumes that the reader has an architectural understanding of the Cisco Unified Computing System, Microsoft Hyper-V, Citrix desktop software, NetApp storage system, and related software.

## Objectives

This document is intended to articulate the design considerations and validation efforts required to design and deploy Citrix XenDesktop on the Cisco Unified Computing System with NetApp storage running in a virtualized environment on top of Microsoft Hyper-V.

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

# Summary of Main Findings

The hosting of Citrix XenDesktop Hosted Virtual Desktops (VDI) and Hosted Shared Virtual Desktops models and FlexCast with Citrix XenServer Hypervisors on Cisco UCS B-Series Blade Servers and NetApp storage were successfully validated.

The Cisco UCS B250 M2 Extended Memory Blade Servers offer an optimal memory configuration that allows virtual desktop hosting servers to use the full CPU capabilities of the servers. The 192 GB of memory allowed a high density of desktop sessions per Cisco UCS B250 M2 Extended Memory Blade Servers while offering 1.5 GB of memory to be allocated per desktop-based virtual machine. We were able to scale to 110 Microsoft Windows 7 desktops while running a knowledge worker load.

The validated environment consisted of a completely virtualized infrastructure with virtual machines hosted by Citrix XenServer. All the virtual desktop and supporting infrastructure components including Active Directory, Citrix Provisioning Server, and the Citrix XenDesktop Desktop Delivery Controllers were hosted in a virtual machine environment on Citrix XenServer 5.6.

The tested design showed linear scalability when expanding from 1 server to 16 servers. The performance testing showed that the same user desktop experience and response times were achieved with 110 desktops running on 1 server as with 1760 desktops running on 16 servers.

The integrated management model and rapid provisioning capabilities of Cisco UCS Manager makes it easy for scaling the number of desktops from small pilots on a single UCS chassis to very large organization-wide deployments running on tens of chassis.

The testing validates that the 10Gbps Unified Fabric provides a high performance; scalable infrastructure and offers deterministic performance with respect to user response times during the load and stress testing.

The testing validates that the tested reference architecture can scale linearly from 1 chassis to 4 chassis and beyond without making any changes to the design or infrastructure components. This also requires the proper backend storage scaling as provided by NetApp storage.

Desktop virtual machine "Boot-up" or "Logon" Storms (from rapid concurrent or simultaneous user logons) need to be considered in the server and storage design as they have largest substantial scalability impact on this solution as well as VDI environments in general. The reference architecture represented in this document was able to handle the additional stresses presented by the most extreme boot-up and log-on storm conditions.

# Infrastructure Components

The following sections detail the infrastructure components used in the reference architecture configuration.

## Cisco Unified Computing System

The Cisco Unified Computing System is a next-generation data center platform that unites compute, network, storage access, and virtualization into a cohesive system designed to reduce total cost of ownership (TCO) and increase business agility. The Cisco Unified Computing System server portfolio consists of the Blade Server platform, B-Series and the C-Series Rack Mount platform. We chose the Cisco UCS B-Series Blade Server platform for this study. The system integrates a low-latency, lossless 10 Gigabit Ethernet unified network fabric with enterprise-class, x86-architecture servers. The system is an integrated, scalable, multi-chassis platform in which all resources participate in a unified management domain.

The main system components include:

- Compute—the system is based on an entirely new class of computing system that incorporates blade servers based on Intel Xeon 5500 Series Processors. The Cisco UCS blade servers offer patented Cisco Extended Memory Technology to support applications with large datasets and allow more virtual machines per server.
- Network—the system is integrated onto a low-latency, lossless, 10-Gbps unified network fabric. This network foundation consolidates what today are three separate networks: LANs, SANs, and high-performance computing networks. The unified fabric lowers costs by reducing the number of network adapters, switches, and cables, and by decreasing power and cooling requirements.
- Virtualization—the system unleashes the full potential of virtualization by enhancing the scalability, performance, and operational control of virtual environments. Cisco security, policy enforcement, and diagnostic features are now extended into virtualized environments to better support changing business and IT requirements.
- Storage access—the system provides consolidated access to both SAN storage and Network Attached Storage (NAS) over the unified fabric. Unifying storage access means that the Cisco Unified Computing System can access storage over Ethernet, Fibre Channel, Fibre Channel over Ethernet (FCoE), and iSCSI, providing customers with choice and investment protection. In addition, administrators can pre-assign storage-access policies for system connectivity to storage resources, simplifying storage connectivity and management while helping increase productivity.
- Management—the system uniquely integrates all the system components, enabling the entire solution to be managed as a single entity through the Cisco UCS Manager software. The Cisco UCS Manager provides an intuitive graphical user interface (GUI), a command-line interface (CLI), and a robust application programming interface (API) to manage all system configuration and operations. The Cisco UCS Manager helps increase IT staff productivity, enabling storage, network, and server administrators to collaborate on defining service profiles for applications. Service profiles are logical representations of desired physical configurations and infrastructure policies. They help automate provisioning and increase business agility, allowing data center managers to provision resources in minutes instead of days.

Working as a single, cohesive system, these components unify technology in the data center. They represent a radical simplification in comparison to traditional systems, helping simplify data center operations while reducing power and cooling requirements. The system amplifies IT agility for improved business outcomes. The Cisco Unified Computing System components illustrated in Figure 1 include, from left to right, fabric interconnects, blade server chassis, blade servers, and in the foreground, fabric extenders and network adapters.

***Figure 1*** ***Cisco Unified Computing System***

## Fabric Interconnect

The Cisco UCS 6100 Series Fabric Interconnects are a core part of the Cisco Unified Computing System, providing both network connectivity and management capabilities for the system (Figure 2). The Cisco UCS 6100 Series offers line-rate, low-latency, lossless 10 Gigabit Ethernet and FCoE functions.

The Cisco UCS 6100 Series provides the management and communication backbone for the Cisco UCS B-Series Blade Servers and Cisco UCS 5100 Series Blade Server Chassis. All chassis, and therefore all blades, attached to the Cisco UCS 6100 Series Fabric Interconnects become part of a single, highly available management domain. In addition, by supporting unified fabric, the Cisco UCS 6100 Series provides both the LAN and SAN connectivity for all blades within its domain.

From a networking perspective, the Cisco UCS 6100 Series uses a cut-through architecture, supporting deterministic, low-latency, line-rate 10 Gigabit Ethernet on all ports, independent of packet size and enabled services. The product family supports Cisco low-latency, lossless 10 Gigabit Ethernet unified network fabric capabilities, which increase the reliability, efficiency, and scalability of Ethernet networks. The fabric interconnect supports multiple traffic classes over a lossless Ethernet fabric from the blade through the interconnect. Significant TCO savings come from an FCoE-optimized server design in which network interface cards (NICs), host bus adapters (HBAs), cables, and switches can be consolidated.

The Cisco UCS 6100 Series is also built to consolidate LAN and SAN traffic onto a single unified fabric, saving the capital and operating expenses associated with multiple parallel networks, different types of adapter cards, switching infrastructure, and cabling within racks. Fibre Channel expansion modules in the interconnect support direct connections from the Cisco Unified Computing System to existing native Fibre Channel SANs. The capability to connect FCoE to native Fibre Channel protects existing storage system investments while dramatically simplifying in-rack cabling.

*Figure 2 Cisco UCS 6120XP 20-Port Fabric Interconnect (Top) and Cisco UCS 6140XP 40-Port Fabric Interconnect*

The Cisco UCS 6100 Series is equipped to support the following module options:

- Ethernet module that provides 6 ports of 10 Gigabit Ethernet using the SFP+ interface
- Fibre Channel plus Ethernet module that provides 4 ports of 10 Gigabit Ethernet using the SFP+ interface; and 4 ports of 1/2/4-Gbps native Fibre Channel connectivity using the SFP interface
- Fibre Channel module that provides 8 ports of 1/2/4-Gbps native Fibre Channel using the SFP interface for transparent connectivity with existing Fibre Channel networks
- Fibre Channel module that provides 6 ports of 1/2/4/8-Gbps native Fibre Channel using the SFP or SFP+ interface for transparent connectivity with existing Fibre Channel networks

*Figure 3 From left to right: 8-Port 1/2/4-Gbps Native Fibre Channel Expansion Module; 4-Port Fibre Channel plus 4-Port 10*

## Cisco UCS 2100 Series Fabric Extenders

The Cisco UCS 2100 Series Fabric Extenders bring the unified fabric into the blade server enclosure, providing 10 Gigabit Ethernet connections between blade servers and the fabric interconnect, simplifying diagnostics, cabling, and management.

The Cisco UCS 2100 Series extends the I/O fabric between the Cisco UCS 6100 Series Fabric Interconnects and the Cisco UCS 5100 Series Blade Server Chassis, enabling a lossless and deterministic FCoE fabric to connect all blades and chassis together. Since the fabric extender is similar to a distributed line card, it does not do any switching and is managed as an extension of the fabric interconnects. This approach removes switching from the chassis, reducing overall infrastructure complexity and enabling the Cisco Unified Computing System to scale to many chassis without multiplying the number of switches needed, reducing TCO and allowing all chassis to be managed as a single, highly available management domain.

The Cisco 2100 Series also manages the chassis environment (the power supply and fans as well as the blades) in conjunction with the fabric interconnect. Therefore, separate chassis management modules are not required.

The Cisco UCS 2100 Series Fabric Extenders fit into the back of the Cisco UCS 5100 Series chassis. Each Cisco UCS 5100 Series chassis can support up to two fabric extenders, enabling increased capacity as well as redundancy.

***Figure 4 Rear view of Cisco UCS 5108 Blade Server Chassis with two Cisco UCS 2104XP Fabric Extenders***

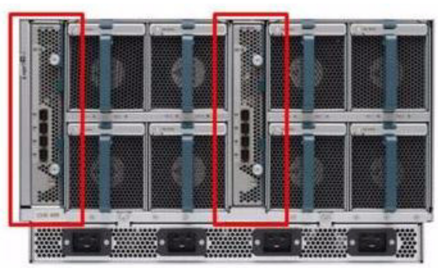

The Cisco UCS 2104XP Fabric Extender has four 10 Gigabit Ethernet, FCoE-capable, Small Form-Factor Pluggable Plus (SFP+) ports that connect the blade chassis to the fabric interconnect. Each Cisco UCS 2104XP has eight 10 Gigabit Ethernet ports connected through the midplane to each half-width slot in the chassis. Typically configured in pairs for redundancy, two fabric extenders provide up to 80 Gbps of I/O to the chassis.

***Figure 5 Cisco UCS 2104XP Fabric Extender***

## Cisco UCS Chassis

The Cisco UCS 5100 Series Blade Server Chassis is a crucial building block of the Cisco Unified Computing System, delivering a scalable and flexible blade server chassis for today's and tomorrow's data center while helping reduce TCO.

Cisco's first blade server chassis offering, the Cisco UCS 5108 Blade Server Chassis, is six rack units (6RU) high and can mount in an industry-standard 19-inch rack. A chassis can house up to eight half-width Cisco UCS B-Series Blade Servers and can accommodate both half- and full-width blade form factors.

Four single-phase, hot-swappable power supplies are accessible from the front of the chassis. These power supplies are 92 percent efficient and can be configured to support non-redundant, N+ 1 redundant and grid-redundant configuration. The rear of the chassis contains eight hot-swappable fans, four power connectors (one per power supply), and two I/O bays for Cisco UCS 2104XP Fabric Extenders.

A passive mid-plane provides up to 20 Gbps of I/O bandwidth per server slot and up to 40 Gbps of I/O bandwidth for two slots. The chassis is capable of supporting future 40 Gigabit Ethernet standards.

***Figure 6 Cisco Blade Server Chassis (front and back view)***

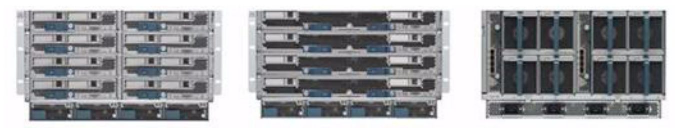

## Cisco UCS B200 M1 Blade Server

The Cisco UCS B200 M1 Blade Server is a half-width, two-socket blade server. The system uses two Intel Xeon 5500 Series Processors, up to 96 GB of DDR3 memory, two optional hot-swappable small form factor (SFF) serial attached SCSI (SAS) disk drives, and a single mezzanine connector for up to 20 Gbps of I/O throughput. The server balances simplicity, performance, and density for production-level virtualization and other mainstream data center workloads.

***Figure 7 Cisco UCS B200 M1 Blade Server***

## Cisco UCS B250 M1 Blade Server

The Cisco UCS B250 M1 Extended Memory Blade Server is a full-width, two-socket blade server featuring Cisco Extended Memory Technology. The system supports two Intel Xeon 5500 Series processors, up to 384 GB of DDR3 memory, two optional SFF SAS disk drives, and two mezzanine connections for up to 40 Gbps of I/O throughput. The server increases performance and capacity for demanding virtualization and large-data-set workloads with greater memory capacity and throughput.

***Figure 8 Cisco UCS B250 M1 Extended Memory Blade Server***

## Intel Xeon 5500 Series Processor

With innovative technologies that boost performance, energy efficiency, and virtualization flexibility, two-processor platforms based on the Intel Xeon 5500 Series Processor make it easier to deliver more business services within existing data center facilities. Data center efficiency starts at the core - with energy-efficient processors and features that help you get the most out of each rack. With a unique combination of performance and energy-efficiency features plus flexible virtualization, the Intel Xeon 5500 Series Processor offers an effective antidote to data center sprawl and improves business competitiveness. The combination of Intel Turbo Boost Technology and Intel Hyper-Threading Technology delivers optimal performance for each enterprise application, and Intel QuickPath Technology dramatically increases application performance and throughput for bandwidth-intensive applications.

Greater per-server performance means that you can do more with fewer servers and potentially save significantly on operating costs. Intel Intelligent Power Technology works alongside these new performance features to deliver better performance with lower power consumption at all operating points, achieving the best available performance/watt. High-performance 95-watt, standard 80-watt and low-power 60-watt versions enable high-density deployments in both rack and blade form factors.

Intel VT with Intel FlexMigration and Intel FlexPriority also gives IT more choice in managing and allocating virtualized workloads across new and existing platforms. Intel Turbo Boost Technology plus hardware assists from Intel VT improves performance for applications running in virtual machines. Intel VT FlexMigration, in combination with virtualization management software, can help IT to conserve power, rebalance workloads and reduce energy consumption.

*Figure 9 Intel Xeon 5500 Series Processor*

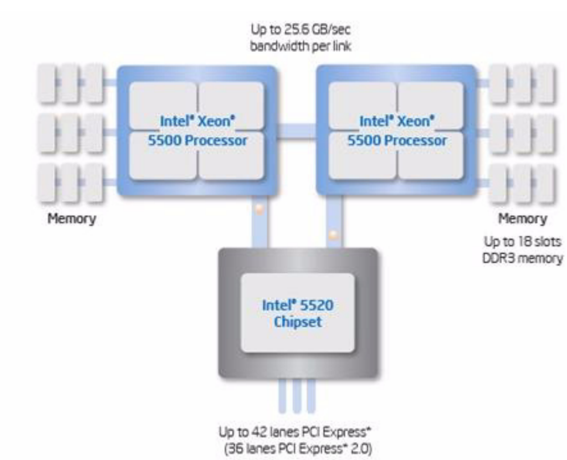


## Intel Xeon 5600 Series Processor

As data centers reach the upper limits of their power and cooling capacity, efficiency has become the focus of extending the life of existing data centers and designing new ones. As part of these efforts, IT needs to refresh existing infrastructure with standard enterprise servers that deliver more performance and scalability, more efficiently. The Intel Xeon 5600 Series Processor automatically regulates power consumption and intelligently adjusts server performance according to your application needs, both energy efficiency and performance. The secret to this compelling combination is Intel's new 32nm Xeon microarchitecture. Featuring Intel Intelligent Power Technology that automatically shifts the CPU and memory into the lowest available power state, while delivering the performance you need, the Intel Xeon

5600 Series Processor with Intel Micro-architecture Xeon delivers the same performance as previous-generation servers but uses up to 30 percent less power. You can achieve up to a 93 percent reduction in energy costs when consolidating your single-core infrastructure with a new infrastructure built on Intel Xeon 5600 Series Processor.

This ground breaking intelligent server technology features:

- Intel's new 32nm Microarchitecture Xeon built with second-generation high-k and metal gate transistor technology.
- Intelligent Performance that automatically optimizes performance to fit business and application requirements and delivers up to 60 percent more performance per watt than Intel Xeon 5500 Series Processor.
- Automated Energy Efficiency that scales energy usage to the workload to achieve optimal performance/watt and with new 40 Watt options and lower power DDR3 memory, you can lower your energy costs even further.
- Flexible virtualization that offers best-in-class performance and manageability in virtualized environments to improve IT infrastructure and enable up to 15:1 consolidation over two socket, single-core servers. New standard enterprise servers and workstations built with this new generation of Intel process technology offer an unprecedented opportunity to dramatically advance the efficiency of IT infrastructure and provide unmatched business capabilities.

*Figure 10 Intel Xeon 5600 Series Processor*


## Intel Xeon Processor 6500 and 7500 Series (Nehalam-EX)

The Intel Xeon processor 7500 series supports up to eight integrated cores and 16 threads, and is available with frequencies up to 2.66 GHz, and 24 MB of cache memory, four Intel QPI links and Intel Turbo boost technology. Thermal design point (TDP) power levels range from 95 watt to 130 watts. The Intel Xeon processor 6500 series provides a lower cost solution for 2-chip servers.

These new Intel processors are packed with more than 20 new features that deliver a leap forward in reliability, availability and serviceability (RAS). These reliability capabilities are designed to improve the protection of data integrity, increase availability and minimize planned downtime.

For example, this is the first Xeon processor to possess Machine Check Architecture (MCA) Recovery, a feature that allows the silicon to work with the operating system and virtual machine manager to recover from otherwise fatal system errors, a mechanism until now found only in Intel® Itanium® processor family and RISC processors.

## Cisco UCS B200 M2 Blade Server

The Cisco UCS B200 M2 Blade Server is a half-width, two-socket blade server. The system uses two Intel Xeon 5600 Series Processors, up to 96 GB of DDR3 memory, two optional hot-swappable small form factor (SFF) serial attached SCSI (SAS) disk drives, and a single mezzanine connector for up to 20 Gbps of I/O throughput. The server balances simplicity, performance, and density for production-level virtualization and other mainstream data center workloads.

*Figure 11 Cisco UCS B200 M2 Blade Server*

## Cisco UCS B250 M2 Extended Memory Blade Server

The Cisco UCS B250 M2 Extended Memory Blade Server is a full-width, two-socket blade server featuring Cisco Extended Memory Technology. The system supports two Intel Xeon 5600 Series Processors, up to 384 GB of DDR3 memory, two optional SFF SAS disk drives, and two mezzanine connections for up to 40 Gbps of I/O throughput. The server increases performance and capacity for demanding virtualization and large-data-set workloads with greater memory capacity and throughput.

*Figure 12 Cisco UCS B250 M2 Extended Memory Blade Server*

## Cisco UCS B440 M1 High-Performance Blade Server

The Cisco UCS B440 M1 High-Performance Blade Server is a full-width, 4-socket system. Two or four Intel Xeon 7500 Series Processors with intelligent performance that automatically adapts to the diverse needs of a virtualized environment and offers advanced reliability for mission-critical workloads. It supports 32 dual in-line memory module (DIMM) slots and up to 256 GB at 1333 MHz based on Samsung's 40 nanometer class (DDR3) technology. There is four optional front-accessible, hot-swappable Small Form-Factor Pluggable (SFFP) drives and an LSI SAS2108 RAID Controller. The Cisco UCS B440 M1 blade server can accommodate two dual-port mezzanine cards for up to 40 Gbps I/O per blade. Options include a Cisco UCS M81KR Virtual Interface Card (VIC) or converged network adapter (Emulex or QLogic compatible).

*Figure 13 Cisco UCS B440 M1 Blade Server*

## Cisco UCS B230 M1 Blade Server

Cisco has expanded the architectural advantages of its Intel Xeon Processor 6500 and 7500 Series-based server platforms with an exceptionally high density blade server. The 2-socket Cisco UCS B230 M1 Blade Server platform delivers high performance and density in a compact, half-width form factor.

In addition, it provides one dual-port mezzanine card for up to 20 Gbps I/O per blade. Options include a Cisco UCS M81KR Virtual Interface Card or converged network adapter (Emulex or QLogic compatible).

Other features include

- 32 dual in-line memory module (DIMM) slots and up to 256 GB at 1066 MHz based on Samsung 40-nanometer class (DDR3) technology
- Two optional front-accessible, hot-swappable solid-state drives (SSDs) and an LSI SAS2108 RAID Controller
- Greatly simplified deployment and systems management with embedded integration into Cisco UCS Manager

Each Cisco UCS 5108 Blade Server Chassis can house up to eight B230 M1 servers (a maximum of 320 per Cisco Unified Computing System).

*Figure 14 Cisco UCS B230 M1 Blade Server*

## Cisco UCS Virtual Interface Card (VIC)

Cisco Virtual Interface Cards were developed ground up to provide acceleration for the various new operational modes introduced by server virtualization. The Virtual Interface Cards are highly configurable and self-virtualized adapters that can create up 128 PCIe endpoints per adapter. These PCIe endpoints are created in the adapter firmware and present fully compliant standard PCIe topology to the host OS or hypervisor.

Each of these PCIe endpoints the Virtual Interface Card creates can be configured individually for the following attributes:

- Interface type: FCoE, Ethernet or Dynamic Ethernet interface device
- Resource maps that are presented to the host: PCIe BAR's, interrupt arrays
- The Network presence and attributes: MTU, VLAN membership
- QoS parameters: 802.1p class, ETS attributes, rate limiting and shaping

*Figure 15 Cisco UCS Virtual Interface Card*

**Note** The Virtual Interface Cards are SR-IOV capable at the hardware level and Cisco will provide a smooth transition to SR-IOV based solution when operating systems and hypervisors support it.

## Cisco Extended Memory Architecture

Modern CPUs with built-in memory controllers support a limited number of memory channels and slots per CPU. The need for virtualization software to run multiple OS instances demands large amounts of memory, and that, combined with the fact that CPU performance is outstripping memory performance, can lead to memory bottlenecks. Even some traditional non-virtualized applications demand large amounts of main memory: database management system performance can be improved dramatically by caching database tables in memory, and modeling and simulation software can benefit from caching more of the problem state in memory.

To obtain a larger memory footprint, most IT organizations are forced to upgrade to larger, more expensive, four-socket servers. CPUs that can support four-socket configurations are typically more expensive, require more power, and entail higher licensing costs. Cisco Extended Memory Technology expands the capabilities of CPU-based memory controllers by logically changing the geometry of main memory while still using standard DDR3 memory. This technology makes every four DIMM slots in the expanded memory blade server appear to the CPU's memory controller as a single DIMM that is four times the size (Figure 15). For example, using standard DDR3 DIMMs, the technology makes four 8-GB DIMMS appear as a single 32-GB DIMM.

This patented technology allows the CPU to access more industry-standard memory than ever before in a two-socket server:

- For memory-intensive environments, data centers can better balance the ratio of CPU power to memory and install larger amounts of memory without having the expense and energy waste of moving to four-socket servers simply to have a larger memory capacity. With a larger main-memory footprint, CPU utilization can improve because of fewer disk waits on page-in and other I/O operations, making more effective use of capital investments and more conservative use of energy.
- For environments that need significant amounts of main memory but which do not need a full 384 GB, smaller-sized DIMMs can be used in place of 8-GB DIMMs, with resulting cost savings: two 4-GB DIMMS are typically less expensive than one 8-GB DIMM.

*Figure 16 Cisco Extended Memory Architecture*

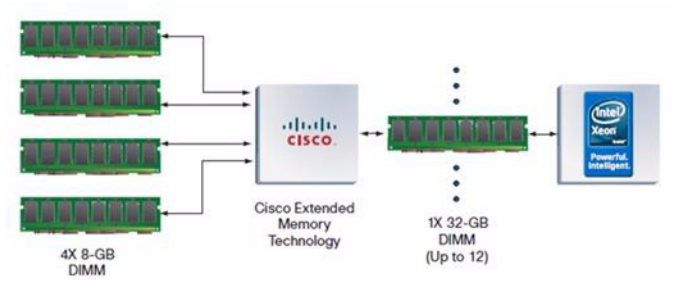


## Cisco UCS C-Series Rack-Mount Servers

The Cisco UCS C-Series Rack-Mount Servers (Figure 17) extend the Cisco Unified Computing System 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 study was carried out on the Cisco UCS B-Series Blade Servers, the C-Series Rack-Mount Servers extend the same benefits to customers. Future desktop virtualization studies are planned on this server platform.

*Figure 17 Cisco UCS C-Series Rack-Mount Servers*


# Citrix XenDesktop

Citrix XenDesktop transforms Windows desktops as an on-demand service to any user, any device, anywhere. XenDesktop quickly and securely delivers any type of virtual desktop or Windows, web and SaaS application to all the latest PCs, Macs, tablets, smartphones, laptops and thin clients - all with a high-definition HDX™ user experience. FlexCast™ delivery technology enables IT to optimize the performance, security and cost of virtual desktops for any type of user, including task workers, mobile workers, power users and contractors. XenDesktop helps IT rapidly adapt to business initiatives, such as offshoring, M&A and branch expansion, by simplifying desktop delivery and enabling user self-service. The open, scalable and proven architecture simplifies management, support and integration.

## Citrix FlexCast Technology

Different types of workers across the enterprise have varying performance and personalization requirements. Some require offline mobility of laptops, others need simplicity and standardization, while still others need high performance and a fully personalized desktop. XenDesktop can meet all these requirements in a single solution with our unique Citrix FlexCast™ delivery technology. With FlexCast, IT can deliver every type of virtual desktop and application, hosted or local, optimized to meet the performance, security and mobility requirements of each individual user. 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 required-or appropriate.
- Hosted VDI Desktops offer a personalized Windows desktop experience for office workers that can be securely delivered over any network to any device.
- Streamed VHD Desktops leverage 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 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 datacenter.
- On-demand applications allow any Microsoft Windows application to be centralized and managed in the datacenter, hosted either on multi-user terminal servers or virtual machines, 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.

A complete overview of the FlexCast technology can be found at Citrix.com, but for our testing purposes and validation represented in this Cisco Validated Design, only the Hosted VDI models were validated on the Cisco UCS hardware in conjunction with NetApp storage solutions. High-Definition User Experience (HDX)Technology.

Citrix XenDesktop delivers an HDX™ user experience on any device, over any network, with better reliability and higher availability than a traditional PC. With Citrix HDX™ technology, users get a rich, high-definition experience, even when using multimedia, real-time collaboration, USB peripherals, and 3D graphics. These user experience technologies must balance performance with low bandwidth—anything else becomes impractical to use and scale. 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 Architecture Overview

The Citrix XenDesktop with FlexCast Delivery Technology can deliver different types of virtual desktops based on the performance, security and flexibility requirements of each individual user. The following section details the Hosted VDI architecture.

### Citrix XenDesktop Hosted VDI Overview

Hosted VDI uses a hypervisor to host all the desktops in the datacenter. 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 reverted back to its clean, original state. Citrix Provisioning Services enables you to stream 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 the figure below:

*Figure 18 Citrix XenDesktop 4 on Microsoft Hyper-V Architecture*

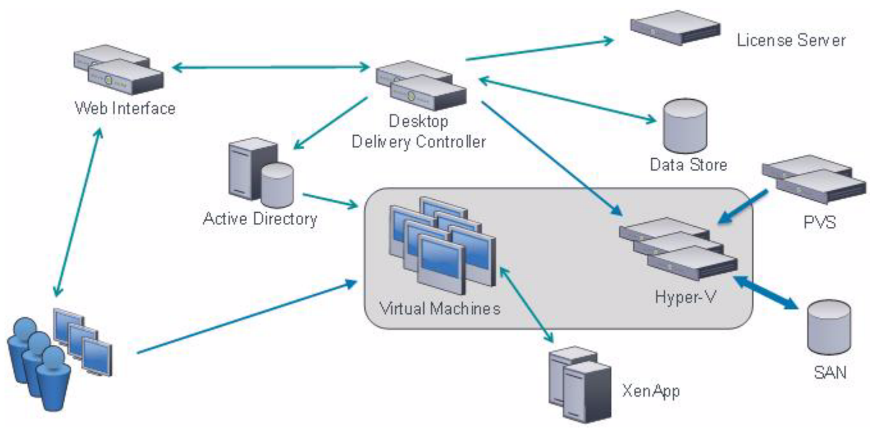


- Web Interface: Web Interface provides the user interface to the XenDesktop environment. Web Interface brokers user authentication, enumerates the available desktops and, upon launch, delivers an .ica file to the Citrix Receiver on the user's local device to initiate a connection. Because Web Interface 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. 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 filename 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.
- Provisioning Services: 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, by providing the best flexibility, and 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 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. It is for this reason why it is recommended to have multiple Controllers per farm
- Data Store: Each XenDesktop farm requires a database called the data store. Citrix XenDesktop uses 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 ICA (Independent Computing Architecture) connections between the virtual desktop and user devices with the Citrix online plug-in.
- Citrix Online Plug-in: Installed on user devices, the Citrix online plug-in enables direct ICA connections from user devices to virtual desktops. The plug-in software is available for a range of different devices so users can connect to published applications from various platforms. You can deploy and update the online plug-in using Citrix Receiver.
- Microsoft Hyper-V:
- Citrix XenApp: Citrix XenApp is an on-demand application delivery solution that enables any Windows application to be virtualized, centralized, and managed in the datacenter, 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 application virtualization and delivery.

All the aforementioned components interact to provide a virtual desktop to an end-user based on the FlexCast Hosted VDI desktop delivery model using 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 outline the sequence of operations executed by XenDesktop to deliver a Hosted VDI desktop to the end user.

*Figure 19 Operational Sequence*

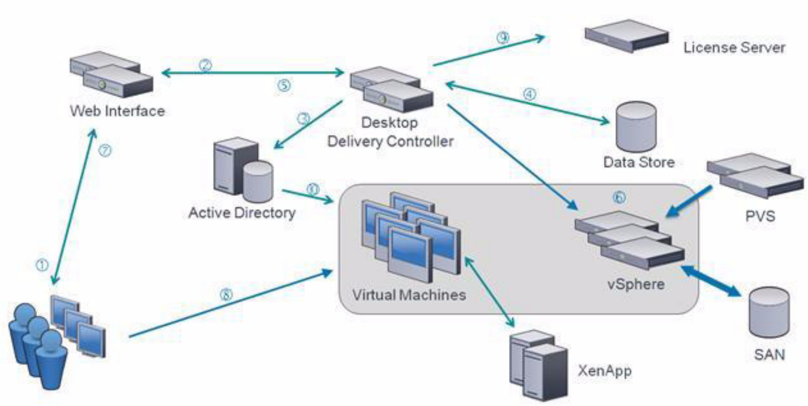

1. The end user launches an internet browser to access Web Interface.
2. Web Interfaces prompts the user for Active Directory credentials and passes the credentials to the Desktop Delivery Controller acting as a dedicated XML server.
3. The XML Service running the dedicated XML server (Desktop Delivery Controller) authenticates the user against Active Directory.
4. After the user is successfully authenticated, the XML Service contacts the Data Store to determine which virtual desktops are available for that user.
5. The virtual desktop information is sent back to Web Interface and Web Interface renders a web page containing a list of available desktops.
6. The user clicks on the desktop icon and Web Interface forwards the request to the Desktop Delivery Controller. If the virtual desktop is powered on, the Desktop Delivery Controller will tell the Virtual Desktop Agent running on the virtual machine to start listening for an incoming session. If the virtual desktop is not powered on, the Desktop Delivery Controller will tell the Hyper-V server to start a new virtual desktop and then notify the Virtual Desktop Agent
   a. In a Hosted VDI configuration with Provisioning Services, the virtual desktop boots through the network PXE boot. The virtual desktop contacts the DHCP server to find an IP address and the location of the boot file. The boot file comes from Provisioning Services and provides instructions for accessing the centralized desktop image.
   b. After the virtual desktop receives the boot file with instructions, it contacts the Provisioning Server and provides its MAC address. Provisioning Server identifies the correct virtual desktop disk based on the MAC address and sends portions of the virtual disk to the virtual desktop required to start-up the machine.
7. The virtual desktop connection information is forwarded onto Web Interface. Web Interface creates a launch file (ICA) for the specific virtual desktop and forwards the launch file to the end user's device.
8. The Virtual Desktop Agent running on the virtual desktop tells the Desktop Delivery Controller that the user has connected. The user's logon information is then sent for validation.
9. The Desktop Delivery Controller validates the login credentials and checks out a license from the Citrix License Server. If the credentials are valid and a license is available, then the credentials, XenDesktop license and policies are sent to the virtual desktop for processing.
10. When the connection has been approved, the Virtual Desktop Agent uses the transferred credentials to logon against Active Directory and applies profile configurations.

## General Citrix XenDesktop Advantages and Value Proposition

Citrix XenDesktop transforms Windows desktops as an on-demand service to any user, any device, anywhere. XenDesktop quickly and securely delivers any type of virtual desktop or Windows, web and SaaS application to all the latest PCs, Macs, tablets, smartphones, laptops and thin clients - all with a high-definition HDX™ user experience.

The following describes the strategic features of Citrix XenDesktop 4:

- Any device, anywhere with Receiver™. Today's digital workforce demands the flexibility to work from anywhere at any time using any device they'd like. Leveraging Citrix Receiver as a lightweight universal client, XenDesktop users can access their desktop and corporate applications from the latest tablets, smartphones, PCs, Macs, or thin client. This enables virtual workstyles, business continuity and user mobility.

- HDX™ user experience. XenDesktop delivers an HDX™ user experience on any device, over any network, while using up to 90% less bandwidth compared to competing solutions. With HDX, the desktop experience rivals a local PC, even when using multimedia, real-time collaboration, USB peripherals, and 3D graphics. Integrated WAN optimization capabilities boost network efficiency and performance even over challenging, high latency links.
- Beyond VDI with FlexCast™. Different types of workers across the enterprise have varying performance and personalization requirements. Some require offline mobility of laptops, others need simplicity and standardization, while still others need high performance and a fully personalized desktop. XenDesktop can meet all these requirements in a single solution with our unique Citrix FlexCast™ delivery technology. With FlexCast, IT can deliver every type of virtual desktop, hosted or local, optimized to meet the performance, security and mobility requirements of each individual user.
- Any Windows, web or SaaS app. With XenDesktop, you can provide your workforce with any type of application they need, including Windows, web and SaaS apps. For Windows apps, XenDesktop includes XenApp™, the on-demand application delivery solution that enables any Windows app to be virtualized, centralized, and managed in the datacenter and instantly delivered as a service to users anywhere on any device. For web and SaaS apps, Receiver seamlessly integrates them into a single interface, so users only need to log on once to have secure access to all their applications.
- Open, scalable, proven. With numerous awards, industry-validated scalability and over 10,000 Citrix Ready products, XenDesktop provides a powerful desktop computing infrastructure that's easier than ever to manage. The open architecture works with your existing hypervisor, storage, Microsoft, and system management infrastructures, with complete integration and automation through the comprehensive SDK.
- Single-instance management. XenDesktop enables IT to separate the device, OS, applications and user personalization and maintain single master images of each. Instead of juggling thousands of static desktop images, IT can manage and update the OS and apps once, from one location. Imagine being able to centrally upgrade the entire enterprise to Windows 7 in a weekend, instead of months. Single-instance management dramatically reduces on-going patch and upgrade maintenance efforts, and cuts datacenter storage costs by up to 90 percent by eliminating redundant copies.
- Data security and access control. With XenDesktop, users can access desktops and applications from any location or device, while IT uses policies that control where data is kept. XenDesktop can prevent data from residing on endpoints, centrally controlling information in the datacenter. In addition, XenDesktop can help ensure that any application data that must reside on the endpoint is protected with XenVault technology. Extensive access control and security policies help ensure that intellectual property is protected, and regulatory compliance requirements are met.

The Cisco Desktop Virtualization Solution with Citrix XenDesktop delivers desktops and applications as an on-demand service to users anywhere, at any time, and on their choice of devices. The solution supports a new balance between IT and users. It empowers users with mobility, flexibility, and productivity on a global scale. It gives IT organizations the tools they need to better meet the changing demands of today's business concerns, including rapidly responding to events ranging from mergers and acquisitions to the opening of a new branch office.

The solution incorporates the most flexible, cost-effective and scalable platform for hosting virtual desktops. Built from the ground up to support virtualization, the solution transforms data center operations by simplifying server and workload management, making IT staff more productive. The Cisco Desktop Virtualization Solution with Citrix XenDesktop protects IT investments by growing and adapting to business needs by incorporating new technologies without forklift upgrades.

The solution delivers an uncompromising user experience that is driven by Citrix HDX technology and can be customized on a per-user basis. The solution extends its reach propelled by Cisco's leadership in enterprise networking and computing. The Cisco Unified Computing System is powered by Intel® Xeon® Series Processors that speed performance with data-center-grade reliability and availability. The solution makes data center operations secure and compliant to a level no other solution can match, helping IT organizations meet regulatory requirements by combining centralized business-critical data with single-instance storage of each OS, application, and user profile.

Cisco and Citrix together deliver a virtual desktop solution that can transform business operations while increasing the productivity of any organization's greatest asset: its people.

# Microsoft Windows 2008 R2 and 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. The adoption of virtualization in the enterprise has increased flexibility in 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 quick disaster recovery. Even so, customers are looking for more flexibility.

Windows Server 2008 R2 Hyper-V provides greater flexibility with live migration. Live migration is integrated with Windows Server 2008 R2 Hyper-V and Microsoft Hyper-V Server 2008 R2. With Hyper-V live migration, you can move running virtual machines (VMs) from one Hyper-V physical host to another, without any disruption or perceived loss of service. IT professionals increasingly look to live migration to create a dynamic and flexible IT environment that responds to emerging business needs. Live migration provides the core technology required for dynamic load balancing, VM placement, high availability for virtualized workloads during physical computer maintenance, and reduced data center power consumption.

## Introduction to Hyper-V

Hyper-V in Windows Server 2008 and Windows Server 2008 R2 enables you to create a virtualized server computing environment. You can use a virtualized computing environment to improve the efficiency of your computing resources by utilizing more of your hardware resources. This is possible because you use Hyper-V to create and manage virtual machines and their resources. Each virtual machine is a virtualized computer system that operates in an isolated execution environment. This allows you to run multiple operating systems simultaneously on one physical computer.

Hyper-V provides software infrastructure and basic management tools that you can use to create and manage a virtualized server computing environment. This virtualized environment can be used to address a variety of business goals aimed at improving efficiency and reducing costs. For example, a virtualized server environment can help you:

- Reduce the costs of operating and maintaining physical servers by increasing your hardware utilization. You can reduce the amount of hardware needed to run your server workloads.
- Increase development and test efficiency by reducing the amount of time it takes to set up hardware and software and reproduce test environments.
- Improve server availability without using as many physical computers as you would need in a failover configuration that uses only physical computers.

## Hyper-V Architecture and Feature Overview

Hyper-V™ is a role in Windows Server® 2008 R2 that provides tools and services that you can use to create and manage a virtualized server computing environment. Within this virtualized environment, you can create and manage virtual machines which allow you to run multiple computers, even with different operating systems, on one physical computer. In this virtualized environment, you can choose to isolate the virtual machines, or specify how they communicate with each other and with an external network. Windows Server 2008 R2 Hyper-V adds new features to the first version of Hyper-V. For example, by using live migration in Windows Server 2008 R2 Hyper-V, you can migrate running VMs from one physical computer to another, and add or remove storage from a VM while it is running. In addition, Windows Server 2008 R2 Hyper-V takes better advantage of physical computer hardware with greater processor support and deeper support for physical computer hardware.

### Hyper-V Architecture

Hyper-V provides a reliable virtualization platform that enables customers to virtualize their infrastructure and reduce costs. It has thin microkernelized hypervisor architecture with minimal attack surface and is available as a Server Core role. With System Center integrated management tools, customers can use a single set of tools to manage both their physical and virtual resources. It easily plugs into customers' IT infrastructure, as they can leverage their existing patching, provisioning, management and support tools, and processes. It provides great value, as it is available as a feature of Windows Server 2008 R2, and customers can leverage the breadth of solutions from Microsoft partners, existing IT Pro skill sets, and comprehensive support from Microsoft.

*Figure 20 Hyper-V Architecture*

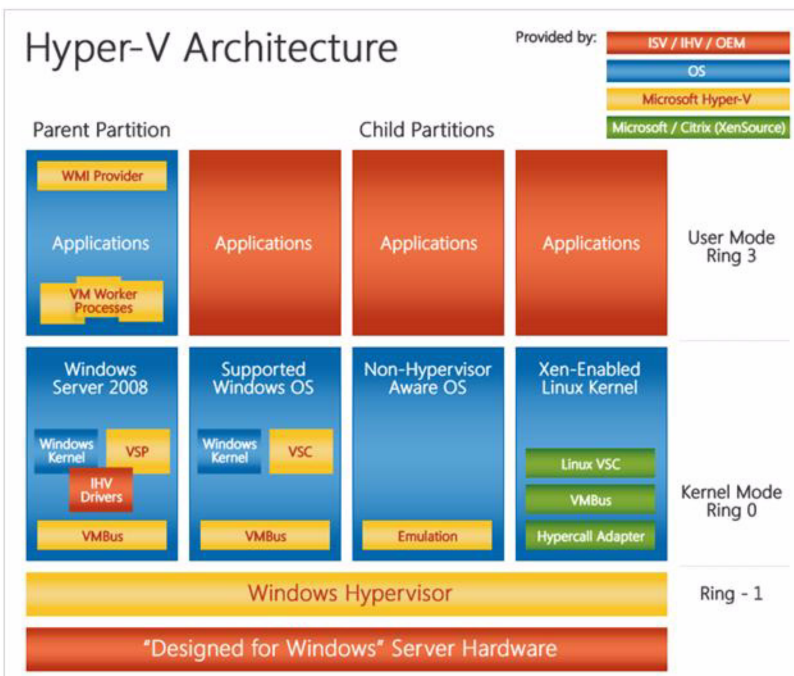

### Key Features of Hyper-V

- Live Migration

  Hyper-V in Windows Server 2008 R2 includes the much-anticipated live migration feature, which allows you to move a virtual machine between two virtualization host servers without any interruption of service. Hyper-V live migration is integrated with Windows Server 2008 R2 Hyper-V and Microsoft Hyper-V Server 2008 R2. With it you can move running VMs from one Hyper-V physical host to another without any disruption of service or perceived downtime. Moving running VMs without downtime using Hyper-V live migration:

  - Provides better agility. 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.
  - Reduces costs and increases productivity. Data centers with multiple Hyper-V physical hosts can service those systems in a controlled fashion, scheduling maintenance during regular business hours. Live migration makes it possible to keep VMs online, even during maintenance, 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.

- Increased Hardware Support for Hyper-V Virtual Machines

  Hyper-V in 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 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 Logical Unit Number (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.

- Improved Cluster Node Connectivity Fault Tolerance

  Because of the architecture of CSV, there is improved cluster node connectivity fault tolerance that directly affects VMs running on the cluster. The CSV architecture implements a mechanism, known as dynamic I/O redirection, where I/O can be rerouted within the failover cluster based on connection availability.

- Enhanced Cluster Validation Tool

  Windows Server 2008 R2 includes a Best Practices Analyzer (BPA) for all major server roles, including Failover Clustering. This analyzer examines the best practices configuration settings for a cluster and cluster nodes.

- Improved Management of Virtual Data Centers

  Even with all the efficiency gained from virtualization, VMs still need to be managed. The number of VMs tends to proliferate much faster than physical computers because machines typically do not require a hardware acquisition. Therefore, management of virtual data centers is even more imperative than ever before.

- Improved Virtual Networking Performance

  The new Hyper-V leverages several new networking technologies contained in Windows Server 2008 R2 to improve overall VM networking performance.

- Increased Performance and Reduce Power Consumption

  Hyper-V in Windows Server 2008 R2 adds enhancements that reduce virtual machine 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 by virtue of the new Core Parking feature implemented in Windows Server 2008 R2.

- Enhanced 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 fewer number of calls from the network stack to the network driver.

  TCP Chimney, which allows the off loading of TCP/IP processing to the network hardware, has been extended to the virtual environment. It improves VM performance by allowing the VM to off load network processing to hardware, especially on networks with bandwidth over 1 GB. This feature is especially beneficial for roles involving large amounts of data transfer, such as the file server role.

  The Virtual Machine Queue (VMQ) feature allows physical computer network interface cards (NICs) to use direct memory access (DMA) to place the contents of packets directly into VM memory, increasing I/O performance.

- Dynamic VM Storage

  Windows Server 2008 R2 Hyper-V supports hot plug-in and hot removal of storage. By supporting the addition or removal of Virtual Hard Drive (VHD) files and pass-through disks while a VM is running, Windows Server 2008 R2 Hyper-V makes it possible to reconfigure VMs quickly to meet changing workload requirements. This feature allows the addition and removal of both VHD files and pass-through disks to existing SCSI controllers for VMs.

- Broad OS Support

  Broad support for simultaneously running different types of operating systems, including 32-bit and 64-bit systems across different server platforms, such as Windows, Linux, and others.

- Network Load Balancing

  Hyper-V includes new virtual switch capabilities. This means virtual machines can be easily configured to run with Windows Network Load Balancing (NLB) Service to balance load across virtual machines on different servers.

- New Hardware Sharing Architecture

  With the new virtual service provider/virtual service client (VSP/VSC) architecture, Hyper-V provides improved access and utilization of core resources, such as disk, networking, and video.

- Virtual Machine Snapshot

  Hyper-V provides the ability to take snapshots of a running virtual machine so you can easily revert to a previous state, and improve the overall backup and recoverability solution.

- Extensible

  Standards-based Windows Management Instrumentation (WMI) interfaces and APIs in Hyper-V enable independent software vendors and developers to quickly build custom tools, utilities, and enhancements for the virtualization platform.

  More comprehensive information on Virtualization and Hyper-V can be found at: http://www.microsoft.com/windowsserver2008/en/us/hyperv-technical-resources.aspx

## Virtual Machine Management

Even with all the efficiency gained from virtualization, virtual machines still need to be managed. The number of virtual machines tends to proliferate much faster than physical computers because machines typically do not require a hardware acquisition. Therefore, management of virtual data centers is even more imperative than ever before.

Windows Server 2008 R2 includes the following improvements that will help you manage your virtual data center:

- Reduced effort for performing day-to-day Hyper-V administrative tasks by using the Hyper-V Management Console.
- Enhanced command-line interface and automated management of Hyper-V administrative tasks by using PowerShell cmdlets.
- Improved management of multiple Hyper-V servers in a virtual data center environment by using System Center Virtual Machine Manager 2008.

## System Center Virtual Machine Manager (Virtual Machine Manager)

System Center Virtual Machine Manager 2008 R2 provides centralized administration and management of your virtual environment, helps to increase physical server utilization, and enables rapid provisioning of new virtual machines by the Virtual Machine Manager administrator and authorized self-service end users.

A Virtual Machine Manager implementation is composed of certain core components that must be implemented and a set of optional components and features that you can implement as needed in your environment.

### Virtual Machine Manager Server

The Virtual Machine Manager server is the hub of a Virtual Machine Manager implementation through which all other Virtual Machine Manager components interact and communicate. Therefore you should install the Virtual Machine Manager server first.

The Virtual Machine Manager server runs the Virtual Machine Manager service, which runs commands, transfers files, and controls communications with other Virtual Machine Manager components and with all virtual machine hosts and Virtual Machine Manager library servers, collectively referred to as managed computers.

The Virtual Machine Manager server also connects to a SQL Server database that stores all Virtual Machine Manager configuration information. You can access this information and configure Virtual Machine Manager by using the Virtual Machine Manager Administrator Console or by using the Windows PowerShell - Virtual Machine Manager command shell. For more information about the Windows PowerShell - Virtual Machine Manager command shell, see Scripting.

By default, the Virtual Machine Manager server is also the default library server. You can use the Virtual Machine Manager library to store file-based resources, such as virtual hard disks (unless attached to a stored virtual machine), templates, ISO images, PowerShell scripts, answer files, and virtual machines. You can set up additional Virtual Machine Manager library servers, which is recommended when you will be managing a large number of hosts.

### Virtual Machine Library Server

The Virtual Machine Manager library is a catalog of resources that can be used to create and configure virtual machines in Virtual Machine Manager.

The library can store the following types of resources:

- File-based resources such as virtual hard disks, virtual floppy disks, ISO images, and scripts. To be used in Virtual Machine Manager, a file must be added to the library. This requires storing the file on a library server on a designated library share. For more information, see Adding File-Based Resources to the Library. For a list of file types that can be stored in the Virtual Machine Manager library, see File Types That Are Indexed During a Library Refresh.
- Virtual machine templates, hardware profiles, and guest operating system profiles, which are configured in Library view for use in creating standard virtual machines. These configurations are stored in the Virtual Machine Manager database but are not represented by physical configuration files.
- Virtual machines that are not in use. The virtual machines are displayed in Library view. However, the files for a stored virtual machine are not displayed in the library because the files cannot be used to create or configure new virtual machines.

The Virtual Machine Manager server is always the default library server. When you install the Virtual Machine Manager server, Setup automatically creates a library share on the Virtual Machine Manager server. You can add additional library shares to the default library server.

After the Virtual Machine Manager server installed, you can add other computers as library servers, as noted in the Optional Virtual Machine Manager Components and Features section later in this topic. However, you cannot remove or relocate the default library server or its library share on the Virtual Machine Manager server. So give careful consideration to the location of the default library share before installing the Virtual Machine Manager server. When you add a computer as a library server, Virtual Machine Manager automatically installs a Virtual Machine Manager agent on the computer.

### Virtual Machine Manager Database

The Virtual Machine Manager database stores all Virtual Machine Manager configuration information, which you can access and modify by using the Virtual Machine Manager Administrator Console. You specify the Virtual Machine Manager database when you are installing the Virtual Machine Manager server.

The Virtual Machine Manager database requires a supported version of Microsoft SQL Server. You can either specify a local or remote instance of an existing Microsoft SQL Server database or have the Setup Wizard install SQL Server 2005 Express Edition SP2 on the Virtual Machine Manager server. The Setup Wizard also installs SQL Server 2005 Tools and creates a SQL Server instance named MICROSOFT$Virtual Machine Manager$ on the local computer.

If you plan to manage more than 150 hosts or integrate Operations Manager 2007 with Virtual Machine Manager to implement Performance and Resource Optimization (PRO) tips and reporting, you should use a Standard or Enterprise version of SQL Server. For more information about supported versions of SQL Server, see System Requirements: Virtual Machine Manager Database.

### Virtual Machine Manger Administrator Console

The Virtual Machine Manager Administrator Console is a graphical user interface (GUI) that you use to:

- Create, deploy, and manage virtual machines.
- Monitor and manage hosts and library servers.
- Manage global configuration settings.

You install the Virtual Machine Manager Administrator Console after installing the Virtual Machine Manager server and then connect it to your Virtual Machine Manager server. You can install the Virtual Machine Manager Administrator Console on the same computer as the Virtual Machine Manager server or on a different computer. You can install the Virtual Machine Manager Administrator Console on multiple computers and use it to connect to and manage any Virtual Machine Manager server. You can connect to and manage only one Virtual Machine Manager server at a time.

When you install a Virtual Machine Manager Administrator Console, Setup also installs Windows PowerShell - Virtual Machine Manager command shell, which makes available the cmdlets you can use from the command line to perform all functions that you can perform within the Virtual Machine Manager Administrator Console.

### Virtual Machine Host

A virtual machine host is a physical computer that hosts one or more virtual machines. You can add one or more hosts to Virtual Machine Manager by using the Add Hosts Wizard in the Virtual Machine Manager Administrator Console. Until you add a host, you cannot use Virtual Machine Manager to create virtual machines and many of the actions in the Virtual Machine Manager Administrator Console are not available. When you add a Windows Server-based host, Virtual Machine Manager automatically installs a Virtual Machine Manager agent on the host and, if it is not already installed or enabled, installs or enables the appropriate virtualization software on the host.

Virtual Machine Manager supports the following types of hosts:

- Windows Server-based hosts that are located in an AD domain that does not have a two-way trust with the Virtual Machine Manager server's AD domain.
- Windows Server-based hosts that are located on a perimeter network.
- Windows Server-based hosts that are in a disjointed namespace, where the host's fully qualified domain name (FQDN) resolved from the domain name service (DNS) is not the same as name obtained from AD.
- VMware ESX Server hosts located anywhere in your environment.

## Optional Virtual Machine Manager Components and Features

### Virtual Machine Manager Self-Service Portal

The Virtual Machine Manager Self-Service Portal is an optional, Web-based component that a Virtual Machine Manager administrator can install and configure to allow users to create and manage their own virtual machines within a controlled environment on a limited group of virtual machine hosts. The Virtual Machine Manager administrator creates self-service user roles which determine the scope of the users' actions on their own virtual machines.

To create, operate, and manage virtual machines, self-service users use the Virtual Machine Manager Self-Service Portal. The administrator determines which host groups self-service users can create virtual machines on. When a self-service user creates a virtual machine, the virtual machine is automatically placed on the most suitable host in the host group based on host ratings.

A Virtual Machine Manager administrator can set a virtual machine quota in a self-service user role and assign quota points to virtual machine templates to limit the number of virtual machines that a user or group can deploy.

### Performance and Resource Optimization (PRO)

Performance and Resource Optimization (PRO) supports workload- and application-aware resource optimization within a virtualized environment. Based on performance and health data provided by PRO-enabled management packs in System Center Operations Manager 2007, PRO can automatically or manually implement PRO tips, recommendations for minimizing downtime and accelerating time to resolution. For more information, see About Performance and Resource Optimization (PRO).

## Failover Clustering and Hyper-V

A failover cluster is a group of independent computers that work together to increase the availability of applications and services. The clustered servers (called nodes) are connected by physical cables and by software. If one of the cluster nodes fails, another node begins to provide service (a process known as failover). Users experience a minimum of disruptions in service. One of the marquee features in Windows Server 2008 R2, Live Migration, uses the new Cluster Shared Volumes (CSV) feature within the failover clustering role in Windows Server 2008 R2.

- Cluster Shared Volumes

  Cluster Shared Volumes (CSV), which functions as a distributed-access file system that is optimized for Hyper-V, is a significant architectural innovation incorporated in failover clustering in Windows Server 2008 R2 Enterprise and in Windows Server 2008 R2 Datacenter, as well as in Microsoft Hyper-V Server 2008 R2.

  Unlike a Clustered File System (CFS), CSV does not use a specialized proprietary file system-it uses the standard NTFS file system, so it requires nothing additional to purchase or support. Regular clustering storage devices can be used—Fibre Channel, iSCSI, or Serial Attached SCSI—as long as these devices receive a logo for Windows Server 2008 R2 and the complete solution passes the Validate a Configuration Wizard.

  With CSV, any node in a failover cluster can access the shared storage and any node can host virtual machines, regardless of which node owns (or manages NTFS on) the storage (see Figure 21). CSV provides many benefits, including easier storage management, greater resiliency against failures, and the ability to store many VMs on a single LUN and have them fail over individually. Most notably, CSV provides the infrastructure to support and enhance live migration of Hyper-V virtual machines.

*Figure 21 Cluster Shared Volumes*

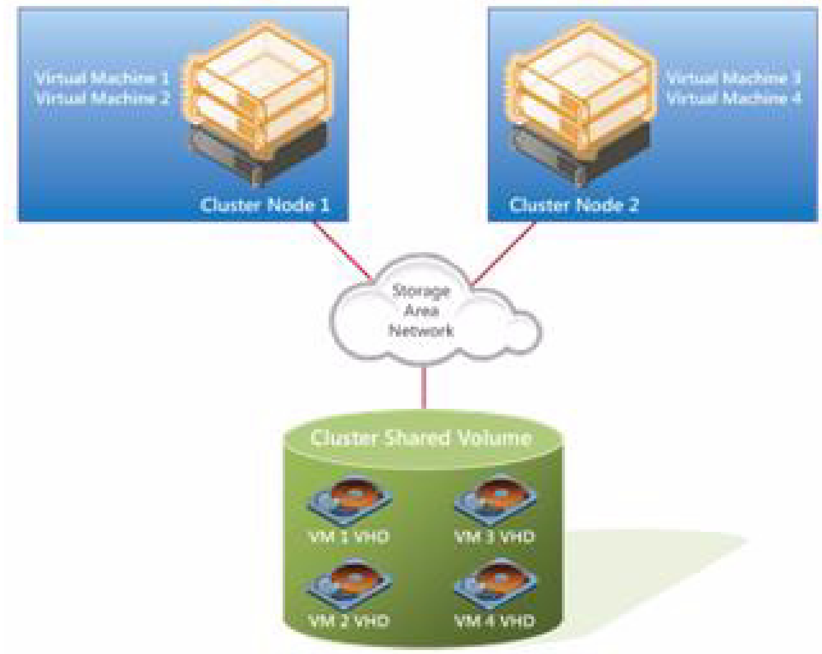


Because CSV provides a consistent file namespace to all the nodes in the failover cluster, any files that are stored on CSV have the same name and path from any node in the failover cluster. CSV volumes are stored as directories and subdirectories beneath the %SystemDrive%\ClusterStorage root folder, as illustrated in Figure 22.

*Figure 22 Single Namespace in CSV*

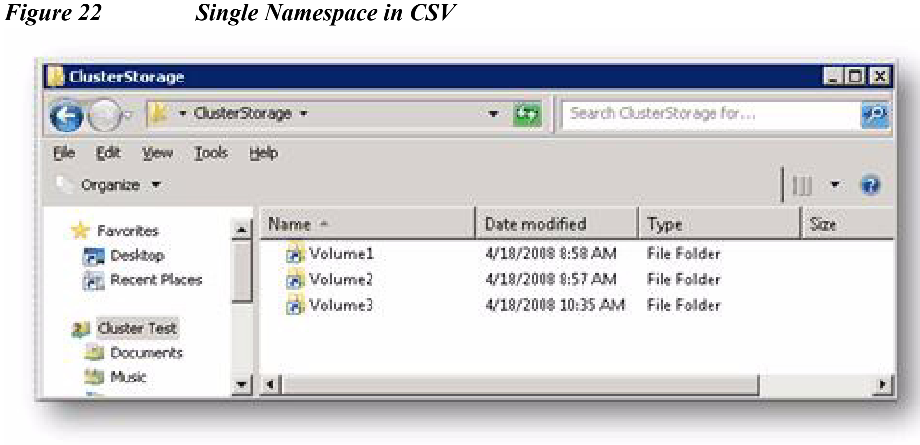


In Windows Server 2008 and in earlier versions of Windows, each resource group was required to have a physical disk resource to manage the cluster's access to the application data residing on the shared storage device. This makes sure that only one node could access the data at a time. However, since only one node could own the LUN at any time, the LUN was the smallest unit of failover. If any application that was running on the LUN needed to move to another node, all the applications on that LUN would also be failed over. This process incurred some downtime during the failover. Therefore, customers frequently ran only a single application from each LUN so that only that one application would become unavailable during a failover. This added significant complexity to storage management—clusters with hundreds of resources needed hundreds of LUNs, which were then challenging to deploy and manage.

CSV breaks the dependency between application resources (the virtual machines) and disk resources; it does not matter where a disk is mounted because it will appear local to all the nodes in the failover cluster. However, the applications on that CSV disk can run on any node at any time.

CSV manages storage access differently than regular clustered disks do (see Figure 23). A node is designated as a coordinator node—this can be any node in a failover cluster. When an application needs to write data to the CSV disk, the application requests permission from the coordinator node. If the data being written does not require changing the file system structure, the coordinator node sends the addresses of the writable blocks, letting the VM write directly to the file system (direct access to the LUN). If the data requires a change to the file system structure—for example, an extension, a move, or a change in the attributes of a file—the coordinator node writes the data to the disk. In this way, the coordinator node helps ensure that access to the NTFS file system is controlled to prevent corruption.

*Figure 23 Coordinator Node in CSV*

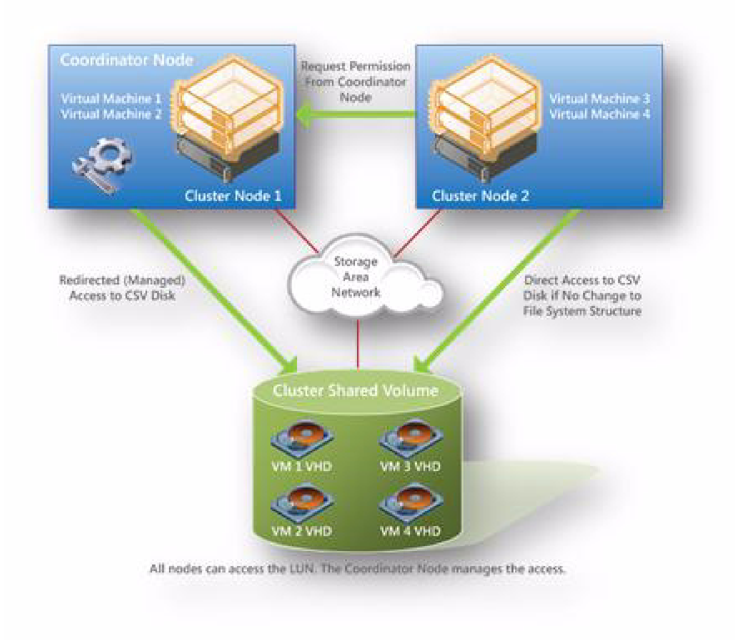


- Live Migration

Live migration is enhanced by CSV within failover clustering in Windows Server 2008 R2. Live migration and CSV are separate but complimentary technologies: they can work independently, but CSV enhances the resilience of live migration and is thus recommended (not required) for live migration. The CSV volumes enable multiple nodes in the same failover cluster to concurrently access the same LUN. From the perspective of the VMs, each appears to actually own a LUN; however, the .vhd files for each VM are stored on the same CSV volume.

CSV is the enabling technology for live migration with Hyper-V. With CSV, moving a virtual machine  between physical servers happen in millisecondsÒwithout dropped network connections or perceived downtime. Migration operations become virtually invisible to connected users.

Once you establish a failover cluster of physical servers that are running Hyper-V, VMs can be live migrated at will. In the event of a node failure, VMs are restarted on another cluster node.

# NetApp Storage

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

- Storage efficiency: Significant cost savings with multiple levels of storage efficiency for all the virtual machine data components.
- Performance: Enhanced user experience with transparent read and write I/O optimization that strongly complements NetApp's storage efficiency capabilities.
- Data protection: Enhanced protection of both the virtual desktop OS data and the user data, with very low overhead for both cost and operations.

## 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 new alternatives and expanded possibilities over traditional storage vendors. All NetApp storage systems utilize the Data ONTAP operating system to provide SAN (FCoE, Fibre Channel, 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 great 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 total cost of ownership.

For large, scalable Citrix XenDesktop environments, the NetApp solution provides the following unique benefits:

- At least 50% savings in storage, power, and cooling requirements
- Most 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, it would require at least 20 B usable data on the shared storage.

Thin provisioning, data deduplication, and FlexClone® 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 on average 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 physically available. With thin provisioning, the storage administrator is able to utilize a pool of physical disks (known as an aggregate) and create logical volumes for different applications to use, while not pre-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.

*Figure 24 Traditional and Thin Provisioning*

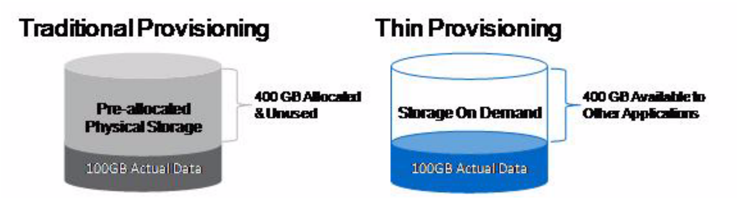


NetApp recommends using thinly provisioned LUNs where possible in the Hyper-V environment for maximum storage efficiency. It is important to note that when using thin provisioning it is important to monitor capacity utilization. In addition administrators should also configure storage management policies 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.

*Figure 25 Increased Disk Utilization With NetApp Thin Provisioning*

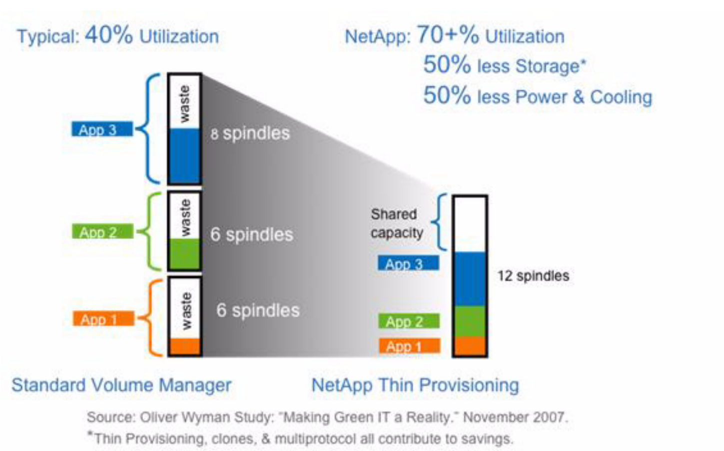


## NetApp Deduplication

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 Citrix XenDesktop 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.

*Figure 26 NetApp Deduplication*

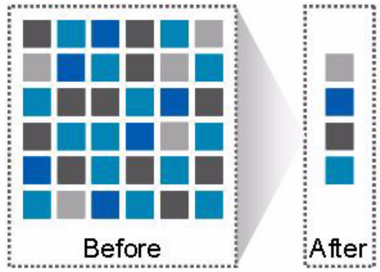


Using NetApp deduplication and file FlexClone not only can reduce the overall storage footprint of Citrix XenDesktop desktops but also can improve performance by using transparent storage cache sharing. Data that is deduplicated or nonduplicated, in the case of file FlexClone data, on disk will only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

Any nondeduplicated 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 deduped data will 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.

*Figure 27 NetApp Deduplication and Flexcone*

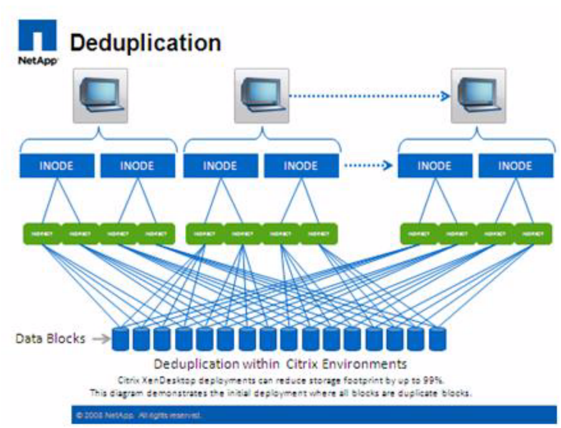


For more information on deduplication, refer to [NetApp TR-3505: NetApp Deduplication for FAS, Deployment and Implementation Guide].

## 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 Citrix XenDesktop 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 nondisruptive 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. Also, 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 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 Citrix XenDesktop 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 increased overall solution cost.

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

- Requiring far less disks and 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 Citrix XenDesktop implementations can see as great as 99 percent 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 10x 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 total cost of ownership.

- Green benefits: Power and cooling costs are reduced as the overall energy needed to run and cool the Flash Cache module is significantly less than even a single shelf of Fibre Channel 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. In contrast, the Flash Cache module consumes only 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. For 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 and PAM

NetApp Flash Cache and PAM are hardware devices that extend the native Data ONTAP TSCS capabilities. Flash Cache increases the amount of available cache which helps reduce virtual desktop storm activities. More details of Flash Cache will be discussed later in this document. For more details on NetApp Flash Cache technology, visit http://www.netapp.com/us/products/storage-systems/flash-cache/flash-cache-tech-specs.html.

**Note** For the remainder of this document, the use of Flash Cache will represent both the Flash Cache and PAM modules.

## NetApp Write Optimization

Virtual desktop I/O patterns are often very 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. 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 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 (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 WAFL® since their invention. NetApp arrays coalesce 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, RAID-DP provides the desired resiliency (protection against dual disk failure) and performance, comparable to RAID 10 but at the cost of RAID 5.

## Flexible Volumes and Aggregates

Flexible volumes (also known as FlexVol volumes) 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. Much like the way that Citrix virtualizes computing resources, NetApp virtualizes the storage resources.

### Data Protection

The availability of thousands of virtual desktops is dependent on the availability of the shared storage on which the virtual desktops are hosted. Thus, using the proper RAID technology is very 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®, 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. RAID DP provides performance that is comparable to that of RAID 10, yet requires fewer disks to achieve equivalent protection. RAID DP provides protection against double disk failure as compared to RAID 5, which can only protect against 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.

## Storage Sizing Best Practices

Storage estimation for deploying Citrix XenDesktop 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/business continuance requirements.
- Size of NAS (CIFS) home directories.

- NetApp strongly recommends storing user data on NAS (CIFS) home drives. 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 back cache size needs to be calculated based on how often the user reboots the desktop and what applications the user uses. We recommend using NFS for write back cache for space efficiency and easy management.
  - NetApp thin provisioning, deduplication, and NetApp snapshot can be used to achieve the desired storage efficiency and data protection for the "user data disk."

## 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 Citrix XenDesktop that can help architect Citrix XenDesktop 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

# Storage Architecture Best Practices

In a Citrix XenDesktop 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 that address the availability and performance for large, scalable Citrix XenDesktop environments. For a complete Citrix XenDesktop deployment guide, refer to NetApp TR-XXXX: XenDesktop on Hyper-V with NetApp.

## Storage System Configuration Best Practices

This section provides 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 as all desktops rely on the shared storage. For more details, visit www.netapp.com/us/products/platform-os/active-active.html.
- 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. For more information, visit http://media.netapp.com/documents/tr-3437.pdf.
- 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 provides 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 only protect against one disk failure. NetApp strongly recommends using RAID-DP on all RAID groups that store Citrix XenDesktop data. For more information on RAID-DP, refer to NetApp white paper 3298 at http://www.netapp.com/us/library/white-papers/wp_3298.html.

- 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 provide proactive alerts to NetApp to help provide faster service. For more details on RLM, visit http://now.netapp.com/NOW/download/tools/rlm_fw/info.shtml.
- Networking infrastructure design (FCoE, FCFibre 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 and 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 NOW site.
- Use the latest storage controller, shelf, and disk firmware available on the NOW 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 unless this is a FAS 2000 series due to lack of disk spindles.
- Replicate data with SnapMirror or SnapVault 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 SyncMirror® and RAID-DP for the highest level of storage resiliency.

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

### Building a High-Performance Storage Architecture

A XenDesktop 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 one 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, 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 a Citrix XenDesktop environments, especially during a boot storm, login storm, or virus storm, as only one copy of deduplicated data will need to be read from the disk (per volume). Each subsequent access of a shared block will be read from Flash Cache and not from disk, increasing performance and decreasing latency and overall disk utilization.

# FlexPod Technical Overview

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

FlexPod™ is a predesigned, base configuration that is built on the Cisco Unified Computing System™, Cisco Nexus® data center switches, NetApp® FAS storage components, and a range of software partners. FlexPod can scale up for greater performance and capacity, or it can scale out for environments that need consistent, multiple deployments. FlexPod is a baseline configuration, but also has the flexibility to be sized and optimized to accommodate many different use cases.

Cisco, NetApp, and VMware have developed FlexPod for VMware® as a platform that can address current virtualization needs and simplify their evolution to an IT as a service (ITaaS) infrastructure. FlexPod for VMware is built on the FlexPod infrastructure stack with added VMware components, including VMware vSphere™ and vCenter™ for virtualized application workloads.

## Audience

This document describes the basic architecture of FlexPod for VMware and also prescribes the procedure for deploying a base FlexPod for VMware configuration. The intended audience of this document includes, but is not limited to, sales engineers, field consultants, professional services, IT managers, partner engineering, and customers who want to deploy the core FlexPod for VMware architecture.

## FlexPod 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 demand change. This includes scaling both up (adding additional resources within a FlexPod unit) and out (adding additional FlexPod units).

Specifically, FlexPod is a defined set of hardware and software that serves as an integrated building block for all virtualization solutions. FlexPod includes NetApp storage, Cisco Nexus networking, the Cisco Unified Computing System (Cisco UCS), and operating system/hypervisor/virtualization software in a single package in which the computing and storage fit as a cohesive pod. . Due to port density, the networking components can accommodate multiple FlexPod configurations. Figure 35 shows the FlexPod components.

***Figure 28 FlexPod for VMware Components***

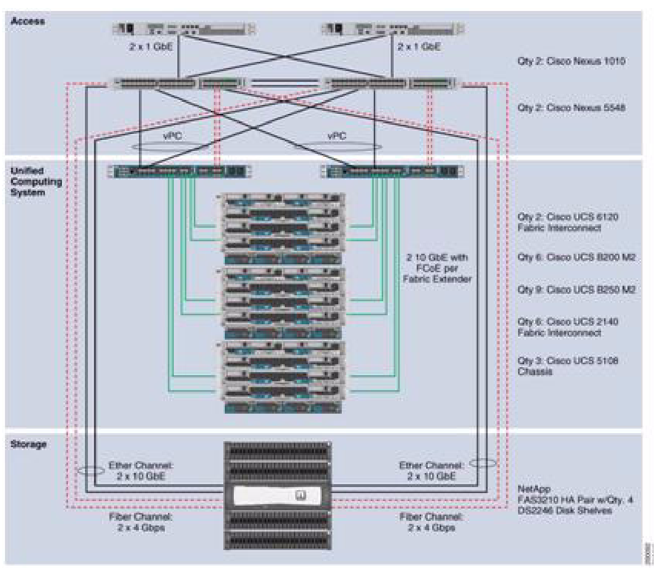


The default hardware includes:

- Two Cisco Nexus 5548 switches
- Two Cisco UCS 6120 fabric interconnects
- Three chassis of Cisco UCS blades with two fabric extenders per chassis

Storage is provided by a NetApp FAS3210CC (HA configuration within a single chassis) with accompanying disk shelves. All systems and fabric links feature redundancy, providing for end-to-end high availability (HA. While this is the default base design, each of the components can be scaled flexibly to support the specific business requirements in question. For example, more (or different) blades and chassis could be deployed to increase compute capacity, additional disk shelves could be deployed to improve I/O capacity and throughput, or special hardware or software features could be added to introduce new features (such as NetApp Flash Cache for dedupe-aware cachingfor VDI deployments).

## FlexPod Market Overview

**The Challenge: Disruptive, Inflexible Transition From Infrastructure Silos**

Today's IT departments are increasingly challenged by the complexity and management of disparate components within their data centers. Rapidly proliferating silos of server, storage, and networking resources combined with numerous management tools and operational processes have led to crippling inefficiencies and costs. Savvy organizations understand the financial and operational benefits of moving from infrastructure silos to a virtualized, shared environment. However, many of them are hesitant to make the transition due to potential short-term business disruptions and long-term architectural inflexibility, which can impede scalability and responsiveness to future business changes. Enterprises and service providers need a tested, cost-effective virtualization solution that can be easily implemented and managed within their existing infrastructures and that scales to meet their future cloud computing objectives.

**The Solution: Unified, pretested, and validated shared infrastructure to simplify your data center transformation**

To meet this challenge NetApp and Cisco have collaborated to create FlexPod™. FlexPod is a proven, long term data center solution built on a flexible, shared infrastructure that can scale easily; be optimized for a variety of mixed application workloads; or be configured for virtual desktop or server infrastructure, secure multi-tenancy and Cloud environments. FlexPod is a prevalidated configuration that delivers a virtualized data center in a rack composed of leading computing, networking, storage, and infrastructure software components.

***Table 1 FlexPod Facilitates a Variety of Virtualized Cloud Environments***

| Environment Examples | Capabilities |
|---|---|
| Exchange 2010, Microsoft SharePoint, or SQL Server | • Start with the FlexPod predesigned, simple, validated data center solution to run on a variety of application workloads<br>• Data center efficiency<br>• Virtualized shared infrastructure |
| Virtual desktop infrastructure | • Deduplication of images<br>• Rapid provisioning through VM cloning<br>• Array-based data protection |
| Development/production | • Space-efficient cloning of datasets and VMs<br>• Data center efficiency<br>• Dynamic infrastructure<br>• Add ITSM and vCloud orchestration for self-service portals<br>• Virtualized shared infrastructure<br>• Optional to add secure multi-tenancy architecture for secure portioning and isolation<br>• Array-based data protection |
| Internal private IaaS cloud/secure multi-tenancy | • Add secure multi-tenancy architecture for secure partitioning and isolation of workloads<br>• Add ITSM/orchestration for automated provisioning, workflows, usage metrics, and reporting |
| Service provider | • Data center efficiency<br>• Dynamic infrastructure for SaaS, PaaS, and IaaS<br>• Add secure multi-tenancy for secure partitioning and isolation of client environments |

# Cisco Networking Infrastructure

## Cisco Nexus 5548 28-Port Switch

The Cisco Nexus® 5548 Switch is a 1RU, 10 Gigabit Ethernet and 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 6 additional 10 Gigabit Ethernet and FCoE ports, up to 8 Fibre Channel 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

### Features and Benefits

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- 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 fiber for long runs.
- Fibre Channel over Ethernet and IEEE Data Center Bridging features supports I/O consolidation, eases management of multiple traffic flows, and optimizes performance. Although implementing SAN consolidation requires only the lossless fabric provided by the Ethernet pause mechanism, the Cisco Nexus 5500 Series provides additional features that create an even more easily managed, high-performance, unified network fabric

### 10 Gigabit Ethernet 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), quality of service (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 network interface card [NIC]). These numbers, together with the congestion management features described next, 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 Fibre Channel over Ethernet, Hardware-Level I/O Consolidation, and End-Port Virtualization. For more information, refer to http://www.cisco.com/en/US/products/ps9670/prod_white_papers_list.html.

# Microsoft Windows 7

Microsoft introduced Windows 7 in fall of 2009 as their next generation desktop operating system to succeed Windows XP, their other flagship software. According to a recent IDC report around 70 percent of the enterprise users are using Windows XP and a majority of them are already looking to migrate to Windows 7. Given the high-level of interest from enterprise customers to migrate to Windows 7 platform, the virtual desktop validation represented in this paper leveraged Windows 7 virtual desktops.

The following section outlines the creation and provisioning of the Windows 7 virtual desktops.

## Microsoft Windows 7 Image Creation and Provisioning

The Microsoft Windows 7 image and additional test workload specific software was initially installed and prepared as a standard Virtual Machine on Microsoft Hyper-V server. Using Citrix Provisioning Services, a virtual disk image (vDisk) is created by uploading the files from the original virtual machine. This single vDisk can serve hundreds of read-only and diskless virtual desktops. Citrix XenDesktop Setup wizard tool is used to create virtual desktop images in bulk.

The XenDesktop Setup Wizard effectively creates virtual machine objects, configures - RAM and assigns network properties. It also creates and configures the relevant Provisioning Services, Desktop Delivery Controller and Active Directory objects associated with these.

While creating virtual desktops on the Hyper-V hosts, the XenDesktop Setup Wizard distributes the virtual desktops on all Hyper-V hosts managed by the Virtual Machine Manager (VMM) instance. In order to only create virtual desktops on the desired Hyper-V hosts, either one of the following two approaches can be used:

Use of a Temporary VMM Server: Build a VMM server with a local SQL Express database and temporarily register the Hyper-V hosts with it. Then run the XenDesktop Setup Wizard and after completion, re-register the Hyper-V hosts with the production VMM server. Modify the XenDesktop group properties appropriately to reflect the correct hosting VMM server.

Use of PowerShell Scripts: Virtual desktops can be created by running PowerShell scripts on the VMM server. These scripts access the VMM server, XenDesktop and Provisioning Services APIs.

PowerShell scripts were implemented during this validation testing primarily because they also facilitate the creation and attachment of the write cache drives to the virtual machines. XenDesktop Setup Wizard discards any drives attached to the virtual machine used as a template, so write-cache drives for all virtual machines need to be added after the desktops are created.

The Appendix contains the scripts used to generate virtual desktops, create write cache drives and attach the drives to the virtual machines. The Appendix also contains the example scripts that can be used to import virtual desktop objects to Provisioning Services and the Desktop Delivery Controller (DDC).

More information as to why the additional virtual disks are needed can be found in the Configuration Topology for Scalability of Citrix XenDesktop on the Cisco Unified Computing System and NetApp Storage section.

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

*Figure 29 Windows 7 Image and vDisk Provisioning Process Overview*

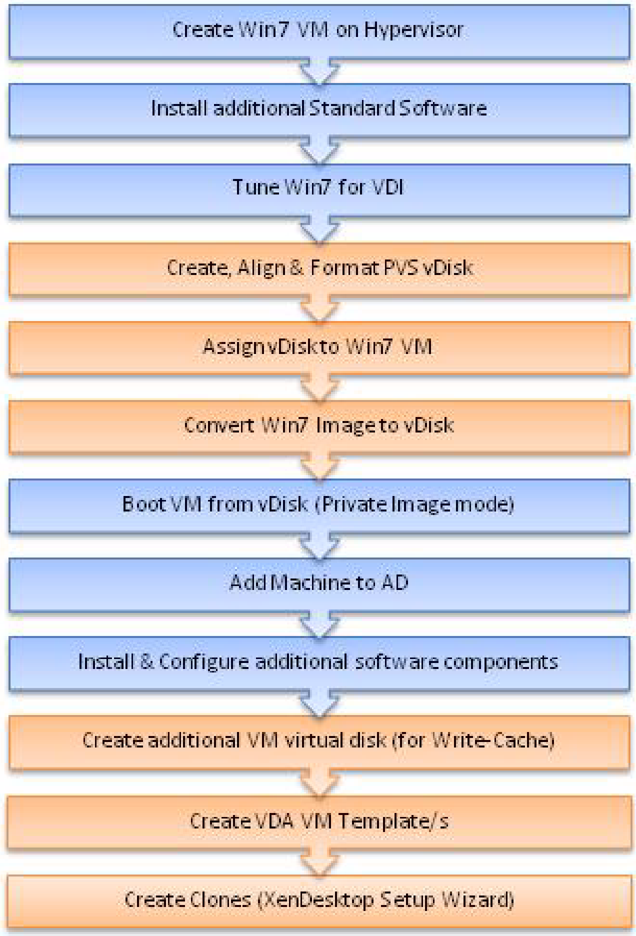


## Create Windows 7 Virtual Machine and Install Standard Software

The following virtual machine configurations and software were used to create the initial Windows 7 virtual machine on the hypervisor which is then later extracted to create a Citrix Provisioning Server vDisk image in .vhd format.

*Table 2 Citrix XenDesktop Configuration*

| **XenDesktop Virtual Desktop Image** | | | |
|---|---|---|---|
| **OS:** | Windows 7 Enterprise 32bit | **Service Pack:** | - |
| **CPU:** | 1 x vCPU | **RAM:** | 1536MB |
| **Disk:** C:\ E:\ | 1 x 16GB (PVS vDisk)<br>1x 3GB Virtual Disk (PVS Write-Cache) | **Network:** | 1 x 1GbE |
| • **Software Installed Prior to cloning to vDisk –**<br>o Citrix Provisioning Server Target Device 5.6 SP1<br>o Microsoft Office Enterprise 2007 SP2<br>o Internet Explorer 8.0.7600.16385<br>o Adobe Reader 9.1.0<br>o Adobe Flash Player 10.0.22 | | | |

### Tuning Microsoft Windows 7 Image for XenDesktop

When many Windows desktops run on a hypervisor it is necessary to try to reduce unnecessary CPU cycles and disk I/O to improve system performance and stability. By turning off unnecessary processes and other unwanted desktop services for instance helps achieve this.

The following configurations were made to the standard image:

- Configure fixed 1.5GB page file
- Configure Networking and Firewall
  - Turn off firewall
  - Set DNS IP addresses for domain
  - Turn off IPV6
- Windows 7 optimization recommendations from the following Citrix blog:

  http://community.citrix.com/pages/viewpage.action?pageId=113247185
  - Recommended "Default User Profile" settings were also applied and copied to "Default User" using the latest Forensic User Profile Manager tool, visit http://www.forensit.com/desktop-management.html
- Citrix PVS TCP Large Send Offload should be disabled on both the PVS server/s and the target device (Windows 7 image). To do this follow the instructions found here:

  http://support.citrix.com/article/CTX117374

### Provisioning Services (PVS) vDisk Creation

When the Windows 7 image has initially been created with the required software, it must be extracted into a Provisioning Server vDisk image. To do this, the Citrix XenConvert 2.1 tool is used which, is part of the PVS Target Device installation.

To create a PVS vDisk:

1. Using the PVS Console (Must use the console from the PVS server)
2. Create new vDisk (16GB) (this may vary depending on requirements).
3. Using Diskpart set the partition offset to 1024. For more information on best practice disk alignment, visit http://support.citrix.com/article/CTX122737.
4. From the PVS server open a command window:

```
C:\>diskpart

DISKPART> list disk

  Disk ###   Status       Size      Free      Dyn   GPT

  --------   ----------   -------   -------   ---   ---

  Disk 0     Online        186 GB       0 B

  Disk 1     Online        16  GB       0 B

DISKPART> select disk 1

Disk 1 is now the selected disk.
```

```
DISKPART> create partition primary align=1024

DiskPart succeeded in creating the specified partition.

DISKPART> Exit
```

To format the vDisk (NTFS):

1. Un-mount the vDisk using the PVS Console.
2. Attach the New vDisk to the Windows 7 Virtual Machine
3. Set the Windows 7 virtual machine to boot from Network.
4. Create a new device in PVS collection and assign MAC address of virtual machine to this PVS object.
5. Assign vDisk and configure following options:
   a. Private Image mode


   b. Manage AD Password


   c. Set device to boot from hard disk

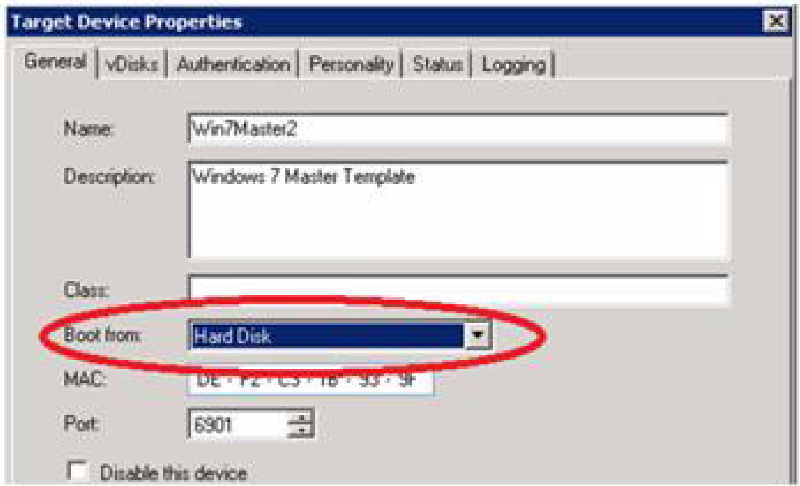


6. Boot Windows 7 virtual machine and check vDisk is attached.

To convert the Windows 7 Virtual Machine Image to Provisioning Services vDisk:

1. To retain the 1024 partition offset in the vDisk the following needs to be added to the C:\Program Files\Citrix\XenConvert.ini:

   [parameters]

   PartitionOffsetBase=1048576

2. Run XenConvert

3. Run PVS Device Optimization Tool by clicking the Optimize button.

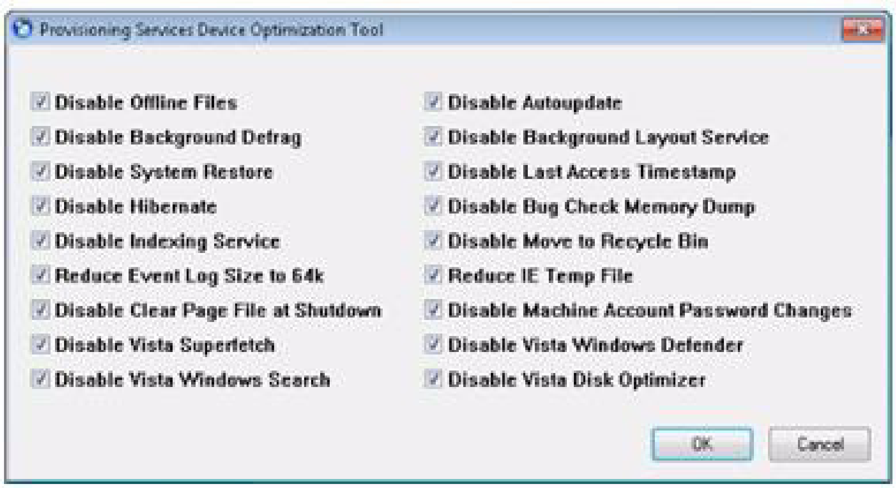


4. Image to assigned vDisk (E:\).

5. When the Imaging process has completed shutdown the virtual machine.

To set the virtual machine to boot from PVS vDisk (rather than vDisk) and start virtual machine:

1. Use the PVS Console to change the Target device options to “boot from vDisk.”

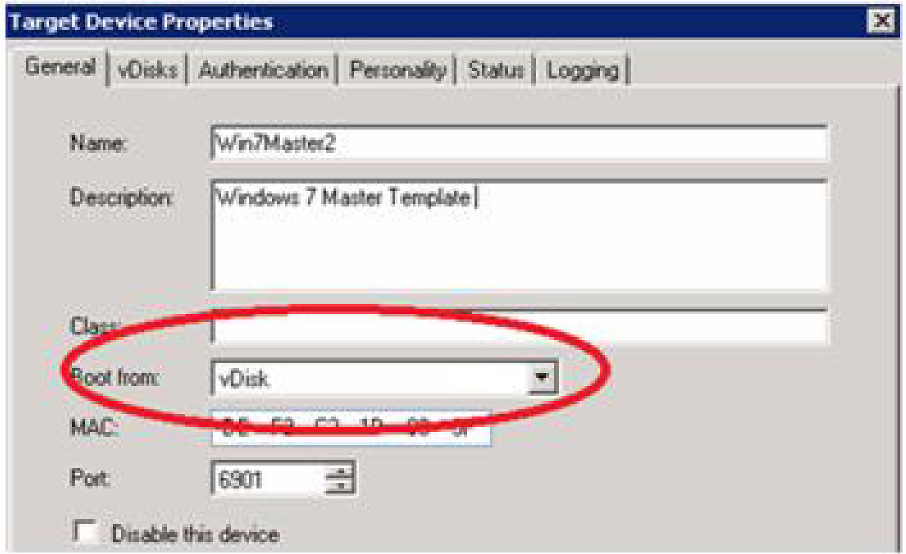


2. Using Hyper-V manager or SCVMM start the virtual machine.
3. Add the host to the domain.
4. Restart Guest OS.

### Install and Configure Additional Software Components

The following software is installed post vDisk cloning:

- Citrix XenDesktop VDA 4.0.5010
- Login VSI 2.1 and STAT Agent (tools used for benchmarking)
- SQL 2K8 Native Client (for STAT agent)

### Add 3-GB Write Cache.VMDK and .VHD to vDisk Image

To match the disk signature, you will need to create and format an additional virtual disk to the Windows 7 image. This will later be detached and used as the default virtual machine template for the vDisk creation process, so that each virtual machine has a unique 3GB virtual disk (E:\ Drive); this is where the client-side PVS Write-Cache will be placed and subsequently all write I/O will be conducted.

To create a new 3-GB Virtual Disk using SCVMM:

1. Create a new virtual disk attached to the Windows 7 virtual machine.
2. Activate the new Disk (Use Standard mode and DO NOT Use Dynamic Mode).
3. Do not format yet.
4. Using Diskpart set the partition offset to 1024.
5. Format the new volume NTFS.
6. Shutdown the virtual machine.
7. Detach the new virtual disk from the virtual machine but do NOT delete it (note where it is stored for next stage below).
8. In the PVS Console change the vDisk Mode to Standard and also change the cache location to be Cache on device's HD.

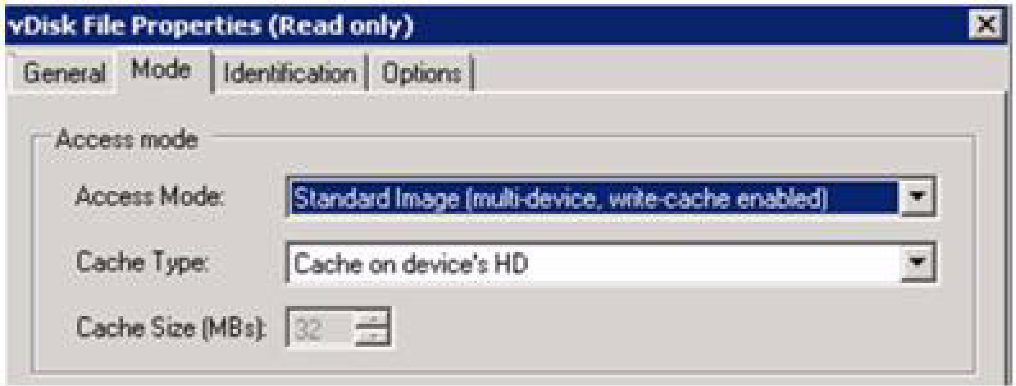


9. Next, the virtual machine templates must be created on the relevant Fibre Channel volume hosted on the NetApp storage. If large numbers of desktops are to be created, it is advisable to mount several FC volumes to the hypervisors balanced between at least 2 NetApp storage controllers.
10. When the FC Volumes have been mounted on the hypervisors, using SCVMM create a Windows virtual machine but do not start it.

To create a new Windows 7 virtual machine (Win7_PVS_Temp):

1. Allocate 1.5 GB RAM.
2. Assign to correct virtual machine Network.
3. Change boot order to Network Boot.
4. Delete assigned Virtual Disk.
5. Attach the Virtual Disk created in the above stage.
6. Convert the virtual machine to a Template.
7. Full Copy - Template to desired NetApp Fibre Channel Volume and name (I.E. Win7PVSTemp (1)).
8. Full Copy - Template to desired NetApp Fibre Channel Volume and name (I.E. Win7PVSTemp (2)).
9. Etc - Until you have a template on each target NFS volume you wish to use.
10. Delete (Win7_PVS_Temp) so that it does not get used accidently.

   Large scale provisioning can be achieved easily by using the XenDesktop Setup Wizard Tool which should be installed on the PVS server.

**Note** The entire XenDesktop infrastructure should be setup and tested prior to converting the base Windows 7 image to a vDisk to help ensure that all the necessary components including active directory membership, are accessibly by this tool.

The aim is to create multiple virtual desktops evenly distributed across all of the available mounted FC data stores, so work out how many you will be creating on each one and then run the XenDesktop Setup Tool.

The XenDesktop Setup Wizard is installed and should be run on the PVS server.

The following steps show how to create the virtual desktop templates by using the XenDesktop setup wizard:

1. Select the XenDesktop Farm.
2. Hosting infrastructure (hypervisor Resource Pool/Cluster).

3. Select the template associated with the Volume you wish to add virtual machine instances. As noted earlier, the XenDesktop setup wizard discards any drives attached to the template.
4. Select the vDisk.
5. Assign Virtual Desktop numbers and Host names.
6. Select desired Organization Unit where machines will be created in AD.
7. Assign Desktops to (existing) Desktop Delivery Controller Group (Group has to be created the first time tool is run).

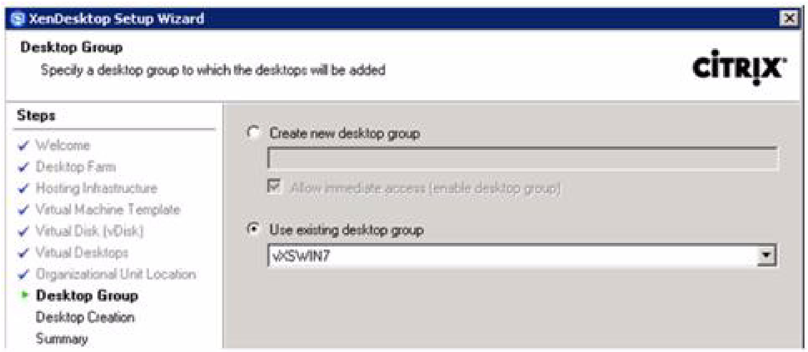


8. Review selections and start creation process
9. When complete, the XenDesktop Setup Wizard should be run again using the same process except a different template should be selected and also start the Virtual desktop numbering from the next available host number (for example, 121 using the example above).

# Citrix XenDesktop Architecture and Design

## Design Fundamentals

There are many reasons for considering 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 are the user classifications:

- 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 workers expect access to all of their applications and data wherever they are.
- External Contractors are increasingly part of your everyday business. They need access to all of your applications and data, yet administrators still have little control over the devices they use and the locations they work from. Consequently, IT needs to adjust 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.

- Road Warriors need access to their virtual desktops 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.

After the user classifications are identified and the business requirements for each user classification are defined, it becomes essential to evaluate the types of virtual desktops that are available based on user requirements.

The following are the potential desktops environments for each user:

- Traditional PC: A traditional PC is what "typically" constituted a desktop environment: physical device with a locally installed operating system.
- Hosted, server-based 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 Desktop: A streamed desktop is a desktop running entirely on the user's local client device. The user interacts with the desktop directly but is 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.

For the purposes of the validation represented in this document, the Hosted VDI desktop model was validated.  The following section provides fundamental design decisions for implementing a Hosted VDI

## Hosted VDI Design Fundamentals

Citrix XenDesktop can be used to deliver a variety of virtual desktop configurations. The following are some high-level design considerations that can be used when evaluating a Hosted VDI deployment:

### Hypervisor Selection

Citrix XenDesktop is hypervisor agnostic, so any of the following hypervisors can be used for hosted VDI desktops:

- XenServer

  Citrix® XenServer® is a complete, managed server virtualization platform built on the powerful Xen® hypervisor. XenServer is designed for efficient management of Windows® and Linux® virtual servers and delivers cost-effective server consolidation and business continuity. More information on XenServer can be obtained at the company website.

- vSphere

  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, HA, Storage vMotion, VMFS, and a mutlipathing storage layer. More information on vSphere can be obtained at the company website.

- 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 website.

### Provisioning Services

Hosted-VDI desktops can be deployed with or without Citrix Provisioning Sevices, 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 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. 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 local 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.
- Ram Cache: 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.
- Server Cache: 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.
- Difference Disk Mode: Difference Cache is in a file on the server, or on a share, SAN, or other. 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. It is slower than RAM cache and Server

### Designing a Citrix XenDesktop Deployment

For detailed information about configurations, architecture, and design recommendations for delivering virtual desktops with XenDesktop, refer to http://support.citrix.com/proddocs/index.jsp?topic=/xendesktop-bdx/cds-admin-deploy-plan-wrapper-bdx.html.

# Solution Validation

This section details the configuration and tuning that was done to various components for a complete solution validation. Some of the design principles behind the architecture are highlighted below:

- The design includes a scalable and modular building block that can be validated internally.
- Run all the XenDesktop infrastructure components, SCVMM, and other supporting services like AD/DNS/DHCP in a Virtual Environment, to highlight "Cisco UCS is the platform for Virtualization."
- The resilient architecture should be able to scale as well as maintain redundancy, maximum bandwidth and robustness.
- For example: Four links from each IOM to FI, Dual SAN Fabric, Multiple links to Network, vPC, Dual paths to Storage, and dual storage controller (for maximum redundancy).
- Architecture based on the already established best practices for all the partner components involved in the reference architecture.

## Reference Architecture Configuration

Figure 29 shows the architecture block diagram for the implementation of Citrix XenDesktops on Cisco UCS and Netapp storage.

*Figure 30 Architecture Block Diagram*

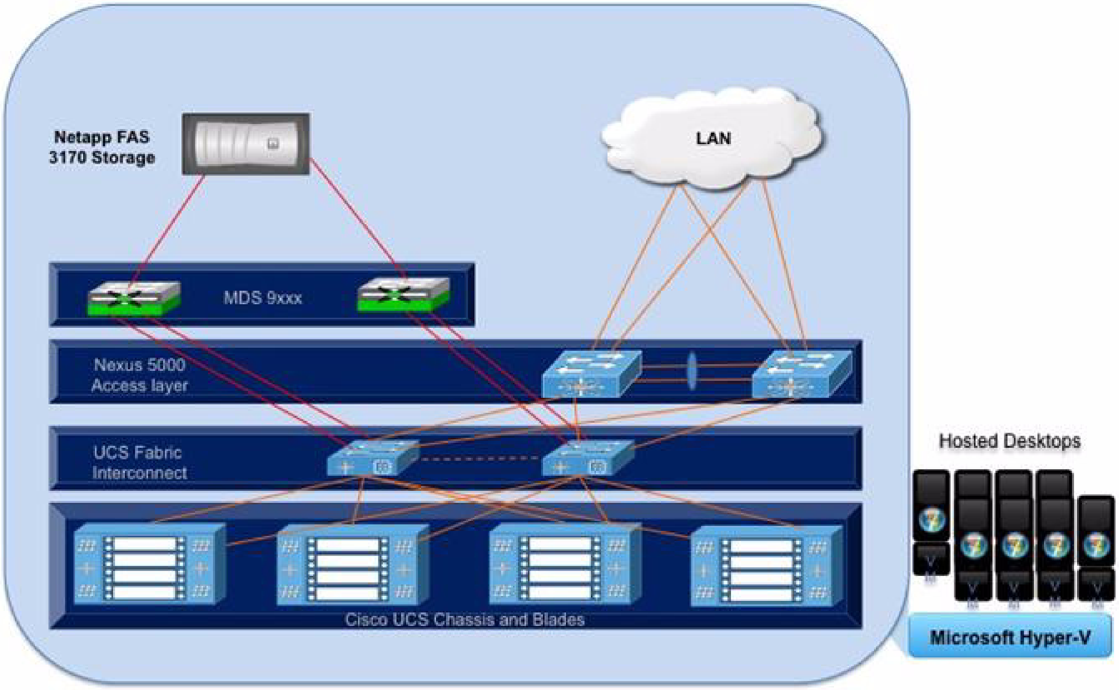


The architecture block diagram above is distinctly divided into four layers:

- Cisco UCS Compute platform
- The virtual desktop infrastructure that runs on Microsoft Hyper-V virtualization platform
- Network Access layer and LAN
- Storage Access (SAN) and Storage array

In Figure 29 , the red lines from the UCS Fabric Interconnect represent the 4 GB FC connectivity to upstream MDS 9x00 switch and the yellow lines represent the 10G Ethernet connectivity.

Figure 30 showcase logical configuration of the test setup for the reference architecture. Some of the key highlights include:

1. For the four chassis setup we are using 4 links from each chassis comprising of B250-M2 with dual Cisco VIC cards. It is possible to run the configuration with two links and get similar results.
2. We are using 4 GB X 48 DIMMS running at 1333 MHz, one could use 8 GB X 24 DIMMS to get 192 GB of memory.
3. Four links per chassis were used in our configuration and four 10G uplink with two ports in a VPC (refer to the N5K config in Appendix A). One could easily connect up to 10 chassis with Nexus 6120 FI with two links per chassis and two 10G uplink ports (1 each for bow-tie configuration) to achieve similar results.

*Figure 31 Detailed Architectural Diagram for Four Chassis Test Setup*

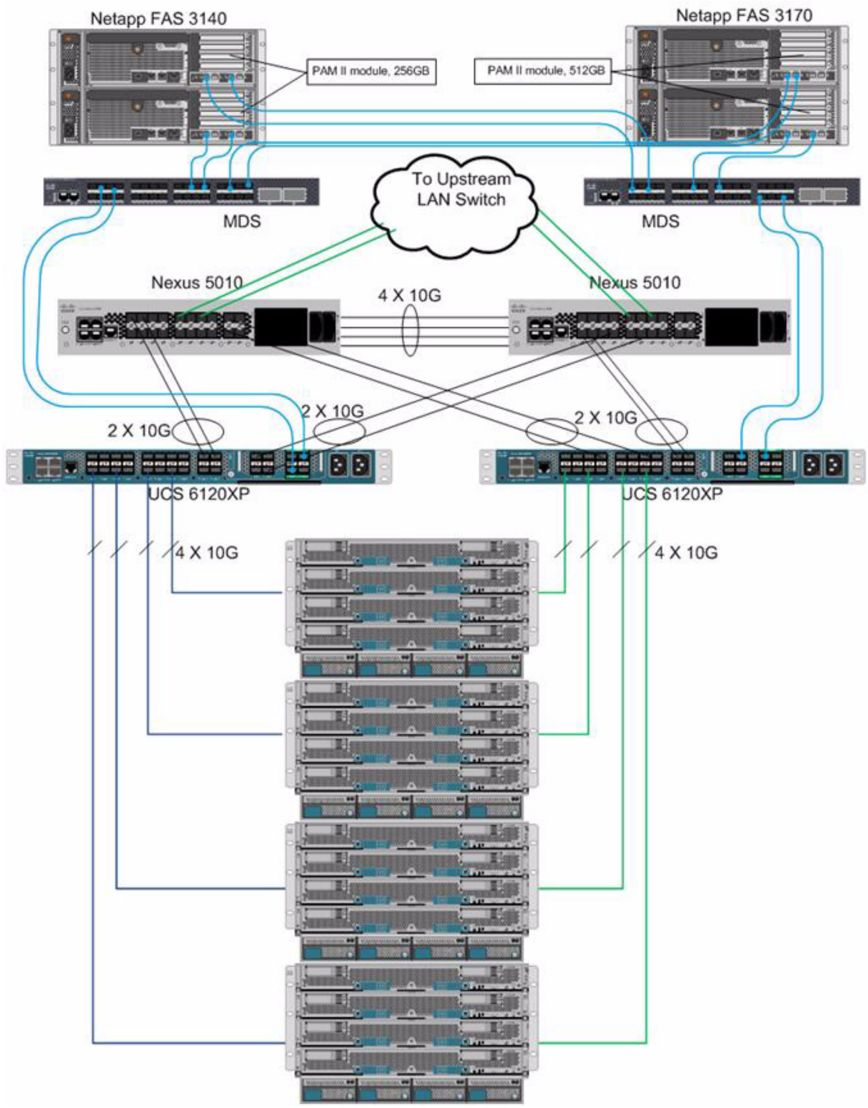


# Cisco Unified Computing System Configuration

This section details the Cisco Unified Computing System configuration that was done as part of the infrastructure build out. The racking, power and installation of the chassis are described in the install guide (refer to http://www.cisco.com/en/US/docs/unified_computing/ucs/hw/chassis/install/ucs5108_install.html) and is beyond the scope of this document. More details on each step can be found in the following documents:

- Cisco Unified Computing System CLI Configuration Guide
http://www.cisco.com/en/US/docs/unified_computing/ucs/sw/cli/config/guide/1.3.1/b_CLI_Config_Guide_1_3_1.html

- Cisco UCS Manager GUI Configuration Guide

  http://www.cisco.com/en/US/docs/unified_computing/ucs/sw/gui/config/guide/1.3.1/b_UCSM_GUI_Configuration_Guide_1_3_1.html

## Cisco UCS General Setup

To configure the Cisco Unified Computing System, do the following:

1. Bring up the Fabric interconnect (FI) 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. Once 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, either 1, 2 or 4 links depending on your application bandwidth requirement. We chose to connect all the four links.
2. Now connect using your favorite 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.
3. Check the firmware on the system and see if it is current. The firmware used in our testing was 1.3(1i).

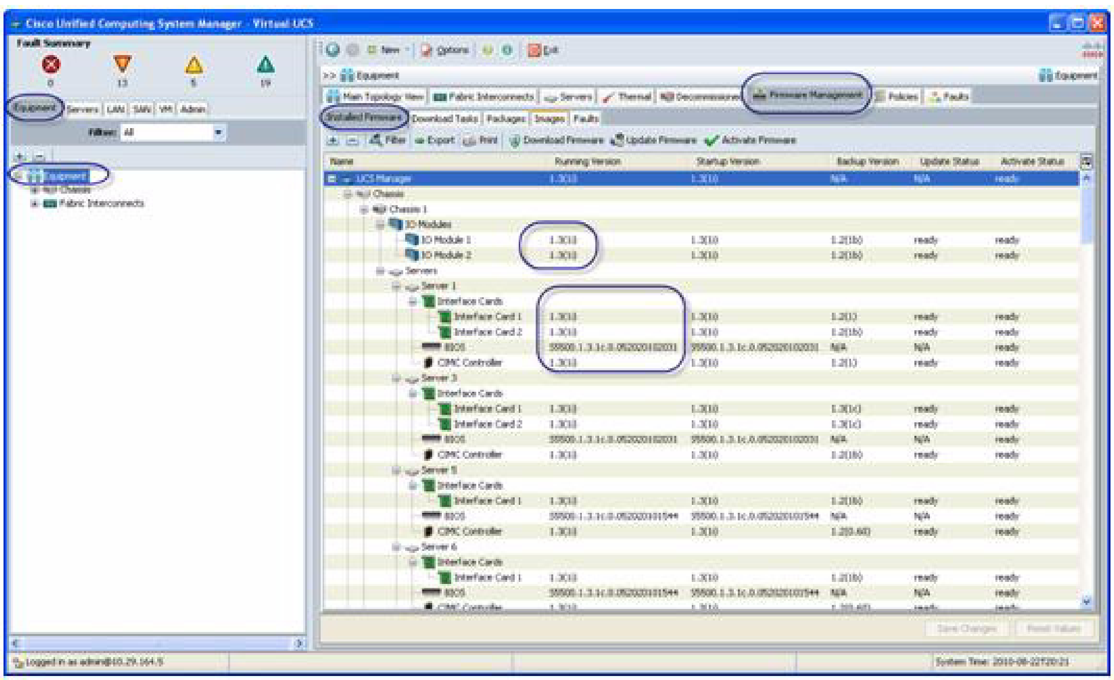

4. If the firmware is not current, follow the installation and upgrade guide to upgrade the Cisco UCS firmware. Also do not forget to upgrade the BIOS to the latest level and associate it with all the blades.

5. Configure and enable the server port on the FI. To bring the chassis online acknowledge the chassis. The Fabric interconnect is configured in End Host Mode.


6. Configure and enable upstream Ethernet links and Fibre Channel links.

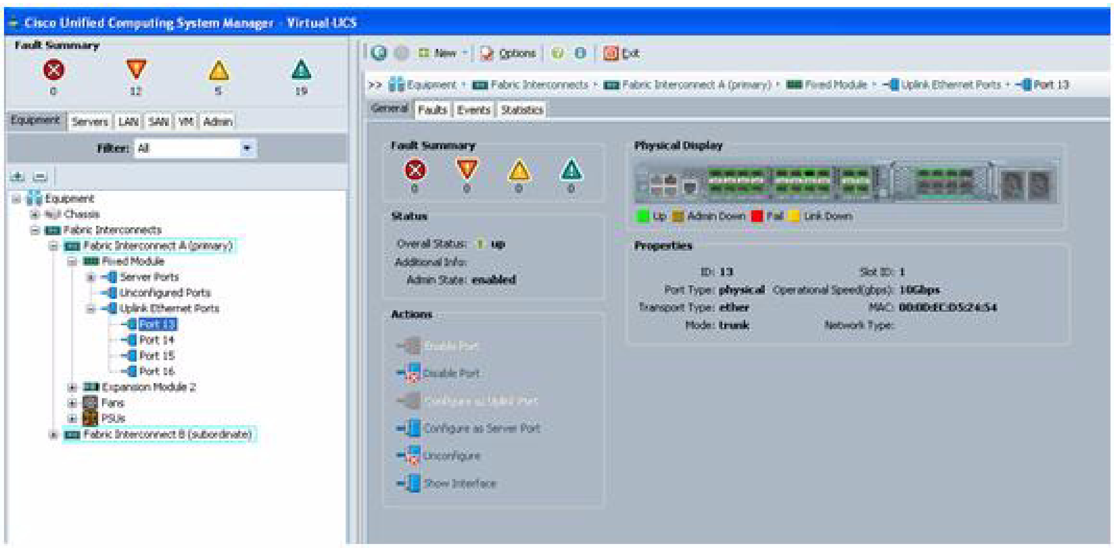


7. When the blades are discovered, it is time to set the KVM IP addresses for each of the blades. This is done through the admin tab ' communication management ' Management IP address pool. One has to make sure we have ample IP address for all the blades and make sure the gateway and netmask is set correctly.

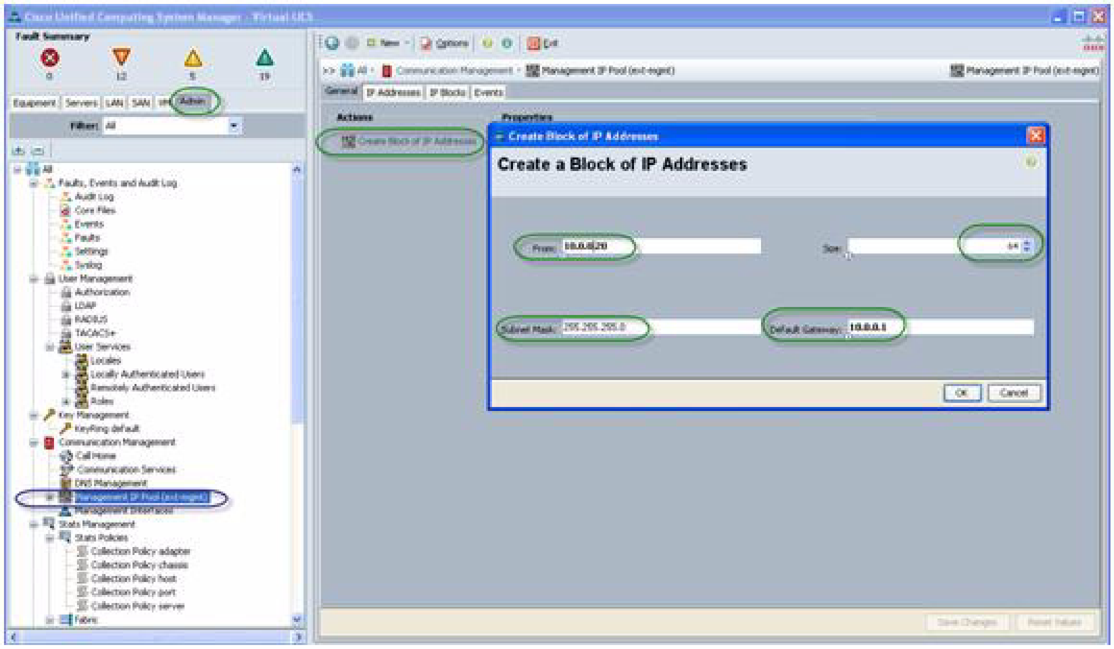


## Cisco UCS Service Profile Template and Service Profile Configuration

To configure the Service template and profile, do the following:

1. Create the pools: MAC pool, WWPN pool, WWNN pool, UUID pool, and Server pool.
   a. MAC pool:

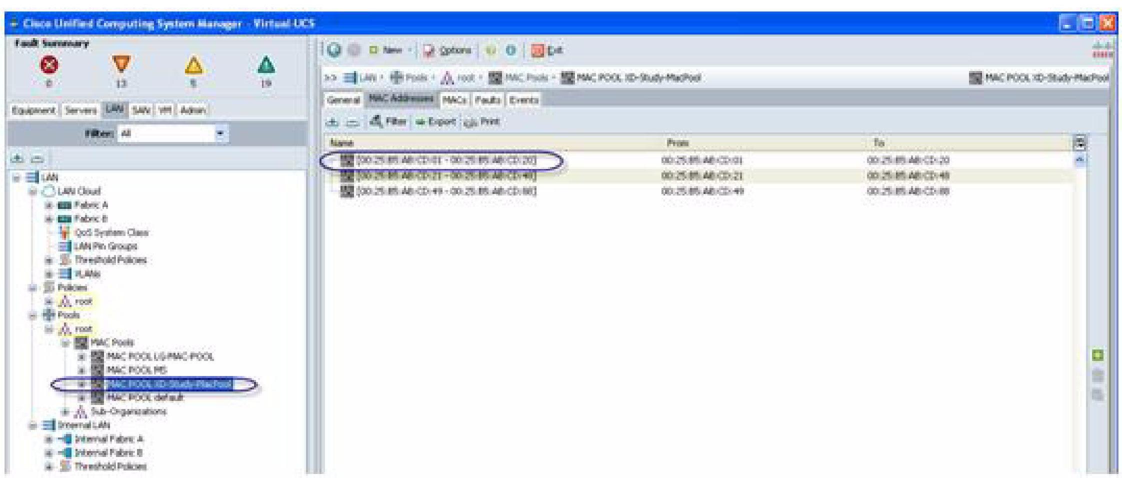

**b.** WWPN pool:

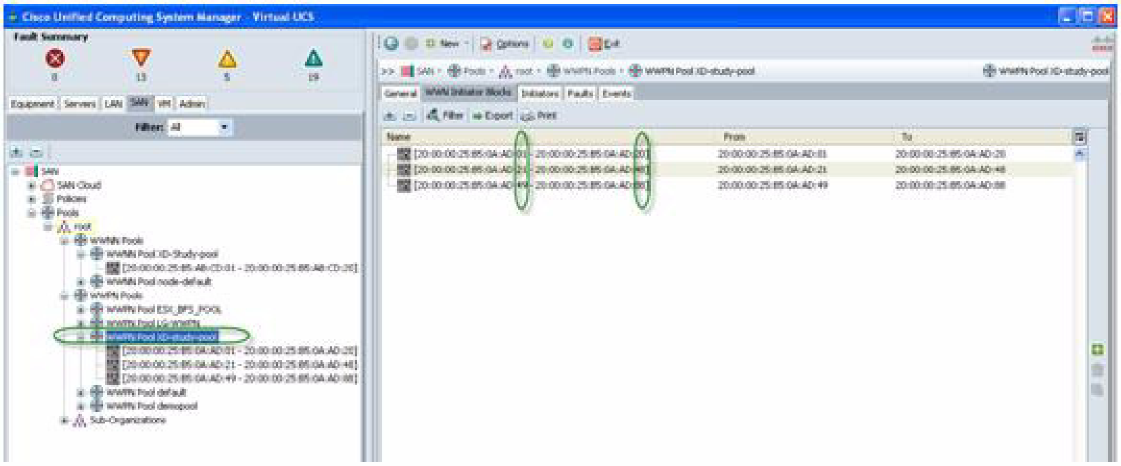

**c.** WWNN pool:

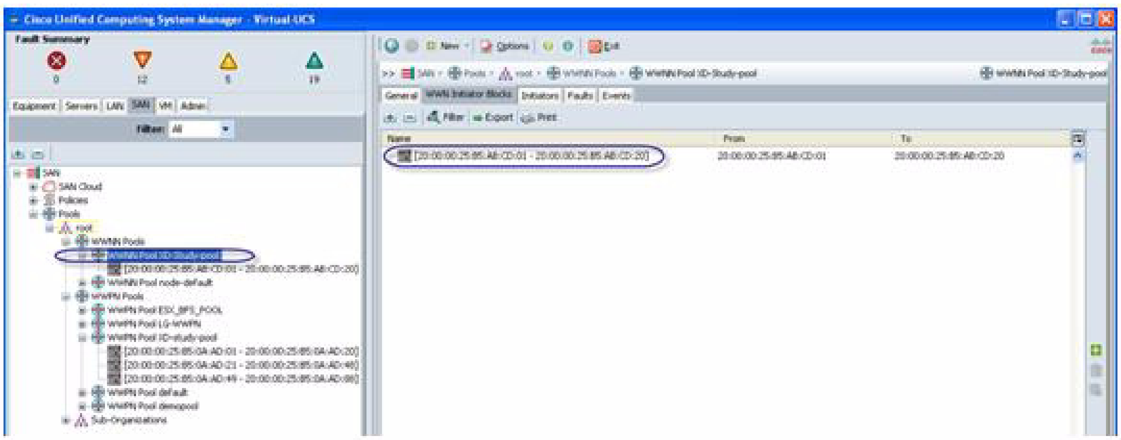

**d.** UUID pool:

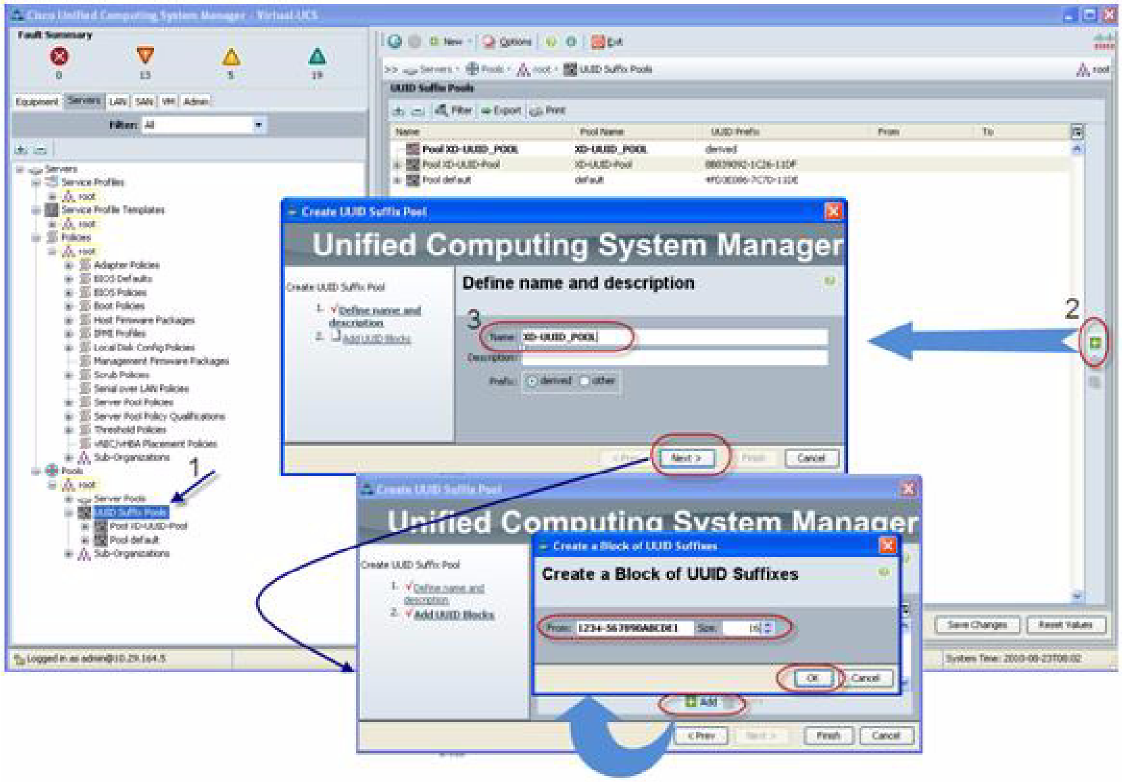
Unified Computing System Manager
Define name and description
Create a Block of UUID Suffixes
Create UUID Suffix Pool
1
2
3

e. Server pool:

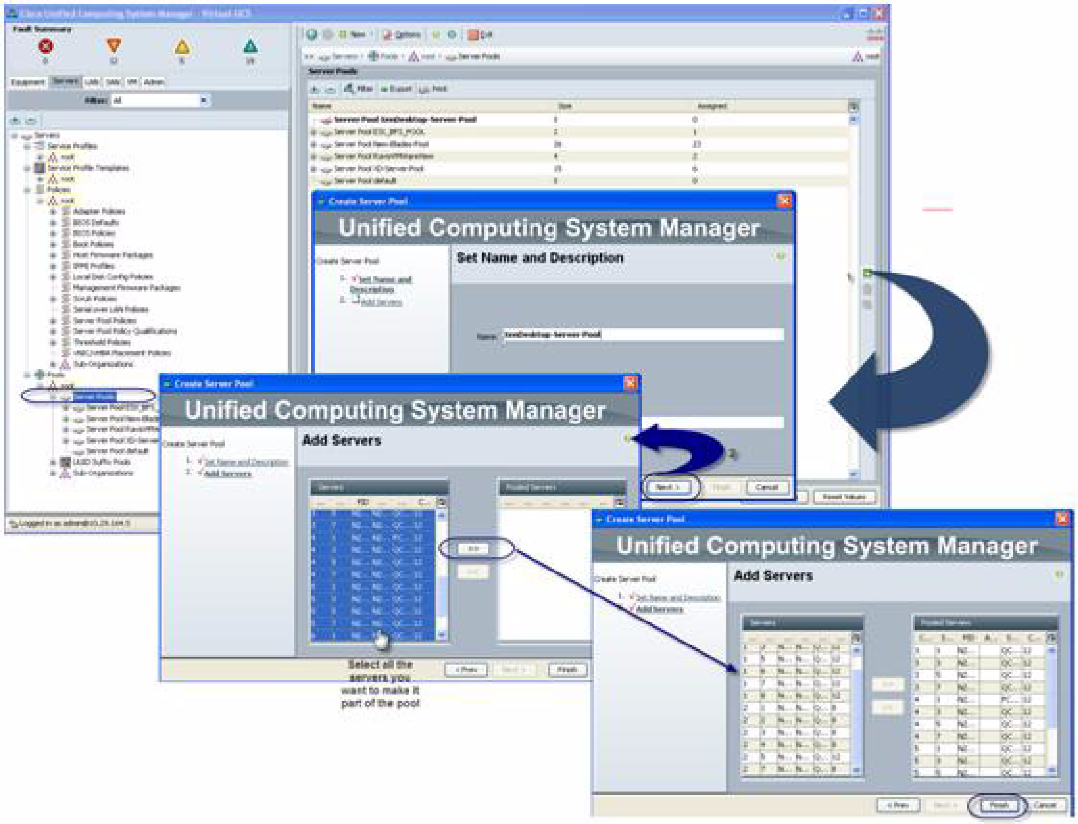


2. Create the vHBA template.

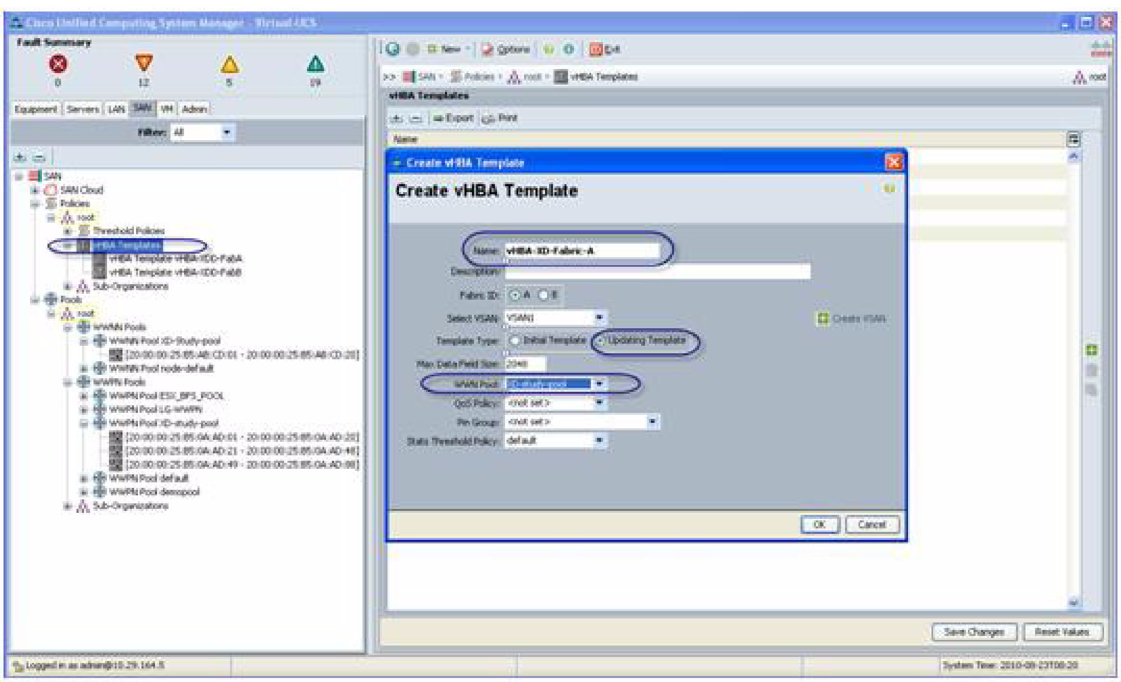

3. Create the vNIC template.

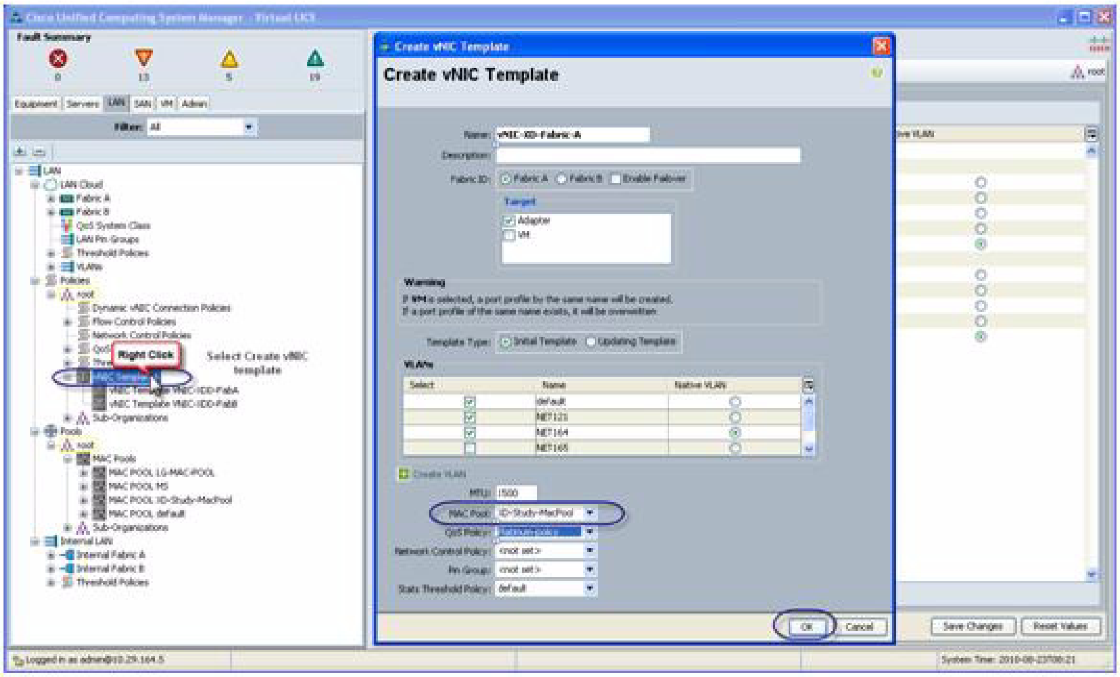


4. Create boot from SAN policies.
5. Create a service profile template using the pools, templates, and policies you previously configured.
6. After associating a server pool to the service profile template, right- click to deploy as many service profiles as you need and Cisco UCS Manager will automatically start configuring these new service profile templates on the selected blade servers.
7. At this point, the servers are ready for OS provisioning. It is recommended to set up a PXE server for the OS install. Virtual media CD based OS installation is also possible.

**Note** Regarding Low Voltage Memory DIMMS: When working with 4 GB -1333MHz DDR3 Low voltage dual rank DIMM it will show up as 1066 if you do not set the performance mode in the policy. To identify the memory DIMM properties go to the equipment tab and follow the steps detailed below.

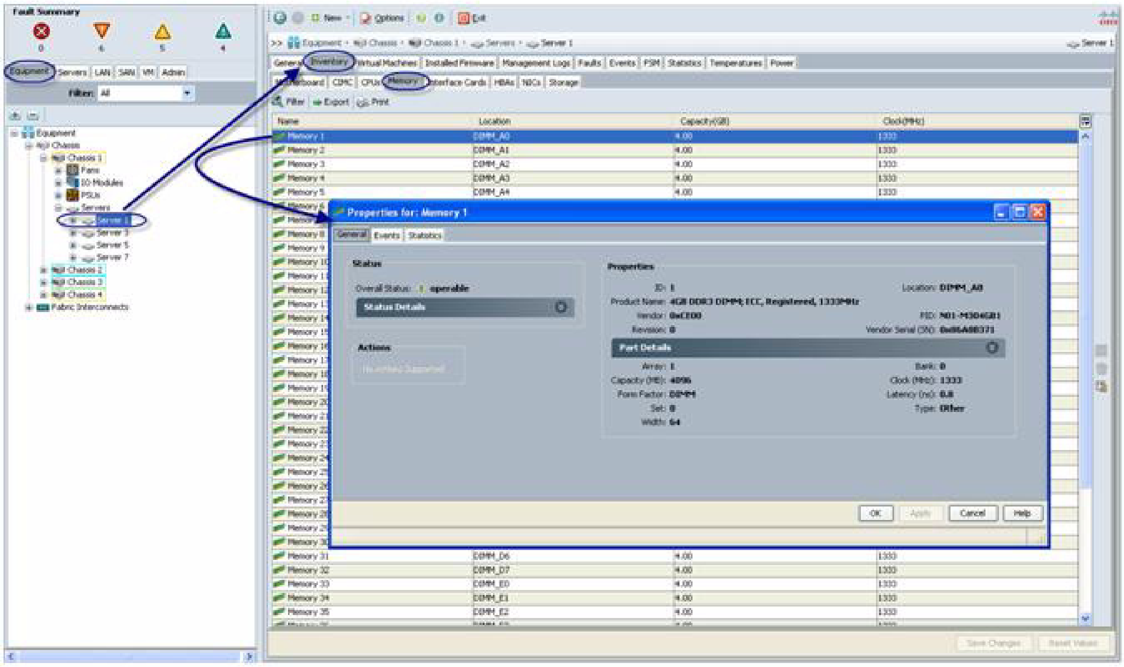

To configure the BIOS policy to make the low voltage DIMMS perform at full 1333 MHz, do the following:

1. Set the performance mode in the BIOS policy (now controlled from the UCSM 1.3(1i)) and reboot the server to take the effect and you will have 1333MHz speed on the memory DIMMs.
2. Configure a BIOS policy for platform-default and change the default power-saving-mode to performance-mode.

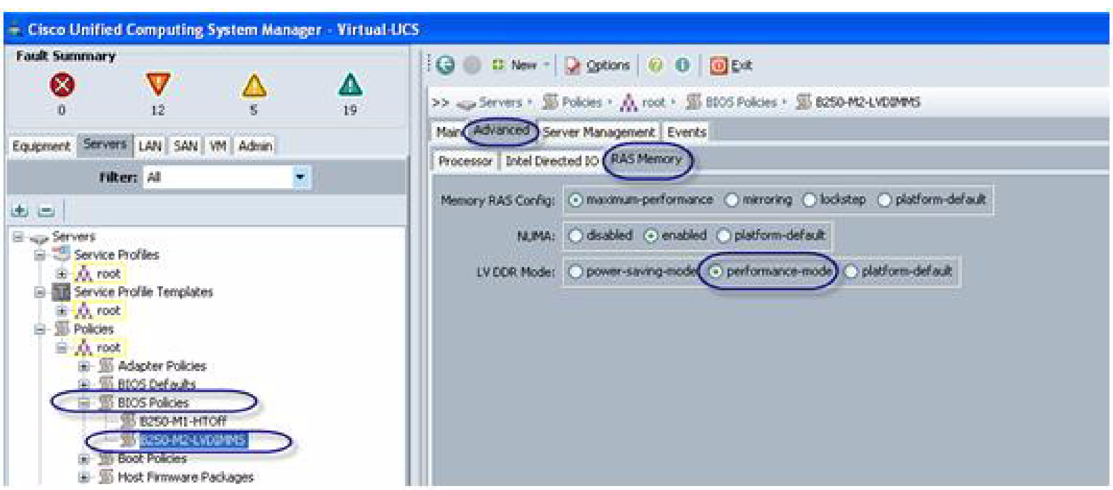


3. Add the following as a policy to the template:

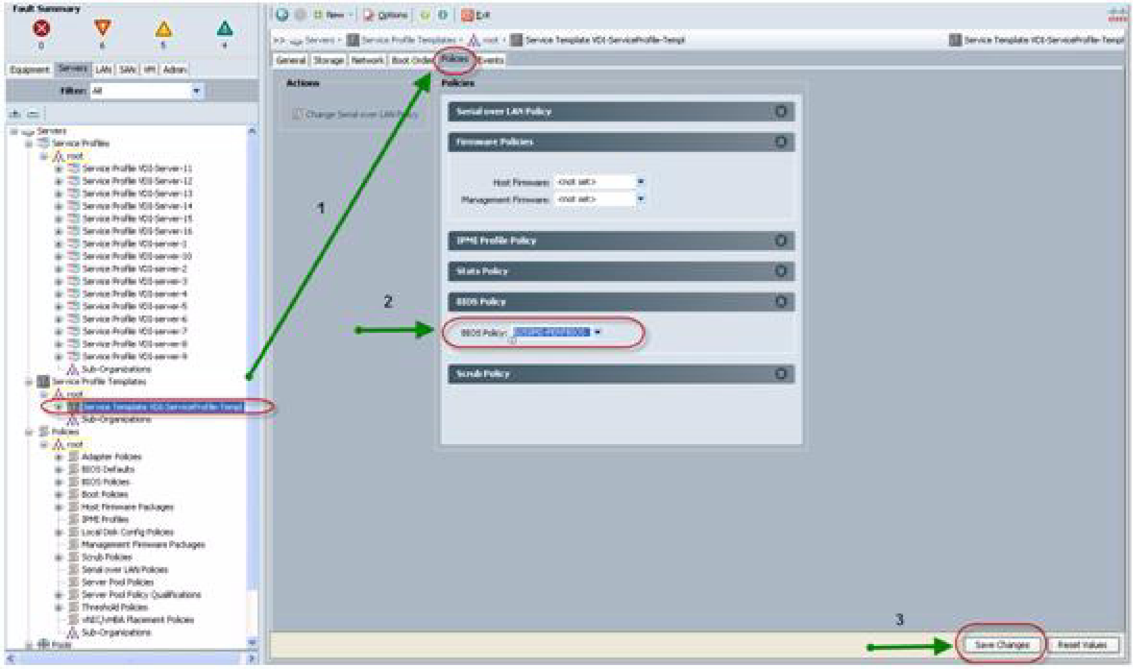


This will reboot the servers and when the servers come back up the memory DIMMs will be in 1333MHz.

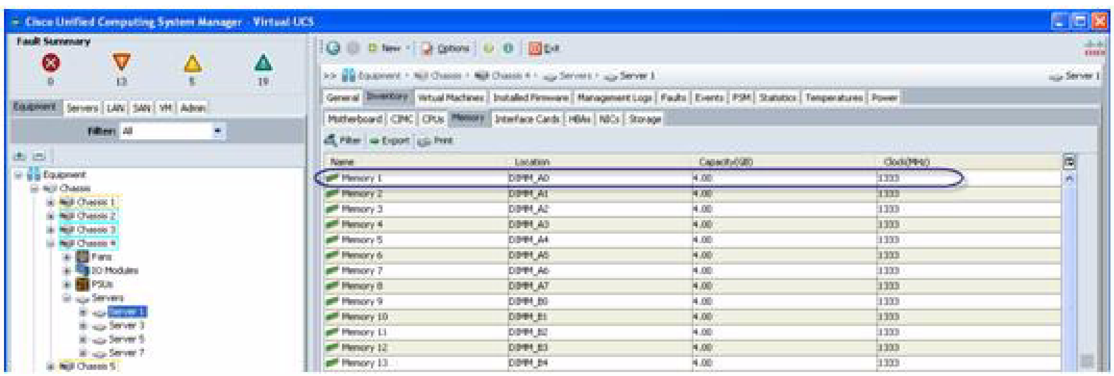

## QOS and COS in Cisco Unified Computing System

Cisco Unified Computing System provides different system classes 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 such as 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 the entire system interfaces have defined QoS rules.

- By default the 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 four or 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)
- Maximum MTU per class allowed is 9216.
- Using the Cisco Unified Computing System, 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):

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

## Cisco UCS System Class Configuration

The Cisco Unified Computing System defines user class names as follows.

- Platinum
- Gold
- Silver
- Bronze

*Table 3 Name Table Map Between Cisco Unified Computing System and Cisco NX-OS Software*

| Cisco UCS Names | Cisco NX-OS Names |
|---|---|
| Best effort | Class-default |
| Fibre Channel | Class-fc |
| Platinum | Class-Platinum |
| Gold | Class-Gold |
| Silver | Class-Silver |
| Bronze | Class-Bronze |

*Table 4 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 5 Default Weight in Cisco Unified Computing System*

| Cisco UCS Class Names | Weight |
|---|---|
| Best effort | 5 |
| Fc | 5 |

To enable QOS on the Cisco Unified Computing System, do the following:

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

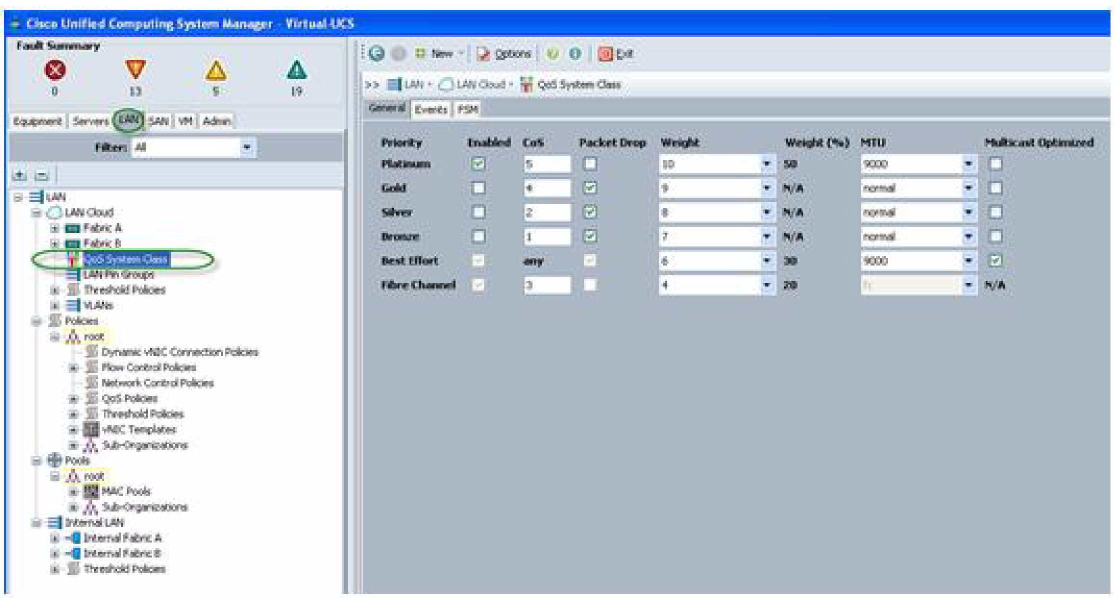


3. In the LAN tab under policies, define a platinum policy and select platinum as the priority.

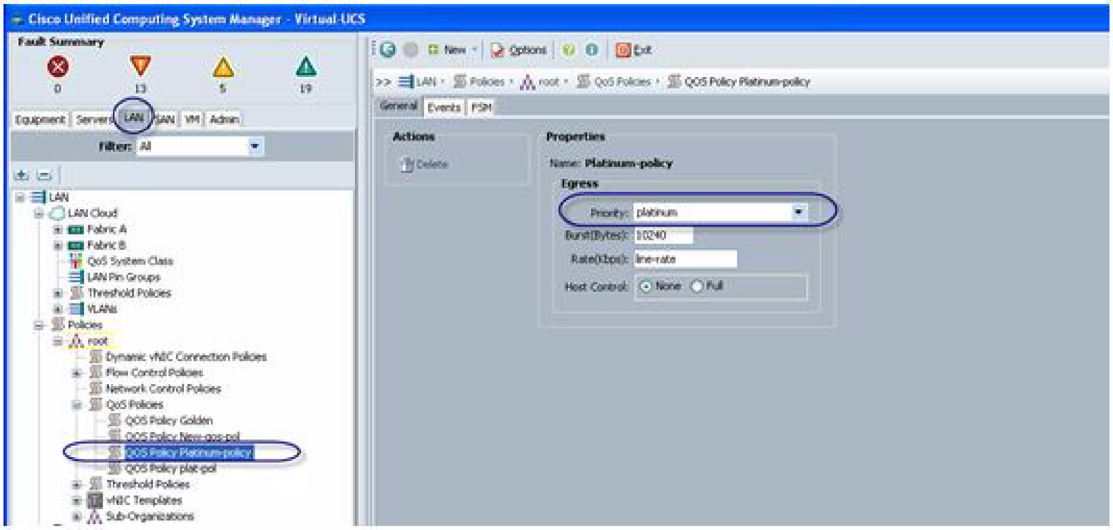


**4.** Include this policy into the vNIC template under the QoS policy.

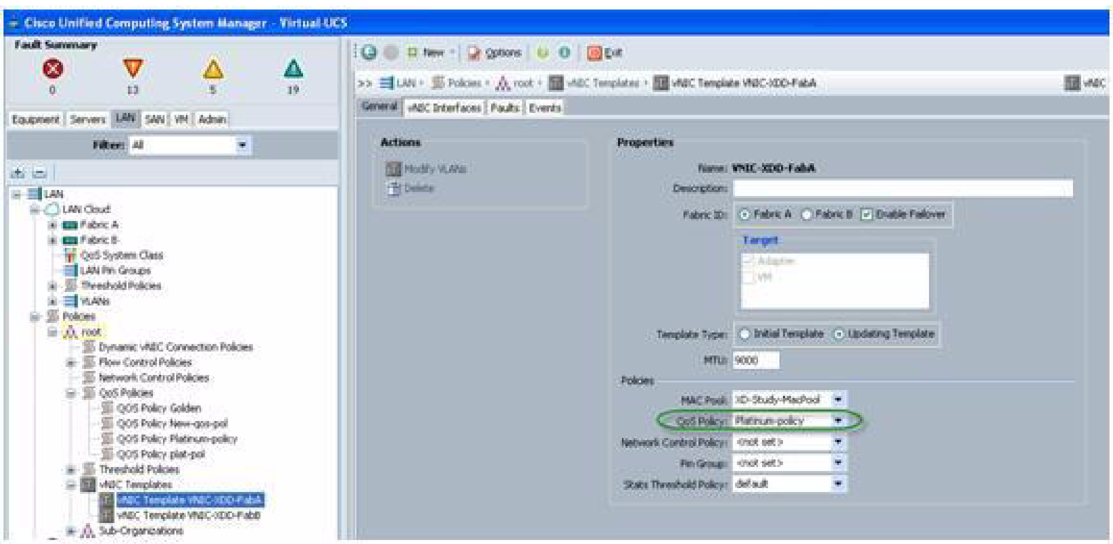


This is a unique value proposition of the Cisco Unified Computing System with respect to end-to-end QOS. This helps ensure the network ports have the highest service level guarantee with Platinum policy and Jumbo frames for obtaining an end-to-end QOS and performance guarantee. Optionally, you can configure the NIC to have a no-drop class along with the platinum policy as shown above.

# Citrix XenDesktop Configuration

Figure 31 shows the Citrix XenDesktop configuration.

*Figure 32 Citrix XenDesktop Configuration*

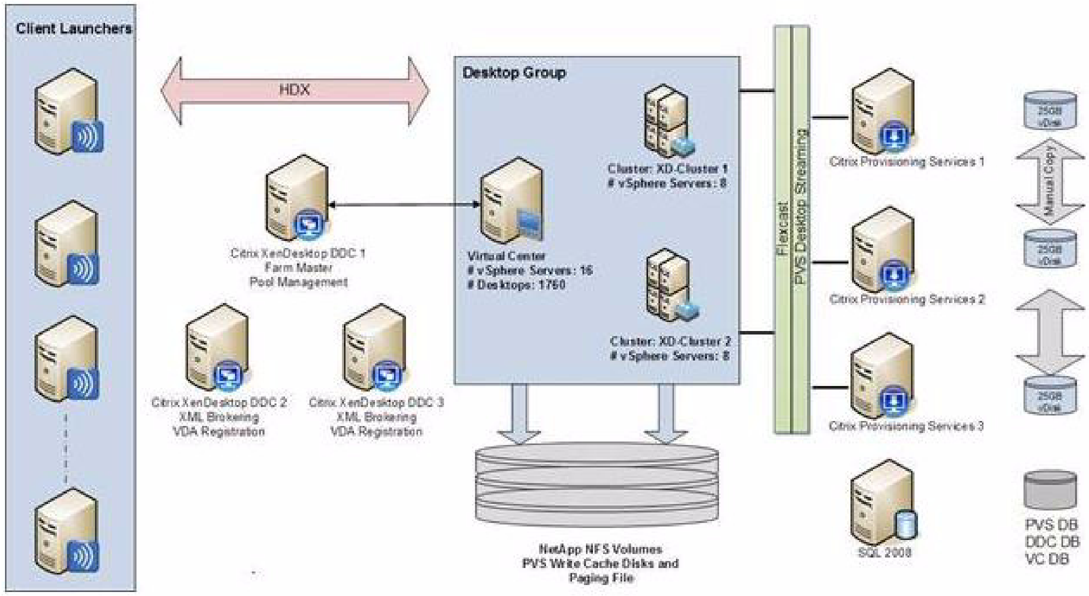


- Environment Summary:
  - 2 Desktop Delivery Controllers
  - 2 Provisioning Services Servers for 8 blade and 3 Provisioning Services Servers for 16 blade test
  - 2 SCVMM Servers
  - 2 Microsoft Hyper-V Clusters
  - 1760 Virtual Desktops
  - 1 Citrix Licensing Server
  - 1 File Server for Roaming Profiles and VSI data
  - 1 SQL 2008 Server for DDC, PVS and VMM database
  - 2 NetApp Storage Controllers,
  - Multiple client launchers

An individual breakdown of the configuration by component is shown in the following tables:

*Table 6 Microsoft Hyper-V Servers*

| Microsoft Server 2008 R2 Hyper-V | | | |
|---|---|---|---|
| **Hardware:** | Cisco UCS B-Series Blade Server | **Model:** | B250 –M2 |
| **OS:** | Microsoft Windows 2008 Server R2 Hyper-V | **Service Pack:** | - |
| **CPU:** | 2 x 6 Core Intel 5680 @ 1333 GHz<br>(24 Logical Cores Total) | **RAM:** | 192 GB @<br>1333 MHz |
| **Disk:** | Boot From SAN | **Network:** | 4 x 10GbE |

*Table 7 Citrix Provisioning Server 5.6*

| Citrix Provisioning Server 5.6 | | | |
|---|---|---|---|
| **OS:** | Windows 2008 Enterprise R2 64bit | **Service Pack:** | - |
| **CPU:** | 4 x vCPU | **RAM:** | 8192MB |
| **Disk:** | 1 x70GB Virtual Disk<br>(hosted on FC volume on NetApp Storage) | **Network:** | 1 x 1GbE |
| • Database for PVS hosted on separate Microsoft SQL Server 2008 64bit | | | |

*Table 8 Citrix XenDesktop Desktop Delivery Controller*

| Citrix XenDesktop DDC | | | |
|---|---|---|---|
| **OS:** | Windows 2003 R2 Enterprise 64bit | **Service Pack:** | 2 |
| **CPU:** | 2 x vCPU | **RAM:** | 4096MB |
| **Disk:** | 1 x50GB Virtual Disk<br>(hosted on FC volume on NetApp Storage) | **Network:** | 1 x 1GbE |
| • Citrix XenDesktop DDC - 400W2K3X64004<br>o Desktop Delivery Controller - Services Hotfix XD*400DDC64007<br>o Pool Management Service Hotfix XD*400AMC002<br>• Database for DDC hosted on separate Microsoft SQL Server 2008 64bit | | | |

*Table 9 Citrix License Server*

| Citrix License Server | | | |
|---|---|---|---|
| **OS:** | Windows 2008 R2 Enterprise 64bit | **Service Pack:** | - |
| **CPU:** | 1 x vCPU | **RAM:** | 2048MB |
| **Disk:** | 1 x50GB Virtual Disk<br>(hosted on FC target volume on NetApp Storage) | **Network:** | 1 x 1GbE |

*Table 10 ICA Client Hosts*

| ICA Client Hosts (VSI Launchers) | | | |
|---|---|---|---|
| **OS:** | Windows 2003 R2 Enterprise 64bit | **Service Pack:** | 2 |
| **CPU:** | 2 x vCPU | **RAM:** | 4096MB |
| **Disk:** | 1 x40GB Virtual Disk<br>(hosted on FC target volume on NetApp Storage) | **Network:** | 1 x 1GbE |

## Citrix XenDesktop Desktop Delivery Controller (DDC)

The DDCs were virtualized on Hyper-V and some of the roles of the DDC were assigned to specific DDC machines, an approach commonly taken in Citrix XenDesktop and XenApp deployments.

The DDCs were configured as follows:

- DDC 1: Farm Master and Pool Management
- DDC 2: VDA Registrations and XML Brokering

In this environment, 2 DDCs (2vCPU, 4GB RAM) easily sustained the farm size of 1760 desktops and proved stable during all stages of testing.

In order to integrate the DDC with the VMM server, PowerShell and the VMM Administrative console were required to be installed on each controller in the farm.

## Farm Configuration

In addition to the standard XenDesktop farm installation, the following additional items were configured or installed:

- Installed Citrix Pool Management Hotfix XDE400PM004
- Installed Citrix Desktop Delivery Controller Hotfix DDCE400W2K3X64007
- Installed Citrix Delivery Services Console Hotfix XDE400AMC002
- Created XenDesktop policy to disable client printer mappings
- Configured DDC1 as Farm Master and Pool Management as per CTX117477
- Configured DDC2 for Registrations and XML Brokering as per CTX117477
- Installed PowerShell and VMM Administrative console on both DDC servers.
- Created two Desktop Groups for two Hyper-V clusters; one each per VMM Server instance.

By default, Pool Management will attempt to start 10 percent of the total pool size. For the 8 blades, 880 desktop tests, the VMM, DDC and PVS servers were able to successfully handle the desktop boot up with the default configuration (88 simultaneous connections). For the 16 blades, 1760 desktops test, this default configuration was changed and the numbers of concurrent requests were throttled to 20 by editing the Pool Management Service configuration file: C:\Program Files\Citrix \VMManagement\CdsPoolMgr.exe.config

- Modify the <appSetting> section by adding the line:
- <add key="MaximumTransitionRate" value="20"/>
- The Pool Management service needs to be restarted to read the new configuration.

Please note that the value of 20 was achieved after attempting higher values of 100, 80 and 40 and observing a few desktops with TFTP time out during boot up, indicating that the environment is unable to handle that many simultaneous requests under high load scenario. To achieve optimal results, customers should fine tune this setting only after conducting tests in their environment.

## Provisioning Services Configuration

For the scaled-out test, two Provisioning Servers supported 880 Windows 7 desktops and three PVS servers supported 1760 desktops. The following items represent additional changes to the environment after the initial default installation:

- Changed the Threads per Port from the default 8 to 32. This is necessary when streaming to high amounts of target devices.
- Configured the bootstrap file to contain the static IP address assigned to each of the provisioning servers.
- Created a local vDisk store for each of the Provisioning Servers and configured it to the E: drive.
- Copied the 25GB Windows7 vDisk to each Provisioning Server's E: drive.

## Citrix Provisioning Services

Citrix Provisioning Server (PVS) is part of the XenDesktop Enterprise and Platinum suites and was used in all tested scenarios, this allows 1000's of virtual machines hosted on hypervisor servers to PXE boot from and share a single gold Windows 7 Image.

## Citrix Provisioning Server (PVS) for Use With Standard Desktops

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

- Virtual desktops are then configured to PXE boot on Hypervisor server
- PVS streams the vDisk image on start to the Hypervisor, and is loaded into RAM
- PVS injects a Security Identifier (SID) and host name as each desktop boots to make them unique in AD. These object mappings are maintained and managed within the PVS server and are visible in the PVS Console under "Collections" view are initially created and mapped by the XenDesktop Setup tool.

**Note** Using CIFS to host the vDisk is not recommended by Citrix; although a Read Only ISCSI target mode can now be used and managed with PVS 5.8, for testing a copy of the vDisk was hosted and maintained on each PVS server's local disk to provide high availability and load balancing by all servers within the farm. As the PVS servers are assigned with 8GB RAM the image will remain persistent and be serviced by RAM after it is initially served for the first time by each server.

PVS servers can be configured in a farm to provide high availability and resilience; connections are automatically failed over to a working server/s within the farm in the event of a failure without interruption to the desktop.

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. This is where ALL write I/O is conducted for the given virtual desktop instance, it is therefore 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 devices HD option was used, a 3GB virtual disk was assigned to the virtual machine templates used in the virtual desktop creation process. By creating the 3GB drives associated with the templates on FC volumes, mounted on the hypervisors; we were then able to create VDI instances each with its own 3GB drive where the PVS Write Cache was placed. In addition the PVS Target device agent was installed in the Windows 7 image and automatically placed the Windows swap file on the same drive when this mode was enabled.

Therefore both the PVS Write Cache and Windows Swap file were hosted on an NFS mounted volumes hosted on NetApp storage. Scalability load balancing across multiple Volumes and storage Controllers was increased by using 4 virtual machine templates (each created on different data stores/storage repositories) and running the XenDesktop Setup Wizard tool 4 times using a different virtual machine template for each process.

Figure 32 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 provided FC volumes where the PVS cache is placed. This helped to ensure that all write I/O took place on the NetApp storage over FC using high performance storage.

*Figure 33 vDisk Hosting on NFS Volumes*

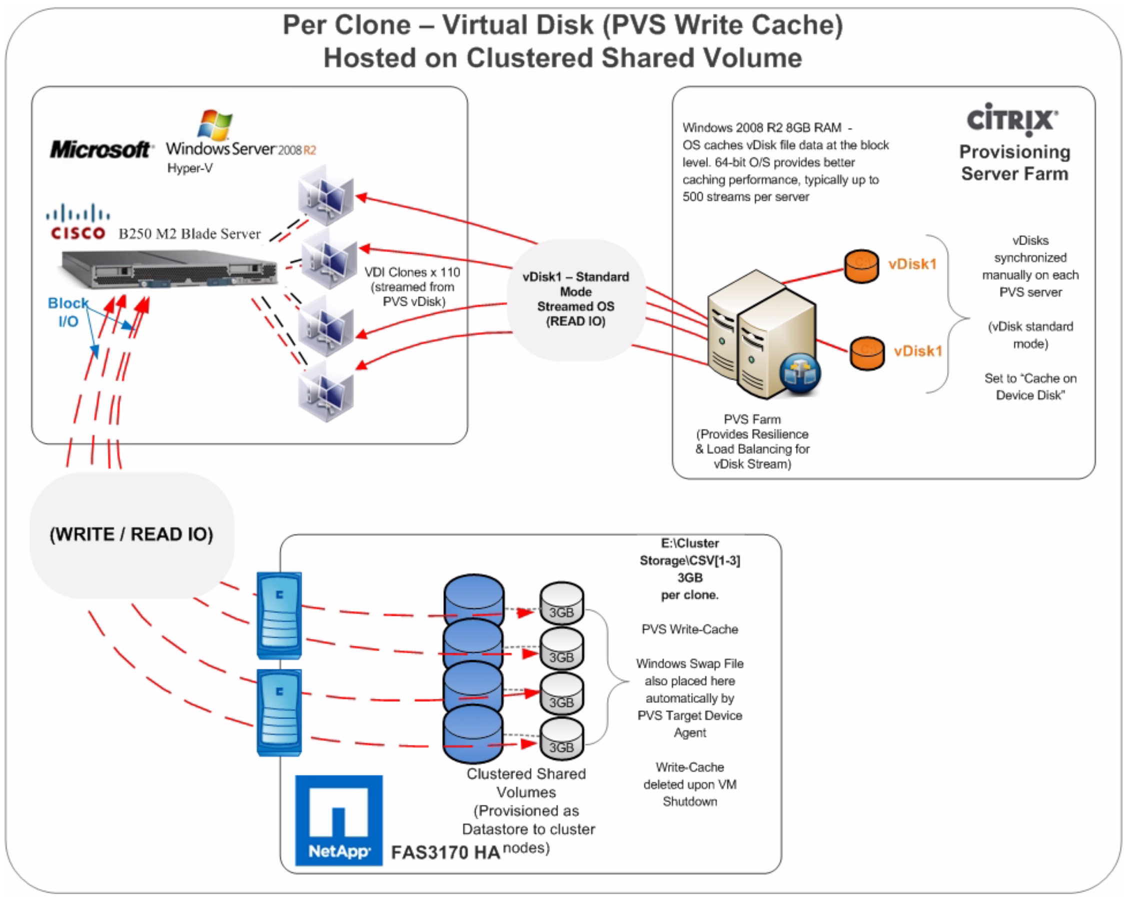


## Infrastructure Components Setup

All the Citrix XenDesktop infrastructure components were hosted on two nodes running Microsoft Hyper-V a Virtual machine. This includes:

- Citrix Provisioning server (PVS)
- Microsoft SCVMM servers

- Citrix Desktop Delivery controller (DDC)
- Licensing server, Active directory, SQL server database server

*Figure 34 Hyper-V Manager on Infrastructure Host 1*

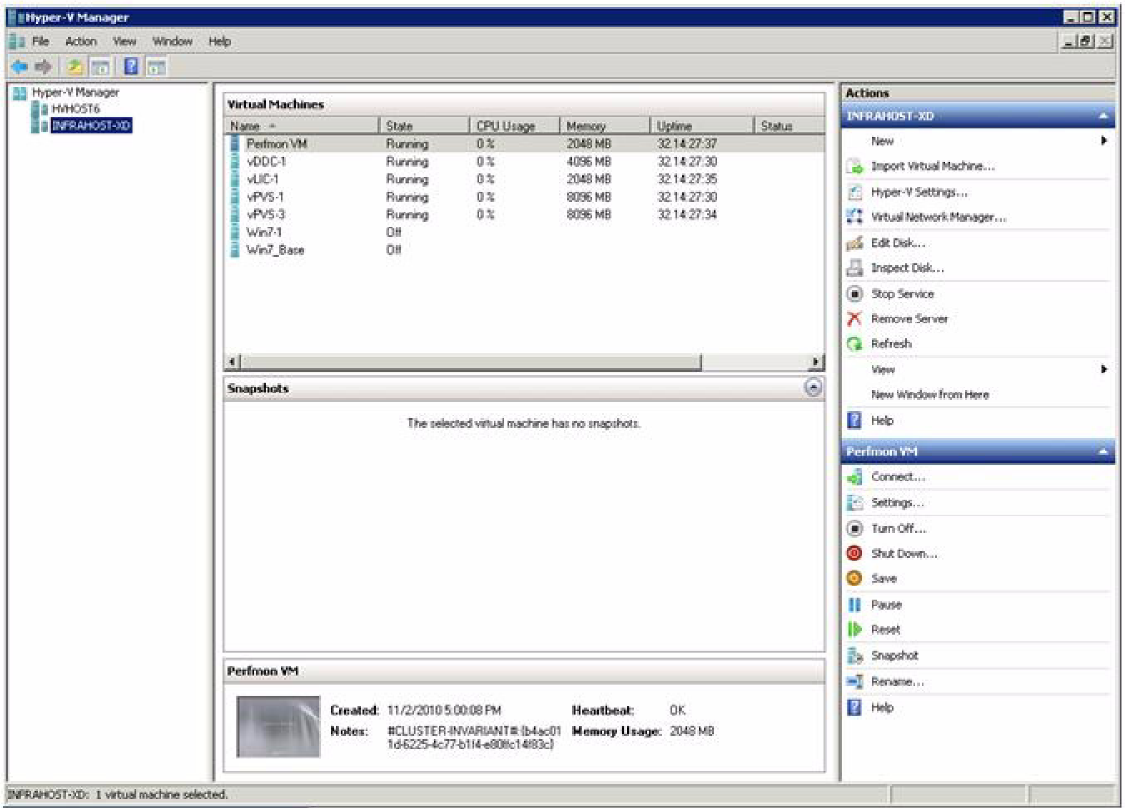

*Figure 35 Hyper-V Manager on Infrastructure Host 2*

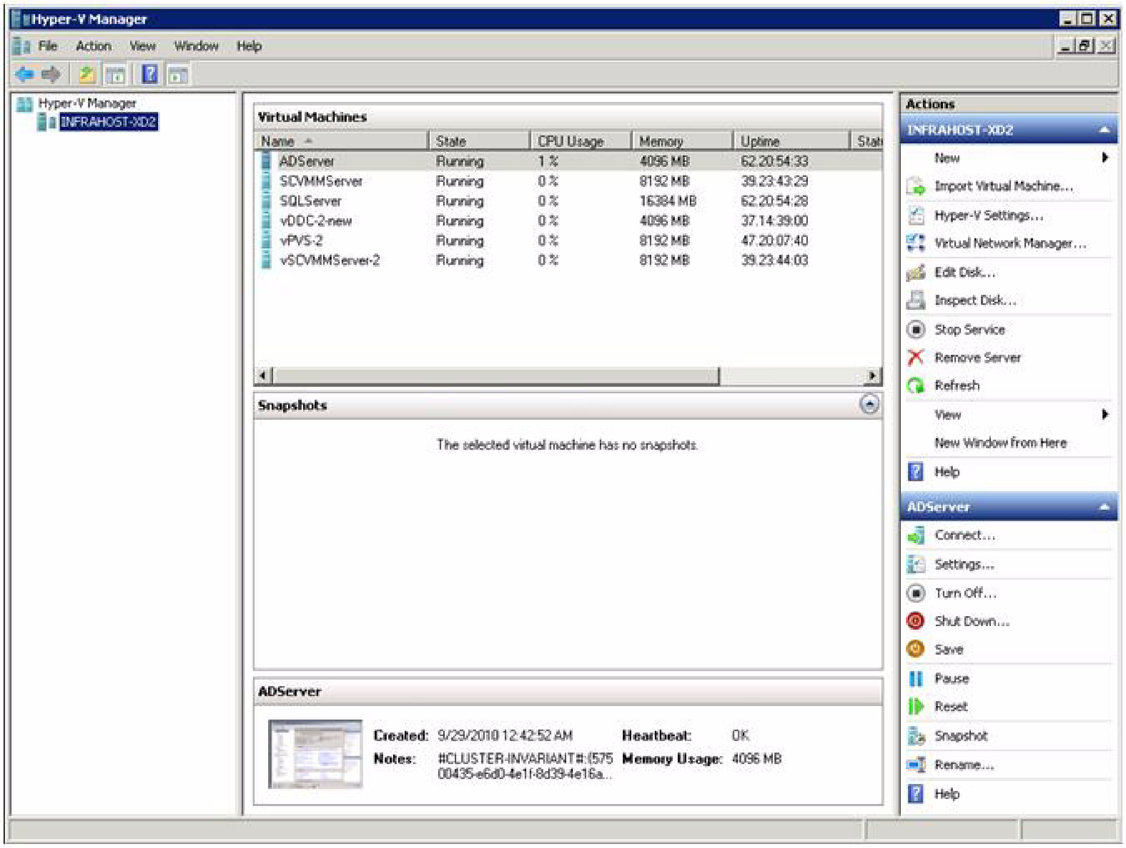

# LAN Configuration

The access layer LAN configuration consists of a pair of Cisco Nexus 5548, a family of low-latency, line-rate, 10 Gigabit Ethernet and FCoE switches for our VDI deployment. Four 10 Gigabit Ethernet uplink ports are configured on each of the Cisco UCS fabric interconnects, and they are connected to the Cisco Nexus 5548 pair in a bow tie manner as shown below. The Fabric interconnect is in End host mode, as we are doing both Fibre Channel as well as Ethernet 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 (Figure 35).

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 5500 and 7000 Series Switches.

*Figure 36 Network Configuration with Upstream Cisco Nexus 5500 Series from the Cisco Unified Computing System*

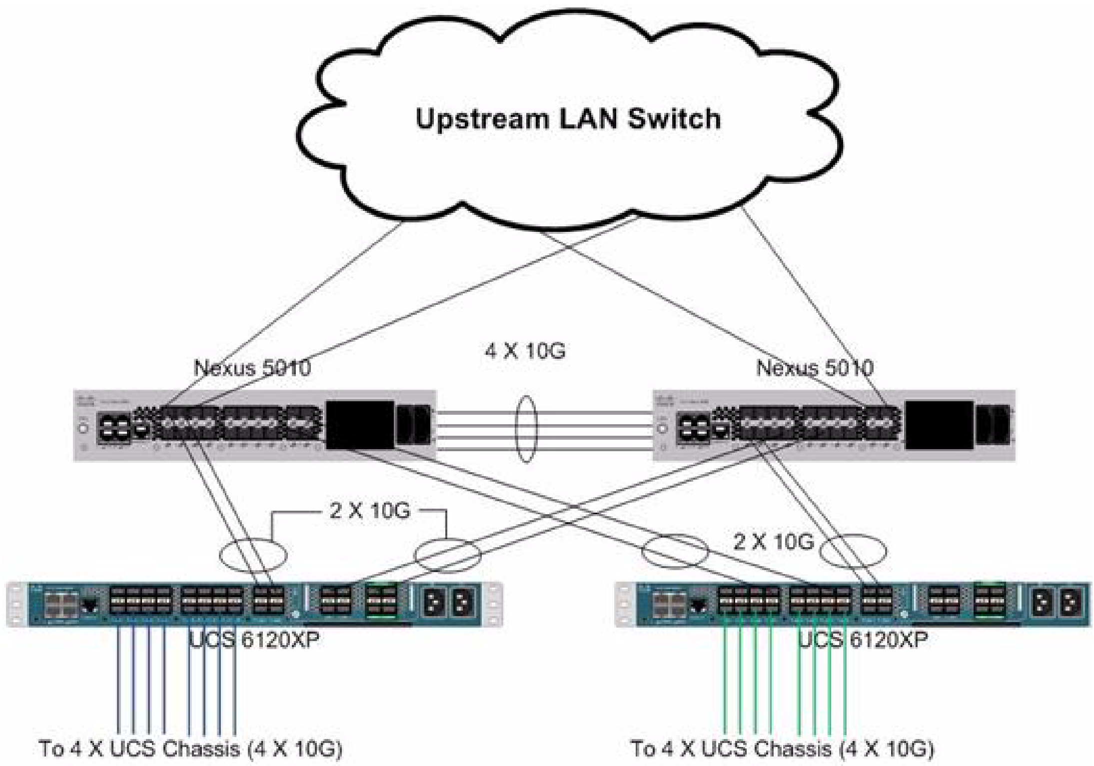


# SAN Configuration

A pair of Cisco MDS 9134 Multilayer Fabric Switches were used in the configuration to connect to the Fibre Channel port of the Cisco UCS fabric interconnect Fibre Channel expansion module ports to a separate NetApp storage (FAS 3140) Fibre Channel ports. Cisco MDS 9000 Family single initiator zone was used to connect to the NetApp Fibre Channel ports. The SAN switch was used for configuring boot from SAN of the Microsoft Hyper-V server blades and also for the cluster shared volume (CSV) to host the Windows 7 desktops. Our recommendation is do single initiator and single target zoning, for small setup like our we did multiple target zoning as shown below. Another recommendation is to configure different target ports into different port groups of the MDS (one target on each port group works very well). Also make sure you do the proper zoning for the different Clustered Share Volumes (CSVs).

The infrastructure volumes (CSVs) were block based and the zoning was done to make those NetApp LUNs visible to the infrastructure and test servers. An example SAN zone configuration is shown below on the Fabric A side:

MDS-A# sh zoneset active vsan 1

zoneset name FAB-A-XD-MSHV-BFS vsan 1

  zone name XD-HyperV-Serv1-fc0 vsan 1

  * fcid 0x470133 [pwwn 20:00:00:25:b5:0a:ad:3e]

  * fcid 0x470200 [pwwn 50:0a:09:83:89:1a:b9:d9]

  * fcid 0x470300 [pwwn 50:0a:09:81:89:1a:b9:d9]

:

  zone name XD-HyperV-Serv2-fc0 vsan 1

* fcid 0x47002e [pwwn 20:00:00:25:b5:0a:ad:3c]

* fcid 0x470200 [pwwn 50:0a:09:83:89:1a:b9:d9]

* fcid 0x470300 [pwwn 50:0a:09:81:89:1a:b9:d9]

Where 20:00:00:25:b5:0a:ad:3e/20:00:00:25:b5:0a:ad:2e are server's pwwn of the CNA that are part of the Fabric A side. Similar zoning is done on the corresponding Cisco MDS 9000 Family switch pair to take care of the Fabric B side as shown below.

MDS-B# sh zoneset active vsan 1

zoneset name FAB-B-XD-MSHV-BFS vsan 1

zone name XD-HyperV-Serv1-fc1 vsan 1

* fcid 0x47002e [pwwn 20:00:00:25:b5:0a:ad:2e]

* fcid 0x470500 [pwwn 50:0a:09:81:99:1a:b9:d9]

* fcid 0x470400 [pwwn 50:0a:09:83:99:1a:b9:d9]

:

zone name XD-HyperV-Serv2-fc1 vsan 1

* fcid 0x470735 [pwwn 20:00:00:25:b5:0a:ad:2c]

* fcid 0x470500 [pwwn 50:0a:09:83:99:1a:b9:d9]

* fcid 0x470400 [pwwn 50:0a:09:81:99:1a:b9:d9]

The NetApp Fibre Channel target ports, 50:0a:09:83:89:1a:b9:d9/50:0a:09:83:99:1a:b9:d9 belong to one controller and 50:0a:09:81:99:1a:b9:d9/50:0a:09:81:89:1a:b9:d9 was part of the second controller. They were spread across the two controllers for redundancy as shown in Figure 36.

*Figure 37 SAN Configuration with Dual MDS and NetApp FAS 3140 Storage*

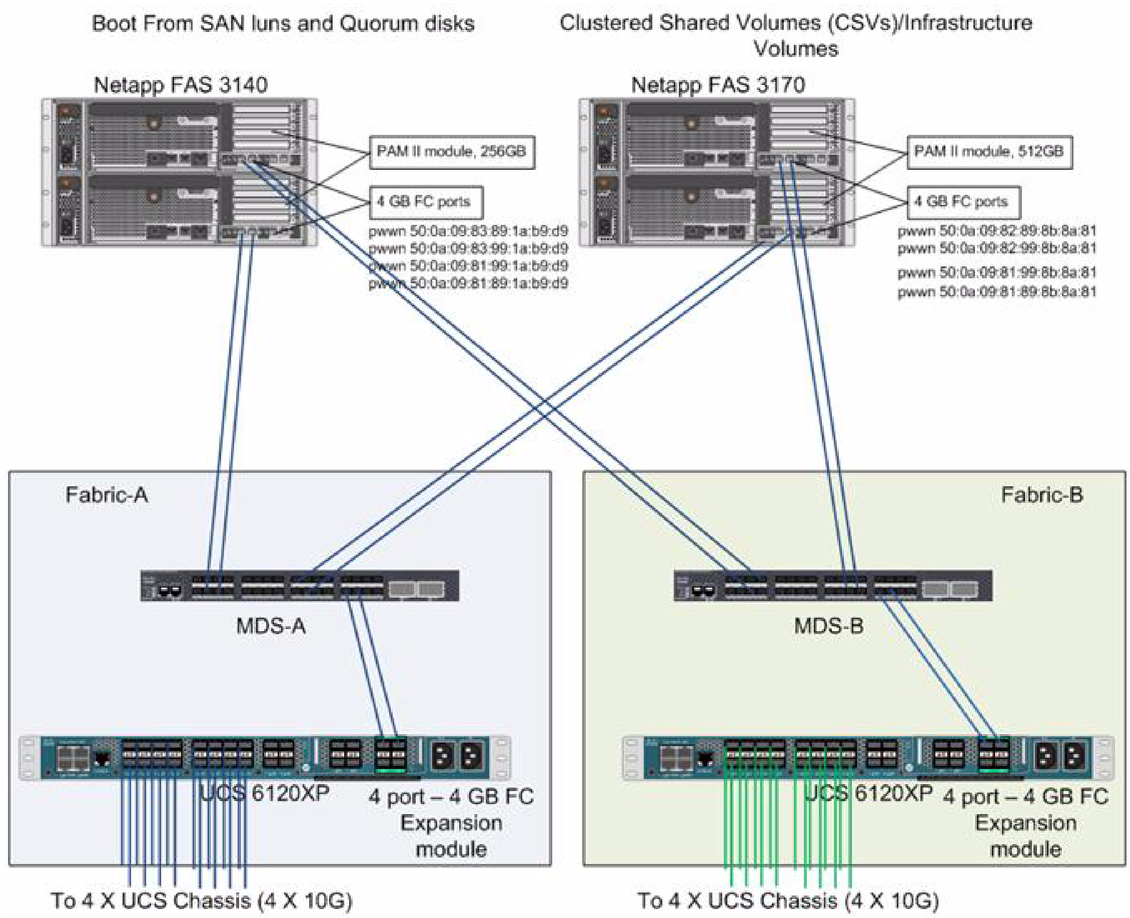


## Boot from SAN

Booting from SAN is another critical 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 supposed 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 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 main 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.

## Configuring Boot from SAN on the Cisco Unified Computing System

With boot from SAN, the image resides on the SAN 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 Fibre Channel capable 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.

**Note** The 2 SAN fabrics are disjoint from data perspective and with the dual port HBA's and storage controller redundancy is provided.

The following are the distinct portions of the BFS procedure:

- Storage array configuration

  First, the storage array admin has to provision LUNs of the required size for installing the OS and to enable the boot from SAN. The boot from SAN LUN is usually LUN 0. The SAN admin also needs to know the port world-wide name of the adapter so that the necessary lun masking is put in place. The lun masking is also a critical step in the SAN LUN configuration.

  For example, in case of NetApp 3140/3170 storage array, the storage admin has to create a BootVolume and then include the blade WWPNs into a initiator group and add them to the port WWPNs where the storage is configured.
- SAN zone configuration
- Cisco UCS configuration of service profile

To configure boot from SAN on the Cisco Unified Computing System, do the following:

1. Create a separate boot from SAN Aggregate.
2. Create a Volume on top of that, call it BFSVOL.

3. Create a LUN by selecting LUN>Add (Add LUN wizard) and set the required fields as shown below.

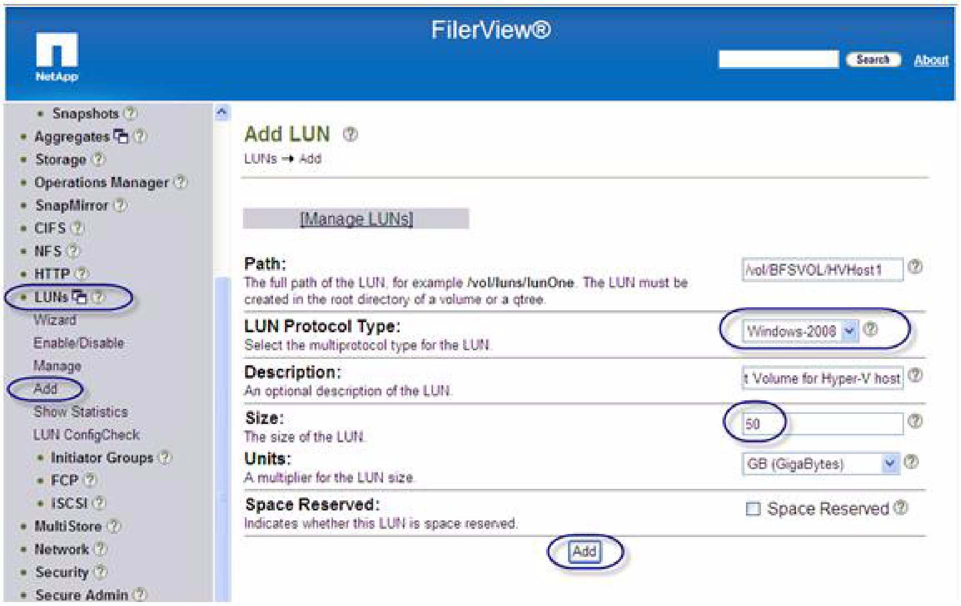


4. Add the LUN to the initiator group.

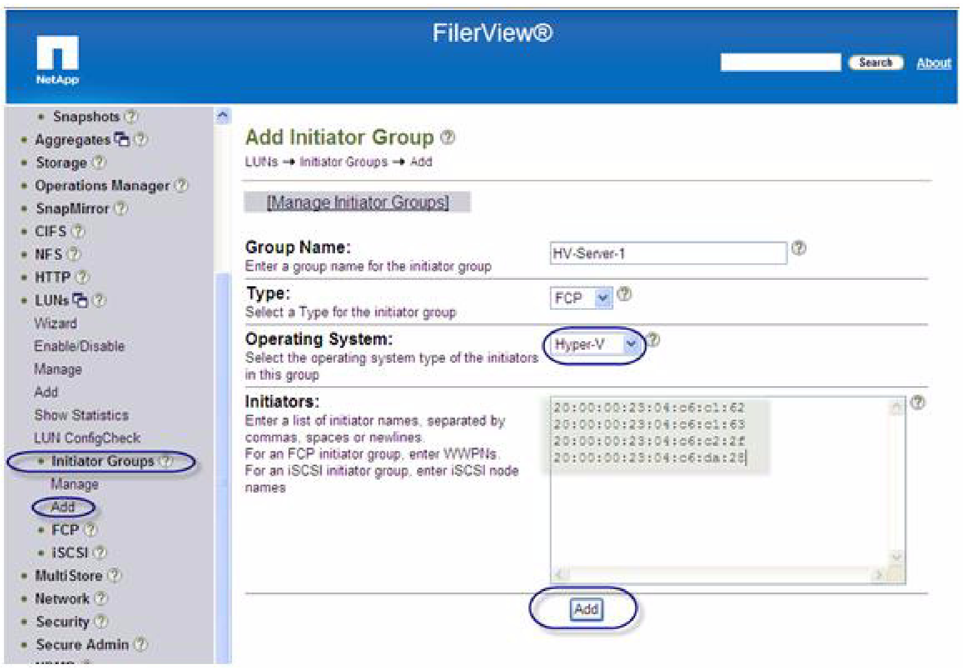


5. Make sure the add initiator group succeeds.

6. Now when you click Manage LUNs, you will notice the Maps group: LUN ID as no maps. For example:


7. Map this LUN to the LUN ID 0 and associate this with the initiator group as shown in the following screenshots:

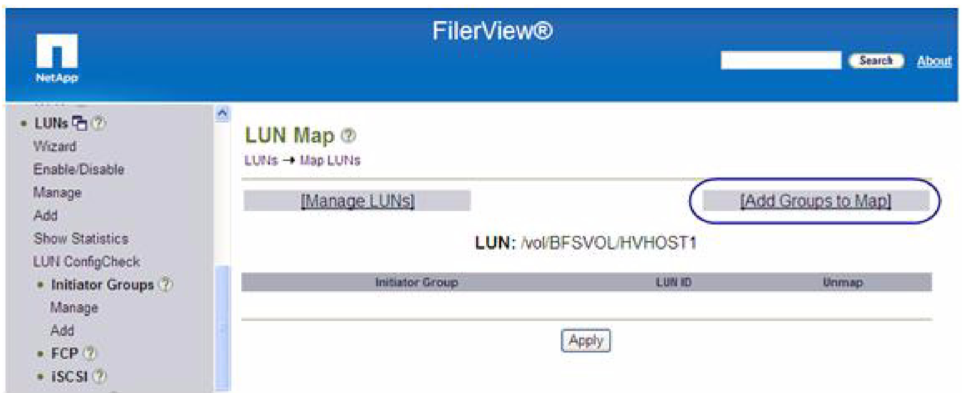


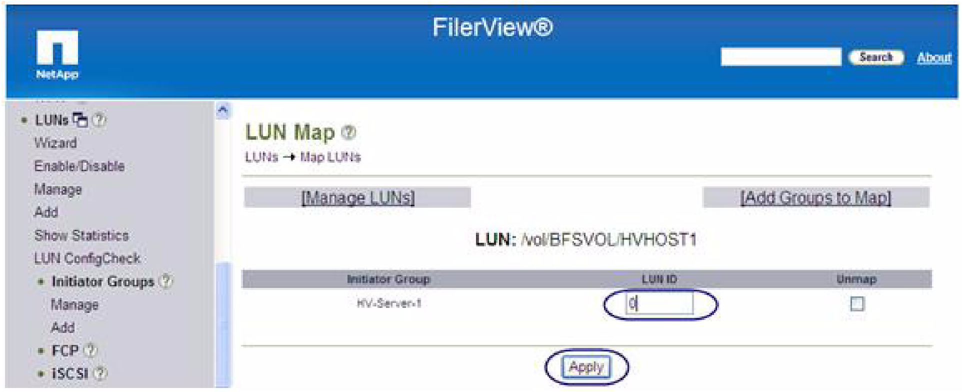


8. Make sure the mapping was successful.

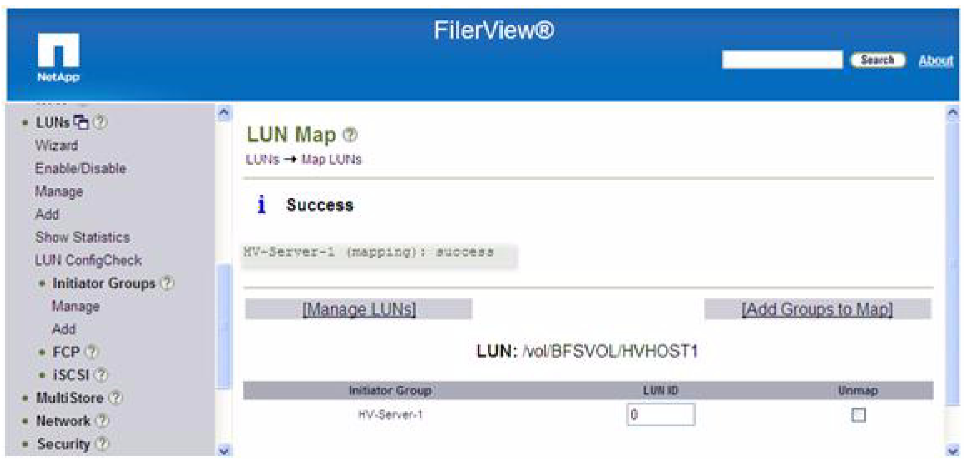


9. Now the LUN will have a map defined with LUN id as 0.


This completes the configuration for the boot from SAN LUN. Windows 2008 R2 install can be done on this LUN using Virtual media.

## SAN Configuration

The NPIV feature has to be turned on in the SAN Switch. Also make sure you have 4 GB SPF+ modules connected to the Cisco UCS 6120 and 6140 XP Fabric Interconnects. 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 can view it on a Cisco MDS Device Manager for a given port (for example, Fc2/16). VSAN configuration can be done either in the SAN switch CLI or the GUI, like the Cisco MDS Device Manager. 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 target pwwn.

```
# show feature | grep npiv

npiv           1      enabled

# show interface br

-------------------------------------------------------------------------------

Interface  Vsan   Admin  Admin   Status        SFP    Oper  Oper   Port

                  Mode   Trunk                        Mode  Speed  Channel

                         Mode                               (Gbps)

-------------------------------------------------------------------------------

fc1/1      1      auto   on      up            swl    F      4     --

fc1/2      1      auto   on      up            swl    F      4     --

fc1/3      1      auto   on      up            swl    F      4     --

fc1/4      1      auto   on      up            swl    F      4     --


# sh int fc1/1 brief

-------------------------------------------------------------------------------

Interface  Vsan   Admin  Admin   Status        SFP    Oper  Oper   Port

                  Mode   Trunk                        Mode  Speed  Channel

                         Mode                               (Gbps)

-------------------------------------------------------------------------------

fc1/1      1      auto   on      up            swl    F      4     --
```

## Cisco UCS Manager Configuration

To enable boot from SAN from a Cisco UCS Manager perspective, do the following:

1. Create a boot policy from the Servers tab; from the right pane, select boot policies then select Add.

2. Enter the name, select reboot on change; do not select enforce vHBA name.

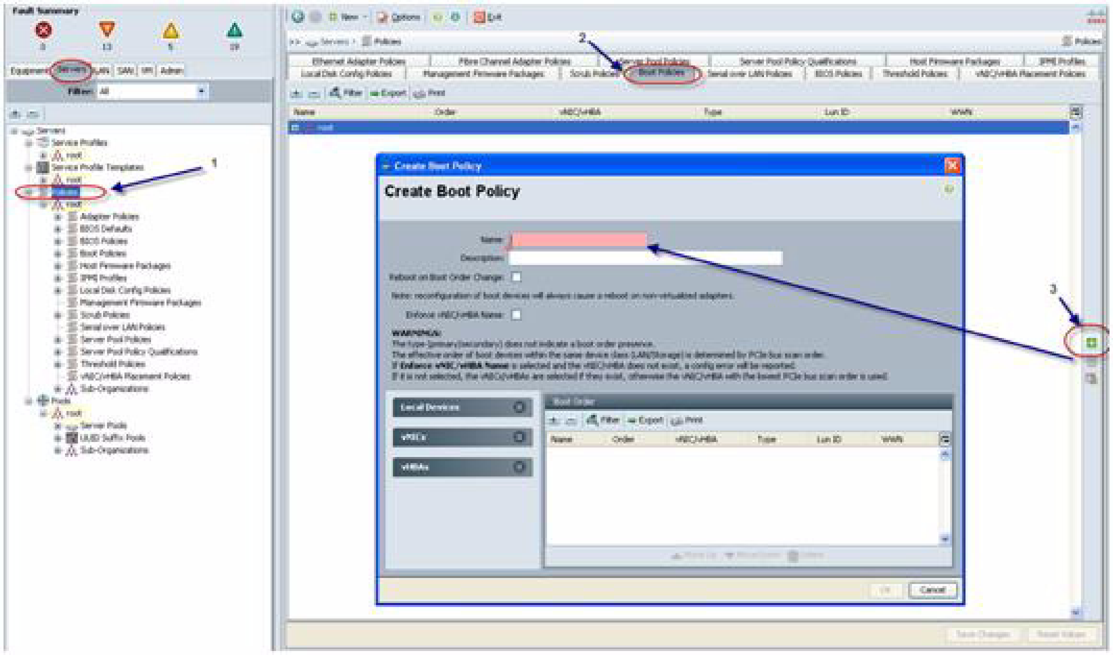


3. Add SAN Boot for primary. The vHBA is optional; it can be left blank and you do not have to enforce the vHBA name.

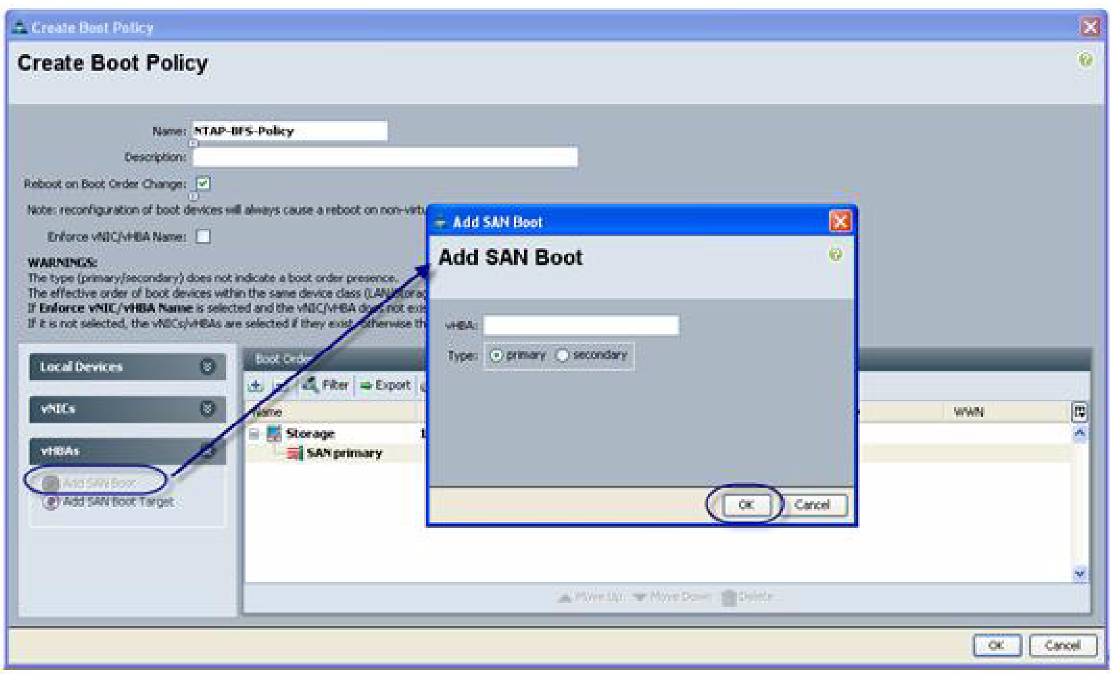

4. Add SAN boot for SAN Secondary.

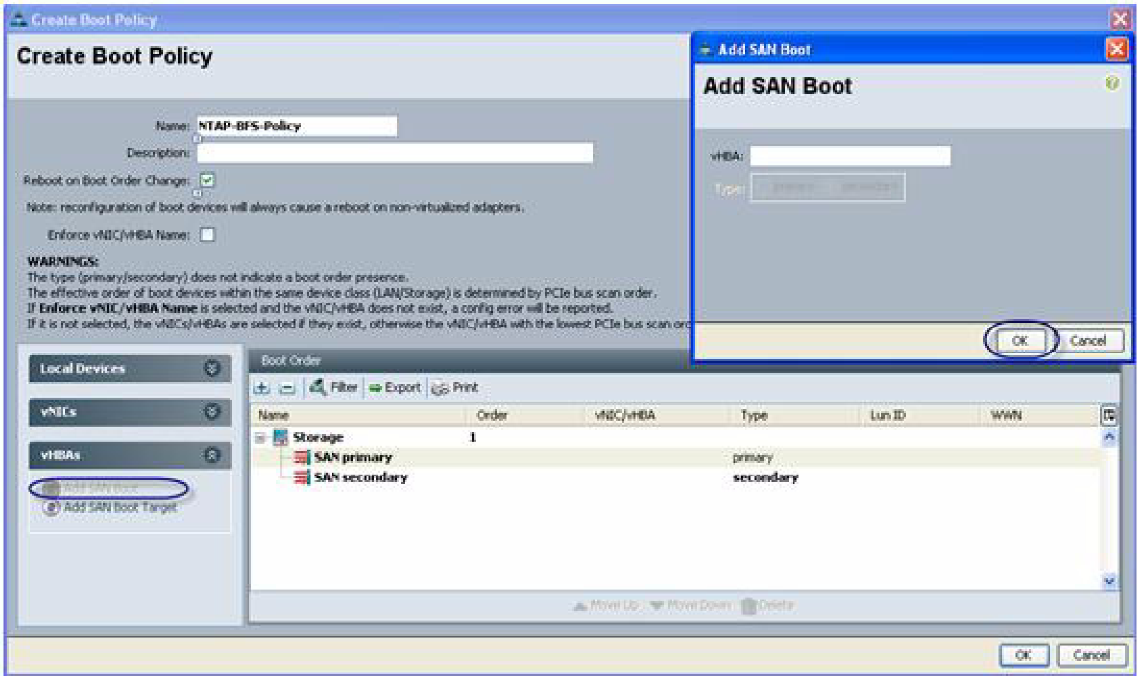


5. Now add Boot target WWPN to the SAN Primary, make sure this is the NetApp FAS 3140 pwwn. Avoid any typos and copy paste from MDS show flogi da.

   MDS-A# sh fcns da vsan 1 | incl Net

   0x470300 N 50:0a:09:81:89:1a:b9:d9 (NetApp) scsi-fcp

   0x470200 N 50:0a:09:83:89:1a:b9:d9 (NetApp) scsi-fcp

   MDS-B # sh fcns da vsan 1 | incl Net

   0x470400 N 50:0a:09:83:99:1a:b9:d9 (NetApp) scsi-fcp

   0x470500 N 50:0a:09:81:99:1a:b9:d9 (NetApp) scsi-fcp

6. Repeat step 5 for SAN primary's SAN Target Secondary.
7. Repeat step 5 for SAN Secondary's - SAN Target Primary.
8. Repeat step 5 for SAN Secondary's - SAN Target Secondary.
9. The Boot from SAN policy will look like the following:

10. Make the association of the service profile template to the Boot from SAN policy during the service profile template configuration. This will apply all the service profiles derived out of the template to have boot from SAN configured to boot from the NetApp Storage.

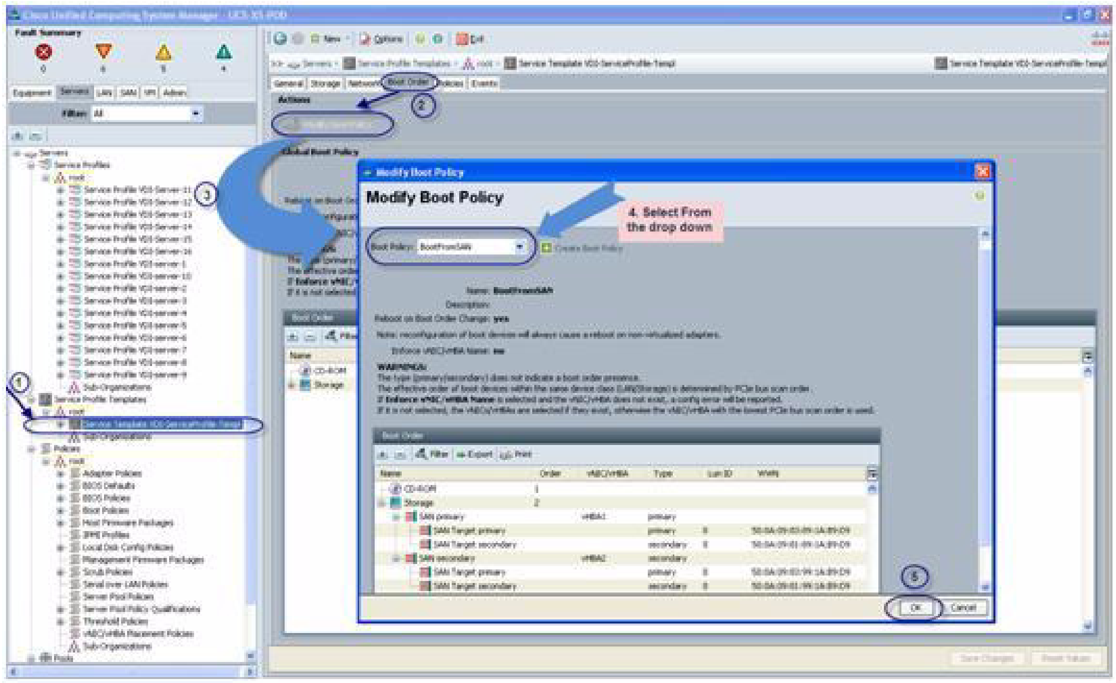


This completes the BFS configuration on Cisco UCS Manager. When the service profile is created from the template, each server will be ready to boot from SAN provided the appropriate OS installation steps has taken place.

# NetApp Storage Configuration

Two NetApp Storage (FAS3140 and FAS3170) were used for single host and eight (8) host testing scenario (Figure 37):

- FAS3140 with PAM-II modules was used to host:
  - Boot volume for 8 Windows 2008 R2 Server Core Hosts
- FAS3170 with PAM-II modules was used to host:
  - Three 2 TB block based Clustered shared volumes for the Virtual desktops on controller 1
  - 2 block based storage volumes for the infrastructure on controller 2

**Note** It is not necessary to separate the boot volumes, infrastructure volumes, and Virtual desktop volumes; all can be hosted on a single FAS3170 or FAS3140 NetApp array (or any similar FAS array).

*Figure 38 NetApp Storage Configuration for Two Chassis*

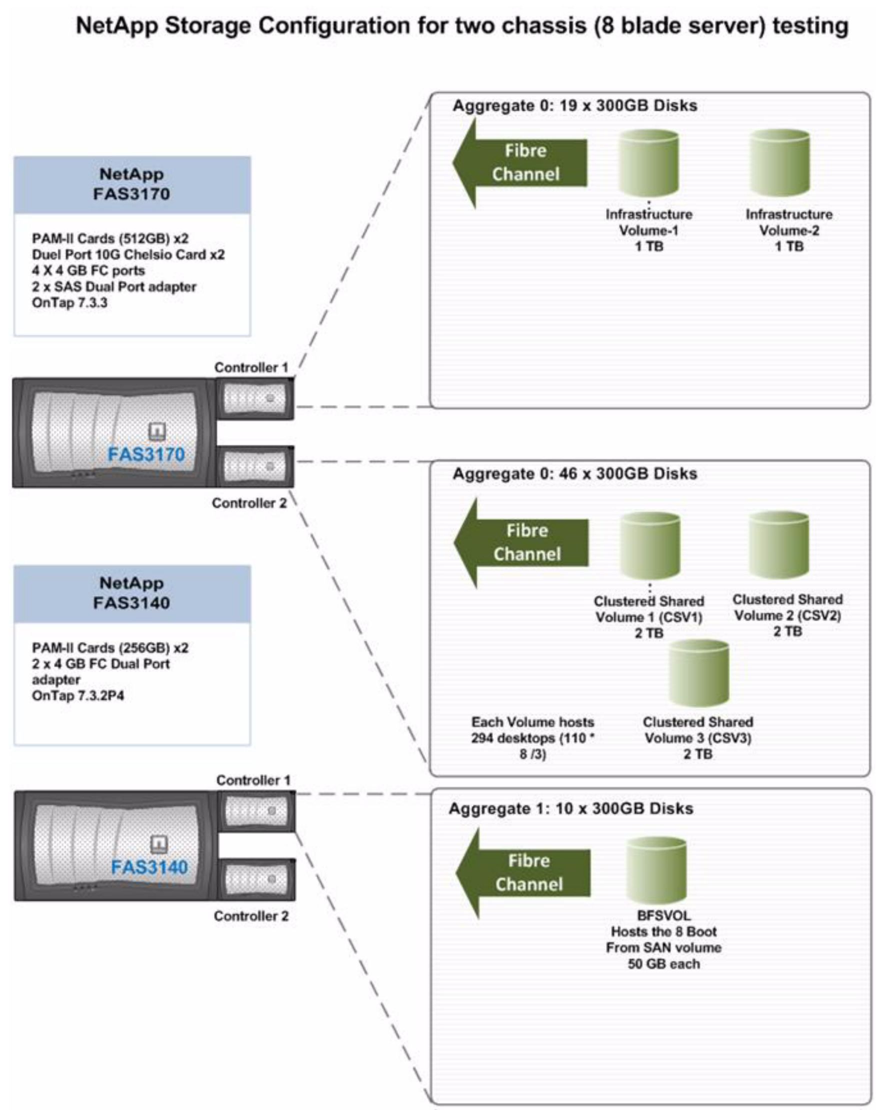


Two NetApp Storage (FAS3140 and FAS3170) were used for large scale 16 host testing scenarios (Figure 36):

- FAS3140 with PAM-II modules was used to host:
  - Boot volume for 16 Windows 2008 R2 Server Core Hosts
- FAS3170 with PAM-II modules was used to host:
  - Six 2 TB Clustered shared volumes for the Virtual desktops on controller 1
  - 3 volumes for the infrastructure
  - NetApp Storage Configuration for Four Chassis (or 16 servers)

**Note** It is not necessary to separate the boot volumes from the infrastructure volumes and Virtual desktop volumes; all volumes can be hosted on a single FAS3170 or FAS3140 NetApp array (or any similar FAS array).

*Figure 39 NetApp Storage Configuration for Four Chassis*

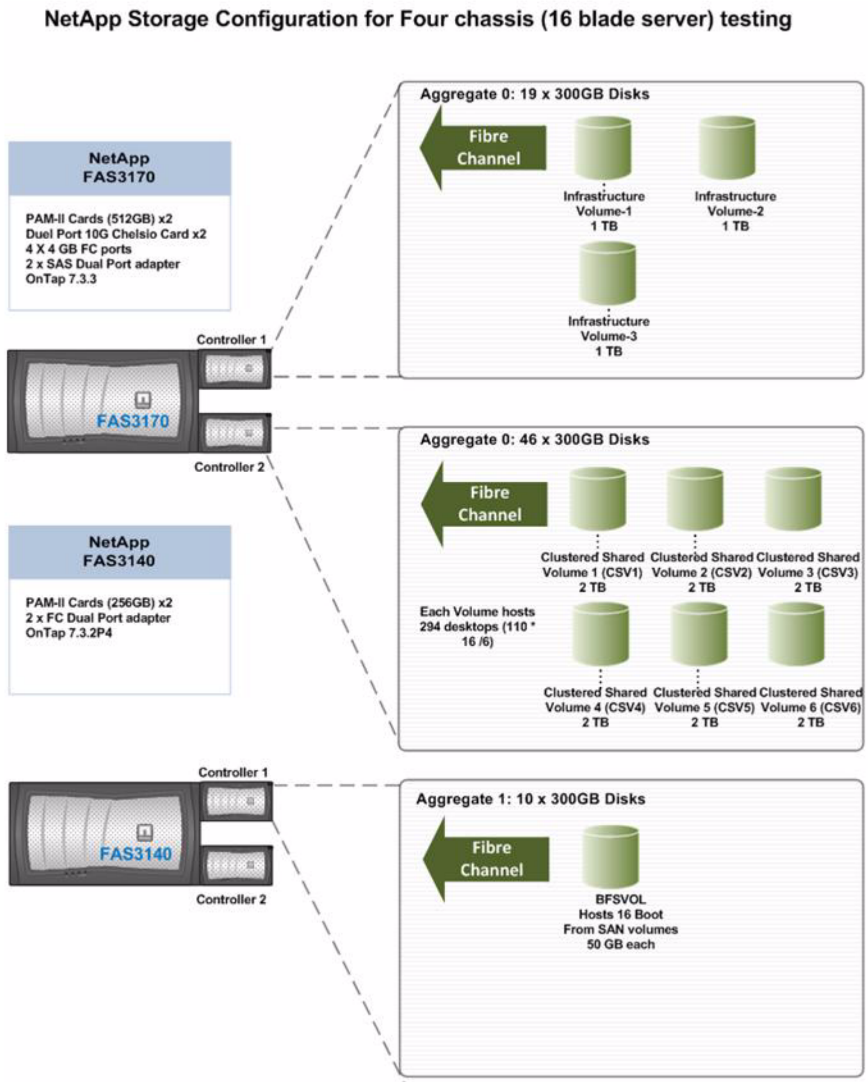


# Example of a NetApp Volume Configuration

This section describes the configuration of an 800 GB FlexVolume using the NetApp FileViewer management utility. The following steps can be done through a command line on the NetApp storage controller CLI.

To configure a NetApp Volume, do the following:

1. Log in to the NetApp storage using a web browser and click FilerView. This starts the NetApp storage configuration application.

2. From the FilerView, select the Aggregates section and click Add to create an aggregate. In this example, create an aggregate out of 46 disks and call it aggr1.


3. From the Volumes section, select Add to add a volume.

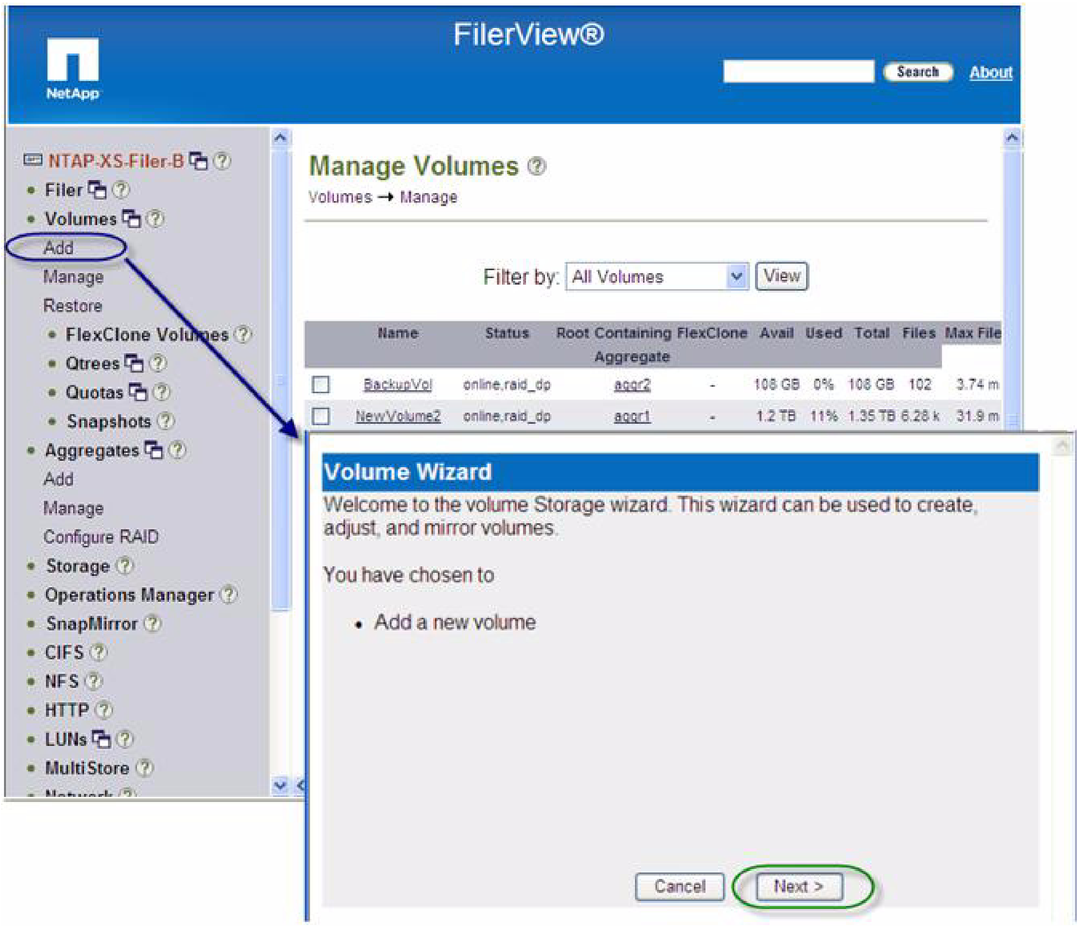


4. Select Flexible for the volume type.

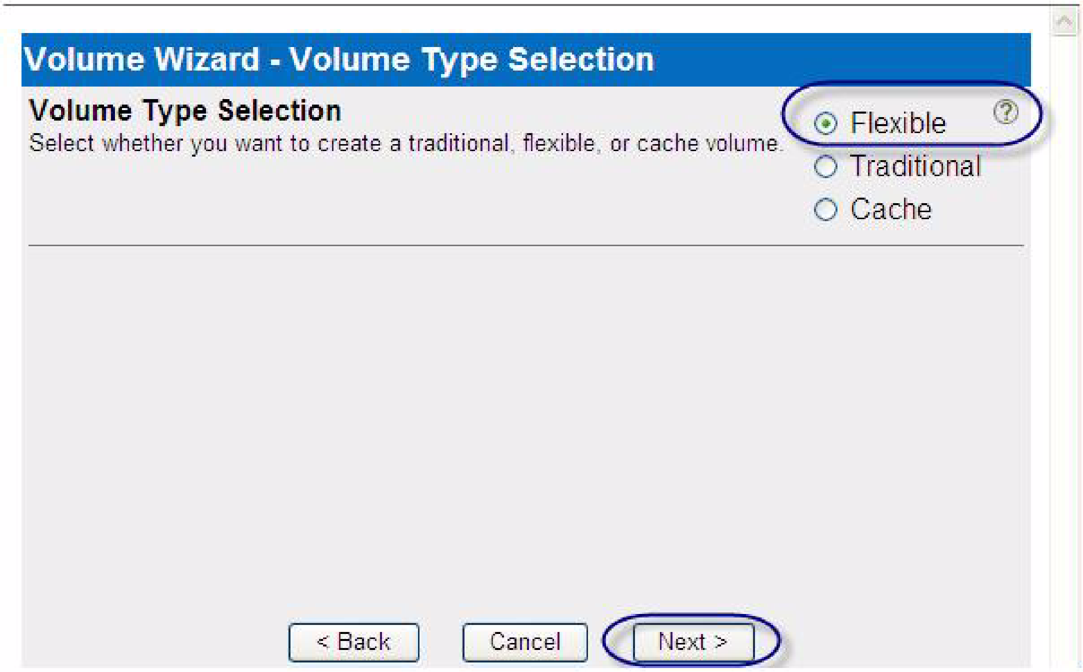


5. Input volume name and language (default is POSIX).

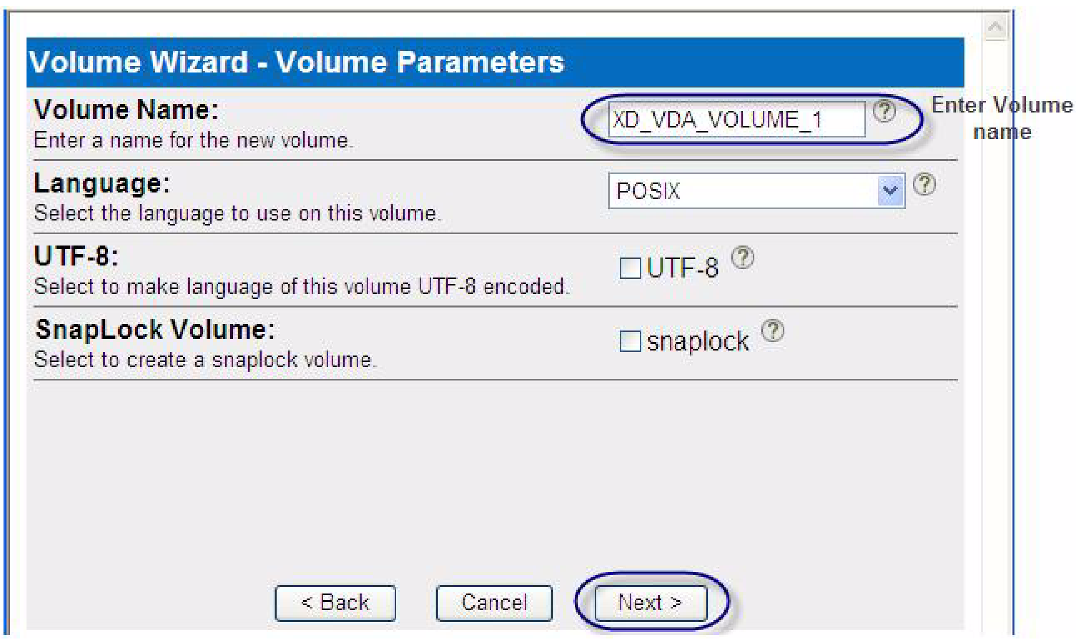


6. Select the aggregate to contain this volume.

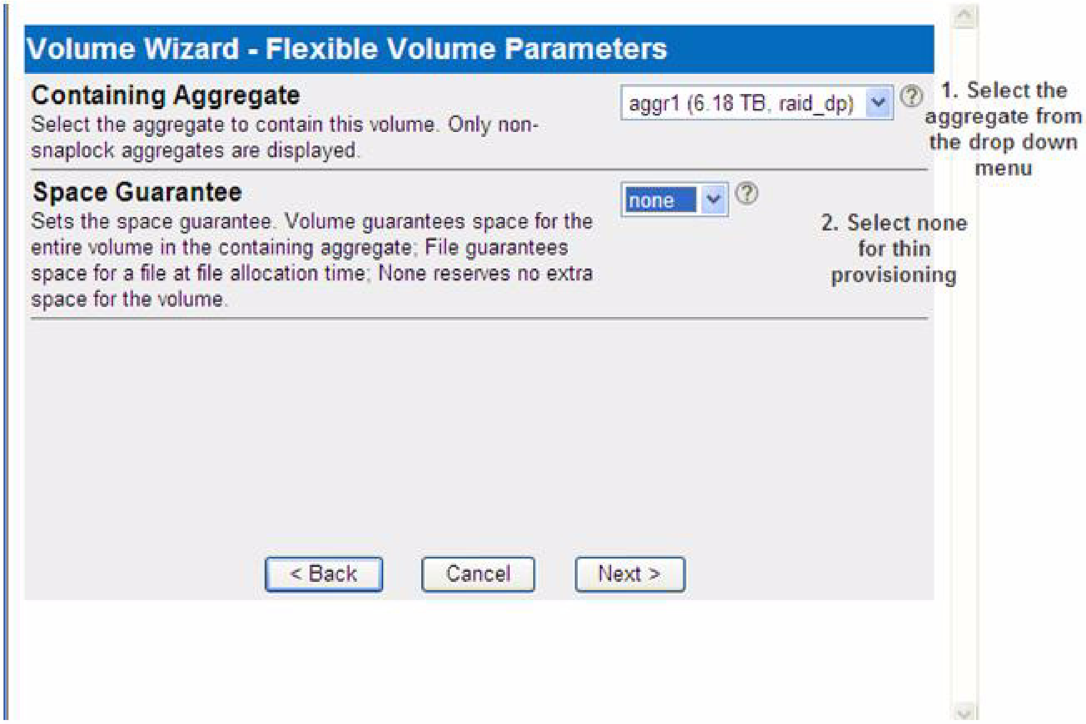


7. Input the volume size and snapshot reserve.

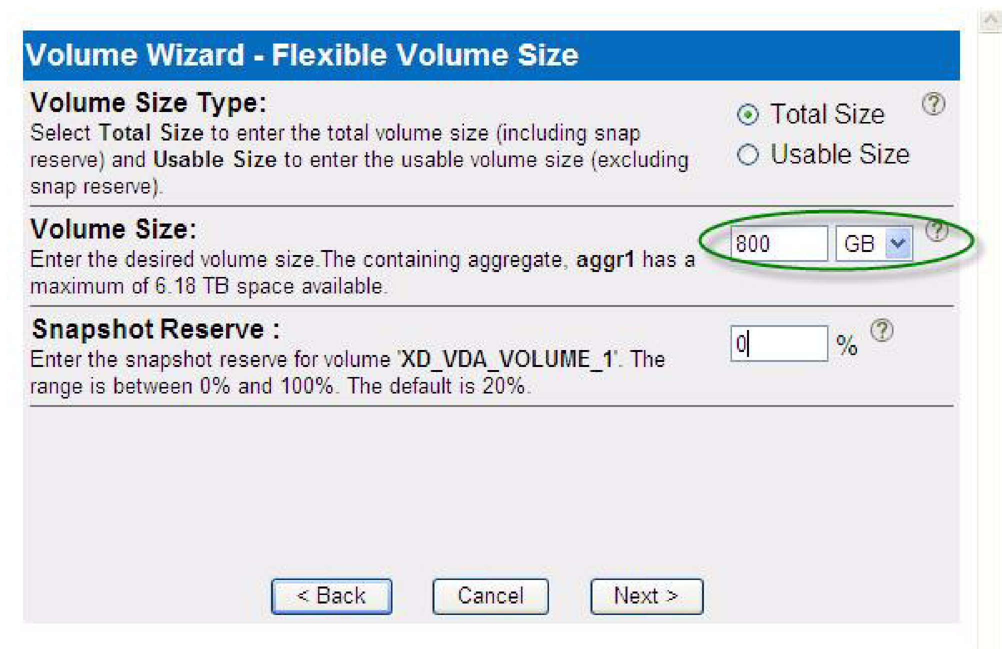


8. Click Commit.

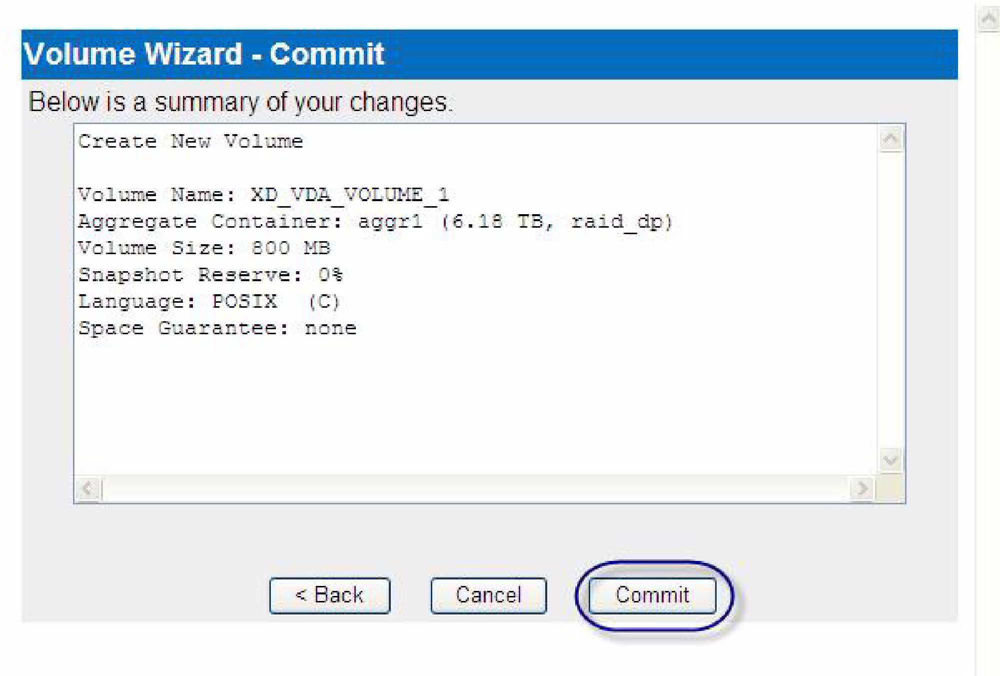


9. After the volume is added, create the Initiator group for each host and map the devices to the newly created volume. When the MDS zone configuration for the initiator and the targets are in place, the LUNS are visible to all he Hyper-V cluster nodes.

## NetApp Storage Configuration for Windows Server 2008 R2 Hyper-V Clusters

One of the major feature in Windows server 2008 R2 is the introduction of Clustered shared volume (CSV) that is only available for Hyper-V role. A CSV is a physical disk connected to the parent partition and shared between multiple Hyper-V server nodes configured as part of a single Windows failover cluster. This is part of "failover clustering" feature. A CSV function as a distributed-access file system optimized for Hyper-V and is based on NTFS. CSV provides the infrastructure to support and enhance live migration of Hyper-V virtual machines. To enable CSV on a cluster, one has to use the Failover Cluster manager MMC snap-in.

In order for the all the hosts in the cluster to see all the CSVs, proper zoning has to be done on the MDS SAN Switch and on the NetApp storage proper masking has to be done. For example, Figure 37 shows how eight Hyper-V hosts which are part of a cluster are configured to see three CSVs. Notice the three CSVs, namely, HyperV_CSV_Vol1, HyperV_CSV_Vol2, and HyperV_CSV_Vol3 are mapped to eight initiator groups corresponding to eight hosts in a Hyper-V cluster and mapped to LUN ID 6,7, and 8 respectively (the LUN ID has to be greater than the quorum disk LUN ID).

*Figure 40 Cluster Shared Volume Configuration on NetApp Storage*

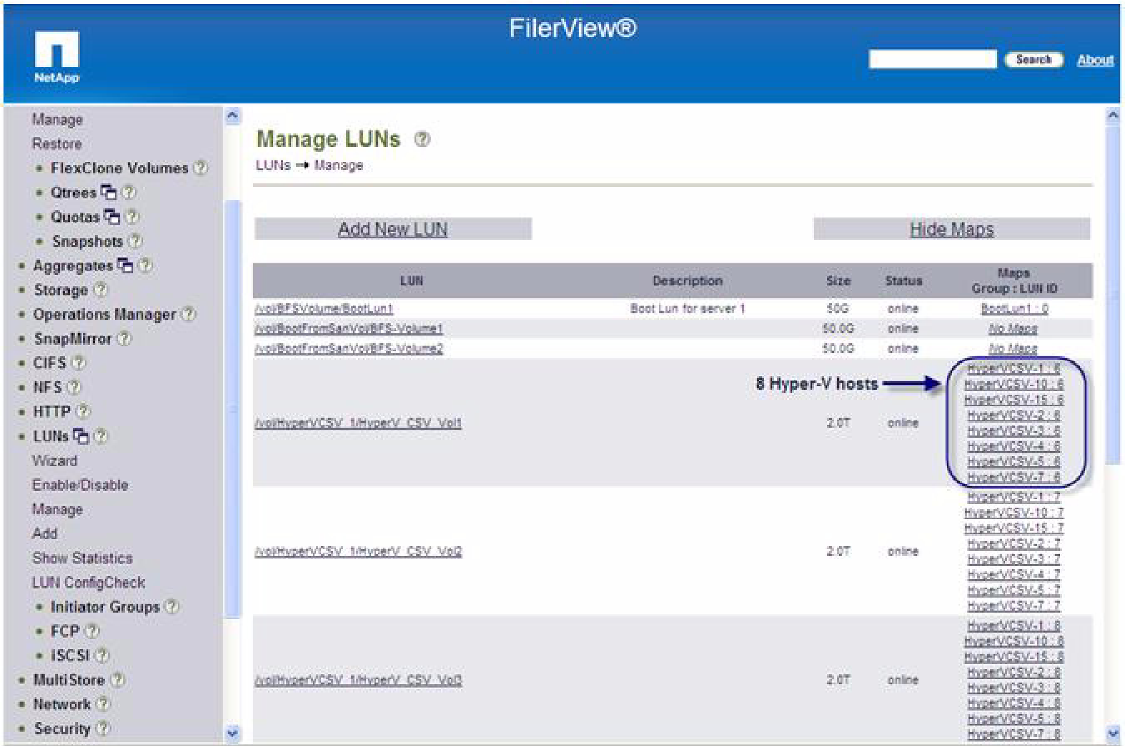


Figure 40 displays a cluster shared volume where three CSVs are shared across eight hosts.

*Figure 41 Cluster Shared Volume Configuration of an Eight Node Hyper-V Cluster*

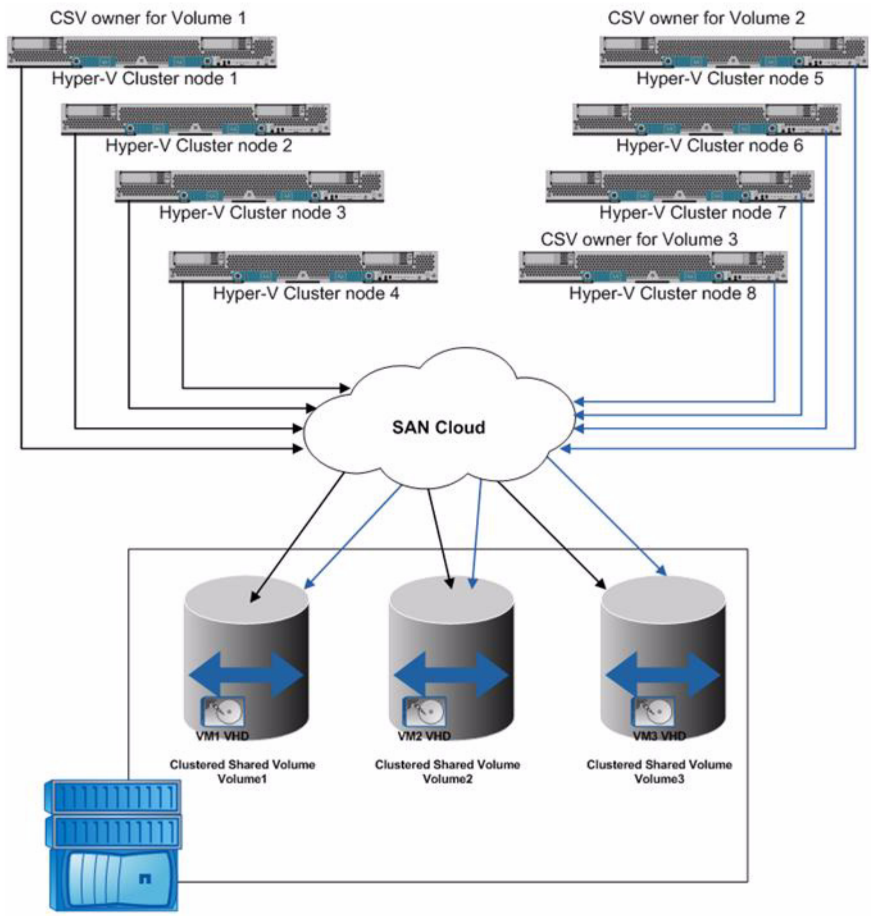


CSV enables single namespace between all Hyper-V hosts or server nodes in the failover cluster. This enables single pathnames across all nodes while accessing any files on any nodes of the failover cluster. CSVs are mounted on %systemdrive%\ClusterStorage root folder as shown in the figure below (All nodes has similar output). The CSVs Volume1, Volume2, and Volume3 are in the C:\ system drive and n the ClusterStorage folder. The paths for each of the volumes are - C:\ClusterStorage\Volume1, C:\ClusterStorage\Volume2, and C:\ClusterStorage\Volume3. This eliminates the requirement of a drive letter assignment and makes management of CSV easy. Further, the VM1 VHD can be accessed by C:\ClusterStorage\Volume1\VM1.vhd.

*Figure 42 Output of Directory Contents of C:\ClusterStorage\ and diskpart*

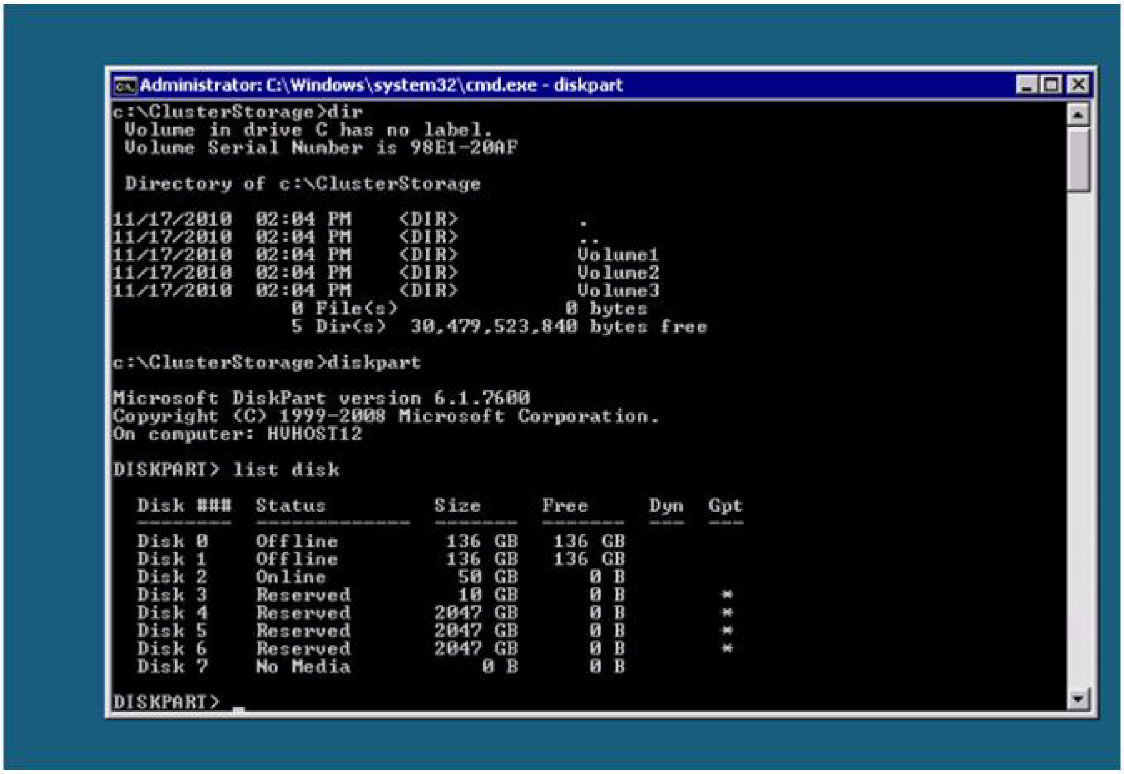


## Quorum Disk Configuration for Windows 2008 R2 Hyper-V Clustering

For failover clustering you need to configure a small LUN (Microsoft recommend at least 500 MB in size) called the quorum disk or a witness disk. This contains all the relevant cluster configuration copy and is used to determine the health and well being of the cluster. A majority of the cluster nodes must remain online and in communication or else the cluster loses quorum and must stop running. Each node must have access to this disk and must be properly zoned and masked from the SAN switch and NetApp storage perspective. Figure 42 shows two quorum disks configured for two clustered groups with eight hosts each.

*Figure 43 Quorum Disks Configured for Two Clustered Groups*

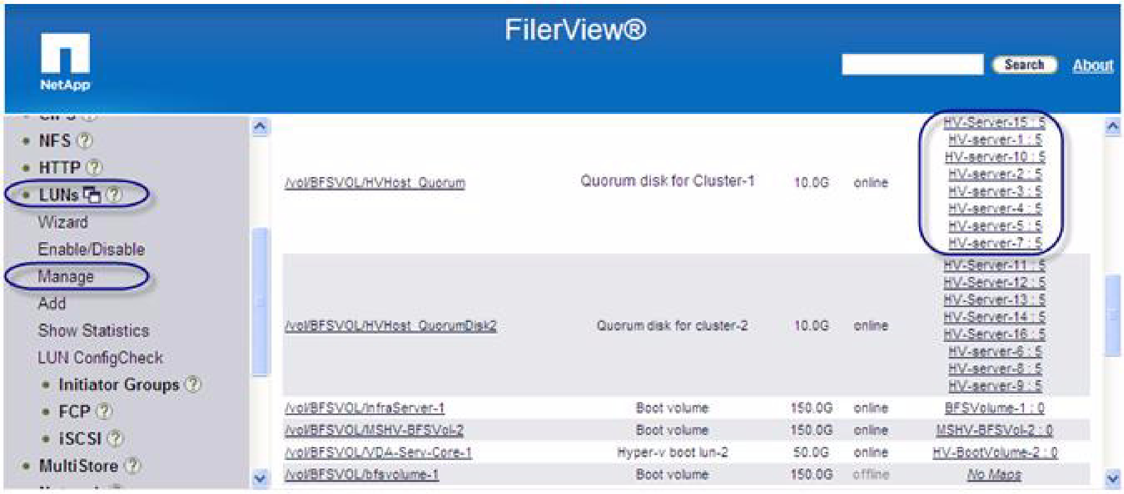


# Microsoft Hyper-V Configuration

This section details the Microsoft Windows 2008 Hyper-V configuration and tuning.

*Table 11 Software Components*

| **Microsoft Windows 2008 R2** | | | |
|---|---|---|---|
| **Hardware:** | Cisco UCS B-Series blade servers | **Model:** | B250 –M2 |
| **OS:** | Microsoft Windows 2008 R2 Server Core | **Service Pack:** | - |
| **CPU:** | 2 x 6 Core Westmere or 5680, 3.33 GHz (24 Logical Cores Total) | **RAM:** | 192 GB (8G X 24 or 4G X 48) @1333 MHz |
| **Boot Disk:** | Boot From SAN | **Local Drive:** | Optional |
| **CNA:** | 2 X Cisco VIC (M81KR) | **Network:** | 2 x 10GbE per VIC |
| **CNA Driver Version:** | Downloaded latest Win2k8 R2 drivers from cco.cisco.com | **FC:** | 2 X 4 G per VIC |
| ○ **Windows 2008 R2 Server Core – Hyper-V**<br>○ **Windows update was run to make sure we had the latest bits** | | | |

## Cisco UCS Configuration for Microsoft Windows 2008 R2 Installation

Boot from SAN was used to install Microsoft Windows 2008 R2 rather than local disk, this is in keeping with the dynamic provisioning of service profiles the Cisco Unified Computing System offers, making them portable and thus allowing physical resources to be reused quickly and easily if required.

The Cisco UCS server policy used also has a consistent Boot Order policy attached and configured for all servers in the Cisco UCS pool.

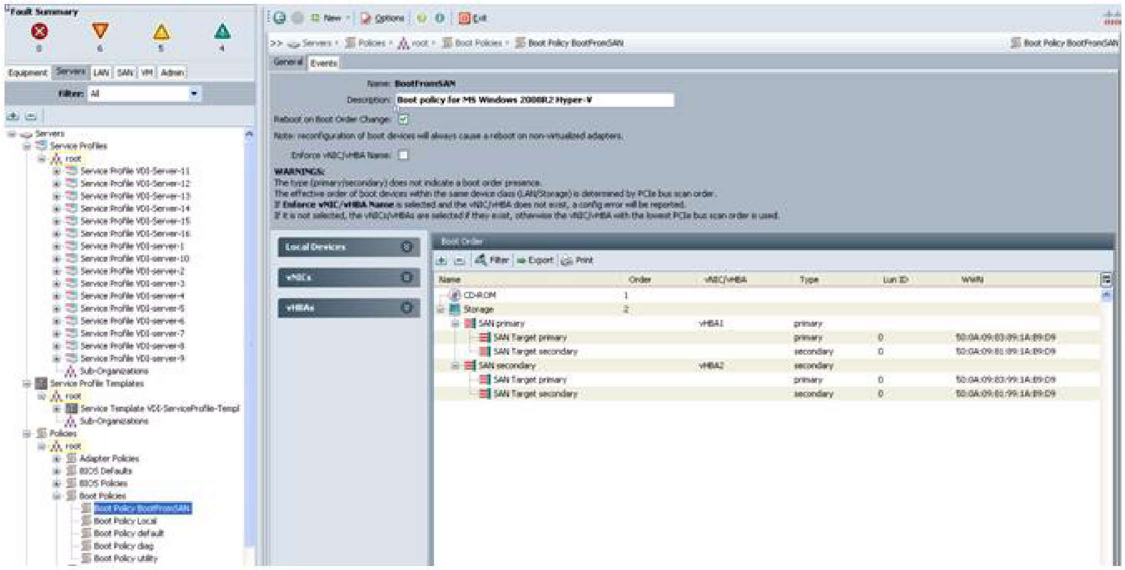

## Cisco UCS VLAN Configuration for Microsoft Windows 2008 R2 Hyper-V Hosts

First step is to configure the UCS manager LAN tab with the VLANs including native VLAN (164 in our case) and the desktop VLAN (122 in our case). For failover cluster configuration add a dummy VLAN 999 as shown in Figure 43.

*Figure 44 VLAN Configuration*

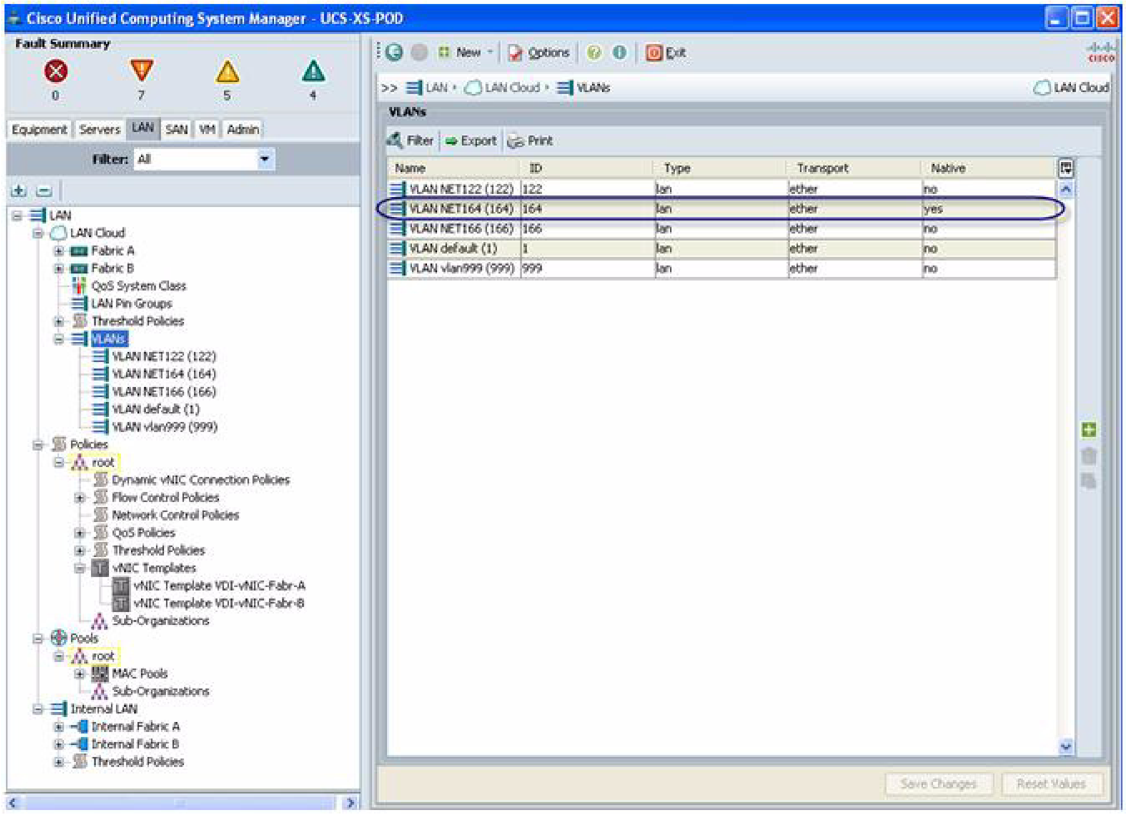


Second step is to configure the Cisco VIC adapters. In our case, we had two Cisco VIC and we configured two NICs each on the card for a total of four physical NICs. From a best practice purpose, we need one for the parent partition, one for the vSwitch, one for backend cluster communication (failover clustering) and one 'reserved' for Live Migration. Technically we could isolate VLANs to each one, but we just put all the VLANs on all the NICs except for the Cluster heartbeat NIC, it should stay on an isolated VLAN. Figure 44 below shows our networking configuration with eth2 and eth3 being used as the parent and vSwitch and eth0 being used for cluster heartbeat.

*Figure 45 Cisco VIC Configuration Showing Various VLAN and vNICS*

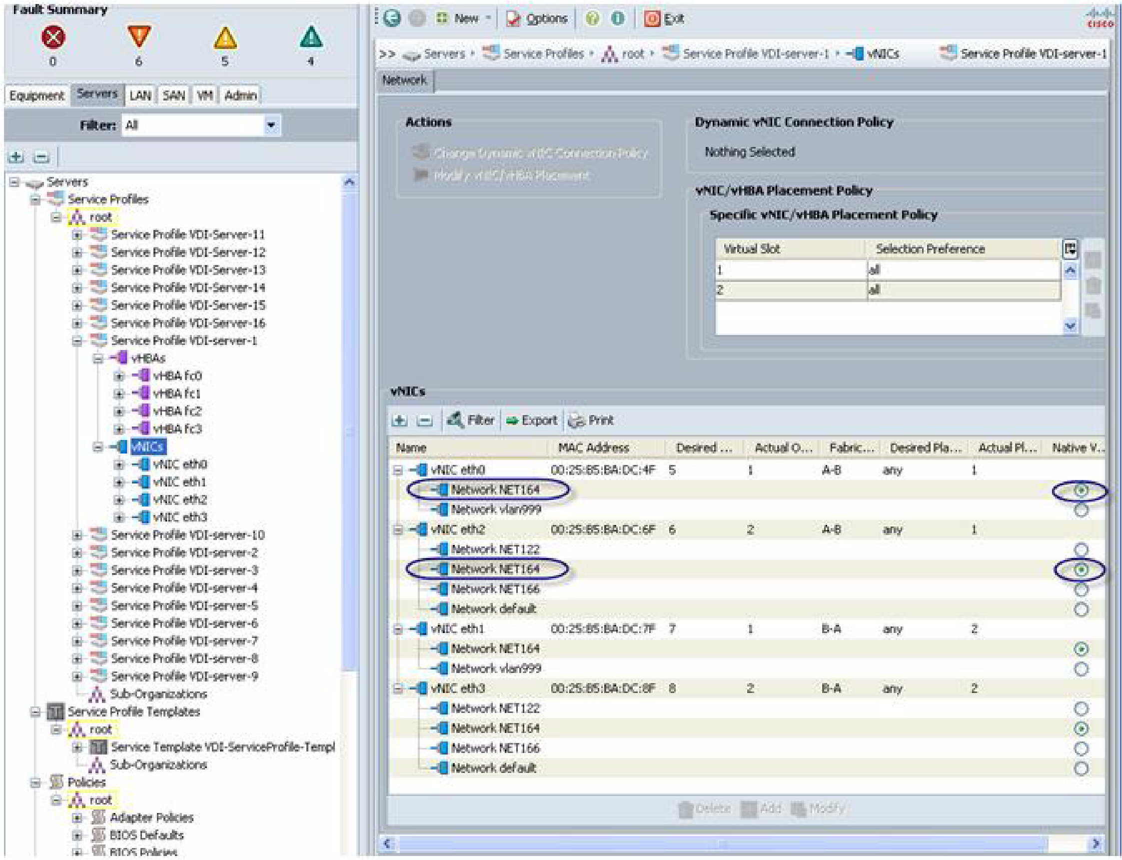

# Windows 2008 R2 Server Core and Hyper-V Configuration

This section details the Windows 2008 R2 Server Core installation; how to turn on the Hyper-V role and configuring the Windows failover clustering.

## Installation of Windows 2008 R2 Server Core

After the service profile is configured on the server and boot from SAN is in place, install the Windows 2008 R2 server, per the instructions as described in the following URLs:

- http://www.cisco.com/en/US/docs/unified_computing/ucs/sw/b/os/windows/install/2008-vmedia-install.html#wp1053494
- http://technet.microsoft.com/en-us/library/dd379511(WS.10).aspx

Important aspects to keep in mind during the installation on B250-M2 blade with Cisco VIC (M81KR) cards:

- Make sure you are installing the latest Cisco VIC driver during the boot process (Step 12.a in the link listed above)
- Make sure during the installation steps you pick "Server core" for the OS type (as shown below)
- During the installation, make sure that the Hyper-V hosts sees only one path. This is achieved by zoning on the MDS; this is required to initiate the installation process.

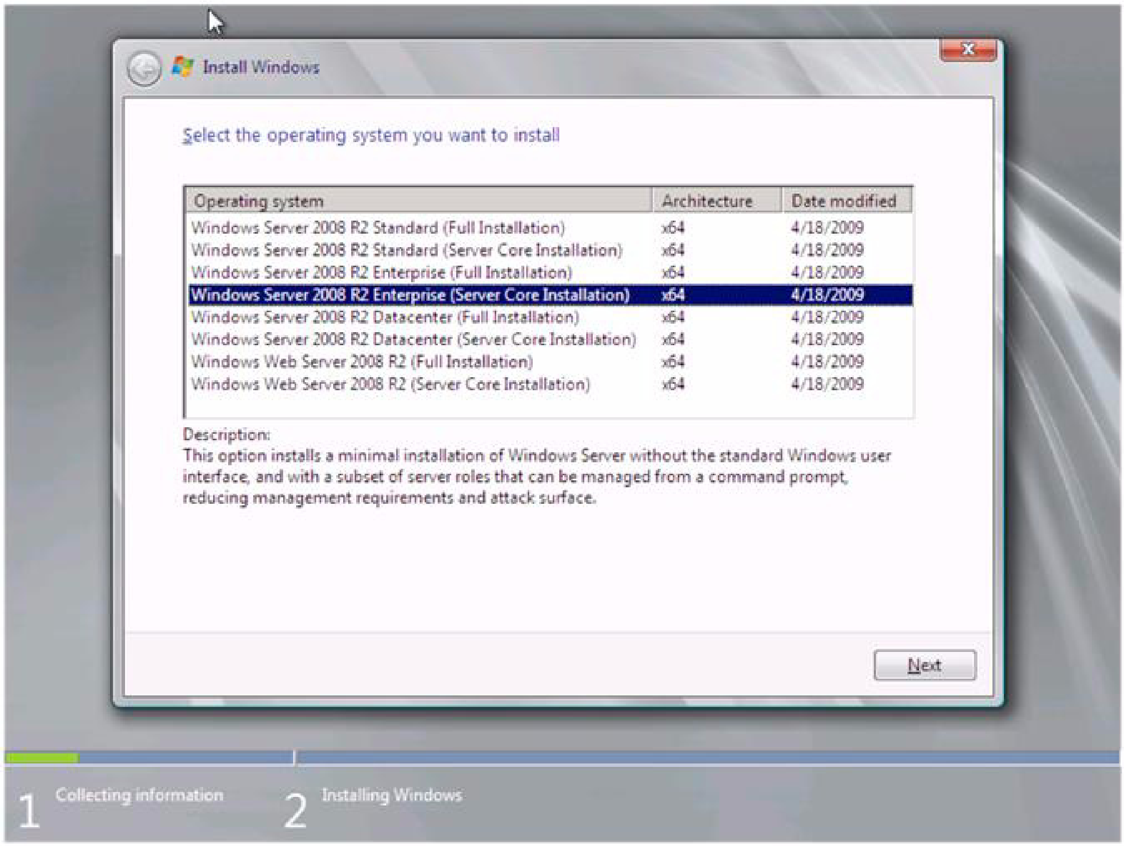


The rest of the install process is straightforward and when the install finishes, reboot the server and login as administrator.

## Add Hyper-V Role

From the command window execute dism /online /enable-feature /featurename:Microsoft-Hyper-V as administrator and restart the system when prompted. This will enable Hyper-V role on the server.

## Enable MPIO on the Golden Image

To enable the MPIO on NetApp storage, do the following operation from the command prompt on the Hyper-V server Core host.

1. dism /online /enable-feature /feature-name:MultipathIo.
2. Launch the MPIO Control panel (from a cmd line).

   mpiocpl
3. From the Discover Devices tab, select the NetApp array and select Add. This will reboot the host.

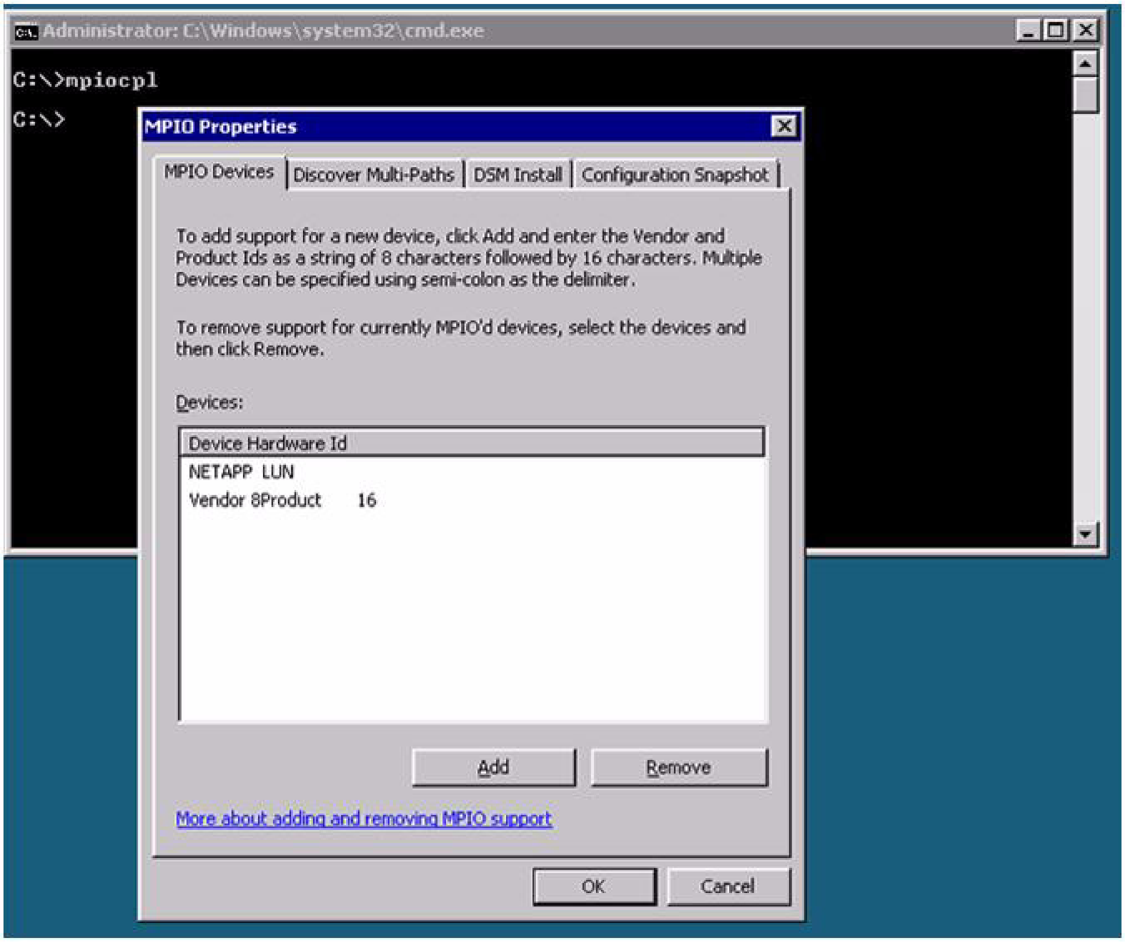


## Deploying Multiple Hyper-V Hosts

As the goal is to get 16 servers configured with Windows Server 2008 R2 with Hyper-V the install process is automated for efficiency and operational reasons. Following procedure is one of the quick ways to deploy Hyper-V servers quickly:

After the first server is installed using VMedia (described above) and Hyper-V role is enabled, login as administrator and run sysprep from the command window; select OOBE and Generalize and shutdown the machine.

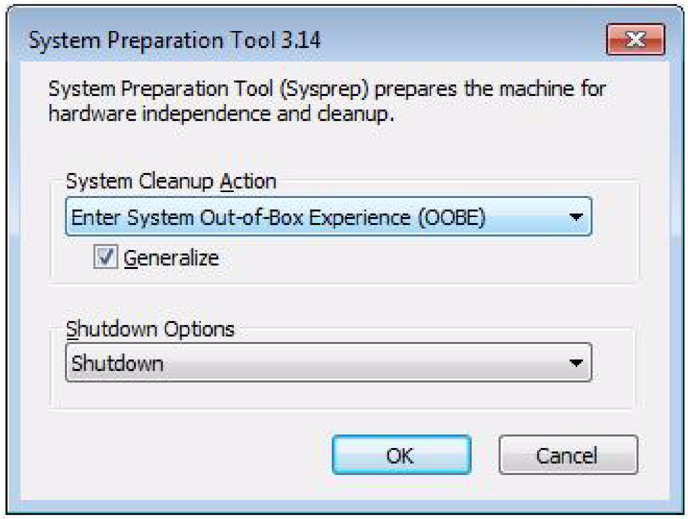


Using the script found in Appendix C, clones were created from the sysprep'ed Windows 2008 R2 golden boot image. A prerequisite is that each Boot LUN was mapped to the corresponding initiator ports of the Cisco UCS Servers on the NetApp storage as described in section and proper zoning is in place.

When all the host boot LUNs are cloned and servers are booted up, setup each server. Configuration steps include:

- Assign a name to the Hyper-V node
- Configure IP address, netmask and gateway
- Setup DNS server, proxy server
- Add license key and activate the key
- Join a domain
- Enable Failover Cluster feature

You can automate this using PowerShell or just run the steps described in Appendix D on each server to configure the hosts.

After the hosts are successfully booted up and has its own individual IP address, the cluster shared volumes and quorum disks are configured the next step is to configure the failover clustering.

## System Center Virtual Machine Manager

System Center Virtual Machine Manager (VMM) 2008 R2 provide centralized administration and management of virtual environment. It enables increased physical server utilization, and enables rapid provisioning of new virtual machines.

Every VMM implementation requires implementation of the following components:

- VMM server and default library server
- VMM database
- VMM Administrator Console

A detailed step-by-step instruction for installing VMM can be found at:

http://technet.microsoft.com/en-us/library/cc764339.aspx

You will install two SCVMM servers for the large scale testing and for each hosted eight Hyper-V hosts with a total of 880 desktop each. Snapshots of the SCVMM servers are shown below.

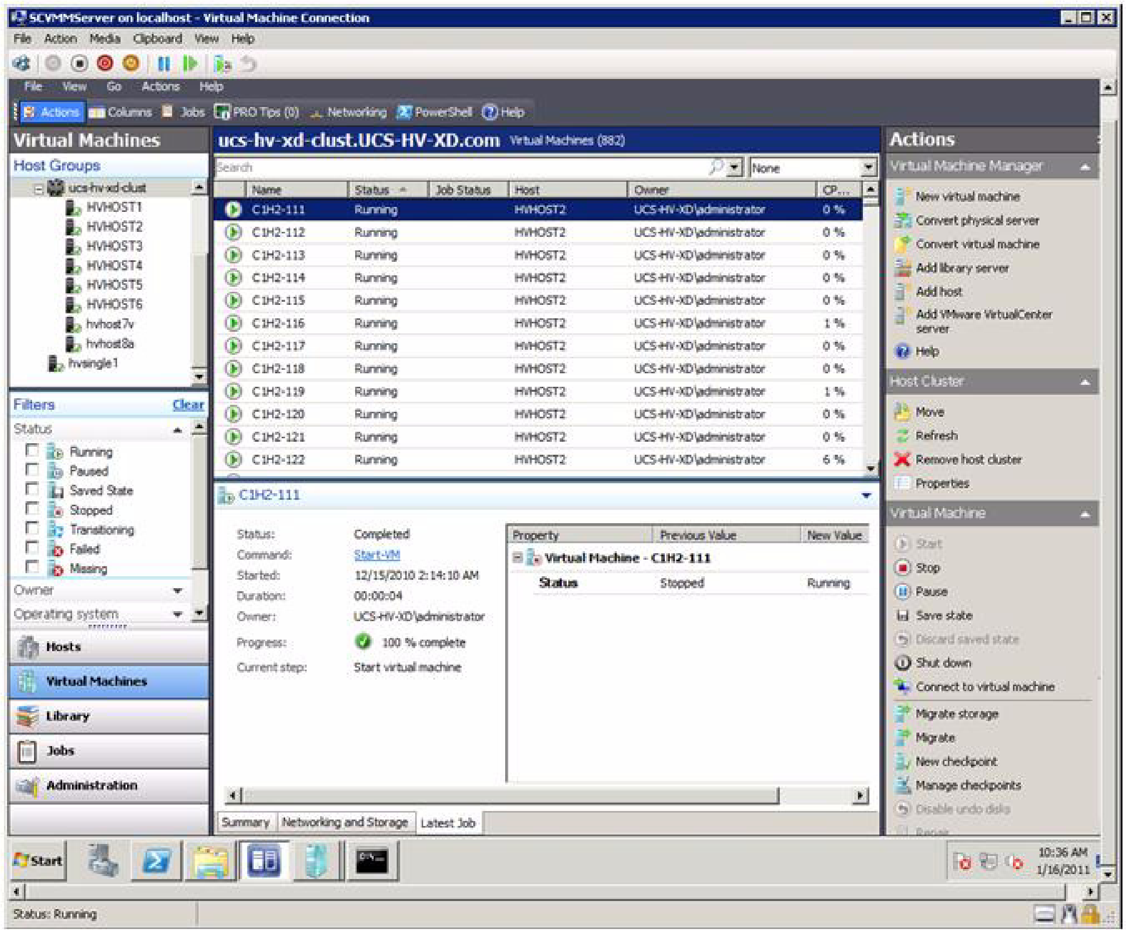

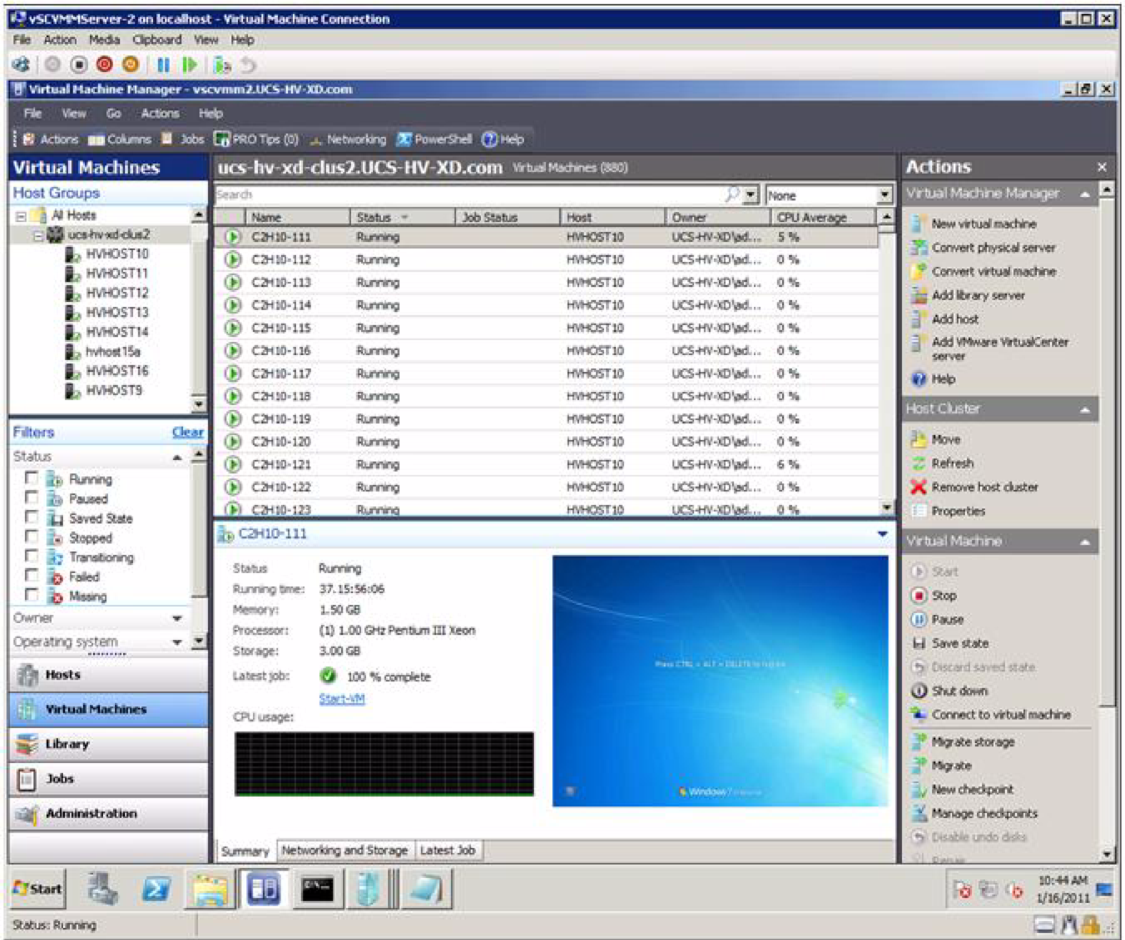


## Networking Configuration

The individual Hyper-V node networking configuration is described in Appendix D. The VNICs are enabled on the Cisco VIC to have the VLAN 122 visibility for the desktop network. This is a private network for the desktops that was enabled on the L3 switch in the configuration as 1800 odd IP addresses are required for the desktops. In the Hyper-V side, add a new network; Desktop-Network and tag it with VLAN 122. Figure 45 shows the System Center Virtual Machine Manager's (SCVMM) view of the networking configuration of HVHOST2 (one of the Hyper-V host) with the VLAN122 configured.

*Figure 46 System Center Virtual Machine Manager*

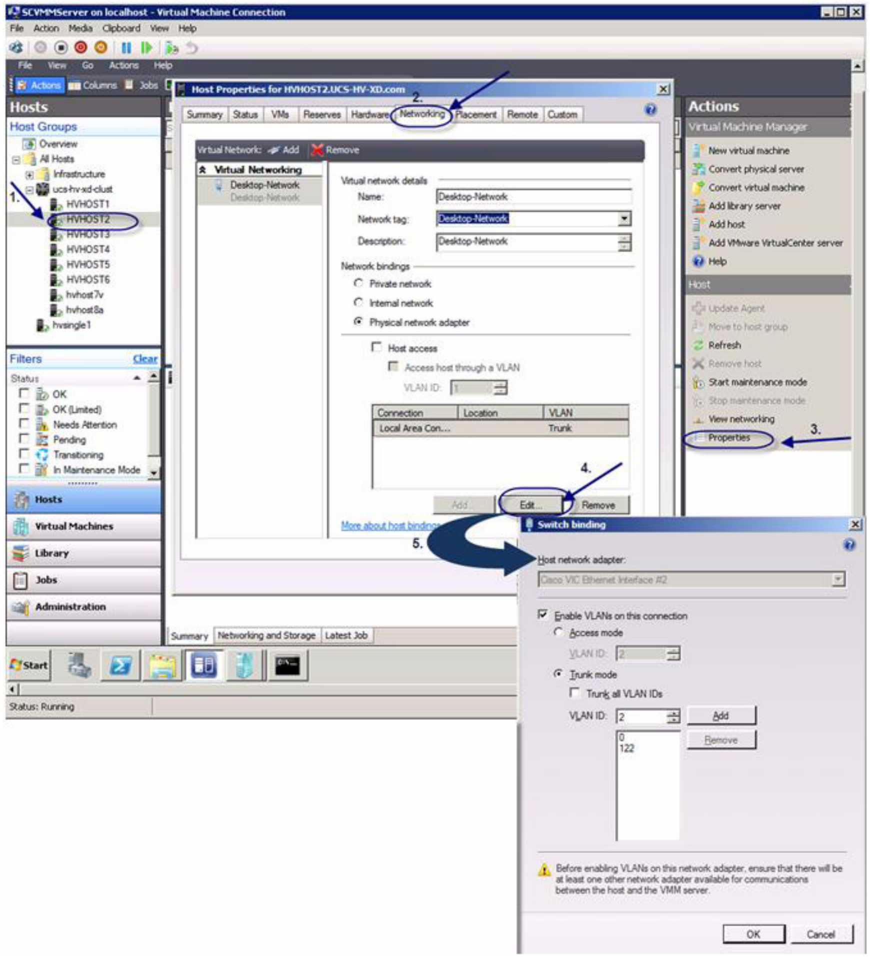


## Configuring Failover Clustering and Failover Cluster Manager

The Failover Clustering feature enables you to create and manage failover clusters. A failover cluster is a group of independent computers that work together to increase the availability of applications and services. The clustered servers or nodes are connected by physical cables and by software. If one of the cluster nodes fails, another node begins to provide service (a process known as failover). Users experience a minimum of disruptions in service.

The failover cluster manager snap-in is used to create clusters, manage the cluster nodes, enable clustered shared volumes, and other clustering configurations.

To open the failover cluster manager snap-in, click start, click administrative tools and then click failover cluster manager. Choose a suitable windows host (in our case, we used the Active Directory VM to run the failover cluster manager).

Follow the required steps to configure Hyper-V and failover clustering feature as described in the Technet documentation:

http://technet.microsoft.com/en-us/library/cc732035(WS.10).aspx

Figure 46 shows how we configured two clusters of eight nodes (Cisco UCS B250-M2 servers). It is recommend to run the validate cluster configuration before deploying the desktops.

*Figure 47 Failover Cluster Manager Showing Two Clusters of Eight Nodes Each*

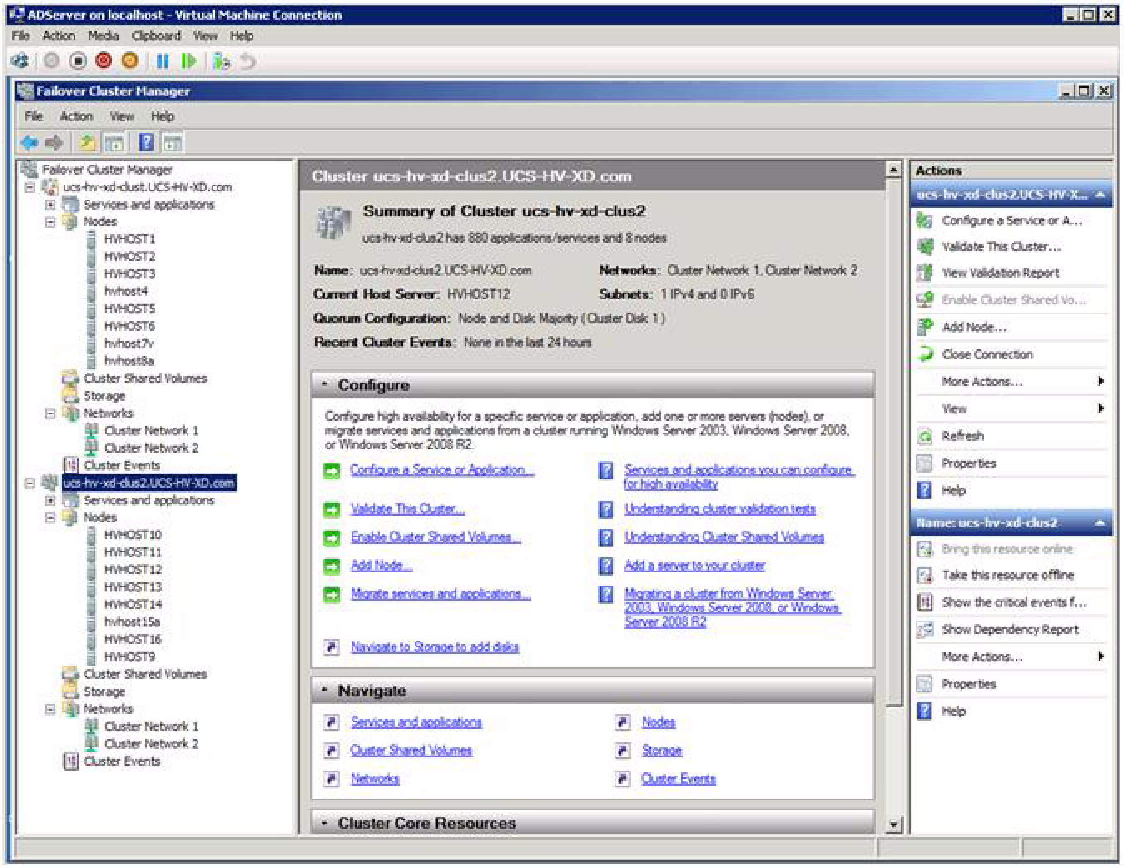


One of the important aspects of the failover clustering is the configuration of the Cluster Shared Volume and as shown in the Figure 47 below. Configure three volumes per cluster to be shared across a cluster. Prior to adding and enabling these cluster volumes, create a primary volume (GPT type), format the volume on the Hyper-V node, and create a NTFS filesystem.

*Figure 48 Configuring the Cluster Shared Volume*

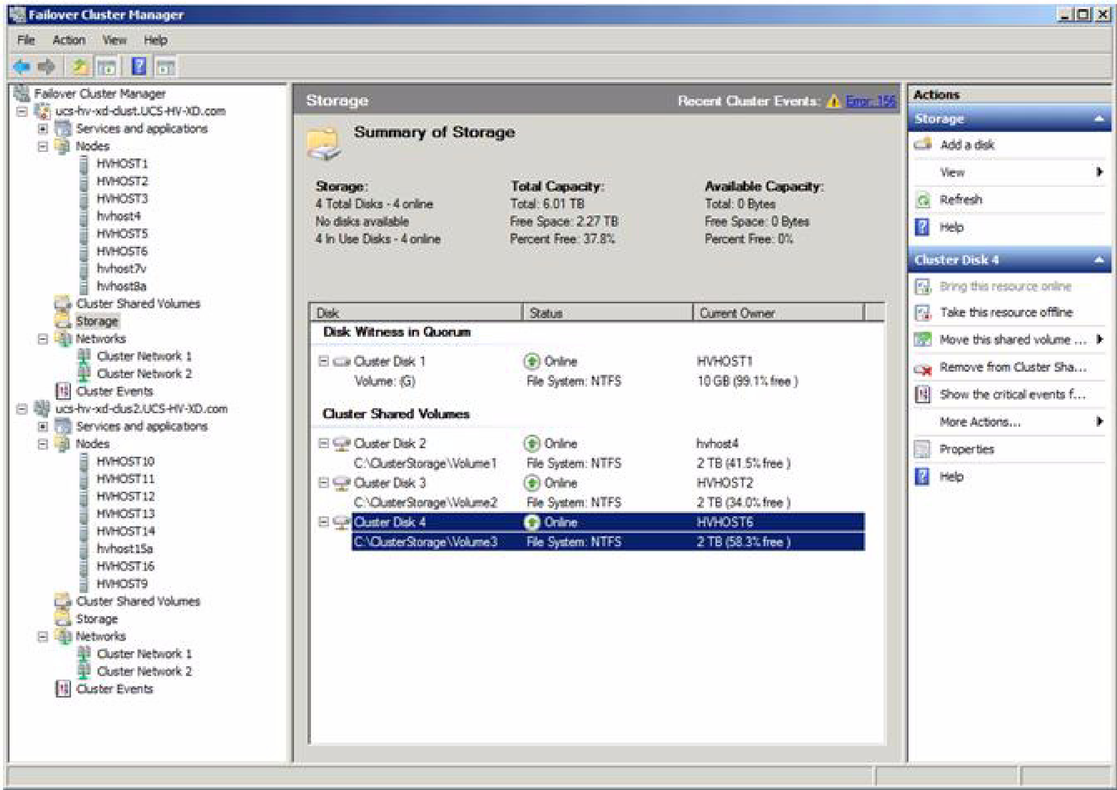


# Test Setup and Configuration

The test strategy is to perform a single server scalability to determine the maximum amount of desktops that can be loaded on a given server within an acceptable user response time based on success criteria parameters defined by LoginVSI. When a single server scalability is successfully completed, scale the environment to two chassis (8 blade servers) and then four chassis (16 blades servers). The following sections depict the single server configuration, detail the two chassis (8 Blade servers), and detail the four chassis (16 blade servers).

# Cisco UCS Test Configuration for Single-Server Scalability Test Setup

*Figure 49 Configuration of Cisco UCS B250 M2 Blade Server for Single-Server Scalability*

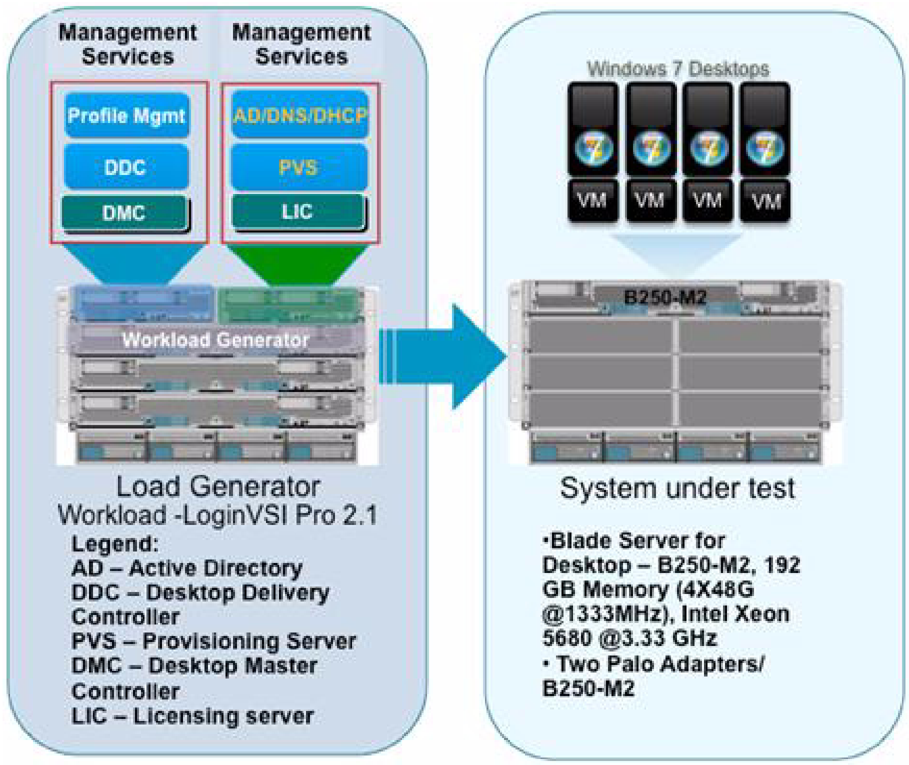


- Hardware components
  - 1 X Cisco UCS B250-M2 (5680 @ 3.33 GHz) blade servers with 192 GB of memory (4 GB X 48 DIMMS @ 1333 MHz)
  - 2 X Cisco UCS B200-M2 (5680 @ 3.33 GHz) blade servers with 48 GB of memory (4 GB X 12 DIMMS @ 1333 MHz)
  - Cisco UCS M81KR (VIC) (two per server)
  - Cisco Nexus 5500 and 7000 Series
  - NetApp FAS 3170 storage array, two controllers
- Software components
  - Cisco UCS firmware 1.3(1i)
  - Microsoft Windows Server 2008 R2 Hyper-V
  - One SCVMM Server
  - XenDesktop 4
    - 2 DDC Servers
    - 2 PVS Servers
- Windows 7—32 bit, 1 vCPU, 1.5 GB of memory, 3GB write cache per virtual machine

# Cisco UCS Configuration for Two-Chassis Test

*Figure 50 Two-Chassis Test Configuraiton—8x Cisco UCS B250 Blade Server*

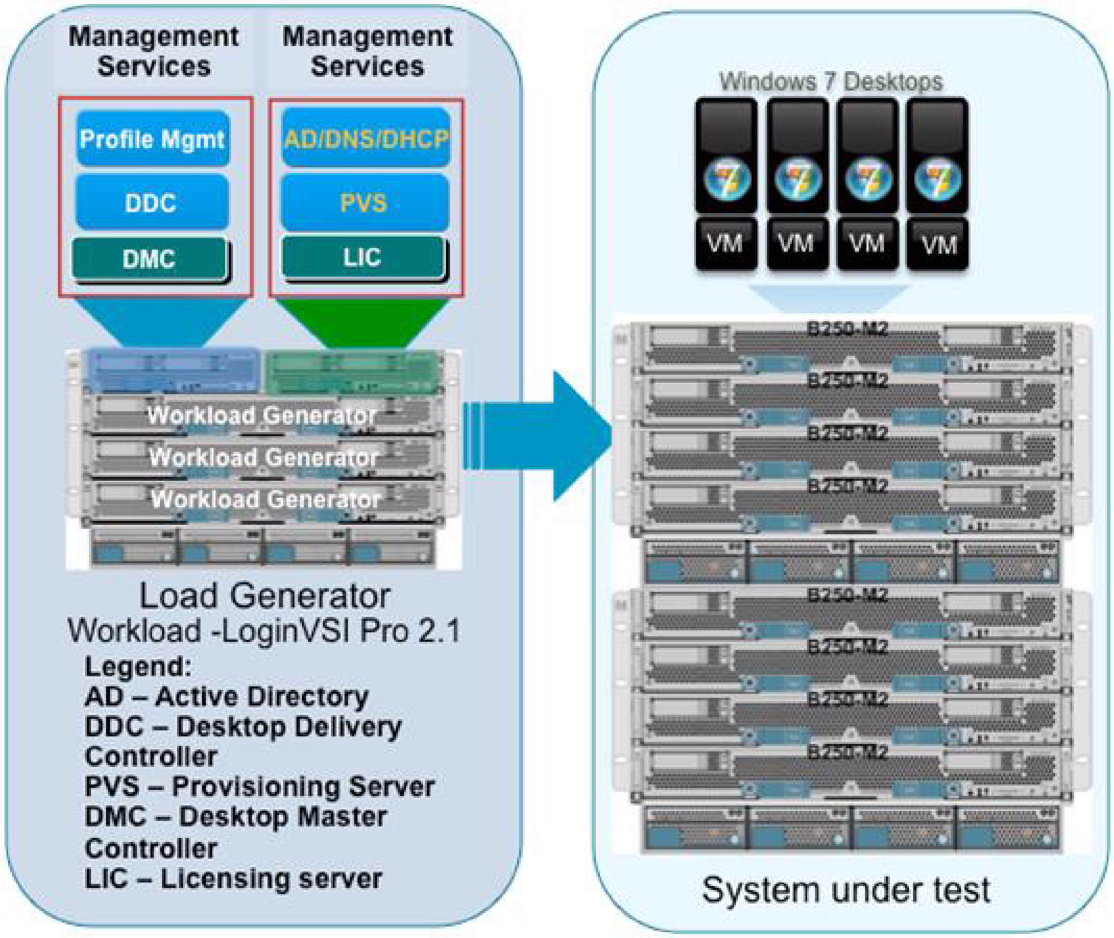


- Hardware components
  - 8 X Cisco UCS B250-M2 (5680 @ 3.33 GHz) blade servers with 192 GB of memory (4 GB X 48 DIMMS @ 1333 MHz)
  - 2 X Cisco UCS B200-M2 (5680 @ 3.33 GHz) blade servers with 48 GB of memory (4 GB X 12 DIMMS @ 1333 MHz)
  - Cisco UCS M81KR (VIC) (two per server)
  - Cisco Nexus 5500 and 7000 Series
  - NetApp FAS 3140 and FAS 3170 storage array, two controllers with 72 SAS drives
- Software components
  - Cisco UCS firmware 1.3(1i)
  - Microsoft Windows Server 2008 R2 Hyper-V
  - One SCVMM Server
  - XenDesktop 4
    - 2 DDC Servers

- 2 PVS Servers
- Windows 7 - 32 bit, 1vCPU, 1.5 GB of memory, 30 GB per virtual machine

# Cisco UCS Configuration for Four-Chassis Test

*Figure 51 Cisco UCS Entry Bundle with Additional Scale Bundles*

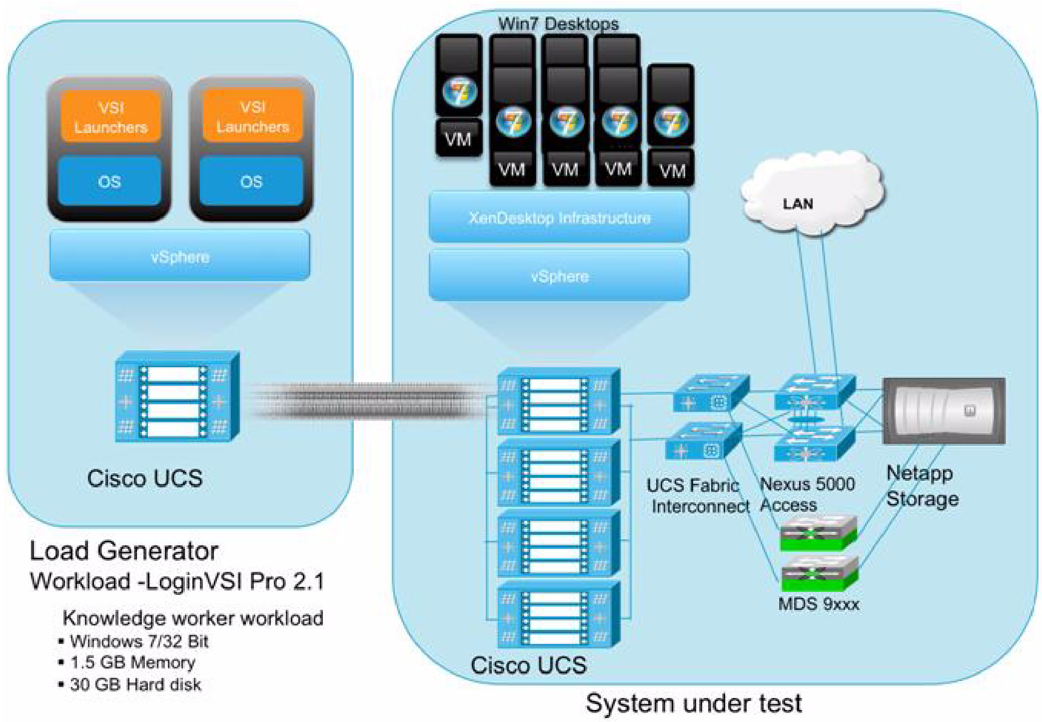


- Hardware components
  - 16 X Cisco UCS B250 M2 (5680 @ 3.33 GHz) blade servers with 192 GB of memory (4 GB X 48 DIMMS @ 1333 MHz)
  - 2 X Cisco UCS B200 M2 (5680 @ 3.33 GHz) blade servers with 48 GB of memory (4 GB X 12 DIMMS @ 1333 MHz)
  - Two Cisco UCS M81KR (VIC) on each B250-M2 blade server
  - Cisco Nexus 5500 and 7000 Series
  - NetApp FAS 3140/FAS 3170 storage array, two controllers with 72 SAS drives
- Software components
  - Cisco UCS firmware 1.3(1i)
  - Microsoft Windows Server 2008 R2 Hyper-V
  - Two SCVMM Servers
  - XenDesktop 4
    - 2 DDC Servers
    - 3 PVS Servers
- Windows 7 - 32 bit, 1vCPU, 1.5 GB of memory, 30 GB per virtual machine

# Testing Methodology

All validation testing was conducted on-site within the Cisco labs with joint support from Citrix, Microsoft 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, user workload execution (also referred to as steady state), and user logoff for the Hosted VDI model. 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 metrics were within the permissible thresholds as noted as success criteria. Test were conducted a total of three times 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 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, provide unique user credentials, initiate the workload, and evaluate the end user experience.

For the Hosted VDI environment an internal Citrix automated test tool was used to generate end user connections into the environment and record performance metrics through an agent running on the core XenDesktop infrastructure components.

## User Workload Simulation—Login VSI 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 platform to allow customers to realistically reference for a variety of worker tasks. To accurately represent a real-world user workload, third-party tools from Login Consultants were used throughout the testing. These tools have the added benefit of taking measurements of the in-session user response time providing an objective way to measure the expected user experience for an individual desktop throughout large scale testing, including login storms.

Login Virtual Session Indexer (Login Consultants VSI 2.1) 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. 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-heavy workload user (intensive knowledge worker) running generic applications such as: Microsoft Office 2007, Internet Explorer including Flash applets and Adobe Acrobat Reader (Note: For the purposes of this test, applications were installed locally, not streamed or hosted on XenApp). Like real users, the scripted session will leave multiple applications open at the same time. Every session will average about 20% minimal user activity, similar to real world usage. Note that during each 12 minute loop users open and close files a couple of time per minutes which is probably more intensive that most users.

The following outline the automated Login VSI simulated user workflows that were used for this validation testing:

- This workload emulates a medium "knowledge worker" using the Office 2007, IE and PDF applications and opens up to 5 applications simultaneously with a type rate of 160ms for each character. The workload observes approximately 2 minutes of idle time which closely simulates real-world users.

- When a session has been started the medium workload will repeat every 12 minutes. During each loop the response time is measured every two minutes.
- Each loop consists of the following operations:
  - Browse and compose Outlook 2007 messages.
  - Open multiple instances of Internet Explorer based browsing sessions including heavy multimedia websites.
  - Open multiple instances of Word 2007 performing open, close and edit operations.
  - Print and review PDF documents using Bullzip PDF Printer and Acrobat Reader.
  - Open, edit and close a randomized large Excel 2007 sheet.
  - Review and edit a PowerPoint 2007 presentation.
  - Perform zip operations using 7-Zip.

## 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 two main metrics, Login VSI Max and Login VSI Correct Optimal Performance Index (COPI). The Login VSI Max evaluates the user response time during increasing user load and the Login VSI COPI score assess the successful start-to-finish execution of all the initiated virtual desktop sessions. These two main metrics are important not only based on the raw data that they provide, but also in their ability to align the test results between multiple vendor and customer environments.

### Login VSI Corrected Optimal Performance Index (COPI)

The Corrected Optimal Performance Index (COPI) is a calculated from specific measurements during each test run to determine how many desktops can be run simultaneously without excessively impacting user experience.

The corrected optimal performance index is based on these measurements:

- The Uncorrected Optimal Performance Index (UOPI) is based on the first 5 consecutive sessions that hit the Optimal Performance Max Reached threshold.
- The "Optimal Performance Max Reached" value is calculated on the response time average of four sessions higher than 2000ms (4 session average response time > 8000 ms).
- The Stuck Session Count (SSC) represents sessions which have become stuck before UOPI, and must therefore be accounted for in the Optimal Performance Index.
- The Lost Session Count (LSC) is a count of completely missing log files; these tests are discarded completely in the corrected index.
- The Corrected Optimal Performance Index (COPI) is then calculated:
- Incorporating the SSC and LSC into a corrected index helps ensure that the test results are fair and comparable. Therefore, the COPI is calculated as:

  COPI=UOPI - (SSC*50%) - LSC

### 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 2000ms and all users response times are expected to be less than 2000ms 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 2000ms for 6 consecutive occurrences. If VSI Max is reached, then the test run is considered a failure given that 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.

# Test Results

The purpose of this testing is to provide the data needed to validate Citrix XenDesktop 4 Hosted VDI FlexCast model with Microsoft Windows Server 2008 R2 Hyper-V virtualizing Microsoft Windows 7 desktops on Cisco UCS blade servers using a NetApp FAS 3140 storage array. 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.

## Citrix XenDesktop Hosted VDI Test Results

This section details the results from the XenDesktop Hosted VDI validation testing. The primary success criteria metrics are provided to validate the overall success of the test cycle. Additional graphs emphasizing the CPU and Memory utilization during peak session load are also present given that Memory consumption was found to be the most limiting factor to prevent further desktops from being hosted in both respective environments. 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.

### Single Cisco UCS Blade Server Validation

The first process in the validation was to ensure that a single Cisco UCS blade server was able to support the desired load of 110 virtual desktops per server. When identifying how many virtual desktops per server, it was important to assess the total available RAM. Each virtual desktop was configured with 1.5 GB of RAM and each blade had 192 GB of RAM available. With 110 virtual desktops on the server, the memory utilized on the environment was slated to be 165 GB of RAM before any hypervisor overhead, therefore making memory ~85% utilized. Based on this analysis the following 110 number of virtual desktops per blade was chosen.

The following table provides the VSI COPI score for the overall single Cisco UCS blade server environment and shows that 100 percent of all the 110 virtual desktop sessions executed without issue.

*Table 12 Single Cisco UCS Blade Server VSi COPI Score*

| | |
|---|---|
| Total Sessions Launched | 110 |
| Uncorrected VSI Max (UOPI) | 110 |
| Stuck Session Count before UVM (SSC) | 0 |
| Lost Session Count before UVM (LSC) | 0 |
| **Correct Optimal Performance Index (COPI = UOPI – (SSC*50%) – LSC)** | **110** |

After it can be confirmed that all 110 sessions executed successfully, it is important to help ensure that the user experience was not degraded as load was increased on the environment. The user response time, as reflected in Login VSI Max Pass or Fail rating, provides the necessary guidance to evaluate the user experience based on workload response time. From the graph below, it can be concluded that the user response time was not affected by the heavy 110 desktop load given that all response times are below the 2000ms threshold.

*Figure 52 110 Desktop Sessions on Microsoft Hyper-V Below 2000ms*

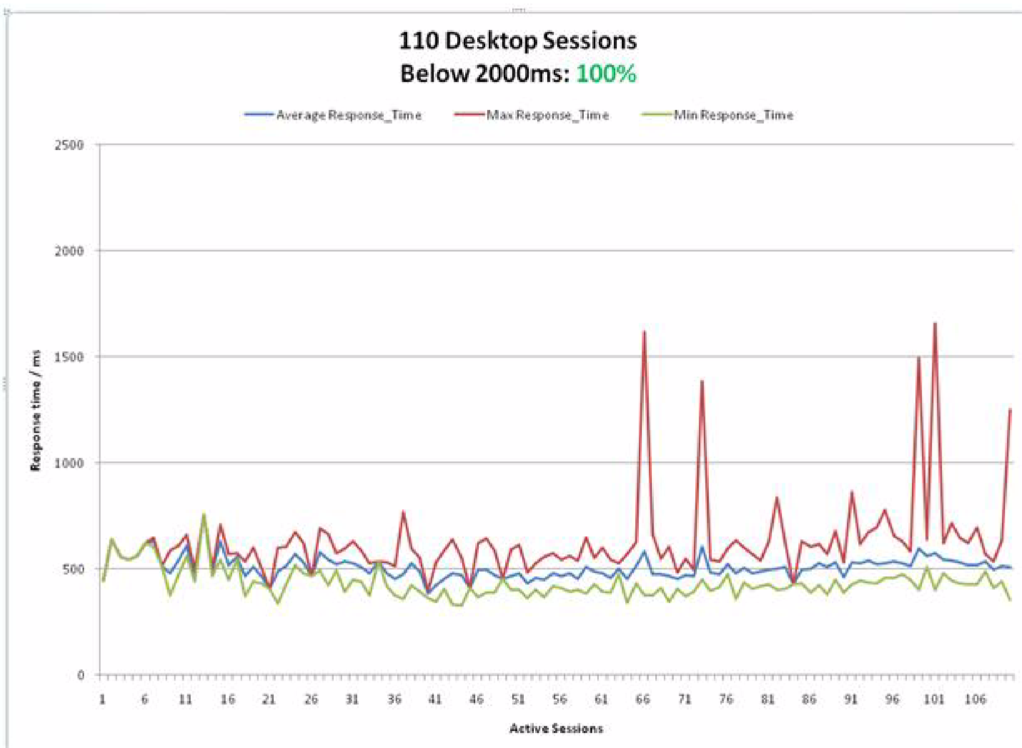


The following graph represents the CPU utilization during the desktop boot up phase for the single server scalability test.

*Figure 53 110 Desktop Boot Up—Microsoft Hyper-V host CPU Utilization*

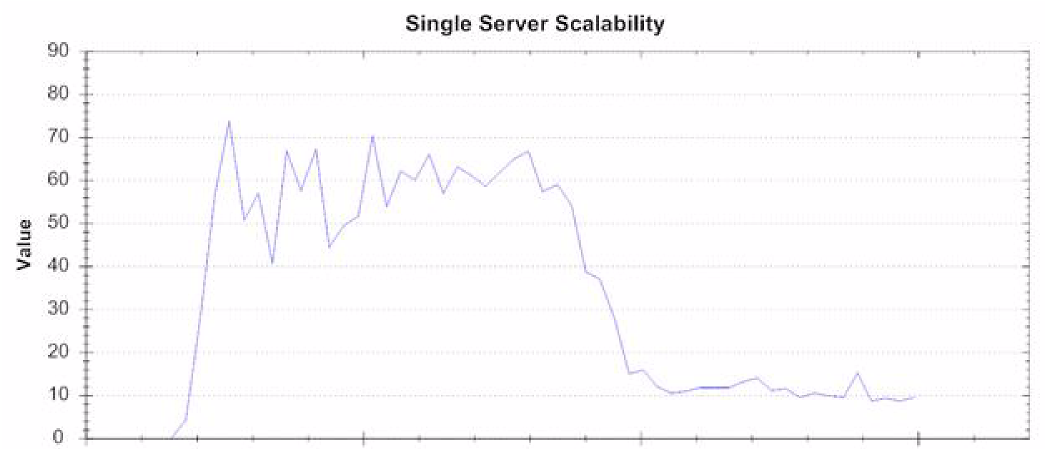


The following graphs represent the CPU utilization during the user workflow steady state execution phase of the single server scalability test.

*Figure 54 110 Desktop Steady State—Microsoft Hyper-V host CPU Utilization*

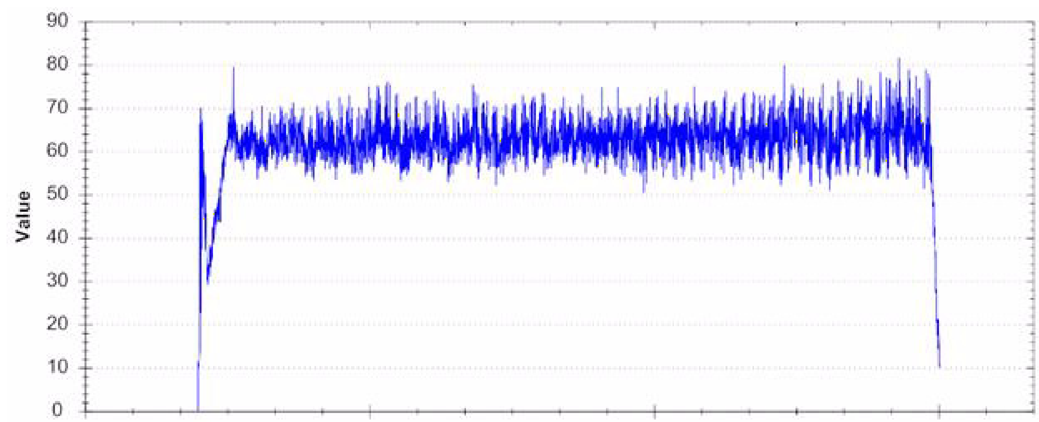


*Figure 55 110 Desktop Steady State—Microsoft Hyper-V Available Memory MBytes*

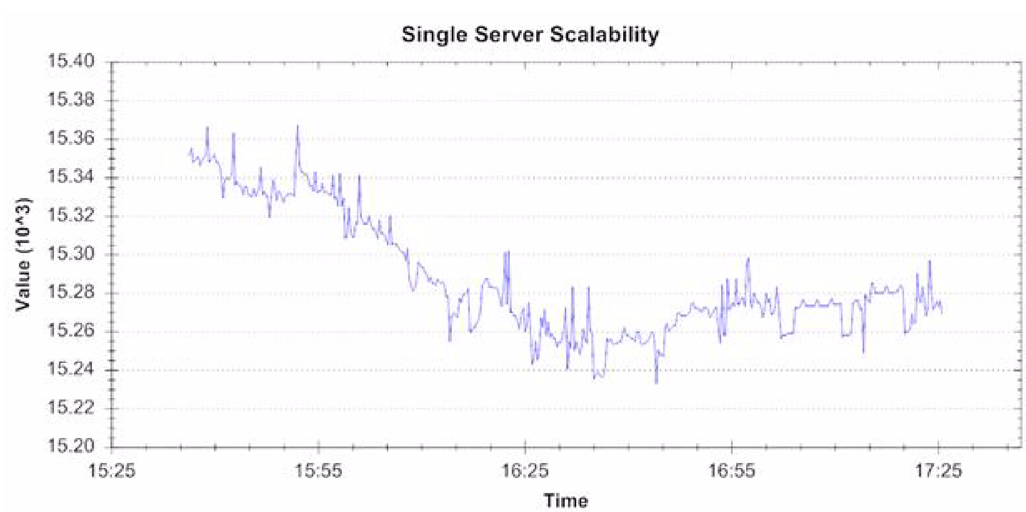

*Figure 56 110 Desktop Steady State—Microsoft Hyper-V Disk Transfers per Second*

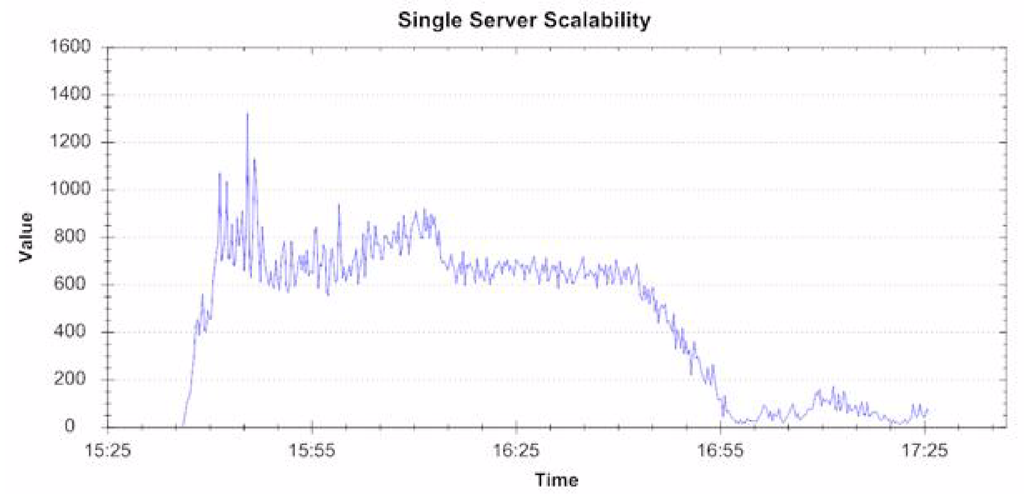


*Figure 57 110 Desktop Steady State—DDC1 CPU Utilization*

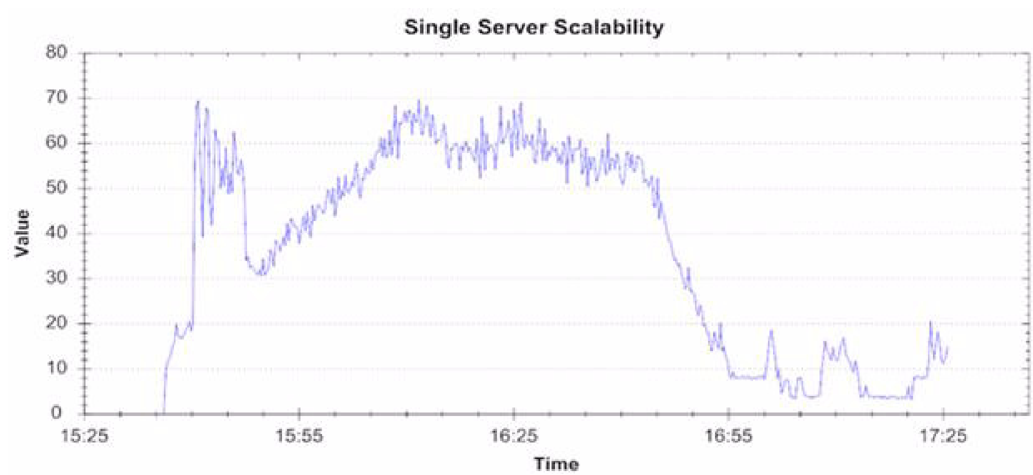


## Two Cisco UCS Blade Chassis Validation

The two Cisco UCS blade chassis environment contained a total of 8 blades with 192 GB of RAM per blade.

The following table provides the VSI COPI score for the overall 8-Cisco UCS blade environment and shows that 100% of all the 880 virtual desktop sessions executed without issue.

*Table 13 Eight Cisco UCS Blade Servers VSI COPI Score*

| | |
|---|---|
| Total Sessions Launched | 880 |
| Uncorrected Optimal Performance Index (UOPI) | 880 |
| Stuck Session Count before UOPI(SSC) | 0 |
| Lost Session Count before UOPI (LSC) | 0 |
| **Corrected Optimal Performance Index (COPI = UOPI – (SSC*50) - LSC)** | **880** |

From the graph below, it can be concluded that the user response time was not affected by the heavy 880 desktop load given that all response times are below the 2000ms threshold.

*Figure 58 880 Desktop Sessions on Microsoft Hyper-V Below 2000ms.*

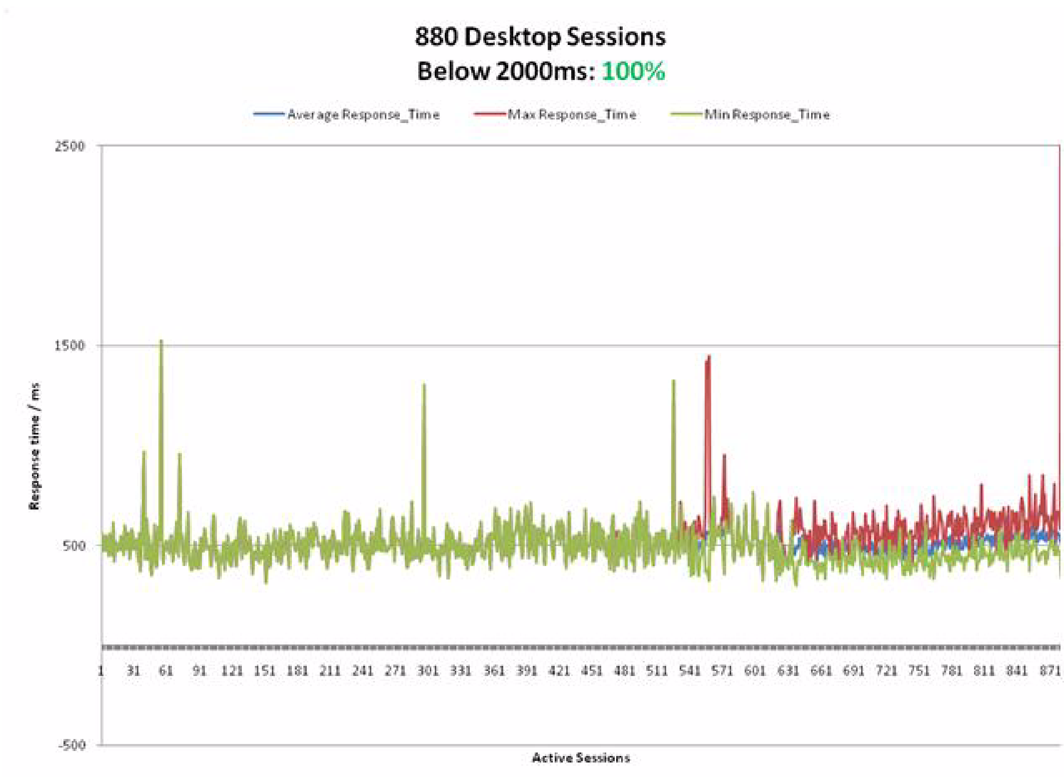


The following graph represents the CPU utilization during the desktop boot up phase of the 8-blade (node) scalability test.

*Figure 59 880 Desktop Boot Up—Microsoft Hyper-V hosts CPU Utilization*

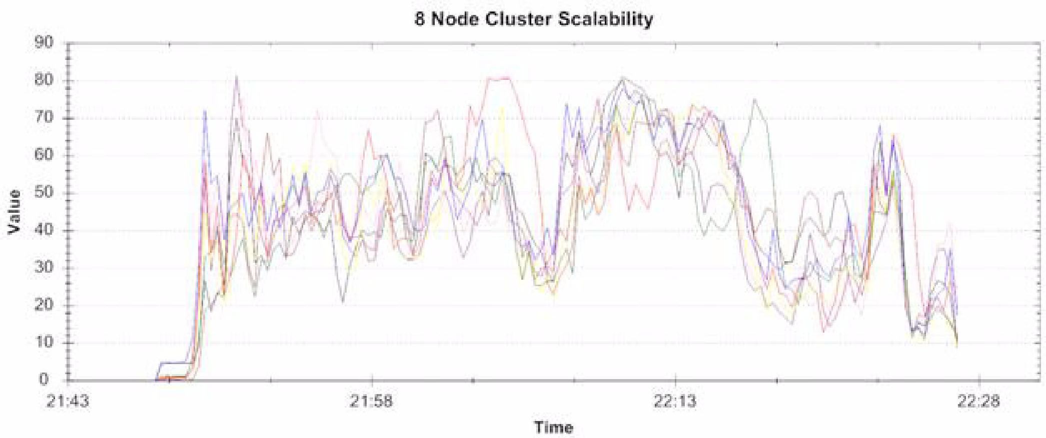


*Figure 60 880 Desktop Boot Up—PVS 1 CPU Utilization*

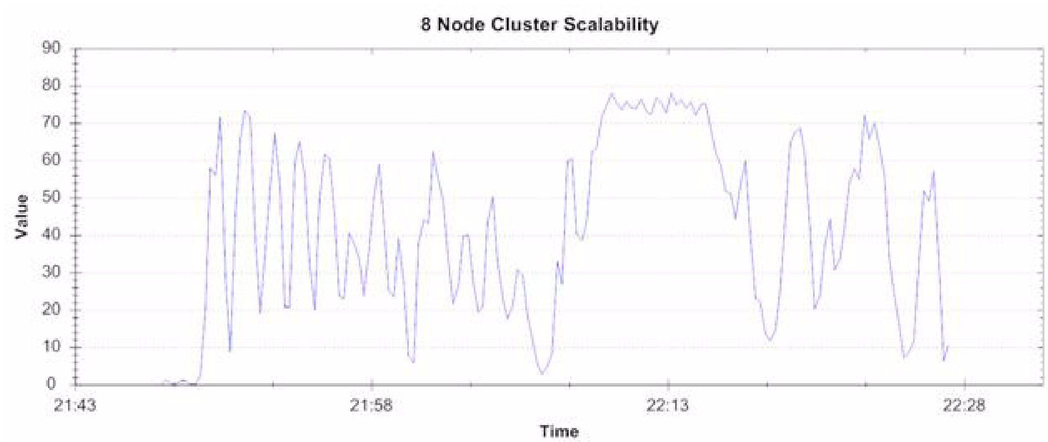


*Figure 61 880 Desktop Boot Up—PVS 1 Network Utilization*

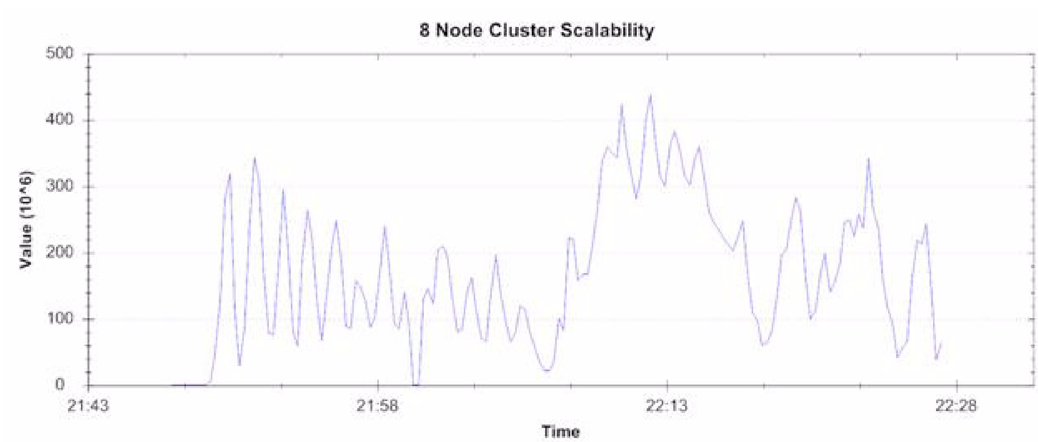

*Figure 62 880 Desktop Boot Up—DDC 1 CPU Utilization*

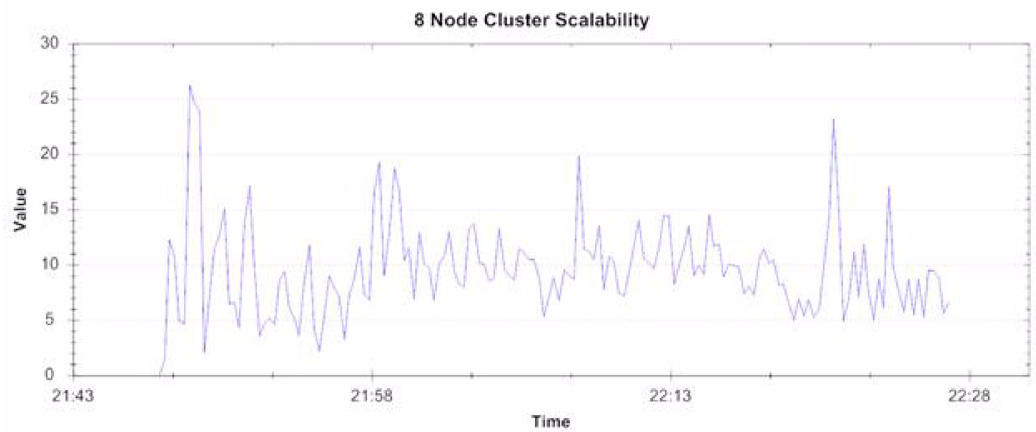


*Figure 63 880 Desktop Boot Up—VMM Server CPU Utilization*

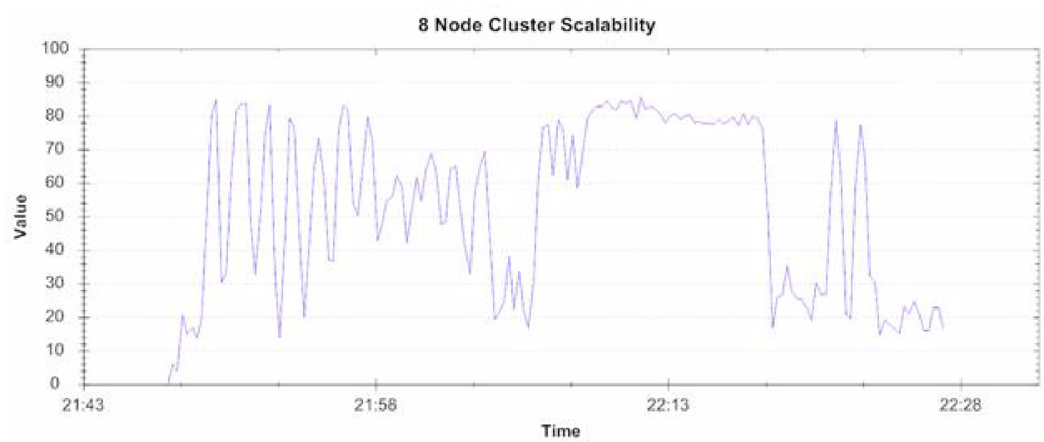


The following graphs represent the CPU utilization during the user workflow steady state execution phase of the 8 node scalability test.

*Figure 64 880 Desktop Steady State—Microsoft Hyper-V hosts CPU Utilization*

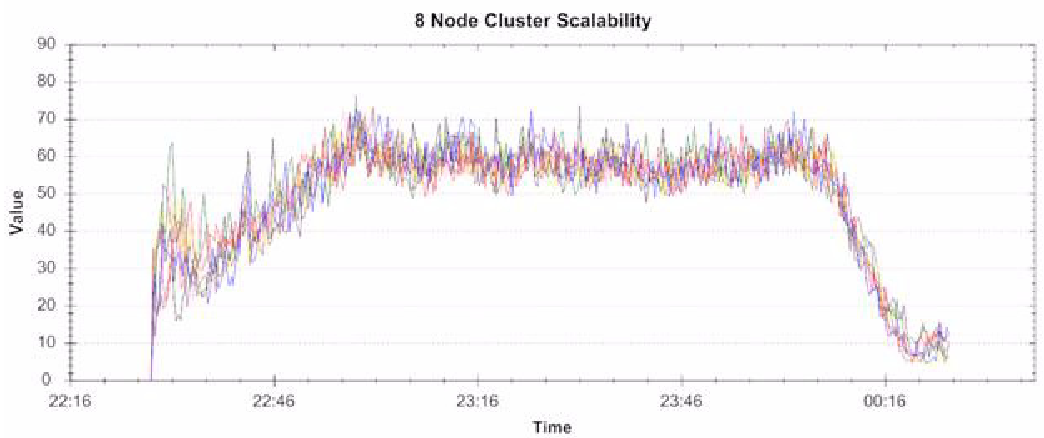


*Figure 65 880 Desktop Steady State—VMM Server CPU Utilization*

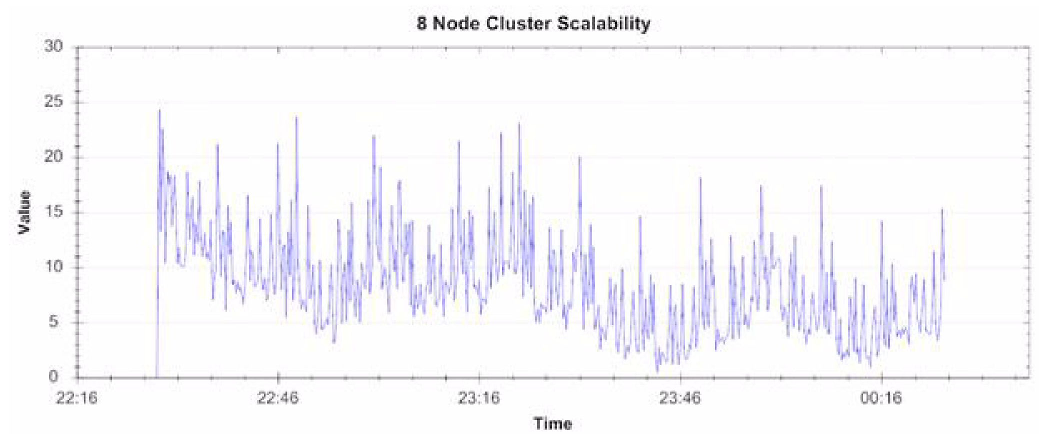


*Figure 66 880 Desktop Steady State—DDC 1 CPU Utilization*

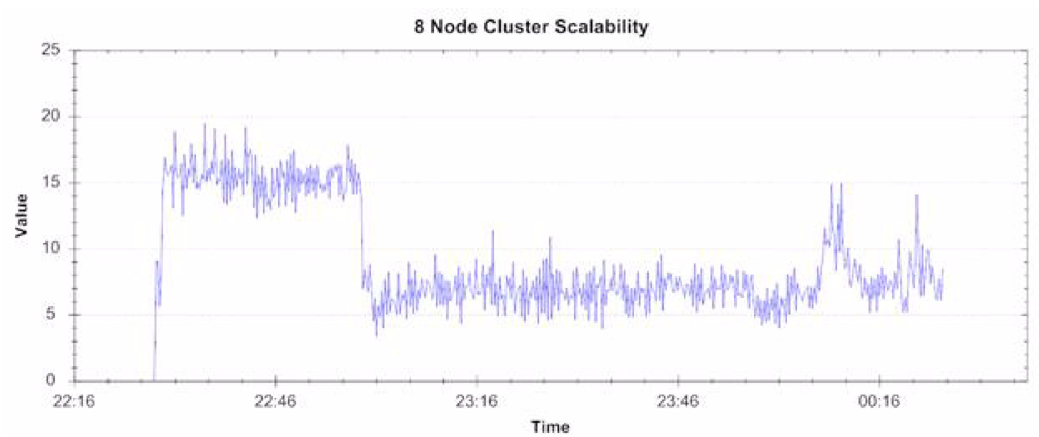


# Four Cisco UCS Blade Chassis Validation

The four Cisco UCS blade chassis environment contained a total of 16 blades with 192GB of RAM per blade. The following table provides the VSI COPI score for the overall 16-UCS blade environment and shows that 1755 of all the 1760 virtual desktop sessions executed without issue.

*Table 14 16 Cisco UCS Blade Servers VSI COPI Score*

| | |
|---|---|
| Total Sessions Launched | 1760 |
| Uncorrected VSI Max (UVM) | 1760 |
| Stuck Session Count before UVM (SSC) | 3 |
| Lost Session Count before UVM (LSC) | 2 |
| **VSI Max (VSIMax = UVM - SSC - LSC)** | **1755** |

From the graph below, it can be concluded that the user response time was not affected by the heavy 1760 desktop load given that all response times are below the 2000ms threshold.

*Figure 67 1760 Desktop Sessions on Microsoft Hyper-V Below 2000ms*

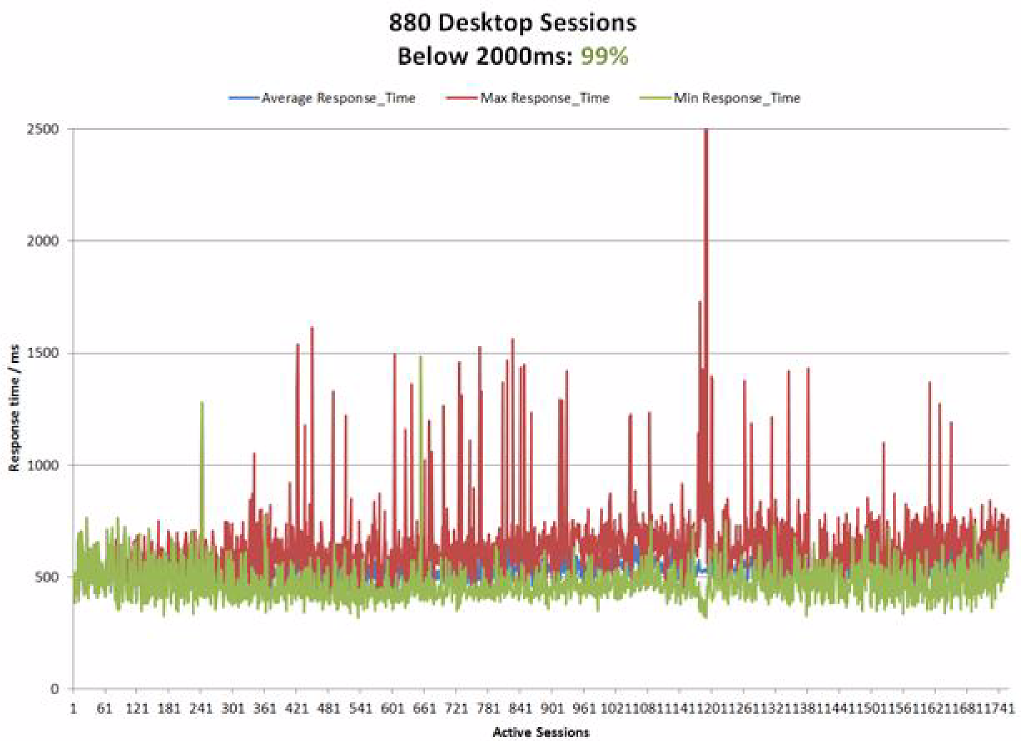


The following graph represents the CPU utilization during the desktop boot up phase of the 16 node scalability test.

*Figure 68 1760 Desktop Boot Up—16 Microsoft Hyper-V hosts CPU Utilization*

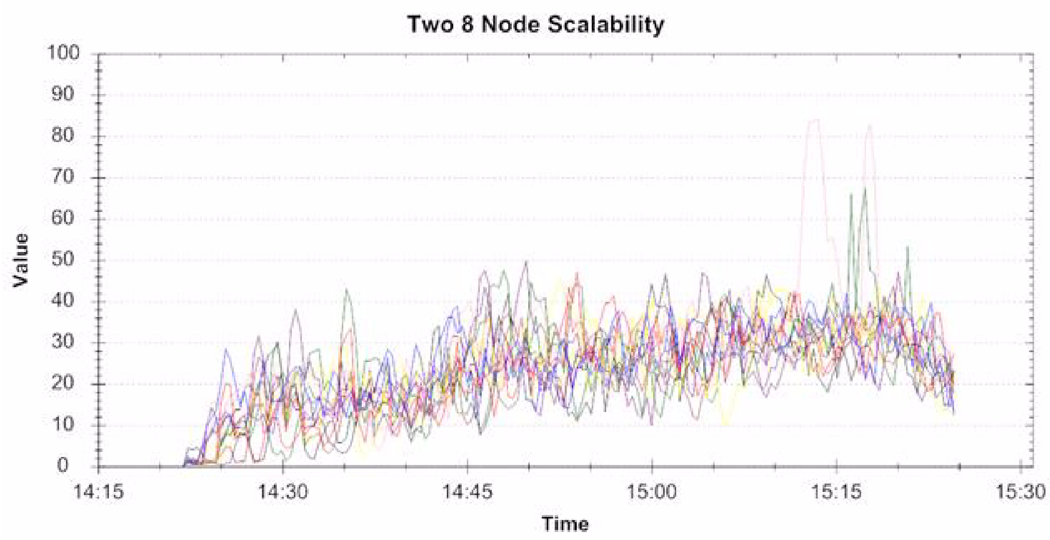


*Figure 69 1760 Desktop Boot Up—3 PVS Servers CPU Utilization*

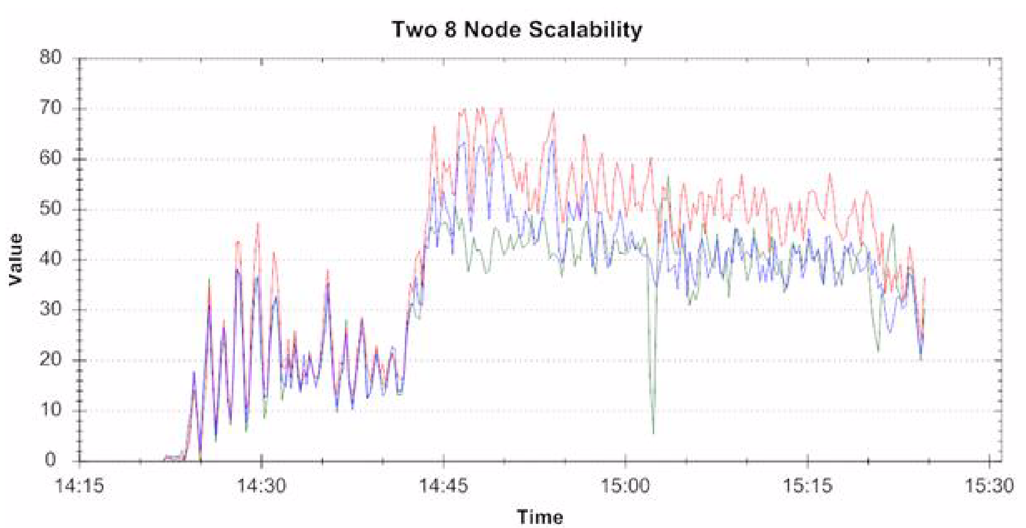

*Figure 70 1760 Desktop Boot Up—3 PVS Servers Network Utilization*

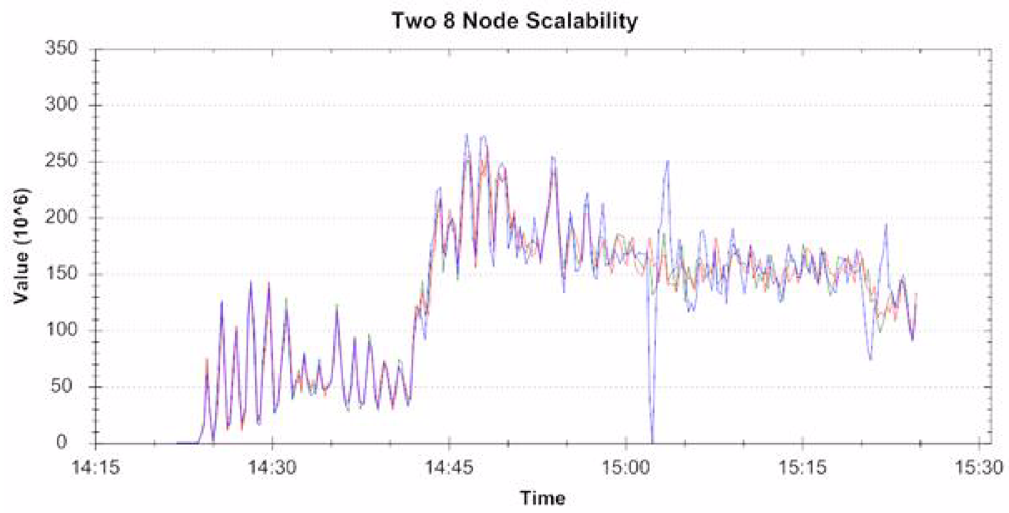


*Figure 71 1760 Desktop Boot Up—2 DDC Servers CPU Utilization*

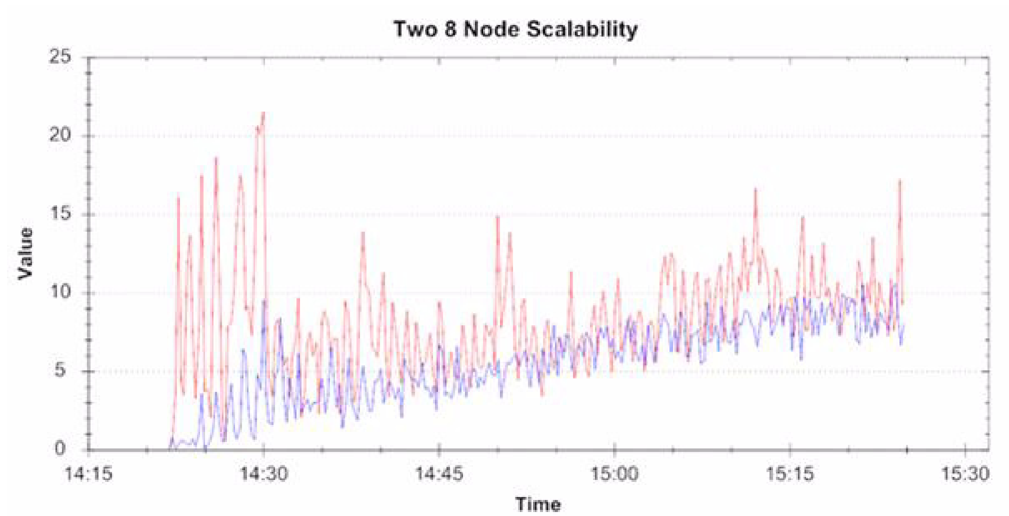

*Figure 72 1760 Desktop Boot Up—2 VMM Servers CPU Utilization*

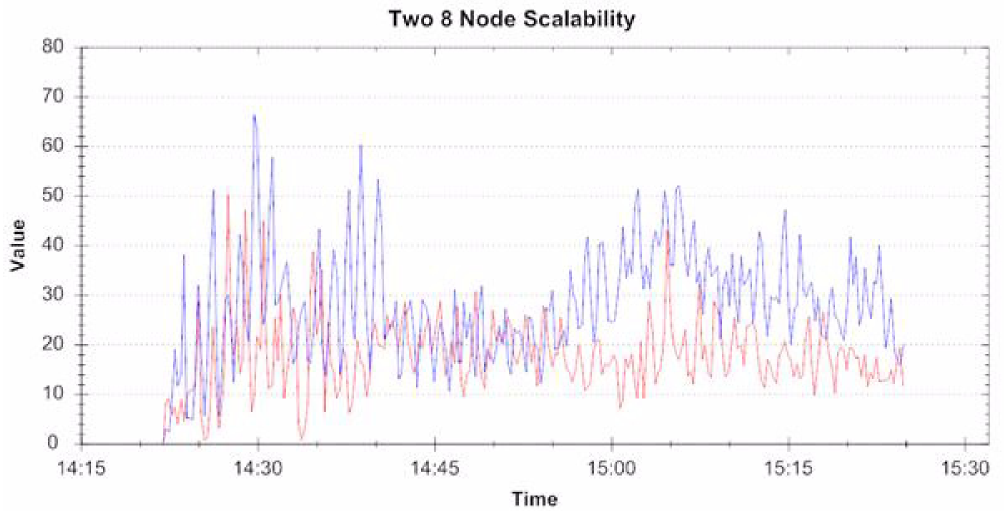


As illustrated below, during the 16 node tests, the average disk transfer per second on the Hyper-V hosts was approximately 274 per second on average with a maximum value of approximately 2140 transfers per second. Please note, the graphs from the NetApp provide more details on IOPS consumed during the tests.

*Figure 73 1760 Desktop Boot Up—16 Microsoft Hyper-V hosts Disk Transfer Per Seconds*

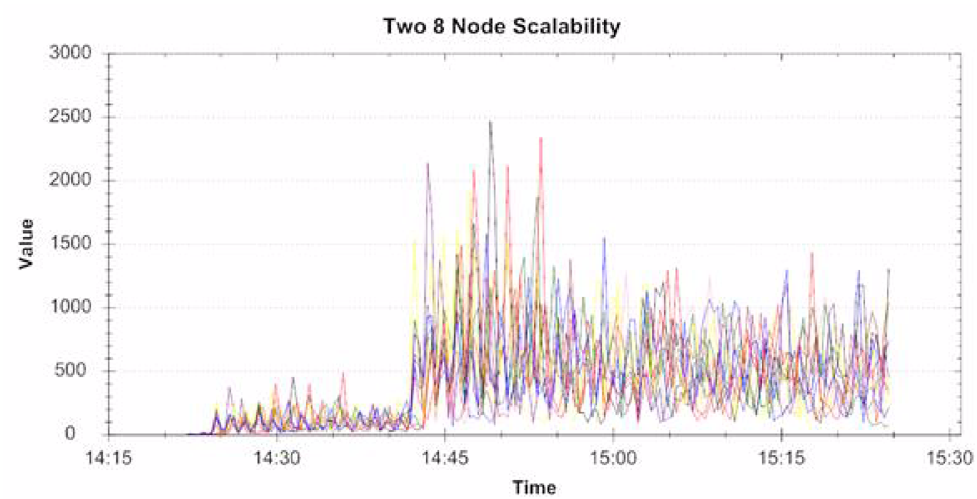


The following graphs represent the utilization during the user workflow steady state execution phase of the 16 node scalability test.

*Figure 74 1760 Desktop Steady State—Microsoft Hyper-V hosts CPU Utilization*

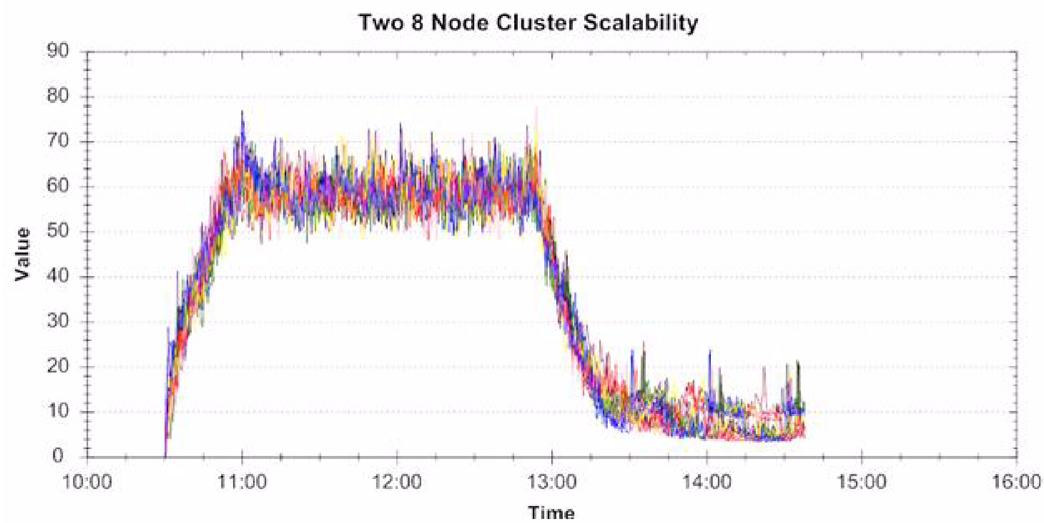


The VSI workflows are executed locally on the virtual desktop and hence the network utilization on the Hyper-V hosts is insignificant during the tests as shown in the figure below.

*Figure 75 1760 Desktop Steady State—Microsoft Hyper-V hosts Network Utilization*

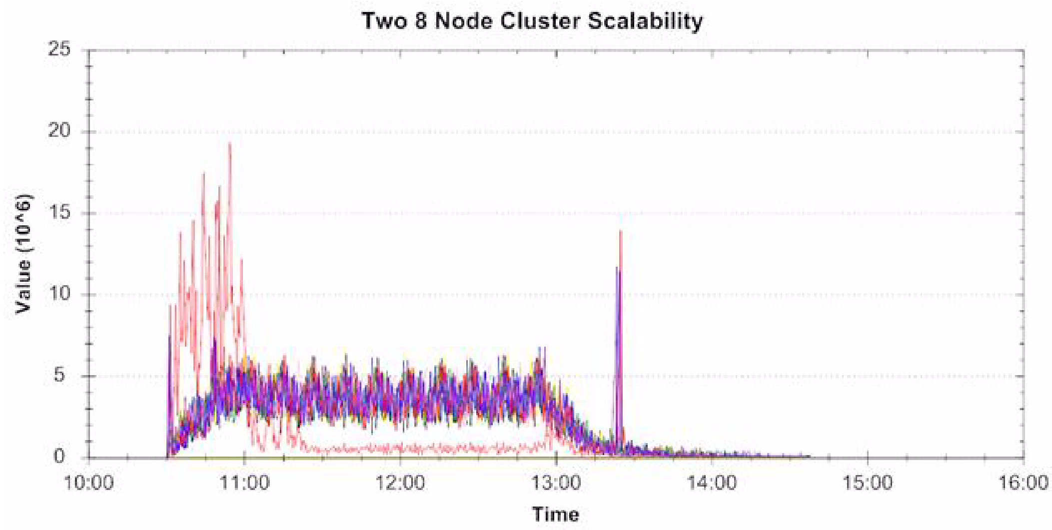


*Figure 76 1760 Desktop Steady State—Microsoft Hyper-V hosts Memory Available (Mbytes)*

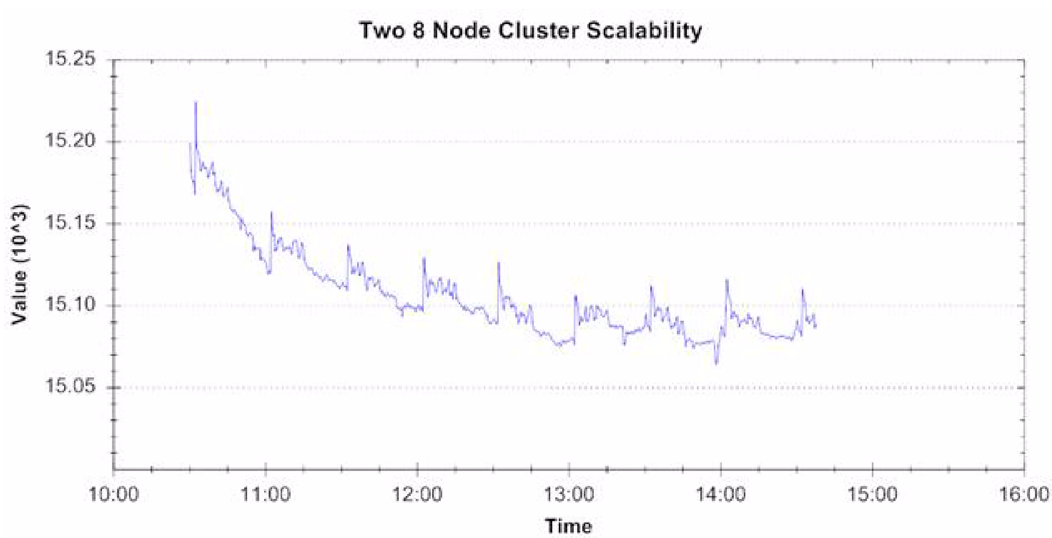


Unlike the virtual desktop boot up phase, the CPU utilization on PVS Servers is significantly low during the workflow execution. This is because the operating system is being streamed to the virtual desktop upon boot up and also due to the write cache drive being attached locally to the virtual desktop.

*Figure 77 1760 Desktop Steady State—3 PVS Servers CPU Utilization*

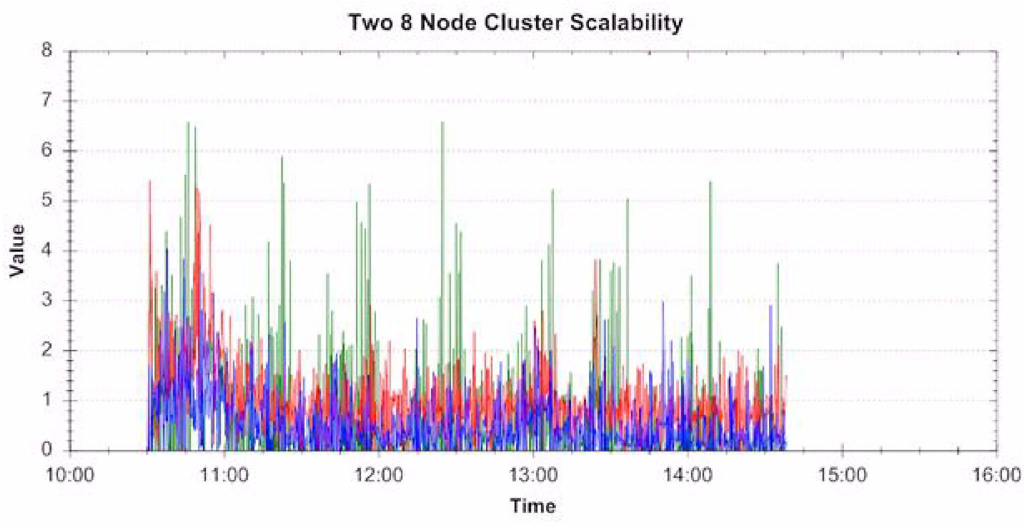


*Figure 78 1760 Desktop Steady State—DDC CPU Utilization*

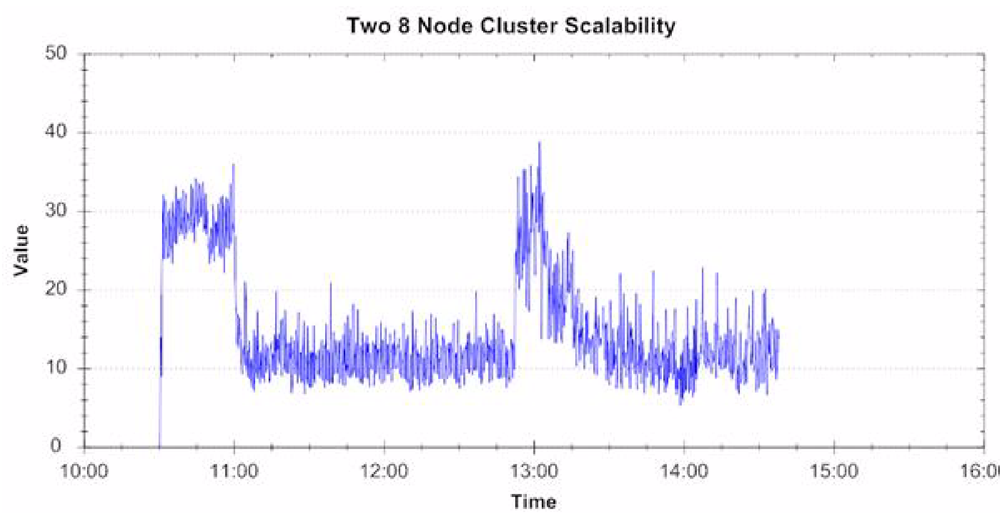

*Figure 79 1760 Desktop Steady State—VMM Server CPU Utilization*

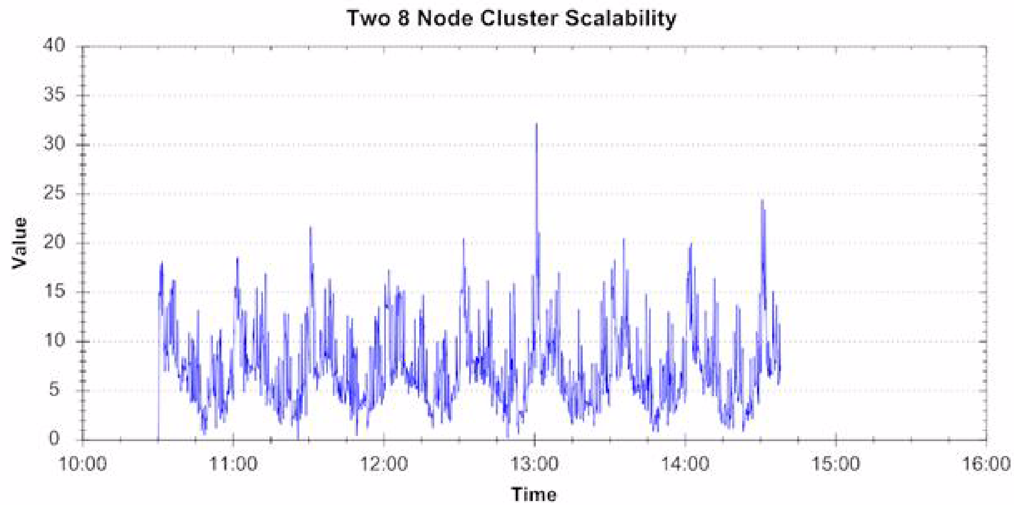


As illustrated below, during the steady state execution of the 16 node tests, the average disk transfer per second on Hyper-V hosts were approximately 391 per second on average and a maximum value of approximately 1100 transfers per second. Please note, the graphs from NetApp provide more details on IOPS consumed during the tests.

*Figure 80 1760 Desktop Steady State—Microsoft Hyper-V Hosts Disk Transfer Per Second*

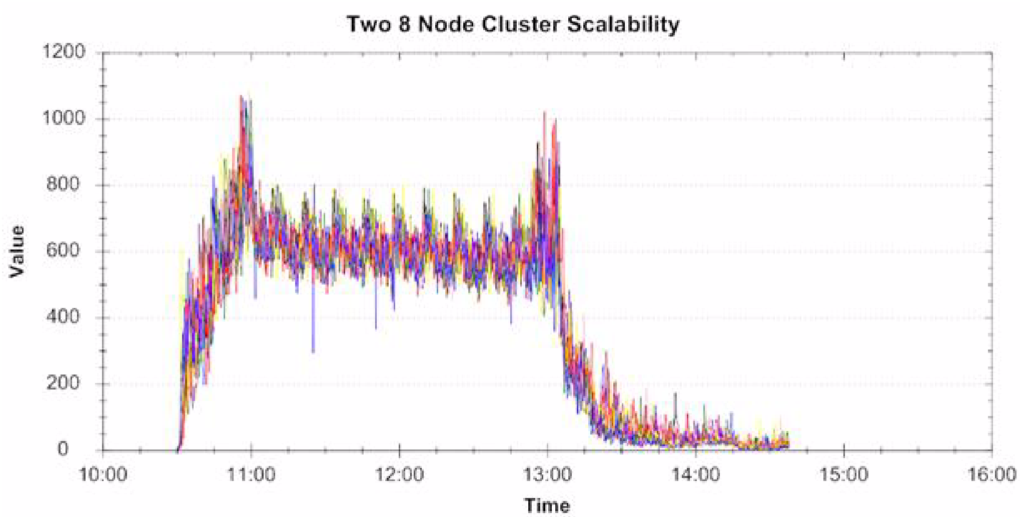


Another key piece of information to note when assessing the overall results of the testing is that the virtual machine per CPU core density was maintained across all test environment configurations. As shown in the table below, the virtual machine density per CPU core was maintained while the number of hosts was increased showing a linear CPU core to VM density ratio.

*Table 15 Virtual Machine Density Per CPU*

| • Windows7 pooled desktops<br>• 1vCPU and 1.5GB RAM.<br>• 3 GB PVS Cache/OS Paging File on NFS Volumes | **Windows 2008 R2 Hyper-V** | | |
|---|---|---|---|
| | No. of Servers Tested | No. of VMs | VMs/Core |
| • Cisco UCS B250 M2s w/ Dual Six Core (3.33GHz) 192GiB RAM | 1 Blade | 110 | **9.16** |
| | 8 Blades | 880 | |
| | 16 Blades | 1760 | |

## Storage Data for Two and Four-Chassis Validation

The following graphs are the throughput IOPs that were observed on the clustered shared volumes. With 880 desktops (110 x 8) hosted on three volumes (around 293 desktop/CSV) we see a peak IOPS of 6250. The ratio of writes to reads is 8:1 and an average of 21 IOPS/desktop during the steady state while the workload was running. Similar observations were made during the four chassis test runs with 1760 desktops.

*Figure 81 Graph of FC OPS on a Volume with Three Clustered Shared Volumes*

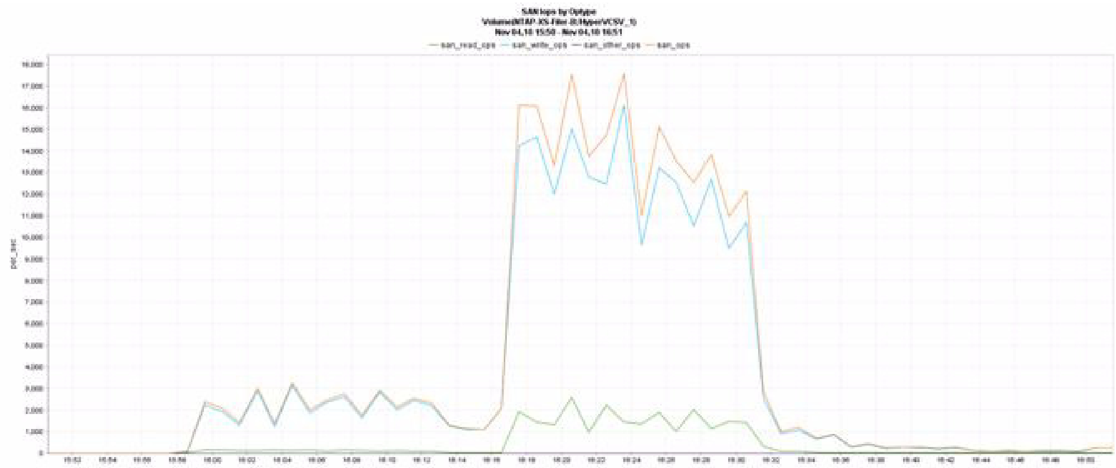

*Figure 82 Graph of FC OPS on Clustered Shared Volumes, Volume1*

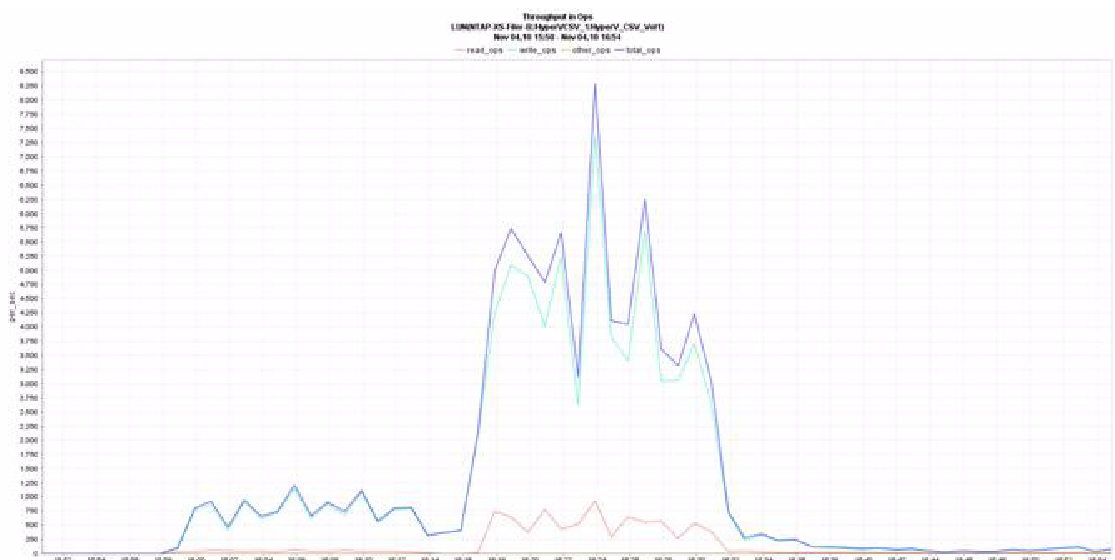

*Figure 83 Graph of FC OPS on Clustered Shared Volumes, Volume2*

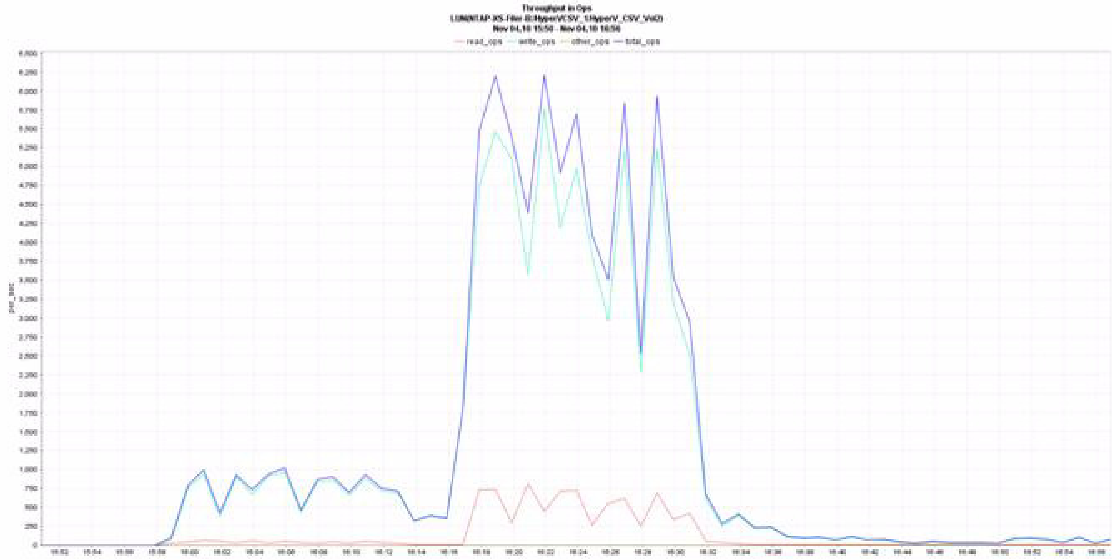

*Figure 84 Graph of FC OPS on Clustered Shared Volumes, Volume3*

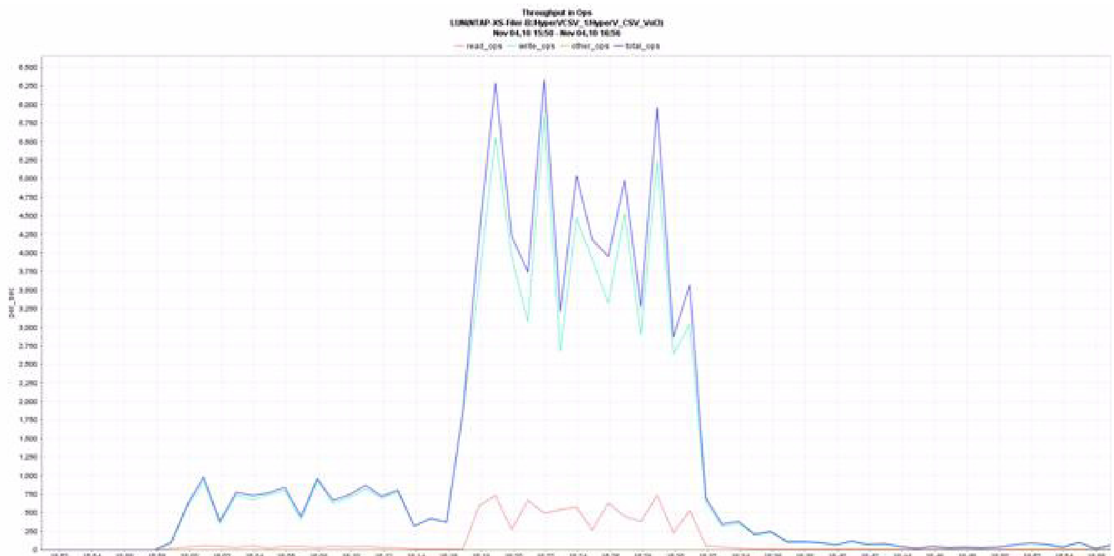

# Scalability Considerations and Guidelines

There are many factors to consider when you begin to scale beyond four chassis or 16 servers, which this reference architecture has successfully tested. This section details how to scale beyond four Cisco UCS chassis.

## Cisco UCS System Configuration

As the results indicate, there is a linear scalability in the Cisco UCS reference architecture implementation.

*Table 16 Hyper-V Number of Chassis, Blade Servers, and Virtual Machines Tested*

| Hyper-V | | | |
|---|---|---|---|
| No. of Chassis | No. of B250-M2 Servers Tested | No. of VMs | VMs/Core |
| 1 | 1 Blade | 110 | **9.16** |
| 2 | 8 Blades | 880 | |
| 4 | 16 Blades | 1760 | |

Cisco Unified Computing System supports up to 20 chassis within a single Cisco UCS domain on a Cisco UCS Fabric interconnect 6120 and up to 40 chassis on a FI 6140, extrapolating the values achieved during the testing, with the following results:

*Table 17 Test Results*

| Hyper-V | | | |
|---|---|---|---|
| **No. of Chassis** | No. of B250-M2 Servers | No. of virtual machines | Virtual machines/Core |
| **8** | 32 Blades | 3520 | **9.16** |
| **12** | 48 Blades | 5280 | |
| **16** | 64 Blades | 7040 | |
| **20** | 80 Blades | 8800 | |

To accommodate the Cisco Nexus 5500 upstream connectivity in the way described in the LAN Configuration section, four Ethernet uplinks are required to be configured on the Cisco UCS Fabric interconnect. Based on the number of uplinks from each chassis, you can calculate how many desktops can be hosted in a single Cisco UCS domain. Assuming two links per chassis, scaling beyond 10 chassis would need a Cisco UCS 6140 fabric interconnect. A 5500 building block can be built out of the RA described in this study with two links per chassis and 12 Cisco UCS chassis comprising of four B250-M2 blades servers each.

The backend storage has to be scaled accordingly, based on the IOP considerations as described in the NetApp scaling section.

Citrix has a modular reference architecture design that details how to scale their components as you scale the number of desktops. Please refer to http://support.citrix.com/article/ctx124087.

# Appendix A—PowerShell Scripts

The following PowerShell Script was used to create the virtual desktops as referenced in Microsoft Windows 7 Image Creation and Provisioning section.

Sample Usage:

PS c:\scripts> .\GenVMs.ps1 VMTargetHost VMBaseName NetworkName LocalVMStoragePath NumberOfDesktopsToCreate StartingAt

Environment Parameters:

PS C:\scripts> .\GenVMs.ps1 HVHOST1 "C1H1-" "Desktop-Network" "C:\clusterstorage\volume" 110 1

Requirements:

Must be run from the VMM Server's PowerShell so the VMM PS libraries are loaded.

The first parameter <VMTargeHost> is the hostname or IP of the Hyper-V host to create the virtual machines on.

The second parameter <VMBaseName> is the virtual machine name.

The third parameter <NetworkName> is the network name which will be attached to the virtual desktop.

The fourth parameter <LocalVMStoragePath> is the VHD storage path. In this case the path of the Cluster Storage Volume mounted on each Hyper-V host.

The fifth parameter <NumberOfDesktopsToCreate> is the number of desktops to be created.

The sixth parameter <StartingAt> is the starting index number for the virtual desktop created. This index will also get appended to the virtual desktop name while creating the desktops.

GenVMs.PS1 contents:

```
# Parse the command-line and verify that the 6 required parameters
   are present, if not display usage info

if ($args -eq $null -or $args.Count -lt 6)

{

    write-output "Usage: GenVMs.ps1 VMTargetHost VMBaseName
   NetworkName LocalVMStoragePath NumberToCreate StartingAt"

    write-output "Example: .\GenVMs.ps1 ""HOST01"" ""HVDesktop""
   ""External"" ""E:\Hyper-V"" 50 1 "

    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-VMMServer -Computername "localhost"



# Place the command-line parameters into named variables for later
   use.
```

```
$VMHost = $args[0]

$VMBaseName = $args[1]

$NetworkName = $args[2]

$VMPath = $args[3]

$VMCount = $args[4]

$StartCount = $args[5]

$EndCount = $StartCount + $VMCount - 1


for ($i=$StartCount; $i -le $EndCount; $i++)

{

    # 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 MAC Address from the pool of available MAC addresses on
   the server. (Alternatively a MAC address could be assigned here.)


    $PooledMACAddress = New-PhysicalAddress -Commit


    # Get a network object for creating the network adapters


    $VNetwork = Get-VirtualNetwork | where {$_.Name -match
   $NetworkName -and $_.VMHost -match $VMHost}


   # Create a Virtual Legacy Network Adapter required for PXE booting
   with Provisioning Services


    New-VirtualNetworkAdapter -JobGroup $JobGroupID
   -PhysicalAddressType Static -PhysicalAddress $PooledMACAddress
   -VirtualNetwork $VNetwork -VLANID 122 -VlanEnabled $True
```

```
# Create a virtual DVD

New-VirtualDVDDrive -JobGroup $JobGroupID -Bus 1 -LUN 0

# Create a new Hardware Profile for a XenDesktop and set the
default values or use the existing profile.

$HWProfile = Get-HardwareProfile | where {$_.Name -eq
"XD4Profile"}

if ($HWProfile -eq $null)
{
   write-output "Hardware profile not found. Creating a default
profile."
    $HWProfile = New-HardwareProfile -Owner
"UCS-HV-XD\Administrator" -Description "Hosted XenDesktop" -Name
"XD4Profile" -CPUCount 1 -MemoryMB 1536 -BootOrder
PXEBoot,IDEHardDrive,CD,Floppy
}

# Create the Virtual Machine and assign the VM Name. This only
works up to 99 virtual machines.
if ($i -lt 10)
{
    $VMName = "{0}0{1}" -f $VMBaseName, $i
}
else
{
    $VMName = "{0}{1}" -f $VMBaseName, $i
}
New-VM -VMMServer $VMMServer -Name $VMName -VMHost $VMHost -Path
  $VMPath -HardwareProfile $HWProfile -JobGroup $JobGroupID
  -RunAsynchronously -RunAsSystem -StartAction NeverAutoTurnOnVM
  -StopAction TurnOffVM
```

```
}
```

The following PowerShell Script was used to copy the write cache drive to the virtual machines as referenced in Microsoft Windows 7 Image Creation and Provisioning section.

Sample Usage:

PS c:\scripts> .\copyvhd.ps1 UNC_VHD_Path VM_Storage_Location VMNameMatch prepend_text append_text

Environment Parameters

PS c:\scripts> .\copyvhd.ps1 \\SCVMM\MSSCVMMLibrary\VHDs\writecachebase.vhd "C:\clusterstorage\volume1" "C1H" "" ""

Requirements:

Must be run from the VMM Server's PowerShell so the VMM PS libraries are loaded.

The first parameter <UNC_VHD_PATH> is the path to the VHD to be copied.

The second parameter <VM_Storage_location> is the destination path to copy the VHD.

The third parameter <VMNameMatch> is the virtual machine name match. Note that in the example above the string 'C1H' will match all desktops stored at "C\clusterstorage\volume1".

The append and prepend text parameters are optional.

COPYVHD.PS1 contents

```
# Check to verify the arguments necessary for the script to run are
   provided.


if ($args -eq $null -or $args.Count -lt 3)

{

    write-output "Usage: copyvhd UNC_VHD_Path VM_Storage_Location
   VMNameMatches prepend_text append_text"

    write-output "Example: .\copyvhd.ps1
   ""\\SCVMM\MSSCVMMLibrary\VHDs\writecachebase.vhd"" ""e:\vms""
   ""XDesktop"" """" """" "

    exit 1

}


# Assign command-line arguments to variables for later use


$VdiskPath = $args[0]

$ClusterStoragePath = $args[1]

$VMNameMatches = $args[2]
```

```
$PrependString = $args[3]

$AppendString = $args[4]


# Grab a handle to the local computer for the VMM Server interface


$VMMServer = Get-VMMServer -Computername "localhost"


# Grab a handle to the VHD specified on the command-line and exit if
   not found


$BaseVdisk = Get-VirtualHardDisk -VMMServer $VMMServer | where {
   $_.Location -eq "$VdiskPath" }


if ($BaseVdisk -eq $null)

{

    write-output "Unable to find vdisk: $VdiskPath"

    exit 1

}


# Get all the VMs that match the name criteria supplied


$VMs = Get-VM | where { $_.Name -match "$VMNameMatches" }

if ($VMs -eq $null)

{

    write-output "No VMs match the pattern: $VMNameMatches"

    exit 1

}

else

{
```

```
    $matchedString = "{0} vms match the pattern: {1}" -f $VMs.Count,
   $VMNameMatches

    write-output $matchedString

}


# The loop below does the following for each VM matched

# 1. Looks for Virtual Disk Drives already attached

# 2. Generates a filename to be used for the new VHD

# 3. Attempts to copy the template VHD from the library and attach

#    it at IDE Bus 0, LUN 0

# 4. Outputs either a "Success" or "Disk Already Attached" message



foreach ($vm in $VMS)

{

    $current_disks = get-VirtualDiskDrive -VM $VM

    if ($current_disks -eq $null -or $current_disks.count -eq 0)

      {

        $filename = "{0}{1}{2}.vhd" -f "$PrependString", $VM.name,
   "$AppendString"

        $cloningMessage = "Attaching {0} to VM {1}" -f $filename,
   $VM.Name

        write-output $cloningMessage

$newvhd = New-VirtualDiskDrive -VM $VM -VirtualHardDisk $BaseVdisk
   -Path "$ClusterStoragePath\$VM" -Filename "$filename" -IDE -Bus
   0 -LUN 0

      }

    else

      {

        $diskattachedmessage = "{0} {1}" -f $VM.Name, "has disk
   already attached"

        write-output $diskattachedmessage
```

```
        }

}
```

The following PowerShell Script was used to mount the write cache drive to the virtual machines as referenced in Microsoft Windows 7 Image Creation and Provisioning section.

Sample Usage:

PS c:\scripts> .\mountvhd.ps1 LocalVMStoragePath VMNameMatch Postpend

Environment Parameters:

PS c:\scripts> .\mountvhd.ps1 "C:\clusterstorage\volume1" "C1H" "_wc"

Requirements:

Must be run from the VMM Server's PowerShell so the VMM PS libraries are loaded.

The first parameter <LocalVMStoragePath> is the path to the locally stored VHD by running the previous 'copyvhd.ps1' script.

The second parameter <VMNameMatch> is the string to match the name of the virtual machine.

The third parameter <Postpend> is an optional postpend string to assist with identifying the VHD purpose.

MOUNTVHD.PS1 contents:

```
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: MountVHD.ps1 LocalVMStoragePath VMNameMatch
   Postpend"
```

```
    write-output "Example: .\MountVHD.ps1 ""E:\Hyper-V""
   ""HVDesktop"" ""_wc"" "

    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

#


$VMHost = Get-VMHost -VMMServer localhost


#

# Loop through VM's on each of the hosts

#


$AllVMs = Get-VM | 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 $AllVMs.Count,
   $VMNameMatches

}

    write-output $matchString

}


foreach ($myVM in $AllVMs)

{

    $LeftToGo = $LeftToGo - 1

    $HyperVGuest = $myVM.Name

    $server = $myVM.hostname

    $vhdToMount = "{0}\{1}\{1}{2}.vhd" -f $VHDPath, $myVM.Name,
   $PostPend

   $Status = "Processing VM:{0} VHD:{1} VMs Left:{2}" -f $myVM.Name,
   $vhdToMount, $LeftToGo

    Write-output $Status


    #

    # try to attach, if that fails... catch error and continue

    #


    try
```

```
 {
     $VMManagementService = Get-WmiObject -computername $server
-class "Msvm_VirtualSystemManagementService" -namespace
"root\virtualization"

     $Vm = Get-WmiObject -computername $server -Namespace
"root\virtualization" -Query "Select * From Msvm_ComputerSystem
Where ElementName='$HyperVGuest'"

     $VMSettingData = Get-WmiObject -computername $server
-Namespace "root\virtualization" -Query "Associators of {$Vm}
Where ResultClass=Msvm_VirtualSystemSettingData
AssocClass=Msvm_SettingsDefineState"

     $VmIdeController = (Get-WmiObject -computername $server
-Namespace "root\virtualization" -Query "Associators of
{$VMSettingData} Where
ResultClass=Msvm_ResourceAllocationSettingData
AssocClass=Msvm_VirtualSystemSettingDataComponent" |
where-object {$_.ResourceSubType -eq "Microsoft Emulated IDE
Controller" -and $_.Address -eq 0})

    $DiskAllocationSetting = Get-WmiObject -computername $server
-Namespace "root\virtualization" -Query "SELECT * FROM
Msvm_AllocationCapabilities WHERE ResourceSubType = 'Microsoft
Synthetic Disk Drive'"

     $DefaultDiskDrive = (Get-WmiObject -computername $server
-Namespace "root\virtualization" -Query "Associators of
{$DiskAllocationSetting} Where
ResultClass=Msvm_ResourceAllocationSettingData
AssocClass=Msvm_SettingsDefineCapabilities" | where-object
{$_.InstanceID -like "*Default"})


     $DefaultDiskDrive.Parent = $VmIdeController.__Path

     $DefaultDiskDrive.Address = 0

     $NewDiskDrive =
($VMManagementService.AddVirtualSystemResources($Vm.__Path,
$DefaultDiskDrive.PSBase.GetText(1)) | ProcessWMIJob
$VMManagementService "AddVirtualSystemResources").NewResources


    $DiskAllocationSetting = Get-WmiObject -computername $server
-Namespace "root\virtualization" -Query "SELECT * FROM
Msvm_AllocationCapabilities WHERE ResourceSubType = 'Microsoft
Virtual Hard Disk'"
```

```
     $DefaultHardDisk = (Get-WmiObject -computername $server
-Namespace "root\virtualization" -Query "Associators of
{$DiskAllocationSetting} Where
ResultClass=Msvm_ResourceAllocationSettingData
AssocClass=Msvm_SettingsDefineCapabilities" | where-object
{$_.InstanceID -like "*Default"})


     $DefaultHardDisk.Parent = $NewDiskDrive

     $DefaultHardDisk.Connection = $vhdToMount


     $VMManagementService.AddVirtualSystemResources($Vm.__Path,
$DefaultHardDisk.PSBase.GetText(1)) | ProcessWMIJob
$VMManagementService "AddVirtualSystemResources"

  }

  catch { }


}
```

The following PowerShell Script was used to generate a list of virtual desktop objects with their respective MAC addresses.This script will generate a list of virtual desktop and its MAC address. This list can then be used to import target devices for the Citrix Provisioning services.

Sample Usage:

PS c:\scripts> .\GenPVS.ps1 SiteName CollectionName Target Device Description ImportFileName VMNameMatches

Environment Parameters:

PS c:\scripts> .\GenPVS.ps1 "Site1" "Collection1" "Cluster1 Desktops" "PVSFile1.csv" "C1H"

Requirements:

Must be run from the VMM Server's PowerShell so the VMM PS libraries are loaded.

The first parameter <SiteName> is Provisioning Services Site name.

The second parameter <CollectionName> is Provisioning Services collection name.

The third parameter <Targe> is target machine description and is optional.

The fourth parameter <Device> is the name of the output text file that will contain the list of target devices. This list can be used to import the target devices on the Provisioning Services.

The fifth parameter is the string that matches the virtual machine name.

GenPVS.PS1 contents:

```
if ($args -eq $null -or $args.Count -lt 5)

{
```

```
    write-output "Usage: GenPVSFile.ps1 SiteName CollectionName
   Description ImportFileName VMNameMatches"

    write-output "Example: .\GenPVSFile.ps1 ""Site"" ""Collection""
   ""XD Desktop"" ""c:\PVSImport.csv"" HVDesktop "

    exit 1

}


$VMNameMatches = $args[4]

$OutFileName = $args[3]


$OutFile = New-Item -type file "$OutFileName"


$VMMServer = Get-VMMServer -Computername "localhost"


$AllVMs = Get-VM | where { $_.Name -match "$VMNameMatches" } | sort
   Name


foreach ($vm in $AllVms)

{

$nicDetails = Get-VirtualNetworkAdapter -VM $vm


$csvString = "{0},{1},{2},{3},{4}" -f $vm.Name,
   $nicDetails.PhysicalAddress, $args[0], $args[1], $args[2]

add-content $OutFile $csvString

}
```

The following PowerShell Script was used to generate a list of virtual desktop objects to import at the Desktop Delivery Controller.

This script will generate a list of virtual desktop and that can be imported at the 'desktop group' on the Desktop Delivery Controller. This is an optional step and the virtual desktops can also be added by browsing the hosting infrastructure and selecting the desktops. The only advantage by running the script is that it includes the virtual desktop fully qualified domain name which can avoid manual addition to each virtual desktop object.

Sample Usage:

PS c:\scripts> .\GenDDC.ps1 FQDNDomainName ImportFileName VMNameMatch

Environment Parameters:

PS c:\scripts> .\GenDDC.ps1 "ucs-hv-xd.com" "DDCFile1.csv" "C1H"

Requirements

Must be run from the VMM Server's PowerShell so the VMM PS libraries are loaded.

The first parameter is the virtual desktop's fully qualified domain name

The second parameter is the name of the output text file that will contain the list of target devices. This list can be used to import the target devices on the Desktop Delivery Controller.

The third parameter is the string that matches the virtual machine name.

GenDDC.PS1 contents:

```
if ($args -eq $null -or $args.Count -lt 3)

{

    write-output "Usage: GenDDCFile.ps1 FQDNDomainName
   ImportFileName VMNameMatches"

    write-output "Example: .\GenDDCFile.ps1 ""xd4.local""
   ""c:\DDCImport.csv"" HVDesktop"

    exit 1

}


$domainname = $args[0]

$ImportFile = $args[1]

$VMNameMatches = $args[2]


$OutFile = New-Item -type file "$ImportFile"


$VMMServer = Get-VMMServer -Computername "localhost"


$allVMs = Get-Vm | where {$_.Name -match $VMNameMatches} |sort Name


# Set Header line of the CSV file


$csvline =
   "[ADComputerAccount],[AssignedUser],[VirtualMachine],[HostId]"

add-content $OutFile $csvline
```

```
# Populate CSV file with information
foreach ($vm in $allVMs)
{
$csvline = "{0}.{1},,{0},{0}" -f $vm.Name, $domainname
add-content $OutFile $csvline
}
```

# Appendix B—Cisco Nexus 5500 Configuration

```
switchname sj2-151-d17-n5548a

system jumbomtu 9000

logging event link-status default

class-map type qos class-platinum

  match cos 5

class-map type queuing class-platinum

match qos-group 2

policy-map type qos system_qos_policy

  class class-platinum

    set qos-group 2

policy-map type queuing system_q_in_policy

  class type queuing class-platinum

    bandwidth percent 50

  class type queuing class-fcoe

    bandwidth percent 20

  class type queuing class-default

    bandwidth percent 30

policy-map type queuing system_q_out_policy

  class type queuing class-platinum

    bandwidth percent 50

  class type queuing class-fcoe

    bandwidth percent 20

  class type queuing class-default

    bandwidth percent 30

class-map type network-qos class-platinum

  match qos-group 2

policy-map type network-qos system_nq_policy

  class type network-qos class-platinum
```

```
    pause no-drop
    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
snmp-server user admin network-admin auth md5 0x6 priv 0x6
   localizedkey

vrf context management
  ip route 0.0.0.0/0 10.29.164.1
vlan 1
vlan 121
  name privateVMDesktop
vlan 122
  name xenDesktop
vlan 164-166
port-channel load-balance ethernet destination-port
vpc domain 2
  role priority 1000
  peer-keepalive destination 10.29.164.3


interface Vlan1

interface port-channel1
```

```
  switchport mode trunk
  vpc peer-link
  spanning-tree port type network
  speed 10000

interface port-channel2
  switchport mode trunk
  vpc 2
  switchport trunk native vlan 164
  switchport trunk allowed vlan 121-122,164-166
  spanning-tree port type edge trunk
  speed 10000

interface port-channel3
  switchport mode trunk
  vpc 3
  switchport trunk native vlan 164
  switchport trunk allowed vlan 121-122,164-166
  spanning-tree port type edge trunk
  speed 10000

interface port-channel4
  switchport mode trunk
  vpc 4
  switchport trunk native vlan 164
  switchport trunk allowed vlan 121-122,164-166
  spanning-tree port type edge
  speed 10000
```

```
interface port-channel5
  switchport mode trunk
  vpc 5
  switchport trunk native vlan 164
  switchport trunk allowed vlan 121-122,164-166
  spanning-tree port type edge
  speed 10000

interface port-channel10
  untagged cos 5
  vpc 10
  switchport access vlan 166
  speed 10000

interface port-channel11
  untagged cos 5
  vpc 11
  switchport access vlan 166
  speed 10000

interface port-channel12
  vpc 12
  switchport access vlan 166
  speed 10000

interface port-channel13
  vpc 13
  switchport access vlan 166
  speed 10000
```

```
interface Ethernet1/1
  switchport mode trunk
  switchport trunk native vlan 164
  switchport trunk allowed vlan 121-122,164-166
  spanning-tree port type edge trunk
  channel-group 4 mode active

interface Ethernet1/2
  switchport mode trunk
  switchport trunk native vlan 164
  switchport trunk allowed vlan 121-122,164-166
  spanning-tree port type edge trunk
  channel-group 4 mode active

interface Ethernet1/3
  switchport mode trunk
  channel-group 1 mode active

interface Ethernet1/4
  switchport mode trunk
  channel-group 1 mode active

interface Ethernet1/5
  switchport mode trunk
  switchport trunk native vlan 164
  switchport trunk allowed vlan 121-122,164-166
  spanning-tree port type edge trunk
  channel-group 5 mode active
```

```
interface Ethernet1/6
  switchport mode trunk
  switchport trunk native vlan 164
  switchport trunk allowed vlan 121-122,164-166
  spanning-tree port type edge trunk
  channel-group 5 mode active

interface Ethernet1/7
  switchport access vlan 166
  spanning-tree port type edge
  channel-group 12

interface Ethernet1/8
  switchport access vlan 166
  spanning-tree port type edge
  channel-group 13

interface Ethernet1/9
  switchport access vlan 166
  spanning-tree port type edge
  channel-group 10

interface Ethernet1/10
  switchport access vlan 166
  spanning-tree port type edge
  channel-group 11

interface Ethernet1/11
```

```
interface Ethernet1/12

interface Ethernet1/13
  switchport mode trunk
  switchport trunk native vlan 164
  switchport trunk allowed vlan 121-122,164-166
  spanning-tree port type edge trunk
  channel-group 2 mode active

interface Ethernet1/14
  switchport mode trunk
  switchport trunk native vlan 164
  switchport trunk allowed vlan 121-122,164-166
  spanning-tree port type edge trunk
  channel-group 2 mode active

interface Ethernet1/15
  switchport mode trunk
  switchport trunk native vlan 164
  switchport trunk allowed vlan 121-122,164-166
  channel-group 3 mode active

interface Ethernet1/16
  switchport mode trunk
  switchport trunk native vlan 164
  switchport trunk allowed vlan 121-122,164-166
  channel-group 3 mode active
```

```
interface Ethernet1/17
  shutdown
  switchport trunk native vlan 164
  switchport trunk allowed vlan 164-166

interface Ethernet1/18
  shutdown
  switchport trunk native vlan 164
  switchport trunk allowed vlan 122,164-166

interface Ethernet1/19

interface Ethernet1/20
  switchport mode trunk
  switchport trunk allowed vlan 121-122,164-166
```

# Appendix C—Cloning Microsoft Windows 2008 R2 Server Core Boot LUN

The following script can be used to create rapid clones of the Windows boot LUN. The requirements are you have installed the Windows 2008 R2 image on one LUN that is part of the volume (in this example BFSVOL) and has been syspreped. The system using the boot LUN is shutdown before taking the snapshot.

```
#1) Check volume status

NTAP-Filer-A> vol status

Volume State  Status Options

vol0 online raid_dp, flex root

BFSVOL online raid_dp, flex     guarantee=none


2) Check LUN status

NTAP-Filer-A> lun show

/vol/BFSVOL/Win8-BootVolume         50.0g (53694627840)   (r/w,
   online, mapped)

:


3) Create Snapshot of sanboot volume which contain LUN I want to
   clone.

NTAP-Filer-A> snap create BFSVOL BFSVOL_SNAP


4) Verify snapshot I just created

NTAP-Filer-A> snap list


5) Clone LUN using following command.

NTAP-Filer-A> lun clone create /vol/BFSVOL/HVHost1 -b
   /vol/BFSVOL/Win8-BootVolume BFSVOL_SNAP


6) Verify clone created as new LUN. In this example it is HVHost1

NTAP-Filer-A> lun show
```

```
7) Verify snapshot status. You will see it as a busy LUN.

NTAP-Filer-A> snap list

Volume BFS_VOL working...

%/used %/total date name

---------- ---------- ----------- --------

0% ( 0%) 0% ( 0%) Dec 19 20:14 BFSVOL_SNAP (busy,LUNs)

10% (10%) 0% ( 0%) Dec 19 20:00 hourly.0


8) Now Split the LUN

NTAP-Filer-A> lun clone split start /vol/BFSVOL/HVHost1


9) Verify status of lun cloning/split.

NTAP-Filer-A> lun clone split status

<should show some percent complete>


10) Verify snapshot. Now you will not see snapshot as busy because
   we just split the LUN.

NTAP-Filer-A> snap list BFSVOL

Volume BFSVOL

working...


  %/used       %/total  date          name

----------  ----------  ------------  --------

  0% ( 0%)     0% ( 0%)  Jan 10 08:01  hourly.0

  0% ( 0%)     0% ( 0%)  Jan 10 00:00  nightly.0

  0% ( 0%)     0% ( 0%)  Jan 09 20:01  hourly.1

  0% ( 0%)     0% ( 0%)  Jan 09 16:01  hourly.2

  0% ( 0%)     0% ( 0%)  Jan 09 12:01  hourly.3

  0% ( 0%)     0% ( 0%)  Jan 09 08:01  hourly.4
```

```
  0% ( 0%)    0% ( 0%)  Jan 09 00:00  nightly.1
  0% ( 0%)    0% ( 0%)  Jan 08 20:01  hourly.5
  1% ( 0%)    0% ( 0%)  Dec 19 20:14  BFSVOL_SNAP


11) Delete the snapshot.
NTAP-Filer-A> snap delete BFSVOL BFSVOL_SNAP


12) Verify LUN
NTAP-Filer-A> lun show
/vol/BFSVOL/Win8-BootVolume 50g (53694627840) (r/w, online, mapped)
/vol/BFSVOL/HVHost1 50g (53694627840) (r/w, online)


13) Unmap the original LUN
NTAP-Filer-A> lun unmap /vol/BFSVOL/Win8-BootVolume HV-SingleServer


14) Offline original LUN
NTAP-Filer-A> lun offline /vol/BFSVOL/Win8-BootVolume


15) Verify LUN status.
NTAP-Filer-A> lun show


16) Make a Comment on Original LUN.
NTAP-Filer-A> lun comment /vol/BFSVOL/Win8-BootVolume "Microsoft
   Hyper-V server Boot Golden LUN"


17) Verfiy LUN status.
NTAP-Filer-A> lun show -v /vol/BFSVOL/Win8-BootVolume
/vol/BFSVOL/Win8-BootVolume    50.0g (53694627840)   (r/w, offline)
Comment: "Microsoft Hyper-V server Boot Golden LUN"
Serial#: W9mfgJZNIAMq
```

```
Share: none

Space Reservation: disabled

Multiprotocol Type: windows


18) Map the new LUN which is cloned

NTAP-Filer-A> lun map /vol/BFSVOL/HVHost1 HV-server-1 0


19) Verify LUN mapping.

NTAP-Filer-A> lun show -m

LUN path                          Mapped to         LUN ID  Protocol

------------------------------------------------------------------
   -----

:

/vol/BFSVOL/HVHost1               HV-server-1            0       FCP

:


NTAP-Filer-A> lun show

/vol/BFSVOL/Win8-BootVolume 50g (53694627840) (r/w, offline,
   unmapped)

/vol/BFSVOL/HVHost1 50g (53694627840) (r/w, online, mapped)
```

# Appendix D—Hyper-V Server Configuration

1. Rename Server

```
netdom renamecomputer %computername% /newname:newcomputername
```

2. Set Static IP Info

First, find the MAC address of vNIC1. Then do an ipconfig /all. Figure out the name of the interface that matches it (like Local Area Connection 3)….then,

```
netsh interface ipv4 set address name="Local Area Connection 3"
   source=static address=x.x.x.x mask=255.255.255.0
   gateway=10.x.y.z
```

3. Set DNS Server

```
netsh interface ipv4 add dnsserver name="Local Area Connection 3"
   address=10.u.v.w
```

4. Set Proxy Info

```
netsh winhttp set proxy proxyserver.com:80
```

5. Set Firewall to Off

```
netsh firewall set opmode mode=disable
```

6. Set Product Key

```
slmgr -ipk <product-license-key>
```

7. Activate Windows

```
slmgr -ato
```

8. Add FailoverCluster Role

```
dism /online /enable-feature /featurename:FailoverCluster-Core
```

9. Join a Domain

```
netdom join %computername% /domain:ucs-hyperv-xd.com
   /userd:administrator /passwordd:yourpassword
```

10. Reboot

```
shutdown -r -t 0
```

# Acknowledgments

This work would not have been possible without the active support and help of Josh Heller (Cisco Systems, UCS TME). The authors would like to thanks Paul Wilson (Citrix) for their support and help in this work.

The authors would like to thank Lisa DeRuyter, Isaac Roybal and Satinder Sethi (Cisco Systems). Brad Garvey, Lee Gates and Nick Rose from NetApp. Samantha Foster from Citrix Systems.

Cisco and Citrix would like to thank Login Consultants for the rights to use the Login VSI benchmarking tool for SBC and VDI environments referenced in this paper. If looking to replicate aspects of the testing represented in this document, please contact Login Consultants to purchase licenses for the Login VSI benchmarking tool.

Key contributors:

Ravindra "Ravi" Venkat (Cisco Systems)

Bhumik Patel (Citrix)

# References

Cisco VIC driver download site

http://www.cisco.com/cisco/software/release.html?mdfid=282724490&flowid=7209&os=Windows&softwareid=282733859&release=1.3(1g)&rellifecycle=&relind=AVAILABLE&reltype=all

NetApp TR-3702: NetApp Storage Best Practices for Microsoft Virtualization http://www.netapp.com/us/library/technical-reports/tr-3702.html

TR-3747: NetApp Best Practices for File System Alignment in Virtual Environments

http://media.netapp.com/documents/tr-3747.pdf

Cisco Nexus 5500 Series Switch CLI Software Configuration Guide

http://www.cisco.com/en/US/docs/switches/datacenter/nexus5500/sw/configuration/guide/cli_rel_4_0_1a/CLIConfigurationGuide.html

Lossless 10 Gigabit Ethernet: The Unifying Infrastructure for SAN and LAN Consolidation. http://www.cisco.com/en/US/prod/collateral/switches/ps9441/ps9670/white_paper_c11-501770.html

Login VSI Benchmarking Tool:

http://www.loginconsultants.com

PVS on XD Best practice:

http://support.citrix.com/servlet/KbServlet/download/19042-102-19576/XenDesktop%20Best%20Practices.pdf

XD Design Handbook:

http://support.citrix.com/article/CTX120760

Citrix eDocs (Citrix Product, Solutions and Technologies Document Library):

http://support.citrix.com/proddocs/index.jsp

Cisco VXI Desktop Virtualization Validated Designs

http://www.cisco.com/en/US/netsol/ns1004/index.html

Cisco VXI End-to-End Validated Designs

http://www.cisco.com/en/US/solutions/ns340/ns414/ns742/ns1100/landing_vxi.html

Cisco Desktop Virtualization Solutions

http://www.cisco.com/en//US/netsol/ns978/index.html

Cloud Computing

http://www.cisco.com/en/US/netsol/ns976/index.html

# 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 www.cisco.com/go/designzone.