



Reference Architecture-Based Design for Implementation of Citrix XenDesktop on Cisco Unified Computing System, VMware vSphere and NetApp Storage

Cisco Validated Design

August 2010



Data Center of the Future





Contents

1.0 Goal	5
1.1 Audience	5
1.2 Objectives	5
2.0 Summary of Main Findings	6
3.0 Infrastructure Components	7
3.1 Cisco Unified Computing System	7
3.2 Cisco Unified Computing System Components	8
3.2.1 Fabric Interconnect	8
3.2.2 Cisco Fabric Extenders Module	9
3.2.3 Cisco UCS Chassis	10
3.2.4 Cisco UCS B200 M1 Blade Server	11
3.2.5 Cisco UCS B250 M1 Blade Server	11
3.2.6 Intel Xeon 5500 Series Processor	11
3.2.7 Intel Xeon 5600 Series Processor	12
3.2.8 Cisco UCS B200 M2 Blade Server	13
3.2.9 Cisco UCS B250 M2 Blade Server	13
3.2.10 Cisco UCS B440 M1 Blade Server	14
3.2.11 Cisco UCS M71KR-Q Converged Network Adapter	14
3.2.12 Cisco UCS Virtual Interface Card (VIC)	15
3.2.13 Cisco VN-Link in Hardware	16
3.2.14 Extended Memory Architecture	17
3.2.15 Cisco UCS C-Series Rack Mount Servers	18
3.3 Citrix XenDesktop	19
3.3.1 FlexCast Technology	19
3.3.2 High-Definition User Experience (HDX) Technology	20
3.3.3 Citrix XenDesktop Architecture Overview	20
3.3.4 XenDesktop Hosted VDI Overview	20
3.3.5 Citrix XenDesktop Hosted Shared Desktops Overview	24
3.3.6 Citrix XenDesktop Hosted Shared Desktops	26
3.3.7 Citrix XenApp Virtual Applications	27
3.3.8 General Citrix XD Advantages and Value Proposition	27
3.4 VMware vSphere4	29
3.4.1 Cisco Nexus 1000v	29
3.5 NetApp Storage Solution and Components	30
3.5.1 Single Scalable Unified Architecture	30
3.5.2 Storage Efficiency	30
3.5.3 Thin Provisioning	31
3.5.4 NetApp Deduplication	31
3.5.5 Deduplication Guidelines	32
3.5.6 NetApp FlexClone	33
3.5.7 Performance	33
3.5.8 Transparent Storage Cache Sharing	33
3.5.9 Flash Cache AND PAM	34
3.5.10 NetApp Write Optimization	34
3.5.11 Flexible Volumes and Aggregates	35
3.5.12 Operational Agility	35
3.5.13 Virtual Storage Console 2.0 (VSC)	35
3.5.14 Provisioning and Cloning Plug-in (V3.1)	35
3.5.15 Backup and Recovery Plug-in	36
3.5.16 SANscreen VM Insight	36
3.5.17 NetApp Operations Manager	37
3.5.18 Data Protection	38
3.5.18.1 RAID-DP	38
3.5.18.2 NetApp Backup and Recovery Plug-in	38



3.5.21 Storage Sizing Best Practices.....	39
3.5.21.1 Gather Essential Solution Requirements.....	39
3.5.21.2 Performance-Based and Capacity-Based Storage Estimation Process.....	40
3.5.21.3 Determine Storage Sizing Building Block	40
3.5.21.4 Decide on Storage Protocol	41
3.5.21.5 Perform Detailed Performance Estimation	41
3.5.21.6 Perform Detailed Capacity Estimation	41
3.5.21.6 Determine Projected Storage Estimation Lifetime	42
3.5.21.7 Getting Recommendations on Storage System Physical and Logical Configuration.....	42
3.5.22 Storage Architecture Best Practices.....	42
3.5.23 Storage System Configuration Best Practices.....	42
3.5.24 Building a Resilient Storage Architecture	43
3.5.25 Top Resiliency Best Practices.....	43
3.5.26 Building a High-Performance Storage Architecture.....	44
3.6 Cisco Networking Infrastructure	46
3.6.1 Cisco Nexus 5010 28-Port Switch.....	46
3.6.2 Cisco Nexus 5000 Series Feature Highlights	46
3.6.2.1 Features and Benefits.....	46
3.6.2.2 10 Gigabit Ethernet and Unified Fabric Features	46
3.6.2.3 Low Latency	46
3.7 Microsoft Windows 7.....	47
3.7.1 Microsoft Windows 7 Image Creation and Provisioning.....	47
3.7.1.1 Create Windows 7 Virtual Machine and Install Standard Software.....	49
3.7.1.2 Tuning Microsoft Windows 7 Image for VDI	49
3.7.1.3 Provisioning Services (PVS) vDisk Creation.....	50
4.0 Architecture and Design of Citrix XenDesktops on Cisco Unified Computing System and NetApp Storage.....	55
4.1 Design fundamentals.....	55
4.1.1 Hosted Shared Design Fundamentals.....	56
4.1.1.1 Citrix XenApp Policies	56
4.1.1.2 Worker Groups.....	56
4.1.1.3 Load Managed Groups.....	56
4.1.1.4 Designing a Citrix XenApp Deployment	56
4.2 Hosted VDI Design Fundamentals	57
4.2.1 Hypervisor Selection.....	57
4.2.2 Provisioning Services	57
4.2.3 Designing a Citrix XenDesktop Deployment.....	58
5.0 Solution Validation	59
5.1 Configuration Topology for Scalability of Citrix XenDesktops on Cisco Unified System and NetApp Storage	59
5.2 Cisco UCS Configuration	61
5.2.1 QOS and COS in Cisco Unified Computing System	68
5.2.2 System Class Configuration.....	68
5.2.3 Cisco UCS System Class Configuration.....	68
5.3 Citrix XenDesktop Configuration	71
5.3.1 XenDesktop Desktop Delivery Controller (DDC).....	73
5.3.2 Farm Configuration.....	73
5.3.3 Provisioning Services Configuration.....	74
5.3.4 Citrix Provisioning Services	74
5.3.5 Citrix Provisioning Server (PVS) for use with Standard Desktops.....	74
5.4 LAN Configuration	77
5.5 SAN Configuration.....	79
5.5.1 Boot from SAN.....	81
5.5.2 Configuring Boot from SAN on Cisco Unified Computing System.....	82
5.5.3 SAN Configuration.....	86
5.5.4 Cisco UCS Manager Configuration	87



5.6 NetApp Storage Configuration	89
5.6.1 Example of a NetApp NFS Volume Configuration	91
5.6.2 NetApp Deduplication in Practice.....	96
5.7 VMware vSphere Configuration	97
5.8 OS Installation	98
5.8.1 Center Configuration for DDC.....	98
5.8.2 Setup NTP Server in VMware vCenter.....	99
5.8.3 VMware vMA or Virtual Management Assistant.....	99
5.8.4 ESX Tuning.....	99
5.8.5 MBR Tool from NetApp.....	100
5.8.6 VMware vSphere Configuration	100
5.8.7 Network configuration.....	101
6.0 Test Setup and Configurations	104
6.1 Cisco UCS Test Configuration for Single Server Scalability Test Setup	104
6.2 Cisco UCS Configuration for Two-Chassis Test	105
6.3 Cisco UCS Configuration for Four-Chassis Test.....	106
6.4 Testing Methodology and Success Criteria	107
6.4.1 Load Generation.....	107
6.4.2 User Workload Simulation – LoginVSI from Login Consultants.....	107
6.4.3 Success Criteria	108
6.4.3.1 Login VSI Corrected Optimal Performance Index (COPi)	108
6.4.3.2 Login VSI Max.....	108
7.0 Test Results	109
7.1 XenDesktop Hosted VDI Test Results.....	109
7.1.1 Results for Single Cisco UCS Blade Server Validation	109
7.2 Results from Two-Chassis Testing.....	113
7.3 Results from Four-Chassis Testing.....	115
7.3.1 Storage Utilization Graphs.....	119
7.4 Citrix XenDesktop with XenApp Hosted Shared Test Results	124
8.0 Scaling and Sizing Guidelines	130
8.1 Scalability Considerations and Guidelines.....	130
8.1.1 Cisco UCS System Configuration	130
8.1.2 vSphere Configuration.....	131
8.2 Sizing Guidelines.....	131
8.2.1 CPU Calculations	131
8.2.2 Memory Calculations.....	131
9.0 Acknowledgements	133
10.0 References	134
APPENDIX A	135



1.0 Goal

This document reports the results of a study evaluating the scalability of the Citrix XenDesktop environment on a Cisco Unified Computing System (UCS) B-series blade servers connected to NetApp Storage array. We also provide the best practice recommendations, sizing guidelines for large-scale customer deployments of XenDesktops on the Cisco Unified Computing System.

1.1 Audience

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

1.2 Objectives

This document is intended to articulate the design considerations and validation efforts required to design, deploy Citrix XenDesktops on Cisco UCS with NetApp storage running in a virtualized environment on top of VMware vSphere.

2.0 Summary of Main Findings

- Scale testing findings: 110 Windows 7 1.5 GB desktops running knowledge worker load.
- With the large memory blade we are seeing the best performance for the price.
- Linear scalability when we went from 1 server to 8 servers to 16 servers, the results were that with 1 server we had 110 desktops running with 16 we got 1760 desktop with the same response time.
- Pure Virtualization: The validated environment consisted of only virtual machines hosted by VMware vSphere. All the virtual desktop and supporting infrastructure components including Active Directory, Citrix Provisioning Server, and the XenDesktop Desktop Delivery Controllers were hosted on VMware vSphere.
- Rapid provisioning with Cisco UCS Manager makes it easy for scaling from 1 chassis to 2 and so on.
- The 10G unified fabric story gets a stronger validation as we see tremendous performance with respect to user response times during the load test.
- The low latency Cisco Virtual Interface (VIC) cards make the configuration more robust in terms of extra vNICs and contribute positively towards the response time equation.
- With proper backend storage scaling we can scale out UCS domain from 4 chassis and beyond without making any changes to the proposed Reference Architecture.
- FlexCast Validated: The Citrix FlexCast models Hosted VDI and Hosted Shared running on UCS blades and NetApp storage were successfully validated.
- Memory Extends Density: With 192GB of RAM, the B250 M2 blades offer an optimal memory configuration for desktop virtualization that allow the hosting infrastructure the ability to fully utilize the CPU capabilities of the servers without restricting the of amount memory allocated to the desktops.
- Storage Simplicity for Desktop Virtualization: NetApp's ability to aggregate multiple storage volumes greatly simplifies the storage needs for desktop virtualization.



3.0 Infrastructure Components

This section describes all the infrastructure components used in the configuration.

3.1 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 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, 5600 and 7500 Series processors. The 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 Cisco UCS Manager software. 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. 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



3.2 Cisco Unified Computing System Components

3.2.1 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 Fibre Channel over Ethernet (FCoE) functions.

The Cisco UCS 6100 Series provides the management and communication backbone for the Cisco UCS B-Series Blade Servers and 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



3.2.2 Cisco Fabric Extenders Module


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 Fibre Channel over Ethernet (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.

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 the 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



3.2.3 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)



3.2.4 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



3.2.5 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




3.2.6 Intel Xeon 5500 Series Processor

With innovative technologies that boost performance, energy efficiency, and virtualization flexibility, two-processor platforms based on the Intel Xeon processor 5500 series 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 processor 5500 series 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 assist from Intel VT improves performance for applications running in virtual machines (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



3.2.7 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® processor 5600 series automatically regulates power consumption and intelligently adjusts server performance according to your application needs, maximizing both energy efficiency and performance. The secret to this compelling combination is Intel's new 32nm Nehalem micro-architecture. 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 processor 5600 series with Intel Micro-architecture Nehalem 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 processor 5600.

This groundbreaking intelligent server technology features:

- Intel's new 32nm Microarchitecture Nehalem 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® processor 5500 series.
- 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



3.2.8 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



3.2.9 Cisco UCS B250 M2 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



3.2.10 Cisco UCS B440 M1 Blade Server

The Cisco UCS B440 M1 blade server is a full-width, 4-socket system with 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 are four optional front-accessible, hot-swappable Small Form-Factor Pluggable (SFFP) drives and an LSI SAS2108 RAID Controller. B440 can accommodate two dual-port mezzanine cards for up to 40 Gbps I/O per blade; Options include a Cisco UCS VIC M81KR Virtual Interface Card or converged network adapter (Emulex or QLogic compatible).

Figure 13. Cisco UCS B440 Blade Server



3.2.11 Cisco UCS M71KR-Q Converged Network Adapter

The Cisco UCS M71KR-Q QLogic Converged Network Adapter (CNA) is a QLogic-based Fibre Channel over Ethernet (FCoE) mezzanine card that provides connectivity for Cisco UCS B-Series Blade Servers in the Cisco Unified Computing System.

Designed specifically for the Cisco UCS blades, the adapter provides a dual-port connection to the midplane of the blade server chassis. The Cisco UCS M71KR-Q uses an Intel 82598 10 Gigabit Ethernet controller for network traffic and a QLogic 4-Gbps Fibre Channel controller for Fibre Channel traffic, all on the same mezzanine card. The Cisco UCS M71KR-Q presents two discrete Fibre Channel host bus adapter (HBA) ports and two Ethernet network ports to the operating system.


Figure 14. Cisco UCS M71KR-Q Network Adapter



The Cisco UCS M71KR-Q provides both 10 Gigabit Ethernet and 4-Gbps Fibre Channel functions using drivers from QLogic, providing:

- Risk mitigation through compatibility with current QLogic adapter-based SAN environments and drivers
- Reduced TCO through consolidation of LAN and SAN traffic over the same mezzanine card and fabric, reducing the overall number of network interface cards (NICs), HBAs, cables, and switches
- Integrated management with Cisco UCS Manager

Figure 15. Cisco UCS M71KR-Q Architecture



3.2.12 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 to 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 16. 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.

3.2.13 Cisco VN-Link in Hardware

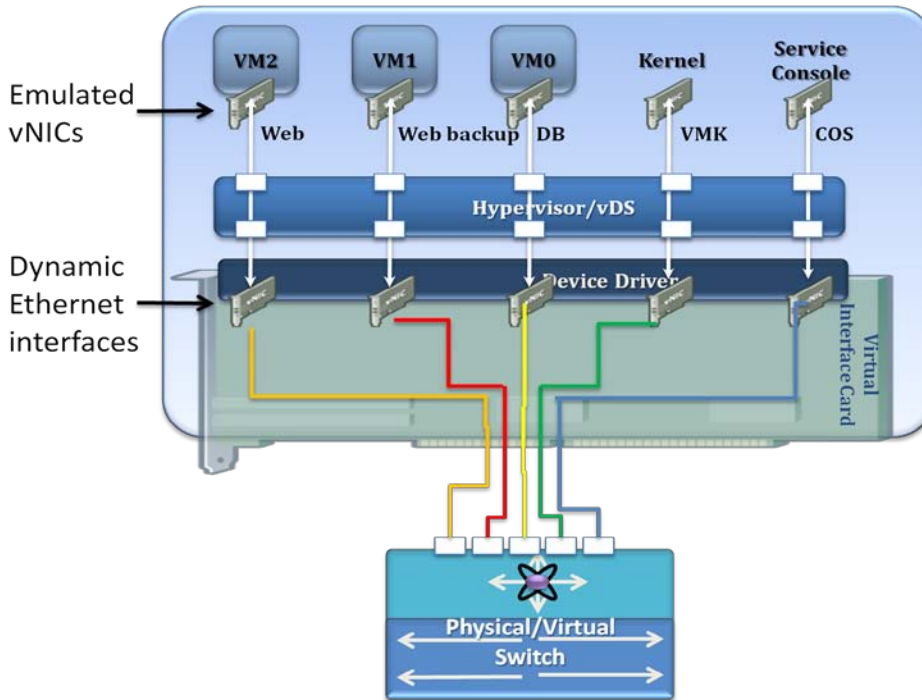
The Virtual Interface Cards are also the first implementations of Cisco's VN-Link in Hardware technology. VN-Link in Hardware eliminates the virtual switch within the hypervisor by providing individual virtual machine virtual ports on the physical network switch. The virtual machine I/O is sent directly to the upstream physical network switch, the Cisco UCS Fabric Interconnect in this case, which takes full responsibility for virtual machine switching and policy enforcement.

In any supported hypervisor environment the Virtual Interface Card presents itself as 3 distinct device types, a FC interface, a standard Ethernet interface and a special Dynamic Ethernet interface. The FC and Ethernet interfaces are consumed by standard vmkernel components and provide standard functionality. The Dynamic interfaces are not visible to vmkernel layers and are preserved as raw PCIe devices.

Using Cisco vDS ESX plugin and VN-Link in Hardware, the Virtual Interface Card provides a solution that is capable of discovering the Dynamic Ethernet interfaces and registering all of them as uplink interfaces for internal consumption of the vDS. As shown in Figure 16, the vDS component on each host discovers the number of uplink interfaces that it has and presents a switch to the virtual machine's running on that host. All traffic from an interface on a virtual machine is sent to the corresponding port of the vDS switch. The traffic is mapped immediately to a unique Dynamic Ethernet interface presented by the Virtual Interface Card. This vDS implementation guarantees the 1:1 relationship with a virtual machine interface and an uplink port. The Dynamic Ethernet interface selected is a precise proxy for the virtual machine's interface.

The Dynamic Ethernet interface presented by the Virtual Interface Card has a corresponding virtual port on the upstream network switch, the Cisco UCS Fabric Interconnect.

Figure 17. Each virtual machine interface has its own virtual port on the physical switch



Cisco UCS Manager running on the Cisco UCS Fabric Interconnect works in conjunction with VMware vCenter software to coordinate the creation and movement of virtual machines. Port profiles are used to describe the virtual machine interface attributes such as VLAN, port security, rate limiting and QoS marking. Port profiles are managed and configured by network administrators using the Cisco UCS Manager. To facilitate integration with the vCenter, the Cisco UCS Manager pushes the catalog of port profiles into vCenter, where they are represented as distinct port groups. This integration allows virtual machine administrators to simply select from a menu of port profiles as they create virtual machines. When a virtual machine is created or moved to a different host, it communicates its port group to the Virtual Interface Card. The Virtual Interface Card asks Cisco UCS Manager for the port profile corresponding to the requested profile, and the virtual port on the Fabric Interconnect switch is configured according to the attributes defined in the port profile.

3.2.14 Extended Memory Architecture

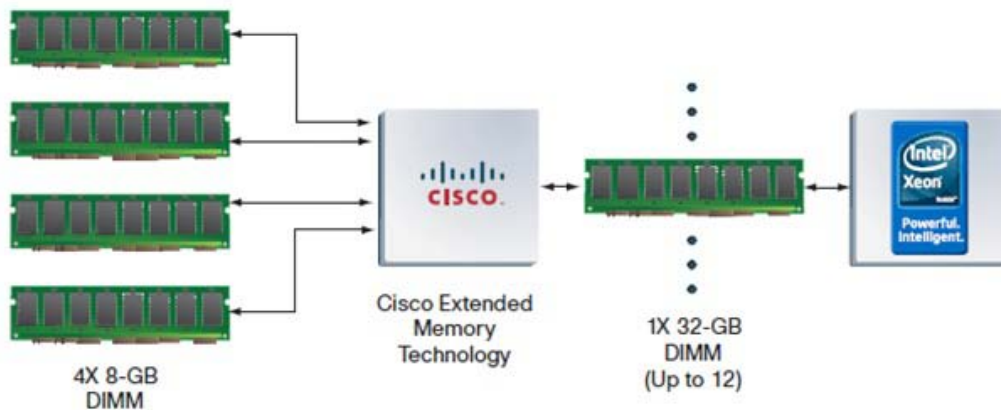
Modern CPUs with built-in memory controllers support a limited number of memory channels and slots per CPU. The need for virtualization software to run multiple OS instances demands large amounts of memory, and that, combined with the fact that CPU performance is outstripping memory performance, can 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. The 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 18). 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 18. Extended Memory Architecture



3.2.15 Cisco UCS C-Series Rack Mount Servers

Cisco UCS C-Series Rack-Mount Servers (Figure 19) extend 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 hetero-geneous data centers or as part of the Cisco Unified Computing.

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 19. Cisco UCS C-Series Rack Mount Servers

**High density
1U Server**



UCS C200 M2




UCS C210 M2

**Economical,
high-capacity
2U Server**

**Memory
Intensive
2U Server**



UCS C250 M2



UCS C460 M1

**Scalability /
Performance
4-Socket/4U**

3.3 Citrix XenDesktop

Citrix XenDesktop is a desktop virtualization solution that delivers Windows desktops as an on-demand service to any user, anywhere. With FlexCast™ delivery technology, XenDesktop can quickly and securely deliver individual applications or complete desktops to the entire enterprise, whether they are task workers, knowledge workers or mobile workers. Users now have the flexibility to access their desktop on any device, anytime, with a high-definition user experience. With XenDesktop, IT can manage single instances of each OS, application and user profile and dynamically assemble them to increase business agility and greatly simplify desktop management. XenDesktop's open architecture enables customers to easily adopt desktop virtualization using any hypervisor, storage or management infrastructure.

3.3.1 FlexCast Technology

[XenDesktop FlexCast](#) is an intelligent delivery technology that recognizes the user, device, and network, and delivers the correct virtual desktop and applications specifically tailored to meet the performance, security, and flexibility requirements of the user scenario. FlexCast for Desktops delivers any type of virtual desktop to any device—and can change this mix at any time. FlexCast for Apps delivers any type of virtual applications to any device. The FlexCast delivery technologies can be broken down into the following categories:

- Hosted Shared Desktops provide a locked down, streamlined and standardized environment with a core set of applications, ideally suited for task workers where personalization is not needed – or allowed.
- Hosted VM Desktops offer a personalized Windows desktop experience, typically needed by office workers, which can be securely delivered over any network to any device.
- Streamed VHD Desktops use the local processing power of rich clients, while providing centralized single-image management of the desktop. These types of desktops are often used in computer labs and training facilities, and when users require local processing for certain applications or peripherals.
- Local VM Desktops 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, apps and user data are automatically synchronized with the datacenter.
- FlexCast for Apps allows any 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 on Citrix.com, but for the purposes of the testing and validation represented in this paper only the Hosted VDI and Hosted Shared models were validated on the Cisco UCS hardware in conjunction with NetApp storage solutions. The Hosted Shared and Hosted VDI models provide a low-cost virtual desktop delivery solution that uses the power of existing PC resources to help customers get started with desktop virtualization.

3.3.2 High-Definition User Experience (HDX) Technology

HDX technology is a set of capabilities that delivers a “high definition” desktop virtualization user experience to end users for any application, device or network. These user experience enhancements balance performance with low bandwidth – anything else becomes impractical to use and scale. HDX technology provides network and performance optimizations to deliver the best user experience over any network, including low bandwidth and high latency WAN connections.


3.3.3 Citrix XenDesktop Architecture Overview

The Citrix XenDesktop Hosted Shared and Hosted VDI FlexCast Delivery Technologies can deliver different types of virtual desktops based on the performance, security and flexibility requirements of each individual user. Although the two desktop delivery models use similar components, the over architecture is distinctly different.

3.3.4 XenDesktop Hosted VDI Overview

Hosted VDI uses a hypervisor to host all the desktops in the data center. Hosted VDI desktops can either be pooled or assigned. Pooled virtual desktops use Citrix Provisioning Services to stream a standard desktop image to each desktop instance upon boot-up therefore the desktop is always 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 detailed below:

Figure 20. Citrix XenDesktop on VMware vSphere 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 4. 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, provides the best flexibility, and offers fewer points of desktop management for both applications and desktops. The Trivial File Transfer Protocol (TFTP) and Pre-boot eXecution Environment (PXE) services are required for the virtual desktop to boot off the network and download the bootstrap file which instructs the virtual desktop to connect to the PVS server for registration and vDisk access instructions.

- **Desktop Delivery Controller:** The XenDesktop controllers are responsible for maintaining the proper level of idle desktops to allow for instantaneous connections, monitoring the state of online and connected virtual desktops and shutting down virtual desktops as needed. The primary XD controller is configured as the farm master server. The farm master is able to focus on its role of managing the farm when an additional XenDesktop Controller acts as a dedicated XML server. The XML server is responsible for user authentication, resource enumeration and desktop launching process. A failure in the XML broker service will result in users being unable to start their desktops. 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 XenDesktops 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 on-demand applications from various platforms. You can deploy and update the online plug-in using Citrix Receiver.
- **Hypervisor:** XenDesktop has an open architecture that supports the use of XenServer, Microsoft Hyper-V or VMware ESX or vSphere. For the purposes of the testing documented in this paper, vSphere 4 was the hypervisor of choice.
- **Citrix XenApp:** Citrix XenApp is an on-demand application delivery solution that enables any Windows application to be virtualized, centralized, 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 on-demand access to streamed and hosted applications.

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

Figure 21. Operational Sequence



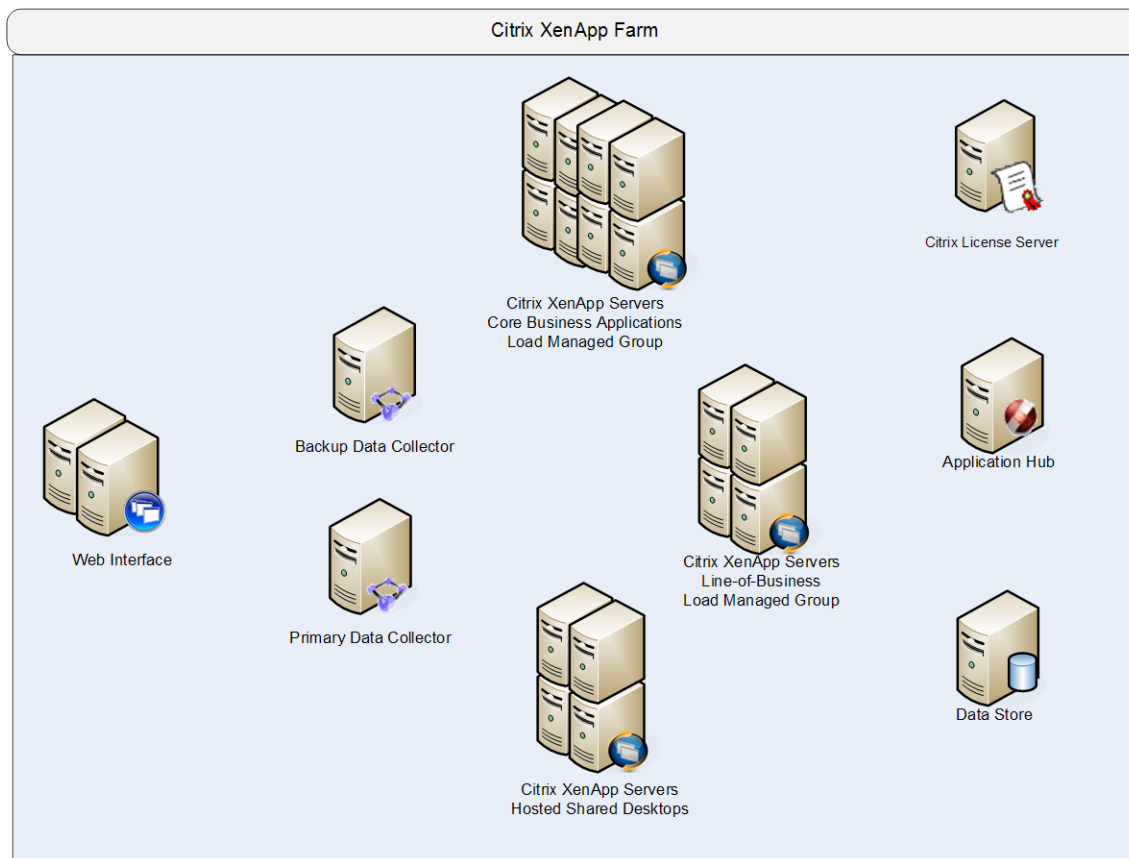
1. The end user launches an internet browser to access Web Interface.
2. Web Interfaces prompts the user for Active Directory credentials and passes the credentials to the Desktop Delivery Controller acting as a dedicated XML server.
3. The XML Service running the dedicated XML server (Desktop Delivery 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 vSphere 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. Once the connection has been approved, the Virtual Desktop Agent uses the transferred credentials to logon against Active Directory and applies profile configurations.

3.3.5 Citrix XenDesktop Hosted Shared Desktops Overview

Hosted Shared desktops use the XenApp feature of XenDesktop to deliver session-based desktops. The Hosted Shared model is built on Microsoft Remote Desktop Services (formerly Terminal Services) platform and end users effectively share one configuration of a Windows Server desktop through independent sessions. The high-level components of the Citrix XenApp feature of XenDesktop architecture for both the Hosted Shared model for desktop delivery and the traditional XenApp model of virtual application delivery are detailed in Figure 22.

Figure 22. Citrix XenApp Architecture



- **Web Interface:** Web Interface provides the user interface for virtual applications and desktops. Web Interface brokers user authentication, enumerates the available desktops and applications. Then upon application or desktop 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.
- **Data Collector:** The data collector is responsible for authenticating users, identifying accessible desktops or applications, and identifying which XenApp server a user should connect. The data collector is the brokering mechanism for requests coming from the end user and Web Interface destined to the XenApp farm. As the size of the XenApp farm increase, the data collector moves from becoming a shared server, responsible for delivering desktops or applications, to a dedicated server. If the primary data collector were

to fail, a backup, with the same hardware and software configuration, should also be available. Similar to Web Interface, providing fault tolerance to the Data Collector servers is recommended.

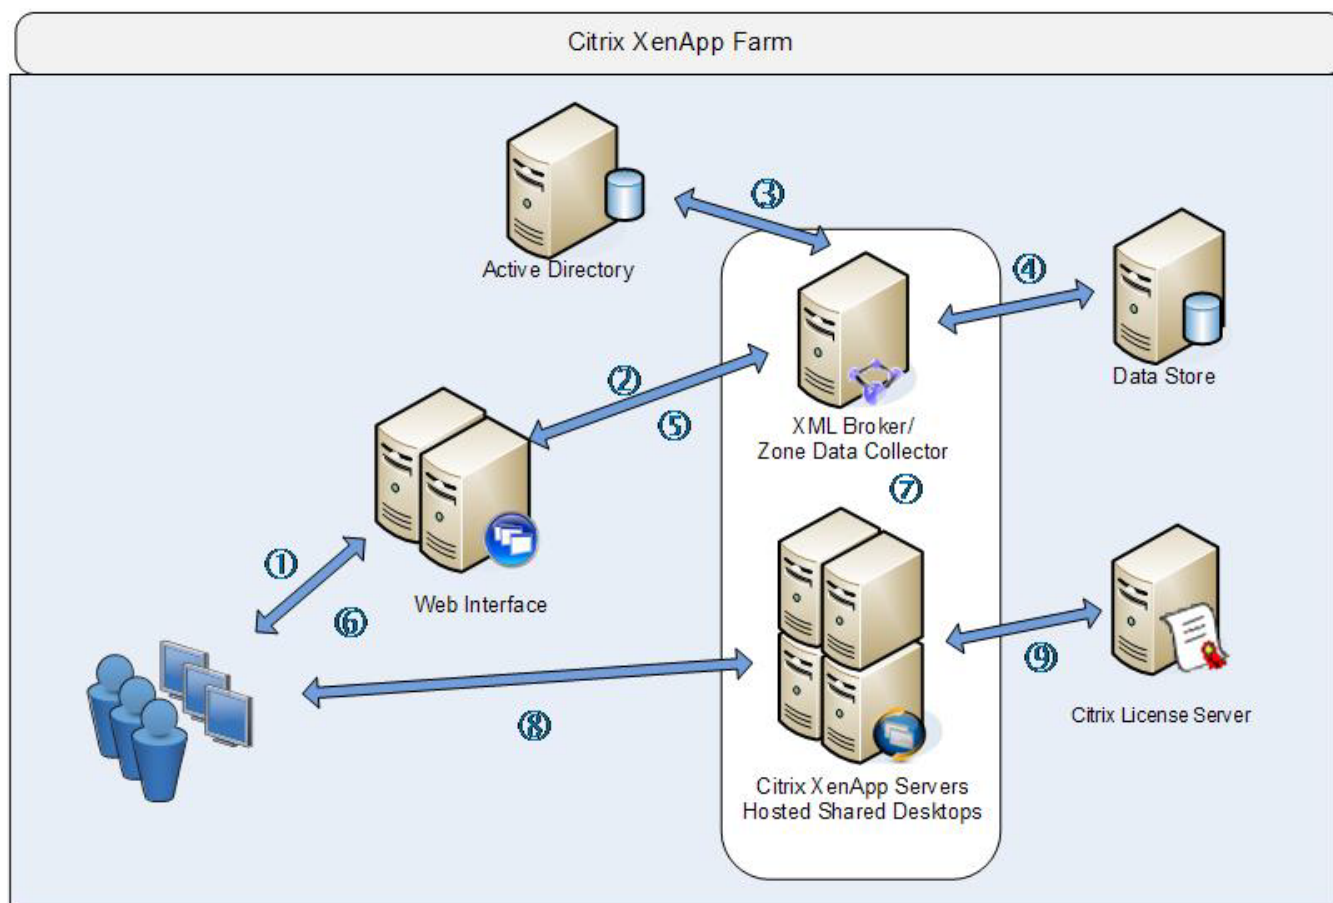
- **Data Collector (Dedicated XML Server):** A Data Collector acting as a dedicated XML server allows the master Data Collector to focus on farm management while directing the Web Interface servers to communicate with the XML servers. The XML broker is responsible for user authentication, resource enumeration and resource launching processes. A failure in the XML broker service will result in users being unable to start their desktop. Due to its criticality it is best to have at least two dedicated XML servers.
- **Load Managed Groups:** Whether delivering applications or desktops, organizations might create load managed groups based on business requirements. Load managed groups are created to focus a set of XenApp servers on a particular set of applications or desktops. This is done for numerous business and technical reasons including update frequency, business unit server ownership, criticality, regional access, and language requirements.
- **When creating a load managed group, each group must provide enough redundancy to be capable of supporting all users in the event of a server failure.** This results in an N+1 scenario where there is at least one additional XenApp server per load managed group. In many situations, organizations implement an N+10% strategy where an additional 10% of XenApp servers per load managed group are allocated in order to allow for multiple server failures or maintenance.
- **License Server:** The license server receives license check-in and check-out requests from the XenApp server in the same fashion as XenDesktop. This service is fairly lightweight and has a grace period for XenApp licenses which allows the system to function normally if the license server becomes unavailable. This grace period offsets the complexity involved with building redundancy into the license server.
- **Data Store:** Each XenApp farm requires a database called a data store. Citrix XenApp uses the data store to centralize configuration information for a farm in one location. The data store maintains all the static information about the XenApp servers, applications and administrators in the server farm.

Citrix XenApp plays a key role in providing an end-to-end virtualization solution. XenApp is fundamentally based on the ability to provide multiple users with access to an independent instance of an application or desktop on a single XenApp server with the popularity previously focused on application virtualization. Before Windows Server 2008 R2, the on-demand XenApp desktop was a server desktop, but now with the release of the Desktop Experience Feature of Windows 2008 R2 a server desktop can be customized with the look and features of a Windows 7 desktop therefore empowering the XenApp virtual desktop delivery model of Hosted Shared desktops.

Given the ability to XenApp to provide both virtual desktops and applications, the following sections outline the order of operations required to access a virtual desktop hosted on XenApp and the ability to launch a virtualized application hosted on XenApp from within a virtual desktop.

3.3.6 Citrix XenDesktop Hosted Shared Desktops

Figure 23. Citrix Hosted Shared Desktop on Citrix XenApp Architecture



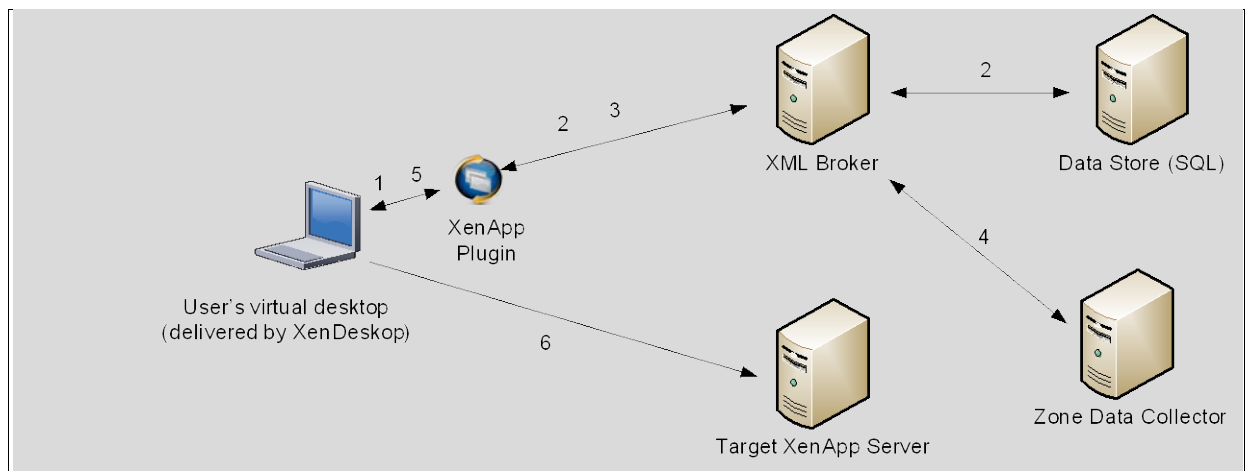
1. The end user launches a browser and enters the URL of the Web Interface site.
2. If using the explicit authentication feature, Web Interfaces prompts the user for Active Directory credentials and passes the credentials to the server acting as the XML Broker. Citrix recommends using the Primary Zone Data Collector as the XML broker server.
3. The XML broker verifies the user's credentials by authenticating the user against Active Directory.
4. After successful verification of the user credentials, the XML broker contacts the Data Store or the locally cached database to determine if the user has permissions to access the published server desktop.
5. The XML broker constructs an XML service response and the icon for that published desktop is populated in the user's Web Interface page.
6. The user clicks on the desktop icon and Web Interface sends a request to the XML broker requesting the address of a XenApp server that can serve the desktop to that user.
7. The XML broker queries the Primary Zone Data Collector (ZDC) to retrieve the address of the appropriate XenApp server. The ZDC returns this address to the XML broker. The XML broker constructs an XML service response and relays the address back to the Web Interfaceserver.
8. The Web Interface server passes the connection information for the assigned XenApp server to the client device in the form of an ICA file. The client device automatically launches the ICA file and connects directly to the desktop of the XenApp server where the Desktop Experience Feature of Windows 2008 R2 is enabled.

9. Before opening the Desktop, the XenApp Server checks out a license from the Citrix License Server on the client's behalf. The client is then connected to the desktop of the XenApp server.

3.3.7 Citrix XenApp Virtual Applications

The following steps outline the order of operations required to access applications virtualized using Citrix XenApp from a Citrix XenDesktop delivered desktop.

Figure 24. Citrix XenApp Application Delivery communication flow



1. The user accesses the XenApp Plug-in within the virtual desktop delivered by XenDesktop. The Plug-in is used in conjunction with its corresponding Web Interface site configured on the Web Interface server.
2. The XenApp Plug-in Web Interface site queries the XML broker to determine a list of applications available to the user. The IMA service on the XML broker queries the local in-memory application cache in order to determine the user's application set. This in-memory application cache is populated from the Local Host Cache. The XML broker constructs an XML service response and relays the application list to the XenApp Plug-In site.
3. The user clicks on the application icon and the XenApp Plug-In site sends a request to the XML broker requesting the address of a XenApp server that can serve that application for the user.
4. The XML broker queries the Zone Data Collector (ZDC) to retrieve the XenApp server address. The ZDC returns this address to the XML broker. The XML broker constructs an XML service response and relays the address to the XenApp Plug-In site.
5. The XenApp Plug-In site on Web Interface server passes the information of the chosen XenApp server to the client device in the form of an ICA file.
6. The client device launches the ICA file connecting directly to the target XenApp server which serves the application.

3.3.8 General Citrix XD Advantages and Value Proposition

Citrix XenDesktop is a desktop virtualization solution that delivers Windows desktops as an on-demand service to any user, anywhere. Whether users are task workers, knowledge workers or mobile workers, XenDesktop can quickly and securely deliver individual applications or complete desktops while providing a high-definition user experience.

The follow statements describe the eight strategic features of Citrix XenDesktop 4:

- Any device, anytime, anywhere. 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 any PC, Mac, thin client or smartphone. This enables complete workplace flexibility, business continuity and user mobility.

- **HDX™ user experience.** XenDesktop 4 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 an experience that rivals a local PC, even when using multimedia, real-time collaboration, USB peripherals, and 3D graphics. XenDesktop 4 offers the best performance while using 90% less bandwidth compared to alternative solutions. New webcam and VoIP support, improved audio, 3D graphics support and branch office WAN optimization help ensure that users can get a high-definition user experience regardless of their location.
- **FlexCast™ delivery technology.** Different types of workers across the enterprise have varying performance and personalization requirements. Some require simplicity and standardization while others need high performance or 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, physical or virtual - each specifically tailored to meet the performance, security and flexibility requirements of each individual user.
- **On-demand apps by XenApp™.** To reduce desktop management cost and complexity, XenDesktop offers the full range of Citrix application virtualization technologies with on-demand apps by XenApp™. This includes integration with Microsoft App-V. With XenApp's virtualization technologies for apps, IT can control data access, manage fewer desktop images, eliminate system conflicts, and reduce application regression testing, making it a requirement for successful desktop virtualization. Adding, updating and removing apps now become simple tasks because users can use a self-service app store, enabling them to access applications instantly from anywhere.
- **Open architecture.** XenDesktop works with your existing hypervisor, storage and Microsoft infrastructures, enabling you to use your current investments – while providing the flexibility to add or change to alternatives in the future. Whether you use XenServer, Microsoft Hyper-V or VMware ESX or vSphere, XenDesktop supports them all and simplifies management of networked storage using StorageLink™ technology. XenDesktop will also closely integrate with Microsoft App-V and System Center for application management.
- **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 data center 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 sets policies that control whether data ever leaves the data center. XenDesktop can dramatically improve endpoint security by eliminating the need for data to reside on the users' devices. Centralized data, encrypted delivery, a hardened SSL VPN appliance and multi-factor authentication further help ensure that only authorized users connect to their desktops, intellectual property is protected, and regulatory compliance requirements are met.
- **Enterprise-class scalability.** XenDesktop includes application, desktop and server virtualization infrastructure that scales to meet the demanding requirements of global enterprises. Pro-active monitoring and reporting enables rapid problem resolution, while intelligent load and capacity management help ensure that problems never arise in the first place. Built-in virtualization management features such as live migration, high availability and bare-metal server provisioning make the infrastructure robust and resilient.



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 increased 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 including mergers and acquisitions, opening a new branch offices, or maintaining business continuity across workplace interruptions.

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, network 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 uncompromised user experience that is driven by Citrix HDX technology and powered by the Cisco Unified Computing System. The solution's Intel® Xeon® processors speed performance and deliver data-center-grade reliability and availability. Local operations are speeded by the integrated 10-Gbps unified fabric, while desktops and applications are streamed to remote users using Cisco wide-area networking technologies. 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.

3.4 VMware vSphere4

Citrix XenDesktop supports all the major hypervisors in the market today. For the purpose of this study we chose VMware vSphere 4 as our hypervisor.

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 multipathing storage layer. More information on vSphere can be obtained at the company website.

3.4.1 Cisco Nexus 1000v

Cisco Nexus 1000V Series Switches are virtual machine access switches that are an intelligent software switch implementation for VMware vSphere environments running the Cisco NX-OS Software operating system. Operating inside the VMware ESX hypervisor, the Cisco Nexus 1000V Series supports Cisco VN-Link server virtualization technology to provide:

- Policy-based virtual machine connectivity
- Mobile virtual machine security and network policy
- Nondisruptive operational model for your server virtualization and networking teams

The Cisco Nexus 1000V Series provides an ideal model in which network administrators define network policy that virtualization or server administrators can use as new similar virtual machines are added to the infrastructure. Policies defined on the Cisco Nexus 1000V Series are exported to VMware vCenter Server to be used and reused by server administrators, as new virtual machines require access to a specific network policy. This concept is implemented on the Cisco Nexus 1000V Series using a feature called port profiles. The Cisco Nexus 1000V Series with the port profile feature eliminates the requirement for the virtualization administrator to create or maintain vSwitch and port group configurations on any of their VMware ESX hosts.

Port profiles create a unique collaborative model, giving server administrators the autonomy to provision new virtual machines without waiting for network reconfigurations to be implemented in the physical network infrastructure. For network administrators, the combination of the Cisco Nexus 1000V Series feature set and the



capability to define a port profile using the same syntax as for existing physical Cisco switches helps ensure that consistent policy is enforced without the burden of managing individual switch ports. The Cisco Nexus 1000V Series solution also provides a consistent network management, diagnostic, and troubleshooting interface to the network operations team, allowing the virtual network infrastructure to be managed like the physical infrastructure.

3.5 NetApp Storage Solution and Components

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.
- **Operational agility:** Enhanced Citrix XenDesktop solution management with tight partner integration.
- **Data protection:** Enhanced protection of both the virtual desktop OS data and the user data, with very low overhead for both cost and operations.

3.5.1 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, FC, 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

3.5.2 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 VM 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 20GB in size and there are supposed to be 1000 virtual machines in the solution, it would require at least 20TB 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% to 90% 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 tradeoffs.

3.5.3 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 preallocating 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 more details on thin provisioning, review [NetApp TR-3563: NetApp Thin Provisioning](#).

Figure 25. Traditional and Thin Provisioning


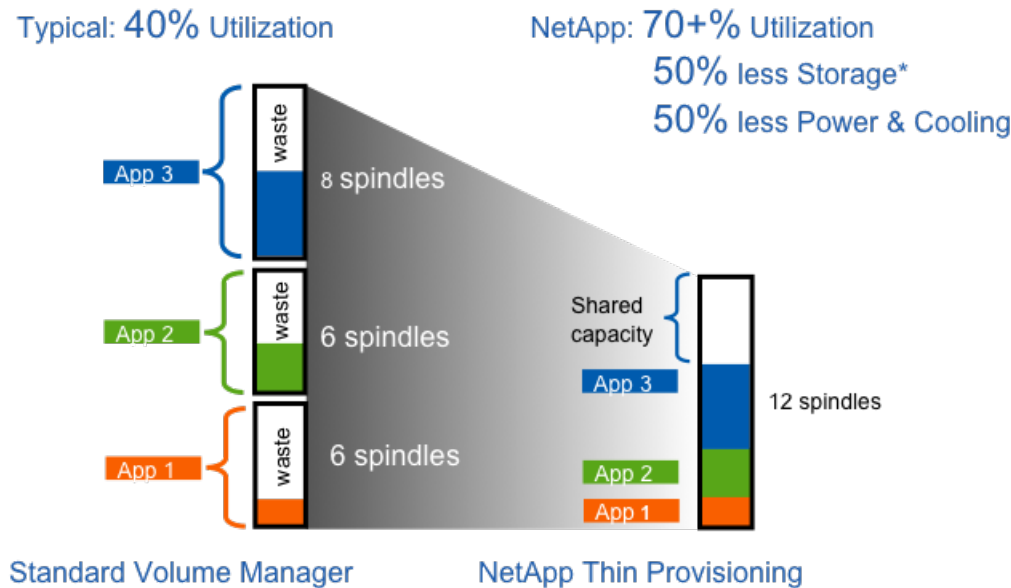


Figure 26. Increased Disk Utilization with NetApp Thin Provisioning




Source: Oliver Wyman Study: "Making Green IT a Reality." November 2007.

*Thin Provisioning, clones, & multiprotocol all contribute to savings.

3.5.4 NetApp Deduplication

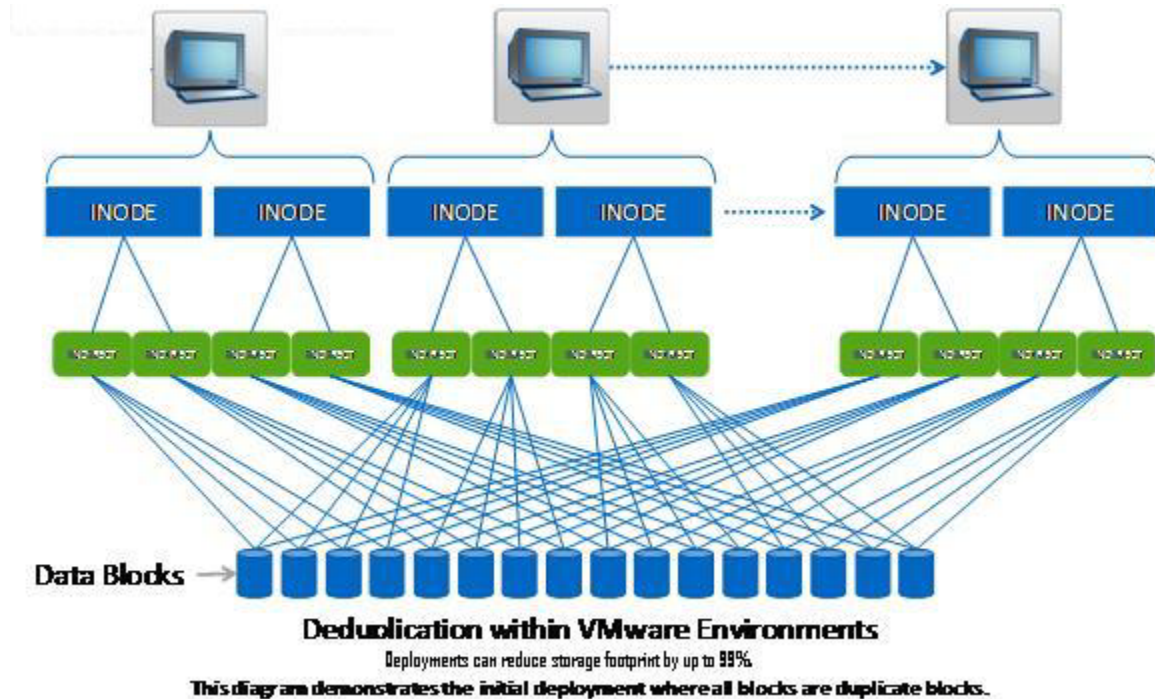
NetApp deduplication saves space on primary storage by removing redundant copies of blocks within a volume hosting hundreds of virtual desktops. This process is transparent to the application and user and can be enabled and disabled on the fly. In a Citrix XenDesktop environment, deduplication provides 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 27. 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 leveraging 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 nonduplicated 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 VMDK 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 28. NetApp Deduplication in VMware Environments



3.5.5 Deduplication Guidelines

- Deduplication is configured and operates on the flexible volumes only.
- Data can be deduplicated up to 255:1 without consuming additional space.

- Each storage platform has different deduplication limits.
- Each volume has dense and nondense size limits.
- Deduplication is configured using the command line.
- Requires Data ONTAP 7.2.5.1, 7.3P1, or later.
- Both a_sis and NearStore® must be licensed for deduplication to work.
- Run deduplication prior to creating Snapshot copies or running SnapMirror or SnapVault updates.

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

3.5.6 NetApp 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 high levels of scalability and significant cost, space, and time savings. Both file-level cloning and volume-level cloning are tightly integrated with the VMware vCenter Server using the NetApp Virtual Storage Console (VSC) Provisioning and Cloning vCenter plug-in. The VSC provides the flexibility to provision and redeploy thousands of virtual machines rapidly with hundreds of virtual machines in each datastore. Note that from a scalability and manageability perspective, this allows for very few datastores to provision and manage as compared to other solutions, which might require one datastore per virtual machine.

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](#).

3.5.7 Performance

Another critical barrier to Citrix XenDesktop adoption is performance issues associated with hosting thousands of virtual machines on shared storage, specifically performance associated with events that produce a large influx of simultaneous I/O such as login storms, boot storms, and antivirus operations. 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

3.5.8 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 leveraging 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% 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 some of the key 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 submillisecond data access. Decreasing the latency results in higher throughput and lower disk utilization, which directly translate into fewer disk 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 DS14mk4 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, each shelf not used saves 3U of rack space. For a real-world deployment, a NetApp solution (with Flash Cache as a key component) would typically replace several such storage shelves; therefore, the savings could be considerably higher than comparing to one disk shelf.

3.5.9 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 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.

3.5.10 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.

3.5.11 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.

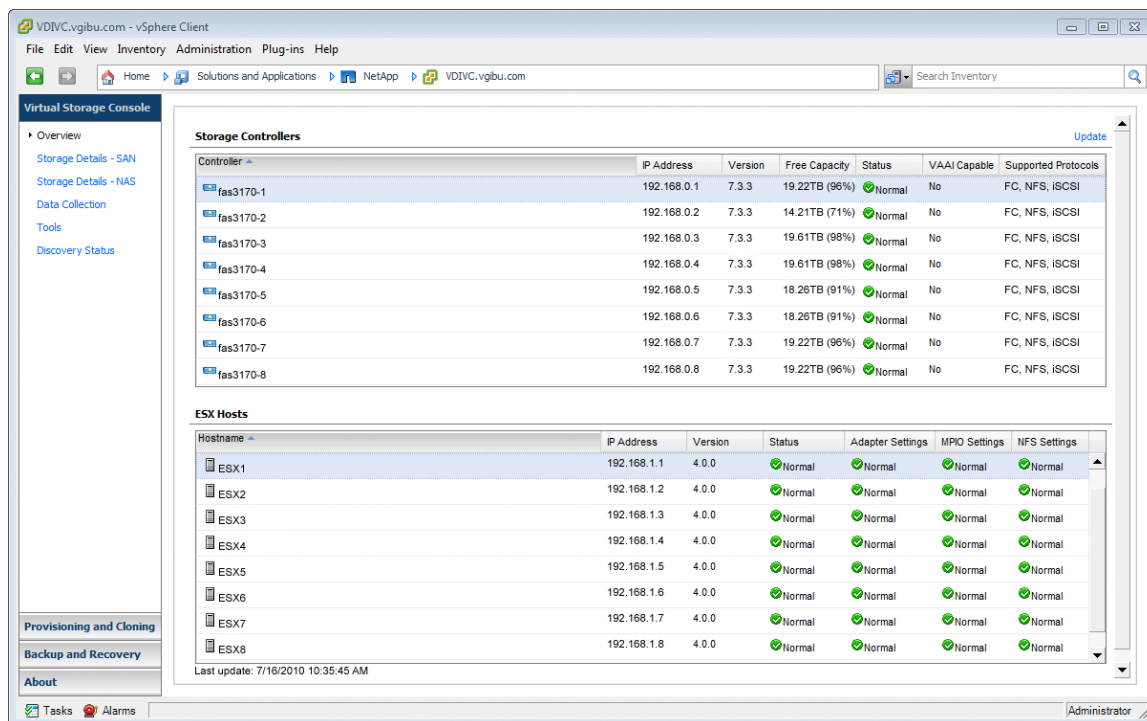
3.5.12 Operational Agility

Implementation and management complexities associated with deploying a Citrix XenDesktop solution are another potential barrier to VDI adoption. The NetApp management solution is operationally agile and provides tight integration with VMware vCenter for rapidly provisioning, managing, configuring, and backing up a Citrix XenDesktop implementation. This section will discuss the components of the NetApp plug-in framework for vSphere, the Virtual Storage Console.

3.5.13 Virtual Storage Console 2.0 (VSC)

The NetApp Virtual Storage Console 2.0 (VSC) for VMware vSphere enables one to manage ESX and ESXi servers connected to NetApp storage systems. VSC is a plug-in to the VMware vCenter that is available to all vSphere clients that connect to the vCenter server. The core function of the VSC allows desktop administrators the ability to provision virtual machines and storage then import virtual machines to XenDesktop.

Figure 29. NetApp VSC Plug-in

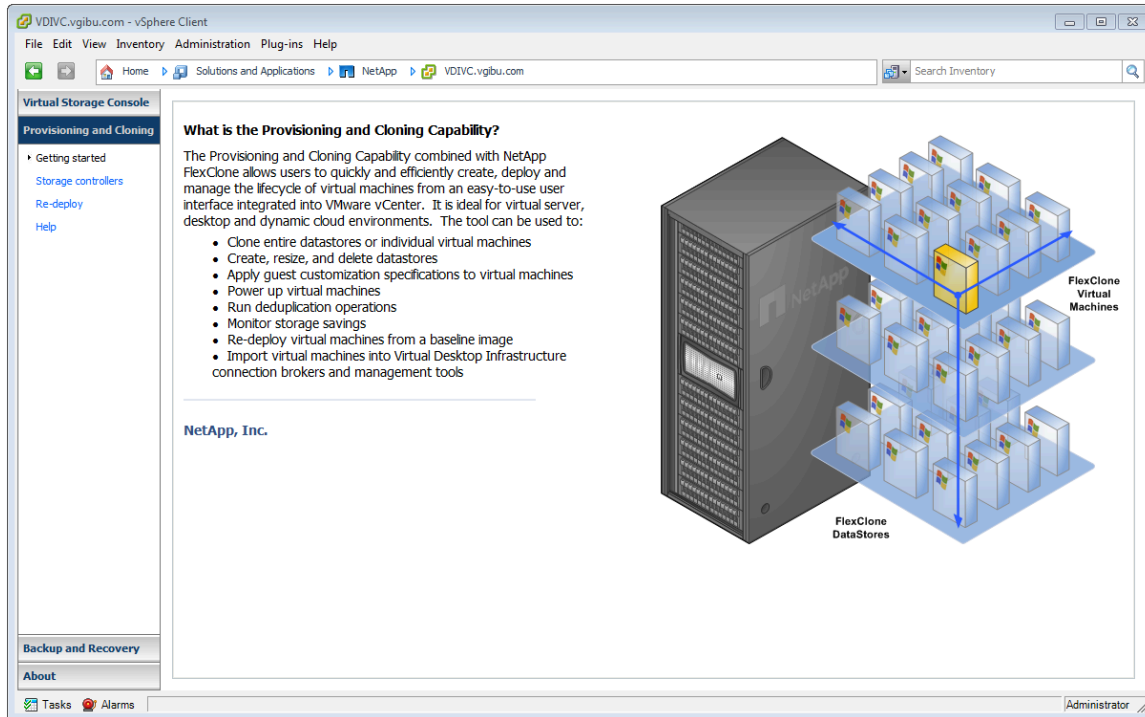


3.5.14 Provisioning and Cloning Plug-in (V3.1)

The NetApp VSC Provisioning and Cloning plug-in v3.1 provides customers with the ability to rapidly provision, manage, import, and redeploy thousands of virtual machines. This plug-in uses file and volume FlexClone technologies, discussed earlier, to create hardware-assisted clones. It is tightly integrated into VMware vCenter as a module of the NetApp VSC 2.0 and can automatically import virtual machines directly into Citrix XenDesktop. Once provisioned, the Provisioning and Cloning plug-in can redeploy virtual machines after patches have been applied to the baseline image. The plug-in also allows desktop administrators the ability to initiate, view, and control deduplication; configure and provision new datastores for NFS, iSCSI, and Fibre Channel and FCoE protocols; resize existing NFS datastores; and destroy datastores. With the provisioning and cloning capabilities

of the VSC 2.0, customers are able to achieve the desired storage efficiency, rapid provisioning, and patch management associated with both dedicated and floating user assignment and persistent and nonpersistent desktops and still be able to achieve performance acceleration with TSCS.

Figure 30. NetApp Provisioning and Cloning Plug-in in VMware vCenter.



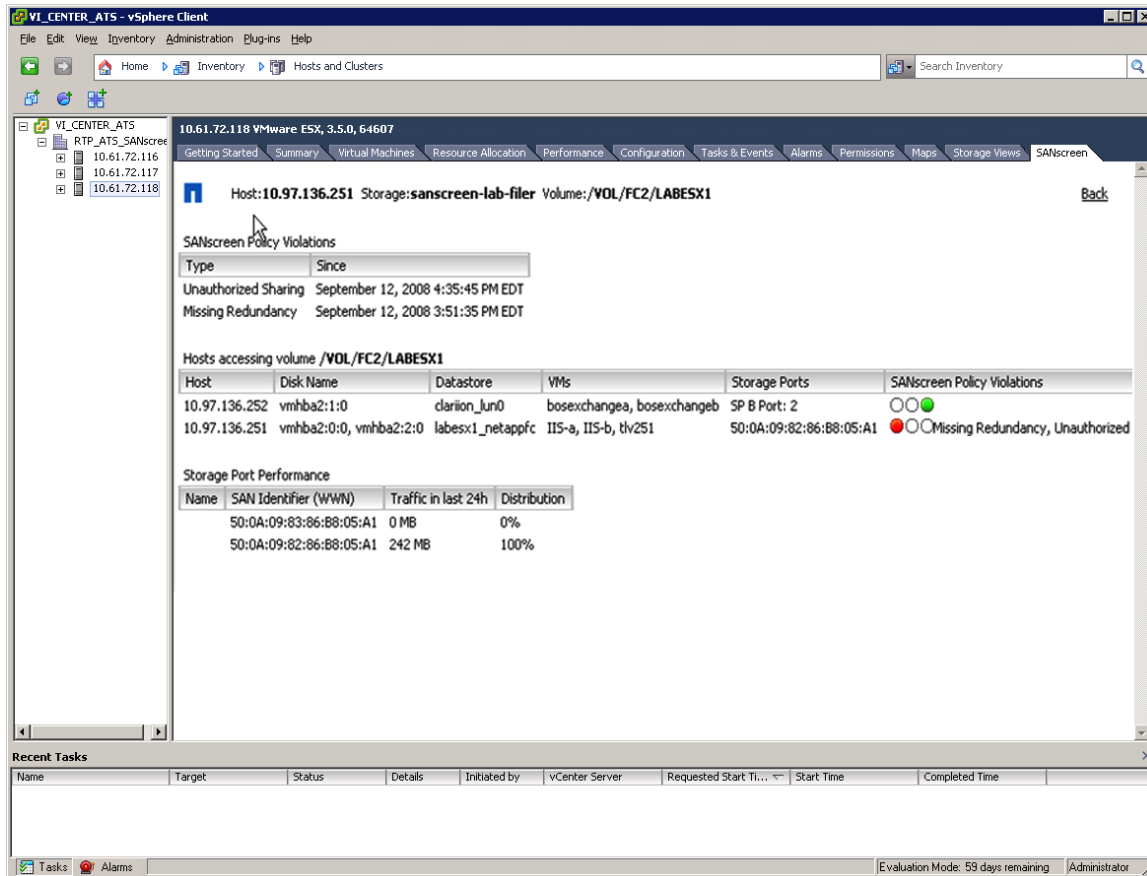
3.5.15 Backup and Recovery Plug-in

The NetApp VSC Backup and Recovery plug-in (formerly SMVI) is a unique, scalable data protection plug-in for VMware vSphere and the Citrix XenDesktop solution. The Backup and Recovery plug-in is now distributed as a part of the Virtual Storage Console 2.0. It integrates VMware snapshot functionality with NetApp Snapshot™ to protect Citrix XenDesktop environments. More information is available in the data protection section of this document.

3.5.16 SANscreen VM Insight

Also available as a VMware vCenter plug-in, NetApp SANscreen® VM Insight also provides cross-domain visibility from the virtual machine to the shared storage, allowing both storage and server administration teams to more easily manage their Citrix XenDesktop storage and server infrastructure. VM Insight provides service-level information for virtual servers, physical servers, and storage devices, as well as VM volume, allocated capacity, and datastore information from VMware vCenter. Based on this visibility, VM Insight shows the actual service paths and server performance information to allow end-to-end monitoring. The enterprise-class data warehouse provides IT the ability to access, query, and analyze virtual machine data; and, when deployed with SANscreen Capacity Manager, it enables capacity planning and chargeback for virtual machine environments. For further details on SANscreen VM Insight, visit www.netapp.com/us/products/management-software/sanscreen/sanscreen-vm-insight.html.

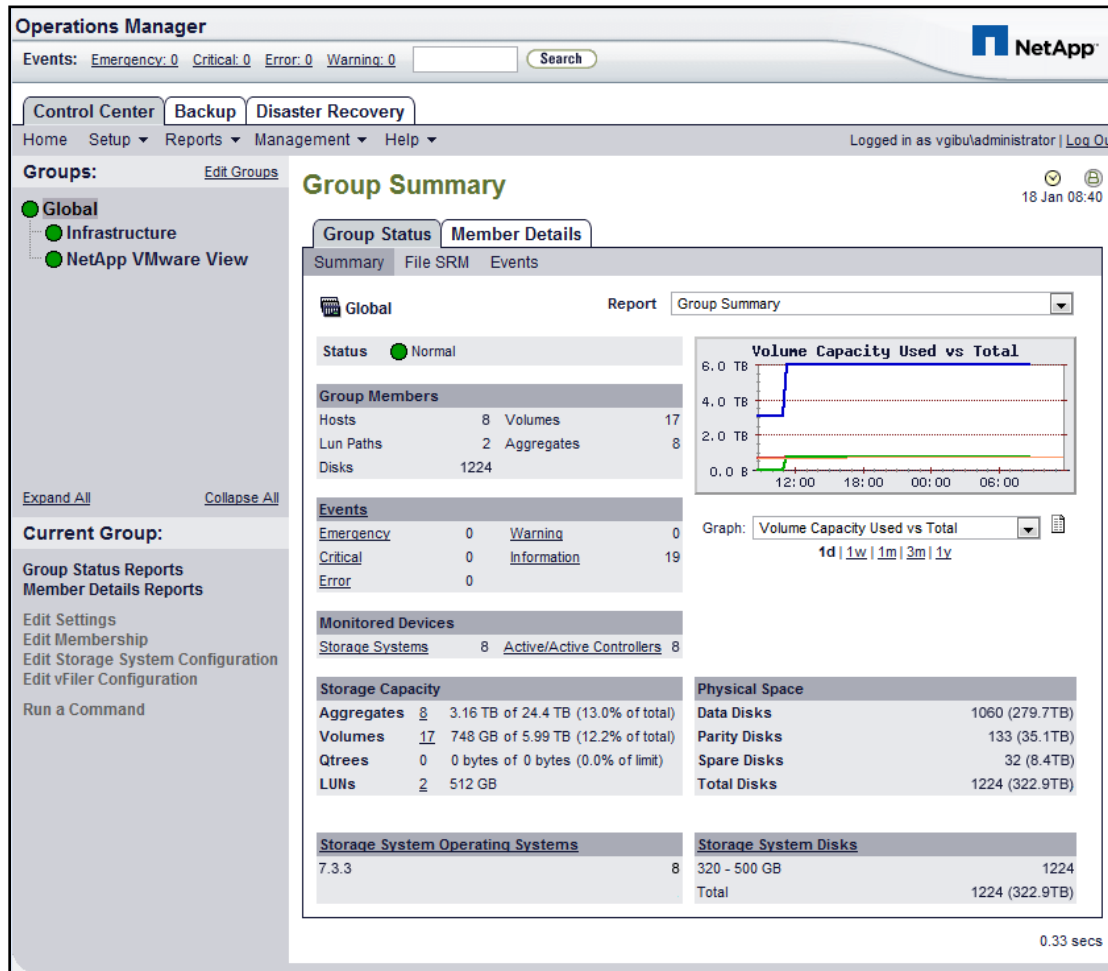
Figure 31. NetApp SANscreen VM Insight Plug-in in VMware vCenter



3.5.17 NetApp Operations Manager

NetApp Operations Manager provides a comprehensive monitoring and management solution for the Citrix XenDesktop storage infrastructure. It provides comprehensive reports of utilization and trends for capacity planning and space usage. It also monitors system performance, storage capacity, and health to resolve potential problems. For further details on Operations Manager, visit www.netapp.com/us/products/management-software/operations-manager.html.

Figure 32. NetApp Operations Manager



3.5.18 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 VSC 2.0's Backup and Recovery plug-in, NetApp SnapMirror®, and NetApp Snapshot copies are critical components of the NetApp solution that help address storage availability.

3.5.18.1 RAID-DP

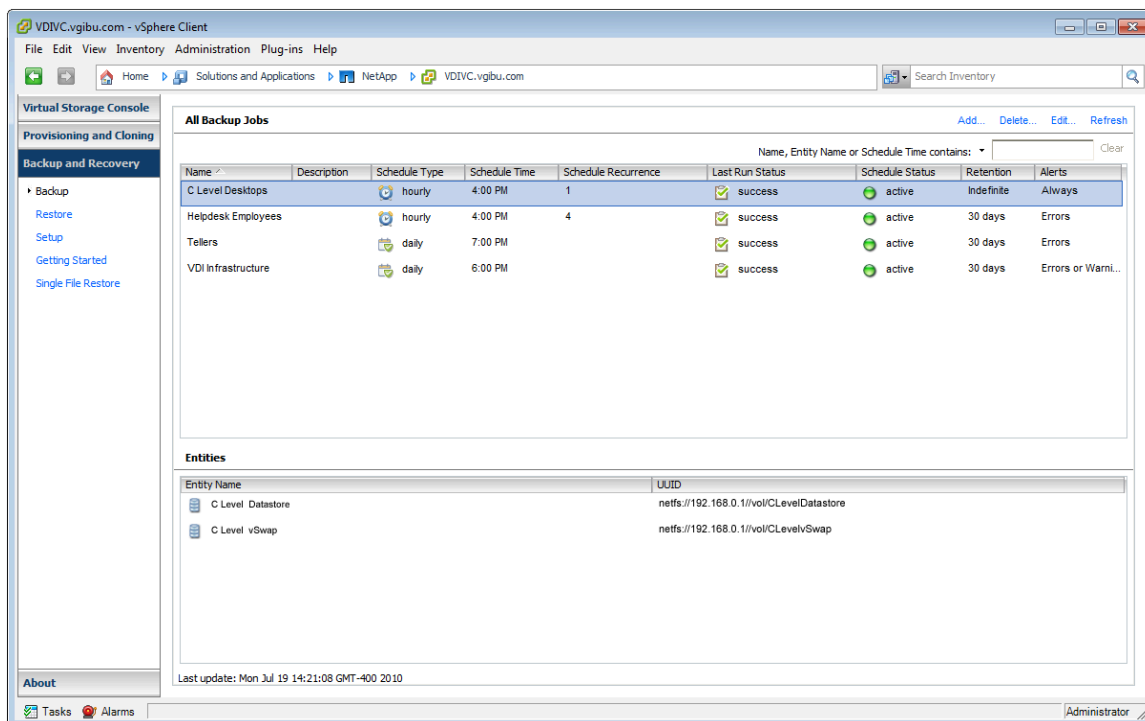
With any Citrix XenDesktop deployment, data protection is critical, because any RAID failure could result in hundreds to thousands of end users being disconnected from their desktops, resulting in lost productivity. RAID-DP provides performance that is comparable to that of RAID 10, yet requires fewer 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 on RAID-DP, see [NetApp TR-3298: RAID-DP: NetApp Implementation of RAID Double Parity for Data Protection](#).

3.5.18.2 NetApp Backup and Recovery Plug-in

The NetApp Backup and Recovery plug-in for VSC 2.0 is a unique, scalable, integrated data protection solution for protecting Citrix XenDesktop environments. The Backup and Recovery plug-in allows customers the ability to use VMware snapshot functionality with NetApp array-based block-level Snapshot copies to provide consistent

backups for the virtual desktops. The Backup and Recovery plug-in is integrated with NetApp SnapMirror replication technology, which preserves the deduplicated storage savings from the source to the destination storage array. Deduplication is then not required to be rerun on the destination storage array. Additionally, when a Citrix XenDesktop environment is replicated with SnapMirror, the replicated data can be quickly brought online to provide production access in the event of a site or data center outage. Also, SnapMirror is fully integrated with VMware Site Recovery Manager (SRM) and NetApp FlexClone technology to instantly create zero-cost writable copies of the replicated virtual desktops at the remote site that can be used for DR testing or for test and development work. For more information on SnapMirror, refer to [NetApp TR-3446: SnapMirror Best Practices Guide](#). For more information on VMware SRM integration, refer to [NetApp TR-3671: VMware Site Recovery Manager](#) in a NetApp Environment. For more detailed information on the Backup and Recovery plug-in, refer to [NetApp TR-3737: SMVI Best Practices](#).

Figure 33. NetApp Backup and Recovery Plug-in



3.5.21 Storage Sizing Best Practices

Storage estimation for deploying Citrix XenDesktop solutions on NetApp includes the following steps:

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

3.5.21.1 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 key 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 user reboot the desktop and what applications the user use. We recommend using NFS for write back cache. NFS is default as thin provisioning and space efficient.
 - NetApp thin provisioning, deduplication, and VSC 2.0 Backup and Recovery data protection solution components can be used to achieve the desired storage efficiency and data protection for the “user data disk.”

3.5.21.2 Performance-Based and Capacity-Based Storage Estimation Process

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:

1. Determine storage sizing building block
2. Perform detailed performance estimation
3. Perform detailed capacity estimation
4. Obtain recommendations on the storage system physical and logical configuration

3.5.21.3 Determine Storage Sizing Building Block

The purpose of this step involves determining the logical storage building block or POD size. This means deciding on the following parameters:

- Storage building block scope: NetApp recommends the Citrix XenDesktop storage sizing building block to be based on the number of datastores required per ESX cluster because it provides benefits of planning and scaling the storage linearly with the number of ESX clusters required for the solution.
- Determine the usable storage required per storage building block (per ESX cluster). For VMFS datastores, there can be multiple LUNs per flexible volume, where each LUN is a datastore. For NFS datastores, each volume can represent a datastore with more virtual machines as compared to VMFS datastores.
- Flexible volume layout across storage controllers: All the flexible volumes belonging to an ESX cluster should be evenly split across the two controllers of the hosting NetApp storage cluster. This is recommended for better performance, as the Citrix XenDesktop deployment scales out from one HA cluster to multiple, hundreds to thousands to tens of thousands of virtual desktops.

3.5.21.4 Decide on Storage Protocol

The two shared storage options available for Citrix XenDesktop are:

- VMFS-based datastores over FCoE, FC, or iSCSI
- NFS-based virtual machines

NetApp is a true unified multiprotocol storage system that has the capability to serve storage for both shared storage options from a single storage cluster without the use of additional SAN or NAS gateway devices.

Both of these are viable and scalable options for Citrix XenDesktop. Consider reading [NetApp TR-3808: VMware vSphere and ESX 3.5 Multiprotocol Performance Comparison Using FC, iSCSI, and NFS](#) for results from a technical performance study conducted by NetApp on different storage protocols. Also, perform a cost benefit analysis for your environment and decide on the storage protocol to be used. The key NetApp value proposition for Citrix XenDesktop holds true across all the protocols.

3.5.21.5 Perform Detailed Performance Estimation

This step involves estimating the total number of disk IOPS and Flash Cache modules required for the Citrix XenDesktop solution based on the requirements. Write I/O optimization with coalescing multiple write operations as single IOPS and transparent storage cache sharing capabilities available in NetApp solution help significantly reduce the amount of data disks required. The calculations are performed based on the IOPS requirement per virtual machine, version of Data ONTAP, transparent storage cache sharing, and customer environment workload characteristics. For getting detailed estimation on savings achieved for your environment, contact your NetApp account team. The output of this step includes:

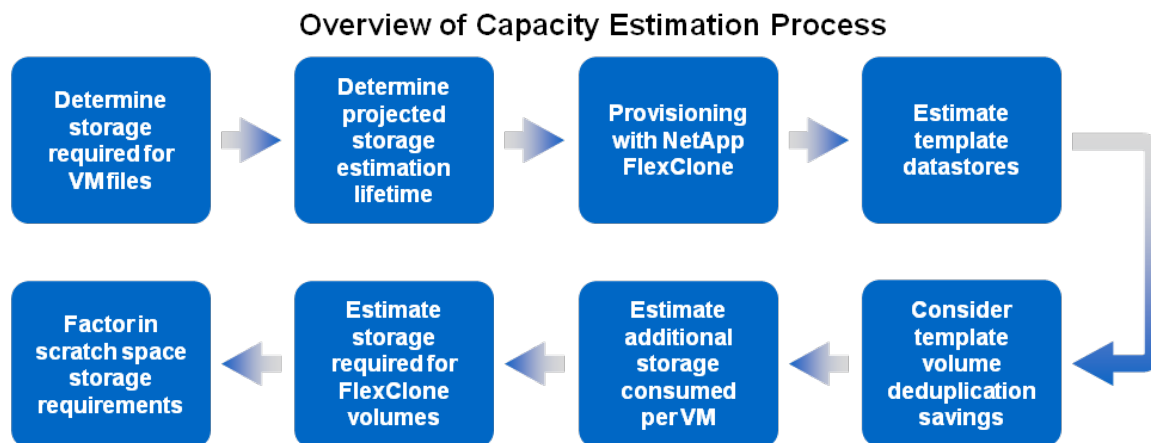
- Total number of Flash Cache modules required
- Total IOPS required in order to meet the performance needs
- Percentage of the total IOPS that require data disks considering the disk savings with write I/O optimization, NetApp transparent storage cache sharing, and Flash Cache capabilities

Citrix XenDesktop considerations: This step is applicable to all the six virtual desktop types available in Citrix XenDesktop 4.5 and can help reduce the total number of spindles required.

3.5.21.6 Perform Detailed Capacity Estimation

Figure 34 describes the various steps involved in the capacity-based storage estimation process. Each of these steps is discussed in detail in the following sections.

Figure 34. Overview of Capacity Estimation Process



3.5.21.6 Determine Projected Storage Estimation Lifetime

Determine the total number of years for which the storage growth has to be factored in. This is important because when using the NetApp FlexClone and deduplication solution components, initially the virtual machines hosted in the FlexClone volumes will not consume any space. But the new writes will require storage to commit the writes.

3.5.21.7 Getting Recommendations on Storage System Physical and Logical 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.

3.5.22 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. A complete Citrix XenDesktop deployment guide can be found [NetApp TR-3795: XenDesktop on ESX with NetApp](#).

3.5.23 Storage System Configuration Best Practices

This section of the solution guide provides a high-level overview of the components and features that should be considered 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](#)

3.5.24 Building a Resilient Storage Architecture

- Active-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 intershell 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 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, FC, or IP). A network infrastructure (FCoE, FC, or IP) should have no single point of failure. A highly available solution includes having two or more FC/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.

For additional information on designing, deploying, and configuring vSphere SAN and IP networks, refer to [TR-3749: NetApp and VMware vSphere Storage Best Practices](#).

3.5.25 Top Resiliency Best 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 deployment release available on the NOW site.
- Use the latest storage controller, shelf, and disk firmware available on the NOW site.
- Disk drive differences are FC, 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.
- 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](#).

3.5.26 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 VMware servers. NetApp recommends a one-to-one alignment of VMware datastores to flexible volumes. This design offers an easy means to understand the VMware 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.

- **LUNS.** LUNs are units of storage provisioned from a NetApp storage controller directly to the ESX Servers. The LUNs presented to the ESX Server are formatted with the VMware File System (VMFS). This shared file system is capable of storing multiple virtual desktops and is shared among all ESX Servers within the HA/DRS cluster. This method of using LUNs with VMFS is referred to as a VMFS datastore. For more information, see the VMware Storage/SAN Compatibility Guide for vSphere ESX 4 and 4i.
- **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. Leveraging 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.

3.6 Cisco Networking Infrastructure

3.6.1 Cisco Nexus 5010 28-Port Switch

The Cisco Nexus 5010 Switch is a 1RU, 10 Gigabit Ethernet/FCoE access-layer switch built to provide more than 500 Gigabits per second (Gbps) throughput with very low latency. It has 20 fixed 10 Gigabit Ethernet/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/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.

3.6.2 Cisco Nexus 5000 Series Feature Highlights

3.6.2.1 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 5000 Series provides additional features that create an even more easily managed, high-performance, unified network fabric.

3.6.2.2 10 Gigabit Ethernet and Unified Fabric Features

The Cisco Nexus 5000 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 5000 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.

3.6.2.3 Low Latency

The cut-through switching technology used in the Cisco Nexus 5000 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 5000 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 5000 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, view http://www.cisco.com/en/US/products/ps9670/prod_white_papers_list.html.

3.7 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 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 (See IDC report titled: Deployment Opportunities for Windows 7).

3.7.1 Microsoft Windows 7 Image Creation and Provisioning

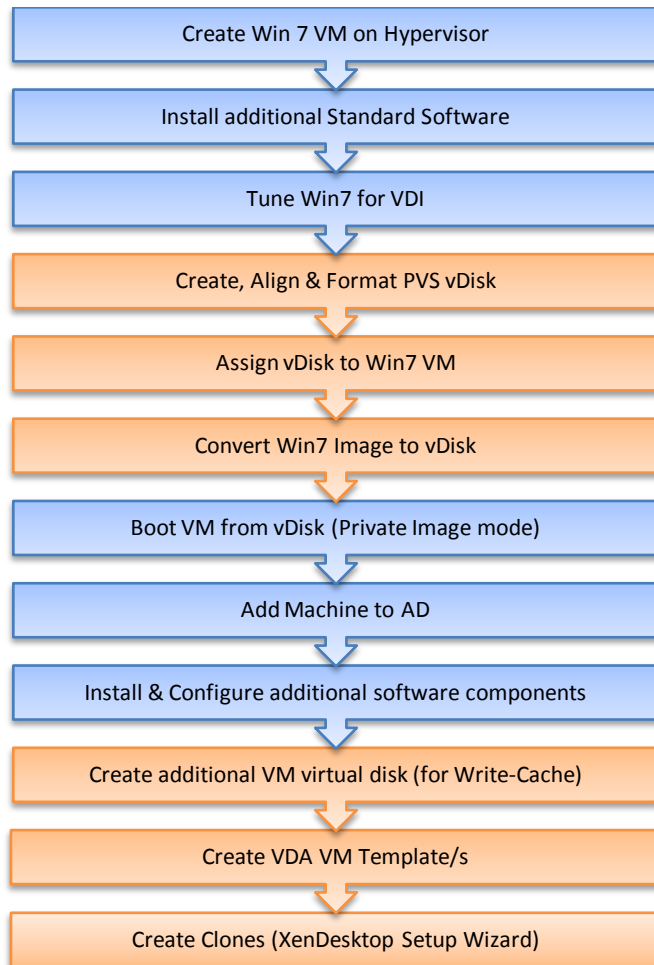
The Microsoft Windows 7 image and additional Software was initially installed and prepared as a standard Virtual Machine on VMware vSphere; prior to each one being converted into separate Citrix Provisioning server vDisk images and then 100's of Vclones being created using the XenDesktop setup wizard tool.

The XenDesktop Setup Wizard effectively creates virtual machine objects, configures - RAM, correct Network assignment and each assigned with a 3GB virtual disk hosted on a datastore mounted on the hypervisor through NFS from a NetAPP provided storage volume. It also creates and configures the relevant PVS, DDC and AD objects associated with these.

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

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

Figure 35. Windows 7 Image and vDisk Provisioning Process Overview



3.7.1.1 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 extracted to create a Citrix Provisioning server vDisk image in .vhd format.

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) 1x 3GB Virtual Disk (PVS Write-Cache)	Network:	1 x 1GbE
<ul style="list-style-type: none"> • Software Installed Prior to cloning to vDisk – <ul style="list-style-type: none"> ◦ VMWare Tools ◦ Citrix Provisioning Server Target Device 5.6.0 ◦ Microsoft Office Enterprise 2001 SP2 ◦ Internet Explorer 8.0.7600.16385 ◦ Adobe Reader 9.1.0 ◦ Adobe Flash Player 10.0.22 			

3.7.1.2 Tuning Microsoft Windows 7 Image for VDI

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

The following tuning/configurations were made to the standard image:

- Configure fixed 1.5GB page file
- Configure Networking & 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?pagelId=113247185>
 - Recommended "Default User Profile" settings were also applied and copied to "Default User" using the latest Forensit 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>

3.7.1.3 Provisioning Services (PVS) vDisk Creation

Once the Windows 7 image has initially been created with the required software it must then 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:

- Using the PVS Console (Must use the console from the PVS server)
- Create new vDisk (16GB) (this may vary depending on requirements).
- Using Diskpart set the partition offset to 1024. For more information on best practice disk alignment, visit <http://support.citrix.com/article/CTX122737>.
- 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
```

- Format the vDisk (NTFS).
- Un-mount the vDisk using the PVS Console.


To attach the New vDisk to the Windows 7 Virtual Machine:

- Set the Windows 7 virtual machine to boot from Network.
- Create a new device in PVS collection & assign MAC address of virtual machine to this PVS object.
- Assign vDisk and configure following options:

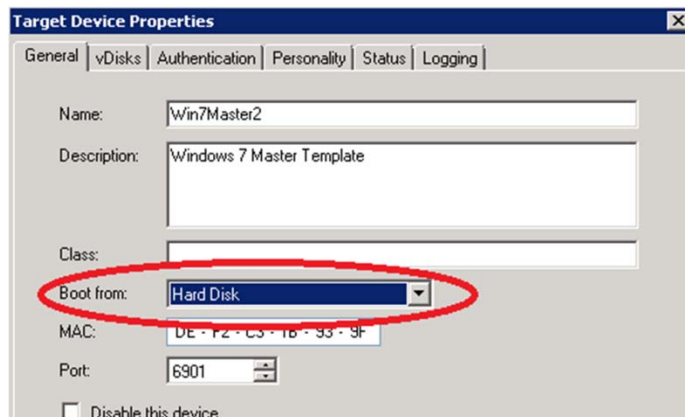
Private Image mode



Manage AD Password



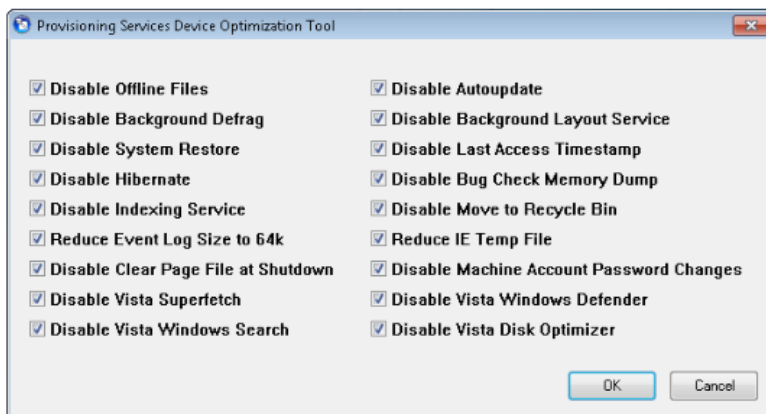
Set device to boot from hard disk



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

Clone Windows 7 Image to vDisk:

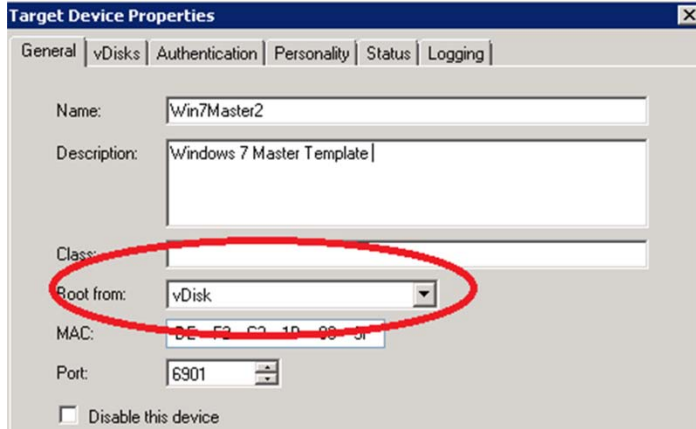
- To retain the 1024 partition offset in the vDisk the following needs to be added to the C:\Program Files\Citrix\XenConvert.ini
[parameters]
PartitionOffsetBase=1048576
- Run XenConvert.
- Run PVS Device Optimization Tool by clicking the Optimize button.



- Image to assigned vDisk (E:\).
- After the Imaging process has completed, shutdown the virtual machine.

Set virtual machine to boot from PVS vDisk (rather than vDisk) and start virtual machine:

- Using the PVS Console change the Target device options to "boot from vDisk."



- Using Virtual Center.

Add Host to Domain:

- Add the host to the Domain.
- Restart Guest OS.

Install and Configure Additional Software Components:

Software Installed Post vDisk Cloning:


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

Add 3GB Writecache.VMDK/.VHD to vDisk image: (so disk signature matches):

It is now necessary 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 cloning process, so that each clone has a unique 3GB virtual disk (E:\ Drive); this is where the per-clone PVS Write-Cache will be placed and subsequently all write IO will be conducted.

Create a new 3GB Virtual Disk using vSphere:

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



Create VDA virtual machine Template: (for cloning):

Next the virtual machine template/s must be created on the relevant NFS data store/s hosted on the NetApp Storage. If large numbers of clones are to be created it is advisable to mount several NFS volumes to the hypervisor/s balanced between at least 2 NetApp Storage controllers.

After the NFS Volumes have been mounted on the hypervisors, using the vSphere client create a Windows virtual machine in the normal way but do not start it:

Create a new Windows 7 virtual machine (Win7_PVS_Temp):

- Allocate 1.5 GB RAM
- Assign to correct virtual machine Network
- Change boot order to Network Boot

Delete assigned Virtual Disk.

- Attach the Virtual Disk created in the above stage.
- Convert the virtual machine to a Template.
- Full Copy – Template to desired NFS Volume and name (I.E. Win7PVSTemp (1)).
- Full Copy – Template to desired NFS Volume and name (I.E. Win7PVSTemp (2)).
- Until you have a template on each target NFS volume you wish to use.
- Delete (Win7_PVS_Temp) so that it does not get used accidentally.

Create VDI Clones:


Large scale cloning 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 creating clones as are registered or configured on each of the components including active directory by this tool.

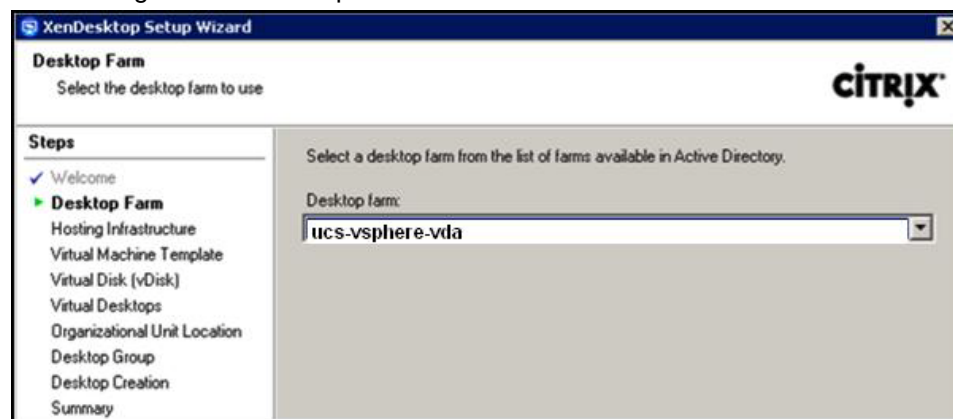
The aim is to create VDI clones evenly distributed across all of the available mounted NFS data stores, so work out how many you will be creating on each one and then run the XenDesktop Setup Tool.

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

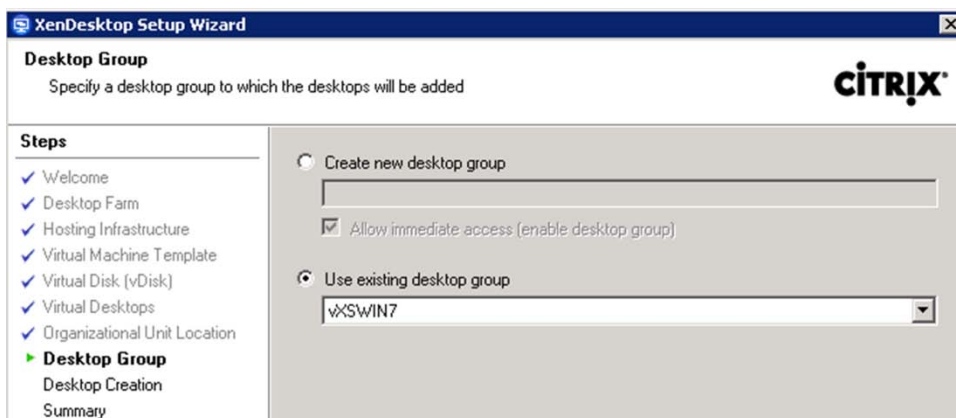
- Selecting the XenDesktop Farm
- Hosting infrastructure (hypervisor Resource Pool/Cluster)
- Select the template associated with the Volume you wish to add virtual machine instances



- Select the vDisk
- Assign Virtual Desktop numbers and Host names



- Select desired Organization Unit where machines will be created in AD
- Assign Desktops to (existing) Desktop Delivery Controller Group (Group has to be created the first time tool is run)



- Review selections and start creation process

After completed, 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).

4.0 Architecture and Design of Citrix XenDesktops on Cisco Unified Computing System and NetApp Storage

4.1 Design fundamentals

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

- **Knowledge Workers** today do not just work in their offices all day – they attend meetings, visit branch offices, work from home, and even coffee shops. These anywhere workers expect access to all of their same applications and data wherever they are.
- **External Contractors** are increasingly part of your everyday business. They need access to 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 is stuck making trade offs on the cost of providing these workers a device vs. the security risk of allowing them access from their own devices.
- **Task Workers** perform a set of well-defined tasks. These workers access a small set of applications and have limited requirements from their PCs. However, since these workers are interacting with your customers, partners, and employees, they have access to your most critical data.
- **Road Warriors** need access to their virtual desktop from everywhere, regardless of their ability to connect to a network. In addition, these workers expect the ability to personalize their PCs, by installing their own applications and storing their own data, such as photos and music, on these devices.
- **Shared Workstation** users are often found in state-of-the-art university and business computer labs, conference rooms or training centers. Shared workstation environments have the constant requirement to re-provision desktops with the latest operating systems and applications as the needs of the organization change, tops the list.

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

- **Traditional PC:** A traditional PC is what “typically” constituted a desktop environment: 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 following two virtual desktop were validated. Each of the sections provides some fundamental design decisions for each respective environment. The aforementioned 'hosted, server-based desktop' is referred to as Hosted Shared, and the 'hosted virtual desktop' as Hosted VDI.

4.1.1 Hosted Shared Design Fundamentals

Citrix XenApp 6 can be used to virtualize both desktops and applications. When evaluating a XenApp 6 deployment, here are some high-level design considerations to evaluate when deploying a server-based desktop.

4.1.1.1 Citrix XenApp Policies

Citrix XenApp 6 policies and server/farm settings have been added to Active Directory group policies enabling administrators to manage XenApp policies using their AD infrastructure. The policies can be created and configured both using Group Policy Managed console and/or directly out of the Citrix Delivery Service Console. This simplifies customer environments and enables administrators to use all the Group Policy features when administering Citrix policies.

Citrix XenApp Policies control configurations such as the ability to map client drives within a virtual desktop session and administrative tasks such as configuring the Citrix License Server FQDN for all servers in the XenApp farm. When deploying a Hosted Shared desktop on XenApp, closely assess the XenApp policies for the following configurations:

- Configure farm settings such as Virtual IP, Health Monitoring and Recovery, and multimedia acceleration
- Control sound quality for client devices
- Allow users to access the Documents folder on their local client device
- Allow or prevent remote users from being able to save to their hard drives from a session
- Allow or prevent users from accessing the Windows clipboard
- Set a required encryption level for Citrix plug-ins
- Set the session importance level, which, along with the application importance level, determines resource allotment for Preferential Load Balancing

4.1.1.2 Worker Groups

Worker groups allow similar XenApp servers to be grouped together to greatly simplify the management of XenApp farms. Worker groups simplify application workload/silo management and help ensure that all the servers in a worker group have the same applications and policies, thus eliminating 'configuration drift'.

4.1.1.3 Load Managed Groups

Load managed groups are created to focus a set of XenApp servers on a particular set of applications or desktops. This is done for numerous business and technical reasons including update frequency, business unit server ownership, criticality, regional access, and language requirements.

When creating a load managed group, each group must provide enough redundancy to be capable of supporting all users in the event of a server failure. This results in an N+1 scenario where there is at least one additional XenApp server per load managed group. In many situations, organizations implement an N+10 percent strategy where an additional 10 percent of XenApp servers per load managed group are allocated in order to allow for multiple server failures or maintenance.

4.1.1.4 Designing a Citrix XenApp Deployment

For a complete overview of configurations, architecture, and design recommendations for publishing both applications and desktops using Citrix XenApp, visit <http://support.citrix.com/proddocs/index.jsp?topic=/xenapp6-w2k8/ps-designing-wrapper.html>

4.2 Hosted VDI Design Fundamentals

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

4.2.1 Hypervisor Selection

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

- XenServer

Citrix® XenServer® is a complete, managed [server virtualization](#) platform built on the powerful Xen® hypervisor. Xen technology is widely acknowledged as the fastest and most secure virtualization software in the industry. XenServer is designed for efficient management of Windows® and Linux® [virtual servers](#) and delivers cost-effective server consolidation and business continuity. More information on Hyper-V can be obtained at the company 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 multipathing 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.

4.2.2 Provisioning Services

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

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

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

- Cache on 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 Cache.

4.2.3 Designing a Citrix XenDesktop Deployment

For a complete overview of configurations, architecture, and design recommendations for delivering virtual desktops with XenDesktop, visit <http://support.citrix.com/proddocs/index.jsp?topic=/xendesktop-bdx/cds-admin-deploy-plan-wrapper-bdx.html>

5.0 Solution Validation

In this section we discuss the configuration and tuning that were done to various components for a complete solution validation.

5.1 Configuration Topology for Scalability of Citrix XenDesktops on Cisco Unified System and NetApp Storage

Figure 36. Architecture Block

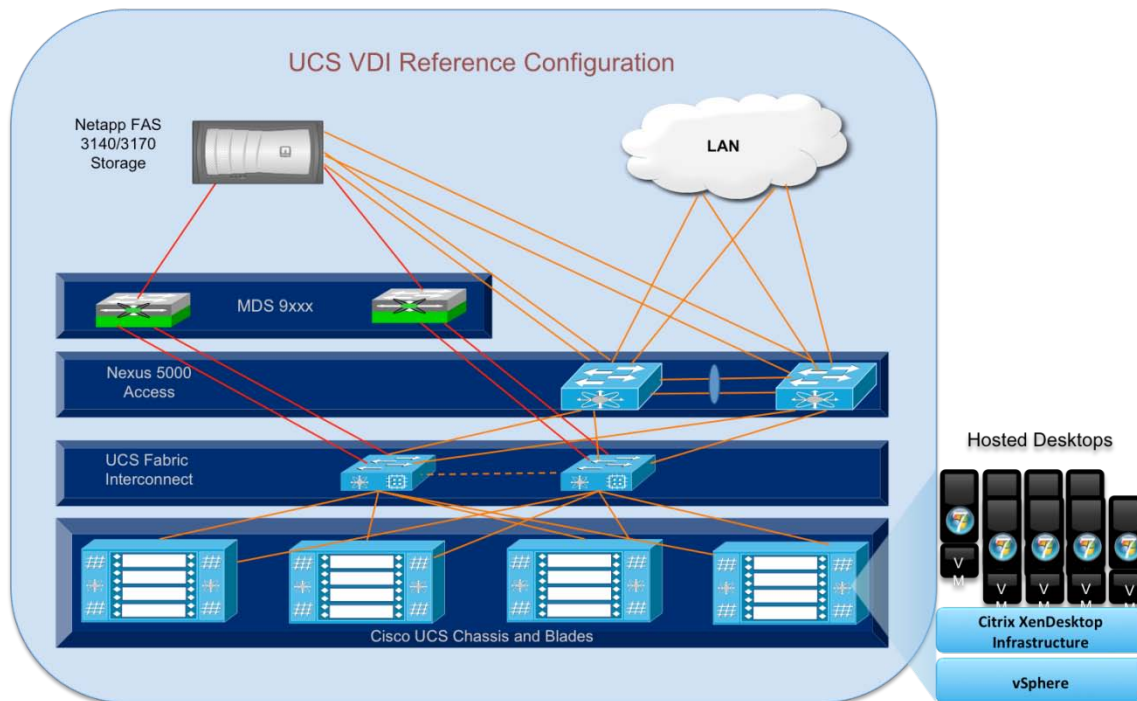


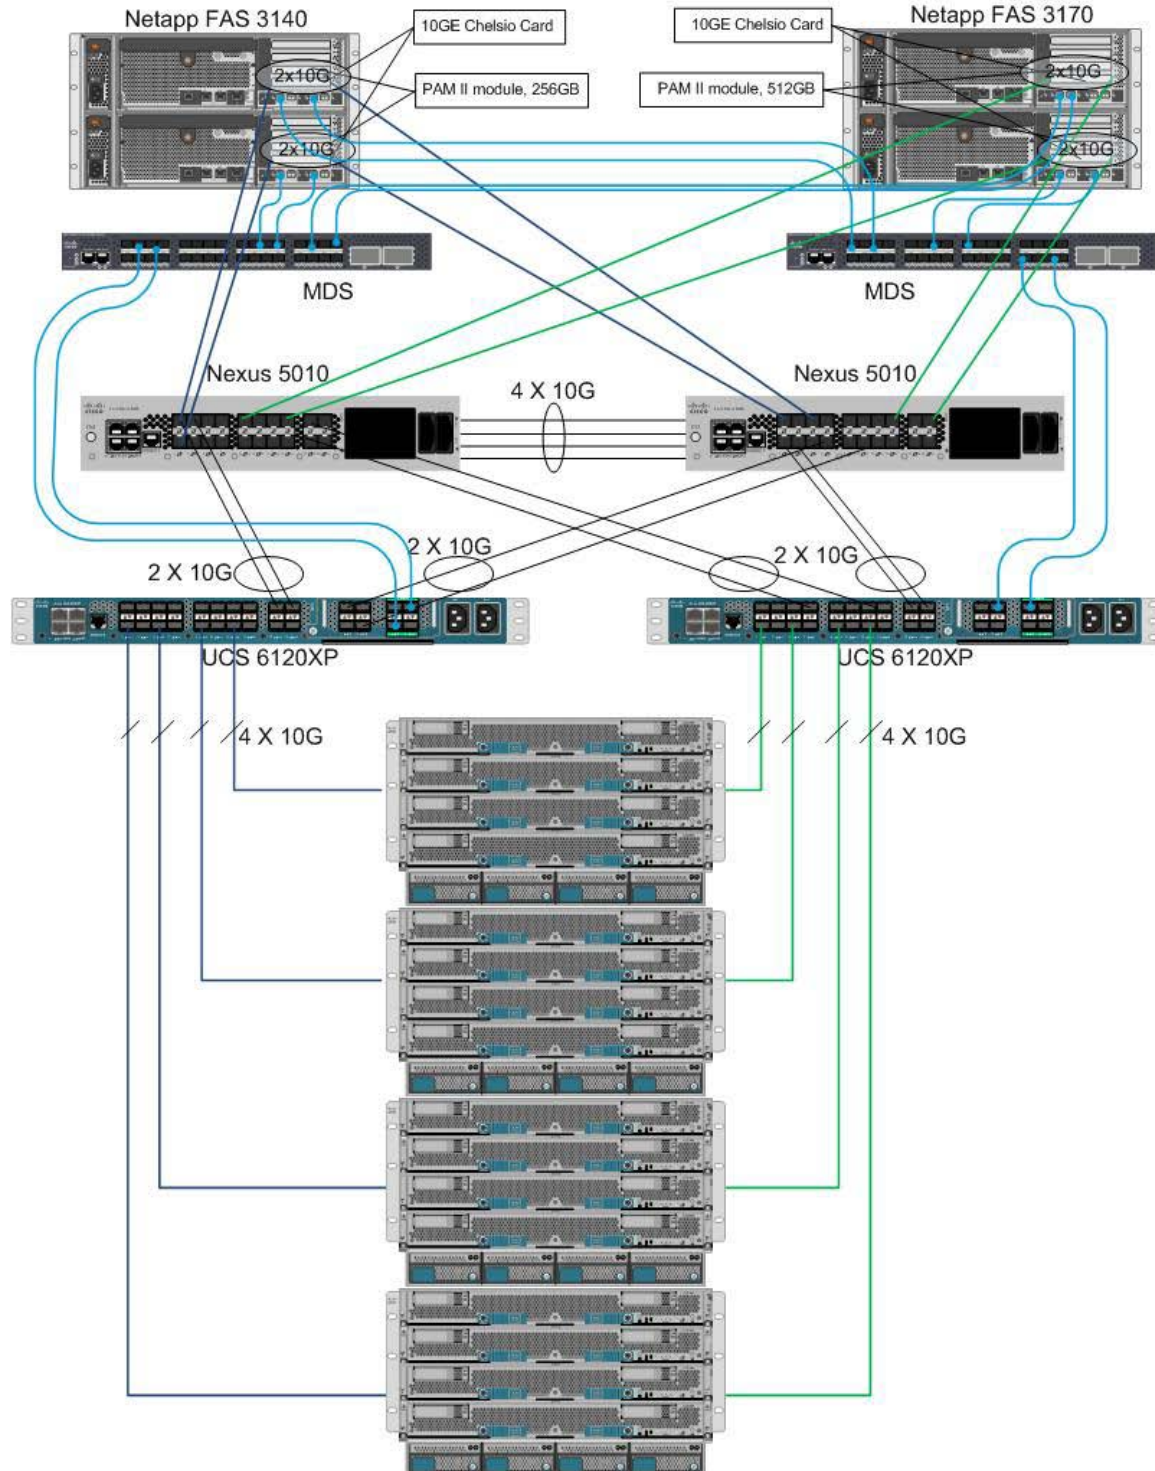
Figure 36 above captures the architecture diagram for purpose of this study. The architecture is distinctly divided into four layers:

- Cisco UCS Compute platform
- The virtual desktop infrastructure that runs on a virtual infrastructure (Hypervisor).
- Network Access layer and LAN
- Storage Access (SAN) and Storage array

Each of these components has a crucial role to play in the desktop deployment in the data center (DC). There are lots of choices for customer to build the DC for the purpose of providing desktops, and each component can make or break a successful user experience.

A detailed architectural diagram is shown in Figure 37 with all the interconnections.

Figure 37. Detailed Architectural Configuration

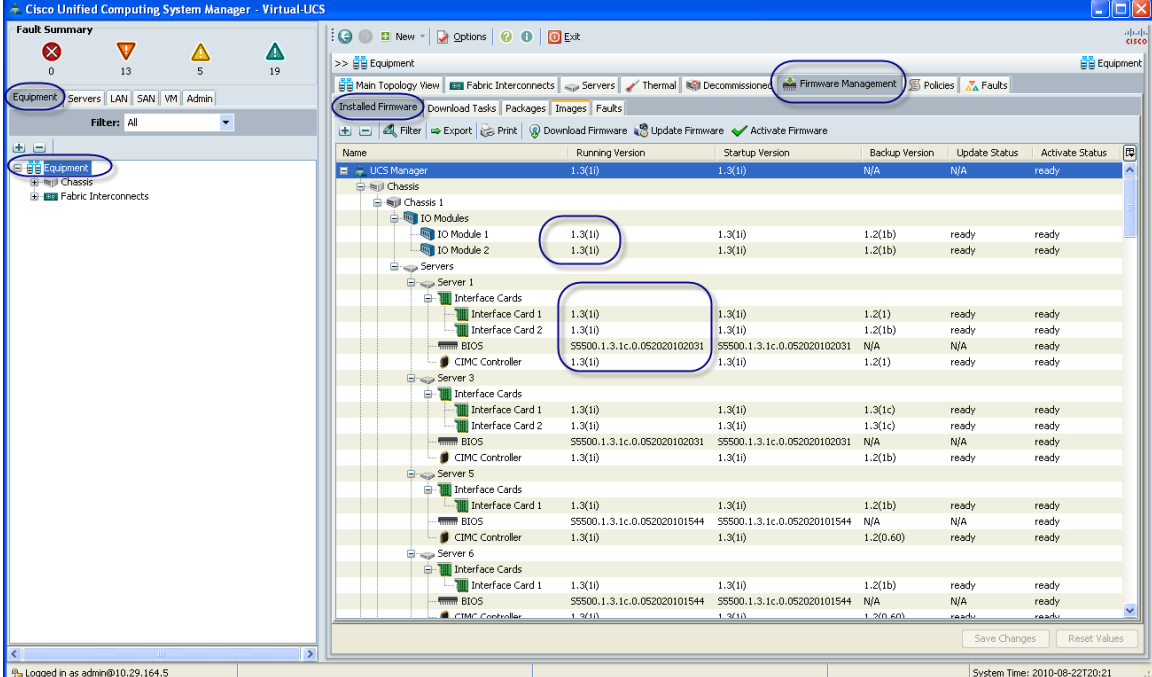


5.2 Cisco UCS Configuration

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

- Cisco UCS CLI Configuration guide
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 M-Series 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

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

1	Bring up the Fabric interconnect and from a Serial Console connection set the IP address, gateway, and the hostname of the primary fabric interconnect. Now bring up the second fabric interconnect after connecting the dual cables between them. The second fabric interconnect automatically recognizes the primary and ask if you want to be part of the cluster, answer yes and set the IP address, gateway and the hostname. 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 UCSM will let you do everything that you could do from the CLI and we will highlight the GUI methodology here.
3	First check the firmware on the system and see if it is current. The latest firmware as of now is 1.3(1i).
 <p>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.</p>	
4	Configure and enable the server port on the FI. To bring the chassis online acknowledge the chassis.

5	<p>Configure and enable upstream Ethernet links and FC links.</p>
6	<p>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. You have to make sure we have ample IP address for all the blades and make sure the gateway and netmask is set correctly.</p>

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

7.1 MAC pool

Name	From	To
[00:25:B5:AB:CD:01 - 00:25:B5:AB:CD:20]	00:25:B5:AB:CD:01	00:25:B5:AB:CD:20
[00:25:B5:AB:CD:21 - 00:25:B5:AB:CD:48]	00:25:B5:AB:CD:21	00:25:B5:AB:CD:48
[00:25:B5:AB:CD:49 - 00:25:B5:AB:CD:88]	00:25:B5:AB:CD:49	00:25:B5:AB:CD:88

7.2 WWPN pool

Name	From	To
[20:00:00:25:B5:0A:AD:01 - 20:00:00:25:B5:0A:AD:20]	20:00:00:25:B5:0A:AD:01	20:00:00:25:B5:0A:AD:20
[20:00:00:25:B5:0A:AD:21 - 20:00:00:25:B5:0A:AD:48]	20:00:00:25:B5:0A:AD:21	20:00:00:25:B5:0A:AD:48
[20:00:00:25:B5:0A:AD:49 - 20:00:00:25:B5:0A:AD:88]	20:00:00:25:B5:0A:AD:49	20:00:00:25:B5:0A:AD:88

7.3 WWNN pool

7.4 UUID pool

7.5 Server pool

The screenshot displays the Cisco Unified Computing System Manager - Virtual UCS interface. The left sidebar shows a tree view with 'Pools' expanded, highlighting 'UUID Suffix Pools' (labeled 1) and 'Server Pools' (labeled 2). The main pane shows the 'UUID Suffix Pools' configuration page with a table of existing pools. Overlaid on this are two configuration windows:

- Create UUID Suffix Pool:** A dialog box with 'Define name and description' section. The 'Name' field is set to 'XD-UUID_POOL' (labeled 3). The 'Prefix' is set to 'derived'.
- Create a Block of UUID Suffixes:** A sub-dialog box with 'From' set to '1234-567890ABCDE1' and 'Size' set to '16'. The 'Add' button is highlighted.

Blue arrows indicate the workflow: from the 'UUID Suffix Pools' list (2) to the 'Create UUID Suffix Pool' dialog (3), and then to the 'Create a Block of UUID Suffixes' sub-dialog.

The screenshot illustrates the process of creating a server pool in the Cisco Unified Computing System Manager. The main window shows the 'Server Pools' list. A 'Create Server Pool' dialog is open, showing 'Set Name and Description' and 'Add Servers' steps. A 'Select all the servers you want to make it part of the pool' dialog is also shown, with a list of servers and a 'Finish' button.

8 Create vHBA template

The screenshot illustrates the process of creating a vHBA template in the Cisco Unified Computing System Manager. The main window shows the 'vHBA Templates' list. A 'Create vHBA Template' dialog is open, showing fields for Name, Description, Fabric ID, Select VSAN, Template Type, Max Data Field Size, WWNN Pool, QoS Policy, Pin Group, and Stats Threshold Policy.

9 Create vNIC template

10 Create boot from SAN policies, adapter policies, Local boot policy for ESX4i install was configured as RAID-Mirrored to avoid any catastrophic disk failures.

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


12 After associating a server pool to the service profile template, just right click to deploy as many service profile as you need and UCSM will automatically start configuring these new service profile templates on the selected blade servers.

13 At this point, the servers are ready for OS provisioning, we would recommend setting up a PXE server to fasten the OS install. Virtual media CD based OS installation is also possible.

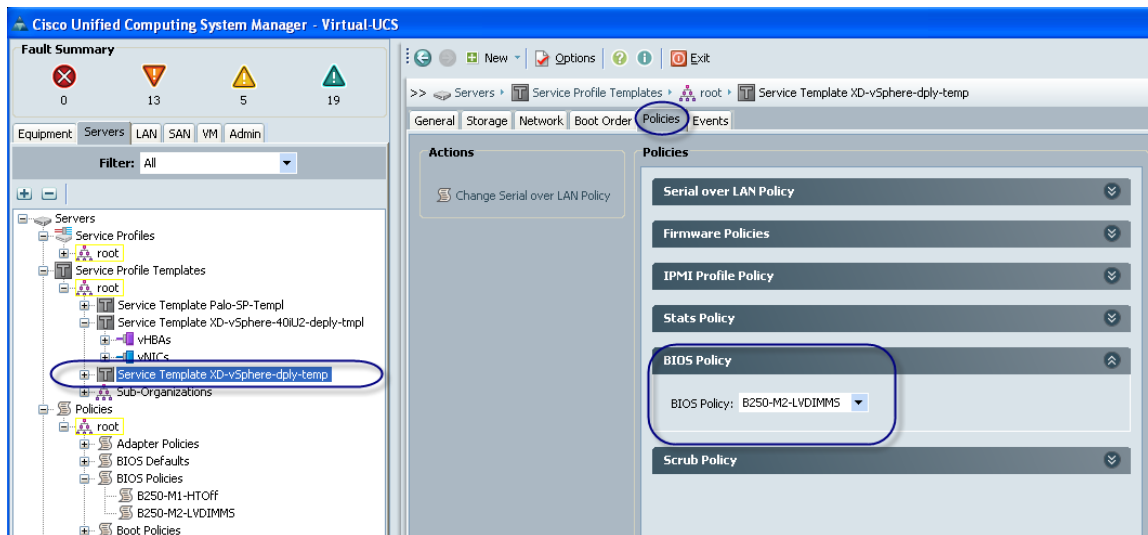
When working with 4 GB -1333MHz DDR3 Low voltage dual rank DIMM is it will show up as 1067 if you do not set the performance mode in the policy. For example:

You will need to 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.

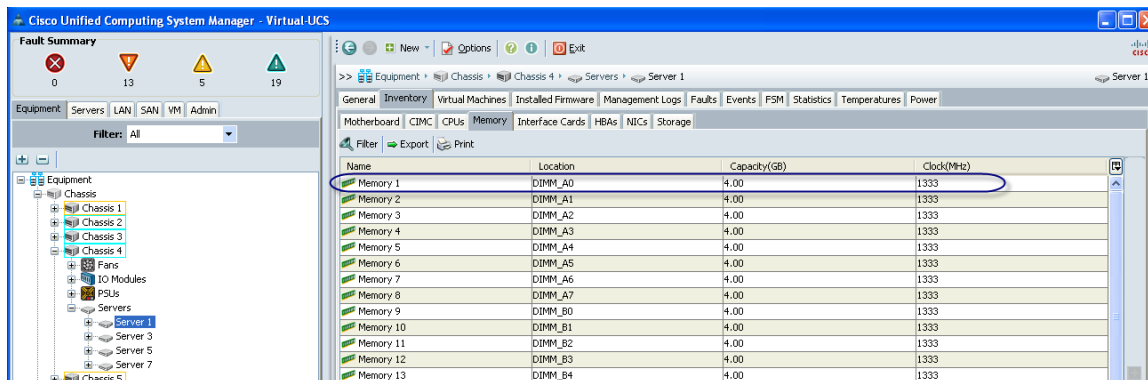
Configure a BIOS policy with everything being in platform-default and just change the default power-saving-mode into performance-mode:



Now add this 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:



5.2.1 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.

5.2.2 System Class Configuration

Systems Class is the global operation where the entire system interfaces with 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 classes allowed with four or more users defined have the following configurable rules:
 - CoS to Class Map
 - Weight: Bandwidth
 - Per class MTU
 - Property of Class (Drop v/s no drop)
- Max MTU per Class allowed is 9216.
- Through 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).

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

5.2.3 Cisco UCS System Class Configuration

Cisco Unified Computing System defines user class names as follows.

- Platinum
- Gold
- Silver
- Bronze

Name Table Map between Cisco Unified Computing System and the NXOS

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

Class to CoS Map by default in Cisco Unified Computing System

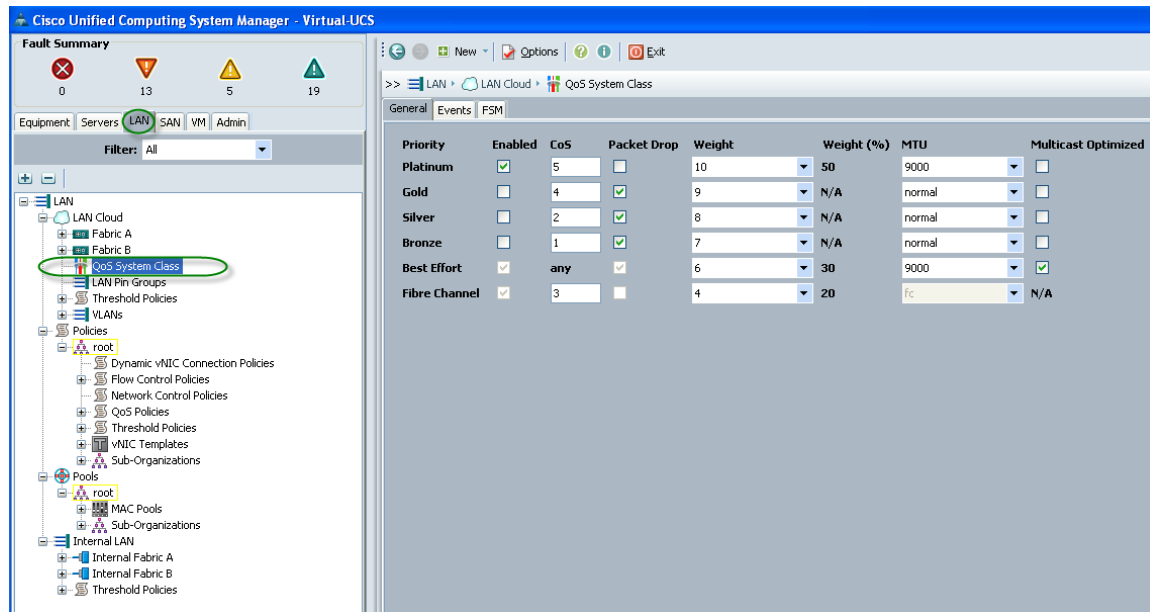
Cisco UCS Class Names	Cisco UCS Default Class Value
Best effort	Match any
Fc	3
Platinum	5
Gold	4
Silver	2
Bronze	1

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:

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




Priority	Enabled	CoS	Packet Drop	Weight	Weight (%)	MTU	Multicast Optimized
Platinum	<input checked="" type="checkbox"/>	5	<input checked="" type="checkbox"/>	10	50	9000	<input type="checkbox"/>
Gold	<input type="checkbox"/>	4	<input checked="" type="checkbox"/>	9	N/A	normal	<input type="checkbox"/>
Silver	<input type="checkbox"/>	2	<input checked="" type="checkbox"/>	8	N/A	normal	<input type="checkbox"/>
Bronze	<input type="checkbox"/>	1	<input checked="" type="checkbox"/>	7	N/A	normal	<input type="checkbox"/>
Best Effort	<input checked="" type="checkbox"/>	any	<input checked="" type="checkbox"/>	6	30	9000	<input checked="" type="checkbox"/>
Fibre Channel	<input checked="" type="checkbox"/>	3	<input type="checkbox"/>	4	20	Fc	N/A

Now in the LAN tab under policies, define a platinum-policy and select platinum as the priority.



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

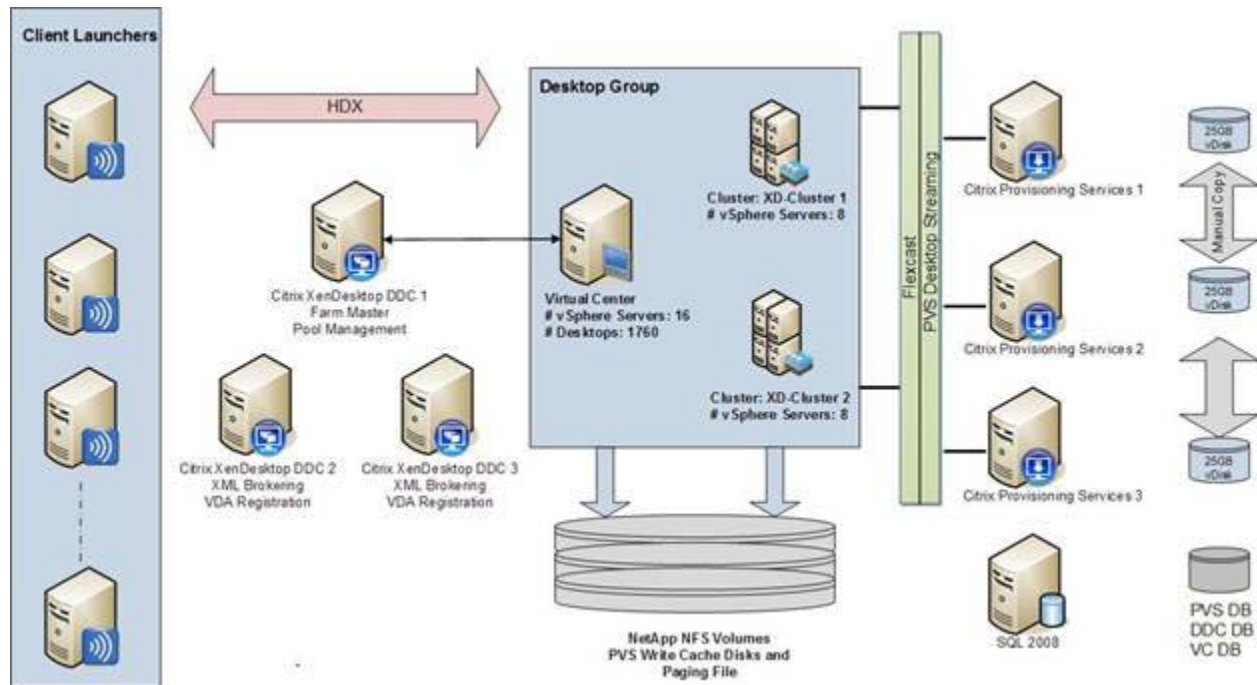


This is a unique value proposition of the Cisco Unified Computing System with respect to end-to-end QoS. For example, you could have a VLAN for the NetApp storage and configure Platinum policy and Jumbo frames and get an end-to-end QoS and performance guarantee. You can configure the NIC to have a no-drop class along with the platinum policy.

5.3 Citrix XenDesktop Configuration

Figure 38 shows the Citrix XenDesktop configuration.

Figure 38. Citrix XenDesktop Configuration



Summary of Environment:

- 3 Desktop Delivery Controllers
- 3 Provisioning Services Servers
- 2 vSphere Clusters
- 1760 Virtual Desktops
- 1 Citrix Licensing Server
- 1 File Server for Roaming Profiles and VSI data
- 1 SQL 2008 Server for DDC and PVS DBs
- 2 NetApp Filers, 6 NFS Volumes
- Multiple client launchers

Configuration by component shown in Tables 1 through 6.

Table 1. VMware vSphere 4

vSphere 4			
Hardware:	Cisco UCS B-series Blade server	Model:	B250 –M2
OS:	Citrix XenServer 5.6	Service Pack:	-
CPU:	2 x 6 Core Intel 5680 @ 1333 GHz (24 Logical Cores Total)	RAM:	192 GB @ 1333 MHz
Disk:	Boot From SAN	Network:	Palo adapter 4 x 10GbE

Table 2. Citrix Provisioning Server 5.6

Citrix Provisioning Server 5.6			
OS:	Windows 2008 Enterprise R2 64bit	Service Pack:	-
CPU:	2 x vCPU	RAM:	8192MB
Disk:	1 x70GB Virtual Disk (hosted on NFS target volume on NetApp Filer)	Network:	1 x 10GbE (VMXNET3)
<ul style="list-style-type: none"> • Database for PVS hosted on separate Microsoft SQL Server 2008 64bit 			

Table 3. Citrix XenDesktop Desktop Delivery Controller

Citrix XenDesktop Desktop Delivery Controller			
OS:	Windows 2003 R2 Enterprise 64bit	Service Pack:	2
CPU:	4 x vCPU	RAM:	4096MB
Disk:	1 x50GB Virtual Disk (hosted on NFS target volume on NetApp Filer)	Network:	1 x 10GbE (VMXNET3)
<ul style="list-style-type: none"> • Citrix XenDesktop DDC - 400W2K3X64004 <ul style="list-style-type: none"> ◦ Desktop Delivery Controller - Services Hotfix XD*400DDC002 ◦ Pool Management Service Hotfix XD*400PM003 • Citrix Web Interface • Database for DDC hosted on separate Microsoft SQL Server 2008 64bit 			

Table 4. Citrix Virtual Desktop Agent

Citrix Virtual Desktop Agent (Virtual Desktops)			
OS:	Windows7 Enterprise 32bit	Service Pack:	1
CPU:	1 x vCPU	RAM:	1536MB
Disk:	C: 1 x 16GB (PVS vDisk) D: 1x 3GB Virtual Disk (PVS Write-Cache and OS paging file hosted on NFS target volume on NetApp Filer)	Network:	1 x 1GbE (E1000)
<ul style="list-style-type: none"> • VDA Service Pack 1 - XDE400VDAWX86SP1 (Version 4.0.5010) 			

Table 5. 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 (hosted on NFS target volume on NetApp Filer)	Network:	1 x 10GbE (VMXNET3)

Table 6. 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 (hosted on NFS target volume on NetApp Filer)	Network:	1 x 10GbE (VMXNET3)

Other dedicated Infrastructure Virtual Machines:

- Active Directory Server (Directory Services, DNS & DHCP)
- User Profile Server (Hosts Roaming Profiles)
- File Server (Collects Login VSI data)
- STAT Console Server (STAT configuration & Data analysis)

5.3.1 XenDesktop Desktop Delivery Controller (DDC)

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

The DDCs were configured such that:

- DDC 1: Farm Master and Pool Management
- DDC 2 & 3: VDA Registrations and XML Brokering

In this environment, 3 DDCs (4vCPU, 4GB RAM) easily sustained the farm size of 1760 desktops and proved stable at all various stages of testing.

5.3.2 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 DDCE400W2K3X64005
- 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 & 3 for Registrations and XML Brokering as per CTX117477

- Created one Desktop Group and used a single Virtual Center that managed the full amount of desktop virtual machines (1760)
- VMware DRS was not used for the testing and 110 desktops were fixed to each host server
- Within Virtual Center, two dedicated clusters were created per two blade chassis and a third cluster for the remaining infrastructure components
 - 2 VC cluster x 880 - Virtual Desktops for the scaled out test of 1760
 - 1 VC cluster - XenDesktop and associated infrastructure

By default, Pool Management will attempt to start 10 percent of the total pool size. In a large environment this may be more than the hosting infrastructure can handle.

- The number of concurrent requests can be throttled 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="40"/>
- The Pool Management service needs to be restarted to read the new configuration.

5.3.3 Provisioning Services Configuration

For the scaled out test, a total of 6 Provisioning Servers supported 1760 Windows 7 desktops. The Provisioning Server streamed to ~586 desktops per virtual machine-based server using a single virtual NIC. Note: it was determined that the PVS farm could have supported the full amount of desktops with two servers.

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

- Changed the Threads per port from the default 8 to 31. 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 D: drive.
- Copied the 25GB Windows7 vDisk to each server's D: drive.

5.3.4 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 1000s of virtual machines hosted on hypervisor servers to PXE boot from and share a single gold Windows 7 Image.

5.3.5 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 IO 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:

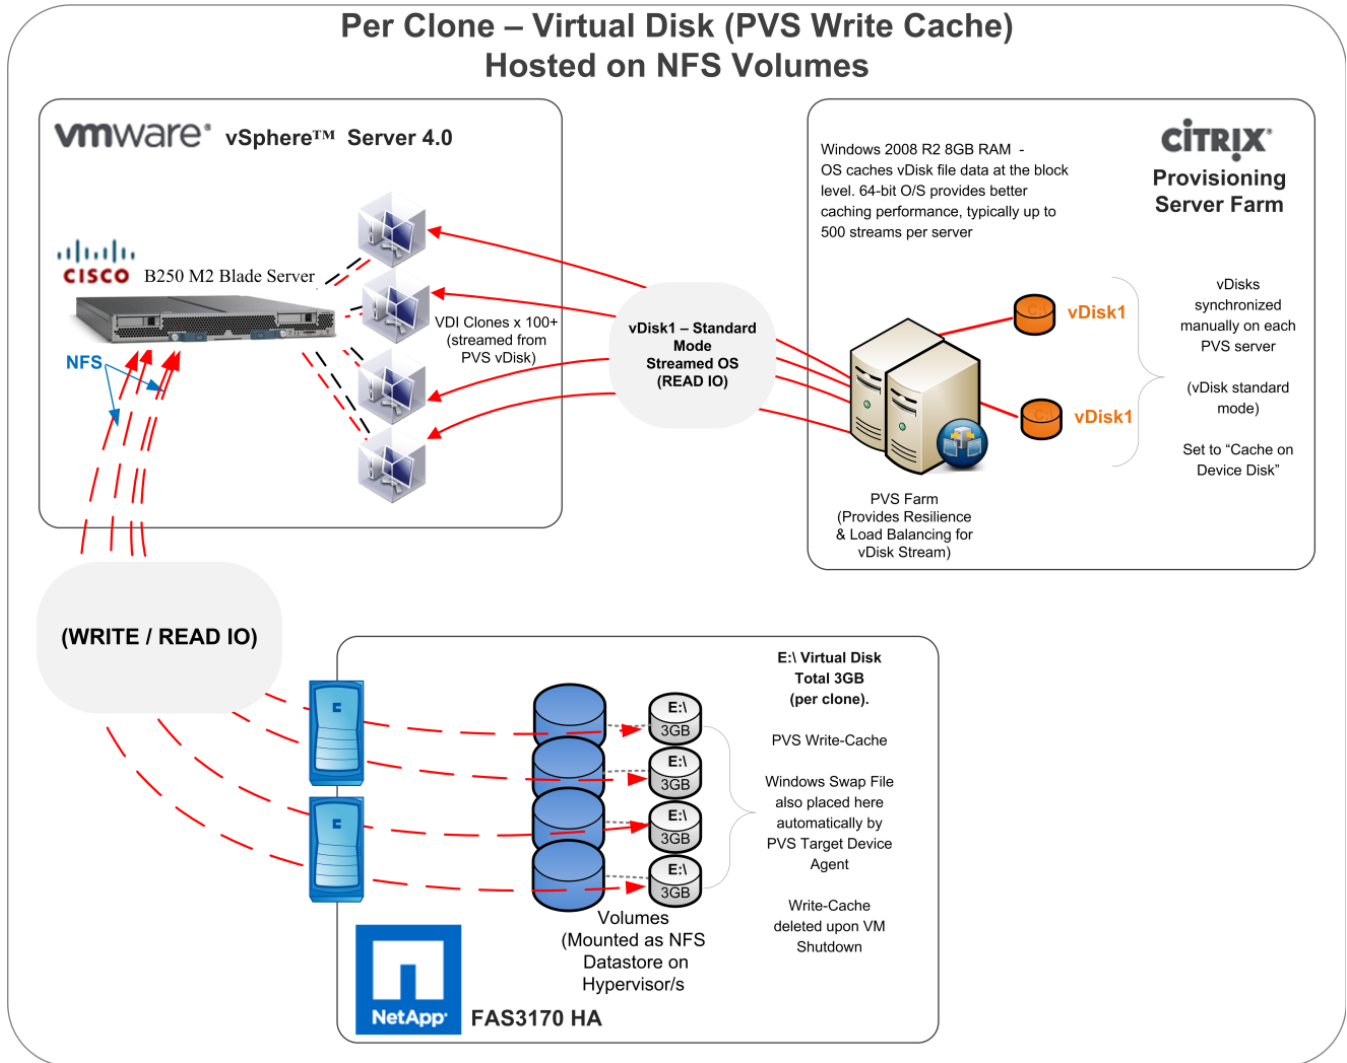
- PVS Server
- Hypervisor RAM
- Device Local Disk (an additional Virtual Disk for VDI instances)

For optimal performance and scalability the "Cache on devices HD" option is used, a 3GB virtual disk is assigned to the virtual machine templates used in the clone creation process (described in section 5.7). By creating the 3GB drives associated with the templates on NFS volumes, mounted on the hypervisors; we are then able to create VDI instances each with its own 3GB drive where the PVS Write Cache will be placed. In addition the PVS Target device agent installed in the Windows 7 image will also automatically place the Windows swap file on the same drive when this mode is enabled.

Therefore both the PVS Write Cache and Windows Swap file is now hosted on an NFS mounted volumes hosted on a NetApp storage. To further increase scalability load balancing across multiple Volumes and storage Controllers was achieved 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 39 below illustrates Multiple virtual machine instances hosted on a hypervisor server booting from a PVS single master image, each one has a virtual disk hosted on different NetApp provided NFS volumes where the PVS cache is placed. This help ensures that ALL Write IO takes place on the NetApp storage over NFS using high performance storage.

Figure 39. vDisk Hosting on NFS Volumes

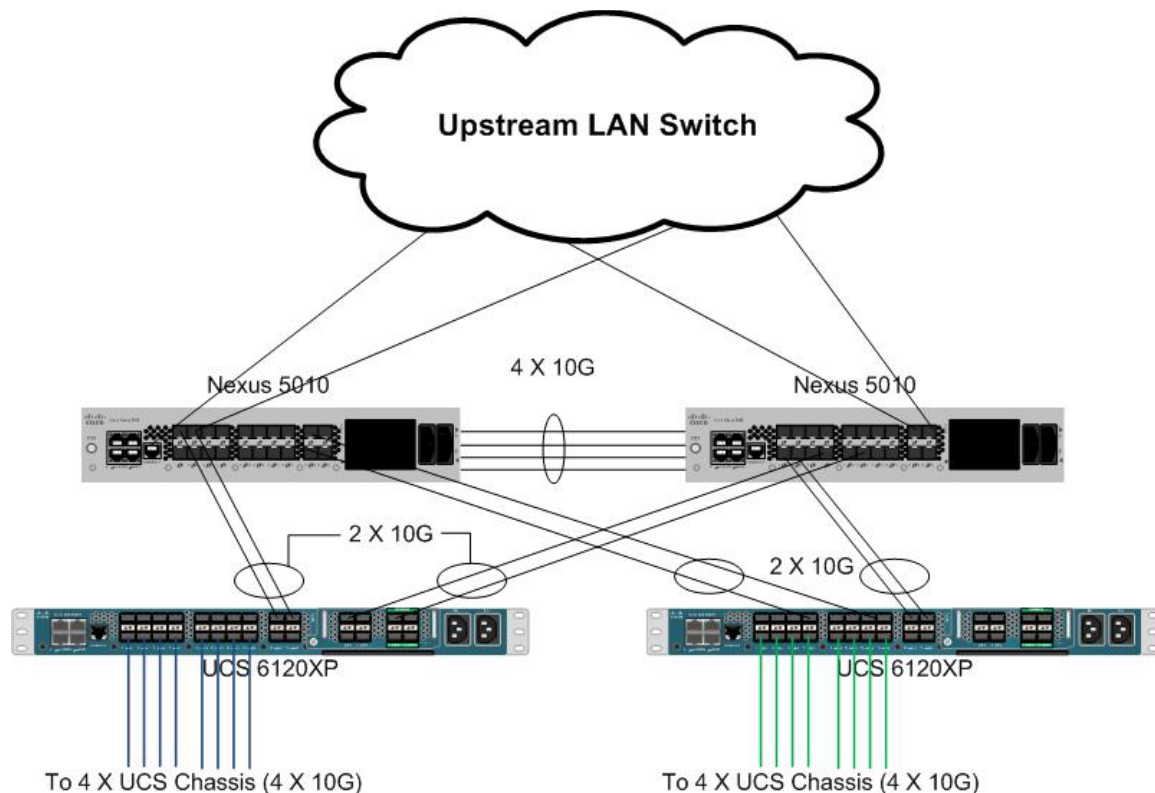


5.4 LAN Configuration

This configuration consists of a pair of Cisco Nexus 5010, a family of low-latency, line-rate, 10 Gigabit Ethernet and Fibre Channel over Ethernet (FCoE) switches for data center applications. Four 10G Ethernet uplink ports are configured on each of the Cisco UCS Fabric interconnect and they are connected to the Nexus 5010 pair in a bow tie manner as shown below. The Fabric interconnect is in End host mode, as we are doing both FC as well as Ethernet data access and as per the recommended best practice of Cisco UCS. We built this out for scale and have provisioned more than 40 G per Fabric interconnect as we are building a scalable and expandable system.

The upstream configuration is beyond the scope of this document; there are some good reference document [4] that talks about best practices of using Cisco Nexus 5000 and 7000.


Figure 40. Network Configuration with Upstream Cisco Nexus 5000 from Cisco Unified Computing System



The Cisco Nexus 5000 is used to connect to the NetApp FAS 3140/3170 storage system for NAS access. NetApp supports dual port 10G Chelsio cards which are configured in a portchannel and connected to the pair of Cisco Nexus 5000 downstream. This allows end-to-end 10G access, we have implemented jumbo frames on the ports and have priority flow control on, with platinum COS for the NetApp storage data access. The NetApp connectivity diagram is shown in Figure 39. Here again, we have a total of 40G bandwidth available for the servers.

Please refer to APPENDIX A for a detailed configuration of one of the Cisco Nexus 5000 used in our setup.

Figure 41. Network Configuration for NetApp NAS or Filer Storage



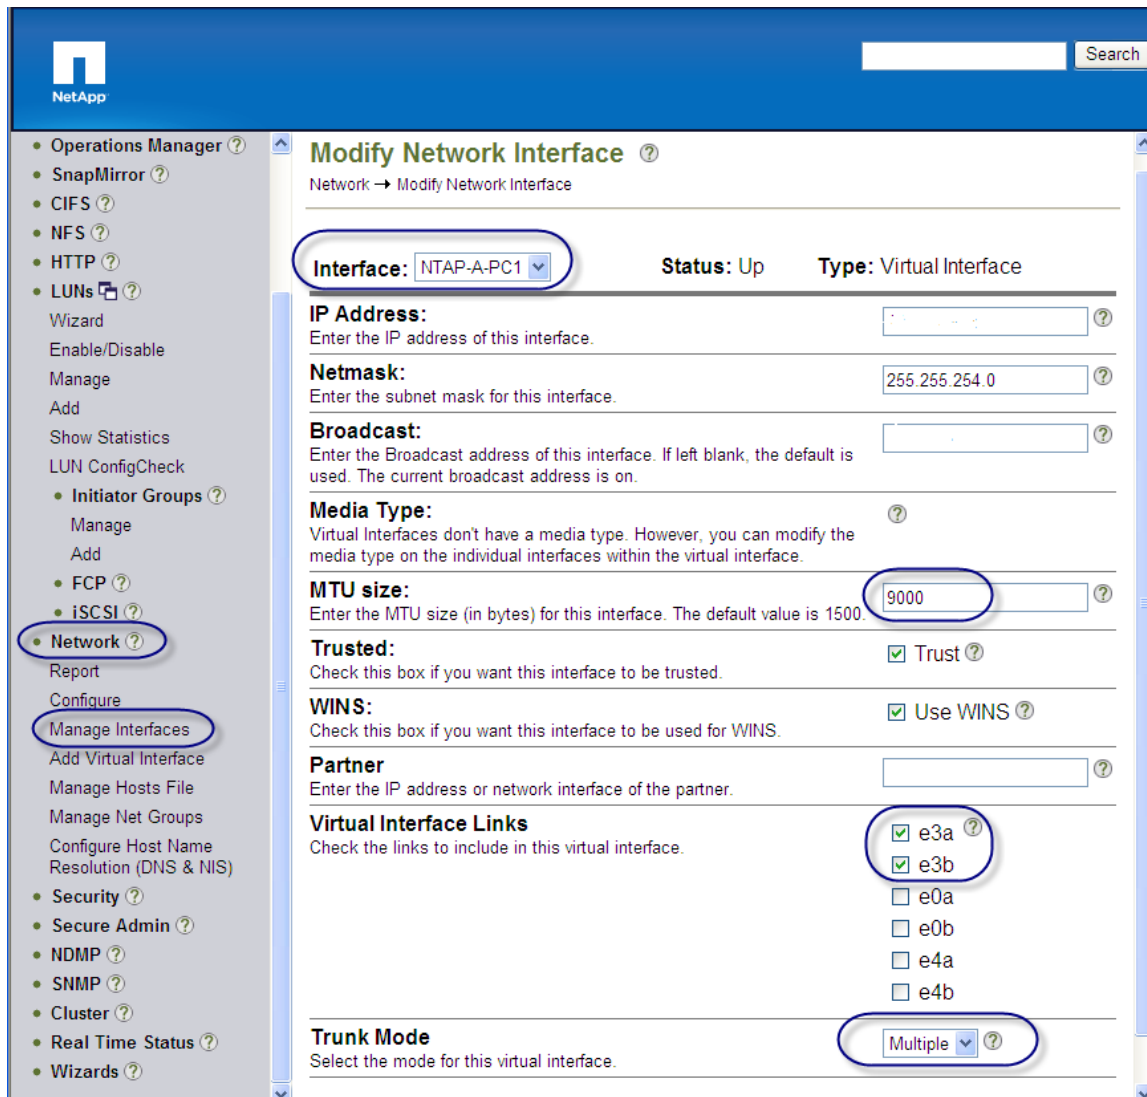
The configuration on the NetApp storage as gathered from the filer view is shown in Figure 42.

Figure 42. Network Configuration on the NetApp Storage Array Side

The screenshot shows the NetApp 'Manage Network Interfaces' web interface. The left sidebar contains a navigation menu with options like Initiator Groups, FCP, iSCSI, Network, Report, Configure, Manage Interfaces (highlighted), Add Virtual Interface, Manage Hosts File, Manage Net Groups, Configure Host Name Resolution (DNS & NIS), Security, Secure Admin, NDMP, SNMP, Cluster, Real Time Status, and Wizards. The main content area is titled 'Manage Network Interfaces' and includes buttons for '[Add a New Virtual Interface]', '[Add a New VLAN]', and '[Show All Interface Details]'. Below these buttons is a table of network interfaces:

Interface	Status	Address	Operations
e0M	Down	--	Modify Up Down --
e0a	Down	--	Modify Up Down --
e0b	Down	--	Modify Up Down --
e4a	Down	--	Modify Up Down --
e4b	Down	--	Modify Up Down --
NTAP-A-PC1	Up	10.255.198.12	Modify Up Down Remove Trunk
- e3a	Up	--	Modify -- --
- e3b	Up	--	Modify -- --
lo	Up	127.0.0.1	-- -- --

A 'Refresh' button is located at the bottom of the interface table.



Modify Network Interface

Network → Modify Network Interface

Interface: NTAP-A-PC1 **Status:** Up **Type:** Virtual Interface

IP Address:
Enter the IP address of this interface.

Netmask:
Enter the subnet mask for this interface. 255.255.254.0

Broadcast:
Enter the Broadcast address of this interface. If left blank, the default is used. The current broadcast address is on.

Media Type:
Virtual Interfaces don't have a media type. However, you can modify the media type on the individual interfaces within the virtual interface.

MTU size:
Enter the MTU size (in bytes) for this interface. The default value is 1500. 9000

Trusted:
Check this box if you want this interface to be trusted. ☒ Trust

WINS:
Check this box if you want this interface to be used for WINS. ☒ Use WINS

Partner
Enter the IP address or network interface of the partner.

Virtual Interface Links
Check the links to include in this virtual interface.
☒ e3a
☒ e3b
☐ e0a
☐ e0b
☐ e4a
☐ e4b

Trunk Mode
Select the mode for this virtual interface. Multiple

5.5 SAN Configuration

A pair of MDS 9134 series was used in the configuration to connect to the FC port of the Cisco UCS Fabric interconnect FC expansion module ports to the NetApp storage FC ports. MDS single initiator zone was done to connect to the NetApp FC ports. The SAN switch was predominantly used for configuring boot from SAN of the ESX server blades.

The infrastructure volumes 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-ESX-BFS vsan 1
```

```
zone name XD-vSphere-Server-1-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-vSphere-Server-2-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 MDS pair to take care of the Fabric B side as shown below.

```
MDS-B# sh zoneset active vsan 1
```

```
zoneset name FAB-B-XD-ESX-BFS vsan 1
```

```
zone name XD-vSphere-Server-1-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-vSphere-Server-2-fc1 vsan 1
```

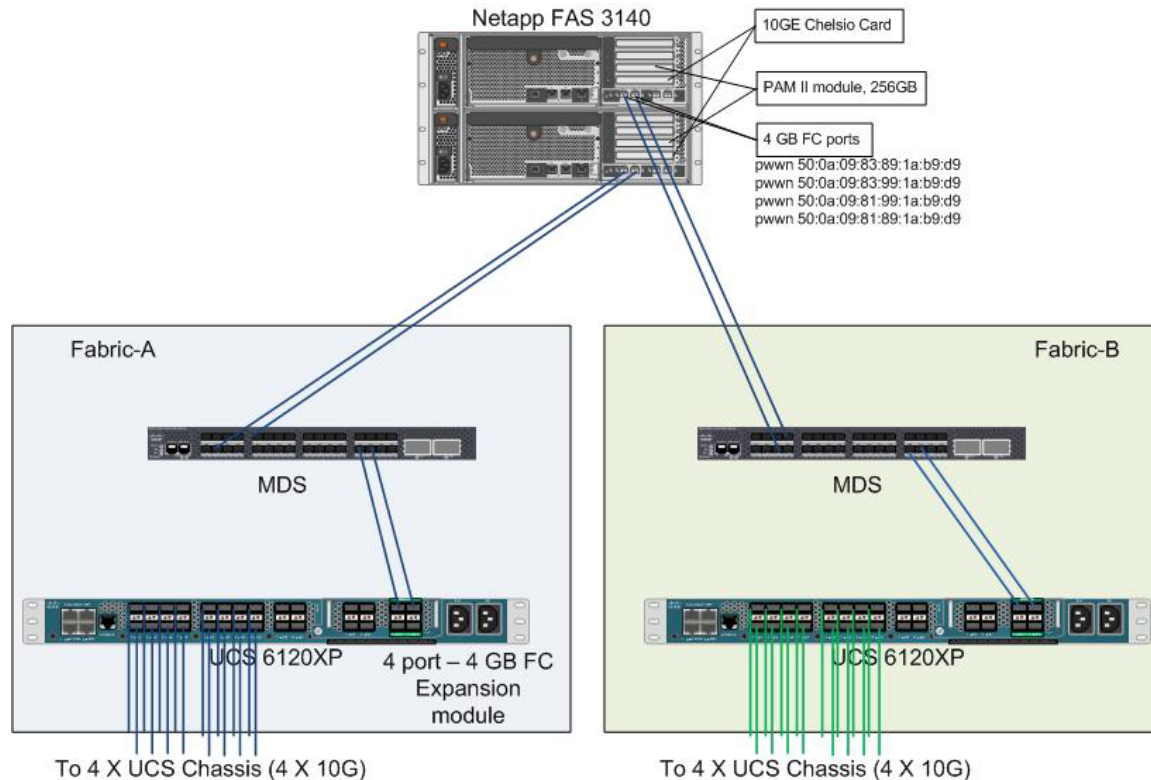
```
* fcid 0x470735 [pwwn 20:00:00:25:b5:0a:ad:2c]
```

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

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

The NetApp FC 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 43.

Figure 43. NetApp FC Target Ports



5.5.1 Boot from SAN

Booting from SAN is another key feature which helps in moving towards stateless computing in which there is no static binding between a physical server and the OS / applications it is 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 key benefits of booting from the network:

- **Reduce Server Footprints:** Boot from SAN alleviates the necessity for each server to have its own direct-attached disk, eliminating internal disks as a potential point of failure. Thin diskless servers also take up less facility space, require less power, and are generally less expensive because they have fewer hardware components.
- **Disaster and Server Failure Recovery:** All the boot information and production data stored on a local SAN can be replicated to a SAN at a remote disaster recovery site. If a disaster destroys functionality of the servers at the primary site, the remote site can take over with minimal downtime.
Recovery from server failures is simplified in a SAN environment. With the help of snapshots, mirrors of a failed server can be recovered quickly by booting from the original copy of its image. As a result, boot from SAN can greatly reduce the time required for server recovery.
- **High Availability:** A typical data center is highly redundant in nature - redundant paths, redundant disks and redundant storage controllers. When operating system images are stored on disks in the SAN, it supports high availability and eliminates the potential for mechanical failure of a local disk.
- **Rapid Redeployment:** Businesses that experience temporary high production workloads can take advantage of SAN technologies to clone the boot image and distribute the image to multiple servers for



rapid deployment. Such servers may only need to be in production for hours or days and can be readily removed when the production need has been met. Highly efficient deployment of boot images makes temporary server usage a cost effective endeavor.

- Centralized Image Management: When operating system images are stored on networked disks, all upgrades and fixes can be managed at a centralized location. Changes made to disks in a storage array are readily accessible by each server.

5.5.2 Configuring Boot from SAN on Cisco Unified Computing System

When you 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 FC 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 BIOS settings. Once the hardware detects the boot device, it follows the regular boot process.

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

There are three distinct portions of BFS procedure namely:

1. Storage array configuration
2. SAN zone configuration
3. Cisco UCS configuration of service profile

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 need 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 an initiator group and add them to the port WWNs where the storage is configured as shown below.

#	Task description
1.	Create a separate boot from SAN Aggregate
2.	Create a Volume on top of that, call it BootVolumes
3.	Add LUN on the BootVolumes, lets call it BFS-Server-9 and 50 GB of space

NetApp3140-A

- Filer
- Volumes
- Aggregates
- Storage
- Operations Manager
- SnapMirror
- CIFS
- NFS
- HTTP
- LUNs
 - Wizard
 - Enable/Disable
 - Manage
 - Add**
 - Show Statistics
 - LUN ConfigCheck
 - Initiator Groups
 - FCP
 - iSCSI
 - Network

Add LUN

LUNs → Add

[Manage LUNs]

Path:
The full path of the LUN, for example /vol/luns/lunOne. The LUN must be created in the root directory of a volume or a qtree.

LUN Protocol Type:
Select the multiprotocol type for the LUN.

Description:
An optional description of the LUN.

Size:
The size of the LUN.

Units:
A multiplier for the LUN size.

Space Reserved:
Indicates whether this LUN is space reserved.

Add

4. Now add the LUN to the initiator group

NetApp3140-A

- Filer
- Volumes
- Aggregates
- Storage
- Operations Manager
- SnapMirror
- CIFS
- NFS
- HTTP
- LUNs
- Wizard
- Enable/Disable
- Manage
- Add
- Show Statistics
- LUN ConfigCheck
- Initiator Groups**
 - Manage
 - Add**
 - FCP
 - iSCSI
 - Network
 - Security

Add Initiator Group

LUNs → Initiator Groups → Add

[Manage Initiator Groups]

Group Name:
Enter a group name for the initiator group

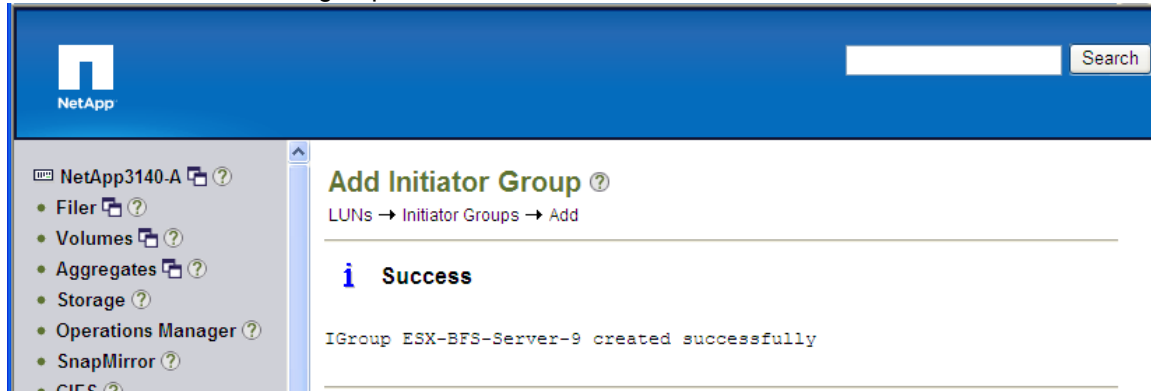
Type:
Select a Type for the initiator group

Operating System:
Select the operating system type of the initiators in this group

Initiators:
Enter a list of initiator names, separated by commas, spaces or newlines.
For an FCP initiator group, enter WWPNs.
For an iSCSI initiator group, enter iSCSI node names

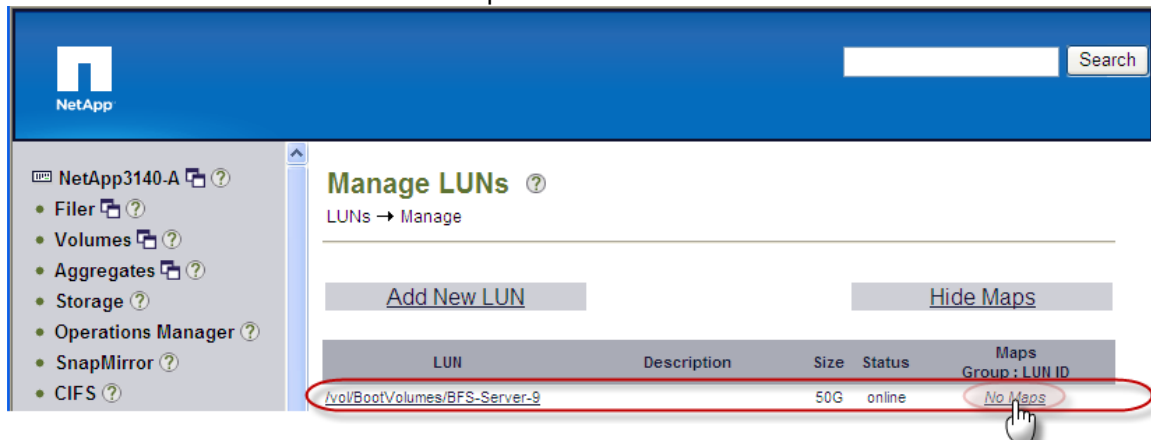
Add

5. Make sure the add initiator group succeeds



The screenshot shows the NetApp GUI for adding an initiator group. The left sidebar shows the navigation menu with 'NetApp3140-A' selected. The main content area is titled 'Add Initiator Group' and shows a success message: 'IGroup ESX-BFS-Server-9 created successfully'.

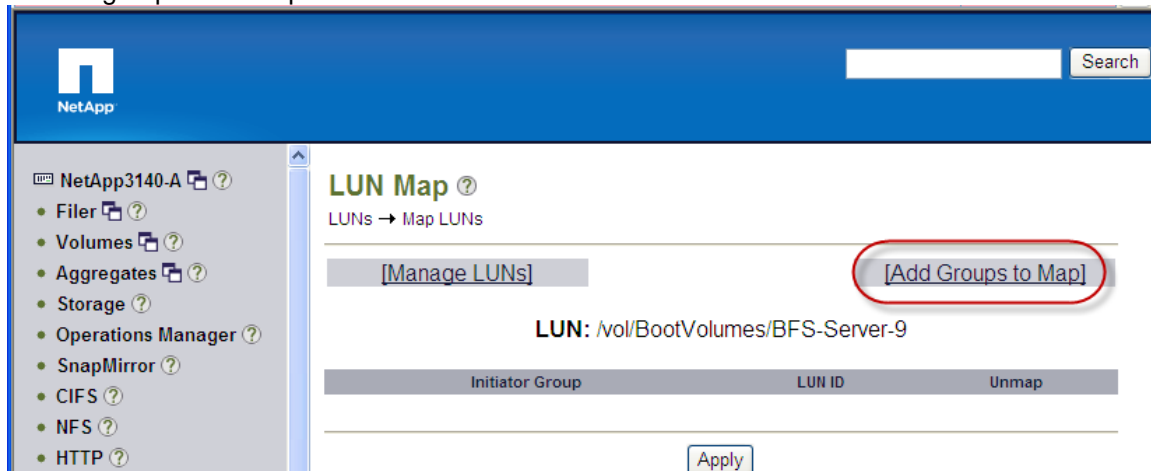
6. Now we need to mask the LUN, proceed to LUN > Manage LUN and select the new LUN which needs to be added and select the "no map" section as shown below.



The screenshot shows the NetApp GUI for managing LUNs. The left sidebar shows the navigation menu with 'NetApp3140-A' selected. The main content area is titled 'Manage LUNs' and shows a table of LUNs. The LUN '/vol/BootVolumes/BFS-Server-9' is selected, and the 'No Maps' button is highlighted.

LUN	Description	Size	Status	Maps Group : LUN ID
/vol/BootVolumes/BFS-Server-9		50G	online	No Maps

7. Add the group to the map



The screenshot shows the NetApp GUI for mapping LUNs. The left sidebar shows the navigation menu with 'NetApp3140-A' selected. The main content area is titled 'LUN Map' and shows the LUN '/vol/BootVolumes/BFS-Server-9'. The 'Add Groups to Map' button is highlighted.

8. Select the new initiator group – ESX-BFS-Server-9 and click add

9.	<p>Now assign a LUN ID to the initiator group</p>
10	<p>Make sure the mapping succeeded.</p>
11	<p>After the LUN map is successfully updated, check to see if the Manage LUNs show a correct mapping.</p>



NetApp3140-A

- File
- Volumes
- Aggregates
- Storage
- Operations Manager
- SnapMirror
- CIFS

Manage LUNs

LUNs → Manage

LUN	Description	Size	Status	Maps Group : LUN ID
/vol/BootVolumes/BFS-Server-9		50G	online	ESX-BFS-Server-9 : 0

12 Repeat the steps 3 through 11 for the number of servers you want to do boot from SAN.

5.5.3 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 61x0 XP fabric interconnect. The port mode is set to AUTO as well as the speed is set to AUTO. Rate mode is "dedicated" and when everything is configured correctly you should see some like the below picture on a MDS switch device manager for a given port (eg. Fc2/16). VSAN configuration can be done either in the SAN Switch CLI or the GUI, like our 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 upfront 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    --
```

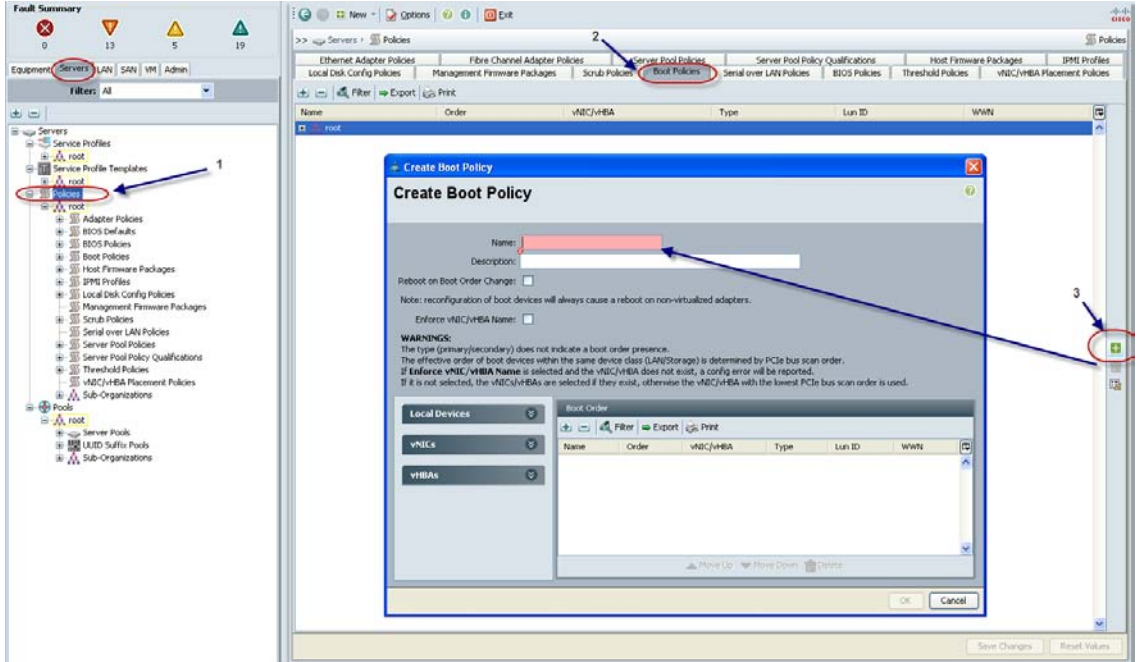
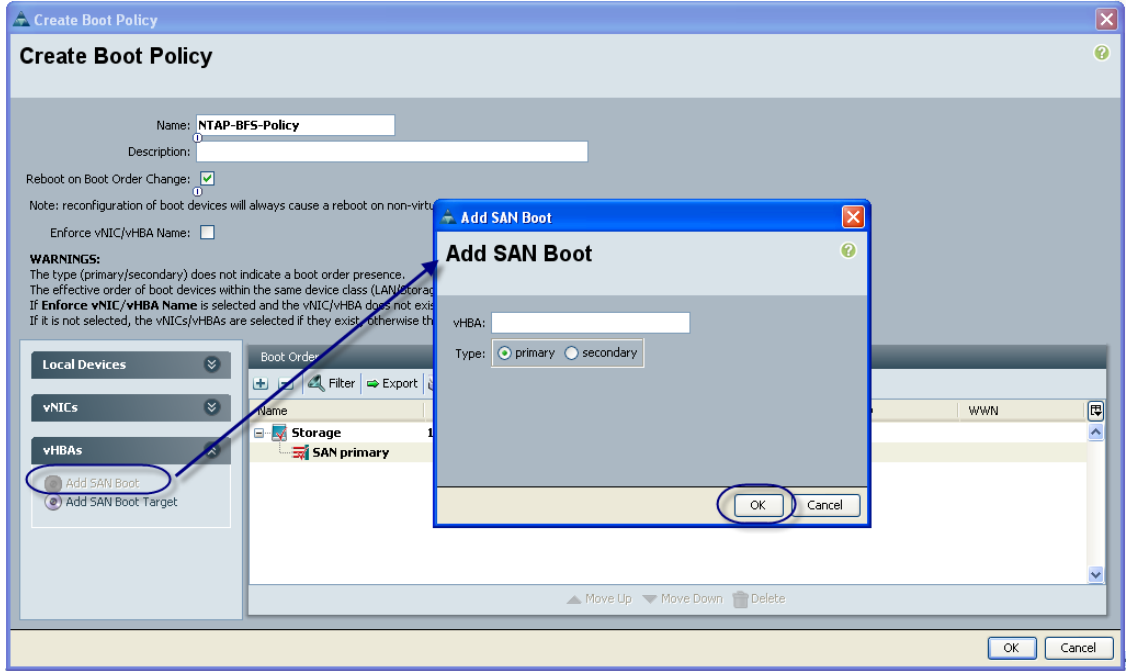
```
# sh int fc1/1 brief
```

```
-----
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    --
```

5.5.4 Cisco UCS Manager Configuration

To enable boot from SAN from a Cisco UCS M-Series perspective, perform the following steps:

Step #	Task description
1.	<p>Create a boot policy in the “Servers” tab. To do this, Select the policies and on the right plane select boot policies and select “Add” button. Enter name, select reboot on change, and do not select “enforce vHBA name”.</p> 
2.	<p>Add SAN Boot for primary. The vHBA is optional, it could be left blank and we do not have to enforce the vHBA name.</p> 
3.	<p>Add SAN boot for SAN Secondary</p>

Create Boot Policy

Name: **NTAP-BFS-Policy**

Description:

Reboot on Boot Order Change: ☒

Note: reconfiguration of boot devices will always cause a reboot on non-virtualized adapters.

Enforce vNIC/vHBA Name: ☐

WARNINGS:
 The type (primary/secondary) does not indicate a boot order presence.
 The effective order of boot devices within the same device class (LAN/Storage) is determined by PCIe bus scan order.
 If **Enforce vNIC/vHBA Name** is selected and the vNIC/vHBA does not exist, a config error will be reported.
 If it is not selected, the vNICs/vHBAs are selected if they exist, otherwise the vNIC/vHBA with the lowest PCIe bus scan order.

Local Devices

- vNICs
- vHBAs
 - Add SAN Boot
 - Add SAN Boot Target

Boot Order

Name	Order	vNIC/vHBA	Type	Lun ID	WWN
Storage	1				
SAN primary			primary		
SAN secondary			secondary		

Add SAN Boot

vHBA:

Type: ☐ primary ☐ secondary

OK Cancel

4. Now add Boot target WWPN to the SAN Primary, make sure this is exactly what 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
```

Create Boot Policy

Name: **NTAP-BFS-Policy**

Description:

Reboot on Boot Order Change: ☒

Note: reconfiguration of boot devices will always cause a reboot on non-virtualized adapters.

Enforce vNIC/vHBA Name: ☐

WARNINGS:
 The type (primary/secondary) does not indicate a boot order presence.
 The effective order of boot devices within the same device class (LAN/Storage) is determined by PCIe bus scan order.
 If **Enforce vNIC/vHBA Name** is selected and the vNIC/vHBA does not exist, a config error will be reported.
 If it is not selected, the vNICs/vHBAs are selected if they exist, otherwise the vNIC/vHBA with the lowest PCIe bus scan order.

Local Devices

- vNICs
- vHBAs
 - Add SAN Boot
 - Add SAN Boot Target
 - Add San Boot Target To SAN primary
 - Add San Boot Target To SAN secondary

Boot Order

Name	Order	vNIC/vHBA	Type	Lun ID	WWN
Storage	1				
SAN primary			primary		
SAN Target primary			primary	0	50:0A:09:81:89:8B:8A:81
SAN secondary			secondary		

Add SAN Boot Target

Boot Target LUN: 0

Boot Target WWPN: **50:0A:09:81:89:8B:8A:81**

Type: ☒ primary ☐ secondary

OK Cancel

- Repeat step 4 for SAN primary's SAN Target Secondary
- Repeat step 4 for SAN Secondary's – SAN Target Primary
- Repeat step 4 for SAN Secondary's – SAN Target Secondary
- At the end your Boot from SAN policy should look like:

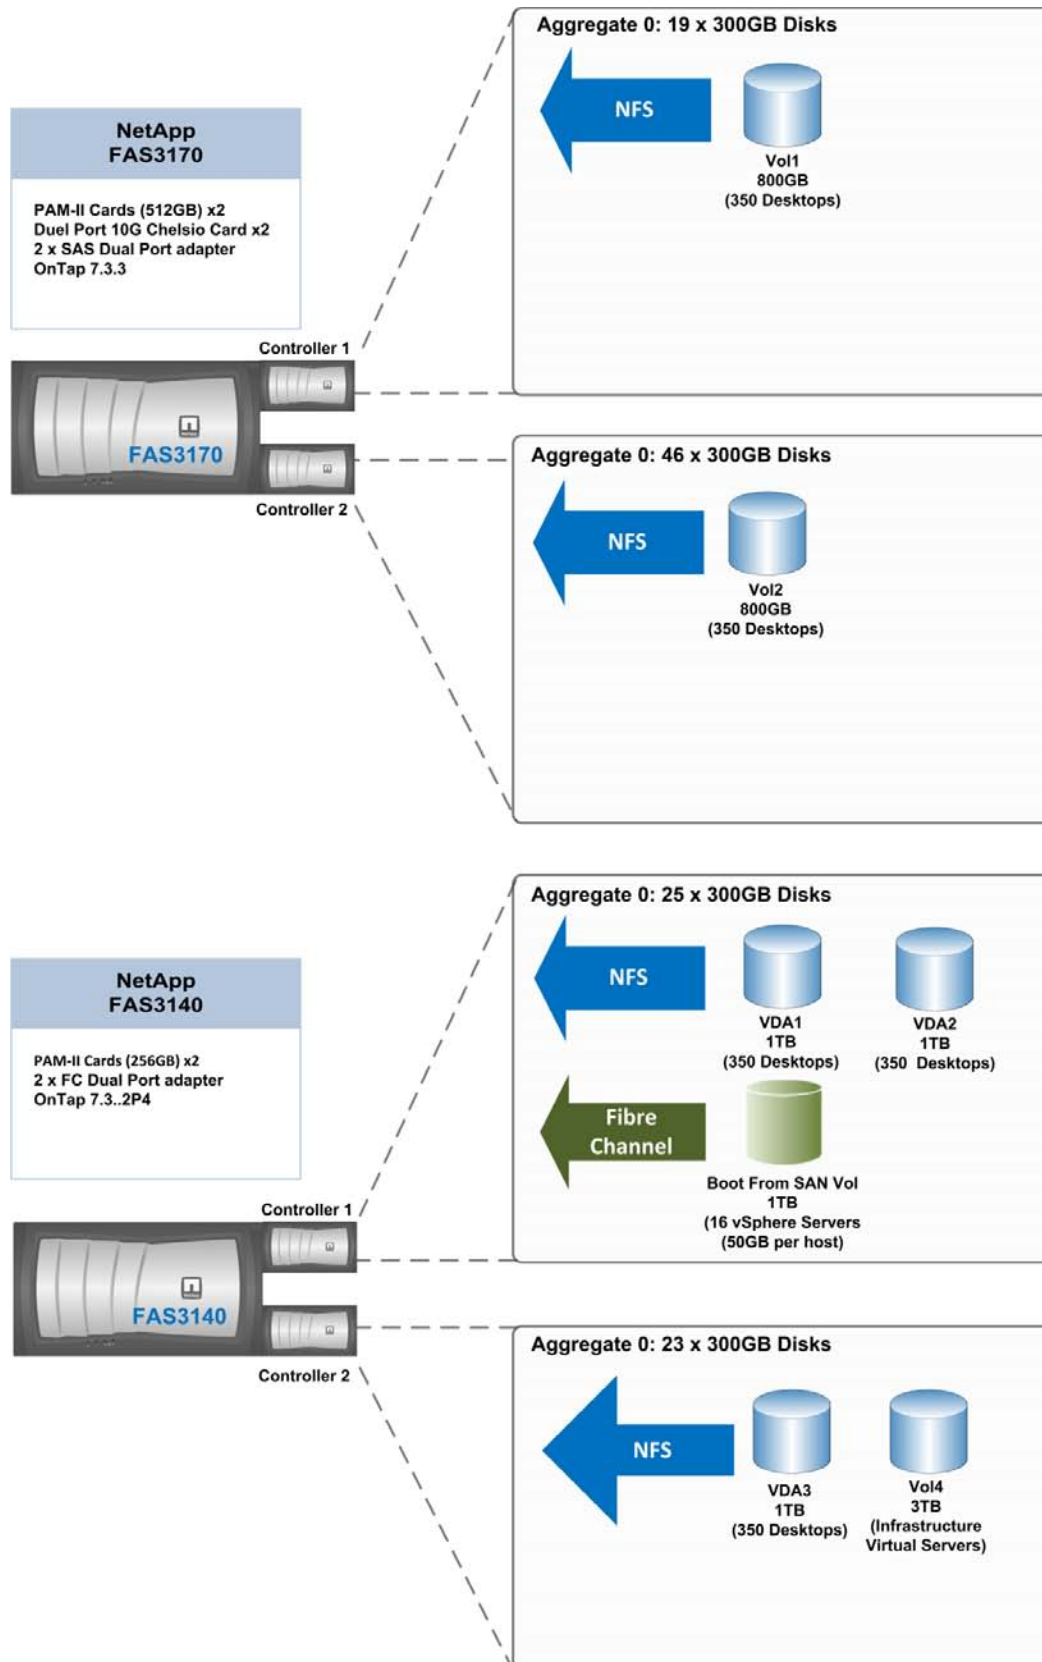
9.	<p>The last step is to make the association of the service profile template to the Boot from SAN policy during the service profile template configuration. You could also modify the Boot order as shown:</p>
10.	<p>This completes the BFS configuration on Cisco UCS-M. When the service profile is created out of the template, each server will be ready to boot from SAN provided the appropriate OS installation steps has taken place.</p>

5.6 NetApp Storage Configuration

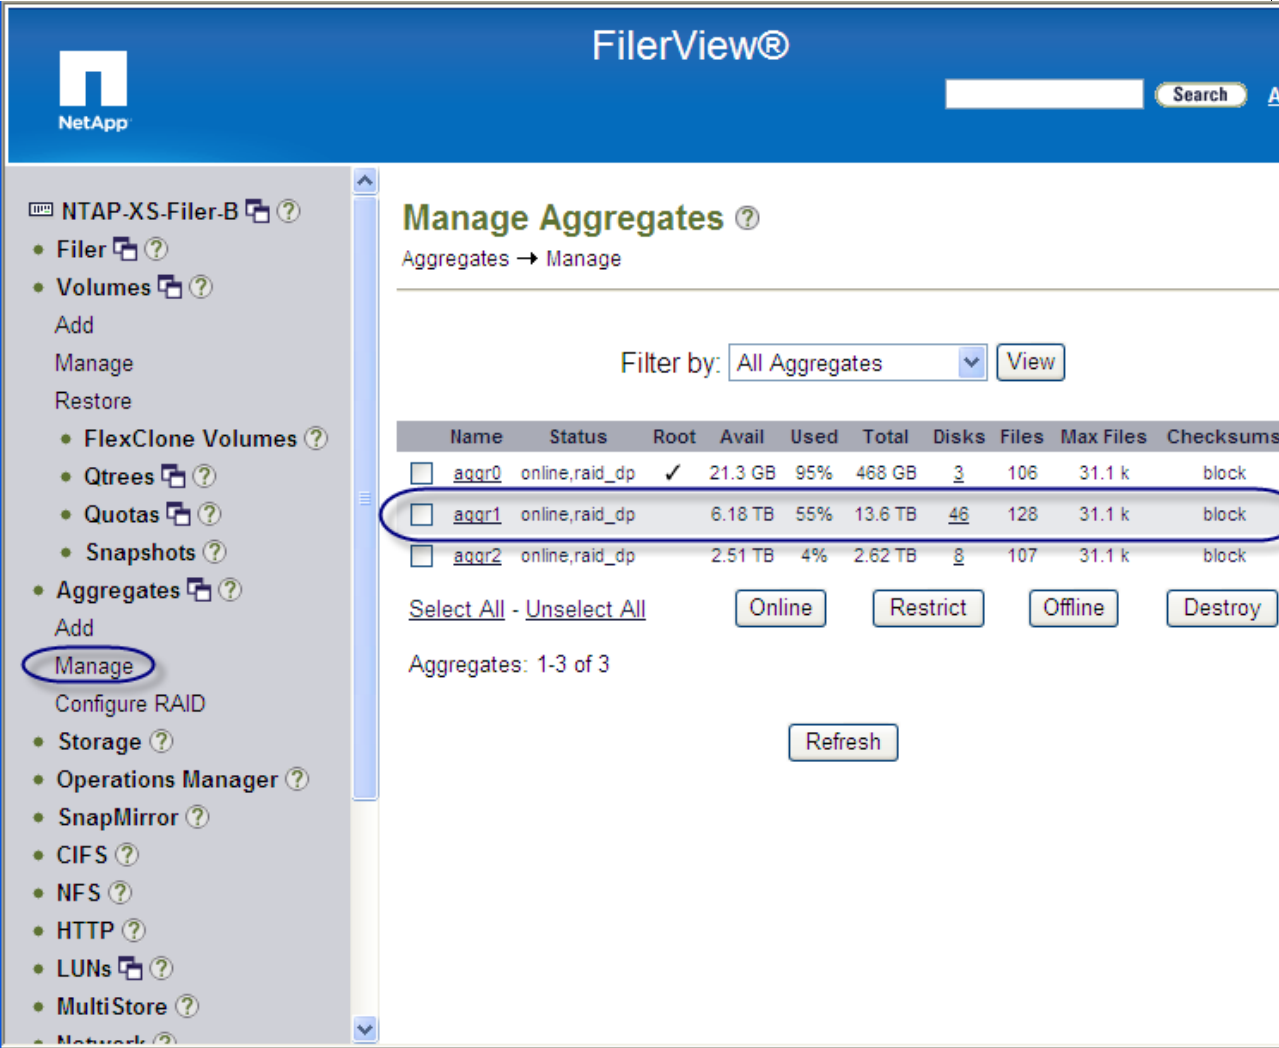
Two NetApp storage array (FAS3140 & FAS3170) were used for large scale 16 host testing scenarios:


- FAS3140 was used to host:
 - All infrastructure Virtual Servers on a single Volume provided by a single controller (2)
 - 3 VDA Volumes 2 from one Controller (1) and 1 from controller (2)
 - Boot volume for 8 vSphere 4.0U2 Hosts
- FAS3170 with PAM-II cards was used to host:
 - 2 VDA volumes, one from each controller

Figure 44. NetApp Storage Configuration



5.6.1 Example of a NetApp NFS Volume Configuration

Task #	Description
1.	Login to the NetApp storage using a web browser and click on filerView. It starts the NetApp filer configuration application.
2.	<p>Once in the filerview select the aggregates section and click add to create a large aggregate. It is best practice to create a large aggregate on top of all the system disks and carve volumes and luns from it. We created a large aggregate out of 46 disks and called it aggr1.</p>  <p>The screenshot displays the 'Manage Aggregates' interface in FilerView. The left-hand navigation pane lists various system components, with 'Aggregates' selected and 'Manage' highlighted. The main pane shows a table of existing aggregates. The aggregate 'aggr1' is highlighted, indicating it is the focus of the configuration. The table columns include Name, Status, Root, Avail, Used, Total, Disks, Files, Max Files, and Checksums. Below the table, there are action buttons: 'Select All - Unselect All', 'Online', 'Restrict', 'Offline', 'Destroy', and a 'Refresh' button. The status of 'aggr1' is 'online,raid_dp' and it consists of 46 disks.</p>
3.	Now from the volumes section, select add to add a volume. An add volume wizard pop up.



FilerView®

NTAP-XS-Filer-B

- Filer
- Volumes
- Add**
- Manage
- Restore
- FlexClone Volumes
- Qtrees
- Quotas
- Snapshots
- Aggregates
- Add
- Manage
- Configure RAID
- Storage
- Operations Manager
- SnapMirror
- CIFS
- NFS
- HTTP
- LUNs
- MultiStore
- Network

Manage Volumes

Volumes → Manage

Filter by: All Volumes

Name	Status	Root	Containing FlexClone	Avail	Used	Total	Files	Max File
Aggregate								
<input type="checkbox"/> BackupVol	online,raid_dp	aggr2	-	108 GB	0%	108 GB	102	3.74 m
<input type="checkbox"/> NewVolume2	online,raid_dp	aggr1	-	1.2 TB	11%	1.35 TB	6.28 k	31.9 m

Volume Wizard

Welcome to the volume Storage wizard. This wizard can be used to create, adjust, and mirror volumes.

You have chosen to

- Add a new volume

4. Select flexible volume for the volume type

Volume Wizard - Volume Type Selection

Volume Type Selection

Select whether you want to create a traditional, flexible, or cache volume.

☒ Flexible ?

☐ Traditional

☐ Cache

< Back

Cancel

Next >

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

Volume Wizard - Volume Parameters

Volume Name:

Enter a name for the new volume.

XD_VDA_VOLUME_1 ?

Enter Volume name

Language:

Select the language to use on this volume.

POSIX

UTF-8:

Select to make language of this volume UTF-8 encoded.

☐ UTF-8 ?

SnapLock Volume:

Select to create a snaplock volume.

☐ snaplock ?

< Back

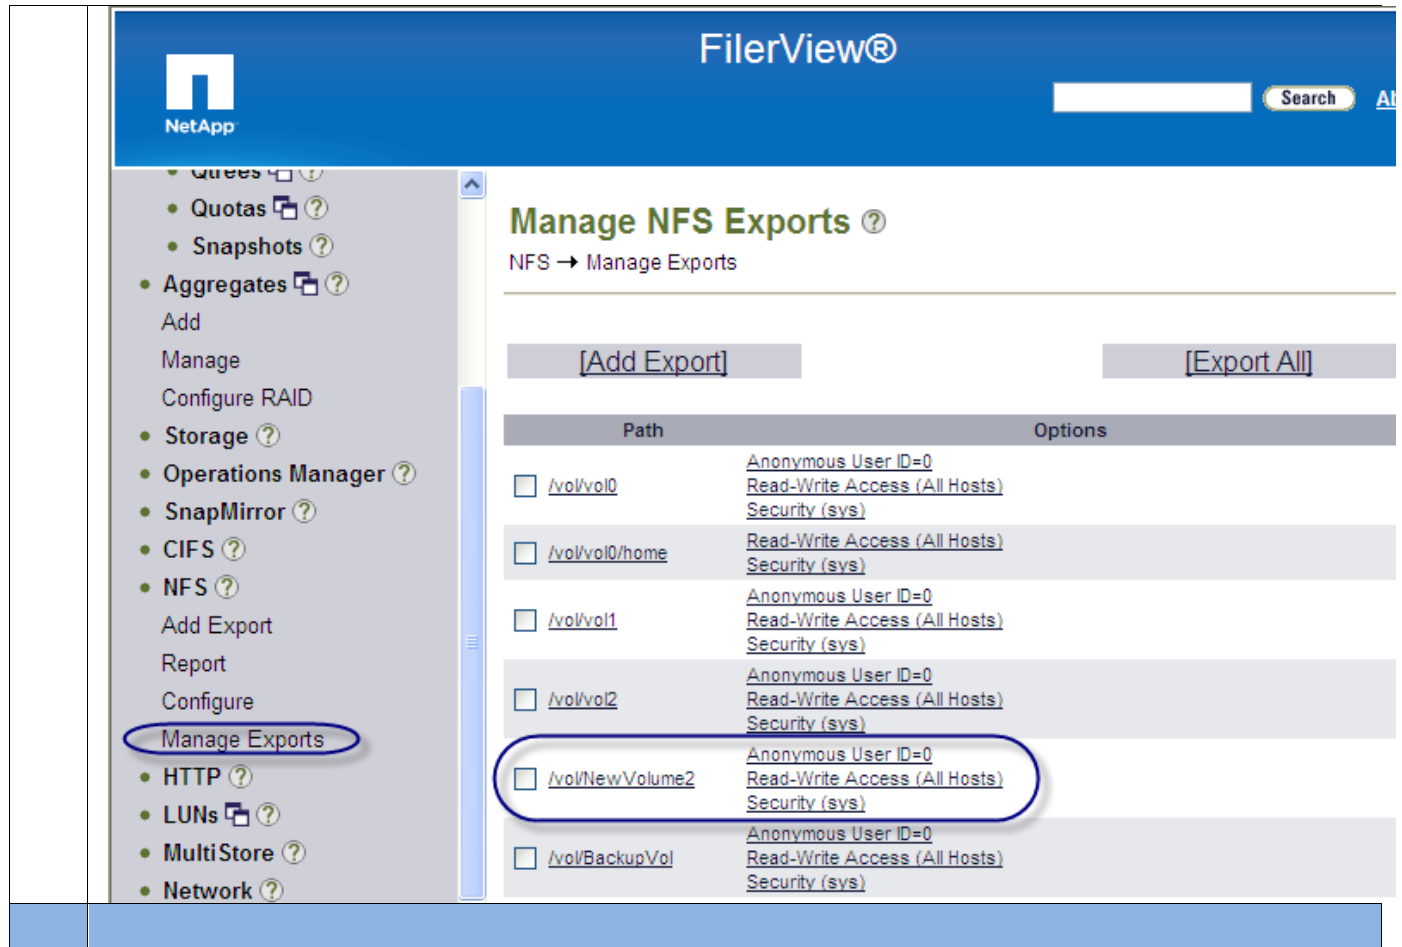
Cancel

Next >

6. Select the aggregate to contain this volume, set space guarantee as none for thin provisioning.

	<div data-bbox="297 237 1417 940"> <h3>Volume Wizard - Flexible Volume Parameters</h3> <div> Containing Aggregate Select the aggregate to contain this volume. Only non-snaplock aggregates are displayed. <div>aggr1 (6.18 TB, raid_dp) ?</div> </div> <div> Space Guarantee Sets the space guarantee. Volume guarantees space for the entire volume in the containing aggregate; File guarantees space for a file at file allocation time; None reserves no extra space for the volume. <div>none ?</div> </div> <div> <div>< Back</div> <div>Cancel</div> <div>Next ></div> </div> </div> <div data-bbox="1404 300 1573 562"> <p>1. Select the aggregate from the drop down menu</p> <p>2. Select none for thin provisioning</p> </div>
7.	<div data-bbox="297 1150 1417 1854"> <h3>Volume Wizard - Flexible Volume Size</h3> <div> Volume Size Type: Select Total Size to enter the total volume size (including snapshot reserve) and Usable Size to enter the usable volume size (excluding snapshot reserve). <div> <input checked="" type="radio"/> Total Size ? <input type="radio"/> Usable Size </div> </div> <div> Volume Size: Enter the desired volume size. The containing aggregate, aggr1 has a maximum of 6.18 TB space available. <div>800 GB ?</div> </div> <div> Snapshot Reserve : Enter the snapshot reserve for volume 'XD_VDA_VOLUME_1'. The range is between 0% and 100%. The default is 20%. <div>0 % ?</div> </div> <div> <div>< Back</div> <div>Cancel</div> <div>Next ></div> </div> </div>
8.	We are all done, press commit.

	
<p>9.</p>	<p>After the volume is added, go to the NFS section, manage exports and click on the volume and configure the volume to give R/W access to hosts, set user ID access, security. After that click on “add export” to make it available to all host. You could also do host based access control instead of all hosts and set root access. For example:</p>



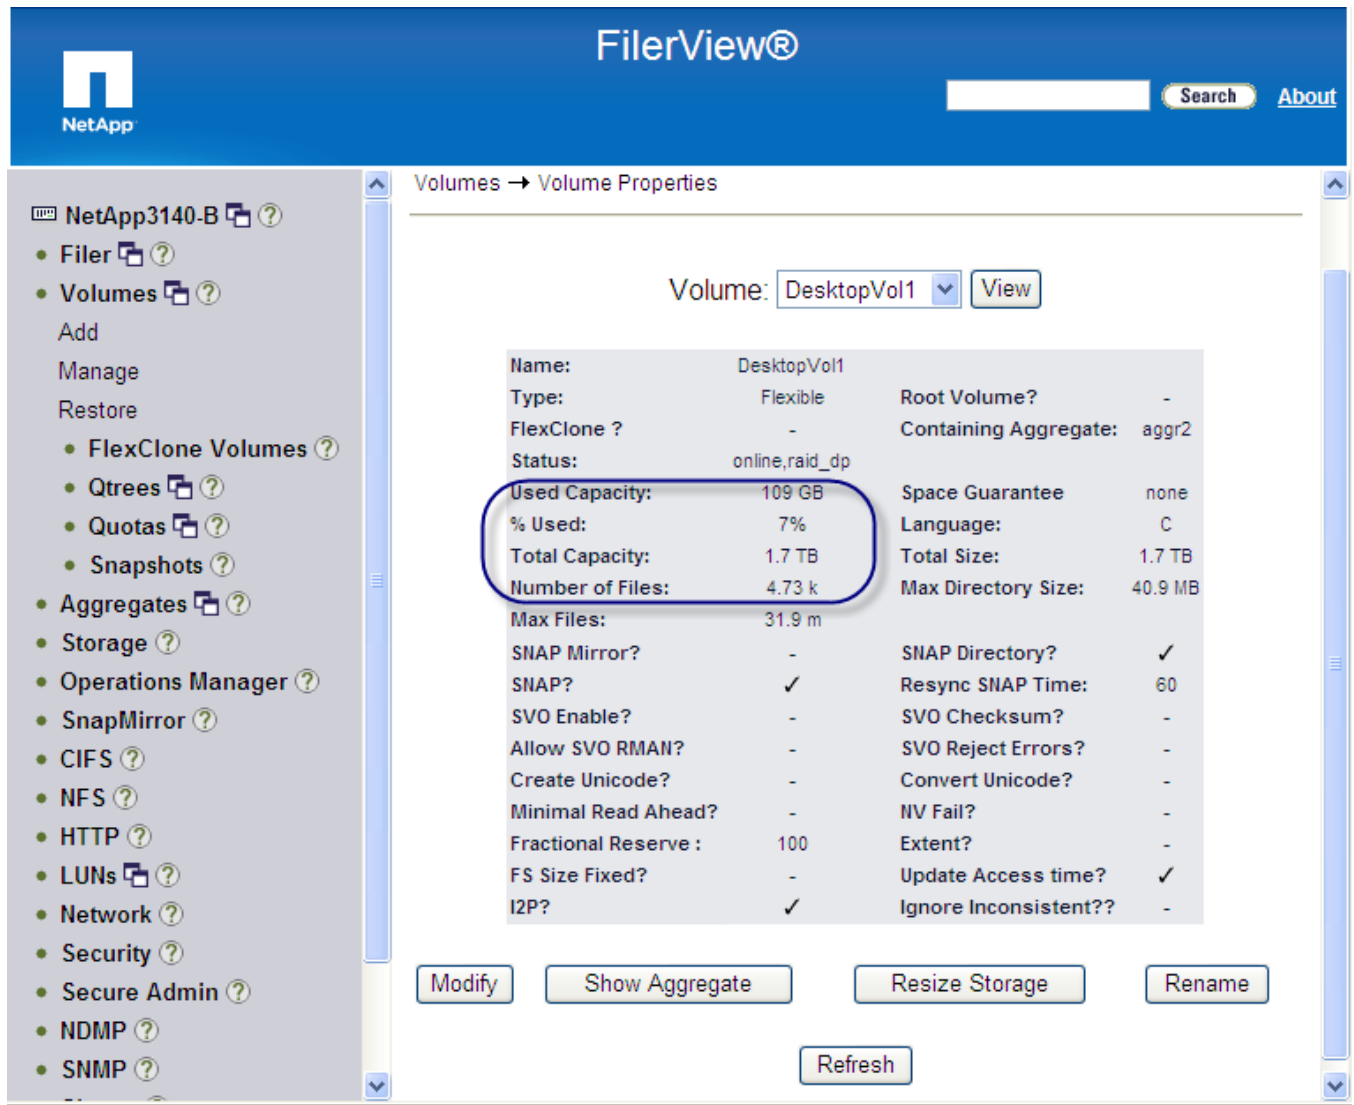
The screenshot shows the NetApp FilerView web interface. The left sidebar contains a navigation menu with items like Quotas, Snapshots, Aggregates, Storage, Operations Manager, SnapMirror, CIFS, NFS, HTTP, LUNs, Multi Store, and Network. The 'Manage Exports' option under NFS is highlighted with a blue oval. The main content area is titled 'Manage NFS Exports' and includes buttons for '[Add Export]' and '[Export All]'. Below these is a table with two columns: 'Path' and 'Options'. The table lists several NFS exports, with the entry for '/vol/NewVolume2' highlighted by a blue oval. The options for this entry are 'Anonymous User ID=0', 'Read-Write Access (All Hosts)', and 'Security (sys)'.

Path	Options
<input type="checkbox"/> /vol/vol0	Anonymous User ID=0 Read-Write Access (All Hosts) Security (sys)
<input type="checkbox"/> /vol/vol0/home	Read-Write Access (All Hosts) Security (sys)
<input type="checkbox"/> /vol/vol1	Anonymous User ID=0 Read-Write Access (All Hosts) Security (sys)
<input type="checkbox"/> /vol/vol2	Anonymous User ID=0 Read-Write Access (All Hosts) Security (sys)
<input type="checkbox"/> /vol/NewVolume2	Anonymous User ID=0 Read-Write Access (All Hosts) Security (sys)
<input type="checkbox"/> /vol/BackupVol	Anonymous User ID=0 Read-Write Access (All Hosts) Security (sys)

5.6.2 NetApp Deduplication in Practice

As described in section 3.5.4, NetApp deduplication saves space on primary storage by removing redundant copies of blocks within a volume hosting hundreds of virtual desktops. An example is shown in Figure 45 for 800 GB volume hosting 428 desktops each with 3 GB capacity.

Figure 45. NetApp Deduplication



5.7 VMware vSphere Configuration

This section details the vSphere configuration and tuning that was done as part of the testing.

The software components details are as follows:

VMware vSphere 4.0U2 Host Build 261974			
Hardware:	Cisco B-Series blade servers	Model:	B250 –M2
OS:	Both VMware vSphere 4.0U2	Service Pack:	-
CPU:	2 x 6 Core Westmere or Intel Xeon 5680, 3.33 GHz (24 Logical Cores Total)	RAM:	196GB
Disk:	Boot From SAN	Network:	4 x 10GbE

VMware vSphere 4.0iU2 Host (ESX4i Installable) Build 261974			
Hardware:	Cisco B-Series blade servers	Model:	B250 –M2
OS:	vSphere 4.0iU2 installable	Service Pack:	-
CPU:	2 x 6 Core Westmere or 5680, 3.33 GHz (24 Logical Cores Total)	RAM:	196GB
Disk:	Local disk install	Network:	4 x 10GbE

VMware vSphere vCenter 4.0U2			
Hardware:	Running in a virtual environment on top of Cisco B-Series blade servers B200 –M2	Model:	B200-M2
OS:	Windows Server 2003 Enterprise Ed. - SP2	Service Pack:	-
CPU:	4 vCPU	RAM:	12 GB
Disk:	C:\ - 50 GB; D:\ - 80 GB	Network:	1 x 10GbE
Network Driver:	VMXNET3	VM Version:	7

One of the goals we set out to test in this exercise was to virtualize the entire infrastructure components including the Citrix XenDesktop management services. We accomplished that with absolutely zero bare metal install of any operating system. All infrastructure components were in a virtual machine including the vCenter.

5.8 OS Installation

We used PXE server to do install of the ESX on the Cisco UCS blade servers. We were using both ESX4 and ESX4i, the ESX4 systems were configured to boot from SAN and the ESX4i was installed on the local SAS drive.

5.8.1 Center Configuration for DDC

XenDesktop 4 requires to enable the Virtual Centre SDK via http. You have to update the proxy.xml file (located on the vcenter server - C:\Documents and Settings\All Users\Application Data\VMware\VMware VirtualCenter\proxy.xml) to allow the SDK via http. Set the accessMode for /sdk to httpAndHttps

```
<e id="5">
<_type>vim.ProxyService.LocalServiceSpec</_type>
<accessMode>httpAndHttps</accessMode>
<port>8085</port>
<serverNamespace>/sdk</serverNamespace>
</e>
```

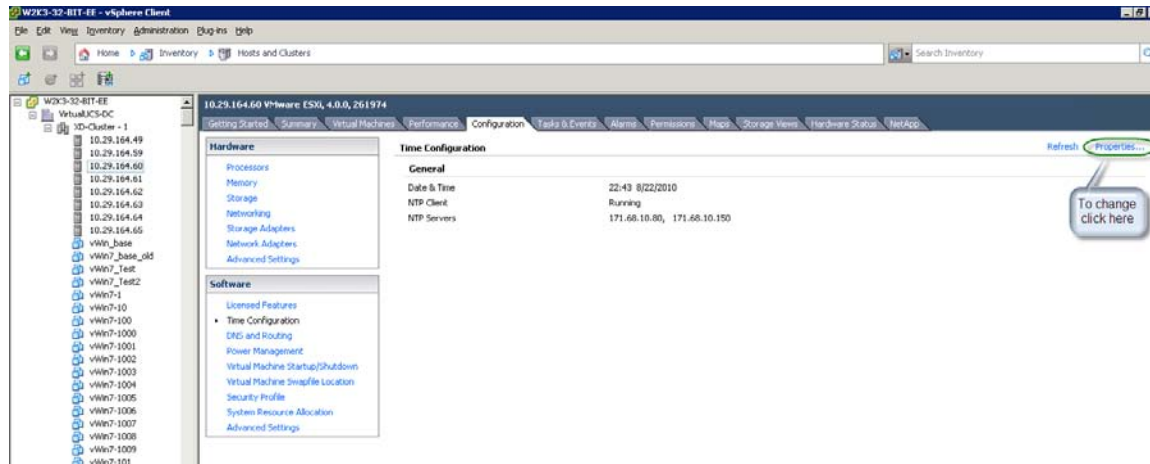
Restart VC Service.

This will take care of the DDC handshake with vCenter.

5.8.2 Setup NTP Server in VMware vCenter

One of the important aspects of running benchmark in a virtualized environment is configuring and setting up a NTP server and configuring from the vCenter for each servers. This is important from a time lag perspective and maintains synchronization of performance data collected across various components.

Figure 46. Configuring an NTP Server



5.8.3 VMware vMA or Virtual Management Assistant

The VMware vSphere Management Assistant (vMA) is a prebuilt Linux virtual machine in which administrators can deploy scripts to manage ESX and ESXi systems. Software included in vMA includes vSphere CLI, an authentication component that supports non-interactive login, and a log collection component. We used this extensively to collect resxtop output from the ESX4i host. This appliance is also used for making configuration changes and the ease of using it very apparent from the time saved in managing 16 hosts.

5.8.4 ESX Tuning

Apart from the NetApp best practices for using NFS storage on vSphere, there was no other tuning that was done on the ESX servers. As described in the NetApp VSC section VSC2.0 was installed and the recommendations were applied to all the servers. The following was the affect of those settings:

NFS performance tuning for vSphere using NetApp VSC2.0.

"QFullSampleSize" = 32, Old Value: 0

"TcpipHeapSize" = 30, Old Value: 16

"MaxVolumes" = 64, Old Value: 8

"HeartbeatMaxFailures" = 10, Old Value: 3

"TcpipHeapMax" = 120, Old Value: 64

"HeartbeatFrequency" = 12, Old Value: 9

Miscellaneous:

Verified the HaltingIdleMsecPenalty value was set to the default which is 100msec.

```
# esxcfg-advcfg -g /Cpu/HaltingIdleMsecPenalty
```

Value of HaltingIdleMsecPenalty is 100



5.8.5 MBR Tool from NetApp

MBR tool is an ESX console-based tool tests and aligns guest file systems on a VMDK for VMFS and NFS datastores. As aligning the file system block boundaries to the underlying NetApp storage system LUN help ensures the best storage performance we gave it a try. We downloaded this tool and ran it from the ESX console and all the VMDKs of the desktops were aligned. For example,

```
# ./mbralign --scan all
```

```
Building file list...
```

```
-----
```

```
/vmfs/volumes/747c046f-1fd1c2af/vWin7-6/vWin7-6-flat.vmdk p1 (NTFS) lba:128 offset:65536 aligned:Yes
```

```
-----
```

```
/vmfs/volumes/747c046f-1fd1c2af/vWin7-1281/vWin7-1281-flat.vmdk p1 (NTFS) lba:128 offset:65536  
aligned:Yes
```

```
-----
```

```
/vmfs/volumes/747c046f-1fd1c2af/vWin7-1282/vWin7-1282-flat.vmdk p1 (NTFS) lba:128 offset:65536  
aligned:Yes
```

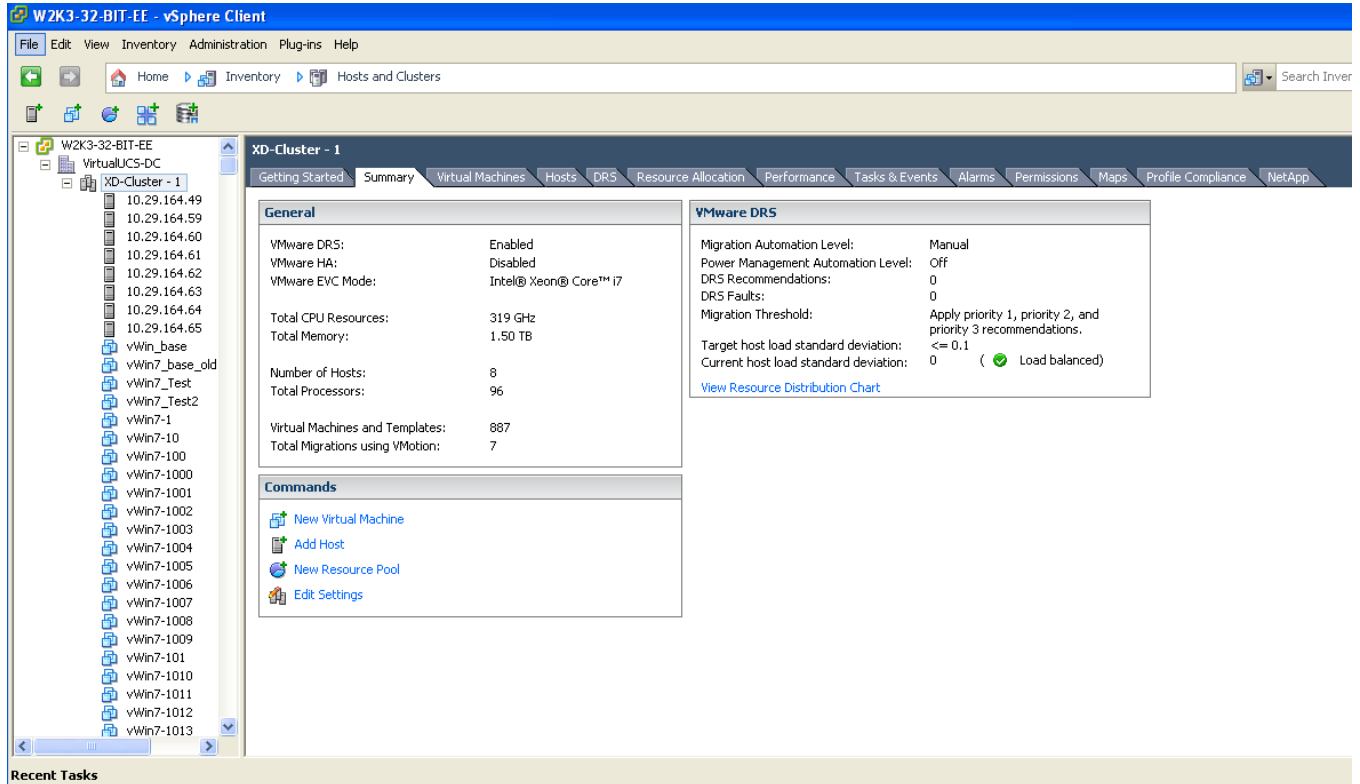
```
-----
```

5.8.6 VMware vSphere Configuration

For single server scale testing, one ESX 4.0U2 server was configured with boot from SAN and above mentioned tuning was applied. One NetApp storage volume was configured for this test.

For two chassis testing a cluster was created and eight servers were made part of that cluster with DRS mode set to manual. Four NAS devices were mounted on the eight servers as NFS mounts and the launcher VMs were used to generate the load to the desktops on the eight servers. Each blade was assigned 110 desktops each for a total of 880 desktops. 220 desktops were distributed on the four volumes.

Figure 47. VMware vSphere Configuration



For the 16 server tests, two clusters were created and cluster-1 had all the eight ESX 4.0U2 and cluster-2 had eight ESX 4.0iU2 installed. A mixture of Cisco UCS VIC as well as Cisco UCS M71KR-Q was used for this testing. An additional nfs volume was added to the setup to lower the load on the FAS 3140. Note the drivers for all the CNA cards used were all built in on ESX4.0U2.

5.8.7 Network configuration

As the configuration consisted of a mixture of cards, Cisco UCS M71KR-Q as well as Cisco VIC (M81KR), different variety of network configuration was created on the ESX host.

For Cisco VIC, we had created four static vNIC and two FC vHBA per card for a total of 8 vNIC and 4 FC ports. The Qlogic CNA on the other hand come with 2 FC ports and 2 vNICS per CNA for a total of 4 port each of FC and vNIC.

Different VLANs were used to separate the desktop traffic from the management traffic, NetApp storage traffic and the vmotion traffic. A simple vswitch configuration for four vnics is done with two active ports and two standby ports. Care should be taken to make sure that the active ports are paired in such a way that they are either part of the Fabric A or Fabric B. You cannot have active-active traffic between Fabric A side ports and Fabric B side ports.

For example on a ESX4i host, the following configuration was done:

The screenshot shows the VMware ESXi 4.0.0 configuration interface. The top navigation bar includes tabs for Getting Started, Summary, Virtual Machines, Performance, Configuration, Tasks & Events, Alarms, Permissions, Maps, Hardware Status, and NetApp. The left sidebar has a 'Hardware' section with 'Networking' highlighted, and a 'Software' section. The main content area shows the 'Virtual Switch: vSwitch0' configuration. It lists several Virtual Machine Port Groups: DesktopPrivNet (VLAN ID: 121), VM Data (VLAN ID: 165), VM-Network (VLAN ID: 164), and Management Network (VLAN ID: 164). On the right, a table lists Physical Adapters: vmnic3 (stand by), vmnic2 (10000 Full), vmnic1 (stand by), and vmnic0 (10000 Full). Buttons for 'Remove...' and 'Properties...' are visible.

The ESX4i console shows the different vmnics with their mac addresses.

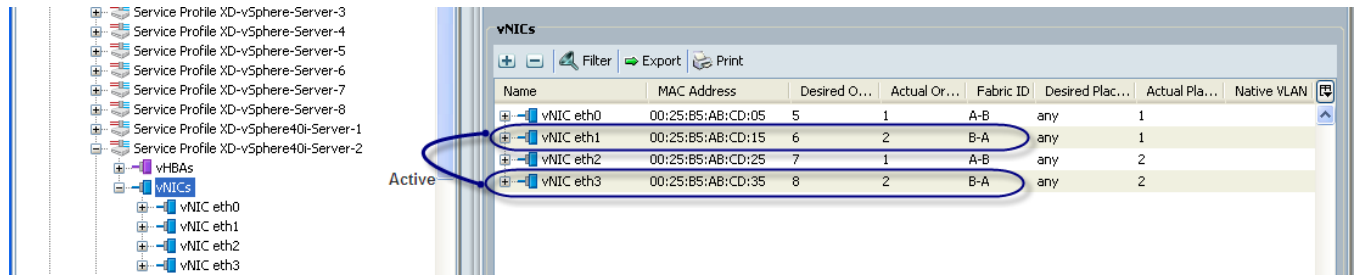
The screenshot shows the ESXi console output for the 'Configure Management Network' command. The output is divided into two main sections: 'Network Adapters' and 'Network Adapters' (repeated). The first section lists the network adapters and their MAC addresses: vmnic0 (00:25:b5:ab:cd:15), vmnic2 (00:25:b5:ab:cd:35), vmnic1 (00:25:b5:ab:cd:05), and vmnic3 (00:25:b5:ab:cd:25). The second section provides a detailed explanation of the adapters and their configuration. The bottom of the console shows the VMware ESXi 4.0.0 Releasebuild-261974 version.

Configure Management Network	Network Adapters
Network Adapters VLAN (optional) IP Configuration IPv6 Configuration DNS Configuration Custom DNS Suffixes	Network Adapters vmnic0 (00:25:b5:ab:cd:15) vmnic2 (00:25:b5:ab:cd:35) vmnic1 (00:25:b5:ab:cd:05) vmnic3 (00:25:b5:ab:cd:25) The adapters listed here provide the default network connection to and from this host. When two or more adapters are used, connections will be fault-tolerant and outgoing traffic will be load-balanced.
<Up/Down> Select	<Enter> Change <Esc> Exit

VMware ESXi 4.0.0 Releasebuild-261974

The VLAN was configured to be the native VLAN on the host that is 164 in our case.

The mac addresses as per Cisco UCS Manager maps to either Fabric-A or Fabric-B.



Name	MAC Address	Desired O...	Actual Or...	Fabric ID	Desired Plac...	Actual Pla...	Native VLAN
vNIC eth0	00:25:B5:AB:CD:05	5	1	A-B	any	1	
vNIC eth1	00:25:B5:AB:CD:15	6	2	B-A	any	1	
vNIC eth2	00:25:B5:AB:CD:25	7	1	A-B	any	2	
vNIC eth3	00:25:B5:AB:CD:35	8	2	B-A	any	2	

The teaming is done such a way that the vNIC on Fabric B (eth1 and eth3) are active and the standby is on Fabric A (eth0 and eth2).

We recommend using Cisco Nexus 1000v for network configuration as it provides advance capability to do DHCP snooping and other smart network switch capabilities in the hypervisor itself. These features have huge benefit in a virtual desktop environment where a vSwitch wouldn't be able to provide such features.

About mounting the NFS datastore: NFS datastores requires a vmkernel port that can connect to the NetApp storage or the NFS server. We use a different VLAN to separate the traffic from other desktop and management network traffic.

After the NetApp is configured with the volume and is exported to all hosts, we need to add the NFS datastore to ESX server. This is easily accomplished by using virtual center, select host, go to the configuration tab, select storage, use the add storage wizard and select Network File system. If everything is configured correctly, after the IP address of the storage array, Datastore name and the volume to mount is entered, the NetApp volume is mounted on that host.

As this task has to be repeated over the 16 servers, it is best to use vMA to do this. The command line to add nfs volumes is:

```
vicfg-nas -a -o 10.0.9.10 -s /vol/NewVolume1 -h $VI_SERVER NTAPVol1
```

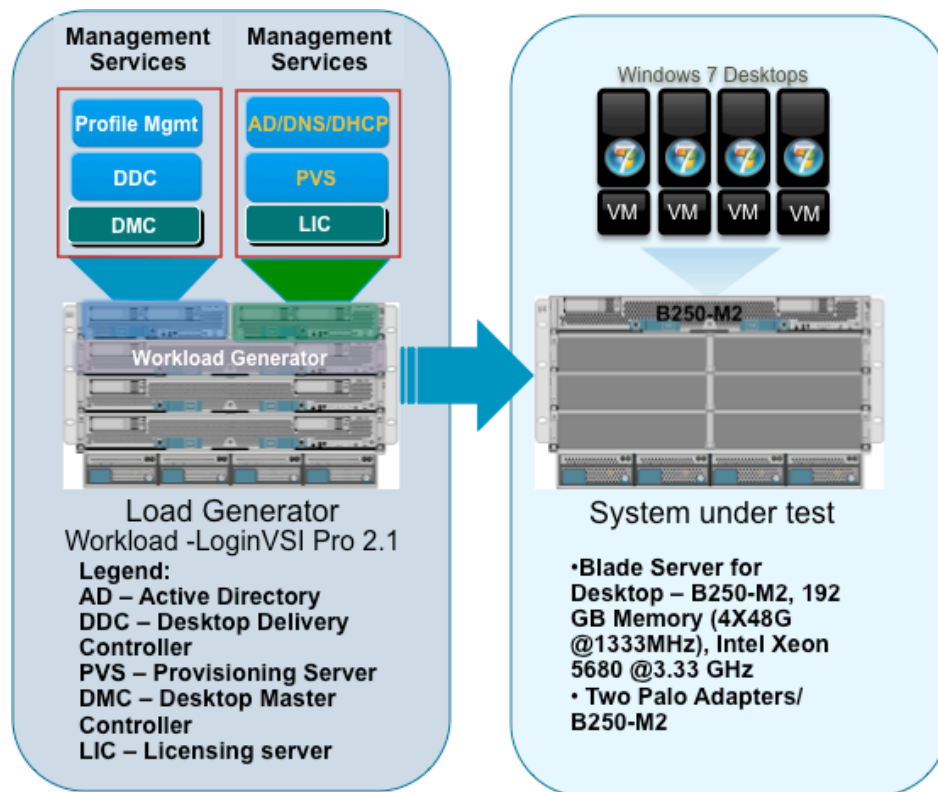
where VI_SERVER can be all the servers we want to mount the volume on.

You could write a small script to iterate on all the available servers and mount the nfs datastore on all ESX hosts.

6.0 Test Setup and Configurations

6.1 Cisco UCS Test Configuration for Single Server Scalability Test Setup

Figure 48. Cisco UCS B250 M2 Blade Servers for Single Server Scalability



Hardware components

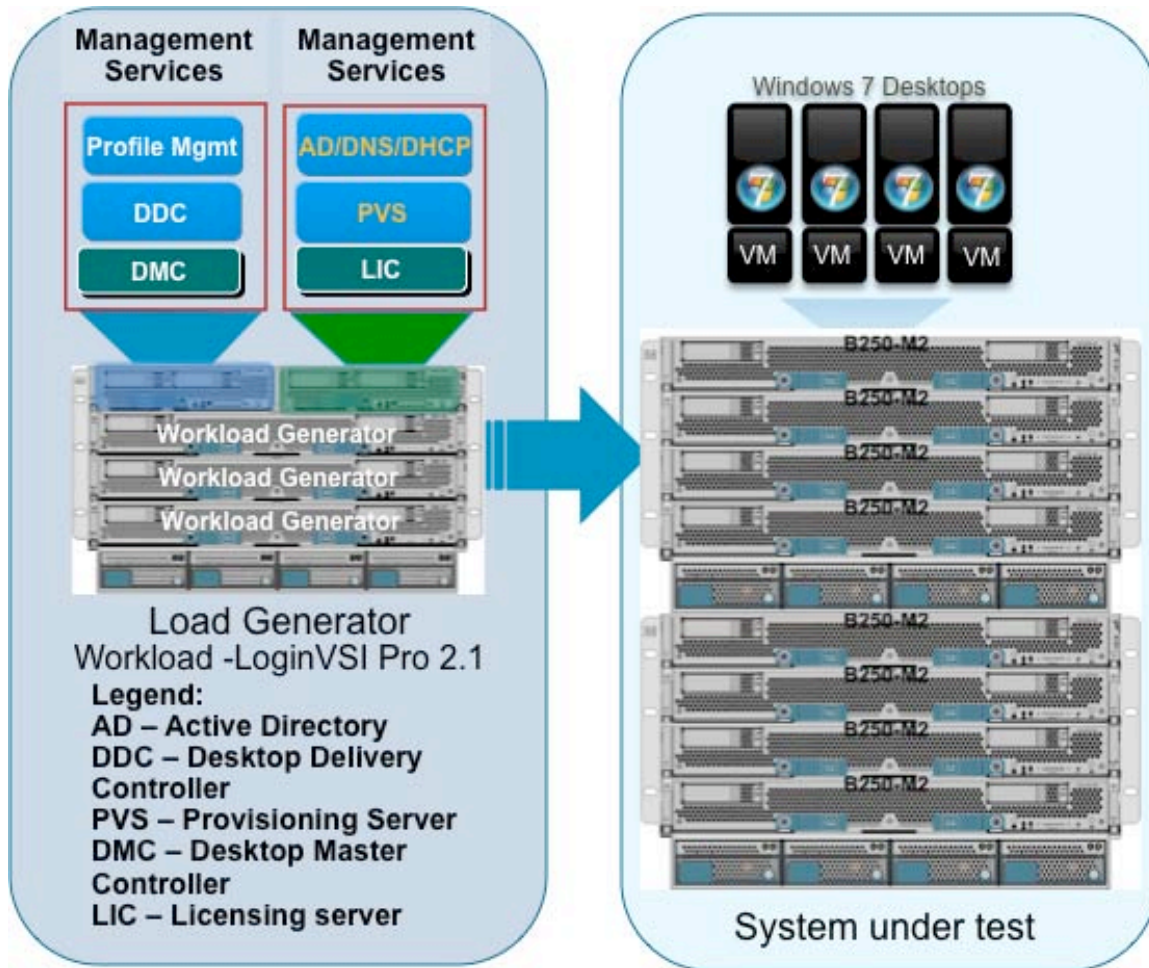
- 1 X Cisco UCS B250-M2 (5680 @ 3.33 GHz) blade servers with 192 GB of memory (4 GB X 48 DIMMS @ 1333 MHz)
- 2 X Cisco UCS B200-M2 (5680 @ 3.33 GHz) blade servers with 48 GB of memory (4 GB X 12 DIMMS @ 1333 MHz)
- Two Cisco Virtual Interface Card Adapters
- Cisco Nexus 5000 and 7000
- NetApp FAS 3140 storage array, two controllers, 2 X Dual port 10G Chelsio cards with 70 SAS drives

Software components

- Cisco UCS firmware 1.3(1i)
- VMware vSphere 4.0 U2, Virtual Center 4.0U2
- XenDesktop 4
- Windows 7 – 32 bit, 1vCPU, 1.5 GB of memory, 30 GB/VM

6.2 Cisco UCS Configuration for Two-Chassis Test

Figure 49. Two-Chassis Test Configuration-8 x Cisco UCS B250 Blade Server



Hardware components

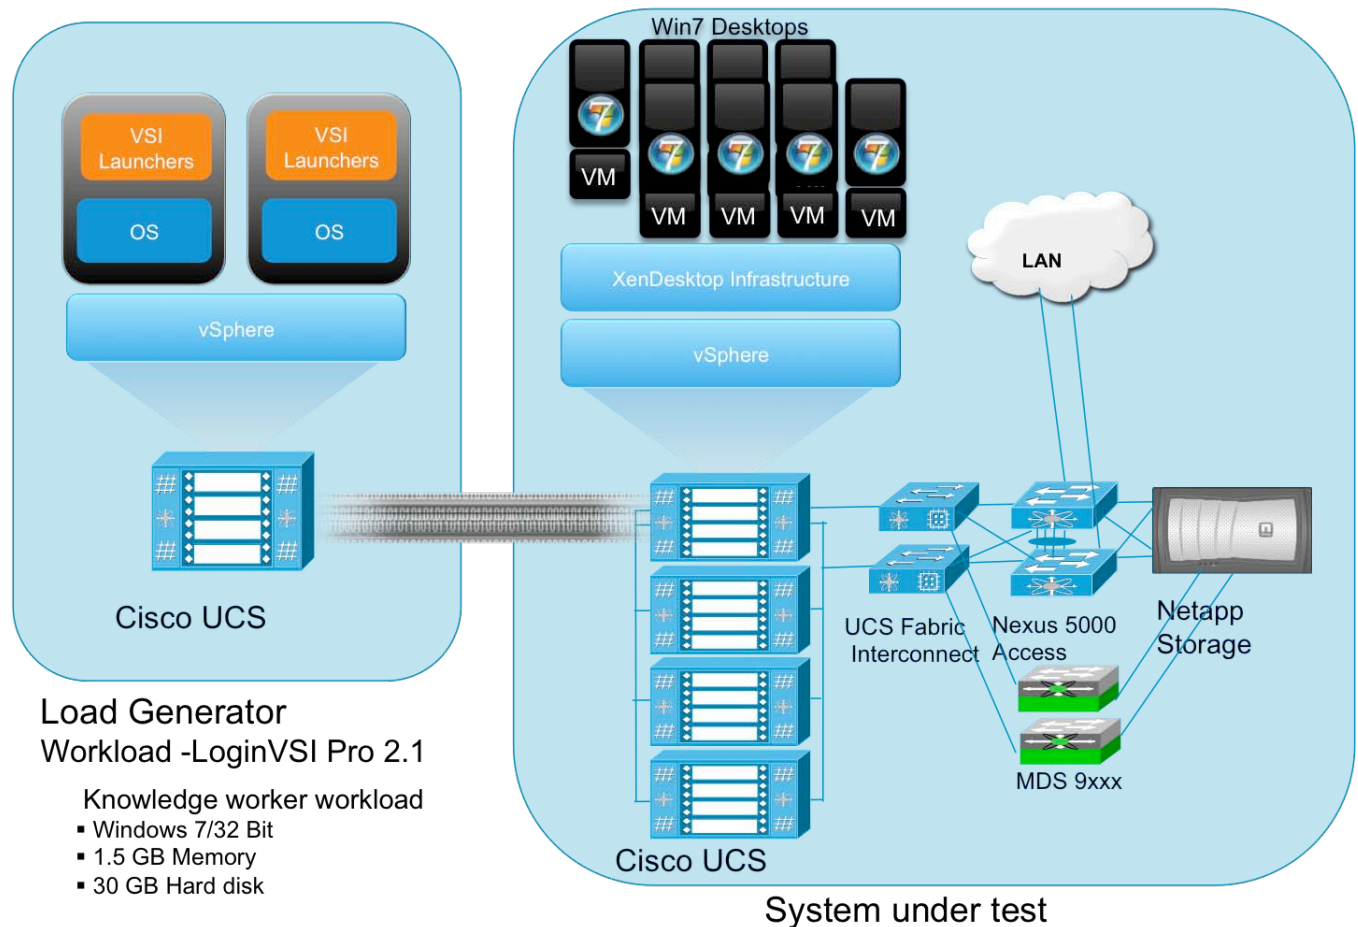
- 8 X Cisco UCS B250-M2 (5680 @ 3.33 GHz) blade servers with 192 GB of memory (4 GB X 48 DIMMS @ 1333 MHz)
- 2 X Cisco UCS B200-M2 (5680 @ 3.33 GHz) blade servers with 48 GB of memory (4 GB X 12 DIMMS @ 1333 MHz)
- Two Cisco Virtual Interface Card Adapters on each blade server
- Cisco Nexus 5000 and 7000
- NetApp FAS 3140 storage array, two controllers, 2 X Dual port 10G Chelsio cards with 70 SAS drives

Software components

- Cisco UCS firmware 1.3(1i)
- VMware vSphere 4.0 U2, vSphere 4.0iU2, Virtual Center 4.0U2
- XenDesktop 4
- Windows 7 – 32 bit, 1vCPU, 1.5 GB of memory, 30 GB/VM

6.3 Cisco UCS Configuration for Four-Chassis Test

Figure 50. Four-Chassis Test Configuration



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 Virtual Interface Card Adapters on each blade servers and combination of Menlo-Q and Menlo-E cards for vSphere.
- Cisco Nexus 5000 and 7000
- NetApp FAS 3140 and FAS 3170 storage array, two controllers, 2 X Dual port 10G Chelsio cards with 70 SAS drives

Software components

- Cisco UCS firmware 1.3(1i)
- VMware vSphere 4.0 U2, VMware vSphere 4iU2, Virtual Center 4.0U2
- XenDesktop 4
- Windows 7 – 32 bit, 1vCPU, 1.5 GB of memory, 30 GB/VM

6.4 Testing Methodology and Success Criteria

All validation testing was conducted on-site within the Cisco labs with joint support from both Citrix and NetApp resources. The testing results focused on the entire process of the virtual desktop lifecycle by capturing metrics during the desktop boot-up, user login, user workload execution (also referred to as steady state), and user logoff for both the Hosted Shared and Hosted VDI models. 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.

6.4.1 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. Based on the environment design, different load generators were used between the Hosted VDI and Hosted Shared environment.

In 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. In the Hosted Shared environment, the standard Login VSI launcher was used simulate multiple users making a direct connection to the shared desktop of the XenApp servers via an ICA connection.

6.4.2 User Workload Simulation – LoginVSI from Login Consultants

One of the most critical factors of validating a XenDesktop deployment is identifying a real-world user workload that is easy for customers to replicate and standardized across 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 Hosted Shared and Hosted VDI testing. These tools have the added benefit of taking measurements of the in-session response time, providing an objective way to measure the expected user experience for individual desktop throughout large scale testing, including login storms.

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

1. 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.
2. Once a session has been started the medium workload will repeat every 12 minutes. During each loop the response time is measured every 2 minutes.
3. Each loop consists of the following operations:

- a. Browse and compose Outlook 2007 messages.
- b. Open multiple instances of Internet Explorer based browsing sessions including heavy multimedia websites.
- c. Open multiple instances of Word 2007 performing open, close and edit operations.
- d. Print and review PDF documents using Bullzip PDF Printer & Acrobat Reader.
- e. Open, edit and close a randomized large Excel 2007 sheet.
- f. Review and edit a PowerPoint 2007 presentation.
- g. Perform zip operations using 7-Zip.

6.4.3 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 key 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 key metrics are important not only based on the raw data that they provide, but also in their ability to align the test results between the Hosted Shared and Hosted VDI models.

6.4.3.1 Login VSICorrected 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 help ensures that the test results are fair and comparable. Therefore, the COPI is calculated as:

$$\text{COPI} = \text{UOPI} - (\text{SSC} * 50\%) - \text{LSC}$$

6.4.3.2 Login VSIMax

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.

7.0 Test Results

The purpose of this testing is to provide the data needed to validate Citrix XenDesktop 4 FlexCast models Hosted VDI and Hosted Shared with VMware vSphere 4 virtualizing Microsoft Windows 7 desktops on Cisco UCS blade servers implemented with NetApp's FAS3170 and 3140 storage systems. The test results are divided into the individual FlexCast models Hosted VDI and Hosted Shared. 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 VMware vSphere scalability.

7.1 XenDesktop Hosted VDI Test Results

This section details the results from the XenDesktop Hosted VDI validation testing. The key success criteria metrics are provided to validate the overall success of each test cycle with all the supporting validation data listed in the validation of the final four chassis configuration. The single server graphs shown in this section are representative of a single vSphere server 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.

7.1.1 Results for Single Cisco UCS Blade Server Validation

The first process in the validation was to find out how many Virtual desktops can be hosted on a single Cisco UCS blade server. When identifying how many virtual desktops per server, it was important to assess the total available RAM, the amount of CPU, IO and Network overheads. Each virtual desktop was configured with 1.5 GB of RAM and each blade had 192 GB of RAM available. We performed a gradual scaling of the desktops till we reached a point where the CPU or memory was maxing out. 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. Another criteria was to avoid any ballooning or swapping, which are usually signs of system degradation.

Table 7 Provides the VSI COPI score for the overall single Cisco UCS blade environment and shows that 100% of all the 110 virtual desktop sessions executed without issue.

Table 7. 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

Once it can be confirmed that all 110 sessions executed successfully, it is imperative 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 51. 110 Desktop Sessions on VMware vSphere Below 2000ms




Figure 52. CPU Utilization on the Server Hosting 110 Desktops




Figure 53. Memory Utilization on Server Hosting 110 Desktops




Figure 54. Network Utilization on Server Hosting 110 Desktops





Figure 55. Vmknics Utilization on the Server Hosting 110 Desktops (NFS traffic)



7.2 Results from Two-Chassis Testing

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

As shown in Figure 56, 99.6% of the measured response times were below the 2000ms threshold to provide the test with a passing score; therefore, it can be concluded that the user response time was not dramatically affected by the heavy 880 desktops load.

Figure 56. 880 Desktop Sessions on VMware vSphere Below 2000ms




Figure 57. Sample CPU Utilization From One of the Servers Hosting 110 Desktops




Figure 58. Sample Memory Utilization From One of the Servers Hosting 110 Desktops





Figure 59. Sample Network Utilization From One of the Servers Hosting 110 Desktops



7.3 Results from Four-Chassis Testing


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-Cisco UCS blade environment and shows that 100% of all the 1760 virtual desktop sessions executed without issue.

Table 9. Overall VSI COPI Score

Total Sessions Launched	1760
Uncorrected Optimal Performance Index (UOPI)	1760
Stuck Session Count before UOPI(SSC)	0
Lost Session Count before UOPI (LSC)	0
Corrected Optimal Performance Index (COPI = UOPI – (SSC*50%) - LSC)	1760


As shown in the figure below, 99.9% of the measured response times were below the 2000ms threshold to provide the test with a passing score; therefore, it can be concluded that the user response time was not dramatically affected by the heavy 1760 desktops load.

Figure 60. 1760 Desktop Sessions on VMware vSphere Below 2000ms




As previously mentioned, the following two graphs are only representative of a single Cisco UCS blade server's 'average CPU utilization' and 'total memory used' to provide a sample of the performance metrics as recorded for the overall 16-blade environment. The average CPU utilization was the most intensive during the workload (steady state) portion of the testing averaging ~95% utilization.

Figure 61. Sample CPU Utilization From Two Servers Hosting 110 Desktops



As seen in the graph below, the majority of the physical memory of the single blade server was consumed by the 110 active desktop sessions. Each virtual desktop was configured with 1.5 GB of RAM. With 110 virtual desktops utilizing 1.5 GB of RAM per virtual desktop, 165GB of the available memory is consumed by virtual desktops. The variance between 165GB and the line shown on the graph is the amount of memory being utilized by the vSphere hypervisor.

Figure 62. Sample Memory Utilization From Two Servers Hosting 110 Desktops



One other 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 VM density per CPU core was maintained while the number of hosts was increased showing a linear CPU core to VM density ratio.

Table 10. VM Density Throughout Testing Phases

<ul style="list-style-type: none"> Windows7 pooled desktops 1vCPU and 1.5GB RAM. 3 GB PVS Cache/OS Paging File on NFS Volumes 	vSphere		
	No. of Servers Tested	No. of VMs	VMs/Core
	1 Blade	110	9.16
	8 Blades	880	
<ul style="list-style-type: none"> Cisco UCS B250 M2s w/ Dual Six Core (3.33GHz) 192GiB RAM 	16 Blades	1760	

7.3.1 Storage Utilization Graphs

There were 350 desktops on each volumes, and number of volumes varied per controller as shown in the figure.

Figure 63. Total Disk Throughput in Kb/sec on 3140 Filer -A




Figure 64. Total Disk Throughput in Kb/sec on 3140 Filer -B




Figure 65. Total Disk Throughput in Kb/sec on 3170 Filer -A




Figure 66. Total Disk Throughput in Kb/sec on 3170 Filer -B




Figure 67. Network Throughput on 3140 Filer-A




Figure 68. Network Throughput on 3140 Filer-B




Figure 69. Network Throughput on 3170 Filer-A




Figure 70. Network Throughput on 3170 Filer-B




Figure 71. One Volume Data for Total Protocol Operations




Figure 72. One Volume Data for Total NFSv3 Read Sizes





Figure 73. One Volume Data for Total NFSv3 Write Sizes




7.4 Citrix XenDesktop with XenApp Hosted Shared Test Results

Customers are looking to virtualize XenApp implementations for a combination of reasons, some which include the flexibility to consolidate under-utilized XenApp servers, application or desktop silos, business continuity planning, etc. So when looking to virtualize XenApp for Hosted Shared desktops, it is important to assess the best virtual machine configurations for optimal performance. Unlike the multi-chassis configuration of the Hosted VDI testing, the key objective for the testing was to obtain an optimal virtualized configuration for multiple XenApp VMs on a single blade. This means finding an optimal balance between number of VMs and number of vCPUs per VM – while successfully supporting the maximum number of user sessions. The following configurations were tested as part of this effort and the Total Number of User Sessions achieved for each configuration is highlighted:

Table 11. Citrix XenApp Virtual Machine Configuration on a Single Cisco UCS Blade Server

XenApp Virtual Machine Configuration on Single Cisco UCS blade server				
Number of XenApp VMs per Blade:	2 VMs/Blade	3 VMs/Blade	4 VMs/Blade	6 VMs/Blade
vCPU per VM	6 vCPU	8 vCPU	6 vCPU	4 vCPU
RAM per VM	16GB RAM	16GB RAM	16 GB RAM	12 GB RAM
Paging File per VM	24 GB	24 GB	24 GB	18 GB
Total Number of User Session across all VMs	120	110	120	120

Figure 74. Total Number of Sessions Per Configuration



- Optimal multiple VM configuration on a single Cisco UCS B200 M2 Blade (maximum scale-out). The maximum number of user sessions supported on a single Cisco UCS B200 M2 blade with multiple XenApp VMs was 120. The optimal virtualized configuration was 2 XenApp VMs each configured with 6vCPU, 16GB RAM and 24GB paging file. To achieve this maximum scale-out value, each XenApp VM supported 60 users. Tests were also conducted with same configurations for 4 XenApp VMs and the maximum scale-out remained at 120.
- Another configuration that achieved 120 user sessions was the 6 XenApp VMs each configured with 4vCPU, 12GB RAM and 18GB of paging file. This configuration is not considered optimal as it requires a higher number of XenApp VMs to achieve the same scale-out value. However, customers implementing XenApp servers in an application based silo fashion might find such configurations optimal for their usage.
- We also conducted testing for 3 XenApp VMs each configured with 8vCPU, 12GB RAM and 18GB of paging file. This configuration yielded the maximum scale-out value of 110 with each VM supporting user sessions.

The maximum scale-out tests were conducted with all 24 logical cores utilized, with hyper-threading enabled and without any vCPU oversubscription. This was based on the general consensus in the industry that for Server Based Computing (SBC) environments, scalability is degraded if more vCPUs are assigned than there are CPU cores.


The following table provides the VSI COPI score for the optimal virtualized configuration consisting of 2 XenApp VMs each configured with 6 vCPU, 16 GB RAM, and 24 GB page file and shows that 100% of all the 120 virtual desktop sessions executed without issue.

Table 12. VSI COPI Score of Two Citrix XenApp Virtual Machines

Total Sessions Launched	120
Uncorrected Optimal Performance Index (UOPI)	120
Stuck Session Count before UOPI(SSC)	0
Lost Session Count before UOPI (LSC)	0
Corrected Optimal Performance Index (COPI = UOPI – (SSC*50%) - LSC)	120


In addition to evaluating the successful completion of the workload within a user's desktop session, it is important to also help ensure that the user experience was not degraded as load was increased on the environment. The following graph provides the VSI Response Time Frequency Distribution which is used to calculate the VSI Max Pass or Fail rating and determines the scalability limits of the system. As seen in the figure, 99.6% of the measured response times were below the 2000ms threshold to provide the test with a passing score proving that the Cisco UCS B200 M2 blade successfully supported 110 Citrix XenApp user sessions without being overloaded.

Figure 75. Login VSI Response Time Frequency Distribution Measure for Executing 120 Citrix XenApp User Sessions on a Cisco UCS B200 M2 Blade Server




When assessing the limit of user sessions per single XenApp VM, it's important to assess both the Memory and CPU of an individual VM. The following graphs provide the 'average CPU utilization' and 'total memory used' for a single XenApp VM from the 2-VM configuration during a steady state of workload execution for the peak of 60 user sessions.

Figure 76. Citrix XenApp Virtual Machine Average CPU Utilization During Steady State Execution of 60 User Sessions




From the CPU utilization graph (above) it can be noted that the CPUs were reaching their limit with 60 user sessions given that the average CPU usage hovered around 80%. In regards to Memory, each XenApp VM was configured with 16GB of RAM per virtual machine, so from the Total Memory Used graph below, it can be concluded that Memory was not a limiting factor.

Figure 77. Citrix XenApp Virtual Machine Total Memory Utilization During Steady State Execution of 60 User Sessions



After concluding that the CPU performance of the individual XenApp VM was the limiting factor, the overall CPU performance of the hosting hypervisor needs to be evaluated. The following graph provides Average CPU Utilization for the Cisco UCS B200 M2 Blade Server during a steady state execution for the two virtual machine configuration with 120 active XenApp desktop sessions.

Figure 78. Cisco UCS B200 M2 Blade Server CPU Utilization During Steady State Execution of 180 Citrix XenApp Sessions



8.0 Scaling and Sizing 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. In this section we give guidance to scale beyond 4 UCS chassis.

8.1 Scalability Considerations and Guidelines

8.1.1 Cisco UCS System Configuration

As the results indicate we are seeing linear scalability in the Cisco UCS reference architecture implementation.

vSphere			
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 we got during the testing we get the following results:

vSphere			
No. of Chassis	No. of B250-M2 Servers	No. of VMs	VMs/Core
8	32 Blades	3520	9.16
12	48 Blades	5280	
16	64 Blades	7040	
20	80 Blades	8800	

To accommodate the Cisco Nexus 5000 upstream connectivity in the way we describe in the lan configuration section, we need four Ethernet uplinks to be configured on the Cisco UCS Fabric interconnect. And based on the number of uplinks from each chassis, we could calculate how many desktops can be hosted in a single UCS domain. Assuming two links per chassis, scaling beyond 10 chassis would need a Cisco UCS 6140 fabric interconnect. A 5000 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>.

8.1.2 vSphere Configuration

On vSphere 4.0U2, the configuration maximums have to be taken in to consideration while doing scaling calculations. Refer to the Configuration Maximums for VMware vSphere 4.0 at: http://www.vmware.com/pdf/vsphere4/r40/vsp_40_config_max.pdf.

The main parameters are:

- DRS Cluster
 - Hosts per DRS cluster
 - Virtual machines per DRS cluster
 - Virtual machines per host in DRS cluster

If you are implementing a HA cluster:

- HA Cluster
 - Hosts per HA cluster
 - Virtual machines per host in HA cluster with 8 or fewer hosts
 - Virtual machines per host in HA cluster with 9 or more hosts

The vCenter scalability is another important criteria as we may have to add another vCenter if we are scaling beyond the capability of single VC. Using Linked vCenter concept one can have more than one VC to manage a large number of ESX hosts and clusters.

- vCenter Server Scalability
 - Hosts (32/64-bit OS server)
 - Powered-on virtual machines (32/64-bit OS server)
 - Registered virtual machines (32/64-bit OS server)

8.2 Sizing Guidelines

One of the key learning from this study was the characterization of Microsoft Windows 7 desktop profile with respect to CPU and Memory. The majority of the physical memory of the single blade server with 192 GB of memory was consumed by the 110 active desktop sessions with each windows 7 virtual desktop configured with 1.5 GB of RAM. With regards to the CPU, the LoginVSI medium workload we noticed around 95% of the CPU was utilized during the steady state or 690 MHz of CPU on the host.

8.2.1 CPU Calculations

CPU in MHz for each desktop = (Number of cores * Frequency in GHZ)*Percentage CPU utilization*1000/Total number of desktops.

CPU in MHZ for each desktop = $(24 * 3.33) * 0.95 * 1000 / 110 = 690 \text{ MHZ}$

If your workload happens to mimic medium workload that we used then each desktop will consume 690 MHZ average.

8.2.2 Memory Calculations

VMware vSphere reserves some 6 percent of the total physical memory available for system memory, and the remaining non-kernel memory can be used by the desktops. This amount has to be subtracted from the total memory during the memory calculations.

For example, consider Windows 7 with 1.5 GB RAM, scalability on a B200-M2 Cisco UCS blade server system with 96 GB of memory and Intel 5680 at 3.33 GHZ speed would be limited by memory and not by the CPU for the knowledge worker load.



Number of desktops/blade = (Total memory on the system * 0.94)/(memory for one Win 7 desktop).

In our example, Number of desktops/blade = (96 * 0.94) / 1.5 = 61 desktops (approximately).

In this case a lower bin CPUs can be used instead of 3.33 GHZ (5680) which will have spare CPU cycles for peak.

Other parameters like overall power consumption of the system, storage IOPS, rack space, cooling considerations also plays an important role in sizing calculations.



9.0 Acknowledgements

Projects of this magnitude could only be done with co-operation of all the parties involved and this work is a clear testimony of that. A lot of people had helped to make this project successful, we would like to acknowledge the contribution of Purnanand for helping out in Networking configuration, Lab guys – TJ and Vincent for accommodating all requests, Lisa DeRuyter for documentation, Joe Vaccaro, and Scott Gainey in SAVBU product marketing organization on the Cisco side. We also thank Rob de Groot, Alfonso Villasenor, Lee Dorrier, and Angela Ge from NetApp for their support of this work. We thank Bhumik Patel for helping out in the XenApp testing, Samantha Foster (business development), and Rana Kannan(PM) from Citrix. Special thanks to Satinder Sethi, SAVBU technical marketing, for being a driving force behind this work.

Key contributors:

Ravindra “Ravi” Venkat (Cisco Systems)

Frank Anderson (Citrix)

Carisa Stringer (Citrix)

Rachel Zhu (NetApp)

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](#) for the Login VSI benchmarking tool.



10.0 References

1. TR-3749: NetApp and VMware vSphere Storage Best Practices

<http://media.netapp.com/documents/tr-3749.pdf>

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

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

3. Cisco Nexus 5000 Series Switch CLI Software Configuration Guide

http://www.cisco.com/en/US/docs/switches/datacenter/nexus5000/sw/configuration/guide/cli_rel_4_0_1a/CLIConfigurationGuide.html

4. Best practices in deploying Cisco Nexus 1000v series switches on Cisco UCS B Series Blade servers

http://www.cisco.com/en/US/prod/collateral/switches/ps9441/ps9902/white_paper_c11-558242.html

5. PVS on XD BP:

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

6. XD Design Handbook:

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

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

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

APPENDIX A

Cisco Nexus 5000 Configuration

```
switchname sj2-151-d17-n5010a
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 0x6ab2f7da5f26e2b1bc37d79438a89bb3 priv
0x6ab2f7da5f26e2b1bc37d79438a89bb3 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
```



Cisco Systems, Inc.

170 West Tasman Drive

San Jose, CA 95134-1706

USA

www.cisco.com

Tel: 408 526-4000

800 553-NETS (6387)

Fax: 408 527-0883

© 2010 Cisco Systems, Inc. All rights reserved. Cisco, the Cisco logo, and Cisco Systems are registered trademarks or trademarks of Cisco Systems, Inc. and/or its affiliates in the United States and certain other countries. All other trademarks mentioned in this document are the property of their respective owners. (0805R)

Document number: UCS-TR100017



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.

ALL DESIGNS, SPECIFICATIONS, STATEMENTS, INFORMATION, AND RECOMMENDATIONS (COLLECTIVELY, "DESIGNS") IN THIS MANUAL ARE PRESENTED "AS IS," WITH ALL FAULTS. CISCO AND ITS SUPPLIERS DISCLAIM ALL WARRANTIES, INCLUDING, WITHOUT LIMITATION, THE WARRANTY OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE AND NONINFRINGEMENT OR ARISING FROM A COURSE OF DEALING, USAGE, OR TRADE PRACTICE. IN NO EVENT SHALL CISCO OR ITS SUPPLIERS BE LIABLE FOR ANY INDIRECT, SPECIAL, CONSEQUENTIAL, OR INCIDENTAL DAMAGES, INCLUDING, WITHOUT LIMITATION, LOST PROFITS OR LOSS OR DAMAGE TO DATA ARISING OUT OF THE USE OR INABILITY TO USE THE DESIGNS, EVEN IF CISCO OR ITS SUPPLIERS HAVE BEEN ADVISED OF THE POSSIBILITY OF SUCH DAMAGES.

THE DESIGNS ARE SUBJECT TO CHANGE WITHOUT NOTICE. USERS ARE SOLELY RESPONSIBLE FOR THEIR APPLICATION OF THE DESIGNS. THE DESIGNS DO NOT CONSTITUTE THE TECHNICAL OR OTHER PROFESSIONAL ADVICE OF CISCO, ITS SUPPLIERS OR PARTNERS. USERS SHOULD CONSULT THEIR OWN TECHNICAL ADVISORS BEFORE IMPLEMENTING THE DESIGNS. RESULTS MAY VARY DEPENDING ON FACTORS NOT TESTED BY CISCO.

CCDE, CCENT, Cisco Eos, Cisco Lumin, Cisco Nexus, Cisco StadiumVision, Cisco TelePresence, Cisco WebEx, the Cisco logo, DCE, and Welcome to the Human Network are trademarks; Changing the Way We Work, Live, Play, and Learn and Cisco Store are service marks; and Access Registrar, Aironet, AsyncOS, Bringing the Meeting To You, Catalyst, CCDA, CCDP, CCIE, CCIP, CCNA, CCNP, CCSP, CCVP, Cisco, the Cisco Certified Internetwork Expert logo, Cisco IOS, Cisco Press, Cisco Systems, Cisco Systems Capital, the Cisco Systems logo, Cisco Unity, Collaboration Without Limitation, EtherFast, EtherSwitch, Event Center, Fast Step, Follow Me Browsing, FormShare, GigaDrive, HomeLink, Internet Quotient, IOS, iPhone, iQuick Study, IronPort, the IronPort logo, LightStream, Linksys, MediaTone, MeetingPlace, MeetingPlace Chime Sound, MGX, Networkers, Networking Academy, Network Registrar, PCNow, PIX, PowerPanels, ProConnect, ScriptShare, SenderBase, SMARTnet, Spectrum Expert, StackWise, The Fastest Way to Increase Your Internet Quotient, TransPath, WebEx, and the WebEx logo are registered trademarks of Cisco Systems, Inc. and/or its affiliates in the United States and certain other countries.

All other trademarks mentioned in this document or website are the property of their respective owners. The use of the word partner does not imply a partnership relationship between Cisco and any other company. (0809R)

© 2010 Cisco Systems, Inc. All rights reserved