VMware View 4.5 on FlexPod for VMware Design Guide

Last Updated: July 6, 2011

Cisco Validated Design

Building Architectures to Solve Business Problems

About the Authors

Nico Donato

Nico Donato, Technical Marketing Engineer, Systems Architecture and Strategy Unit, Cisco Systems

Nico Donato is a Technical Marketing Engineer in the Cisco Data Center Partner Applications Innovation Center in the Systems Architecture and Strategy Unit (SASU). He previously supported the MDS product line for Storage Area Networks in Cisco's Technical Assistance Center for the past 5 years. Nico's work prior to joining Cisco included systems administration and level III/IV support in Windows environments and Storage Networking in the NAS and SAN space. His tenure also includes two years with Network Appliance. Nico has been in the IT industry for the last 14 years with interests in security and volunteers time with InfraGard.

Ramesh Issac

Ramesh Isaac, Technical Marketing Engineer, Systems Architecture and Strategy Unit, Cisco Systems

Ramesh has worked in data center and mixed use lab settings over the past 15 years. He started in information technology supporting Unix environments, with the last couple of years focused on designing and implementing multi-tenant virtualization solutions in Cisco labs. Ramesh holds certifications from Cisco, VMware, and Red Hat.

Chris O'Brien

Chris O'Brien, Solutions Architect, Systems Architecture and Strategy Unit, Cisco Systems

Chris O'Brien is a Solutions Architect for data center technologies in Cisco's Systems Architecture and Strategy Unit (SASU). He is currently focused on data center design validation and application optimization. Previously, O'Brien was an application developer and has been working in the IT industry for more than 15 years.

The authors would like to give special thanks to Ravi Venkat (Cisco), Jack McLeod (NetApp), Mac Binesh (VMware), Ravi Neelakant (VMware), and Fred Schimscheimer the author of RAWC (VMware) for contributions and assistance without which this paper would not have been possible.

About Cisco Validated Design (CVD) Program

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

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.

Any Internet Protocol (IP) addresses and phone numbers used in this document are not intended to be actual addresses and phone numbers. Any examples, command display output, network topology diagrams, and other figures included in the document are shown for illustrative purposes only. Any use of actual IP addresses or phone numbers in illustrative content is unintentional and coincidental.

VMware View 4.5 on FlexPod for VMware Design Guide

VMware View 4.5 on FlexPod for VMware Design Guide

Introduction

This document reports the results of a study evaluating the scalability of VMware® View 4.5™ on FlexPod™, a best-of-breed bundling of Cisco® Unified Computing System™ (UCS) B-series blade servers connected to the NetApp® FAS Storage array. It is an update of a previously published document that includes new scaling findings, which were achieved with performance enhancements that are included in this document. Best practice design recommendations and sizing guidelines for large-scale customer deployments are also provided.

Audience

This document is intended to assist solution architects, sales engineers, field engineers, and consultants in planning, designing, and deploying VMware View hosted VDI solutions on the Cisco UCS. This document assumes the reader has an architectural understanding of the Cisco UCS, VMware View 4.5 software, NetApp storage system, and related software.

Updated Content in this Document Version

- Summary of Main Findings
- FlexPod from Cisco and NetApp
- VMware vSphere Kernel Adjustments for High CPU Environments
- Test Results

Summary of Main Findings

- Scale test findings—160 Windows 7 desktops (1.5 GB) running knowledge worker load were supported with one blade server.

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

- Scale findings increased from the previous study through the use of kernel adjustments detailed in VMware vSphere Kernel Adjustments for High CPU Environments.
- The large memory blade (B250 M2) delivers the best price performance.
- Linear scalability from 1 to 8 servers, with the following results:
  - One server supported 160 desktops
  - Eight servers supported 1280 desktops (with similar response times)
- Pure Virtualization—The validated environment consists of all virtual machines hosted by VMware vSphere. All the virtual desktop and supporting infrastructure components, including Active Directory and vCenter Server, are hosted on VMware vSphere.
- VMware View allows simplified management of large numbers of automatically standardized desktop resources.
- Rapid provisioning with Cisco UCS Manager makes it easy to scale from one chassis to two and so on.
- The 10G unified fabric delivers tremendous performance with respect to user response times during the load test.
- The low latency Cisco Virtual Interface (VIC) cards enable more robust configurations with virtual NICs and contributes to the excellent response time.
- The validated design provides linear scalability without changing the reference architecture.
- The B250 M2 blade with 192 GB of memory delivers optimal memory for desktop virtualization that allows full CPU utilization of the server without restricting the amount memory allocated to the desktops.
- Advanced Storage Technologies simplify management, enable scalable designs, and reduce TCO:
  - Storage efficiency with multiple levels of storage (thin provisioning, data deduplication, and FlexClone®)
  - Performance enhanced with Transparent Storage Cache Sharing (TSCS) and extended with Flash Cache (PAM II)

Note For the remainder of this document, the term Flash Cache will represent the Flash Cache (PAM II) module.

  - Enhanced VDI solution management for operational efficiency

Infrastructure Components

This section describes the infrastructure components used in the system design.

FlexPod from Cisco and NetApp

Figure 1 FlexPod

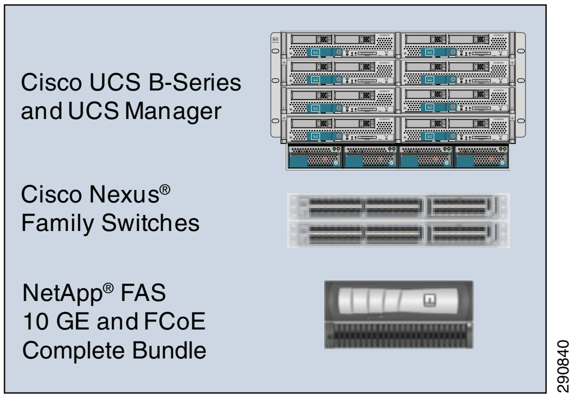

FlexPod is a predesigned, base configuration that is built on the Cisco Unified Computing System (UCS), Cisco Nexus® data center switches, NetApp FAS storage components, and a range of software partners. FlexPod can scale up for greater performance and capacity or it can scale out for environments that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

Benefits

- Low-risk, standardized, shared infrastructure supporting a wide range of environments
- Highest possible data center efficiency
- IT flexibility providing the business agility to scale out or up and manage resource pools

Features

- Complete data center in a single rack
- Performance-matched stack
- Step-by-step deployment guides
- Solutions guide for multiple environments
- Multiple classes of computing and storage supported in a single FlexPod
- Centralized management with NetApp OnCommand and Cisco UCS Manager

As the name suggests, the FlexPod architecture is highly modular or “pod” like. While each customer’s FlexPod may vary in its exact configuration, once a FlexPod unit is built it can easily be scaled as requirements and demand change. This includes scaling both up (adding additional resources within a FlexPod unit) and out (adding additional FlexPod units). FlexPods are created to be “flexed” to the customer’s application needs.

FlexPod for VMware

- FlexPod infrastructure with VMware vSphere and vCenter to support virtualized application workloads

- Complex solutions can be layered on top of FlexPod for VMware.
  - With this environment the ESMT reference architecture was implemented on FlexPod for VMware (http://www.cisco.com/en/US/docs/solutions/Enterprise/Data_Center/Virtualization/securecldg_V2.html).
- Optional components include the Nexus 1010 and the Nexus 1000v
- This FlexPod for VMware was "flexed" with:
  - High memory availability within the Cisco UCS B250 M2 blades
  - Storage acceleration achieved through NetApp Flash Cache

For more information on FlexPod, see: http://www.cisco.com/en/US/docs/solutions/Enterprise/Data_Center/Virtualization/flexpod_vmware.html.

Cisco Unified Computing System

The Cisco UCS 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® 5600 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 UCS 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 UCS components illustrated in Figure 2 include, from left to right, Fabric Interconnects, blade server chassis, blade servers, and in the foreground, fabric extenders and network adapters.

Figure 2 Cisco Unified Computing System

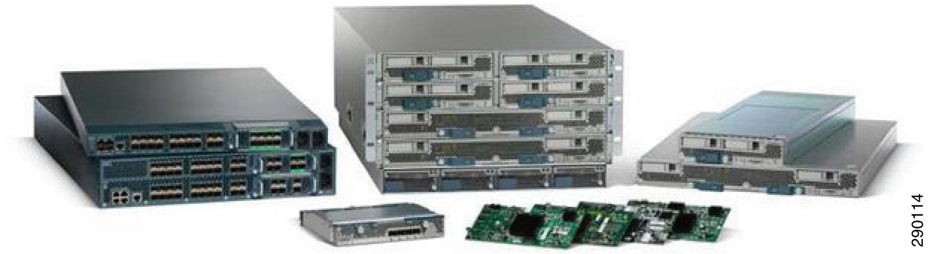

Cisco Unified Computing System Components

Fabric Interconnect

The Cisco UCS 6100 Series Fabric Interconnects are a core part of the Cisco UCS, providing both network connectivity and management capabilities for the system (Figure 3). 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 UCS 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 3 Cisco UCS 6120XP 20-Port Fabric Interconnect (Top) and Cisco UCS 6140XP 40-Port Fabric Interconnect

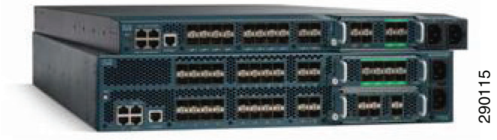

The Cisco UCS 6100 Series is equipped to support the following module options:

- 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 4 From Left to Right—8-Port 1/2/4-Gbps Native Fibre Channel Expansion Module; 4-Port Fibre Channel plus 4-Port 10 Gigabit Ethernet Module; 6-Port 10 Gigabit Ethernet Module; and 6-Port 1/2/4/8-Gbps Native Fibre Channel Expansion Module

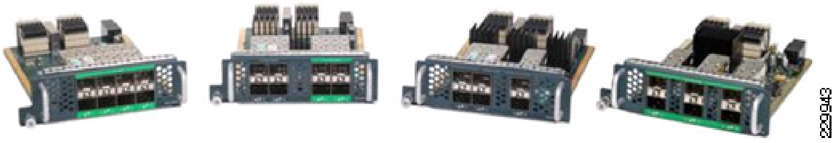

Cisco Fabric Extenders Module

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 UCS 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 5 Rear View of the Cisco UCS 5108 Blade Server Chassis with Two Cisco UCS 2104XP Fabric Extenders

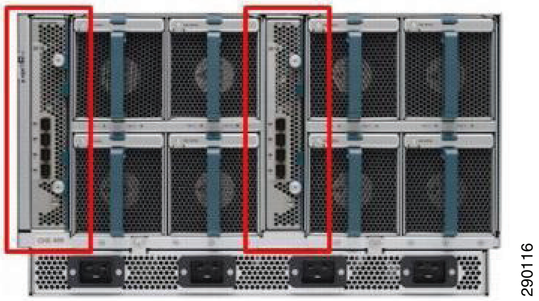

The Cisco UCS 2104XP Fabric Extender has four 10 Gigabit Ethernet, FCoE-capable, Small Form-Factor Pluggable Plus (SFP+) ports that connect the blade chassis to the Fabric Interconnect. Each Cisco UCS 2104XP has eight 10 Gigabit Ethernet ports connected through the midplane to each half-width slot in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

Figure 6 Cisco UCS 2104XP Fabric Extender

Cisco UCS Chassis

The Cisco UCS 5100 Series Blade Server Chassis is a crucial building block of the Cisco UCS, 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 half-width server slot and up to 40 Gbps of I/O bandwidth per full-width server slot. The chassis is capable of supporting future 40 Gigabit Ethernet standards.

Figure 7 Cisco Blade Server Chassis (Front and Back View)

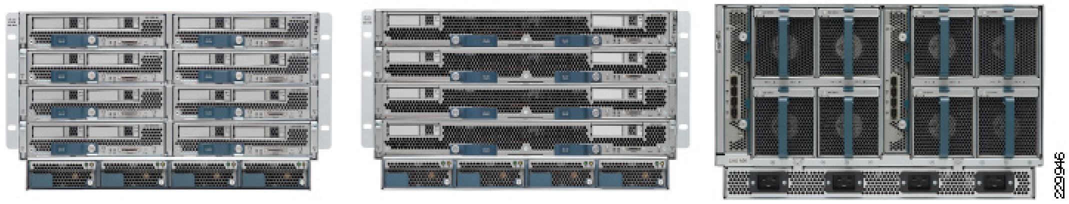

Intel Xeon 5600 Series Processor

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 ground breaking 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 8 Intel Xeon 5600 Series Processor

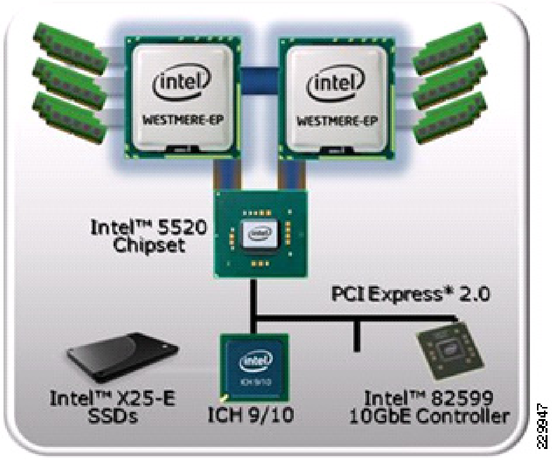

Cisco UCS B200 M2 Blade Server

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 9 Cisco UCS B200 M2 Blade Server

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 dataset workloads with greater memory capacity and throughput.

Figure 10 Cisco UCS B250 M2 Extended Memory Blade Server

Cisco UCS Virtual Interface Card (VIC)

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

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

- Interface type—FCoE, Ethernet, or Dynamic Ethernet interface device
- Resource maps that are presented to the host—PCIe BARs, interrupt arrays
- The Network presence and attributes—MTU, VLAN membership
- Quality of Service (QoS) parameters—802.1p class, ETS attributes, rate limiting, and shaping

Figure 11 Cisco UCS Virtual Interface Card

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

Cisco 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 three 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 plug-in 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 12, the vDS component on each host discovers the number of uplink interfaces that it has and presents a switch to the virtual machines 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 12 Each Virtual Machine Interface has Its Own Virtual Port on the Physical Switch

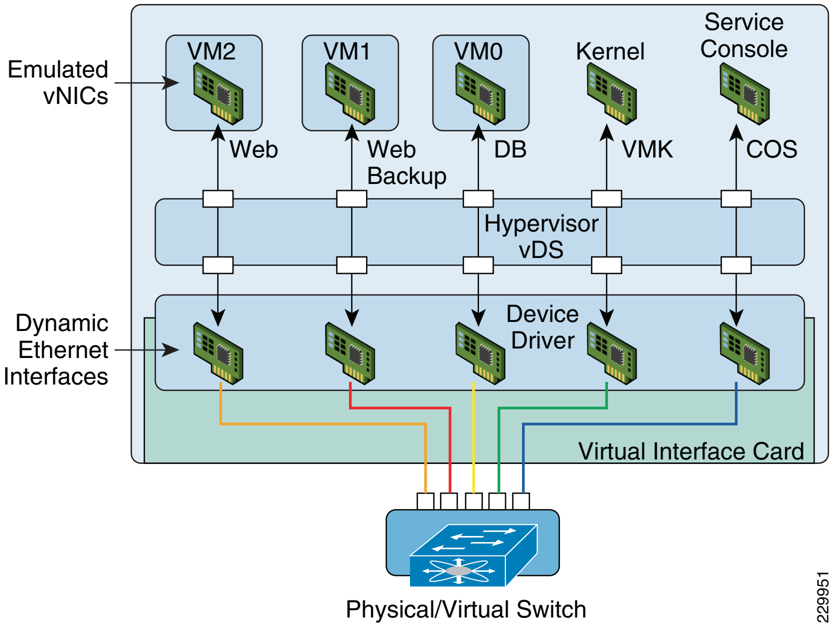

Cisco UCS Manager running on the Cisco UCS Fabric Interconnect works in conjunction with VMware vCenter software to coordinate the creation and movement of virtual machines. Port profiles are used to describe the virtual machine interface attributes such as VLAN, port security, rate limiting, and QoS 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.

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 13). 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 13 Extended Memory Architecture

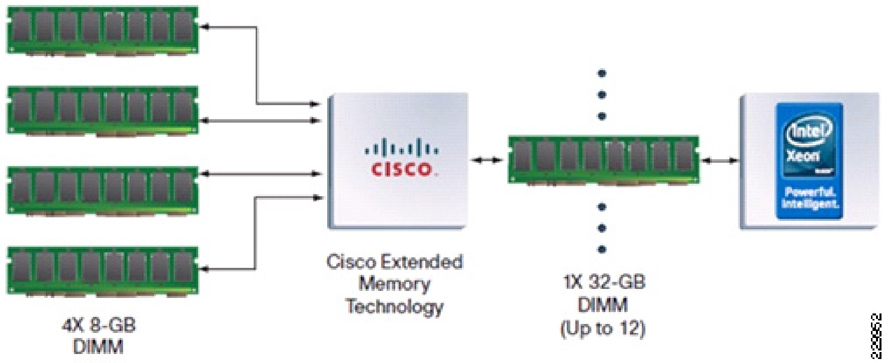

VMware vSphere4 and VCenter Server

VMware vSphere is a virtualization platform for holistically managing large collections of infrastructure-CPUs, storage, networking-as a seamless, flexible, and dynamic operating environment. Unlike traditional operating systems that manage an individual machine, VMware vSphere aggregates the infrastructure of an entire data center to create a single powerhouse with resources that can be allocated quickly and dynamically to any application in need.

VMware vSphere provides revolutionary benefits, but with a practical, non-disruptive evolutionary process for legacy applications. Existing applications can be deployed on VMware vSphere with no changes to the application or the OS on which they are running.

VMware vSphere delivers the performance required to run business-critical applications in large-scale environments. VMware vSphere provides two-four times the performance of the previous generation platform (VMware Infrastructure 3) while keeping virtualization overhead at a very limited 2-10 percent. With these performance numbers, VMware vSphere is able to run large, resource-intensive databases and, in many cases, enables applications to scale better on newer multicore servers.

VMware vSphere provides a set of application services that enable applications to achieve unparalleled levels of availability, security, and scalability. For example, with VMware vSphere, all applications can be protected from downtime with VMware High Availability (HA) and VMware Fault Tolerance (FT), without the complexity of conventional clustering. In addition, applications can be scaled dynamically to meet changing loads with capabilities such as Hot Add and VMware Distributed Resource Scheduler (DRS).

The VMware vCenter Product Family is an advanced virtualization management platform, which unlocks the power of virtualization through proactive management and centralized control of virtual infrastructure. For example, VMware vCenter AppSpeed enables IT operations to monitor and ensure the service levels of distributed multi-tier applications running on VMware vSphere. VMware vCenter Lab Manager 4 provides developers and application owners on-demand, self-service access to a library of application and development environments to accelerate develop and test cycles.

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.

VMware View

VMware View 4.5 desktop virtualization platform enables you to run virtual desktops in the data center and deliver desktops to employees as a secure managed service. End users gain a familiar, personalized environment that they can access from any number of devices anywhere throughout the enterprise or from home. Administrators gain centralized control, efficiency, and security by having desktop data in the data center.

Types of Users

There are many reasons to consider a virtual desktop solution, such as an ever growing and diverse base of user devices, management complexity of traditional desktops, security, and user owned/non-IT supported devices. It is important to understand the requirements of the user community to design and deploy a successful Virtual Desktop environment. Following are some typical types of users:

- Knowledge workers today do not just work in their offices all day; they attend meetings, visit branch offices, and 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 everyday business. They need access to many applications and data, yet administrators have little control over the devices they use or the locations from which they work. Consequently, IT must trade off the cost of providing these workers a device versus 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. Since these workers interact with customers, partners, and employees, they often have access to critical data.
- Road warriors need access to their virtual desktop from everywhere, regardless of how they are connected to a network. 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 typically found in university and business computer labs, in conference rooms, and in training centers. Shared workstation environments require desktop re-provisioning with the latest operating systems or applications, as the needs of the organization change.

The Virtual Desktop user community requirements will drive system design decisions.

VMware View consists of the following major components which work together to deliver a flexible and robust Virtual Desktop environment.

View Connection Server

This software service acts as a broker for client connections. The View Connection Server authenticates users through Windows Active Directory and directs the request to the appropriate virtual machine, physical or blade PC, or Windows Terminal Services server.

View Connection Server provides the following management capabilities:

- Authenticating users

- Entitling users to specific desktops and pools
- Assigning applications packaged with VMware ThinApp to specific desktops and pools
- Managing local and remote desktop sessions
- Establishing secure connections between users and desktops
- Enabling single sign-on
- Setting and applying policies

Inside the corporate firewall, you install and configure a group of two or more View Connection Server instances. Their configuration data is stored in an embedded LDAP directory and is replicated among members of the group.

Outside the corporate firewall, in the DMZ, you can install and configure View Connection Server as a security server. Security servers in the DMZ communicate with View Connection Servers inside the corporate firewall. Security servers offer a subset of functionality and are not required to be in an Active Directory domain.

View Agent

You install the View Agent service on all virtual machines, physical systems, and Terminal Service servers that you use as sources for View desktops. This agent communicates with View Client to provide features such as connection monitoring, virtual printing, and access to locally connected USB devices.

If the desktop source is a virtual machine, you first install the View Agent service on that virtual machine and then use the virtual machine as a template or as a parent of linked clones. When you create a pool from this virtual machine, the agent is automatically installed on every virtual desktop.

You can install the agent with an option for single sign-on. With single sign-on, users are prompted to log in only when they connect to View Connection Server and are not prompted a second time to connect to a virtual desktop.

View Client

The client software for accessing View desktops runs either on a Windows or Mac PC as a native application or on a thin client if you have View Client for Linux.

After logging in, users select from a list of virtual desktops that they are authorized to use. Authorization can require Active Directory credentials, a UPN, a smart card PIN, or an RSA SecurID token.

An administrator can configure View Client to allow end users to select a display protocol. Protocols include PCoIP, Microsoft RDP, and HP RGS. The speed and display quality of PCoIP rival that of a physical PC.

View Client with Local Mode (formerly called Offline Desktop) is a version of View Client that has been extended to allow end users to download virtual machines and use them on their local systems regardless of whether they have a network connection.

View Administrator

This Web-based application allows administrators to configure View Connection Server, deploy and manage View desktops, control user authentication, troubleshoot end user issues, initiate and examine system events, and carry out analytical activities.

When you install a View Connection Server instance, the View Administrator application is also installed. This application allows administrators to manage View Connection Server instances from anywhere without having to install an application on their local computer.

vCenter Server

This service acts as a central administrator for VMware ESX servers that are connected on a network. vCenter Server, formerly called VMware VirtualCenter, provides the central point for configuring, provisioning, and managing virtual machines in the data center.

In addition to using these virtual machines as sources for View desktop pools, you can use virtual machines to host the server components of VMware View, including Connection Server instances, Active Directory servers, and vCenter Server instances.

View Composer

View Composer installs as a software service on a vCenter Server instance to manage the virtual machines. View Composer can then create a pool of linked clones from a specified parent virtual machine. This strategy reduces storage costs by up to 90 percent.

Each linked clone acts like an independent desktop, with a unique host name and IP address, yet the linked clone requires significantly less storage because it shares a base image with the parent.

Because linked-clone desktop pools share a base image, you can quickly deploy updates and patches by updating only the parent virtual machine. End users' settings, data, and applications are not affected. As of View 4.5, you can also use linked-clone technology for View desktops that you download and check out to use on local systems.

VMware View also provides the option to use tiered storage. This allows placement of important data on high I/O storage and infrequently used data on less expensive drives. Configuring View Composer replicas to a SSDD can yield high read performance for linked clone provisioning and concurrent references from linked clones.

View Transfer Server

This software manages and streamlines data transfers between the data center and View desktops that are checked out for use on end users' local systems. View Transfer Server is required to support desktops that run View client with Local Mode (formerly called Offline Desktop).

Several operations use View Transfer Server to send data between the View desktop in vCenter Server and the corresponding local desktop on the client system.

View Transfer Server synchronized local desktops with the corresponding desktops in the data center by replicating user-generated changes to the data center. Replications occur at intervals that you specify in local-mode policies. You can also initiate replications in View Administrator. You can set a policy that allows users to initiate replications from their local desktops.

View Transfer Server keeps local desktops up-to-date by distributing common system data from the data center to local clients. View Transfer Server download View Composer base images from the image repository to local desktops.

If a local computer is corrupted or lost, View Transfer Server can provision the local desktop and recover the user data by downloading the data and system image to the local desktop.

NetApp Storage Solution and Components

NetApp provides a scalable, unified storage and data management solution for VDI. 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 VDI 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.

Single Scalable Unified Architecture

The NetApp unified storage architecture provides customers with an agile and scalable storage platform. 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 reduction in total cost of ownership.

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. With traditional storage, the total storage required equals the sum of the storage required by each virtual machine. For example, if each virtual machine is 20GB in size and there are 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).

Thin Provisioning

Thin provisioning is a way of logically presenting more storage to hosts than is physically available. With thin provisioning, the storage administrator is able to utilize a pool of physical disks (known as an aggregate) and 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 14 Traditional and Thin Provisioning

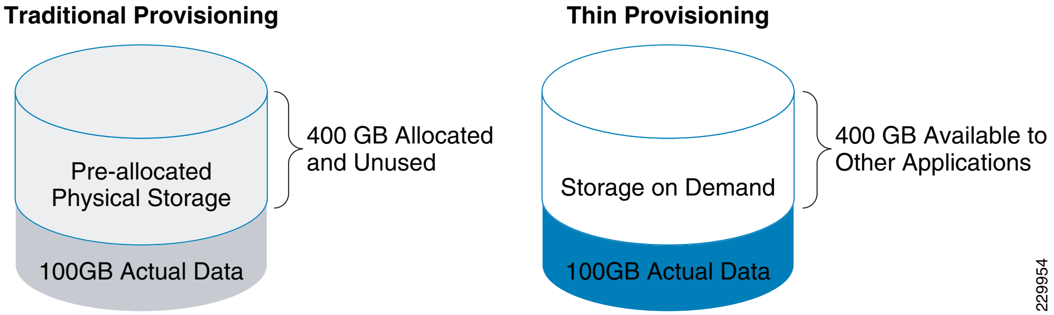

Figure 15 Increased Disk Utilization with NetApp Thin Provisioning

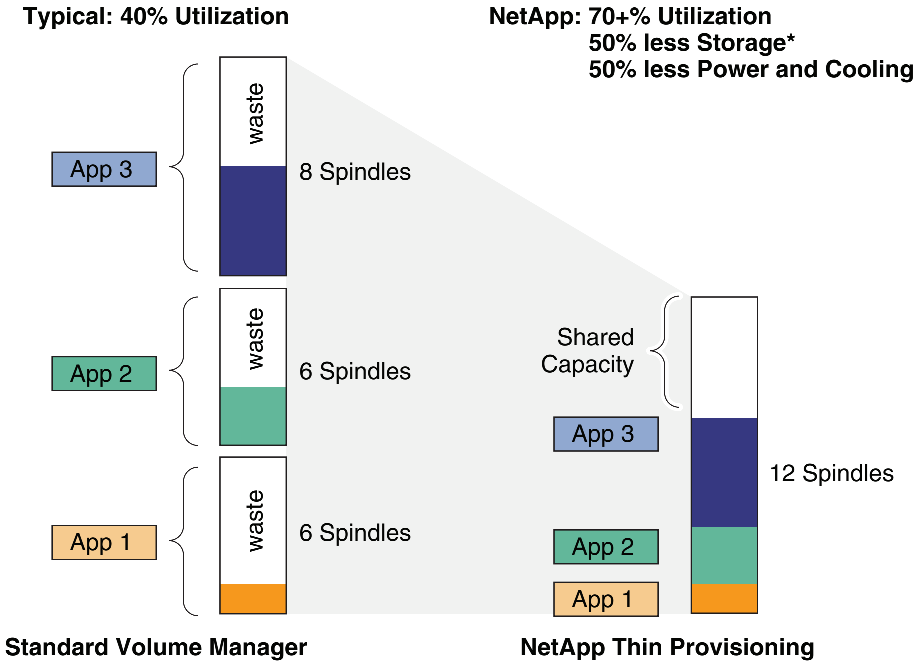

Source: Oliver Wyman Study: "Making Green IT a Reality." November 2007.
*Thin Provisioning, clones, and multiprotocol all contribute to savings.

NetApp Deduplication

NetApp deduplication saves space on primary storage by removing redundant copies of blocks within a volume hosting hundreds of virtual desktops. This process is transparent to the application and user and can be enabled and disabled on the fly. In a VDI environment, deduplication provides significant space savings, given that each virtual machine has 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 16 NetApp Deduplication

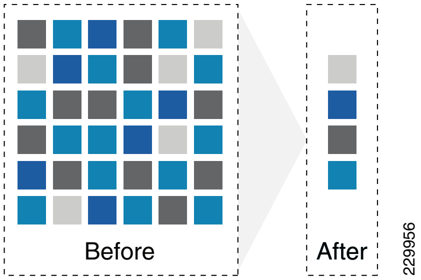

Using NetApp deduplication and file FlexClone can reduce the overall storage footprint of VDI desktops and improve performance by leveraging transparent storage cache sharing. Data that is deduplicated or nondeduplicated, in the case of file FlexClone data, on disk will only exist in storage array cache once per volume. All subsequent reads from any of the virtual machine disks of a block that is already in cache will be read from cache and not from disk, therefore improving performance by 10X. Any nondeduplicated data that is not in cache must be read from disk. Data that is deduplicated but does not have as many block references as a heavily deduped 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 17 NetApp Deduplication in VMware Environments

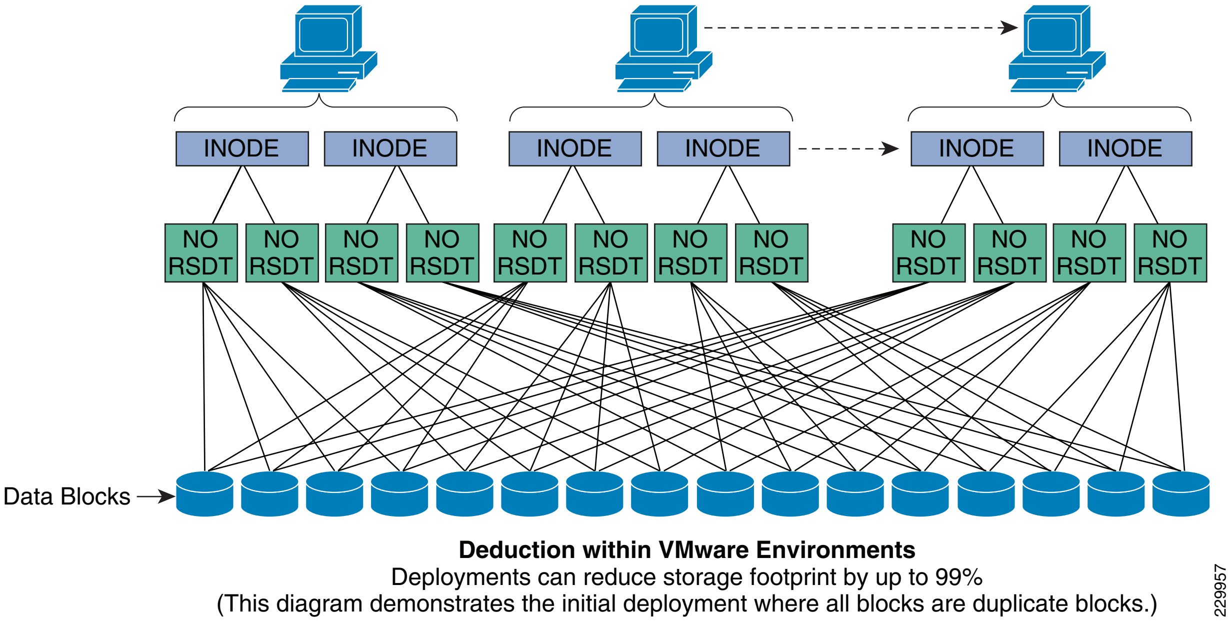
INODE
INODE
INODE
INODE
INODE
INODE
NO RSDT
NO RSDT
NO RSDT
NO RSDT
NO RSDT
NO RSDT
NO RSDT
NO RSDT
NO RSDT
NO RSDT
NO RSDT
NO RSDT
Data Blocks
Deduction within VMware Environments
Deployments can reduce storage footprint by up to 99%
(This diagram demonstrates the initial deployment where all blocks are duplicate blocks.)
229957

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 non-dense 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.

Data Protection via RAID-DP

With any VMware View 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.

Performance

Another critical barrier to VDI 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 VDI using a shared storage infrastructure, significant performance issues might arise during these critical operations. This means the solution could 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. These solution components save customers money by:

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

Transparent Storage Cache Sharing

Transparent storage cache sharing (TSCS) allows customers to benefit from NetApp's storage efficiency and at the same time significantly increase I/O performance. TSCS is natively built into the Data ONTAP operating system and works by 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. VDI 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 18W/hr. and generates 90BTU/hr. By not deploying an additional 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 one disk shelf.

Flash Cache

NetApp Flash Cache 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, see:
http://www.netapp.com/us/products/storage-systems/flash-cache/flash-cache-tech-specs.html.

FlexScale

FlexScale is the tunable software component to Flash Cache. It is a licensed feature of Data ONTAP 7.3 or greater. FlexScale allows different caching modes to be used based on the type of workload. The different modes of caching are metadata only, normal user data, and low-priority blocks. FlexScale will allow system administrators to tune their NetApp controllers for VMware View environments.

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.

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. For more information on RAID DP, see NetApp TR-3298: RAID-DP: NetApp Implementation of RAID Double Parity for Data Protection.

Flexible Volumes and Aggregates

Aggregates are NetApp's virtualization layer, which abstracts physical disks from logical datasets, which are referred to as flexible volumes (also known as FlexVol volumes). Aggregates offer pooled disk resources including IOPS to storage administrators. The flexible volumes each aggregate contains has its own unique logical capacity. Flexible volumes can be thin-provisioned and the logical capacity resized as needed by the storage administrator.

FlexVols are shared out as file level (NFS or CIFS) mount points or are further allocated as LUNs for block level (iSCSI or FCP) access. FlexVols can be readily offered to the VMware environment as datastores; it is recommended that there be a one-to-one alignment between the FlexVol and ESX datastore assets to provide an easy mapping between the VMware server administrator's view and the storage administrator's view of the virtual environment. NetApp also suggests large aggregates to support most VMware environments. VMware environments often have random I/O requirements; large aggregates provide maximum flexibility to VMware administrators as a large pool of I/O resources is readily made available. With NetApp's inherent storage virtualization techniques, all data sets or virtual machines housed within a shared storage infrastructure take advantage of RAID-DP from a performance and protection standpoint.

NetApp Operations Manager

Implementation and management complexities associated with deploying a VDI 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 VDI implementation.

NetApp Operations Manager provides a comprehensive monitoring and management solution for the VDI 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, see: http://www.netapp.com/us/products/management-software/operations-manager.html.

Figure 18 NetApp Operations Manager

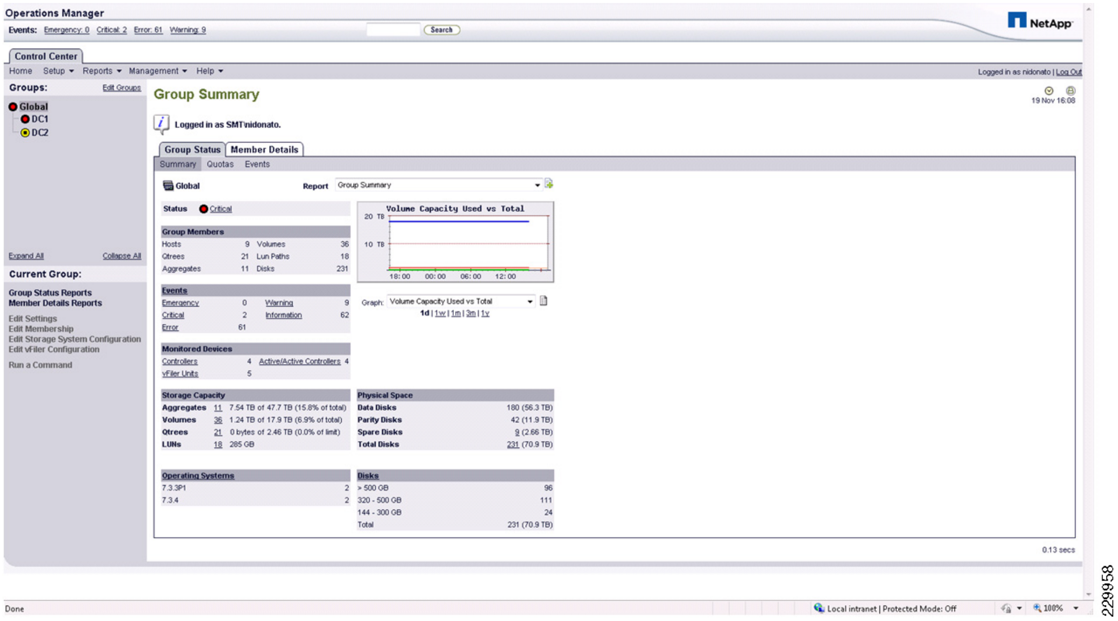

Group Summary

Logged in as SMT\nidonato.

Global

Report: Group Summary

Graph: Volume Capacity Used vs Total

Status: Critical

Group Members			
Hosts	9	Volumes	36
Qtrees	21	Lun Paths	18
Aggregates	11	Disks	231

Events			
Emergency	0	Warning	9
Critical	2	Information	62
Error	61		

Monitored Devices			
Controllers	4	Active/Active Controllers	4
vFiler Units	5		

Storage Capacity		
Aggregates	11	7.54 TB of 47.7 TB (15.8% of total)
Volumes	36	1.24 TB of 17.9 TB (6.9% of total)
Qtrees	21	0 bytes of 2.46 TB (0.0% of limit)
LUNs	18	285 GB

Physical Space	
Data Disks	180 (56.3 TB)
Parity Disks	42 (11.9 TB)
Spare Disks	9 (2.66 TB)
Total Disks	231 (70.9 TB)

Operating Systems	
7.3.3P1	2
7.3.4	2

Disks	
> 500 GB	96
320 - 500 GB	111
144 - 300 GB	24
Total	231 (70.9 TB)

Storage Architecture Best Practices

In a VDI environment, the availability and performance of the storage infrastructure is critical because thousands of users will be affected by storage outages or performance issues. The storage architecture must provide a high level of availability and performance. For detail deployment guidance for VMware View, and highly resilient storage designs, see:

- NetApp TR-3770: VMware View on NetApp Deployment Guide
- NetApp TR-3437: Storage Best Practices and Resiliency Guide
- NetApp TR-3450: Active-Active Controller Overview and Best Practices Guidelines
- NetApp TR-3749: NetApp and VMware vSphere Storage Best Practices
- Netapp TR-3705: Netapp and VMware View Solution Guide

Cisco Networking Infrastructure

Cisco Nexus 5020 28-Port Switch

The Cisco Nexus 5020 Switch is a 1RU, 10 Gigabit Ethernet/FCoE access layer switch built to provide more than 1 Terabit per second (Gbps) throughput with very low latency. It has 40 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 six additional 10 Gigabit Ethernet/FCoE ports, up to eight Fibre Channel ports, or a combination of both. The switch has a single serial console port and a single out-of-band 10/100/1000-Mbps Ethernet management port. Two N+1 redundant, hot-pluggable power supplies and five N+1 redundant, hot-pluggable fan modules provide highly reliable front-to-back cooling.

Cisco Nexus 5000 Series Feature Highlights

Features and Benefits

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

- The combination of high port density, wire-speed performance, and extremely low latency makes the switch an ideal product to meet the growing demand for 10 Gigabit Ethernet at the rack level. The switch family has sufficient port density to support single or multiple racks fully populated with blade and rack-mount servers.
- Built for today's data centers, the switches are designed just like the servers they support. Ports and power connections are at the rear, closer to server ports, helping keep cable lengths as short and efficient as possible. Hot-swappable power and cooling modules can be accessed from the front panel, where status lights offer an at-a-glance view of switch operation. Front-to-back cooling is consistent with server designs, supporting efficient data center hot- and cold-aisle designs. Serviceability is enhanced with all customer-replaceable units accessible from the front panel. The use of SFP+ ports offers increased flexibility to use a range of interconnect solutions, including copper for short runs and fiber for long runs.
- Fibre Channel over Ethernet and IEEE Data Center Bridging features supports I/O consolidation, eases management of multiple traffic flows, and optimizes performance. Although implementing SAN consolidation requires only the lossless fabric provided by the Ethernet pause mechanism, the Cisco Nexus 5000 Series provides additional features that create an even more easily managed, high-performance, unified network fabric.

10 Gigabit Ethernet and Unified Fabric Features

The Cisco Nexus 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.

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), 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, see: http://www.cisco.com/en/US/products/ps9670/prod_white_papers_list.html.

Microsoft Windows 7

Microsoft introduced Windows 7 in the Autumn 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 the IDC report *Deployment Opportunities for Windows 7*).

Microsoft Windows 7 Golden Image Creation and Provisioning

Microsoft Windows 7 can be provisioned for View 4.5 with two methods:

- The traditional guest OS install and application configuration.
- Using the Microsoft Deployment Toolkit (MDT).

Each of these methods provide different optimization modes and configurations. Detailed step-by-step configuration information for both methods can be found at: http://www.vmware.com/files/pdf/VMware-View-OptimizationGuideWindows7-EN.pdf. This paper describes the two methods in detail and explains optimization techniques for each method.

Solution Validation

This section discusses the preparation for the creation of the validation environment.

Topology Configuration for Scalability of VMware View 4.5 on Cisco Unified System and NetApp Storage

The View scalability testing was conducted as a tenant on a buildout of the Enhanced Secure Multi-Tenancy (ESMT) architecture: http://www.cisco.com/en/US/solutions/ns340/ns414/ns742/ns743/ns1050/landing_dcVDDC.html.

Figure 19 Enhanced Secure Multi-Tenancy Components

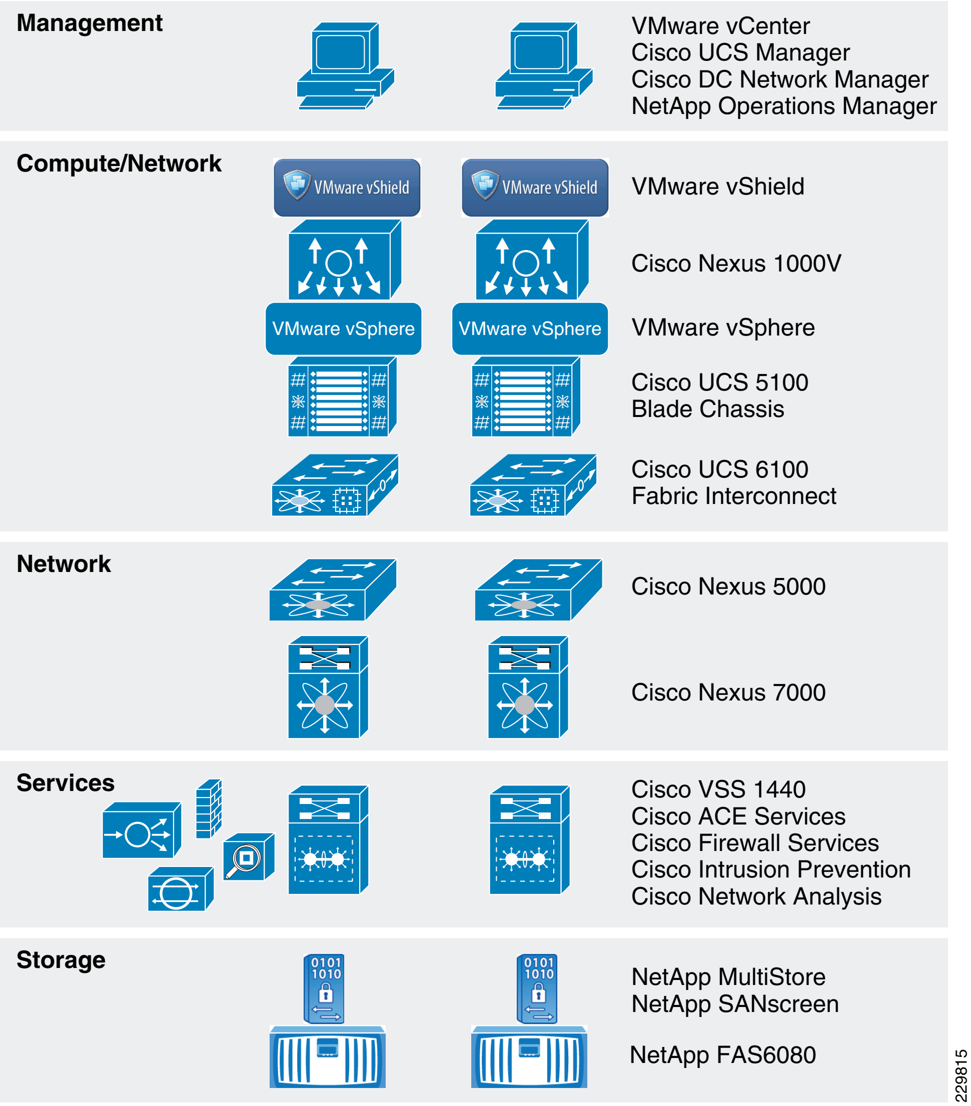

The Enhanced Secure Multi-Tenancy architecture was created through a design partnership between Cisco, NetApp, and VMware to provide secure isolation and high performance within an IaaS setting.

A detailed architectural diagram is shown in Figure 20 with all the interconnections.

Figure 20 Enhanced Secure Multi-Tenancy Architecture

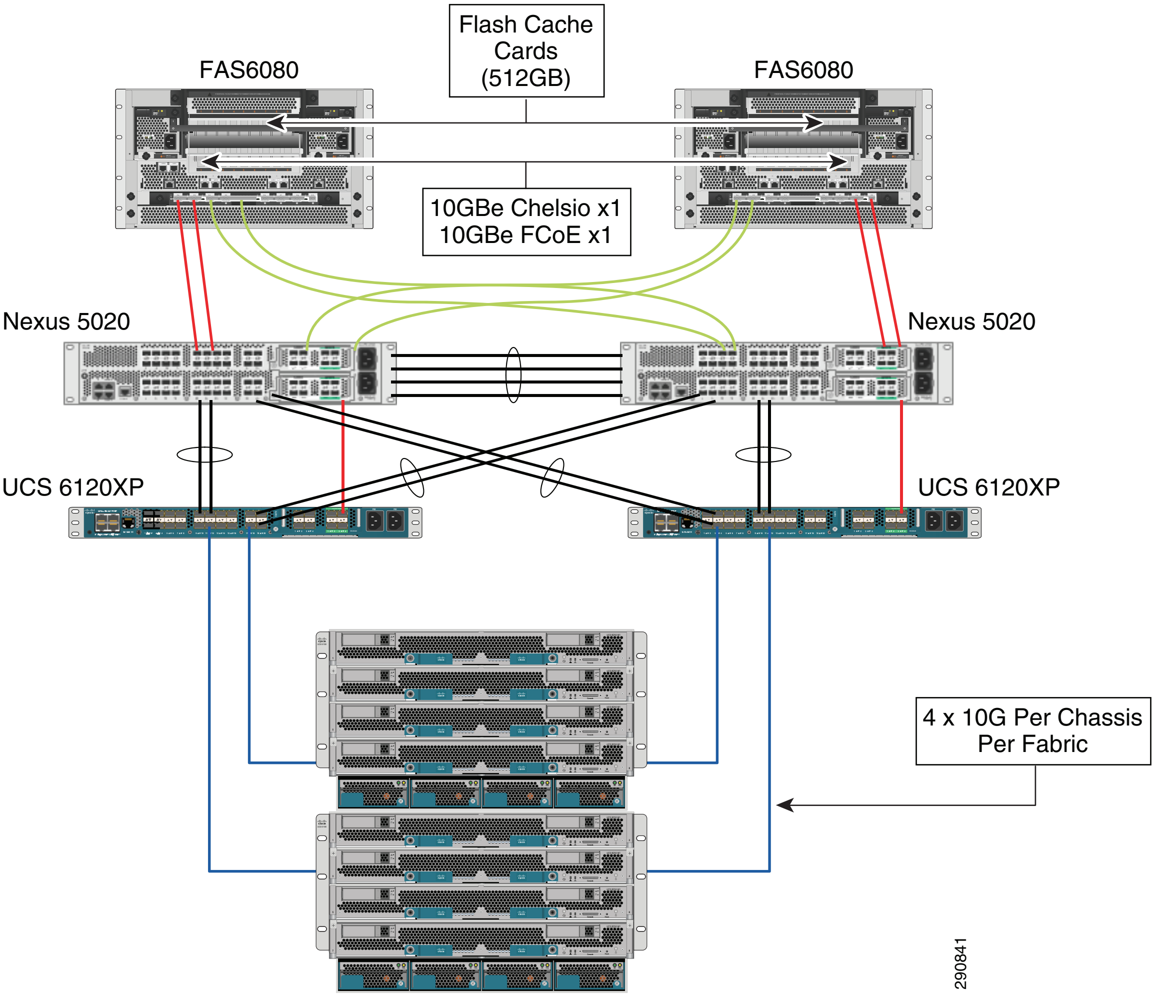

Cisco UCS Configuration

This section discusses 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 are beyond the scope of this document. More details on each step can be found in:

- 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

The configuration of the Cisco UCS Service Profiles was implemented following the steps detailed in the Secure Multi-Tenancy Deployment Guide: http://www.cisco.com/en/US/docs/solutions/Enterprise/Data_Center/Virtualization/securecldeployg.html#wp99931. Figure 21 shows the high level flow of the configuration process.

Figure 21 UCS Manager High Level Flow

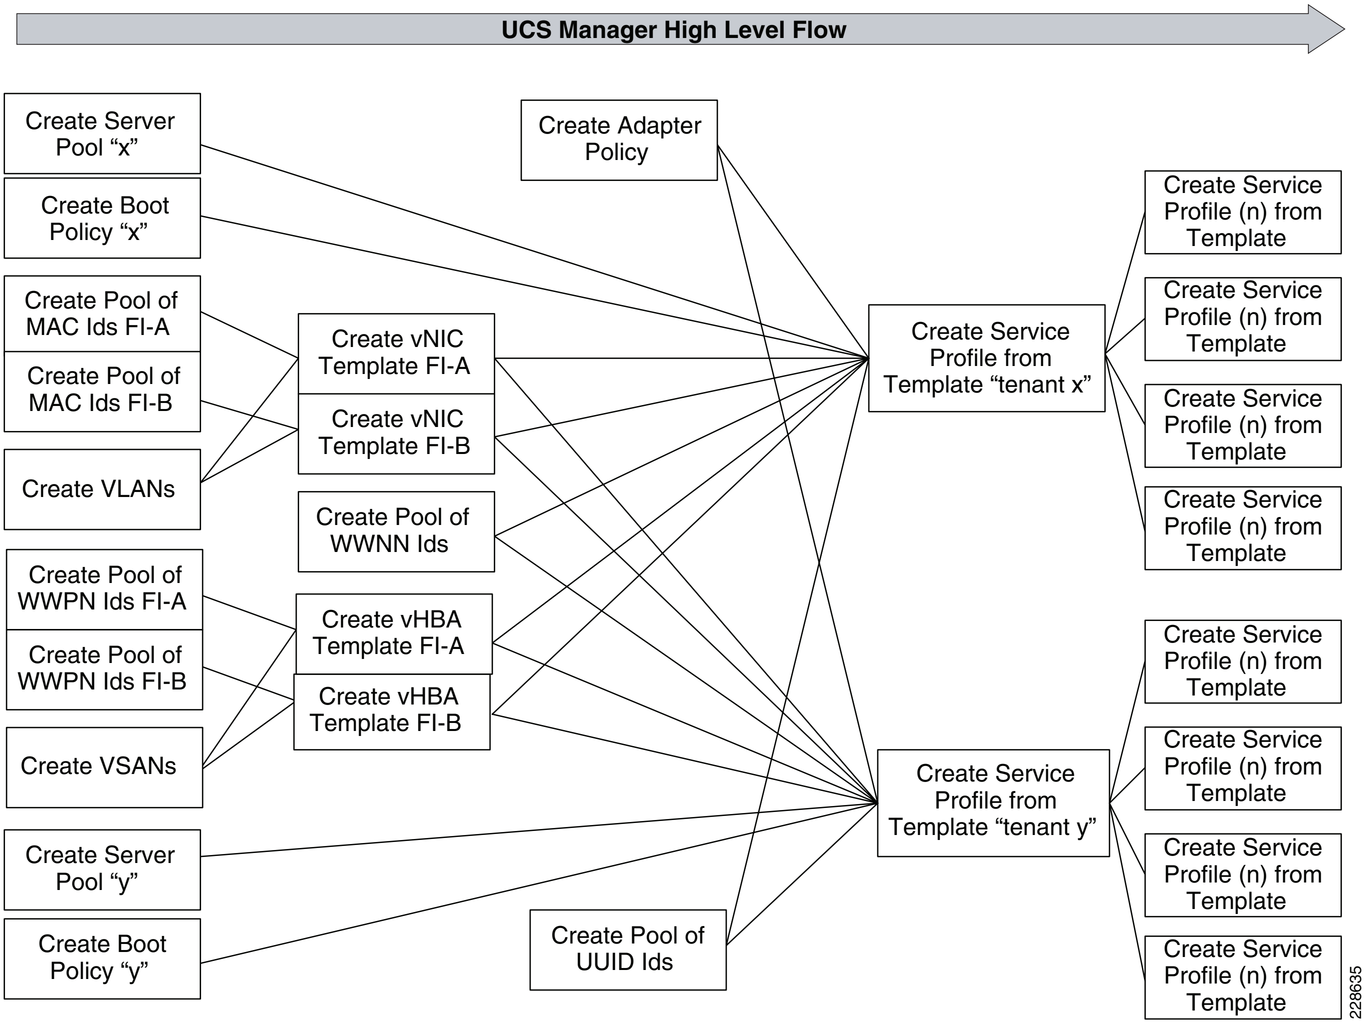
UCS Manager High Level Flow
Create Server Pool "x"
Create Boot Policy "x"
Create Pool of MAC Ids FI-A
Create Pool of MAC Ids FI-B
Create VLANs
Create Pool of WWPN Ids FI-A
Create Pool of WWPN Ids FI-B
Create VSANs
Create Server Pool "y"
Create Boot Policy "y"
Create vNIC Template FI-A
Create vNIC Template FI-B
Create Pool of WWNN Ids
Create vHBA Template FI-A
Create vHBA Template FI-B
Create Adapter Policy
Create Pool of UUID Ids
Create Service Profile from Template "tenant x"
Create Service Profile from Template "tenant y"
Create Service Profile (n) from Template
Create Service Profile (n) from Template
Create Service Profile (n) from Template
Create Service Profile (n) from Template
Create Service Profile (n) from Template
Create Service Profile (n) from Template
Create Service Profile (n) from Template
Create Service Profile (n) from Template
228635

The UCS Service Profile allows hardware to be presented in a stateless manner that is completely transparent to the OS and the applications that run on it. The Service Profile boots from a LUN that is tied to the WWPN specified, allowing an installed OS instance to be locked with the Service Profile. The independence from server hardware allows installed systems to be re-deployed between blades. Through the use of pools and templates, UCS hardware can be deployed quickly to scale.

QoS and CoS in Cisco Unified Computing System

Cisco UCS provides different system classes of service to implement QoS 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.

Voice, video, and other time-sensitive applications experience optimal performance where QoS policies are enforced across the fabric.

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 versus no-drop)
- Max MTU per Class allowed is 9216.
- Through the Cisco UCS we can map one CoS value to a particular class.
- Apart from FcoE class there can be only one more class configured as no-drop property.
- Weight can be configured based on 0 to 10 numbers. Internally the system will calculate the bandwidth based on the following equation (there will be rounding off of the number):

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

Cisco UCS System Class Configuration

Cisco UCS defines user class names as follows.

- Platinum
- Gold
- Silver
- Bronze

Table 1 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

Table 2 Class to CoS Map by Default in Cisco Unified Computing System

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

Table 3 Default Weight in Cisco Unified Computing System

Cisco UCS Class Names	Weight
Best effort	5
Fc	5

To enable QoS on the Cisco UCS:

Step 1 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.

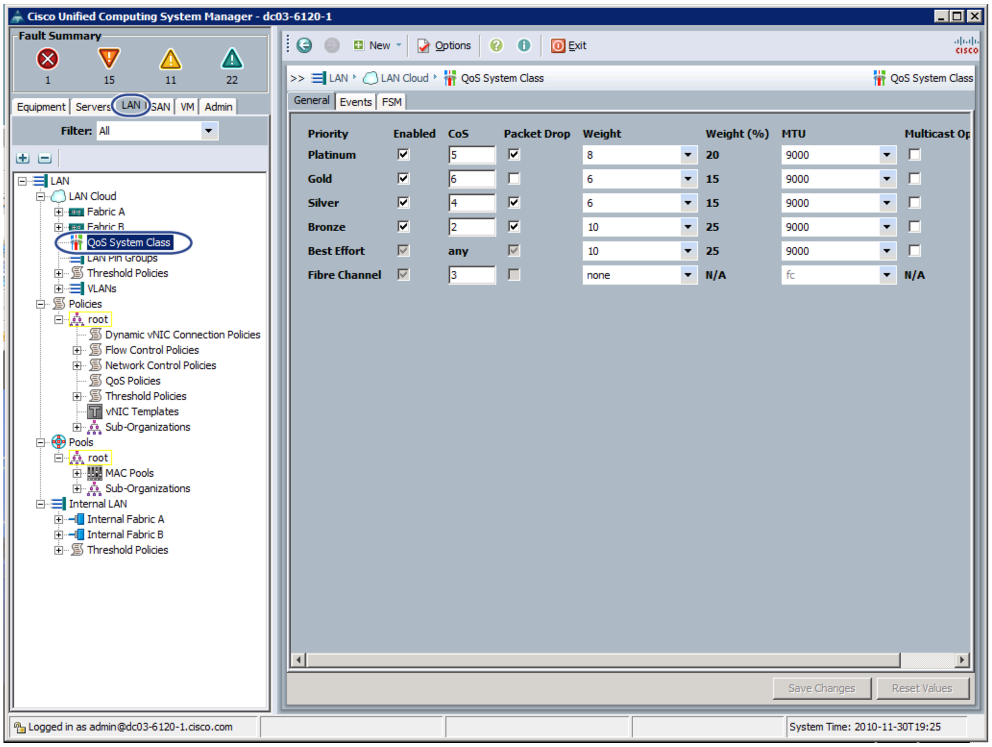

Priority	Enabled	CoS	Packet Drop	Weight	Weight (%)	MTU	Multicast Op
Platinum	☑	5	☑	8	20	9000	□
Gold	☑	6	□	6	15	9000	□
Silver	☑	4	☑	6	15	9000	□
Bronze	☑	2	☑	10	25	9000	□
Best Effort	☑	any	☑	10	25	9000	□
Fibre Channel	☑	3	□	none	N/A	fc	N/A

229973

Step 2 In the LAN tab under policies, define a platinum-policy and select Platinum as the priority.

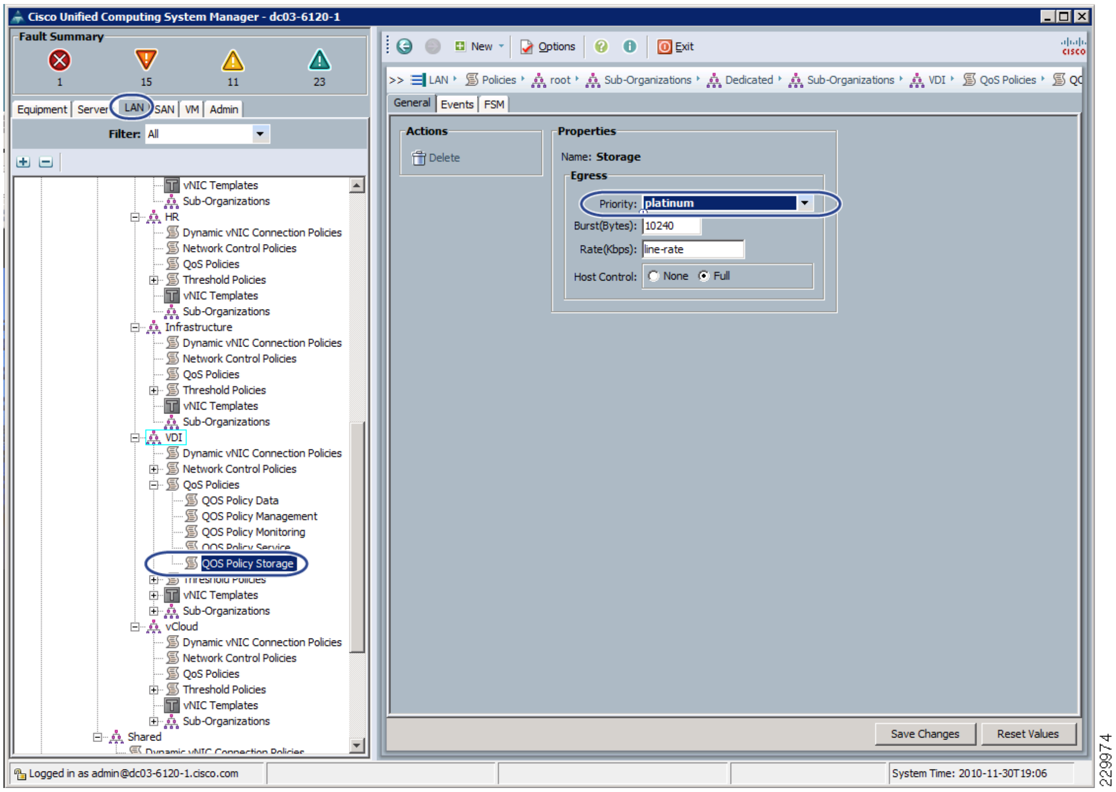

Name: Storage

Priority: platinum

Burst(Bytes): 10240

Rate(Kbps): line-rate

Host Control: ○ None ● Full

Step 3 Include this policy into the vNIC template under the QoS policy.

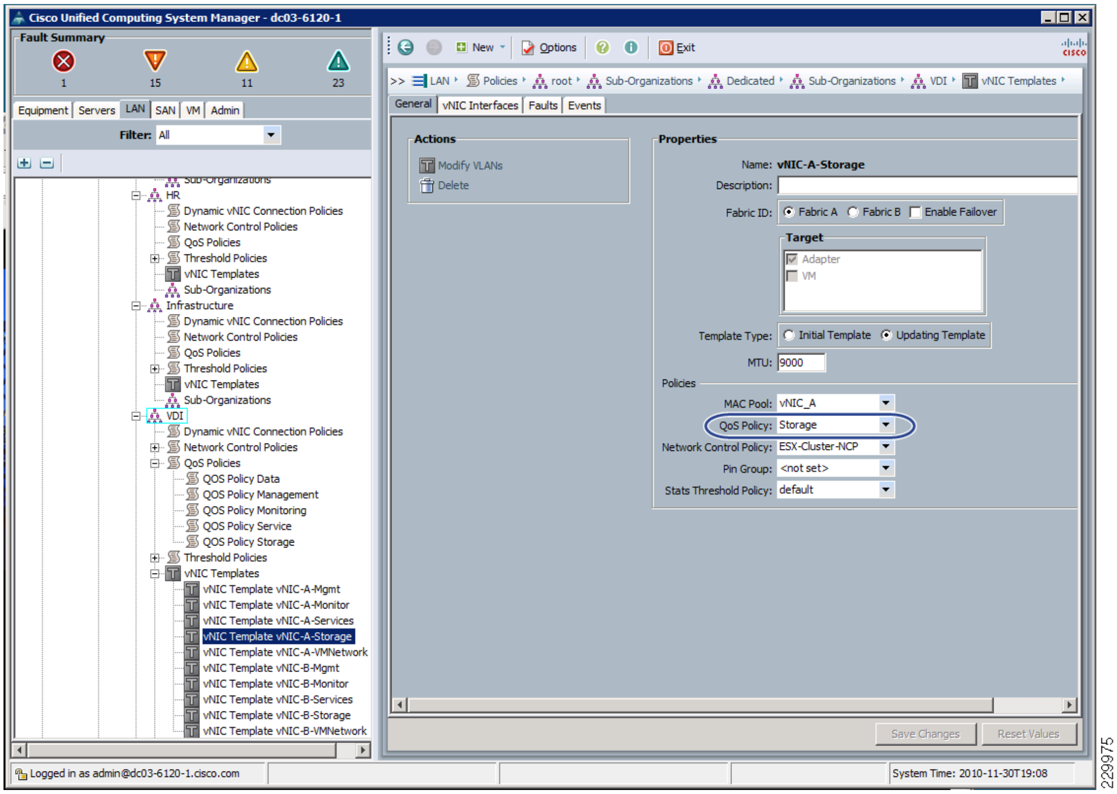

Name: vNIC-A-Storage

Description:

Fabric ID: ● Fabric A ○ Fabric B □ Enable Failover

Target

☑ Adapter

□ VM

Template Type: ○ Initial Template ● Updating Template

MTU: 9000

Policies

MAC Pool: vNIC_A

QoS Policy: Storage

Network Control Policy: ESX-Cluster-NCP

Pin Group: <not set>

Stats Threshold Policy: default

Logged in as admin@dc03-6120-1.cisco.com

System Time: 2010-11-30T19:08

This is a unique value proposition of the Cisco UCS 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.

VMware View 4.5 Configuration

The topology of the network in the test environment is shown in Figure 22.

Figure 22 Test Network Layer 3 Topology

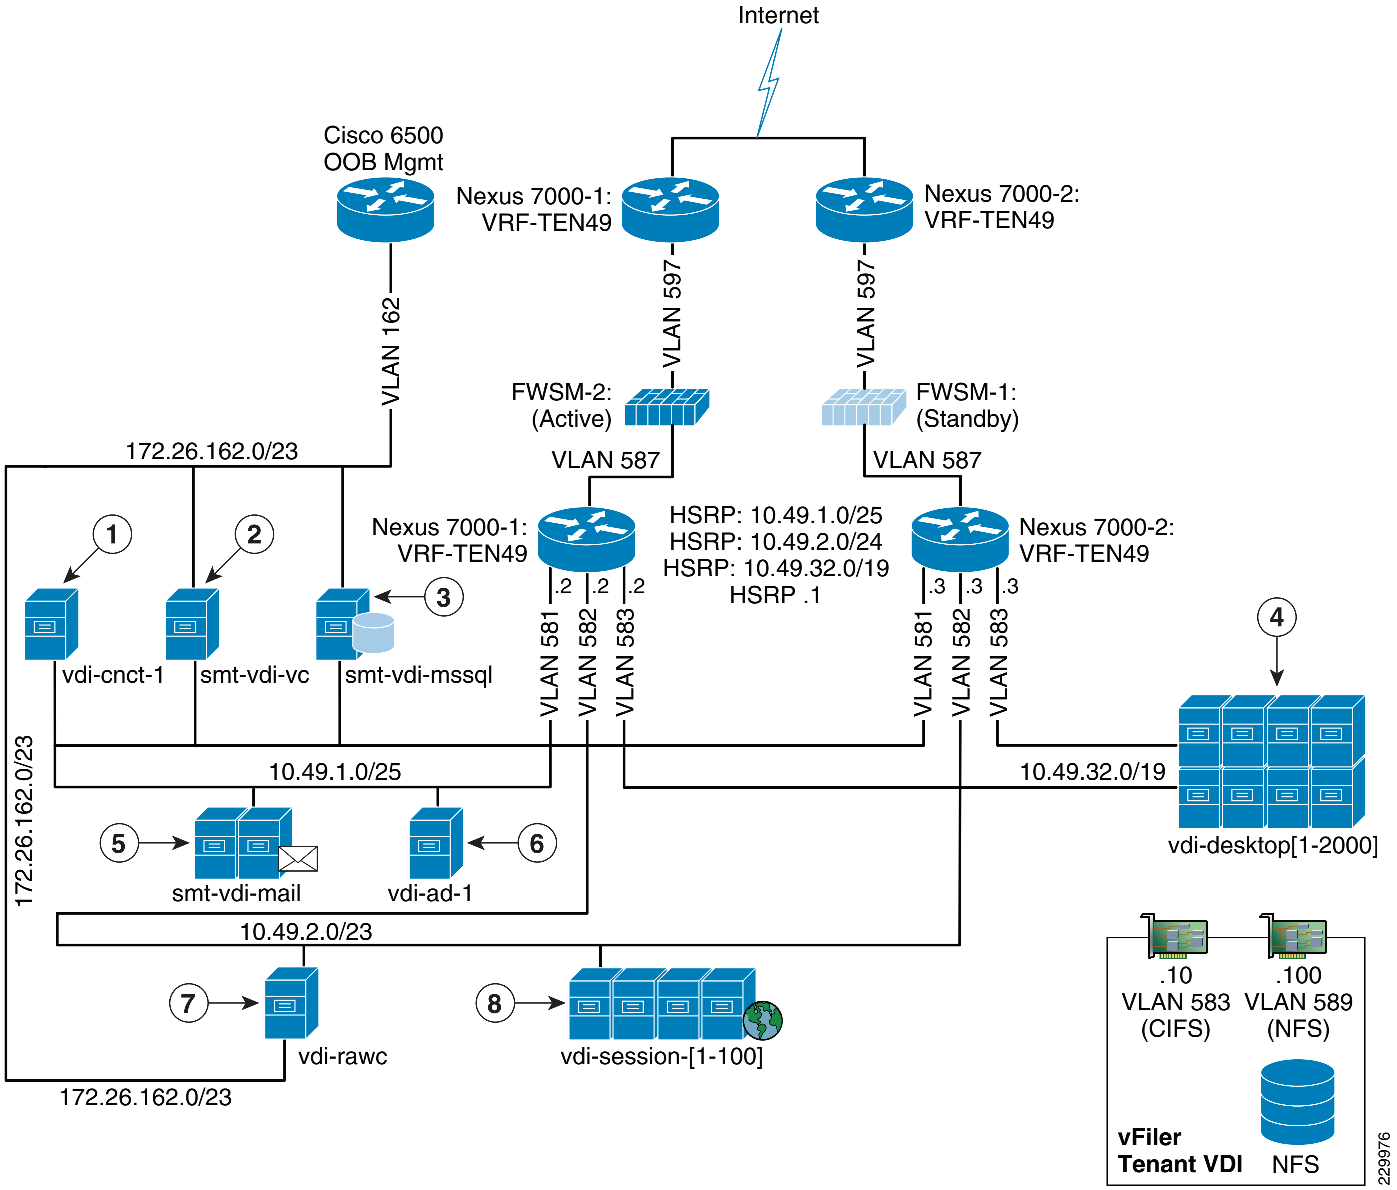

Figure 22 legend:

1. View Connection Server
2. vCenter Server
3. SQL 2008 Server for vCenter, View Composer, and View Events
4. 1 vSphere Cluster holding 1280 Virtual Desktops
5. CentOS IMAP mailserver
6. AD/DHCP server
7. RAWC Controller
8. RAWC Session Launchers

The component configurations are shown in Table 4 through Table 12.

Table 4 VMware vSphere 4.1

vSphere 4.1 ESXi 260247			
Hardware:	Cisco UCS B-series Blade server	**Model:**	B250 –M2
OS:	VMware ESXi 4.1.0	**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 5 VMware vCenter 4.1

VMware vSphere vCenter 4.1.0 Build 258902			
OS:	Windows 2008 Enterprise R2 64bit	**Service Pack:**	-
CPU:	2 x vCPU	**RAM:**	4096MB
Disk:	2 x 40GB Virtual Disk	**Network:**	2 x 10GbE (VMXNET3)
View Composer 2.5.0-291081			

Table 6 VMware vCenter, View, VUM Database Server

MSSQL 2008 R2			
OS:	Windows 2008 Enterprise R2 64bit	**Service Pack:**	-
CPU:	2 x vCPU	**RAM:**	4096MB
Disk:	2 x 40GB Virtual Disk	**Network:**	1 x 10GbE (VMXNET3)

Table 7 VMware View Connection Server

VMware View 4.5 Connection Server			
OS:	Windows 2008 Enterprise R2 64bit	**Service Pack:**	-
CPU:	2 x vCPU	**RAM:**	8192MB
Disk:	1 x 40GB Virtual Disk	**Network:**	1 x 10GbE (VMXNET3)
View Connection Server 4.5.0-293049			

Table 8 VMware View Desktop Agent

VMware View Desktop Agent (Virtual Desktops)			
OS:	Windows7 Enterprise 32bit	**Service Pack:**	1
CPU:	1 x vCPU	**RAM:**	1536MB
Disk:	1 x 40GB	**Network:**	1 x 1GbE (VMXNET3)
View Agent 4.5.0-293049			

Table 9 Active Directory/DHCP Server

MS Active Directory/DHCP Server			
OS:	Windows 2008 Enterprise R2 64bit	**Service Pack:**	-
CPU:	2 x vCPU	**RAM:**	4096MB
Disk:	1 x 40GB Virtual Disk	**Network:**	1 x 10GbE (VMXNET3)

Table 10 Mail Server

CentOS 5.3 Dovecot IMAP Server			
OS:	CentOS 5.3_i386	**Service Pack:**	-
CPU:	1 x vCPU	**RAM:**	1024MB
Disk:	1 x 15GB Virtual Disk	**Network:**	1 x 10GbE (Flexible)

Table 11 RAWC Controller

VMware RAWC 1.2			
OS:	Windows 2003 R2 Enterprise 64bit	**Service Pack:**	2
CPU:	4 x vCPU	**RAM:**	2048MB
Disk:	1 x 20GB Virtual Disk	**Network:**	2 x 10GbE (VMXNET3)

Table 12 Session Launcher

VMware RAWC 1.2 Session Launcher			
OS:	Windows XP 32bit	**Service Pack:**	1

Table 12 ***Session Launcher***

CPU:	2 x vCPU	**RAM:**	8192MB
Disk:	1 x 8GB	**Network:**	1 x 1GbE (VMXNET3)

VMware View components were configured following the guidelines specified in the VMware View 4.5 Installation Guide: http://www.vmware.com/pdf/view45_installation_guide.pdf.

LAN Configuration

This configuration consists of a pair of Cisco Nexus 5020s, 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 Interconnects and they are connected to the Nexus 5020 pair in a bow tie manner as shown in Figure 23. The Fabric Interconnect is in End host mode, as we are doing both FC as well as Ethernet data access per the recommended best practice for Cisco UCS. We have provisioned more than 40G per Fabric Interconnect as we are building a scalable and expandable system.

The upstream configuration is beyond the scope of this document, but is discussed in the Secure Multi-Tenancy Deployment Guide: http://www.cisco.com/en/US/docs/solutions/Enterprise/Data_Center/Virtualization/securecldeployg.html#wp998369.

Figure 23 ***Network Configuration with Upstream Cisco Nexus 5000 from Cisco Unified Computing System***

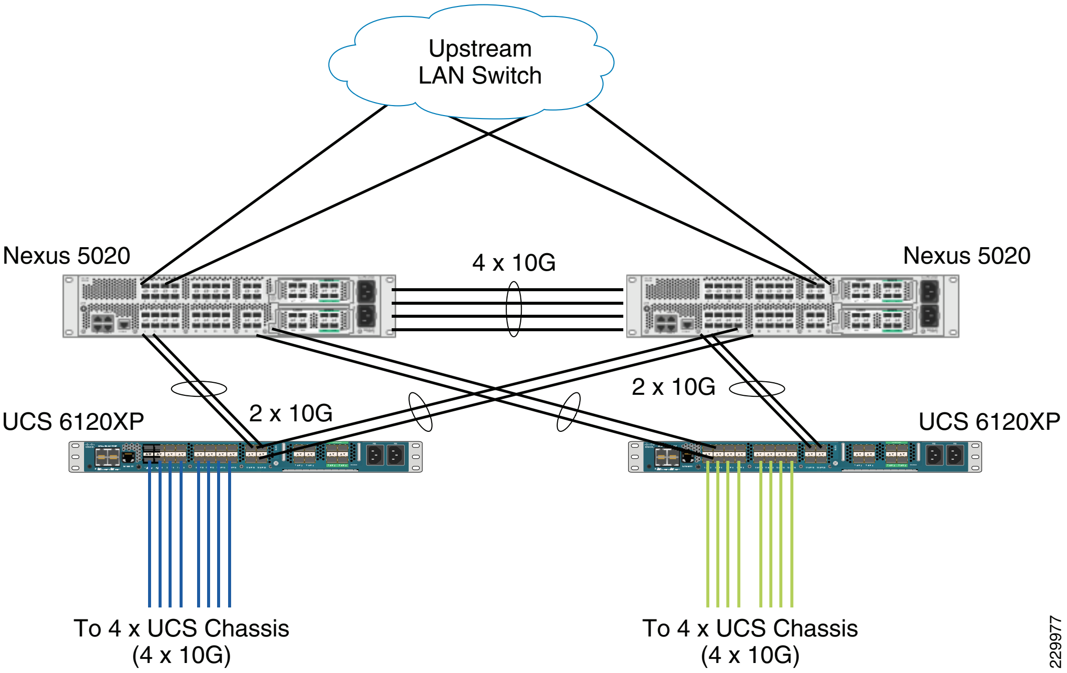

The Cisco Nexus 5000 is used to connect to the NetApp FAS 6080 storage system for NAS access. NetApp supports dual-port 10G Chelsio cards which are configured in a portchannel and connected to the pair of Cisco Nexus 5000s downstream. This allows end-to-end 10G access; we have implemented jumbo frames on the ports and have priority flow control on with platinum CoS for NetApp storage data access. The NetApp connectivity diagram is shown in Figure 24. Again, we have a total of 40G bandwidth available for the servers.

Refer to Appendix A—Select Configurations for a detailed configuration of one of the Cisco Nexus 5000s used in this setup.

Figure 24 Network Configuration for NetApp NAS or Filer Storage

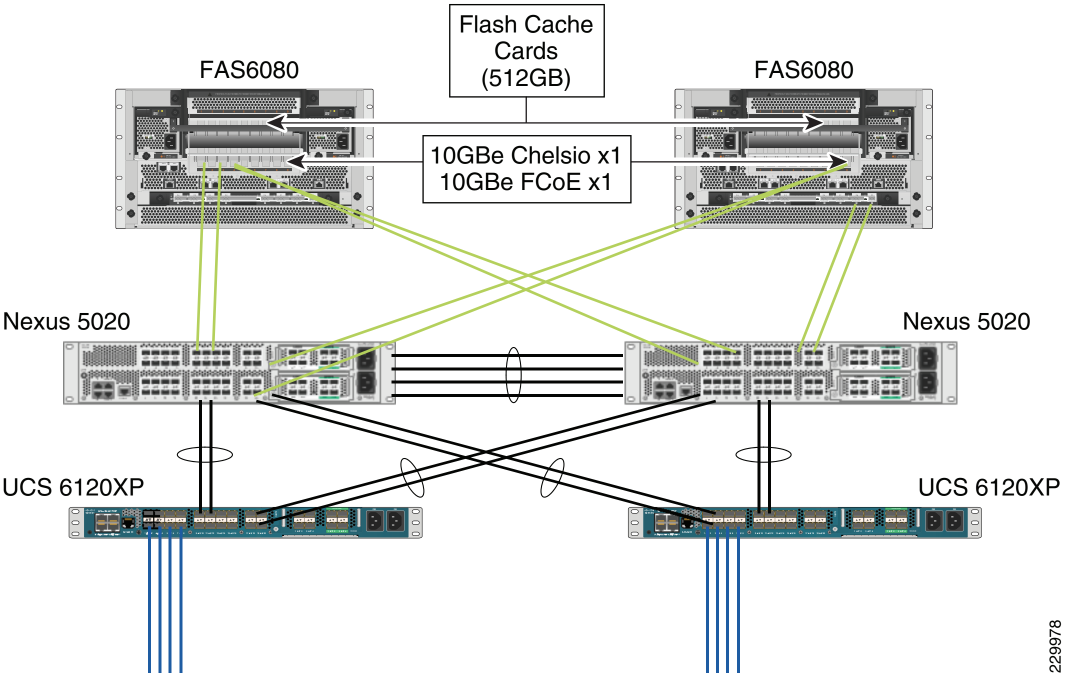

The configuration on the NetApp storage as gathered from the filer view is shown in Figure 25.

Figure 25 Network Configuration on the NetApp Storage Array Side—1

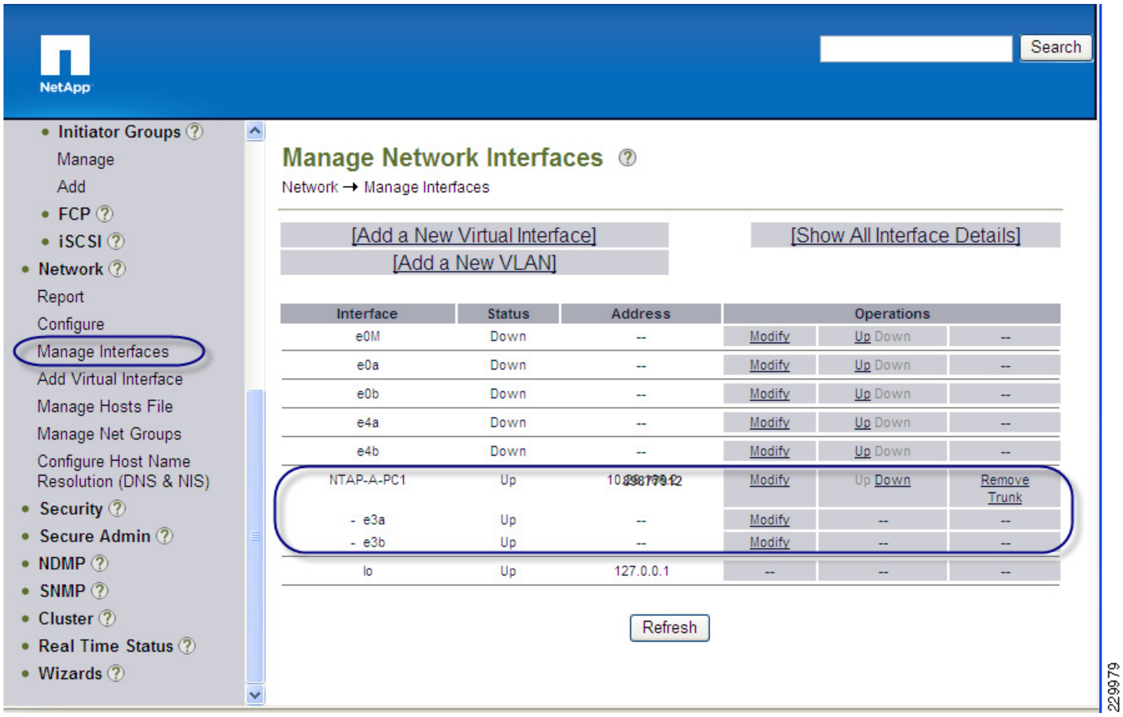

Manage Network Interfaces ?

Network → Manage Interfaces

[Add a New Virtual Interface] [Show All Interface Details]

[Add a New VLAN]

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	[illegible]	Modify	Up Down	Remove Trunk
- e3a	Up	--	Modify	--	--
- e3b	Up	--	Modify	--	--
lo	Up	127.0.0.1	--	--	--

Refresh

Figure 26 Network Configuration on the NetApp Storage Array Side—2

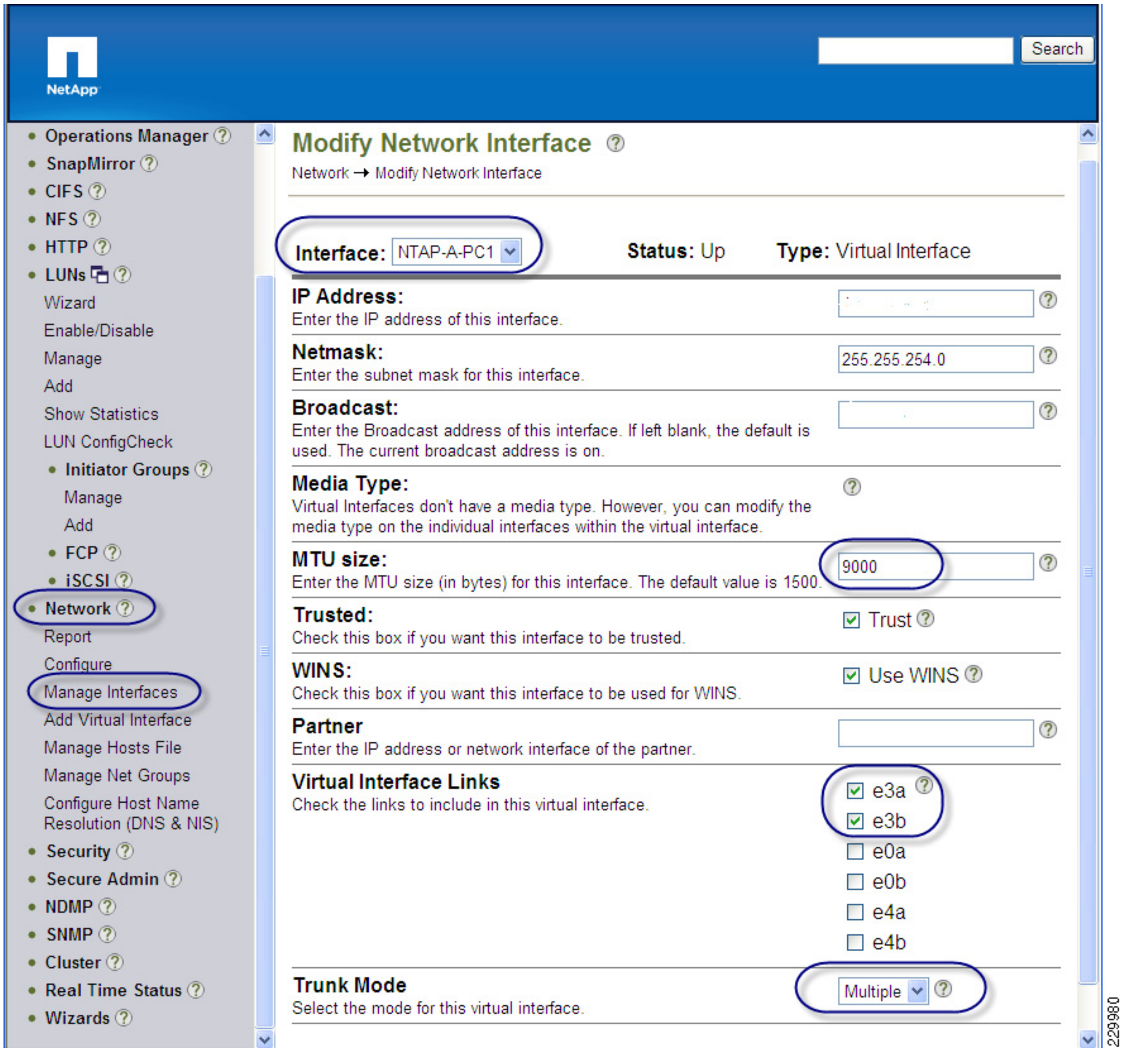

Modify Network Interface

Network → Modify Network Interface

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

IP Address:
Enter the IP address of this interface.

Netmask: 255.255.254.0
Enter the subnet mask for this interface.

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

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

WINS: ☑ Use WINS
Check this box if you want this interface to be used for 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 Multiple
Select the mode for this virtual interface.

SAN Configuration

In this section we discuss stateless environments. A stateful environment is when a blade boots from local disk. The stateless environment starting with boot from SAN allows a more robust fix on a failure state. If a blade fails, you simply swap it out for a new one since the profile is tagged to the slot(s) and the newly inserted hardware will simply take the configuration that the previous failed blade had.

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 Port World Wide Name (pWWN) of the HBAs and the BFS (Boot From SAN) policy would also move 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 are:

- Reduce server footprints—Boot from SAN alleviates the need 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 the functionality of the servers at the primary site, the remote site can take over with minimal downtime.

  Recovery from server failures is simplified in a SAN environment. With the help of snapshots, mirrors of a failed server can be recovered quickly by booting from the original copy of its image. As a result, boot from SAN can greatly reduce the time required for server recovery.
- High availability—A typical data center is highly redundant in nature—redundant paths, redundant disks, and redundant storage controllers. When operating system images are stored on disks in the SAN, it supports high availability and eliminates the potential for mechanical failure of a local disk.
- Rapid redeployment—Businesses that experience temporary high production workloads can take advantage of SAN technologies to clone the boot image and distribute the image to multiple servers for rapid deployment. Such servers may only need to be in production for hours or days and can be readily removed when the production need has been met. Highly efficient deployment of boot images makes temporary server usage a cost effective endeavor.
- Centralized image management—When operating system images are stored on networked disks, all upgrades and fixes can be managed at a centralized location. Changes made to disks in a storage array are readily accessible by each server.

Configuring Boot from SAN on 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 two SAN fabrics are disjoint from a data perspective and, with the dual port HBAs and storage controller, redundancy is provided.

There are three distinct portions of the BFS procedure:

- Storage array configuration
- SAN zone configuration
- 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 a NetApp 6080 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.

Step 1 Create a separate boot from SAN Aggregate.

Step 2 Create a Volume on top of that named BootVolumes.

Step 3 Add LUN on the BootVolumes (named, for example, BFS-Server-9) and 50 GB of space:

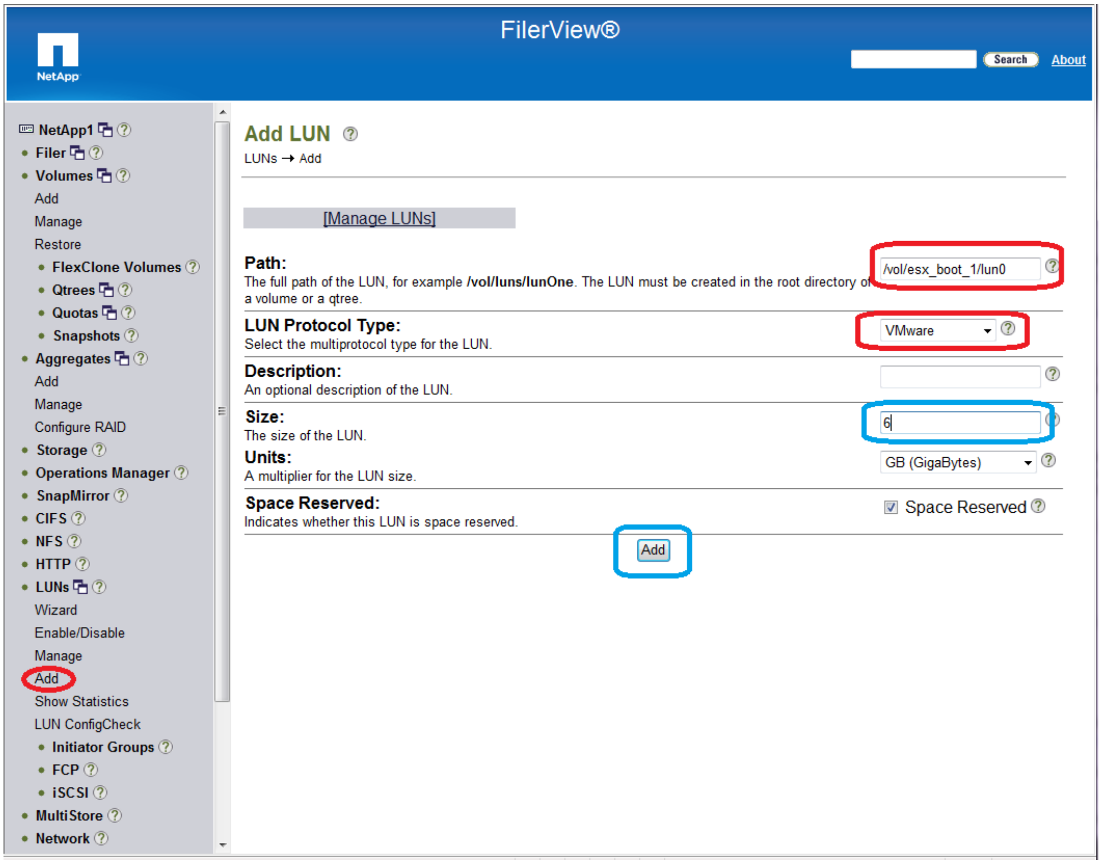

Add LUN

LUNs → Add

[Manage LUNs]

Path: /vol/esx_boot_1/lun0
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: VMware
Select the multiprotocol type for the LUN.

Description:
An optional description of the LUN.

Size: 6
The size of the LUN.

Units: GB (GigaBytes)
A multiplier for the LUN size.

Space Reserved: ☑ Space Reserved
Indicates whether this LUN is space reserved.

Add

Step 4 Now add the LUN to the initiator group:

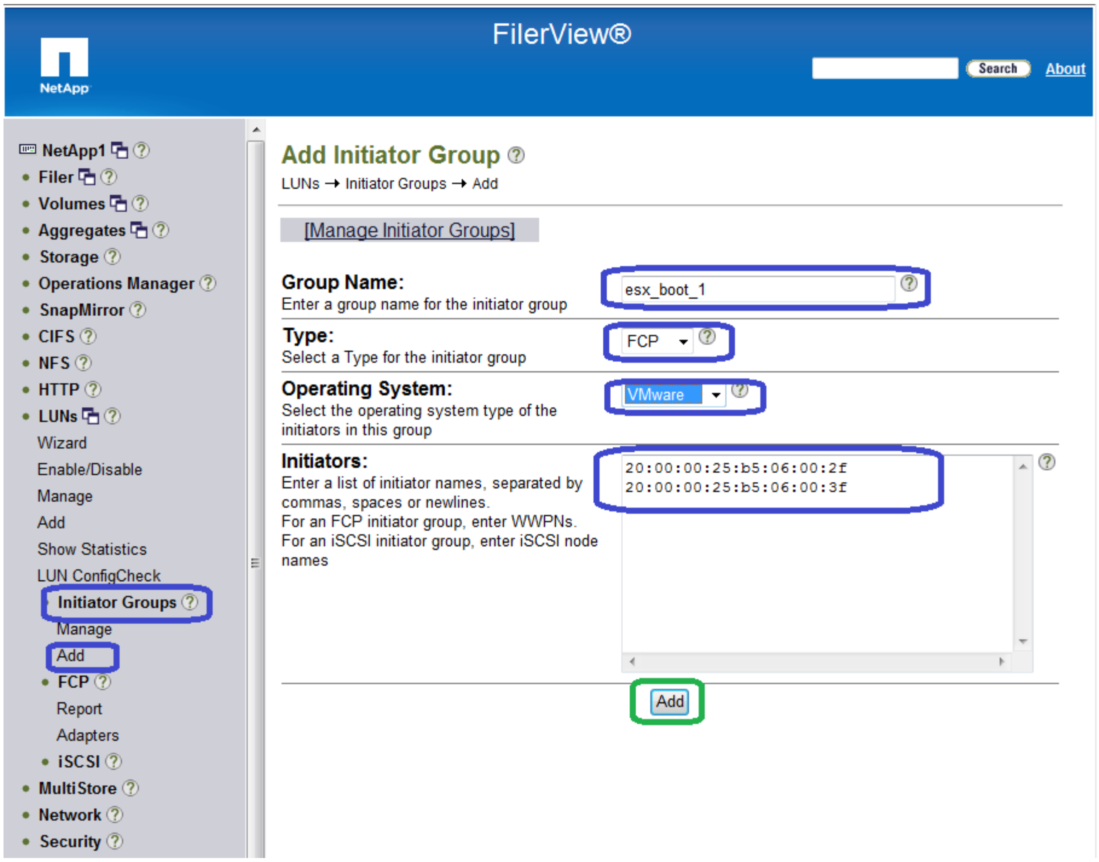

Add Initiator Group

LUNs → Initiator Groups → Add

Group Name: esx_boot_1
Enter a group name for the initiator group

Type: FCP
Select a Type for the initiator group

Operating System: VMware
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

```
20:00:00:25:b5:06:00:2f
20:00:00:25:b5:06:00:3f
```

229982

Step 5 Make sure the add initiator group succeeds:

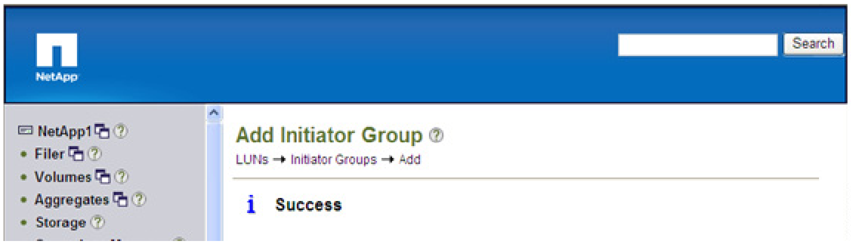

Add Initiator Group

LUNs → Initiator Groups → Add

i Success

229983

Step 6 Next, mask the LUN. Select LUN—>Manage LUN and select the new LUN which needs to be added. Select the “no map” section as shown below:

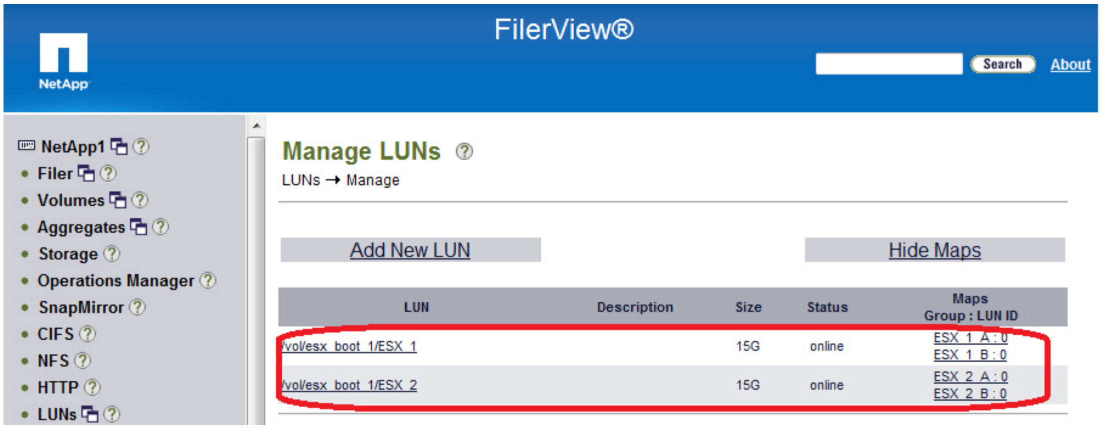

Manage LUNs

LUNs → Manage

Add New LUN | Hide Maps

LUN	Description	Size	Status	Maps Group : LUN ID
/vol/esx_boot_1/ESX_1		15G	online	ESX_1_A : 0 ESX_1_B : 0
/vol/esx_boot_1/ESX_2		15G	online	ESX_2_A : 0 ESX_2_B : 0

229984

Step 7 Add the group to the map:

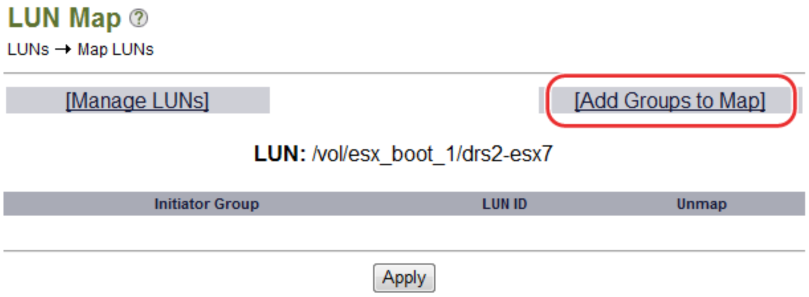

LUN Map

LUNs → Map LUNs

LUN: /vol/esx_boot_1/drs2-esx7

Initiator Group	LUN ID	Unmap

229985

Step 8 Select the new initiator group, drs2-esx7, and click **Add**:

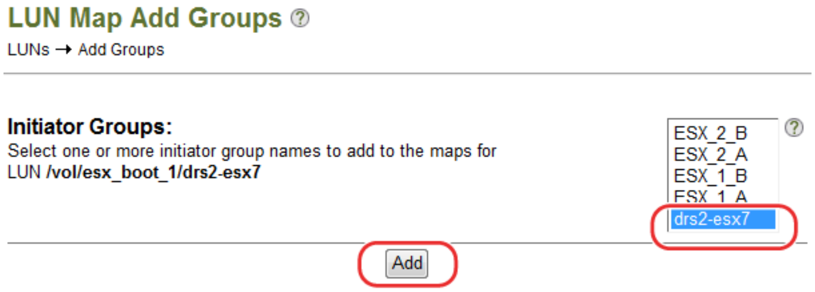

LUN Map Add Groups

LUNs → Add Groups

Initiator Groups:
Select one or more initiator group names to add to the maps for LUN **/vol/esx_boot_1/drs2-esx7**

ESX_2_B
ESX_2_A
ESX_1_B
ESX_1_A
drs2-esx7

229986

Step 9 Now assign a LUN ID to the initiator group:

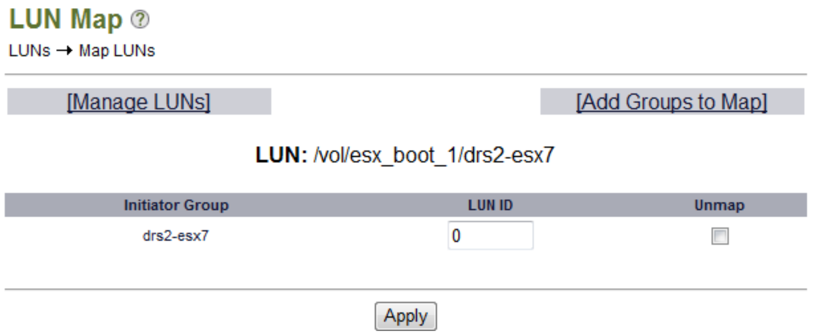

LUN: /vol/esx_boot_1/drs2-esx7

Initiator Group	LUN ID	Unmap
drs2-esx7	0	□

229987

Step 10 Make sure the mapping succeeded:

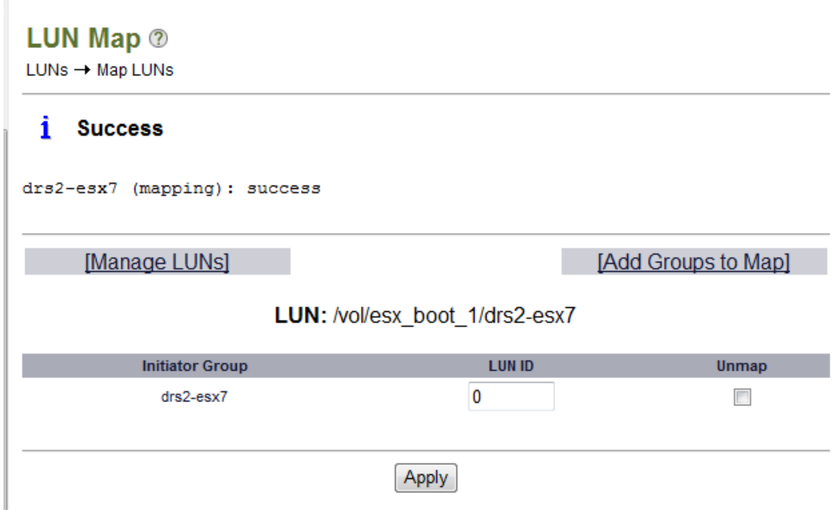

ℹ **Success**

```
drs2-esx7 (mapping): success
```

LUN: /vol/esx_boot_1/drs2-esx7

Initiator Group	LUN ID	Unmap
drs2-esx7	0	□

229988

Step 11 After the LUN map is successfully updated, check to see if the Manage LUNs show a correct mapping:

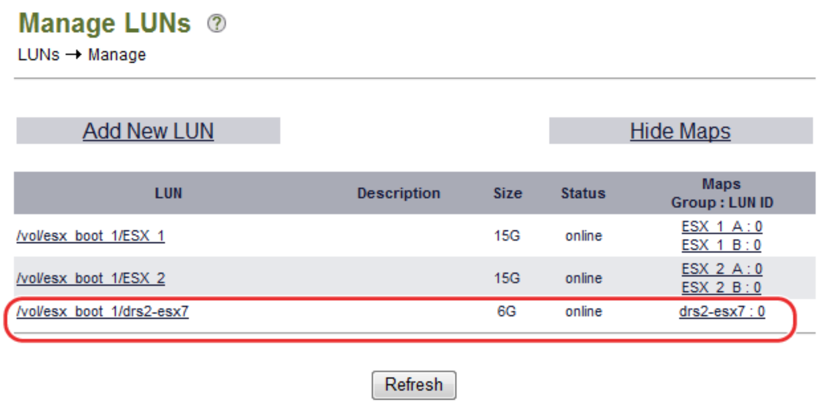

LUN	Description	Size	Status	Maps Group : LUN ID
/vol/esx_boot_1/ESX_1		15G	online	ESX_1_A : 0 ESX_1_B : 0
/vol/esx_boot_1/ESX_2		15G	online	ESX_2_A : 0 ESX_2_B : 0
/vol/esx_boot_1/drs2-esx7		6G	online	drs2-esx7 : 0

229989

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

SAN Zone Configuration

The NPIV feature has to be turned on in the SAN switch. Also make sure you have 4 GB SPF+ modules connected to the Cisco UCS 6120 XP Fabric Interconnect. The port mode is set to AUTO and the speed is set to AUTO. VSAN configuration can be done either in the SAN switch CLI or the GUI (Cisco Nexus 5020 Device Manager). Cisco Fabric Manager can also be used to get an overall picture of the SAN configuration and zoning information. As discussed earlier, SAN zoning is done upfront for all the pWWN of the initiators with the NetApp target pWWN.

Cisco UCS Manager Configuration

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

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

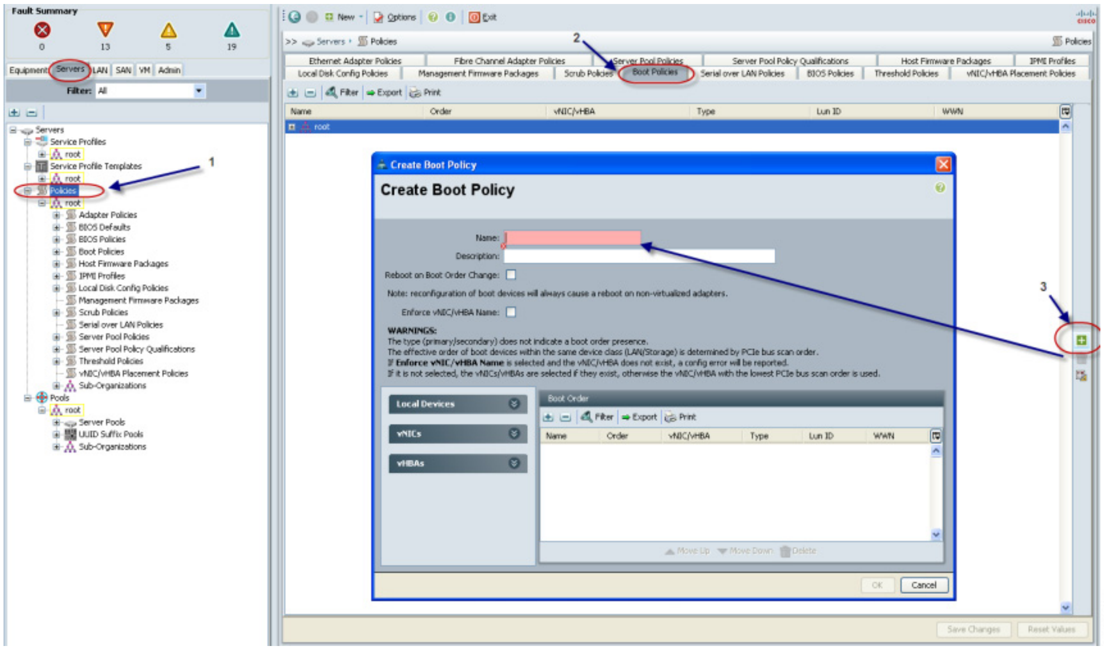

Create Boot Policy

Name:
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 is used.

Step 2 Add SAN boot for primary. The vHBA is optional; it could be left blank and we do not have to enforce the vHBA name:

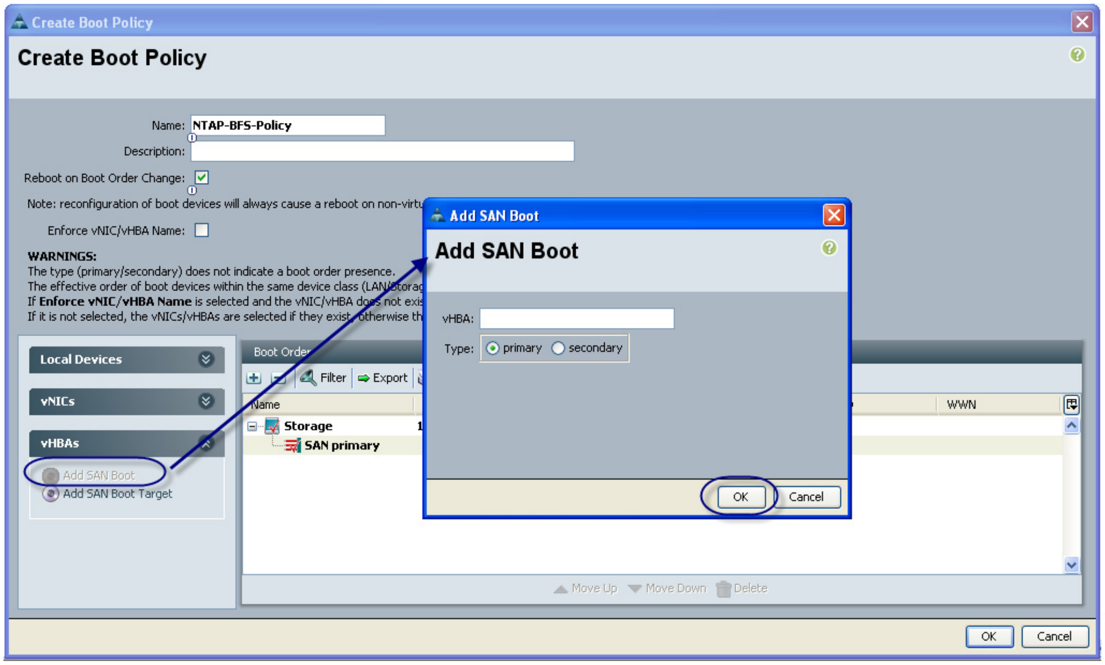

Create Boot Policy

Name: NTAP-BFS-Policy

Description:

Reboot on Boot Order Change: ☑

Note: reconfiguration of boot devices will always cause a reboot on non-virtu

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/Storag
If **Enforce vNIC/vHBA Name** is selected and the vNIC/vHBA does not exis
If it is not selected, the vNICs/vHBAs are selected if they exist, otherwise th

Add SAN Boot

vHBA:

Type: ● primary ○ secondary

Name		WWN
Storage	1	
SAN primary		

229991

Step 3 Add SAN boot for SAN secondary:

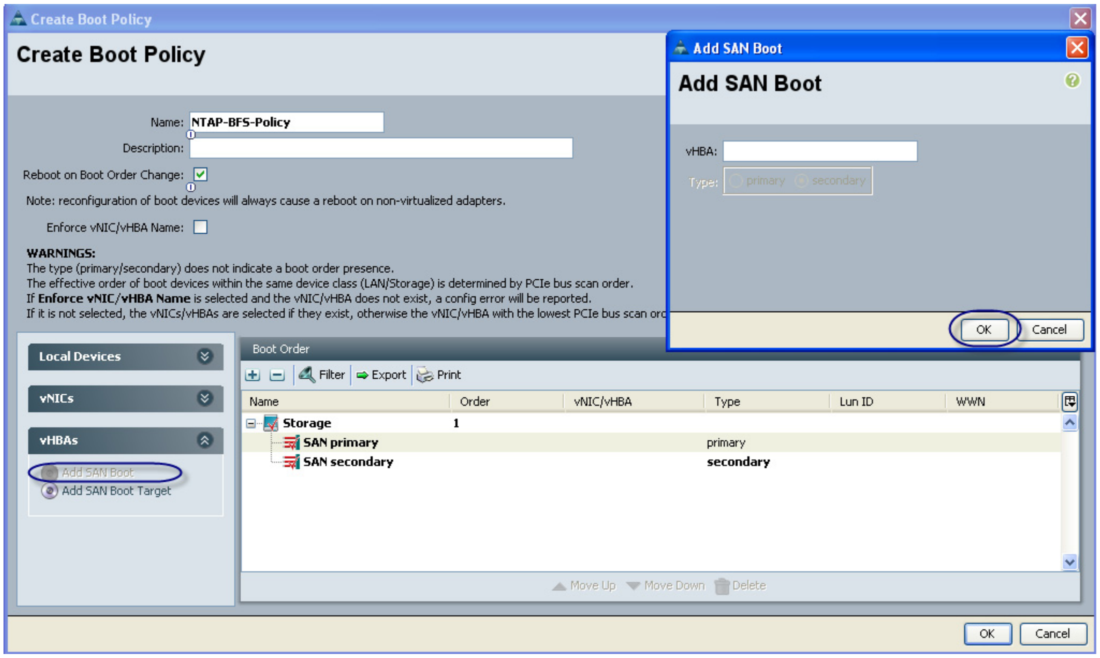

Create Boot Policy

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 ord

Add SAN Boot

vHBA:

Type: ○ primary ○ secondary

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

229992

Step 4 Now add Boot target WWPN to the SAN primary, ensuring that this is the exact pWWN for the NetApp FAS 6080. To avoid any typos, copy and paste from Nexus 5000 "show flogi da".

```
DC04-N5K-1#  sh fcns da vsan 906 | incl Net
0x860002    N     50:0a:09:87:87:49:9a:b4 (NetApp)        scsi-fcp
0x860005    N     50:0a:09:83:89:49:9a:b4 (NetApp)        scsi-fcp
DC05-N5K-1#  sh fcns da vsan 907 | incl Net
0x950002    N     50:0a:09:88:87:49:9a:b4 (NetApp)        scsi-fcp
0x950005    N     50:0a:09:88:97:49:9a:b4 (NetApp)        scsi-fcp
```

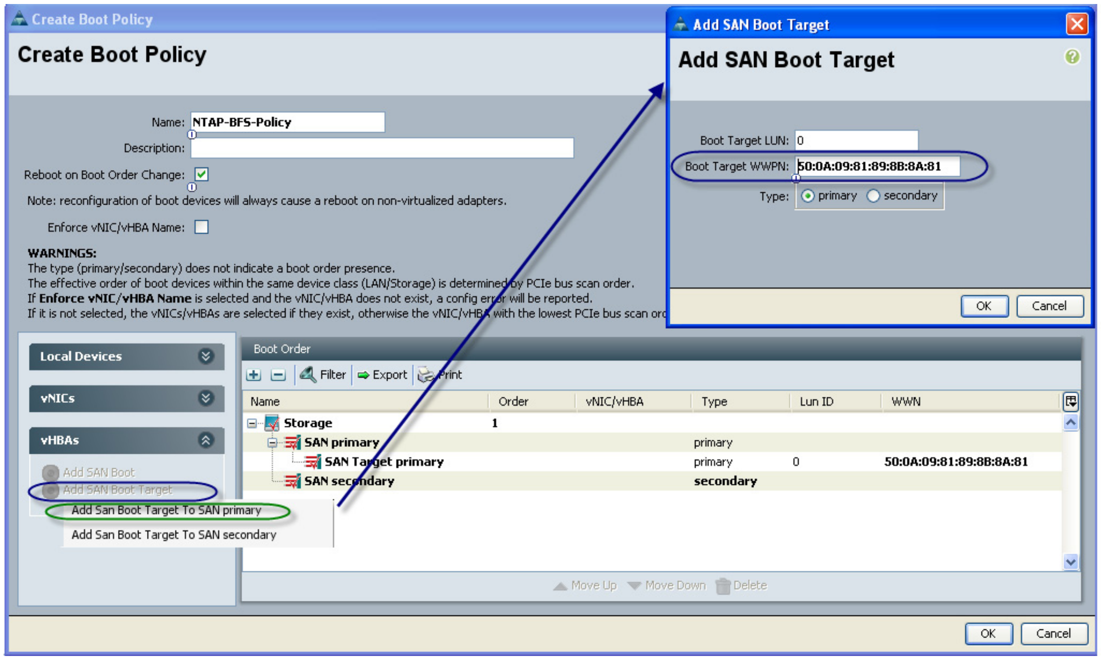

Create Boot Policy

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 ord

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

Step 5 Repeat step 4 for SAN primary's SAN Target Secondary.

Step 6 Repeat step 4 for SAN Secondary's SAN Target Primary.

Step 7 Repeat step 4 for SAN Secondary's SAN Target Secondary.

Step 8 At the end your Boot from SAN policy should look like:

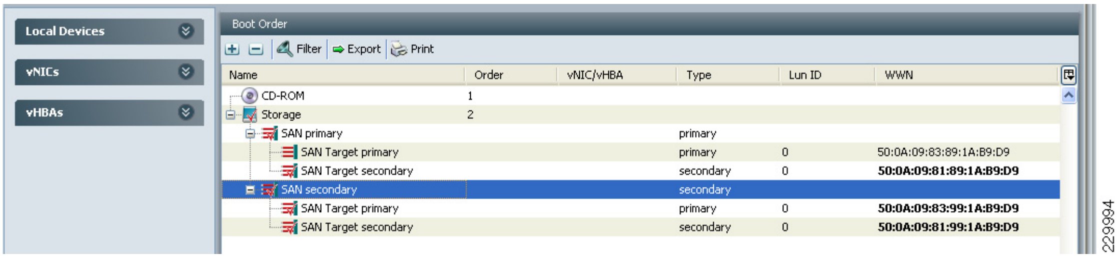

Boot Order

Name	Order	vNIC/vHBA	Type	Lun ID	WWN
CD-ROM	1				
Storage	2				
SAN primary			primary		
SAN Target primary			primary	0	50:0A:09:83:89:1A:B9:D9
SAN Target secondary			secondary	0	50:0A:09:81:89:1A:B9:D9
SAN secondary			secondary		
SAN Target primary			primary	0	50:0A:09:83:99:1A:B9:D9
SAN Target secondary			secondary	0	50:0A:09:81:99:1A:B9:D9

Step 9 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:

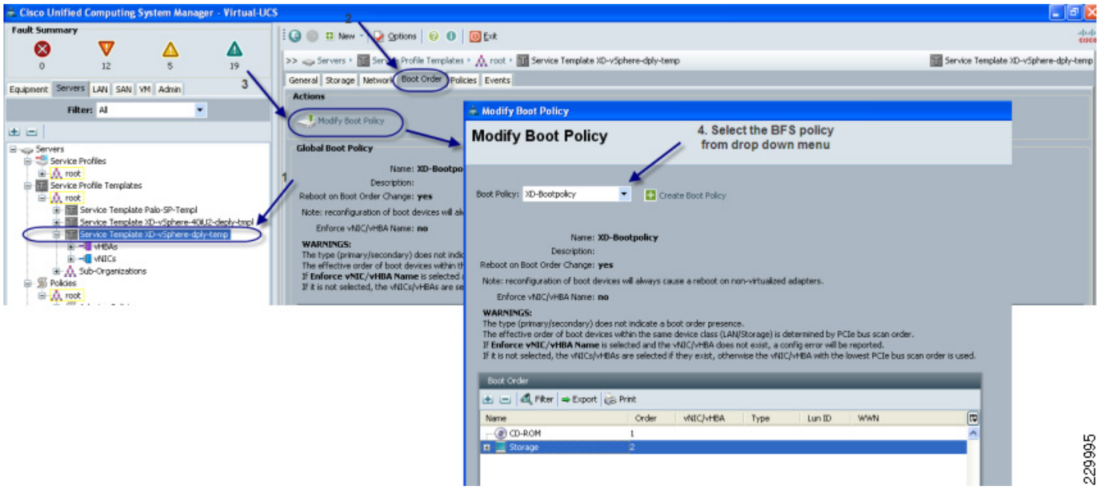
Cisco Unified Computing System Manager - Virtual-UCS
Fault Summary
0
12
5
19
Equipment
Servers
LAN
SAN
VM
Admin
Filter: All
Servers
Service Profiles
root
Service Profile Templates
root
Service Template Palo-SP-Templ
Service Template XD-vSphere-40U2-deply-tmpl
Service Template XD-vSphere-dply-temp
vHBAs
vNICs
Sub-Organizations
Policies
root
New
Options
Exit
>> Servers ▸ Service Profile Templates ▸ root ▸ Service Template XD-vSphere-dply-temp
Service Template XD-vSphere-dply-temp
General
Storage
Network
Boot Order
Policies
Events
Actions
Modify Boot Policy
Global Boot Policy
1
2
3
Modify Boot Policy
4. Select the BFS policy from drop down menu
Create Boot Policy
Boot Order
Filter
Export
Print
229995

Modify Boot Policy

Boot Policy: XD-Bootpolicy

Name: **XD-Bootpolicy**
Description:
Reboot on Boot Order Change: **yes**
Note: reconfiguration of boot devices will always cause a reboot on non-virtualized adapters.
Enforce vNIC/vHBA Name: **no**

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 is used.

Name	Order	vNIC/vHBA	Type	Lun ID	WWN
CD-ROM	1				
Storage	2				

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 have taken place.

NetApp Storage Configuration

A pair of NetApp FAS 6080 storage arrays were used for test scenarios scaling up to 14 hosts:

FAS6080 was used to host:

- All infrastructure Virtual Servers on a single Volume provided by a single controller (2)
- Three VDA Volumes—Two from one Controller (1) and one from controller (2)
- Boot volume for 4 vSphere 4.1 Hosts

Figure 27 NetApp Storage Configuration

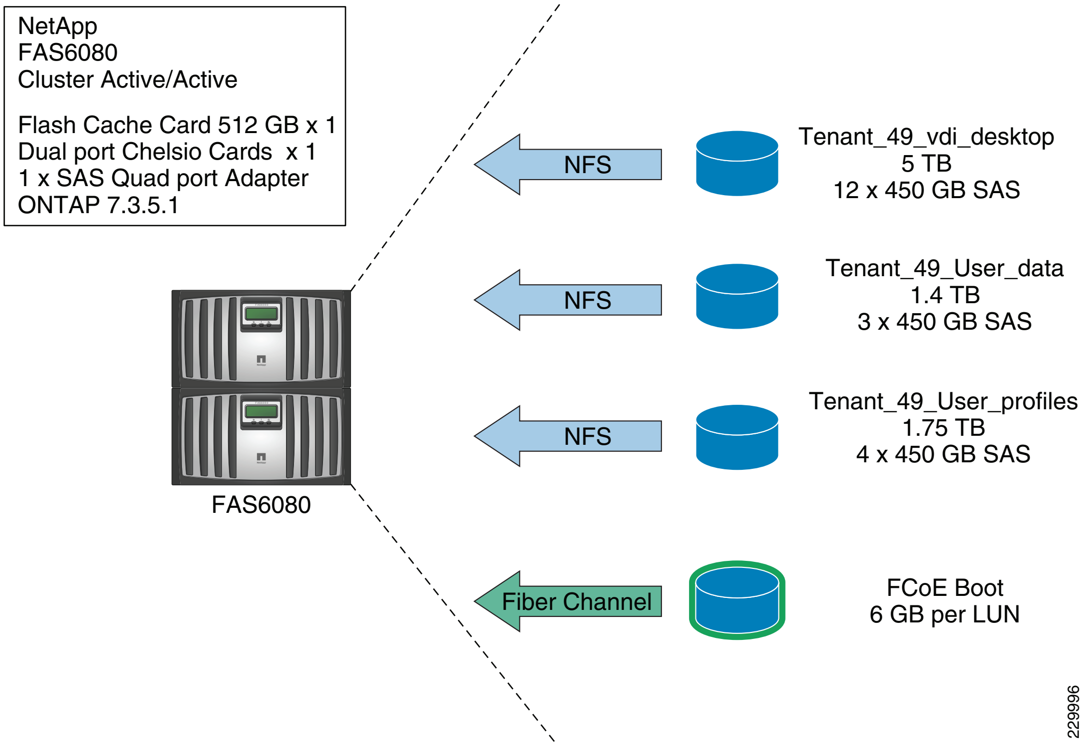

Example of a NetApp NFS Volume Configuration

Step 1 Login to the NetApp storage using a Web browser and click **FilerView**. It starts the NetApp filer configuration application.

Step 2 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.

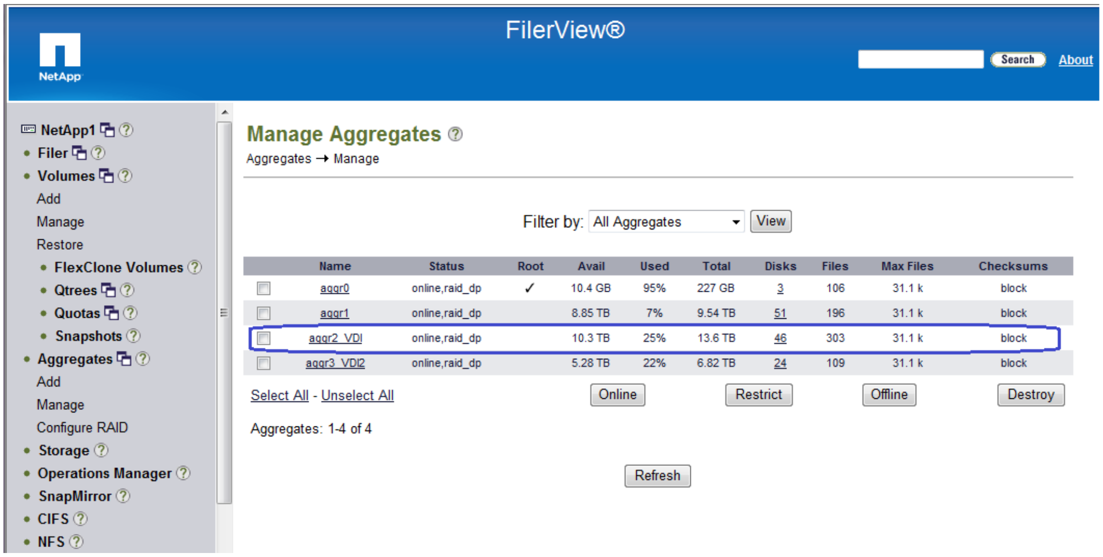

Manage Aggregates

Aggregates → Manage

Filter by: All Aggregates View

	Name	Status	Root	Avail	Used	Total	Disks	Files	Max Files	Checksums
□	aggr0	online,raid_dp	✓	10.4 GB	95%	227 GB	3	106	31.1 k	block
□	aggr1	online,raid_dp		8.85 TB	7%	9.54 TB	51	196	31.1 k	block
□	aggr2_VDI	online,raid_dp		10.3 TB	25%	13.6 TB	46	303	31.1 k	block
□	aggr3_VDI2	online,raid_dp		5.28 TB	22%	6.82 TB	24	109	31.1 k	block

Select All - Unselect All Online Restrict Offline Destroy

Aggregates: 1-4 of 4

Refresh

Step 3 Now from the volumes section, click **Add** to add a volume. An add Volume Wizard pops up:

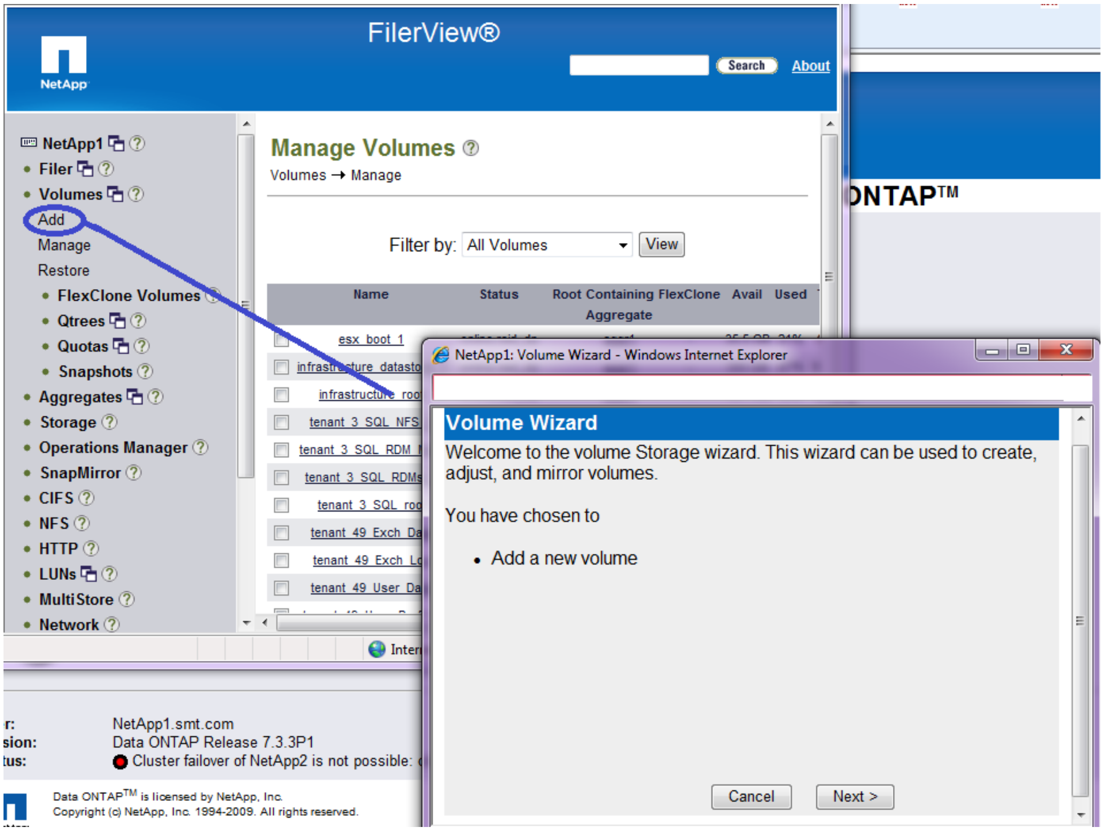

Manage Volumes

Volumes → Manage

Filter by: All Volumes View

	Name	Status	Root Containing FlexClone Aggregate	Avail	Used
□	esx_boot_1				
□	infrastructure_datasto				
□	infrastructure_roo				
□	tenant_3_SQL_NFS				
□	tenant_3_SQL_RDM_				
□	tenant_3_SQL_RDMs				
□	tenant_3_SQL_roo				
□	tenant_49_Exch_Da				
□	tenant_49_Exch_Lo				
□	tenant_49_User_Da				

r: NetApp1.smt.com
sion: Data ONTAP Release 7.3.3P1
tus: ● Cluster failover of NetApp2 is not possible: (

Data ONTAP™ is licensed by NetApp, Inc.
Copyright (c) NetApp, Inc. 1994-2009. All rights reserved.

NetApp1: Volume Wizard - Windows Internet Explorer

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

Step 4 Select Flexible for the Volume Type Selection:

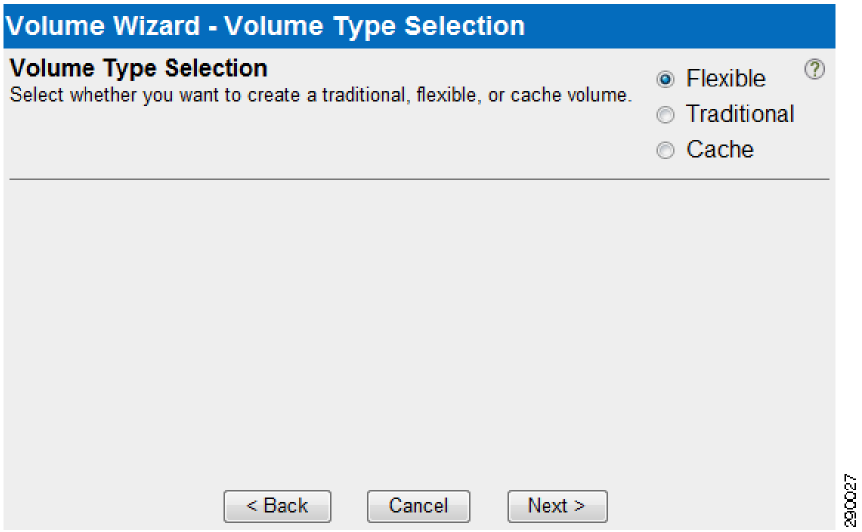

Volume Wizard - Volume Type Selection

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

● Flexible
○ Traditional
○ Cache

Step 5 Input volume name and language (the default POSIX is fine):

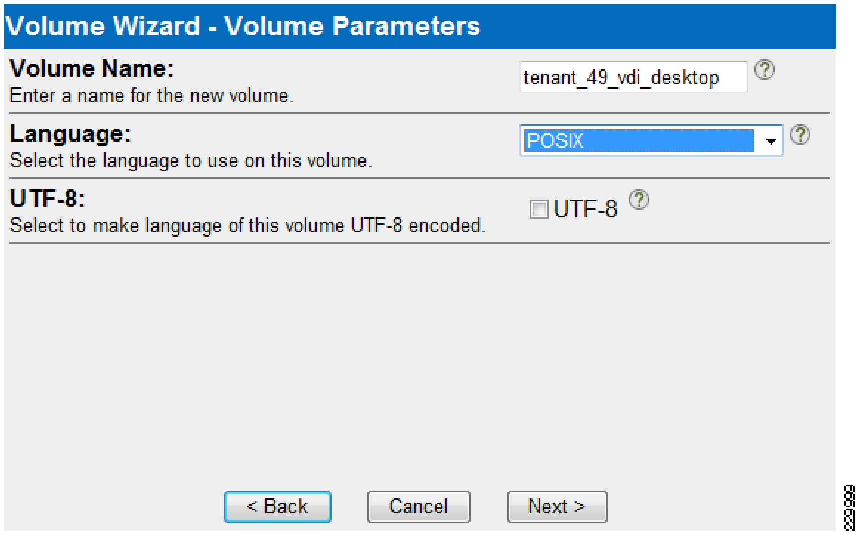

Volume Wizard - Volume Parameters

Volume Name: tenant_49_vdi_desktop
Enter a name for the new volume.

Language: POSIX
Select the language to use on this volume.

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

Step 6 Select the aggregate to contain this volume and set space guarantee as none for thin provisioning:

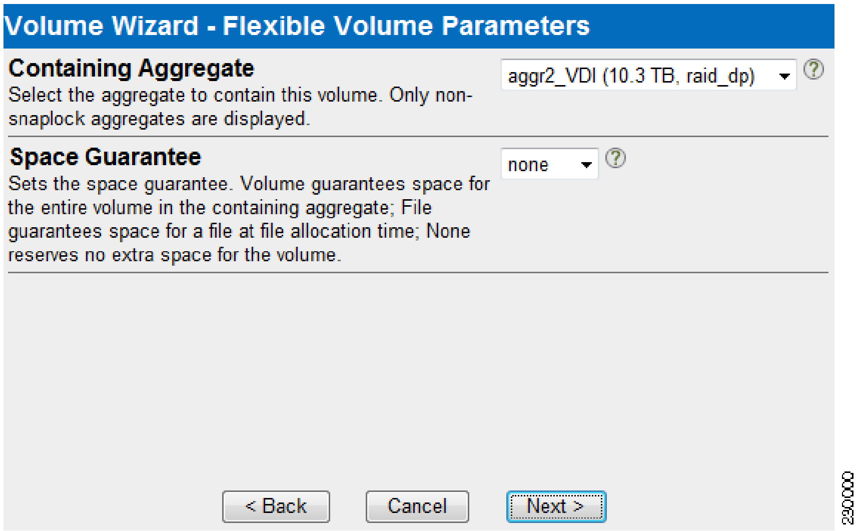

Volume Wizard - Flexible Volume Parameters

Containing Aggregate
Select the aggregate to contain this volume. Only non-snaplock aggregates are displayed.
aggr2_VDI (10.3 TB, raid_dp)

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.
none

Step 7 Input the volume size and set snapshot reserve to 0:

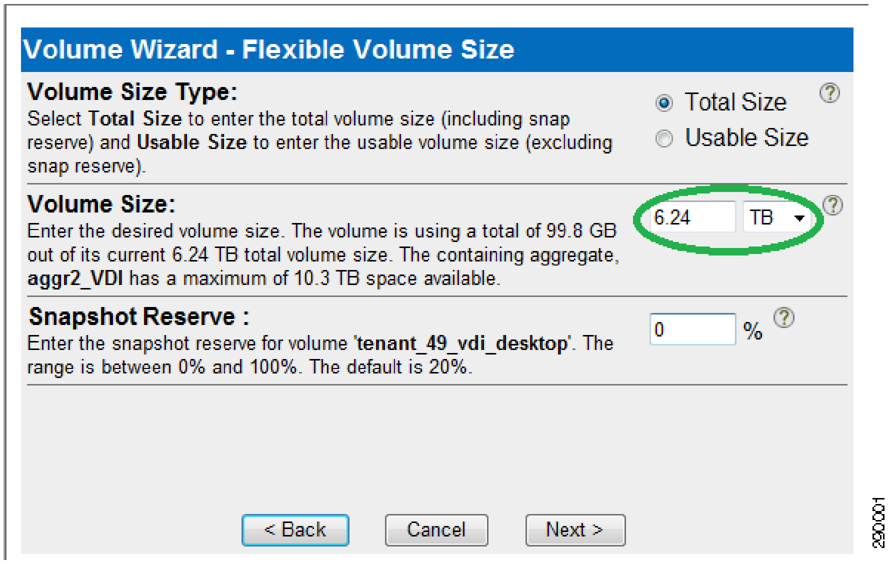

Volume Wizard - Flexible Volume Size

Volume Size Type:
Select **Total Size** to enter the total volume size (including snap reserve) and **Usable Size** to enter the usable volume size (excluding snap reserve).
● Total Size
○ Usable Size

Volume Size:
Enter the desired volume size. The volume is using a total of 99.8 GB out of its current 6.24 TB total volume size. The containing aggregate, **aggr2_VDI** has a maximum of 10.3 TB space available.
6.24 TB

Snapshot Reserve :
Enter the snapshot reserve for volume '**tenant_49_vdi_desktop**'. The range is between 0% and 100%. The default is 20%.
0 %

Step 8 You are now finished, so click **Commit**:

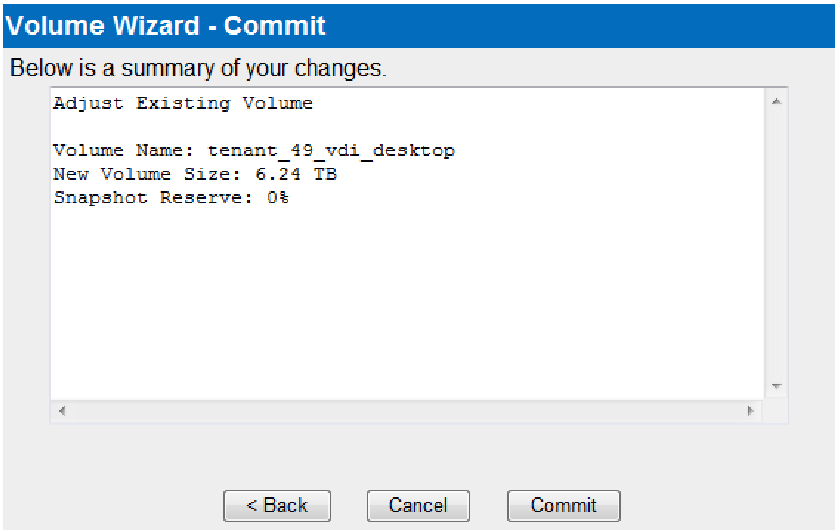

Below is a summary of your changes.

```
Adjust Existing Volume

Volume Name: tenant_49_vdi_desktop
New Volume Size: 6.24 TB
Snapshot Reserve: 0%
```

290002

Step 9 After the volume is added, go to the Multi-store, click **Manage Vfilers**, and add the path for the volume you just created. For more on Vfilers in the ESMT solution, see the ESMT CVD (http://www.cisco.com/en/US/solutions/ns340/ns414/ns742/ns743/ns1050/landing_dcVDDC.html). In our example below, our Vfiler is tenant 49 VDI.

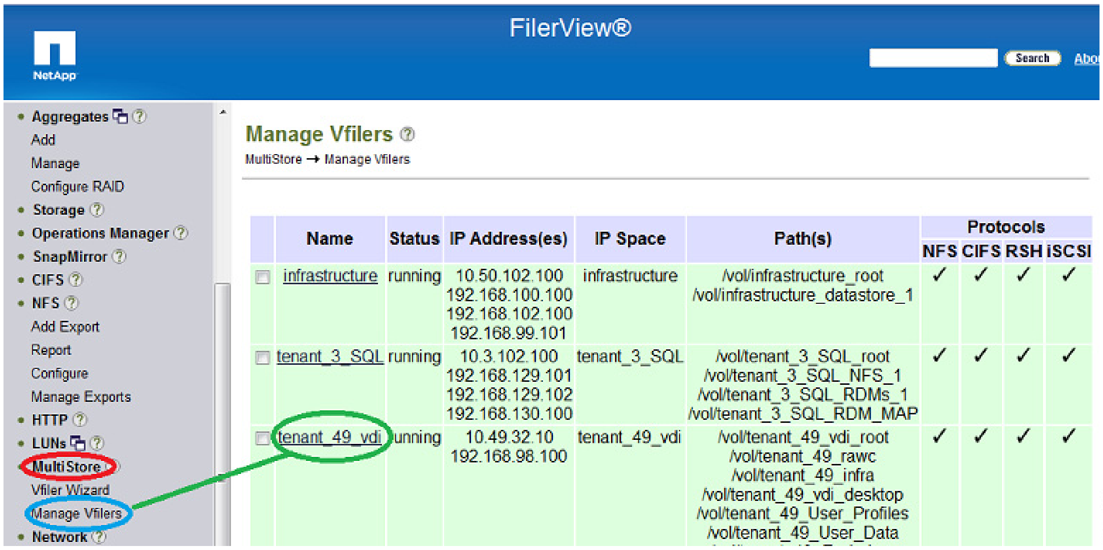

Manage Vfilers

MultiStore → Manage Vfilers

	Name	Status	IP Address(es)	IP Space	Path(s)	Protocols			
						NFS	CIFS	RSH	iSCSI
□	infrastructure	running	10.50.102.100 192.168.100.100 192.168.102.100 192.168.99.101	infrastructure	/vol/infrastructure_root /vol/infrastructure_datastore_1	✓	✓	✓	✓
□	tenant_3_SQL	running	10.3.102.100 192.168.129.101 192.168.129.102 192.168.130.100	tenant_3_SQL	/vol/tenant_3_SQL_root /vol/tenant_3_SQL_NFS_1 /vol/tenant_3_SQL_RDMs_1 /vol/tenant_3_SQL_RDM_MAP	✓	✓	✓	✓
□	tenant_49_vdi	running	10.49.32.10 192.168.98.100	tenant_49_vdi	/vol/tenant_49_vdi_root /vol/tenant_49_rawc /vol/tenant_49_infra /vol/tenant_49_vdi_desktop /vol/tenant_49_User_Profiles /vol/tenant_49_User_Data	✓	✓	✓	✓

290003

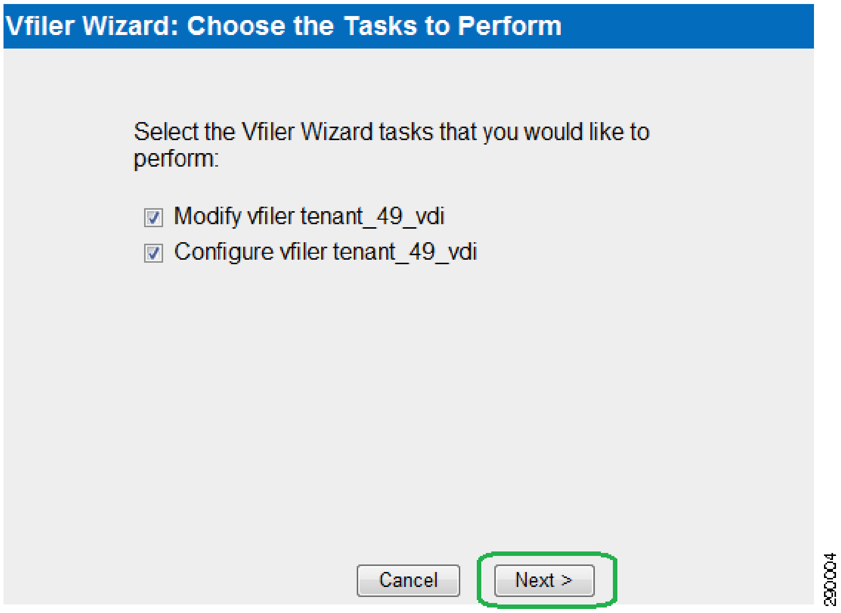

Vfiler Wizard: Choose the Tasks to Perform

Select the Vfiler Wizard tasks that you would like to perform:

☑ Modify vfiler tenant_49_vdi

☑ Configure vfiler tenant_49_vdi

Step 10 After clicking on the vfiler name circled above, on specific tenant add your path:

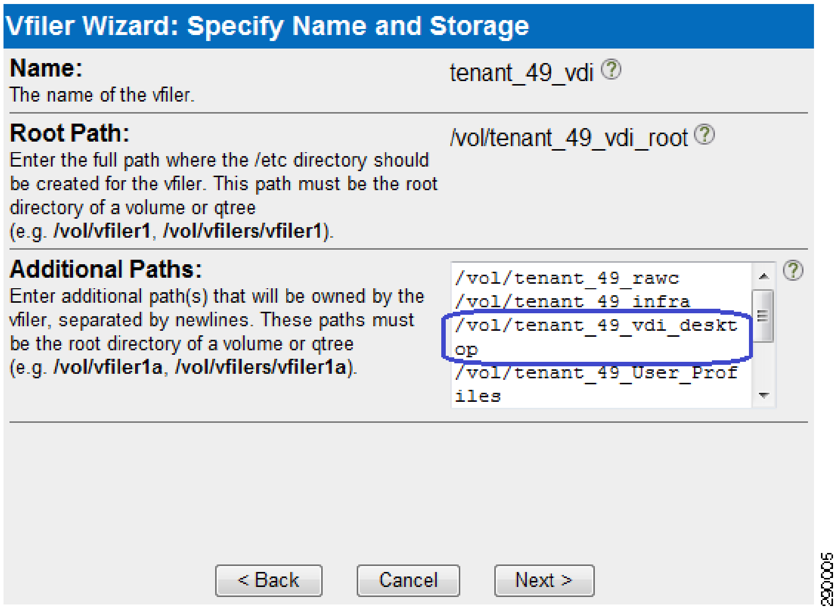

Vfiler Wizard: Specify Name and Storage

Name: tenant_49_vdi
The name of the vfiler.

Root Path: /vol/tenant_49_vdi_root
Enter the full path where the /etc directory should be created for the vfiler. This path must be the root directory of a volume or qtree (e.g. **/vol/vfiler1**, **/vol/vfilers/vfiler1**).

Additional Paths:
Enter additional path(s) that will be owned by the vfiler, separated by newlines. These paths must be the root directory of a volume or qtree (e.g. **/vol/vfiler1a**, **/vol/vfilers/vfiler1a**).

```
/vol/tenant_49_rawc
/vol/tenant_49_infra
/vol/tenant_49_vdi_deskt
op
/vol/tenant_49_User_Prof
iles
```

Step 11 Click **Next** all the way through to complete the wizard.

NetApp Deduplication in Practice

As described in NetApp Deduplication, 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 28 for 5 TB volume hosting 1540 desktops each with 40 GB capacity.

Figure 28 NetApp Deduplication

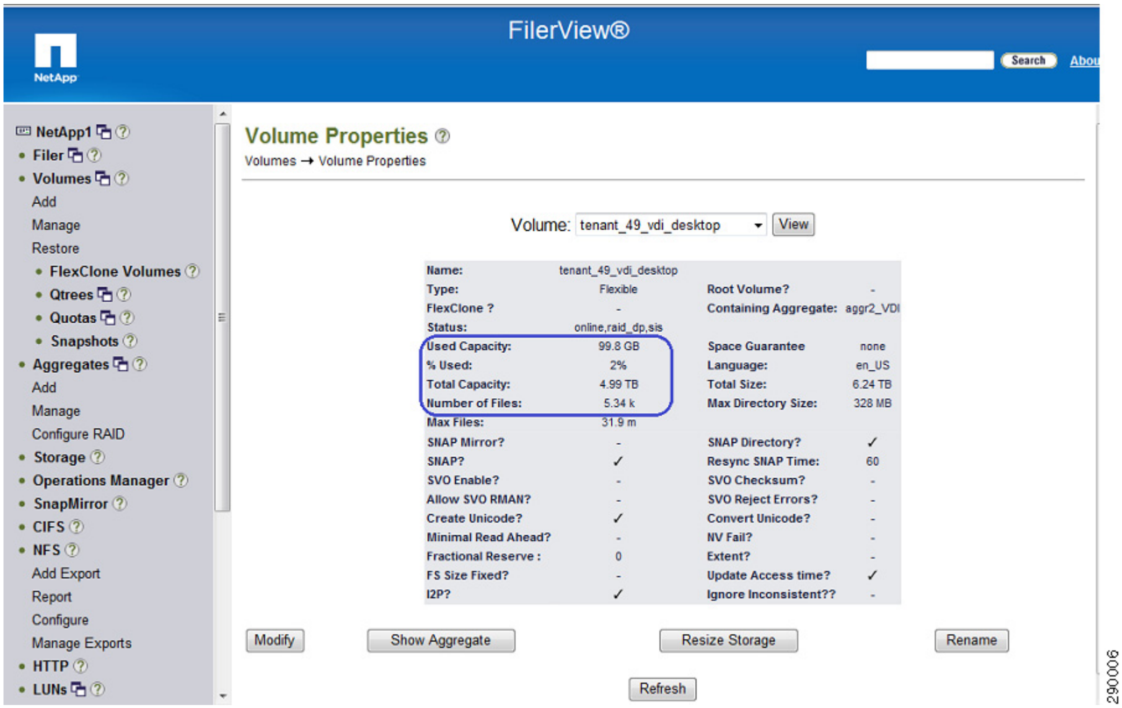

Volume: tenant_49_vdi_desktop

Name:	tenant_49_vdi_desktop		
Type:	Flexible	Root Volume?	-
FlexClone ?	-	Containing Aggregate:	aggr2_VDI
Status:	online,raid_dp,sis		
Used Capacity:	99.8 GB	Space Guarantee	none
% Used:	2%	Language:	en_US
Total Capacity:	4.99 TB	Total Size:	6.24 TB
Number of Files:	5.34 k	Max Directory Size:	328 MB
Max Files:	31.9 m		
SNAP Mirror?	-	SNAP Directory?	✓
SNAP?	✓	Resync SNAP Time:	60
SVO Enable?	-	SVO Checksum?	-
Allow SVO RMAN?	-	SVO Reject Errors?	-
Create Unicode?	✓	Convert Unicode?	-
Minimal Read Ahead?	-	NV Fail?	-
Fractional Reserve :	0	Extent?	-
FS Size Fixed?	-	Update Access time?	✓
I2P?	✓	Ignore Inconsistent??	-

NetApp FlexScale

As described in FlexScale, NetApp FlexScale allows administrators to customize the Flash Cache operations to meet their environment requirements. Extensive scalable VMware View testing within the NetApp solution labs has shown that significant performance improvements can be gained by turning on metadata and normal user data caching modes in FlexScale. The following configuration is recommended:

```
options flexscale.enable on
options flexscale.normal_data_blocks on
```

Note Configuration recommendations are based on TR-3705.

OS Installation

ESXi Installation

ESXi installs were performed via Virtual Media through KVM from UCSM. UCS blades mounted the ESXi ISO and were installed to SAN boot targets.

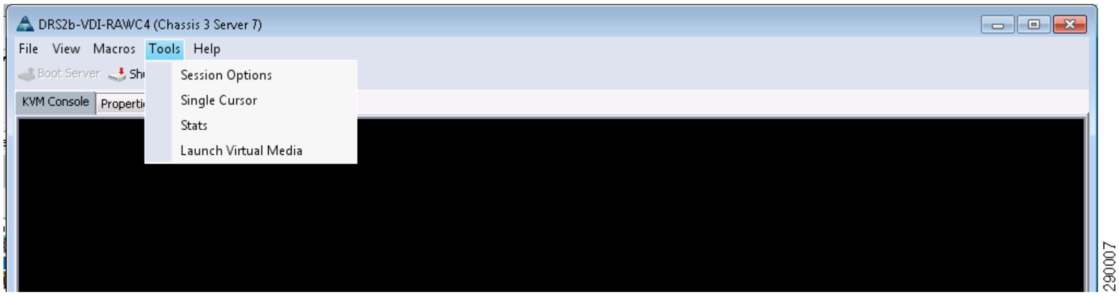

PXE installation can be used as a more efficient alternative. When installing from PXE, the management VLAN will need to be set as the default VLAN of the first vNIC.

VMware vSphere Kernel Adjustments for High CPU Environments

The default CPU fairness algorithm in vSphere tries to help VMs catch up by setting aside the other logical processor in a hyperthreaded environment. This is configured through a parameter called HaltingIdleMsecPenalty (HIMP). HIMP is a number of milliseconds that is multiplied by the number of vCPUs. The derived number is used as a cumulative value across vCPUs that the vCPUs can fall behind before logical processors become reserved to help a VM catch up.

The default implementation is to start reserving logical processors after a vCPU falls behind more than 100 milliseconds. This may be more aggressive than is needed in certain environments. These reserved logical CPUs can lead to excessive amounts of dormant CPU threads in systems that:

- Have more than 50% CPU utilization
- Are very close to exactly committed (number of vCPUs = number of pCPUs +/- 25%)
- Have particular kinds of bursty CPU usage patterns

To allow more flexibility, an adjustment to HaltingIdleMsecPenalty was made:

```
# vicfg-advcfg --set 2000 /Cpu/HaltingIdleMsecPenalty
```

And:

```
# vicfg-advcfg --set 80000 /Cpu/HaltingIdleMsecPenaltyMax
```

Where HaltingIdleMsecPenaltyMax is an upper level of cumulative milliseconds that HIMP multiplied by the number of vCPUs is allowed to reach. For more information, see: http://kb.vmware.com/selfservice/microsites/search.do?language=en_US&cmd=displayKC&externalId=1020233.

Optional Memory Reclamation in VMware vSphere

There is a default of 6% of physical memory reserved for the ESX/ESXi host in vSphere 4.1. This is needed in systems with less than 64GB of memory, but can be an over commitment in a system like the B250 M2 with 192GB or more of memory. This reservation of memory can be reduced with the **vsish** command within the tech support mode of ESXi 4.1. A reduction to 2% memory reserved on a host could be achieved with the command:

```
vsish -e set /sched/freeMemoryState/minFreePct 2
```

This example memory reclamation would free up almost 8GB of memory on a B250 with 192GB of RAM, allowing a VDI deployment to support 5-10 more desktops per blade, amounting to a significant increase in large-scale deployments.

The **vsish** command is available along with the general distribution of ESXi and can be used to tune the memory as per KB article 1033687 (see URL below) and this is the recommended best practices to achieve better utilization of memory and more than 160 desktops per UCS blade.

The **vsish** command is not available in the general distribution of the ESX server package and the debug packages have to be installed in order to obtain **vsish** and is not recommended for production use.

The need for memory and CPU tuning (HIMP settings) will not be necessary in the newer releases of vSphere as this is being handled internally by the kernel.

The **vsish** command was not implemented in the View testbed from which these results were generated, but is mentioned as an option to consider to reduce host ballooning in high utilization environments. The current implementation of **vsish** would not persist through a reboot of the host and would need to be added to the /etc/rc.local file of the system. For more information, see: http://kb.vmware.com/selfservice/microsites/search.do?language=en_US&cmd=displayKC&externalId=1033687.

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. While vMA is optional, 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.

By registering each installed host with the VIFastPass through vifpinit on the vMA, redundant tasks can be looped through with the default bash shell of the vMA.

Nexus 1000V VSM and VEM Module Installation

ESXi hosts were installed with the Nexus 1000V Virtual Ethernet Module (VEM) and added to the Nexus 1000V Virtual Supervisor Module (VSM) that was created and joined to vCenter as a vNetwork Distributed Switch.

For more information on the installation and configuration of the Cisco Nexus 1000V and VEM, see:

- http://www.cisco.com/en/US/partner/docs/switches/datacenter/nexus1000/sw/4_2_1_s_v_1_4/install/software/guide/n1000v_install_software.html
- http://www.cisco.com/en/US/products/ps9902/products_installation_and_configuration_guides_list.html
- http://www.cisco.com/en/US/docs/switches/datacenter/nexus1000/sw/4_2_1_s_v_1_4/install/vem/guide/n1000v_vem_install.html
- For a quick video walkthrough of the Nexus VSM installation, watch: http://www.youtube.com/watch?v=-sxWiz7S-z0

Network Configuration

The Cisco VIC vNIC interfaces were created as discussed in the Enhanced Secure Multi-Tenancy Design Guide (http://www.cisco.com/en/US/solutions/ns340/ns414/ns742/ns743/ns1050/landing_dcVDDC.html).

Different VLANS were used to separate the desktop traffic from the management traffic, NetApp storage traffic, and the vMotion traffic. Hosts were installed with the Cisco Virtual Ethernet Module (VEM), joined to the Nexus 1000V VSM, and VMkernel NICs were created for NFS and vMotion.

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

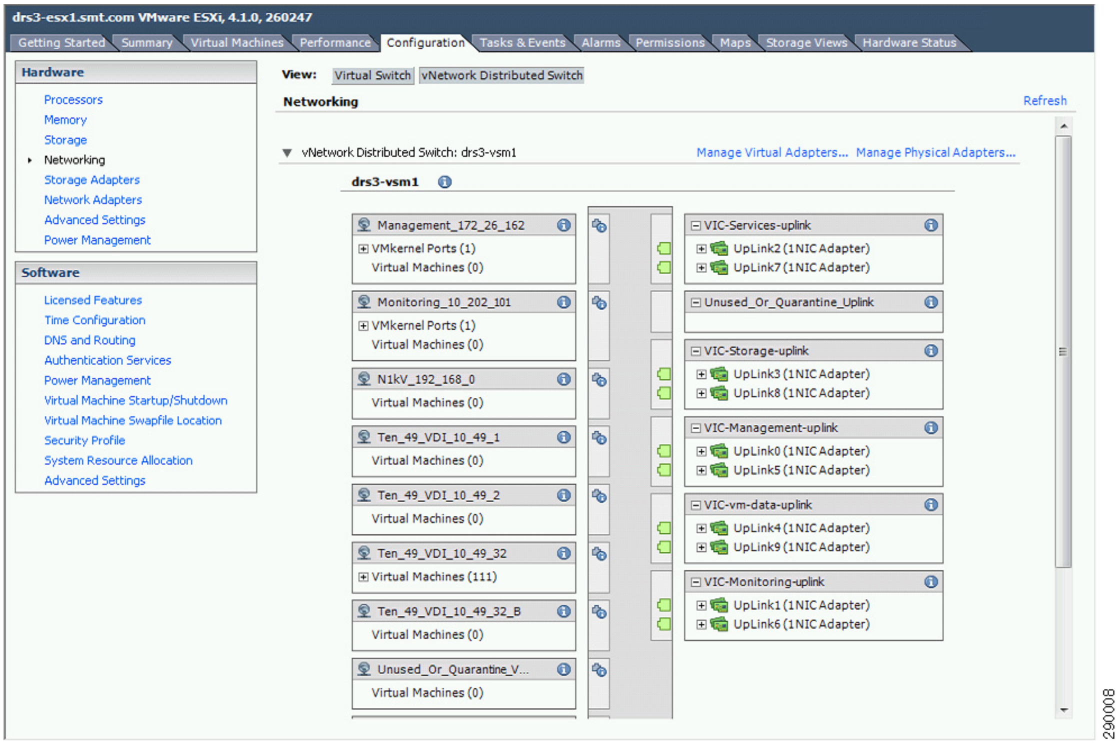

Using the Cisco Nexus 1000V for network configuration provides advanced capabilities to do DHCP snooping and other smart network switch capabilities in the hypervisor itself. These features have a huge benefit in a virtual desktop environment where a vSwitch would not be able to provide such features.

Setup NTP Server and NAS Datastores with vMA

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 server. This is important from a time lag perspective and maintains synchronization of performance data collected across various components.

The configuration of NTP and NAS resources can be set through the vMA:

```
esxcfg-ntp -a x.x.x.6
esxcfg-ntp -a x.x.x.9
esxcfg-ntp --stop
esxcfg-ntp --start
esxcfg-nas -a -o 192.168.98.100 -s  /vol/tenant_49_infra tenant_49_infra
esxcfg-nas -a -o 192.168.98.100 -s  /vol/tenant_49_rawc tenant_49_rawc
esxcfg-nas -a -o 192.168.98.100 -s  /vol/tenant_49_vdi_desktop tenant_49_vdi_desktop
```

VMware vSphere Configuration

For single server scale testing, one ESXi 4.1.0 server was configured with boot from SAN. 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. One NAS device was 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.

Figure 29 VMware vSphere Configuration

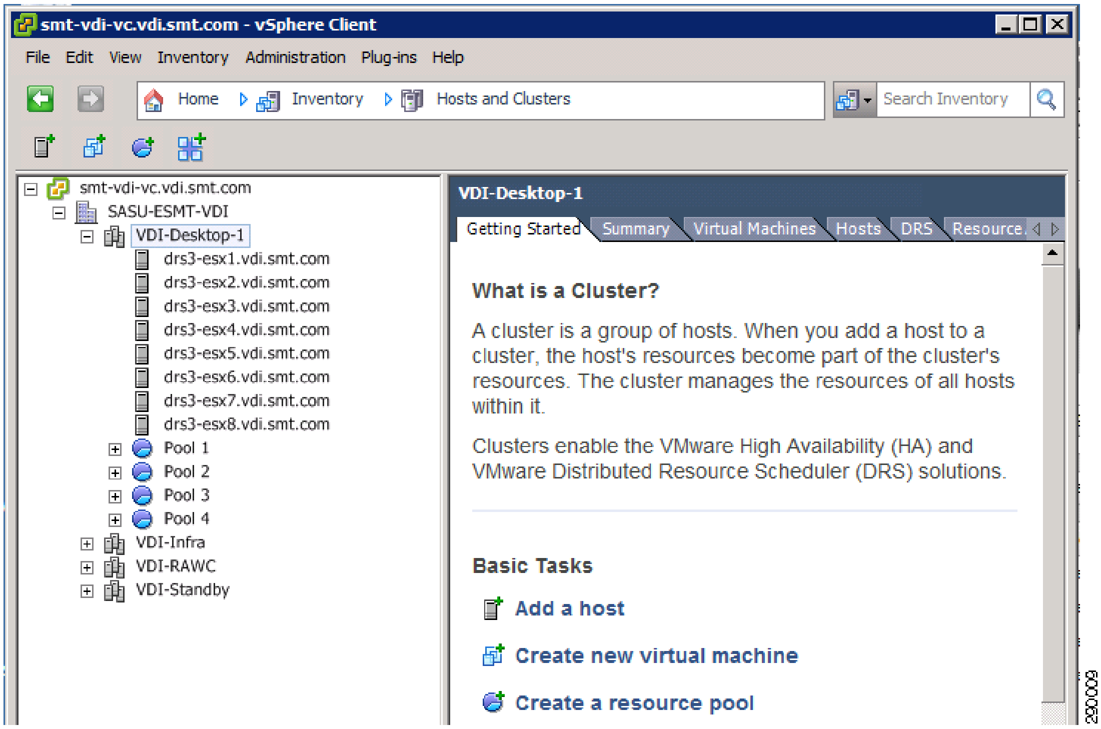

What is a Cluster?

A cluster is a group of hosts. When you add a host to a cluster, the host's resources become part of the cluster's resources. The cluster manages the resources of all hosts within it.

Clusters enable the VMware High Availability (HA) and VMware Distributed Resource Scheduler (DRS) solutions.

Figure 29 shows the vSphere environment described above as seen through the vSphere Client.

VMware View Configuration

VMware View is composed of:

- View Composer
- View Connection Server
- View Transfer Server
- View Security Server
- View Agent

The Transfer and Security servers were not used in our test environment.

For more detailed instructions and requirements, refer to the VMware View 4.5 Installation Guide at: http://www.vmware.com/pdf/view45_installation_guide.pdf. For scaling beyond what was demonstrated in this environment, refer to the VMware View 4.5 Architecture Planning Guide at: http://www.vmware.com/pdf/view45_architecture_planning.pdf.

VMware View Composer

View Composer is an application that will run as a service on the hosting vCenter(s). A View Composer SQL database must be created for Composer, which can be MS SQL or Oracle.

VMware View Connection Server

View Connection Server installs to a dedicated system and requires an Event Database that can be hosted from an MS SQL or Oracle instance. The Connection server will need an SSL certificate that can be self signed for testing purposes.

The first Connection server will be a Standard installation (supporting up to 2000 desktops). Subsequent Connection servers can be configured as Replicas that pull their configuration from a designated View Connection Server Standard install.

VMware View Agent

View Agent installs to the VM configured to be the Client Desktop Golden Image to provide access to enabled users through the Connection Server.

Once the Golden image has been optimized according the VMware recommendations specified in Microsoft Windows 7 Golden Image Creation and Provisioning, a snapshot is taken of the VM. This snapshot will be referenced by any pools created from it. Pool creation and administration occurs through the VMware View Administrator interface, which can be found at https://<View Connection Server>admin/.

The pool is created from the Pools option under the Inventory tab of View Administrator:

Figure 30 Pool Creation

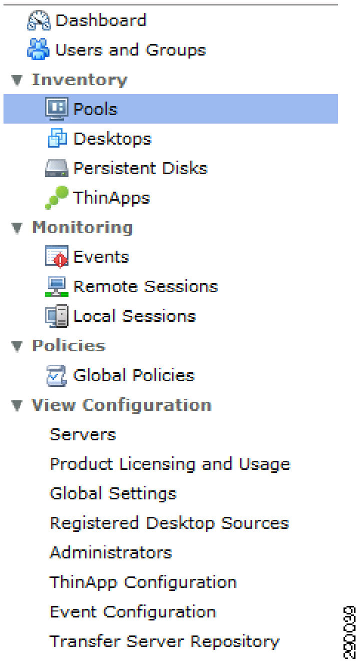

The type of pool selected is “Automated Pool”:

Figure 31 Automated Pool

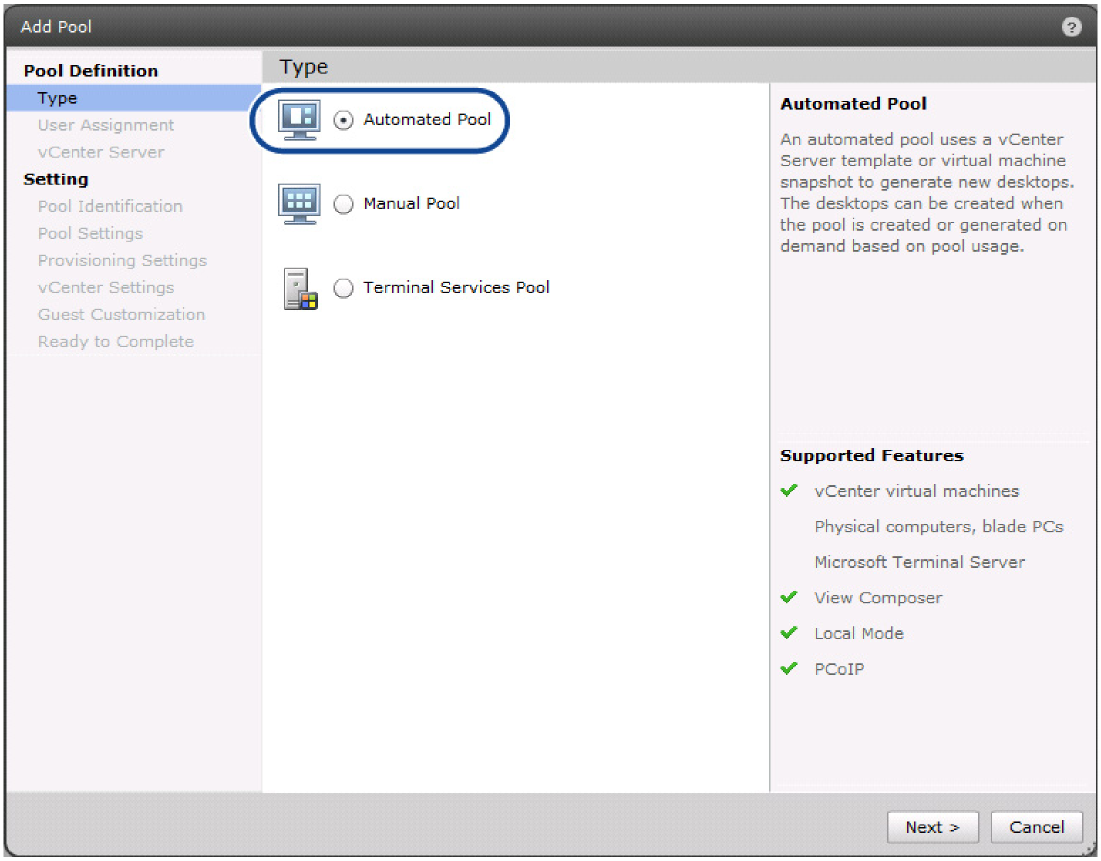

Type

● Automated Pool

○ Manual Pool

○ Terminal Services Pool

Automated Pool

An automated pool uses a vCenter Server template or virtual machine snapshot to generate new desktops. The desktops can be created when the pool is created or generated on demand based on pool usage.

Supported Features

- ✓ vCenter virtual machines
- Physical computers, blade PCs
- Microsoft Terminal Server
- ✓ View Composer
- ✓ Local Mode
- ✓ PCoIP

User Assignment is left as “Dedicated”:

Figure 32 User Assignment—Dedicated

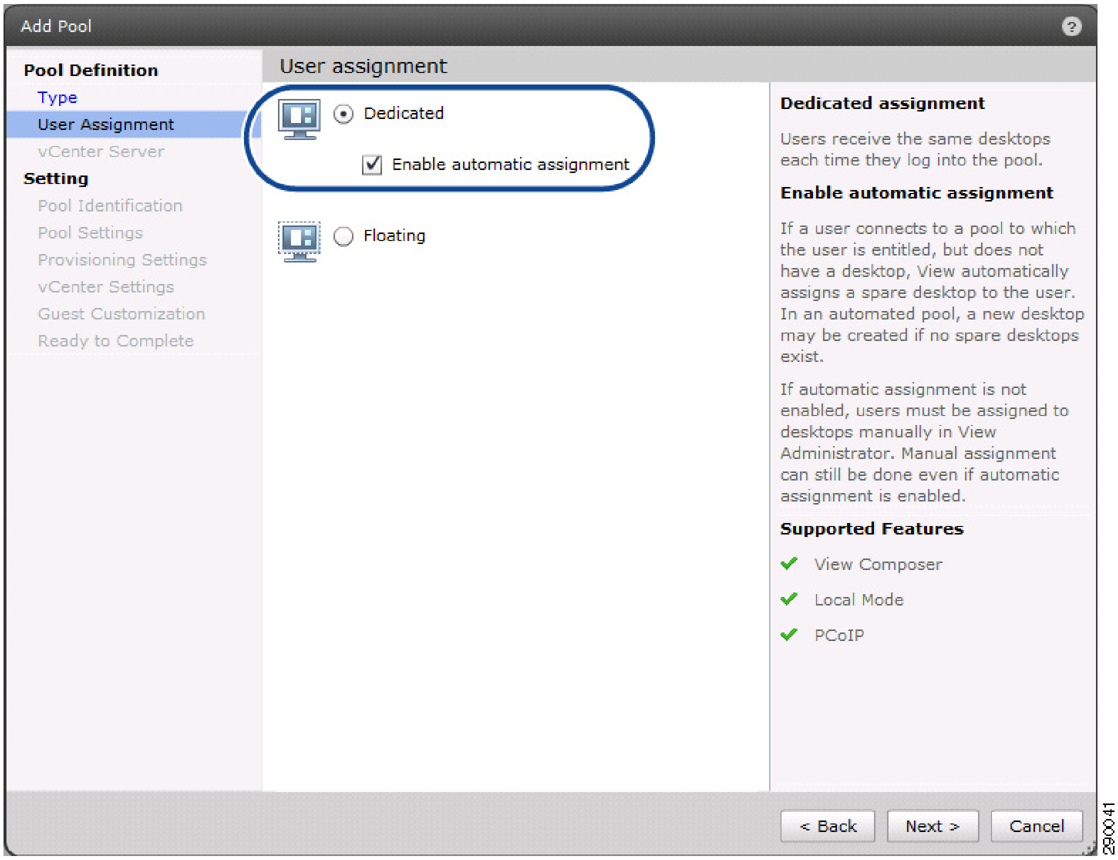

● Dedicated

☑ Enable automatic assignment

○ Floating

Dedicated assignment

Users receive the same desktops each time they log into the pool.

Enable automatic assignment

If a user connects to a pool to which the user is entitled, but does not have a desktop, View automatically assigns a spare desktop to the user. In an automated pool, a new desktop may be created if no spare desktops exist.

If automatic assignment is not enabled, users must be assigned to desktops manually in View Administrator. Manual assignment can still be done even if automatic assignment is enabled.

Supported Features

✓ View Composer

✓ Local Mode

✓ PCoIP

vCenter Server specification is changed from “Full virtual machines” to “View Composer linked clones”:

Figure 33 vCenter Server Specification

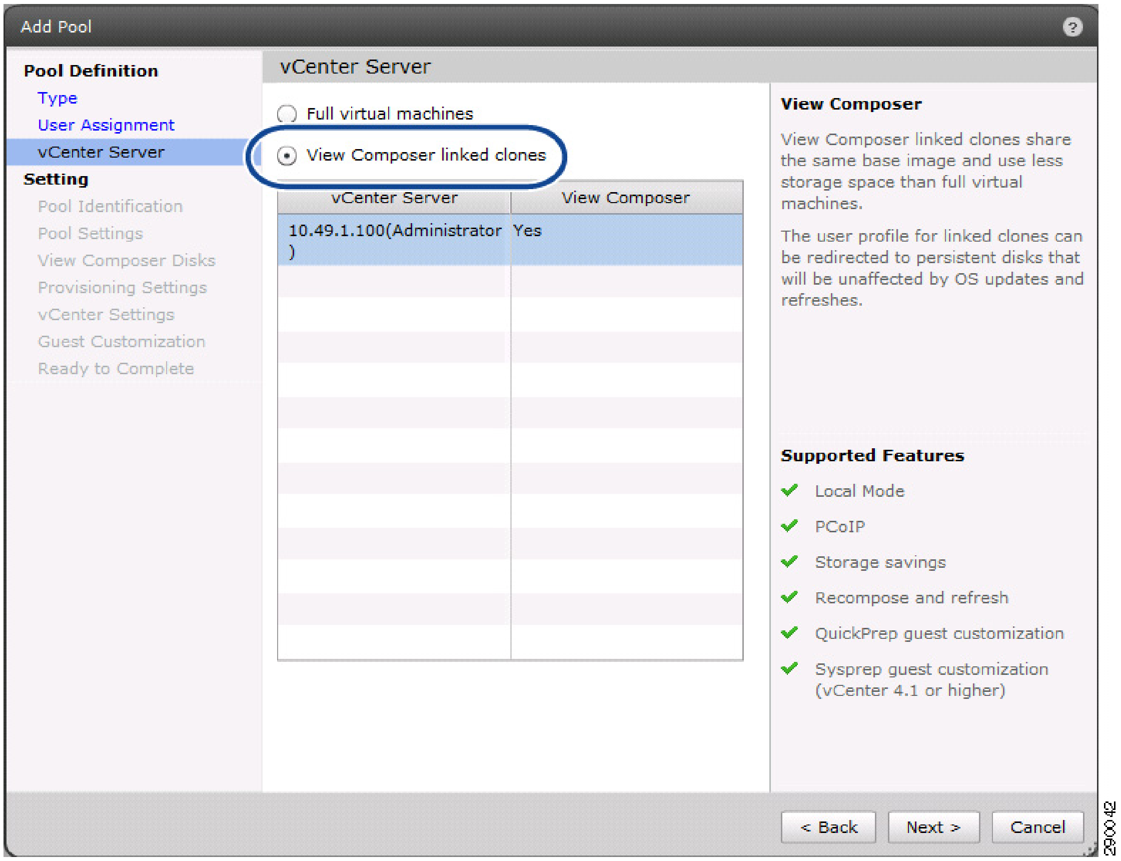

vCenter Server

○ Full virtual machines

● View Composer linked clones

vCenter Server	View Composer
10.49.1.100(Administrator)	Yes

View Composer

View Composer linked clones share the same base image and use less storage space than full virtual machines.

The user profile for linked clones can be redirected to persistent disks that will be unaffected by OS updates and refreshes.

Supported Features

- ✓ Local Mode
- ✓ PCoIP
- ✓ Storage savings
- ✓ Recompose and refresh
- ✓ QuickPrep guest customization
- ✓ Sysprep guest customization (vCenter 4.1 or higher)

The Pool Identification ID and Display name was set to pool1 to match a resource pool that had been set within the cluster:

Figure 34 Pool Identification ID and Display Name

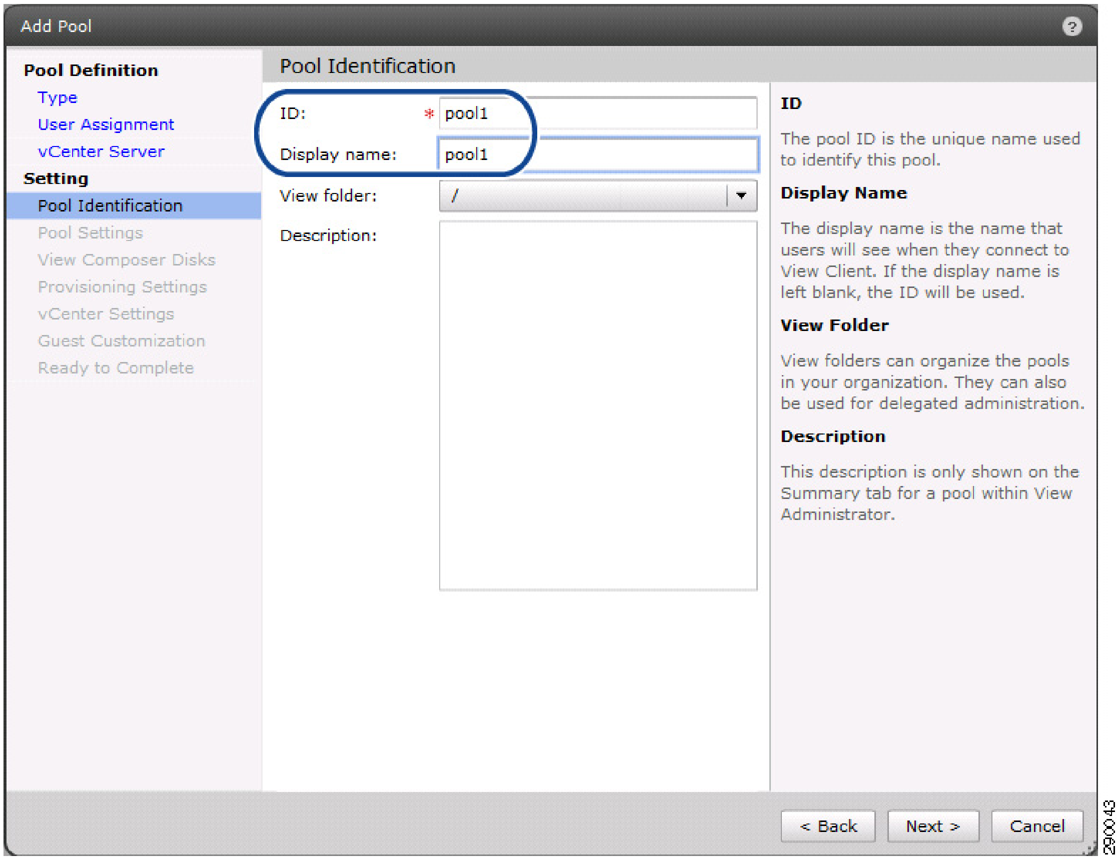

Add Pool

Pool Identification

ID: * pool1

Display name: pool1

View folder: /

Description:

ID

The pool ID is the unique name used to identify this pool.

Display Name

The display name is the name that users will see when they connect to View Client. If the display name is left blank, the ID will be used.

View Folder

View folders can organize the pools in your organization. They can also be used for delegated administration.

Description

This description is only shown on the Summary tab for a pool within View Administrator.

Pool Settings were left unchanged with the exception of “Automatically logoff after disconnect:”, which was changed to “Immediate”:

Figure 35 Pool Settings

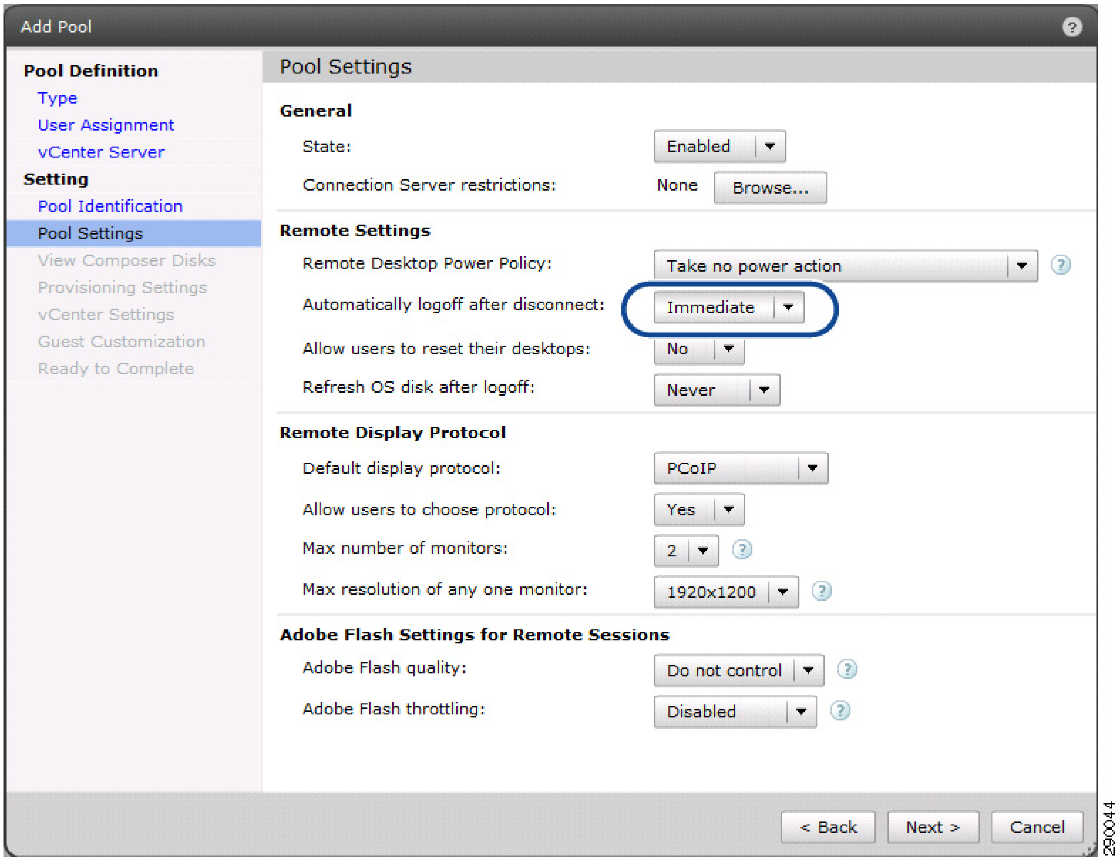

Pool Settings

General

State: Enabled

Connection Server restrictions: None Browse...

Remote Settings

Remote Desktop Power Policy: Take no power action

Automatically logoff after disconnect: Immediate

Allow users to reset their desktops: No

Refresh OS disk after logoff: Never

Remote Display Protocol

Default display protocol: PCoIP

Allow users to choose protocol: Yes

Max number of monitors: 2

Max resolution of any one monitor: 1920x1200

Adobe Flash Settings for Remote Sessions

Adobe Flash quality: Do not control

Adobe Flash throttling: Disabled

View Composer Disks was set to "Do not redirect Windows profile" and the '"Disposable File Redirection" was left to "Redirect disposable files to a non-persistent disk":

Figure 36 Settings for View Componser Disks and Disposable File Redirection

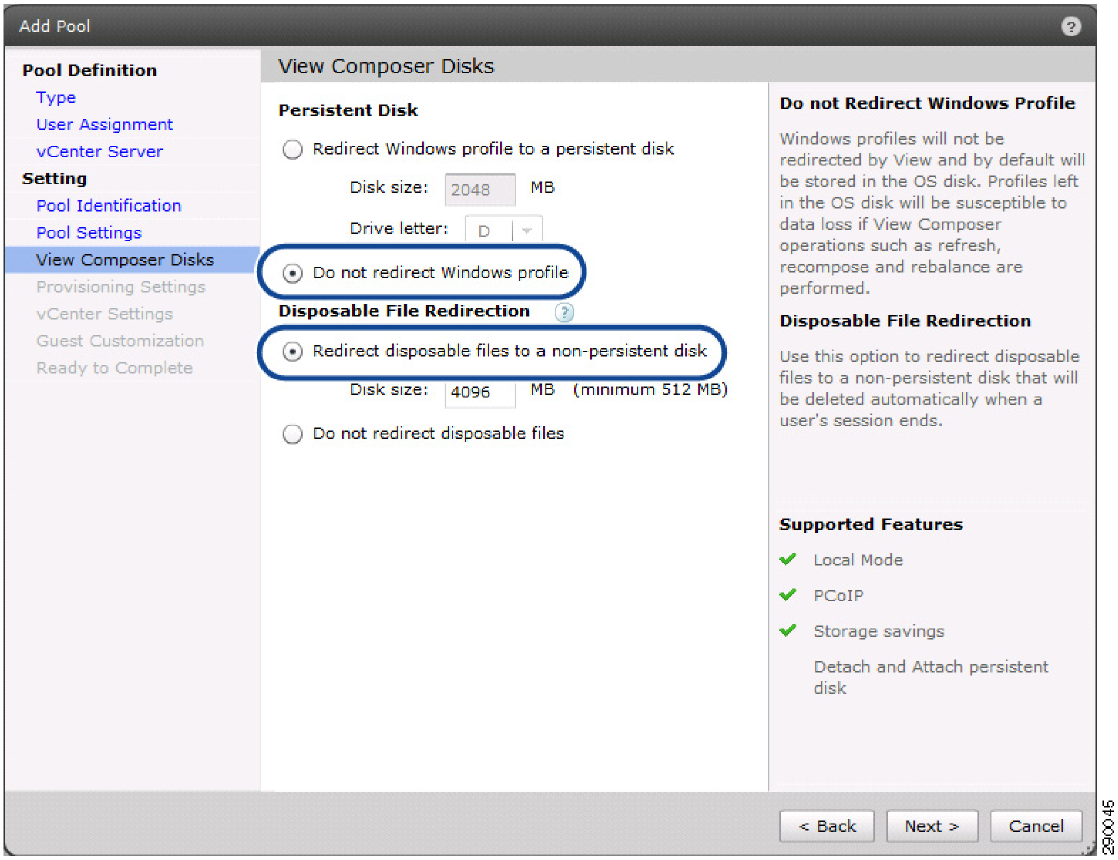

View Composer Disks

Persistent Disk

○ Redirect Windows profile to a persistent disk

Disk size: 2048 MB

Drive letter: D

● Do not redirect Windows profile

Disposable File Redirection

● Redirect disposable files to a non-persistent disk

Disk size: 4096 MB (minimum 512 MB)

○ Do not redirect disposable files

Do not Redirect Windows Profile

Windows profiles will not be redirected by View and by default will be stored in the OS disk. Profiles left in the OS disk will be susceptible to data loss if View Composer operations such as refresh, recompose and rebalance are performed.

Disposable File Redirection

Use this option to redirect disposable files to a non-persistent disk that will be deleted automatically when a user's session ends.

Supported Features

✓ Local Mode

✓ PCoIP

✓ Storage savings

Detach and Attach persistent disk

Provisioning Settings specified a naming pattern of “vdi-pool1-” to help correlate the created linked clones to the desktop pool with which they are associated. “Max number of desktops:” specifies the number of VMs to create for the pool:

Figure 37 Provisioning Settings and Maximum Number of Desktops

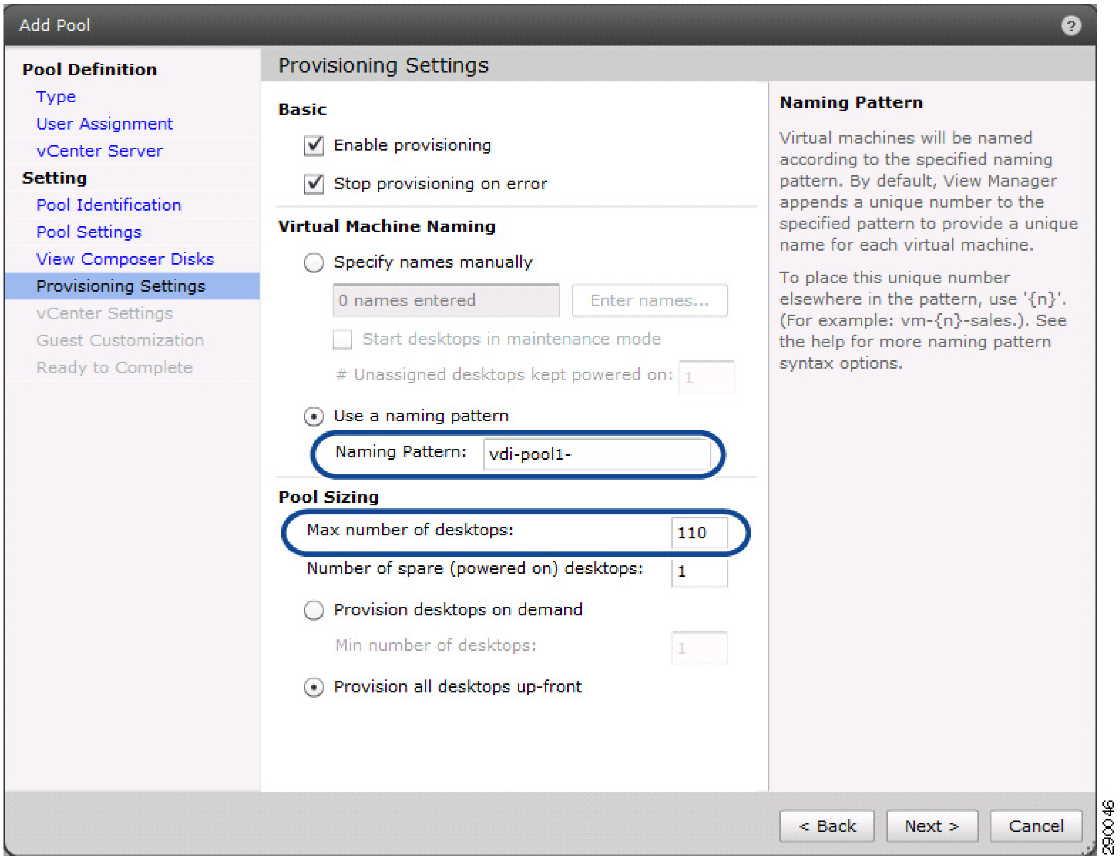

Provisioning Settings

Basic

☑ Enable provisioning

☑ Stop provisioning on error

Virtual Machine Naming

○ Specify names manually

0 names entered | Enter names...

□ Start desktops in maintenance mode

# Unassigned desktops kept powered on: 1

● Use a naming pattern

Naming Pattern: vdi-pool1-

Pool Sizing

Max number of desktops: 110

Number of spare (powered on) desktops: 1

○ Provision desktops on demand

Min number of desktops: 1

● Provision all desktops up-front

Naming Pattern

Virtual machines will be named according to the specified naming pattern. By default, View Manager appends a unique number to the specified pattern to provide a unique name for each virtual machine.

To place this unique number elsewhere in the pattern, use '{n}'. (For example: vm-{n}-sales.). See the help for more naming pattern syntax options.

vCenter Settings specified the:

- Default Image and the appropriate snapshot from which to create the linked clones.
- VM folder on vCenter to place the pool VMs
- Host or Cluster on which to create the VMs
- Resource Pool if created on a cluster
- Datastore and the over subscription policy, which was selected as Aggressive

Figure 38 vCenter Settings

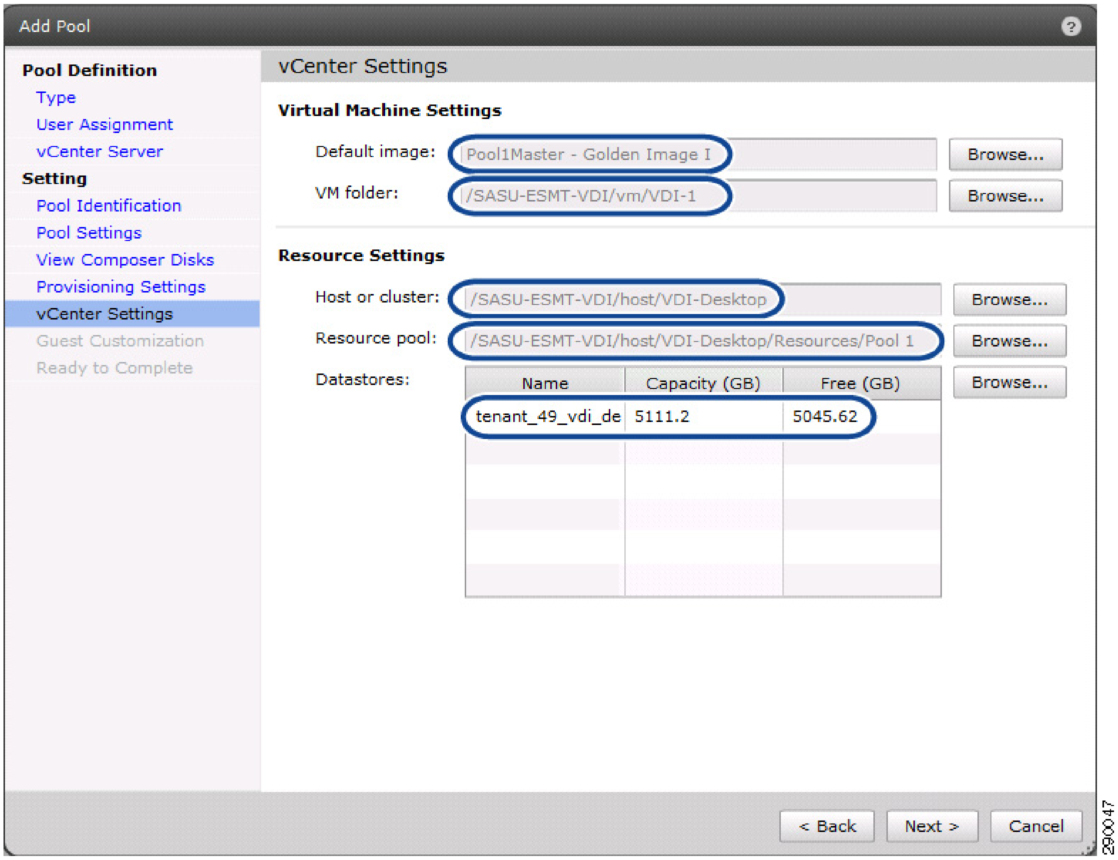

vCenter Settings

Virtual Machine Settings

Default image: Pool1Master - Golden Image I

VM folder: /SASU-ESMT-VDI/vm/VDI-1

Resource Settings

Host or cluster: /SASU-ESMT-VDI/host/VDI-Desktop

Resource pool: /SASU-ESMT-VDI/host/VDI-Desktop/Resources/Pool 1

Datastores:

Name	Capacity (GB)	Free (GB)
tenant_49_vdi_de	5111.2	5045.62

Guest Customization requires the selection of the Domain to which the VMs will be joined and the AD container in which they will be placed:

Figure 39 Guest Customization

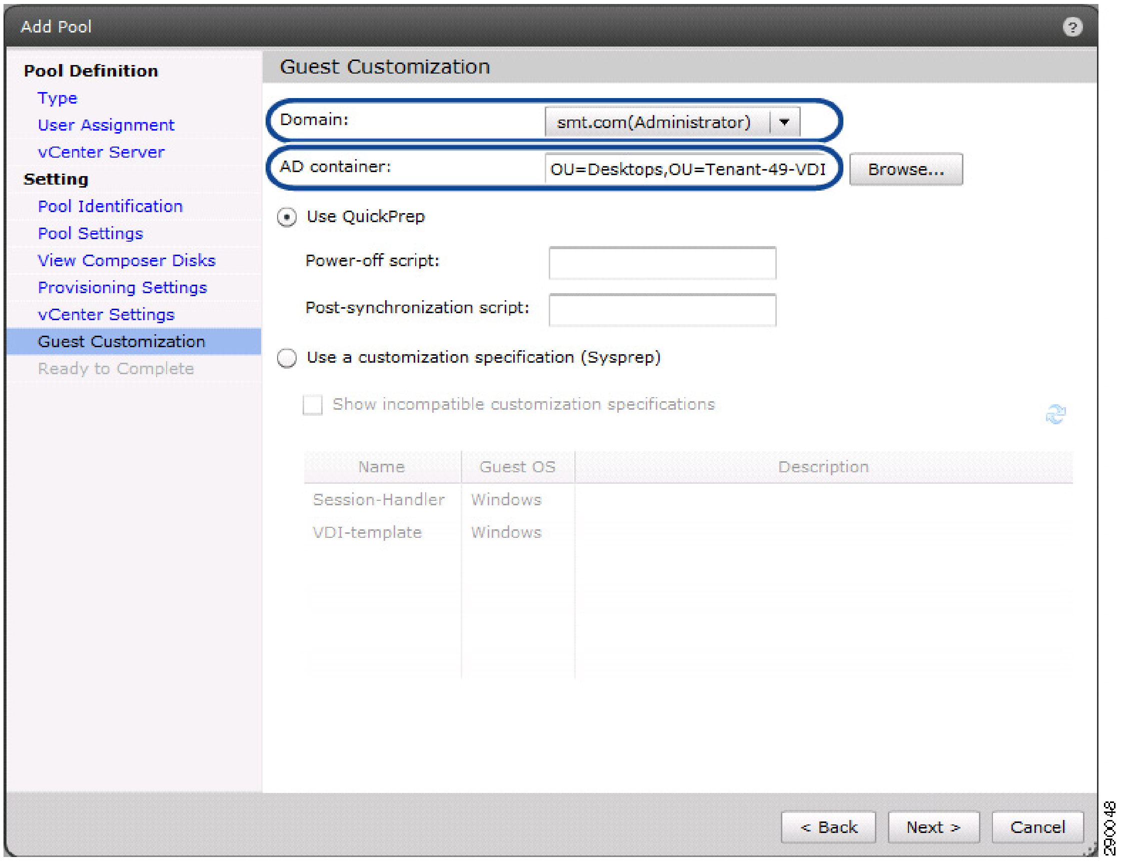

Add Pool

Pool Definition
- Type
- User Assignment
- vCenter Server

Setting
- Pool Identification
- Pool Settings
- View Composer Disks
- Provisioning Settings
- vCenter Settings
- Guest Customization
- Ready to Complete

Guest Customization

Domain: smt.com(Administrator)

AD container: OU=Desktops,OU=Tenant-49-VDI

● Use QuickPrep

Power-off script:

Post-synchronization script:

○ Use a customization specification (Sysprep)

□ Show incompatible customization specifications

Name	Guest OS	Description
Session-Handler	Windows	
VDI-template	Windows	

A summary of the selected options is presented before the pool is created:

Figure 40 Summary of Selected Options

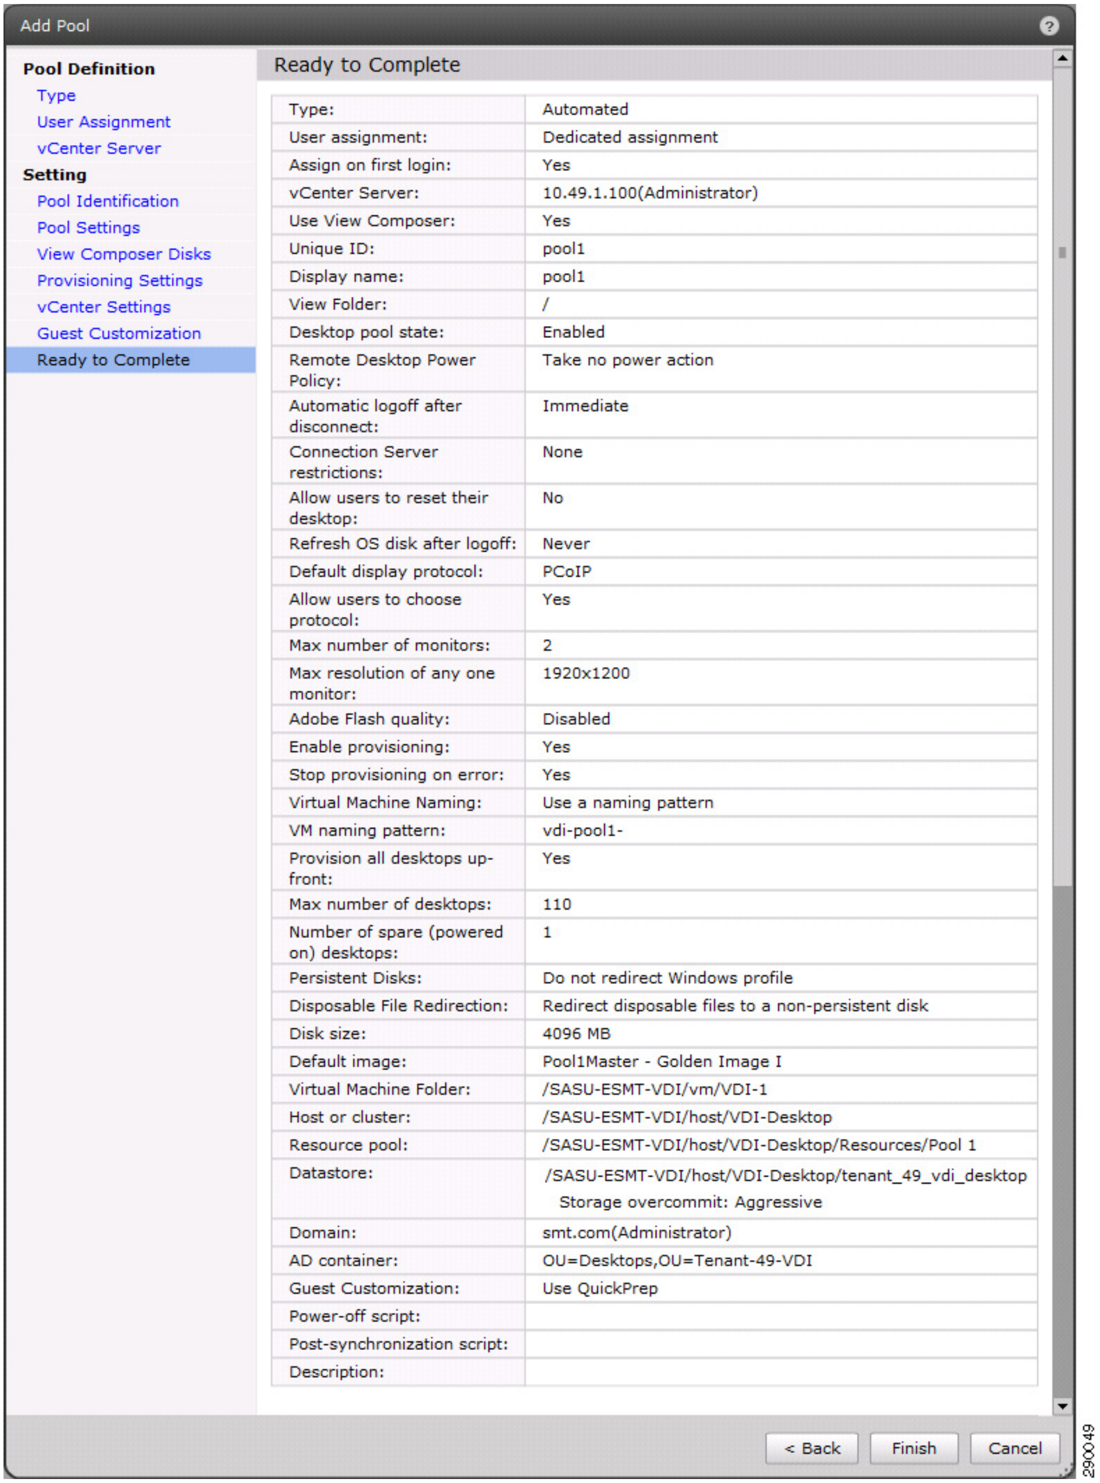

Ready to Complete

Type:	Automated
User assignment:	Dedicated assignment
Assign on first login:	Yes
vCenter Server:	10.49.1.100(Administrator)
Use View Composer:	Yes
Unique ID:	pool1
Display name:	pool1
View Folder:	/
Desktop pool state:	Enabled
Remote Desktop Power Policy:	Take no power action
Automatic logoff after disconnect:	Immediate
Connection Server restrictions:	None
Allow users to reset their desktop:	No
Refresh OS disk after logoff:	Never
Default display protocol:	PCoIP
Allow users to choose protocol:	Yes
Max number of monitors:	2
Max resolution of any one monitor:	1920x1200
Adobe Flash quality:	Disabled
Enable provisioning:	Yes
Stop provisioning on error:	Yes
Virtual Machine Naming:	Use a naming pattern
VM naming pattern:	vdi-pool1-
Provision all desktops up-front:	Yes
Max number of desktops:	110
Number of spare (powered on) desktops:	1
Persistent Disks:	Do not redirect Windows profile
Disposable File Redirection:	Redirect disposable files to a non-persistent disk
Disk size:	4096 MB
Default image:	Pool1Master - Golden Image I
Virtual Machine Folder:	/SASU-ESMT-VDI/vm/VDI-1
Host or cluster:	/SASU-ESMT-VDI/host/VDI-Desktop
Resource pool:	/SASU-ESMT-VDI/host/VDI-Desktop/Resources/Pool 1
Datastore:	/SASU-ESMT-VDI/host/VDI-Desktop/tenant_49_vdi_desktop Storage overcommit: Aggressive
Domain:	smt.com(Administrator)
AD container:	OU=Desktops,OU=Tenant-49-VDI
Guest Customization:	Use QuickPrep
Power-off script:	
Post-synchronization script:	
Description:	

After the pool has been created, it must be entitled to an AD group of intended users:

Figure 41 Group of Intended Users

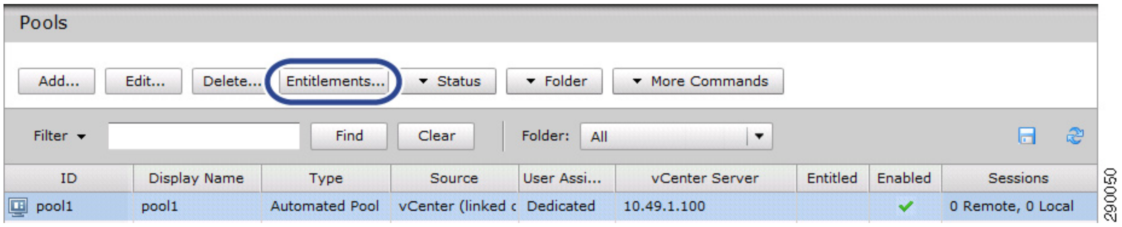

ID	Display Name	Type	Source	User Assi...	vCenter Server	Entitled	Enabled	Sessions
pool1	pool1	Automated Pool	vCenter (linked c	Dedicated	10.49.1.100		✓	0 Remote, 0 Local

With the pool created and entitled, desktops are ready for access by specified users with the View Client, which can be downloaded from the base URL of the Connection Server.

Test Setup and Configurations

Figure 42 Test Environment Used for Server Scalability Testing

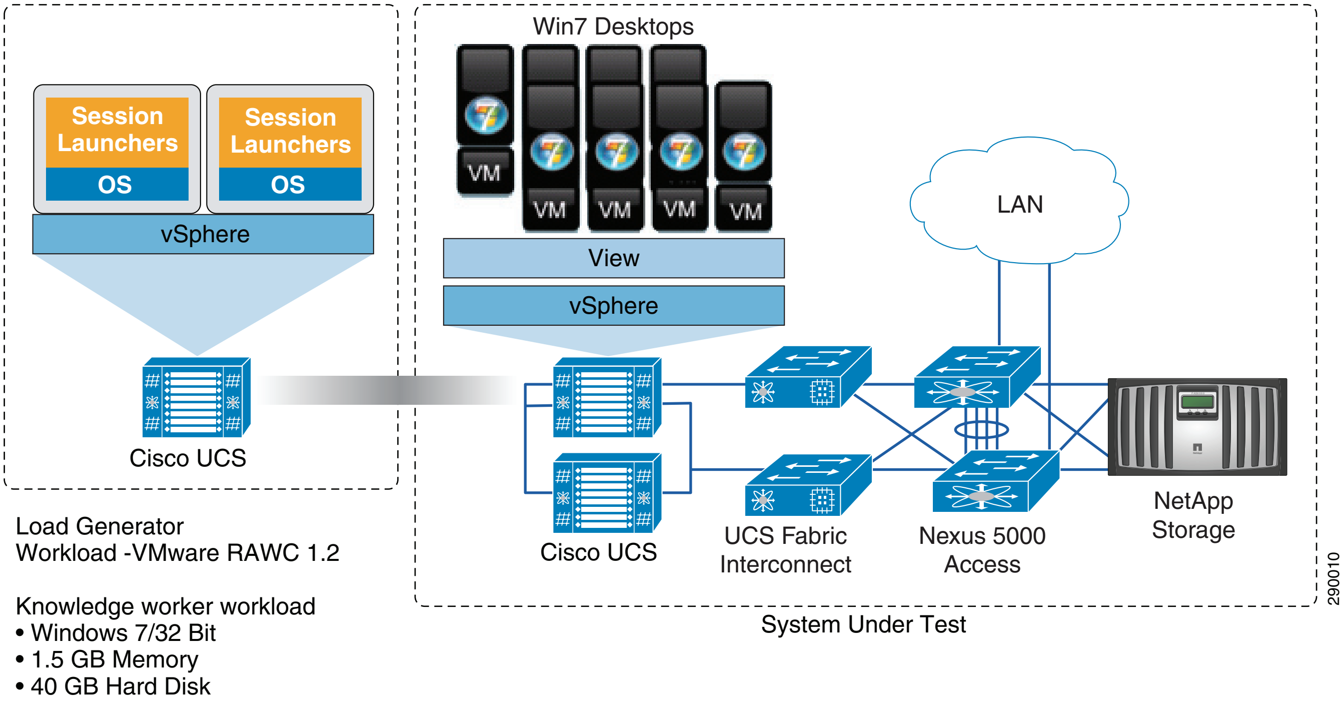

Cisco UCS Test Configuration for Single Server Scalability Test Setup

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)
- Two Cisco Virtual Interface Card Adapters
- Cisco Nexus 5000 and 7000

- NetApp FAS 6080 storage array, two controllers, 2 X Dual port 10G Chelsio cards with 70 SAS drives

Software Components

- Cisco UCS firmware 1.3(1n)
- VMware vSphere 4.1, Virtual Center 4.1
- VMware View 4.5
- Windows 7—32 bit, 1vCPU, 1.5 GB of memory, 40 GB/VM

Cisco UCS Configuration for Two-Chassis Test

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)
- Two Cisco Virtual Interface Card Adapters on each blade server
- Cisco Nexus 5000 and 7000
- NetApp FAS 6080 storage array, two controllers, 2 X Dual port 10G Chelsio cards with 70 SAS drives

Software Components

- Cisco UCS firmware 1.3(1n)
- VMware vSphere 4.1 (ESXi), Virtual Center 4.1
- VMware View 4.5
- Windows 7—32 bit, 1vCPU, 1.5 GB of memory, 40 GB/VM

Testing Methodology and Success Criteria

VMware's RAWC tool was used to simulate Virtual Desktop Users. This tool is composed of three main components:

- RAWC Controller—Built on Windows 2003 R2 and includes a package known as RAWC controller interface. The interface enables selecting the workload for each desktop login.
- Session Launcher—Built on Windows XP SP1 and .net platforms. The session launcher is installed on and accesses a share on the RAWC Controller. The Session Launcher serves two functions:
  - Captures the login information.
  - Records successful login by writing a log file of the session instance which is started.
- Workload Simulator—Starts running on the Desktop at logon from the View Client and launches applications in real time. The simulator begins reading, writing, and producing data through the specified applications to simulate the activity which a real user would produce.

Load Generation

VMware has developed a user profile for the RAWC tool to simulate a knowledge worker. This user profile includes typical office applications, as shown in Figure 43.

Figure 43 User Profile for RAWC Tool

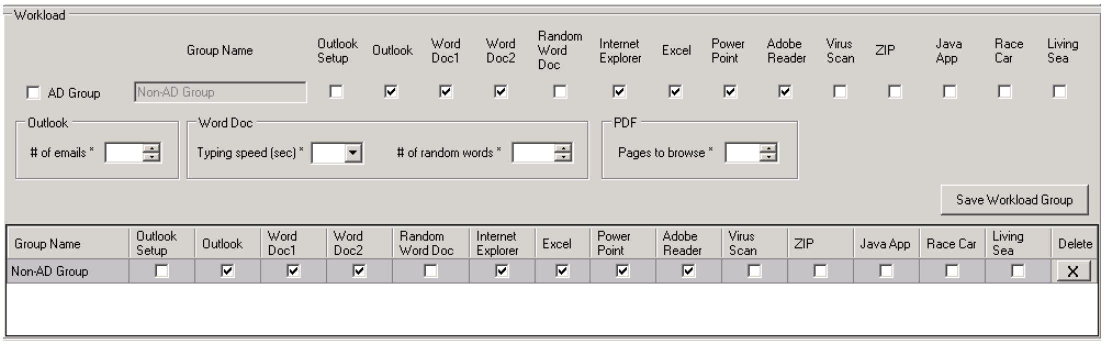

Workload

Group Name	Outlook Setup	Outlook	Word Doc1	Word Doc2	Random Word Doc	Internet Explorer	Excel	Power Point	Adobe Reader	Virus Scan	ZIP	Java App	Race Car	Living Sea
□ AD Group Non-AD Group	□	☑	☑	☑	□	☑	☑	☑	☑	□	□	□	□	□

Outlook: # of emails *

Word Doc: Typing speed (sec) * # of random words *

PDF: Pages to browse *

Group Name	Outlook Setup	Outlook	Word Doc1	Word Doc2	Random Word Doc	Internet Explorer	Excel	Power Point	Adobe Reader	Virus Scan	ZIP	Java App	Race Car	Living Sea	Delete
Non-AD Group	□	☑	☑	☑	□	☑	☑	☑	☑	□	□	□	□	□	X

Each Session Launcher can handle up to 20 user sessions. To begin operation, the Session Launcher starts sessions to the VMware View Server. Each session launches with a specific user account. A process starts on each user session which loads the script to launch specified application activity. The VDI desktop machine mounts a share on the RAWC controller which serves two functions. First the Microsoft Windows 7 Desktop logs a successful login, then it reads the configuration from the output file and begins the workload process.

User Workload Simulation

The RAWC tool is designed to simulate the activity of typical desktop users. Most users check their E-mail first after logging in; hence Outlook is the first application to run. Thereafter, applications are run in random order and with various operations, including opening and closing files, creating files, entering data, and saving files. The applications included in the workload are Microsoft Office 2007 Professional Plus products: Outlook, two Microsoft Word documents, Multiple Microsoft Excel Spreadsheets, and Microsoft PowerPoint. It also includes opening and browsing a document with Adobe Reader and surfing the Web with Microsoft Internet Explorer.

For more information about the VMware RAWC tool, see: http://www.vmware.com/files/pdf/VMware-WP-WorkloadConsiderations-WP-EN.pdf.

Workload Verification Criteria

The test procedure determined a successful workload run by evaluating three criteria:

- Workload completion and timing written to log files on a share from the RAWC Controller
- Visual verification of View Client logins within the RDP sessions initiated to the Session Launchers by the RAWC Controller
- Data collected from remote esxtop generated from a vMA appliance

Test Results

The purpose of this testing is to provide the data needed to validate VMware View 4.5 in VMware vSphere 4.1 virtualizing Microsoft Windows 7 desktops on Cisco UCS blade servers implemented with NetApp's FAS6080 storage system. The test results are divided into two sections:

- Results for Single Cisco UCS Blade Server Validation
- Results for Eight Cisco UCS Blade Server Validation

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 VMware View 4.5 with VMware vSphere 4.1 scalability.

VMware View 4.5 VDI Test Results

This section details the results from the VMware View 4.5 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 eight blade (two chassis) configuration.

Results for Single Cisco UCS Blade Server Validation

The first process in the validation was to find out how many Virtual desktops could 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 the blade had 192 GB of RAM available.

The three criteria used for validation were:

- Application Open time, measured in RAWC, is less than 2.5 seconds average for each application in the workload.
- Scaling of the desktops until a point was reached where the average Physical CPU Percent Utilization was less than 80%. This was determined between the steady state period where all desktops were running the workload.
- System Ballooning is less than 20% of aggregate memory.

The Physical CPU Percent Utilization of the server hosting each of the desktop configurations was captured. As shown in Figure 44, the processor utilization pattern is consistent across each of the test iterations. It was determined that the 160 desktop count allowed approximately 30% of the physical processor capacity to be available for workload spikes and unforeseen failure conditions.

Figure 44 ESX Host Physical CPU Percentage Utilization Time

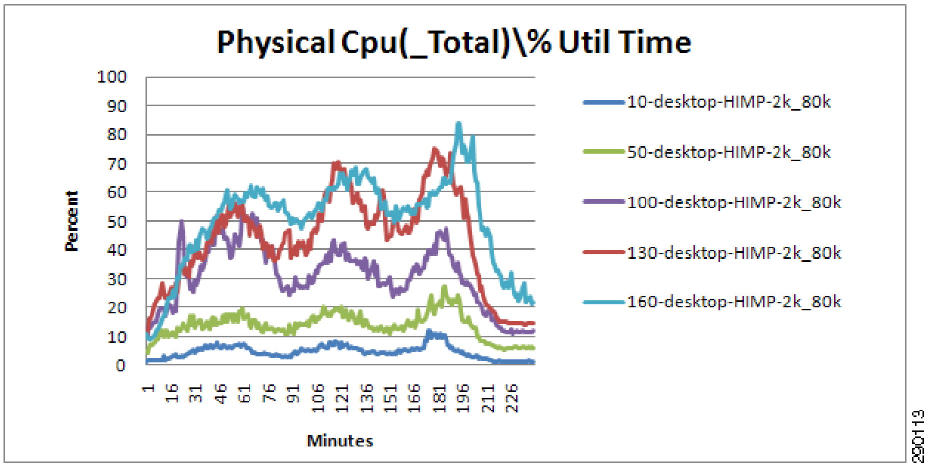

Figure 45 shows the physical hardware execution context or Physical CPU Percent Utilization (PCPU) values versus the Physical CPU % Core Utilization (Core Util) values captured during testing. The Core Util is the percent of time that at least one of the threads on a hyperthreaded core is used, averaged over all cores. The PCPU is the percent of time each individual thread is used, averaged over all threads. The PCPU may refer to a physical CPU core if hyperthreading is unavailable or a logical CPU (LCPU) if hyperthreading is enabled, which makes PCPU a better metric for what is commonly thought of as CPU utilization.

The Intel Westmere processors employ hyperthreading, so the B250-M2 platforms with two six core Westmere processors will have a total of 24 LCPU. The PCPU utilization percentage is available for each LCPU and represents the percentage of real time that the PCPU (LCPU) was not idle. Figure 45 indicates the percentage CPU utilization averaged over all Physical CPUs for the 160 desktop workload. The values observed in Figure 45 confirm the LCPU processors have resources available to manage workload spikes and unforeseen failure conditions.

Figure 45 ESX Host Physical CPU Percent Utilization for Single B250 160 User Workload

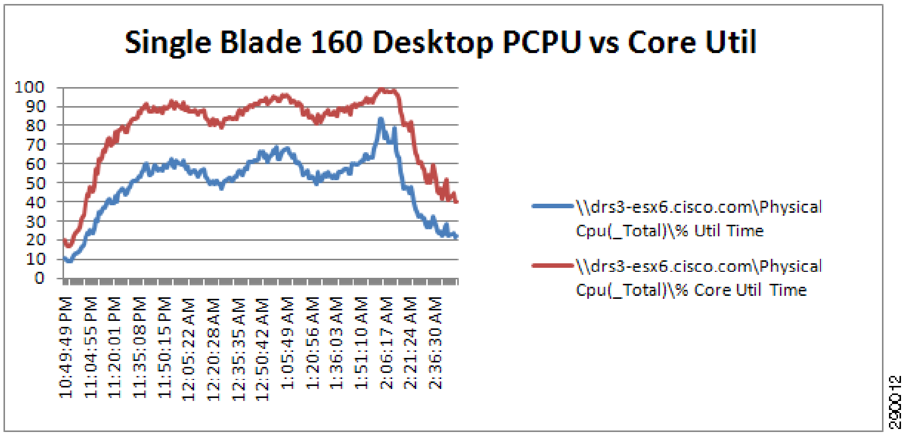

Memory utilization on the ESX host falls into two general categories, either it is consumed by the ESX host kernel or it is dedicated to the virtual machines, i.e., desktops. Without vsish memory reclamation in place, the default of 6% leaves approximately 180GB of memory available to the virtual desktops.

With 1.5 GB allocated per desktop, little to no ballooning can be seen up to the 120 desktop level. At the 160 desktop load level, swapping was observed, but at a level that accounted for less than 1% of aggregate memory.

For more information on memory resource management in vSphere, see: http://www.vmware.com/pdf/usenix_resource_mgmt.pdf.

Figure 46 shows that at and beyond 130 desktops, all of the available non-kernel memory is being used by the desktop VMs. vSphere responds to this memory pressure by using the ballooning mechanism to reclaim memory from the VMs, as shown in Figure 47.

Figure 46 Non-Kernel Memory Utilization

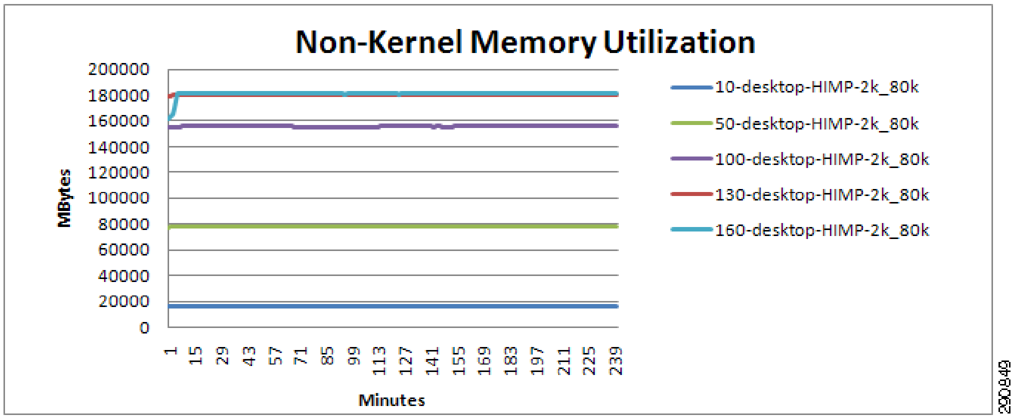

Figure 47 shows the amount of memory ballooning that was observed during differing workloads. With 160 desktops, the peak ballooning was 7-11GB (4-6%), which is well below the 20% criteria.

Figure 47 Memory Ballooning

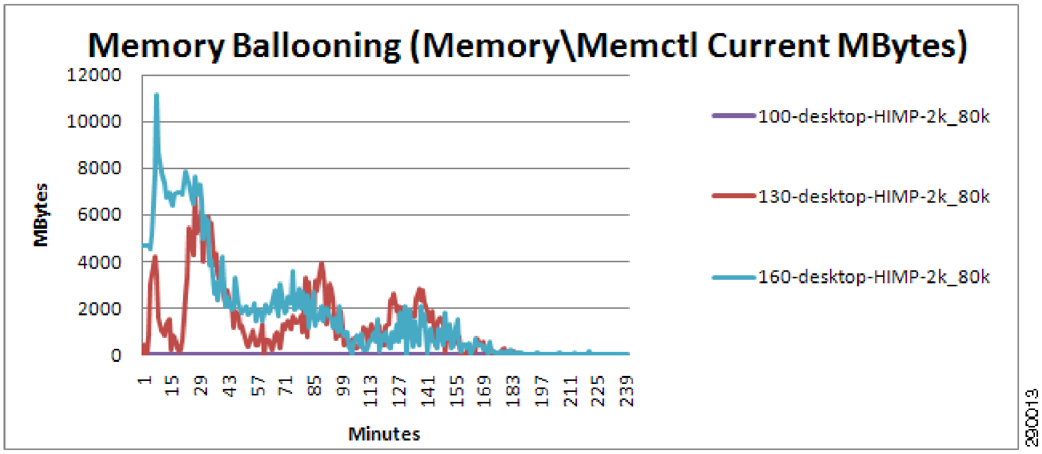

User experience was analyzed as a metric of average open times of the applications within the workload, with an objective of keeping application access consistent with a non-constrained environment that stayed under 2.5 seconds for each application. As shown in Figure 48, application open times remained consistent through the 160 desktop runs.

Figure 48 Average Application Open Time

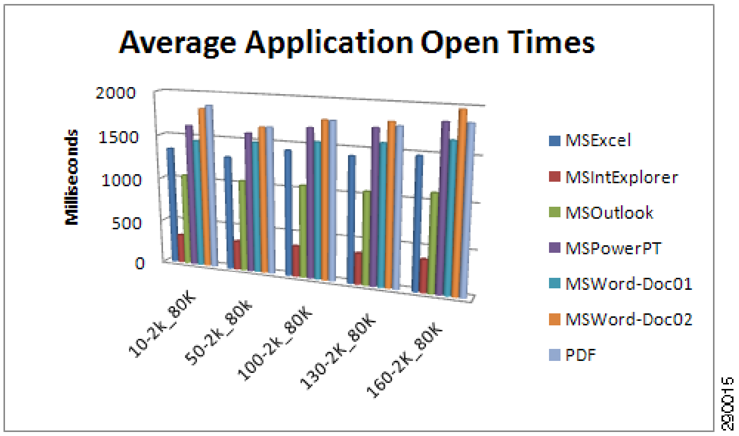

The NetApp FAS controllers record the number of IOPS which occurred during each test case. The NFS-based IOPS for the 160 desktop iteration are shown in Figure 49; this chart depicts the NFS IOPS for the VDI vFiler. This number represents the mix of reads, writes, and other operations taking place during the knowledge worker workload.

Figure 49 Single Blade NFS IOPS for the VDI vFiler

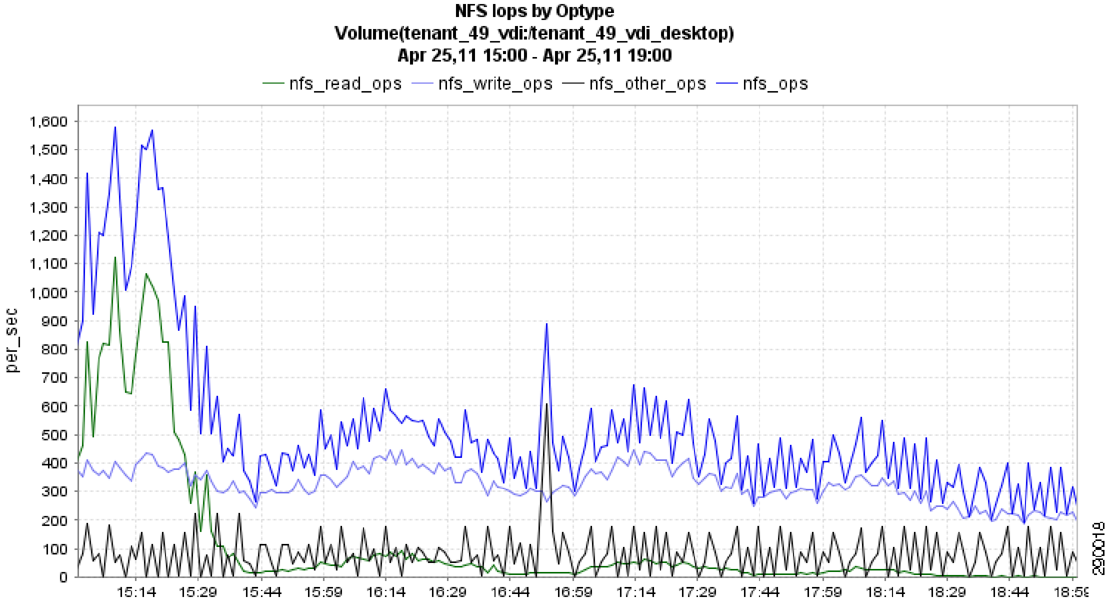

Results for Eight Cisco UCS Blade Server Validation

The eight blade load test focused on the scalability and predictability of the VDI environment. The eight blade test scenario leveraged the findings of the single blade tests by deploying 160 virtual desktops per blade. The following test conditions were created:

- Each host was provisioned with 160 Desktops VMs.
- Four pools of 320 VMs were used within the cluster of eight hosts.

- Automated vMotion was turned off within DRS.
- Test workloads were continued at a 0-24 minute delay.
- Session Launchers were increased to 20 users per Launcher.
- Delay of the RDP connection between each Session Launcher was increased from 3 seconds to 20 seconds.

The eight blade tests showed predictable resource utilization well aligned with the single blade test results. The processor and memory utilization of each blade performed within the performance expectations derived from the single blade tests. The network utilization and IOPS rate grew in a predictable manner.

Note Each of the combined performance charts depicted below correspond to the single blade server charts previously described.

Figure 50 Eight Host Physical CPU Percentage Utilization Time—1280 Desktop Run

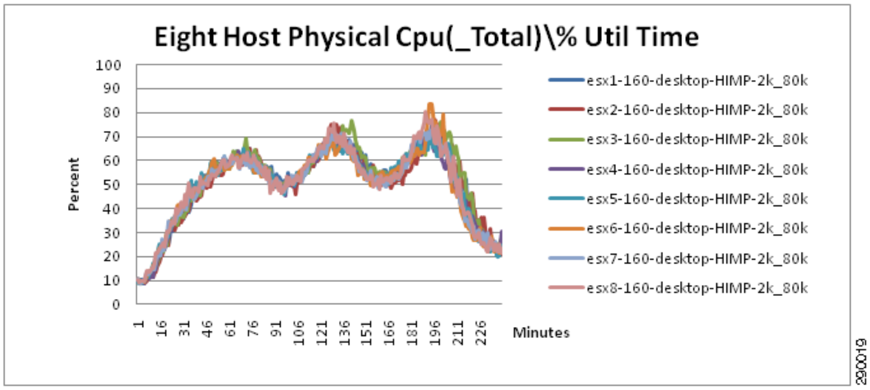

Figure 51 Eight Blade NFS IOPS for the VDI vFiler

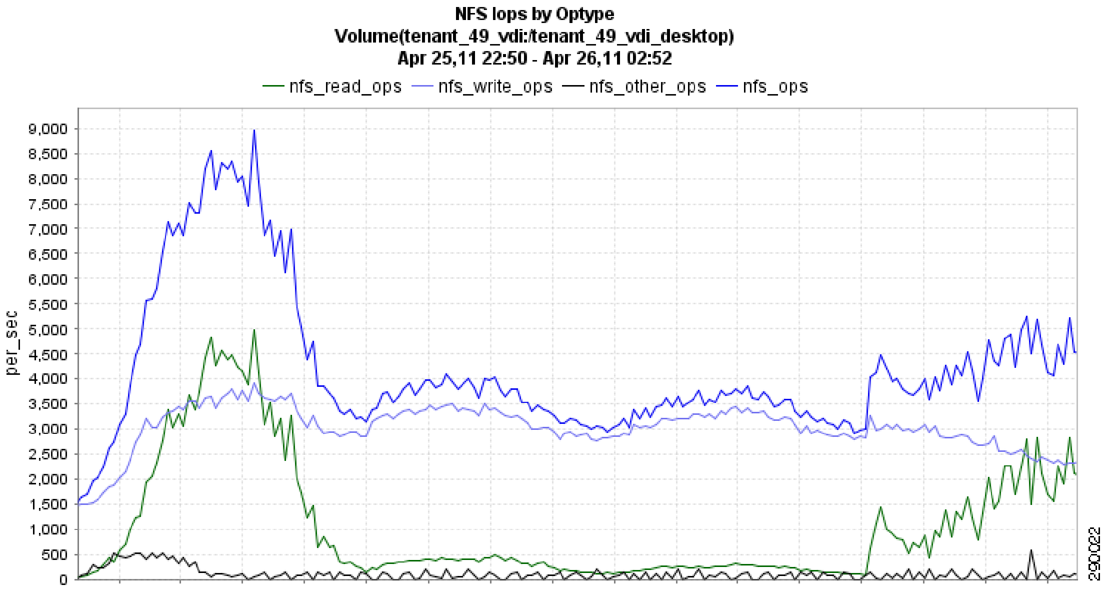

Scaling and Sizing Guidelines

There are many factors to consider when you begin to scale beyond two chassis or eight servers, which this reference architecture has successfully tested. In this section we give guidance to scale beyond two UCS chassis.

Scalability Considerations and Guidelines

The 160 Desktop load for the single blade characterization was selected as our baseline because it met the three criteria of:

- Application Open time, measured in RAWC, was less than 2.5 seconds average for each application in the workload.
- The desktops were scaled until a point was reached where the average Physical CPU Percent Utilization was less than 80%. This was determined between the steady state period where all desktops were running the workload.
- System Ballooning was less than 20% of aggregate memory.

Cisco UCS System Configuration

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

Table 13 UCS System Configuration—Servers Tested

vSphere			
No. of Chassis	**No. of B250-M2 Servers Tested**	**No. of VMs**	**VMs/Core**
1	1 Blade	160	13.33
2	8 Blades	1280	

Extrapolating the values we got during the testing onto a dedicated desktop chassis, we get the following results:

Table 14 UCS System Configuration—Server Extrapolation

vSphere			
No. of Chassis	**No. of B250-M2 Servers**	**No. of VMs**	**VMs/Core**
4	16 Blades	2560	13.33
8	32 Blades	5120	
12	48 Blades	7680	
16	64 Blades	10240	

The backend storage has to be scaled accordingly, based on the IOP considerations described in NetApp TR-3770, referenced in Storage Architecture Best Practices.

VMware has additional references in its architecture planning guide that details how to scale their components as you scale the number of desktops: http://www.vmware.com/pdf/view45_architecture_planning.pdf. You should also refer to Storage Architecture Best Practices.

vSphere Configuration

On vSphere 4.1, the configuration maximums have to be taken in to consideration while doing scaling calculations. Refer to the Configuration Maximums for VMware vSphere 4.1: http://www.vmware.com/pdf/vsphere4/r41/vsp_41_config_max.pdf.

The main parameters are:

- DRS Cluster
  - Hosts per DRS cluster[1]
  - 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[1]

[1] vSphere supports up to 32 hosts in a DRS/HA cluster, while View 4.5 does not allow more than eight hosts in a 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.

Sizing Guidelines

One of the key findings from this study was the characterization of Microsoft Windows 7 desktop profile with respect to CPU and memory. All of the physical memory of the single blade server with 192 GB of memory was consumed by the 160 active desktop sessions with each Windows 7 virtual desktop configured with 1.5 GB of RAM.

Extending beyond the physical memory required over-subscription. VMware supports a ratio of up to 1.5 memory oversubscription. The 160 desktop load ran at a 1.33 memory oversubscription ratio. This was primarily resolved through memory ballooning, though less than 1% of the memory required swapping. With this ballooning and swapping there was little to no user experience degradation observed.

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 = (12 * 3.33)*0.80*1000/160 = 199.8 MHz

If your workload is similar to the knowledge worker workload used, then each desktop will consume approximately 200 MHz on average.

Memory Calculations

VMware vSphere reserves approximately six 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 memory calculations.

Let us consider the scalability of Windows 7 with 1.5 GB RAM on a different UCS blade server. The Cisco B200-M2 blade server with 96 GB of memory and Intel 5680 processor at 3.33 GHz speed will be limited by memory and not by CPU when running the knowledge worker load.

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

In our example, the recommended load of desktops/blade to avoid memory oversubscription = (96 * 0.94) / 1.5 = 61 desktops (approximately).

In this case a lower series CPUs could be used instead of the 3.33 GHz (5680) or retained to have spare CPU cycles for load spikes.

Sizing Summary

This testing provides sizing recommendations based on a specific, controlled workload. The results were focused on keeping an environment that provided peak performance to desktop users.

Your workload will vary from the workload used in this testing. The 160 desktops per blade allows a good rule of thumb for sizing, but the scalability of the VDI solution brought together with Cisco, NetApp, and VMware allows you to easily expand your environment when needed.

References

FlexPod

FlexPod for VMware Deployment Model:
http://www.cisco.com/en/US/docs/solutions/Enterprise/Data_Center/Virtualization/flexpod_vmware.html

Cisco

Cisco Enhanced Secure Multi-Tenancy (ESMT) architecture:
http://www.cisco.com/en/US/solutions/ns340/ns414/ns742/ns743/ns1050/landing_dcVDDC.html.

Cisco Secure Multi-Tenancy Deployment Guide:
http://www.cisco.com/en/US/docs/solutions/Enterprise/Data_Center/Virtualization/securecldeployg.html#wp99931

Cisco UCS chassis install guide:
:http://www.cisco.com/en/US/docs/unified_computing/ucs/hw/chassis/install/ucs5108_install.html)

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

Cisco Nexus 5000 Series features:
http://www.cisco.com/en/US/products/ps9670/prod_white_papers_list.html

Installation and configuration of the Cisco Nexus 1000V and VEM:

- http://www.cisco.com/en/US/partner/docs/switches/datacenter/nexus1000/sw/4_2_1_s_v_1_4/install/software/guide/n1000v_install_software.html
- http://www.cisco.com/en/US/products/ps9902/products_installation_and_configuration_guides_list.html
- http://www.cisco.com/en/US/docs/switches/datacenter/nexus1000/sw/4_2_1_s_v_1_4/install/vem/guide/n1000v_vem_install.html

NetApp

NetApp Flash Cache technology:
http://www.netapp.com/us/products/storage-systems/flash-cache/flash-cache-tech-specs.html

NetApp Operations Manager:
http://www.netapp.com/us/products/management-software/operations-manager.html

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

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

NetApp TR-3298: RAID-DP: NetApp Implementation of RAID Double Parity for Data Protection: http://media.netapp.com/documents/tr-3298.pdf

NetApp TR-3770: VMware View on NetApp Deployment Guide: http://media.netapp.com/documents/tr-3770.pdf

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

NetApp TR-3450: Active-Active Controller Overview and Best Practices Guidelines: http://media.netapp.com/documents/tr-3450.pdf

NetApp TR-3749: NetApp and VMware vSphere Storage Best Practices: http://media.netapp.com/documents/tr-3749.pdf

Netapp TR-3705: Netapp and VMware View Solution Guide: http://media.netapp.com/documents/tr-3705.pdf

VMware

Configuration Maximums for VMware vSphere 4.1: http://www.vmware.com/pdf/vsphere4/r41/vsp_41_config_max.pdf.

Provisioning Microsoft Windows 7 for View 4.5: http://www.vmware.com/files/pdf/VMware-View-OptimizationGuideWindows7-EN.pdf

VMware architecture planning guide: http://www.vmware.com/pdf/view45_architecture_planning.pdf.

VMware View 4.5 Installation Guide: http://www.vmware.com/pdf/view45_installation_guide.pdf.

VMware RAWC tool: http://www.vmware.com/files/pdf/VMware-WP-WorkloadConsiderations-WP-EN.pdf.

Appendix A—Select Configurations

Cisco Nexus 5000 Configuration

```
!Command: show running-config
!Time: Tue Nov 23 16:49:14 2010

version 4.2(1)N1(1)
feature fcoe
feature npiv
feature telnet
feature tacacs+
feature udld
feature lacp
feature vpc
feature lldp
logging level aaa 5
logging level cdp 6
logging level vpc 6
logging level radius 5
logging level tacacs 5
logging level monitor 6
logging level port-channel 6
logging level spanning-tree 6
role feature-group name Network-Operator
  feature zone
  feature cdp
  feature arp
  feature ping
  feature vsan
  feature vlan
  feature fdmi
  feature fcns
  feature fspf
  feature rscn
role feature-group name Network-Admin
  feature aaa
  feature arp
  feature cdp
  feature l3vm
  feature ping
  feature snmp
  feature radius
  feature tacacs
  feature install
  feature license
  feature callhome
  feature platform
  feature access-list
  feature svi
  feature eth-span
  feature ethanalyzer
  feature spanning-tree
  feature acl
  feature sfm
  feature fcns
  feature fcsp
  feature fdmi
  feature fspf
  feature rlir
  feature rscn
  feature span
  feature vsan
  feature wwnm
  feature zone
  feature fcanalyzer
role name Network-Operator
  rule 1 permit read feature-group Network-Operator
role name Network-Admin
  rule 1 permit read-write feature-group Network-Admin
username admin password 5 $1$43gsLyBb$Xsht0lHW7wbbq9nMGo1HX.  role network-admin
ip domain-lookup
ip host DC05-N5K-1 x.x.x.28
tacacs-server key 7 "K1kmN0gy"
ip tacacs source-interface mgmt0
tacacs-server host x.x.x.214
tacacs-server host x.x.x.216 timeout 3
aaa group server tacacs+ TacacsServer
    server x.x.x.216
    deadtime 1
    use-vrf management
    source-interface mgmt0
hostname DC05-N5K-1
logging event link-status default
```

```
service unsupported-transceiver
ip access-list classify_COS_5
  10 permit ip 192.168.98.100/32 any
  20 permit ip 10.49.32.10/32 any
  30 permit ip any 192.168.98.100/32
  40 permit ip any 10.49.32.10/32
class-map type qos Platinum_Traffic
  match access-group name classify_COS_5
class-map type qos Gold_Transactional
  match cos 6
class-map type qos Bronze_Transactional
  match cos 2
class-map type qos Silver_Transactional
  match cos 4
class-map type qos Platinum_Transactional
  match cos 5
class-map type queuing Gold_Traffic_Q
match qos-group 3
class-map type queuing Bronze_Traffic_Q
match qos-group 5
class-map type queuing Silver_Traffic_Q
match qos-group 4
class-map type queuing Platinum_Traffic_Q
match qos-group 2
policy-map type qos Global_Classify_NFS_Application
  class Platinum_Traffic
    set qos-group 2
  class Platinum_Transactional
    set qos-group 2
  class Gold_Transactional
    set qos-group 3
  class Silver_Transactional
    set qos-group 4
  class Bronze_Transactional
    set qos-group 5
policy-map type queuing Global_BW_Queuing
  class type queuing Platinum_Traffic_Q
    priority
    bandwidth percent 20
  class type queuing Gold_Traffic_Q
    bandwidth percent 20
  class type queuing Silver_Traffic_Q
    bandwidth percent 20
  class type queuing Bronze_Traffic_Q
    bandwidth percent 15
  class type queuing class-fcoe
    bandwidth percent 15
  class type queuing class-default
    bandwidth percent 10
class-map type network-qos Gold_Traffic_NQ
  match qos-group 3
class-map type network-qos Bronze_Traffic_NQ
  match qos-group 5
class-map type network-qos Silver_Traffic_NQ
  match qos-group 4
class-map type network-qos Platinum_Traffic_NQ
  match qos-group 2
policy-map type network-qos Netapp_Qos
  class type network-qos Platinum_Traffic_NQ
    set cos 5
    queue-limit 30000 bytes
    mtu 9000
  class type network-qos Gold_Traffic_NQ
    set cos 6
    mtu 9000
  class type network-qos Silver_Traffic_NQ
    set cos 4
    mtu 9000
  class type network-qos Bronze_Traffic_NQ
    set cos 2
    mtu 9000
  class type network-qos class-default
    mtu 9000
system qos
  service-policy type qos input Global_Classify_NFS_Application
  service-policy type queuing output Global_BW_Queuing
  service-policy type network-qos Netapp_Qos
snmp-server user admin network-admin auth md5 0xfcb3ba35cacbdb0b2b72f5ea4dea6364 priv
0xfcb3ba35cacbdb0b2b72f5ea4dea6364 localizedkey
snmp-server host x.x.x.204 traps version 2c public  udp-port 2162
snmp-server host x.x.x.34 traps version 2c public  udp-port 1163
snmp-server host x.x.x.34 traps version 2c public  udp-port 1164
snmp-server enable traps entity fru
snmp-server community public group network-operator
aaa authentication login default group TacacsServer local
vrf context management
  ip route 0.0.0.0/0 x.x.x.1
vlan 1
vlan 2
  name default-native
vlan 101
  name ten1-101
vlan 105
  name ten1-105
vlan 106
  name ten1-106
vlan 107
  name ten1-107
vlan 109
  name ten1-109
vlan 110
  name ten1-110
vlan 113
  name ten2-113
vlan 119
  name ten2-119
vlan 120
  name ten2-120
vlan 122
  name ten3-122
vlan 123
  name ten3-123
vlan 124
  name ten3-124
vlan 129
  name ten3-129
vlan 130
  name ten3-130
vlan 132
  name ten4-132
vlan 140
  name ten4-140
vlan 162
  name mgmt
vlan 581
  name ten49-VDI-Infra-581
vlan 582
  name ten49-VDI-RAWC-582
```

```
vlan 583
  name ten49-VDI-Desktop-583
vlan 584
  name ten49-584
vlan 585
  name ten49-585
vlan 586
  name ten49-586
vlan 587
  name ten49-587
vlan 588
  name ten49-588-iscsi
vlan 589
  name ten49-589-nfs
vlan 590
  name ad-services
vlan 592
  name ten50-592
vlan 599
  name ten50-599
vlan 600
  name ten50-600
vlan 900
  name drs1-ctrl-unlimited
vlan 901
  name drs1-ctrl-limited
vlan 902
  name drs1-monitoring
vlan 905
  name drs1-vc-snap
vlan 907
  fcoe
  name drs1-fabric-b
spanning-tree pathcost method long
spanning-tree port type edge bpduguard default
spanning-tree port type edge bpdufilter default
vpc domain 100
  peer-keepalive destination x.x.x.25
vsan database
  vsan 907 name "drs1-fabric-b"
cfs ipv4 distribute
cfs eth distribute
device-alias database
  device-alias name ESX1_B pwwn 20:00:00:25:b5:06:b1:2f
  device-alias name ESX2_B pwwn 20:00:00:25:b5:06:b1:3f
  device-alias name ESX31_B pwwn 20:00:00:25:b5:03:b2:2f
  device-alias name ESX32_B pwwn 20:00:00:25:b5:03:b2:3f
  device-alias name NetApp1_7b pwwn 50:0a:09:88:87:49:9a:b4
  device-alias name NetApp2_7b pwwn 50:0a:09:88:97:49:9a:b4

device-alias commit

fcdomain fcid database
  vsan 907 wwn 20:41:00:05:9b:74:b5:80 fcid 0x950000 dynamic
  vsan 907 wwn 20:41:00:0d:ec:cf:e5:00 fcid 0x950001 dynamic
  vsan 907 wwn 50:0a:09:88:87:49:9a:b4 fcid 0x950002 dynamic
!              [NetApp1_7b]
  vsan 907 wwn 20:00:00:25:b5:0b:01:0f fcid 0x950003 dynamic
  vsan 907 wwn 20:00:00:25:b5:0b:01:0e fcid 0x950004 dynamic
  vsan 907 wwn 50:0a:09:88:97:49:9a:b4 fcid 0x950005 dynamic
!              [NetApp2_7b]
  vsan 907 wwn 20:00:00:25:b5:06:b1:2f fcid 0x950006 dynamic
!              [ESX1_B]
  vsan 907 wwn 20:00:00:25:b5:06:b2:2f fcid 0x950007 dynamic
  vsan 907 wwn 20:00:00:25:b5:06:b1:3f fcid 0x950008 dynamic
!              [ESX2_B]
  vsan 907 wwn 20:00:00:25:b5:03:b2:2f fcid 0x950009 dynamic
!              [ESX31_B]
  vsan 907 wwn 20:00:00:25:b5:03:b2:3f fcid 0x95000a dynamic
!              [ESX32_B]
  vsan 907 wwn 20:00:00:25:b5:13:b2:0e fcid 0x95000b dynamic
  vsan 907 wwn 20:00:00:25:b5:13:b2:0a fcid 0x95000c dynamic
  vsan 907 wwn 20:00:00:25:b5:13:b2:0c fcid 0x95000d dynamic
  vsan 907 wwn 20:00:00:25:b5:13:b2:0f fcid 0x95000e dynamic
  vsan 907 wwn 20:00:00:25:b5:13:b2:0b fcid 0x95000f dynamic
  vsan 907 wwn 20:00:00:25:b5:13:b2:0d fcid 0x950010 dynamic
  vsan 907 wwn 20:00:00:25:b5:13:b2:19 fcid 0x950011 dynamic


interface port-channel1
  switchport mode trunk
  vpc peer-link
  switchport trunk native vlan 2
  switchport trunk allowed vlan 1-103,105-579,581-905,908-3967,4048-4093
  spanning-tree port type network
  speed 10000

interface port-channel11
  description NetApp1 vPC 11 link Eth1/33 NetApp2 6/a
  switchport mode trunk
  vpc 11
  switchport trunk allowed vlan 106,109-110,115,119-120,124,129-130
  switchport trunk allowed vlan add 140,582-583,589-590,592,599-600
  switchport trunk allowed vlan add 1101-1112,1199-1200
  spanning-tree port type edge trunk
  spanning-tree bpdufilter enable
  speed 10000

interface port-channel12
  description NetApp2 vPC 12 link Eth1/35 NetApp2 6/b
  switchport mode trunk
  vpc 12
  switchport trunk allowed vlan 106,109-110,115,119-120,124,129-130
  switchport trunk allowed vlan add 140,582-583,589-590,592,599-600
  switchport trunk allowed vlan add 1101-1112,1199-1200
  spanning-tree port type edge trunk
  spanning-tree bpdufilter enable
  speed 10000

interface port-channel31
  description ** vPC 31 Eth2/2,Eth3/2 dc03-6120-1-A Eth 2/3,Eth2/4 **
  switchport mode trunk
  vpc 31
  switchport trunk native vlan 2
  switchport trunk allowed vlan 2,101,105-107,109-110,113,115,119-120
  switchport trunk allowed vlan add 122-124,129-130,132,140,162,581-590
  switchport trunk allowed vlan add 592,599-600,900-902,905
  spanning-tree port type edge trunk
  spanning-tree bpdufilter enable
  speed 10000

interface port-channel32
  description ** vPC 32 to dc03-6120-1-B Po32 Eth2/1,Eth2/2 **
  switchport mode trunk
  vpc 32
  switchport trunk native vlan 2
  switchport trunk allowed vlan 2,101,105-107,109-110,113,115,119-120
  switchport trunk allowed vlan add 122-124,129-130,132,140,162,581-590
  switchport trunk allowed vlan add 592,599-600,900-902,905
  spanning-tree port type edge trunk
  spanning-tree bpdufilter enable
  speed 10000

interface port-channel45
  description ** Port Channel to Aggregation Nexus 7000 dc09 , dc10 **
  switchport mode trunk
  vpc 45
  switchport trunk native vlan 2
  switchport trunk allowed vlan 2,101,105-107,109-110,113,115,119-120
  switchport trunk allowed vlan add 122-124,129-130,132,140,162,581-590
  switchport trunk allowed vlan add 592,599-600,900-902,905
  speed 10000

interface port-channel61
  description ** vPC61 Eth2/4,Eth3/4 dc06-6120-1-A Eth2/3-4 **
  switchport mode trunk
  vpc 61
  switchport trunk native vlan 2
  switchport trunk allowed vlan 2,101,105-107,109-110,113,115,119-120
  switchport trunk allowed vlan add 122-124,129-130,132,140,162,581-590
  switchport trunk allowed vlan add 592,599-600,900-902,905
  spanning-tree port type edge trunk
  spanning-tree bpdufilter enable
  speed 10000

interface port-channel62
  switchport mode trunk
  vpc 62
  switchport trunk native vlan 2
  switchport trunk allowed vlan 2,101,105-107,109-110,113,115,119-120
  switchport trunk allowed vlan add 122-124,129-130,132,140,162,581-590
  switchport trunk allowed vlan add 592,599-600,900-902,905
  spanning-tree port type edge trunk
  spanning-tree bpdufilter enable
  speed 10000

interface vfc34
  bind interface Ethernet1/34
  no shutdown

interface vfc36
  bind interface Ethernet1/36
  no shutdown
vsan database
  vsan 907 interface vfc34
  vsan 907 interface vfc36
  vsan 907 interface fc2/1
  vsan 907 interface fc2/2
  vsan 907 interface fc2/3
  vsan 907 interface fc2/4
  vsan 907 interface fc3/1
  vsan 907 interface fc3/2
  vsan 907 interface fc3/3
  vsan 907 interface fc3/4

interface fc2/1
  switchport trunk mode auto
  switchport trunk allowed vsan 907
  no shutdown

interface fc2/2
  switchport trunk mode auto
  switchport trunk allowed vsan 907
  no shutdown

interface fc2/3
  switchport trunk mode auto
  switchport trunk allowed vsan 907
  no shutdown

interface fc2/4
  switchport trunk mode auto
  switchport trunk allowed vsan 907
  no shutdown

interface fc3/1
  switchport trunk mode auto
  switchport trunk allowed vsan 907
  no shutdown

interface fc3/2
  switchport trunk mode auto
  switchport trunk allowed vsan 907
  no shutdown

interface fc3/3
  switchport trunk mode auto
  switchport trunk allowed vsan 907
  no shutdown

interface fc3/4
  switchport trunk mode auto
  switchport trunk allowed vsan 907
  no shutdown

interface Ethernet1/1
  description mgmt12-c250-3 nic4
  switchport mode trunk
  switchport trunk allowed vlan 581-582
  speed 1000

interface Ethernet1/2
  description mgmt12-c250-4 nic4
  switchport mode trunk
  switchport trunk allowed vlan 581-582
  spanning-tree port type edge trunk
  speed 1000

interface Ethernet1/3
  description smt05-n1010-1 port 1
  switchport mode trunk
  switchport trunk allowed vlan 162
  spanning-tree port type edge trunk
  spanning-tree bpdufilter enable
  speed 1000

interface Ethernet1/4
  description smt04-n1010-1 port 2
  switchport mode trunk
  switchport trunk allowed vlan 162
  spanning-tree port type edge trunk
  spanning-tree bpdufilter enable
  speed 1000

interface Ethernet1/5
  description smt05-n1010-1 port 3
  switchport mode trunk
  switchport trunk allowed vlan 900
  spanning-tree port type edge trunk
  spanning-tree bpdufilter enable
  speed 1000

interface Ethernet1/6
  description smt04-n1010-1 port 4
  switchport mode trunk
  switchport trunk allowed vlan 900
  spanning-tree port type edge trunk
  spanning-tree bpdufilter enable
  speed 1000

interface Ethernet1/7

interface Ethernet1/8

interface Ethernet1/9

interface Ethernet1/10

interface Ethernet1/11

interface Ethernet1/12

interface Ethernet1/13

interface Ethernet1/14

interface Ethernet1/15

interface Ethernet1/16

interface Ethernet1/17

interface Ethernet1/18

interface Ethernet1/19

interface Ethernet1/20

interface Ethernet1/21

interface Ethernet1/22

interface Ethernet1/23

interface Ethernet1/24

interface Ethernet1/25

interface Ethernet1/26

interface Ethernet1/27

interface Ethernet1/28

interface Ethernet1/29

interface Ethernet1/30

interface Ethernet1/31
interface Ethernet1/32

interface Ethernet1/33
  switchport mode trunk
  switchport trunk allowed vlan 106,109-110,115,119-120,124,129-130
  switchport trunk allowed vlan add 140,582-583,589-590,592,599-600
  switchport trunk allowed vlan add 1101-1112,1199-1200
  spanning-tree port type edge trunk
  channel-group 11 mode active

interface Ethernet1/34
  description ** Eth 1/34,Eth1/36 Netapp1 Controller ports 7/0-1 **
  switchport mode trunk
  switchport trunk allowed vlan 1,907
  spanning-tree port type edge trunk

interface Ethernet1/35
  switchport mode trunk
  switchport trunk allowed vlan 106,109-110,115,119-120,124,129-130
  switchport trunk allowed vlan add 140,582-583,589-590,592,599-600
  switchport trunk allowed vlan add 1101-1112,1199-1200
  spanning-tree port type edge trunk
  channel-group 12 mode active

interface Ethernet1/36
  description ** Eth 1/34,Eth1/36 Netapp1 Controller ports 7/0-1 **
  switchport mode trunk
  switchport trunk allowed vlan 1,907
  spanning-tree port type edge trunk

interface Ethernet1/37
  switchport mode trunk
  switchport trunk native vlan 2
  switchport trunk allowed vlan 2,101,105-107,109-110,113,115,119-120
  switchport trunk allowed vlan add 122-124,129-130,132,140,162,581-590
  switchport trunk allowed vlan add 592,599-600,900-902,905
  channel-group 45 mode passive

interface Ethernet1/38
  switchport mode trunk
  switchport trunk native vlan 2
  switchport trunk allowed vlan 1-103,105-579,581-905,908-3967,4048-4093
  channel-group 1 mode active

interface Ethernet1/39
  switchport mode trunk
  switchport trunk native vlan 2
  switchport trunk allowed vlan 2,101,105-107,109-110,113,115,119-120
  switchport trunk allowed vlan add 122-124,129-130,132,140,162,581-590
  switchport trunk allowed vlan add 592,599-600,900-902,905
  channel-group 45 mode passive

interface Ethernet1/40
  switchport mode trunk
  switchport trunk native vlan 2
  switchport trunk allowed vlan 1-103,105-579,581-905,908-3967,4048-4093
  channel-group 1 mode active

interface Ethernet2/1
  description ** vPC 32 Eth2/1,Eth3/1 dc03-6120-1-B Eth 2/1,Eth2/2 **
  switchport mode trunk
  switchport trunk native vlan 2
  switchport trunk allowed vlan 2,101,105-107,109-110,113,115,119-120
  switchport trunk allowed vlan add 122-124,129-130,132,140,162,581-590
  switchport trunk allowed vlan add 592,599-600,900-902,905
  channel-group 32 mode active

interface Ethernet2/2
  description ** vPC 31 Eth2/2,Eth3/2 dc03-6120-1-A Eth 2/3,Eth2/4 **
  switchport mode trunk
  switchport trunk native vlan 2
  switchport trunk allowed vlan 2,101,105-107,109-110,113,115,119-120
  switchport trunk allowed vlan add 122-124,129-130,132,140,162,581-590
  switchport trunk allowed vlan add 592,599-600,900-902,905
  channel-group 31 mode active

interface Ethernet2/3
  description ** vPC62 Eth2/3,Eth3/3 dc06-6120-1-B Eth2/1,Eth2/2 **
  switchport mode trunk
  switchport trunk native vlan 2
  switchport trunk allowed vlan 2,101,105-107,109-110,113,115,119-120
  switchport trunk allowed vlan add 122-124,129-130,132,140,162,581-590
  switchport trunk allowed vlan add 592,599-600,900-902,905
  channel-group 62 mode active

interface Ethernet2/4
  description ** vPC61 Eth2/4,Eth3/4 dc06-6120-1-A Eth2/3-4 **
  switchport mode trunk
  switchport trunk native vlan 2
  switchport trunk allowed vlan 2,101,105-107,109-110,113,115,119-120
  switchport trunk allowed vlan add 122-124,129-130,132,140,162,581-590
  switchport trunk allowed vlan add 592,599-600,900-902,905
  channel-group 61 mode active

interface Ethernet3/1
  description ** vPC 32 Eth2/1,Eth3/1 dc03-6120-1-B Eth 2/1,Eth2/2 **
  switchport mode trunk
  switchport trunk native vlan 2
  switchport trunk allowed vlan 2,101,105-107,109-110,113,115,119-120
  switchport trunk allowed vlan add 122-124,129-130,132,140,162,581-590
  switchport trunk allowed vlan add 592,599-600,900-902,905
  channel-group 32 mode active

interface Ethernet3/2
  description ** vPC 31 Eth2/2,Eth3/2 dc03-6120-1-A Eth 2/3,Eth2/4 **
  switchport mode trunk
  switchport trunk native vlan 2
  switchport trunk allowed vlan 2,101,105-107,109-110,113,115,119-120
  switchport trunk allowed vlan add 122-124,129-130,132,140,162,581-590
  switchport trunk allowed vlan add 592,599-600,900-902,905
  channel-group 31 mode active

interface Ethernet3/3
  description ** vPC62 Eth2/3,Eth3/3 dc06-6120-1-B Eth2/1,Eth2/2 **
  switchport mode trunk
  switchport trunk native vlan 2
  switchport trunk allowed vlan 2,101,105-107,109-110,113,115,119-120
  switchport trunk allowed vlan add 122-124,129-130,132,140,162,581-590
  switchport trunk allowed vlan add 592,599-600,900-902,905
  channel-group 62 mode active

interface Ethernet3/4
  description ** vPC61 Eth2/4,Eth3/4 dc06-6120-1-A Eth2/3-4 **
  switchport mode trunk
  switchport trunk native vlan 2
  switchport trunk allowed vlan 2,101,105-107,109-110,113,115,119-120
  switchport trunk allowed vlan add 122-124,129-130,132,140,162,581-590
  switchport trunk allowed vlan add 592,599-600,900-902,905
  channel-group 61 mode active

interface mgmt0
  description *** DC05-N5k-1 Mgmt0 To zone10-6504 G3/8
  ip address x.x.x.28/23
line console
boot kickstart bootflash:/n5000-uk9-kickstart.4.2.1.N1.1.bin
boot system bootflash:/n5000-uk9.4.2.1.N1.1.bin
interface fc2/1
interface fc2/2
interface fc2/3
interface fc2/4
interface fc3/1
interface fc3/2
interface fc3/3
interface fc3/4
logging logfile messages 6
zoneset distribute full vsan 907
!Full Zone Database Section for vsan 907
zone name ESX3_B vsan 907
    member pwwn 20:00:00:25:b5:03:b2:2f
!               [ESX31_B]
    member pwwn 50:0a:09:88:97:49:9a:b4
!               [NetApp2_7b]

zone name ESX1_B vsan 907
    member pwwn 20:00:00:25:b5:06:b1:2f
!               [ESX1_B]
    member pwwn 50:0a:09:88:87:49:9a:b4
!               [NetApp1_7b]

zone name ESX2_B vsan 907
    member pwwn 20:00:00:25:b5:06:b1:3f
!               [ESX2_B]
    member pwwn 50:0a:09:88:87:49:9a:b4
!               [NetApp1_7b]

zone name ESX4_B vsan 907
    member pwwn 20:00:00:25:b5:03:b2:3f
!               [ESX32_B]
    member pwwn 50:0a:09:88:97:49:9a:b4
!               [NetApp2_7b]

zone name drs3-esx8 vsan 907
    member pwwn 50:0a:09:88:97:49:9a:b4
!               [NetApp2_7b]
    member pwwn 20:00:00:25:b5:13:b2:0f

zone name drs3-esx9 vsan 907
    member pwwn 50:0a:09:88:97:49:9a:b4
!               [NetApp2_7b]
    member pwwn 20:00:00:25:b5:13:b2:0e

zone name drs3-esx10 vsan 907
    member pwwn 50:0a:09:88:97:49:9a:b4
!               [NetApp2_7b]
    member pwwn 20:00:00:25:b5:13:b2:0d

zone name drs3-esx11 vsan 907
    member pwwn 50:0a:09:88:97:49:9a:b4
!               [NetApp2_7b]
    member pwwn 20:00:00:25:b5:13:b2:0c

zone name drs3-esx12 vsan 907
```

```
    member pwwn 50:0a:09:88:97:49:9a:b4
!                [NetApp2_7b]
    member pwwn 20:00:00:25:b5:13:b2:0b

zone name drs3-esx13 vsan 907
    member pwwn 50:0a:09:88:97:49:9a:b4
!                [NetApp2_7b]
    member pwwn 20:00:00:25:b5:13:b2:0a

zone name drs3-esx14 vsan 907
    member pwwn 20:00:00:25:b5:13:b2:19
    member pwwn 50:0a:09:88:97:49:9a:b4
!                [NetApp2_7b]

zoneset name DC05-N5K-1_active vsan 907
    member ESX3_B
    member ESX1_B
    member ESX2_B
    member ESX4_B
    member drs3-esx8
    member drs3-esx9
    member drs3-esx10
    member drs3-esx11
    member drs3-esx12
    member drs3-esx13
    member drs3-esx14

zoneset name DC-5-N5K-1_active vsan 907
    member drs3-esx8
    member drs3-esx9
    member drs3-esx10
    member drs3-esx11
    member drs3-esx12
    member drs3-esx13
    member drs3-esx14

zoneset name dc05-n5k-1_active vsan 907
    member drs3-esx8

zoneset activate name DC05-N5K-1_active vsan 907
```

Cisco Nexus 1000 Configuration

```
version 4.0(4)SV1(3a)
username admin password 5 $1$hyEqTfoB$YfrRymcg.TgMBa.RMN0Ae/  role network-admin
ssh key rsa 2048
ntp server x.x.x.9
ip domain-lookup
ip host drs3-vsm1 172.26.163.143
kernel core target 0.0.0.0
kernel core limit 1
system default switchport
logging event link-status default
policy-map type qos Gold_C0S_6
  class class-default
    set cos 6
policy-map type qos Bronze_CoS_2
  class class-default
    set cos 2
policy-map type qos Silver_CoS_4
```

```
  class class-default
    set cos 4
policy-map type qos Platinum_CoS_5
  class class-default
    set cos 5
vem 3
  host vmware id 6a40d3aa-7346-11df-0003-10000000000f
vem 4
  host vmware id 6a40d3aa-7346-11df-0003-10000000000d
vem 5
  host vmware id 6a40d3aa-7346-11df-0003-10000000001f
vem 6
  host vmware id 6a40d3aa-7346-11df-0003-10000000001d
vem 7
  host vmware id 6a40d3aa-7346-11df-0003-10000000001e
vem 8
  host vmware id 6a40d3aa-7346-11df-0003-10000000000c
vem 9
  host vmware id 6a40d3aa-7346-11df-0003-10000000000e
vem 10
  host vmware id 6a40d3aa-7346-11df-0003-100000000019
vem 11
  host vmware id 6a40d3aa-7346-11df-0003-100000000009
vem 12
  host vmware id 6a40d3aa-7346-11df-0003-10000000001a
vem 13
  host vmware id 6a40d3aa-7346-11df-0003-100000000028
vem 14
  host vmware id 6a40d3aa-7346-11df-0003-10000000001b
vem 15
  host vmware id 6a40d3aa-7346-11df-0003-10000000000a
vem 16
  host vmware id 6a40d3aa-7346-11df-0003-100000000018
snmp-server user admin network-admin auth md5 0x7ccf323f71b74c6cf1cba6d255e9ded9
priv 0x7ccf323f71b74c6cf1cba6d255e9ded9 localizedkey
snmp-server enable traps license
snmp-server community public group vdc-operator
vrf context management
  ip route 0.0.0.0/0 x.x.x.1
hostname drs3-vsm1
flow exporter namsm-1
  description dc08-namsm-1 netflow collector
  destination x.x.x.63
  transport udp 3000
  source mgmt0
  version 9
flow monitor NFMonitor
  record netflow-original
  exporter namsm-1
  timeout active 1800
  cache size 4096
vlan 1
vlan 162
  name Flash-Management
vlan 581
  name ten49-VDI_Mgmt
vlan 582
  name ten49-VDI_RAWC
vlan 583
  name ten49-VDI_Desktop
vlan 589
  name ten49-VDI_NFS
vlan 900
  name n1k-ctrl-pkt
```

```
vlan 901
  name vmotion
vlan 902
  name monitoring
vdc drs3-vsm1 id 1
  limit-resource vlan minimum 16 maximum 513
  limit-resource monitor-session minimum 0 maximum 64
  limit-resource vrf minimum 16 maximum 8192
  limit-resource port-channel minimum 0 maximum 256
  limit-resource u4route-mem minimum 32 maximum 80
  limit-resource u6route-mem minimum 16 maximum 48
port-profile type vethernet Management_172_26_162
  vmware port-group
  vmware max-ports 64
  switchport mode access
  switchport access vlan 162
  service-policy type qos input Gold_CoS_6
  no shutdown
  system vlan 162
  state enabled
port-profile type vethernet Monitoring_10_202_101
  capability l3control
  vmware port-group
  switchport mode access
  switchport access vlan 902
  no shutdown
  system vlan 902
  state enabled
port-profile type vethernet N1kV_192_168_0
  vmware port-group
  switchport mode access
  switchport access vlan 900
  service-policy type qos input Platinum_CoS_5
  no shutdown
  system vlan 900
  state enabled
port-profile type vethernet Ten_49_VDI_10_49_1
  vmware port-group
  switchport mode access
  switchport access vlan 581
  service-policy type qos input Gold_CoS_6
  no shutdown
  state enabled
port-profile type vethernet Ten_49_VDI_10_49_2
  vmware port-group
  switchport mode access
  switchport access vlan 582
  service-policy type qos input Bronze_CoS_2
  no shutdown
  state enabled
port-profile type vethernet Ten_49_VDI_10_49_32
  vmware port-group
  vmware max-ports 1024
  switchport mode access
  switchport access vlan 583
  service-policy type qos input Bronze_CoS_2
  no shutdown
  state enabled
port-profile type vethernet Ten_49_VDI_10_49_32_B
  vmware port-group
  vmware max-ports 1024
  switchport mode access
  switchport access vlan 583
  service-policy type qos input Bronze_CoS_2
  no shutdown
  state enabled
port-profile type ethernet Unused_Or_Quarantine_Uplink
  description Port-group created for Nexus1000V internal usage. Do not use.
  vmware port-group
  shutdown
  state enabled
port-profile type vethernet Unused_Or_Quarantine_Veth
  description Port-group created for Nexus1000V internal usage. Do not use.
  vmware port-group
  shutdown
 state enabled
port-profile type vethernet VDI_NFS_192_168_98
  vmware port-group
  switchport mode access
  switchport access vlan 589
  service-policy type qos input Platinum_CoS_5
  no shutdown
  state enabled
port-profile type ethernet VIC-Management-uplink
  vmware port-group
  switchport mode trunk
  switchport trunk allowed vlan 162,581
  channel-group auto mode on mac-pinning
  no shutdown
  system vlan 162
  state enabled
port-profile type ethernet VIC-Monitoring-uplink
  vmware port-group
  switchport mode trunk
  switchport trunk allowed vlan 902
  system mtu 9000
  channel-group auto mode on mac-pinning
  no shutdown
  state enabled
port-profile type ethernet VIC-Services-uplink
  vmware port-group
  switchport mode trunk
  switchport trunk allowed vlan 900-901
  system mtu 9000
  channel-group auto mode on mac-pinning
  no shutdown
  system vlan 900-901
  state enabled
port-profile type ethernet VIC-Storage-uplink
  vmware port-group
  switchport mode trunk
  switchport trunk allowed vlan 589
  system mtu 9000
  channel-group auto mode on mac-pinning
  no shutdown
  state enabled
port-profile type ethernet VIC-vm-data-uplink
  vmware port-group
  switchport mode trunk
  switchport trunk allowed vlan 582-583
  channel-group auto mode on mac-pinning
  no shutdown
  state enabled
port-profile type vethernet vMotion_192_168_1
  vmware port-group
  switchport mode access
  switchport access vlan 901
  service-policy type qos input Silver_CoS_4
  no shutdown
  state enabled
…
…
interface control0
logging logfile messages 6
boot kickstart bootflash:/nexus-1000v-kickstart-mz.4.0.4.SV1.3a.bin sup-1
boot system bootflash:/nexus-1000v-mz.4.0.4.SV1.3a.bin sup-1
boot kickstart bootflash:/nexus-1000v-kickstart-mz.4.0.4.SV1.3a.bin sup-2
boot system bootflash:/nexus-1000v-mz.4.0.4.SV1.3a.bin sup-2
logging level cdp 6
monitor session 1 type erspan-source
  description ** VDI desktops **
  destination ip 10.202.101.200
  erspan-id 49
  ip ttl 64
  ip prec 0
  ip dscp 0
  mtu 1500
  header-type 2
  no shut
logging level eth_port_channel 6
svs-domain
  domain id 103
  control vlan 900
  packet vlan 900
  svs mode L2
svs connection vcenter
  protocol vmware-vim
  remote ip address x.x.x.221 port 80
  vmware dvs uuid "f0 3b 0e 50 b2 0a d9 99-1b 68 e2 bf 6f 2d f4 f0" datacenter-name 
SASU-ESMT-VDI
  connect
```