VMware View 4.5 on Cisco Unified Computing System and EMC Unified Storage Design Guide

Last Updated: May 13, 2011

Cisco Validated Design

Building Architectures to Solve Business Problems

About the Authors

Gabe Dixon, Technical Marketing Engineer, Systems Architecture and Strategy Unit, Cisco Systems

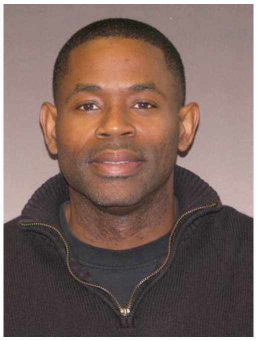

Gabe Dixon

Gabriel Dixon is a technical marketing engineer for data center technologies in Cisco's Systems Architecture and Strategy Unit (SASU). He is currently focused on partner solution validation with VMware and EMC. Dixon has been at Cisco for more than 10 years, and his roles have included systems and solutions testing positions for the Cisco Catalyst 4000 and 6000 series of switches. Prior to Cisco, he worked at Bay Networks and Sun Microsystems, delivering network management solutions as a systems test engineer and consulting engineer. Dixon holds a bachelor of science degree in management information systems from San Jose State University and a master of science degree in technology management from the University of San Francisco.

Aaron Linn, Technical Marketing Engineer, Systems Architecture and Strategy, Cisco Systems

Aaron Linn

Aaron Linn is a technical marketing engineer in Cisco's Systems Architecture and Strategy Unit (SASU), where he focuses on data center and virtualization technologies. He has over 14 years experience in systems networking, troubleshooting, and design. In more than 11 years at Cisco, Linn has also been a customer support engineer. Before joining Cisco, he was a network manager at a Silicon Valley semiconductor company for 3 years. Linn holds a bachelor of science degree in business management from San Diego State University and technical certifications from Cisco, EMC, VMware, and the Storage Networking Industry Association.

The authors woujld like to give special thanks to Ravi Venkat (Cisco), Clinton Kitson (EMC), 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 Cisco Unified Computing System and EMC Unified Storage Design Guide

VMware View 4.5 on Cisco Unified Computing System and EMC Unified Storage Design Guide

Introduction

This document reports the results of a study evaluating the scalability of VMware® View 4.5™ on Cisco® Unified Computing System™ (UCS) B-series blade servers connected to EMC® Celerra® 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, EMC Celerra storage system, and related software.

Updated Content in this Document Version

- Summary of Main Findings
- 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.
- Scale findings increased from the previous study through the use of kernel adjustments detailed in VMware vSphere Kernel Adjustments for High CPU Environments.

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

- 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 of amount memory allocated to the desktops.
- Advanced Storage Technologies simplify management, enable scalable designs, and reduce TCO:
  - De-duplication
  - Compression
  - VAAI
  - Fully automated storage-tiering
  - Enterprise Flash drives to improve performance and reduce cost of VDI deployment

Infrastructure Components

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

Cisco Unified Computing System

The Cisco Unified Computing System is a next-generation data center platform that unites compute, network, storage access, and virtualization into a cohesive system designed to reduce total cost of ownership (TCO) and increase business agility. The system integrates a low-latency, lossless 10 Gigabit Ethernet unified network fabric with enterprise-class, x86-architecture servers. The system is an integrated, scalable, multi-chassis platform in which all resources participate in a unified management domain.

The main system components include:

- Compute—The system is based on an entirely new class of computing system that incorporates blade servers based on Intel® Xeon® 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 Unified Computing System can access storage over Ethernet, Fibre Channel, Fibre Channel over Ethernet (FCoE), and iSCSI, providing customers with choice and investment protection. In addition, administrators can pre-assign storage access policies for system connectivity to storage resources, simplifying storage connectivity and management while helping increase productivity.
- Management—The system uniquely integrates all the system components, enabling the entire solution to be managed as a single entity through Cisco UCS Manager software. Cisco UCS Manager provides an intuitive graphical user interface (GUI), a command-line interface (CLI), and a robust application-programming interface (API) to manage all system configuration and operations. Cisco UCS Manager helps increase IT staff productivity, enabling storage, network, and server administrators to collaborate on defining service profiles for applications. Service profiles are logical representations of desired physical configurations and infrastructure policies. They help automate provisioning and increase business agility, allowing data center managers to provision resources in minutes instead of days.

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

Figure 1 Cisco Unified Computing System

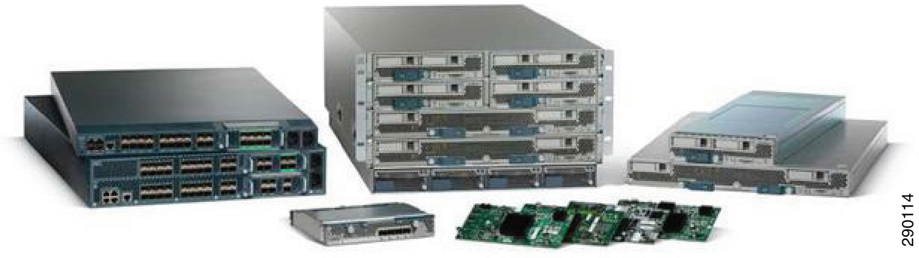

Cisco Unified Computing System Components

Fabric Interconnect

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

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

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

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

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

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

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

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

Figure 3 From Left to Right—8-Port 1/2/4-Gbps Native Fibre Channel Expansion Module; 4-Port Fibre Channel plus 4-Port 10 Gigabit Ethernet Module; 6-Port 10 Gigabit Ethernet Module; and 6-Port 1/2/4/8-Gbps Native Fibre Channel Expansion Module

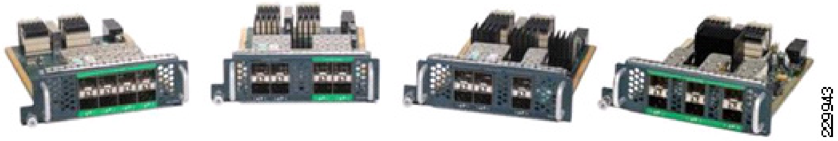

Cisco Fabric Extenders Module

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 Unified Computing System to scale to many chassis without multiplying the number of switches needed, reducing TCO and allowing all chassis to be managed as a single, highly available management domain.

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

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

Figure 4 Rear View of the Cisco UCS 5108 Blade Server Chassis with Two Cisco UCS 2104XP Fabric Extenders

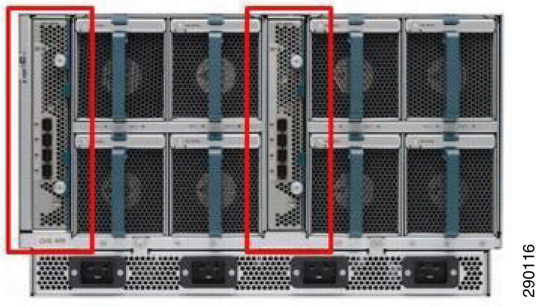

The Cisco UCS 2104XP Fabric Extender has four 10 Gigabit Ethernet, FCoE-capable, Small Form-Factor Pluggable Plus (SFP+) ports that connect the blade chassis to the fabric interconnect. Each Cisco UCS 2104XP has eight 10 Gigabit Ethernet ports connected through the midplane to each half-width slot in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

Figure 5 Cisco UCS 2104XP Fabric Extender

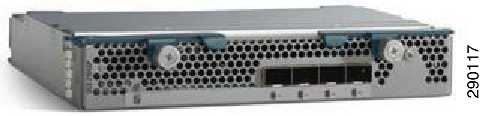

Cisco UCS Chassis

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

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

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

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

Figure 6 Cisco Blade Server Chassis (Front and Back View)

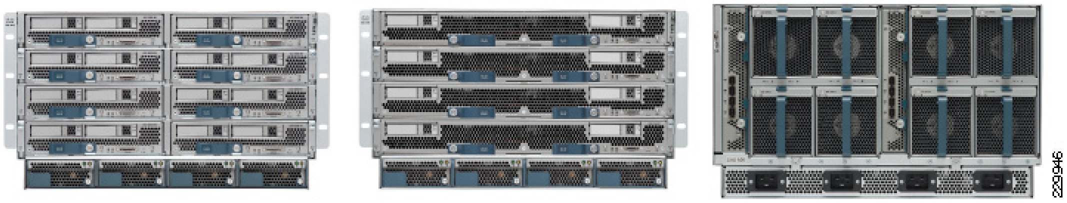

Intel Xeon 5600 Series Processor

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 7 Intel Xeon 5600 Series Processor

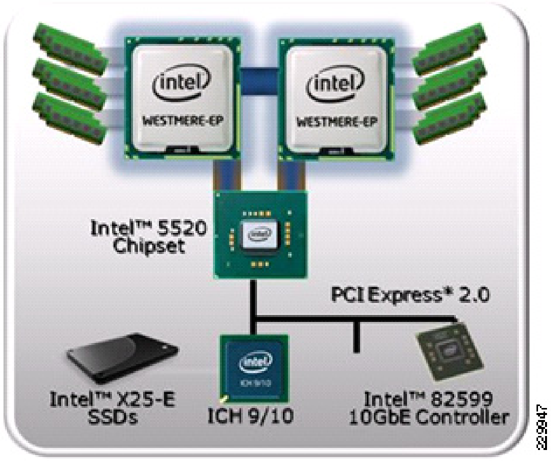

Cisco UCS B200 M2 Blade Server

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 8 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 9 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—HBA, 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
- QoS parameters—802.1p class, ETS attributes, rate limiting, and shaping

Figure 10 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 11, 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 11 Each Virtual Machine Interface has Its Own Virtual Port on the Physical Switch

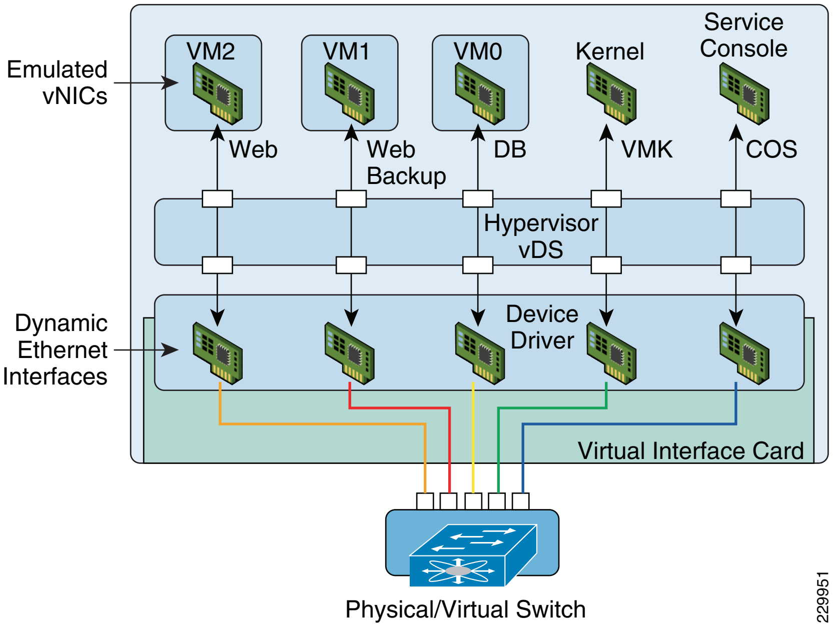

Cisco UCS Manager running on the Cisco UCS Fabric Interconnect works in conjunction with VMware vCenter software to coordinate the creation and movement of virtual machines. Port profiles are used to describe the virtual machine interface attributes such as VLAN, port security, rate limiting, and QoS 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 12). 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 12 Extended Memory Architecture

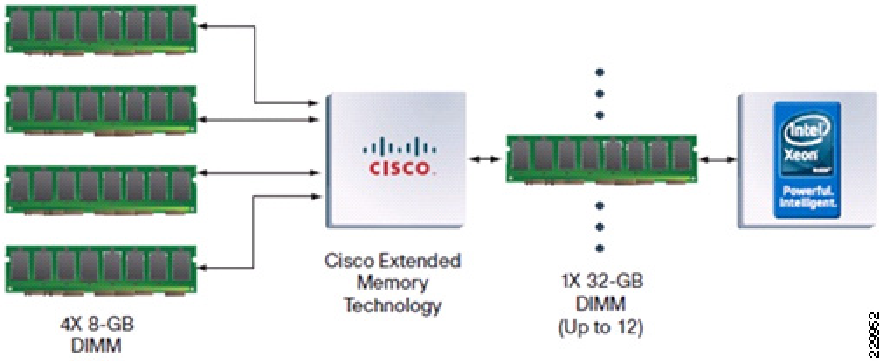

VMware vSphere4 and VCenter Server

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

You install this 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.

EMC Unified Storage Solution and Components

EMC Unified Storage

EMC unified storage brings flexibility to multi-protocol environments. With EMC unified storage you can connect to multiple storage networks using NAS, iSCSI, and Fibre Channel SAN in an integrated package. EMC unified storage leverages advanced technologies like EMC FAST and EMC FAST Cache in the latest release of FLARE (release 30) to optimize performance for the virtual desktop environment, helping support service level agreements. EMC unified storage supports vStorage APIs for Array Integration (VAAI) which was introduced in VMware vSphere 4.1. VAAI enables quicker Virtual Desktop provisioning and start-up.

The following three key EMC technologies can be enabled independently or together to minimize the cost of virtualized desktops:

- EMC FAST (Fully Automated Storage Tiering)—EMC has enhanced its FAST technology to be more automated with sub-LUN tiering. This feature works at the storage pool level, below the LUN abstraction. Where earlier versions of FAST operated above the LUN level, FAST now analyzes data patterns at a far more granular level. As an example, rather than move an 800 GB LUN to enterprise Flash drives, FAST now identifies and monitors the entire storage pool in 1 GB chunks. If data becomes active, then FAST automatically moves only these "hot" chunks to a higher tier like Flash.

As data cools, FAST also correctly identifies which chunks to migrate to lower tiers and proactively moves them. With such granular tiering, it is now possible to reduce storage acquisition while at the same time improve performance and response time. Since FAST is fully automated and policy driven, there is no manual intervention required to make this happen, so you save on operating costs as well. EMC Fast can benefit all virtual desktop deployments by leveraging different tiers of storage automatically. EMC FAST is of particular value in virtual desktop deployments using thickly provisioned desktops or VMware View Composer environments leveraging user data disks.

- EMC FAST Cache—A new feature introduced in FLARE release 30 that allows utilizing Enterprise Flash Drives (EFD) as an expanded cache layer for the array. FAST Cache is an array-wide feature that can be enabled for any LUN or storage pool. FAST Cache works by examining 64 KB chunks of data in FAST Cache enabled objects on the array. Any 64 KB chunk that has data accessed more than two times will have the chunk copied to the FAST Cache. Subsequent accesses to that data chunk will be serviced from the flash drives backing the FAST Cache. This allows promotion of very active data to flash drives which dramatically improves response times for very active data and reduces data "hot spots" that can occur within the LUN.

  FAST Cache is both an extended read and write cache absorbing read heavy activity like boot storms and antivirus scans as well as write heavy workloads like patch and application updates.

- Block Data Compression—EMC unified storage introduces Block Data Compression which allows customers to save and reclaim space anywhere in their production environment with no restrictions. This capability makes storage even more efficient by compressing data and reclaiming valuable storage capacity. Data compression works as a background task to minimize performance overhead. Block Data Compression also supports thin LUNs, and EMC unified storage integrated with VMware View 4.5 drives efficiency in virtual desktop environments with support for tiered storage. EMC FAST Cache mitigates the impact of log-in storms, AV scanning, and recompose events by absorbing both read-intensive and write-intensive I/O spikes. EMC unified storage allows VMware View 4.5 storage to be placed across the optimal disk tiers. In this tiered storage environment, replicas can be provisioned to Enterprise Flash Drives (EFDs) while linked-clones and user data can be placed respectively on Fibre Channel and SATA drives.

Figure 13 shows how EMC unified storage technologies are deployed to optimize VMware View 4.5 environment.

Figure 13 EMC Unified Storage Optimizes VMware View 4.5

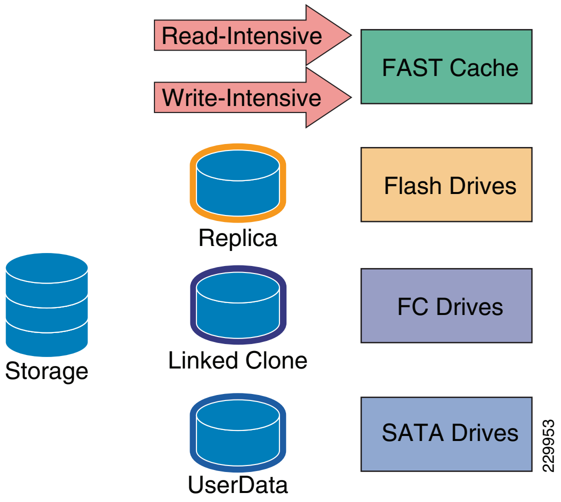

Cisco Networking Infrastructure

Cisco Nexus 5010 28-Port Switch

The Cisco Nexus 5010 Switch is a 1RU, 10 Gigabit Ethernet/FCoE access layer switch built to provide more than 500 Gigabits per second (Gbps) throughput with very low latency. It has 20 fixed 10 Gigabit Ethernet/FCoE ports that accept modules and cables meeting the Small Form-Factor Pluggable Plus (SFP+) form factor. One expansion module slot can be configured to support up to 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), quality of service (QoS), and all other data path features turned on. The low latency on the Cisco Nexus 5000 Series enables application-to-application latency on the order of 10 microseconds (depending on the network interface card [NIC]). These numbers, together with the congestion management features described next, make the Cisco Nexus 5000 Series a great choice for latency-sensitive environments.

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

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

Figure 14 High Level Architecture

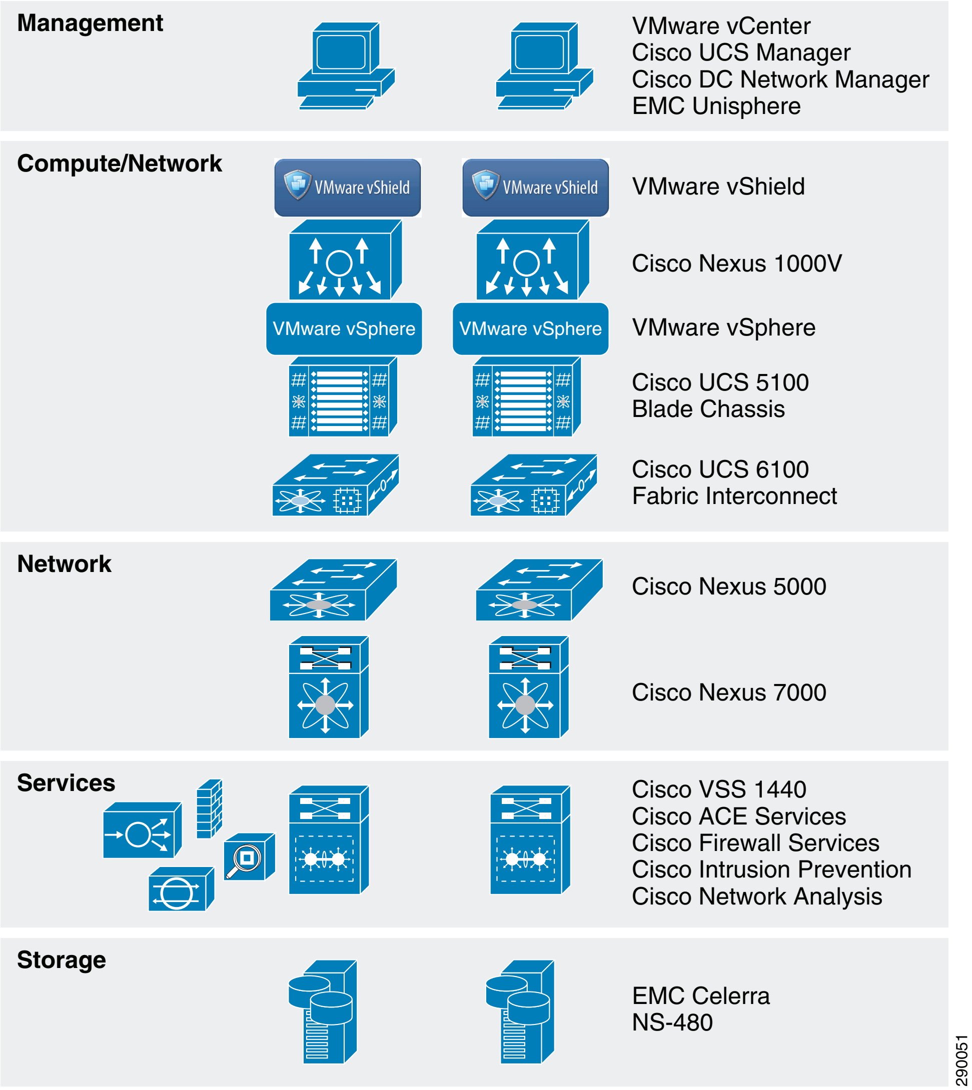

Figure 15 Detailed Architecture

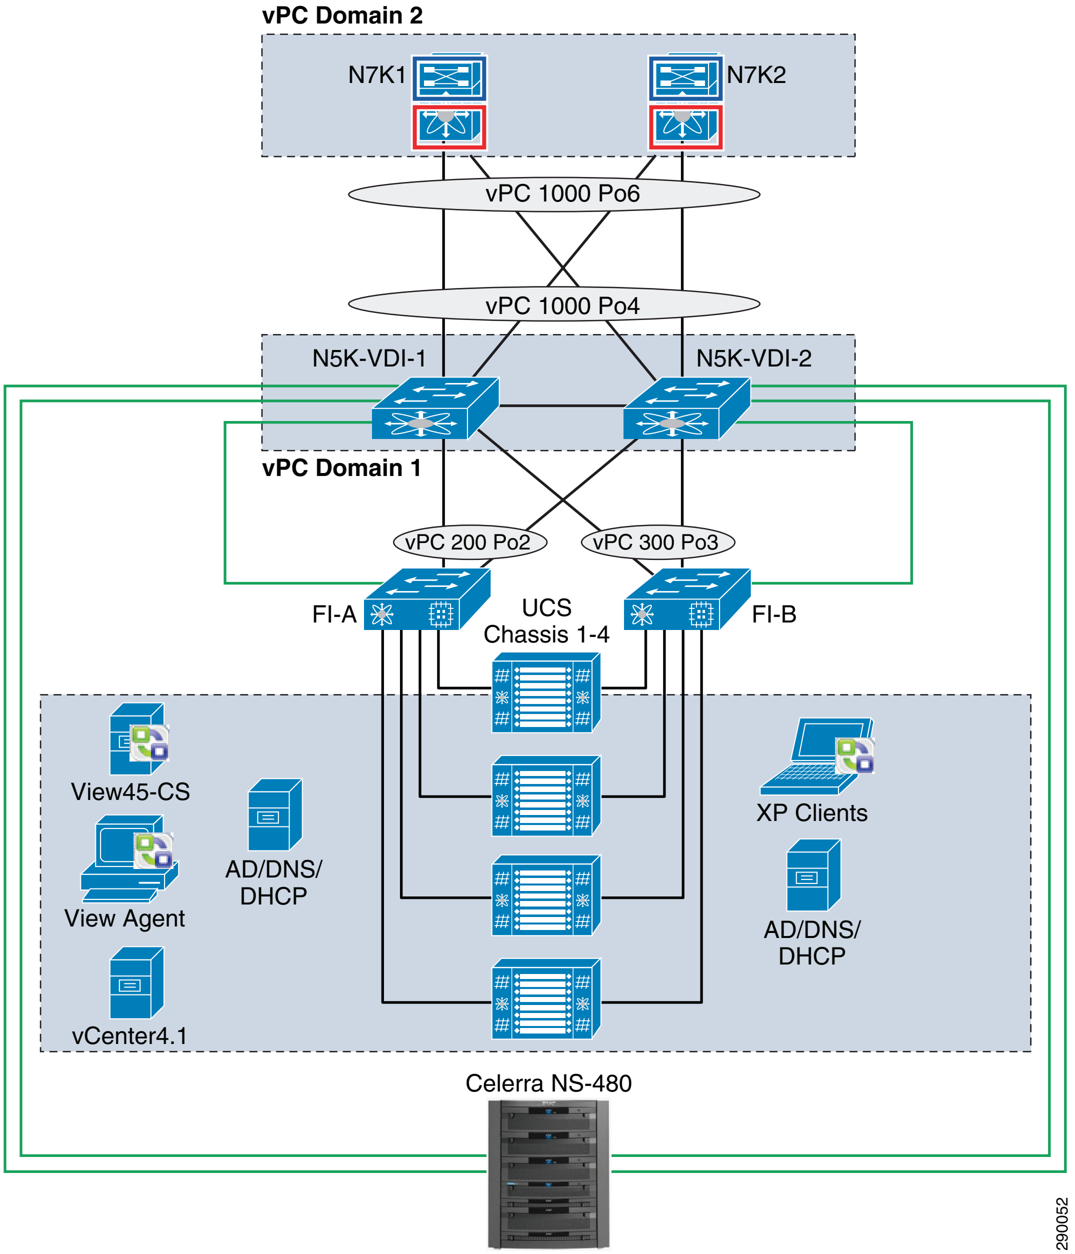

Cisco UCS Configuration

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 Service Profile Configuration Guide: http://www.cisco.com/en/US/products/ps10281/products_configuration_example09186a0080af7515.shtml. Figure 16 shows the high level flow of the configuration process.

Figure 16 UCS Manager High Level Flow

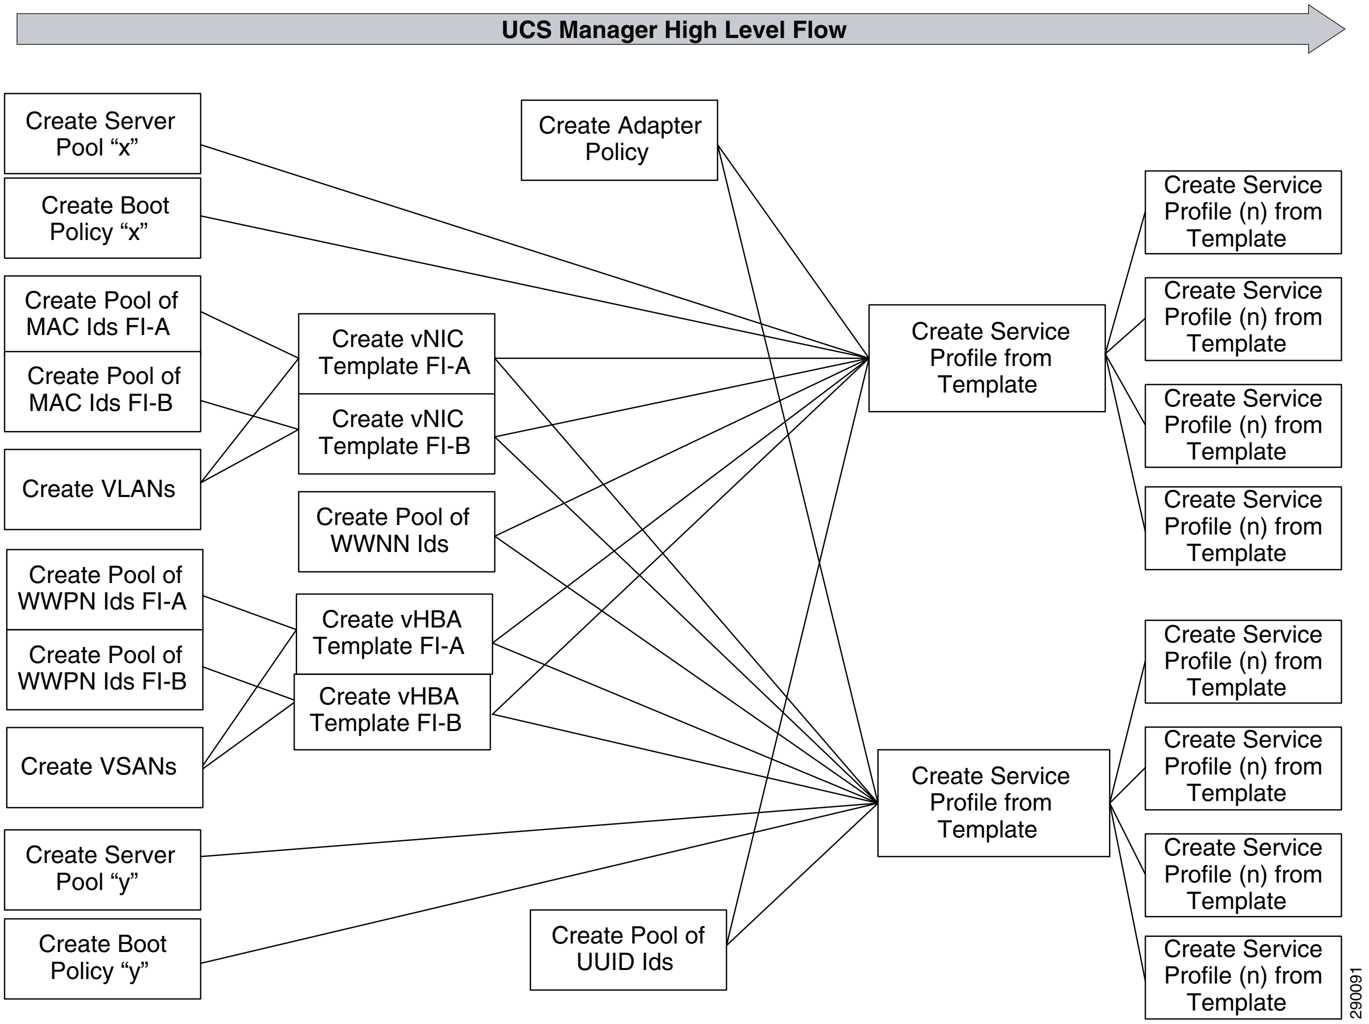
UCS Manager High Level Flow
Create Server Pool "x"
Create Boot Policy "x"
Create Pool of MAC Ids FI-A
Create Pool of MAC Ids FI-B
Create VLANs
Create Pool of WWPN Ids FI-A
Create Pool of WWPN Ids FI-B
Create VSANs
Create Server Pool "y"
Create Boot Policy "y"
Create vNIC Template FI-A
Create vNIC Template 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
Create Service Profile 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
Create Service Profile (n) from Template
290091

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.

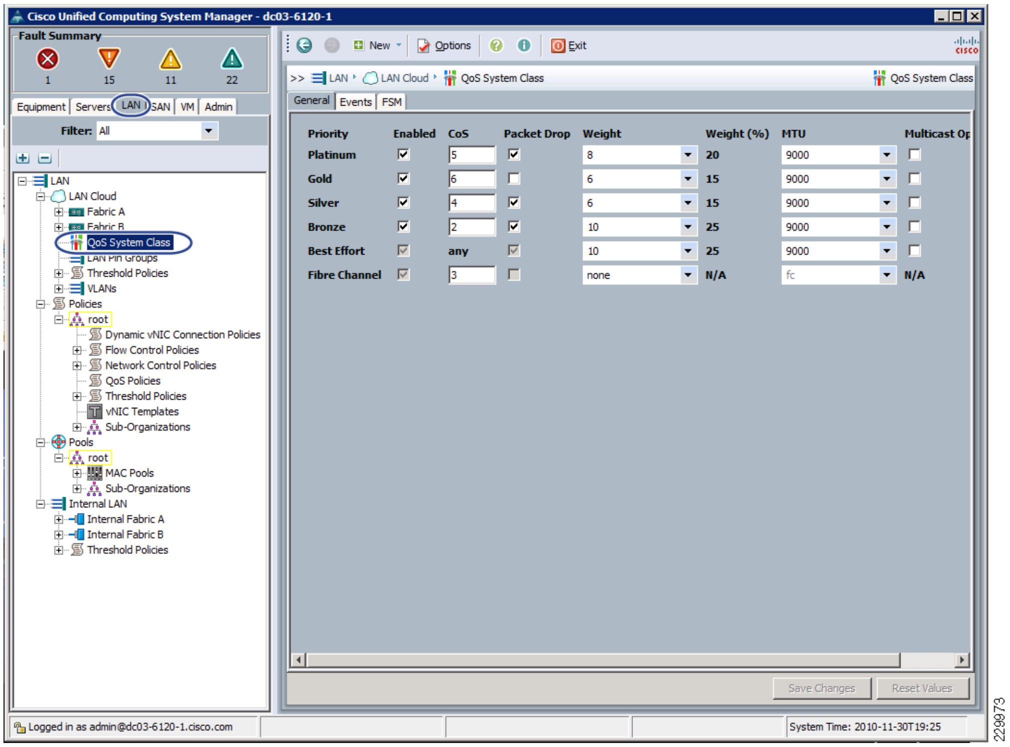

Priority	Enabled	CoS	Packet Drop	Weight	Weight (%)	MTU	Multicast Op
Platinum	☑	5	☑	8	20	9000	□
Gold	☑	6	□	6	15	9000	□
Silver	☑	4	☑	6	15	9000	□
Bronze	☑	2	☑	10	25	9000	□
Best Effort	☑	any	☑	10	25	9000	□
Fibre Channel	☑	3	□	none	N/A	fc	N/A

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

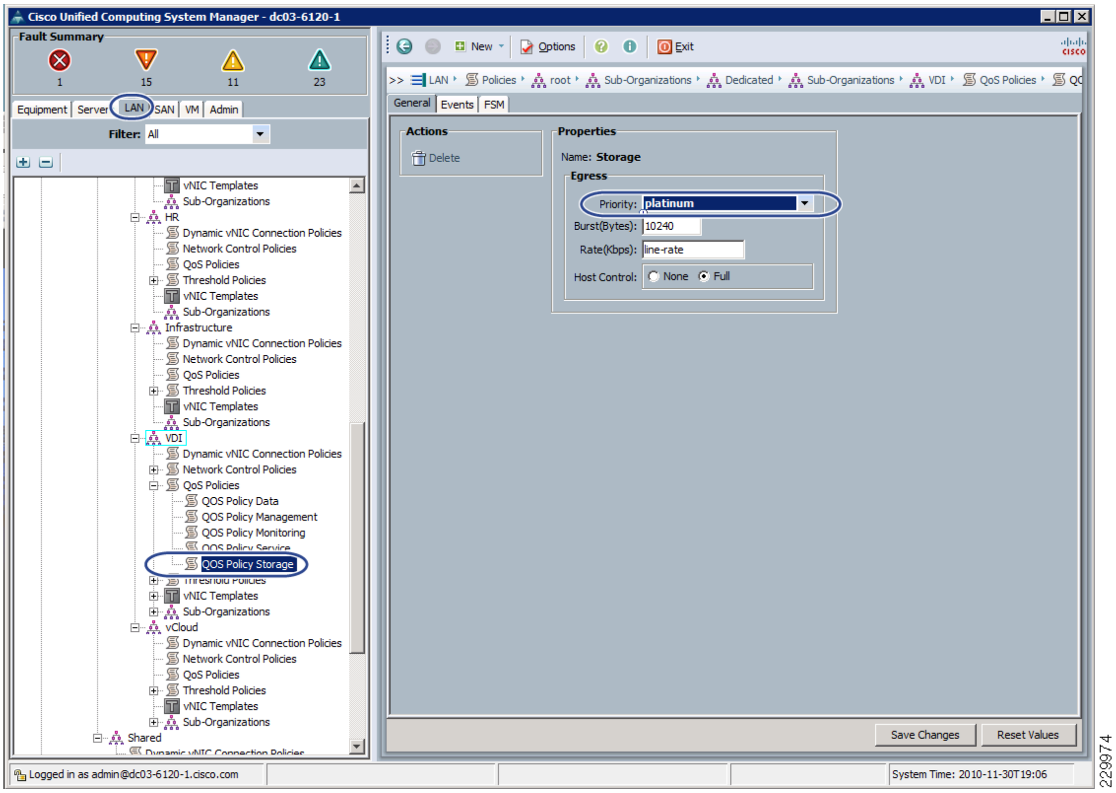

Properties

Name: Storage

Egress

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.

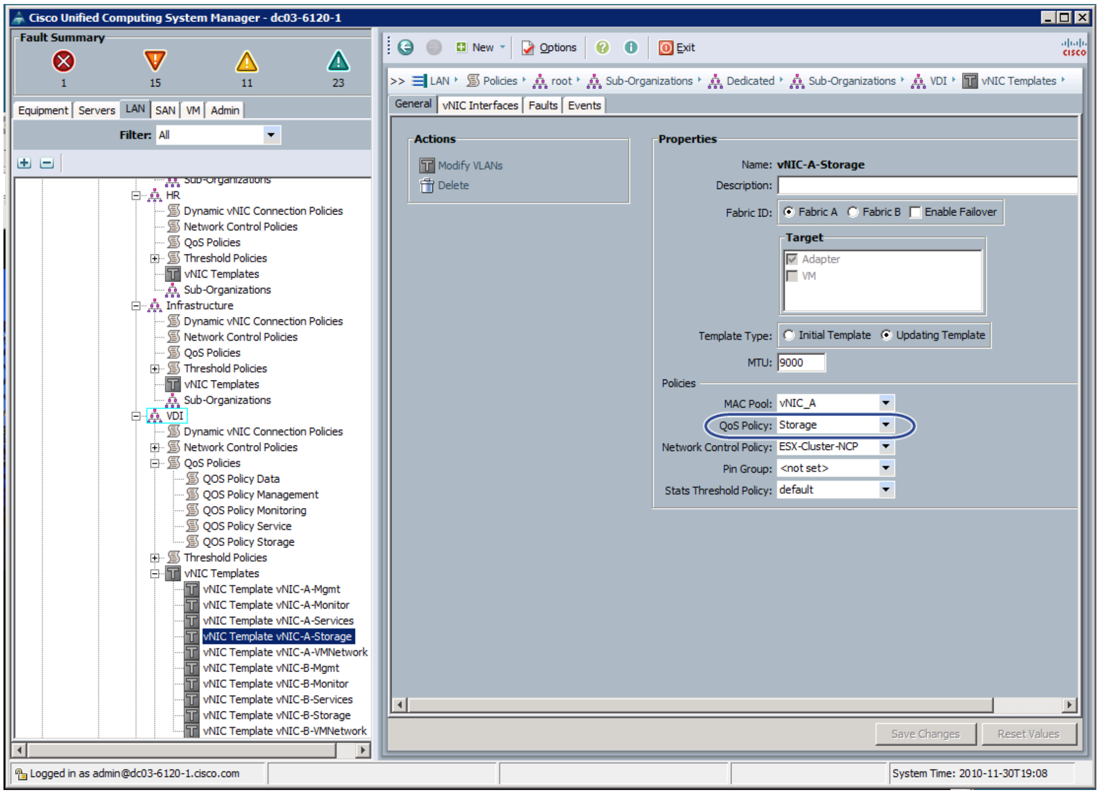
Cisco Unified Computing System Manager - dc03-6120-1
Fault Summary
1
15
11
23
Equipment
Servers
LAN
SAN
VM
Admin
Filter: All
Sub-Organizations
HR
Dynamic vNIC Connection Policies
Network Control Policies
QoS Policies
Threshold Policies
vNIC Templates
Sub-Organizations
Infrastructure
Dynamic vNIC Connection Policies
Network Control Policies
QoS Policies
Threshold Policies
vNIC Templates
Sub-Organizations
VDI
Dynamic vNIC Connection Policies
Network Control Policies
QoS Policies
QOS Policy Data
QOS Policy Management
QOS Policy Monitoring
QOS Policy Service
QOS Policy Storage
Threshold Policies
vNIC Templates
vNIC Template vNIC-A-Mgmt
vNIC Template vNIC-A-Monitor
vNIC Template vNIC-A-Services
vNIC Template vNIC-A-Storage
vNIC Template vNIC-A-VMNetwork
vNIC Template vNIC-B-Mgmt
vNIC Template vNIC-B-Monitor
vNIC Template vNIC-B-Services
vNIC Template vNIC-B-Storage
vNIC Template vNIC-B-VMNetwork
Logged in as admin@dc03-6120-1.cisco.com
New
Options
Exit
>> LAN › Policies › root › Sub-Organizations › Dedicated › Sub-Organizations › VDI › vNIC Templates ›
General
vNIC Interfaces
Faults
Events
Actions
Modify VLANs
Delete
Properties
Save Changes
Reset Values
System Time: 2010-11-30T19:08
229975

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

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 EMC 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 17.

Figure 17 Test Network Topology

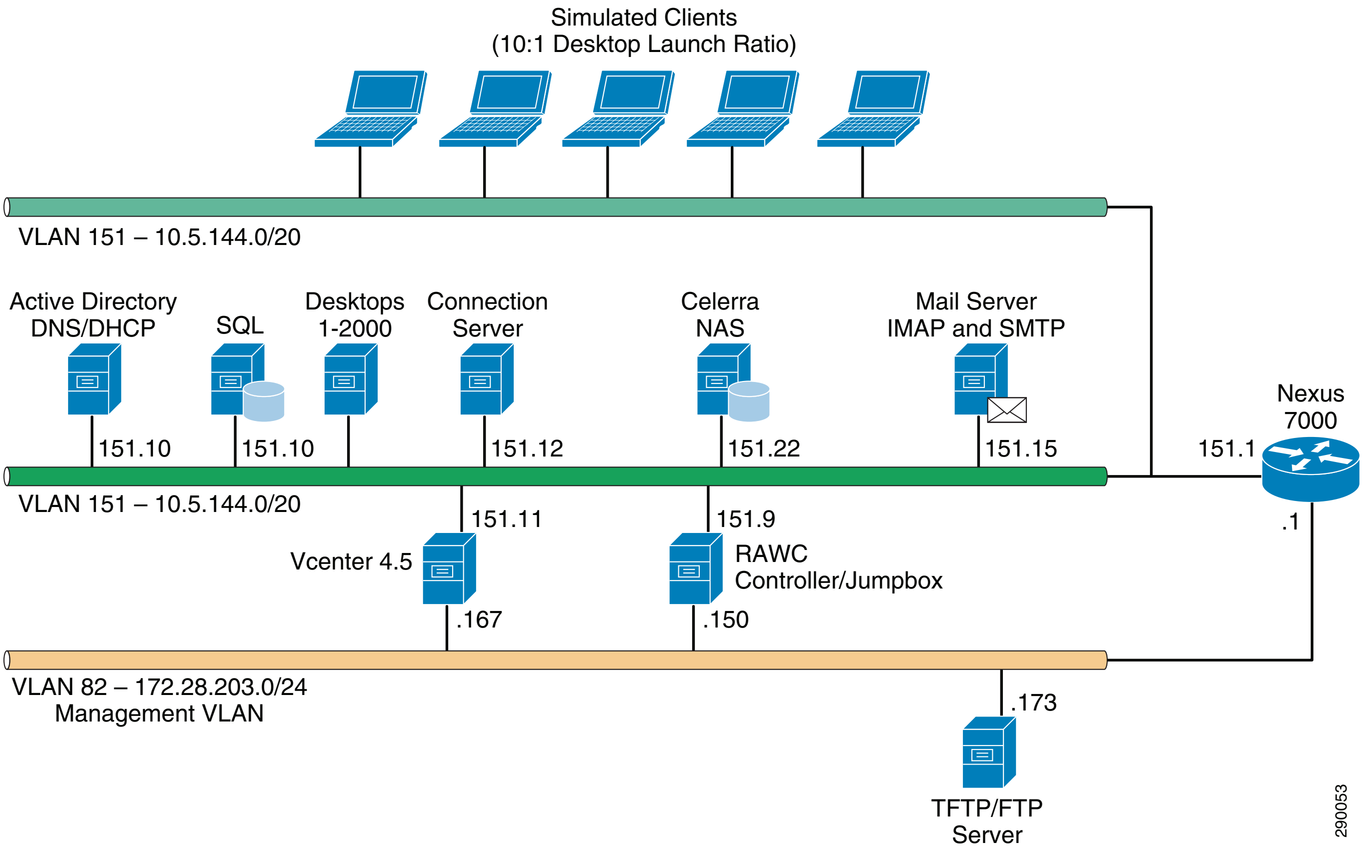

Summary of test environment:

- 2 vSphere Clusters
- 1 vCenter Server
- 1 SQL 2008 Server for vCenter and View Events
- 1 View Connection Server
- 1540 Virtual Desktops
- 1 AD Server
- 1 DHCP Server
- 1 CentOS IMAP mailserver
- 2 Celerra NAS Heads, 2 NFS Volumes
- 1 RAWC Controller
- 100 RAWC Session Launchers

The component configurations are shown in Table 4 through Table 13.

Table 4 VMware vSphere 4.1

vSphere 4.1 ESXi 260247			
Hardware:	Cisco UCS B-series Blade server	**Model:**	B250 –M2

Table 4 ***VMware vSphere 4.1***

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:**	VIC 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 x40GB 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 x40GB Virtual Disk	**Network:**	2 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:	1x40GB Virtual Disk	**Network:**	2 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

Table 8 VMware View Desktop Agent

Disk:	1 x 40GB	**Network:**	1 x 1GbE (VMXNET3)
View Agent 4.5.0-293049			

Table 9 Active Directory Server

MS Active Directory Server			
OS:	Windows 2008 Enterprise R2 64bit	**Service Pack:**	-
CPU:	2 x vCPU	**RAM:**	4096MB
Disk:	1x40GB Virtual Disk	**Network:**	2 x 10GbE (VMXNET3)

Table 10 DHCP Server

MS DHCP Server			
OS:	Windows 2008 Enterprise R2 64bit	**Service Pack:**	-
CPU:	1 x vCPU	**RAM:**	4096MB
Disk:	1x40GB Virtual Disk	**Network:**	2 x 10GbE (VMXNET3)

Table 11 Mail Server

CentOS 5.3 Dovecot IMAP Server			
OS:	CentOS 5.3_i386	**Service Pack:**	-
CPU:	1 x vCPU	**RAM:**	1024MB
Disk:	1x15GB Virtual Disk	**Network:**	1 x 10GbE (Flexible)

Table 12 RAWC Controller

VMware RAWC 1.2			
OS:	Windows 2003 R2 Enterprise 64bit	**Service Pack:**	2
CPU:	4 x vCPU	**RAM:**	2048MB
Disk:	1 x20GB Virtual Disk	**Network:**	2 x 10GbE (VMXNET3)

Table 13 *Session Launcher*

VMware RAWC 1.2 Session Launcher			
OS:	Windows XP 32bit	**Service Pack:**	1
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 5010s, 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 5010 pair in a bow tie manner as shown in Figure 18. 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 40Gbps per Fabric Interconnect as we are building a scalable and expandable system.

Figure 18 ***Network Configuration with Upstream Cisco Nexus 5000 from Cisco Unified Computing System***

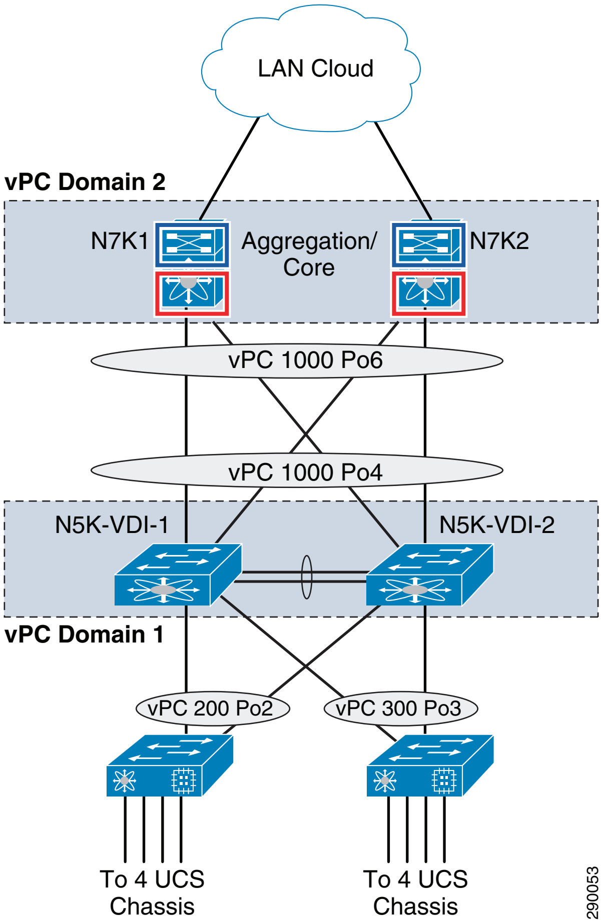

The Cisco Nexus 5000 is used to connect to the EMC Celerra NS-480 storage system for NAS access. EMC supports dual port 10G Chelsio cards which are configured in a portchannel and connected to the pair of Cisco Nexus 5000s downstream. Jumbo frames are configured on all the ports.

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

Figure 19 Network Configuration for EMC Celerra NS-480

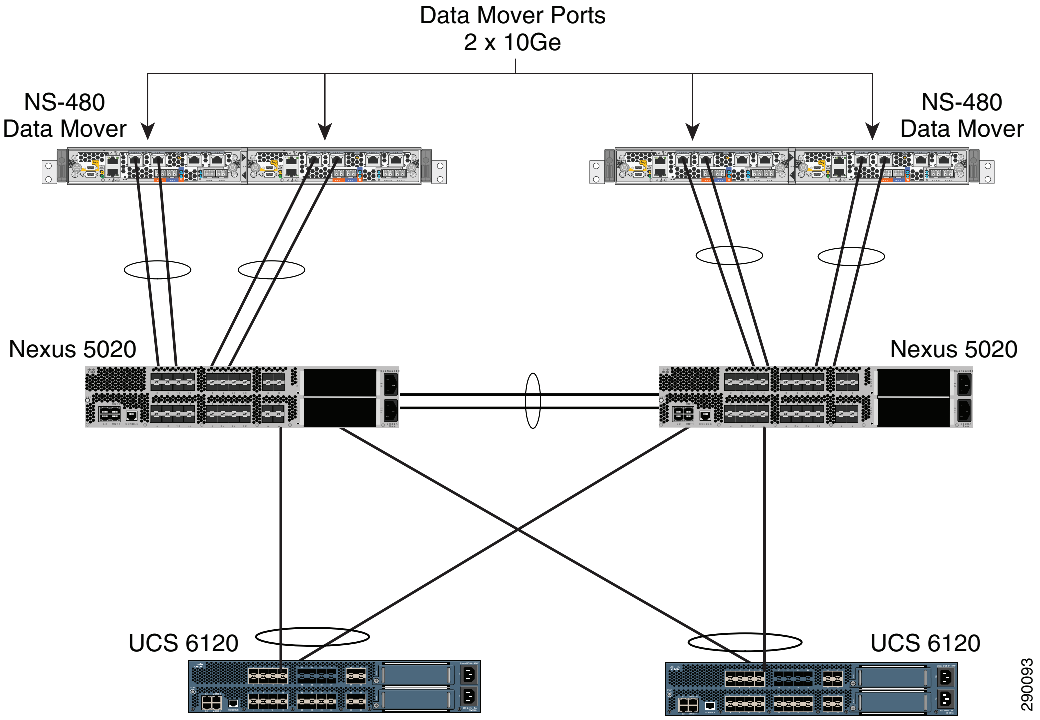

EMC NS-480 LAN Configuration

Physical design considerations:

- EMC recommends that switches support gigabit Ethernet (GbE) connections and Link Aggregation Control Protocol (LACP), and the ports on switches support copper-based media.

Logical design considerations:

- This validated solution uses virtual local area networks (VLANs) to segregate network traffic of various types to improve throughput, manageability, application separation, high availability, and security.
- The IP scheme for the virtual desktop network must be designed such that there are enough IP addresses in one or more subnets for the DHCP Server to assign them to each virtual desktop.

Link aggregation:

- Celerra unified storage provides network high availability or redundancy by using link aggregation. This is one of the methods used to address the problem of link or switch failure.
- A link aggregation is a high-availability feature that enables multiple active Ethernet connections to appear as a single link with a single MAC address and potentially multiple IP addresses.
- In this solution, LACP is configured on Celerra, which combines two 10 GbE ports into a single virtual device. If a link is lost in the Ethernet port, the link fails over to another port. All traffic is distributed across the active links.

Celerra Network Configuration Data Mover Ports

The Celerra NS-4800 consists of two blades. These blades can be configured in an active/active or active/passive configuration. In the active/passive configuration, the passive blade serves as a failover device for the active blade. In this solution, the blades operate in the active/passive mode.

The Celerra NS-480 blade consists of four 10 GbE controller ports. Ports fxg0 and fxg1 are configured by using LACP to support virtual machine traffic, home folder access, and external access for roaming profiles. Ports fxg2 and fxg3 are left free for further expansion.

The external_interface device is used for administrative purposes to move data in and out of the private network on VLAN 274. Both interfaces exist on the LACP1 device configured on fxg0 and fxg1.

The ports are configured as follows:

```
external_interface protocol=IP device=lacp1
inet=10.6.121.55 netmask=255.255.255.0
broadcast=10.6.121.255
UP, Ethernet, mtu=1500, vlan=521,
macaddr=0:60:16:26:19:0
lacp1_int protocol=IP device=lacp1
inet=192.168.80.5 netmask=255.255.240.0
broadcast=192.168.95.255
UP, Ethernet, mtu=9000, vlan=274,
macaddr=0:60:16:26:19:0
```

LACP Configuration on the Data Mover

To configure the link aggregation that uses fxg0 and fxg1 on server_2, type the following at the command prompt:

```
$ server_sysconfig server_2 -virtual -name <Device Name> -
create trk
-option "device=fxg0,fxg1 protocol=lacp"
```

To verify if the ports are channeled correctly, type:

```
$ server_sysconfig server_2 -virtual -info lacp1
server_2:
*** Trunk lacp1: Link is Up ***
*** Trunk lacp1: Timeout is Short ***
*** Trunk lacp1: Statistical Load C is IP ***
Device Local Grp Remote Grp Link LACP Duplex Speed
--------------------------------------------------------------
fxg0 100000 5888 Up Up Full 10000 Mbs
fxg1 100000 5888 Up Up Full 10000 Mbs
```

The remote group number must match for both cge ports and the LACP status must be “Up.” Verify if the appropriate speed and duplex are established as expected.

SAN Configuration

Storage Components

Storage pools Storage pools in FLARE release 30 support heterogeneous drive pools. In this solution, four 20-disk pools were configured from 14 FC disks and 6 SATA drives. Four thick LUNs, each 1.25 TB in size, were created from this storage pool, as shown in Figure 20. FAST Cache is enabled for the pool.

Figure 20 LUNs for Storage Pool

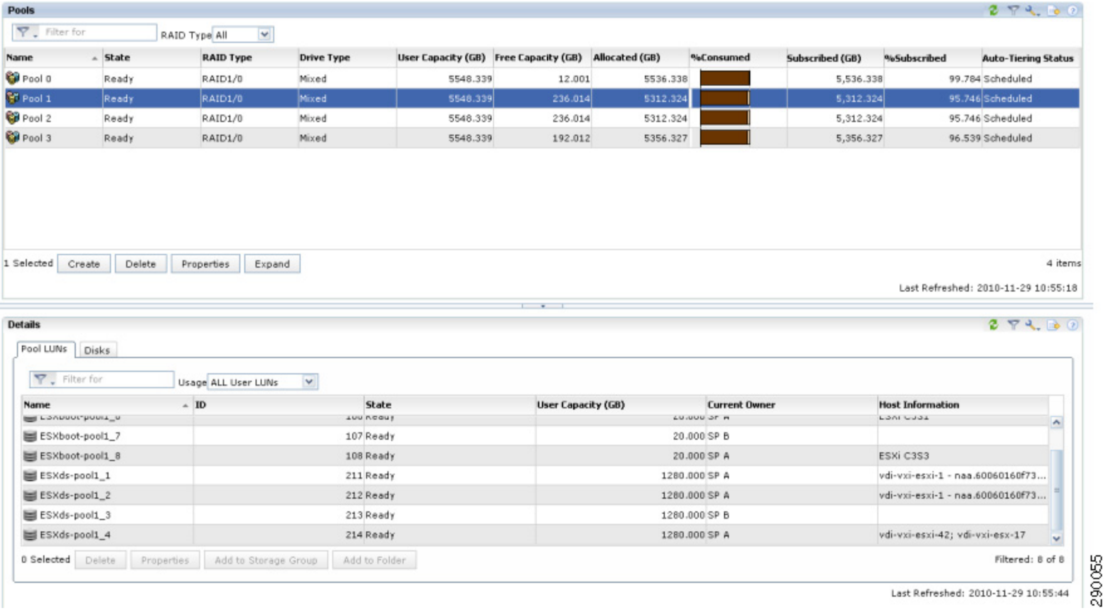

Pools

Filter for RAID Type All

Name	State	RAID Type	Drive Type	User Capacity (GB)	Free Capacity (GB)	Allocated (GB)	%Consumed	Subscribed (GB)	%Subscribed	Auto-Tiering Status
Pool 0	Ready	RAID1/0	Mixed	5548.339	12.001	5536.338		5,536.338	99.784	Scheduled
Pool 1	Ready	RAID1/0	Mixed	5548.339	236.014	5312.324		5,312.324	95.746	Scheduled
Pool 2	Ready	RAID1/0	Mixed	5548.339	236.014	5312.324		5,312.324	95.746	Scheduled
Pool 3	Ready	RAID1/0	Mixed	5548.339	192.012	5356.327		5,356.327	96.539	Scheduled

1 Selected 4 items

Last Refreshed: 2010-11-29 10:55:18

Details

Filter for Usage ALL User LUNs

Name	ID	State	User Capacity (GB)	Current Owner	Host Information
ESXboot-pool1_6	106	Ready	20.000	SP A	ESXi C3S1
ESXboot-pool1_7	107	Ready	20.000	SP B	
ESXboot-pool1_8	108	Ready	20.000	SP A	ESXi C3S3
ESXds-pool1_1	211	Ready	1280.000	SP A	vdi-vxi-esxi-1 - naa.60060160f73...
ESXds-pool1_2	212	Ready	1280.000	SP A	vdi-vxi-esxi-1 - naa.60060160f73...
ESXds-pool1_3	213	Ready	1280.000	SP B	
ESXds-pool1_4	214	Ready	1280.000	SP A	vdi-vxi-esxi-42; vdi-vxi-esx-17

0 Selected Filtered: 8 of 8

Last Refreshed: 2010-11-29 10:55:44

For each LUN in the storage pool, the tiering policy is set to Highest Available Tier to ensure that all frequently accessed desktop data remains on the FC disks. As data ages and is used infrequently, it is moved to the SATA drives in the pool.

Enable FAST Cache

FAST Cache is enabled as an array-wide feature in the system properties of the array in Unisphere™. Click the **FAST Cache** tab, click the **Create** button, and select the eligible four EFDs to create the FAST Cache. There are no user-configurable parameters for the FAST Cache.

Figure 21 FAST Cache

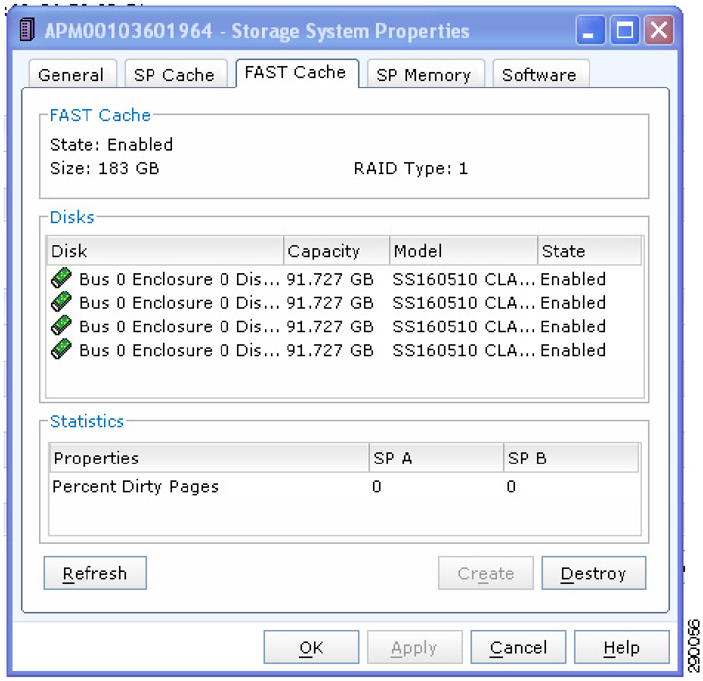

FAST Cache

State: Enabled

Size: 183 GB RAID Type: 1

Disks

Disk	Capacity	Model	State
Bus 0 Enclosure 0 Dis...	91.727 GB	SS160510 CLA...	Enabled
Bus 0 Enclosure 0 Dis...	91.727 GB	SS160510 CLA...	Enabled
Bus 0 Enclosure 0 Dis...	91.727 GB	SS160510 CLA...	Enabled
Bus 0 Enclosure 0 Dis...	91.727 GB	SS160510 CLA...	Enabled

Statistics

Properties	SP A	SP B
Percent Dirty Pages	0	0

FAST Cache is not enabled for the replica storage in this solution. The replica images are serviced from the EFDs. Enabling FAST Cache for these LUNs causes additional overhead without added performance.

If the replica images are stored on FC disks, enable FAST Cache for those LUNs.

To enable FAST Cache for any LUNs in a pool, go to the properties of the pool and click the **Advanced** tab. Select Enabled to enable FAST Cache, as shown in Figure 22.

Figure 22 Enabling FAST Cache

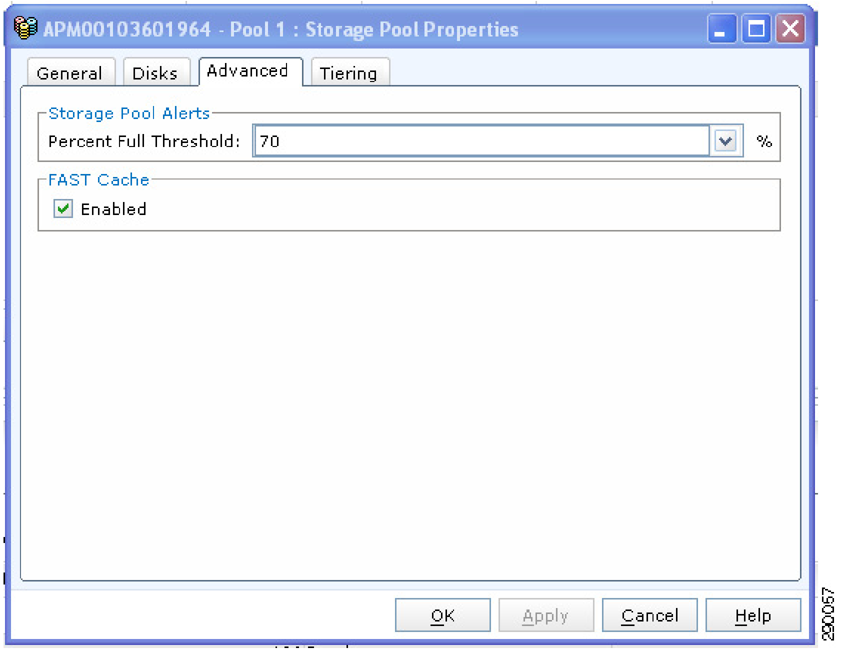

Storage Pool Alerts

Percent Full Threshold: 70 %

FAST Cache

☑ Enabled

Configure FAST

To configure the FAST feature for a pool LUN, go to the properties for a pool LUN and click the **Tiering** tab. From this location, set the tiering policy for the LUN.

Figure 23 Tiering Policy for LUN

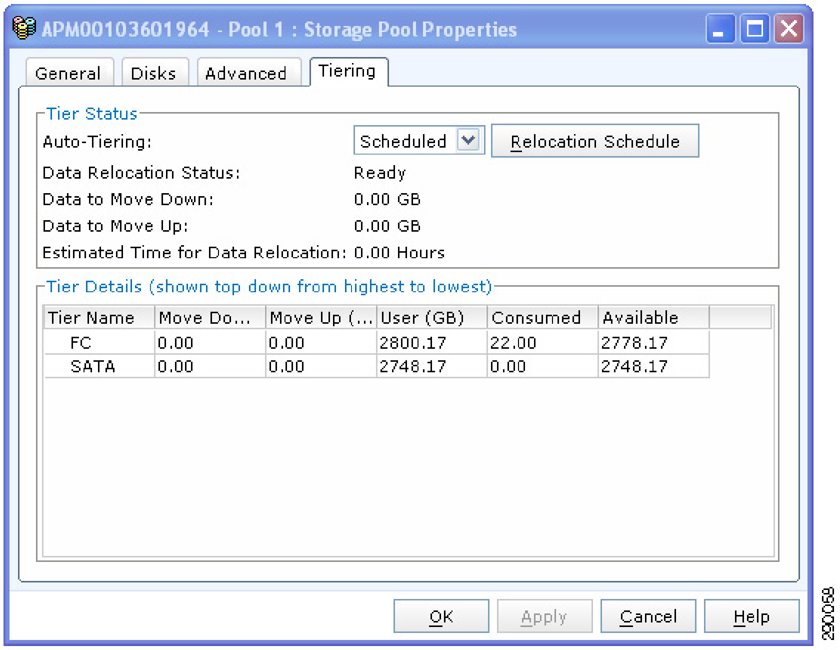

Tier Status

Auto-Tiering: Scheduled

Data Relocation Status: Ready

Data to Move Down: 0.00 GB

Data to Move Up: 0.00 GB

Estimated Time for Data Relocation: 0.00 Hours

Tier Details (shown top down from highest to lowest)

Tier Name	Move Do...	Move Up (...	User (GB)	Consumed	Available
FC	0.00	0.00	2800.17	22.00	2778.17
SATA	0.00	0.00	2748.17	0.00	2748.17

Celerra Home Directory feature

The Celerra Home Directory installer is available on the NAS Tools and Apps CD for each DART release and can be downloaded from Powerlink at the following location:

Home > Support > Software Downloads and Licensing > Downloads C > Celerra Software.

Instructions to install the Celerra Home Directory feature are located on the EMC Celerra Network Server Documentation CD available on Powerlink.

After the feature is installed, use the Celerra Management Microsoft Management Console (MMC) snap-in to configure the Home Directory feature. A sample configuration is shown in Figure 24 and Figure 25.

Figure 24 Sample Configuration —1

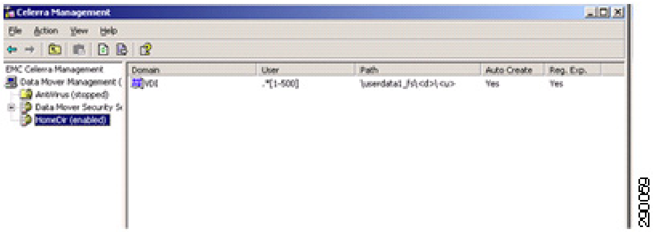

Domain	User	Path	Auto Create	Reg. Exp.
VDI	.*[1-500]	\userdata1_fs\<d>\<u>	Yes	Yes

For any user account that ends with a suffix between 1 and 500, the sample configuration shown in Figure 25 automatically creates a user home directory on the \userdata1_fs file system in the format \userdata1_fs\<domain>\<user> and maps the H:\ drive to this path. Each user has exclusive rights to the folder.

Figure 25 Sample Configuration —2

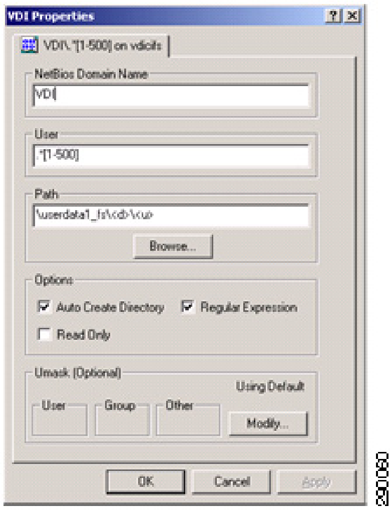

NetBios Domain Name: VDI

User: .*[1-500]

Path: \userdata1_fs\<d>\<u>

Options

☑ Auto Create Directory ☑ Regular Expression

☐ Read Only

Umask (Optional)

Using Default

User Group Other

The *EMC Celerra Home Directory-A Detailed Review* white paper available on Powerlink provides more details and advanced configuration examples.

CIFS Configuration

Example of a EMC CIFS/NFS Volume Configuration

This procedure explains how to create a CIFS share via EMC Unishsphere on your EMC Celerra System and map a drive to the share from your VDI desktop.

Step 1 Log into the Control Station IP address:

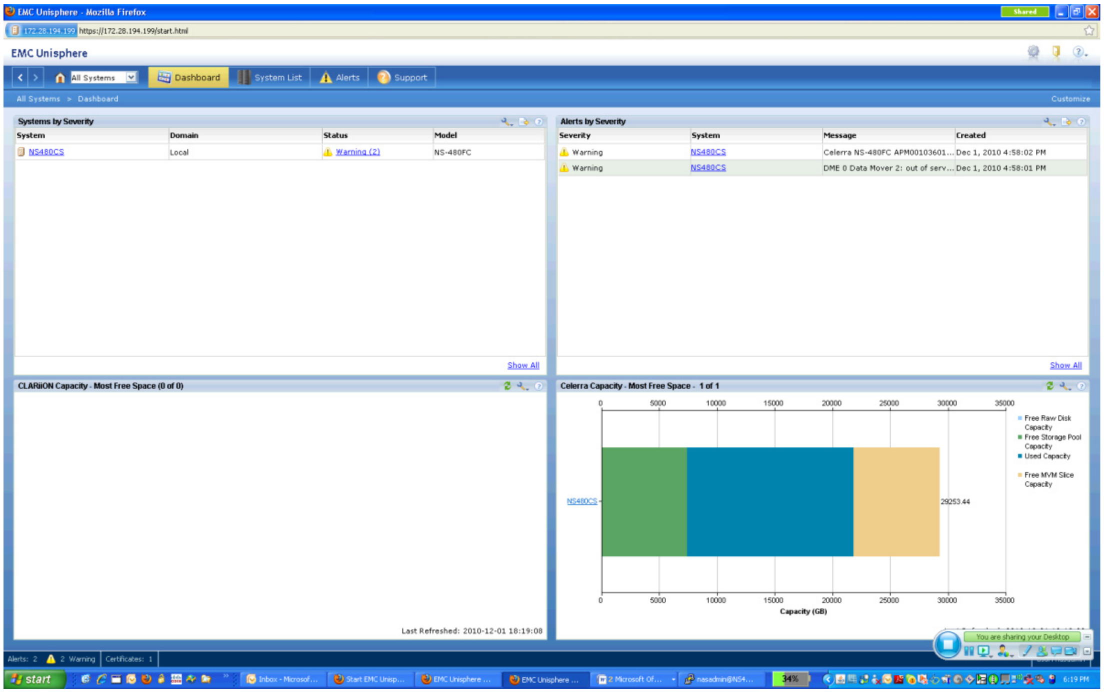

Systems by Severity

System	Domain	Status	Model
NS480CS	Local	Warning (2)	NS-480FC

Alerts by Severity

Severity	System	Message	Created
Warning	NS480CS	Celerra NS-480FC APM00103601...	Dec 1, 2010 4:58:02 PM
Warning	NS480CS	DME 0 Data Mover 2: out of serv...	Dec 1, 2010 4:58:01 PM

290094

Step 2 Click on the sharing as shown below, which also shows any pre existing shares. You can also see that we have a CIFs called fs1 already configured.

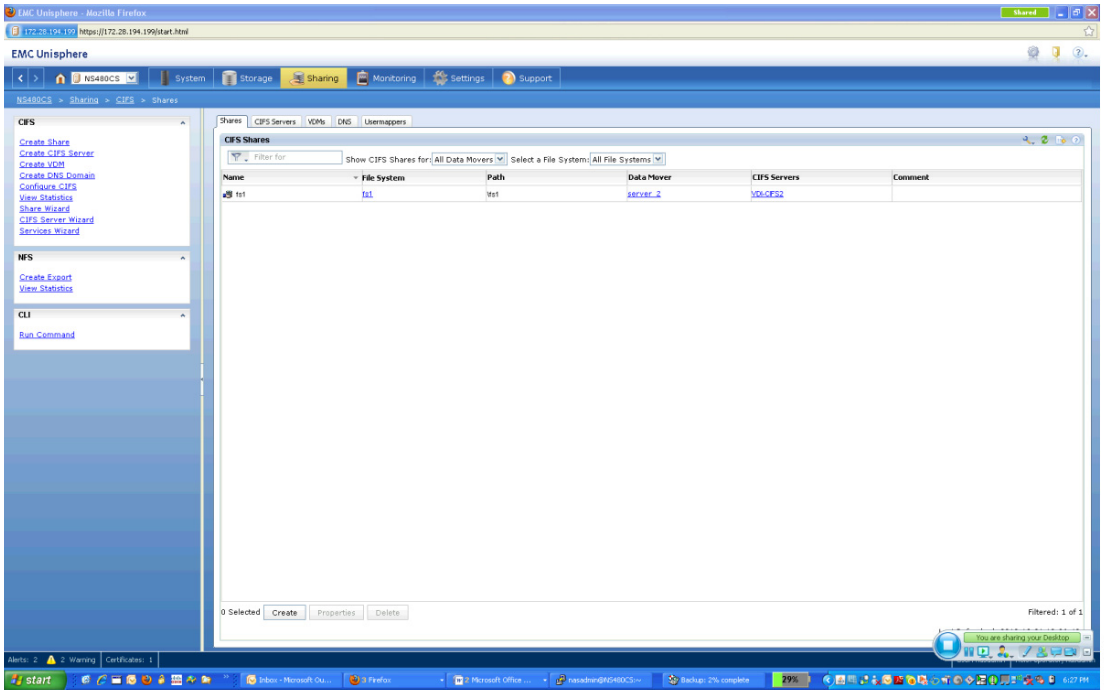

Show CIFS Shares for: All Data Movers · Select a File System: All File Systems

Name	File System	Path	Data Mover	CIFS Servers	Comment
fs1	fs1	\fs1	server_2	VDI-CIFS2	

0 Selected · Filtered: 1 of 1

290095

Step 3 From the Sharing tab use the pull down menu. Click the **CIFS Share Wizard**, which brings up the Share Wizard dialog boxes that will guide you through your NFS—or in this case CIFS—setups.

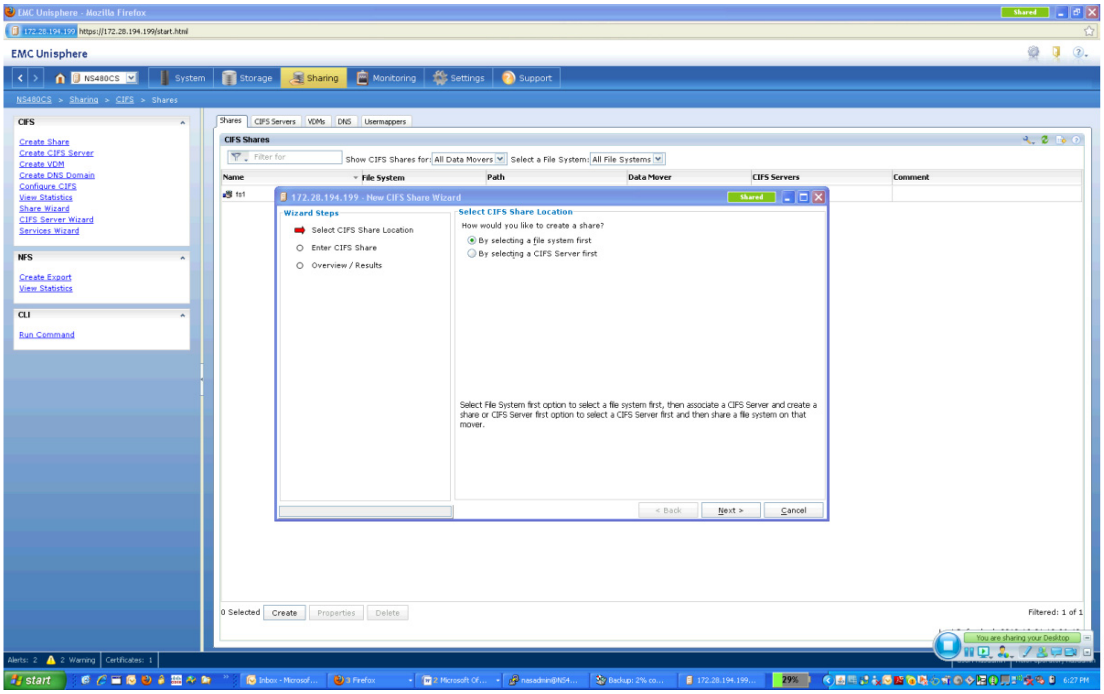

Select CIFS Share Location

How would you like to create a share?

● By selecting a file system first

○ By selecting a CIFS Server first

Select File System first option to select a file system first, then associate a CIFS Server and create a share or CIFS Server first option to select a CIFS Server first and then share a file system on that mover.

290096

Step 4 Here we will create the CIFs by selecting a file system, which is the default setting. Click **Next**.

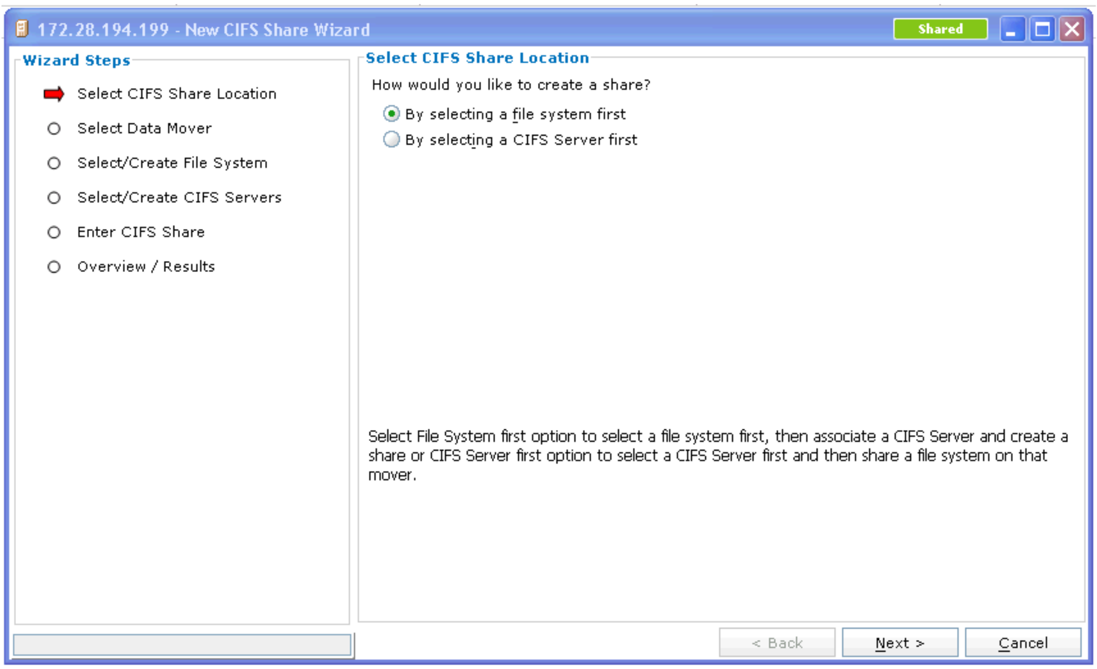

Wizard Steps

- Select CIFS Share Location
- ○ Select Data Mover
- ○ Select/Create File System
- ○ Select/Create CIFS Servers
- ○ Enter CIFS Share
- ○ Overview / Results

Select CIFS Share Location

How would you like to create a share?

● By selecting a file system first
○ By selecting a CIFS Server first

Select File System first option to select a file system first, then associate a CIFS Server and create a share or CIFS Server first option to select a CIFS Server first and then share a file system on that mover.

Step 5 Select the desired data mover on which you want to create/configure your CIFS share on. Here will select the data mover named server_2. Click **Next**.

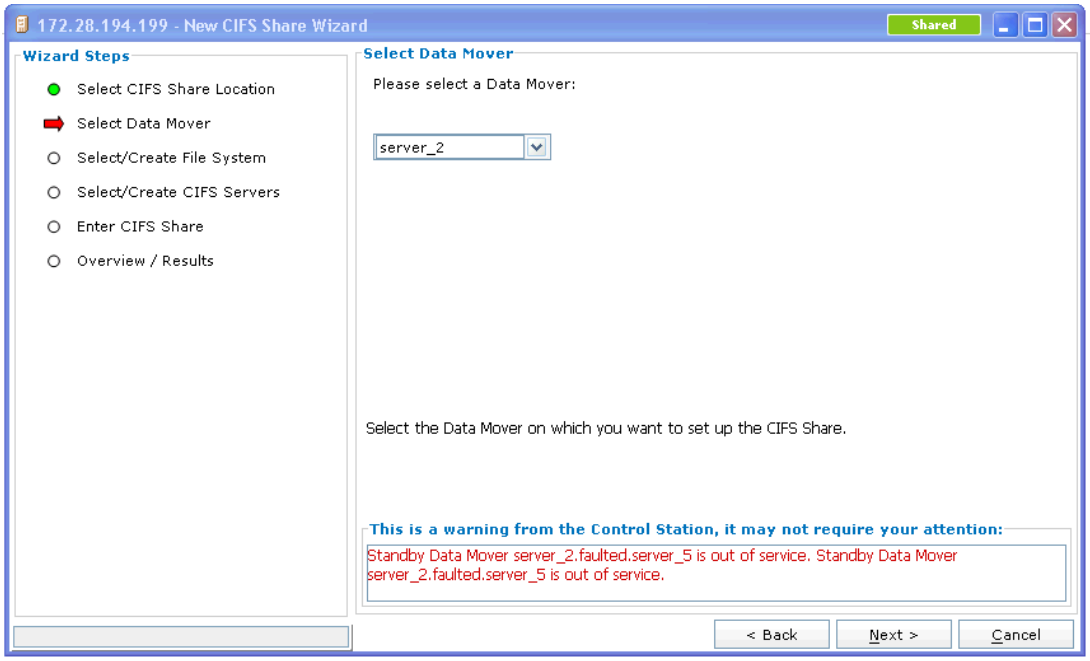

Wizard Steps

- ● Select CIFS Share Location
- Select Data Mover
- ○ Select/Create File System
- ○ Select/Create CIFS Servers
- ○ Enter CIFS Share
- ○ Overview / Results

Select Data Mover

Please select a Data Mover: server_2

Select the Data Mover on which you want to set up the CIFS Share.

This is a warning from the Control Station, it may not require your attention:

Standby Data Mover server_2.faulted.server_5 is out of service. Standby Data Mover server_2.faulted.server_5 is out of service.

Step 6 In this step you select the data mover on which the CIFS will be created. You can use the pull down to see all your available data movers. After selecting the desired data mover, click **Next**.

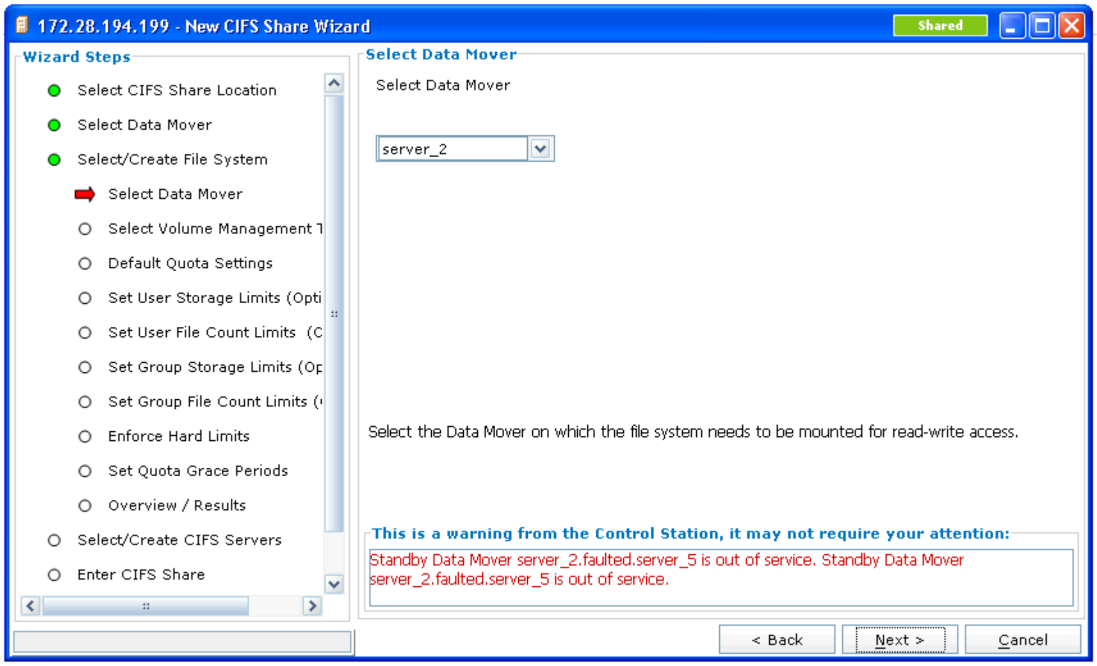
172.28.194.199 - New CIFS Share Wizard
Shared
Wizard Steps
● Select CIFS Share Location
● Select Data Mover
● Select/Create File System
➡ Select Data Mover
○ Select Volume Management T
○ Default Quota Settings
○ Set User Storage Limits (Opti
○ Set User File Count Limits (C
○ Set Group Storage Limits (Op
○ Set Group File Count Limits (
○ Enforce Hard Limits
○ Set Quota Grace Periods
○ Overview / Results
○ Select/Create CIFS Servers
○ Enter CIFS Share
< Back
Next >
Cancel
290099

Select Data Mover

Select Data Mover: server_2

Select the Data Mover on which the file system needs to be mounted for read-write access.

This is a warning from the Control Station, it may not require your attention:

Standby Data Mover server_2.faulted.server_5 is out of service. Standby Data Mover server_2.faulted.server_5 is out of service.

Step 7 Here you select whether you are creating the the share using a volume or a pre-fined storage pool. We select storage pool and click **Next**.

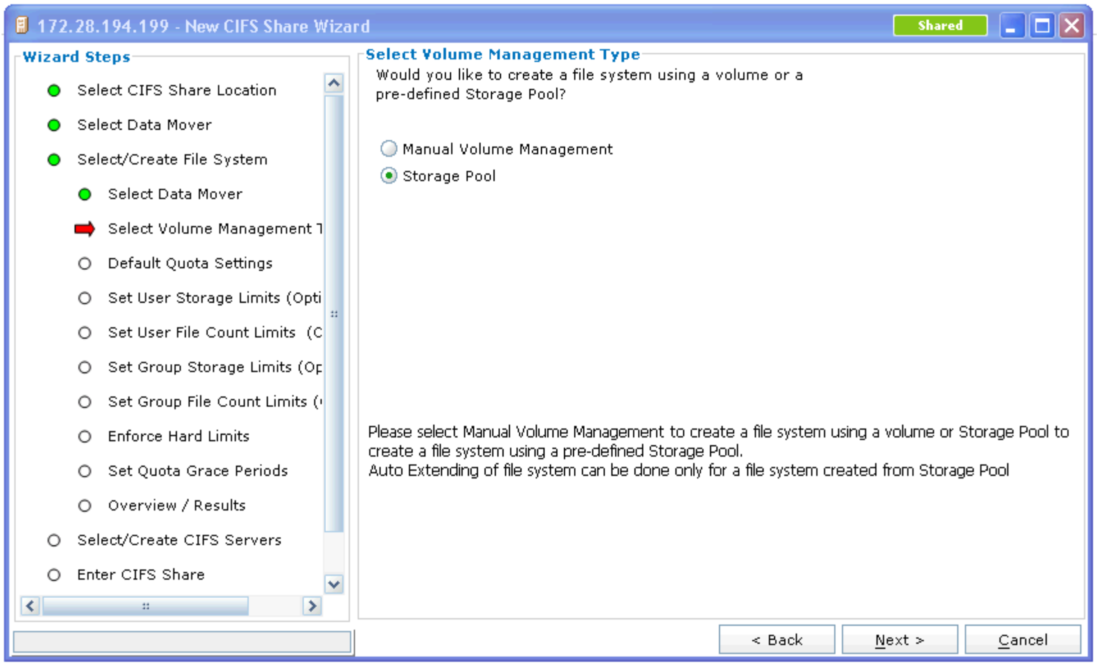
172.28.194.199 - New CIFS Share Wizard
Shared
Wizard Steps
● Select CIFS Share Location
● Select Data Mover
● Select/Create File System
● Select Data Mover
➡ Select Volume Management T
○ Default Quota Settings
○ Set User Storage Limits (Opti
○ Set User File Count Limits (C
○ Set Group Storage Limits (Op
○ Set Group File Count Limits (
○ Enforce Hard Limits
○ Set Quota Grace Periods
○ Overview / Results
○ Select/Create CIFS Servers
○ Enter CIFS Share
< Back
Next >
Cancel
290100

Select Volume Management Type

Would you like to create a file system using a volume or a pre-defined Storage Pool?

○ Manual Volume Management
● Storage Pool

Please select Manual Volume Management to create a file system using a volume or Storage Pool to create a file system using a pre-defined Storage Pool.
Auto Extending of file system can be done only for a file system created from Storage Pool

Step 8 Select the desired prefined pool.

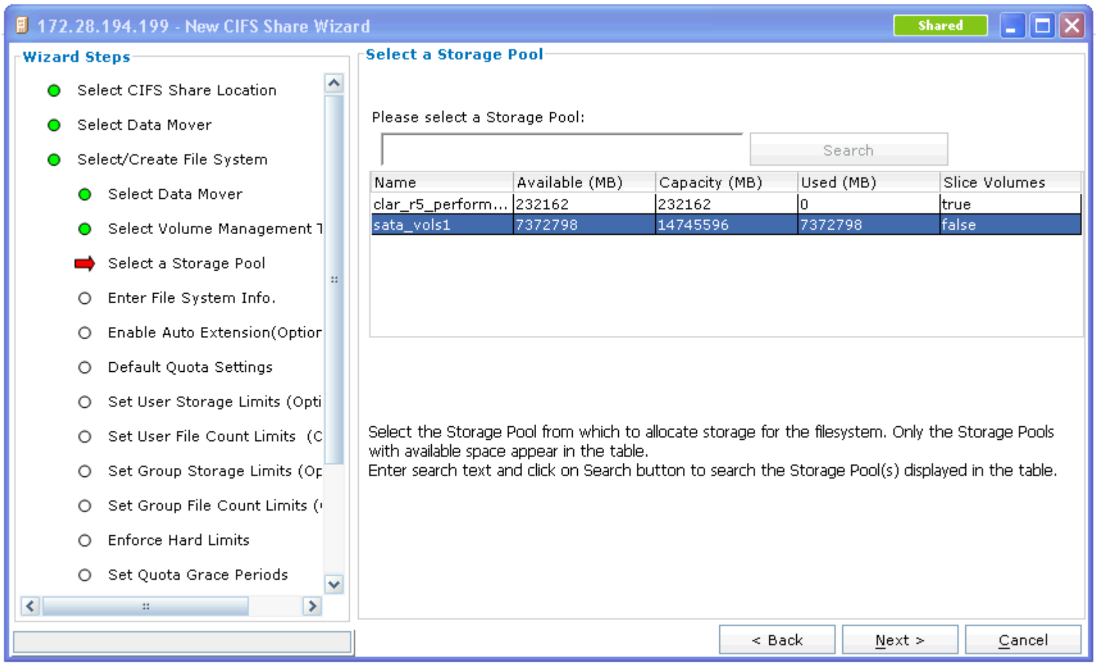
172.28.194.199 - New CIFS Share Wizard
Shared
Wizard Steps
● Select CIFS Share Location
● Select Data Mover
● Select/Create File System
● Select Data Mover
● Select Volume Management T
Select a Storage Pool
○ Enter File System Info.
○ Enable Auto Extension(Optior
○ Default Quota Settings
○ Set User Storage Limits (Opti
○ Set User File Count Limits (C
○ Set Group Storage Limits (Op
○ Set Group File Count Limits (
○ Enforce Hard Limits
○ Set Quota Grace Periods
Search
< Back
Next >
Cancel
290101

Select a Storage Pool

Please select a Storage Pool:

Name	Available (MB)	Capacity (MB)	Used (MB)	Slice Volumes
clar_r5_perform...	232162	232162	0	true
sata_vols1	7372798	14745596	7372798	false

Select the Storage Pool from which to allocate storage for the filesystem. Only the Storage Pools with available space appear in the table.
Enter search text and click on Search button to search the Storage Pool(s) displayed in the table.

Step 9 Define the CIFS share name. We used VDICIFS in this example.

Also here is where we select the check box for the EMC Data Deduplication feature.

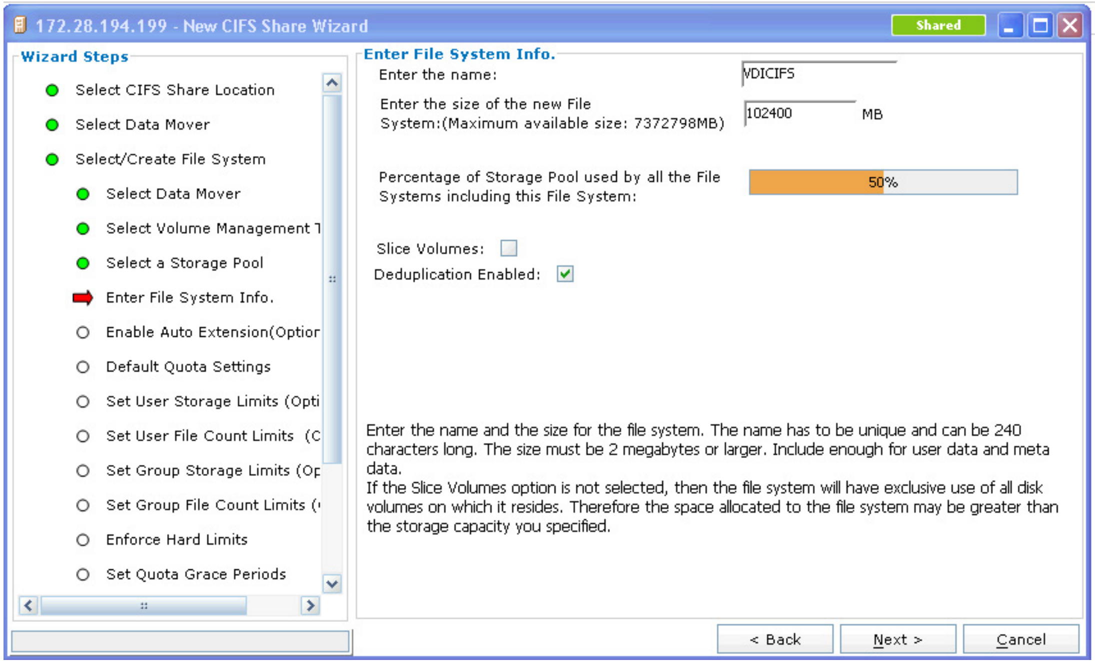
172.28.194.199 - New CIFS Share Wizard
Shared
Wizard Steps
● Select CIFS Share Location
● Select Data Mover
● Select/Create File System
● Select Data Mover
● Select Volume Management T
● Select a Storage Pool
Enter File System Info.
○ Enable Auto Extension(Optior
○ Default Quota Settings
○ Set User Storage Limits (Opti
○ Set User File Count Limits (C
○ Set Group Storage Limits (Op
○ Set Group File Count Limits (
○ Enforce Hard Limits
○ Set Quota Grace Periods
< Back
Next >
Cancel
290102

Enter File System Info.

Enter the name: VDICIFS

Enter the size of the new File System:(Maximum available size: 7372798MB) 102400 MB

Percentage of Storage Pool used by all the File Systems including this File System: 50%

Slice Volumes: □

Deduplication Enabled: ☑

Enter the name and the size for the file system. The name has to be unique and can be 240 characters long. The size must be 2 megabytes or larger. Include enough for user data and meta data.
If the Slice Volumes option is not selected, then the file system will have exclusive use of all disk volumes on which it resides. Therefore the space allocated to the file system may be greater than the storage capacity you specified.

Step 10 Here is where you define the high water mark and maximum capacity.

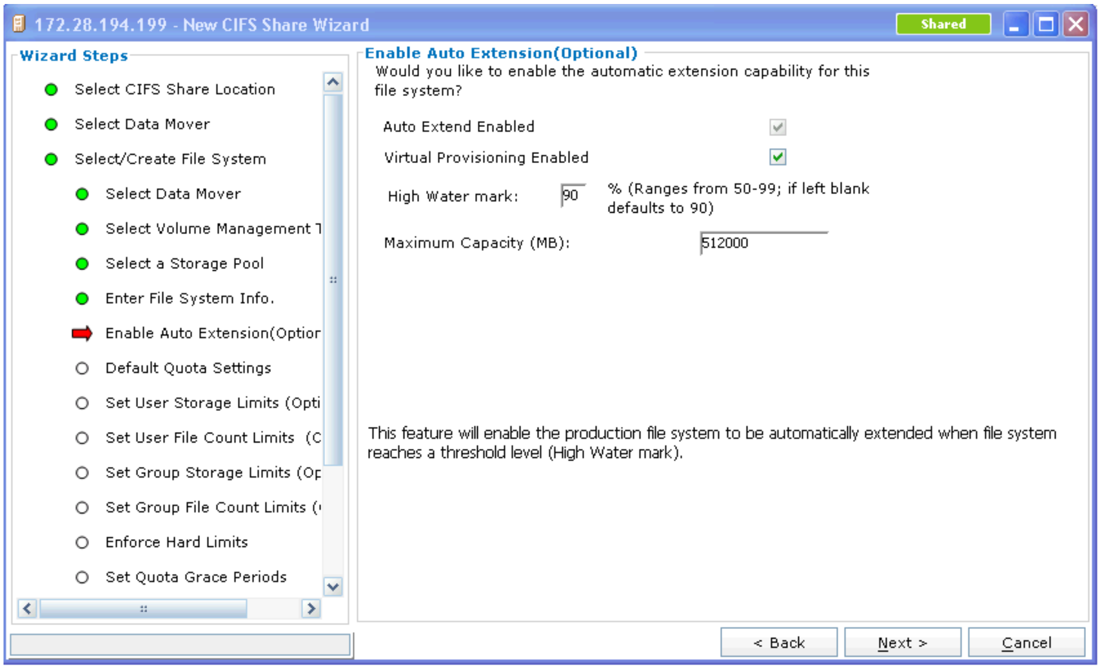
172.28.194.199 - New CIFS Share Wizard
Shared
Wizard Steps
Select CIFS Share Location
Select Data Mover
Select/Create File System
Select Data Mover
Select Volume Management T
Select a Storage Pool
Enter File System Info.
Enable Auto Extension(Optior
Default Quota Settings
Set User Storage Limits (Opti
Set User File Count Limits (C
Set Group Storage Limits (Op
Set Group File Count Limits (
Enforce Hard Limits
Set Quota Grace Periods
< Back
Next >
Cancel

Enable Auto Extension(Optional)

Would you like to enable the automatic extension capability for this file system?

Auto Extend Enabled ☑

Virtual Provisioning Enabled ☑

High Water mark: 90 % (Ranges from 50-99; if left blank defaults to 90)

Maximum Capacity (MB): 512000

This feature will enable the production file system to be automatically extended when file system reaches a threshold level (High Water mark).

290103

Step 11 Default Quota Setting was not used in this testing. Click **Next**.

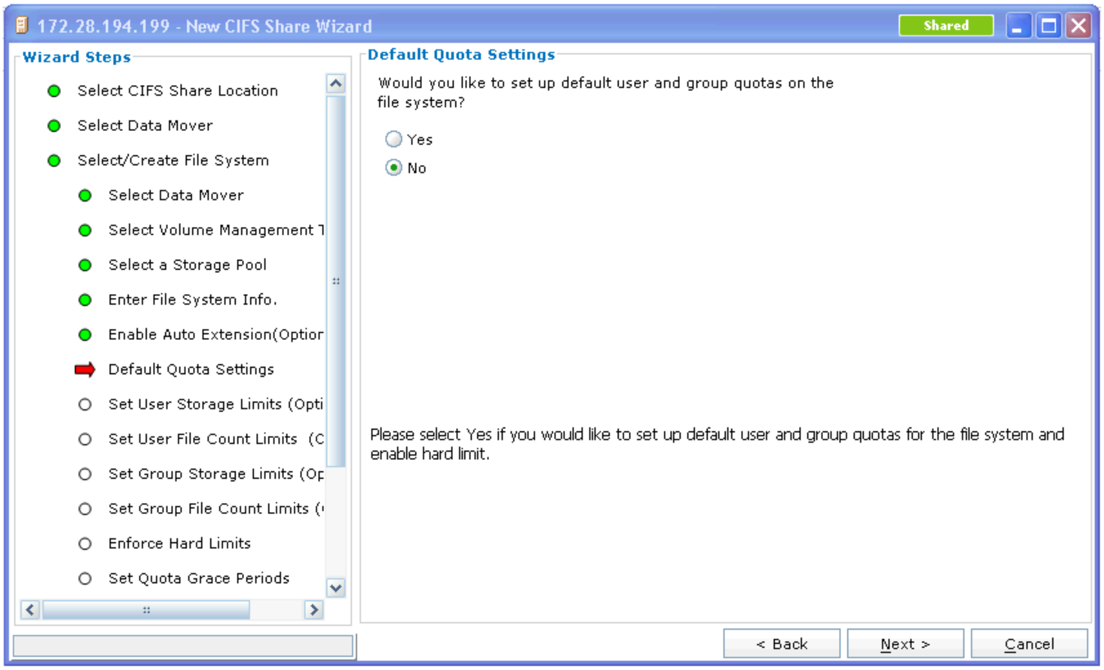
172.28.194.199 - New CIFS Share Wizard
Shared
Wizard Steps
Select CIFS Share Location
Select Data Mover
Select/Create File System
Select Data Mover
Select Volume Management T
Select a Storage Pool
Enter File System Info.
Enable Auto Extension(Optior
Default Quota Settings
Set User Storage Limits (Opti
Set User File Count Limits (C
Set Group Storage Limits (Op
Set Group File Count Limits (
Enforce Hard Limits
Set Quota Grace Periods
< Back
Next >
Cancel

Default Quota Settings

Would you like to set up default user and group quotas on the file system?

○ Yes

● No

Please select Yes if you would like to set up default user and group quotas for the file system and enable hard limit.

290104

Step 12 Review the configuration of the CIFS before you submit for the CIFS creation.

Once reviewed click **Submit**.

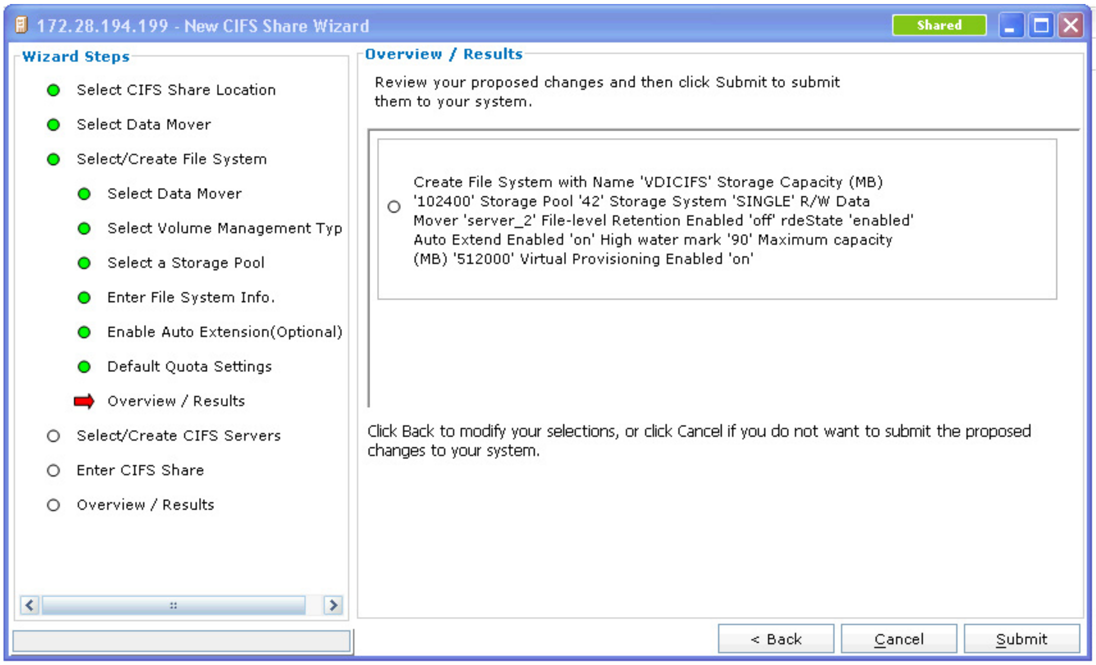
172.28.194.199 - New CIFS Share Wizard
Shared
Wizard Steps
● Select CIFS Share Location
● Select Data Mover
● Select/Create File System
● Select Data Mover
● Select Volume Management Typ
● Select a Storage Pool
● Enter File System Info.
● Enable Auto Extension(Optional)
● Default Quota Settings
➡ Overview / Results
○ Select/Create CIFS Servers
○ Enter CIFS Share
○ Overview / Results
Overview / Results
< Back
Cancel
Submit

Review your proposed changes and then click Submit to submit them to your system.

○ Create File System with Name 'VDICIFS' Storage Capacity (MB) '102400' Storage Pool '42' Storage System 'SINGLE' R/W Data Mover 'server_2' File-level Retention Enabled 'off' rdeState 'enabled' Auto Extend Enabled 'on' High water mark '90' Maximum capacity (MB) '512000' Virtual Provisioning Enabled 'on'

Click Back to modify your selections, or click Cancel if you do not want to submit the proposed changes to your system.

290105

Step 13 Once created you will see the new CIFS share in the available shares. Select the one you just create to complete the configuration process.

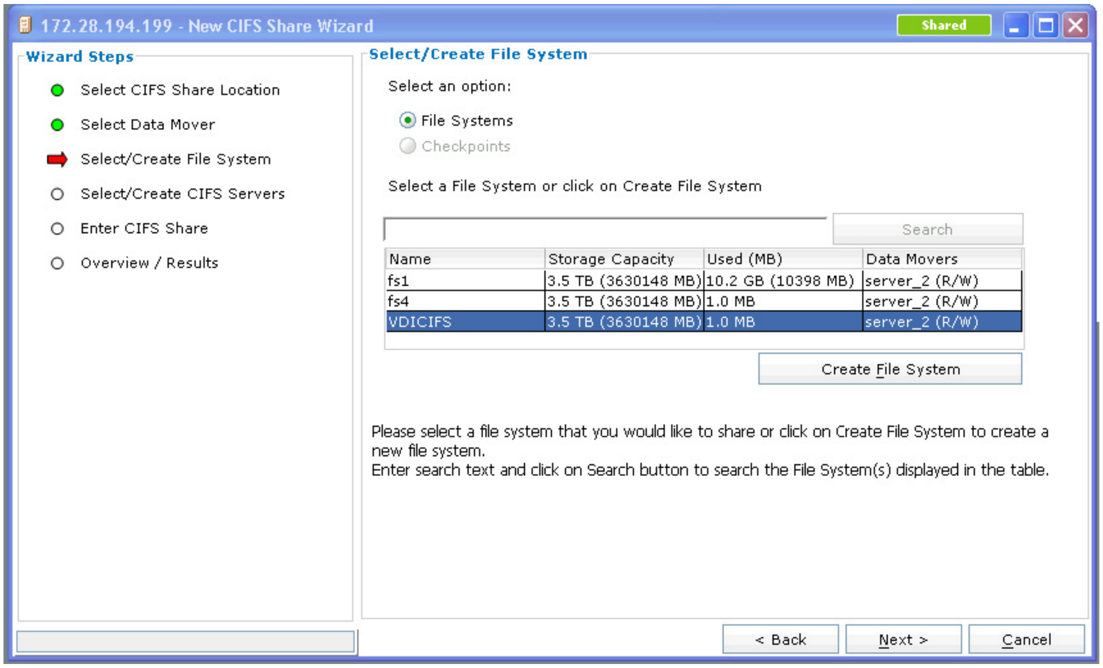
172.28.194.199 - New CIFS Share Wizard
Shared
Wizard Steps
● Select CIFS Share Location
● Select Data Mover
➡ Select/Create File System
○ Select/Create CIFS Servers
○ Enter CIFS Share
○ Overview / Results
Select/Create File System
Search
Create File System
< Back
Next >
Cancel

Select an option:

● File Systems
○ Checkpoints

Select a File System or click on Create File System

Name	Storage Capacity	Used (MB)	Data Movers
fs1	3.5 TB (3630148 MB)	10.2 GB (10398 MB)	server_2 (R/W)
fs4	3.5 TB (3630148 MB)	1.0 MB	server_2 (R/W)
VDICIFS	3.5 TB (3630148 MB)	1.0 MB	server_2 (R/W)

Please select a file system that you would like to share or click on Create File System to create a new file system.
Enter search text and click on Search button to search the File System(s) displayed in the table.

290106

Step 14 Select the CIFS share and click **Next**.

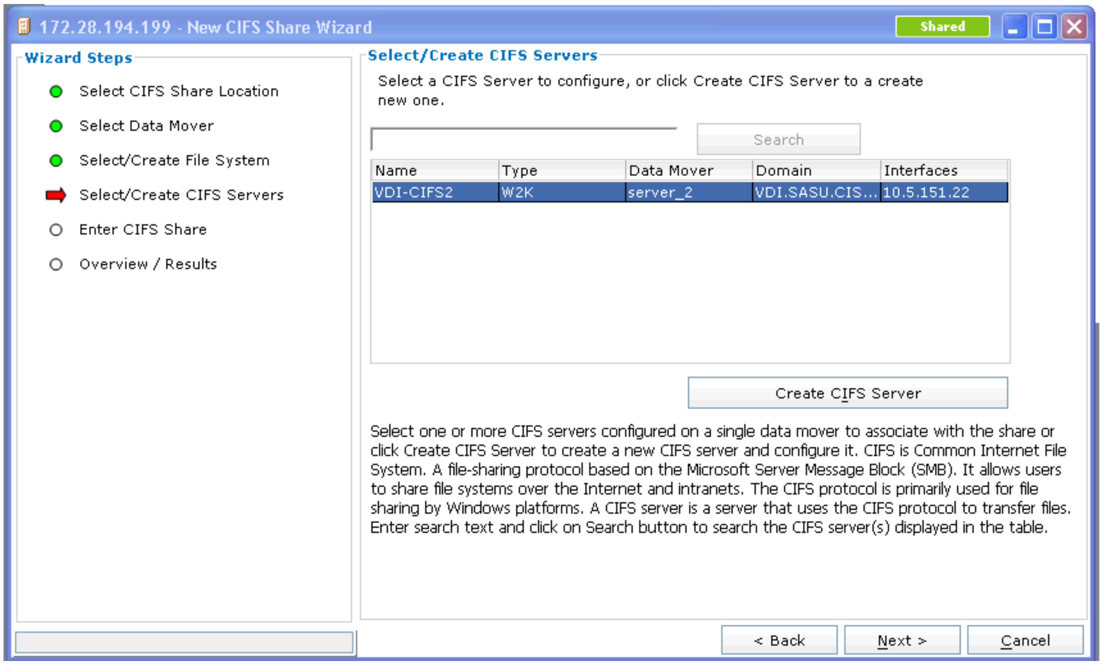
172.28.194.199 - New CIFS Share Wizard
Shared
Wizard Steps
Select CIFS Share Location
Select Data Mover
Select/Create File System
Select/Create CIFS Servers
Enter CIFS Share
Overview / Results
Search
Create CIFS Server
< Back
Next >
Cancel

Select/Create CIFS Servers

Select a CIFS Server to configure, or click Create CIFS Server to a create new one.

Name	Type	Data Mover	Domain	Interfaces
VDI-CIFS2	W2K	server_2	VDI.SASU.CIS...	10.5.151.22

Select one or more CIFS servers configured on a single data mover to associate with the share or click Create CIFS Server to create a new CIFS server and configure it. CIFS is Common Internet File System. A file-sharing protocol based on the Microsoft Server Message Block (SMB). It allows users to share file systems over the Internet and intranets. The CIFS protocol is primarily used for file sharing by Windows platforms. A CIFS server is a server that uses the CIFS protocol to transfer files. Enter search text and click on Search button to search the CIFS server(s) displayed in the table.

290107

Step 15 Here is where you giving the CIFS share the name that host will see it as.

Path name will be shown once the CIFs name is entered.

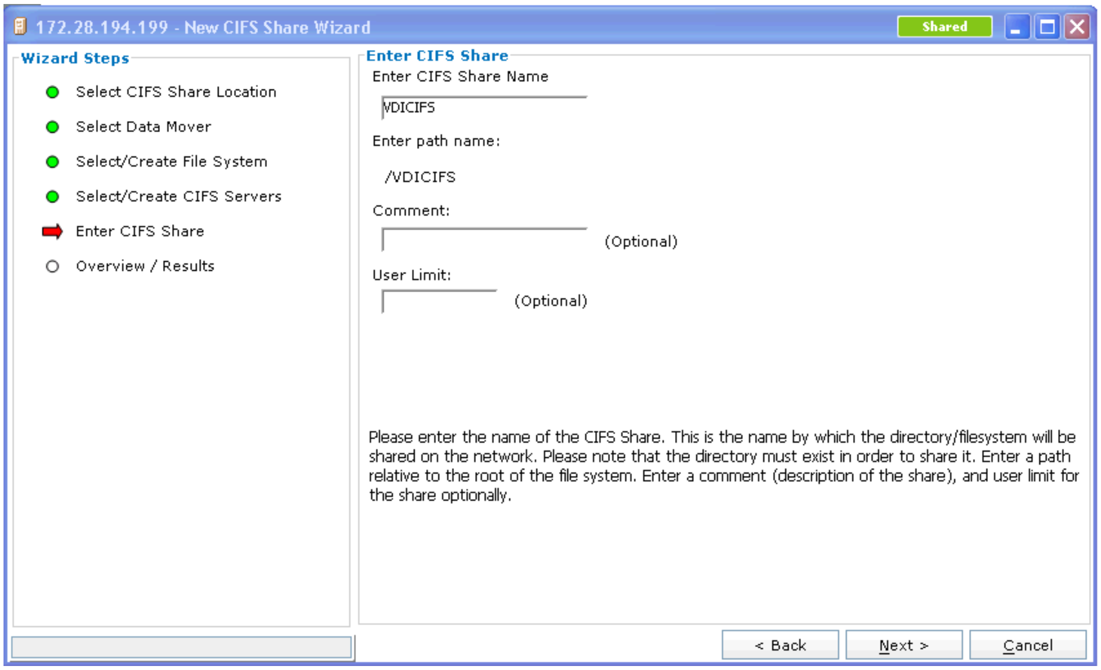
172.28.194.199 - New CIFS Share Wizard
Shared
Wizard Steps
Select CIFS Share Location
Select Data Mover
Select/Create File System
Select/Create CIFS Servers
Enter CIFS Share
Overview / Results
< Back
Next >
Cancel

Enter CIFS Share

Enter CIFS Share Name: VDICIFS

Enter path name: /VDICIFS

Comment: (Optional)

User Limit: (Optional)

Please enter the name of the CIFS Share. This is the name by which the directory/filesystem will be shared on the network. Please note that the directory must exist in order to share it. Enter a path relative to the root of the file system. Enter a comment (description of the share), and user limit for the share optionally.

290108

Step 16 Final Overview/Results page. After you review naming and path, click **Finish** to complete the process.

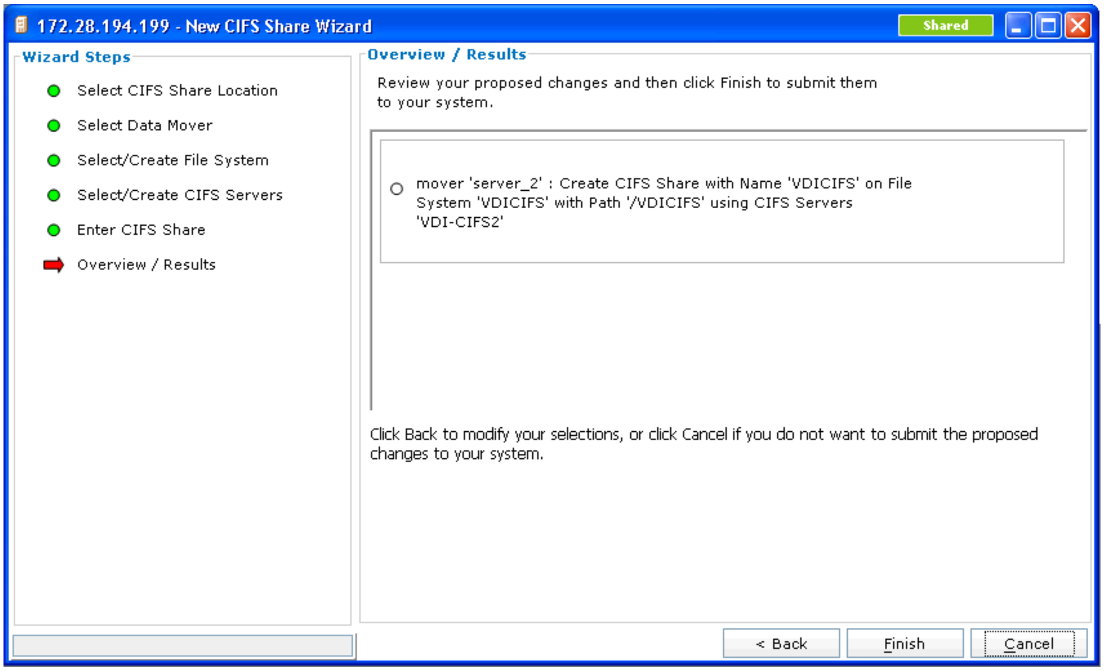
172.28.194.199 - New CIFS Share Wizard
Shared
Wizard Steps
Select CIFS Share Location
Select Data Mover
Select/Create File System
Select/Create CIFS Servers
Enter CIFS Share
Overview / Results
Overview / Results
< Back
Finish
Cancel

Review your proposed changes and then click Finish to submit them to your system.

○ mover 'server_2' : Create CIFS Share with Name 'VDICIFS' on File System 'VDICIFS' with Path '/VDICIFS' using CIFS Servers 'VDI-CIFS2'

Click Back to modify your selections, or click Cancel if you do not want to submit the proposed changes to your system.

290109

Step 17 Here we see the complete and the successful message.

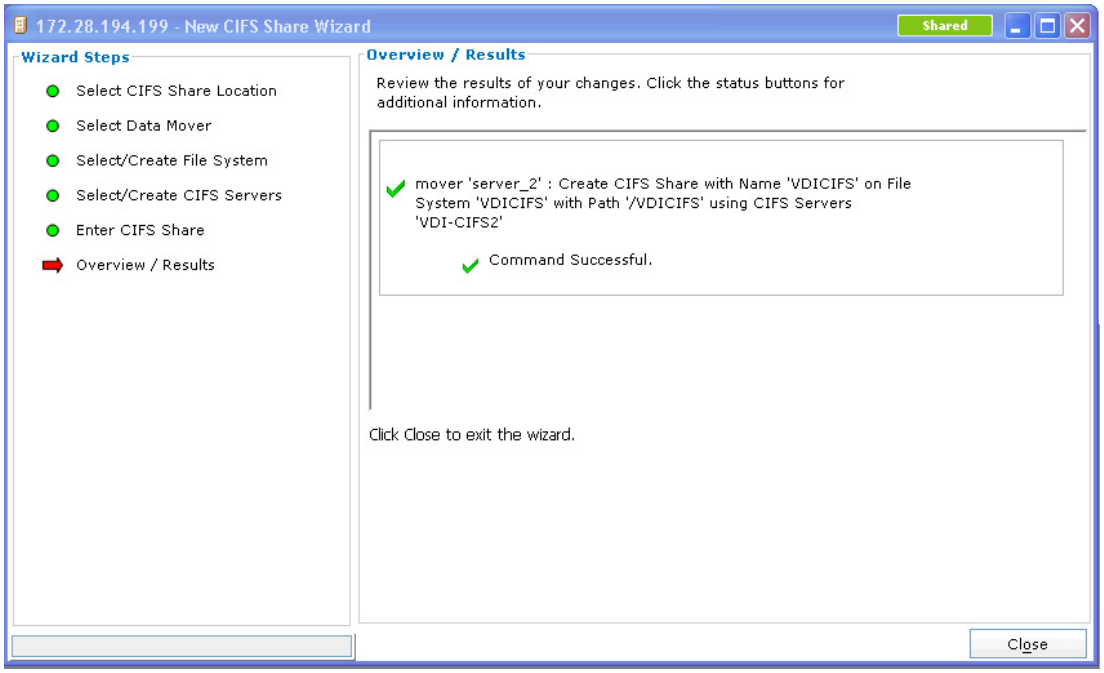
172.28.194.199 - New CIFS Share Wizard
Shared
Wizard Steps
Select CIFS Share Location
Select Data Mover
Select/Create File System
Select/Create CIFS Servers
Enter CIFS Share
Overview / Results
Overview / Results
Close

Review the results of your changes. Click the status buttons for additional information.

✓ mover 'server_2' : Create CIFS Share with Name 'VDICIFS' on File System 'VDICIFS' with Path '/VDICIFS' using CIFS Servers 'VDI-CIFS2'

✓ Command Successful.

Click Close to exit the wizard.

290110

Step 18 After those steps are completed the CIFs share will now appear in the main plane for the Share. VDICIFS is no listed as a available share.

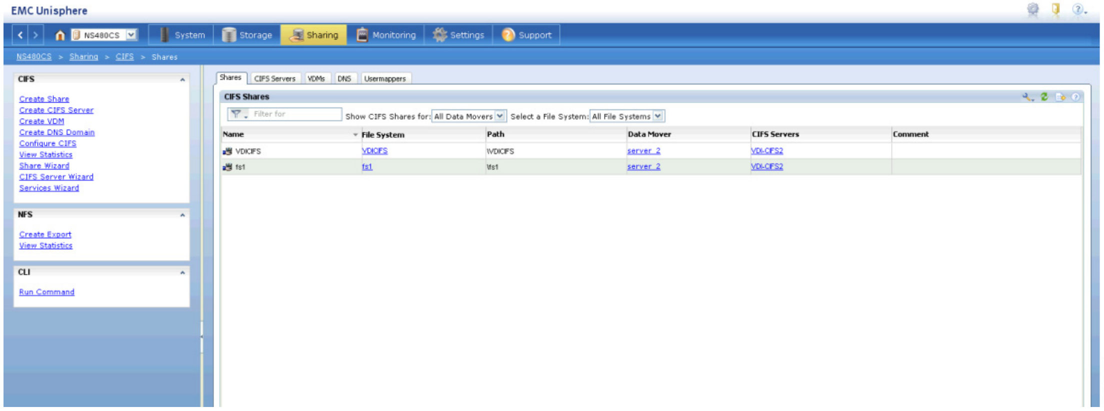

Show CIFS Shares for: All Data Movers Select a File System: All File Systems

Name	File System	Path	Data Mover	CIFS Servers	Comment
VDICFS	VDICFS	\VDICFS	server_2	VDI-CIFS2	
fs1	fs1	\fs1	server_2	VDI-CIFS2	

290111

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.

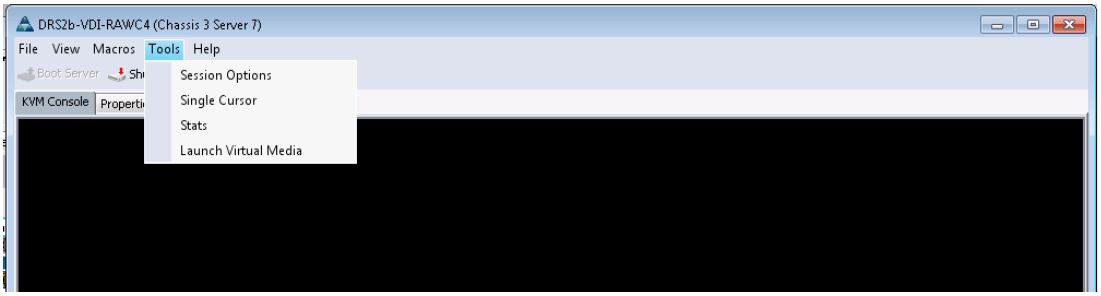

290007

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 command would not persist through a reboot of the host and would need to be added to the /etc/rc.local file of the system.

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. 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 VEM Module Installation

ESXi hosts were installed with the Nexus 1000V Virtual Ethernet Module (VEM) and added to vCenter.

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

- http://www.cisco.com/en/US/docs/switches/datacenter/nexus1000/sw/4_0_4_s_v_1_3_b/install/vsm/guide/n1000v_vsm_install.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_0_4_s_v_1_3_b/install/vem/guide/n1000v_vem_install.html

Network Configuration

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

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

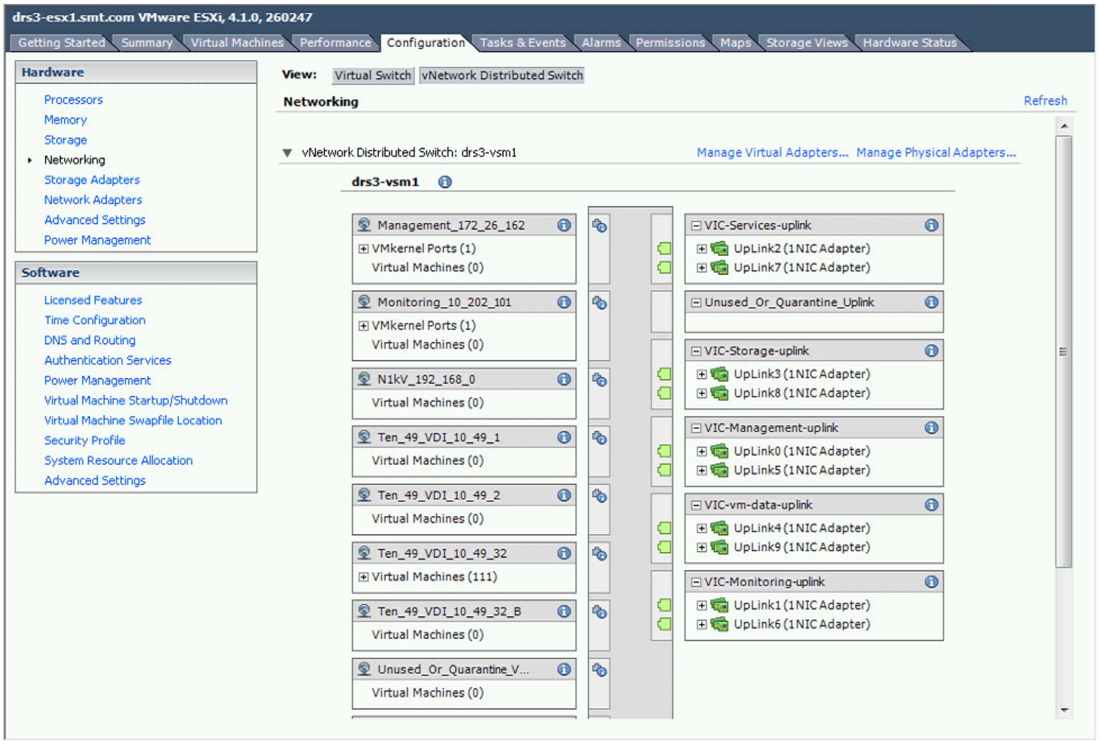

290008

Using the Cisco Nexus 1000V for network configuration provides advance 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 EMC 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 seven servers as NFS mounts and the launcher VMs were used to generate the load to the desktops on the eight servers. Each blade was assigned 160 desktops each for a total of 1280 desktops.

Figure 26 VMware vSphere Configuration

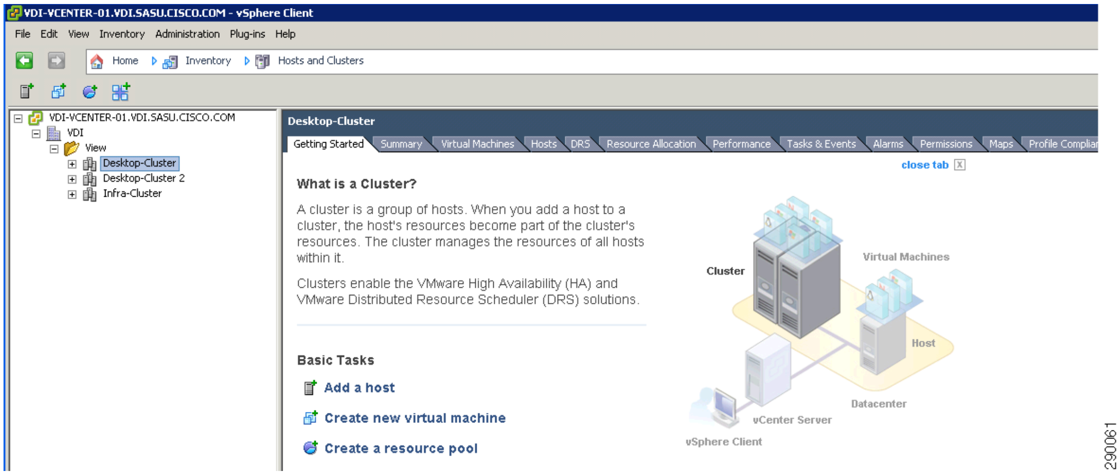

What is a Cluster?

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.

Basic Tasks

- Add a host
- Create new virtual machine
- Create a resource pool

VMware View 4.5 Components and Setup

VMware View Installation Overview

The *VMware View Installation Guide* available on the VMware Website has detailed procedures about installing View Connection Server and View Composer 2.5. There are no special configuration instructions required for this solution.

The *ESXi Installation and vCenter Server Setup Guide* available on the VMware Website has detailed procedures about installing vCenter Server and ESXi and is not covered in further detail in this paper. There are no special configuration instructions required for this solution.

VMware View Setup

Before deploying desktop pools, ensure that the following steps from the *VMware View Installation Guide* have been completed:

- Prepare Active Directory
- Install View Composer 2.5 on vCenter Server
- Install View Connection Server
- Add a vCenter Server instance to View Manager

VMware View Desktop Pool Configuration

VMware recommends using a maximum of 250 desktops per replica image, which requires creating a unique pool for every 250 desktops. In this solution, persistent automated desktop pools were used.

To create a persistent/dedicated automated desktop pool as configured for this solution, complete the following steps:

Step 1 Log in to the VMware View Administration page, which is located at https://server/admin, where "server" is the IP address or DNS name of the View Manager server.

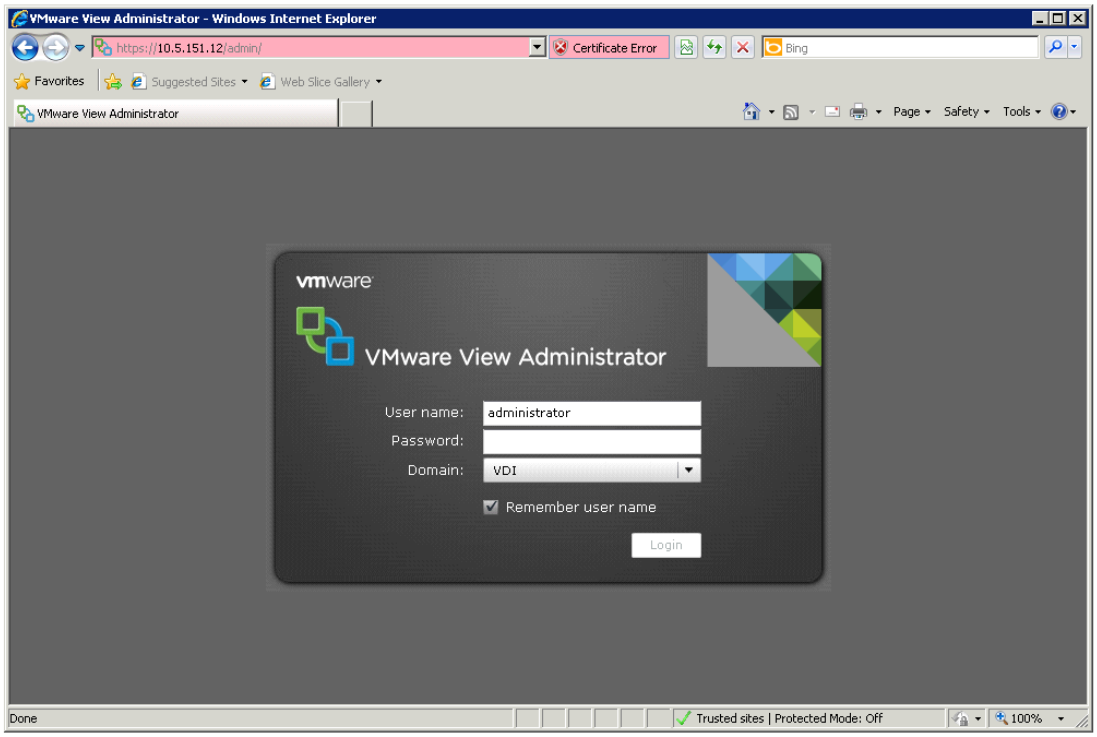

VMware View Administrator - Windows Internet Explorer

https://10.5.151.12/admin/

VMware View Administrator

User name: administrator

Password:

Domain: VDI

☑ Remember user name

Step 2 From the main VMWAare View Administrator Dashboard, under Inventory, click the **Pools** link in the left pane.

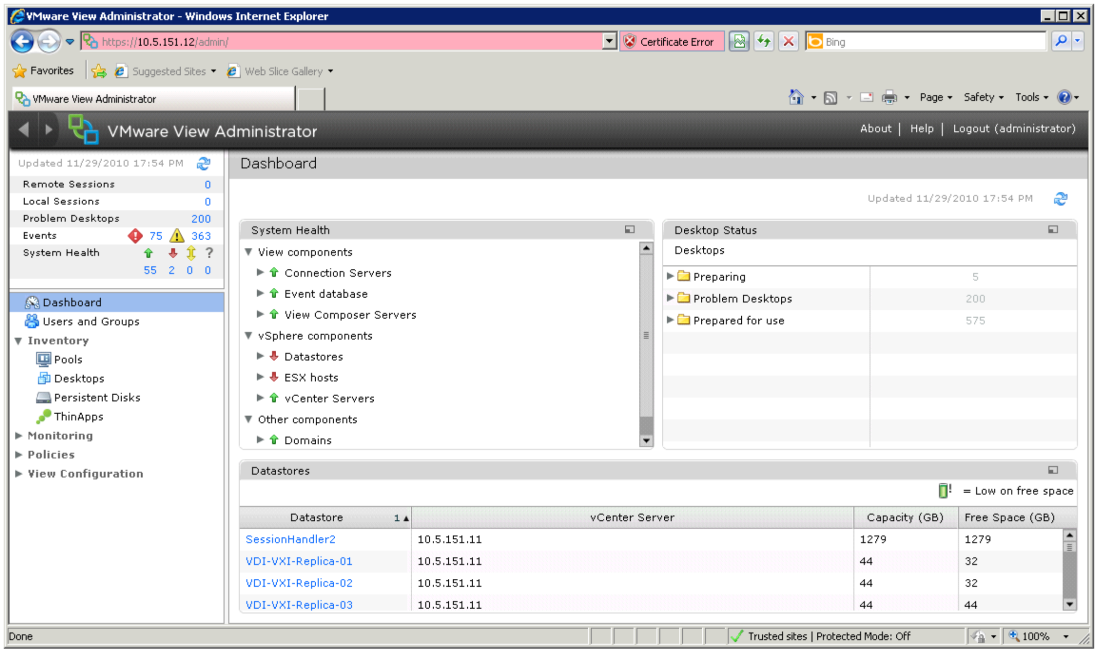

Remote Sessions		0
Local Sessions		0
Problem Desktops		200
Events	75	363
System Health	55 2	0 0

Desktops	
Preparing	5
Problem Desktops	200
Prepared for use	575

Datastore	vCenter Server	Capacity (GB)	Free Space (GB)
SessionHandler2	10.5.151.11	1279	1279
VDI-VXI-Replica-01	10.5.151.11	44	32
VDI-VXI-Replica-02	10.5.151.11	44	32
VDI-VXI-Replica-03	10.5.151.11	44	44

290063

Step 3 Click the **Add** button under the Pools banner.

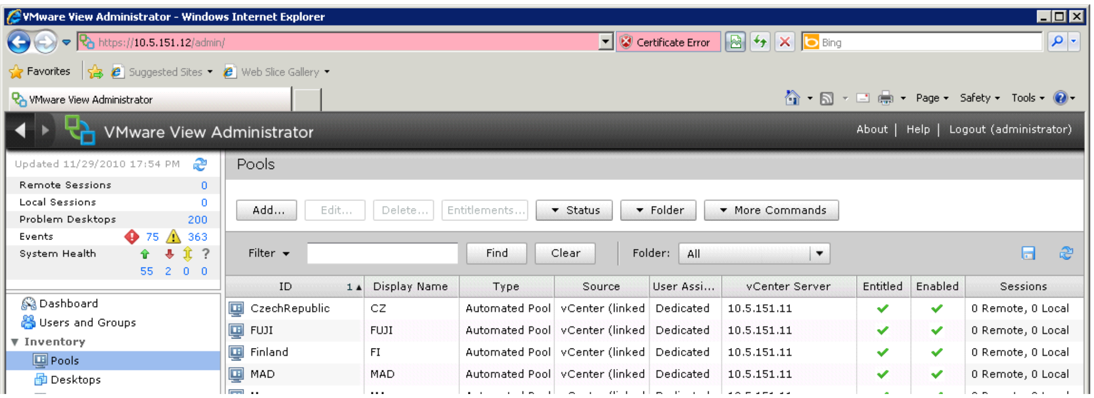

Remote Sessions		0
Local Sessions		0
Problem Desktops		200
Events	75	363
System Health	55 2	0 0

ID	Display Name	Type	Source	User Assi...	vCenter Server	Entitled	Enabled	Sessions
CzechRepublic	CZ	Automated Pool	vCenter (linked	Dedicated	10.5.151.11	✓	✓	0 Remote, 0 Local
FUJI	FUJI	Automated Pool	vCenter (linked	Dedicated	10.5.151.11	✓	✓	0 Remote, 0 Local
Finland	FI	Automated Pool	vCenter (linked	Dedicated	10.5.151.11	✓	✓	0 Remote, 0 Local
MAD	MAD	Automated Pool	vCenter (linked	Dedicated	10.5.151.11	✓	✓	0 Remote, 0 Local

290064

Step 4 In the Type page, select Automated Pool to define the pool type and click **Next**.

Three pool types exist:

- Automated—Uses a vCenter virtual machine snapshot or server template to generate the desktops.
- Manual—Uses an existing set of machine for the desktop provisioning; this can any type of physical machine.
- Terminal Services pool—Provides offline desktop access.

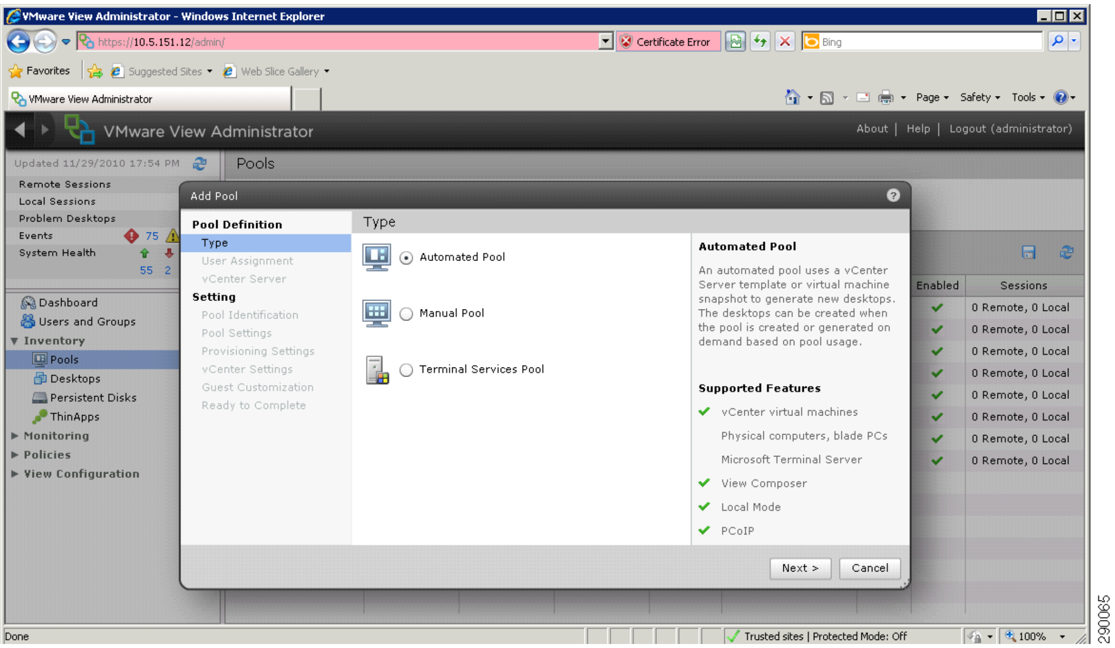

Add Pool

Pool Definition

- Type
- User Assignment
- vCenter Server

Setting

- Pool Identification
- Pool Settings
- Provisioning Settings
- vCenter Settings
- Guest Customization
- Ready to Complete

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

Next > | Cancel

Step 5 In the User Assignment page, select the Dedicated radio button and ensure that the Enable automatic assignment checkbox is selected. Click **Next**. The two options for user assignment method are Dedicated and floating. We selected a Dedicated user assignment in which each user will be assigned to a specific desktop resource from within the pool. A floating pool will provision any available desktop from the pool whereas dedicated will assign the a desktop resource to a user.

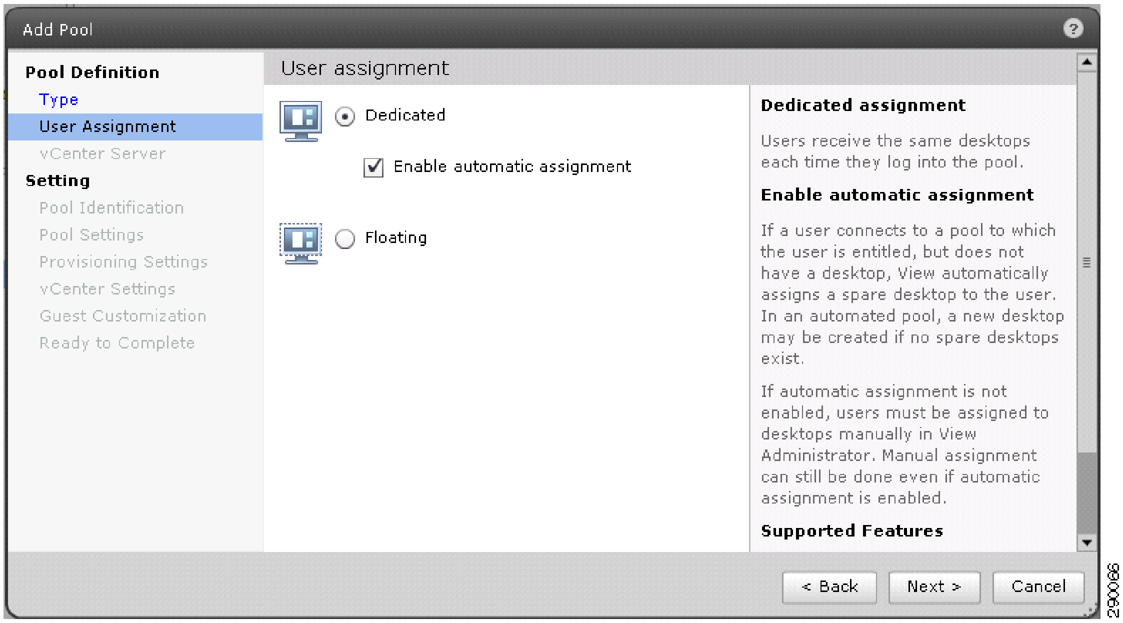

Add Pool

Pool Definition

- Type
- User Assignment
- vCenter Server

Setting

- Pool Identification
- Pool Settings
- Provisioning Settings
- vCenter Settings
- Guest Customization
- Ready to Complete

User assignment

● 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

< Back | Next > | Cancel

Step 6 In the vCenter Server page, select View Composer linked clones and select a vCenter Server that supports View Composer, as shown in the following figure. Click **Next**.

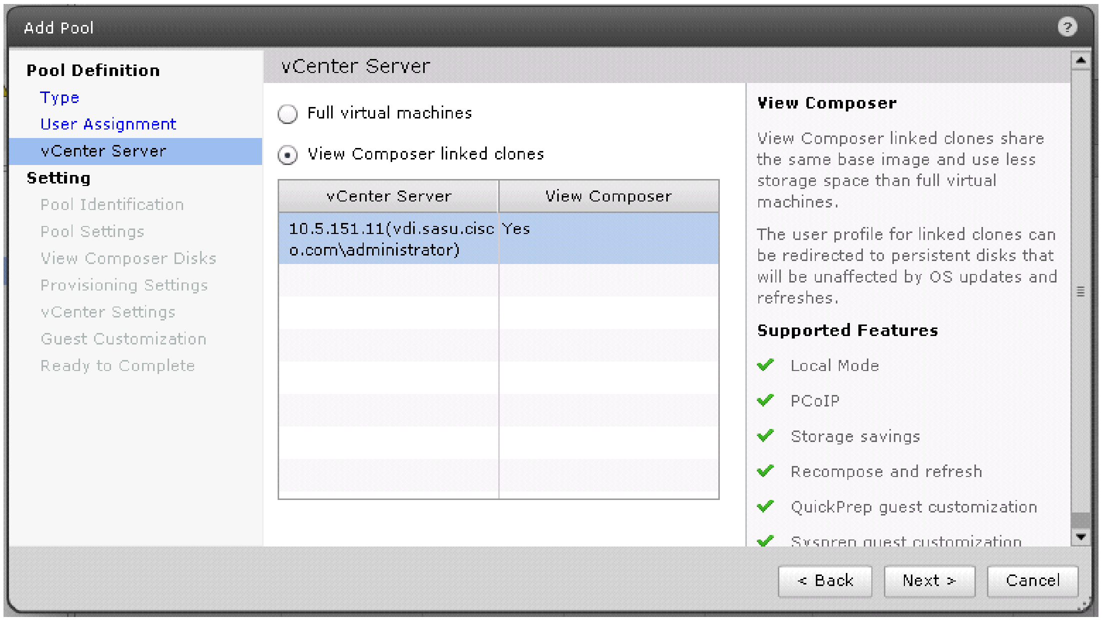

vCenter Server

○ Full virtual machines

● View Composer linked clones

vCenter Server	View Composer
10.5.151.11(vdi.sasu.cisco.com\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

290067

Step 7 In the vCenter Server page, select View Composer linked clones and select a vCenter Server that supports View Composer, as shown in the following figure. Click **Next**.

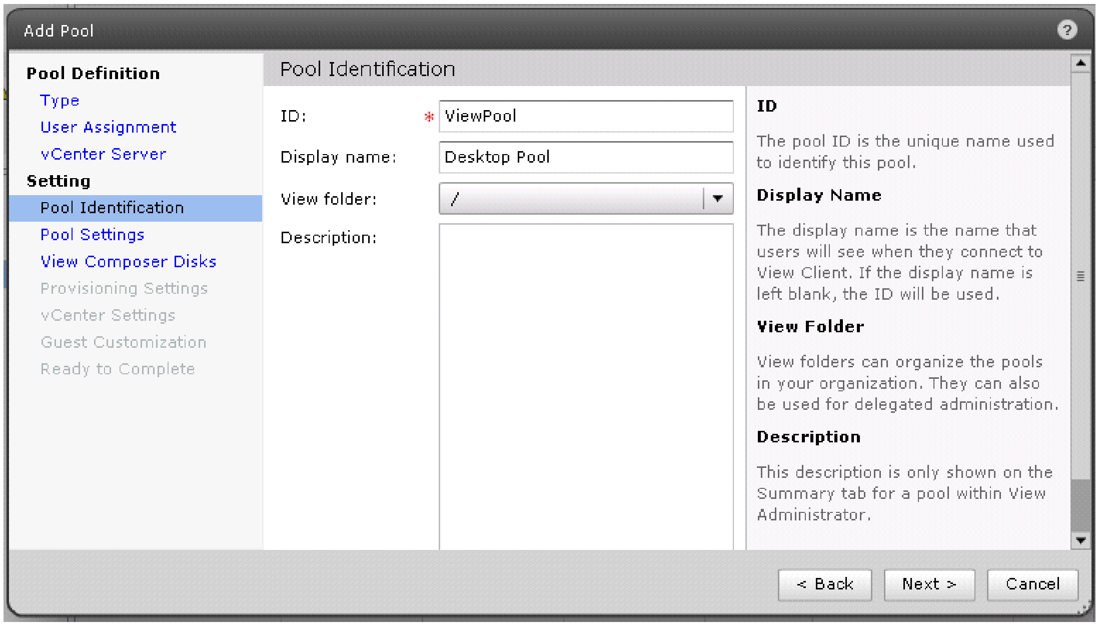

Pool Identification

ID: * ViewPool

Display name: Desktop Pool

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.

290068

Step 8 In the Pool Identification page, enter the required information and click **Next**.

Make sure the pool name is unique. The Display name will be shown in the border banner of the View 4.5 desktop. Display name and Description fields are optional.

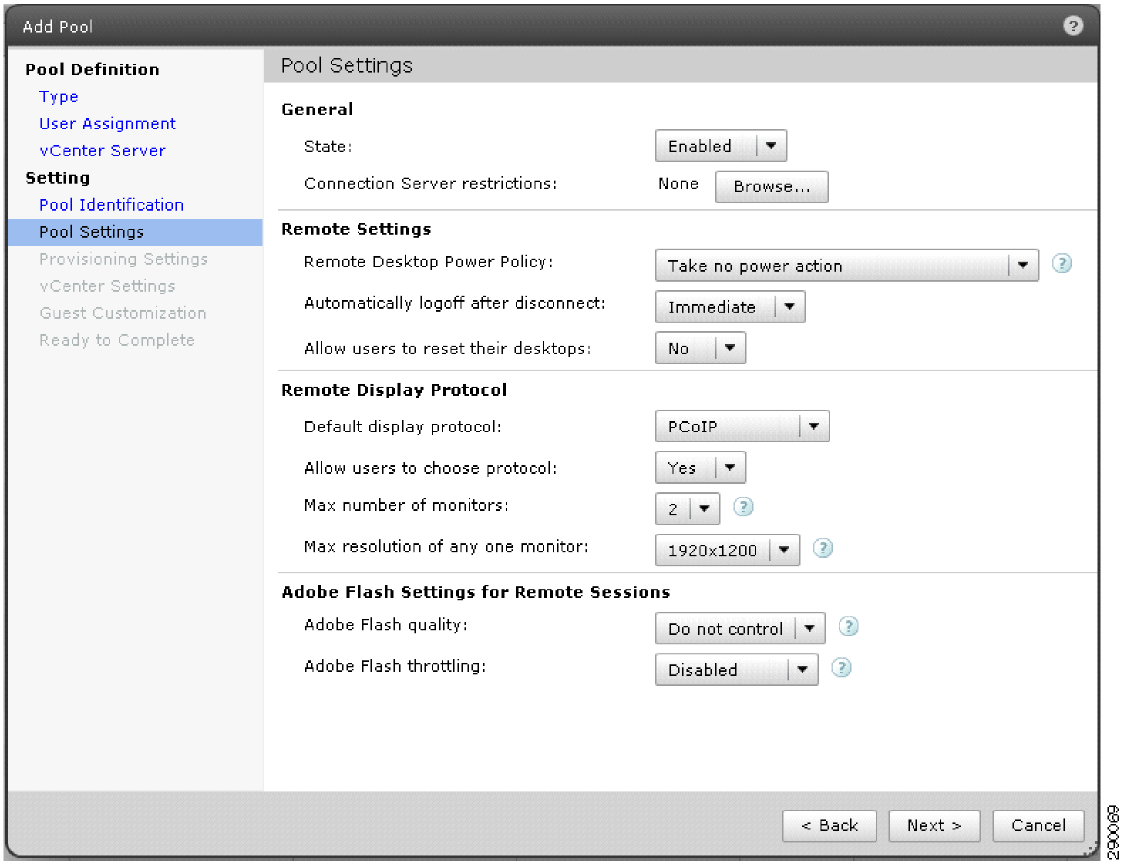

Pool Settings

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

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

Step 9 In the Pool Settings page, make any required changes and click **Next**.

For the testing done in this paper we changed only one setting, the automatically logoff after disconnect, which was changed to Immediate; all other settings were left as default.

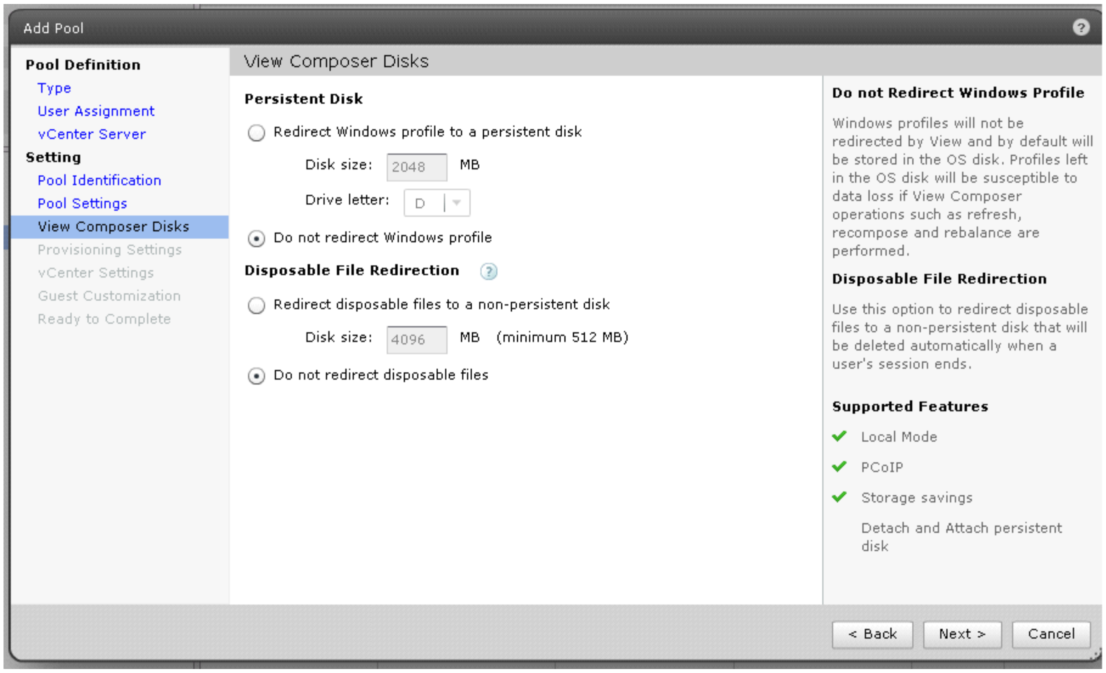

Add Pool

Pool Definition
- Type
- User Assignment
- vCenter Server

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

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

< Back | Next > | Cancel

290070

Step 10 In the View Composer Disk page, select both the Do not redirect Windows profile and Do not redirect disposal files radio buttons and click **Next**.

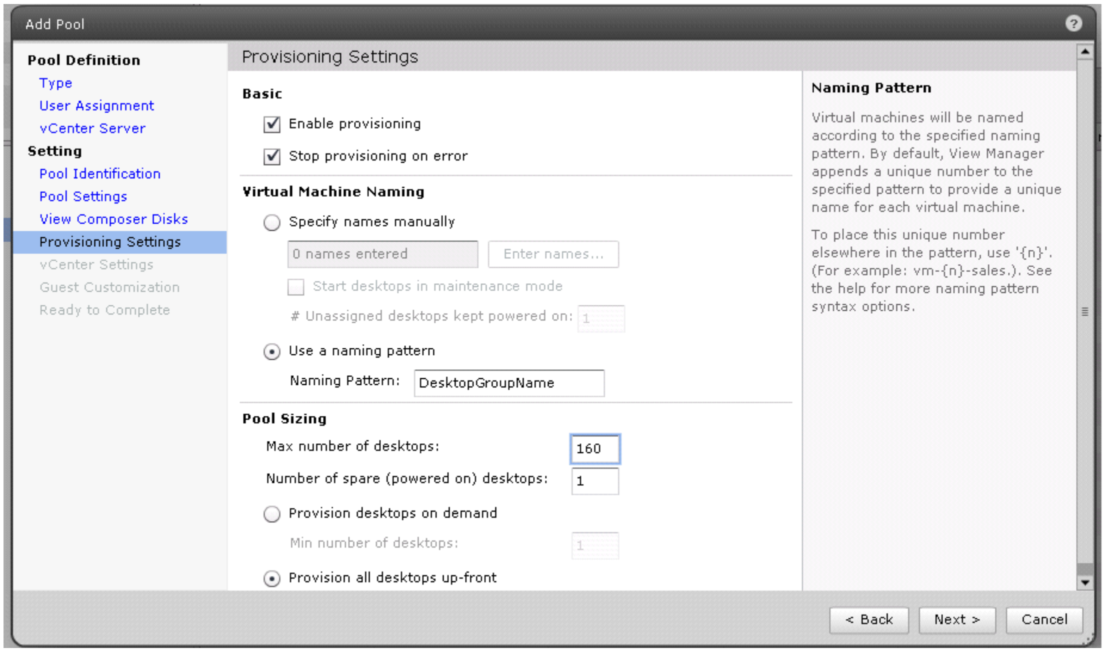

Add Pool

Pool Definition
- Type
- User Assignment
- vCenter Server

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

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: DesktopGroupName

Pool Sizing

Max number of desktops: 160

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.

< Back | Next > | Cancel

290071

Step 11 In the Provisioning Settings page, select a name for the desktop pool and enter the number of desktops to provision, as shown in the following figure. This example shows a pool of 160 desktops will be generated with the naming pattern Desktop-1 through Desktop-160. Note that the naming convention must be under 14 characters. Click **Next**.

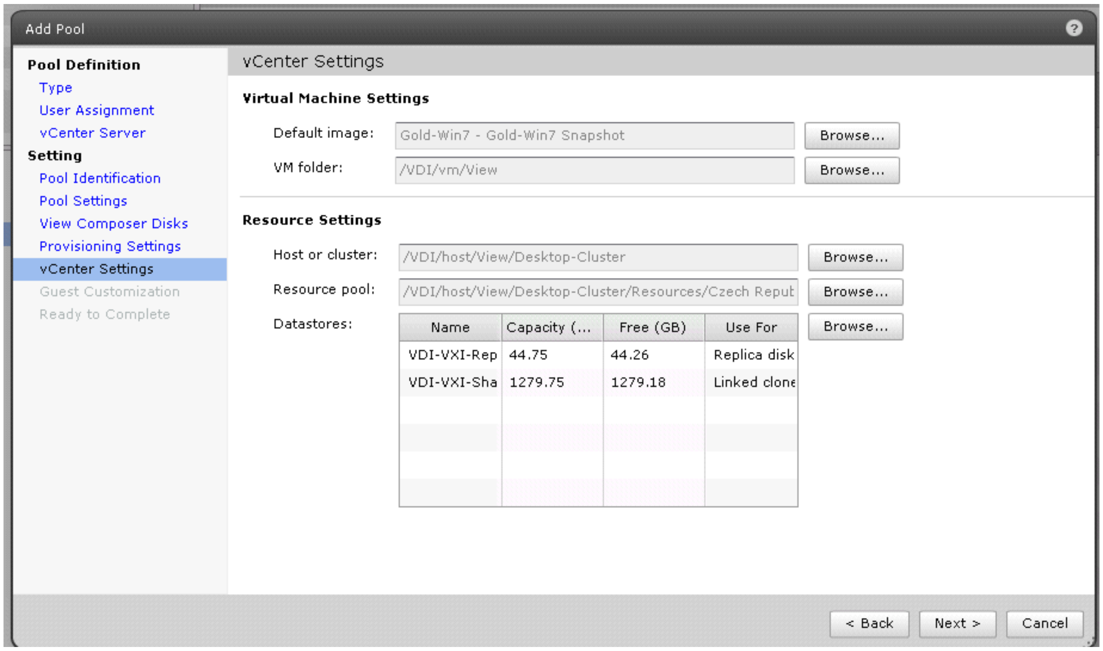

vCenter Settings

Virtual Machine Settings

Default image: Gold-Win7 - Gold-Win7 Snapshot

VM folder: /VDI/vm/View

Resource Settings

Host or cluster: /VDI/host/View/Desktop-Cluster

Resource pool: /VDI/host/View/Desktop-Cluster/Resources/Czech Reput

Datastores:

Name	Capacity (...	Free (GB)	Use For
VDI-VXI-Rep	44.75	44.26	Replica disk
VDI-VXI-Sha	1279.75	1279.18	Linked clone

290072

Step 12 In the vCenter Settings page, browse to select a default image and folder for the virtual machines, the cluster hosting the virtual desktops, the resource pool to hold the desktops, and the datastores that will be used to deploy the desktops, as shown in the following figure. Below we have selected one 1.2 TB FC LUN for the clones to reside on and one 45GB EFD LUN for the replica to reside on.

Note For a virtual machine to be eligible and show up as a possible master image, it needs to have a snapshot taken when the VM is powered off.

Click **Next**.

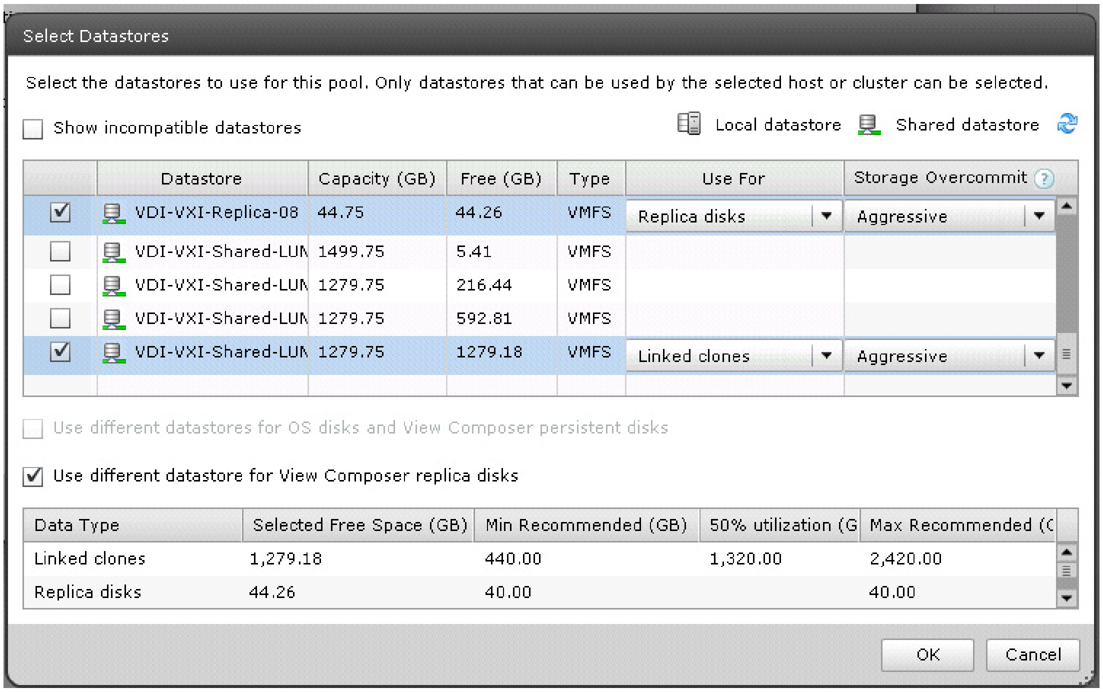

Select Datastores

Select the datastores to use for this pool. Only datastores that can be used by the selected host or cluster can be selected.

□ Show incompatible datastores

	Datastore	Capacity (GB)	Free (GB)	Type	Use For	Storage Overcommit
☑	VDI-VXI-Replica-08	44.75	44.26	VMFS	Replica disks	Aggressive
□	VDI-VXI-Shared-LUN	1499.75	5.41	VMFS		
□	VDI-VXI-Shared-LUN	1279.75	216.44	VMFS		
□	VDI-VXI-Shared-LUN	1279.75	592.81	VMFS		
☑	VDI-VXI-Shared-LUN	1279.75	1279.18	VMFS	Linked clones	Aggressive

□ Use different datastores for OS disks and View Composer persistent disks

☑ Use different datastore for View Composer replica disks

Data Type	Selected Free Space (GB)	Min Recommended (GB)	50% utilization (G	Max Recommended ((
Linked clones	1,279.18	440.00	1,320.00	2,420.00
Replica disks	44.26	40.00		40.00

290073

Step 13 In the Select Datastores page, select Use different datastore for View Composer replica disks and select the datastores for replica and linked clone images, as shown in the following figure.

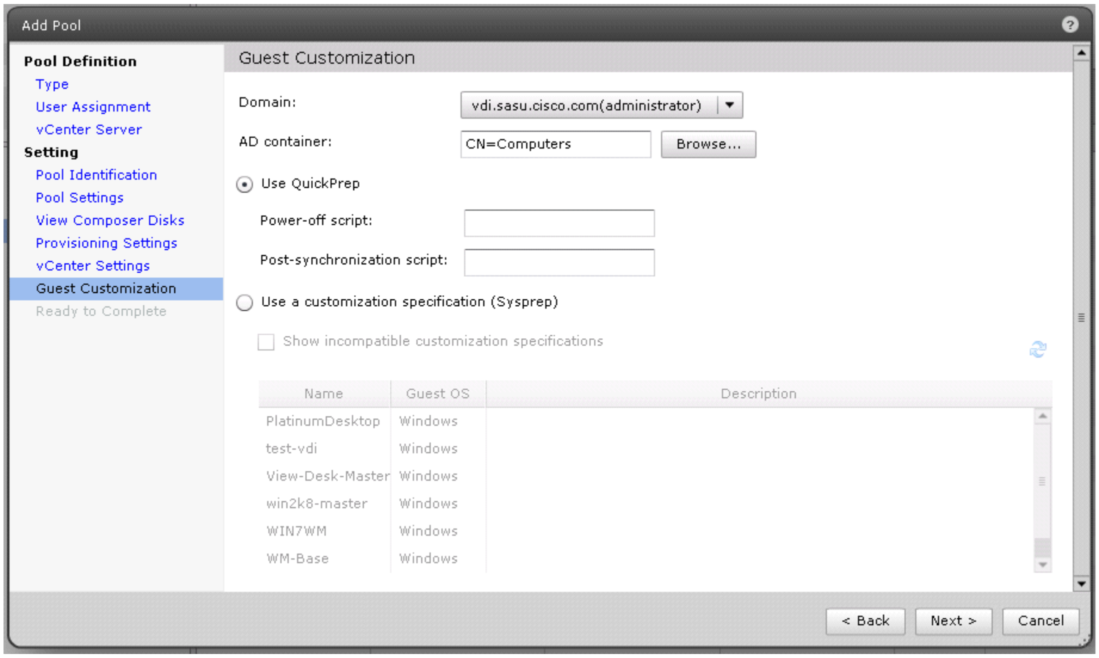

Add Pool

Guest Customization

Domain: vdi.sasu.cisco.com(administrator)

AD container: CN=Computers

● Use QuickPrep

Power-off script:

Post-synchronization script:

○ Use a customization specification (Sysprep)

□ Show incompatible customization specifications

Name	Guest OS	Description
PlatinumDesktop	Windows	
test-vdi	Windows	
View-Desk-Master	Windows	
win2k8-master	Windows	
WIN7WM	Windows	
WM-Base	Windows	

290074

Step 14 In the Guest Customization page, select the domain and AD container and then select the **Use QuickPrep** radio button. Click **Next**.

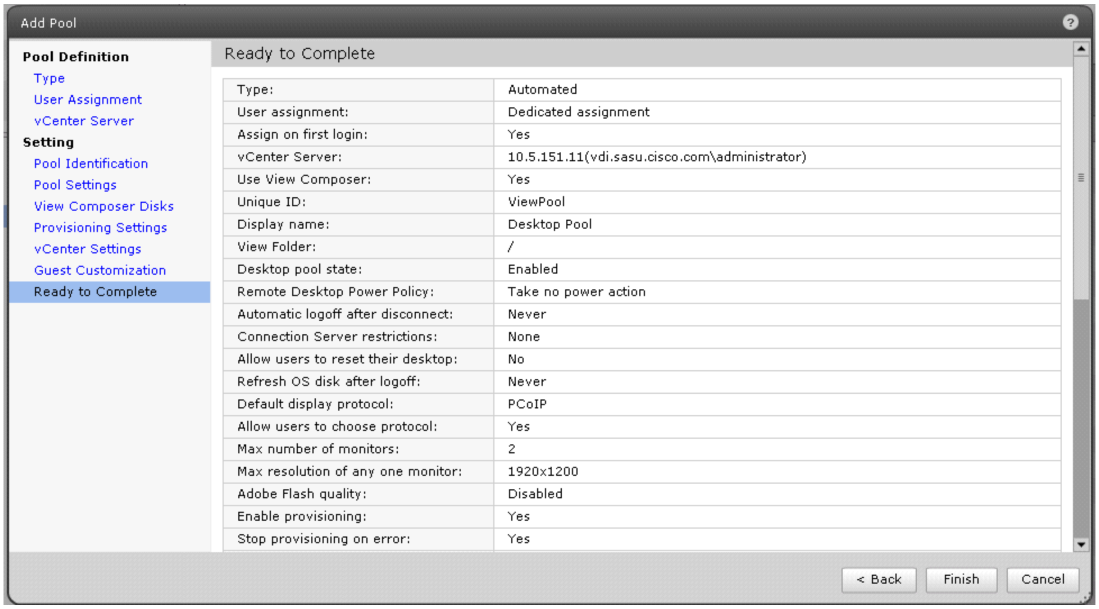

Ready to Complete

Type:	Automated
User assignment:	Dedicated assignment
Assign on first login:	Yes
vCenter Server:	10.5.151.11(vdi.sasu.cisco.com\administrator)
Use View Composer:	Yes
Unique ID:	ViewPool
Display name:	Desktop Pool
View Folder:	/
Desktop pool state:	Enabled
Remote Desktop Power Policy:	Take no power action
Automatic logoff after disconnect:	Never
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

Step 15 In the Ready to Complete page, verify the settings for the pool and then click the **Finish** button to start the deployment of the virtual desktops.

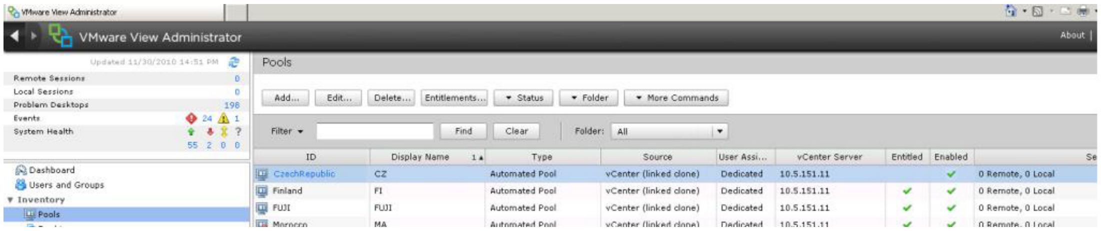

Updated 11/30/2010 14:51 PM

Remote Sessions	0
Local Sessions	0
Problem Desktops	198
Events	24 1
System Health	55 2 0 0

Pools

Folder: All

ID	Display Name	Type	Source	User Assi...	vCenter Server	Entitled	Enabled	Se
CzechRepublic	CZ	Automated Pool	vCenter (linked clone)	Dedicated	10.5.151.11		✓	0 Remote, 0 Local
Finland	FI	Automated Pool	vCenter (linked clone)	Dedicated	10.5.151.11	✓	✓	0 Remote, 0 Local
FUJI	FUJI	Automated Pool	vCenter (linked clone)	Dedicated	10.5.151.11	✓	✓	0 Remote, 0 Local
Morocco	MA	Automated Pool	vCenter (linked clone)	Dedicated	10.5.151.11	✓	✓	0 Remote, 0 Local

Step 16 One the pool generation is completed you will need to add the entitlements to the pool. This will entitle certain users to login to this pools desktop resources. The figure below shows that the CzechRepublic pool does not have any assigned users entitled. To set up entitlements select the pool by clicking on it and then clicking on the **Entitlements** button.

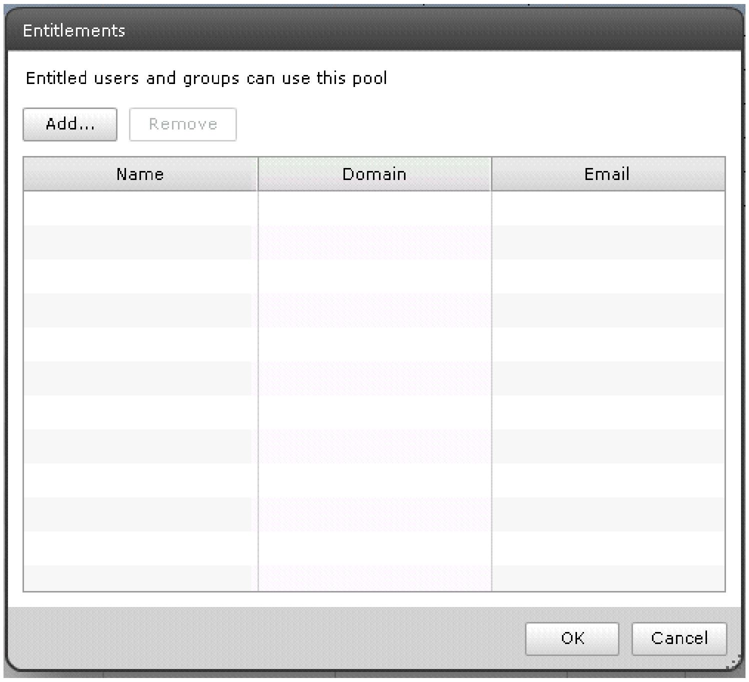

Entitlements

Entitled users and groups can use this pool

Name	Domain	Email

Step 17 After clicking the **Entitlements** button, you see the following window. Click **Add**.

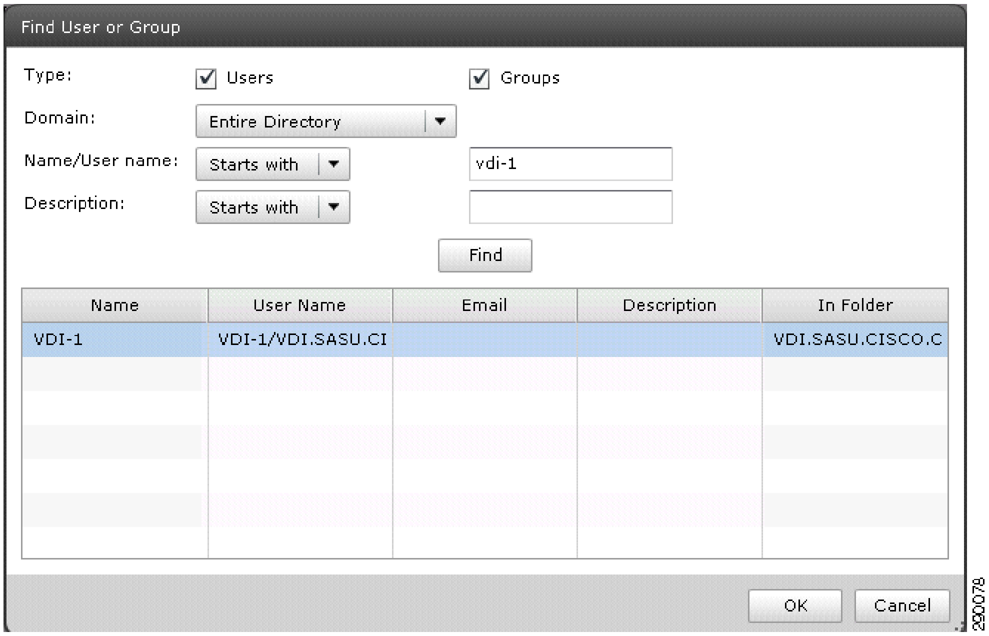

Find User or Group

Type: ☑ Users ☑ Groups

Domain: Entire Directory

Name/User name: Starts with vdi-1

Description: Starts with

Name	User Name	Email	Description	In Folder
VDI-1	VDI-1/VDI.SASU.CI			VDI.SASU.CISCO.C

Step 18 After clicking **Add**, a Find User or Group box will pop up. In the figure below, we search for our VDI user group. Enter the name in the Start with field and click **Find**. After your user or group is found, select it by clicking on it and then select OK to have it added to the entitlements list for the pool.

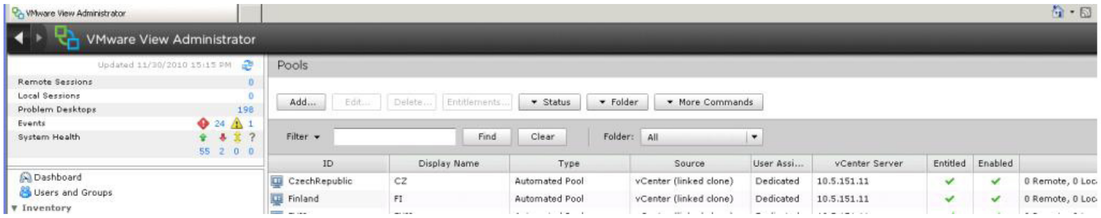

ID	Display Name	Type	Source	User Assi...	vCenter Server	Entitled	Enabled	
CzechRepublic	CZ	Automated Pool	vCenter (linked clone)	Dedicated	10.5.151.11	✓	✓	0 Remote, 0 Loc
Finland	FI	Automated Pool	vCenter (linked clone)	Dedicated	10.5.151.11	✓	✓	0 Remote, 0 Loc

290079

The added group will appear in the Entitlements window. At this point desktops from the pool are accessible. You will now see a green check mark under the Entitled column for this pool.

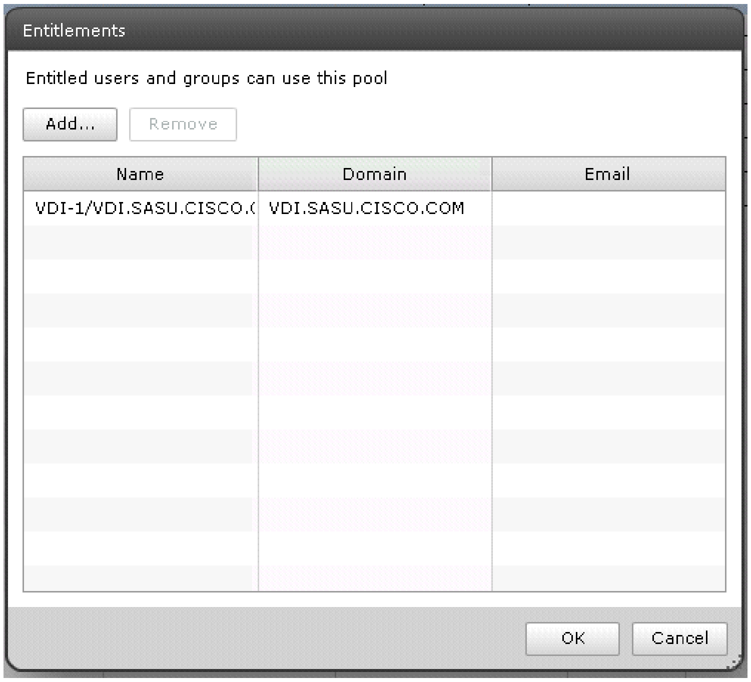

Name	Domain	Email
VDI-1/VDI.SASU.CISCO.(	VDI.SASU.CISCO.COM	

290080

Test Setup and Configurations

Figure 27 Test Environment Used for Server Scalability Testing

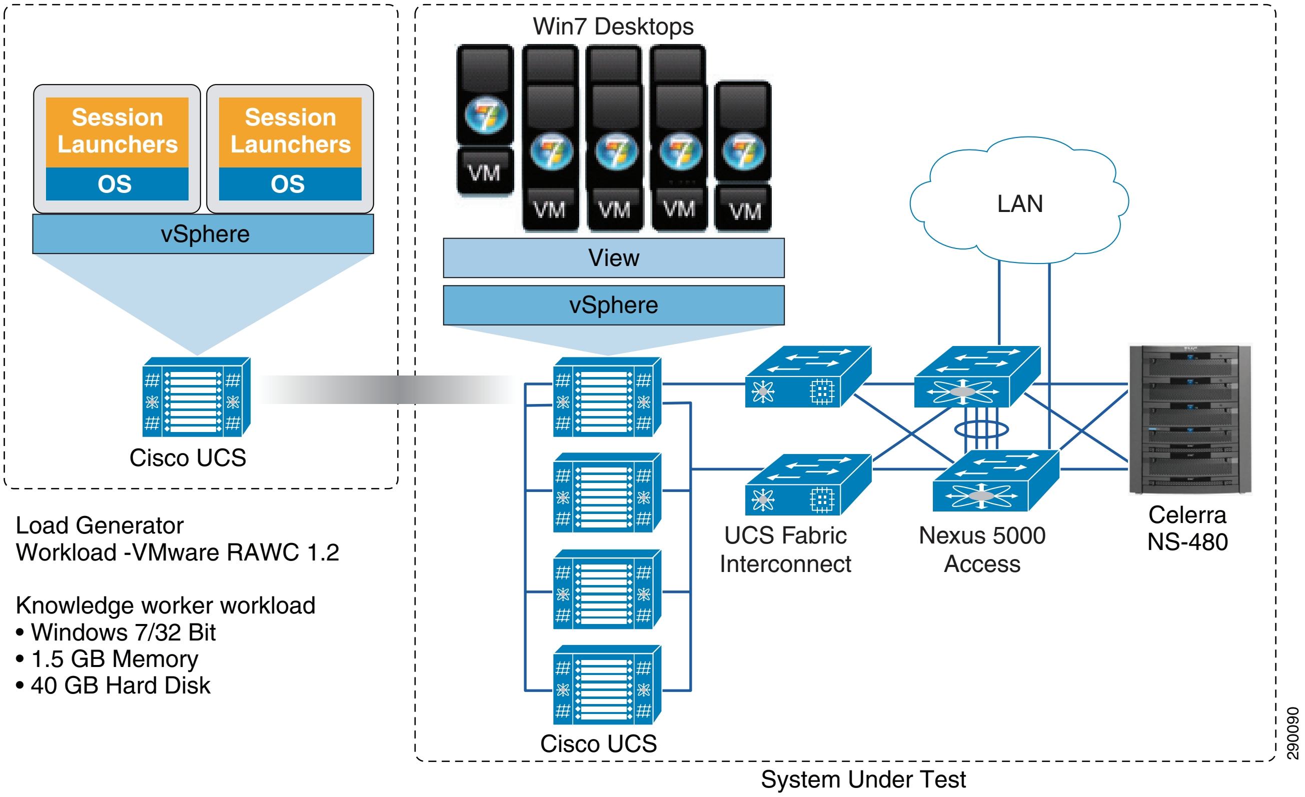

Cisco UCS Test Configuration for Single Server Scalability Test Setup

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 on each B250-M2
- Cisco Nexus 5000 and 7000
- EMC Celerra NS 480 storage array, two controllers, 2 X 10G Optical Ethernet with 12 EFD, and 63 Fibre Channel and 60 SATA II 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 B250-M2
- Cisco Nexus 5000 and 7000
- EMC Celerra NS-480 storage array with Dual 2.8 GHz IV Intel Xeon processors and 8GB of memory per storage processor

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 logfile of the session instance which is started.
- Workload Simulator—Starts running on the Desktop at logon from the View Client and launches applications in realtime. 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 28.

Figure 28 User Profile for RAWC Tool

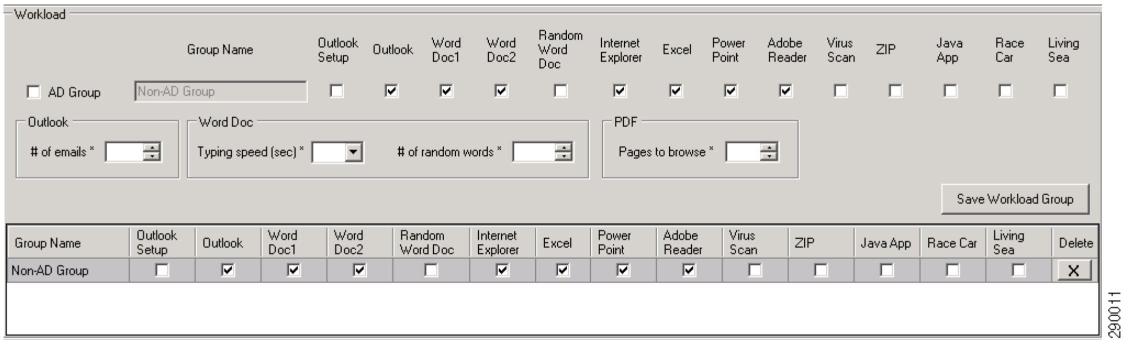

Workload

	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 *

Save Workload Group

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	☐	☑	☑	☑	☐	☑	☑	☑	☑	☐	☐	☐	☐	☐	☒

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.

Success Criteria

Determining a successful workload test is done by evaluating the following criteria:

- Test run completed successfully as determined by:
  - Workload completion and timing written to log files on a share from the RAWC Controller
  - Data collected from remote esxtop generated from a vMA appliance
  - Visually verify that the desktop login groups have launched correctly from each of the session launchers. This is done by monitoring the RDP sessions initiated from the RAWC controller.
- Application open time, measured in RAWC, is less than 2.5 seconds.
- System Ballooning is less than 20% of aggregate memory.
- 80% average CPU load during the workload period for all Desktops under test.

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 the EMC Celerra NS-480 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 CPU utilization of the server hosting each of the desktop configurations was captured. As shown in Figure 29, the processor utilization pattern is consistent across each of the test iterations. CPU utilization consistently exceeded 80% above 160 desktops using the knowledge worker profile workload. More importantly, Application Open Time as measured by RAWC exceeded 2.5 seconds, which becomes a noticeable delay in the user experience. Therefore it was determined that the 160 desktop count would be the maximum allowable number with acceptable levels of user experience.

Figure 29 ESX Host CPU Percentage Core Utilization Time

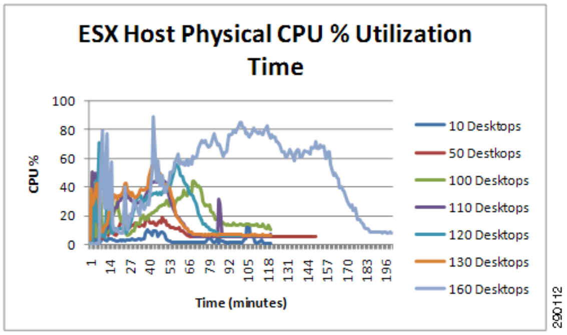

The use of EFD (Solid State) drives in the tiered storage architecture enables scaling the overall Desktop architecture to accommodate more desktops, since the overall CPU utilization is lower when compared to conventional storage architectures. It can be observed from Figure 29 that the average PCPU is approximately 80% with 160 desktops; it was determined that this is the maximum level to handle other processing tasks should the need arise and thus the single UCS blade can host more desktops.

Note The 160 Desktop Run was for three iterations as opposed to two to ensure accuracy. This explains the longer runtime.

Figure 30 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 30 indicates the percentage CPU utilization averaged over all Physical CPUs for the 160 desktop workload. The values observed in Figure 30 confirm the LCPU processors have resources available to manage workload spikes and unforeseen failure conditions.

Figure 30 ESX Host CPU Utilization Percentage

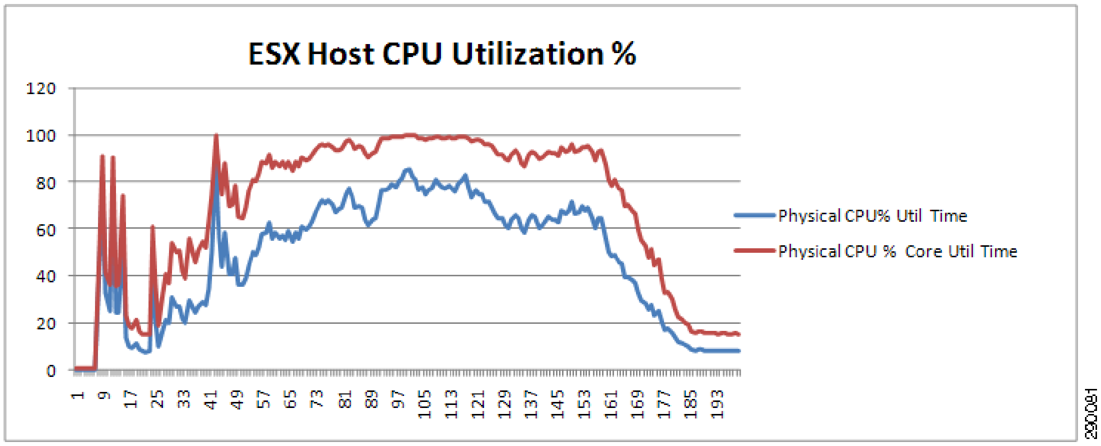

Figure 31 shows that at and beyond 130 desktops, all of the available non-kernel memory is being used by the desktop VMs. vSphere responds to this memory pressure by using the ballooning mechanism to reclaim memory from the VMs, as shown in Figure 32.

Figure 31 Non-Kernal Memory Utilization

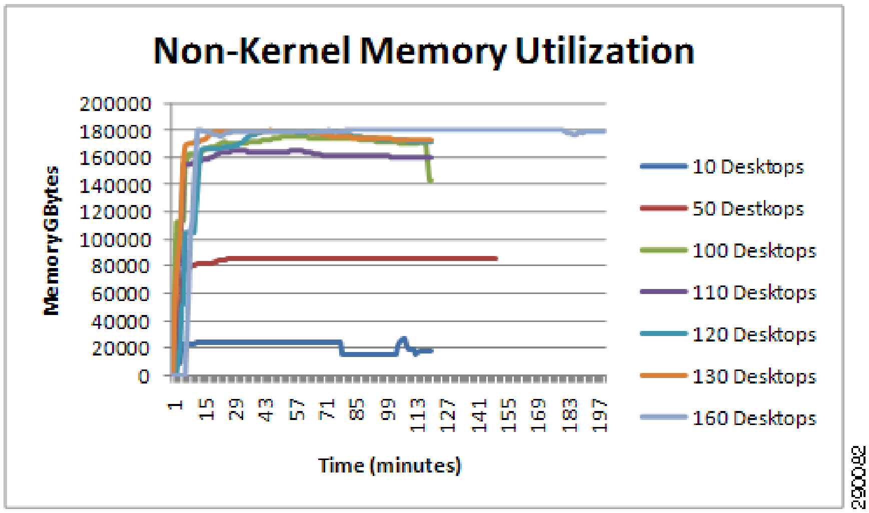

Figure 32 Memory Ballooning in MBytes

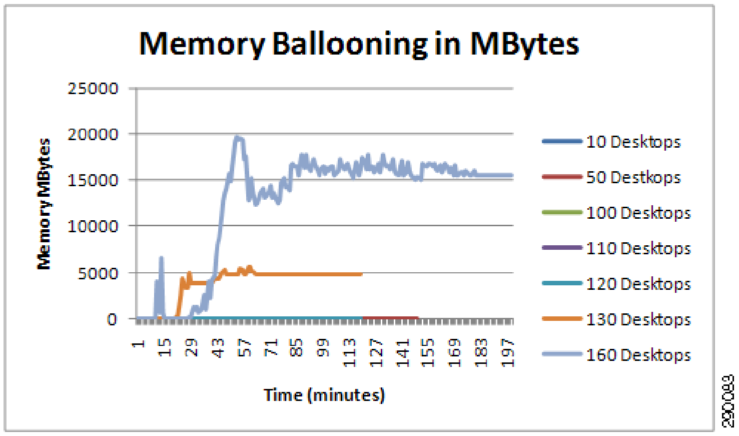

For more information on ballooning, see:
http://www.usenix.org/events/osdi02/tech/full_papers/waldspurger/waldspurger_html/node6.html.

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. As shown in Figure 33, application open times remained consistent throughout the desktop runs. Our validation revealed that application open times began to slightly degrade above the 160 desktop number. Maximum open times were under 2.5 seconds for all runs up to 160 desktops.

Figure 33 Average Application Open Time

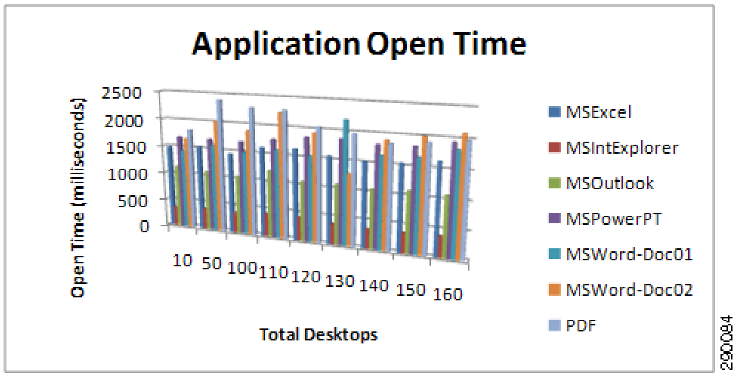

After determining the number of acceptable virtual desktops on a single blade server given the aforementioned memory and CPU utilization requirements, the amount of network traffic traversing the UCS fabric interconnects to the Nexus 5000 was captured. It is important to note these are 10 Gigabit Ethernet connections supporting 160 desktops executing a knowledge worker workload. The captures were taken from the port channel downlinks of each Nexus 5000 to the Fabric Interconnects, which includes all Ethernet-based traffic. From a Fibre Channel utilization perspective, all IO traffic initiated from the UCS Chassis was captured from the FC ports on the Fabric Interconnects.

Figure 34 n5k-vdi-1-port-channel2—VPC to FI-A

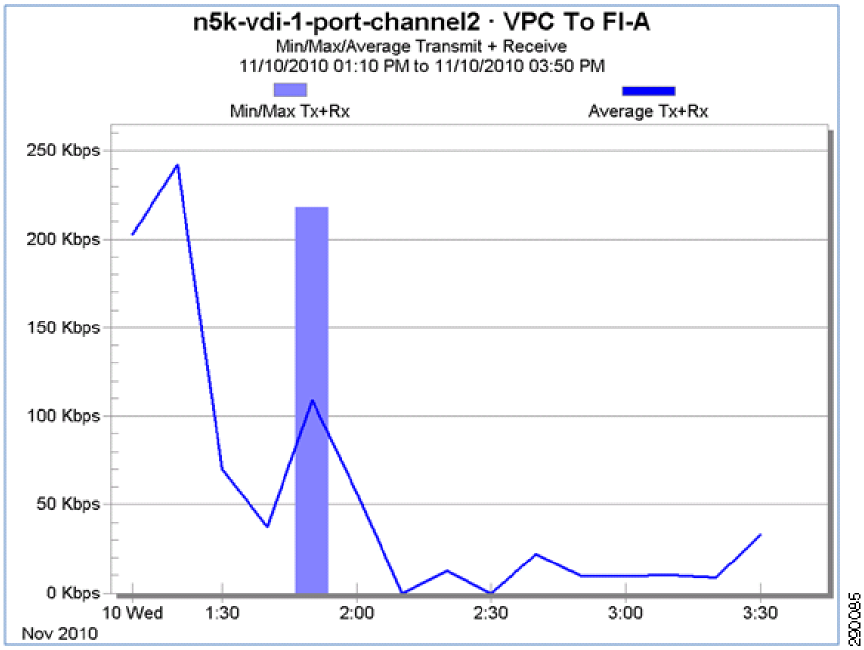

Figure 35 n5k-vdi-1-port-channel3—VPC to FI-B

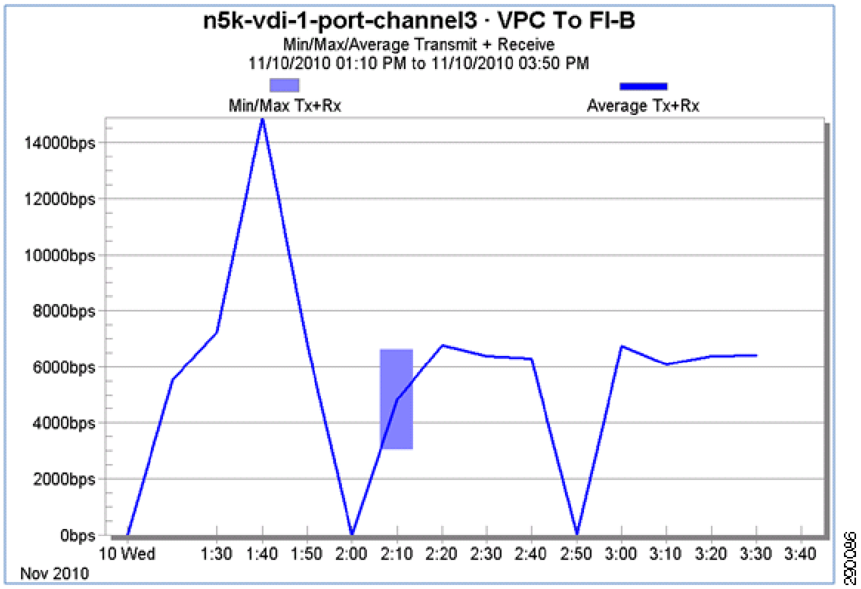

Figure 36 DC1-VDI1-FI-A-fc2/3

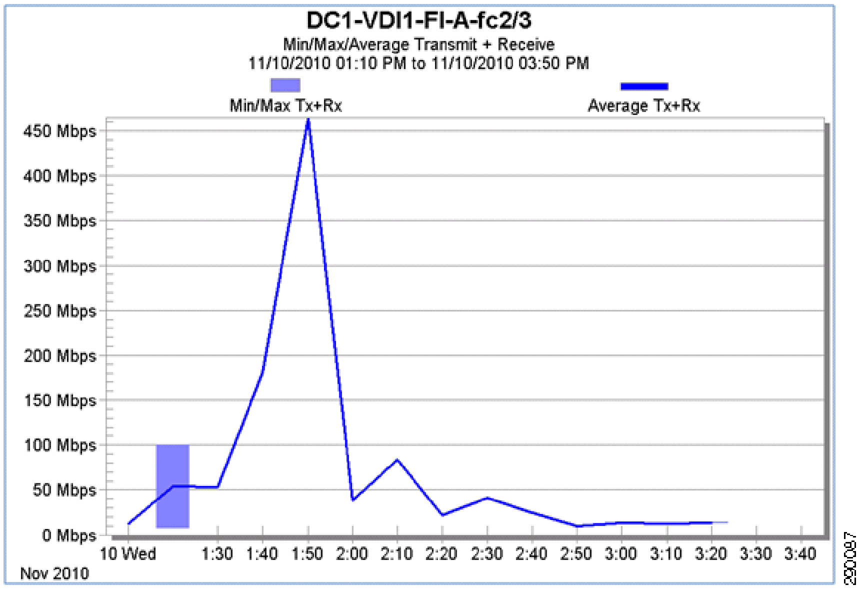

Figure 37 DC1-VDI1-FI-A-fc2/4

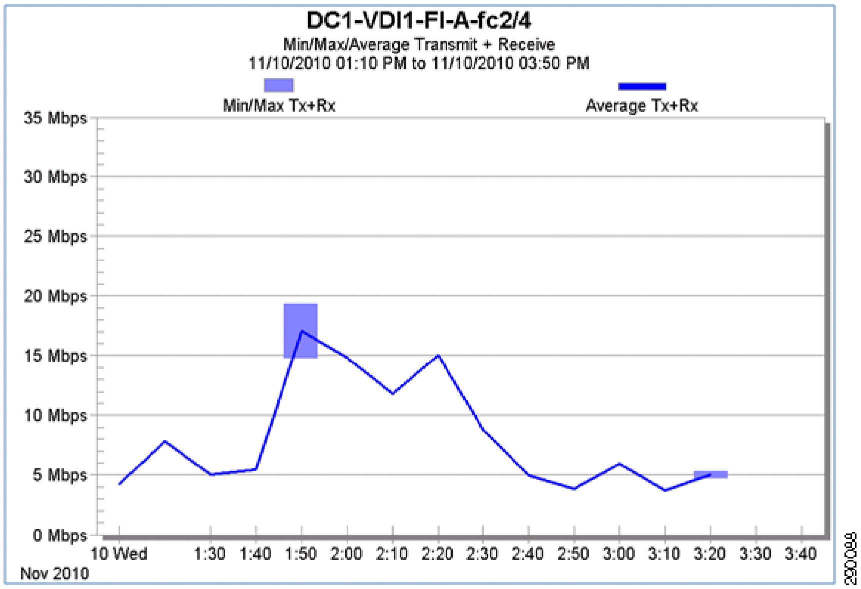

Results for Eight Cisco UCS Blade Server Validation

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.
- Two pools of 385 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 equate to the single blade server charts previously described.

Figure 38 Eight Host CPU Percentage Core Utilization Time

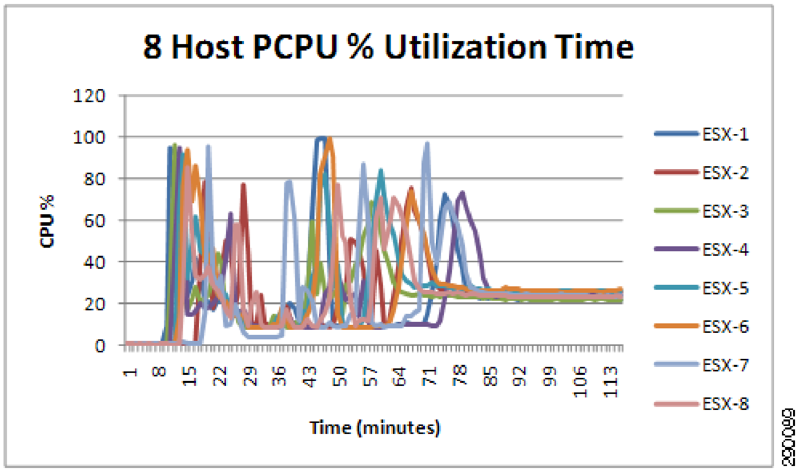

Scaling and Sizing Guidelines

There are many factors to consider when you begin to scale beyond four chassis or 14 servers, which this reference architecture has successfully tested. In this section we give guidance to scale beyond four UCS chassis.

Scalability Considerations and Guidelines

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

- Application open time, measured in RAWC, is less than 2.5 seconds.
- System Ballooning is less than 20% of aggregate memory.
- 80% average CPU load during the workload period for all Desktops under test.

Cisco UCS System Configuration

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

Table 14 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	

UCS uplink density configuration is at the customer's discretion. Extrapolating the tested values, we get the results in Table 15.

Table 15 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	3520	
12	48 Blades	5280	
16	64 Blades	7040	
20	80 Blades	8800	

The backend storage has to be scaled accordingly, based on the IOP considerations described in EMC Unified Storage Solution and Components.

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 EMC Unified Storage Solution and Components.

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]
  - 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.

[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.

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. The majority 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 this will require memory over-subscription. VMware supports a ratio of up to 1.5 memory oversubscription, but depending upon load, this can create memory ballooning and swapping which may lead to performance degradation.

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 = 200 MHZ

If your workload is similar to the knowledge worker workload used, then each desktop will consume approximately 400 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, EMC, and VMware allows you to easily expand your environment when needed.

References

Cisco

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/docs/switches/datacenter/nexus1000/sw/4_0_4_s_v_1_3_b/install/vsm/guide/n1000v_vsm_install.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_0_4_s_v_1_3_b/install/vem/guide/n1000v_vem_install.html

EMC

EMC Celerra NS-480: http://www.emc.com/products/detail/hardware/celerra-ns480.htm

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

Nexus 5000 Configuration

```
n5k-vdi-1# show run
version 4.1(3)N2(1)
feature fcoe
feature npiv
feature telnet
feature lacp
feature vpc
vpc domain 1
  peer-keepalive destination x.x.x.194
role name default-role
  description This is a system defined role and applies to all users.
  rule 5 permit show feature environment
  rule 4 permit show feature hardware
  rule 3 permit show feature module
  rule 2 permit show feature snmp
  rule 1 permit show feature system
username admin password 5 $1$.wPOYrbg$bcOUHottwMShE.BHGw9sy.  role network-admin
no password strength-check
ip host n5k-vdi-1 x.x.x.193
switchname n5k-vdi-1
snmp-server user admin network-admin auth md5 0x1738580ade8f22ca2c065e2d4c5dc5d7 priv
0x1738580ade8f22ca2c065e2d4c5dc5d7
 localizedkey
snmp-server host 172.28.203.171 traps version 2c public  udp-port 2162
snmp-server enable traps entity fru
snmp-server community public group network-admin
vrf context management
  ip route 0.0.0.0/0 x.x.x.1
vlan 1,82,150-152
vsan database
  vsan 100 name "SAN-Fabric-A"
device-alias database
  device-alias name VDI-VXI-ESX1-HBA0 pwwn 20:05:04:00:0a:00:00:0f
  device-alias name VDI-VXI-ESX1-HBA1 pwwn 20:05:08:00:0a:01:00:0f
  device-alias name VDI-VXI-ESX3-HBA0 pwwn 20:05:04:00:0a:00:00:0d
  device-alias name VDI-VXI-ESX4-HBA0 pwwn 20:05:04:00:0a:00:00:0c
  device-alias name VDI-VXI-ESX5-HBA0 pwwn 20:05:04:00:0a:00:00:0b
  device-alias name VDI-VXI-ESX11-HBA0 pwwn 20:05:04:00:0a:00:00:07
  device-alias name VDI-VXI-ESX11-HBA1 pwwn 20:05:08:00:0a:01:00:07
  device-alias name VDI-VXI-ESX12-HBA0 pwwn 20:05:04:00:0a:00:00:08
  device-alias name VDI-VXI-ESX12-HBA1 pwwn 20:05:08:00:0a:01:00:08
  device-alias name VDI-VXI-ESX13-HBA0 pwwn 20:05:04:00:0a:00:00:09
  device-alias name VDI-VXI-ESX15-HBA0 pwwn 20:05:04:00:0a:00:00:0a
  device-alias name VDI-VXI-ESX17-HBA0 pwwn 20:05:04:00:0a:00:00:05
  device-alias name VDI-VXI-ESX17-HBA1 pwwn 20:05:08:00:0a:01:00:05
  device-alias name VDI-VXI-ESX18-HBA0 pwwn 20:05:04:00:0a:00:00:04
  device-alias name VDI-VXI-ESX18-HBA1 pwwn 20:05:08:00:0a:01:00:04
  device-alias name VDI-VXI-ESX33-HBA0 pwwn 20:05:04:00:0a:00:00:03
  device-alias name VDI-VXI-ESX33-HBA1 pwwn 20:05:08:00:0a:01:00:03
  device-alias name VDI-VXI-ESX41-HBA0 pwwn 20:05:04:00:0a:00:00:06
  device-alias name VDI-VXI-ESX41-HBA1 pwwn 20:05:08:00:0a:01:00:06
  device-alias name VDI-VXI-ESX42-HBA0 pwwn 20:05:04:00:0a:00:00:02
  device-alias name VDI-VXI-ESX42-HBA1 pwwn 20:05:08:00:0a:01:00:02
  device-alias name VDI-VXI-ESX43-HBA0 pwwn 20:05:04:00:0a:00:00:01
  device-alias name VDI-VXI-ESX43-HBA1 pwwn 20:05:08:00:0a:01:00:01
  device-alias name VDI-VXI-ESX44-HBA0 pwwn 20:05:04:00:0a:00:00:00
  device-alias name VDI-VXI-ESX44-HBA1 pwwn 20:05:08:00:0a:01:00:00
  device-alias name VDI-VXI-FI-A-FC2-3 pwwn 20:43:00:05:9b:7f:55:00
  device-alias name VDI-VXI-FI-A-FC2-4 pwwn 20:44:00:05:9b:7f:55:00
  device-alias name VDI-VXI-FI-B-FC2-3 pwwn 20:43:00:05:73:a3:ef:40
  device-alias name VDI-VXI-FI-B-FC2-4 pwwn 20:44:00:05:73:a3:ef:40
  device-alias name VDI-VXI-ESX16a-HBA1 pwwn 20:05:08:00:0a:01:00:0e
  device-alias name VDI-VXI-NS-480-SPA1 pwwn 50:06:01:67:46:e0:0b:e1
  device-alias name VDI-VXI-NS-480-SPB1 pwwn 50:06:01:6f:46:e0:0b:e1
  device-alias name CLARIION-NS-480-SPA0 pwwn 50:06:01:66:46:e0:0b:e1
  device-alias name CLARIION-NS-480-SPB0 pwwn 50:06:01:6e:46:e0:0b:e1

device-alias commit

fcdomain fcid database
  vsan 1 wwn 50:06:01:6e:46:e0:0b:e1 fcid 0x4f00ef dynamic
!            [CLARIION-NS-480-SPB0]
  vsan 1 wwn 20:43:00:05:9b:7f:55:00 fcid 0x4f0000 dynamic
!            [VDI-VXI-FI-A-FC2-3]
  vsan 1 wwn 50:06:01:66:46:e0:0b:e1 fcid 0x4f02ef dynamic
!            [CLARIION-NS-480-SPA0]
  vsan 1 wwn 20:44:00:05:9b:7f:55:00 fcid 0x4f0001 dynamic
!            [VDI-VXI-FI-A-FC2-4]
  vsan 100 wwn 50:06:01:66:46:e0:0b:e1 fcid 0x4400ef dynamic
!              [CLARIION-NS-480-SPA0]
  vsan 100 wwn 50:06:01:6e:46:e0:0b:e1 fcid 0x4401ef dynamic
!              [CLARIION-NS-480-SPB0]
  vsan 100 wwn 20:43:00:05:9b:7f:55:00 fcid 0x440000 dynamic
!              [VDI-VXI-FI-A-FC2-3]
  vsan 100 wwn 20:44:00:05:9b:7f:55:00 fcid 0x440001 dynamic
!              [VDI-VXI-FI-A-FC2-4]
  vsan 100 wwn 20:05:04:00:0a:00:00:0f fcid 0x440002 dynamic
!              [VDI-VXI-ESX1-HBA0]
  vsan 100 wwn 20:05:04:00:0a:00:00:0e fcid 0x440003 dynamic
  vsan 100 wwn 20:05:04:00:0a:00:00:0d fcid 0x440004 dynamic
!              [VDI-VXI-ESX3-HBA0]
  vsan 100 wwn 20:05:04:00:0a:00:00:0a fcid 0x440005 dynamic
!              [VDI-VXI-ESX15-HBA0]
  vsan 100 wwn 20:05:04:00:0a:00:00:0b fcid 0x440006 dynamic
!              [VDI-VXI-ESX5-HBA0]
  vsan 100 wwn 20:05:04:00:0a:00:00:0c fcid 0x440007 dynamic
!              [VDI-VXI-ESX4-HBA0]
  vsan 100 wwn 20:05:04:00:0a:00:00:09 fcid 0x440008 dynamic
!              [VDI-VXI-ESX13-HBA0]
  vsan 100 wwn 20:05:04:00:0a:00:00:04 fcid 0x440009 dynamic
!              [VDI-VXI-ESX18-HBA0]
  vsan 100 wwn 20:05:04:00:0a:00:00:05 fcid 0x44000a dynamic
!              [VDI-VXI-ESX17-HBA0]
  vsan 100 wwn 20:05:04:00:0a:00:00:08 fcid 0x44000b dynamic
!              [VDI-VXI-ESX12-HBA0]
  vsan 100 wwn 20:05:04:00:0a:00:00:07 fcid 0x44000c dynamic
!              [VDI-VXI-ESX11-HBA0]
  vsan 100 wwn 20:05:04:00:0a:00:00:06 fcid 0x44000d dynamic
!              [VDI-VXI-ESX41-HBA0]
  vsan 100 wwn 20:05:04:00:0a:00:00:02 fcid 0x44000e dynamic
!              [VDI-VXI-ESX42-HBA0]
  vsan 100 wwn 20:05:04:00:0a:00:00:01 fcid 0x44000f dynamic
!              [VDI-VXI-ESX43-HBA0]
  vsan 100 wwn 20:05:04:00:0a:00:00:00 fcid 0x440010 dynamic
!              [VDI-VXI-ESX44-HBA0]
  vsan 100 wwn 20:05:04:00:0a:00:00:03 fcid 0x440011 dynamic
!              [VDI-VXI-ESX33-HBA0]


interface port-channel1
```

```
    description To n5k-vdi-2
    switchport mode trunk
    vpc peer-link
    spanning-tree port type network
    speed 10000

  interface port-channel2
    description VPC To FI-A
    switchport mode trunk
    switchport trunk allowed vlan 82,150-152
    vpc 200
    speed 10000

  interface port-channel3
    description VPC To FI-B
    switchport mode trunk
    switchport trunk allowed vlan 82,150-152
    vpc 300
    speed 10000

  interface port-channel4
    description VPC To 7K VPC
    switchport mode trunk
    switchport trunk allowed vlan 82,150-152
    vpc 1000
    speed 10000

  interface port-channel100
    switchport mode trunk
    switchport trunk allowed vlan 82,150-152
    speed 10000

  interface port-channel101
    switchport mode trunk
    switchport trunk allowed vlan 82,150-152
    speed 10000
  vsan database
    vsan 100 interface fc2/1
    vsan 100 interface fc2/2
    vsan 100 interface fc2/3
    vsan 100 interface fc2/4

  interface fc2/1
    no shutdown

  interface fc2/2
    no shutdown

  interface fc2/3
    no shutdown

  interface fc2/4
    no shutdown

  interface fc2/5

  interface fc2/6

  interface fc2/7

  interface fc2/8

  interface fc3/1
```

```
interface fc3/2

interface fc3/3

interface fc3/4

interface fc3/5

interface fc3/6

interface Ethernet1/1
  description To n5k-vdi-2
  switchport mode trunk
  channel-group 1 mode active

interface Ethernet1/2
  description To n5k-vdi-2
  switchport mode trunk
  channel-group 1 mode active

interface Ethernet1/3
  switchport mode trunk
  switchport trunk allowed vlan 82,150-152
  channel-group 4 mode active

interface Ethernet1/4
  switchport mode trunk
  switchport trunk allowed vlan 82,150-152
  channel-group 4 mode active

interface Ethernet1/5

interface Ethernet1/6

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
  switchport mode trunk
  switchport trunk allowed vlan 82,150-152
  channel-group 2 mode active

interface Ethernet1/18
  switchport mode trunk
  switchport trunk allowed vlan 82,150-152
  channel-group 3 mode active
```

```
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
  description To DM2
  switchport mode trunk
  switchport trunk allowed vlan 82,150-152
  channel-group 100 mode active

interface Ethernet1/34
  description To DM2
  switchport mode trunk
  switchport trunk allowed vlan 82,150-152
  channel-group 100 mode active

interface Ethernet1/35
  description To DM4
  switchport mode trunk
  switchport trunk allowed vlan 82,150-152
  channel-group 101 mode active

interface Ethernet1/36
  description To DM4
  switchport mode trunk
  switchport trunk allowed vlan 82,150-152
  channel-group 101 mode active

interface Ethernet1/37

interface Ethernet1/38

interface Ethernet1/39

interface Ethernet1/40

interface mgmt0
  ip address x.x.x.193/24
line console
boot kickstart bootflash:/n5000-uk9-kickstart.4.1.3.N2.1.bin
boot system bootflash:/n5000-uk9.4.1.3.N2.1.bin
cfs eth distribute
interface fc2/1
interface fc2/2
interface fc2/3
interface fc2/4
interface fc2/5
interface fc2/6
interface fc2/7
interface fc2/8
interface fc3/1
interface fc3/2
interface fc3/3
interface fc3/4
interface fc3/5
interface fc3/6
!Full Zone Database Section for vsan 100
zone name vdi-vxi-esx1-hba0-emc-ns480-spa0 vsan 100
    member pwwn 20:05:04:00:0a:00:00:0f
!               [VDI-VXI-ESX1-HBA0]
    member pwwn 50:06:01:66:46:e0:0b:e1
!               [CLARIION-NS-480-SPA0]

zone name vdi-vxi-esx1-hba0-emc-ns480-spb0 vsan 100
    member pwwn 20:05:04:00:0a:00:00:0f
!               [VDI-VXI-ESX1-HBA0]
    member pwwn 50:06:01:6e:46:e0:0b:e1
!               [CLARIION-NS-480-SPB0]

zone name vdi-vxi-esx16a-hba1-emc-ns480-spa0 vsan 100
    member pwwn 20:05:04:00:0a:00:00:0e
    member pwwn 50:06:01:66:46:e0:0b:e1
!               [CLARIION-NS-480-SPA0]

zone name vdi-vxi-esx16-hba1-emc-ns480-spa0 vsan 100
    member pwwn 50:06:01:66:46:e0:0b:e1
!               [CLARIION-NS-480-SPA0]
    member pwwn 20:05:04:00:0a:00:00:00
!               [VDI-VXI-ESX44-HBA0]

zone name vdi-vxi-esx16-hba0-emc-ns480-spb0 vsan 100
    member pwwn 20:05:04:00:0a:00:00:00
!               [VDI-VXI-ESX44-HBA0]
    member pwwn 50:06:01:6e:46:e0:0b:e1
!               [CLARIION-NS-480-SPB0]

zone name vdi-vxi-esx16a-hba0-emc-ns480-spb0 vsan 100
    member pwwn 20:05:04:00:0a:00:00:0e
    member pwwn 50:06:01:6e:46:e0:0b:e1
!               [CLARIION-NS-480-SPB0]

zone name vdi-vxi-esx16a-hba1-emc-ns480-spb0 vsan 100
zone name vdi-vxi-esx3-hba0-emc-ns480-spa0 vsan 100
    member pwwn 20:05:04:00:0a:00:00:0d
!               [VDI-VXI-ESX3-HBA0]
    member pwwn 50:06:01:66:46:e0:0b:e1
!               [CLARIION-NS-480-SPA0]

zone name vdi-vxi-esx3-hba0-emc-ns480-spb0 vsan 100
    member pwwn 20:05:04:00:0a:00:00:0d
!               [VDI-VXI-ESX3-HBA0]
    member pwwn 50:06:01:6e:46:e0:0b:e1
!               [CLARIION-NS-480-SPB0]
```

```
zone name vdi-vxi-esx15-hba0-emcns480-spa0 vsan 100
    member pwwn 20:05:04:00:0a:00:00:0a
!               [VDI-VXI-ESX15-HBA0]
    member pwwn 50:06:01:66:46:e0:0b:e1
!               [CLARIION-NS-480-SPA0]

zone name vdi-vxi-esx15-hba0-emcns480-spb0 vsan 100
    member pwwn 20:05:04:00:0a:00:00:0a
!               [VDI-VXI-ESX15-HBA0]
    member pwwn 50:06:01:6e:46:e0:0b:e1
!               [CLARIION-NS-480-SPB0]

zone name vdi-vxi-esx5-hba0-emc-ns480-spa0 vsan 100
    member pwwn 20:05:04:00:0a:00:00:0b
!               [VDI-VXI-ESX5-HBA0]
    member pwwn 50:06:01:66:46:e0:0b:e1
!               [CLARIION-NS-480-SPA0]

zone name vdi-vxi-esx4-hba0-emc-ns480-spa0 vsan 100
    member pwwn 20:05:04:00:0a:00:00:0c
!               [VDI-VXI-ESX4-HBA0]
    member pwwn 50:06:01:66:46:e0:0b:e1
!               [CLARIION-NS-480-SPA0]

zone name vdi-vxi-esx4-hba0-emc-ns480-spb0 vsan 100
    member pwwn 20:05:04:00:0a:00:00:0c
!               [VDI-VXI-ESX4-HBA0]
    member pwwn 50:06:01:6e:46:e0:0b:e1
!               [CLARIION-NS-480-SPB0]

zone name vdi-vxi-esx5-hba0-emc-ns480-spb0 vsan 100
    member pwwn 20:05:04:00:0a:00:00:0b
!               [VDI-VXI-ESX5-HBA0]
    member pwwn 50:06:01:6e:46:e0:0b:e1
!               [CLARIION-NS-480-SPB0]

zone name vdi-vxi-esx13-hba0-emc-ns480-spa0 vsan 100
    member pwwn 20:05:04:00:0a:00:00:09
!               [VDI-VXI-ESX13-HBA0]
    member pwwn 50:06:01:66:46:e0:0b:e1
!               [CLARIION-NS-480-SPA0]

zone name vdi-vxi-esx17-hba0-emc-ns480-spa0 vsan 100
    member pwwn 20:05:04:00:0a:00:00:05
!               [VDI-VXI-ESX17-HBA0]
    member pwwn 50:06:01:66:46:e0:0b:e1
!               [CLARIION-NS-480-SPA0]

zone name vdi-vxi-esx18-hba0-emc-ns480-spa0 vsan 100
    member pwwn 20:05:04:00:0a:00:00:04
!               [VDI-VXI-ESX18-HBA0]
    member pwwn 50:06:01:66:46:e0:0b:e1
!               [CLARIION-NS-480-SPA0]

zone name vdi-vxi-esx12-hba0-emc-ns480-spa0 vsan 100
    member pwwn 20:05:04:00:0a:00:00:08
!               [VDI-VXI-ESX12-HBA0]
    member pwwn 50:06:01:66:46:e0:0b:e1
!               [CLARIION-NS-480-SPA0]

zone name vdi-vxi-esx12-hba0-emc-ns480-spb0 vsan 100
    member pwwn 20:05:04:00:0a:00:00:08
!               [VDI-VXI-ESX12-HBA0]
    member pwwn 50:06:01:6e:46:e0:0b:e1
```

```
!               [CLARIION-NS-480-SPB0]

zone name vdi-vxi-esx11-hba0-emc-ns480-spa0 vsan 100
    member pwwn 20:05:04:00:0a:00:00:07
!               [VDI-VXI-ESX11-HBA0]
    member pwwn 50:06:01:66:46:e0:0b:e1
!               [CLARIION-NS-480-SPA0]

zone name vdi-vxi-esx11-hba1-emc-ns480-spb0 vsan 100
    member pwwn 20:05:04:00:0a:00:00:07
!               [VDI-VXI-ESX11-HBA0]
    member pwwn 50:06:01:6e:46:e0:0b:e1
!               [CLARIION-NS-480-SPB0]

zone name vdi-vxi-esx17-hba1-emc-ns480-spa0 vsan 100
    member pwwn 20:05:08:00:0a:01:00:05
!               [VDI-VXI-ESX17-HBA1]
    member pwwn 50:06:01:66:46:e0:0b:e1
!               [CLARIION-NS-480-SPA0]

zone name vdi-vxi-esx18-hba1-emc-ns480-spa0 vsan 100
    member pwwn 20:05:08:00:0a:01:00:04
!               [VDI-VXI-ESX18-HBA1]
    member pwwn 50:06:01:66:46:e0:0b:e1
!               [CLARIION-NS-480-SPA0]

zone name vdi-vxi-esx17-hba0-emc-ns480-spb0 vsan 100
    member pwwn 20:05:04:00:0a:00:00:05
!               [VDI-VXI-ESX17-HBA0]
    member pwwn 50:06:01:6e:46:e0:0b:e1
!               [CLARIION-NS-480-SPB0]

zone name vdi-vxi-esx18-hba0-emc-ns480-spb0 vsan 100
    member pwwn 20:05:04:00:0a:00:00:04
!               [VDI-VXI-ESX18-HBA0]
    member pwwn 50:06:01:6e:46:e0:0b:e1
!               [CLARIION-NS-480-SPB0]

zone name vdi-vxi-esx17-hba1-emc-ns480-spb0 vsan 100
    member pwwn 20:05:08:00:0a:01:00:05
!               [VDI-VXI-ESX17-HBA1]
    member pwwn 50:06:01:6e:46:e0:0b:e1
!               [CLARIION-NS-480-SPB0]

zone name vdi-vxi-esx18-hba1-emc-ns480-spb0 vsan 100
    member pwwn 20:05:08:00:0a:01:00:04
!               [VDI-VXI-ESX18-HBA1]
    member pwwn 50:06:01:6e:46:e0:0b:e1
!               [CLARIION-NS-480-SPB0]

zone name vdi-vxi-esx41-hba0-emc-ns480-spa0 vsan 100
    member pwwn 20:05:04:00:0a:00:00:06
!               [VDI-VXI-ESX41-HBA0]
    member pwwn 50:06:01:66:46:e0:0b:e1
!               [CLARIION-NS-480-SPA0]

zone name vdi-vxi-esx41-hba0-emc-ns480-spb0 vsan 100
    member pwwn 20:05:04:00:0a:00:00:06
!               [VDI-VXI-ESX41-HBA0]
    member pwwn 50:06:01:6e:46:e0:0b:e1
!               [CLARIION-NS-480-SPB0]

zone name vdi-vxi-esx42-hba0-emc-ns480-spa0 vsan 100
    member pwwn 20:05:04:00:0a:00:00:02
```

```
!                [VDI-VXI-ESX42-HBA0]
    member pwwn 50:06:01:66:46:e0:0b:e1
!                [CLARIION-NS-480-SPA0]

zone name vdi-vxi-esx42-hba0-emc-ns480-spb0 vsan 100
    member pwwn 20:05:04:00:0a:00:00:02
!                [VDI-VXI-ESX42-HBA0]
    member pwwn 50:06:01:6e:46:e0:0b:e1
!                [CLARIION-NS-480-SPB0]

zone name vdi-vxi-esx43-hba0-emc-ns480-spa0 vsan 100
    member pwwn 20:05:04:00:0a:00:00:01
!                [VDI-VXI-ESX43-HBA0]
    member pwwn 50:06:01:66:46:e0:0b:e1
!                [CLARIION-NS-480-SPA0]

zone name vdi-vxi-esx43-hba0-emc-ns480-spb0 vsan 100
    member pwwn 20:05:04:00:0a:00:00:01
!                [VDI-VXI-ESX43-HBA0]
    member pwwn 50:06:01:6e:46:e0:0b:e1
!                [CLARIION-NS-480-SPB0]

zone name vdi-vxi-esx44-hba0-emc-ns480-spa0 vsan 100
    member pwwn 20:05:04:00:0a:00:00:00
!                [VDI-VXI-ESX44-HBA0]
    member pwwn 50:06:01:66:46:e0:0b:e1
!                [CLARIION-NS-480-SPA0]

zone name vdi-vxi-esx44-hba0-emc-ns480-spb0 vsan 100
    member pwwn 20:05:04:00:0a:00:00:00
!                [VDI-VXI-ESX44-HBA0]
    member pwwn 50:06:01:6e:46:e0:0b:e1
!                [CLARIION-NS-480-SPB0]

zone name vdi-vxi-esx13-hba0-emc-ns480-spb0 vsan 100
    member pwwn 20:05:04:00:0a:00:00:09
!                [VDI-VXI-ESX13-HBA0]
    member pwwn 50:06:01:6e:46:e0:0b:e1
!                [CLARIION-NS-480-SPB0]

zone name vdi-vxi-esx33-hba0-emc-ns480-spa0 vsan 100
    member pwwn 20:05:04:00:0a:00:00:03
!                [VDI-VXI-ESX33-HBA0]
    member pwwn 50:06:01:66:46:e0:0b:e1
!                [CLARIION-NS-480-SPA0]

zone name vdi-vxi-esx33-hba0-emc-ns480-spb0 vsan 100
    member pwwn 20:05:04:00:0a:00:00:03
!                [VDI-VXI-ESX33-HBA0]
    member pwwn 50:06:01:6e:46:e0:0b:e1
!                [CLARIION-NS-480-SPB0]

zoneset name VDI-VXI-SAN-A vsan 100
    member vdi-vxi-esx1-hba0-emc-ns480-spa0
    member vdi-vxi-esx1-hba0-emc-ns480-spb0
    member vdi-vxi-esx16a-hba1-emc-ns480-spa0
    member vdi-vxi-esx16a-hba0-emc-ns480-spb0
    member vdi-vxi-esx3-hba0-emc-ns480-spa0
    member vdi-vxi-esx3-hba0-emc-ns480-spb0
    member vdi-vxi-esx15-hba0-emcns480-spa0
    member vdi-vxi-esx15-hba0-emcns480-spb0
    member vdi-vxi-esx5-hba0-emc-ns480-spa0
    member vdi-vxi-esx4-hba0-emc-ns480-spa0
    member vdi-vxi-esx4-hba0-emc-ns480-spb0
    member vdi-vxi-esx5-hba0-emc-ns480-spb0
    member vdi-vxi-esx13-hba0-emc-ns480-spa0
    member vdi-vxi-esx17-hba0-emc-ns480-spa0
    member vdi-vxi-esx18-hba0-emc-ns480-spa0
    member vdi-vxi-esx12-hba0-emc-ns480-spa0
    member vdi-vxi-esx12-hba0-emc-ns480-spb0
    member vdi-vxi-esx11-hba0-emc-ns480-spa0
    member vdi-vxi-esx11-hba1-emc-ns480-spb0
    member vdi-vxi-esx17-hba1-emc-ns480-spa0
    member vdi-vxi-esx18-hba1-emc-ns480-spa0
    member vdi-vxi-esx17-hba0-emc-ns480-spb0
    member vdi-vxi-esx18-hba0-emc-ns480-spb0
    member vdi-vxi-esx17-hba1-emc-ns480-spb0
    member vdi-vxi-esx18-hba1-emc-ns480-spb0
    member vdi-vxi-esx41-hba0-emc-ns480-spa0
    member vdi-vxi-esx41-hba0-emc-ns480-spb0
    member vdi-vxi-esx42-hba0-emc-ns480-spa0
    member vdi-vxi-esx42-hba0-emc-ns480-spb0
    member vdi-vxi-esx43-hba0-emc-ns480-spa0
    member vdi-vxi-esx43-hba0-emc-ns480-spb0
    member vdi-vxi-esx44-hba0-emc-ns480-spa0
    member vdi-vxi-esx44-hba0-emc-ns480-spb0
    member vdi-vxi-esx33-hba0-emc-ns480-spa0
    member vdi-vxi-esx33-hba0-emc-ns480-spb0

zoneset activate name VDI-VXI-SAN-A vsan 100

n5k-vdi-1#
```