



Driving Efficiency. Control. Choice

Scalability Study for Deploying VMware View on Cisco UCS and EMC Symmetrix V-Max Systems

Introduction

Goal of This Document

This document reports the results of a study evaluating the scalability of the VMware View virtual desktop environment on a Cisco Unified Computing System (UCS) connected to EMC Symmetrix V-Max Series arrays with Enginuity via a Cisco MDS director class SAN Switch. We also provide guidance and sizing recommendation for customer deployments for Cisco UCS Server, EMC Symmetrix V-Max Storage, and VMware View.

Audience

The target audience for this document includes sales engineers in the field, field consultants, advanced services specialists, and customers who want to deploy VMware View on UCS connected to enterprise class Symmetrix V-Max storage arrays. The document also explores potential business benefits of interest to senior executives.

Objectives

This document is intended to demonstrate the scalability of an enterprise class VMware View virtual desktop environment on Cisco UCS servers connected to EMC Symmetrix V-Max storage systems. The document also provides guidelines for deploying a large scale VMware View solution using innovative products such as the Cisco UCS server and EMC Symmetrix V-Max storage systems. This paper also highlights the collaborative efforts of three partner companies—EMC, VMware, and Cisco—working together on a common goal of providing proven technology to customers.

Summary of Findings

This study accomplished the objective of scaling a large number of desktops, 160 in all, on a single UCS blade server connected to a Symmetrix V-Max storage system. A fully-populated UCS chassis with eight half slot UCS blades and with 96 GB of memory can run up to 1,280 desktops (160 desktops per blade). Furthermore, the utilization rates seen on the system back plane clearly showed that the linear scaling seen on a single chassis can be extended to include multiple chassis. Given the fact that UCS scales to 40 chassis or 320 UCS blade servers with a Nexus 6140 XP fabric interconnect, we clearly see a linear scalability of desktops with a large multichassis configuration. However, the larger UCS configurations would require an appropriately scaled up configuration of the Symmetrix V-Max storage array to support the larger deployment. The Symmetrix V-Max storage array, as discussed in [EMC Symmetrix V-Max Storage System](#), can scale to configurations much larger than that used for the current study.

This study also determined that a 128 virtual desktop environment can be deployed on a single UCS blade with 48 GB of memory. The joint study between Cisco, EMC, and VMware clearly achieved the objective of scaling the UCS platform and determining the safe limits for VMware View deployments using VMware View and an EMC Symmetrix V-Max storage system. The linear scaling and capabilities exhibited by the UCS servers connected to the Symmetrix V-Max storage system makes it an ideal platform for deploying virtual desktop environment using VMware View for the enterprise.

Summary of UCS Sizing Recommendations

The study demonstrated that a UCS blade with 48 GB of memory (6, quad rank, 8 GB DIMMs) easily supported a VMware View virtual desktop environment consisting of 128 users with a “knowledge worker” profile. The configuration also had enough resources to withstand the boot storms that might occur if an unplanned event resulted in the reboot of a VMware ESX Server. IT managers should, therefore, plan on no more than 128 desktops on a single UCS blade configured with 48 GB of memory.

The 48 GB UCS blade configuration, however, is not expected to fully utilize the compute capabilities provided by the Xeon 5500 processors. The capabilities of the processor can be fully utilized by increasing the memory on each blade to 96 GB via the use of 12, quad rank, 8 GB DIMMs. In this configuration, one can easily achieve the maximum number of virtual machines than can be executed by the VMware ESX Server version 3.5 kernel. This limit, 170, described in the VMware knowledge base article KB1006393, allows full utilization of the system resources during the steady state execution of a knowledge worker workload. However, to support the possibility of additional workload during workload ramp-up and unexpected spikes, users should plan on deploying no more than 160 desktops on a single fully-populated UCS blade. It is important to note that in a VMware View deployment that utilizes VMware HA, the maximum number of virtual desktops on a single, fully-populated blade should not exceed $\text{TRUNC}(170 \times (n - 1)/n)$, where n is the number of UCS blades in a single cluster. Adhering to the recommended maximum in a VMware View deployment that includes VMware HA ensures that the theoretical maximum discussed earlier is not exceeded if one of the VMware ESX Servers experiences an unplanned outage.

The resources provided by the system backplane of the UCS system can easily support the workload generated by a fully-populated UCS chassis. Therefore, for larger deployments that require multiple UCS blades, users can linearly scale their environment and execute a maximum of 1,280 virtual desktops on a single UCS chassis. Furthermore, given the fact that UCS scales to 40 chassis or 320 UCS blade servers with a single Nexus 6140 XP fabric interconnect, we clearly see a linear scalability of desktops with a large multichassis configuration.

Summary of Symmetrix V-Max Storage Array Recommendations

The VMware View environment deployed on Cisco UCS blades should be connected to a minimum of four front end Fibre Channel ports on the Symmetrix V-Max storage system. The I/O workload generated from a single, fully-populated UCS chassis with virtual desktops executing a workload similar to a knowledge worker can be handled by four ports on the Symmetrix V-Max system. For maximum availability, the four ports should be equally spread across two different directors on the Symmetrix V-Max storage array. In Symmetrix V-Max storage arrays that have multiple storage engines, the two directors should reside on different storage engines. The number of physical disks required to support a VMware View deployment depends on a number of factors. A 1,280 desktop deployment anticipated to execute a workload similar to that generated by a knowledge worker can be easily handled by approximately 100, 15,000 RPM physical disks configured with RAID-5 protection. This configuration is expected to have enough resources to support the additional workload generated by unplanned events.

The storage from the Symmetrix V-Max array should be presented to the UCS blades executing VMware View desktops as thin devices using Symmetrix Virtual Provisioning with the data across all physical disks in the thin pool. The thin devices can either be fully allocated during the binding process or can be configured to allocate storage on demand. The fully allocated thin devices are advisable in deployments that have well known virtual desktop workloads which require all of the resources professed by the physical disks. The on-demand allocation solution allows customers to optimally use the storage resources. Other applications can utilize the physical media in case the peak workload generated by the VMware View desktops do not fully utilize the resources of the physical disks.

VMware View

VMware View is a robust desktop virtualization solution that allows IT organizations to reap the benefits of traditional server-based computing without the challenges that often accompany server-based solutions. By leveraging the benefits and advantages of VMware View, IT organizations can take the first steps toward transitioning away from the era of distributed PC computing towards cloud computing and the delivery of user desktops as a service.

VMware View 3 is a universal client solution built on VMware's industry-leading virtualization platform that lets IT administrators manage operating systems, hardware, applications, and users independent of each other wherever they may reside. VMware View streamlines desktop and application management, thus reducing costs while increasing data security through centralization. This in turn results in greater user flexibility and IT control.

VMware View centralizes the access control and confines the corporate data to the data center with encryption and provides advanced security measures. VMware View thus enables customers to extend the value of VMware Infrastructure and desktop virtualization environments to encompass not only desktops in the data center, but also applications and the delivery of these environments securely to remote clients, online or off, anywhere.

The VMware View segregation of operating environments from applications provides an inherent layer of security. IT professionals can consolidate and maintain a small number of core desktop images or pools that everyone can use. Implementing a standardized solution drives down the total cost of ownership (TCO) and helps minimize the complexity and unpredictability of a customized solution. By implementing a standardized solution that integrates with established IT processes and procedures requiring minimal change, IT organizations can build the foundation on which they can easily and effectively adapt to changing user requirements and technology trends.

Cisco Unified Computing System

The Cisco Unified Computing System (UCS) is a next-generation data center platform that unites compute, network, and storage access. The platform, optimized for virtual environments, is designed within open industry standard technologies and aims to reduce 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, multichassis platform in which all resources participate in a unified management domain.

UCS Components

The main system components include:

- UCS 5108 blade server chassis that fits on a standard rack and is 6RU high. Each UCS chassis can hold either eight half slot or four full slot blade servers, two redundant fabric extenders, eight cooling fans, and four power supply units. The cooling fans and power supply are hot swappable and redundant. The chassis requires only two power supplies for normal operation; the additional power supplies are for redundancy. The highly-efficient (in excess of 90%) power supplies, in conjunction with the simple chassis design that incorporates front to back cooling, makes the UCS system very reliable and energy efficient.
- UCS B-Series blade servers are based on an entirely new class of computing system that incorporates blade servers based on Intel Xeon 5500 Series processors. The blade servers offer patented Cisco Extended Memory Technology to support applications with large data sets and allow more virtual machines per server. The blade servers comes in two form factors:
 - UCS B200-M1—Half slot two socket blade server
 - UCS B250-M1—Full slot Extended Memory Server

[Table 1](#) details the features of the blade server.

Table 1 **UCS B-Series Blade Server Features**

UCS B200-M1	UCS B250-M1
Eight blades per chassis	Four blades per chassis
12 DIMM slots (up to 96 GB)	48 DIMM slots (up to 384 GB)
One dual port Converged Network Adapter (CNA)	Two dual port Converged Network Adapter (CNA)

Common system features:

- Two Intel Xeon 5500 series processors (quad cores)
- Two optional SAS hard drives
- Hot pluggable blades and hard disk drive support
- Blade service processor
- Stateless blade design
- 10 Gb/s CNA and 10 Gb/s Ethernet adapter options

Cisco Extended Memory Technology

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 nonvirtualized 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. 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; eight 1-GB DIMMS are typically less expensive than one 8-GB DIMM.

I/O Adapters

The blade server has various Converged Network Adapters (CNA) options. Customers can select the most appropriate option based on their needs. The following options are offered:

- Efficient, High-Performance Ethernet with the Cisco UCS 82598KR-CI 10 Gigabit Ethernet Adapter

The Cisco UCS 82598KR-CI 10 Gigabit Ethernet Adapter is designed to deliver efficient, high-performance Ethernet connectivity. This adapter uses Intel silicon to present two 10 Gigabit Ethernet NICs to the peripheral component interconnect (PCI) device tree, with each NIC connected to one of the two fabric extender slots on the chassis. Like all the mezzanine cards available for the Cisco Unified Computing System, this card supports the Cisco DCE features needed to manage multiple independent network traffic streams over the same link. The adapter's MAC addresses are just-in-time configured by Cisco UCS Manager and the adapter is designed for:

- Network-intensive workloads, such as Web servers, in which all content is accessed over Network File System (NFS) or iSCSI protocols
 - Environments in which efficiency and performance are important considerations
- QLogic converged network adapter (M71KR-Q) and Emulex converged network adapter (M71KR-E)

For organizations needing compatibility with existing data center practices that rely on Emulex or QLogic Fibre Channel HBAs, the Cisco UCS M71KR-E Emulex and UCS M71KR-Q QLogic Converged Network Adapters provide compatibility with interfaces from Emulex and QLogic, respectively. These CNAs use Intel silicon to present two 10 Gigabit Ethernet NICs and either two Emulex or two QLogic HBAs to the PCI device tree. The operating system sees two NICs and two HBAs and the existence of the unified fabric is completely transparent. A Cisco application-specific integrated circuit (ASIC) multiplexes one Ethernet and one Fibre Channel traffic stream onto each of the two midplane connections to the fabric extender slots. These CNAs are most appropriate for:

- Organizations that want to continue to use Emulex or QLogic drivers in the Cisco Unified Computing System
 - Organizations that want to streamline the qualification process for new Fibre Channel hardware; use of standard HBA silicon allows use of HBA vendor-provided drivers
 - Both traditional physical and virtualized environments
- Cisco UCS VIC M81KR Virtual Interface card

The full benefits of the Cisco Unified Computing System are delivered through the use of Cisco UCS M81KR Virtual Interface Cards in blade servers. This card presents up to 128 virtual interfaces to the PCI device tree in compliance with PCI Express (PCIe) standards.

Unlike other network adapters, in which only the identity of each NIC and HBA is programmable through Cisco UCS Manager, both the type and the identity of each of the Cisco UCS M81KR Virtual Interface Card's virtual NICs are programmable. Any combination of Ethernet NICs and Fibre Channel HBAs and their corresponding MAC addresses and WWNs can be programmed onto the card, making the I/O architecture of the server programmable using a dynamic provisioning model.

This card, combined with service profiles, supports a very powerful and flexible, wire-once environment. The service profile defines virtual NICs when it is applied to the server. For example, a Cisco UCS M81KR Virtual Interface Card could be configured to have four Ethernet NICs supporting network-attached storage (NAS) connections with one service profile and an FCoE card with four Fibre Channel HBAs and six Ethernet NICs when the next service profile is applied.

The programmable nature of the virtual interface card makes it ideal for:

- Service provider and enterprise data centers needing the utmost flexibility in the way they deploy server resources
- Companies needing strong investment protection; one adapter can satisfy both Ethernet and Fibre Channel connectivity requirements and a single I/O interface can be deployed onto systems throughout the data center.
- Fixed server environments in which the capability to rehost or scale applications is needed
- Virtualized environments in which:
 - Each virtual NIC can appear as a physical NIC to virtual machines, eliminating the overhead of the virtual switch
 - Per-virtual machine network policies need to be applied, including security and QoS
 - The I/O performance boost achieved through pass-through switching is important

The CNA reduces the number of adapters, cables, and access layer switches that are required for LAN and SAN connectivity. This significantly reduces capital and operational expenditure, power and cooling, and administrative overheads.

Cisco UCS 2100 Series Fabric Extenders

The Cisco UCS 2104XP Fabric Extender brings the I/O fabric into the blade server chassis and supports up to four 10-Gbps connections between blade servers and the parent fabric interconnect, simplifying diagnostics, cabling, and management. The fabric extender multiplexes and forwards all traffic using a cut-through architecture over one to four 10-Gbps unified fabric connections. All traffic is passed to the parent fabric interconnect, where network profiles are managed efficiently and effectively by the fabric interconnects. Each of up to two fabric extenders per blade server chassis has eight 10GBASE-KR connections to the blade chassis midplane, with one connection to each fabric extender from each of the chassis' eight half slots. This configuration gives each half-width blade server access to each of two 10-Gbps unified fabric connections for high throughput and redundancy.

The benefits of the fabric extender design include:

- **Scalability**—With up to four 10-Gbps uplinks per fabric extender, network connectivity can be scaled to meet increased workload demands simply by configuring more uplinks to carry the additional traffic.
- **High availability**—Chassis configured with two fabric extenders can provide a highly available network environment.
- **Reliability**—The fabric extender manages traffic flow from network adapters through the fabric extender and onto the unified fabric. The fabric extender helps create a lossless fabric from the adapter to the fabric interconnect by dynamically throttling the flow of traffic from network adapters into the network.
- **Manageability**—The fabric extender model extends the access layer without increasing complexity or points of management, freeing administrative staff to focus more on strategic than tactical issues. Because the fabric extender also manages blade chassis components and monitors environmental conditions, fewer points of management are needed and cost is reduced.
- **Virtualization optimization**—The fabric extender supports Cisco VN-Link architecture. Its integration with VN-Link features in other Cisco UCS components, such as the fabric interconnect and network adapters, enables virtualization-related benefits including virtual machine-based policy enforcement, mobility of network properties, better visibility, and easier problem diagnosis in virtualized environments.
- **Investment protection**—The modular nature of the fabric extender allows future development of equivalent modules with different bandwidth or connectivity characteristics, protecting investments in blade server chassis.
- **Cost savings**—The fabric extender technology allows the cost of the unified network to be accrued incrementally, helping reduce costs in times of limited budgets. The alternative is to implement and fund a large, fixed-configuration fabric infrastructure long before the capacity is required.

UCS 6100 XP Series Fabric Interconnect

The UCS 6100 XP fabric interconnect is based on the Nexus 5000 product line. However, unlike the Nexus 5000 products, it provides additional functionality of managing the UCS chassis with the embedded UCS manager. A single 6140 XP switch can support up to 40 chassis or 320 servers with half slot blades or 160 full slot blades.

Some of the salient features provided by the switch are:

- 10 Gigabit Ethernet, FCoE capable, SFP+ ports
- 20 and 40 fixed port versions with expansion slots for additional Fiber Channel and 10 Gigabit Ethernet connectivity

- Up to 1.04 Tb/s of throughput
- Hot pluggable fan and power supplies, with front to back cooling system
- Hardware based support for Cisco VN-Link technology
- Can be configured in a cluster for redundancy and failover capabilities

UCS Manager

Data centers have become complex environments with a proliferation of management points. From a network perspective, the access layer has fragmented, with traditional access layer switches, switches in blade servers, and software switches used in virtualization software all having separate feature sets and management paradigms. Most current blade systems have separate power and environmental management modules, adding cost and management complexity. Ethernet NICs and Fibre Channel HBAs, whether installed in blade systems or rack-mount servers, require configuration and firmware updates. Blade and rack-mount server firmware must be maintained and BIOS settings must be managed for consistency. As a result, data center environments have become more difficult and costly to maintain, while security and performance may be less than desired. Change is the norm in data centers, but the combination of x86 server architectures and the older deployment paradigm makes change difficult:

- In fixed environments in which servers run OS and application software stacks, rehosting software on different servers as needed for scaling and load management is difficult to accomplish. I/O devices and their configuration, network configurations, firmware, and BIOS settings all must be configured manually to move software from one server to another, adding delays and introducing the possibility of errors in the process. Typically, these environments deploy fixed spare servers already configured to meet peak workload needs. Most of the time these servers are either idle or highly underutilized, raising both capital and operating costs.
- Virtual environments inherit all the drawbacks of fixed environments, and more. The fragmentation of the access layer makes it difficult to track virtual machine movement and to apply network policies to virtual machines to protect security, improve visibility, support per-virtual machine QoS, and maintain I/O connectivity. Virtualization offers significant benefits; however, it adds more complexity.

Programmatically Deploying Server Resources

Cisco UCS Manager provides centralized management capabilities, creates a unified management domain, and serves as the central nervous system of the Cisco Unified Computing System. Cisco UCS Manager is embedded device-management software that manages the system from end to end as a single logical entity through an intuitive GUI, CLI, or XML API. Cisco UCS Manager implements role- and policy-based management using service profiles and templates. This construct improves IT productivity and business agility. Now infrastructure can be provisioned in minutes instead of days, shifting IT's focus from maintenance to strategic initiatives.

Dynamic Provisioning with Service Profiles

Cisco Unified Computing System resources are abstract in the sense that their identity, I/O configuration, MAC addresses and WWNs, firmware versions, BIOS boot order, and network attributes (including QoS settings, ACLs, pin groups, and threshold policies) all are programmable using a just-in-time deployment model. The manager stores this identity, connectivity, and configuration information in service profiles that reside on the Cisco UCS 6100 Series Fabric Interconnect. A service

profile can be applied to any resource to provision it with the characteristics required to support a specific software stack. A service profile allows server and network definitions to move within the management domain, enabling flexibility in the use of system resources.

Service profile templates allow different classes of resources to be defined and applied to a number of resources, each with its own unique identities assigned from predetermined pools. The same management techniques apply whether the server is physical or whether it is a virtual machine directly connected to one or more of the 128 virtual devices provided by the Cisco UCS M81KR Virtual Interface Card.

Role-Based Administration and Multi-Tenancy Support

The UCS Manager offers role-based management that helps organizations make more efficient use of their limited administrator resources. Cisco UCS Manager allows organizations to maintain IT disciplines while improving teamwork, collaboration, and overall effectiveness. Server, network, and storage administrators maintain responsibility and accountability for their domain policies within a single integrated management environment. Compute infrastructure can now be provisioned without the time-consuming manual coordination between multiple disciplines previously required. Roles and privileges in the system can be easily modified and new roles quickly created.

Administrators focus on defining policies needed to provision compute infrastructure and network connectivity. Administrators can collaborate on strategic architectural issues and implementation of basic server configuration can now be automated. Cisco UCS Manager supports multi-tenant service providers and enterprise data centers serving internal clients as separate business entities. The system can be logically partitioned and allocated to different clients or customers to administer as their own.

[Table 2](#) summarizes the various components that constitute UCS.

Table 2 UCS System Components


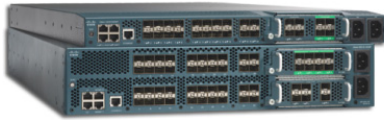




<p>UCS Manager Embedded- manages entire system</p>	 227092
<p>UCS Fabric Interconnect 20 Port 10Gb FCoE 40 Port 10Gb FCoE</p>	 227093
<p>UCS Fabric Extender Remote line card</p>	 227094
<p>UCS Blade Server Chassis Flexible bay configurations</p>	 227095

Table 2 UCS System Components

UCS Blade Server Industry-standard architecture	 A photograph of a UCS Blade Server, showing a rack-mounted unit with multiple bays.
UCS Virtual Adapters Choice of multiple adapters	 A photograph of several UCS Virtual Adapters, which are circuit boards with various components and connectors.

EMC Symmetrix V-Max Storage System

The EMC Symmetrix V-Max Series with Enginuity provides an extensive offering of new features and functionality for the next era of high-availability virtual data centers. With advanced levels of data protection and replication, the Symmetrix V-Max system is at the forefront of enterprise storage area network (SAN) technology. Additionally, the Symmetrix V-Max array has the speed, capacity, and efficiency to transparently optimize service levels without compromising its ability to deliver performance on demand. These capabilities are of greatest value for large virtualized server deployments such as VMware Virtual Data Centers.

Symmetrix V-Max system is EMC's high-end storage array that is purpose-built to deliver infrastructure services within the Next Generation data center. Built for reliability, availability, and scalability, Symmetrix V-Max uses specialized engines, each of which includes two redundant director modules providing parallel access and replicated copies of all critical data.

Symmetrix V-Max's Enginuity operating system provides several advanced features, such as Auto-provisioning Groups for simplification of storage management, Virtual Provisioning for ease of use and improved capacity utilization, and Virtual LUN technology for non-disruptive mobility between storage tiers. All of the industry-leading features for Business Continuation and Disaster Recovery have been the hallmarks of EMC Symmetrix storage arrays for over a decade and continues in the Symmetrix V-Max system. These are further integrated into the VMware Virtual Infrastructure for disaster recovery with EMC's custom Site Recovery Adapter for VMware's Site Recovery Manager.

Combined with the rich capabilities of EMC ControlCenter and EMC's Storage Viewer for vCenter, administrators are provided with end-to-end visibility and control of their virtual data center storage resources and usage. EMC's new PowerPath/VE support for ESX provides optimization of usage on all available paths between virtual machines and the storage they are using, as well as proactive failover management.

Introduction

The Symmetrix V-Max system (see [Figure 1](#)) with the Enginuity operating environment is a new enterprise-class storage array that is built on the strategy of simple, intelligent, modular storage. The array incorporates a new high-performance fabric interconnect designed to meet the performance and scalability demands for enterprise storage within the most demanding virtual data center installations. The storage array seamlessly grows from an entry-level configuration with a single, highly available Symmetrix V-Max Engine and one storage bay into the world's largest storage system with eight engines


and 10 storage bays. The largest supported configuration is shown in [Figure 2](#). When viewing [Figure 2](#), refer to the following list that indicates the range of configurations supported by the Symmetrix V-Max storage array:

- 2-16 director boards
- 48-2,400 disk drives
- Up to 2 PB usable capacity
- Up to 128 Fibre Channel ports
- Up to 64 FICON ports
- Up to 64 Gig-E/iSCSI ports

Figure 1 **Symmetrix V-Max System**



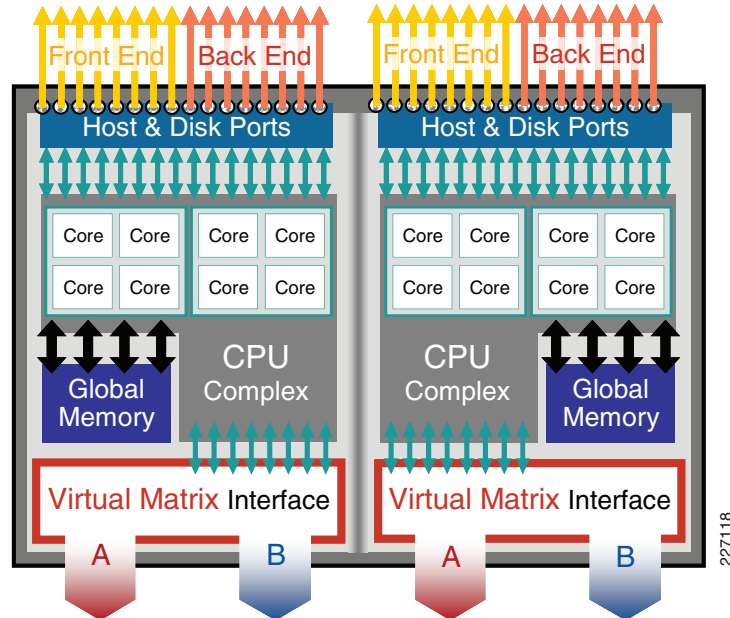
Figure 2 Symmetrix V-Max System Features



The enterprise-class deployments in a modern data center are expected to be always available. The design of the Symmetrix V-Max storage array enables it to meet this stringent requirement. The replicated components which comprise every Symmetrix V-Max configuration assure that no single point of failure can bring the system down. The hardware and software architecture of the Symmetrix V-Max storage array allows capacity and performance upgrades to be performed online with no impact to production applications. In fact, all configuration changes, hardware and software updates, and service procedures are designed to be performed online and non-disruptively. This ensures that customers can consolidate without compromising availability, performance, and functionality, while leveraging true pay-as-you-grow economics for high-growth storage environments.

The Symmetrix V-Max system can include two to 16 directors inside one to eight Symmetrix V-Max Engines. Each Symmetrix V-Max Engine has its own redundant power supplies, cooling fans, SPS Modules, and Environmental Modules. Furthermore, the connectivity between the Symmetrix V-Max array engines provides direct connections from each director to every other director, creating a redundant and high-availability Virtual Matrix. Each Symmetrix V-Max Engine has two directors that can offer up to eight host access ports each, therefore allowing up to 16 host access ports per Symmetrix V-Max Engine.

Figure 3 shows a schematic representation of a single Symmetrix V-Max storage engine.

Figure 3 Symmetrix V-Max Storage Engine


The powerful, high-availability Symmetrix V-Max Engine provides the building block for all Symmetrix V-Max systems. It includes four quad-core Intel Xeon processors, 64-128 GB of Global Memory; 8-16 ports for front-end host access or Symmetrix Remote Data Facility channels using Fibre Channel, FICON, or Gigabit Ethernet; and 16 back-end ports connecting to up to 360 storage devices using 4Gb Fibre Channel SATA or Enterprise Flash Drives.

Each of the two integrated directors in a Symmetrix V-Max Engine has main three parts: the back-end director, the front-end director, and the cache memory module. The back-end director consists of two back-end I/O modules with four logical directors that connect directly into the integrated director. The front-end director consists of two front-end I/O modules with four logical directors that are located in the corresponding I/O annex slots. The front-end I/O modules are connected to the director via the midplane.

The cache memory modules are located within each integrated director, each with eight available memory slots. Memory cards range from 2 to 8 GB, consequently allowing anywhere between 16 to 64 GB per integrated director. For added redundancy, the Symmetrix V-Max system uses mirrored cache, so memory is mirrored across engines in a multi-engine setup. In case of a single-engine configuration, the memory is mirrored inside the engine across the two integrated directors.

EMC's Virtual Matrix interconnection fabric permits the connection of up to 8 Symmetrix V-Max Engines together to scale out total system resources and flexibly adapt to the most demanding virtual data center requirements.

Figure 4 shows a schematic representation of a maximum Symmetrix V-Max configuration.

Figure 4 *Fully Configured Symmetrix V-Max Storage System*

227112

The Symmetrix V-Max system supports the drive types in [Table 3](#).

Table 3 *Symmetrix V-Max System Supported Drive Types*

Drive type	Rotational speed	Capacity
4 Gb/s FC	15k	146 GB, 300 GB, 450 GB
4 Gb/s FC	10k	400 GB
SATA	7.2k	1 TB
4 Gb/s Flash (SSD)	N/A	200 GB, 400 GB

Symmetrix V-Max Software Features

Auto-provisioning Groups

The expansion of VMware Virtual Infrastructure implementations in current data centers drives the need for larger consolidated storage arrays. As demand increases, storage administrators need the usual tasks involved with managing storage to be consolidated, simplified, and less time consuming. One task in the process of managing the storage environment is the provisioning of storage. The redundant and common steps associated with this process are simplified in the Symmetrix V-Max system with the introduction of a new method for device masking and mapping named Auto-provisioning Groups. This feature greatly reduces the time it takes to map and mask Symmetrix devices and present them to a VMware Virtual Infrastructure. The time saved with Auto-provisioning Groups is magnified when allocating large amounts of shared storage to large virtual server farms, thus allowing more time for storage administrators to focus on more important, critical tasks. In addition to Auto-provisioning Groups, the

Symmetrix V-Max system provides support for larger Symmetrix devices and automatic creation of metavolumes. These enhancements reduce the instances in which the storage administrator is forced to perform additional activities when provisioning storage to a VMware Virtual Infrastructure environment.

Virtual Provisioning

EMC Symmetrix Virtual Provisioning provides for significant savings in the amount of physical storage capacity required to support large numbers of virtual machines that are configured with common operating system and application software components. In combination with VMware View linked clone technology, Symmetrix Virtual Provisioning allows the storage required for virtual desktop deployments to be provisioned at the levels actually required and reduces the amount of allocated but unused storage that is common in many organizations.

Virtual LUN Technology

The value of business data normally changes with time. Thus the importance of the virtual machines utilizing the data also changes with time. The VMware Virtual Infrastructure provides users with mature tools to manage the lifecycle of VMware virtual machines. The Symmetrix V-Max system introduces the ability to non-disruptively and transparently move data volumes within the storage array, empowering the storage administrator to provide lifecycle management for business data. Enhanced Virtual LUN technology allows Symmetrix devices to be moved between RAID levels or disk types with no disruption to the VMware Virtual Infrastructure accessing the device. This feature allows the Symmetrix V-Max system to support a wide range of service levels within a single storage array and provides a key building block for implementing Information Lifecycle Management.

Disaster Recovery and Symmetrix Remote Data Facility (SRDF)

The combination of EMC's V-Max Symmetrix Remote Data Facility (SRDF) capability with VMware's Site Recovery Manager (SRM) provides very robust and high performance support for protection from any disaster that might disable a primary production site. The SRDF capability to provide continuous off-site replication of all critical data enables VMware's SRM to provide integrated disaster recovery management and automation. This includes integrated management of disaster recovery plans, non-disruptive testing of disaster recovery plans, and automated failover and recovery. These advanced tools simplify design, development, and testing of DR plans and ensure rapid and automatic execution of those plans when business needs require it.

For additional information about utilizing VMware Virtual Infrastructure with EMC Symmetrix storage arrays, refer to Using EMC Symmetrix Storage in VMware Virtual Infrastructure Environments—Solution Guide, available at:
<http://www.emc.com/collateral/hardware/solution-overview/h2529-vmware-esx-svr-w-symmetrix-wp-1dv.pdf>.

Sizing a VMware View Environment

The resource utilization in the physical desktop environment has to be studied before a VMware View environment can be sized appropriately. The following performance counters should be observed over a period of time:

- CPU—Administrators need to understand both average and peak CPU utilization levels.

- **Memory**—The amount of memory used by the operating system and applications should be determined.
- **Storage**—Administrators need to understand the amount of storage being utilized and the performance characteristics (IOPS, type of IOs, average size of IO, etc.).

Estimating CPU

The CPU utilization on a physical desktop can be observed by collecting the perfmon data, % Processor Time. The data should be collected over several days on statistically significant number of physical desktops to size the VMware View environment. Additional counters maybe used to get better understanding of the physical desktop environment.

The average and peak values of CPU utilization should be utilized to determine the CPU requirements for virtualizing the physical desktops. For example, if the monitored physical PC showed an average of 2.79% CPU utilization on a single core 2.2 GHz processor, the virtualized version of the physical desktop will require approximately 130MHz of CPU. Furthermore, if all of the sampled desktops show similar CPU requirement, the data can be extrapolated. Therefore, to support 64 virtual machines on a single VMware ESX Server, a minimum of 8.3GHz of CPU is needed.

It is important to note that virtualization of workload imposes a measurable amount of overhead. Administrators must consider the additional processing power needed to accommodate the work associated with virtualizing the environment, handling the display protocol, and ensuring that there is enough headroom to accommodate any utilization spikes. As a guideline, administrators should plan on adding anywhere from 10% (representing a lesser degree of protection from processor spikes) to 25% (representing a higher degree of protection) to the CPU requirement estimated from the exercise discussed earlier.

Estimating Memory

As is the case in estimating CPU, there is not a straight correlation between the amount of physical RAM used and the amount of virtual RAM. When a physical desktop environment is virtualized, memory requirements actually decrease. This reduction is the result of a feature in the VMware ESX Server commonly referred to as transparent memory sharing.

Transparent Memory Sharing

Transparent page sharing in the VMware ESX server is a feature that greatly reduces physical memory requirements. As operating systems and applications are loaded into memory, VMware ESX's memory manager looks for redundant blocks of code. Rather than have each virtual machine load its own copy of the redundant code, VMware ESX creates a master copy of the memory pages and then redirects each virtual machine to those shared blocks.

This memory sharing can result in each virtual machine requiring less physical RAM than what has been allocated to the virtual machine allowing users to “overcommit” the RAM in their servers. This in turn allows for greater scalability and reduced hardware costs.

Mapping Physical to Virtual Memory Requirements

The perfmon counter, “Memory Available Mbytes”, should be used to understand the memory requirements of the operating system and applications that are in use on a physical desktop. The performance data should be collected over a significant period of time over a statistically significant number of physical desktops.

Studies conducted by VMware have shown that most users PC targeted for virtualization on average utilized 254 MB of RAM. However, VMware recommends provisioning a minimum of 512 MB RAM for a Windows XP virtual machine. Therefore, without transparent memory sharing, a deployment of 64 virtual machines configured with 512 MB of RAM on a single VMware ESX Server requires a total of 32 GB of RAM. However, these same studies showed that the transparent memory sharing feature reduces the physical memory footprint for a Windows XP to approximately 125 MB per virtual machine. Similar reductions can be observed for common applications.

Table 4 shows the average estimate of the optimized memory footprint for a set of Microsoft Office applications. Table 4 also shows the anticipated memory optimized footprint for a single virtualized desktop.

Table 4 *Optimized Memory Utilization for Common Microsoft Applications*

Application/Program/Operating System	Optimized Memory Usage
Windows XP	125 MB
Microsoft Word	15 MB
Microsoft Excel	15 MB
Microsoft PowerPoint	10 MB
Microsoft Outlook	10 MB
Total	175 MB

Based on the data above, the total memory necessary to support the applications running on Windows XP in 64 virtual machines would be approximately 11.2 GB (64 virtual machines x 175 MB/virtual machine). Adding additional memory to handle spikes and accommodate any extra applications would bring the total estimate to 16 GB.

The discussion above does not consider other applications that maybe used by desktop users. These applications could be common off the shelf applications such as Adobe Acrobat or custom applications developed by the users’ organizations. The additional applications require extra memory over and above the numbers cited above. The additional memory requirement should be accounted for when estimating the overall memory requirement for a VMware View deployment.

Estimating Storage Requirements

Performance

The perfmon counters, Bytes Transferred/Sec and Transfers/Sec, provide guidance on the performance requirements of a physical desktop. The data should be collected over several days and on statistically significant number of desktops to understand the performance requirements for a virtual desktop environment.

The storage throughput requirement for a virtual desktop environment can be easily estimated by multiplying the average I/O operations and throughput per desktop by the total number of virtual desktops in the VMware View deployment. For example, if the average IOPS and throughput per desktop is determined to be 5 and 115 Kbs, a 64 virtual desktop deployment requires:

IOPS—5 IOPS x 64 virtual machines = 320 IOPS

Throughput—115 Kbps x 64 virtual machines = 7360 Kbps

Unlike CPU and memory, storage performance tends to be bursty with significant increases in both I/O rate and throughput for a very short duration. Furthermore, the I/O rates increase significantly during the boot up process of desktops. These behaviors need to be accounted for when designing the storage environment for VMware View environments. The ability of the storage array to absorb unexpected spikes in performance requirements becomes important when designing the storage for VMware View deployments.

Capacity

Calculating the amount of storage required to support a VMware View environment is a fairly straightforward process. Administrators need to take the base size of their proposed virtual machine (accounting for the operating system, applications, and local data) and add additional storage to account for suspend/resume and page log files, as well as an amount to cover any extra headroom.

For example, consider a virtual desktop with 512 MB of memory and a 10 GB virtual disk. To accommodate the page file and suspend/resume, administrators need to add disk space equal to the RAM allocated to the virtual machine, in this case 512 MB. Similarly, a good estimate for the size of log files is 100 MB per virtual machine. Therefore, the storage requirement for the virtual desktop can be estimated to be 11.1 GB.

The storage estimate for a larger environment can be estimated by multiplying the storage requirement per virtual desktop by the total number of desktops in the environment. For example, if 64 virtual desktops discussed above are deployed, the total amount of storage required can be estimated to be 710 GB. An additional 15% is recommended to allow for any unplanned storage requirement. It is important to note that these estimates are independent of any storage reduction capabilities.

Administrators must also decide on how many virtual machines will be deployed on each LUN. This decision is usually dependent on the particular storage array that will be used. For a more detailed discussion on this topic, see:

- <http://www.vmware.com/resources/techresources/1059>
- http://www.vmware.com/files/pdf/VMware_VDI_Server_and_Storage_Sizing_120508.pdf

VMware View Deployment on UCS Systems

Innovations such as extended memory, virtualization optimization with VN-Link technology, unified and centralized management, and use of service profiles make the Cisco UCS systems ideally suited for deploying VMware View environments. The following sections discuss the advantages of deploying a VMware View environment on Cisco UCS system.

UCS Technology Can Benefit Total Cost of Ownership and Manageability

VMware View desktops run in a virtualized data center environment which significantly improves manageability and security. However, the execution of many desktops on a single physical server can require large amounts of memory. Since CPU performance can outstrip memory performance, memory bottlenecks can arise. Many enterprises today either deploy 4-socket servers or multiple 2-socket servers to address this issue. This approach, however, can increase the total cost of ownership.

Cisco UCS's innovative Extended Memory Technology expands the capability of a 2-socket server to support up to 48 DIMM slots or 384 GB of memory using standard memory DIMMs. This innovative technology thus allows VMware View virtual desktop environments (that may require larger memory footprints) to run on a 2-socket x86 server. This in turn reduces the TCO for the solution with lower capital costs as well as lower operating costs.

Each desktop in a VMware View solution runs inside a virtual machine (VM) on the server. IT managers expect to apply network policies (e.g., security or compliance policy) per desktop and expect the policies to persist as the virtual machines move between servers to balance workloads. Cisco's VN-LINK technology in the Cisco UCS provides these capabilities.

The VN-LINK technology can be either provided through the software-based Nexus 1000v or with Cisco's virtual interface card, the Cisco UCS VIC M81KR. The VIC M81KR has significant benefits in a VMware View environment. It provides up to 128 virtual interfaces that can be connected directly to virtual desktop virtual machines by using pass-through switching or hypervisor bypass technology. The hypervisor bypass technology helps reduce the CPU overhead dedicated to networking, thereby providing additional CPU cycles to run more virtual desktops on the same server.

The use of the VN-LINK technology, however, does not increase the administrative overhead. The provisioning and management of all the virtual interfaces and policies in the UCS system are done centrally from UCS Manager. The centralized approach provided by UCS manager not only simplifies "desktop" management, but also provides visibility and policy enforcement to each virtual desktop to meet regulatory compliance.

Dynamic VMware View Provisioning and Deployment

IT managers expect to dynamically provision compute and storage resources for various operations. For example, an IT manager in a financial institution may choose to add compute resources to the VMware View environment by adding physical servers as employees show up in the morning. However, when employees go home in the evening, the IT managers may choose to repurpose the servers run analytics useful for the next day. Similar patterns can be found in other industries, as in large retail environments where the shift is from transactions to analytics and back.

The server infrastructure needs to be highly agile and offer dynamic provisioning capabilities to quickly accommodate changes in the environment. The Cisco UCS system facilitates dynamic provisioning with Service Profiles. The Cisco UCS "Service Profile" abstracts compute and network attributes like identity (UUID, MAC address and WWN), I/O configuration, firmware versions, and BIOS boot order. The service profile can quickly be applied to a hardware blade when a server needs to be deployed.

The use of a service profile in a VMware View environment can be elucidated by considering the use case presented earlier. To support that use case, the IT manager can predefine two service profile templates, View-Service-Profile and Analytic-Service-Profile. The first service profile has all of the compute and network attributes necessary for a physical server to run virtual desktops. The second service profile has the definitions to support the analytics workload.

With these definitions in place, in the morning when additional resources are required to support the virtual desktops, the IT manager instantiates a new service profile from the “View-Service-Profile” template. The server has all the LAN and SAN properties like VLAN, VSAN required to deploy the desktops. Furthermore, tying the profiles to support boot from SAN, the physical server can be quickly powered on with the appropriate operating system required to support the VMware View workload. At the end of the workday, the IT manager can instantiate the “Analytic-Service-Profile” template on the same set of servers. Similar to the “View-Service-Profile” the second template has all of the LAN and SAN properties required to run the analytics workload. If the profile also includes the boot from SAN attributes, the physical server can be rebooted with an operating system required to run the analytics workload.

It is important to note that the provisioning and management of the service profiles are conducted through the UCS manager. Therefore, the added functionality provided by the UCS platform does not increase the management overhead.

Focus on Reducing TCO

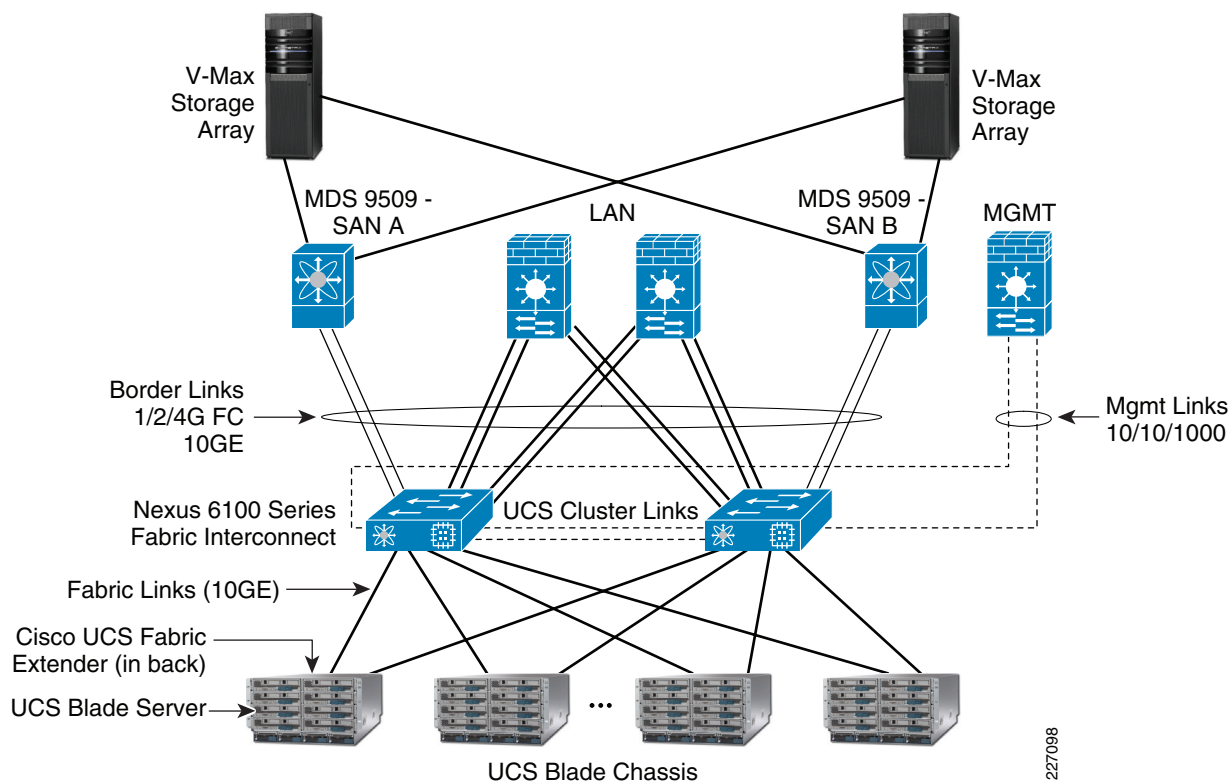
IT managers are deploying virtual desktop solutions for both security and cost (CAPEX and OPEX) benefits. The Cisco Unified Computing System represents a radical simplification of traditional architectures, dramatically reducing the number of devices that must be purchased, cabled, configured, powered, cooled, and secured. Instead of using multiple different types of adapters, switches, and management tools, it uses 10 Gbps converged network adapters and unified fabric with an innovative Fabric Extender IO architecture to unify and simplify the infrastructure. Less infrastructure has facilitated a unique internal design, making the UCS system highly energy efficient. Unified IO with the innovative Fabric Extender IO architecture helps cut down significantly on cabling and further simplifies management.

In addition to the savings derived from the physical design, the UCS system helps decrease TCO and improves business agility by automating just-in-time provisioning through use of “Service Profiles”. UCS Manager allows IT managers from server, network, and storage disciplines to discover, provision, and monitor all the system components from within a single tool. This reduces operational complexity, which both decreases risk and enables better IT responsiveness. The system is based on open standards and can interoperate with existing system management tools.

Solution Overview

[Figure 5](#) shows a multichassis UCS configuration with the UCS Nexus 6140 XP fabric interconnect to the SAN and LAN infrastructure. One of the key value propositions of UCS is that it plugs right in to the existing data center SAN and LAN infrastructure. This can be seen in [Figure 5](#), where everything above the Nexus 61x0 fabric interconnect already exists in the data center or can be quickly deployed as the core components. As the uplink Ethernet ports coming off the Nexus 6100 switch are 10 gigabit Ethernet, to fully utilize the capabilities it is important to connect the Ethernet ports to a 10 gigabit Ethernet LAN switch, such as Cisco Nexus 5000 or a Cisco CAT 6x00 series 10 gigabit Ethernet switch. Similarly, since the Nexus 6100 switch supports VSAN technology, it is important to connect the SAN components to a switch that supports VSAN technology, such as the Cisco MDS 9509. The core fabric should provide connectivity to the enterprise storage array, such as the EMC Symmetrix V-Max system. The management links could be a 10/100/1000 Mb/s link connected to LAN switch.

Figure 5 *High-Level Topological View of VMware View Deployment on UCS and Symmetrix V-Max Array*

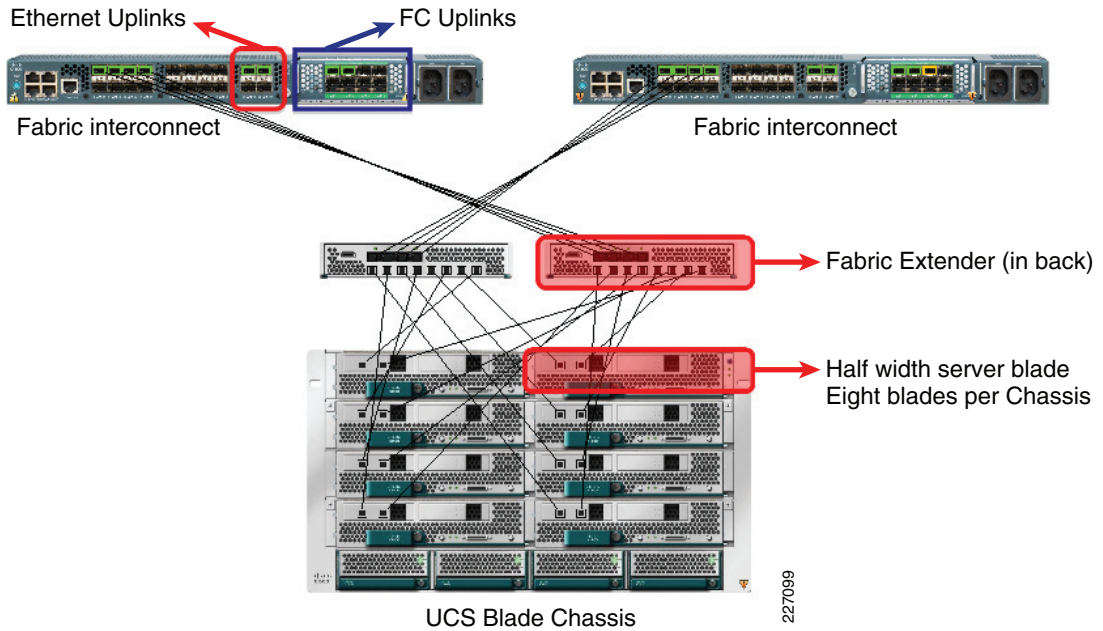


One interesting point to note in [Figure 5](#) is the fully redundant nature of the setup. All connections out of the UCS Blade chassis, including the blade connectivity in to the UCS chassis, the dual fabric interconnect connection from the fabric extender, and the SAN and LAN connectivity are fully redundant. This builds resiliency to the data center infrastructure and addresses any issues that can arise due to failure of a single component in the system. The management of the UCS Nexus 61x0 XP fabric interconnect can be on a separate LAN and avoid data traffic for out-of-band management. Each of the links from the blade, from the chassis fabric extender and the fabric interconnect, is 10 GB for FCOE and Ethernet and 4 Gb for FC.

Detailed Topology

[Figure 6](#) shows the Cisco UCS configuration for the VMware View testing. The UCS chassis comprised of eight fully populated memory blades (96 GB each). The four ports of the fabric extenders were connected to the upstream fabric interconnect. This provided 40 Gbps line rate and a oversubscription ratio of 1:2, meaning the 80 Gbps (8 X 10 Gbps) from the blades are funneled through the 40 Gbps upstream connection. Two ports on each fabric interconnect were configured as uplink Ethernet ports and four ports each were configured as downstream server ports. The expansion port of the fabric interconnect was an 8-port FC card (1/2/4 Gb). Two of the ports in the FC card were connected with 4 Gb FC SFP modules to a Cisco MDS SAN switch that leveraged VSAN technology to segregate the I/O workload generated by the VMware View deployment from other storage area network activities. The Symmetrix V-Max ports are then connected to the Cisco MDS SAN switch in a highly redundant fashion.

Figure 6 UCS Configuration Used for the Scalability Testing of the VMware View Solution



The tables below list the configuration details of all of the server, LAN, and SAN components that were used for the study.

Table 5 Physical Network Details

Networking Components	
Quantity	Description
2	Nexus 5000 10 GE Layer 2 Switch Modules
2	10 Gigabit Ethernet Modules
Building Block Network Components	
1	48 Port Network Switch
VLAN Configuration	
VLAN ID	Description
16	VMware View Desktops—Infrastructure
20	Management

Table 6 SAN Switch Details

Quantity	Description
2	Cisco MDS 9509 Director class switch

Table 7 **Software Components**

Description
VMware ESX 3.5 U4
VMware Virtual Center 2.5 for ESX 3.5U4 running on Windows 2003 server SP3
Windows 2003 server SP3 running DHCP, Active directory services, SQL server

Table 8 **View Desktop Configuration**

Virtual Machine Configuration	
Quantity	Description
640	Windows XP SP2—Virtual Machines
	512 GB RAM
	8 GB Virtual Disk

Table 9 **Physical UCS Server Configuration**

Quantity	Description
1	UCS 5108 Blade Server Chassis/4 PSU/8 fans/2 fabric extender
8	UCS B200 M1 Blade Server
2	Quad Core Xeon 5500 series 2.93 GHz processors per blade server
12	8 GB DDR3 DIMMS 1066 MHz per blade totaling 96 GB/Blade server
1	UCS M71KR-Q QLogic Converged Network Adapter/PCIe/2port 10Gb per blade server
2	73GB SAS 15K RPM SFF HDD/hot plug per blade
2	UCS 6120XP 20-port Fabric Interconnect/2 PSU/2 fans
2	8-port, 4 Gb FC Expansion port module
4	4 Gbps Fibre Channel-SW SFP, LC
20	10GBASE-SR SFP Module
8	10GBASE-CU SFP+ Cable 5 Meter
8	Fiber cables for connectivity to FC and 10GE

Configuration of UCS System for VMware View Solution

The hardware configuration presented in [Figure 6](#) was configured through a series of steps to support the VMware View scalability study. Specifically, the following configuration steps were performed:

1. Configuration of Nexus 6120 Fabric Interconnect.
 - a. Configuration of the cluster pair.
 - b. Configuration of KVM.
2. Configuration of service profiles to bring up the UCS Blade servers.

3. Configuration of Boot from SAN and boot policy.
4. Installation of ESX 3.5U4 on all the blades.
5. Configuring VC and deployment of VMs using VMware View 3.0.

The following sections discuss in detail the configuration process.

Configuration of the Nexus 6120 XP Fabric Interconnect

The Cisco Nexus 6120 XP fabric interconnect is always configured in a cluster pair for redundancy and provides resiliency and access to the system configuration data in a rare case of a switch failure. The configuration database is replicated from the primary switch to the standby switch. All operations are transaction based, keeping the data on both switches in lock step.

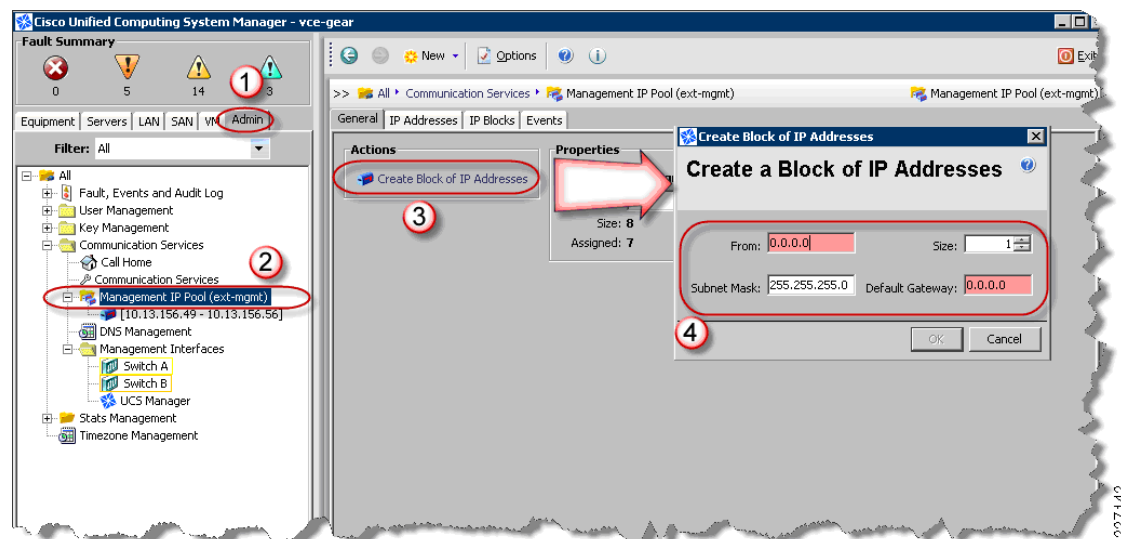
The fabric interconnect comes preinstalled with a software bundle. When the switch is powered on, at the loader prompt, minimal configuration, such as setting the IP address of the switch, netmask, and the gateway has to be performed using the serial port connection. After configuring the switch, the kickstart image and system image are loaded using the boot command. When the switch boots up, the setup command is run, which guides you through the configuration steps. There are two options to configure, namely cluster or standalone. While bringing up the primary switch, select the standalone mode. Two IP addresses are assigned at this time, one for the switch itself and one for the virtual IP address. The second switch is similarly booted off and assigned the third IP address. However, instead of performing a standalone configuration, the “cluster” mode should be selected. At this point, the two switches communicate and the second switch becomes part of the cluster in the standby mode.

Once the primary switch comes online, the UCS system can be configured by pointing a browser to the IP address of the primary switch. Alternatively, the Cisco UCS CLI commands can be used to configure the system.

KVM Configuration

After the Nexus 6120 XP switches have been configured, the first operation would be to block off eight IP addresses for the KVM console for each blade. These IP addresses have to be part of the management domain to which the switch belongs. This could be done in the “Admin” tab of the Cisco UCS manager software by following the numbered process shown in [Figure 7](#).

Figure 7 *Configuring Pool of IP Addresses for KVM Access to the UCS Blade Servers*



Service Profile Configuration

There are a number of steps for creating a service profile using UCS manager. The following operations need to be performed before a service profile can be created.

- **SAN Configuration** (in the SAN tab)—In this tab, users can set the VSANs to use in the SAN (if any) and set up WWN pools and WWN pools for assignment to the blade servers virtual HBAs (vHBA).
- **LAN Configuration** (in the LAN tab)—Users can set up the VLAN assignments to the vNICs in this tab. In addition, users can also set up MAC address pools for assignment to vNICs.
- **Boot Policy** (in the Servers tab)—The boot policy based on each individual operating system is created in this tab.
- **Firmware policy** (in the Servers tab)—To ensure the blade servers run the appropriate firmware at all times, it is important to configure the firmware policy. If this step is not performed, the factory default firmware is used.

Once the preparatory steps have been completed, a service profile template can be generated for a given type of workload. As shown in [Figure 8](#), this activity is performed in the “Servers” tab. During the creation of the service profile template using the numbered procedure shown in [Figure 8](#), appropriate policies created earlier for the SAN and LAN clouds should be applied.

The service profile template can be used to create service profiles to be used with the blade servers. It is important to ensure that the WWN and MAC addresses for the service profiles are from the pool and are not derived. The service profile can be associated with a blade server when needed. The status of the association can be observed in the server’s finite state machine (FSM) monitor.

Booting UCS Blade Servers from SAN

There are three distinct steps that have to be performed to allow the UCS blades to boot from the storage area network:

- Storage array configuration

- SAN configuration
- UCS configuration of service profile

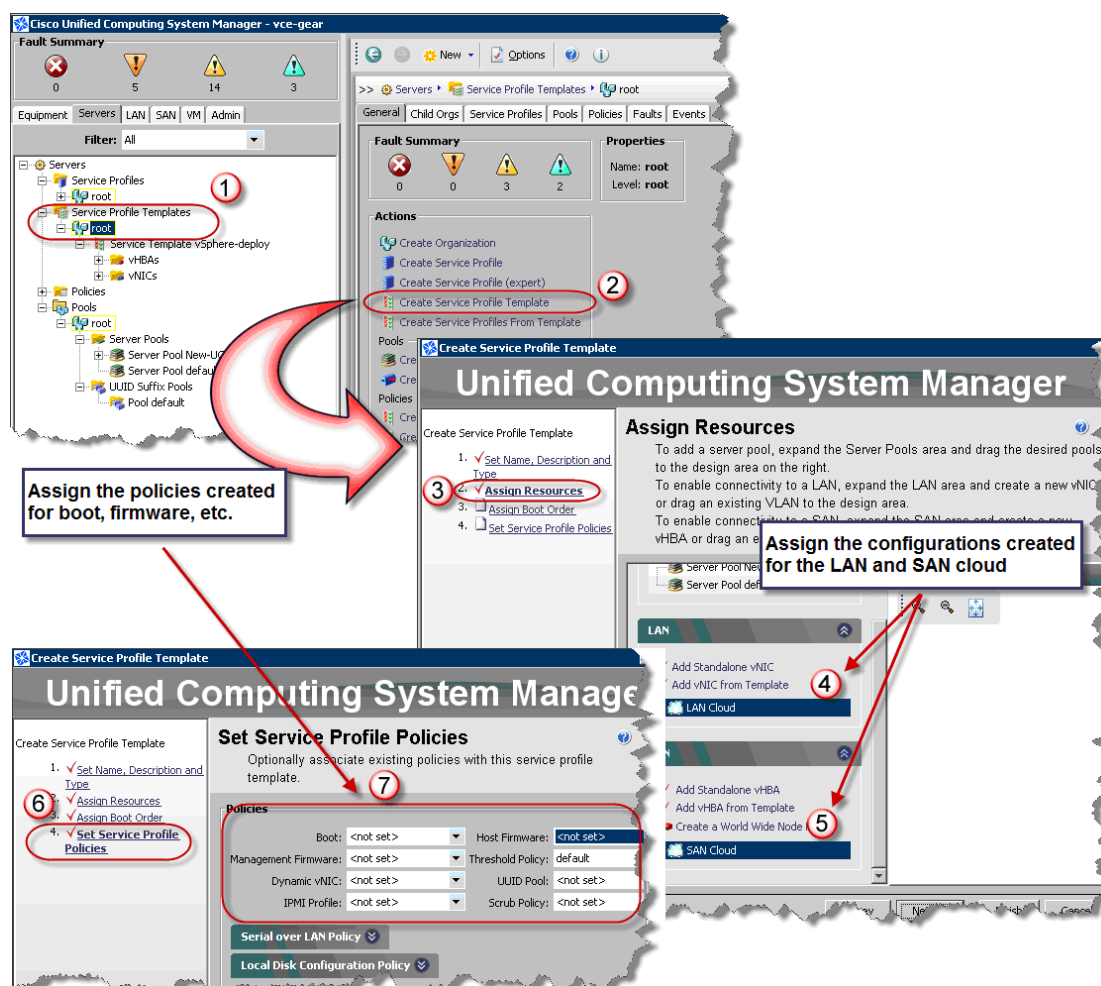
Storage Array Configuration

The configuration process to provision storage to enable boot from SAN is no different than that for other requirements. First, the storage array administrator has to create at least a 10 GB device for installing the operating system. Although not a requirement, it is advisable to present this device at address 0. To perform this activity on the Symmetrix V-Max storage array, the storage administrator needs to create a view that is a conjunction of an initiator group, port group, and a storage group.

The initiator group object contains the WWPN assigned to the UCS blade by the active service profile discussed in the following section. The port group is a construct that includes the front-end ports of the Symmetrix V-Max storage array from which the UCS blades would boot. The WWPN of the front-end ports of the directors should be utilized when creating the boot policy component of the UCS service profile. The storage group definition would include the device that was created to enable the UCS blade to boot from the SAN.

There are two different choices when deploying a boot from SAN solution. The first option segregates the front-end ports used for boot activity from those used for data access. The second option uses the same front-end ports for accessing both boot and data devices. In general, EMC recommends the separation of the ports from which the boot and data devices are accessed. However, for VMware View environments the I/O workload is normally not significant enough to justify the separation of the boot and data devices. Therefore, in VMware View environment, EMC recommends utilizing the same set of Fibre Channel ports on the Symmetrix V-Max storage array for both boot and data devices.

Figure 8 Creating New Service Profile Template for VMware View Deployment



SAN Configuration

The NPIV feature has to be enabled in the SAN switch connected to the UCS system. In addition, 4 GB SPF+ modules are connected to the Nexus 61x0 XP fabric interconnect with the port mode and speed set to AUTO.

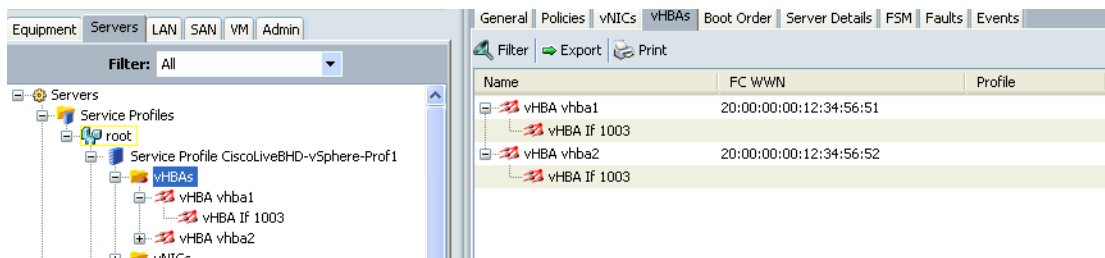
The SAN fabric has to be zoned to allow the vHBA assigned to the UCS blade to access the target Fibre Channel ports on the storage array (the Symmetrix V-Max ports in this study). In addition, as discussed in the previous section, appropriate views have to be created on the Symmetrix V-Max storage array that allows the WWN associated with the vHBA access to the boot device.

Service Profile Configuration on UCS

The following steps need to be performed to enable boot from SAN. The procedure is the same as that presented in [Service Profile Configuration on UCS](#), however it is presented in greater detail.

- Step 1** When creating the service profile configuration from the service profile template, as shown in [Figure 9](#), create two vHBA and assign a WWN to each interface. In addition, the vHBA should be assigned to the appropriate VSAN. In [Figure 9](#), vhma1 and vhma2 are created in VSAN 1003.

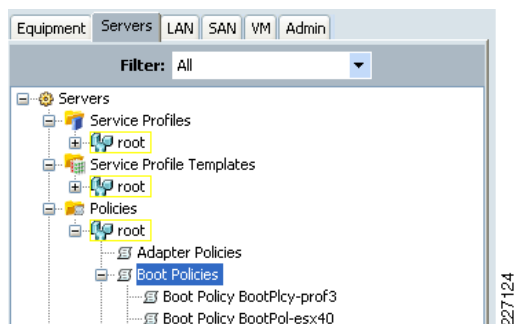
Figure 9 *Creating The Service Profile Configuration From the Service Profile Template*



227123

- Step 2** The boot policy should be created next. The location of this configuration option is shown in [Figure 10](#).

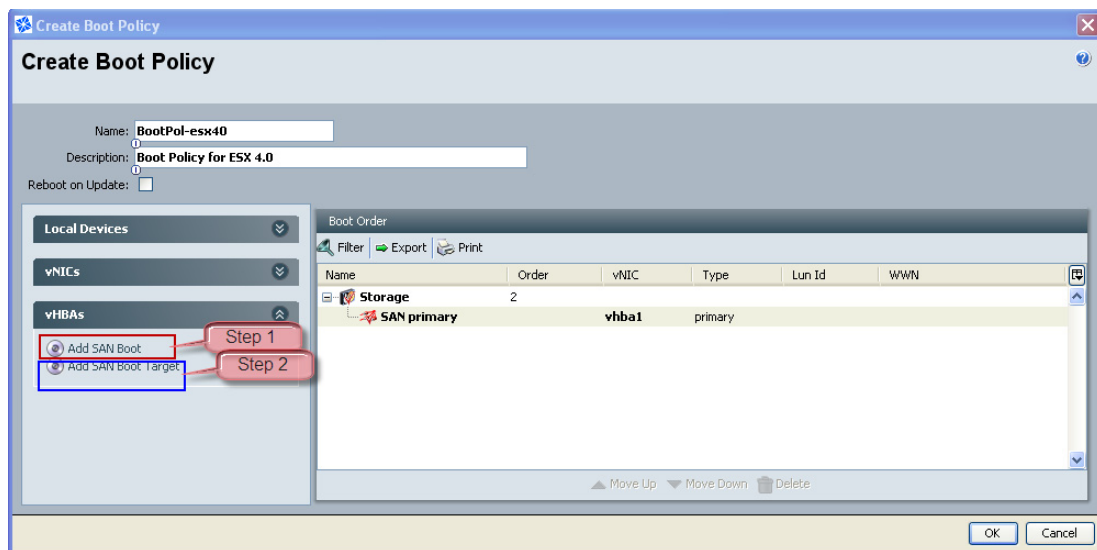
Figure 10 *Boot Policy Configuration Option*



227124

- Step 3** Define a name for the boot policy in the “Create boot policy” wizard along with a description. Then under vHBA, select “Add SAN Boot”. This is highlighted as step 1 in [Figure 11](#).

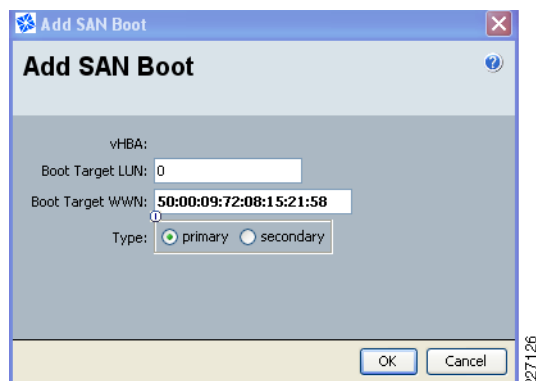
Figure 11 *Create Boot Policy*



227125

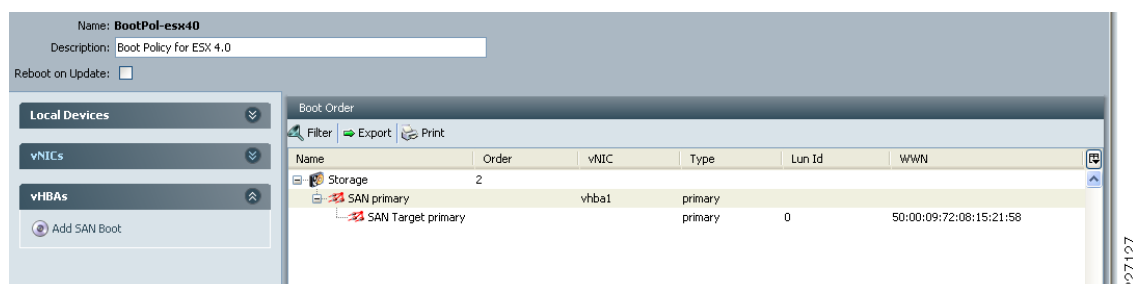
- Step 4** Selecting the “Add SAN Boot” option opens a new window in which the WWN name of the Symmetrix V-Max storage port should be entered. The LUN address of the boot device should also be entered. As discussed earlier, a LUN address of 0 is recommended for accessing the boot device. The process is depicted in [Figure 12](#).

Figure 12 Add SAN Boot




- Step 5** An example boot policy for booting a UCS blade from an EMC Symmetrix V-Max storage array is shown in [Figure 13](#).

Figure 13 Example Boot Policy for Booting a UCS Blade from an EMC Symmetrix V-Max Storage Array



- Step 6** The boot policy that was created in the previous step should be assigned to the boot order of the appropriate service profile, as shown in [Figure 14](#).

Figure 14 Assign Boot Policy to the Boot Order of the Appropriate Service Profile



- Step 7** The newly created or modified service profile can now be associated to any blade in the environment. When the service profile is applied the blade is automatically rebooted. The FSM on the server where the service profile is applied can be monitored to observe the progress.

Once the service profile with the appropriate boot policy has been assigned to a UCS blade server, the installation of VMware ESX Server can be performed using the Virtual Media facility of the UCS manager.

Installation and Configuration of VMware ESX Server

The installation of the VMware ESX Server on the UCS blade can be performed using the virtual media facility functionality provided by the KVM associated with the UCS blade server. The KVM console provides a utility that allows users to mount a physical CD/DVD or an ISO file on the client executing the UCS manager GUI. The CD/DVD or ISO is presented to the blade as a virtual CD drive. The installation of the VMware ESX Server can then be performed similar to that performed on any other physical server.

After the installation of the VMware ESX Server is complete, the following configurations should be performed to reap the full benefits offered by the UCS system:

-
- Step 1** Configure the vSwitch with appropriate VLAN information for separating the VM traffic and workload traffic.
 - Step 2** Enable jumbo frame support on the virtual switch configured on the ESX Server. The MTU of the virtual switch should be changed to 9000 bytes. This can be achieved, for example, by executing the command **esxcfg-vswitch -m 9000 vSwitch0**.
 - Step 3** By default the virtual switch on VMware ESX Server version 3.5 supports only 56 ports. This number should be increased to 248 ports. The change can be performed utilizing the VMware vCenter client.
 - Step 4** The vCPU limit should be increased to 192 (the maximum for ESX 3.5U4). This is essential to support the large number of virtual desktops that can be supported on a single UCS blade server.
 - Step 5** If needed, root login from outside can be enabled by configuring the ssh_config file and restarting the ssh daemon.
-

Symmetrix V-Max Configuration for VMware View Solution

The Symmetrix V-Max storage array used for the testing had two storage engines with 64 GB of cache and 240, 300 GB, 15,000 RPM disk. The array also had 32 Fibre Channel ports configured either for host access or for Symmetrix Remote Data Facility (SRDF). A subset of the resources available on the Symmetrix V-Max storage arrays were utilized for the solution presented in this article. The configuration of the storage array for the testing was selected to ensure enough resources were available to satisfy the anticipated maximum number of desktops on a single UCS chassis.

Storage was provided to the UCS chassis as eight, 512 GB virtually provisioned logical devices. A datastore was created on each one of the logical devices. The virtually provisioned devices utilize very little physical storage initially (to store the VMware file system and Symmetrix V-Max storage array metadata) and consume storage only when the desktops are deployed and they start actively modifying data. This approach also has the additional advantage of spreading the workload generated from the desktops across larger number of physical hard disk than otherwise possible.

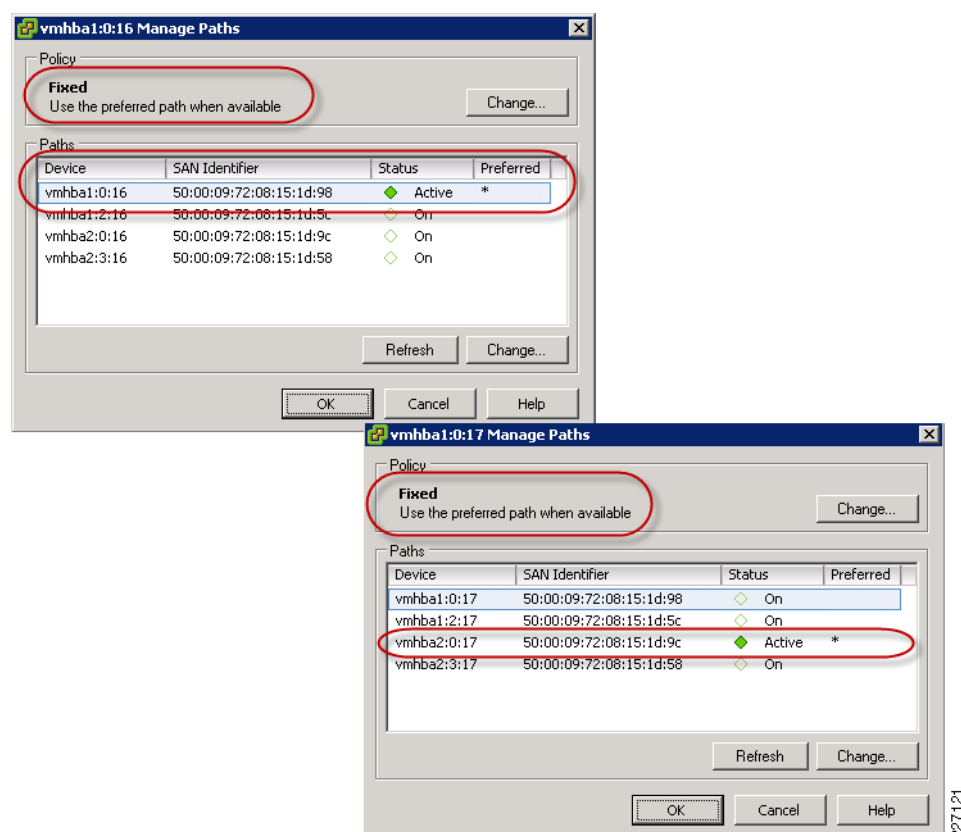
Each logical device was presented to the UCS blades on four independent Fibre Channel ports. The connectivity was provided following the best practices when utilizing Symmetrix V-Max storage systems in a VMware Virtual Infrastructure environment. Further details can be found in the EMC

Symmetrix V-Max and VMware Virtual Infrastructure whitepaper at

<http://www.emc.com/collateral/hardware/white-papers/h6209-symmetrix-v-max-vmware-virtual-infrastructure-wp.pdf>.

VMware ESX Server version 3.5 that was utilized in the testing does not perform any load balancing across all possible active paths. To ensure utilization of the four Fibre Channel ports presented from the Symmetrix V-Max storage system, static load balancing was utilized. The preferred path for each LUN presented to the UCS blade was changed so that the IOs to the LUNs used either a different HBA or a Symmetrix V-Max Fibre Channel port. Figure 15 shows the preferred path for two of the logical devices that were utilized for the testing. It can be seen from the figure that all IOs to the logical device at address 16 is sent through vmhba1 on one of the Symmetrix V-Max FC ports, whereas the IOs to the logical device at address 17 is channeled through the second HBA to a different Symmetrix V-Max FC port. The static load balancing technique discussed above is strongly recommended for all VMware View deployments on Symmetrix V-Max storage systems.

Figure 15 Static Load Balancing When Using VMware ESX Server with Symmetrix V-Max Array

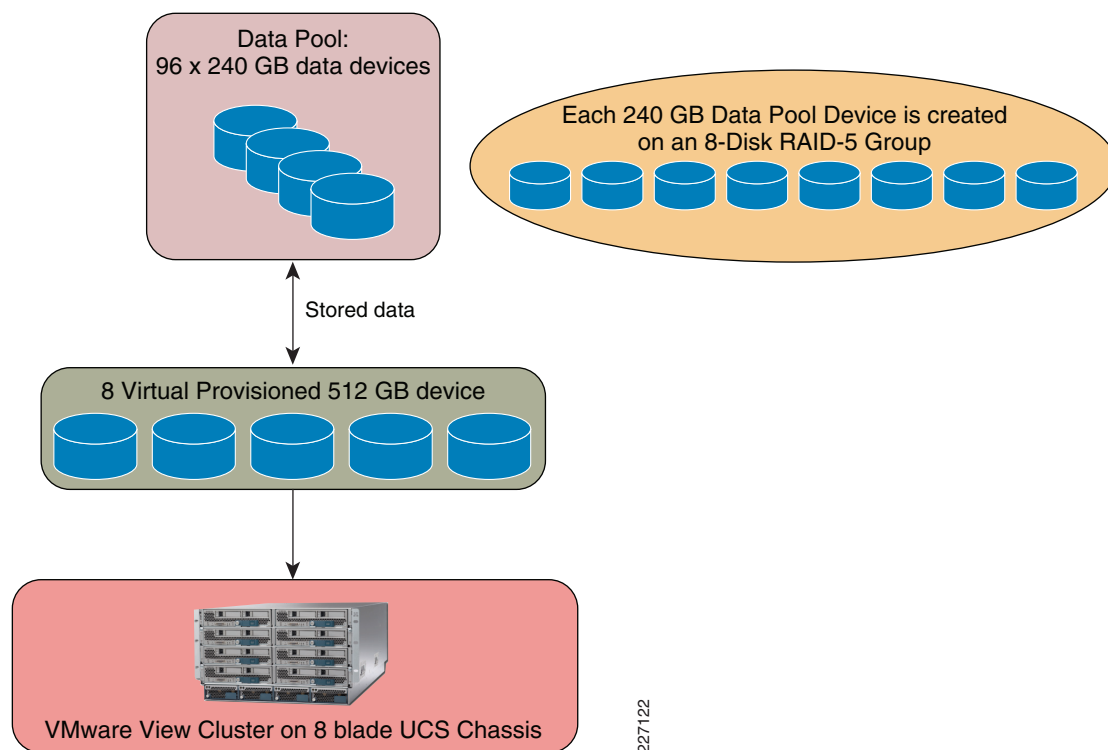


The logical devices were bound to a storage pool spread across 128 of the 240, 300 GB hard disk drives available on the Symmetrix V-Max array. The data devices created on the physical devices were protected with a RAID-5 scheme with 7 data members.

Figure 16 shows the logical view of the data stores as configured in the EMC Symmetrix V-Max array used for the solution presented in this paper. Each virtual desktop is provided with a VMware View Composer linked clone of a master boot disk image. This configuration provides a very high ratio of virtual storage space to used physical disk space. If desktop users need more space, you can add data pool devices. If more bandwidth is needed, you can configure the new data pool devices across additional physical spindles. There is a tradeoff between saving storage space and conserving storage bandwidth

and the optimal balance depends on the deployment. The combination of EMC Virtual Provisioning with VMware View Composer allows users to gain an optimal balance by providing sufficient physical spindles to handle the workload while allowing the flexibility of utilizing the storage only when required.

Figure 16 *Logical View of Storage Provided in Support of VMware View Environment*



EMC provides flexible storage options for virtual desktop deployments to meet the scale and configuration requirements for a wide range of environments (Table 1). These options all offer high performance, 99.999 percent availability, and thin provisioning to allow storage consolidation and optimization.

Solution Validation Testing

Testing the Scalability of VMware View on UCS and Symmetrix V-Max systems

Determining the scalability of the VMware View environment deployed on UCS system connected to an EMC Symmetrix V-Max system can be subjective. The VMware tool, `esxtop`, was utilized to understand the behavior and the performance of the solution. The VMware `esxtop` tool provides a real-time view of ESX Server system, service console, and virtual machines sorted by CPU usage. It lists CPU utilization for each physical processor, memory utilization, and disk and network bandwidth for each network and disk device available to the ESX Server machine.

The data from the esxtop tool was collected at one-minute intervals while the scripted simulated desktop workload was executed on the virtual desktops repeatedly. The server was considered to have reached its maximum capacity if the time taken for the simulated workload iteration is more than certain percent above that required for the previous iteration. The server was also considered to have reached the maximum capacity if any portion of the scripted workload failed as a result of predefined timeouts.

User Workload Simulation

The simulated desktop workload was based on scripts created using AutoIT (see [References](#) for a link). The scripts used simulate a “knowledge worker” user profile, with the following applications:

- AutoIT v3—Software used to run the workload.
- Microsoft Word—Open, minimize, maximize, and close the application, write words and numbers, and save modifications.
- Microsoft Excel—Open, minimize, maximize, and close the application, write random numbers, insert and delete columns and rows, copy and paste formulas, and save modifications.
- Microsoft PowerPoint—Open, minimize, maximize, and close the application and conduct a slide show presentation.
- Microsoft Outlook—Open, minimize, maximize, and close the application and send E-mail messages.
- Internet Explorer—Open, minimize, maximize, and close the application.
- Windows Media Player 11—Open and close the application and play back WMV files.
- Adobe Reader 8.0—Open, minimize, maximize, and close the application and browse PDF files.
- McAfee Virus Scanning—Real-time scanning.
- PKZIP—Open and close the application and compress large files.

By cross-referencing the result with esxtop utility server activities, test were run to generate the office work sequences to stress the server scalability, performance, and stability as the virtual desktop load was increased on the UCS server blade. The storage system, as discussed earlier, was configured a priori to handle the maximum anticipated workload expected from a fully loaded UCS chassis.

The workload intelligently incorporates grouping with Active Directory to match a real-world profile and policy. The data collected during the study allows analysis of the following parameters:

- Application response time
- Average CPU utilization
- Average I/O rate for disk volume
- Average available memory usage
- Percentage of CPU performance on ESX servers

Interpreting the Results

Since VMware does not provide the remote display protocol (RDP), and since the latencies will vary based on the protocol used and the physical distance between the clients and the desktops (LAN or WAN deployment), we do not monitor the latencies for specific GUI operations during the workload test.

Any interpretation of the results and analyses presented in this paper must take into account the nature of the test workload. Workloads or user environments that involve different applications, different data sets, different activity profiles, and different infrastructure conditions may yield different results.

Testing Methodology

The first goal of the testing strategy was to determine the largest possible deployment on a single UCS blade server with 48 GB of memory as well as 96 GB of memory. This was achieved by incrementing the number of deployed desktops running the simulated workload discussed above by 32. The study started with a small deployment of 32 virtual machines and expanded to a maximum of 160 desktops.

The workload was run in three modes, ramp up, steady state, and disruptive mode. Ramp up mode involved VMware View gradually booting the virtual desktops as the client requests for a desktop session were received. The amount of time required for ramp-up depended on the number of virtual machines on the VMware ESX Server and the rate at which the client requests for a desktop arrived. The performance data was recorded until the last virtual desktop in the environment was booted up and had started executing the workload.

The disruptive mode involved the reboot of all virtual desktops that were previously deployed. The goal of this test was to evaluate the impact of an unplanned event like VMware ESX Server outage on the UCS system and the Symmetrix V-Max storage array. The performance data during the reboot of the virtual desktops was collected and analyzed.

The last mode was the steady state mode in which case all of the virtual desktops were executing the simulated user workload described earlier. The resource utilization during the steady state execution was also collected and analyzed.


Test Results

In this section we present the results from the study with an emphasis on CPU, memory, and disk utilizations. The three cases that highlighted the various scalability limits of the solutions are discussed below.

VMware View Deployment with 128 Desktops on UCS Server with 48 GB of Memory


[Figure 17](#) shows the performance of CPU utilization of a VMware ESX Server on UCS server with 48 GB of memory as the virtual desktops are gradually powered on by VMware View servers. The graph also shows the CPU utilization during the steady state execution of the simulated workload. It can be seen from [Figure 17](#) that during the ramp up process the maximum CPU utilization is approximately 90% with a steady state CPU utilization of 50%. It is also interesting to note that there is unequal CPU utilization of each logical processor in the CPU complex with the maximum utilization exhibited by CPU 0. Although the reason for the skewed performance is uncertain, it is possible the higher utilization of CPU 0 could be due to the processing performed by the management daemons running on the service console accepting requests from the VMware View servers requesting power-on operations of individual desktops.

Figure 17 CPU Utilization During the Ramp-Up of 128 Virtual Desktops on a Single UCS Server




The memory utilization of the UCS blade during the ramp up process of 128 virtual desktops is shown in [Figure 18](#). It can be seen from [Figure 18](#) that the memory utilization rapidly increases and peaks at 90% as virtual desktops are instantiated and powered on by the VMware ESX Server kernel. It can also be observed that the memory utilization reduces over period of observation to a steady state value of 70%. The drop in memory utilization can be attributed to the transparent memory sharing feature of the VMware ESX kernel.

Figure 18 *Memory Utilization During the Ramp-Up of 128 Virtual Desktops on a Single UCS Server with 48 GB of Memory*



The intense I/O generated during the ramp up process can be seen in [Figure 19](#). The throughput reaches 60 MB/s during the ramp up in the relatively granular observation interval. The peak I/O rate and throughput is anticipated to be more intensive than that shown in [Figure 19](#). The capability of the Symmetrix V-Max storage array to absorb bursty I/O traffic resulted in no perceptible delays or timeouts on the VMware ESX Server during the ramp up process. The steady state throughput settles at approximately 10 MB/s with occasional bursts reaching 30 MB/s.

Figure 19 *Disk Usage During the Ramp-Up of 128 Virtual Desktops on a Single UCS Server with 48 GB of Memory*



Storage System Performance with 128 Virtual Desktops on Single Blade

Figure 20 shows the average utilization of the four ports of the Symmetrix V-Max storage array supporting the VMware View environment running 128 desktops on a single UCS blade. Figure 21 displays the utilization of the front-end Fibre Channel ports when the desktop virtual machines are rapidly booted after a planned or unplanned outage of the ESX Server. It can be seen from the figure that the average utilization of the ports does not exceed 5% during steady state workload processing and does not exceed 10% during the boot storm that occurs after a reboot. The low utilization rates seen on the Fibre Channel directors is not surprising since the environment was sized to support the maximum number of desktops that can be simultaneously executed on a single UCS chassis. Furthermore, it can be seen from the figures that the Symmetrix V-Max configuration can easily support sudden burst in activity that can occur in VMware View environments.

Figure 20 Utilization of the Front End Ports on Symmetrix V-Max Storage Array






Figure 21 Utilization Rate of the Font-End Ports After an Unplanned ESX Server Reboot



The average I/O rate on the four Fibre Channel ports used for testing with a 128 desktop VMware View environment on a single UCS blade is shown in [Figure 22](#). The average I/O rate the Fibre Channel ports were subject to during a boot storm in the environment is shown in [Figure 23](#). The results displayed in the figure are similar to those presented earlier in [Figure 20](#) and [Figure 21](#). The peak I/O rate generated during a boot storm is approximately 5 times that generated by steady state workload similar to the ratio


between the utilization of the ports discussed in [Figure 20](#) and [Figure 21](#). The I/O rate generated by the 128 VMware View desktops running a workload representative of a knowledge worker is well within the capabilities of the resources allocated from the Symmetrix V-Max arrays.

Figure 22 *I/O Rate on Front -End Symmetrix V-Max Ports During Execution of Steady State Workload*



227102


Figure 23 *I/O Rate on Front -End Symmetrix V-Max Ports After an Unplanned ESX Server Reboot*



VMware View Deployment with 160 Desktops on UCS Server

The initial study with 128 desktops executing the knowledge worker profile (discussed in [User Workload Simulation](#)) on a single UCS blade server clearly showed that the CPU utilization during normal operations were well below the server's capabilities. However, the memory utilization with 48 GB was much higher than the CPU utilization, indicating a potential for memory starvation if more desktops are deployed on the single UCS blade. Therefore, the study was repeated with the number of virtual desktops on a single UCS blade server increased to 160, in conjunction with an increase in the amount of memory on a single blade to 96 GB. The number was selected due to the limit on the maximum number of virtual machines that can be executed on a single VMware ESX Server version 3.5 when the virtual CPU count is increased to 192. Consult VMware KB article 1006393 for further details. In this section we present the storage array statistics observed when a single UCS blade was executing close to the theoretical limit of virtual desktops on a single VMware ESX Server.

[Figure 24](#) shows the CPU utilization of a fully loaded UCS blade server during the gradual boot up of 160 virtual desktops by VMware View server. It can be seen from [Figure 24](#) that the peak CPU utilization approaches 80%. However, the ramp up occurs over a longer period of time when compared to the boot up process of 128 virtual desktops on a smaller memory configuration UCS blade server (see [Figure 17](#)). The figure also shows that the average CPU utilization rate during the steady state operation (70%) is higher than that observed with 128 virtual desktop (50%) indicating a better utilization of the resources provided by the UCS blade server.

Figure 24 CPU Utilization During the Ramp-Up of 160 Virtual Desktops on a Single UCS Server

The memory utilization of the UCS blade server during the ramp-up of 160 virtual desktops is shown in [Figure 25](#). As expected the utilization rate of the memory drops dramatically since the amount of memory on the UCS blade server was increased to 96 GB during this phase of the testing. It can be seen from [Figure 25](#) that the peak memory utilization does not exceed 60% as the 160 virtual desktops are booted. The transparent memory sharing reclaims memory pages and reduces the utilization to approximately 40% during the steady-state execution of the workload.

[Figure 26](#) and [Figure 27](#) shows the CPU utilization and memory utilization, during the steady state operation of the 160 virtual desktops executing the workload described earlier. The CPU utilization is approximately 70% during the period of observation with the one of the logical processor in each core performing more operations than the other. Although the memory utilization seems to fluctuate tremendously, it should be noted that the utilization stays within a narrow band between 32% to 35%. The minor changes in the memory usage can be expected as shared pages are changed by one of the virtual desktops resulting in new allocation of memory page or when the VMware kernel determines common memory page and reduces the allocation.

Figure 25 *Memory Utilization During the Ramp-Up of 160 Virtual Desktops on a Single UCS Blade Server*



Figure 26 CPU Utilization During Steady State Execution of 160 Virtual Desktops on a Single UCS Server






Figure 27 *Memory Utilization During Steady State Execution of 160 Virtual Desktops on a Single UCS Blade Server*

Storage System Performance with 160 Virtual Desktops on Single Blade


Figure 28 shows the average utilization rate of the four Symmetrix V-Max storage ports when executing the 160 virtual desktops on a single UCS blade. It can be seen from Figure 28 that the average utilization rate rarely exceed 2% on any of the Fibre Channel ports. A quick comparison with the utilization rate of the ports shown in Figure 20 indicates an insignificant change in the utilization rate between the 128 to 160 virtual desktops. The workload generated by each VMware View desktop includes simulation of breaks taken by workers during normal business hours. The duration and when these breaks occur during the execution of the workload are determined randomly. Therefore, for small increases in number of desktops simultaneously running in VMware View environment, the probability of multiple desktops sleeping at the same time increases proportionally. This results in statistically insignificant change in the I/O workload and, therefore, utilization percentages of the resources available on the storage array.

Figure 28 Utilization of the Front End Ports on Symmetrix V-Max Storage Array



The utilization rate during a boot storm event, however, is dependent on the number of virtual machines on a single server that have to be powered on. This increase can be observed in [Figure 29](#), which shows an approximately 25% increase in the utilization rate of the Fibre Channel ports on the Symmetrix V-Max storage array during the reboot of 160 virtual desktops on a single UCS blade when compared to the utilization rate of the same Fibre Channel ports observed during the reboot of 128 virtual desktops (see [Figure 21](#)). It is important to note that the increases in the utilization rate is bound by the vCenter Server that controls the number of simultaneous virtual machines that are restarted on the VMware ESX Servers managed by that instance.


Figure 29 *Utilization Rate of the Front End Ports on Symmetrix V-Max Storage Array after ESX Server Reboot*



The non-linear and linear increases discussed in the earlier paragraphs during steady state workload generation and boot storm, respectively, is obvious when the I/O rate generated on the Symmetrix V-Max storage ports is observed. This is shown in [Figure 30](#) and [Figure 31](#).


The increase in the number of deployed virtual desktops on a single UCS blade from 128 to 160 results in the average I/O rate increase of approximately 50 operations per second. The increase in the average I/O rate activity during a boot storm of 500 operations per second, however, is much more significant. Nevertheless, the recommended minimum configuration of using 4 Fibre Channel ports when connecting a VMware ESX Server cluster group to a Symmetrix V-Max storage array can easily accommodate the workload generated by the largest possible virtual desktop deployment on a single VMware ESX Server version 3.5.

Figure 30 *I/O Rate on the Front End Symmetrix V-Max Ports During Execution of Steady State Workload*



227106

Figure 31 *I/O Rate on Front End Symmetrix V-Max Ports After an Unplanned ESX Server Reboot*




227104

VMware View Deployment with 640 Desktops on Four UCS Servers

With 640 desktops across four blades, the resource utilization on every blade server was similar to that described in the 160 desktops on a single UCS blade server section. The observation clearly indicates that the UCS chassis and properly configured Symmetrix V-Max storage array can scale linearly in support of a VMware View environment.

Figure 32 shows the average CPU utilization on each blade server during the steady state execution of a workload simulating a knowledge worker (see [User Workload Simulation](#) for a detailed description of the workload). The 640 desktops during the tests were spread equally across the available servers to ensure optimal use of the available resources. It can be seen from the figure that the CPU utilization of all four blades is approximately equal and at the same levels (70%) as that discussed in the previous section.

Figure 32 *Average CPU Utilization on Four UCS Blade Server with Each Running 160 Virtual Desktops*



The average memory utilization on each blade during the steady state execution of the 640 virtual desktops is shown in Figure 33. Similar to the results discussed in Figure 27, the memory utilization hovered between 32% and 35% indicating a well balanced workload that scales linearly with increasing UCS blade servers.

Figure 33 *Average Memory Utilization During Steady State Execution of 640 Virtual Desktops*





Figure 34 reflects the average application execution time from all virtual desktops running on the UCS blade servers. The application times represent the amount of time it took to open, close, or save a document that was created. The figure represents an average of the time data that was sampled out of the number of iterations that were conducted on the virtual desktops. Figure 34 clearly shows that the performance experienced by the simulated worker (less than 5 seconds) is within the range considered as acceptable performance.

Figure 34 *Average Application Execution Time in Seconds*




227153

Storage System Performance with 640 Virtual Desktops on Four Blade UCS System

Figure 35 shows the utilization of the Fibre Channel ports during the gradual power-on operation of the 640 virtual desktop performed by VMware View servers in response to client requests for desktop sessions. It can be seen from Figure 35 that the utilization of the Fibre Channel resources rapidly increases to approximately 10% before dropping down to less than 2% after all of the 640 virtual desktops have been powered on.


Figure 35 Utilization of the Front End Port on Symmetrix V-Max Storage Array During Ramp-Up of 640 Virtual Desktops



After the power-on activity, the simulated workload is automatically initiated on each of the virtual desktop and the utilization of the Fibre Channel ports, as shown in [Figure 36](#), settles down to an average value of approximately 4%. The average utilization of the ports is significantly larger than that for 160 virtual desktop environment executing on a single blade. Once again, due to the reasons cited earlier, the utilization rate does not increase linearly with the number of virtual desktop deployment.


An unplanned or a planned event that results in a restart of all of the virtual desktops in a VMware View environment can generate intense I/O workload to the storage environments. [Figure 37](#) shows the average utilization of the Fibre Channel front end ports of the Symmetrix V-Max storage array in response to the restart of the entire VMware View environment hosting 640 virtual desktops across four UCS blades.

Figure 36 Utilization of the Front End Port on Symmetrix V-Max Storage Array During Steady State Operation of 640 Virtual Desktops



227109


Figure 37 Utilization Rate of the Front End Port on Symmetrix V-Max Storage Array When Rebooting 640 Virtual Desktops



It can be seen from the figure that although the utilization of the Fibre Channel ports increases significantly in response to a major boot storm, that a properly configured Symmetrix V-Max storage array can mitigate the impact of a major event. The figure also shows that the front-end ports allocated to the VMware View environment have enough resources to handle a fully loaded UCS chassis.


Figure 38 shows the maximum utilization of the physical disks allocated to the VMware View environment (128, 300 GB, 15 K RPM drives) during the restart of the 640 virtual desktops. It is important to note that the maximum utilization of the physical disks can occur at different times during the observation period, in this case the restart of the 640 virtual desktops. The benefit of using EMC Virtual Provisioning in VMware View environment to provide wide striping across large number of physical disks is clearly visible in Figure 38. In addition, the maximum utilization of any of the physical disk utilized in the environment does not exceed 50% indicating sufficient resources to handle the worst possible scenario in a VMware View environment.

Figure 38 *Disk Utilization Indicating Maximum Utilization of Physical Disk During Reboot of 640 Virtual Desktops*



The maximum utilization of the physical disks allocated to the VMware View environment (128, 300 GB, 15 K RPM drives) during the steady state operation of 640 virtual desktops is shown in [Figure 39](#). This is the workload one should design the storage environment for while ensuring enough head room to accommodate growth in the environment or unexpected burst in the activity due to unplanned events. [Figure 39](#) clearly shows that the maximum utilization of any of the physical disk does not exceed 30% indicating that the storage array can easily handle additional workload that would be generated from the VMware View desktops deployed on the four additional blades available in the UCS chassis.

Figure 39 *Disk Utilization Indicating Maximum Utilization of Physical Disk During Steady State Operation*



Conclusion

The architecture and large memory capabilities of the UCS system connected to a modular and scalable Symmetrix V-Max storage system enables customers to scale a VMware View deployment in ways not previously possible. During the testing, the dual quad core Xeon 5500 series processors barely reached 75% of their capacity with 160 virtual desktops running a steady-state workload similar to that anticipated from a knowledge worker. The memory utilization during the execution of the workload on a single blade was approximately 50%. The study was able to demonstrate the scalability of the UCS chassis by running 640 desktops with simulated workloads on four blades with fully-populated memory (96 GB achieved through the use of 12, quad rank, 8 GB DIMMS). The study did not reveal any stress on the chassis backplane indicating a linear scaling as the workload and the number of blades in the environment was increased from a single blade. Therefore, we do not expect any degradation in the performance if the deployment is expanded to an eight blade configuration executing 1,280 virtualized desktops. Furthermore, given the fact that UCS scales to 40 chassis or 320 UCS blade servers with a Nexus 6140 XP fabric interconnect, we clearly see a linear scalability of desktops with a large multichassis configuration.

The linear scaling exhibited by the UCS system was fully supported by the capabilities of the EMC Symmetrix V-Max system. The minimally configured Symmetrix V-Max storage array was able to accommodate the I/O workload generated by the largest deployment configured in this study. The utilization of the Fibre Channel front end ports of the EMC Symmetrix V-Max storage array rarely exceeds 10% during steady state workload processing and 30% during the boot storm that occurs after an ESX Server reboot. The average utilization of the physical disks configured to support the VMware View deployment did not exceed 30% during the execution of the steady state workload and 50% during the boot storm event referred to earlier. The non-linear increase in the physical disk utilization rate from steady state to an unplanned event, such as VMware ESX Server reboot, highlights the capability of the Symmetrix V-Max storage system to use its large cache to buffer unexpected, short duration spikes in the I/O workload. The low utilization rates seen on the storage system also clearly showed that increasing the size of the deployment from 640 virtualized desktop to 1,280 virtualized desktop would be easily handled by the Symmetrix V-Max system.

Furthermore, the study clearly showed that a multichassis UCS environment, with appropriately designed VMware Virtual Infrastructure environment and properly configured Symmetrix V-Max storage system, can be anticipated to host in excess of several 1000s of virtual desktops.

References

- VMware View Reference Architecture
<http://www.vmware.com/resources/techresources/1084>
- VMware 3.0
<http://www.vmware.com/products/view/>
- Cisco UCS
<http://www.cisco.com/go/unifiedcomputing>
- Cisco Data Center Solutions
<http://www.cisco.com/go/datacenter>
- Cisco Validated Designs
<http://www.cisco.com/go/designzone>
- EMC Symmetrix V-Max System
<http://www.emc.com/products/detail/hardware/symmetrix-v-max.htm>

- EMC Symmetrix V-Max System and VMware Virtual Infrastructure
<http://www.emc.com/collateral/hardware/white-papers/h6209-symmetrix-v-max-vmware-virtual-infrastructure-wp.pdf>
- Using EMC Symmetrix Storage in VMware Virtual Infrastructure
<http://www.emc.com/collateral/hardware/solution-overview/h2529-vmware-esx-svr-w-symmetrix-wp-ldv.pdf>



VMware, Inc.
 3401 Hillview Ave
 Palo Alto, CA 94304 USA
 Tel: 650-427-5000 or 877-486-9273
 Fax: 650-427-5001
www.vmware.com



Cisco Systems, Inc.
 170 West Tasman Drive
 San Jose, CA 95134 USA
 Tel: 408-526-4000 or 800-553-6387 (NETS)
 Fax: 408-527-0883
www.cisco.com



EMC Corporation
 176 South Street
 Hopkinton, MA 01748 USA
 Tel: 508-435-1000
www.emc.com

Copyright © 2009. VMware, Inc. All rights reserved. Protected by one or more U.S. Patent Nos. 6,397,242, 6,496,847, 6,704,925, 6,711,672, 6,725,289, 6,735,601, 6,785,886, 6,789,156, 6,795,966, 6,880,022, 6,944,699, 6,961,806, 6,961,941, 7,069,413, 7,082,598, 7,089,377, 7,111,086, 7,111,145, 7,117,481, 7,149, 843, 7,155,558, 7,222,221, 7,260,815, 7,260,820, 7,269,683, 7,275,136, 7,277,998, 7,277,999, 7,278,030, 7,281,102, 7,290,253, 7,356,679 and patents pending.

EMC believes the information in this publication is accurate as of its publication date. The information is subject to change without notice. VMware and Vmotion are registered trademarks of VMware, Inc. For the most up-to-date listing of EMC product names, see EMC Corporation Trademarks on EMC.com. All other trademarks used herein are the property of their respective owners. © Copyright 2009 EMC Corporation. All rights reserved.

© 2009 Cisco Systems, Inc. All rights reserved. Cisco, the Cisco logo, and Cisco Systems are registered trademarks or trademarks of Cisco Systems, Inc. and/or its affiliates in the United States and certain other countries. All other trademarks mentioned in this document or Website are the property of their respective owners. The use of the word partner does not imply a partnership relationship between Cisco and any other company. (0805R) 06/09