Cisco UCS and Application Delivery for Microsoft Hyper-V Virtualization of Exchange 2010 with NetApp Storage

Last Updated: July 13, 2010

Cisco Validated Design

Building Architectures to Solve Business Problems

About the Authors

Karen Chan

Karen Chan, Solutions Architect, CMO Enterprise Solutions Engineering (ESE), Cisco Systems

Karen is currently a Data Center Solutions Architect in the Enterprise Solutions Engineering team at Cisco. Prior to this role, she was an Education Architect in the Industry Solutions Engineering team at Cisco. She has also worked as a technical marketing engineer in the Retail group on the Cisco PCI for Retail solution as well as in the Financial Services group on the Cisco Digital Image Management solution. Prior to Cisco, she spent 11 years in software development and testing, including leading various test teams at Citrix Systems, Orbital Data, and Packeteer, to validate software and hardware functions on their WAN optimization and application acceleration products. She holds a bachelor of science in electrical engineering and computer science from the University of California, Berkeley.

Brad Reynolds

Brad Reynolds, Technical Marketing Engineer, CMO Enterprise Solutions Engineering (ESE), Cisco Systems

Brad is a Technical Marketing Engineer for Data Center technologies on Cisco's Enterprise Solutions Engineering (ESE) team. He is currently focused on application solution validation based on virtualization technologies such as VMWare and HyperV. Brad has been at Cisco for over 9 years and prior to joining the ESE Datacenter team he was a TME for the CMO Demonstrations team and was the demonstrations lead engineer focusing on Security, Mobility and Service Provider Demonstrations. He has over 6 years experience in creating compelling demonstrations of Cisco's technologies starting with his first position at Cisco in the Executive Briefing Center in 2001. In the EBC Brad's team evolved in to not only supporting SP demonstrations for the EBC but also for events and tradeshows. Brad also spent several years focusing on WLAN and Mobility and was a contributor and co-author to technical papers including Mobility Design Guides, and the Mobile Access Router and Mesh Networks Design Guide.

Cisco UCS and Application Delivery for Microsoft Hyper-V Virtualization of Exchange 2010 with NetApp Storage

Introduction

This design guide presents an end-to-end solution architecture that demonstrates how enterprises can virtualize their Exchange 2010 environment on Cisco Unified Computing System. This is accomplished by using Microsoft Hyper-V virtualization and NetApp storage in a Cisco Data Center 3.0 network architecture with application optimization services provided by Cisco Application Control Engine and Cisco Wide Area Application Services.

As enterprises face the critical problem of growing their data center when they are already reaching maximum capacity for space, cabling, power, cooling, and management, they should consider consolidating and virtualizing their application servers into their data center. At the top of the list of servers to consolidate and virtualize are those mission-critical enterprise applications that typically require deployment of a large number of physical servers. Microsoft Exchange 2010 is one such application as it requires four different server roles for a complete system and more often requires at least double that number of servers to provide minimal redundancy in the server roles. A virtualization infrastructure design that supports an enterprise-level Exchange 2010 deployment also provides the necessary services to address manageability of solution components, disaster recovery, high availability, rapid provisioning, and application optimization and delivery across the Wide Area Network.

Audience

This document is primarily intended for enterprises interested in virtualizing their Exchange 2010 servers using Microsoft Hyper-V on the Cisco Unified Computing System. This document is also directed at readers who would like to understand the issues that were considered in the design of the end-to-end solution architecture and those who would like to better understand the features and architectures implemented in the Cisco, Microsoft, and NetApp solution components.

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

Document Objectives

This document provides guidance on how to design a data center architecture that provides network, server, and application-level services that are needed when an enterprise is trying to consolidate and virtualize its Exchange 2010 server farm on the Cisco Unified Computing System using Microsoft Hyper-V. CXO and architect-level decision makers can use the details of the solution components in this design guide to implement this solution in their enterprise. The following categories of information are provided:

- Network topology
- Exchange 2010 application architecture
- Data Center 3.0 network architecture
- Microsoft Hyper-V features and architecture
- Cisco Wide Area Application Services for Exchange 2010
- Cisco Application Control Engine for Exchange 2010
- Bill of Materials
- URLs to other relevant design and deployment guides

Solution Architecture

Architecture Overview

This solution provides an end-to-end architecture with Cisco, Microsoft, and NetApp technologies that addresses challenges that IT organizations encounter in various places in their network when they are virtualizing their Microsoft Exchange server farms, whether it is within their server farms at the server platform and application layers or within their remote branch offices. The issues that must be addressed are not limited to the server hardware or the virtualization platform; they extend beyond the data center server farm through the data center LAN and SAN to the branch and remote office locations. As more servers are consolidated at the data center, there are questions of how to manage the growing number of application servers and physical servers and how to deliver applications to users at the branch offices or remote locations given the bandwidth limitations on WAN links. In addition, it is important that the architecture can support the addition of security and other services like QoS to the data center, branch, and remote sites as network requirements evolve.

Architecture Components

Exchange 2010 Application Architecture

Exchange 2010 Server and Client Scenario for This Solution

While Exchange deployments come in different combinations and sizes, the basic methodology and best practices documented in this solution can be applied.

A sample Exchange deployment was defined to validate the design and deployment guidelines presented in this Cisco Validated Design. Refer to Figure 1 as you read this section about the details of this scenario.

The Exchange servers and clients deployed consisted of the following:

- All Exchange server roles are virtualized on Hyper-V. A pair of client access server and hub transport server virtual machines reside on one blade on one UCS chassis, while the other pair resides on another blade on the second UCS chassis.
  - One Database Availability Group (DAG) consisting of four mailbox servers for redundancy and failover.
  - There were a total of eight database copies in the DAG. Each mailbox server hosts two databases total: its own primary active database and a redundant passive copy of the active database from another server.
  - Microsoft Windows 2008 R2 with Active Directory and DNS roles enabled deployed as a virtual machine on its own blade.
- 5000 mailboxes to support 5000 campus users simulated with Microsoft's Exchange LoadGen 2010 Beta tool. The user profile is 100 messages sent/received a day (where message size is about 75 KB) and the client type is Outlook 2007 online mode.
- Four Outlook 2010 clients located in the branch office were used to test Cisco Wide Area Acceleration Services across the WAN. These four clients were configured with Outlook 2010 MAPI/RPC for one scenario and then configured with Outlook Anywhere for the second scenario.
- A single Outlook Web Access remote user was used to test traffic optimization provided by the Cisco WAAS Mobile server and client software for optimizing HTTPS traffic.
- A single Outlook 2010 user configured with Outlook Anywhere was used to test traffic optimization provided by Cisco WAAS Mobile for optimization of RPC/HTTPS traffic.

Figure 1 shows the placement of virtual machines on the Cisco Unified Computing System blades.

Figure 1 Hyper-V Virtual Machines on Cisco UCS Blades

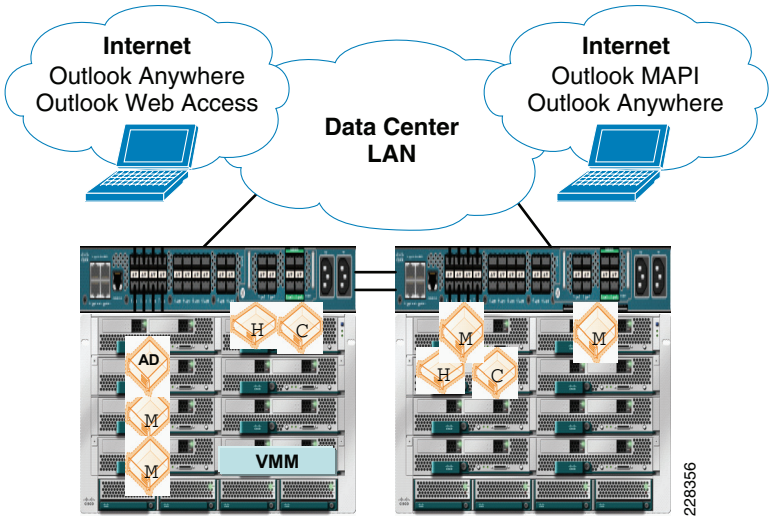

Legend for Figure 1: C = Client Access Server, H = Hub Transport Server, M = Mailbox Server, AD = Active Directory, VMM = Hyper-V Virtual Machine Manager

Server Roles Overview

This solution involved implementing three of the following five Exchange 2010 server roles as Hyper-V virtual machines on the Cisco Unified Computing System:

- Hub Transport Server
- Client Access Server
- Mailbox Server
- Edge Transport Server[1]
- Unified Messaging Server[2]

Hub Transport Server Role

For those familiar with earlier versions of Exchange Server 2007, the Hub Transport server role replaces what was formerly known as the bridgehead server in Exchange Server 2003. The function of the Hub Transport server is to intelligently route messages within an Exchange Server 2010 environment. By default, SMTP transport is very inefficient at routing messages to multiple recipients because it takes a message and sends multiple copies throughout an organization. As an example, if a message with a 5 MB attachment is sent to 10 recipients in an SMTP network, typically at the sendmail routing server, the 10 recipients are identified from the directory and 10 individual 5 MB messages are transmitted from the sendmail server to the mail recipients, even if all of the recipients' mailboxes reside on a single server.

The Hub Transport server takes a message destined to multiple recipients, identifies the most efficient route to send the message, and keeps the message intact for multiple recipients to the most appropriate endpoint. Hence, if all of the recipients are on a single server in a remote location, only one copy of the 5 MB message is transmitted to the remote server. At that server, the message is then broken apart with a copy of the message dropped into each of the recipient's mailboxes at the endpoint.

The Hub Transport server in Exchange Server 2010 does more than just intelligent bridgehead routing; it also acts as the policy compliance management server. Policies can be configured in Exchange Server 2010 so that after a message is filtered for spam and viruses, the message goes to the policy server to be assessed whether the message meets or fits into any regulated message policy and appropriate actions are taken. The same is true for outbound messages; the messages go to the policy server, the content of the message is analyzed, and if the message is determined to meet specific message policy criteria, the message can be routed unchanged or the message can be held or modified based on the policy. As an example, an organization might want any communications referencing a specific product code name or a message that has content that looks like private health information, such as Social Security number, date of birth, or health records of an individual, to be held or encryption to be enforced on the message before it continues its route.

1. Since the Edge Transport role is an optional role not mandatory for an Exchange deployment, it was not included in this solution. For guidance on deploying an Edge Transport server in a Data Center 3.0 architecture with ACE load balancing, see: http://www.cisco.com/en/US/docs/solutions/Verticals/mstdcmsftex.html#wp623892.
2. Unified Messaging was not included in this solution since this solution focuses on the design and deployment of Exchange 2010 E-mail services.

Client Access Server Role

The Client Access Server role in Exchange Server 2010 (as was also the case in Exchange Server 2007) performs many of the tasks that were formerly performed by the Exchange Server 2003 front-end server, such as providing a connecting point for client systems. A client system can be an Office Outlook client, a Windows Mobile handheld device, a connecting point for OWA, or a remote laptop user using Outlook Anywhere to perform an encrypted synchronization of their mailbox content.

Unlike a front-end server in Exchange Server 2003 that effectively just passed user communications on to the back-end Mailbox server, the CAS does intelligent assessment of where a user's mailbox resides and then provides the appropriate access and connectivity. This is because Exchange Server 2010 now has replicated mailbox technology where a user's mailbox can be active on a different server in the event of a primary mailbox server failure. By allowing the CAS server to redirect the user to the appropriate destination, there is more flexibility in providing redundancy and recoverability of mailbox access in the event of a system failure.

Mailbox Server Role

The Mailbox server role is merely a server that holds users' mailbox information. It is the server that has the Exchange Server EDB databases. However, rather than just being a database server, the Exchange Server 2010 Mailbox server role can be configured to perform several functions that keep the mailbox data online and replicated. For organizations that want to create high availability for Exchange Server data, the Mailbox server role systems would likely be clustered, and not just a local cluster with a shared drive (and, thus, a single point of failure on the data), but rather one that uses the new Exchange Server 2010 Database Availability Groups. The Database Availability Group allows the Exchange Server to replicate data transactions between Mailbox servers within a single-site data center or across several data centers are multiple sites. In the event of a primary Mailbox server failure, the secondary data source can be activated on a redundant server with a second copy of the data intact. Downtime and loss of data can be drastically minimized, if not completely eliminated, with the ability to replicate mailbox data on a real-time basis.

Microsoft eliminated single copy clusters, Local Continuous Replication, Clustered Continuous Replication, and Standby Continuous Replication in Exchange 2010 and substituted in their place Database Availability Group (DAG) replication technology. The DAG is effectively CCR, but instead of a single active and single passive copy of the database, DAG provides up to 16 copies of the database and provides a staging failover of data from primary to replica copies of the mail. DAGs still use log shipping as the method of replication of information between servers. Log shipping means that the 1 MB log files that note the information written to an Exchange server are transferred to other servers and the logs are replayed on that server to build up the content of the replica system from data known to be accurate. If during a replication cycle a log file does not completely transfer to the remote system, individual log transactions are backed out of the replicated system and the information is re-sent.

Unlike bit-level transfers of data between source and destination used in Storage Area Networks (SANs) or most other Exchange Server database replication solutions, if a system fails, bits do not transfer and Exchange Server has no idea what the bits were, what to request for a resend of data, or how to notify an administrator what file or content the bits referenced. Microsoft's implementation of log shipping provides organizations with a clean method of knowing what was replicated and what was not. In addition, log shipping is done with small 1 MB log files to reduce bandwidth consumption of Exchange Server 2010 replication traffic. Other uses of the DAG include staging the replication of data so that a third or fourth copy of the replica resides "offline" in a remote data center; instead of having the data center actively be a failover destination, the remote location can be used to simply be the point where data is backed up to tape or a location where data can be recovered if a catastrophic enterprise environment failure occurs.

A major architecture change with Exchange Server 2010 is how Outlook clients connect to Exchange Server. In previous versions of Exchange Server, even Exchange Server 2007, RPC/HTTP and RPC/HTTPS clients would initially connect to the Exchange Server front end or Client Access Server to reach the Mailbox servers while internal MAPI clients would connect directly to their Mailbox Server. With Exchange Server 2010, all communications (initial connection and ongoing MAPI communications) go s through the Client Access Server, regardless of whether the user was internal or external. Therefore, architecturally, the Client Access Server in Exchange Server 2010 needs to be close to the Mailbox server and a high-speed connection should exist between the servers for optimum performance.

Edge Transport Server Role

The Edge Transport server role is a dedicated server function that performs spam and virus filtering as the first point of entry of messages into an Exchange Server environment. Rather than having unwanted messages go directly to an Exchange Server back-end server taxing the database server with filtering of messages, the Edge Transport server offloads this task. Rather than spam and virus filtering thousands of messages for recipients who do not exist in the environment, the Edge Transport server accesses and stores a copy of certain Active Directory data, such as all valid Exchange Server 2010 E-mail recipient mail addresses. This is so incoming messages can be checked against this Active Directory Application Mode (ADAM) directory and messages for recipients who do not exist in the organization are immediately deleted. In addition, the Edge Transport server role performs a safelist aggregation function as well, where it gathers safelist information from Outlook clients and brings the safelists out to the edge. By bringing individual users' safelists to the edge, the Edge Transport server can now take into consideration individual user preferences to receive certain types of messages so that the messages are forwarded on to the recipient even if the Edge Server's antispam software would normally quarantine or delete the message. After a first pass at deleting messages for nonexistent recipients, a spam and virus scan can then be run on the remaining messages. After being determined to be clean, the messages can be forwarded on to either a Hub Transport server or to an Exchange Server 2010 back-end server.

Data Center Server Platform—Cisco Unified Computing System

Cisco Unified Computing System Overview

Today, IT organizations assemble their data center environments from individual components. Their administrators spend significant amounts of time manually accomplishing basic integration tasks rather than focusing on more strategic, proactive initiatives. The industry is in a transition away from the rigid, inflexible platforms that result and moving toward more flexible, integrated, and virtualized environments.

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

Managed as a single system whether it has one server or 320 servers with thousands of virtual machines, the Cisco Unified Computing System decouples scale from complexity. The Cisco Unified Computing System accelerates the delivery of new services simply, reliably, and securely through end-to-end provisioning and migration support for both virtualized and nonvirtualized systems. It provides the following benefits:

- Embedded system management—Management is uniquely integrated into all the components of the system, enabling the entire solution to be managed as a single entity through Cisco UCS Manager. Cisco UCS Manager provides an intuitive GUI, a command-line interface (CLI), and a robust API to manage all system configuration and operations. Cisco UCS Manager enables IT managers of storage, networking, and servers to collaborate easily on defining service profiles for applications.
- Just-in-time provisioning with service profiles—Cisco UCS Manager implements role- and policy-based management using service profiles and templates. Infrastructure policies—such as power and cooling, security, identity, hardware health, and Ethernet and storage networking—needed to deploy applications are encapsulated in the service profile. 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.
- Unified fabric—Cisco's unified fabric technology reduces cost by eliminating the need for multiple sets of adapters, cables, and switches for LANs, SANs, and high-performance computing networks. The system's fabric extenders pass all network traffic to parent fabric interconnects, where it can be processed and managed centrally, improving performance and reducing points of management. The unified fabric is a low-latency lossless 10-Gbps Ethernet foundation that enables a "wireonce" deployment model in which changing I/O configurations no longer means installing adapters and recabling racks and switches.
- State-of-the-art performance—Intel® Xeon® 5500 series processors automatically and intelligently adjust server performance according to application needs, increasing performance when needed and achieving substantial energy savings when not.

  Performance and power settings can also be manually configured.
- Energy efficiency—The system is designed for energy efficiency. Power supplies are 92 percent efficient and the Intel Xeon 5500 series processors use automated low-power states to better match power consumption with workloads. The simplified design of the Cisco UCS B-Series Blade Servers improves airflow efficiency and can reduce the number of components that need to be powered and cooled by more than 50 percent compared to traditional blade server environments; similar component reduction can be achieved with the Cisco UCS C-Series Rack-Mount Servers.

Cisco Unified Computing System Component Details

Figure 2 Cisco Unified Computing System

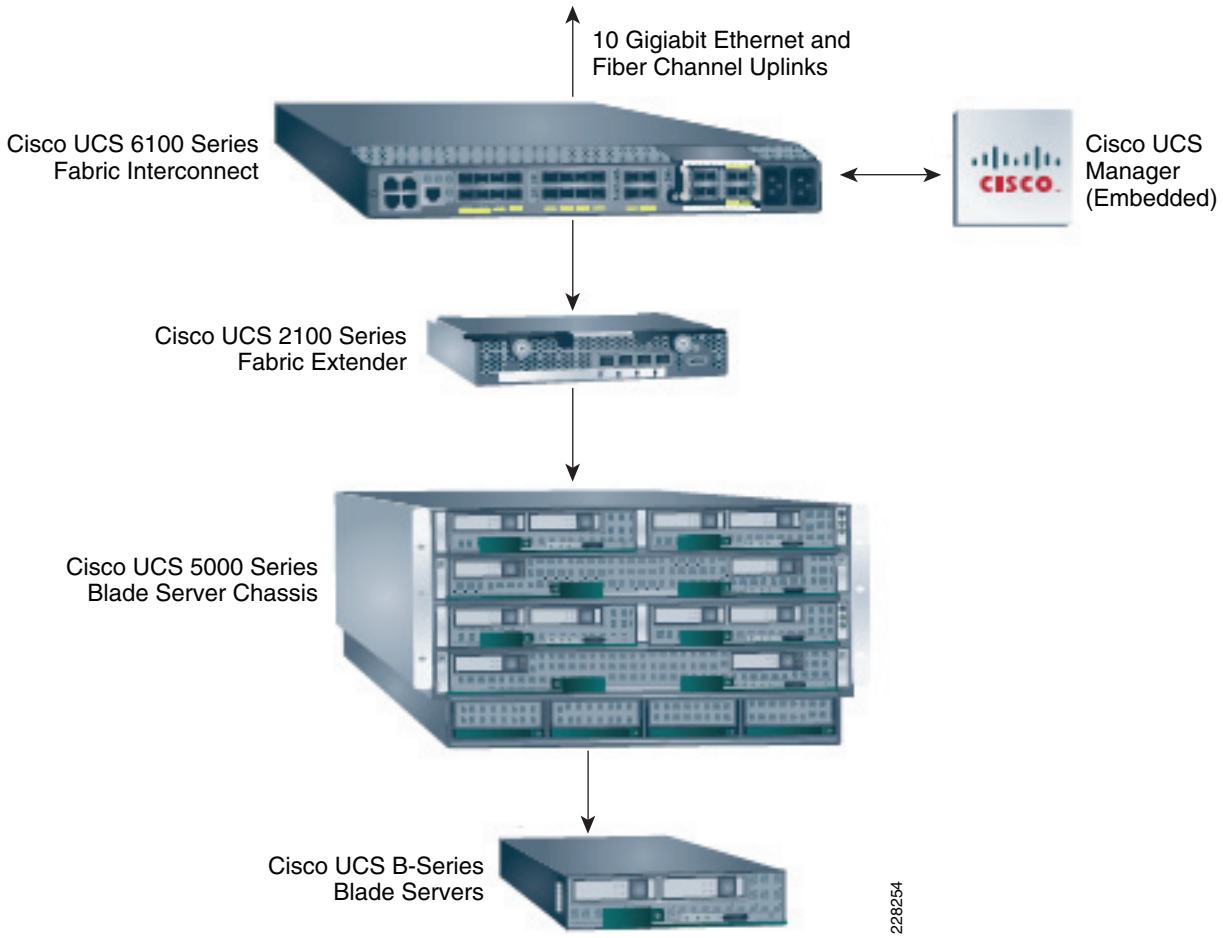

The main system components include:

- 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. This design uses the UCS B200-M1 half-slot two socket blade servers, which has 12 DIMM slots (up to 96GB) per blade server and one dual-port Converged Network Adapter (CNA). Each UCS B200-M1 also has the following system features:
  - Two Intel Xeon 5500 series processors (quad cores)
  - Two optional SAS/SATA 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

I/O Adapters

The blade server has various Converged Network Adapters (CNA) options. The following two CNA options were used in this Cisco Validated Design:

- 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 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 supports up to 40 chassis or 320 servers with half 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

In this solution, two UCS 6120 Fabric Interconnects were configured in a cluster pair for redundancy.

UCS Service Profiles for Quick Provisioning and Server Portability

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 blade server 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.

Service Profiles and SAN Boot for Server Portability

UCS Service Profiles combined with Windows SAN boot allow for application servers to be moved from server blade to server blade for quick provisioning and rapid recovery in cases of hardware or network failure. When a new server needs to be brought online quickly, an existing service profile can be associated with the new blade hardware without having to move any internal disks and without having to change the zoning or LUN masking of the SAN or VSANs and VLANs. The parameters set in the service profile can remain unchanged as well.

The deployment of an architecture consisting of SAN booted physical resources provides great flexibility and resiliency to a multi-tenant infrastructure. A SAN booted deployment consists of hosts in the environment having a converged network adapter (CNA) capable of translating SCSI commands via fibre channel (or FCoE, which is not in scope for this Cisco Validated Design). Hosts then access their boot OS via logical unit number (LUN) or storage container mapped on an external storage array. This boot methodology can be accomplished with software or hardware initiators and, for the purposes of this document, local HBAs are discussed. When using NetApp controllers, SAN booted hosts have superior RAID protection and increased performance when compared to traditional local disk arrays. Furthermore, SAN booted resources can easily be recovered, are better utilized, and scale much quicker than local disk installs. Operating systems and hypervisors provisioned via NetApp storage controllers take advantage of storage efficiencies such as Thin Provisioning, Deduplication, and Snapshot Technology inherent in NetApp products. Another major benefit of SAN-booted architectures is that they can be deployed and recovered in minutes dependent on the OS to be installed.

SAN booted resources can easily be recovered, are better utilized, and scale much quicker than local disk installs. Operating systems and hypervisors provisioned on the SAN can take advantage of storage efficiencies and data protection provide by storage vendors like NetApp. Note that other storage vendor can be used to provide the SAN support for this solution. NetApp was chosen because of its prevalence in the industry.

Design Considerations for Exchange 2010 Virtualization on UCS

This section covers information that is useful when determining which Exchange roles to virtualize, how many virtual machines you need, and whether there are enough hardware resources to support the virtual deployment.

Virtualization Software

Microsoft supports virtualization of Exchange 2010 on hardware virtualization software, with certain caveats:

- The hardware virtualization software is running Windows Server 2008 or 2008 R2 with Hyper-V technology, or Microsoft Hyper-V Server 2008/2008 R2, or any other third-part hypervisor that has been validated under the "Windows Server Virtualization Validation Program." For additional details on this program, see: http://www.windowsservercatalog.com/svvp.aspx?svvppage=svvp.htm.

Roles That Can be Virtualized

- The Unified Messaging role cannot be virtualized because of its real-time response requirements for voice communications. All of the other roles can be virtualized. Note that Unified Messaging is outside the scope of this solution.
- The Edge Transport role can be run on a Unified Computing System B-Series or C-Series server. Since this server role belongs in the DMZ, VLANs would be used for an isolated connection to the DMZ. The role can run in a VM or on a physical server. Note that the Edge Transport role was not validated in this solution.
- Depending on the number of mailboxes and the user load profile required, it may be possible to support a multiple role server configuration that is virtualized. Multiple role configurations are preferable for smaller deployments to simplify server management. The following TechNet posting on "Understanding the Impact of Multiple Role Configurations on Capacity Planning" gives more details on how CPU and memory requirements of multiple role deployments: http://technet.microsoft.com/en-us/library/dd298121.aspx#WM.
- Microsoft does not support combining Exchange high availability solutions (database availability groups (DAGs)) with hypervisor-based clustering, high availability, or migration solutions. As a result, Hyper-V live migration of the VMs can only be used for the other Exchange roles: CAS, Edge Transport, and Hub Transport. Also, if live migration of CAS, Edge Transport, and/or Hub Transport VMs is desired, be sure you do not install these roles with the mailbox role on any given VM.
- Active Directory and DNS servers can be virtualized, taking into account the CPU and memory requirements to determine how may domain controller VM instances you will need. Microsoft provides performance and scaling information for Active Directory based on the processor type and number of Exchange users that need to be supported.
  - For information on how to design an Active Directory deployment, see: http://technet.microsoft.com/en-us/library/bb123715.aspx.
  - For information on specific sizing information to determine the number of servers needed for certain capacity requirements (the information applies to both Windows 2003 Active Directory and Windows 2008 Active Directory), see: http://technet.microsoft.com/en-us/library/cc728303%28WS.10%29.aspx.

  The scope of this solution did not entail a detailed exploration of Active Directory deployment. Active Directory was installed as a virtual machine for the WAAS and ACE testing. For the 5000 Exchange mailbox testing, Active Directory was installed natively on a separate UCS blade to utilize

up to eight cores as recommended by Microsoft for the mailbox core to Active Directory core ratio of 8:1, since eight mailbox cores were implemented in that 5000 mailbox test. Hyper-V limits the number of cores you can assign to a given VM to four.

Hyper-V Host Server

- Only management software (for example, antivirus software, backup software, or virtual machine management software) should be deployed on host servers. No other server-based applications (for example, Exchange, Microsoft SQL Server, Active Directory, or SAP) should be installed on the host server. The host servers should be dedicated to running guest virtual machines.
- Some hypervisors include features for taking snapshots of virtual machines. Virtual machine snapshots capture the state of a virtual machine while it is running. This feature enables you to take multiple snapshots of a virtual machine and then revert the virtual machine to any of the previous states by applying a snapshot to the virtual machine. However, virtual machine snapshots are not application aware and using them can have unintended and unexpected consequences for a server application that maintains state data, such as Exchange. As a result, making virtual machine snapshots of an Exchange guest virtual machine is not supported.

Server Processor Requirements

Hyper-V allows you to decide how many virtual processors you want to assign to a guest operating system. A virtual processor does not necessarily have to correspond to a physical processor or to a physical CPU core.

Different operating systems installed on a virtual machine have different requirements for the number of virtual processors they can support. Windows Server 2008 64-bit can take up to four virtual processors.

When determining how many virtual processors to assign to each Exchange server role (e.g., Client Access Server, Hub Transport Server, and Mailbox Server), Microsoft recommendations are based on using the number of mailbox server cores as a basis (Table 1).

Table 1 Virtual Processors Assigned to Each Exchange Server Role

Role:Role	Core Ratio
Mailbox Role: Client Access Server Role	4:3
Mailbox Role: Hub Transport Server Role with Antivirus	5:1
Mailbox Role: Active Directory Server	8:1

Microsoft recommends that server core ratios are maintained as closely as possible when determining how many instances of each role to deploy, though it is expected that only the largest deployments will be able to scale out their number of roles to match the ratios exactly. For more detailed information on calculating server core ratios for the different Exchange server roles, see Microsoft documentation at: http://technet.microsoft.com/en-us/library/dd346701.aspx.

For further details and more examples of calculating processor capacity for your mailbox servers, refer to Microsoft documentation at: http://technet.microsoft.com/en-us/library/ee712771.aspx. There is also a Microsoft Excel spreadsheet that allows you to input the specifics of your deployment and calculates the processor, memory, and storage requirements for your mailbox servers; that calculator is available at: http://msexchangeteam.com/archive/2009/11/09/453117.aspx.

Memory Requirements

There are several components to consider when determining how many virtual machines you can support on a given Hyper-V server and how much memory needs to be allocated for a given Exchange virtual machine.

The first set of considerations concern how much memory must be left available for the root partition to use to support the virtual machines and the hypervisor:

- There needs to be enough memory to accommodate the host operating system running on the root partition itself, which for Windows 2008 R2 is a minimum of 512 MB.
- There needs to be 300 MB of memory for running the hypervisor.
- There needs to be 32 MB of memory to support the first GB of RAM allocated to a virtual machine and an additional 8 MB of memory for every extra GB of RAM allocated to that virtual machine. (e.g., For a virtual machine given 8 GB of RAM would actually require a total of 8 GB + 32 MB + (8MB * 7) or 88 MB in addition to the 8 GB assigned to the virtual machine.

The second set of considerations concern how much memory is required by each Exchange server role:

- Microsoft recommends that 2 GB of RAM be allocated for each processor core assigned to the Client Access server.
- Microsoft recommends that 1 GB of RAM be allocated for each processor core assigned to the Hub Transport server.
- For the Mailbox server role, the determination takes a little more work. It involves considering the load profile for the user mailboxes, the database cache size required for that load profile (see Table 2), and the amount of RAM needed to support the database cache size (see Table 3).
- For example, for a 150 messages/day load profile, 9 MB of database cache per mailbox is required. For a server hosting up to 2500 active mailboxes, 2500 * 9 or 21.97GB of is required for the total database cache. The closest physical memory configuration would be 24 GB that should be allocated to that server.

Table 2 Database Cache Size for Given Mailbox Load Profile

Messages sent/received per mailbox per day	Database cache per user (MB)
50	3
100	6
150	9

Table 3 Server RAM for Given Database Cache Size

Server physical memory (RAM)	Database cache size: (Mailbox role only)
2GB	512 MB
4GB	1 GB
8GB	3.6 GB
16GB	10.4 GB
24GB	17.6 GB
32GB	24.4 GB
48GB	39.2 GB

For detailed information on how to size memory on the different Exchange server roles, see: http://technet.microsoft.com/en-us/library/dd346700.aspx.

Virtual Processor Allocation for Exchange Servers

Mailbox Servers

We first have to figure out how many virtual processors to assign to each mailbox server virtual machine. Based on the design considerations for the number of virtual processors to assign to each virtual machine, the following calculation was done in the lab:

1. Since the UCS 5108 B200-M1 blade has two Intel Xeon 5570 2.93GHz quad-core processors, the number of megacycles per core was determined to be 6600 megacycles.
2. The number of mailboxes maximum that can be active on one mailbox server (after failover) is 2500 (i.e., 5000 mailboxes divided by four mailbox servers in the DAG = 1250; 1250 x 2 = 2500 active and passive mailboxes total.)
3. Microsoft's LoadGen 2010 Beta load generator for Exchange 2010 was used for load testing the servers in the lab. The Outlook 2007 Online user profile was selected. It was discovered that the "Heavy" user load profile for Outlook 2007 online generated 138 tasks per user day. Rounding this up, the 150 messages/day was used to determine that each user mailbox would require three megacycles.
4. If a DAG node failure require the passive database on this server to be active, this server no longer hosts any passive mailboxes so no megacycles need to be added for this.
5. To calculate the total number of megacycles needed to support the maximum number of active mailboxes after failover: 2500 active mailboxes maximum multiplied by three (for 150 msgs/day profile) gives 7500. Since there is one passive copy of this server's primary database on another server, multiple that number by 1.10 (to add 10%) to give 8250.
6. We do not need to calculate the total number of megacycles to support the number of passive databases due to 4. above.
7. The total required megacycles for this server is 8250.
8. We start off planning to assign two virtual processors to this mailbox server. 2 x 6600 (number of megacycles per core) gives 13200. Divide 8250 from 7. above by 13200 to give 62.5% peak CPU utilization on this server after failover has occurred. This is acceptable, so we can assign two virtual processors to each mailbox server of the four mode DAG.

Client Access Servers

A 4:3 mailbox to client access server core ratio is recommended by Microsoft. Since there are a total of eight mailbox cores (four servers with two cores each), Microsoft would recommend six client access cores or three Client Access Servers with two virtual processors each. For this solution validation that involved no more than 5000 mailbox users and a user profile of 100 messages sent/received per day, two Client Access servers with two virtual processors were implemented as that was shown to be sufficient to handle the client load.

Configuring each of the two Client access servers we need for server redundancy with two virtual processors gives us the flexibility to load balancing between them with the ACE hardware load balancer or to assign connections strictly to a single client access server. For ease of testing, LoadGen also can be used to load balance client sessions between both Client Access servers.

That was shown to be sufficient for this 5000 mailbox test. Depending on the performance monitoring and other management tools running on the Client Access server, a LoadGen test should be run on the configuration to determine whether the processor count assignment is sufficient for CPU utilization.

Hub Transport Servers

Since the recommended Mailbox server core to Hub Transport server core ratio is 5:1, a total of two Hub Transport server cores are recommended to support the eight Mailbox server cores. For server redundancy, it is recommended that two Hub Transport servers be deployed.

Validating Processor Utilization on Exchange Servers

After the design considerations discussed above were applied, LoadGen 2010 Beta was used to generate a Heavy usage profile for the 5000 mailbox users, distributed across the four mailbox nodes of the DAG. Windows Performance Monitoring was run on each server role to validate CPU utilization.

The Performance Monitoring Tool is available through Exchange Management Console on each of the Exchange servers. Performance counters to monitor RPC latency and outstanding requests, message recipients delivery and submission statistics, for example, can be enabled through the Performance Monitor to give a more detailed view of how well the server is handling the client and mail submissions. See the following Microsoft documentation on further details on how to user Windows 2008 Performance Monitoring: http://technet.microsoft.com/en-us/library/cc770309(WS.10).aspx.

See Figure 3 and Figure 4 for the LoadGen setup used.

Figure 3 LoadGen Setup—1

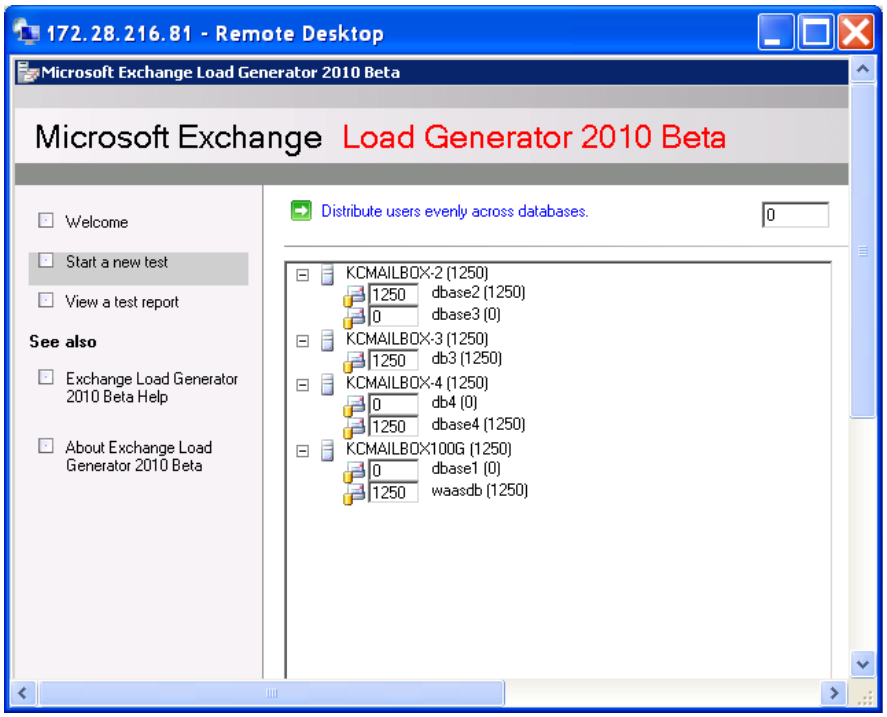

Figure 4 LoadGen Setup—2

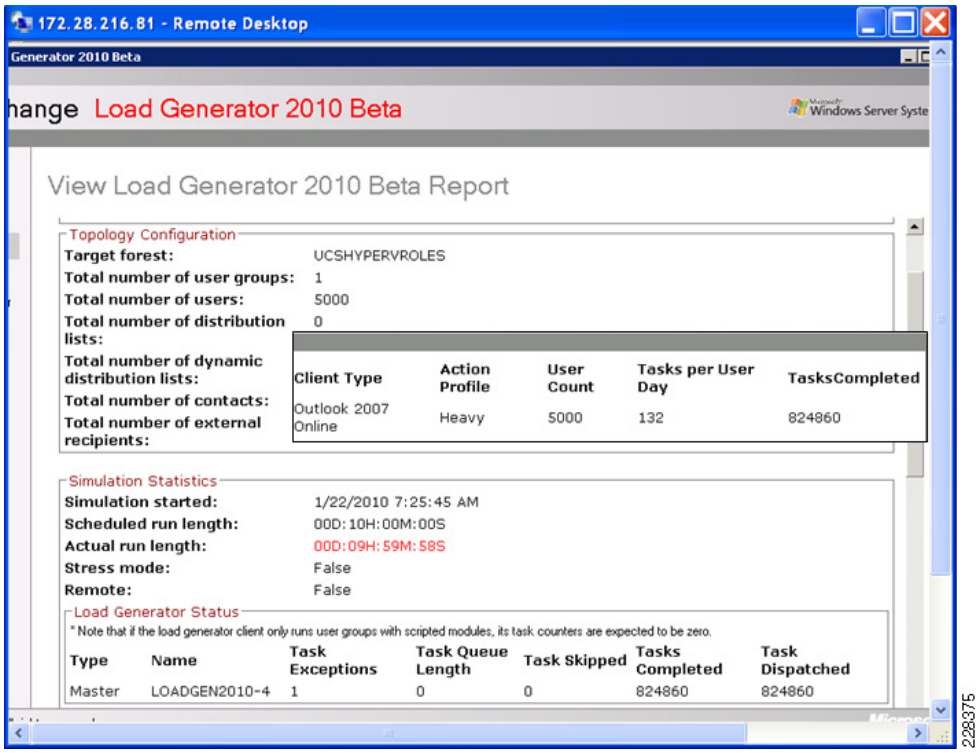

Figure 5 shows CPU utilization captured for Performance Monitor on one of the mailbox servers when only one of its two database copies is active. The graphs shows CPU utilization hovering around 30-40% on average with peaks around 75%, providing enough capacity for activating the passive mailboxes in a DAG node failover.

Figure 5 CPU Utilization on a Mail Server with One Active and One Passive Database

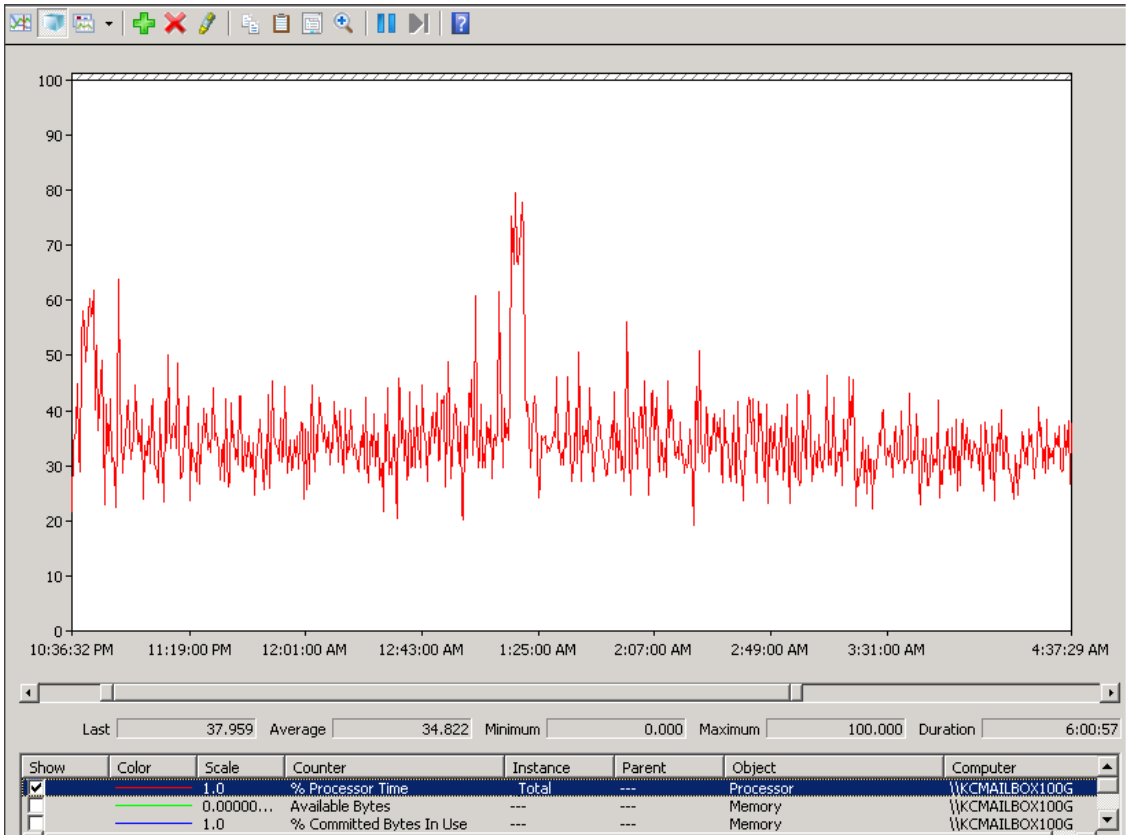

Figure 6 shows the CPU utilization on one of the Client Access servers hovering around 30% and peaking at around 80% at times.

Figure 6 CPU Utilization on a Client Access Server

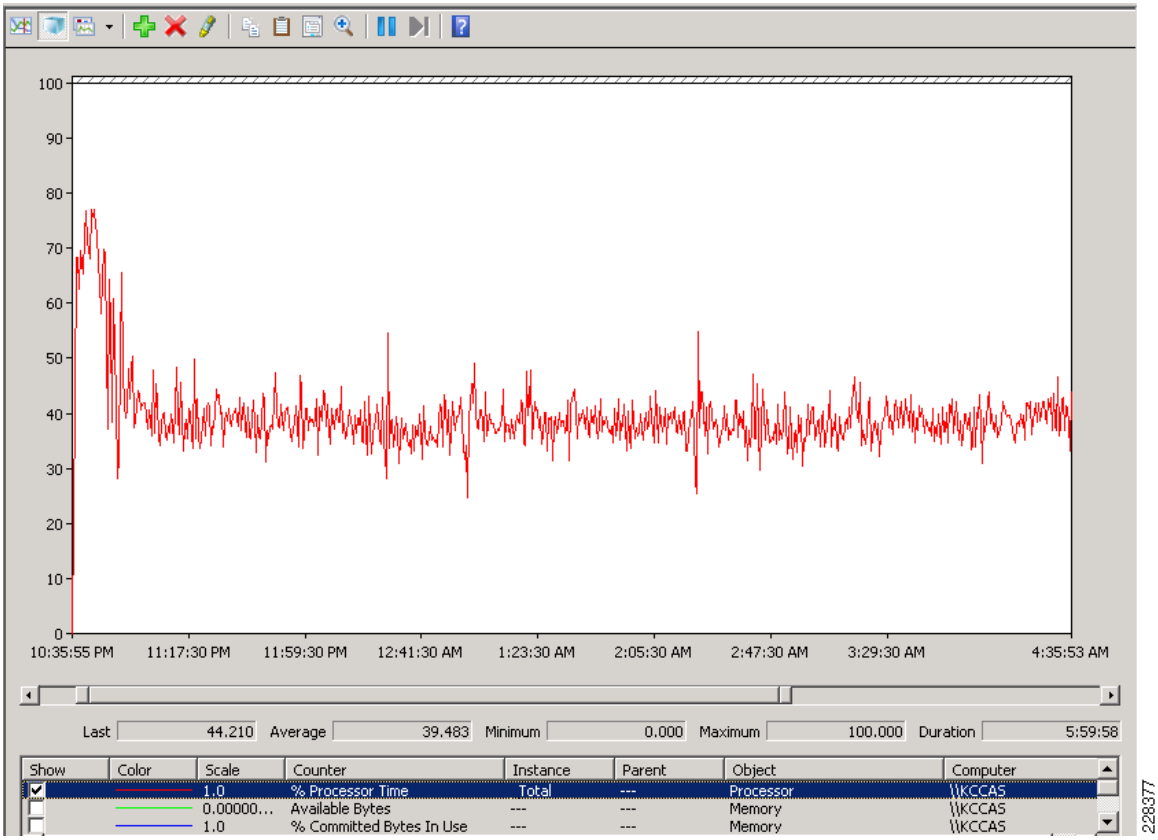

Microsoft Windows 2008 R2 Hyper-V for Virtualization and Live Migration

Microsoft Exchange 2010 server roles are virtualized as individual Hyper-V virtual machines deployed on Windows 2008 R2 Hyper-V servers on the Unified Computing System server blades. See Exchange 2010 Application Architecture for details on the application architecture.

Windows Server 2008 R2 was designed to provide features that would improve the performance of virtual networking to support Hyper-V virtual machines. The following highlights these features:

- Support for Jumbo frames, previously available in non-virtual environments, has been extended to work with VMs. This feature enables VMs to use Jumbo Frames up to 9014 bytes if the underlying physical network supports it. Supporting Jumbo frames reduces the network stack overhead incurred per byte and increases throughput. In addition, there is a significant reduction of CPU utilization due to the fewer number of calls from the network stack to the network driver.
- TCP Chimney, which allows the offloading of TCP/IP processing to the network hardware, has been extended to the virtual environment. It improves VM performance by allowing the VM to offload network processing to hardware, especially on networks with bandwidth over 1 Gigabit. This feature is especially beneficial for roles involving large amounts of data transfer, such as the file server role.
- The Virtual Machine Queue (VMQ) feature allows physical computer Network Interface Cards (NICs) to use direct memory access (DMA) to place the contents of packets directly into VM memory, increasing I/O performance.

Microsoft Hyper-V R2

Microsoft Hyper-V is offers as a server role packaged into the Windows Server 2008 R2 installation or as a standalone server. In either case, it is a hypervisor-based virtualization technology for x64 versions of Windows Server 2008. The hypervisor is a processor-specific virtualization platform that allows multiple isolated operating systems to share a single hardware platform. In order to run Hyper-V virtualization, the following system requirements must be met:

- An x86-64-capable processor running an x64 version of Windows Server 2008 Standard, Windows Server 2008 Enterprise, or Windows Server 2008 Datacenter.
- Hardware-assisted virtualization—This is available in processors that include a virtualization option; specifically, Intel VT, which is an extension to the x86 architecture. The processor extension allows a virtual machine hypervisor to run an unmodified operating system without incurring significant performance penalties within OS emulation.
- An NX bit-compatible CPU must be available and the Hardware Data Execution Prevention (DEP) bit must be enabled in the BIOS. For the Unified Computing System, these are offered/enabled by default.
- Memory minimum 2 GB with attention to providing more memory based on the virtual OS and application requirements. The stand-alone Hyper-V Server does not require an existing installation of Windows Server 2008 and minimally requires 1GB of memory and 2GB of disk space.

Hyper-V provides isolation of operating systems running on the virtual machines from each other through partitioning or logical isolation by the hypervisor. Each hypervisor instance has at least one parent partition that runs Windows Server 2008. The parent partition houses the virtualization stack that has direct access to the hardware devices such as the network interface cards and are responsible for creating the child partitions that host the guest operating systems. The parent partition creates these child partitions using the hypercall API, an application programming interface exposed by Hyper-V.

A virtualized partition does not have access to the physical processor, nor does it handle its real interrupts. Instead, it has a virtual view of the processor and runs in Guest Virtual Address, which (depending on the configuration of the hypervisor) might not necessarily be the entire virtual address space. A hypervisor could choose to expose only a subset of the processors to each partition. The hypervisor handles the interrupts to the processor and redirects them to the respective partition using a logical Synthetic Interrupt Controller (SynIC). Hyper-V can hardware accelerate the address translation between various Guest Virtual Address spaces by using an IOMMU (I/O Memory Management Unit) which operates independent of the memory management hardware used by the CPU.

Child partitions do not have direct access to hardware resources, but instead have a virtual view of the resources, in terms of virtual devices. Any request to the virtual devices is redirected via the VMBus to the devices in the parent partition, which manages the requests. The VMBus is a logical channel which enables inter-partition communication. The response is also redirected via the VMBus. If the devices in the parent partition are also virtual devices, it is redirected further until it reaches the parent partition, where it gains access to the physical devices. Parent partitions run a Virtualization Service Provider (VSP), which connects to the VMBus and handles device access requests from child partitions. Child partition virtual devices internally run a Virtualization Service Client (VSC), which redirect the request to VSPs in the parent partition via the VMBus. This entire process is transparent to the guest OS.

Virtual Devices can also take advantage of a Windows Server Virtualization feature, named Enlightened I/O, for storage, networking, and graphics subsystems, among others. Enlightened I/O is specialized virtualization-aware implementation of high-level communication protocols like SCSI to take advantage of VMBus directly, that allows bypassing any device emulation layer. This makes the communication more efficient, but requires the guest OS to support Enlightened I/O. Windows 2008, Windows Vista, and SUSE Linux are currently the only operating systems that support Enlightened I/O, allowing them therefore to run faster as guest operating systems under Hyper-V than other operating systems that need

to use slower emulated hardware. Hyper-V enlightened I/O and a hypervisor-aware kernel are provided via installation of Hyper-V integration services. Integration components, which include virtual server client (VSC) drivers, are also available for other client operating systems.

See Appendix B—Hyper-V Logical Diagram for a logical overview of the Hyper-V architecture components.

Overview of Hyper-V Live Migration of Exchange Virtual Machines

Hyper-V built into Windows Server 2008 R2 supports Live Migration with the use of Cluster Shared Volumes (CSVs). This allows for an individual virtual machine that is still online and actively supporting user sessions to be moved seamlessly to a different physical server (Hyper-V server) without disrupting user services.

Live Migration

For this solution, Hyper-V Live Migration is supported on the Client Access and Hub Transport server roles. Because the Mailbox servers are participating in a Database Availability Group cluster and because Live Migration requires its own Windows failover cluster and the two cannot co-exist, the solution relies on the redundancy and failover benefits of the DAG feature to move the active Mailbox services from one UCS blade hardware to another.

Hyper-V live migration moves running VMs with no impact on VM availability to users. By pre-copying the memory of the migrating VM to the destination physical host, live migration minimizes the transfer time of the VM. A live migration is deterministic, meaning that the administrator, or script, that initiates the live migration can control which computer that is the destination for the live migration. The guest operating system of the migrating VM is unaware the migration is happening, so no special configuration for the guest operating system is needed.

During the live migration setup stage, the source physical host creates a TCP connection with the destination physical host. This connection transfers the VM configuration data to the destination physical host. A skeleton VM is set up on the destination physical host and memory is allocated to the destination VM.

In the second stage of a live migration, the memory assigned to the migrating VM is copied over the network to the destination physical host This memory is referred to as the working set of the migrating VM. A page of memory is 4 kilobytes.

In addition to copying the working set of the migrating VM to the destination physical host, Hyper-V on the source physical host monitors the pages in the working set for that migrating VM. As memory pages are modified by the migrating VM, they are tracked and marked as being modified. The list of modified pages is simply the list of memory pages the migrating VM has modified after the copy of its working set has begun.

During this phase of the migration, the migrating VM continues to run. Hyper-V iterates the memory copy process several times, each time a smaller number of modified pages need to be copied to the destination physical computer. After the working set is copied to the destination physical host, the next stage of the live migration begins.

The next stage is a memory copy process that duplicates the remaining modified memory pages for the migrating VM to the destination physical host. The source physical host transfers the register and device state of the VM to the destination physical host.

During this stage, the network bandwidth available between the source and destination physical hosts is critical to the speed of the live migration and using a 1 Gigabit Ethernet or faster is important. In fact, Microsoft recommends that a 1 Gigabit Ethernet connection be dedicated for the live migration network between cluster nodes to transfer the large number of memory pages typical for a virtual machine.

The faster the source physical host transfers the modified pages from the migrating VMs working set, the more quickly the live migration completes.

The number of pages transferred in this stage is dictated by how actively the VM is accessing and modifying memory pages. The more modified pages, the longer the VM migration process takes for all pages to be transferred to the destination physical host.

After the modified memory pages are copied completely to the destination physical host, the destination physical host has an up-to-date working set for the migrating VM. The working set for the migrating VM is present on the destination physical host in the exact state it was in when the migrating VM began the migration process.

During the next stage of a live migration, control of the storage associated with the migrating VM, such as any VHD files or pass-through disks, is transferred to the destination physical host. The destination server now has the up-to-date working set for the migrating VM as well as access to any storage used by that migrating VM. At this point the VM is running on the destination server and has been fully migrated.

At this point a message is sent to the physical network switch which causes it to re-learn the MAC addresses of the migrated VM so that network traffic to and from the migrated VM can use the correct switch port.

Cluster Shared Volumes

With Windows Server 2008 R2, Hyper-V uses Cluster Shared Volumes (CSV) storage to support Live Migration of Hyper-V virtual machines from one Hyper-V server to another. CSV enables multiple Windows Servers to access SAN storage using a single consistent namespace for all volumes on all hosts. Multiple hosts can access the same Logical Unit Number (LUN) on SAN storage so that VHD files (virtual hard drives) can be migrated from one hosting Hyper-V server to another. Cluster Shared Volumes is available as part of the Windows Failover Clustering feature of Windows Server® 2008 R2.

Multiple VLANs for Different Traffic Types

Microsoft recommends that dedicated links be assigned to various traffic types that are critical to the operation of given applications and network services. For this solution, there are five critical traffic types that should therefore be isolated from each other. Since this solution involves the use of UCS B200-M1 blades with two physical 10 Gigabit Ethernet interfaces per blade, logical VLAN interfaces can be defined on each physical 10GE interface to provide this traffic isolation. Looking forward, Microsoft Quality-of-Service can be applied with the VLANs to provide bandwidth guarantees to Hyper-V network traffic; however, that is outside the scope of this solution.

There are five different traffic types involved in this solution architecture.

Some of the traffic types involve communication to the Hyper-V Server (parent OS) and some involve traffic up to the virtual guest operating systems themselves:

- Live Migration
- VM Management (Virtual Machine management by System Center Operations Manager Virtual Machine Manager and Hyper-V Manager)
- Cluster Interconnect and Cluster Shared Volumes
- Virtual machine access
  - Exchange Client Access and application data traffic
  - Database Availability Group clustering
- Blade server management

The dual-port Cisco UCS 82598KR-CI 10 Gigabit Ethernet Adapter offers sufficient network bandwidth over its two network interfaces to each UCS server blade. Each traffic type can then be put on its own VLAN and the appropriate Windows QoS policy applied for bandwidth guarantees to each.

Microsoft recommends two specific Windows QoS policies for limiting bandwidth utilization by Live Migration traffic and by blade server management traffic. Figure 7 shows these two QoS policies can be configured locally on each of the UCS blade servers using Windows 2008 Local Group Policy Editor.

The QoS policy for limiting bandwidth utilization by Live Migration traffic would be given the following non-default parameters through the Local Group Policy Editor UI:

- No DSCP value setting is needed.
- Thottle rate is set to 50% of a 10GE link or 5120Mbps.
- TCP destination port is set to 6600 for Live Migration traffic.

The QoS policy for limiting bandwidth utilization by blade server management traffic would be given the following non-default parameters through the Local Group Policy Editor UI:

- No DSCP value setting is needed.
- Throttle rate is set to 1% of a 10GE link or 102Mbps.
- Source IP address/subnet is set to the subnet for remote management of the UCS blade servers (e.g., 10.7.53.0/24). The destination IP address is left at the default setting of “any”.
- TCP or UDP protocol is selected for protocol type.

Figure 7 Configuring QoS Policies Using Windows 2008 Local Group Policy Editor

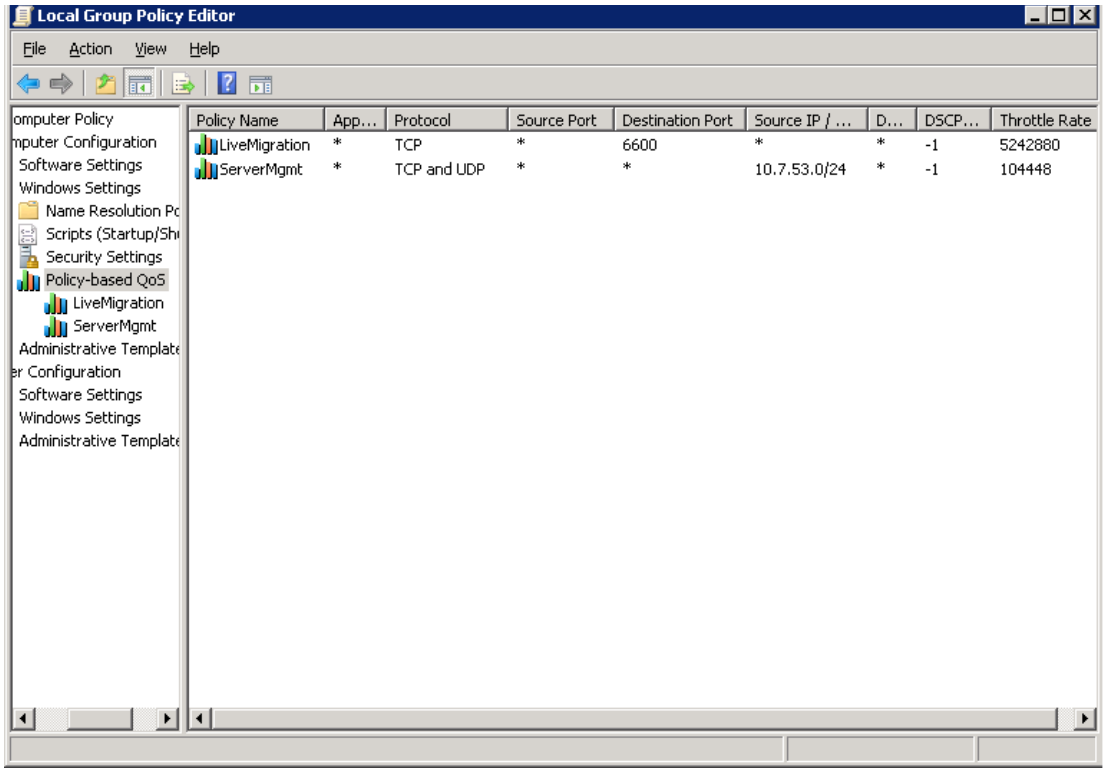

Designing Application Delivery for Branch and Mobile Users

Application delivery to users at the branch and remote locations involves optimizing user sessions across the WAN to avoid WAN congestion issues and load balancing user sessions among multiple Exchange server role instances to better utilize available servers in the farm and avoid server overload. WAN optimization and application acceleration is achieved through deployment of Cisco Wide Area Application Services while load balancing is provided by the Cisco Application Control Engine service module deployed in the Catalyst 6500 Services Chassis at the data center aggregation layer.

Load Balancing for an Exchange 2010 Environment

In Microsoft's Exchange 2010 there have been some architectural changes that have increased the importance of load balancing. Although there are by default some software based load balancing technologies used in the exchange 2010 architecture a hardware based load balancer such as the Cisco Application Control Engine (ACE) can be beneficial.

Two changes to Exchange 2010 solution in regards to the Client Access Server (CAS) are the RPC Client Access Service and the Exchange Address Book Service. In previous iterations, Exchange Outlook clients would connect directly to the mailbox server and in Exchange 2010 all client connections internal and external are terminated at the CAS or by a pool of CAS servers.

Commonly Windows Network Load Balancing (WNLB) is used to load balance different roles of the Exchange 2010 environment. A hardware load balancer such as ACE can be used in place to eliminate some limitations such as scalability and functionality. Microsoft suggests a hardware load balancer is needed when deployments have more then eight CAS servers. However in some scenarios with deployments of less then eight CAS servers, a hardware load balancer can also be used. One specific limitation of WNLB is it is only capable of session persistence based on the client's IP address.

ACE can be used to offload CUP-intensive tasks such as SSL encryption and decryption processing and TCP session management, this greatly improves server efficiency. More complex deployments also dictate the use of a hardware load balancer, for example if CAS server role is co-located on servers running the mailbox server role in a database availability configuration (DAG). This requirement is due to a incompatibility with the Windows Failover clustering and WNLB. For more information, see Microsoft Support Knowledge Base article 235305.

This document and its technology recommendations are intended to be used in a pure Exchange 2010 deployment, however some may work in Exchange 2007 or mixed deployments. Such deployments are beyond the scope of this document as they were not tested and validated. For more information on load balancing an Exchange 2007 deployment with ACE, see: http://www.cisco.com/en/US/docs/solutions/Verticals/mstdcmsftex.html.

Exchange 2010 MAPI Client Load Balancing Requirements

As previously mentioned there are some changes in Exchange 2010 from Exchange 2007, including how Exchange clients communicate. Outlook 2010 clients use MAPI over RPC to the Client Access Servers, therefore IP-based load balancing persistence should be used. However with RPC the TCP port numbers are dynamically derived when the service is started and configurations for a large range of dynamically created ports is not desirable. Because of this a catch all can be used for RPC traffic at the load balancer for any TCP based traffic destined for the server farm.

Microsoft does give a work around to this and allow static port mapping to simplify load balancing; however in this design our testing was limited to the use of a catch all for Exchange 2010 client testing.

Configuring Static Port Mapping For RPC-Based Services

If the ports for these services are statically mapped, then the traffic for these services is restricted to port 135 (used by the RPC portmapper) and the two specific ports that have been selected for these services.

The static port for the RPC Client Access Service is configured via the registry. The following registry key should be set on each Client Access Server to the value of the port that you wish to use for TCP connections for this service.

Key:
HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\MSExchangeRPC\Parameters System

Value: TCP/IP Port

Type: DWORD

The static ports for the two RPC endpoints maintained by the Exchange Address Book Service are set in the Microsoft.Exchange.Address Book.Service.Exe.config file which can be found in the bin directory under the Exchange installation path on each Client Access Server. The "RpcTcpPort" value in the configuration file should be set to the value of the port that you wish to use for TCP connections for this service. This port handles connections for both the Address Book Referral (RFR) interface and the Name Service Provider Interface (NSPI).

Note This only affects connections for "internal" connections via TCP and does not affect Outlook Anywhere connections that take advantage of RPC/HTTP tunneling. Outlook Anywhere connections to the RPC Client Access Service occur on port 6001 and this is not configurable.

Exchange 2010 Outlook Web Access Load Balancing Requirements

In regards to Outlook Web App (OWA), client session information is maintained on the CAS for the duration of a user's session. These Web sessions are either terminated via a timeout or logout. Access permissions for sessions may be affected by client session information depending on the authentication used. If an OWA client session request is load balanced between multiple CAS servers, requests can be associated with different sessions to different servers. This causes the application to be unusable due to frequent authentication requests.

IP-based session persistence is a reliable and simple method of persistence, however if the source of the session is being masked at any given point in the connection by something such as network address translation (NAT), it should not be used. Cookie-based persistence is the optimum choice for OWA sessions. Cookie-based persistence can either be achieved by insertion of a new cookie in the HTTP stream by the load balancer or by using the existing cookie in the HTTP stream that comes from the CAS server.

If forms-based authentication is being used there is a timing concern that has to be addressed. SSL ID-based persistence as a fall back can be used to address this timing constraint.

The following scenario depicts the sequence of events that require SSL IP-based persistence as a fallback:

1. Client sends credentials via an HTTP POST action to /owa/auth.owa.
2. Server responds authentication request from a client and redirect to /owa/ and sets two cookies, "cadata" and "sessionid".
3. Client follows redirect and passes up the two new cookies.
4. Server sets "UserContext" cookie.

It is critical that the server which issued the two cookies in step 2 “cadata” and “sessionid” is the same server that is accessed in step 3. By using SSL ID-based persistence for this part of the transaction you can maintain that the requests are sent to the same server in both steps.

It is also important to understand the browser process model in relation to OWA session workloads if SSL ID-based persistence is configured. With Internet Explorer 8 as an OWA client, some operations may open a new browser window or tab such as opening a new message, browsing address lists, or creating a new message. When a new window or tab is launched, so is a new SSL ID and therefore a new logon screen because the session can possibly be going to a different CAS server that is unaware of the previous session. By using client certificates for authentication, IE8 does not spawn new processes and subsequently client traffic remains on the same CAS server.

Outlook Anywhere Load Balancing Requirements

Outlook Anywhere Clients use a RPC Proxy component to proxy RPC calls to the RPC Client Access Service and Exchange Address Book Service. If the real Client IP is available to the CAS, IP-based persistence can be used with a load balancer for these connections. However commonly in the case with remote clients, the client IP is most likely using NAT at one or more points in the network connection. Therefore a less basic persistence method is needed.

If basic authentication is being used persistence can be based on the Authorization HTTP header. If your deployment is a pure Outlook 2010 environment using the Outlook 2010 client you can also use a cookie with the value set to “OutlookSession” for persistence.

Cisco Application Control Engine Overview

The Cisco ACE provides a highly available and scalable data center solution from which the Microsoft Exchange Server 2010 application environment can benefit. Currently, the Cisco ACE is available as an appliance or integrated service module in the Cisco Catalyst 6500 platform. The Cisco ACE features and benefits include the following:

- Device partitioning (up to 250 virtual ACE contexts)
- Load balancing services (up to 16 Gbps of throughput capacity and 325,000 Layer 4 connections/second)
- Security services via deep packet inspection, access control lists (ACLs), unicast reverse path forwarding (uRPF), Network Address Translation (NAT)/Port Address Translation (PAT) with fix-ups, syslog, etc.
- Centralized role-based management via Application Network Manager (ANM) GUI or CLI
- SSL-offload (up to 15,000 SSL sessions via licensing)
- URL rewrite
- Support for redundant configurations (intra-chassis, inter-chassis, and inter-context)

ACE can be configured in the following modes of operation:

- Transparent
- Routed
- One-Armed

The following sections describe some of the Cisco ACE features and functionalities used in the Microsoft Exchange Server 2010 application environment.

ACE One-Armed Mode Design

One-armed configurations are used when the device that makes the connection to the virtual IP address (VIP) enters the ACE on the same VLAN that the server resides. The servers must traverse back through the ACE before reaching the client. This is done with either source NAT or policy-based routing. Because the network design for this document has ACE VIP on the same VLAN as the actual servers being load balanced, a One-armed mode was used.

In the One-armed mode of operation clients send application requests through the multilayer switch feature card (MSFC), which routes them to a virtual IP address (VIP) within the Application Control Engine (ACE), which is configured with a single VLAN to handle client and server communication. Client requests arrive at the VIP and the ACE picks the appropriate server and then uses the destination Network Address Translation (NAT) to rewrite the destination IP to that of the rserver and rewrite the source IP with one from the nat-pool. Once the client request is fully NAT'd, it is sent to the server over the same VLAN which it was originally received. The server responds to the Cisco ACE based on the source IP of the request. The Cisco ACE receives the response. The ACE then changes the source IP to be the VIP and routes the traffic to the MSFC.

The MSFC then forwards the response to the client. Figure 8 displays these transactions at a high level.

Figure 8 One-Armed Mode Design Transactions

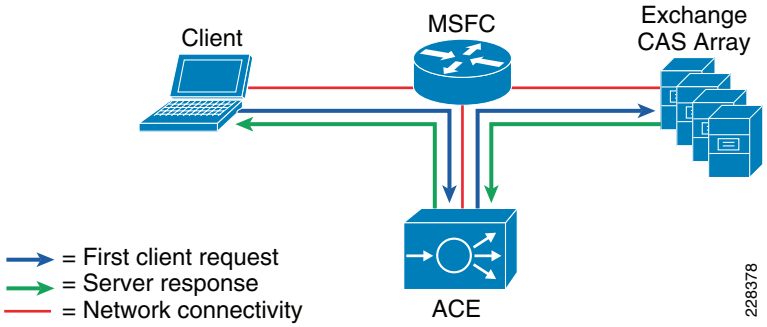

ACE Virtualization

Virtualization is a prevalent trend in the enterprise today. From virtual application containers to virtual machines, the ability to optimize the use of physical resources and provide logical isolation is gaining momentum. The advancement of virtualization technologies includes the enterprise network and the intelligent services it offers.

The Cisco ACE supports device partitioning where a single physical device may provide multiple logical devices. This virtualization functionality allows system administrators to assign a single virtual ACE device to a business unit or application to achieve application performance goals or service-level agreements (SLAs). The flexibility of virtualization allows the system administrator to deploy network-based services according to the individual business requirements of the customer and technical requirements of the application. Service isolation is achieved without purchasing another dedicated appliance that consumes more space and power in the data center.

SSL Offload

The Cisco ACE is capable of providing secure transport services to applications residing in the data center. The Cisco ACE implements its own SSL stack and does not rely on any version of OpenSSL. The Cisco ACE supports TLS 1.0, SSLv3, and SSLv2/3 hybrid protocols. There are three SSL relevant deployment models available to each ACE virtual context:

- SSL termination—Allows for the secure transport of data between the client and ACE virtual context. The Cisco ACE operates as an SSL proxy, negotiating and terminating secure connections with a client and a non-secure or clear text connection to an application server in the data center. The advantage of this design is the offload of application server resources from taxing the CPU and memory demands of SSL processing, while continuing to provide intelligent load balancing.
- SSL initiation—Provides secure transport between the Cisco ACE and the application server. The client initiates an unsecure HTTP connection with the ACE virtual context, the Cisco ACE acting as a client proxy negotiates an SSL session to an SSL server.
- SSL end-to-end—Provides a secure transport path for all communications between a client and the SSL application server residing in the data center. The Cisco ACE uses SSL termination and SSL initiation techniques to support the encryption of data between client and server. Two completely separate SSL sessions are negotiated, one between the ACE context and the client, the other between the ACE context and the application server. In addition to the intelligent load balancing services the Cisco ACE provides in an end-to-end SSL model, the system administrator may choose to alter the intensity of data encryption to reduce the load on either the front-end client connection or back-end application server connection to reduce the SSL resource requirements on either entity.

SSL URL Rewrite Offload

The Cisco ACE is capable of inserting or deleting HTTP header information for connections it is sustaining. This capability is highly useful when an application server responds with a HTTP 302 or "Moved Temporarily" response to a client's HTTP GET or HEAD request. The HTTP 302 response usually indicates a new HTTP LOCATION URL for the client to access. Modifying the HTTP LOCATION value for a secure connection is known as SSL URL rewrite. The SSL URL Rewrite feature allows the system administrator to alter the HTTP LOCATION value returned to the client resulting in granular control of the application's session flow and persistence in the data center.

Session Persistence

Session persistence is the ability to forward client requests to the same server for the duration of a session. Microsoft supports session persistence for their Microsoft Exchange environment via the following methods:

- Source IP sticky
- Cookie sticky

The Cisco ACE supports each of these methods, but given the presence of proxy services in the enterprise, Cisco recommends using the cookie sticky method to guarantee load distribution across the server farm wherever possible as session-based cookies present unique values to use for load balancing.

In addition, the Cisco ACE supports the replication of sticky information between devices and their respective virtual contexts. This provides a highly available solution that maintains the integrity of each client's session.

Health Monitoring

The Cisco ACE device is capable of tracking the state of a server and determining its eligibility for processing connections in the server farm. The Cisco ACE uses a simple pass/fail verdict but has many recovery and failures configurations, including probe intervals, timeouts, and expected results. Each of these features contributes to an intelligent load-balancing decision by the ACE context.

The predefined probe types currently available on the ACE module are:

- ICMP

- TCP
- UDP
- Echo (TCP/UDP)
- Finger
- HTTP
- HTTPS (SSL Probes)
- FTP
- Telnet
- DNS
- SMTP
- IMAP
- POP
- RADIUS
- Scripted (TCL support)

Note that the potential probe possibilities available via scripting make the Cisco ACE an even more flexible and powerful application-aware device. In terms of scalability, the Cisco ACE module can support 1000 open probe sockets simultaneously.

Cisco Wide Area Application Services Between Branch and Data Center

To provide optimization and acceleration services between the branch and data center, a Cisco WAAS appliance was deployed at the data center WAN aggregation tier in a one-armed deployment and a WAAS network module was deployed in the Integrated Services Router at the branch edge.

To appreciate how the Cisco Wide Area Application Services (Cisco WAAS) provides WAN optimization and application acceleration benefits to the enterprise, consider the basic types of centralized application messages that are transmitted between remote branches. For simplicity, two basic types are identified:

- Bulk transfer applications—Transfer of files and objects, such as FTP, HTTP, and IMAP. In these applications, the number of round-trip messages might be few and might have large payloads with each packet. Examples include Web portal or thin client versions of Oracle, SAP, Microsoft (SharePoint, OWA) applications, E-mail applications (Microsoft Exchange, Lotus Notes), and other popular business applications.
- Transactional applications—High numbers of messages transmitted between endpoints. Chatty applications with many round-trips of application protocol messages that might or might not have small payloads. Examples include CIFS file transfers.

The Cisco WAAS uses the technologies described in the following sections to enable optimized Exchange Outlook communication between the branch office outlook clients and Exchange servers in the data center by providing TFO optimization, LZ compression, DRE caching, MAPI acceleration, and SSL acceleration.

Advanced Compression Using DRE and LZ Compression

Data Redundancy Elimination (DRE) is an advanced form of network compression that allows the Cisco WAAS to maintain an application-independent history of previously-seen data from TCP byte streams. Lempel-Ziv (LZ) compression uses a standard compression algorithm for lossless storage. The combination of using DRE and LZ reduces the number of redundant packets that traverse the WAN, thereby conserving WAN bandwidth, improving application transaction performance, and significantly reducing the time for repeated bulk transfers of the same application.

TCP Flow Optimization

The Cisco WAAS TCP Flow Optimization (TFO) uses a robust TCP proxy to safely optimize TCP at the Cisco WAE device by applying TCP-compliant optimizations to shield the clients and servers from poor TCP behavior due to WAN conditions. The Cisco WAAS TFO improves throughput and reliability for clients and servers in WAN environments through increases in the TCP window sizing and scaling enhancements—as well as through the implementation of congestion management and recovery techniques—to ensure that the maximum throughput is restored in the event of packet loss. By default, Cisco WAAS provides only TFO for RDP. If RDP compression and encryption are disabled, then full optimization (TFO+ DRE/LZ) can be enabled for RDP flows.

Messaging Application Programming Interface (MAPI) Protocol Acceleration

The MAPI application accelerator accelerates Microsoft Outlook Exchange 2010 traffic that uses the Messaging Application Programming Interface (MAPI) protocol. Microsoft Outlook 2010 clients are supported. Clients can be configured with Outlook in cached or non-cached mode; both modes are accelerated. Secure connections that use message authentication (signing) or encryption or Outlook Anywhere connections (MAPI over HTTP/HTTPS) are not accelerated by the MAPI application accelerator. To allow internal Outlook users to benefit from application acceleration of their Exchange traffic, these users can leverage the Outlook Anywhere option to run MAPI over HTTPS, in which case SSL acceleration of their E-mail traffic can be leveraged.

Secure Socket Layer (SSL) Protocol Acceleration

Cisco WAAS provides the widest set of customer use cases for deploying WAN optimization into SSL-secured environments. For example, many industries and organizations use Web proxy servers to front-end their key business process applications in the data center and encrypt with SSL from the proxy server to remote users in branch sites. Cisco WAAS provides optimized performance for delivering these applications, while preserving the SSL encryption end-to-end—something competitive products' SSL implementations do not support. To maximize application traffic security, WAAS devices provide additional encryption for data stored at rest and can be deployed in an end-to-end SSL-optimized architecture with Cisco ACE application switches for SSL offload.

Cisco WAAS provides SSL optimization capabilities that integrate fully with existing data center key management and trust models and can be used by both WAN optimization and application acceleration components. Private keys and certificates are stored in a secure vault on the Cisco WAAS Central Manager. The private keys and certificates are distributed in a secure manner to the Cisco WAAS devices in the data center and stored in a secure vault, maintaining the trust boundaries of server private keys. SSL optimization through Cisco WAAS is fully transparent to end users and servers and requires no changes to the network environment.

Among the several cryptographic protocols used for encryption, SSL/TLS is one of the most important. SSL/TLS-secured applications represent a growing percentage of traffic traversing WAN links today. Encrypted secure traffic represents a large and growing class of WAN data. Standard data redundancy elimination (DRE) techniques cannot optimize this WAN data because the encryption process generates an ever-changing stream of data, making even redundant data inherently non-reducible and eliminating the possibility of removing duplicate byte patterns. Without specific SSL optimization, Cisco WAAS can still provide general optimization for such encrypted traffic with transport flow optimization (TFO). Applying TFO to the encrypted secure data can be helpful in many situations in which the network has a high bandwidth delay product (BDP)1 and is unable to fill the pipe.

Termination of the SSL session and decryption of the traffic is required to apply specific SSL optimizations such as Cisco WAAS DRE and Lempel-Ziv (LZ) compression techniques to the data. Minimally, SSL optimization requires the capability to:

- Decrypt traffic at the near-side Cisco WAAS Wide Area Application Engine (WAE) and apply WAN optimization to the resulting clear text data.
- Re-encrypt the optimized traffic to preserve the security of the content for transport across the WAN.
- Decrypt the encrypted and optimized traffic on the far-side Cisco WAAS WAE and decode the WAN optimization.
- Re-encrypt the resulting original traffic and forward it to the destination origin server.

The capability to terminate SSL sessions and apply WAN optimizations to encrypted data requires access to the server private keys. Further, the clear-text data received as a result of decryption must be stored on the disk for future reference to gain the full benefits of DRE. These requirements pose serious security challenges in an environment in which data security is paramount. Security by itself is the most important and sensitive aspect of any WAN optimization solution that offers SSL acceleration.

Figure 9 SSL Protocol Acceleration

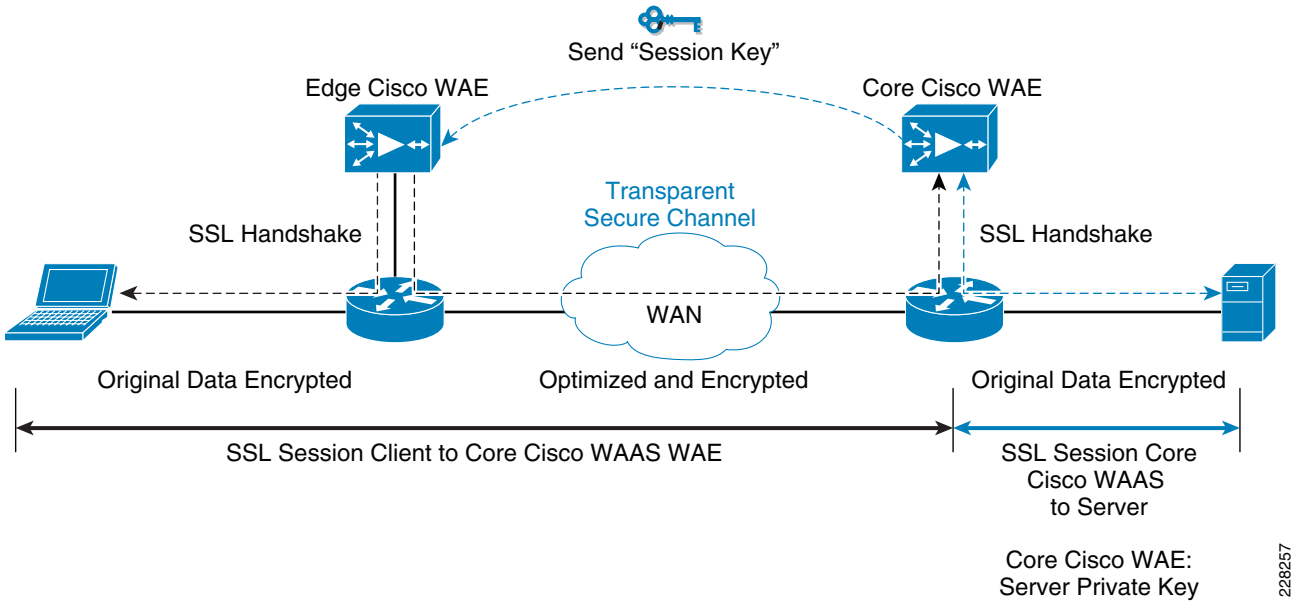

During initial client SSL handshake, the core Cisco WAE in the data center participates in the conversation.

The connection between the Cisco WAEs is established securely using the Cisco WAE device certificates and the Cisco WAEs cross-authenticate each other:

- After the client SSL handshake occurs and the data center Cisco WAE has the session key, the data center Cisco WAE transmits the session key (which is temporary) over its secure link to the edge Cisco WAE so that it can start decrypting the client transmissions and apply DRE.
- The optimized traffic is then re-encrypted using the Cisco WAE peer session key and transmitted, in-band, over the current connection, maintaining full transparency, to the data center Cisco WAE.
- The core Cisco WAE then decrypts the optimized traffic, reassembles the original messages, and re-encrypts these messages using a separate session key negotiated between the server and the data center Cisco WAE.
- If the back-end SSL server asks the client to submit an SSL certificate, the core Cisco WAE requests one from the client. The core Cisco WAE authenticates the client by verifying the SSL certificate using a trusted Certificate Authority (CA) or an Online Certificate Status Protocol (OCSP) responder.

Cisco WAAS Mobile

In addition to Cisco WAAS for branch optimization, Cisco offers Cisco WAAS Mobile for telecommuters, mobile users, and small branch and home office users who access corporate networks and need accelerated application performance. Cisco WAAS Mobile client is purpose-built for Microsoft Windows PCs and laptops. To provide WAAS Mobile services to remote users, a Windows 2008 WAAS server was deployed as a virtual machine on the Unified Computing System to support user connections into the data center Exchange server farm.

Advanced Data Transfer Compression

Cisco WAAS Mobile maintains a persistent and bi-directional history of data on both the mobile PC and the Cisco WAAS Mobile server. This history can be used in current and future transfers, across different VPN sessions, or after a reboot, to minimize bandwidth consumption and to improve performance. In addition, instead of using a single algorithm for all file types, Cisco WAAS Mobile uses a file-format specific compression to provide higher-density compression than generic compression for Microsoft Word, Excel, and PowerPoint files, Adobe Shockwave Flash (SWF) files, ZIP files, and JPEG, GIF, and PNG files.

Application-Specific Acceleration

Cisco WAAS Mobile reduces application-specific latency for a broad range of applications, including Microsoft Outlook Messaging Application Programming Interface (MAPI), Windows file servers, or network-attached storage using CIFS, HTTP, HTTPS and other TCP-based applications, such as RDP.

Transport Optimization

Cisco WAAS Mobile extends Cisco WAAS technologies to handle the timing variations found in packet switched wireless networks, the significant bandwidth latency problems of broadband satellite links, and noisy Wi-Fi and digital subscriber line (DSL) connections. The result is significantly higher link resiliency.

Management

Cisco Unified Computing System Manager

Virtual Machine Manager 2008 R2

Microsoft's Virtual Machine Manager 2008 R2 allows administrators to deploy and manage a large number of virtual machines across the network. It is a series of components that include Windows Server, SQL Server, a local agent, an Administrative console, and a self-service console.

The SQL database supports the VMM library that contains all VMM objects; such objects include:

- Hardware profiles—These profiles make up the virtual hardware components of a VM, which include BIOS boot order, CPU count and type, physical RAM, floppy drive, serial ports, IDE/SCSI adapters, and network adapters.
- Guest OS profiles—These profiles contain information about the name, administrator password, Windows product key, time zone, and Windows operating system type of the VM.
- Virtual disks and ISOs—The VMM library stores the virtual hard disks (VHD files) that are created by the administrator when either a new virtual machine is created or when a new virtual hard drive is created for an existing virtual machine. ISO images are saved in the VMM library so they can be mounted anytime to any virtual machine to simulate a DVD/CD-ROM.
- VM Templates—Templates can be created once so that future VMs can be created using those templates. They usually consist of a VHD, a hardware profile, and an OS profile.

Figure 10 Virtual Machine Manager

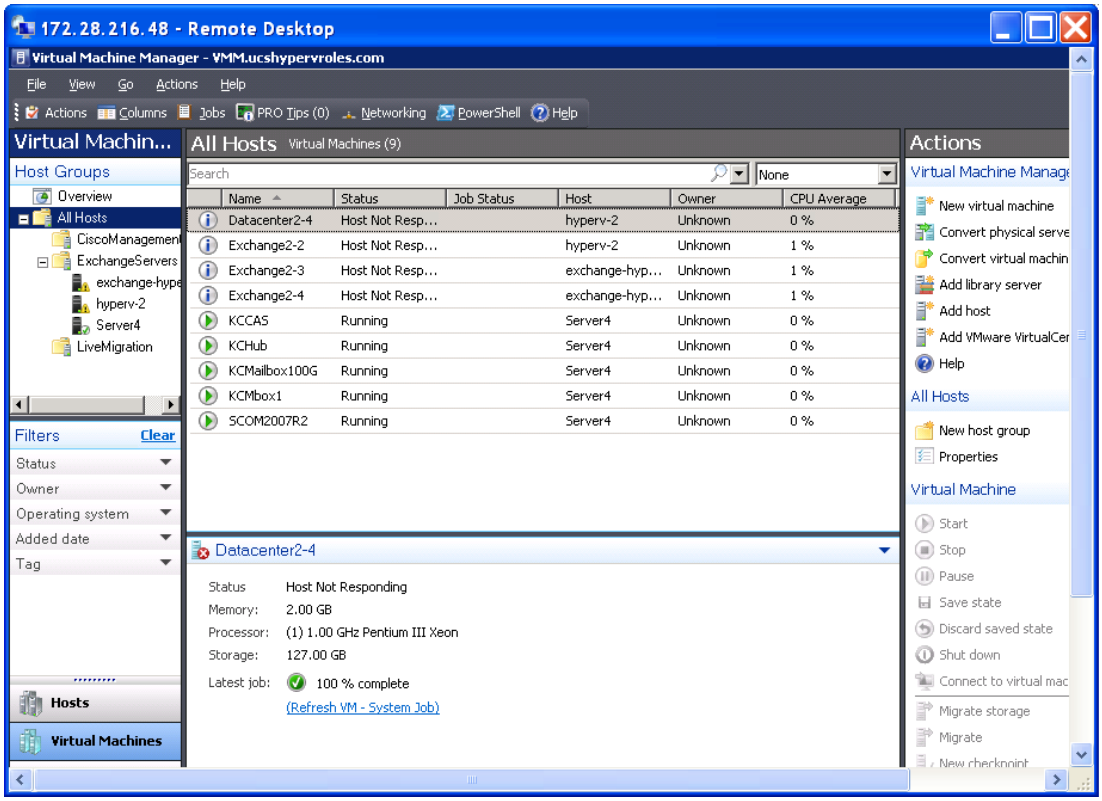

WAAS Central Management

The Cisco WAAS Central Manager provides a centralized mechanism for configuring features, reporting, and monitoring. It can manage a topology containing thousands of Cisco WAE nodes. The Cisco WAAS Central Manager can be accessed from a Web browser, allowing management from essentially anywhere in the world. Access to the Cisco WAAS Central Manager is secured and encrypted with Secure Sockets Layer (SSL) and users can be authenticated through a local database or a third-party authentication service such as RADIUS, TACACS, or Microsoft Active Directory. Any Cisco WAE device can be converted into a WAAS Central Manager through a setting in the configuration, not requiring any change in software version.

Within a Cisco WAAS topology, each Cisco WAE runs a process called central management system (CMS). The CMS process provides SSL-encrypted bidirectional configuration synchronization of the Cisco WAAS Central Manager and the Cisco WAE devices. The CMS process is also used to exchange reporting information and statistics at a configurable interval. When the administrator applies configuration or policy changes to a Cisco WAE device or a group of Cisco WAE devices (a device group), the Cisco WAAS Central Manager automatically propagates the changes to each of the managed Cisco WAE devices. Cisco WAE devices that are not available to receive the update receive the update the next time they become available.

Cisco WAAS Central Manager User Interface

The Cisco WAAS Central Manager is used to centrally configure, manage, and monitor a topology of Cisco WAE devices. The Cisco WAAS Central Manager device mode can be configured on a Cisco WAE to provide scalable, secure, robust, centralized Web management for all Cisco WAE devices in the

deployment. The Cisco WAAS Central Manager provides device-specific and system-wide configuration, monitoring, and reporting capabilities, including policy configuration and distribution within the Cisco WAAS deployment.

The Cisco WAAS Central Manager GUI allows administrators to easily perform the following tasks:

- Configure system and network settings for an individual Cisco WAAS device or device group.
- Create, edit, and delete application policies that determine the actions (optimizations) that a Cisco WAAS device performs when it intercepts specific types of traffic.
- Create device groups to allow concurrent management and configuration of large groups of Cisco WAE devices.
- Configure role-based access control (RBAC) for management separation.
- Manage and acknowledge system and device alarms.
- View reports detailing the effectiveness of acceleration and optimization within a Cisco WAAS network.
- Examine per-connection statistics on an individual Cisco WAE device, including connection details, applied optimization policy, and detailed reduction statistics.

Figure 11 shows some of the reporting capabilities of the Cisco WAAS Central Manager—providing an initial dashboard view of application traffic types and bandwidth optimization upon login by the administrator. Deployment and Configurations does not go into the details of setting up a WAAS Central Manager; for more information, see the WAAS Configuration Guide at: http://www.cisco.com/en/US/docs/app_ntwk_services/waas/waas/v413/configuration/guide/maint.html#wp1159476.

Figure 11 Reporting Capabilities of the Cisco WAAS Central Manager

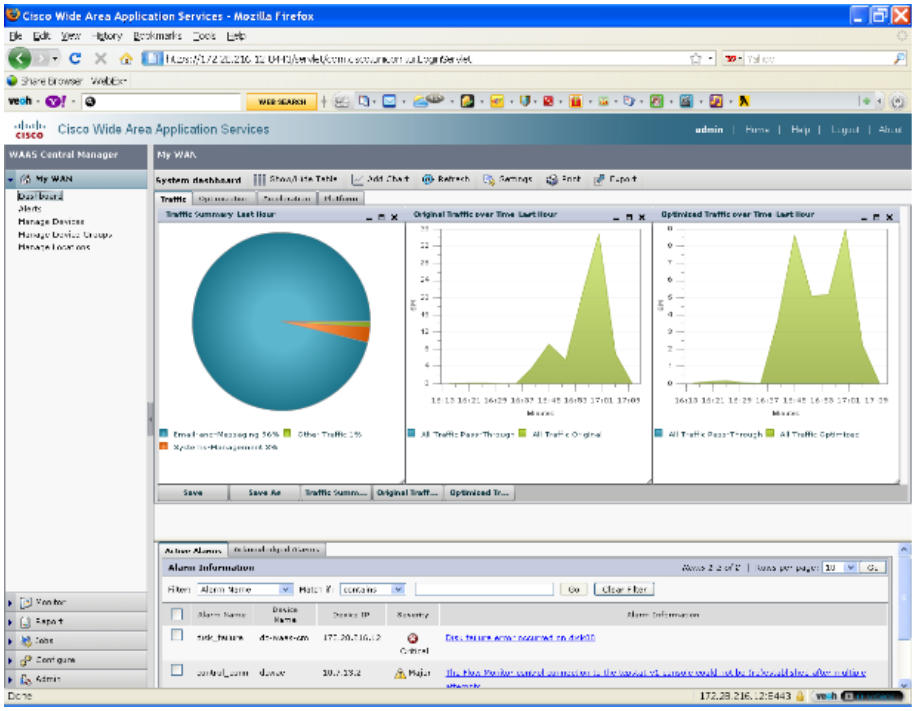

ACE Management

Application Network Manager

The Cisco Application Network Manager (ANM) is an appliance that offers GUI-based management and monitoring for ACE devices as well as the Cisco Content Services Switch (CSS), Content Switching Module (CSM), and Content Switching Module with SSL (CSM-S). Cisco ANM uses highly secure and logged Web-based access to facilitate delegation of tasks, such as service activation and suspension, to different users through administratively defined role-based access control (RBAC).

Cisco ANM 3.0 Guided Setup allows you to quickly perform the following tasks:

- Establish communication between Cisco ANM and Cisco ACE devices.
- Configure Cisco ACE devices that are new to the network by establishing network connectivity in either standalone or high-availability deployments.
- Create and connect to a Cisco ACE virtual context.
- Set up a load-balancing application from Cisco ACE to a group of back-end servers.

Direct configuration of contexts with out the use of the guided setup is also supported.

The Cisco ANM also supports advanced real time monitoring and historical graphing though its Monitoring Dashboards. Figure 12 shows an example of the ANM monitoring dash board.

Figure 12 ANM Monitoring Dash Board

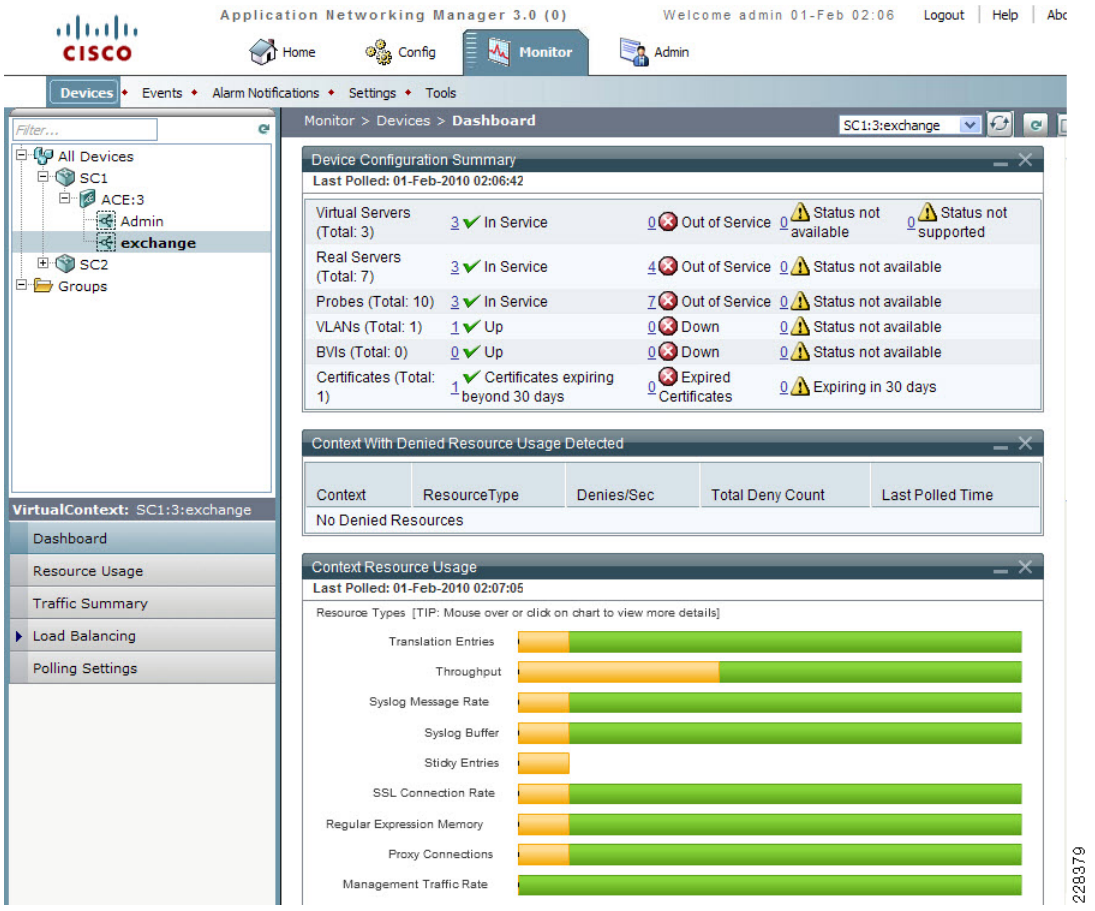

For more information on ANM, see:
http://www.cisco.com/en/US/prod/collateral/contnetw/ps5719/ps6904/data_sheet_c78-572610.html.

Data Center Network Architecture

This Cisco Validated Design uses the Cisco Data Center 3.0 Architecture as its data center design because of its strengths and features. The following highlights features of the Cisco Data Center 3.0 Architecture that were implemented in this Cisco Validated Design.

- Separate core, aggregation, and access layers for a multi-tier data center network architecture
- Virtual Device Contexts on the Nexus 7000 series routers in the core and aggregation
- Virtual Route Forwarding tables on the core and aggregation Nexus 7000 series routers
- Redundancy throughout the core, aggregation, and access layers with HSRP
- Leveraging Catalyst 6500 series switches as a Services Chassis to provide application optimization services with the Cisco Application Control Engine service module and to serve as the platform for adding on additional data center services.

For information on the Cisco Data Center 3.0 Architecture with details on leveraging the Catalyst 6500 series as the Service Chassis at the aggregation layer, refer to the following Cisco Validated Design:
http://www.cisco.com/en/US/solutions/ns340/ns414/ns742/ns743/ns994/landing_service_patterns.html

Figure 13 Data Center Multi-Tier Network Architecture Diagram

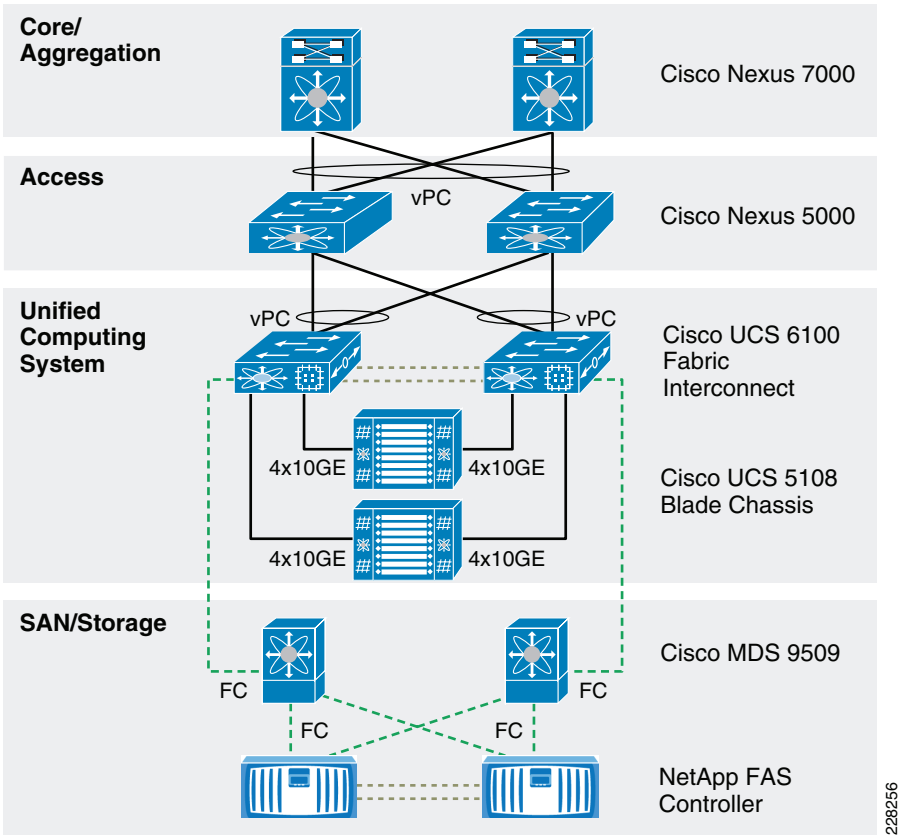

Data Center 10 Gigabit Ethernet LAN

Today's data centers are being transformed as the most recent server technologies are deployed in them. Multisocket servers with multicore processors deliver immense amounts of computing power in rack-mount servers and blade systems. This vast computing power has stimulated deployment of virtualization software that can more effectively put today's server capacity to use. The raw computing power of today's servers, combined with the higher utilization enabled by virtualization technology, places greater demands on IP networks than ever before.

The alignment of these two trends is encouraging the rapid adoption of 10 Gigabit Ethernet. Today, IT departments deploy 10 Gigabit Ethernet through the use of third-party adapter cards installed in server expansion slots. As server manufacturers respond to the demand for this level of bandwidth, an increasing number of rack-mount and blade systems will have multiple 10 Gigabit Ethernet ports integrated on the motherboard. Some blade systems now accommodate 10 Gigabit Ethernet switches integrated into the chassis. Other blade systems are likely to offer 10 Gigabit pass-through connections that make each blade's 10 Gigabit Ethernet ports available to access layer switches, eliminating oversubscription within the blade chassis.

Regardless of the mechanism by which 10 Gigabit Ethernet enters the data center, IT departments must consider a new set of trade-offs when deploying it in data center racks. 10 Gigabit Ethernet is more than just a means to Gigabit Ethernet; planning is needed to use it as effectively and efficiently as possible. Factors to consider include the following:

- As server density increases, so does I/O bandwidth. 10 Gigabit Ethernet provides the bandwidth needed by servers with powerful multicore processors.
- Gigabit Ethernet is often supported by access layer switches in intermediate distribution frames (IDFs) integrated into rows of server racks. With Gigabit Ethernet, the low cost of copper cabling made end-of-row switch configurations practical. With 10 Gigabit Ethernet, the cost of fiber transceivers and cabling is substantial. The alternative, 10GBASE-T, uses more expensive cabling, imposes higher latencies, and uses considerably more power. These constraints suggest use of a top-of-rack configuration with a better cabling solution.
- 10 Gigabit Ethernet may be accompanied with an increased demand for uplink capacity from the access to the aggregation layer. Indeed, as virtualization software simplifies the movement of virtual machines from server to server, IP networking requirements become more fluid and less predictable, requiring top-of-rack switching technology that can scale.

The Cisco Nexus 7000 Series Switches comprise a modular data center-class product line designed for highly-scalable 10 Gigabit Ethernet networks with a fabric architecture that scales beyond 15 terabits per second (Tbps). Designed to meet the requirements of the most mission-critical data centers, it delivers continuous system operation and virtualized pervasive services. The Cisco Nexus 7000 Series is based on a proven operating system, with enhanced features to deliver real-time system upgrades with exceptional manageability and serviceability. Its innovative design is purpose built to support end-to-end data center connectivity, consolidating IP, storage, and interprocess communication (IPC) networks onto a single Ethernet fabric.

The Cisco Nexus™ 5000 Series is a family of line-rate, low-latency, low-cost-per-port 10 Gigabit Ethernet/Fibre Channel over Ethernet (FCoE) switches for data center access-layer applications. The Cisco® Nexus 5000 Series supports a cost-effective top-of-rack deployment model that delivers both the port density and the bandwidth needed by servers in both top-of-rack and pod-style environments. This document presents several scenarios that illustrate the effectiveness of these switches in supporting both rack-mount and blade systems in high-density configurations.

Through the user of the Nexus 7000 and 5000 series routes and switches in the data center, a fully Ten Gigabit Ethernet LAN is implemented in this design. Given the elevated levels of application and management traffic that is generated as more application and management servers are consolidated in the data center, Ten Gigabit Ethernet in the data center provides sufficient bandwidth as QoS is designed and implemented throughout the data center to remote offices.

Data Center LAN Availability

Access layer is designed with the following key design attributes in Nexus 5000:

- Enables loop-less topology via vPC (Virtual Port-Channel) technology. The two-tier vPC design is enabled such that all paths from end-to-end are available for forwarding.
- Nexus 7000 to Nexus 5000 is connected via a single vPC between redundant devices and links. In this design four 10Gbps links are used, however for scalability one can add up to eight vPC member in current Nexus software release.
- The design recommendation is that any edge layer devices should be connected to Nexus 5000 with port-channel configuration.
- For details about the configuration and options to enable vPC, see: http://www.cisco.com/en/US/prod/collateral/switches/ps9441/ps9670/configuration_guide_c07-543563.html.
- RPVST+ is used as the spanning tree protocol. Redundant Nexus 7000s are the primary and secondary root for all VLANs with matching redundant default gateway priority.

Data Center Fiber—Channel SAN

The SAN consists of core and edge SAN storage layers to facilitate high-speed data transfers between hosts and storage devices. SAN designs are based on the FiberChannel (FC) protocol. Speed, data integrity, and high availability are key requirements in an FC network. In some cases, in-order delivery must be guaranteed. Traditional routing protocols are not necessary on FC.

Fabric Shortest Path First (FSFP), similar to OSPF, runs on all switches for fast fabric convergence and best path selection. Redundant components are present from the hosts to the switches and to the storage devices. Multiple paths exist and are in use between the storage devices and the hosts. Completely separate physical fabrics are a common practice to guard against control plane instability, ensuring high availability in the event of any single component failure.

SAN Core Layer

The SAN core layer provides high-speed connectivity to the edge switches and external connections. Connectivity between core and edge switches are 10 Gbps links or trunking of multiple full rate links for maximum throughput. Core switches also act as master devices for selected management functions, such as the primary zoning switch and Cisco fabric services. In addition, advanced storage functions such as virtualization, continuous data protection, and iSCSI reside in the SAN core layer.

SAN Edge Layer

The SAN edge layer is analogous to the access layer in an IP network. End devices such as hosts, storage, and tape devices connect to the SAN edge layer. Compared to IP networks, SANs are much smaller in scale, but the SAN must still accommodate connectivity from all hosts and storage devices in the data center. Over-subscription and planned core-to-edge fan out ratio result in high port density on SAN switches. On larger SAN installations, it is common to segregate the storage devices to additional edge switches.

Multilayer Director Switch 9509

High Availability

The Cisco MDS 9500 Series of Multilayer Directors was designed from the beginning for high availability. Beyond meeting the basic requirements of nondisruptive software upgrades and redundancy of all critical hardware components, the Cisco MDS 9500 Series software architecture offers an unparalleled level of availability. The Cisco MDS 9500 Series Supervisor-2 Module has the unique ability to automatically restart failed processes, making it exceptionally robust. In the rare event that a supervisor module is reset, complete synchronization between the active and standby supervisors ensures stateful failover with no disruptions to traffic.

The Cisco MDS 9500 Series includes 1+1 redundant crossbars. Each crossbar provides the necessary bandwidth to deliver full system performance ensuring that loss or removal of a single crossbar has no impact on system performance. Cisco MDS 9500 Series directors deliver maximum system performance, even in the event of a crossbar failure.

High availability is implemented at the fabric level with the industry's most robust and highest-performance ISLs. PortChannel capability allows users to aggregate up to 16 physical links into one logical bundle. The PortChannel can consist of any speed-matched ports in the chassis, ensuring that the bundle remains active in the event of a port, application-specific integrated circuit (ASIC), or module failure. The bundle can sustain the failure of any physical link without causing a reset. The Cisco MDS 9500 Series of multilayer directors takes high availability to a new level, ensuring solutions that exceeds the 99.999 percent uptime requirements of today's most demanding environments.

Virtual SANs

Ideal for efficient, secure SAN consolidation, VSANs allow more efficient SAN utilization by creating hardware-based isolated environments within a single SAN fabric or switch. Each VSAN can be zoned as a typical SAN and maintains its own fabric services for added scalability and resilience. VSANs allow the cost of SAN infrastructure to be shared among more users, while ensuring absolute segregation of traffic and retaining independent control of configuration on a VSAN-by-VSAN basis.

Integrated SAN Routing

In another step toward deploying efficient, cost-effective, consolidated storage networks, the Cisco MDS 9500 Series of multilayer directors supports IVR, the industry's first routing functionality for Fibre Channel. IVR allows selective transfer of data traffic between specific initiators and targets on different VSANs while maintaining isolation of control traffic within each VSAN. With IVR, data can transit VSAN boundaries while maintaining control plan isolation, thereby maintaining fabric stability and availability. Integrated IVR eliminates the need for external routing appliances, greatly increasing routing scalability while delivering line-rate routing performance, simplifying management, and eliminating the challenges associated with maintaining separate systems. Integrated IVR means lower total cost of SAN ownership.

Multiprotocol Intelligence

The Cisco MDS 9500 Series architecture enables multilayer and multiprotocol functionality, allowing it to transparently integrate multiple transport technologies for maximum flexibility. Beginning with Fibre Channel, FICON, iSCSI, and FCIP, the Cisco MDS 9500 Series is a robust, multiprotocol platform designed for deployment of cost-optimized storage networks. Today, users can implement up to 10-Gbps Fibre Channel or FICON for high-performance applications, iSCSI over Ethernet for cost-effective connectivity to shared storage pools, and FCIP for connectivity between data centers.

Ease of Management

To meet the needs of all users, the Cisco MDS 9500 Series provides three principal modes of management: Cisco MDS 9000 Family CLI, Cisco Fabric Manager, and integration with third-party storage management tools.

The Cisco MDS 9500 Series presents the user with a consistent, logical CLI. Adhering to the syntax of the widely known Cisco IOS® CLI, the Cisco MDS 9000 Family CLI is easy to learn and delivers broad management functionality. The Cisco MDS 9000 Family CLI is an extremely efficient and direct interface designed to provide optimal functionality to administrators in enterprise environments.

Cisco Fabric Manager is a responsive, easy-to-use Java application that simplifies management across multiple switches and fabrics. Cisco Fabric Manager enables administrators to perform vital tasks such as topology discovery, fabric configuration and verification, provisioning, monitoring, and fault resolution. All functions are available through a secure interface, enabling remote management from any location.

Cisco Fabric Manager may be used independently or in conjunction with third-party management applications. Cisco provides an extensive API for integration with third-party and user-developed management tools.

Data Center SAN Availability

Some issues to consider when designing a fibre channel SAN booted fabric include, but are not limited to, virtual SANs (VSANs), zone configurations, n-port virtualization, fan in/fan out ratios, high availability, and topology size. Each of these components, when not configured correctly, can lead to a fabric that is not highly available due to fibre channel requiring a loss-less nature. In this multi-tenant architecture, an improperly configured SAN impacts the boot OS and in turn tenant VMs and data sets.

A basic understanding of fibre channel fabrics is required for design of the SAN booted environment. Cisco VSANs are a form of logically partitioning a physical switch to segment traffic based on design needs. By deploying VSANs, an administrator can separate primary boot traffic from secondary traffic, ensuring reliability and redundancy. Additionally, as deployments grow, subsequent resources can be placed in additional VSANs to further aide in any segmentation needs from a boot or data access perspective. For instance, as a mutli-tenant environment grows beyond the capacity of a single UCSM, additional SAN booted hosts can be added without impacting existing compute blades or deploying new switches dependent upon port counts. Furthermore, the use of interVSAN routing or IVR enables and administrator to securely and logically associate resources even if they are not in the same VSAN.

Zoning within a fabric is used to prevent extraneous interactions between hosts and storage ports which can lead to very “chatty” fabric in which there is an abundance of initiator cross-talk. Through the creation of zones which exist in a given VSAN, a single port of an initiator can be grouped with the desired storage port to increase security, improve performance, and aid with the troubleshooting of the fabric. A typical SAN booted architecture consists of redundant fabrics (A and B) with primary and secondary boot paths constructed via zones in each fabric. Traditionally as SANs grow, the switches required increases to accommodate the port count needed. This is particularly true in legacy blade center environments as each fibre channel I/O module would constitute another switch to be managed with its own security implications. Additionally, from a performance perspective, this is a concern as each switch or VSAN within an environment has its own domain ID, adding another layer of translation. N-port ID Virtualization or NPIV is a capability of the fibre channel protocol that allows multiple N-ports to share a single physical port. NPIV is particularly powerful in large SAN environments as hosts that log into an NPIV-enabled device would actually be presented directly to the north-bound fabric switch. This improves performance and ease of management. NPIV is a component of the Fabric Interconnect within a UCS deployment and a requirement of any northbound FC switch.

The fan-in characteristics of a fabric is defined as the ratio of host ports that connect to a single target port, while fan-out is the ratio of target ports or LUNs that are mapped to a given host. Both are performance indicators, with the former relating to host traffic load per storage port and the latter relating storage load per host port. The optimum ratios for fan-in and fan-out are dependent on the switch, storage array, HBA vendor, and the performance characteristics of I/O workload. High availability within a FC fabric is easily attainable via the configuration of redundant paths and switches. A given host is deployed with a primary and redundant initiator port which is connected to the corresponding fabric. With a UCS deployment, a dual port mezzanine card is installed in each blade server and a matching vHBA and boot policy are setup providing primary and redundant access to the target device. These ports access the fabric interconnect as N-ports which are passed along to a northbound FC switch. Zoning within the redundant FC switches is done such that if one link fails then the other handles data access. Multipathing software is installed dependent on the operating system which ensures LUN consistency and integrity.

When designing SAN booted architectures, considerations are made regarding the overall size and number of hops that an initiator would take before it is able to access its provisioned storage. The fewer hops and fewer devices that are connected across a given interswitch link, the greater the performance of a given fabric. A common target ratio of hosts across a given switch link would be between 7:1 or 10:1, while an acceptable ratio may be as high as 25:1. This ratio can vary greatly depending on the size of the architecture and the performance required.

SAN Connectivity should involve or include:

- The use of redundant VSANs and associated zones
- The use of redundant interswitch links ISLs where appropriate
- The use of redundant target ports
- The use of redundant fabrics with failover capability for fiber channel SAN booted infrastructure

Data Center Storage

SAN Storage with NetApp FAS 6070 and Data ONTAP

The NetApp FAS controllers share a unified storage architecture based on the Data ONTAP® 7G operating system and use an integrated suite of application-aware manageability software. This provides efficient consolidation of SAN, NAS, primary, and secondary storage on a single platform while allowing concurrent support for block and file protocols using Ethernet and Fibre Channel interfaces, including FCoE, NFS, CIFS, and iSCSI. This common architecture allows businesses to start at an entry level storage platform and easily migrate to the higher-end platforms as storage requirements increase, without learning a new OS, management tools, or provisioning processes.

To provide resilient system operation and high data availability, Data ONTAP 7G is tightly integrated to the hardware systems. The FAS systems use redundant, hot-swappable components, and with the patented dual-parity RAID-DP (high-performance RAID 6), the net result can be superior data protection with little or no performance loss. For a higher level of data availability, Data ONTAP provides optional mirroring, backup, and disaster recovery solutions. For more information, see: http://www.netapp.com/us/products/platform-os/data-ontap/.

While this solution focuses on specific hardware, including the FAS6070, any of the FAS platforms, including the FAS 6080, FAS6040, FAS3140, and FAS3170, are supported based on your sizing requirements and expansion needs with all of the same software functionality and features. Similarly, the quantity, size, and type of disks used within this environment may also vary depending on storage and performance needs. NetApp generally recommends NetApp FAS 3000 series for Exchange

deployments. Additional add-on cards, such as the Performance Accelerator Modules (PAM II), can be utilized in this architecture to increase performance by adding additional system cache for faster data access.

For the scope of this design, the NetApp FAS 6070 with FC disks was used to provide three primary functionalities:

- Boot from SAN capability to each physical server
- Storage for Windows 2008 R2 Cluster Shared Volumes to support the Failover Cluster configuration in this design
- Storage for Mailbox database and log files

Further details on these functionalities are provided in Data Center Server Platform—Cisco Unified Computing System and "Server Virtualization Platform: Microsoft Hyper-V R2".

NetApp Storage Availability

RAID groups are the fundamental building block when constructing resilient storage arrays containing any type of application data set or virtual machine deployment. There exists a variety of levels of protection and costs associated with different RAID groups. A storage controller that offers superior protection is an important consideration to make when designing a multi-tenant environment as hypervisor boot, guest VMs, and application data sets are all deployed on a a shared storage infrastructure. Furthermore, the impact of multiple drive failures is magnified as disk size increases. Deploying a NetApp storage system with RAID DP offers superior protection coupled with an optimal price point.

RAID-DP is a standard Data ONTAP feature that safeguards data from double disk failure by means of using two parity disks. With traditional single-parity arrays, adequate protection is provided against a single failure event such as a disk failure or error bit error during a read. In either case, data is recreated using parity and data remaining on unaffected disks. With a read error, the correction happens almost instantaneously and often the data remains online. With a drive failure, the data on the corresponding disk has to be recreated, which leaves the array in a vulnerable state until all data has been reconstructed onto a spare disk.

With a NetApp array deploying RAID-DP, a single event or second event failure is survived with little performance impact as there exists a second parity drive. NetApp controllers offer superior availability with less hardware to be allocated. Aggregates are concatenations of one or more RAID groups that are then partitioned into one or more flexible volumes. Volumes are shared out as file level (NFS or CIFS) mount points or are further allocated as LUNs for block level (iSCSI or FCP) access. With NetApp's inherent storage virtualization, all data sets or virtual machines housed within a shared storage infrastructure take advantage of RAID-DP from a performance and protection standpoint. For example, with a maximum UCS deployment there could exist 640 local disks (two per blade) configured in 320 independent RAID-1 arrays all housing the separate hypervisor OS. Conversely, using a NetApp array deploying RAID-DP, these OSes could be within one large aggregate to take advantage of pooled resources from a performance and availability perspective.

Much as an inferior RAID configuration is detrimental to data availability, the overall failure of the storage controller serving data can be catastrophic. Combined with RAID-DP, NetApp HA pairs provide continuous data availability for multi-tenant solutions. The deployment of an HA pair of NetApp controllers ensures the availability of the environment both in the event of failure and in the event of upgrades.

Storage controllers in an HA pair have the capability to seamlessly take over its partner's roles in the event of a system failure. These include controller personalities, IP addresses, SAN information, and access to the data being served. This is accomplished using cluster interconnections, simple administrative setup, and redundant paths to storage. In the event of an unplanned outage, a node

assumes the identity of its partner with no reconfiguration required by any associated hosts. HA pairs also allow for non-disruptive upgrades for software installation and hardware upgrades. A simple command is issued to takeover and giveback identity.

The following considerations should be made when deploying an HA pair:

- Best practices should be deployed to ensure any one node can handle the total system workload.
- Storage controllers communicate heartbeat information using a cluster interconnect cable.
- Takeover process takes seconds.
- TCP sessions to client hosts are re-established following a timeout period.
- Some parameters must be configure identically on partner nodes.

For additional information regarding NetApp HA pairs, see:
http://media.netapp.com/documents/clustered.pdf.

Deployment and Configurations

This section describes how to configure the solution components and shows the specific configurations that were implemented during the validation of this solution. The information is presented in the order in which the solution should be implemented based on dependencies among the solution components. For example, the configuration of the Service Profiles on the Unified Computing System requires specifying parameters specific to your SAN storage—so it is important to have the SAN storage components finalized first. Where there are no such dependences, that is specifically mentioned so that the user of this deployment guide can know when they can jump back from one section to another without being concerned about sequence. While this section provides details on CLI commands and GUI steps, it is not meant to be a configuration guide for any of the products in this solution as existing product documentation details how to configure the various features of each of these products. Therefore, this section details commands and GUI screenshots for configuration steps that will save the reader time, but where the instructions would be too lengthy for this document, URLs are provided for existing documentation.

Deployment Phases

These deployment sections assume that the Data Center 3.0 topology has been built out following design guidance provided in existing Data Center 3.0 design guides (see References). It also assumes that the design of the Exchange 2010 server roles and virtual machines has been done.

Deployment of this solution can be organized into the following major phases, in the logical order of deployment:

- Connect and configure Unified Computing System
  - System Name, Security, and Management IP Address
  - Fabric Interconnects Ports
  - Boot Policy
  - UUID, WWPN, and MAC Address Pools
  - VLANs
  - vNIC and vHBA Templates
  - Service Profiles

- LAN switches
- Configure NetApp storage array
  - Aggregates, Volumes, LUNs, and initiator groups
- Configure SAN zones
- Provisioning Hyper-V Servers with UCS Service Profiles
- Setting up Hyper-V Servers to support Live Migration of Virtual Machines
- Creating Virtual Machines and installing Exchange 2010 Server Roles
- Adding in Application Control Engine (ACE) hardware load balancing
- Adding in Wide Area Application Services between branch and data center
- Adding in Wide Area Application Services between mobile users and data center

Network Topology

The network topology in Figure 14 was built in the lab to represent a branch, a data center, and a remote user/mobile worker.

Figure 14 Network Topology—Branch, Data Center, and Remote User

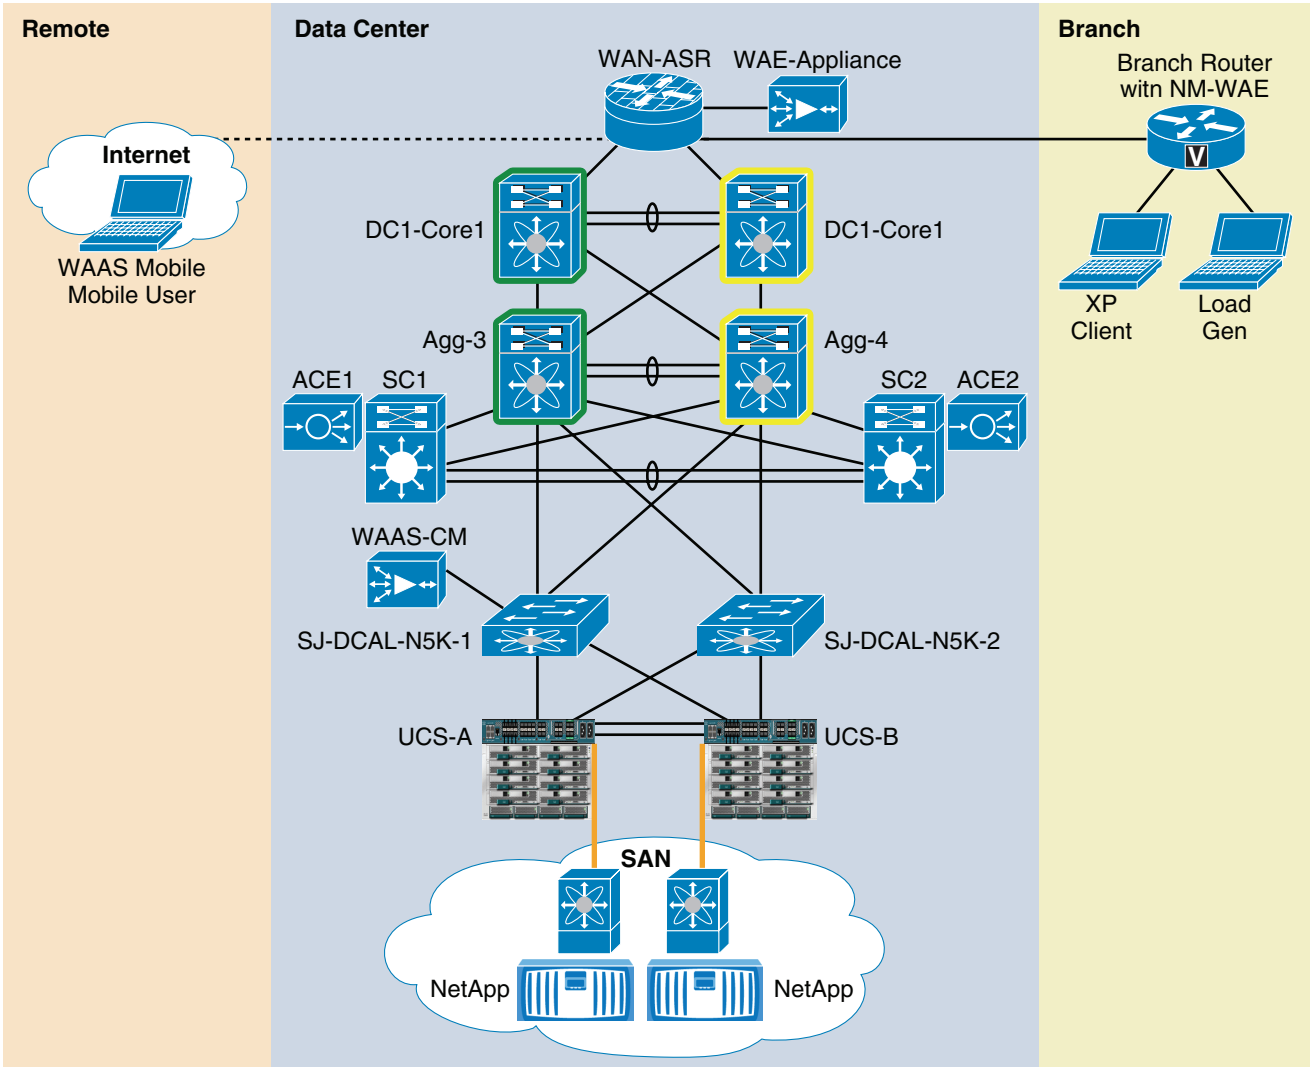

Use the following key to map the labels in Figure 14 to the information provided in this section:

- WAN-ASR = WAN Aggregation Services Router
- WAE Appliance = Wide Area Application Services appliance
- NM-WAE = Wide Area Application Services network module
- Branch-Router = Integrated Services Router
- DC1-Core1 = virtual device context on the Nexus 7000 for implementing redundant core switches
- Agg-3, Agg-4 = virtual device contexts on the Nexus 7000 for implementing redundant aggregation switches
- Green color lines around DC1-Core1 and Agg-3 = two virtual device contexts defined on the first physical Nexus 7000 switch to implement the first set of core and aggregation layer switches
- Yellow color lines around the DC1-Core1 and Agg-4 = two virtual device contexts defined on the second physical Nexus 7000 switch to implement the second set of core and aggregation layer switches.

- SC1 = First Catalyst 6500 series switch configured as a Services Chassis
- ACE1 = First Application Control Engine service module in the SC1 Services Chassis
- SC2 = Second Catalyst 6500 series switch configured as a Services Chassis
- ACE2 = Second Application Control Engine service module in the SC2 Services Chassis
- SJ-DCAL-N5k-1 = First Nexus 5000 series access switch
- SJ-DCAL-N5k-2 = Second Nexus 5000 series access switch
- UCS-A, UCS-B = Each is a Unified Computing System 5108 chassis fully loaded with 8 half-width server blades connected to the access layer by a Unified Computing System 6120 Fabric Interconnect.
- Orange line to SAN = Fiber Channel connectivity to MDS 9500 series fabric switches connected to NetApp storage. For further details on SAN design guidance, see this storage networking design site: http://www.cisco.com/en/US/netsol/ns747/networking_solutions_sub_program_home.html.

WAN simulation of bandwidth size, latency, and packet loss was provided through use of the open source WAN Bridge simulation tool, also available for software download with a valid user account from the Cisco Partner Tools and Resources—>Software Downloads site at http://www.cisco.com/en/US/partner/support/tsd_most_requested_tools.html. The WAN Bridge machine was placed between the WAN Aggregation Services router and the branch Integrated Services router for branch to data center testing and then between the remote user and the data center for remote (Internet) user testing.

Note The WAN aggregation layer of the data center should include separate Aggregation Services Routers for Internet and branch access, which is not provided in this solution as security is outside this scope. For best practices on building out a secure Enterprise WAN edge network, see the following design guide: http://www.cisco.com/en/US/docs/solutions/Enterprise/Security/SAFE_RG/chap7.html#wp1052953.

The branch included an Integrated Services Router (ISR) equipped with a Wide Area Application Services (WAAS) Engine network module. For details on configuring a WAAS network module on an ISR, see “Configuring Cisco WAAS for Network Modules for Cisco Access Routers” at http://www.cisco.com/en/US/docs/app_ntwk_services/waas/waas/v403/module/configuration/guide/wsnmecfg.html. Connected to the ISR on the Gigabit LAN is a standalone server that hosted several Hyper-V virtual machines for testing branch to data center traffic:

- One Windows XP client with Outlook 2010
- Four 64-bit Windows 2008 Enterprise servers with Exchange LoadGen 2010 Beta installed

The remote user consisted of a One Windows XP client with Outlook 2010 and Wide Area Application Services Mobile client software installed.

The data center consisted of core, aggregation, and access tiers—with the core and aggregation tiers implemented in the form of virtual device contexts on two physical Nexus 7000 series switches. The Layer 2 access tier consisted of redundant Nexus 5000 series switches. For details on this Data Center 3.0 network architecture, see the Data Center 3.0 IP Network Infrastructure design guide at http://www.cisco.com/en/US/solutions/ns340/ns414/ns742/ns743/ns994/landing_dc_ip_ntwk_infra.html.

To provide Application Control Engine (ACE) hardware load-balancing, a Catalyst 6500 series switch housing an Application Control Engine service module was configured as a Services Chassis connected to each Nexus 7000 aggregation switch through port channels. For further details on implementing the

Services Chassis, see the Data Center Services Patter design guide at http://www.cisco.com/en/US/docs/solutions/Enterprise/Data_Center/DC_3_0/dc_serv_pat.html#wp1023222.

To provide Wide Area Application Services from data center to branch users, a Wide Area Application Services Engine (WAAS) appliance was deployed at the WAN aggregation layer connected one-armed to the WAN Aggregation Services Router (ASR). The WAN ASR is configured with Layer 2 Web Communication Cache Protocol (WCCP) to redirect traffic to the WAAS appliance for compression, optimization, and application protocol acceleration. For further details on configuring the Aggregation Services Router with WCCP in the deployment of the WAAS appliance, see the following section in the Cisco IOS XE IP Application Services Configuration Guide: http://www.cisco.com/en/US/netsol/ns747/networking_solutions_sub_program_home.html.

Data Center 3.0 Configuration

Data center networking technology is currently an area of rapid change. Higher-performance end nodes and the migration to 10-Gigabit Ethernet for edge connectivity and Fiber Channel over Ethernet (FCoE) for convergence are changing design standards while virtualization capabilities are expanding the tools available to the network architect. When designing the data center network, a solid hierarchical foundation provides for high availability and continued scalability. This foundation also provides the flexibility to create different logical topologies utilizing device virtualization, the insertion of service devices, as well as traditional Layer 3 and Layer 2 network configuration.

A structured data center environment uses a physical layout that correlates tightly to the hierarchy of the network topology. Decisions on cabling types and the placement of patch panels and physical aggregation points must match the interface types and densities of the physical switches being deployed. In a new data center build out, the two can be designed simultaneously, also taking into consideration the constraints of power and cooling resources. When seeking to avoid significant new investment within an existing data center facility, the pre-existing physical environment of cabling, power, and cooling can strongly influence the selection of switching platforms. Careful planning in conjunction with networking requirements and an eye toward flexibility for the future is critical when designing the physical data center environment. Taking a modular approach to data center design provides flexibility and scalability in both network topology design and utilization of physical resources.

This document focuses on the Cisco Unified Computing System, Exchange 2010 components Application Control Engine (ACE), and the Wide Area Application Services (WAAS), and not on the Datacenter Network infrastructure design. The network design is based on the Cisco Data Center 3.0 Network design. In depth design and deployment recommendation detail on this architecture is already covered in the document Data Center Design-IP Network Infrastructure at: http://www.cisco.com/en/US/docs/solutions/Enterprise/Data_Center/DC_3_0/DC-3_0_IPInfra.pdf.

If you are not already familiar with Cisco's DC 3.0 architecture, refer to this document and review the sections Data Center Network Topologies, Layer 3 Design and Features, and Layer 2 Design and Features.

Cisco Unified Computing System Configuration

UCS 6120 Fabric Interconnects

The UCS 6120 Fabric Interconnects are first put into clustered mode. This can be done if the UCS service profiles have already been created without requiring the service profiles to be reconfigured.

To enable cluster pairing on each UCS Fabric Interconnect:

1. Identify four static management IP addresses: one cluster management IP address, one IP address for mgmt0 on each interconnect, and one for the gateway address for each interconnect.
2. Make sure you connect your two fabric interconnects together through their Layer 1 and Layer 2 ports so they can detect each other as peers after clustering is enabled.
3. If your fabric interconnects are unconfigured, you access each one via their serial port console to go through the initial setup wizard that configures some basic parameters, including the static IP addresses mentioned above. The wizard prompts you to specify if you want the fabric interconnect as part of a cluster. Be sure to answer "yes" to that prompt. The following is an example of the prompts you see in the wizard:

```
Enter the installation method (console/gui)?  console
Enter the setup mode (restore from backup or initial setup) [restore/setup]? setup
You have chosen to setup a new switch.  Continue? (y/n): y
Enter the password for "admin": adminpassword%958
Confirm the password for "admin": adminpassword%958
Do you want to create a new cluster on this switch (select 'no' for standalone setup
or if you want this switch to be added to an existing cluster)? (yes/no) [n]: yes
Enter the switch fabric (A/B): A
Enter the system name: interconnectA
Mgmt0 IPv4 address: 192.168.10.10
Mgmt0 IPv4 netmask: 255.255.255.0
IPv4 address of the default gateway: 192.168.10.1
Virtual IPv4 address : 192.168.10.12
Configure the DNS Server IPv4 address? (yes/no) [n]: yes
  DNS IPv4 address: 20.10.20.10
Configure the default domain name? (yes/no) [n]: yes
  Default domain name: domainname.com
Following configurations will be applied:
  Switch Fabric=A
  System Name=foo
  Management IP Address=192.168.10.10
  Management IP Netmask=255.255.255.0
  Default Gateway=192.168.10.1
  Cluster Enabled=yes
  Virtual Ip Address=192.168.10.12
  DNS Server=20.10.20.10
  Domain Name=domainname.com
Apply and save the configuration (select 'no' if you want to re-enter)? (yes/no): yes
```

   After applying the configuration, you must then connect to your second fabric interconnect through its console port and go through basic setup.

   For more information on the initial setup configuration, see: http://www.cisco.com/en/US/docs/unified_computing/ucs/sw/cli/config/guide/1.0.2/CLI_Config_Guide_1.0.2_chapter4.html#concept_4E8AA2C1ED584D808E830EB6D352DCF6.
4. If your fabric interconnects have already gone through initial setup via the console, you have to open up an SSH or Telnet session to the IP address you already assigned to the UCS Manager. This gives you access to CLI commands to use to complete the cluster pairing.
5. In the CLI, issue the following command to put both your fabric interconnects into cluster mode if they were originally configured for standalone:

```
UCS-A# connect local-mgmt
UCS-A(local-mgmt)# enable cluster 192.168.1.101
This command will enable cluster mode on this setup. You cannot change it
back to stand-alone. Are you sure you want to continue? (yes/no): yes
UCS-A(local-mgmt)#
```

UCS B200-M1 Server Blades and Service Profiles

Pre-Configuration Checklist

- Verify you have upgraded your UCS to the latest firmware version. This design was validated with 1.0.1(e) and 1.0.2(d), but it is recommended that the latest version be used for the most recent bug fixes and features. This deployment guide does not cover the steps on how to upgrade the UCS firmware. To update the Cisco UCS firmware, refer to the "Firmware Management" chapter in either of the following documents or the appropriate version more recent than 1.0.2 for your firmware version:
  - Cisco UCS Manager CLI Configuration Guide: http://www.cisco.com/en/US/docs/unified_computing/ucs/sw/cli/config/guide/1.0.2/b_CLI_Config_Guide_1.0.2.html
  - Cisco UCS Manager GUI Configuration Guide: http://www.cisco.com/en/US/docs/unified_computing/ucs/sw/gui/config/guide/1.0.2/b_GUI_Config_Guide.html
- Determine whether you have Qlogic or Emulex Converged Network Adapters on your blades. You need to know this so you load the proper drivers during different steps in the installation. Obtain a copy of the Qlogic and/or Emulex CNA drivers for Windows 2008 64-bit from the Cisco Software Download site for Unified Computing System and the tools and drivers bundles for the UCS adapters: http://tools.cisco.com/support/downloads.
- If you have a mix of Qlogic and Emulex CNAs on your set of blades, be sure you take note of this when moving service profiles from blade to blade, so that you do not accidentally move a SAN Boot-based Windows installation to a blade that has the wrong HBA drivers.
- Create a Management IP Pool of addresses that is on the same subnet as the IP addresses you selected for the fabric interconnect cluster and for the individual interconnects. This pool of management IP addresses are needed for UCS Manager to manage each of the blade servers brought online when associated with service profiles.
- Determine the range of values you will specify for the following parameters you need in your Service Profiles: WWPNs for the vHBAs, MAC addresses for the vNICs, UUIDs.
- Create pools of resources for your service profiles for the following attribute types: Server blades, WWNN (UUID), WWPN, MAC address. If you decide to assign WWPNs to the vHBAs using a WWPN pool to randomly assign values to the Service Profiles, then you should configure your NetApp storage (see Setting Up NetApp Storage) after you finish create your Service Profiles, as you need to make sure the WWPNs you specify on your fabric switches and in your NetApp storage system Initiator Groups all match.
- Through the LAN tab of UCS Manager, create the VLANs that you will use for your service profiles beforehand for easier configuration. Do the same for your VSANs using the SAN tab of UCS Manager. You have the option of creating them later during the defining of your service profiles.
- Obtain the Target WWPN(s) for your NetApp storage system(s) and the LUN IDs for the storage partitions you created previously for the following purposes:
  - Cluster Shared Volumes for Windows Failover Clustering and Live Migration
  - Boot LUNs so servers installed on a UCS blade can be portable to any other UCS blade

Defining Your Service Profiles for Portability and High Availability

This section of the document is not meant to give step-by-step instructions on how to configure a Service Profile. It is meant to point out the important configuration steps tested during solution validation that provided Fabric Interconnect failover for LAN uplink availability to each server blade and that allowed the Service Profiles (and hence the entire server installation on the blade) to be portable to other blades in the system.

For detailed instructions on configuring Service Profiles and associating them with the UCS blades using the UCS Manager GUI interface, see "Cisco UCS Manager GUI Configuration Guide, 1.0(2)" at: http://www.cisco.com/en/US/docs/unified_computing/ucs/sw/gui/config/guide/1.0.2/b_GUI_Config_Guide.html.

It is useful to note that the service profiles are not specific to the Microsoft Exchange 2010 or other enterprise applications that need to be virtualized on UCS. Defining a service profile and assigning it to a UCS blade is equivalent to choosing a standalone hardware platform, installing the appropriate firmware, and installing the appropriate network interfaces on it for network connectivity required.

There are a couple of things you can do to simplify your configuration of multiple UCS Service Profiles.

The factors that need to be considered when defining the service profiles are related to LAN and SAN connectivity and boot order for your server blade.

- Since each server blade has an uplink to each of the two UCS 6120 Fabric Interconnects to failover to a secondary LAN uplink, the Service Profile should have one Ethernet vNIC set to "Fabric A" and the second Ethernet vNIC set to "Fabric B" for the "Fabric ID" parameter in the "General"—>"Properties" section of the Service Profile right-side pane (Figure 15).

Figure 15 Fabric Interconnect Settings on vNIC

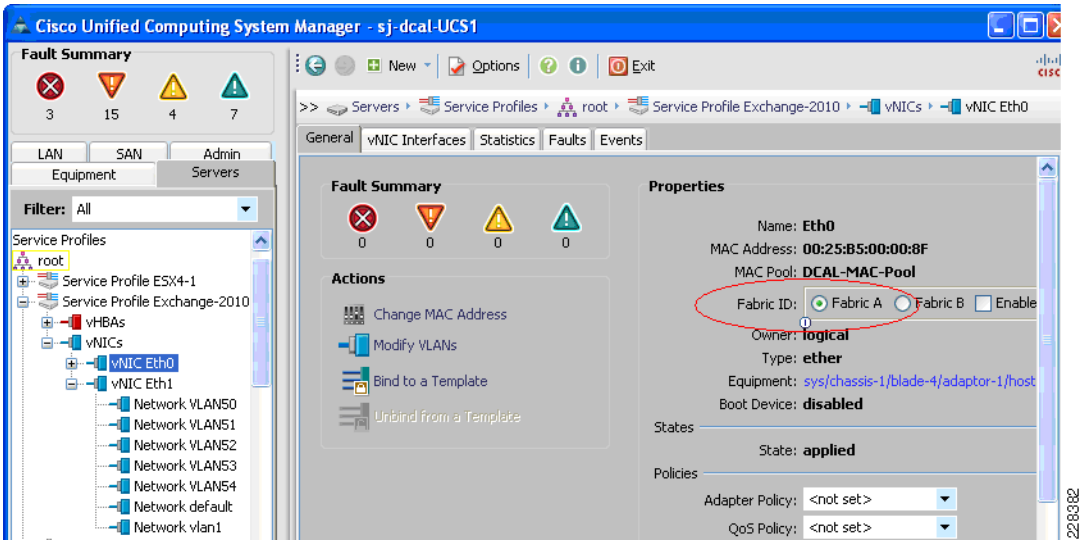

- Create the blade server management VLAN as the native VLAN for each vNIC.
- The VSAN that has been defined on the SAN fabric for connectivity of the Exchange server farm to the NetApp LUNs needs to be set in the service profile.
- The World Wide Port Numbers (WWPNs) configured for the Initiator Groups on the NetApp storage system need to be assigned to the vHBAs in the UCS Service Profiles.

- Since each UCS blade bootd from a NetApp LUN, boot from SAN needs to be configured in the server blade boot policy. Figure 16 shows the UCS Manager window for configuring a boot policy order for SAN boot. Note that boot from virtual media should always be configured first so there is always the option to boot from CD/DVD to recover a server. Also note that boot policy includes the WWPN of the SAN Target and LUN ID 0 matching the LUN 0 assignment on the NetApp storage system.

Figure 16 Configuring Boot Policy Order for SAN Boot

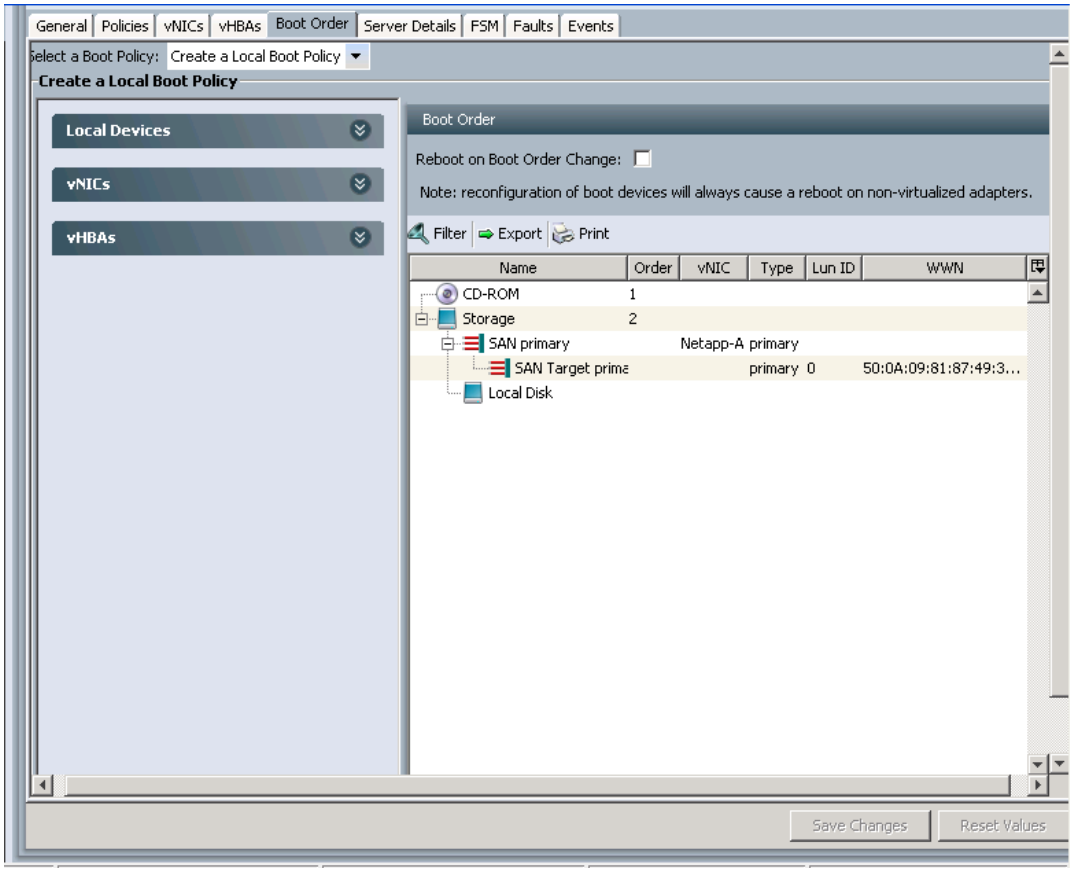

After the above items have been configured, the UCS service profile view of the vNIC interfaces with the required VLANs and the vHBA interface associated with the correct VSAN look like what is shown in Figure 17.

Figure 17 UCN Service Profile View of vNIC Interfaces

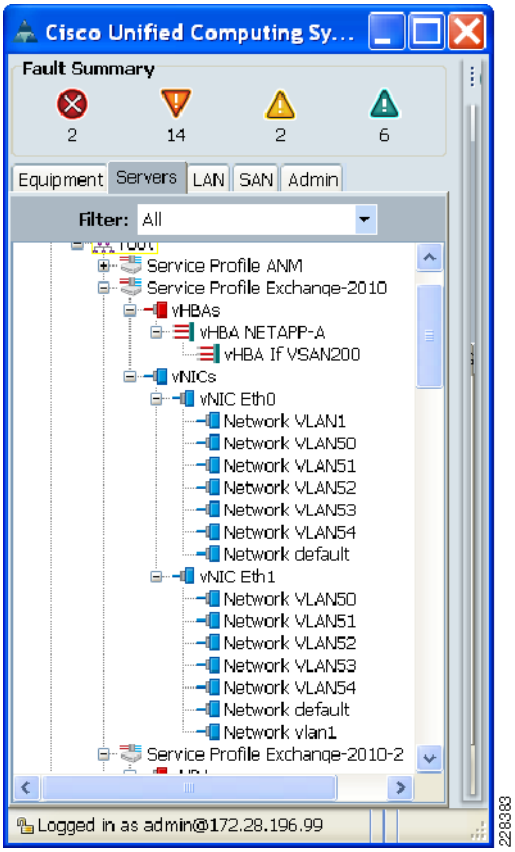

Since the Service Profiles for the seven blades needed to validate the Exchange Server scenario all require the same LAN and SAN connectivity and they are targeting the NetApp storage system's WWPN through VSAN 200 to boot off of their LUN 0, it is possible to create just one Service Profile from scratch using the "Create Service Profile (expert)" option on UCS Manager and clone this one service profile six more times using the "Create a Clone" option.

Note If you decide to clone your existing service profile, you need to have created pools of resources for the profile parameters of UUID/WWNN, WWPN, and MAC address and you need to assign your existing service profile to use the pools of resources for those parameters. When you clone that profile, the creation process of the clone automatically assigns a new set of those parameters from the available resources in your pools. If you manually assign a specific value for any of those parameters (e.g., when you specify the UUID, you choose the option to do "Manual using OUI"), then you may end up with a NULL/invalid value for your parameters.

As each Service Profile is created, you can select "Change Service Profile Association" in the "General" right-side pane of the Service Profile UCS Manager window and assign the service profile to a blade—either dynamically from a pool of available blades or a specific blade of choice.

LAN Switches

The Nexus 5000 access switches northbound of the UCS Fabric Interconnects are configured with the VLANs to support the different traffic types mentioned earlier (i.e., Live migration, cluster interconnect, virtual machine management, virtual machine access, server blade management).

```
vlan 50
   name live-migration
vlan 51
   name cluster
vlan 52
   name vmm
vlan 53
   name server-mgmt
vlan 54
   name vm-access
```

The 10GE ports for the fabric interconnect uplinks are configured as trunks to allow the VLANs defined above.

```
interface Ethernet1/1
   switchport mode trunk
   switchport trunk allowed vlan 50-54
   spanning-tree port type edge trunk
   logging event port link-status
   logging event port trunk-status

interface Ethernet1/2
   switchport mode trunk
   switchport trunk allowed vlan 50-54
   spanning-tree port type edge trunk
   logging event port link-status
   logging event port trunk-status

interface Ethernet1/3
   switchport mode trunk
   switchport trunk allowed vlan 50-54
   spanning-tree port type edge trunk
   logging event port link-status
   logging event port trunk-status

interface Ethernet1/4
   switchport mode trunk
   switchport trunk allowed vlan 50-54
   spanning-tree port type edge trunk
   logging event port link-status
   logging event port trunk-status
```

Setting Up NetApp Storage

Before beginning to install Windows OS on each UCS blade, NetApp storage needs to be set up with the proper number of LUNs to support the requirements of the Exchange servers. This document assumes the following:

- The user has a basic knowledge of using the NetApp FilerView browser-based utility for allocate storage on the NetApp storage system.

- The appropriate number of disks on the NetApp storage system has been assigned to an aggregate of disks that were used to build this solution. It is assumed that two volumes of sufficient size have been created from that disk aggregate, one for each type of LUN that needs to be created, described below.
- There is only a single path available to the LUN during the SAN boot installation. Control this with zoning or HBA activation. After the OS is installed, install the MPIO feature and make sure the device-specific module (DSM) claims the LUNs. Re-enable the redundant paths at this point.

The following types of LUNs must be created on the NetApp storage system. (Note that the NetApp storage system itself does not have awareness of these types. The distinction of these two types of LUNs is only important for how the LUNs are used in this solution.)

- **Boot LUNs**—Since there are seven UCS blades that needs to be set up for the scenario described in Exchange 2010 Server and Client Scenario for This Solution, there need to be seven LUNs dedicated as boot LUNs for the Windows operating system installations. Each boot LUN should be assigned as LUN 0 through the NetApp "Manage LUNs" interface in the FilerView browser interface. As you can see in Figure 18, the two servers that are being assigned LUNs (dcal_ucs_svr3 and dcal_ucs_svr 4) are assigned LUN 0 as each of their boot LUNs.
- **Cluster Shared Volume LUNs**—Microsoft Cluster Shared Volumes are required to support Hyper-V live migration of a virtual machine from one physical server to another. These Cluster Shared Volume LUNs on the NetApp storage system should be given access to those machines that are participating in live migration. As you can see in Figure 18, in the "Manage LUNs" window of the browser-based NetApp FilerView, the two servers that are participating in Cluster Shared Volumes (dcal_ucs_svr3 and dval_ucs_svr4) are both mapped to /vol/dcal_cluster_luns/winshare1 and /vol/dcal_cluster_luns/winshare2, which means both servers can access those two LUNs.

  Microsoft has detailed documentation on setting up Cluster Shared Volumes for Hyper-V virtual machines in their "Using Cluster Shared Volumes in a Failover Cluster in Windows Server 2008 R2" at: http://technet.microsoft.com/en-us/library/ff182346(WS.10).aspx.
- **Mailbox Database LUNs**—Microsoft recommends that the mailbox database files and transaction log files be stored on separate disk spindles from the Windows operating system for optimal performance. Therefore, it is recommended that a third set of LUNs be allocated for this purpose.

Figure 18 Manage LUNs Window

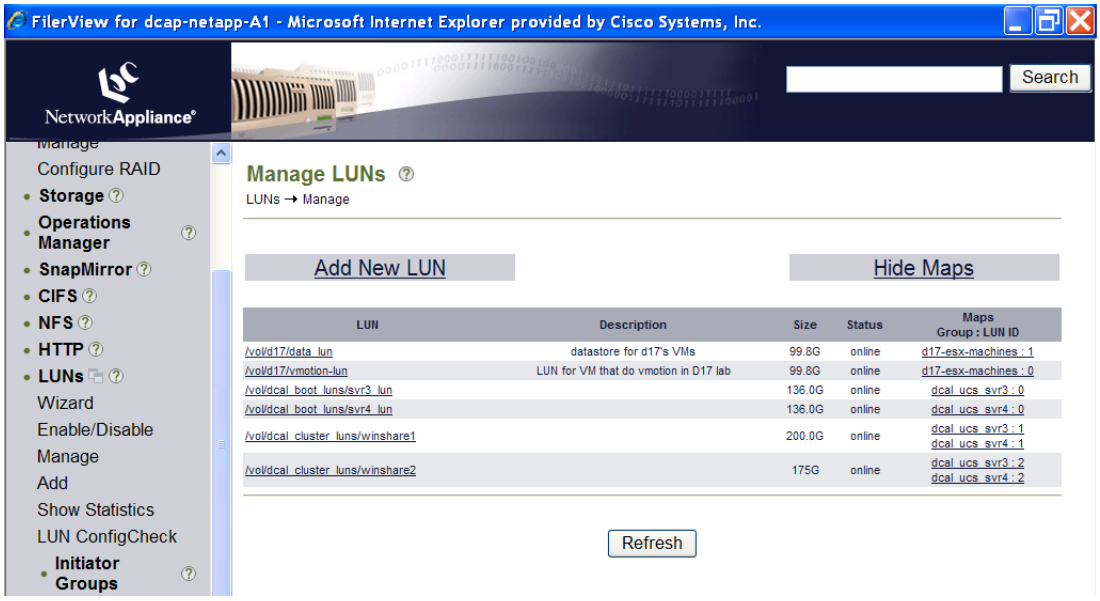

The mapping of LUNs to Initiator Groups on the NetApp storage system is relevant to the UCS Service Profile you create for each server blade later. To map the LUNs, you first need to create Initiator groups that specify the World-Wide Port Numbers (WWPN) of the UCS server blades, then you map the LUNs to the appropriate Initiator Groups. In Figure 18, the servers identified as dcal_ucs_svr3 and dcal_ucs_svr4 are defined in LUN Initiator Groups on the NetApp storage system; these initiator groups specify the WWPNs of the servers that needs to access the LUNs.

Since UCS Service Profiles allow you to assign arbitrary WWPNs to the UCS blade hardware, you can configure your NetApp storage system beforehand with these WWPNs and then make sure you assign them in your UCS service profiles to the blade HBAs at a later step.

Figure 19 Initiator Groups Defined for NetApp LUNs

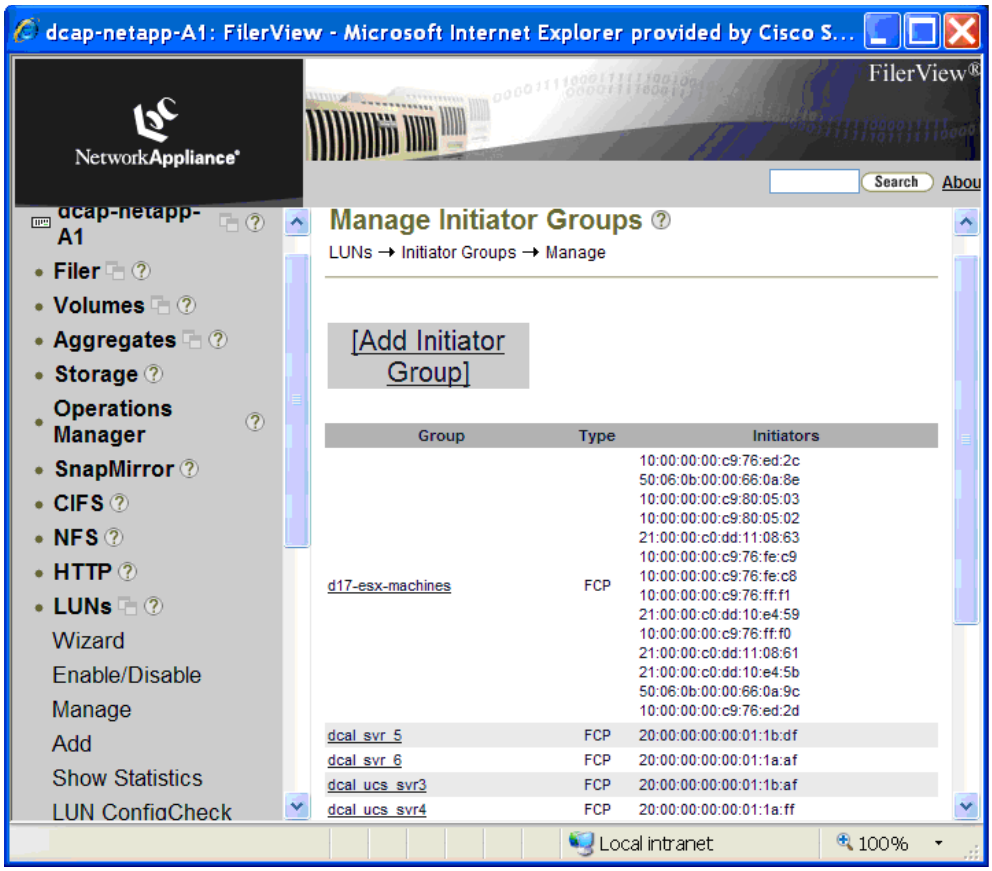

SAN Fabric Connectivity

This document does not detail how to configure a secure and highly-available SAN fabric. However, it is important to note that in addition to the LUN masking implemented on the NetApp storage system in the previous section, it is important to create the proper soft zoning in your SAN fabric based on the security and availability requirements of your enterprise SAN. For this solution validation, VSAN 200 was created on the MDS 9509 Fabric switch to connect the UCS server blade initiators to the NetApp storage target.

Figure 20 shows that Fabric Manager was used to configure the soft zoning and VSAN on the MDS fabric switch—the WWPNs for the vHBAs of the service profiles and the target WWPN on the NetApp storage system are configured as zone members in the same zoneset for the VSAN 200.

For further details on configuring the MDS 9500 series fabric switches with Fabric Manager, refer to "Cisco MDS 9000 Family Fabric Manager Configuration Guide, Release 4.x" at http://www.cisco.com/en/US/docs/switches/datacenter/mds9000/sw/4_1/configuration/guides/fm_4_1/fmguide.html.

Figure 20 Using Fabric Manager to Configure Soft Zoning and VSAN

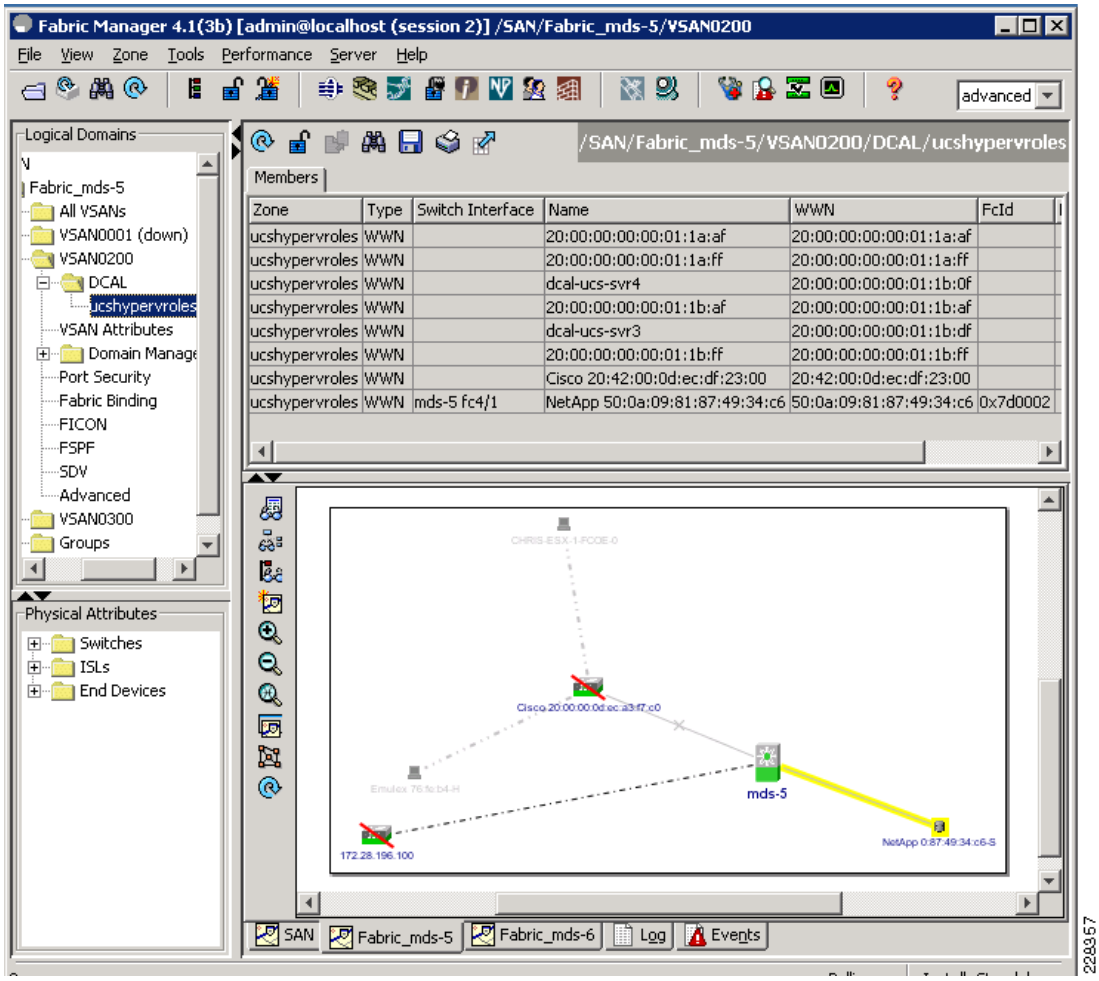

Setting Up Windows Hyper-V Servers on UCS Server Blades

The next phase of deployment is setting up the Windows Hyper-V servers to run on your UCS blades. The Windows 2008 operating system with the Hyper-V role will be installed on the NetApp LUNs set up previously. Having the UCS blades set up to SAN boot with the NetApp LUNs allows for the Windows Hyper-V servers to be moved to any other blade in the case where a given blade experiences a hardware failure. The installation is done using the KVM Console and KVM Console "Launch Virtual Media". As mentioned previously, be sure there is only a single path to the boot LUN through the SAN fabric before proceeding with the installation.

Useful Notes on Windows SAN Boot Installation

To install Windows 2008 for SAN Boot on a UCS blade, refer to the instructions provided by Microsoft in "Windows 2008 - Windows Boot from Fiber-Channel SAN- An Executive Overview and Detailed Technical Instructions for the System Administrator" at http://download.microsoft.com/download/f/9/7/f9775acc-baa6-45cc-9dec-b82983705620/Boot%20fro

m%20SAN%20in%20Windows.doc. This Microsoft document provides instructions on using the BIOS utility for Emulex HBAs to ensure the HBA BIOS is enabled and to verify that the appropriate LUN is being targeted for the installation.

Where QLogic HBAs are used, if you need to troubleshoot your SAN connectivity, the following notes may be useful:

- Make sure your service profile for your blade has the following boot order: Virtual CD/DVD first and SAN second. Make sure the primary LUN target WWPN shown in the boot policy matches the target WWPN for your SAN storage. (This would apply even if you are using Emulex HBAs.)
- Qlogic FastUtil! BIOS utility can be accessed by hitting "Ctrl-Q" while the blade is booting. Quiet Boot may need to be disabled in the BIOS so that the HBA portion of the boot process is displayed and "Ctrl-Q" can be hit at the right point in time.
- In the FastUtil! utility, verify that your desired LUN is listed as LUN 0 under the "Scan Fibre Channel Loop" window.
- Be sure that the ROM BIOS is installed and enabled for the HBA that your blade is using to reach the SAN. If the BIOS is disabled for your given HBA, the boot process reports that the ROM BIOS is disabled for that adapter. On any HBAs for which the BIOS is enabled, the boot process reports that it is installed (Figure 21).

Figure 21 ROM BIOS Enabled for FC SAN Boot

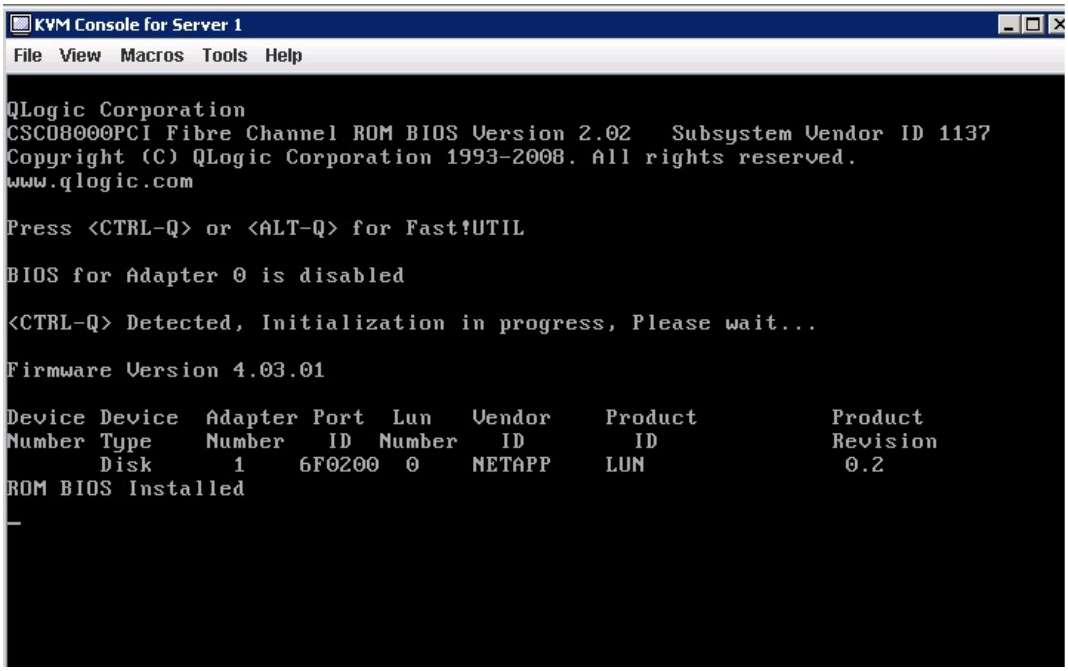

- If the BIOS on your HBA is not installed, you can go to the "Configuration Settings" for the correct HBA and set the BIOS to "enabled" to correct this situation after reboot of your blade.

Installing Windows to Boot from SAN

1. Launch the KVM console for your blade server from UCS Manager. Use "Launch Virtual Media" available through the KVM console to map your Windows installation ISO as a virtual drive (Figure 22).

Figure 22 KVM Console

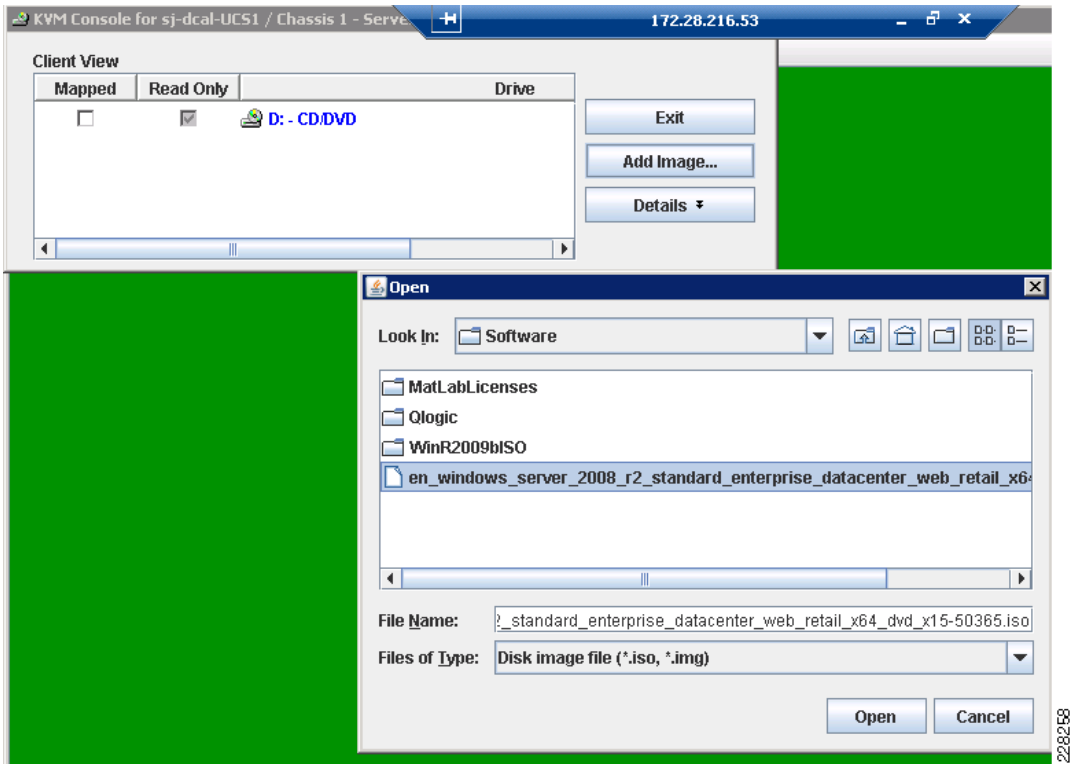

2. The installation wizard brings up a list of drives and prompt you to select the drive on which to install Windows (Figure 23). Note that the drives listed are all the internal drives on the blade server. You need to install the HBA drivers next in order for the blade server to connect to the boot LUN on the SAN.

Figure 23 Internal Disks on Blade Server Before HBA Drivers Loaded

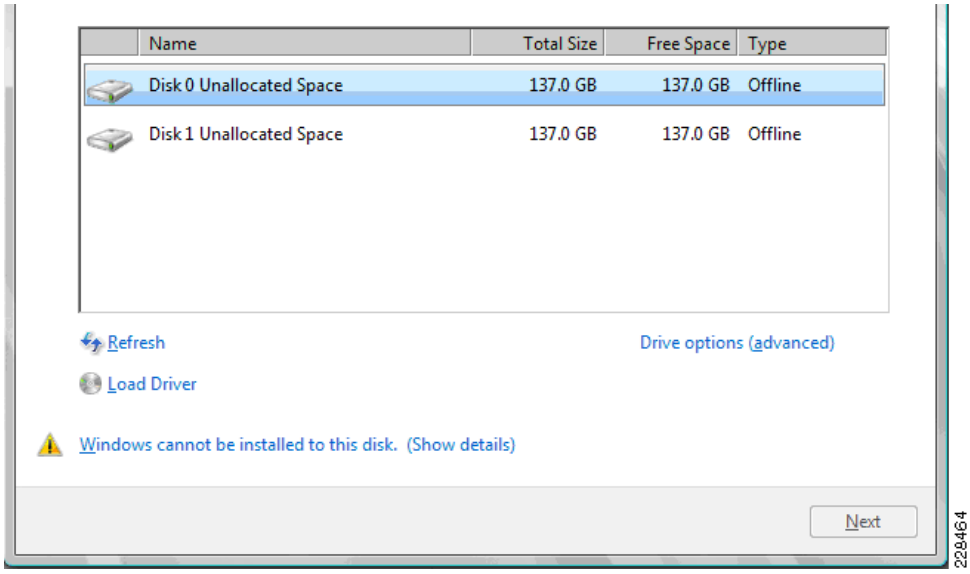

3. Use the "Launch Virtual Media" interface to unmap the Windows ISO and to map the ISO with your HBA drivers. Then select "Load Driver" in the Windows menu to install the HBA drivers (Figure 24).

Figure 24 Installing HBA Drives

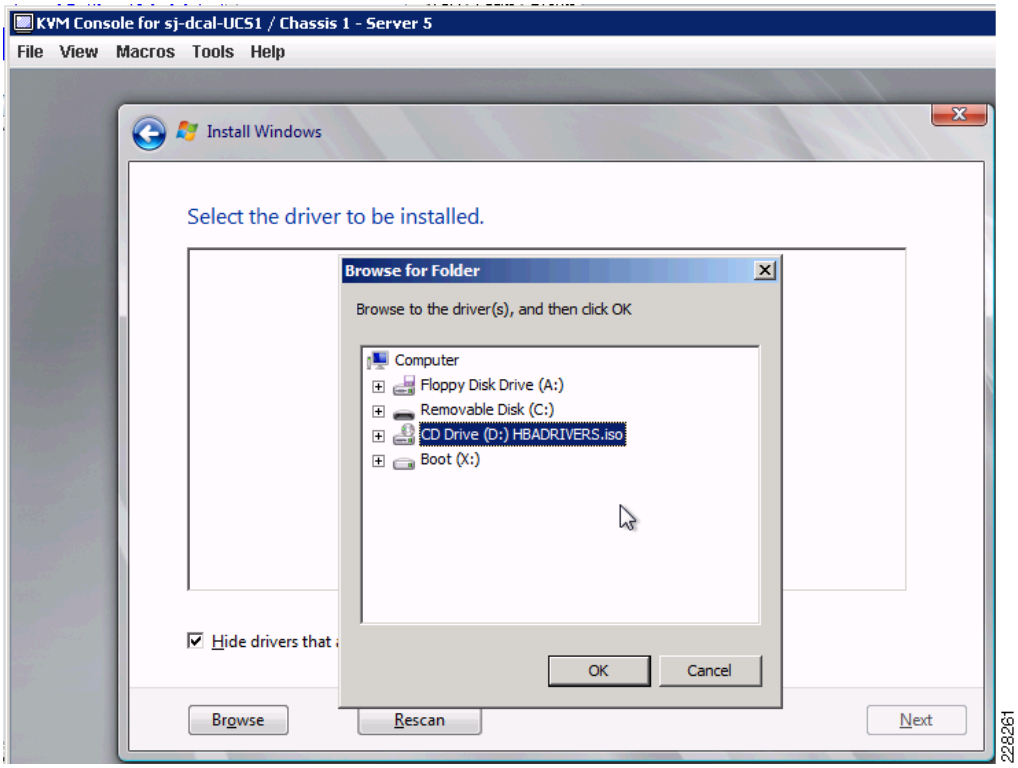

4. After the drivers are installed, the new list of drives includes the boot LUN (Figure 25). Select the LUN and continue with the installation of Window.

Figure 25 List of Drives Including Boot LUN

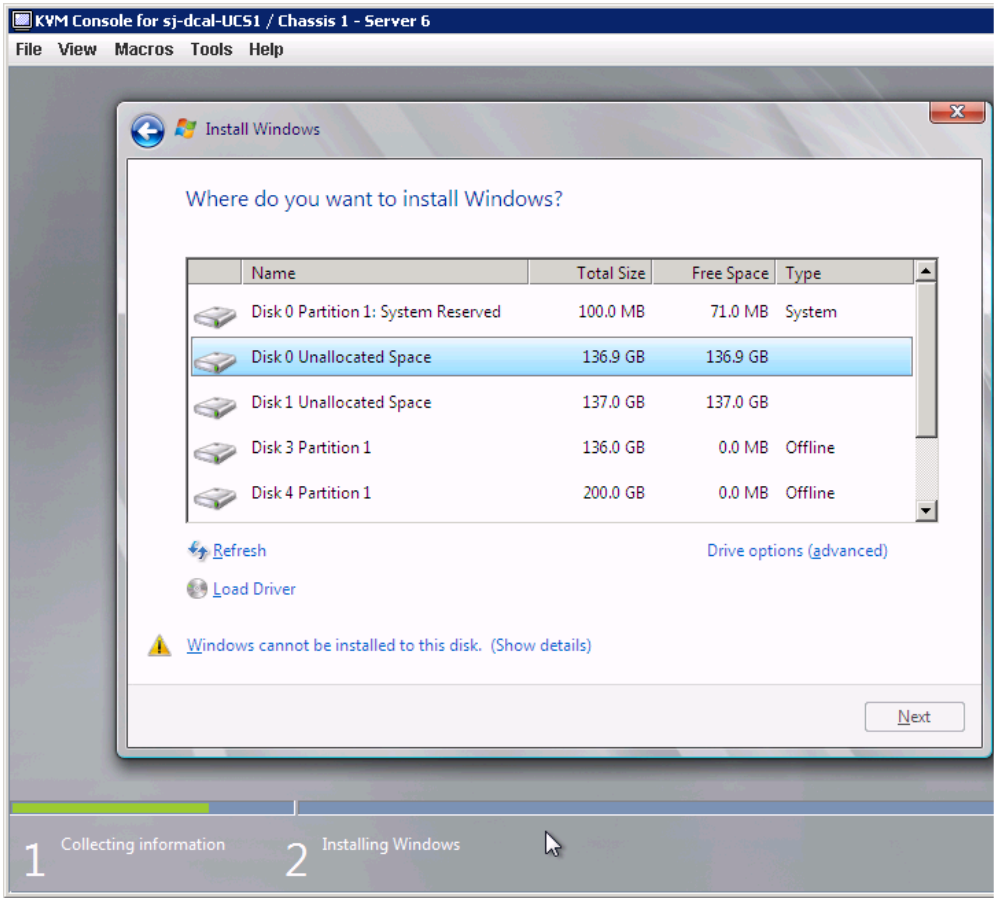

Note If you have internal drives on your UCS blade server, you may run into the case when the Master Boot Record (MBR) gets written to an internal drive even though you selected the proper LUN on SAN storage for your Windows installation. In those cases, it is better to remove the internal drive during the SAN boot installation.

Demonstration of Server Portability Through Service Profiles

The following example demonstrates how a server installation can be moved from one UCS blade server (Server 1/1) to a second blade server (Server 1/6) by moving the service profile association from the first blade to the second. VLAN and VSAN memberships stay the same, as do the server parameters IP address, UUID, MAC addresses, and WWPNs. LUN zoning and masking for server access to SAN storage remain the same as well.

Server 1/1 is assigned the service profile “NICTeam”. The UUID that has been assigned to this server through the service profile, as shown in Figure 26, can be verified as the same as the system UUID displayed through Windows PowerShell on the server (see Figure 27).

Figure 26 UUID Assignment in Service Profile

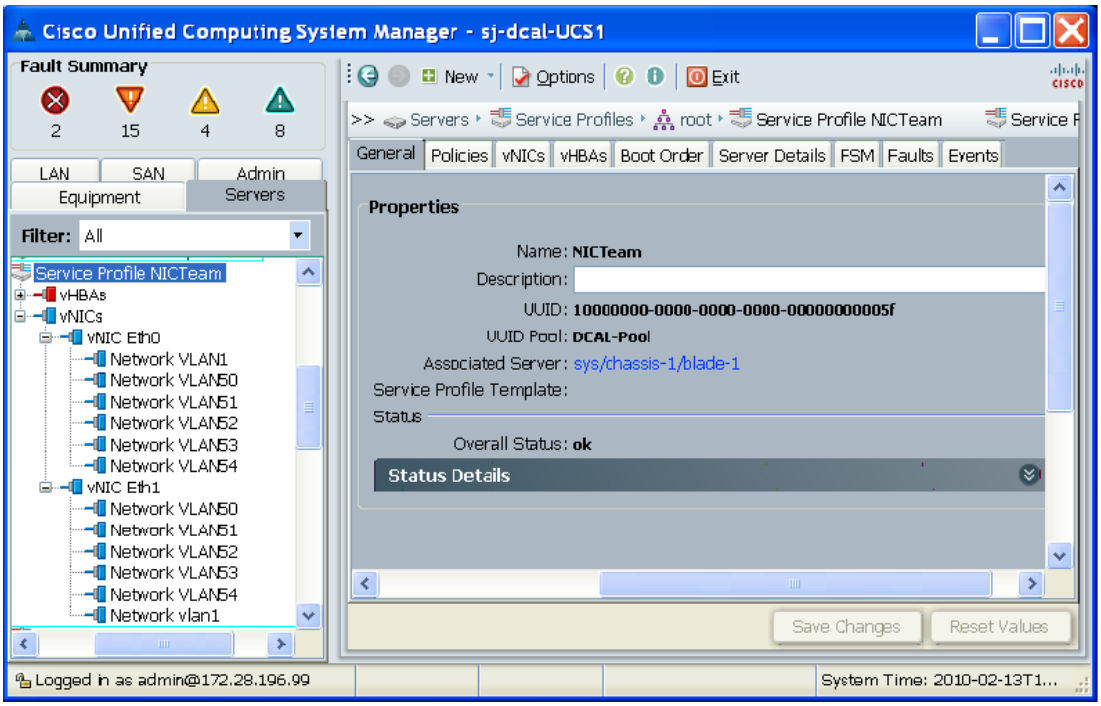

Figure 27 UUID on Server 1/1 Shown on Windows PowerShell

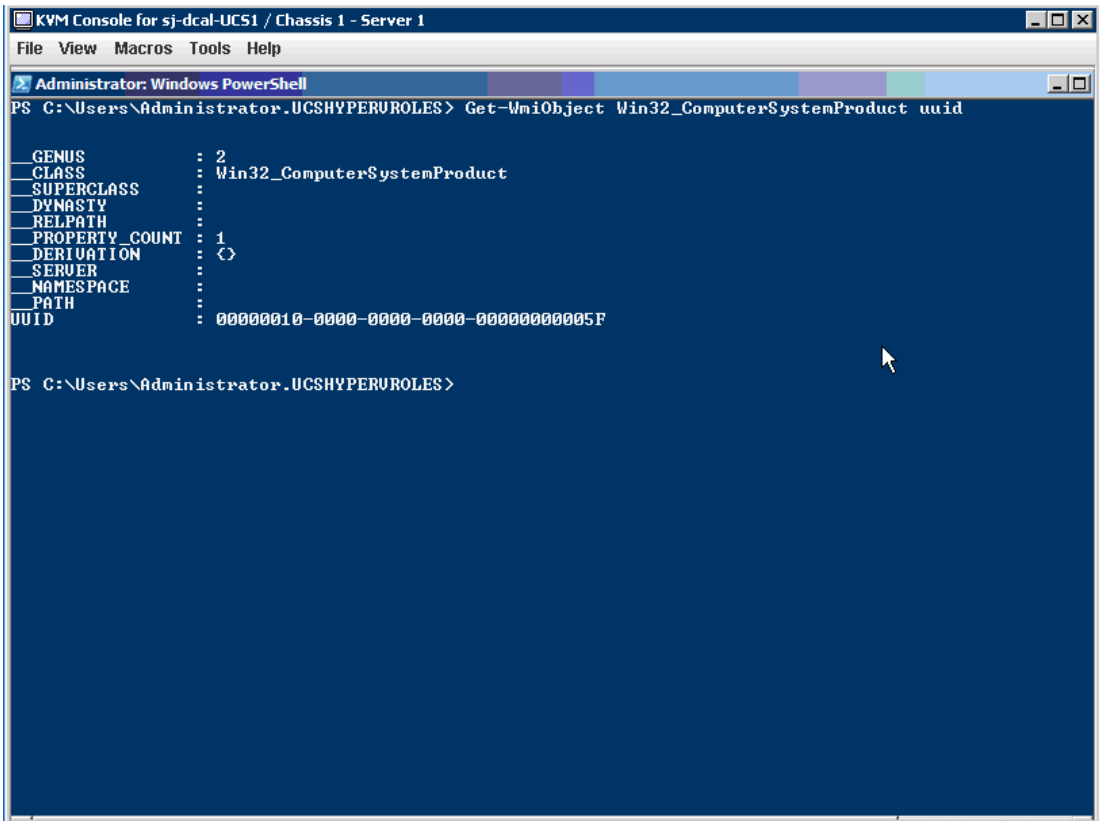

The MAC addresses assigned to the Ethernet vNICs through the service profile, as shown in Figure Figure 28, can be verified as being the same MAC addresses for the LAN connections displayed by the **ipconfig /all** DOS command in Figure 29.

Figure 28 MAC Addresses Assigned in Service Profile

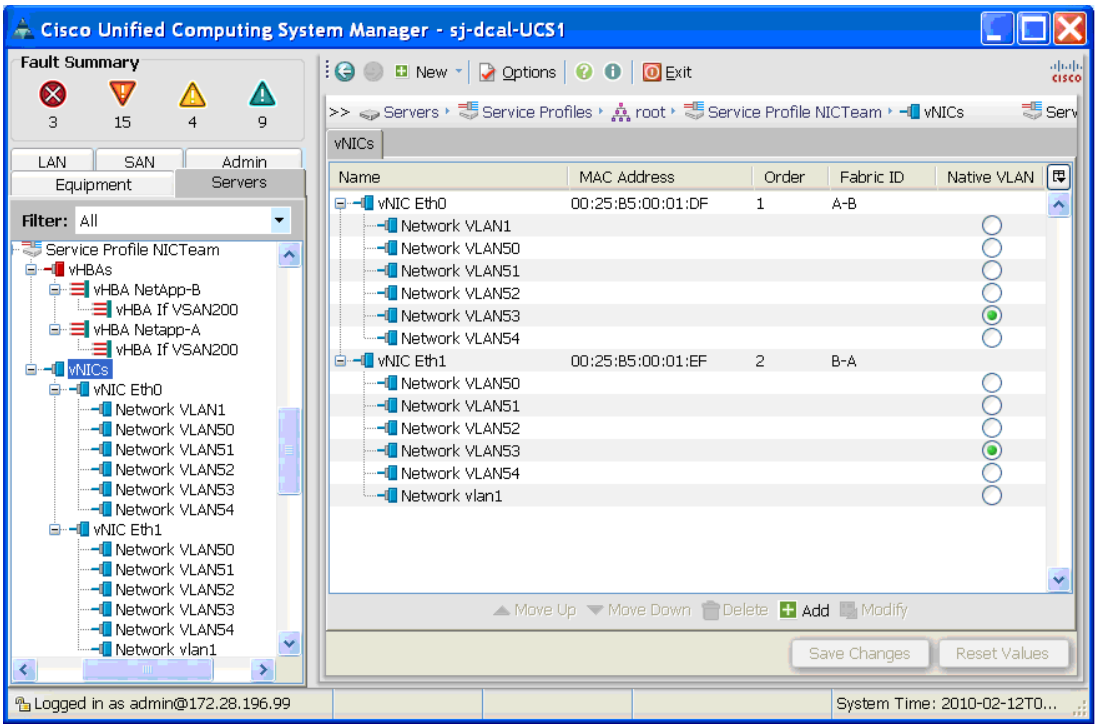

Figure 29 MAC Addresses Shown on Server 1/1 by ipconfig /all

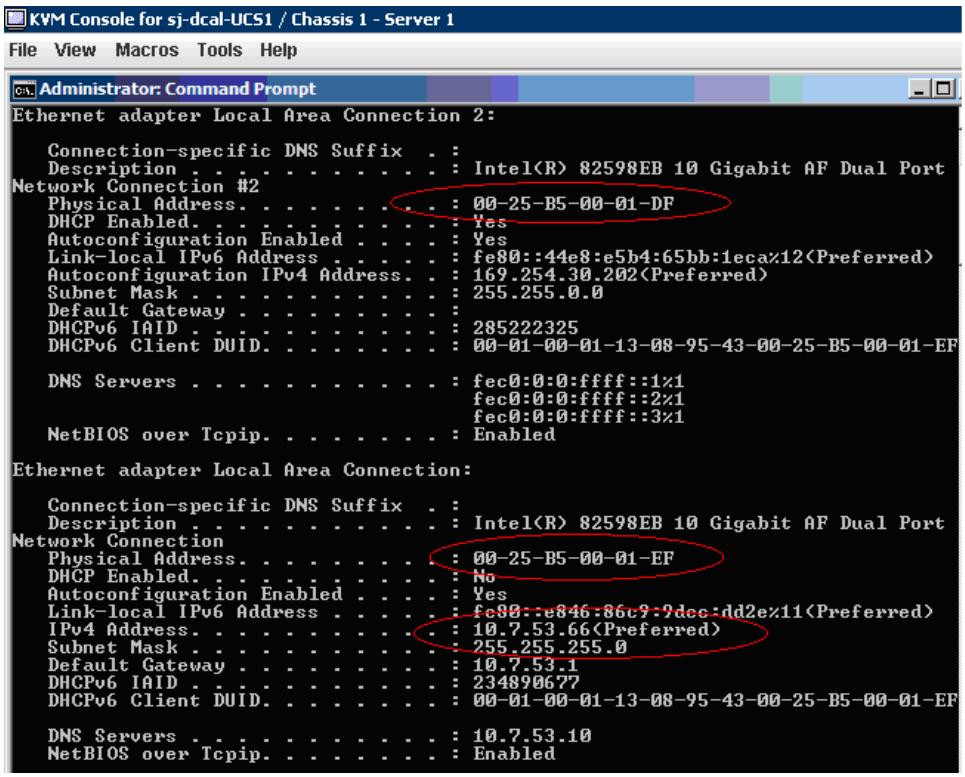

The WWPNs and VSAN assignment set in the service profile are shown in Figure 30. Figure 31 shows that same WWPN assignment on Server 1/1 through Windows 2008 Storage Explorer.

Figure 30 WWPN and VSAN Assignment in Service Profile

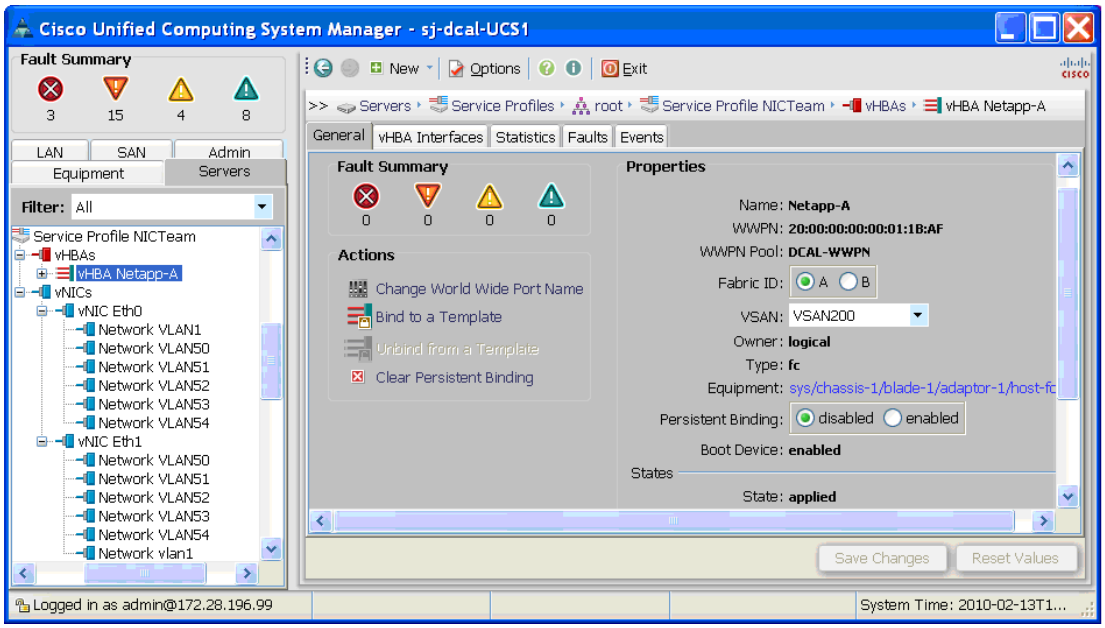

Figure 31 WWPN of Server 1/1 HBA Through Windows Storage Explorer

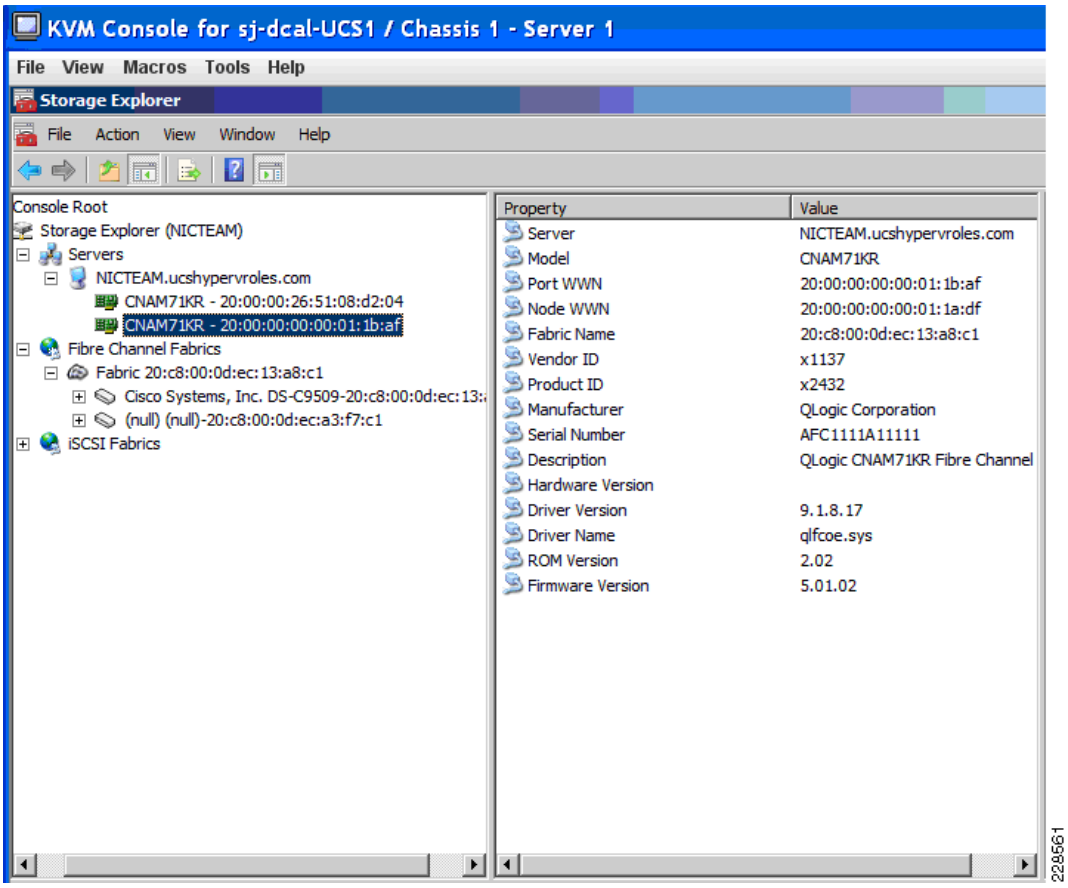

Through FC SAN connectivity provided by the HBA, two LUNs on the NetApp storage show up in Windows Server Device Manager on Server 1/1—one as the boot LUN 0 and the second LUN 1 for data storage. These LUNs similarly automatically show up on Server 1/6 after the service profile is associated with this second server (see Figure 32).

Figure 32 LUNs for Server 1/1 and for Server 1/6 After Service Profile Move

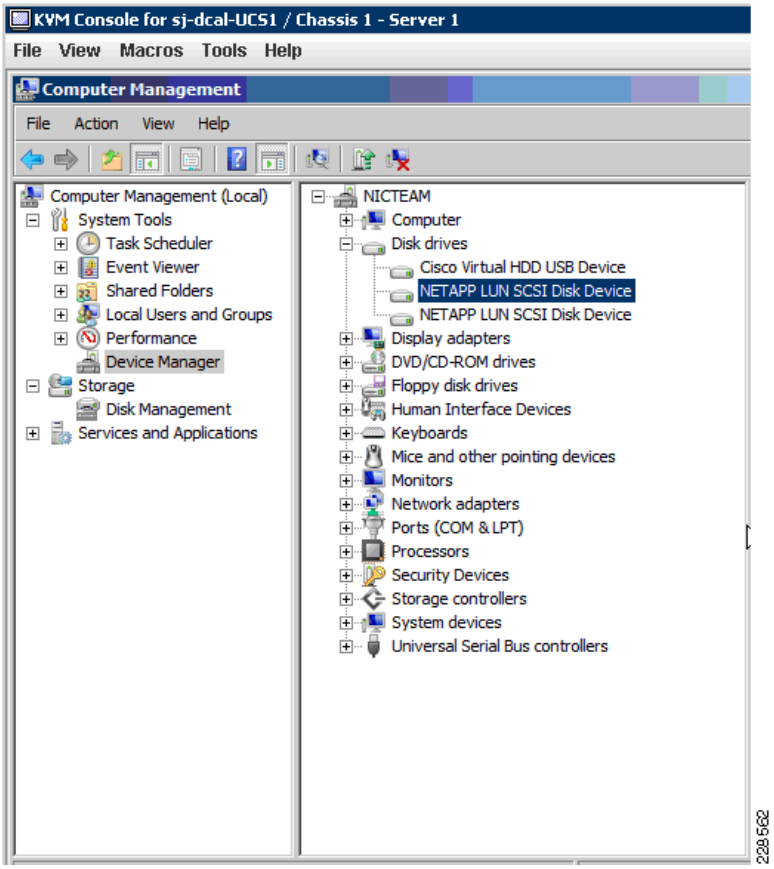

Before and after the service profile is moved from Server 1/1 to Server 1/6, the SAN fabric switch always shows that the same WWPN assigned to each server is active on the SAN fabric and is a member of VSAN 200 (see Figure 33).

Figure 33 MDS Fabric Switch Shows WWPN as Active Zone Member of VSAN 200

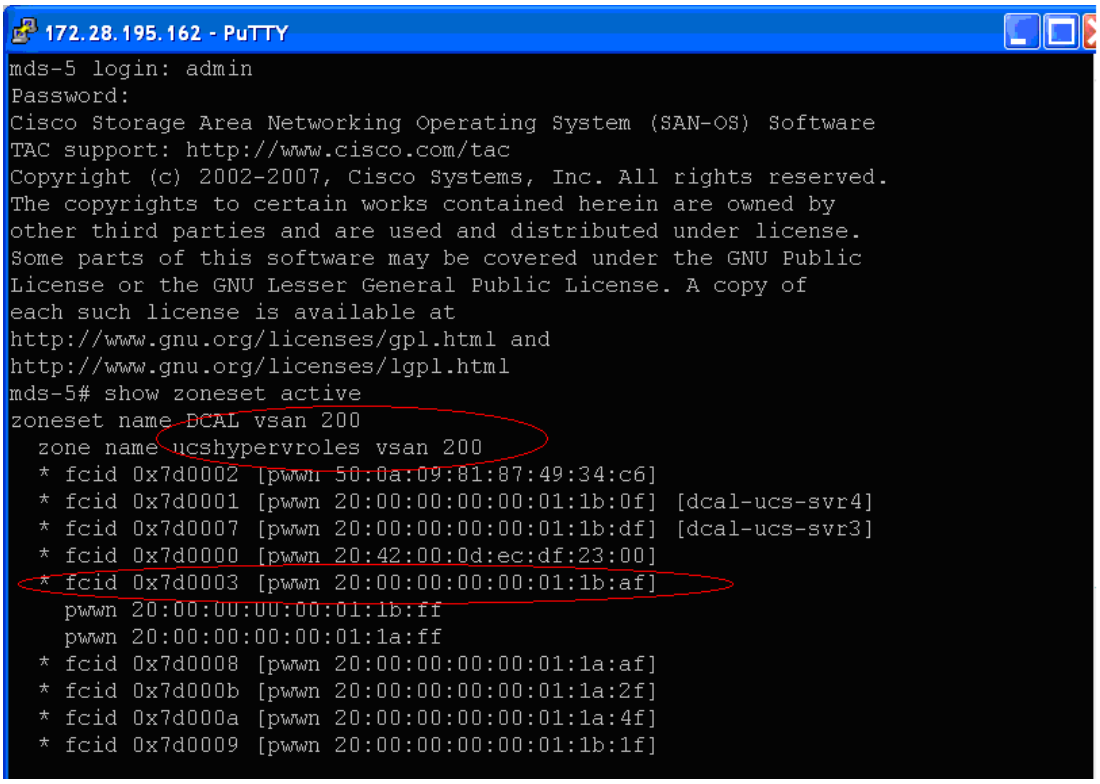

Before moving the service profile from Server 1/1 to Server 1/6, we can examine Server 1/6. A review of the Server 1/6 hardware parameters by selecting "Equipment"—>"Server 6" on UCS Manager and browsing through the Server 1/6 configuration. As Figure 34, Figure 35, and Figure 36 show, Server 1/6 has its own settings for the following parameters before the service profile is associated to it: UUID, Ethernet NIC MAC addresses, and HBA WWPNs.

Figure 34 Derived or Burned-in UUID on Server 1/6

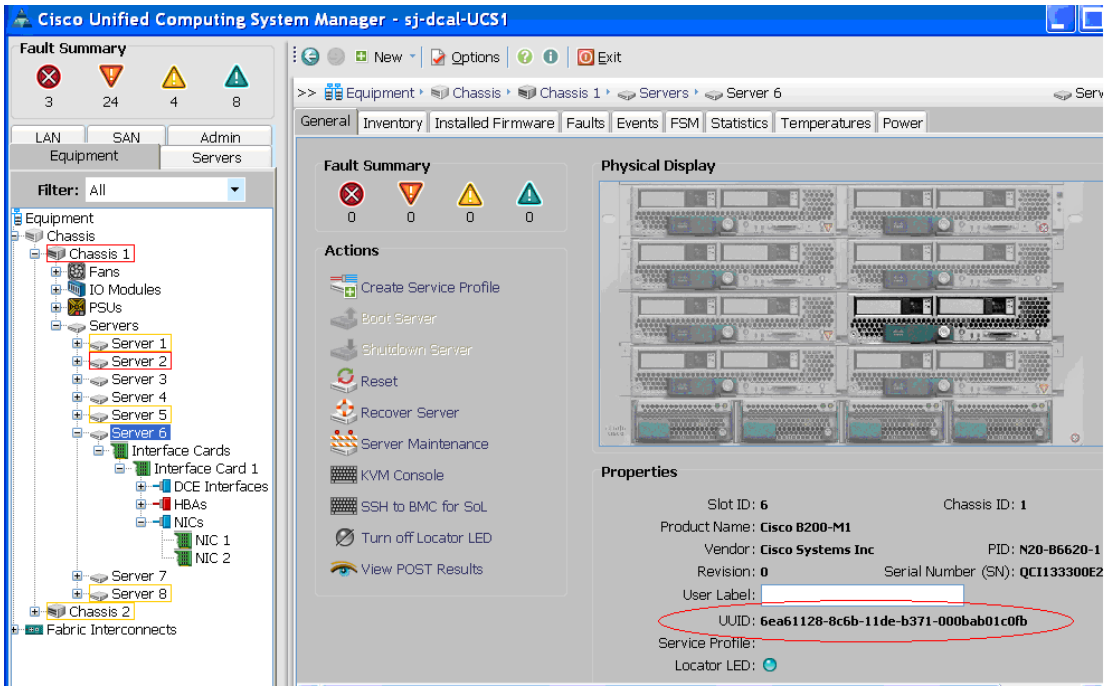

Figure 35 Derived or Burned-in MAC Addresses on NICs

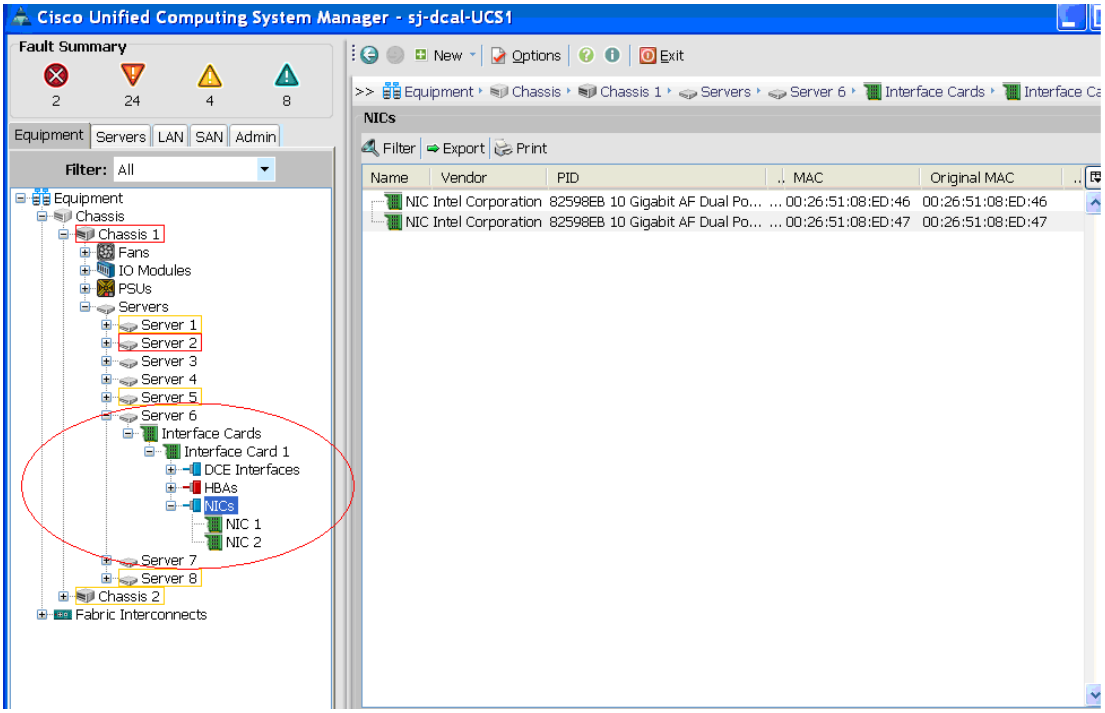

Figure 36 Derived or Burned-in WWPNs on HBAs

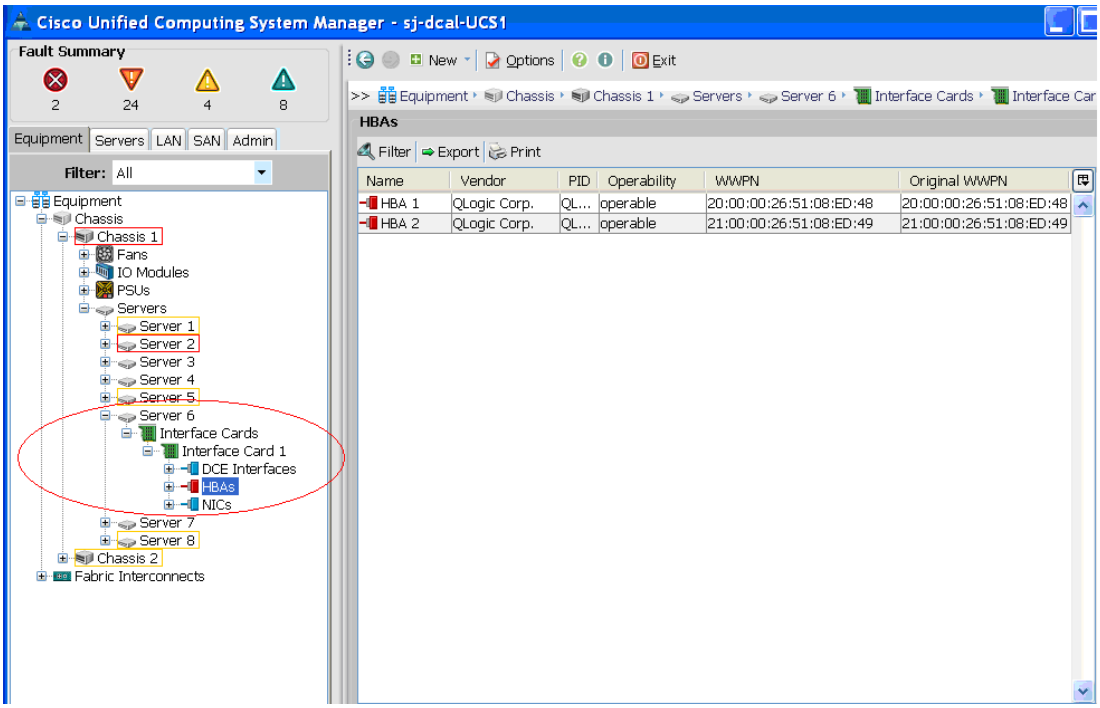

Now, we can move the service profile using the "Change Service Profile Association" option shown in Figure 37.

Figure 37 Change Service Profile Association

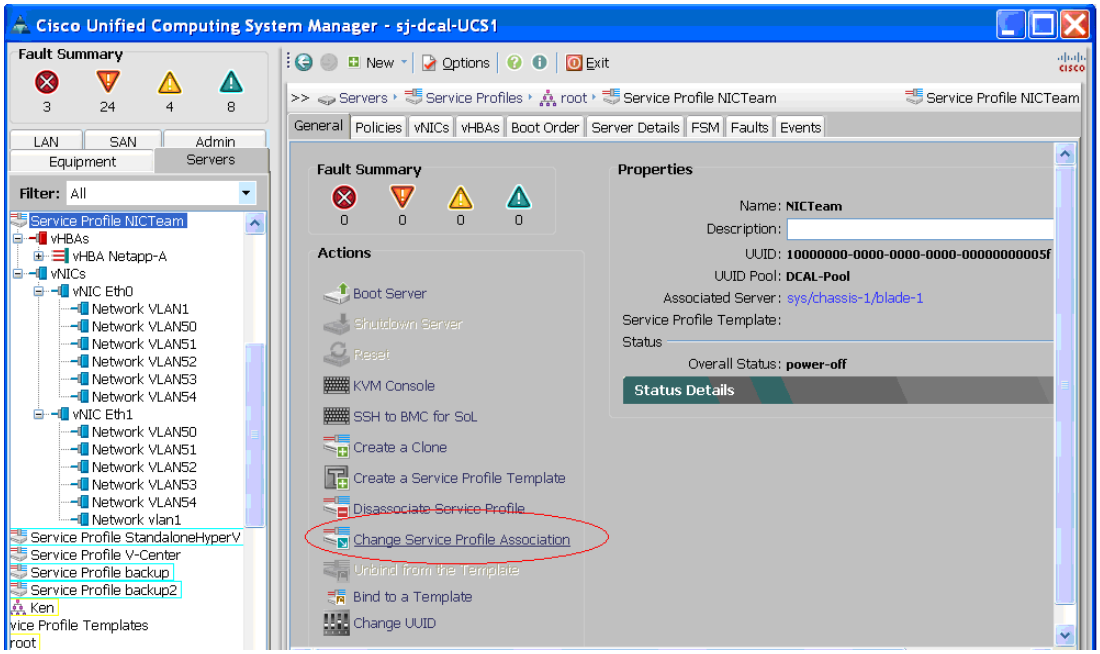

Select Server 1/6 as shown in Figure 38.

Figure 38 Select Server for Service Profile Association

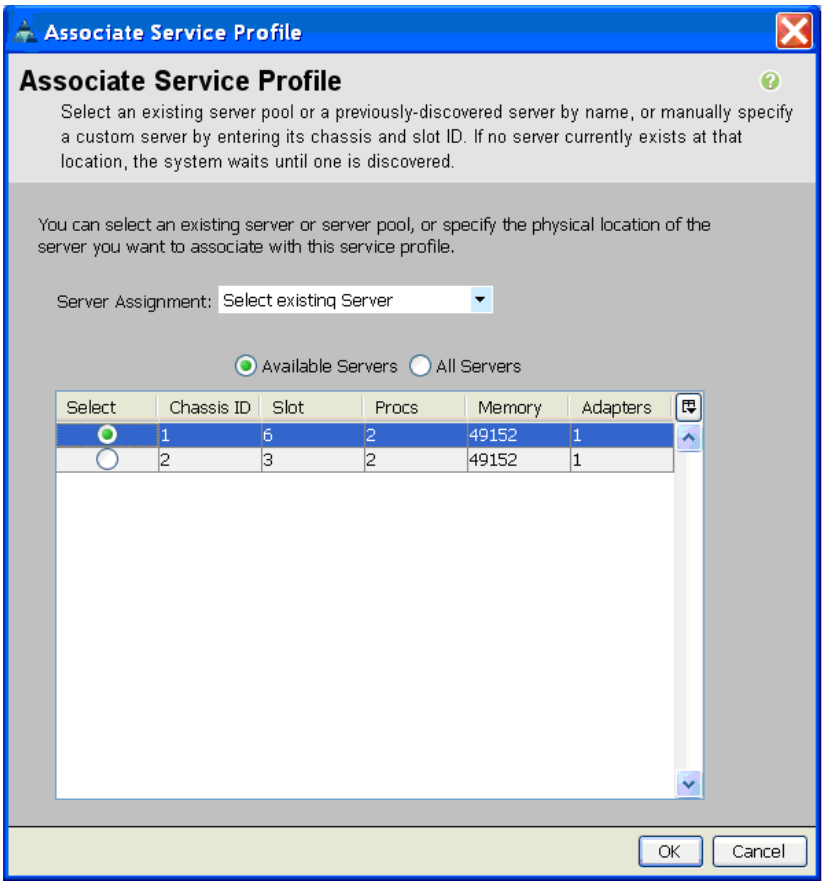

After the association of the service profile to Server 1/6 is done, no other configuration changes are required to bring up the Server 1/6 with the same VLAN and VSAN assignment and network connectivity as Server 1/1 had. As with Server 1/1, the same MAC addresses, WWPN, and UUID can be checked on Server 1/6 in Windows to verify they are the same ones defined in the service profile. As with Server 1/1, the same two NetApp LUNs can be accessed by Server 1/6, since the LUN masking and zone membership are the same as before. See Figure 39 through Figure 44 with screenshots for Server 1/6, similar to what was taken for Server 1/1 above.

Figure 39 Same UUID of Service Profile Now Associated with Server 1/6

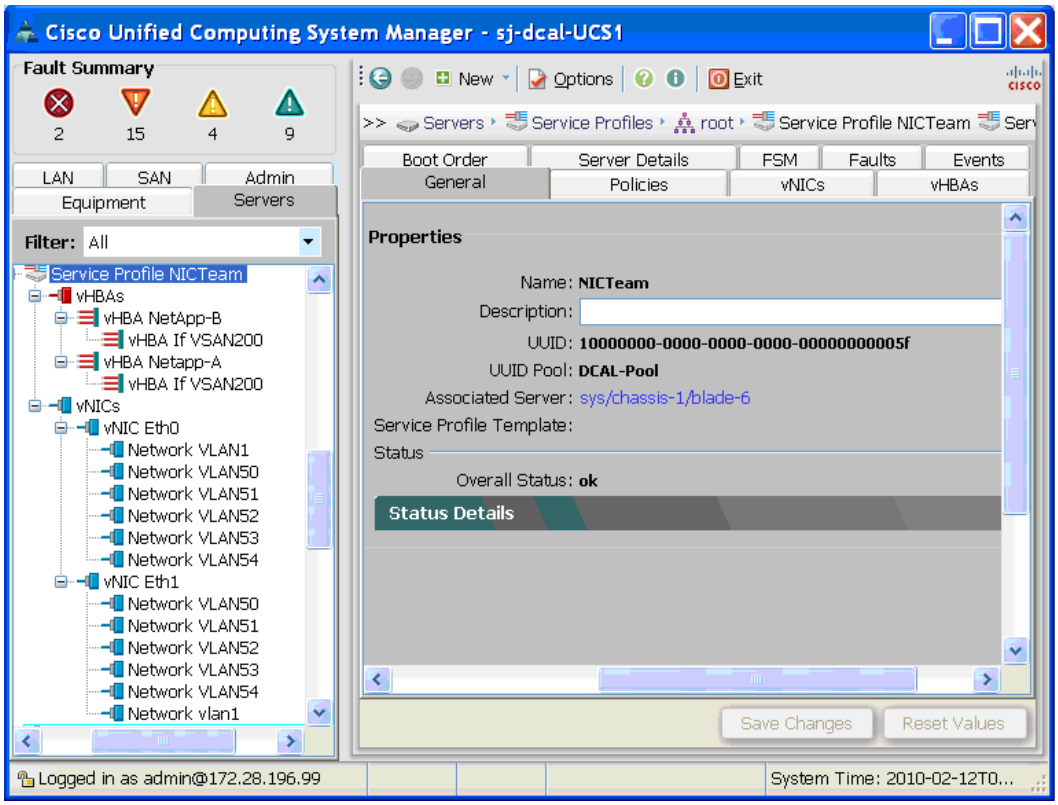

Figure 40 Windows PowerShell on Server 1/6 Reports Same UUID as in Service Profile

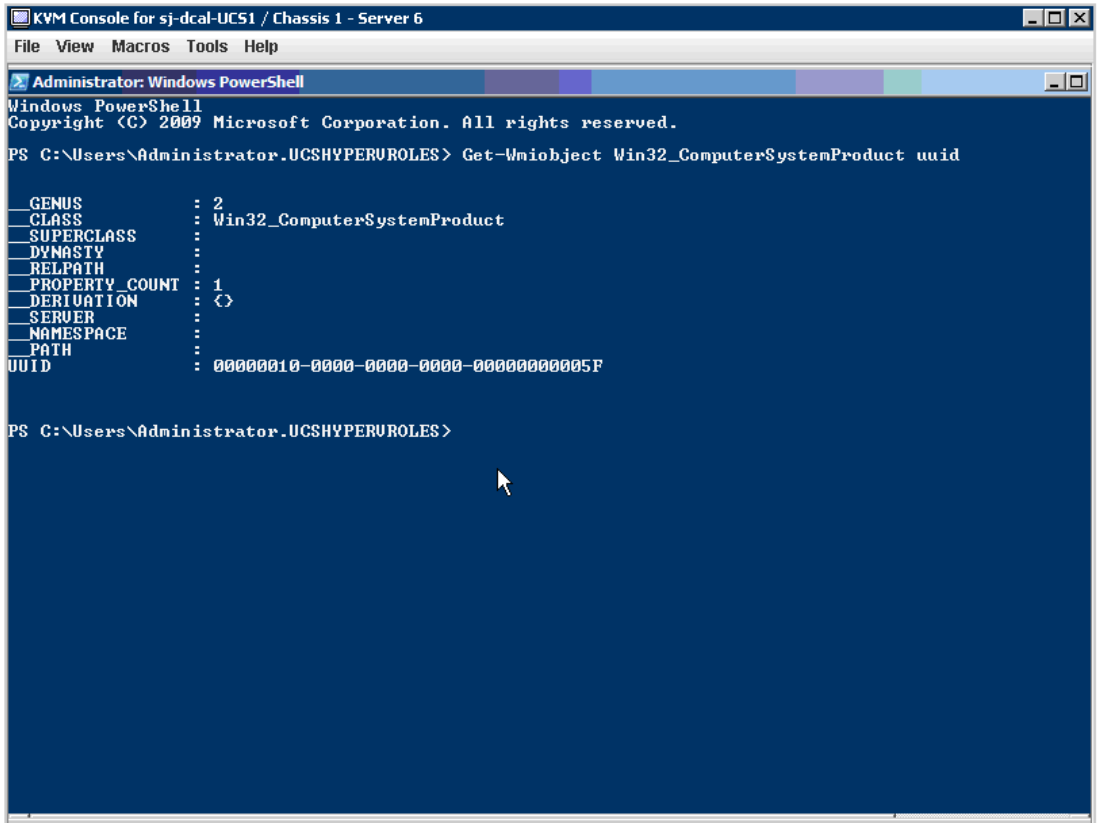

Figure 41 Same MAC Addresses in Service Profile Shown on Server 1/6

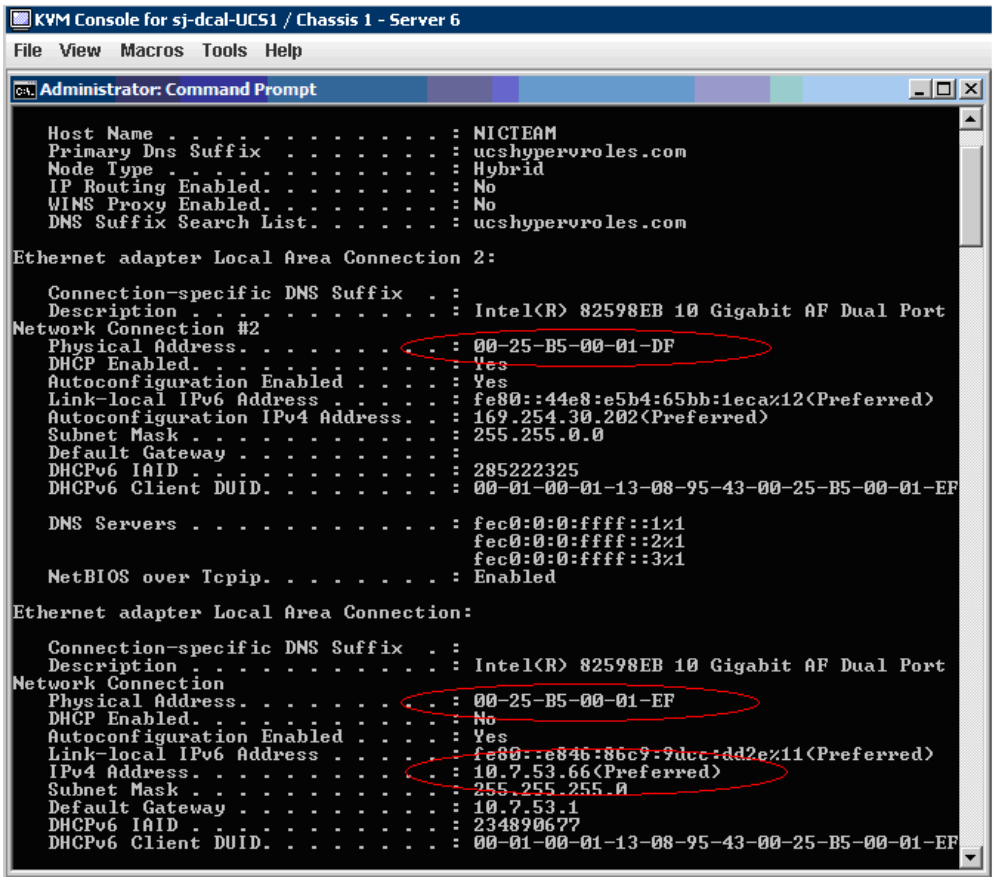

Figure 42 VSAN 200 and WWPN of Service Profile Assigned to Server 1/6

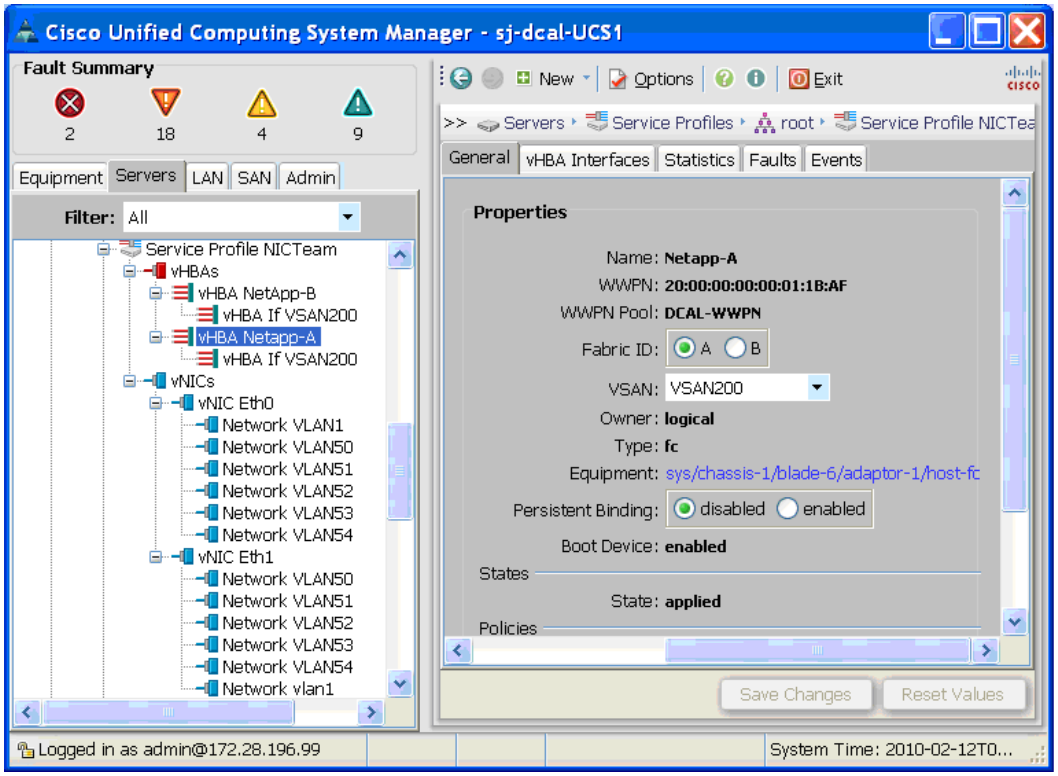

Figure 43 Storage Explorer on Server 1/6 Shows Same WWPN as in Profile

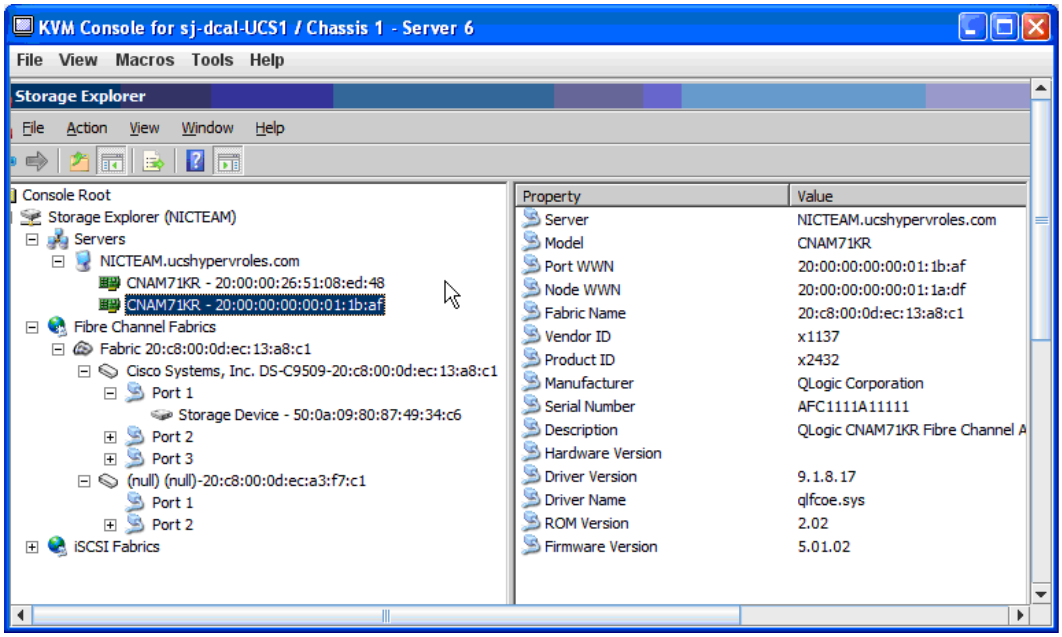

Figure 44 Same Two NetApp LUNs on Server 1/1 Now Appear on Server 1/6

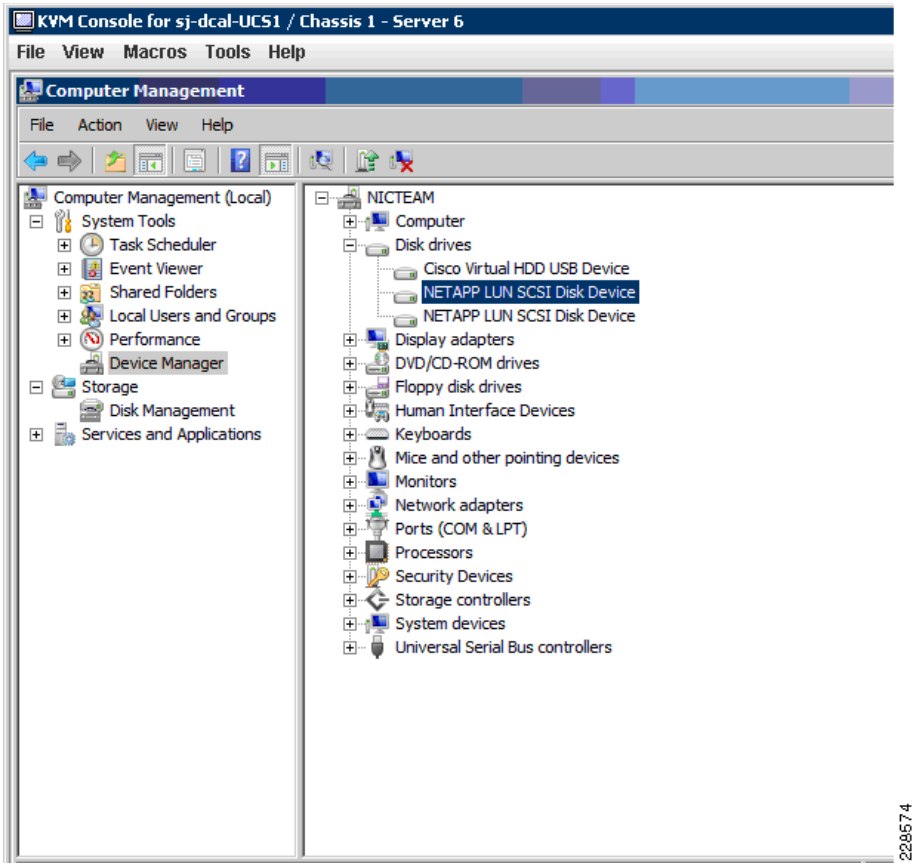

Table 4 shows which VLANs must be assigned to the blade servers used in this solution, which are shown in Figure 45.

Figure 45 Cisco UCS Blade Servers and VMs

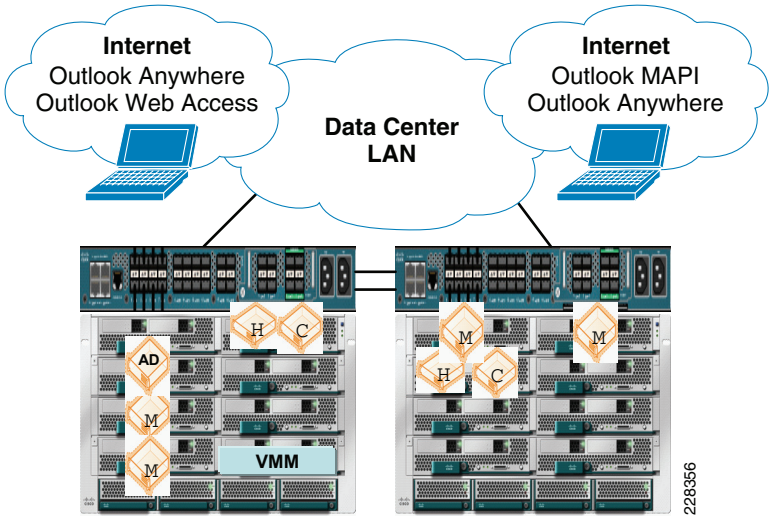

Table 4 Blade Server Assignment to VLANs

Blade	Live Migration	Virtual machine management	Cluster interconnect	Virtual machine access	Blade server management
With AD VM	X	X	X	X	X
With Mailbox VM		X		X	X
With CAS and HT VMs	X	X	X	X	X

Figure 46 Exchange Server Virtual Machines Logically Connected to Data Center

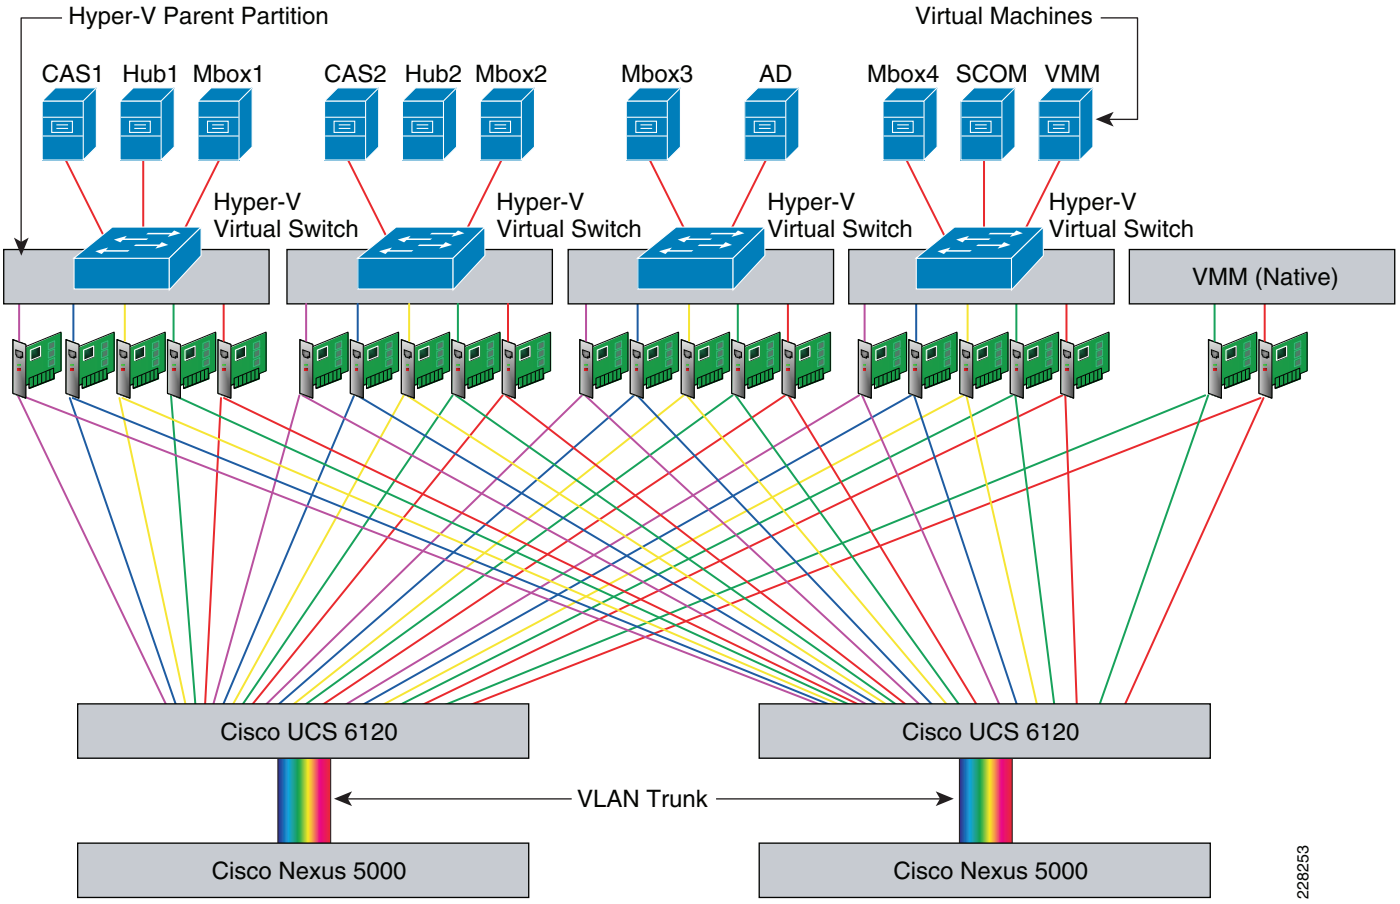

Figure 46 illustrates how the Exchange server virtual machines are logically connected to the rest of the data center. The UCS 6120 Fabric Interconnects trunk all the VLANs up to the Nexus 5000 access switch to the data center. Note that the number of server blades shown and the placement of the Exchange server roles is for illustrating the virtual networking; Figure 46 does not reflect the actual deployment of Exchange virtual machines on server blades in the solution.

Windows Failover Clustering and Live Migration Configuration

In order for Live Migration of the Exchange server virtual machines from one Hyper-V server to another, the following must be configured:

- All of the Hyper-V servers must be on the VLAN for Windows Failover Cluster traffic and on the VLAN for Live Migration traffic.
- Each Hyper-V server must have the Cluster Shared Volumes LUNs connected online and mapped as drives.
- Windows Failover Clustering must be enabled on all Hyper-V servers that participate in Live Migration. Windows Failover Clustering can be enabled through Windows Server Manager—>Add Features. The Failover Cluster Manager software that is installed at that point can be used on any Hyper-V server to create the cluster and add the other Hyper-V servers. This section does not go into the details of setting up the failover cluster for Live Migration. Microsoft has detailed documentation on how to set up Windows Failover Clustering with Cluster Shared Volumes—"Hyper-V: Using Live Migration with Cluster Shared Volumes in Windows Server 2008 R2" at http://technet.microsoft.com/en-us/library/dd446679(WS.10).aspx.

Figure 47 shows the creation of a cluster in process with the current two nodes in the cluster validated and sharing the two Cluster Shared Volumes to store the VHD files of any virtual machines created later.

Figure 47 Creation of a Cluster

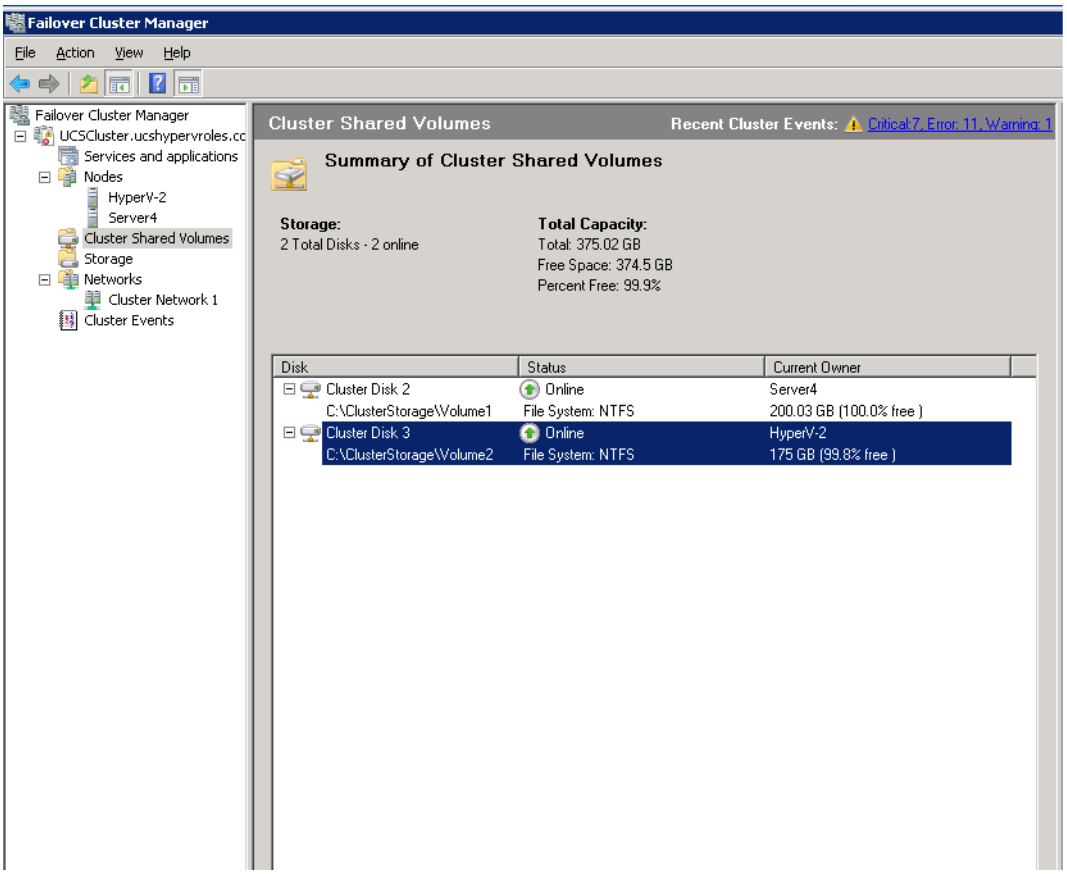

Installation of Exchange 2010 Server Roles on Hyper-V Virtual Machines

This document does not cover the installation and configuration steps to deploy each of the Exchange server roles on Windows 2008 R2. There is documentation available on the Microsoft site to guide you through the steps of installing new Exchange 2010 servers: http://technet.microsoft.com/en-us/library/aa998619.aspx.

Reminder: Validation of this solution involves setting up a total of eight virtual machines for the Exchange server roles and one virtual machine for the Active Directory server/domain controller. Each virtual machine hosts a single Exchange server role.

Table 5 enumerates the roles (table columns) that were created as virtual machines on the UCS server blades (table rows) that were used for this solution validation.

Table 5 Roles Created as Virtual Machines on UCS Server Blades

	Client Access	Hub Transport	Mailbox	Active Directory
Blade 1/3				X
Blade 1/4	X	X		
Blade 1/5			X	
Blade 1/7			X	
Blade 2/1			X	
Blade 2/2			X	
Blade 2/3	X	X		

The Client Access server and Hub Transport server virtual machines could be hosted on the same UCS server blade because there were enough virtual processors and memory that could be assigned to each virtual machine per Microsoft recommendations on CPU and memory sizing for these two server roles. (see Design Considerations for Exchange 2010 Virtualization on UCS). The two Client Access servers and Hub Transport servers are hosted on separate physical blades for hardware redundancy. The Client Access servers rely on the Application Control Engine (ACE) load balancer to load balance all client requests between them. Which of the two Hub Transport servers would receive mail submissions depends on selection by Active Directory. For further details, see Microsoft documentation on messaging routing at: http://technet.microsoft.com/en-us/library/aa998825.aspx#Exchange2010.

Note The creation of the virtual machines can be done locally on each Hyper-V Server through Windows Server Manager. However, with a large number of virtual machines to manage, it may easier to use the System Center Operations Manager Virtual Machine Manager R2 (VMM R2) management software. VMM R2 allows you to connect to all of your Hyper-V servers from one centralized platform and create, delete, and Live Migrate virtual machines from server to server. For detailed instructions on how to install and use Hyper-V VMM R2, refer to the Microsoft documentation at: http://technet.microsoft.com/en-us/scvmm/default.aspx.

Client Access Server Configuration

For detailed instructions on setting up a Client Access server, refer to the following online Microsoft documentation: http://technet.microsoft.com/en-us/library/bb124915.aspx. Be sure to select the Custom Installation option when installing in order to install the CAS role only. The default installation of a single CAS server role was used for this solution.

Mailbox Database Availability Group Configuration

For detailed instructions on setting up a mailbox DAG, refer to the following online Microsoft documentation: “Managing Database Availability Groups” at: http://technet.microsoft.com/en-us/library/dd298065.aspx. Be sure to select the Custom Installation option when installing in order to install the Mailbox role only. The default installation of a single Mailbox server role was used for this solution.

Application Control Engine Configuration for Load-Balancing Exchange

ANM Guided Setup

The easiest way to configure the Application Control Engine is by using the Application Network Monitoring software. By the use of ANMs guided setup you can easily import an ACE device, configure the ace hardware, and create and manage Virtual Device contexts using a guided setup. Figure 48 displays the starting screen for ANM. Start by importing your current ACE device by clicking on the Import a Device link; this synchronizes any existing configurations to ANM, such as the management IP and VLAN information.

Figure 48 Starting Screen for ANM

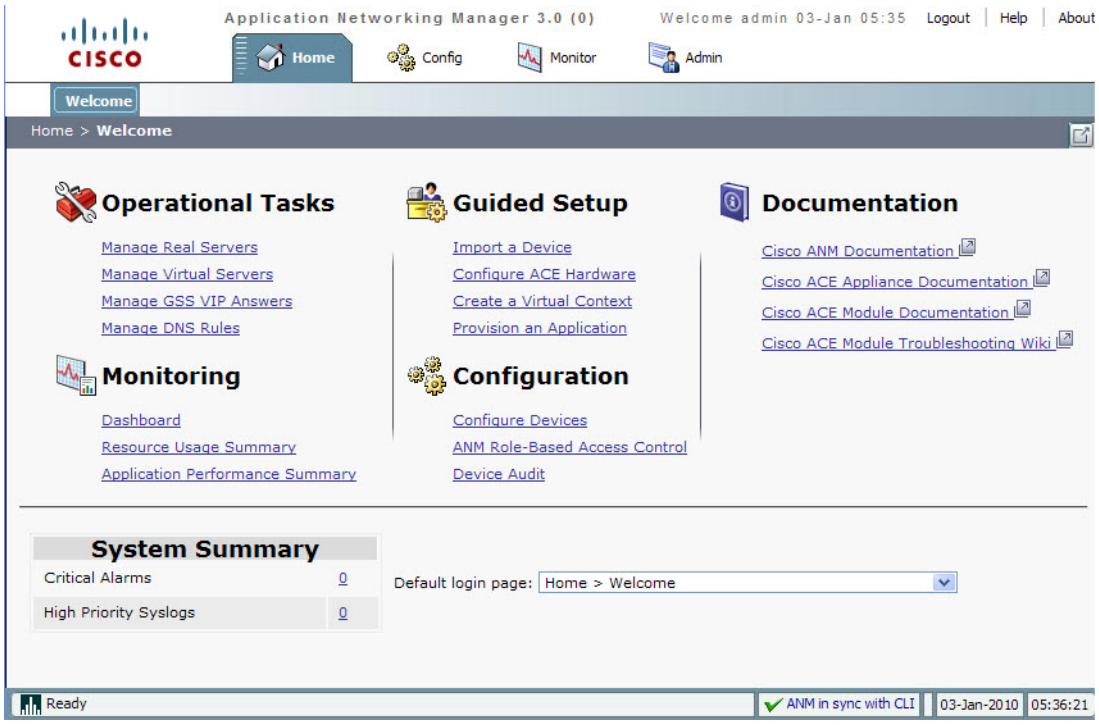

Basic Hardware Configuration

Next you can continue with the hardware setup by clicking the hardware setup link and install any needed licenses. ACE's functionality can be extended by installing licenses. Depending on what tasks you need to do with your ACE you many need separate licensing. To install any licenses you have, click Install... for each one and specify the file's location (Figure 49).

Because each ACE license is only valid on a single hardware device, licenses are not synchronized between HA peer devices. You must install an appropriate version of each license independently on both the primary and secondary ACE devices.

After you are done installing licenses, or if you do not have any licenses to add, continue to the next step.

Figure 49 Installing Licenses

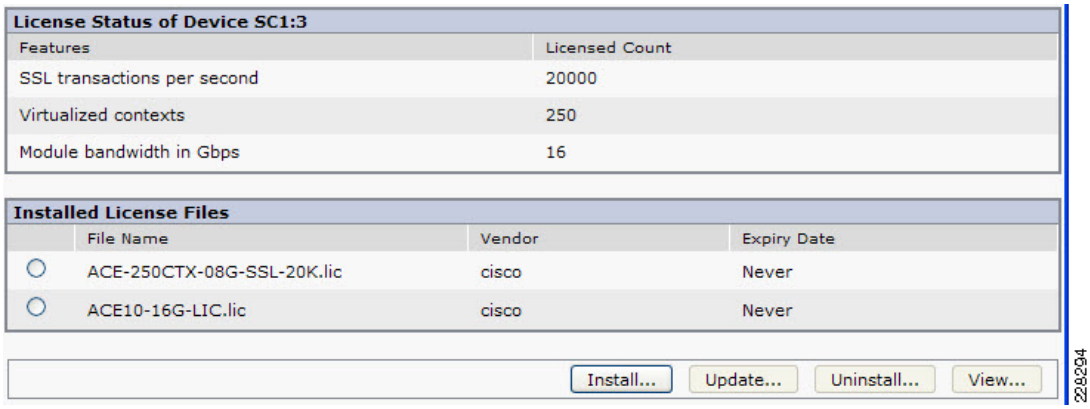

Resource Classes

The capacity of a virtual context is determined by its resource class, which is specified as percentages of available device resources.

Make sure an appropriate resource class is defined for your planned virtual context. The "default" resource class works well for basic configurations that do not need sticky session persistence. For sticky session persistence, you need a resource class that allocates more than 0% of the "Sticky" resources. Note that you can always start with "default" and change to a different resource class later.

Click on the link to add a recourse class and create a new rescore class that we can use with sticky persistence later. In the example in Figure 50, a resource class has been created called gold with the depicted settings. As mentioned the sticky persistence setting is set as a minimum of 10% of the resources available and a maximum of unlimited.

Figure 50 Defining a Resource Class

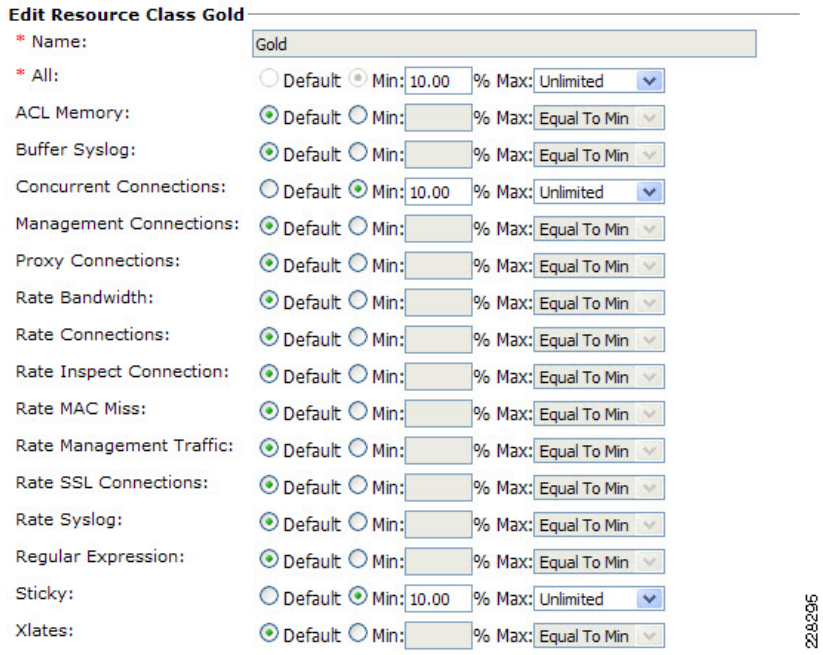

Virtual Context Creation

You are now ready to add an additional context that is used to manage your exchange environment. Click the link to create virtual context, then to create a new virtual context, click the + button.

Configure the context's basic settings such as VLANs, a name, gateway, and resource class. In the example in Figure 51, we are using the gold resource class we created earlier as we are using sticky persistence. So make sure to choose the resource class that was created earlier.

Figure 51 Virtual Context Creation

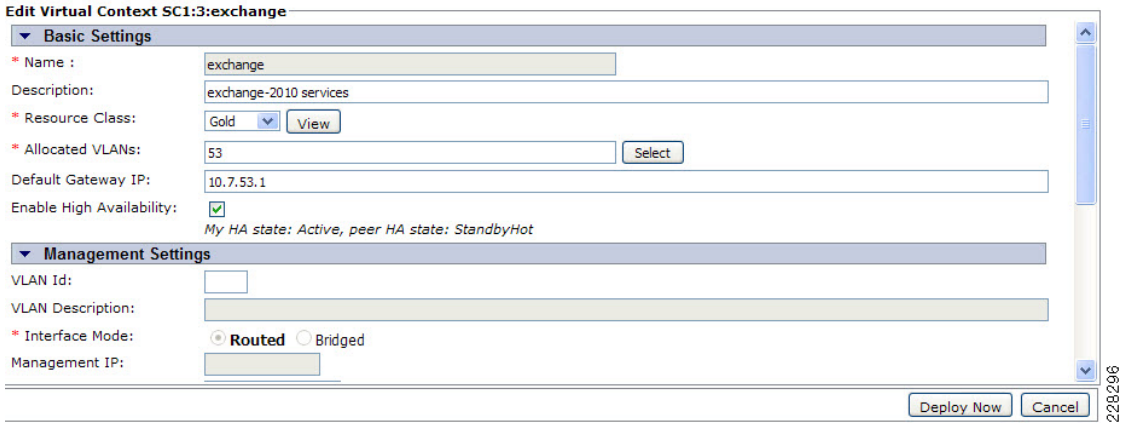

Application Setup

Now that you have added the Virtual context to be used with Exchange, you are ready to begin your application setup. Click the link to navigate your browser to the application setup section. Select the newly created virtual context and for your topology type choose one-armed as shown in Figure 52.

In a one-armed topology, the ACE virtual context is connected to the real servers via an independent router and acts as neither a switch nor a router for the real servers. Clients send requests to the virtual IP (VIP) on the ACE. The ACE then uses network address translation (NAT) to send the requests to the real servers. Server responses return through the ACE rather than directly to the original clients. This topology is convenient, since the ACE can be almost anywhere on the network, but its reliance on NAT makes it impractical in some situations.

Figure 52 Selecting One-Armed Topology

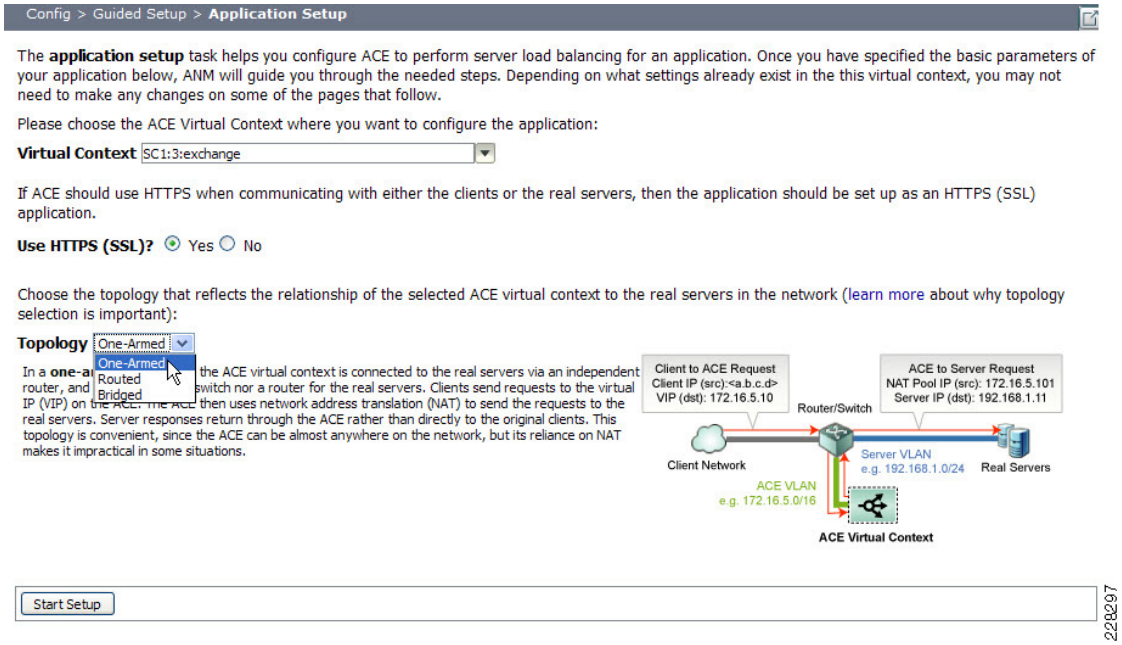

Before you can begin configuring your Virtual servers you must first finalize a few more steps, such as adding any required VLANs, NAT pools, ACLs, or SSL proxies.

Creating VLANs

Click the link to add VLANs to your context. In order to communicate with the client and real servers, a VLAN interface must be specified to carry client and server traffic to and from the router/switch. If that VLAN interface is not already listed below, add it by clicking the + button. Figure 53 shows the required fields for the VLAN settings of the virtual context.

Figure 53 Required Fields for VLAN Settings of Virtual Context

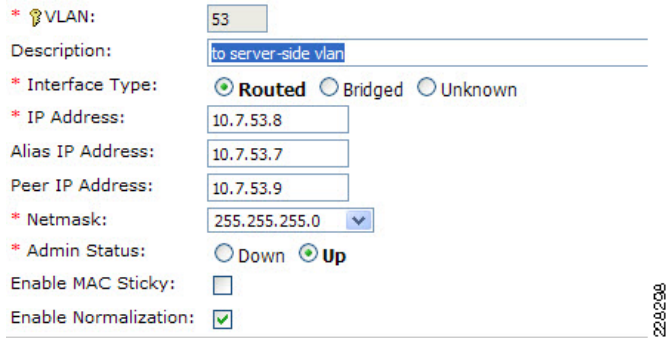

Adding NAT Pools

Next create a NAT pool by clicking on the NAT pools link. In a one-armed configuration, you need a NAT pool as it provides a set of IP addresses that ACE can use as source addresses when sending requests to the real servers. The NAT pool must be configured on the same VLAN interface that you identified or created in the previous step.

Figure 54 shows the NAT pool that is configured for this example. For VLAN ID, specify the VLAN number from the VLAN Interfaces step. Use the NAT Pool ID set in the field or enter your own. Then enter the IP address range for the NAT pool using the Start IP Address and End IP Address or Netmask fields.

The NAT pool can be as small as single IP address and it can even re-use the virtual IP (VIP) you plan to use for the application. However, NAT pools with more IP addresses allow more concurrent requests to the real servers. ACE allows several thousand concurrent connections per NAT pool IP address, as long as port address translation (PAT) is enabled. With PAT disabled, ACE can only handle one connection per NAT pool IP address.

Figure 54 Adding a NAT Pool

Managing ACLs

To continue, click the ACL link. An access control list (ACL) is a packet-level security policy that applies to one or more VLAN interfaces. Each ACL has a list of entries and each entry defines a source, a destination, and whether to permit or deny traffic. ACE denies incoming traffic by default, so ACLs must be explicitly defined to permit traffic to applications.

Specifically, for traffic to reach your application, an Extended ACL must be applied to the client interface in the Input direction. It must have an entry whose:

- Action is Permit
- Source definition includes the expected client IP addresses (or is set to “Any”)
- Destination definition includes the expected ACE virtual server IP addresses (or is set to “Any”)

Use the small + button by the ACL name to show its entries and verify they allow traffic. If needed, add or edit an ACL to permit traffic to your application. In the example in Figure 55, we have one default ACL with the indicated settings.

Figure 55 Managing ACLs

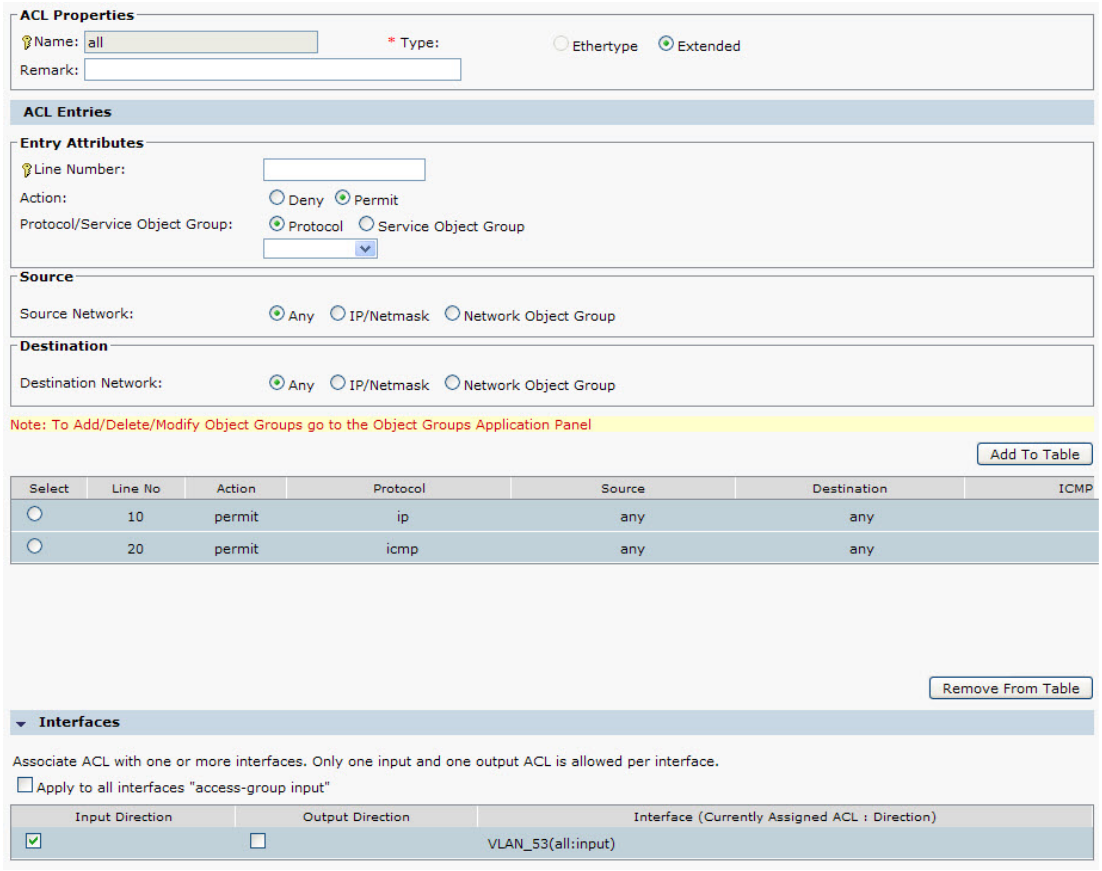

SSL Proxy Configuration

To terminate or initiate HTTPS connections with ACE, the virtual context must have at least one SSL proxy service. An SSL proxy contains the certificate and key information needed to terminate HTTPS connections from the client or initiate them to the servers.

Create a new SSL proxy by clicking the SSL Proxy Setup... button at the bottom of the page. As shown in Figure 56, you may use the sample key and certificate.

Figure 56 SSL Proxy Configuration

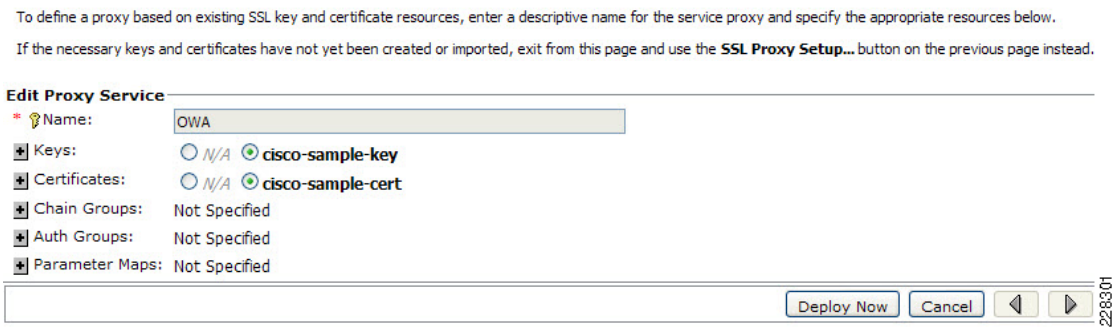

Virtual Servers

Virtual Servers are used to define the load-balancing configuration for a specific application. To add a new load-balanced application, click the + button to create a virtual server. For our exchange example we are creating three Virtual servers, one each for:

- IMAPI over RPC
- OWA and Outlook anywhere
- SSL redirection

MAPI-RPC Virtual Server Configuration

First create the IMAPI-RPC virtual server using the settings shown in Figure 57.

In the Layer 7 load balancing section, create a sticky group. Under the action, select sticky and configure the sticky group settings. Give the group a name and for the type choose IP Netmask. For the actual net mask setting, use a 32 bit netmask (255.255.255.255). For the address type, select source as the type.

Now create the sticky farm which will be associated with your real CAS servers. Figure 58 shows the IP addresses and names of the real CAS servers that can be selected to be associated with the server farm.

In the NAT configuration section, use the NAP pool that was created earlier. This is also shown at the bottom Figure 58.

Figure 57 Creating the IMAPI-RPC Virtual Server

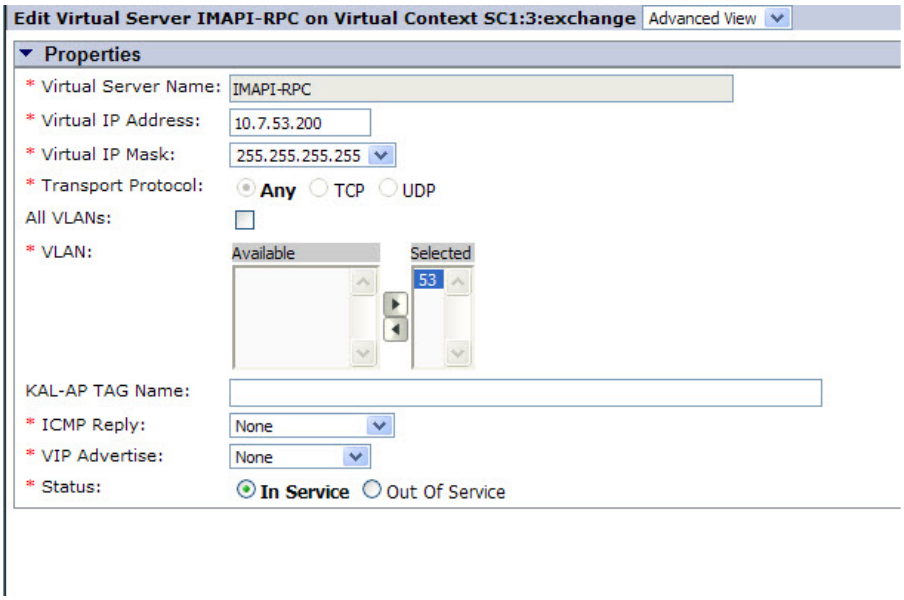

Figure 58 Creating the Sticky Farm Asociated with Real CAS Servers

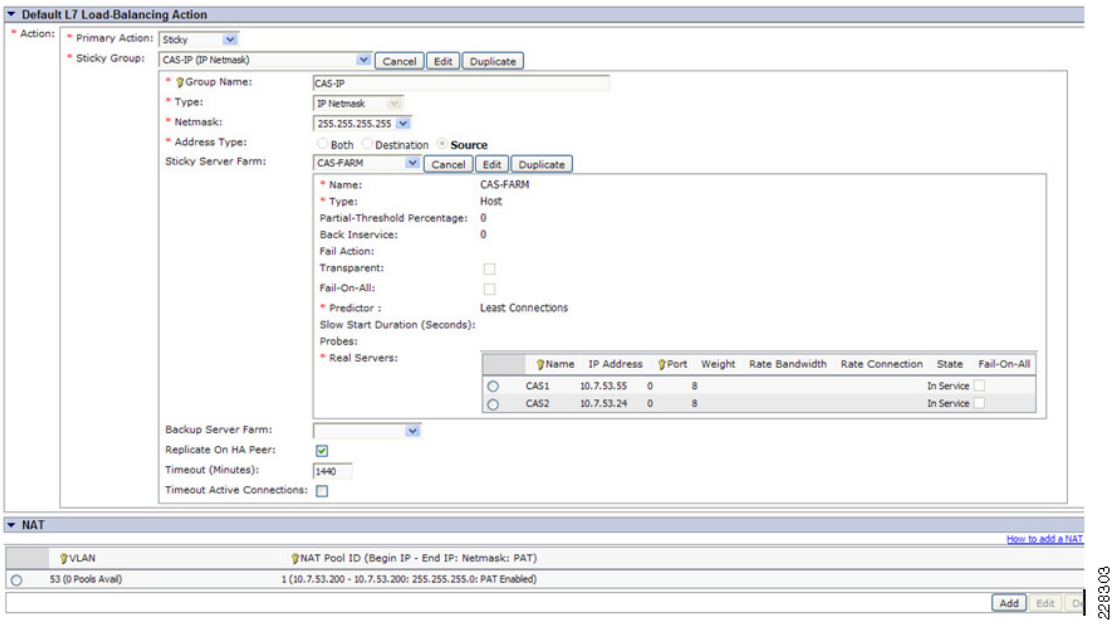

Outlook Web Access/Outlook Anywhere Virtual Server Setup

Now add another virtual server to support Outlook Web access and Outlook Anywhere. For the application protocol, select HTTPS and use 443 for the port. Under the SSL termination settings section, choose the OWA proxy service that was created earlier (shown in Figure 59). In the Layer 7 load balancing section for the rule match, select inline match and for the conditions type choose HTTP header. For the header name, select User-Agent and for the header value use MSRPC as the value.

Now that you have the inline match set, you can choose an action. Set the primary action to sticky and create a sticky group. In the example shown in Figure 60, you see that we created a sticky group called CAS-RPC-HTTP. Its type is set as an HTTP Header and the header name is set as Authorization. For the Sticky Farm, create a sticky farm to be used. In this example there is a sticky server farm called CAS-FARM-80. The sticky server farm settings can be seen in Figure 61.

In the Default L7 Load-Balancing Action section, choose Sticky as the primary action. For the sticky group, create a sticky group with the type of HTTP cookie and a cookie name of Cookie. Check the enable insert box and Browser Expire box and choose the CAS-FARM-80 Sticky Server Farm. For the time out setting option, set it to 60 minutes. These details can be seen in Figure 62.

For the Virtual servers NAT settings, you may share the same NAT pool that is used with the MAPI-RPC virtual server.

Figure 59 Specifying the OWA Proxy Service Name

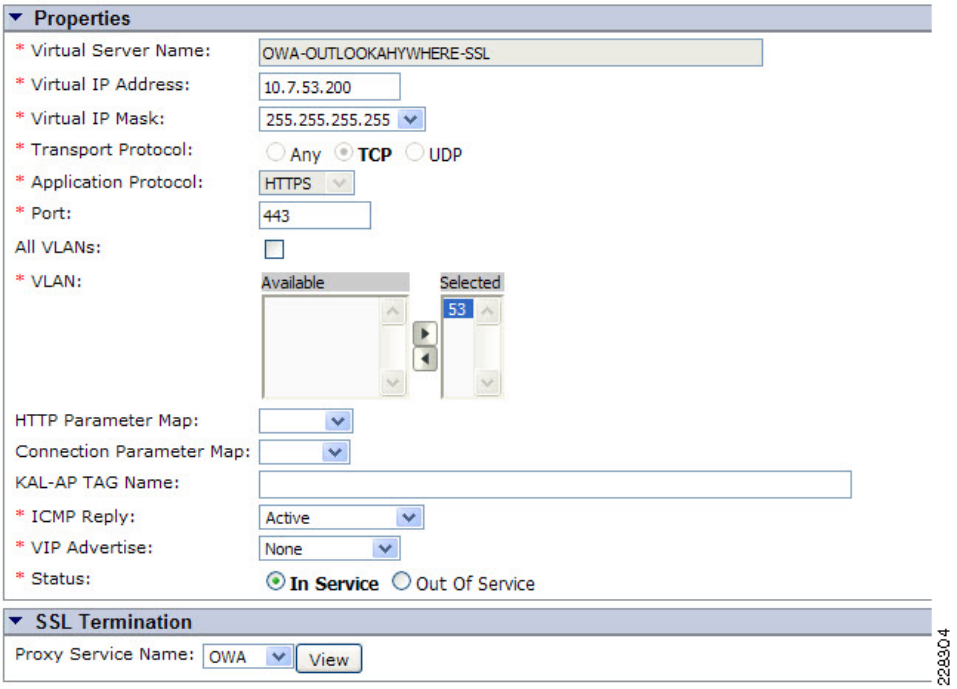

Figure 60 Creating the Sticky Group

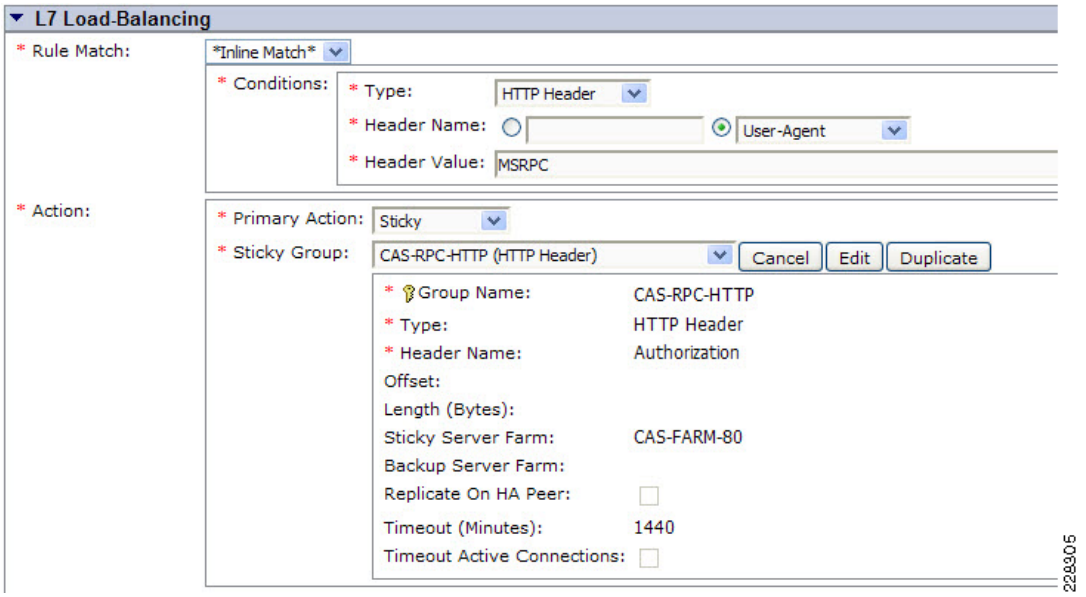

Figure 61 Sticky Server Farm Settings

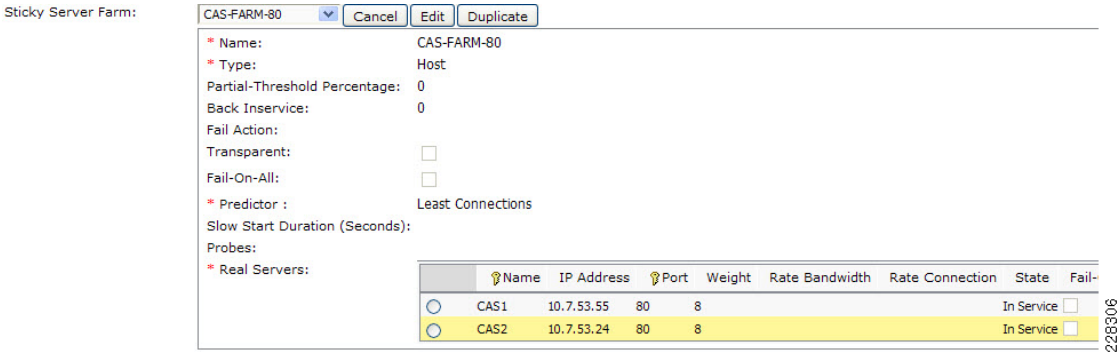

Figure 62 Creating the Layer 7 Load Balancing Action

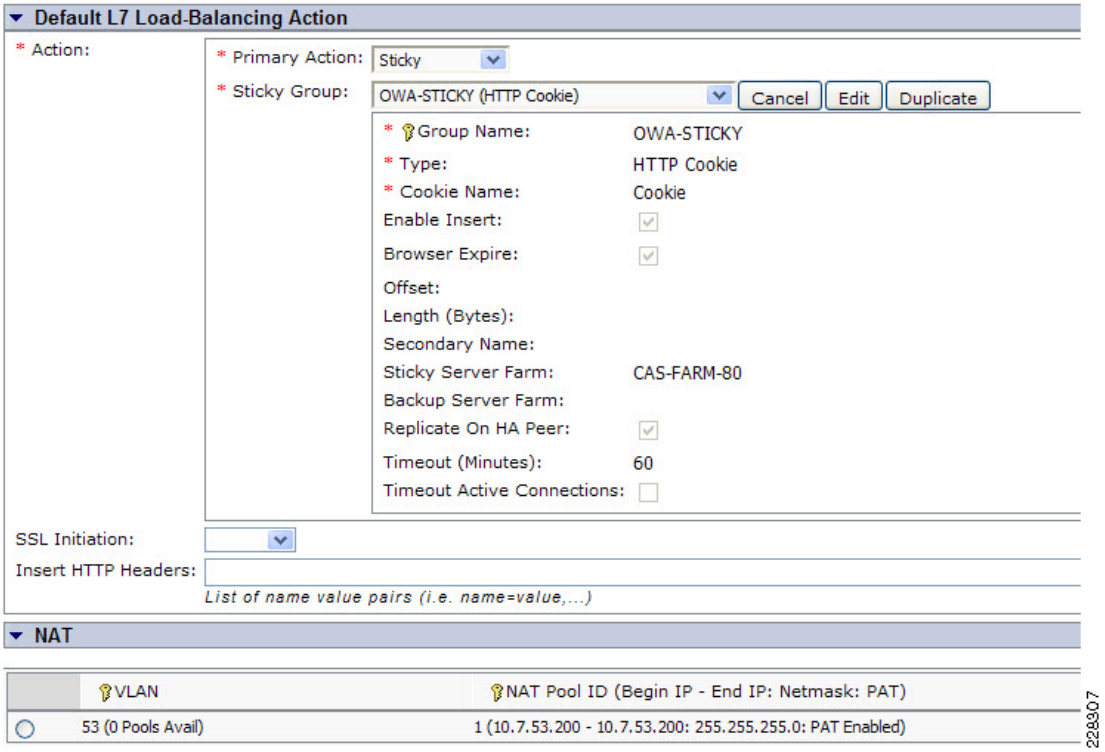

SSL—Redirect Virtual Server Setup

Another Virtual Server that is needed for the Exchange load balancing configuration is one that is set to redirect traffic from port 80 to port 443. To do this, create a virtual server as shown in Figure 63 and in the Layer 7 load balancing section of the configuration, set the primary action as load balance. For the server farm, create one and set the type as redirect. For the predictor option, choose Round Robin, then add a real server and for the name, select the previous name given in the server farm section. In our example shown in Figure 64, this is SSLREDIRECT. In the redirection code section, set it to 302 and for the URL, use the URL of your OWA server farm, such as https://aceexchange-vip.ucshypervroles.com/owa.

In this example, the machine name aceexchange-vip.ucshypervroles.com points to the VIP of the server farm in DNS 10.7.53.200. Therefore all traffic that comes in destined for http://aceexchange-vip.ucshypervroles.com/owa is redirected to https://aceexchange-vip.ucshypervroles.com/owa and load balanced according to the rules in the previous Virtual server configuration sections.

Figure 63 Creating a Virtual Server

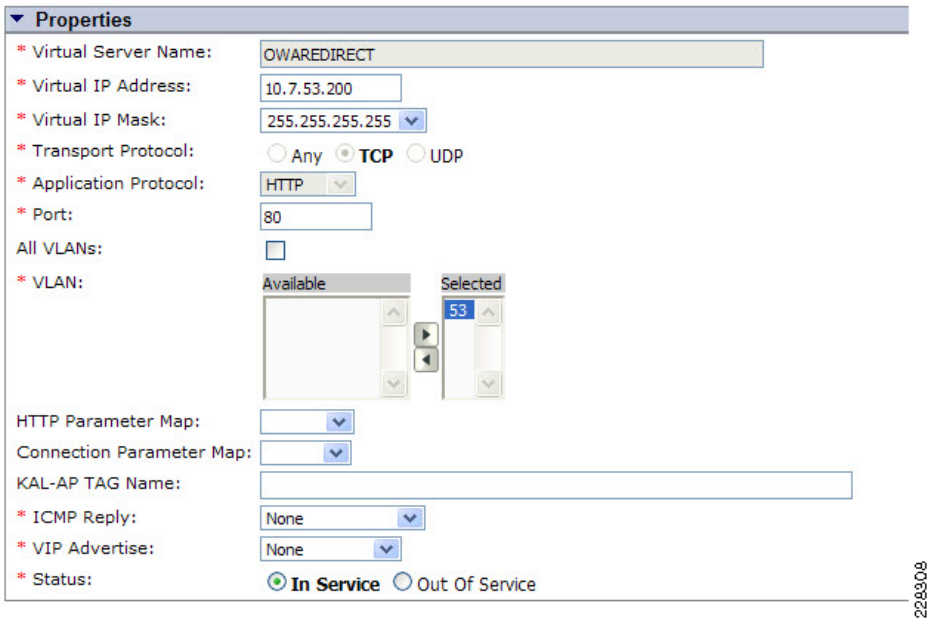

Figure 64 SSL Redirection

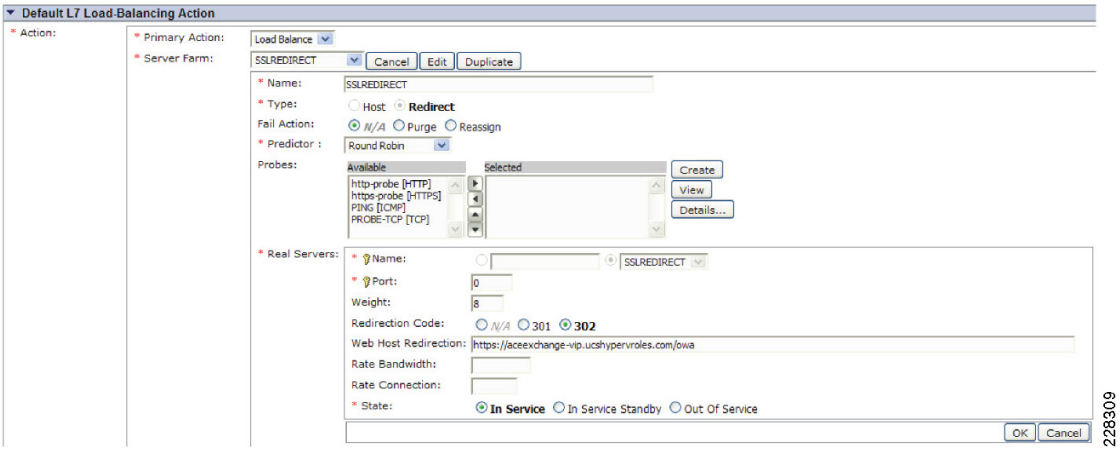

Exchange SSL Settings

As a requirement of SSL offload to the ACE, a minor change needs to be made to the CAS servers. When a SSL connection is terminated on the ACE, all connections from the ACE to the CAS servers are made via standard http. To support these non-SSL connections from the ACE to the CAS, you need to modify the default SSL settings in the Internet Information Services Manager. As shown in Figure 65, select the default site within Internet Information Manager on your CAS servers. Next uncheck the box so that the site no longer requires SSL. This is shown in Figure 66.

Figure 65 IIS SSL Setting

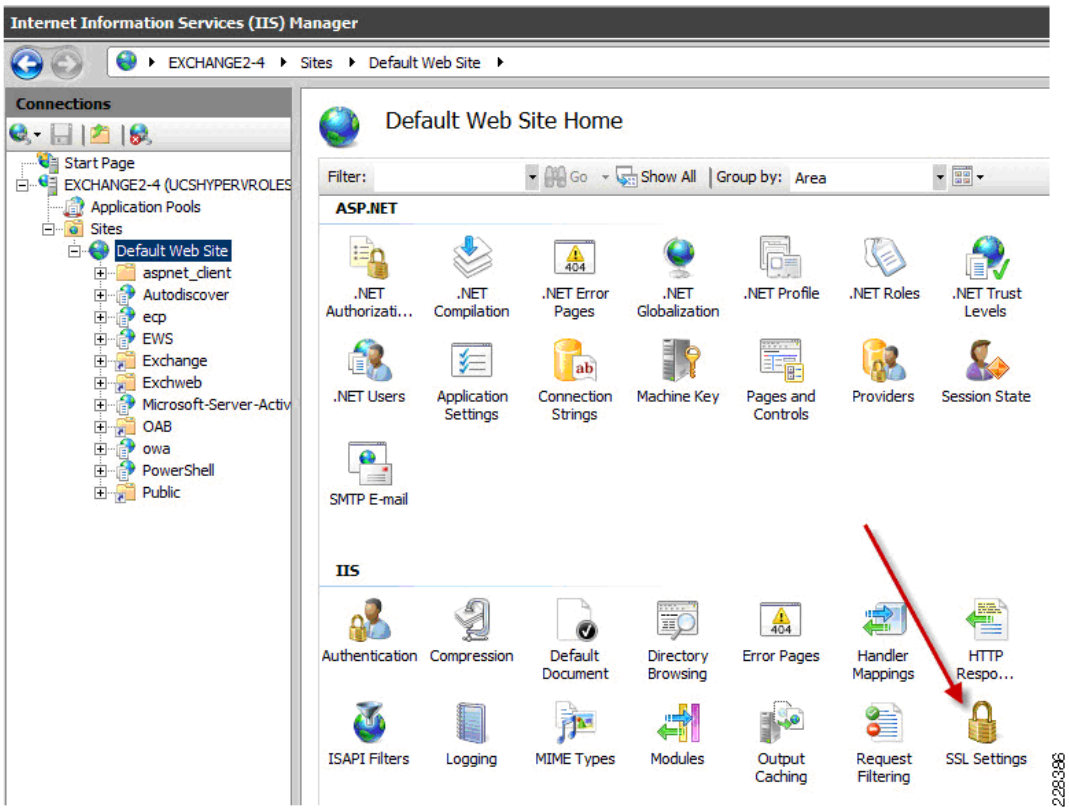

Figure 66 Unselecting SSL Requirement

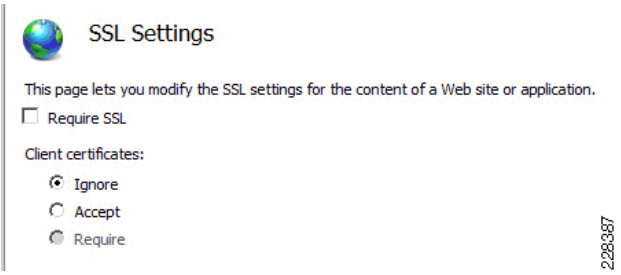

Exchange CAS Array Configuration

Another requirement to support hardware load balancing for your CAS servers is the configuration of a CAS array. In its configuration you must create a CAS array through the Exchange Management Shell that groups your CAS servers together. Then make sure the name of your CAS array is a DNS registered name that points to the ACE load balancer virtual IP.

1. On your domain controller, create a forward lookup entry that maps your ACE load balancer virtual IP to the fully-qualified domain name that you assign to your CAS array in the next step.
2. Then, create the CAS array with the command "New-ClientAccessArray -Name <NameoftheArray> -FQDN <NameoftheArray.Fully-qualified domain name> -Site <ADSiteName>".

3. For each of the mailbox databases that will be front-ended by your CAS servers, you need to assign the CAS array as the RPC Client Access server setting for each mailbox database. To see a list of all your mailbox database names, use the "Get-MailboxDatabase" command.
4. If you need to find out what each of your Mailbox databases is using as its CAS server, use the command "GET-MailboxDatabase "dbase4" | f1 RpcClientAccessServer".
5. Set the CAS server setting for each mailbox database to be the newly created CAS array with the command "Set-MailboxDatabase <database name/id> -RpcClientAccessServer <NameoftheArray.Fully-qualified domain name>".
6. When you initiate a connection from your Exchange user, be sure to specify the CAS array fully-qualified domain name (which is the same as the ACE load balancer VIP fully-qualified domain name) as the server to which you are connecting. The connection then goes to the ACE to be load-balanced amongst the CAS servers in your CAS array.

Wide Area Application Services

Validating WAAS Between Branch and Data Center

Figure 67 WAAS Appliances and WAAS Mobile in Test Topology

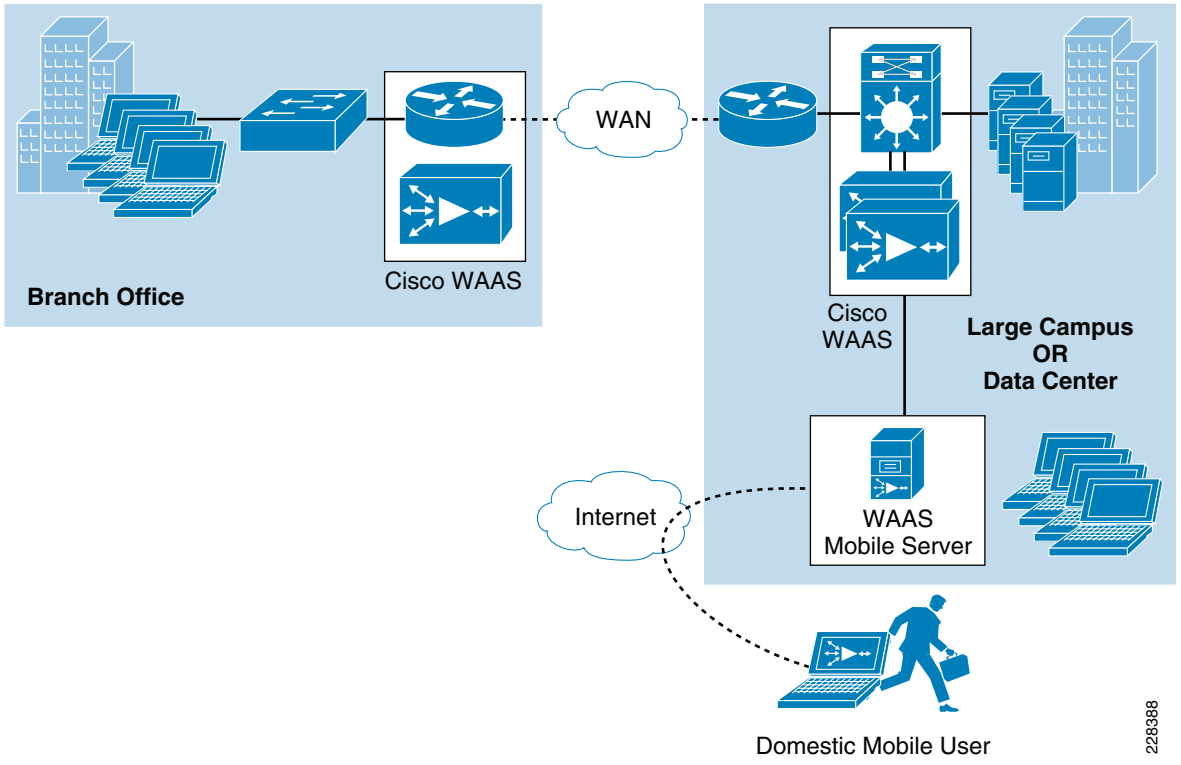

Two scenarios were tested in the lab to demonstrate Cisco WAAS performance improvements, one for users at the branch office and one for remote/mobile users, as shown in Figure 67. Two Exchange client types were used in testing WAAS performance between a branch office and the data center—Outlook Anywhere and Outlook 2010 RPC/MAPI client. Exchange online mode was chosen for the Outlook clients to minimize background traffic so that it would be possible to focus the performance measurement on the E-mail attachment download. Exchange client types were used to test WAAS

performance for a remote/mobile user: Outlook Anywhere and Outlook Web Access. Given the flexibility and security of Outlook Anywhere for client connectivity into the data center from either a branch office or a remote location, many enterprises may choose to deploy this client configuration at their branches to support the mobile and work-at-home lifestyles of their employees.

The scenario for branch office users consists of the following workflow:

1. An employee on the data center campus network sends an E-mail with an attached 7 MB Word doc to four employees located at the branch office.
2. The first branch office employee receives the Word document. Performance numbers are measured for this transaction. SSL acceleration performance can be shown for Outlook Anywhere traffic and MAPI acceleration performance can be shown for Outlook MAPI/RPC (unencrypted) traffic.
3. The rest of the branch office employees receive the Word document. Performance numbers are taken to show this subsequent optimization after the first file download has populated the DRE cache.
4. The employee on the campus network at corporate edits the document. The document is now a little less than 7 MB. The employee resends the document to the four branch employees.
5. The first employee at the branch receives the Word document. Performance numbers are measured for this transaction. This transaction still benefits from the DRE caching because the file is similar to the one sent previously.
6. The remaining three employees receive the Word document and performance numbers indicate their file downloads receive even higher performance optimization because the cache has just previously been updated with this exact file.

This simulates a very typical scenario in any enterprise environment, for example, a manager at the corporate office sends out a Word document for review by several employees at a branch office. The caching, LZ compression, SSL acceleration, and MAPI acceleration capabilities of the WAAS devices at the data center and the branch work to minimize the amount of redundant data traversing the WAN. This is especially important for branch offices since their WAN link sizes are typically limited to 3Mbps or smaller. For these scenarios, several link sizes ranging from 768Kbps to 3Mbps were tested. Latency and packet loss numbers were chosen to represent different geographical distances between branch office and data center.

The scenario for the remote/mobile user consists of the following workflow:

1. An employee on the data center campus network sends an E-mail with an attached 7 MB Word doc to a remote user.
2. The remote user opens the E-mail attachment.
3. Performance numbers are taken to show the optimization achieved with this first “cold” E-mail download where the local delta cache on the client has no data that can be leveraged to minimize redundant bytes across the WAN.
4. The user at the data center resends the E-mail attachment. For this second E-mail download, the local cache on the client is “hot” or has been populated by the first download. The remote user opens the E-mail attachment a second time and performance numbers are taken showing the benefits of the caching.

Setting Up the Outlook Anywhere Server and Client

The following configurations need to be done on the Exchange Client Access server:

1. Configure the Client Access Server so that it does not require its Outlook clients to use MAPI encryption. This allows clients that have encryption disabled to connect to the server and mailboxes and receive greater WAAS benefits. This is done through the Exchange Management Shell on the Client Access Server with the following series of commands:

```
[PS] C:\Windows\system32>Set-RpcClientAccess -Server KCCAS  -EncryptionRequired $false
[PS] C:\Windows\system32>Set-RpcClientAccess -Server KCMAILBOX-2 -EncryptionRequired
$false
[PS] C:\Windows\system32>Set-RpcClientAccess -Server KCMAILBOX-3 -EncryptionRequired
$false
[PS] C:\Windows\system32>Get-RpcClientAccess

Server          Responsibility            MaximumCo Encryption
BlockedClientVersions
                                              nnections    Required
------                ----------------          ---------          ---------
---------------------
KCCAS               Mailboxes             65536        False
KCMAILBOX-2     PublicFolders      65536        False
KCMAILBOX-3     PublicFolders      65536        False
```

2. Outlook Anywhere was enabled on the Client Access Server using the Exchange Management Console and the external FQDN and the authentication method are configured in this step.

Figure 68 Enabling Outlook Anywhere in Exchange Management Console

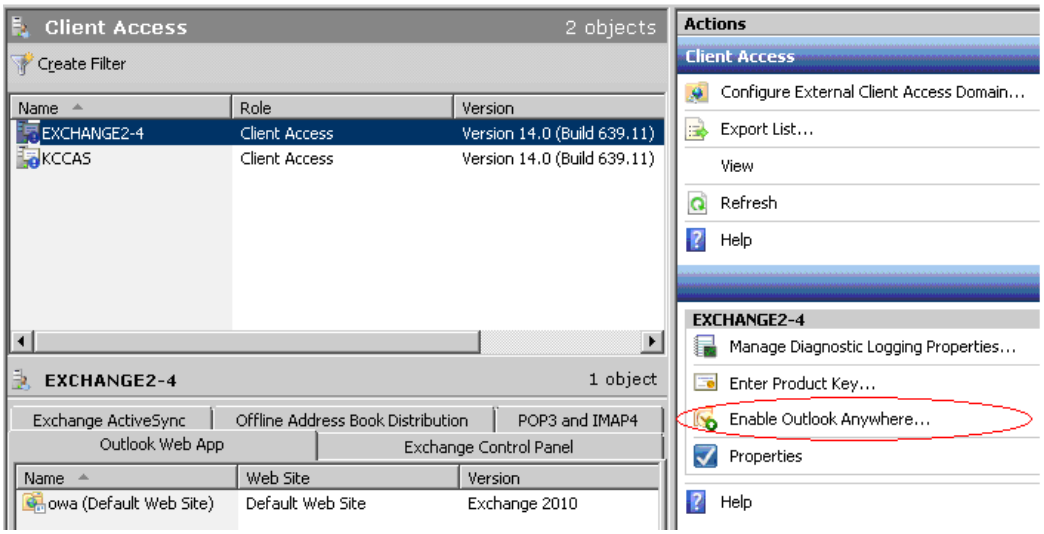

The following configuration was done on each of the Outlook 2010 clients at the branch office:

1. Since Outlook Anywhere is secured through SSL, it is not necessary to enable MAPI encryption on the Outlook 2010 client. WAN optimization benefits are greater with MAPI encryption disabled. The Outlook 2010 client enables encryption by default so that option needs to be de-selected. See Figure 69 for the location of that setting.
2. Since NTLM authentication is configured on the Client Access Server (see Figure 70), NTLM authentication is chosen on the Outlook client as shown in Figure 69.

Figure 69 Disabling MAPI Encryption and NTLM Authentication Setting

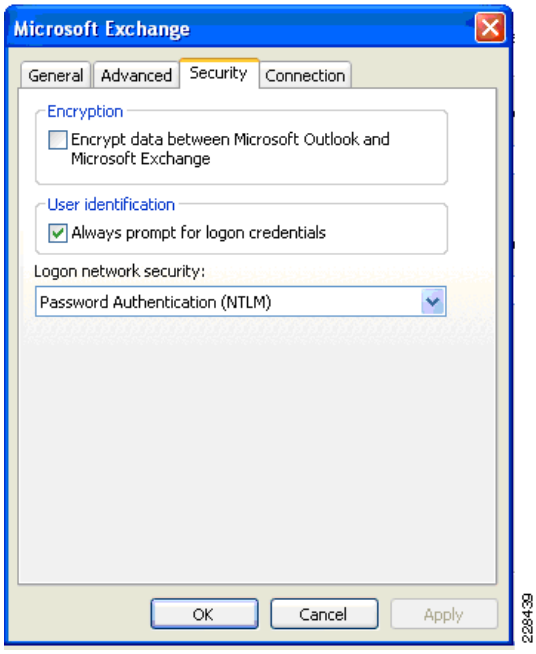

Figure 70 NTLM Authentication for Outlook Anywhere on Server

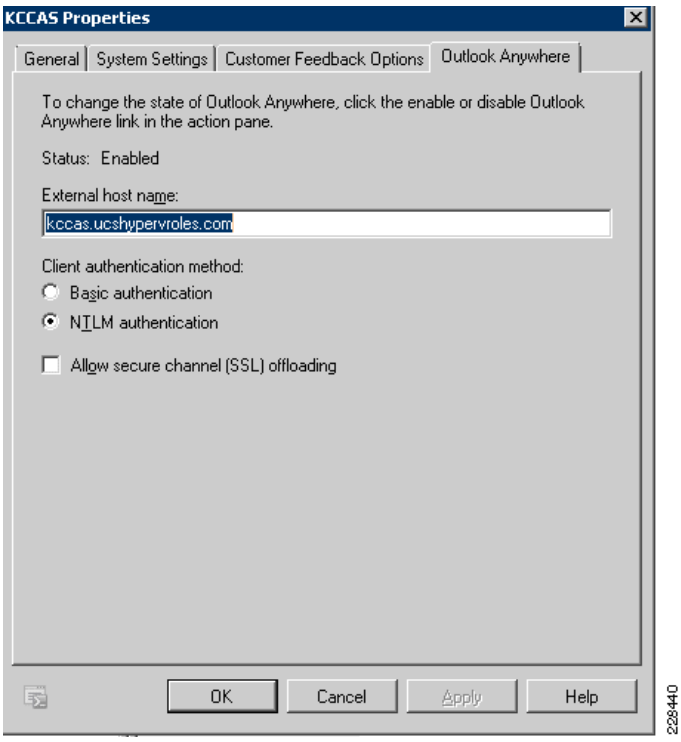

3. The Outlook 2010 client must be set to use HTTPS instead of MAPI/RPC. This setting is configured under the "Connection" tab by selecting "Connect to Microsoft Exchange using HTTP" and going into the "Exchange Proxy Settings" window to configure the following parameters (also see Figure 71 and Figure 72).
   - URL that the client should use to connect to the SSL proxy server—This URL would be set to the FQDN of the internal host name of your CAS server. In the case where an ACE device is deployed to load-balance CAS servers, the FQDN for the ACE VIP would be used in this URL.
   - Select "Connect Using SSL only"—The "Only connect to proxy servers that have this principal name in their certificate" field can be left blank or filled in with the common name of your SSL certificate, which in this test was the FQDN of the Client Access server.
   - NTLM authentication is chosen to match the server setting.
   - Select the two options that allow the Outlook client to fail over to using TCP/IP (i.e., RPC/MAPI) should HTTPS fail to connect.

Figure 71 Enabling Outlook Anywhere on Outlook Client

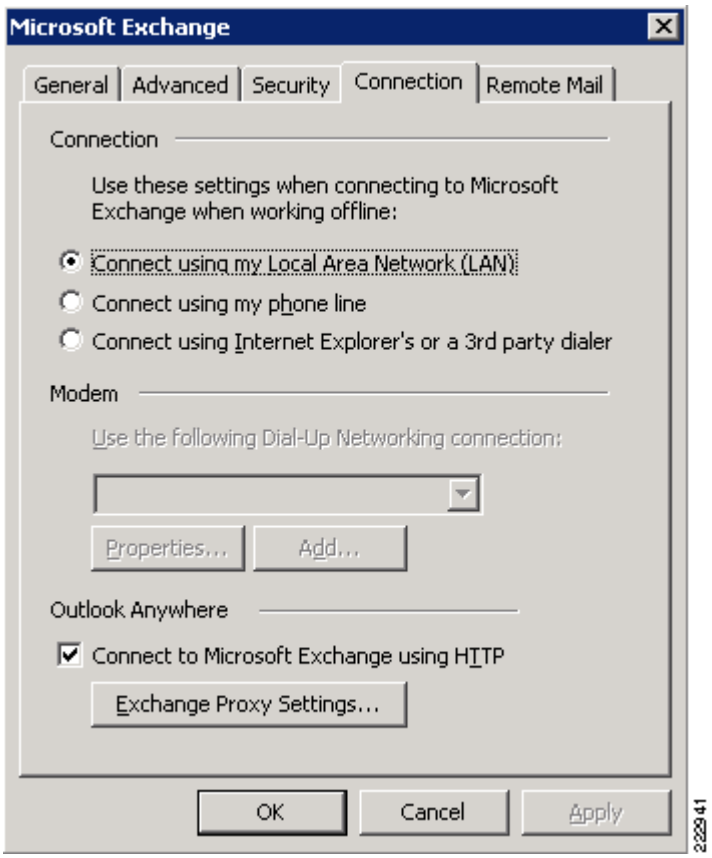

Figure 72 SSL Proxy Settings for Outlook Anywhere Client

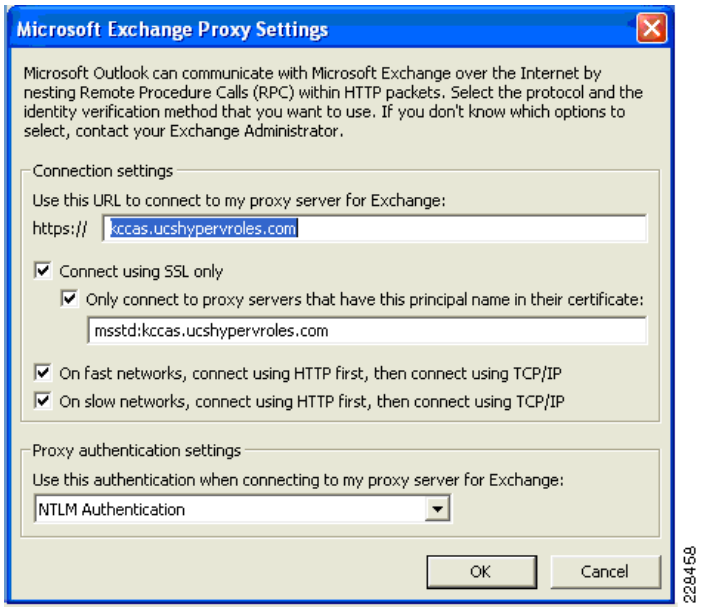

Creating an SSL Certificate for Outlook Anywhere and Outlook Web Access

Note Outlook Web Access does not require a trusted third-party certificate. However, since Outlook Anywhere does require that, the same certificate is used for both types of SSL sessions.

The steps for setting up the SSL certificate on the Exchange Client Access Server and on each Outlook 2010 client machine are:

1. Select the Client Access Server in the Exchange Management Console and select "New Exchange Certificate" (see Figure 73). This utility helps you generate a certificate request that you can submit to a Certificate Authority to create the certificate or certificate chain file.

Figure 73 New Certificate Request Through Exchange Management Console

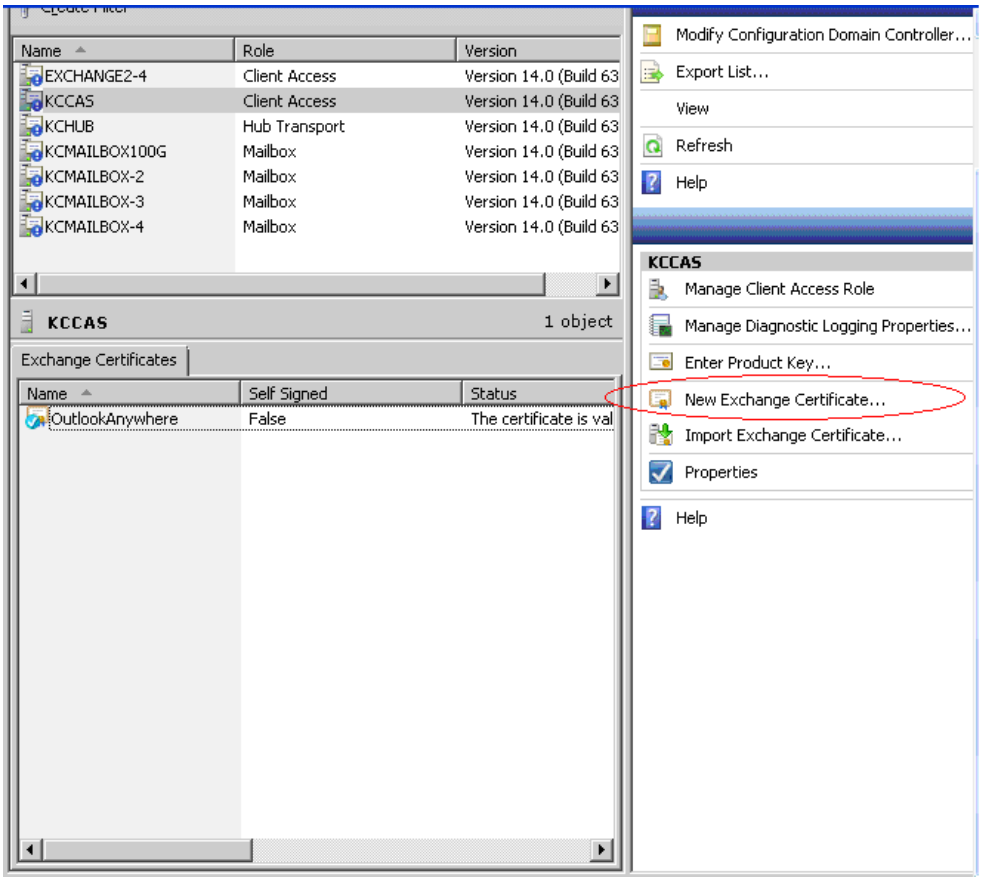

2. Specify the following values in the wizard:
   - Select the option to make this a wildcard certificate.
   - Specify the root domain for the wildcard (e.g., *.mycompany.com).
   - Fill in the organization, country, city, and state that are appropriate.
   - Specify the location and filename (*.cer) of the new certificate request you are generating with this wizard.
3. Once the certificate request (*.cer) file is created, it must be submitted to a Certificate Authority (CA). In this solution testing, the Microsoft Certificate Server available on Windows 2008 R2 Data Center was used. To do that, first open up the certificate request file in Notepad.

Figure 74 Input for Certificate Request to CA

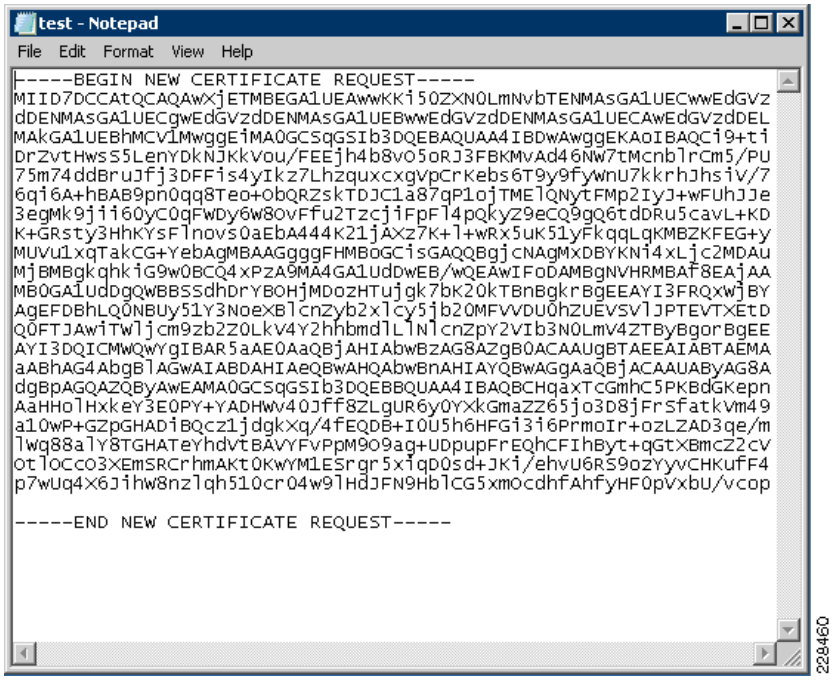

4. Copy all the contents of the request file, which you need when you submit your request to the CA.
5. Open a browser to the Microsoft CA server. You need to use two functions available—requesting a new certificate and downloading the CA certificate (see Figure 75).

Figure 75 Windows Certificate Server

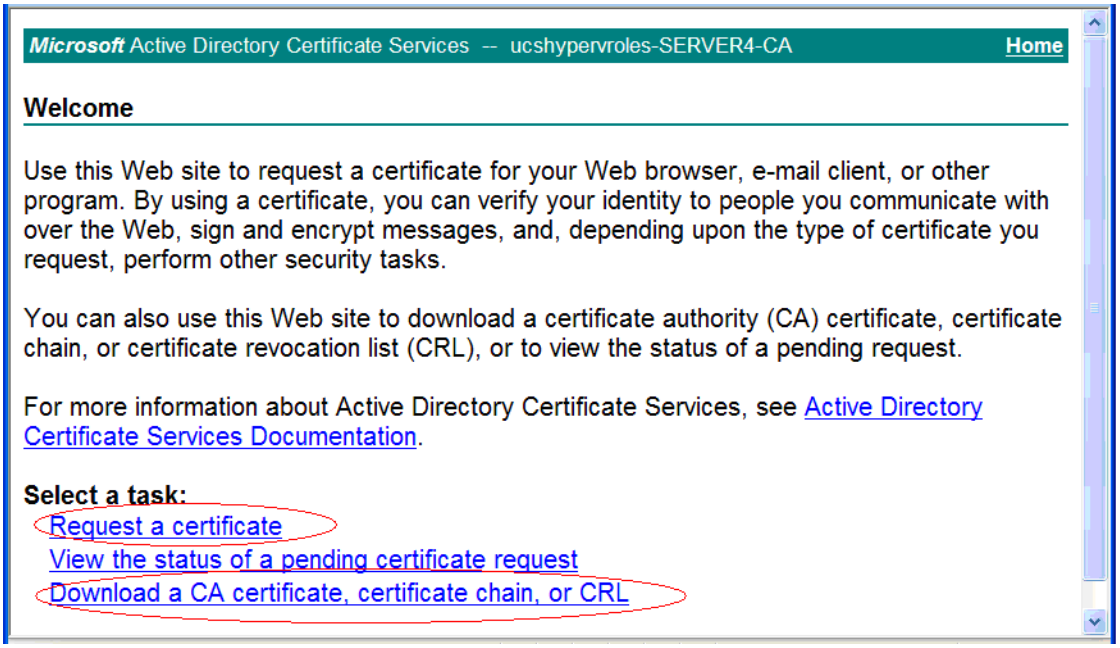

6. First, download the CA certificate and install it into the Trusted Root Certification Authorities Certificate folder in the client's local machine store. This allows the Outlook Anywhere client to trust all certificates issued by this CA.
7. Next, use the "Request a Certificate"—>"advanced certificate request"—>"Submit a certificate request by using a base-64-encoded…". Paste in the contents you copied from step 3 above into the "Saved Request" box, select "Web Server" as the certificate template, and submit the request.
8. Download the certificate chain that is generated.
9. In the Exchange Management Console, use the "Import Exchange Certificate" to import the certificate chain.
10. The last step is to assign the "IIS" service to the certificate as shown in Figure 76.

Figure 76 Assign Services to Certificate in Exchange Management Console

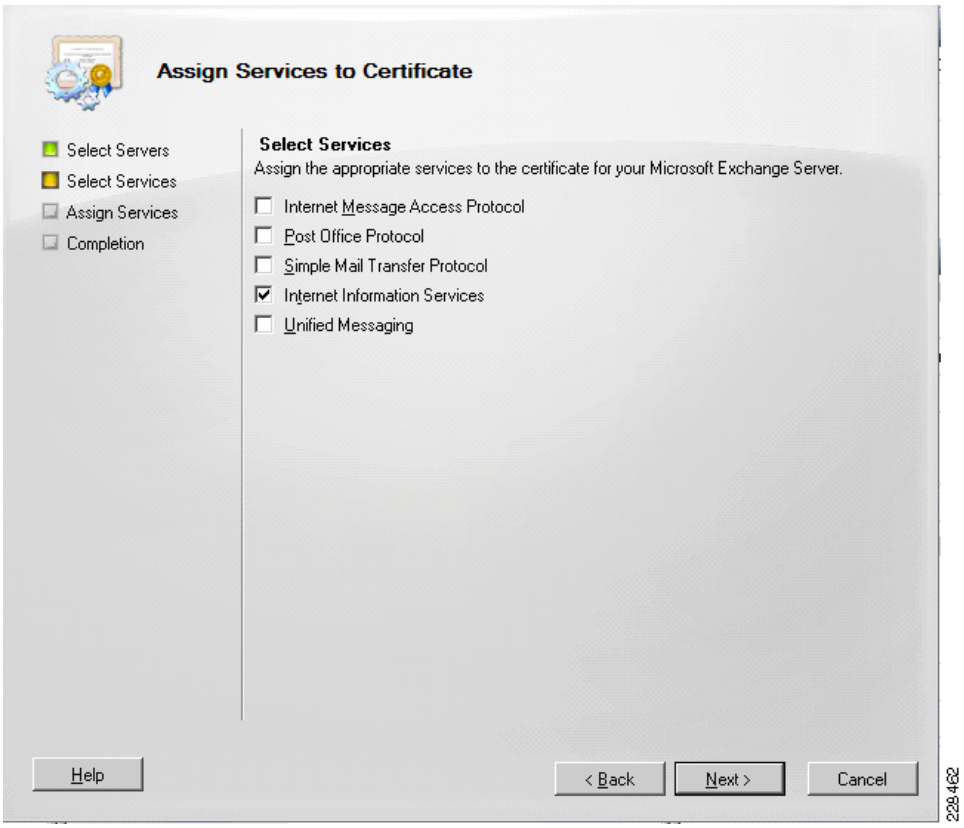

Branch to Data Center

Figure 77 Branch to Data Center

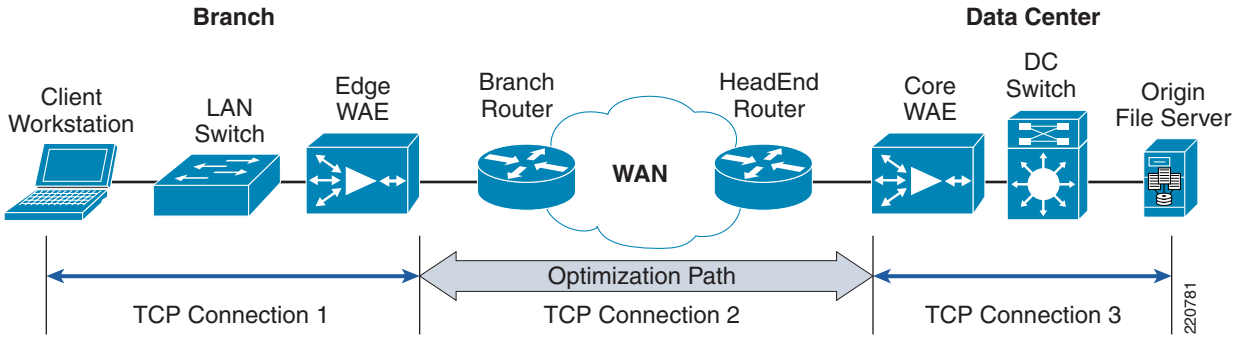

WAN optimization and acceleration is provided for client sessions from the branch to the data center using the Wide Area Application Engine network module on the branch Integrated Services router and the Wide Area Application Services Engine WAE-7371 appliance at the data center WAN edge. This section is meant to highlight how these devices were configured to support the application acceleration involved in this solution, specifically SSL acceleration of Outlook Anywhere traffic and MAPI acceleration of Outlook 2010 RPC (MAPI) client. To demonstrate MAPI acceleration of a branch user that is typically connected to the data center through a secure VPN tunnel, native encryption on the Outlook client was disabled.

Device Configurations for Branch to Data Center WAN Optimization

Data Center WAE-7371 and Branch NM-WAE-502 Configuration

This section shows the parts of the WAE-7371 and NM-WAE-502 configuration that were relevant to the solution for network connectivity and management, for traffic flow optimization and compression, and for the appropriate application acceleration to be applied to the Exchange client traffic types. Note that the WAAS Central Manager can be used to configure, edit, and review the WAE device configurations. The command line output of the device configurations are given below instead of the WAAS Central Manager GUI screenshots to keep the information more concise. The lines of configuration shown in these sections do not necessarily correspond to how they are displayed when a command line user issues a **show running-config** command on the WAE device.

The WAE-7371 was connected to the ASR 1002 WAN edge router in the data center, so its default gateway was set to the interface on the ASR to which it was connected. Also, the WAE device is registered to the WCCP service group on the ASR WAN router and the router uses Layer 2 redirection to redirect traffic to the WAE device.

```
interface GigabitEthernet 1/0
 description To WAN-ASR GE3
 ip address 10.7.13.2 255.255.255.0
 exit
ip default-gateway 10.7.13.1
!
wccp router-list 1 10.7.13.1
wccp tcp-promiscuous router-list-num 1 l2-redirect mask-assign
wccp version 2
```

The NM-WAE module was installed in the branch router and connected to it through its internal interface; it had to be configured to point to the IP address of that internal interface. It is also registered in the WCCP service group on the branch router. For this branch configuration, redirection is done through GRE.

```
interface GigabitEthernet 1/0
 ip address 10.7.12.2 255.255.255.0
 no autosense
 bandwidth 1000
 full-duplex
 exit
interface GigabitEthernet 2/0
 shutdown
 exit
!
!
ip default-gateway 10.7.12.1


!
wccp router-list 1 10.7.12.1
wccp tcp-promiscuous router-list-num 1
wccp version 2
```

Both WAE devices at the data center and the branch had to be configured with the following configurations.

In order to ensure date and time are synchronized with the WAAS network module at the branch and WAAS Central Manager for accurate reporting and statistics, the Network Time Protocol server on the device was set to the enterprise NTP server.

```
ntp server 171.68.10.80
```

```
Note that the MAPI/RPC Outlook client initially targets port 135 on the Client Access
Server for its  Microsoft End Point Mapper (EPM) traffic.  This initial port can be
classified as EPM on the WAE device and optimized. Note that this classification of EPM
traffic is needed only for the MAPI/RPC traffic and is not necessary for proper
classification and optimization of Outlook Anywhere or Outlook Web Access traffic.
   classifier MS-EndPointMapper
      match dst port eq 135
   exit
      name Other classifier MS-EndPointMapper action optimize DRE no compression
```

The dynamically-negotiated ports in the end point mapping process between Exchange Client Access server and Outlook 2010 RPC/MAPI client had to be classified correctly as MAPI. This required a special map adaptor to be used to classify this traffic which would then have MAPI acceleration applied to it.

```
   map adaptor EPM mapi
      name Email-and-Messaging All action optimize full accelerate mapi
   exit
   map adaptor EPM ms-ad-replication
      name Replication All action optimize full
   exit
   map adaptor EPM ms-frs
      name Replication All action optimize full
   exit
   map adaptor EPM f5cc5a18-4264-101a-8c59-08002b2f8426
      name Email-and-Messaging All action pass-through
   exit
   map other optimize full
exit
```

The classification of HTTP over SSL is done on port 443.

```
classifier HTTPS
   match dst port eq 443
exit
```

The SSL acceleration service is enabled on each WAE device, as shown with the **show accelerator ssl** CLI command:

```
dc-wae1#show accelerator ssl

Accelerator     Licensed        Config State    Operational State
-----------     --------        ------------    -----------------
ssl             Yes             Enabled         Running

SSL:
   Policy Engine Config Item            Value
   -------------------------            -----
   State                                Registered
   Default Action                       Use Policy
   Connection Limit                     12000
   Effective Limit                      12000
   Keepalive timeout                    5.0 seconds
```

The WAE device has to report optimization and acceleration statistics to the WAAS Central Manager, which is enabled with:

```
central-manager address 10.7.53.9
cms enable
```

The data center WAE device must be configured with a certificate for the SSL handshake to support acceleration of Outlook Web Access and Outlook Anywhere traffic. Since Outlook Anywhere will not accept self-signed certificates, the self-signed certificate generated by the WAE device cannot be used; instead, the trusted certificate generated using the Exchange Management Console is used (described in Creating an SSL Certificate for Outlook Anywhere and Outlook Web Access).

The trusted certificate is imported into the data center WAE device using the WAAS Central Manager in the configuration where the SSL application acceleration service is defined and enabled for the Exchange Client Access Servers.

The following **show crypto certificates** shows the trusted third-party certificate installed in the device's managed store after it was imported. Note that the common name of the certificate can be a wildcard URL instead of the specific FQDN of the Exchange server or proxy server through which the Outlook Anywhere clients connects. This gives greater flexibility so that the same certificate can support multiple servers.

```
dc-wae1#show crypto certificates

Certificate Only Store:
-----------------------
File: dc-ca-01-CA.ca             Format: PEM
Subject: CN=dc-ca-01-CA
Issuer: CN=dc-ca-01-CA
--------------------------------------------------------------------------------
File: ESEDataCenterCA.ca                 Format: PEM
Subject: O=ESE/OU=Data Center/emailAddress=chrobrie@cisco.com/L=RTP/ST=North Car
olina/C=US/CN=ESEDataCenterCA
Issuer: O=ESE/OU=Data Center/emailAddress=chrobrie@cisco.com/L=RTP/ST=North Caro
lina/C=US/CN=ESEDataCenterCA
--------------------------------------------------------------------------------

Managed Store:
```

```
--------------
File: CAS.p12            Format: PKCS12
EEC: Subject: C=US/ST=CA/L=San Jose/O=SJDCAL/OU=ESE/CN=*.ucshypervroles.com
     Issuer: DC=com/DC=ucshypervroles/CN=ucshypervroles-SERVER4-CA
-------------------------------------------------------------------------------

Local Store:
------------
Machine Self signed Certificate
-------------------------------
Format: PKCS12
Subject: C=US/ST=California/L=San Jose/OU=ADBU/O=Cisco Systems/CN=NO-HOSTNAME/em
ailAddress=tac@cisco.com
Issuer: C=US/ST=California/L=San Jose/OU=ADBU/O=Cisco Systems/CN=NO-HOSTNAME/ema
ilAddress=tac@cisco.com

Management Service Certificate
------------------------------
Format: PKCS12
EEC:Subject: C=US/ST=California/L=San Jose/OU=ADBU/O=Cisco Systems/CN=NO-HOSTNAM
E/emailAddress=tac@cisco.com
    Issuer: C=US/ST=California/L=San Jose/OU=ADBU/O=Cisco Systems/CN=NO-HOSTNAME
/emailAddress=tac@cisco.com
The WAAS Self Signed Certificate is being used as the Management Service Certifi
cate
```

Remote User to Data Center

Cisco WAAS Mobile consists of the WAAS Mobile server software deployed on the Exchange application data VLAN behind the data center Nexus 5000 access layer and the WAAS mobile client software installed on the Exchange user's laptop. No changes are required on the Exchange 2010 servers and Outlook 2010 users and Outlook Web Access users can quickly download and install the pre-configured WAAS Mobile client configuration from the WAAS Mobile Server and be up and running with the optimization and acceleration benefits within minutes. This section explains what was configured to enable and optimize SSL connections into the Exchange 2010 server farm. It also gives performance test results.

WAAS Mobile Server and Client Configuration

The WAAS Mobile Server was configured to support Outlook Anywhere connections from the Outlook 2010 client and to support HTTPS connections from Internet Explorer to the Exchange Web service. The following are the steps involved. Note that all configuration of the WAAS Mobile Server and WAAS Mobile client distribution is done by opening a browser to the WAAS Mobile Server (e.g., http://<WAAS Mobile Server FQDN>/WAASMobile). Refer to the following Cisco WAAS Mobile Server documentation for basic setup information:
http://www.cisco.com/en/US/docs/app_ntwk_services/waas/waas_mobile/v3.4/configuration/administration/guide/CiscoWAASMobile_AG3.4.pdf.

1. Since the solution is using a trusted root CA, set up the WAAS Mobile Server to be a subordinate CA.
2. Create the WAAS Mobile client distribution for the Outlook Anywhere and Outlook Web Access client types.
3. On the Outlook client machine, download the client distribution and install it.

Configure WAAS Mobile Server as Subordinate CA

1. Stop the WAAS Mobile Server.
2. Create or Set the reg value HKEY_LOCAL_MACHINE\SOFTWARE\ICT\AcceleNetServer\Options\HTTPS\UseSelfSigned CACert to 0 (DWORD).
3. Start the WAAS Mobile Server.
4. This should force it to create a certificate request file that gets placed in: C:\WINDOWS\system32. The file name should be WAASMobileCA.req.
5. You must submit that file to your Enterprise Certificate Authority (Enterprise CA) to get a certificate file.
6. It is important that the entire certificate chain be gathered from the CA (the file type is p7b) and then placed on the WAAS Mobile Server.
7. Import the certificate into the personal machine store on the WAAS Mobile server machine using mmc and the Certificate Snap-in.
8. It is also essential that the CA root certificate be installed in the Trusted Root Certification Authorities—>Certificate folder on the WAAS Mobile server if it is not already (i.e., the WAAS Mobile server must trust the CA).
9. Create key (STRING): “ClientCertStoreForCACert”=”CA” on the server under:

   Software\ICT\AcceleNetServer\Options\Version\<Client-DistributionLabel>\Options\HTTPS

   Where <Client-DistributionLabel> is the label of the client distribution. Click “Apply Changes” on the same HTTPS page in the client configuration page in the server Web interface to ensure that the configuration change is pushed out to clients.
10. Restart the WAAS Mobile server.

Create Client Distribution for WAAS Mobile Clients

Create the client distribution for your Outlook Anywhere and Outlook Web Access connections. For basic information on creating a client distribution and configuring basic parameters, refer to the WAAS Mobile Server Administration Guide: http://www.cisco.com/en/US/docs/app_ntwk_services/waas/waas_mobile/v3.4/configuration/administration/guide/CiscoWAASMobile_AG3.4.pdf.

The following highlights the steps that are important to this solution:

1. Go to the Proxied Process List Web page. Since the process version of outlook.exe for Outlook 2010 is higher than 12.9, a new proxied process must be manually defined. For flexibility for future version changes, use a wildcard for the maximum version setting. First, delete the existing outlook.exe process in the list for version 12.0 to 12.9.
2. Create a new process as shown in Figure 78, with minimum version set to 12.0 and maximum version left as a wildcard. Keep the defaults and add the process. Apply change to save it to the client distribution.

Figure 78 Proxied Process List

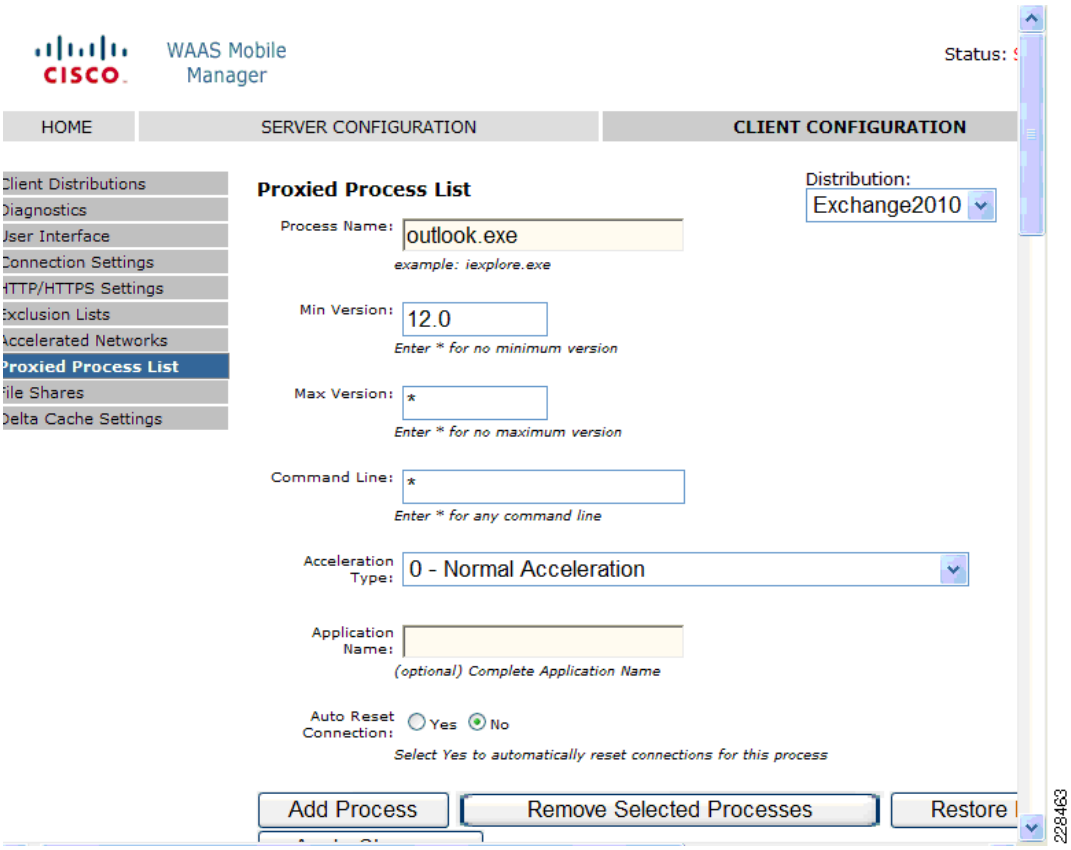

3. Configure the HTTP/HTTPS Settings as shown in Figure 79. Note that a host inclusion list can be used or the broader setting of accelerating all HTTPS sites can be chosen. Since the new outlook.exe process for Outlook 2010 was created earlier, it can now be chosen from the "Process Acceleration List" choices and added to the HTTPS settings.

Figure 79 HTTPS Settings for Client Distribution

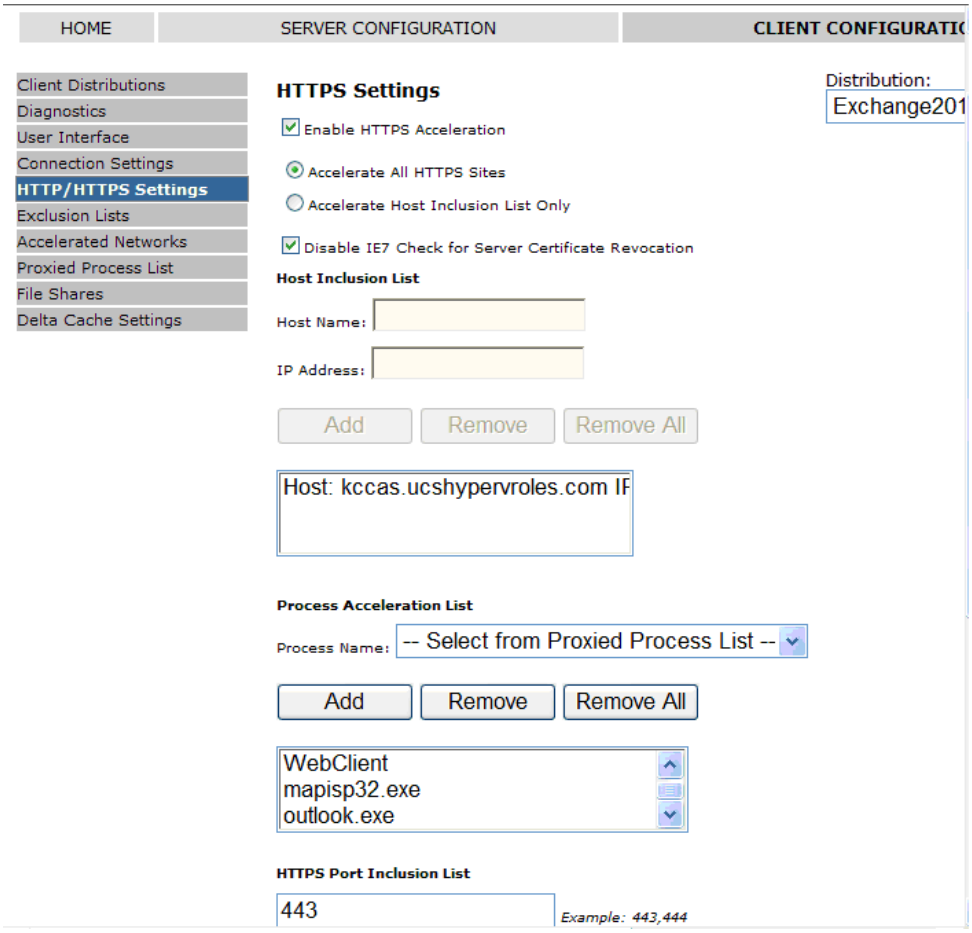

4. The "Delta Cache Settings" can be left at their default values except that the HTTPS Caching should be enabled to realize the full benefits of HTTPS acceleration.

Figure 80 Delta Cache Settings for Client Distribution

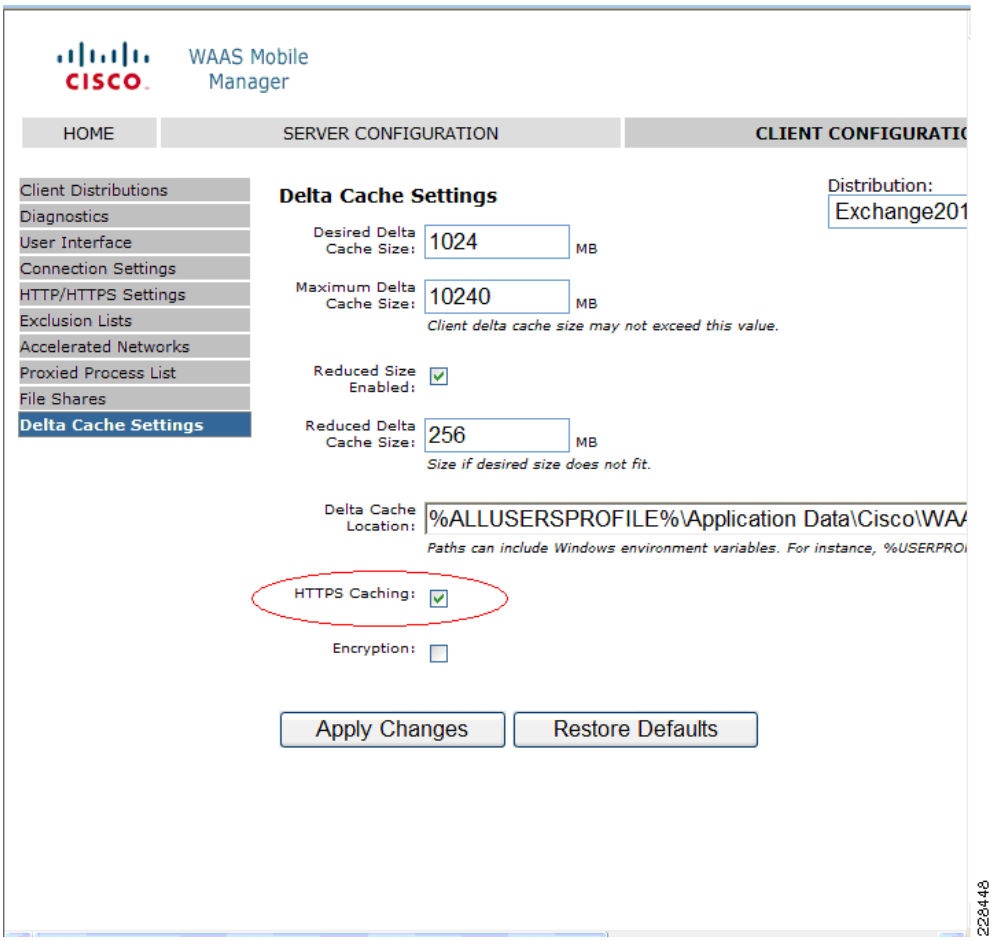

Installing WAAS Mobile Client

The client distribution that was created previously can now be downloaded by browsing to the client distribution links from the client machine and installing the executable (see Figure 81).

Figure 81 Client Distribution URLs for Download

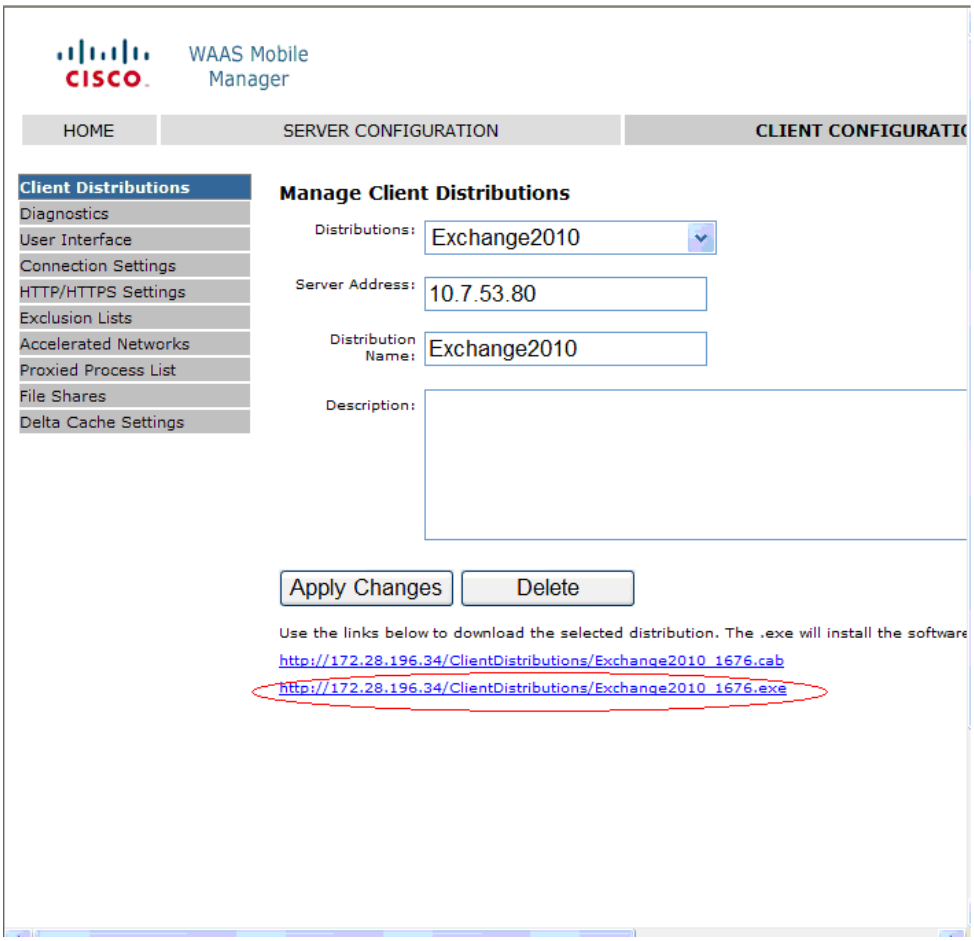

WAAS Performance Results

Table 6 lists the WAN link profiles that were used in the tests.

Table 6 WAN Link Profiles Used in Testing

Link Profile	Bandwidth	Latency	Packet Loss
Regional Office 1	3Mbps	40ms	0.1%
Regional Office 2	1.44 Mbps	40ms	0.1%
Other Coast Office	768Kbps	60ms	0.1%
Remote User	1.5Mbps	50ms	0.5%

The Remote User profile was used for WAAS Mobile testing. The other link profiles were used for branch to data center WAAS testing.

Branch Users

For the data center to branch scenario, where the campus user sends the same E-mail with attachment to all four users at the branch office, a WAAS appliance at the WAN edge of the data center and a WAAS network module on the branch Integrated Services Router provide TCP optimization, LZ compression, and SSL acceleration benefits. The following optimization numbers were measured using the reporting available on WAAS Central Manager. Examples of measurements that were provided on the data center WAE device command line are given in Appendix A—WAE Device CLI Statistics.

Outlook Anywhere Results

Native MAPI encryption is disabled to realize full optimization benefits.

Figure 82 Traffic Reduction for Outlook Anywhere Branch Users

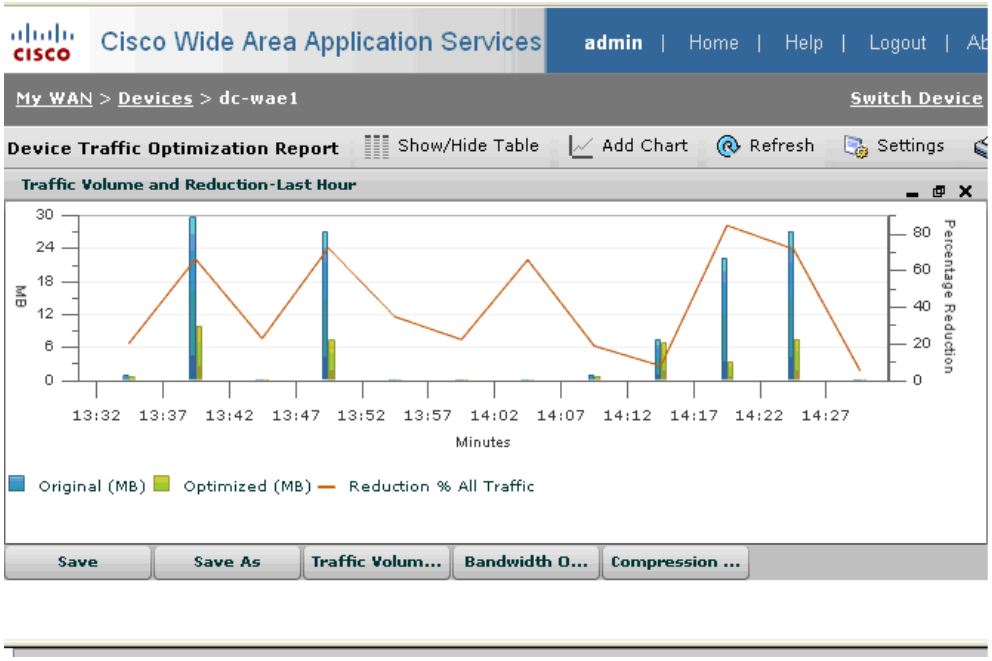

The bars at times 14:14, 14:19, and 14:24 in the graph in Figure 82 correspond to the performance numbers given in Table 7.

Table 7 Performance Numbers for Outlook Anywhere Branch to DC

	% reduction in traffic	Effective link capacity
Cold (no cache data)	9%	1.1
Hot (in cache)	85%	6.6
Warm (file is slight modified)	39%	2
Subsequent Hot (in cache)	70%	3.5

- Cold = One branch user opens E-mail attachment first before anyone else.
- Hot = Second and subsequent users open same E-mail attachment. Cache is hot.
- Warm = Sender modifies file slightly and sends it out. One branch user opens it first.

- Subsequent Hot = The rest of the branch users open up the modified file.

The graph in Figure 83 shows the effective link capacity for the Outlook Anywhere traffic that resulted from caching and compression (see the following Figure 84).

Figure 83 Effective Link Capacity for Outlook Anywhere

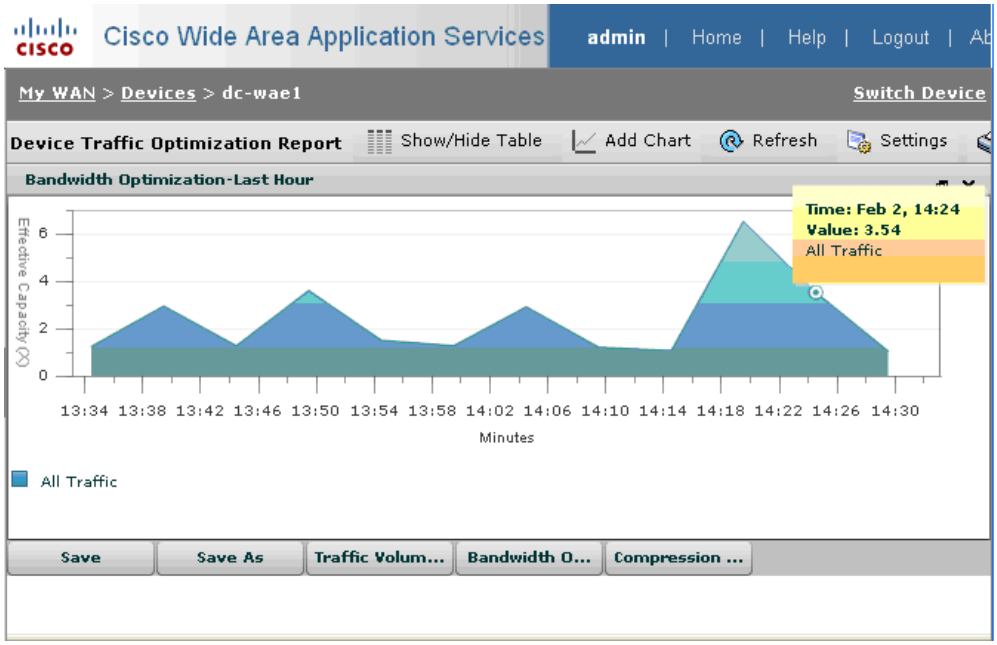

Figure 84 Percentage of Traffic Compressed for Outlook Anywhere

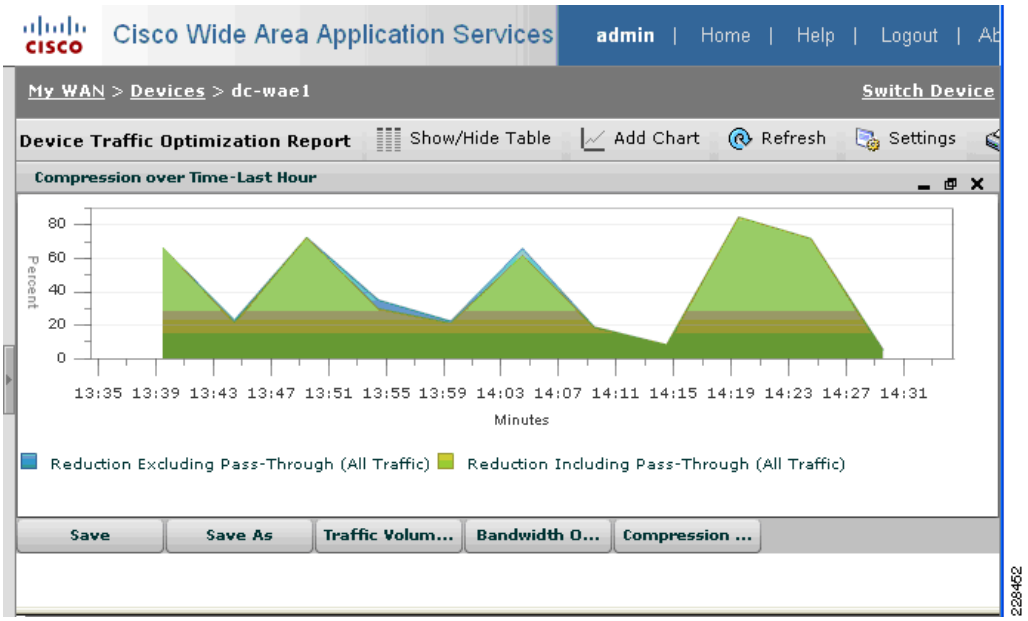

The SSL acceleration graphs shows that SSL acceleration was applied to provide the performance benefits.

Figure 85 Increased Effective Link Capacity from SSL Acceleration of Outlook Anywhere

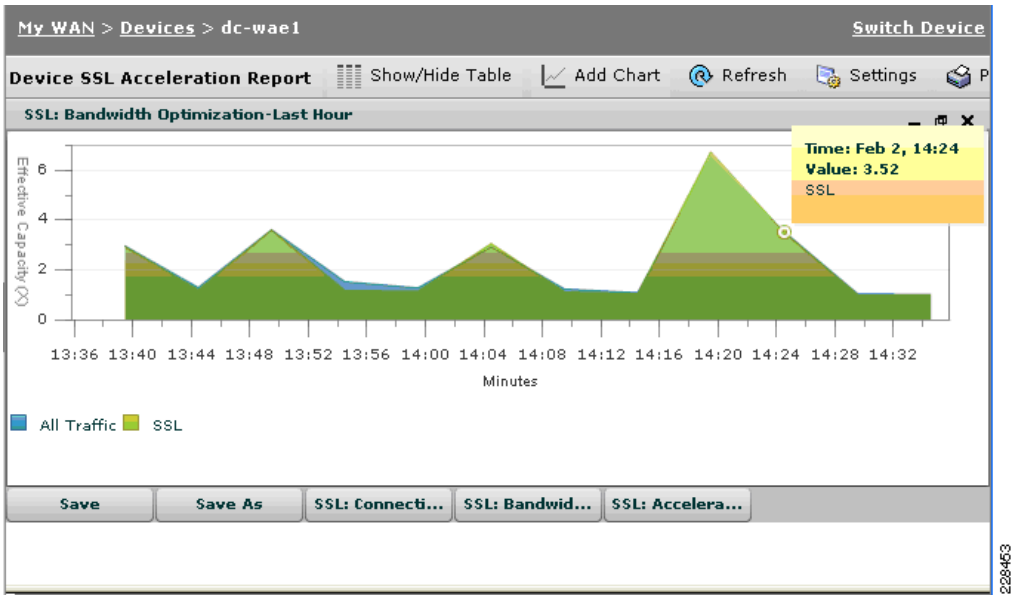

Table 8 shows the performance benefits for unencrypted RPC/MAPI Outlook client.

Table 8 RPC/MAPI Outlook Client Results

	% reduction in traffic	Effective link capacity
Cold	6.5%	1.06
Hot	91.3%	11.45
Warm	42.8	1.74
Subsequent Hot	88.2	8.46

As Table 9 shows, WAAS provides performance for reducing E-mail download times using Outlook MAPI/RPC (unencrypted) clients. Download times of an E-mail with a 7 MB attachment are significantly reduced once the cache is “hot” (where E-mail has been sent previously) or “warm” (in the case where the file attachment is modified and resent). The Regional Office 2 profile was used for this test.

Table 9 E-mail Download Times in Seconds

	No WAAS	Cold Cache	HotCache	Warm Cache
Outlook MAPI/RPC	40	40	14	25

MAPI/RPC Branch Users—Load Test

Exchange LoadGen 2010 was used to simulate 50 simultaneous MAPI/RPC Outlook 2007 users (online mode) at the branch. The load profile that was selected with the heavy load profile, which simulates 100 sent/received 75KB messages a day. The test was configured to run for a total of four hours, with a simulation time of three hours to give about 30 minutes of ramp up time at the beginning and 30 minutes

of ramp down time. Stress mode was chosen to increase the number of transactions generated by each user. Throughout the duration of the four hour test, screenshots of the performance results reported by WAAS Central Manager were taken. The following test results were captured.

The aggregate traffic from the 50 MAPI/RPC Outlook users simulates a variety of Outlook tasks, such as browsing calendars, making appointments, sending/receiving E-mails, and browsing contacts.

The following graphs were taken from WAAS Central Manager about half-way into the test at around two hours. As Figure 86 shows, a bit more than 221,000 tasks had been completed at that point.

Figure 86 Number of Tasks Completed After Two Hours

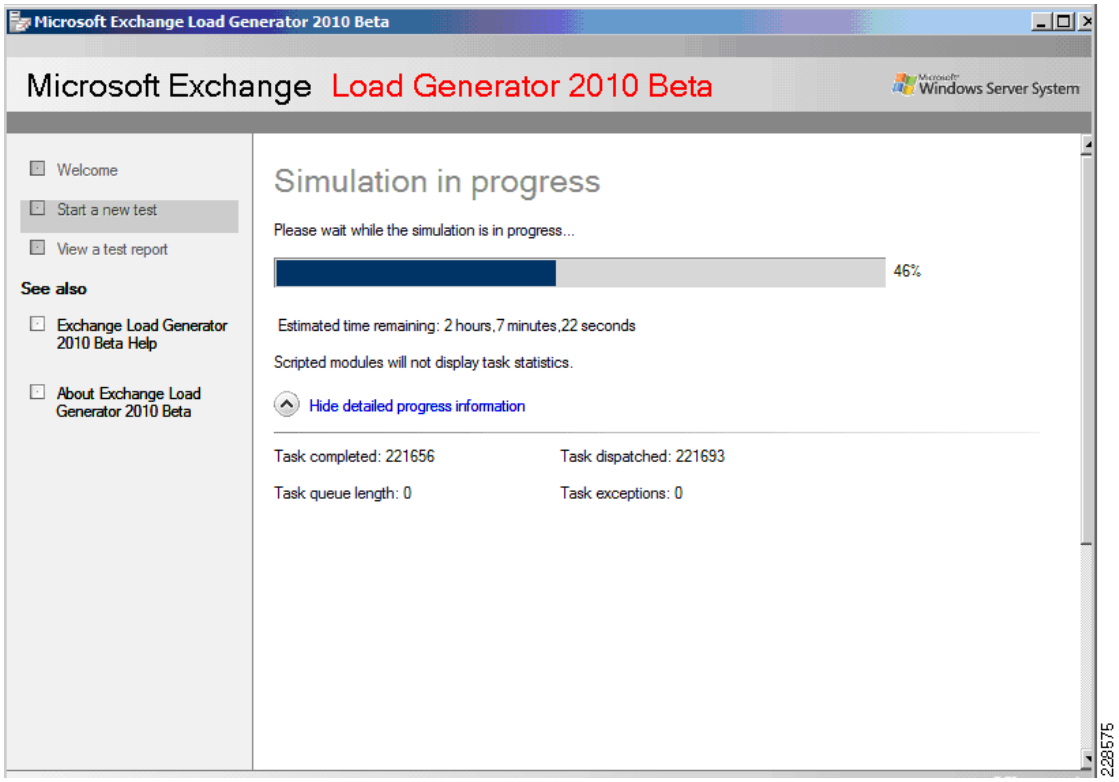

Figure 87 shows significant traffic reduction during the second hour of the four hour test, when ramp up has completed and all 50 users are fully active. The average percentage of volume reduction is around 75%.

Figure 87 Traffic Reduction During the Second Hour for 50 Simultaenous MAPI/RPC Branch Users

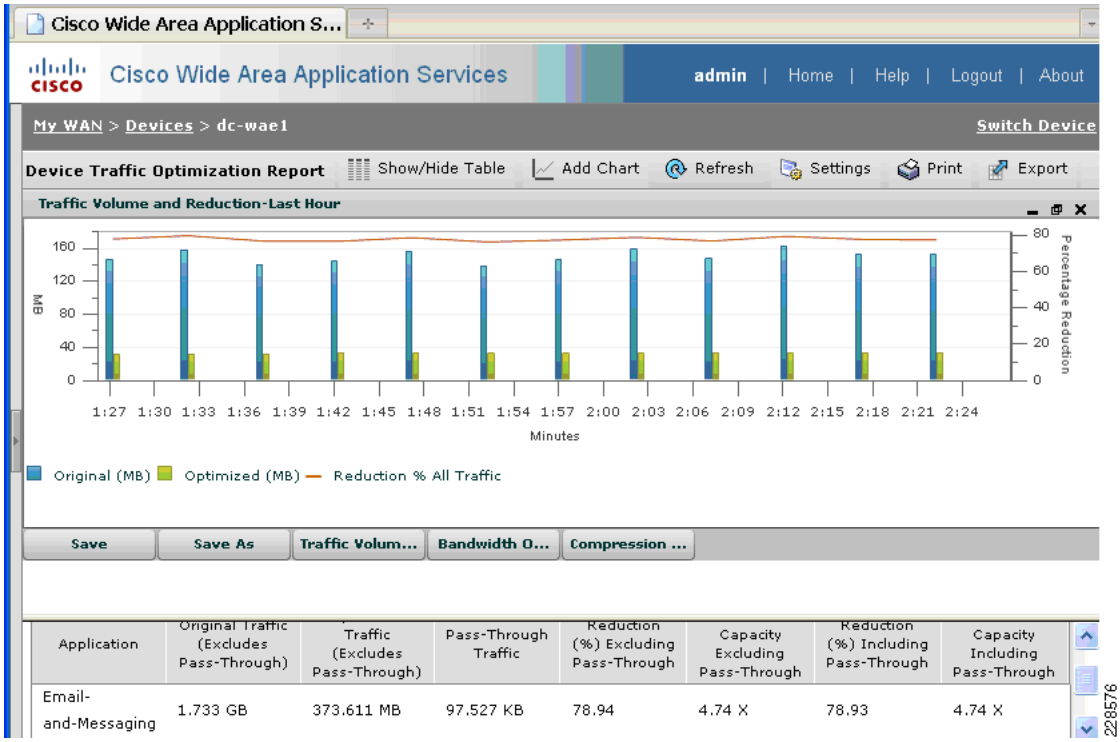

Figure 88 shows the compression ratio during this second hour. Note that the compression ratio stays around 75% for the remainder of the test.

Figure 88 Compression Ratio During Second Hour of Test

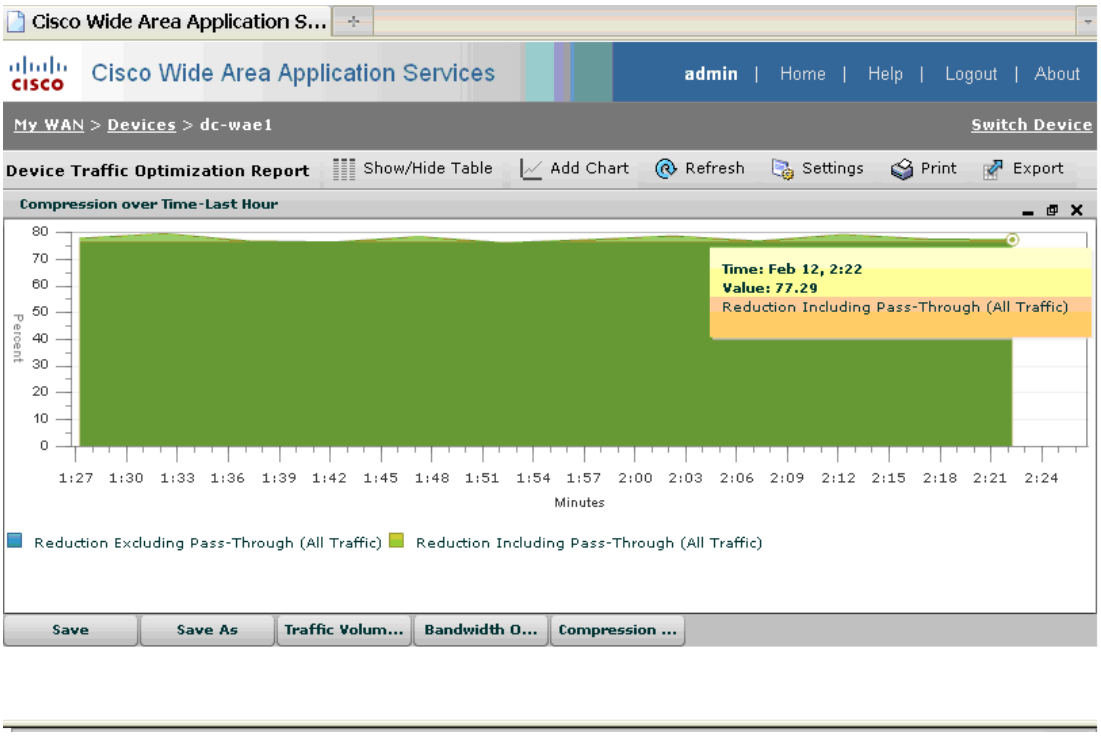

Figure 89 shows the average effective capacity during the second hour of the test is around 4.4. The effective capacity stays at this average for the remainder of the four hour test.

Figure 89 Effective Capacity Due to Traffic Optimization and Acceleration During Second Hour of Test

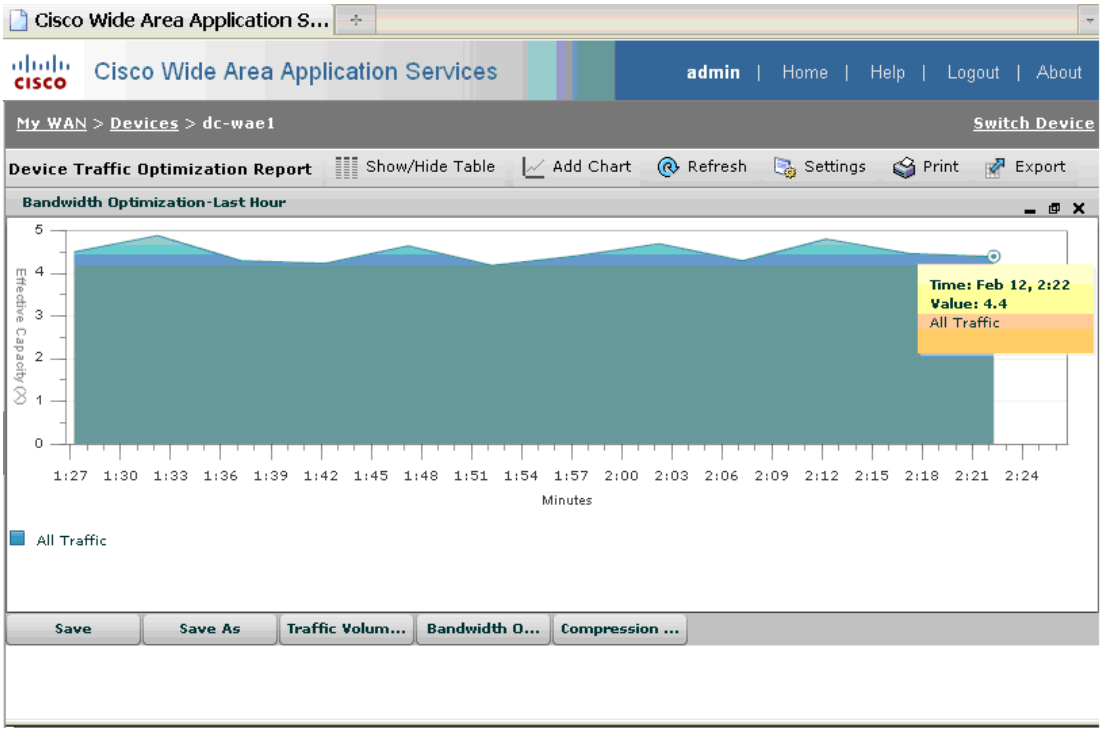

Figure 90 shows that MAPI acceleration contributed to the 75% traffic reduction seen in this test.

Figure 90 MAPI Acceleration of E-mail Traffic

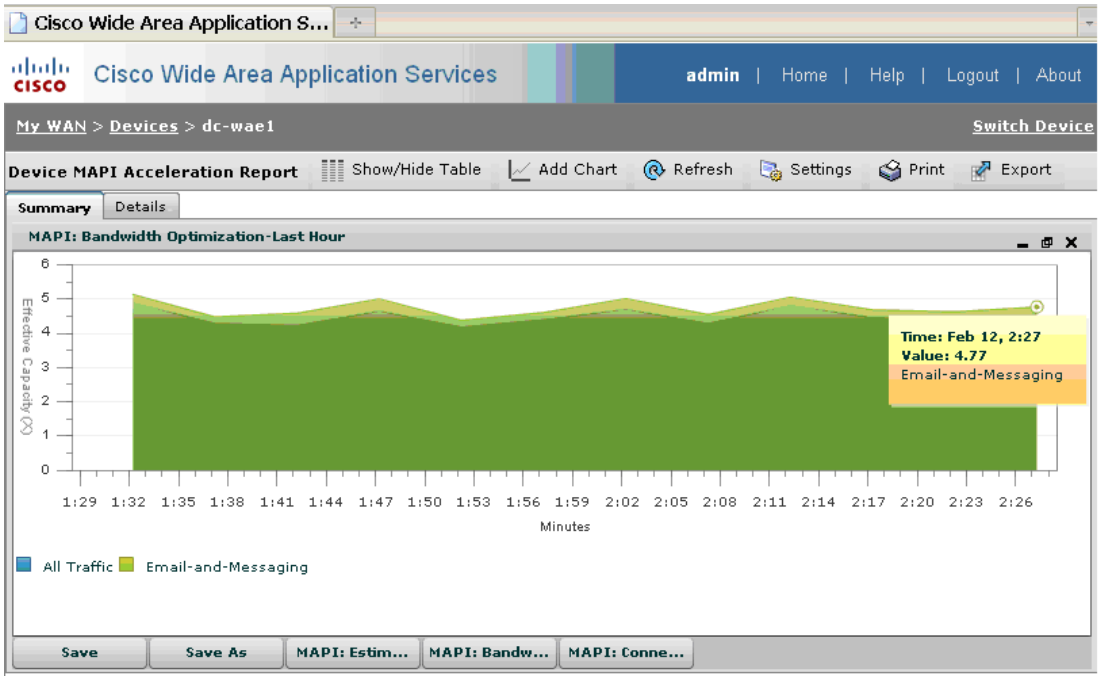

At the end of four hours, a total of 452,283 tasks were completed as shown in Figure 91. The distribution of tasks during the four hours for the 50 users is as shown in Table 10.

Figure 91 Total Tasks Completed at End of Test

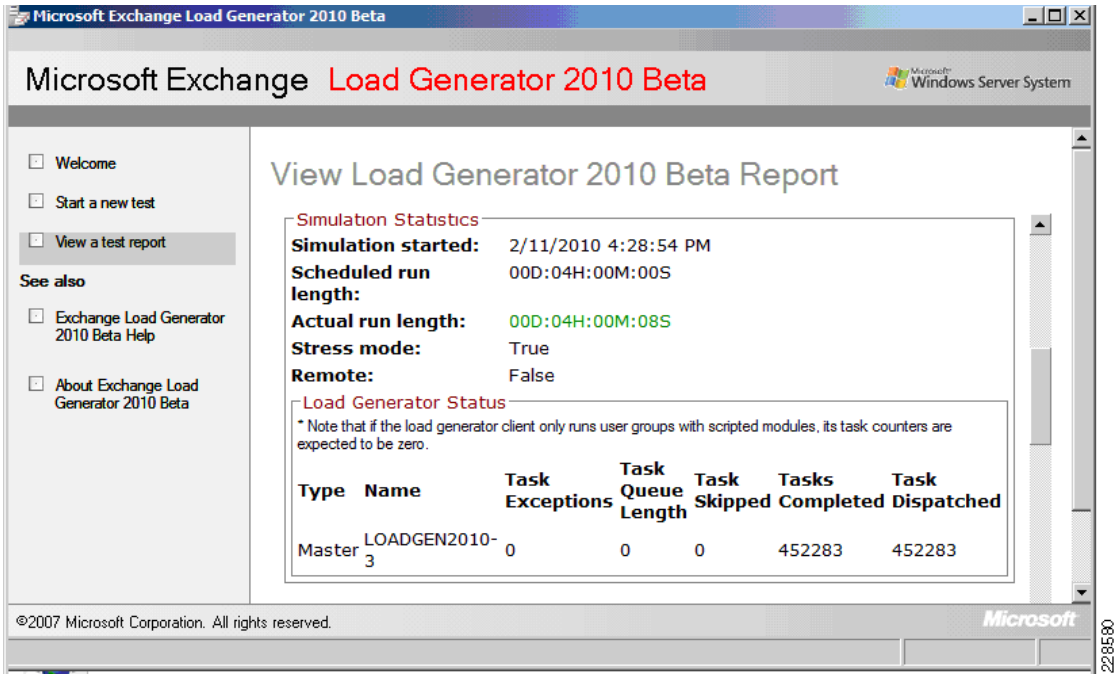

Table 10 List of Tasks Simulated by LoadGen

Task Name	Count	Actual Distribution (%)	Configured Distribution (%)
BrowseCalendarTask	41139	9	9
BrowseContactsTask	34231	7	7
BrowsePublicFolderTask	0	0	0
BrowseTasksTask	0	0	0
CreateContactTask	3398	0	0
CreateFolderTask	0	0	0
CreateTaskTask	3491	0	0
DeleteMailTask	0	0	0
EditRulesTask	0	0	0
EditSmartFolderTask	0	0	0
ExportMailTask	0	0	0
InitializeMailboxTask	0	0	0
LogoffTask	10152	2	2
LogonTask	0	0	0
MakeAppointmentTask	6846	1	1
ModuleInitTask	1	0	0

Table 10 List of Tasks Simulated by LoadGen

Task Name	Count	Actual Distribution (%)	Configured Distribution (%)
MoveMailTask	0	0	0
PostFreeBusyTask	14030	3	3
PublicFolderPostTask	0	0	0
ReadAndProcessMessagesTask	273751	60	60
RequestMeetingTask	6880	1	1
SearchTask	0	0	0
SendMailTask	54915	12	12

Remote Users

Performance numbers for the download of a 7 MB E-mail attachment by a remote user were measured using a combination of the Client Manager panel on the WAAS Mobile Client and the reports available on WAAS Mobile Server. The two client types Outlook Anywhere and Outlook Web Access were tested. The Remote User link profile was used.

Table 11 E-mail Attachment Download Times in Seconds

	No WAAS	Cold	Hot
Outlook Web Access	63	42	8
RPC/MAPI[1]	44	35	11

1. Native MAPI encryption disabled.

Figure 92 from the WAAS Mobile Server was captured after a cold file download by an Outlook Anywhere client followed by a hot file download of the same E-mail attachment, with some additional background traffic captured. As can be seen, Outlook Anywhere HTTPS traffic with a hot cache is reduced by 75% or 3X-4X effective capacity.

Figure 92 Compression Summary on WAAS Mobile Server for Outlook Anywhere

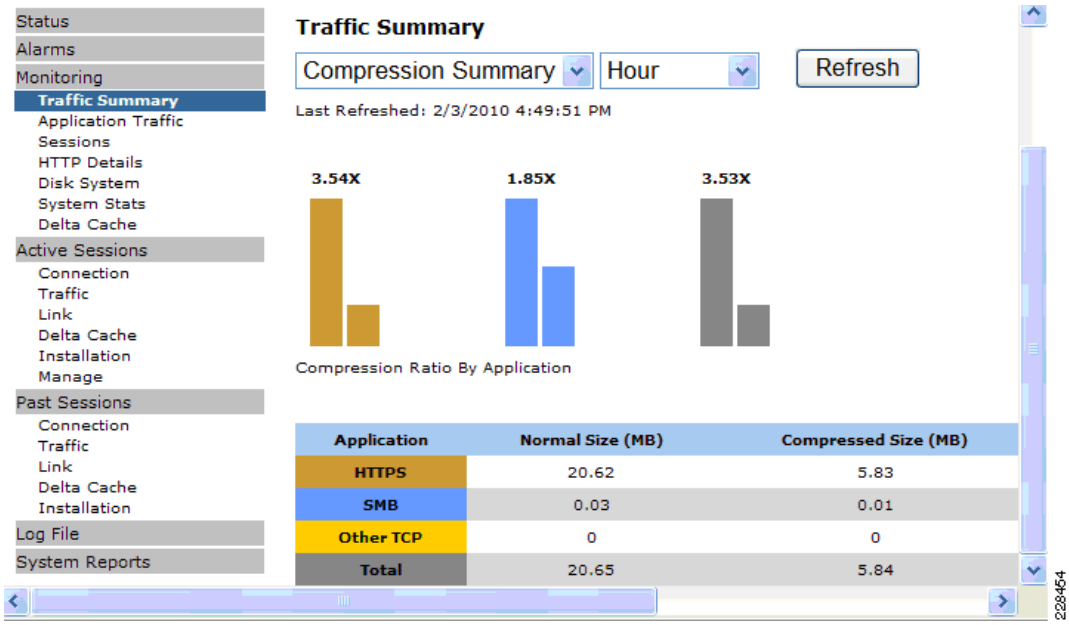

Similar numbers are given on the WAAS Client Manager panel of the Outlook Anywhere PC, as Figure 93 shows.

Figure 93 Compression Ratio on Client Manager for Outlook Anywhere

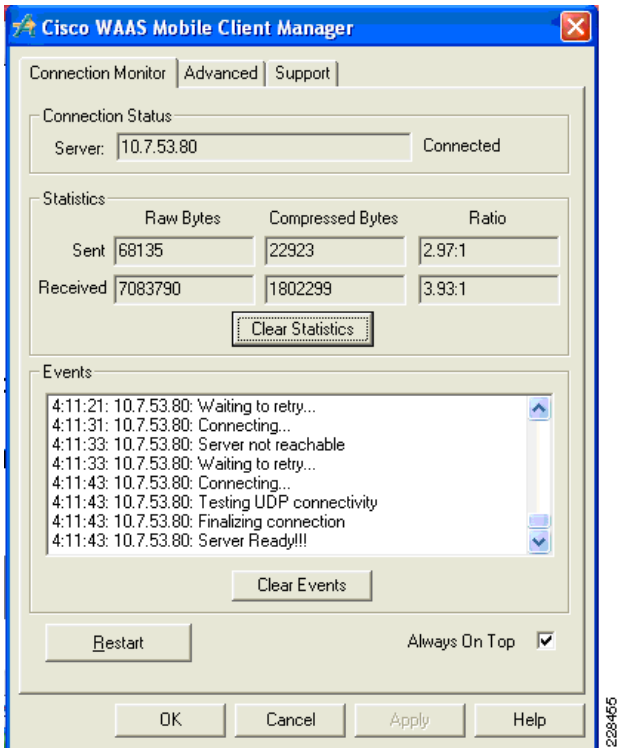

Cumulatively, after the consecutive E-mail attachment downloads have been run for the previous tests, the cumulative compression summary (which includes non-E-mail background traffic to the Exchange server) for Outlook Anywhere shows an 8X compression ratio. In that same graph, the cold and hot E-mail attachment downloads for RPC/MAPI client resulted in a cumulative compression ratio of 2X-3X (see Figure 94).

Figure 94 Compression Summary for Outlook Anywhere and RPC/MAPI Client

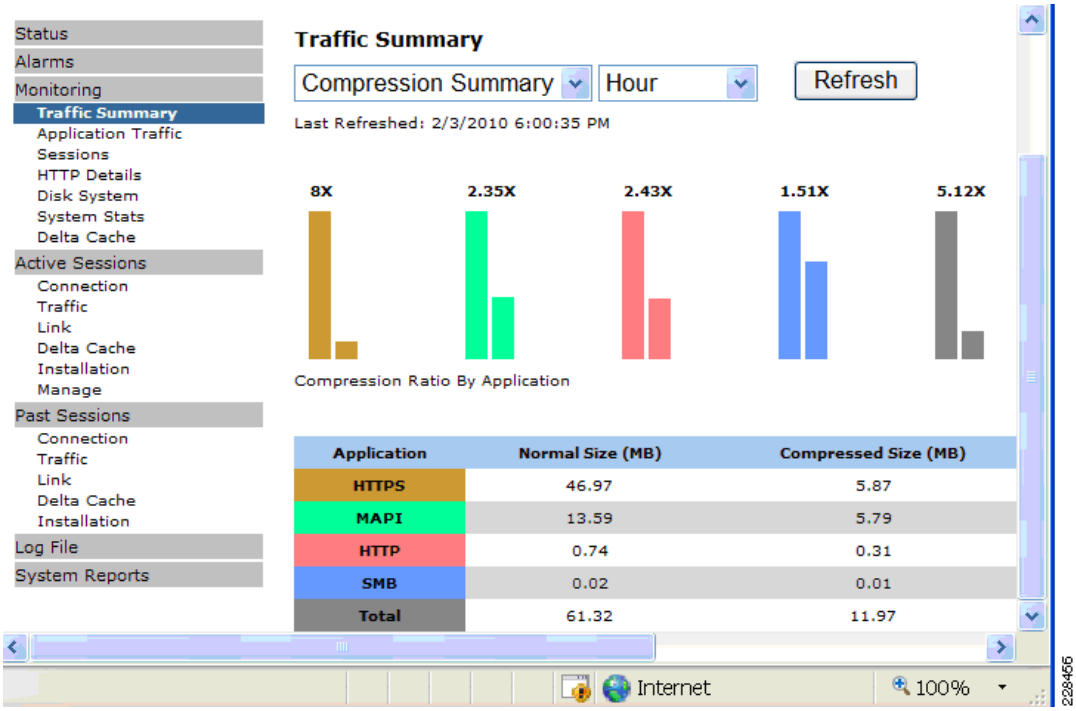

Bill of Materials

Solution Components

Platform	Version
ASR1002 (RP1)	12.2(33)XNB1
Nexus 7000	4.1(4)
C6509-E	12.2(18)
ACE10-6500-K9	A2(3.0)
C3825	12.4(11)XJ
WAE-7371	4.1.5.a.4
NM-WAE-502	4.1.5.a.4
WAAS Mobile Server	3.4.2

Platform	Version
WAAS Mobile Client	3.4.2
WAAS Central Manager (enabled on WAE-502-K9)	4.1.5.a.4
UCS 6120XP UCS 5108 Blade Server Chassis UCS 2104XP Fabric Extender UCS B200 M1 Blade Server	1.0.2(d)
MDS 9509	3.2.2(c)
Fabric Manager	4.1.3(b)
Windows 2008 Server R2 with Hyper-V	DataCenter and Enterprise editions, x64
Exchange 2010 Release Candidate, x64	en_exchange_server_2010_rc_x64_417957
Outlook 2010- Technical Preview	O2010_ProfessionalVL_volume_ship_x86_en-us
Windows XP Professional with SP3	en_windows_xp_professional_with_service_pack_3_x86_cd_x14-80428
System Center Operations Manager 2007 R2 Virtual Machine Manager	

References

In addition to the references cited in the document, there is additional information on the data center architecture and data center virtualization available in the following references:

- Data Center Service Patterns
  http://www.cisco.com/en/US/docs/solutions/Enterprise/Data_Center/DC_3_0/dc_serv_pat.html#wp1037942
- Security and Virtualization in the Data Center
  http://www.cisco.com/en/US/docs/solutions/Enterprise/Data_Center/DC_3_0/dc_serv_pat.html#wp1037942

Appendix A—WAE Device CLI Statistics

WAAS Between Data Center and Branch

Outlook Anywhere

The following **show stat conn optimized** output from the command line of the WAAS appliance at the data center edge shows that the WAAS appliances are able to apply SSL acceleration, LZ compression, and DRE caching to the multiple SSL connections generated from each Outlook Anywhere client.

```
D:DRE,L:LZ,T:TCP Optimization RR:Total Reduction Ratio
A:AOIM,C:CIFS,E:EPM,G:GENERIC,H:HTTP,M:MAPI,N:NFS,S:SSL,V:VIDEO
```

```
ConnID        Source IP:Port          Dest IP:Port            PeerID Accel RR
205166      10.7.10.102:3843       10.7.53.55:443 00:21:55:86:1e:77 TSDL  21.9%
205167      10.7.10.102:3845       10.7.53.55:443 00:21:55:86:1e:77 TSDL  23.4%
205168      10.7.10.102:3846       10.7.53.55:443 00:21:55:86:1e:77 TSDL  14.7%
205169      10.7.10.102:3847       10.7.53.55:443 00:21:55:86:1e:77 TSDL  16.0%
205172      10.7.10.102:3851       10.7.53.55:443 00:21:55:86:1e:77 TSDL  36.9%
205173      10.7.10.102:3853       10.7.53.55:443 00:21:55:86:1e:77 TSDL  42.2%
205174      10.7.10.102:3854       10.7.53.55:443 00:21:55:86:1e:77 TSDL  21.3%
205175      10.7.10.102:3855       10.7.53.55:443 00:21:55:86:1e:77 TSDL  31.3%
205178      10.7.10.112:1347       10.7.53.55:443 00:21:55:86:1e:77 TSDL  21.7%
205179      10.7.10.112:1349       10.7.53.55:443 00:21:55:86:1e:77 TSDL  21.9%
205180      10.7.10.112:1350       10.7.53.55:443 00:21:55:86:1e:77 TSDL  14.4%
205181      10.7.10.112:1351       10.7.53.55:443 00:21:55:86:1e:77 TSDL  16.0%
205186      10.7.10.112:1355       10.7.53.55:443 00:21:55:86:1e:77 TSDL  32.8%
205187      10.7.10.112:1357       10.7.53.55:443 00:21:55:86:1e:77 TSDL  46.7%
205188      10.7.10.112:1358       10.7.53.55:443 00:21:55:86:1e:77 TSDL  21.3%
205189      10.7.10.112:1359       10.7.53.55:443 00:21:55:86:1e:77 TSDL  34.1%
205192      10.7.10.111:1305       10.7.53.55:443 00:21:55:86:1e:77 TSDL  21.8%
205193      10.7.10.111:1307       10.7.53.55:443 00:21:55:86:1e:77 TSDL  22.2%
205194      10.7.10.111:1308       10.7.53.55:443 00:21:55:86:1e:77 TSDL  14.9%
205195      10.7.10.111:1309       10.7.53.55:443 00:21:55:86:1e:77 TSDL  16.1%
205198      10.7.10.111:1313       10.7.53.55:443 00:21:55:86:1e:77 TSDL  33.1%
205200      10.7.10.111:1315       10.7.53.55:443 00:21:55:86:1e:77 TSDL  47.4%
205204      10.7.10.111:1316       10.7.53.55:443 00:21:55:86:1e:77 TSDL  21.0%
205205      10.7.10.111:1317       10.7.53.55:443 00:21:55:86:1e:77 TSDL  34.6%
205208      10.7.10.110:2491       10.7.53.55:443 00:21:55:86:1e:77 TSDL  20.3%
205210      10.7.10.110:2493       10.7.53.55:443 00:21:55:86:1e:77 TSDL  17.5%
205211      10.7.10.110:2494       10.7.53.55:443 00:21:55:86:1e:77 TSDL  14.7%
205212      10.7.10.110:2495       10.7.53.55:443 00:21:55:86:1e:77 TSDL  16.2%
205215      10.7.10.110:2499       10.7.53.55:443 00:21:55:86:1e:77 TSDL  33.9%
205216      10.7.10.110:2501       10.7.53.55:443 00:21:55:86:1e:77 TSDL  45.2%
205218      10.7.10.110:2502       10.7.53.55:443 00:21:55:86:1e:77 TSDL  20.9%
205219      10.7.10.110:2503       10.7.53.55:443 00:21:55:86:1e:77 TSDL  34.4%
205222      10.7.10.110:2507       10.7.53.55:443 00:21:55:86:1e:77 TSDL  27.5%
```

After the Exchange user at the data center campus has sent the 7 MB E-mail attachment to all four users in the branch, the first user opens up the E-mail attachment and experiences some minor optimization benefits from the LZ compression. Negligible benefits are reaped from the DRE caching since this is the first download of this file and the cache is still cold.

After the remaining three users download the E-mail attachment, it is clear the DRE cache has been leveraged through the **show stat dre** output. 3142 messages have been processed through the DRE cache, where 64.89% of those messages were already in the cache.

```
dc-wae1#show stat dre
;
;
Connections:   Total (cumulative): 68   Active: 33

Encode:
   Overall: msg:       3213, in:  27729 KB, out:   8930 KB, ratio:  67.79%
       DRE: msg:       3142, in:  27727 KB, out:   9736 KB, ratio:  64.89%
DRE Bypass: msg:         71, in:   1604 B
        LZ: msg:       2719, in:   6562 KB, out:   5755 KB, ratio:  12.31%
```

Outlook MAPI/RPC (Unencrypted)

The following **show stat conn optimized** output from the command line of the WAAS appliance at the data center edge shows that the WAAS appliances are able to apply MAPI acceleration, LZ compression, and DRE caching to the MAPI connections generated from each Outlook RPC client.

```
dc-wae1#show stat conn optimized
 Current Active Optimized Flows:                       6
   Current Active Optimized TCP Plus Flows:          6
   Current Active Optimized TCP Only Flows:          0
   Current Active Optimized TCP Preposition Flows:   0
Current Active Auto-Discovery Flows:                 24
Current Reserved Flows:                              108
Current Active Pass-Through Flows:                   0
Historical Flows:                                    519

D:DRE,L:LZ,T:TCP Optimization RR:Total Reduction Ratio
A:AOIM,C:CIFS,E:EPM,G:GENERIC,H:HTTP,M:MAPI,N:NFS,S:SSL,V:VIDEO

ConnID         Source IP:Port           Dest IP:Port             PeerID Accel RR
205508      10.7.10.102:4011      10.7.53.55:34052 00:21:55:86:1e:77 TMDL  07.1%
205514      10.7.10.112:1465      10.7.53.55:34052 00:21:55:86:1e:77 TMDL  92.8%
205522      10.7.10.111:1422      10.7.53.55:34052 00:21:55:86:1e:77 TMDL  92.8%
205528      10.7.10.110:2690      10.7.53.55:34052 00:21:55:86:1e:77 TMDL  92.8%
205546      10.7.10.102:4024        10.7.53.10:139 00:21:55:86:1e:77 TGDL  37.6%
205548       10.7.12.2:56445        10.7.53.5:7878 00:21:55:86:1e:77 TL    77.1%
```

After the Exchange user at the data center campus has sent the 7 MB E-mail attachment to all four users in the branch, the first user opens up the E-mail attachment and experiences some minor optimization benefits from the LZ compression. Negligible benefits are reaped from the DRE caching since this is the first download of this file and the cache is still cold.

After the remaining three users download the E-mail attachment, it is clear the DRE cache has been leveraged through the **show stat dre** output. 12428 messages have been processed through the DRE cache, where 72.92% of those messages were already in the cache.

```
dc-wae1#show stat dre
 Connections:   Total (cumulative): 61   Active: 6

Encode:
   Overall: msg:      13796, in:  52556 KB, out:  13083 KB, ratio:  75.11%
       DRE: msg:      12428, in:  52368 KB, out:  14179 KB, ratio:  72.92%
DRE Bypass: msg:      10351, in:    188 KB
        LZ: msg:      10650, in:  10936 KB, out:   9589 KB, ratio:  12.32%
 LZ Bypass: msg:       3146, in:   3431 KB
    Avg latency:      0.238 ms     Delayed msg:       9605
  Encode th-put:  16009 KB/s
  Message size distribution:
    0-1K=11%  1K-5K=29%  5K-15K=58%  15K-25K=0%  25K-40K=0%  >40K=0%
```

Appendix B—Hyper-V Logical Diagram

Figure 95 provides a high-level overview of the architecture of a Hyper-V environment running on Windows Server 2008.

Figure 95 Overview of Hyper-V Architecture

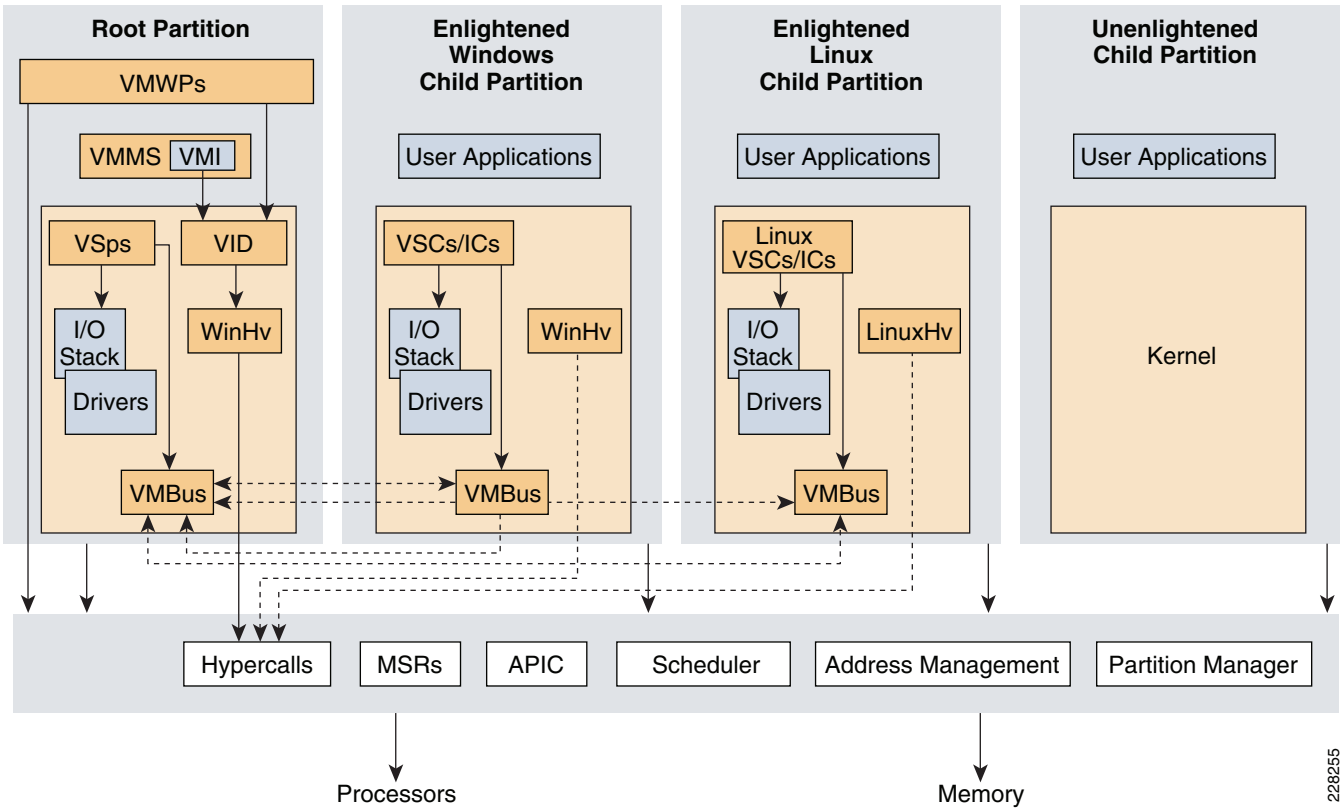

- APIC (Advanced Programmable Interrupt Controller)—A device which allows priority levels to be assigned to its interrupt outputs.
- Child partition—Partition that hosts a guest operating system - All access to physical memory and devices by a child partition is provided via the Virtual Machine Bus (VMBus) or the hypervisor.
- Hypercall— Interface for communication with the hypervisor. The hypercall interface accommodates access to the optimizations provided by the hypervisor.
- Hypervisor—A layer of software that sits between the hardware and one or more operating systems. Its primary job is to provide isolated execution environments called partitions. The hypervisor controls and arbitrates access to the underlying hardware.
- IC (integration component)—Component that allows child partitions to communication with other partitions and the hypervisor.
- I/O stack—Input/output stack
- MSR (Memory Service Routine)
- Root partition—Manages machine-level functions such as device drivers, power management, and device hot addition/removal. The root (or parent) partition is the only partition that has direct access to physical memory and devices.
- VID (Virtualization Infrastructure Driver)—Provides partition management services, virtual processor management services, and memory management services for partitions.
- VMBus—Channel-based communication mechanism used for inter-partition communication and device enumeration on systems with multiple active virtualized partitions. The VMBus is installed with Hyper-V Integration Services.

- VMMS (Virtual Machine Management Service)—Responsible for managing the state of all virtual machines in child partitions.
- VMWP (Virtual Machine Worker Process)—A user mode component of the virtualization stack. The worker process provides virtual machine management services from the Windows Server 2008 instance in the parent partition to the guest operating systems in the child partitions. The Virtual Machine Management Service spawns a separate worker process for each running virtual machine.
- VSC (Virtualization Service Client)—A synthetic device instance that resides in a child partition. VSCs utilize hardware resources that are provided by Virtualization Service Providers (VSPs) in the parent partition. They communicate with the corresponding VSPs in the parent partition over the VMBus to satisfy a child partitions device I/O requests.
- VSP (Virtualization Service Provider)—Resides in the root partition and provides synthetic device support to child partitions over the Virtual Machine Bus (VMBus).
- WinHv (Windows Hypervisor Interface Library)—WinHv is essentially a bridge between a partitioned operating system's drivers and the hypervisor which allows drivers to call the hypervisor using standard Windows calling conventions.
- WMI—The Virtual Machine Management Service exposes a set of Windows Management Instrumentation (WMI)-based APIs for managing and controlling virtual machines.

Appendix C—Device Configurations

ACE Exchange Virtual Context Configuration

```
access-list all line 10 extended permit ip any any
access-list all line 20 extended permit icmp any any


probe icmp PING
  interval 5
  passdetect interval 2
  passdetect count 1
probe tcp PROBE-TCP
  port 5007
  interval 2
  faildetect 2
  passdetect interval 10
  passdetect count 2
probe http http-probe
  interval 60
  passdetect interval 60
  passdetect count 2
  request method get url /exchweb/bin/auth/owalogon.asp
  expect status 400 404
probe https https-probe
  interval 60
  passdetect interval 60
  passdetect count 2
  request method get url /owa/auth/login.aspx
  expect status 400 404


rserver host CAS1
  ip address 10.7.53.55
  probe http-probe
```

```
  probe https-probe
  inservice
rserver host CAS2
  ip address 10.7.53.24
  probe http-probe
  probe https-probe
  inservice
rserver host EdgeTransport1
  ip address 10.7.53.22
  probe http-probe
  inservice
rserver host EdgeTransport2
  ip address 10.7.53.23
  probe http-probe
  inservice
rserver redirect SSLREDIRECT
  webhost-redirection https://aceexchange-vip.ucshypervroles.com/owa 302
  inservice

serverfarm host CAS-FARM
  predictor leastconns
  rserver CAS1
    inservice
  rserver CAS2
    inservice
serverfarm host CAS-FARM-80
  predictor leastconns
  rserver CAS1 80
    inservice
  rserver CAS2 80
    inservice
serverfarm host EdgeTransportFarm
  predictor leastconns
  probe PING
  probe PROBE-TCP
  rserver EdgeTransport1
    inservice
  rserver EdgeTransport2
    inservice
serverfarm redirect SSLREDIRECT
  rserver SSLREDIRECT
    inservice

parameter-map type http STICKY
  persistence-rebalance

sticky ip-netmask 255.255.255.255 address source CAS-IP
  replicate sticky
  serverfarm CAS-FARM
sticky http-cookie sessionid exchange-sticky-sessionid-grp
  timeout 20
  serverfarm CAS-FARM-80
sticky http-cookie Cookie OWA-STICKY
  cookie insert browser-expire
  timeout 60
  replicate sticky
  serverfarm CAS-FARM-80
sticky http-header Authorization CAS-RPC-HTTP
  serverfarm CAS-FARM-80

ssl-proxy service OWA
  key cisco-sample-key
  cert cisco-sample-cert
```

```
class-map match-any IMAPI-RPC
  2 match virtual-address 10.7.53.200 any
class-map match-all OWA-OUTLOOKAHYWHERE-SSL
  2 match virtual-address 10.7.53.200 tcp eq https
class-map match-all OWAREDIRECT
  2 match virtual-address 10.7.53.200 tcp eq www

policy-map type management first-match mgmt-pm
  class class-default
    permit

policy-map type loadbalance first-match IMAPI-RPC
  class class-default
    sticky-serverfarm CAS-IP
policy-map type loadbalance first-match OWA-OUTLOOKANYWHERE
  match OUTLOOK_ANYWHERE http header User-Agent header-value "MSRPC"
    sticky-serverfarm CAS-RPC-HTTP
  class class-default
    sticky-serverfarm OWA-STICKY
policy-map type loadbalance http first-match SSLREDIRECT
  class class-default
    serverfarm SSLREDIRECT


policy-map multi-match int53
  class OWAREDIRECT
    loadbalance vip inservice
    loadbalance policy SSLREDIRECT
  class OWA-OUTLOOKAHYWHERE-SSL
    loadbalance vip inservice
    loadbalance policy OWA-OUTLOOKANYWHERE
    loadbalance vip icmp-reply active
    nat dynamic 1 vlan 53
    ssl-proxy server OWA
  class IMAPI-RPC
    loadbalance vip inservice
    loadbalance policy IMAPI-RPC
    nat dynamic 1 vlan 53

interface vlan 53
  description to server-side vlan
  ip address 10.7.53.8 255.255.255.0
  alias 10.7.53.7 255.255.255.0
  peer ip address 10.7.53.9 255.255.255.0
  access-group input all
  nat-pool 1 10.7.53.200 10.7.53.200 netmask 255.255.255.0 pat
  service-policy input int53
  no shutdown

ip route 0.0.0.0 0.0.0.0 10.7.53.1
```

About Cisco Validated Design (CVD) Program

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