Integrating Microsoft Exchange Server 2007 in a Cisco Multisite Data Center Design

Cisco Validated Design

March 4, 2009

Contents

About the Document 1-3

Audience 1-3

Document Objectives 1-3

Document Format and Naming Conventions 1-3

Solution Overview 1-3

Solution Topology 1-4

Cisco Technology Overview 1-6

ACE Virtualization 1-6

Application Control Engine Global Site Selector 1-10

Cisco Content Network Registrar 1-11

Wide Area Application Engine 1-12

Microsoft Exchange Server 2007 Overview 1-13

Microsoft Exchange 2007 Server Roles 1-13

Microsoft Active Directory and Multisite Data Centers 1-16

Tested Microsoft Exchange Server 2007 Deployment Models 1-20

Microsoft Exchange Server 2007 Layout 1-20

Single-Site AD with Stretched CCR 1-21

Multisite Active Directory—Local CCR + Remote SCR 1-32

Optimization and Availability Support for Microsoft Exchange Server 2007 in a Cisco Multisite Data Center 1-37

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

Enterprise Network Architecture 1-38

- Data Center Network Components 1-38
- Front-End Network 1-38
  - Core Layer 1-39
  - Aggregation Layer 1-40
  - Access Layer 1-40
- Back-End Network 1-41
  - SAN Core Layer 1-41
  - SAN Edge Layer 1-41
- Branch Network Components 1-42
- Multisite Data Center Components 1-43

Design and Implementation Details 1-44

- Design Goals 1-44
- Enterprise Data Center Design 1-44
  - Site Selection 1-46
  - Route Health Injection 1-51
  - Layer 2 Extension 1-51
- Enterprise Edge Design 1-53
  - Client Access Server Role 1-55
  - Edge Server Role 1-72

Appendix 1-77

- ACE SSL Proxy Configuration 1-77
- Outlook Anywhere Configuration 1-79
  - Client Access Server (CAS) 1-79
  - Outlook Client 1-81

Cisco Validated Design 1-84

About the Document

This document provides design and configuration guidance for site and server load balancing, Secure Sockets Layer (SSL)- offload and WAN optimization in a Microsoft Exchange Server 2007 environment when it is deployed into a Cisco multisite data center architecture. An overview of the various Microsoft Exchange Server 2007 roles and operations will be given to provide the reader some context as to how the application environment is impacted in a multisite data center design.

Audience

This document is intended for network engineers and architects who need to understand both the basics of a Microsoft Exchange environment and the design and configuration options for providing advanced network services for Microsoft Exchange Server 2007.

Document Objectives

The objective of this document is to provide customers guidance on how to leverage a Cisco multisite data center design to support a Microsoft Exchange Server 2007 environment. The document is not meant to introduce the reader to basic Cisco data center design configurations nor is it meant to be a resource to learn the details of Microsoft Exchange Server 2007. The reader must be familiar with the basic Cisco data center concepts and products as well as the basics of Microsoft Exchange Server 2007 components, roles, and deployment scenarios as documented by Microsoft Corporation. The prerequisite knowledge can be acquired through many documents and training opportunities available both through Cisco and Microsoft. Below are a few recommended information resources that readers would find useful in these areas of interest:

Cisco Connection Online – Data Center:

http://www.cisco.com/go/dc

Cisco Solution Reference Network Designs (SRND):

http://www.cisco.com/go/srnd

Microsoft Exchange Server 2007:

http://www.microsoft.com/exchange/default.mspx

Document Format and Naming Conventions

User-defined properties such as access control list names and policy definitions are shown in ALL CAPS to assist the reader in understanding what is user-definable versus command specific. All commands are shown in `Courier` font. All commands that are applicable to the section covered will be in **`BOLD`**.

Solution Overview

The multisite solution described in this document equally applies across financial, manufacturing, consumer or information-based industries interested in constructing and deploying efficient and productive data centers. Data centers house the applications and information critical to the business, whatever that may be. Today, enterprises recognize that a data center is more than racks of compute power, but an asset with the potential to provide a competitive edge. As a result, industries are reevaluating their data center deployments with an interest to consolidate or expand where necessary to address the following:

- New infrastructure including network and compute resources (64-bit platforms, blade servers, switches, and routers)
- Regulatory compliance (typically resulting in expanded security and storage infrastructure)
- Facility space, power, and cooling to support new infrastructure
- New application environments and performance expectations
- Disaster recovery

The multisite solution described in this document focuses on the expectations of the application of four fundamental design goals:

- Application high availability
- Application scalability
- Data and application security
- Application performance

This document highlights network-based technologies used within and between data centers to achieve these objectives.

Solution Topology

Figure 1 depicts the Microsoft Exchange Server 2007 solution topology tested, where two distinct data centers (Data Center 1 and Data Center 2) are deployed leveraging Cisco's infrastructure design best practices. Note that each site provides local redundancy, scalability, and security for the applications it hosts. A multisite solution should simply extend the functionality of a single-site and should not compromise the integrity of either.

At each site in Figure 1, the hub and mailbox servers leverage the Layer 2 and 3 services of a well designed access and aggregation layer. The access and aggregation layers consist of the Cisco Catalyst 6500s with Sup720s. In the aggregation layer of each site, a pair of Cisco 7200 routers with NPE-G2s provide an L2TPv3 tunnel. This tunnel establishes Layer 2 adjacency between sites on a per-VLAN basis, efficiently meeting the requirements of our Exchange Server 2007 environment while controlling spanning tree domain creep. The L2TPv3 tunnel traverses the core layer, which is a high-speed Layer 3 fabric consisting of the Cisco Catalyst 6500s with Sup720s. The red lines indicate the use of 10 GigabitEthernet throughout the access, aggregation, and core layers.

Figure 1 defines two points of access into the data center for remote users via the WAN or the Internet. The remote branch users in the WAN benefit from the transparent and symmetric application optimization services of the Cisco Wide Area Application Services (WAAS). Cisco Wide Area Application Engines (WAEs) are located at each site and at the remote branch. Users originating from the Internet connect via a DMZ local to each data center site. The DMZ consists of Cisco Catalyst 6500s with Sup720s housing the Cisco Application Control Engine (ACE) service module, which provides application and security services. The Exchange edge and CAS roles reside in this location. In addition, the Internet edge houses a cluster of Cisco ACE Global Site Selectors (GSS), which monitor the state of each data center's Exchange application environment and uses this knowledge to provide intelligent selection between sites.

This document discusses each of the areas defined in Figure 1 to provide a better understanding of the application and the network deployed to support it.

Figure 1 Solution Topology

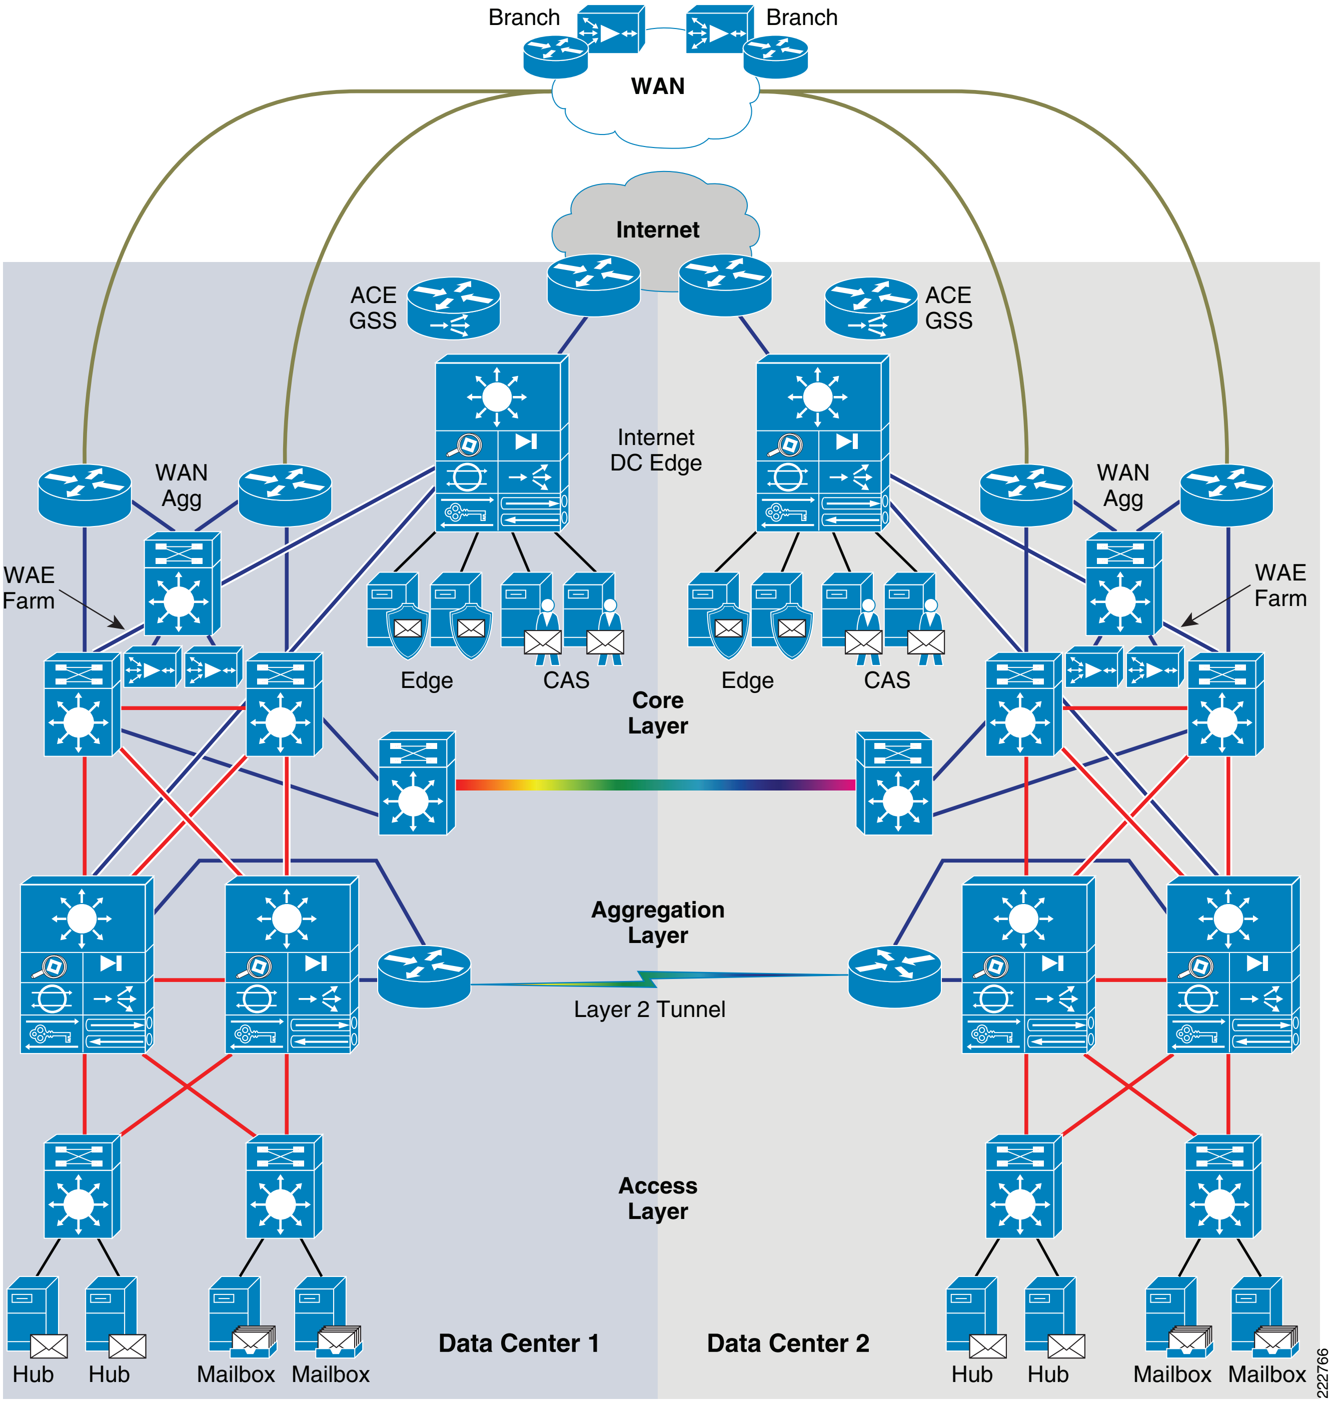

Cisco Technology Overview

This section provides an overview of the main Cisco products and technologies used in this design. The following products are addressed:

- Cisco Application Control Engine (ACE)
- Cisco ACE Global Site Selector (ACE GSS)
- Cisco Wide Area Application Engine (WAE)

The Cisco ACE provides a highly available and scalable data center solution from which the Microsoft Exchange Server 2007 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 345,000 L4 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, and so on
- Centralized role-based management via Application Network Manager (ANM) GUI or CLI
- SSL-offload (up to 15,000 SSL sessions via licensing)
- Support for redundant configurations (intra-chassis, inter-chassis, and inter-context)

The following sections describe some of the Cisco ACE features and functionalities used in the Microsoft Exchange Server 2007 application environment.

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.

Figure 2 shows the use of virtualized network services afforded via the Cisco ACE and Cisco Firewall Services Module (FWSM). In Figure 2, a Cisco Catalyst 6500 housing a single Cisco ACE and FWSM supports the business processes of five independent business units. The system administrator determines the application requirements and assigns the appropriate network services as virtual contexts. Each context contains its own set of policies, interfaces, resources, and administrators. The Cisco ACE and FWSMs allow routed, one-arm, and transparent contexts to co-exist on a single physical platform.

Figure 2 Service Chaining via Virtualized Network Services

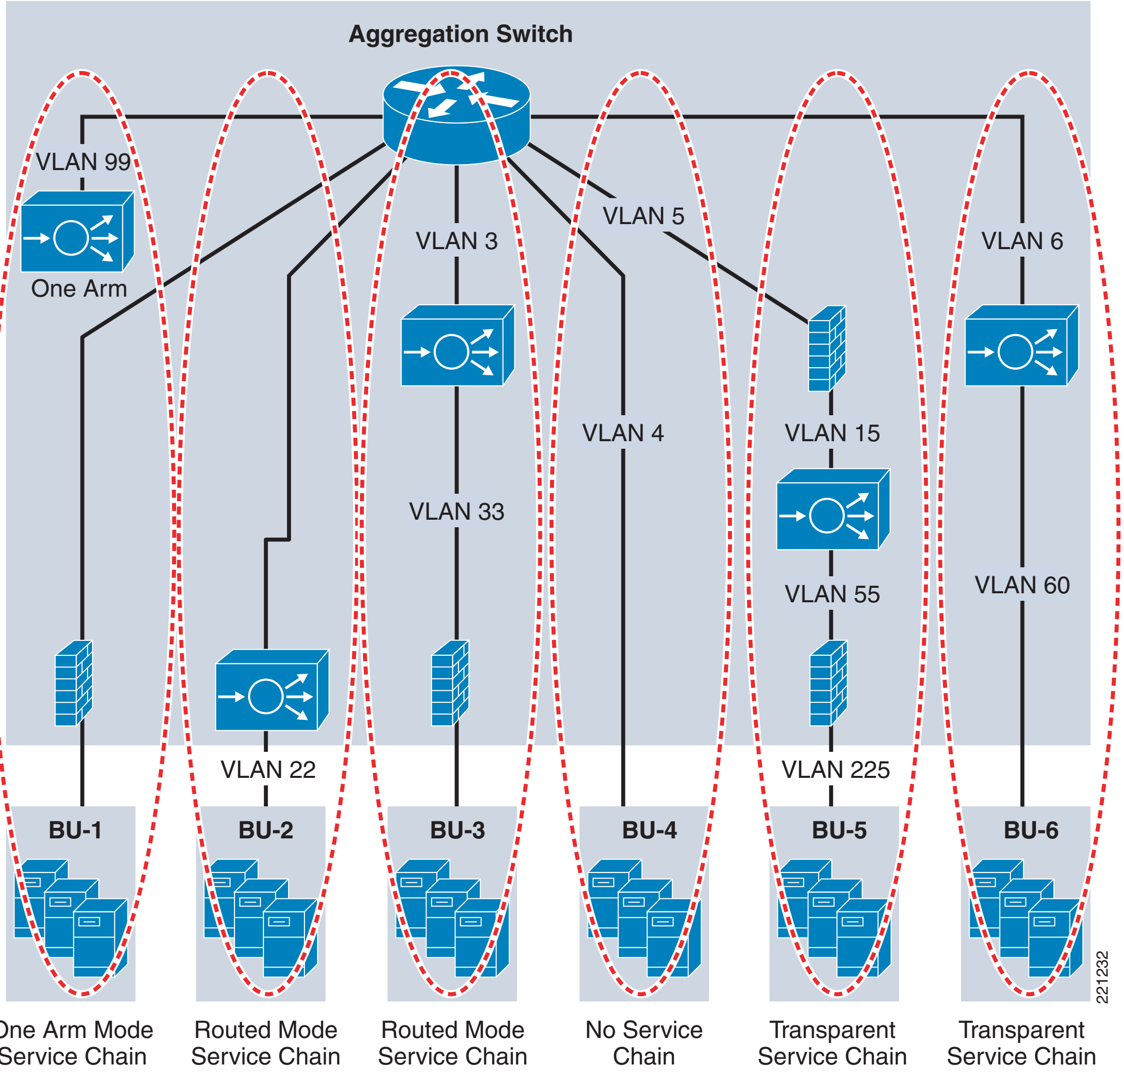

Note For more information on ACE virtualization, see the *Application Control Engine Module Virtualization Configuration Guide* at the following URL: http://www.cisco.com/en/US/products/hw/modules/ps2706/products_configuration_guide_book09186a00806882c6.html

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.

SSL Session ID Reuse

SSL session ID reuse allows the client and server to reuse the secret key negotiated during a previous SSL session. This feature generally improves the volume of SSL sessions that an SSL server or SSL proxy can effectively maintain. Clients residing with remote connectivity, for instance across a WAN, generally benefit from this feature. The SSL negotiation load is effectively reduced on the SSL proxy server while simultaneously improving the user experience as key negotiation is a rather lengthy process. The Cisco ACE may maintain the SSL session ID indefinitely or up to 20 hours with a timeout configuration.

It should be noted that SSL ID reuse does not compromise the security of the data center. The ID reuse feature only acknowledges that a secret key already exists between the client and server. Nonetheless the client must leverage this key for the application server to receive data from the client. The security resides in the secret key not the SSL session ID.

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.

The following example shows the **sessionid** cookie inserted into the client's Microsoft Exchange request via the **Set-Cookie** command from the server. It is also possible to insert cookies into the HTTP header via the Cisco ACE.

```
(Status-Line):HTTP/1.1 302 Moved Temporarily
Set-Cookie:aceoptimized=R3191602213; path=/
Location:http://owa.ese.cisco.com/owa/auth/logon.aspx?url=http://owa.ese.cisco.com/owa&rea
son=0
Set-Cookie:sessionid=; path=/; expires=Thu, 01-Jan-1970 00:00:00 GMT
Set-Cookie:cadata=; path=/; expires=Thu, 01-Jan-1970 00:00:00 GMT
Connection:close
Content-Length:0
```

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.

Allowed Server Connections

Enterprise data centers should perform due diligence on all deployed server and network devices, determining the performance capabilities to create a more deterministic, robust, and scalable application environment. The Cisco ACE allows the system administrator to establish the maximum number of active connections values on a per-server basis and/or globally to the server farm. This functionality protects the end device, whether it is an application server or network application optimization device such as the WAE.

Route Health Injection

Route Health Injection (RHI) allows the Cisco ACE to advertise host routes associated with any number of virtual IP addresses hosted by the device. The injection of the host route to the remaining network offers Layer 3 availability and convergence capabilities to the application environment.

KAL-AP UDP Agent

The Cisco ACE supports the KeepAlive-Appliance Protocol (KAL-AP) via a local UDP agent. This agent responds to KAL-AP queries from site selectors, such as the Cisco Global Site Selector, to provide the status and workload associated with one or more virtual IP addresses maintained by an ACE virtual context. The KAL-AP agent supports both domain and tagged formed queries. Tagged formed queries allow the verification of VIP state across NAT devices, such as firewalls or routers, and multiple ports for the same virtual IP address. This real-time information provides a more robust and accessible application as load and availability information may be leveraged to distribute traffic intelligently across multiple enterprise sites.

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.

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

- 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.

Application Control Engine Global Site Selector

Overview

The Cisco Application Control Engine Global Site Selector (Cisco ACE GSS) is an appliance that offers failover protection via Global Server Load Balancing (GSLB). The Cisco GSS device allows the enterprise to distribute and balance workload across multiple sites, providing the following benefits:

- Work-load distribution
- Disaster recovery and failover protection
- Improved user experience
- DNS offload

The Cisco GSS becomes part of the enterprise's DNS routing hierarchy as the authoritative DNS server for those services under its domain. The Cisco GSS intelligently resolves DNS requests with the additional knowledge of the site's availability and the associated application's state. This knowledge is gained from tight integration with load-balancers such as the Cisco Content Services Switch (CSS), Cisco Content Switch Module (CSM), and the Cisco ACE. Each of these load-balancers monitor the state of local application servers and communicate this information to the Cisco GSS where a global enterprise aware decision can be made. Currently, the Cisco GSS can support approximately 4,000 virtual IP addresses. The Cisco GSS includes the following factors prior to responding to a DNS request:

- Availability
- Proximity
- Load
- Source of the request (DNS proxy)
- Preference

Note The Cisco GSS device may also monitor individual servers, IOS SLB devices, DRP-enabled routers, Cisco's Local Director, and Cisco cache engines.

Keepalives

The Cisco GSS leverages keepalives to determine the state of a particular VIP under its domain. The Cisco GSS supports the following keepalive types:

- ICMP
- TCP
- HTTP HEAD
- KeepAlive-Appliance Protocol (KAL-AP)
- Scripted Keepalives
- Name Server

These keepalive types can be used individually or in a multiport group to determine the status of a virtual IP address. As a rule, the Cisco GSS does not respond to a DNS query with a VIP that has been declared inactive.

The KAL-AP keepalive is particularly useful when the Cisco network load-balancing technology is present. The Cisco GSS queries the load-balancer at each site for VIP state and load information. The detailed response received by the Cisco GSS from the network load-balancer can be used to distribute load efficiently across sites.

Note The keepalive timers may be adjusted to establish an acceptable failure window for the enterprise.

Cisco Content Network Registrar

The Cisco Content Network Registrar (CNR) is a separate process running on the GSS appliance that provides both DNS and DHCP support. As a full-featured DNS server, the CNR maintains the resource records (RR) within each enterprise DNS zone it supports. Mail Exchange (MX) resource records are of particular importance for an enterprise messaging application. MX records provide a list of hostnames providing mail exchange services within a domain. The CNR subsystem provides the MX functionality required for successful messaging.

Note For more information on the Cisco Content Network Registrar, refer to: http://www.cisco.com/en/US/products/sw/netmgtsw/ps1982/index.html

For more information on the Cisco Global Site Selector, refer to: http://www.cisco.com/en/US/products/hw/contnetw/ps4162/tsd_products_support_series_home.html

Wide Area Application Engine

To appreciate how the Cisco Wide Area Application Services (WAAS) provides WAN and application optimization 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 roundtrip messages may be few, and may 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 number of messages transmitted between endpoints. Chatty applications with many round-trips of application protocol messages that may or may not have small payloads.

The Cisco WAAS uses the technologies described in the following subsections to provide a number of features, including application acceleration, file caching, print service, and DHCP to benefit both types of applications.

Advanced Compression Using DRE and Lempel-Ziv 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.

Transport File Optimizations

The Cisco WAAS Transport File Optimizations (TFO) uses a robust TCP proxy to safely optimize TCP at the WAE device by applying TCP-compliant optimizations to shield the clients and servers from poor TCP behavior because of 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 implementing congestion management and recovery techniques to ensure that the maximum throughput is restored if there is packet loss.

Common Internet File System Caching Services

Common Internet file system (CIFS), used by Microsoft applications, is inherently a highly chatty transactional application protocol where it is not uncommon to find several hundred transaction messages traversing the WAN just to open a remote file. The Cisco WAAS provides a CIFS adapter that can inspect and to some extent predict what follow-up CIFS messages are expected. By doing this, the local WAE caches these messages and sends them locally, significantly reducing the number of CIFS messages traversing the WAN.

Print Services

The Cisco WAAS provides native SMB-based Microsoft print services locally on the WAE device. Along with CIFS optimizations, this allows for branch server consolidation at the data center. Having full-featured local print services means less traffic transiting the WAN. Without the Cisco WAAS print services, print jobs are sent from a branch client to the centralized server(s) across the WAN, and then back to the branch printer(s), thus transiting the WAN twice for a single job. The Cisco WAAS eliminates the need for either WAN trip.

Note For more information on these enhanced services, see the *Cisco Wide Area Application Services (WAAS) V4.0 Technical Overview* at the following URL:
http://www.cisco.com/en/US/products/ps6870/products_white_paper0900aecd8051d5b2.shtml

Microsoft Exchange Server 2007 Overview

The Microsoft Exchange Server 2007 offers many advantages to customers in the form of built-in protection, flexible access methods and operational efficiency. Customers are looking for ways to cut cost and increase productivity while ensuring that there is high availability. Microsoft Exchange Server 2007 was designed to offer solutions to these most demanding customer messaging requirements and do so for a variety of endpoints, from any location and to provide access to messaging resources in a secure and highly available manner.

Some of these customer requirements are met by enabling the following:

- Integrated message filtering
- Business continuance via several clustering and disaster recovery options
- Endpoint security for a variety of access methods which include a web client, Outlook, mobile, and POP/IMAP
- Flexible policy creation, management and reporting for legal compliance needs
- Streamlined setup, administration and management via the Microsoft Exchange Management Console, Exchange Management Shell, and Systems Center products
- Scalability and performance improvements through a x64-based architecture, increased memory support, and more intelligent message routing

There are many feature improvement and advantages of using Microsoft Exchange Server 2007 as well as comparisons with Microsoft Exchange Server 2003. Additional information on these features, advantages and comparisons can be found at:

http://www.microsoft.com/exchange/evaluation/default.mspx

Microsoft Exchange Server 2007 requires an existing Microsoft Active Directory (AD) deployment and leverages AD as a means to store and share information within the Exchange environment. More information regarding the planning and deployment of Microsoft Active Directory in support of Exchange Server 2007 can be found here: http://technet.microsoft.com/en-us/library/bb123715.aspx.

Note All references to Exchange Server 2007 used in testing imply the most up-to-date version of Exchange at time of validation, which is Exchange Server 2007 Service Pack 1 (SP1).

Microsoft Exchange 2007 Server Roles

There are five roles in Microsoft Exchange Server 2007. Each role serves a unique purpose within the Microsoft Exchange architecture and is flexible enough to be deployed in various sized organizations with varying requirements.

All roles (except Edge Transport) can be installed together on a single platform or can be deployed completely independent of one another. Small-medium customers can leverage the diverse number of Microsoft Exchange Server 2007 features while limiting the amount of hardware required for deployment by deploying the roles on the same server. Large organizations can leverage having multiple roles deployed in a redundant fashion on independent hardware platforms in geographically dispersed locations.

The five roles in Microsoft Exchange Server 2007 are:

- Client Access Server (CAS)
- Hub Transport (HT)
- Mailbox Server (MBX)
- Edge Transport (ET)
- Unified Messaging (UM)

The following sections will describe four of the five roles at a high-level and is not meant to be a full tutorial on the architecture, design, and operation of each role. The UM role is the only role that was not tested in the Cisco multisite data center design due to time constraints. A future version of this document will include the UM role in the Cisco multisite data center design. Detailed information on the Microsoft Exchange Server 2007 product, architecture, and design is found at:
http://www.microsoft.com/exchange or http://technet.microsoft.com/en-us/library/bb124558.aspx

Client Access Server

The client access server (CAS) provides access for a variety of client endpoints. The CAS role was formerly known as the Exchange front-end server. The CAS role supports access via the following methods:

- Microsoft Outlook Web Access (OWA)
- Post Office Protocol Version 3 (POP3)
- Internet Message Access Protocol Version 4 (IMAP4)
- Microsoft Exchange ActiveSync client
- Microsoft Outlook Anywhere

The CAS role also supports various other web services such as the offline address book (OAB) distribution and the autodiscover service. The list above shows that the CAS role can provide access to messaging services via many different endpoint types such as computers with web browsers, Outlook outside of the corporate firewall, email clients using POP3/IMAP4 and even mobile devices. Endpoints using another method of access such as Messaging Application Programming Interface (MAPI) most often connect directly to the mailbox server (MBX) role while within the corporate firewall (see Mailbox Server, page 15).

In the simplest terms, the CAS role provides a front-end service for the MBX role for non-MAPI connections. The CAS communicates directly with the MBX. The CAS role is optional if there are no requirements to use non-MAPI clients.

Microsoft recommends to deploy multiple CAS for performance, scalability, and availability purposes. The Microsoft Exchange Server 2007 fully supports multiple CAS role servers to be active simultaneously. This is ideal for an active/active multisite data center design.

Hub Transport Server

The Hub Transport (HT) role, formerly known as the Bridgehead server, is the central role for intelligent message routing delivery and policy control. Unlike the CAS and Edge Transport (ET) roles, the HT is required.

All mail flow external to the organization and internal within the organization is handled by the HT role. The HT role can use the ET as an SMTP relay for messages going to/from the Internet or it can handle the SMTP relay role on its own. The HT communicates directly with the MBX, other HT roles, and the ET.

Messaging routing within the Microsoft Exchange environment is requires the configuration of Active Directory (AD). AD is used to ensure that optimal message routing is accomplished within and between AD sites. This is quite different from previous Microsoft Exchange versions where routing groups were the primary method for messaging routing.

As was the case with the CAS role, it is recommended by Microsoft to deploy multiple HT roles for performance, scalability and availability purposes. Microsoft Exchange Server 2007 fully supports for the HT role to have multiple servers active simultaneously. This is ideal for an active/active multisite data center design.

Mailbox Server

The mailbox server (MBX) role is the database for all user messaging data. Users are homed to a particular MBX and associated storage group. As mentioned before, MAPI-based clients such as those running Microsoft Outlook connect directly to the MBX while within the corporate firewall. The MBX role is a required component of an Exchange Server 2007 deployment.

Microsoft Exchange Server 2007 has several options for maintaining high availability (HA) of the MBX role to include Local Continuous Replication (LCR), Cluster Continuous Replication (CCR), Standby Continuous Replication (SCR – Service Pack 1-only) and Single Copy Cluster (SCC). For more information on these solutions refer to the following URL:
http://technet.microsoft.com/en-us/library/bb124721.aspx

The two HA solutions that are discussed in this document are CCR and SCR. CCR is used to provide a two-node cluster of the mailbox role that allows for either automatic failover or manual failover of the cluster nodes. SCR allows for multiple standby nodes to pull mailbox logs from the primary MBX to provide disaster recovery (DR) and also mailbox database portability. SCR is a great choice for geographically dispersed data centers as well as for providing a way to replicate mailbox data to multiple data centers simultaneously. The two solutions (CCR and SCR) can be used together.

The MBX role is the only Exchange Server 2007 role that does not support an active/active configuration. However, the MBX role is also the only role that supports clusters. Therefore, more than one MBX can be deployed for scalability and availability purposes, but a user can only be connected to a single MBX that user is associated with. As will be discussed later on, if the primary MBX is unavailable, a standby server located within the same or different data center can take over the role.

The MBX communicates directly with the CAS, HT and, if deployed, the standby node in a clustered mailbox server (CMS).

Edge Transport Server

The Edge Transport (ET) role is used as a dedicated Internet SMTP relay as well as a means to provide message hygiene. The ET can be used to filter messages (SPAM) and also provide virus protection at the initial ingress point of the messaging system.

The ET leverages Active Directory but in a different way than the other roles. Active Directory Application Mode (ADAM) is used to store recipient information for the exchange organization so that the ET can know for which users it can accept mail. Since the ET is deployed at the network edge, it should be deployed as securely as possible. In an effort to secure the internal AD information, the ET has a one-way connection with the internal HT roles and uses an EdgeSync subscription as a method to replicate internal AD information with the ADAM instance running on each ET. This allows recipient information to be stored on the ET for mail acceptance purposes without exposing the internal AD topology and objects to an attacker if the server is compromised. Microsoft recommends that a "perimeter" AD environment be deployed to help facilitate the management of common policies and operations for the ET roles.

Microsoft recommends deploying multiple ET roles for performance, scalability and availability purposes. Microsoft Exchange Server 2007 fully supports for the ET role to have multiple servers active simultaneously. This is ideal for an active/active multisite data center design.

Figure 3 shows a high-level view of the four tested Microsoft Exchange 2007 Server roles and a basic traffic flow between each role.

Figure 3 High-level view of Microsoft Exchange Server 2007 roles

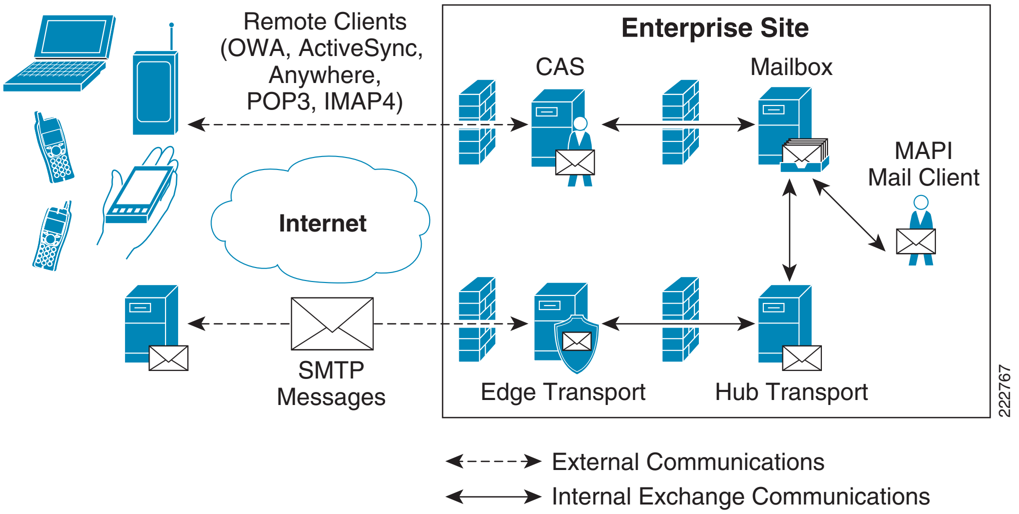

Microsoft Active Directory and Multisite Data Centers

As mentioned before, Microsoft Active Directory plays a critical and required role in the Microsoft Exchange Server 2007 environment. In the testing conducted by Cisco, there were two AD deployment options that were used between data centers. The first was using a single AD site for two active data center locations and the second was using an AD site for each data center location by using the Microsoft Active Directory Sites and Services capability to create and manage AD replication between sites.

Note All designs and references in this document are based on using Microsoft Windows Server 2003 R2 SP2. Microsoft Exchange Server 2007 with SP1 supports the upcoming release of Microsoft Windows Server 2008. However, at the time of publication of this document, Windows Server 2008 is not shipping. Future updates to this document will include the architectural changes to the presented designs when Windows Server 2008 is commercially available and has been tested by Cisco.

Single AD Site — Multiple Data Center Locations

There are many things to consider in a "normal" AD deployment model that will determine the success or failure of a scalable and available AD implementation. Adding the additional issues involved with now spanning a single AD site to multiple physical locations that can be geographically dispersed by great distance may be too great for many customers to undergo. Some, but certainly not all, of the considerations that a customer needs to account for are:

- Available network bandwidth and latency between each data center
- Suitable AD replication schedule between domain controllers/global catalog servers
- Contention between AD replication and other application/network traffic between data centers
- Containment of AD objects to a local site for management and security purposes

The considerations listed above will most often dictate that the data centers are close enough to each other to provide adequate bandwidth and low latency.

Note This document is not intended to provide the required knowledge for AD planning and implementation for Microsoft Exchange Server 2007. Additional information related to AD requirements for Exchange Server 2007 can be found at: http://technet.microsoft.com/en-us/library/bb123715.aspx.

The single AD site model was used and tested as it was the best model to allow for nearly all Microsoft Exchange Server 2007 components to function in an active/active role. As mentioned before, the mailbox server role is the only role that cannot support load balancing and/or active/active configurations. The CAS, HT and ET roles can support an active/active data center deployment. The reader must research and understand the Microsoft AD and Exchange Server 2007 implications of such a design before considering it for deployment.

Figure 4 shows a high-level overview of the single AD site model as it resides within two active data centers. The dashed box indicates that both DC locations are within the same single AD site. The only role in this scenario that cannot support an active/active configuration is the mailbox server role. In this example, the Microsoft Exchange Server 2007 Continuous Cluster Replication (CCR) feature is used to cluster the mailbox server role with the active Clustered Mailbox Server (CMS) in the primary DC and the standby CMS in the secondary DC. All other roles shown can be active in both DC locations simultaneously.

Figure 4 Single Active Directory Site with Stretched CCR—Two Data Center Locations

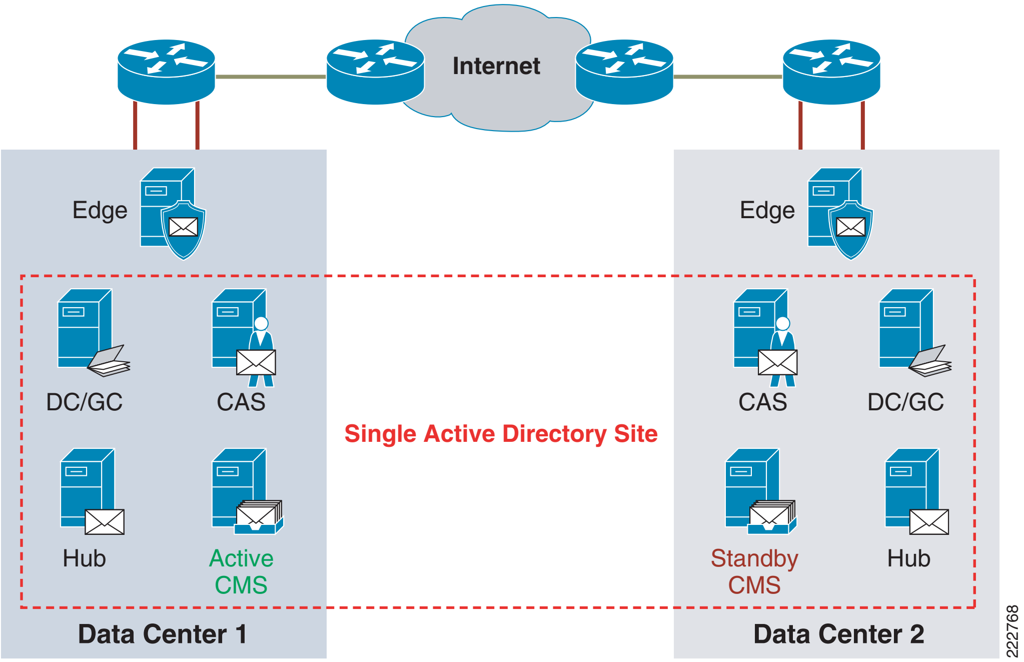

Multiple AD Sites—Multiple Data Centers

While the single AD site model allows for the ability to have most Exchange Server 2007 roles in an active/active configuration, the requirements for supporting such a design may outweigh the advantages. As discussed in Single AD Site — Multiple Data Center Locations, there are many considerations to plan for when dealing with a single AD site model for Exchange Server 2007. The AD, Exchange, and network administrators must balance the active use of resources in all data center locations against the management and cost associated with the support of full active-use of each resource in each location.

The model of supporting at least one AD site per data center location is easier to plan and deploy as well as support, especially when the data centers are geographically dispersed. If the primary goal is that of site-wide disaster recovery versus load balancing between sites, the multiple AD site model is more appropriate. With that said, it is possible to have some roles of the Exchange Server 2007 deployment be active/active in the multiple AD site model. The following are two examples of using an active/active configuration with multiple AD sites:

- CAS Deployment—A deployment may have multiple CAS roles per DC location and each DC has one or more AD sites. If site load balancing directs a user request to a CAS role located in a different DC (subsequently a different AD site) than the user belongs to, a feature known as CAS-CAS proxying can still connect the user to the correct CAS role for their site which then connects to the correct mailbox server. This feature allows for the CAS roles to be active at both DC locations. More information can be found on CAS-CAS proxying at: http://technet.microsoft.com/en-us/library/bb310763.aspx
- Edge Transport Deployment—Using the above CAS deployment scenario where there are multiple ET roles that are deployed in multiple DC locations, it is possible to allow all ET roles to be operational at all DC locations. EdgeSync subscriptions are used to establish connectors between HT and ET roles. The EdgeSync subscription connects the ET role to the HT role located in a specific site. Based on this process, if a ET role receives mail that is meant for a mail recipient located in different AD site than that ET role is subscribed to (via the Hub), the message is routed to the local Hub which routes the message to the HT role in the destination AD site. This configuration is described in more detail here: http://technet.microsoft.com/en-us/library/bb266920.aspx

Similar types of considerations exist for both single AD and multiple AD site models but are less stringent for the multiple AD site model. Microsoft Active Directory Sites and Services is designed to implement and deploy multiple AD sites, their resources and schedules for AD replication. As they apply to AD and Exchange, bandwidth and latency requirements for the network are also less stringent because the links between DC locations are mostly used for AD and Exchange Server 2007 Mailbox replication versus full-time use for replication in addition to active/active traffic flow.

Depending on how the Exchange Server 2007 mailbox clustering is deployed, there are two common ways to implement multiple AD sites between data centers:

- Stretched CCR—AD site per DC with the primary AD site stretched to include the passive node mailbox server located in the second DC.
- Local CCR + Remote Standby Continuous Replication (SCR)—AD site per DC with both CCR nodes at the primary DC and SCR node in the secondary DC.

There is more discussion on CCR and SCR in upcoming sections of this document.

Figure 5 illustrates using Microsoft Exchange Server 2007 with a stretched CCR design between two AD sites. There is one AD site per DC location but with the AD site from the primary location being stretched across the Layer 2 DC-to-DC interconnect (not shown) to the standby DC location. This stretched AD site design is required as Exchange CCR nodes must exist in the same logical AD site regardless of which physical DC they reside in.

Figure 5 Multiple Active Directory Sites with Stretched CCR—Two Data Center Locations

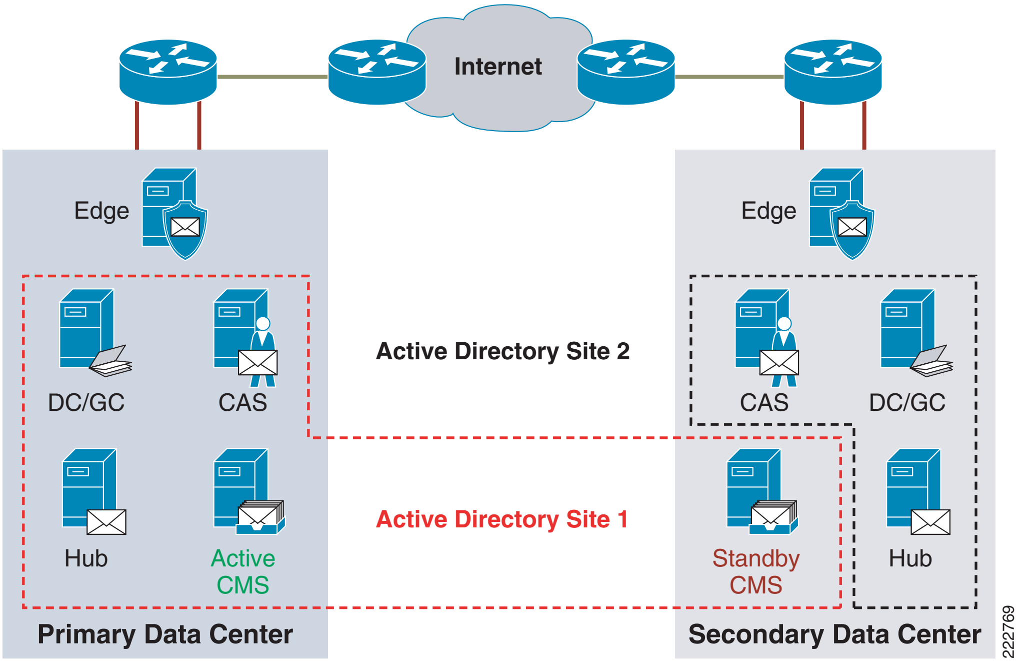

Figure 6 illustrates the use of Microsoft Exchange Server 2007 with a local CCR and remote SCR implementation. There is one AD site per DC location and since the CCR implementation is local to the primary site and no longer stretched, there is no need to also stretch the AD site for CCR between physical DC locations. SCR offers an excellent way to provide mailbox server availability without requiring alternative AD site designs.

Figure 6 Multiple Active Directory Sites with Local CCR + Remote SCR — Two Data Center Locations

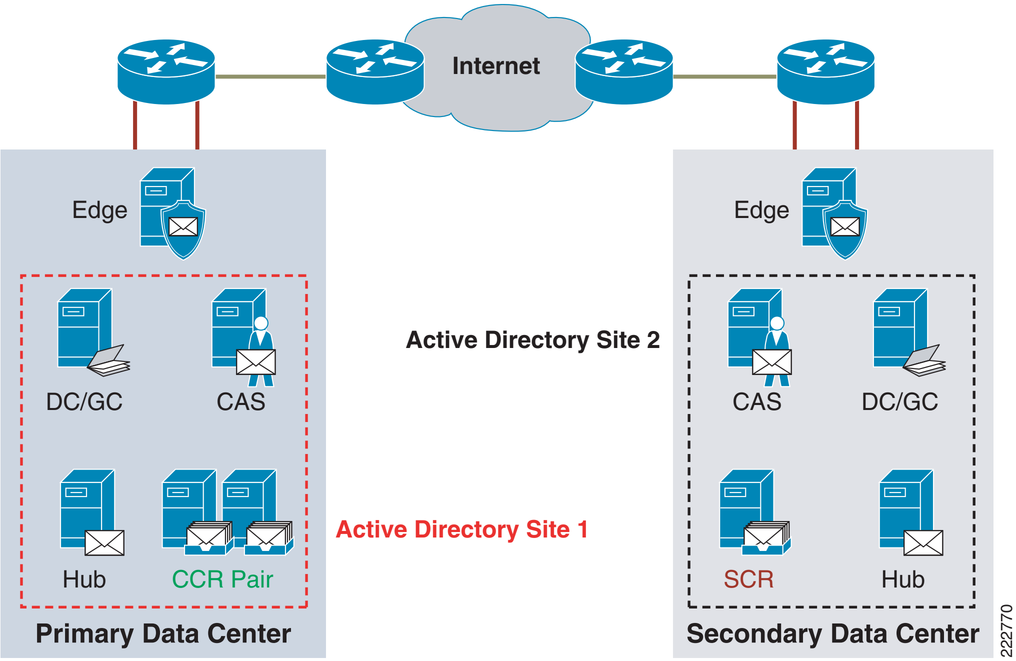

There are many decisions that need to be made in correct order when a server and/or site failure occurs. Microsoft has a well documented flowchart that discusses what to do in the event of a resource or site failure with Exchange Server 2007. The documentation can be found here: http://msexchangeteam.com/archive/2007/10/08/447211.aspx

Tested Microsoft Exchange Server 2007 Deployment Models

Microsoft Exchange Server 2007 Layout

There are many possible combinations of Exchange Server 2007 implementations. In this document, two implementation examples are explored in more depth and have specific Cisco product, feature, and design elements associated with both implementation examples. The two AD and Exchange Server 2007 implementation examples discussed in this document are:

- Single-Site AD with Stretched CCR—Two active/active data centers
- Multisite Active Directory—Local CCR + Remote SCR— Active/standby data centers

Single-Site AD with Stretched CCR

As discussed earlier, the goal of the single AD site with stretched CCR design is to support an active/active data center design for Microsoft Exchange Server 2007. Having the Exchange roles in a single logical AD site eliminates the complexity and delay of having to perform an AD "fix up" on Exchange roles in the event of a site failure at the primary site. Since each Exchange role is within a single AD site, nothing within AD has to be done in the event of failure at either site to allow Exchange to continue operating.

The AD layout of this design is discussed in the previous section and illustrated in Figure 4. The following section is more focused on the Exchange Server 2007 roles, their locations within the two data centers, and specific Exchange considerations for supporting the stretched CCR design.

Client Access Server—Active/Active DC

Multiple Microsoft Exchange 2007 servers running the CAS role are used not only to provide fault tolerance for individual server failures and scalability to support larger volumes of sessions, but also to provide a means for supporting local site load balancing as well as geographical load balancing between sites.

In addition to being an ideal candidate for server and site load balancing, the CAS role can additionally take advantage of network optimization services and SSL-offloading.

In Figure 7, a total of four Exchange are running the CAS role are deployed at the two DC locations. In this example, the CAS role has been deployed at the Internet DC (IDC) edge in a DMZ context that is specifically configured for the CAS role and services both internal and external client connections. Optionally, CAS roles can be deployed within the internal enterprise DC for internal CAS services while Internet-facing CAS roles service clients that are externally located. Both deployment options are supported in a Cisco multisite data center solution.

Figure 7 CAS Deployment – Active/Active Data Center

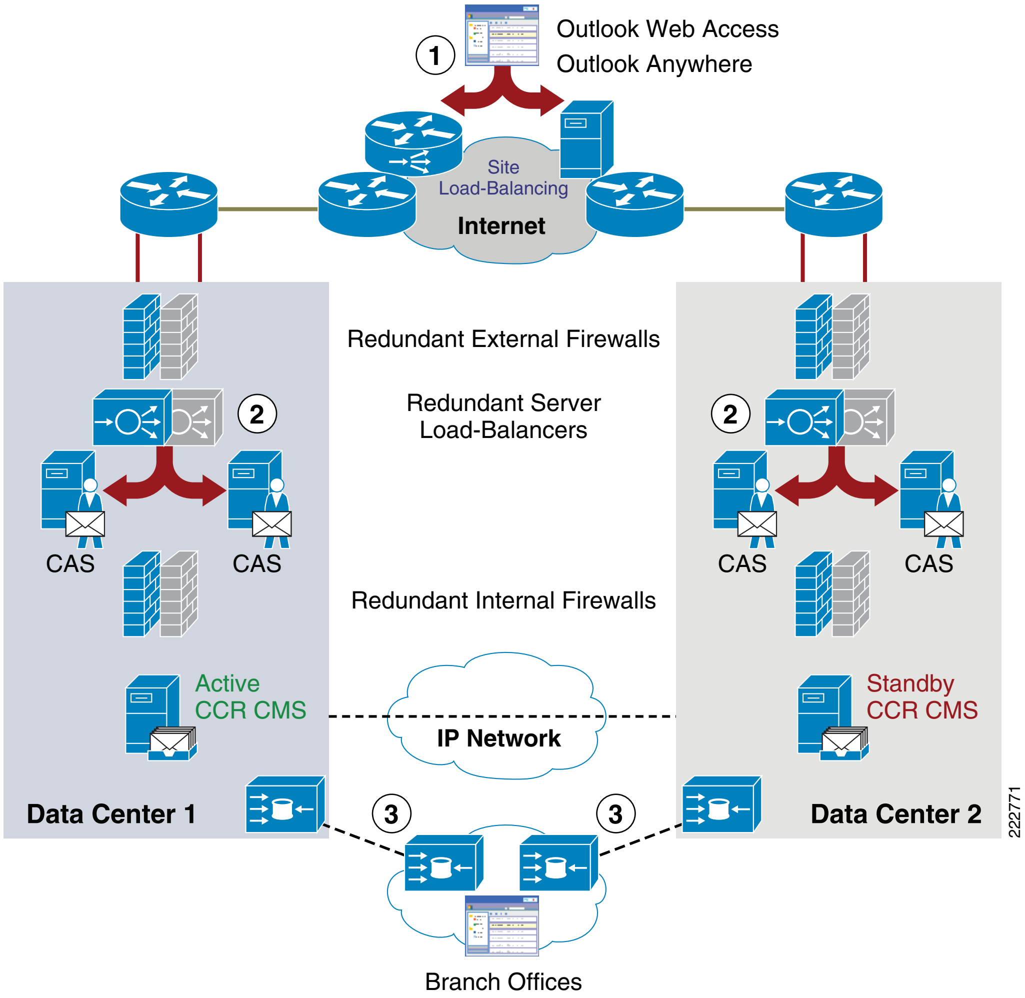

The numbered objects in Figure 7 correspond to the areas where the CAS role can interoperate with networking services.

1. Site selection and load balancing for each of the CAS Web (OWA, Outlook Anywhere, Autodiscover, etc…) and non-Web (POP3/IMAP4) services via the Cisco Global Site Selector product or generic DNS round-robin.
2. The Cisco ASA or FWSM can be used to provide firewall services. The Cisco ACE module can be deployed for Layer 4 through Layer 7 load balancing and can monitor the health of the CAS services and intelligently balance traffic amongst multiple CAS roles as well as report availability to the Cisco GSS. Also, at the same location, SSL-offload can be performed on the CAS role to help scale services such as OWA which uses HTTPS. The SSL-offload features of the Cisco ACE can help reduce CPU utilization on the CAS role by offloading the encryption/decryption process for each individual HTTPS session.
3. If branch office users connect to the CAS services located at either of the active data center locations, the Cisco WAE product can perform WAN optimization on the sessions to reduce bandwidth utilization, optimize the TCP sessions and reduce or eliminate duplicate data being transmitted between sites. It is important to note that the Microsoft Exchange and network administrators work together to understand the PROS and CONS of optimizing CAS services by

using a WAN optimizer. A decision needs to be made on whether or not client-to-server encryption of data in transit is more or less important than network optimization of the same traffic. End-to-end encryption and optimization of the payload are generally mutually exclusive.

Some customers will disable encryption for the OWA, Outlook Anywhere, or Outlook MAPI connections and rely on site-to-site VPN for encryption between the branch and DC. This still allows for encryption across the WAN while enabling optimization of the MAPI or HTTP flows.

Figure 8 shows an example traffic flow for the CAS role when an external user is connecting to CAS services, in this case OWA.

Figure 8 CAS Traffic Flow Example

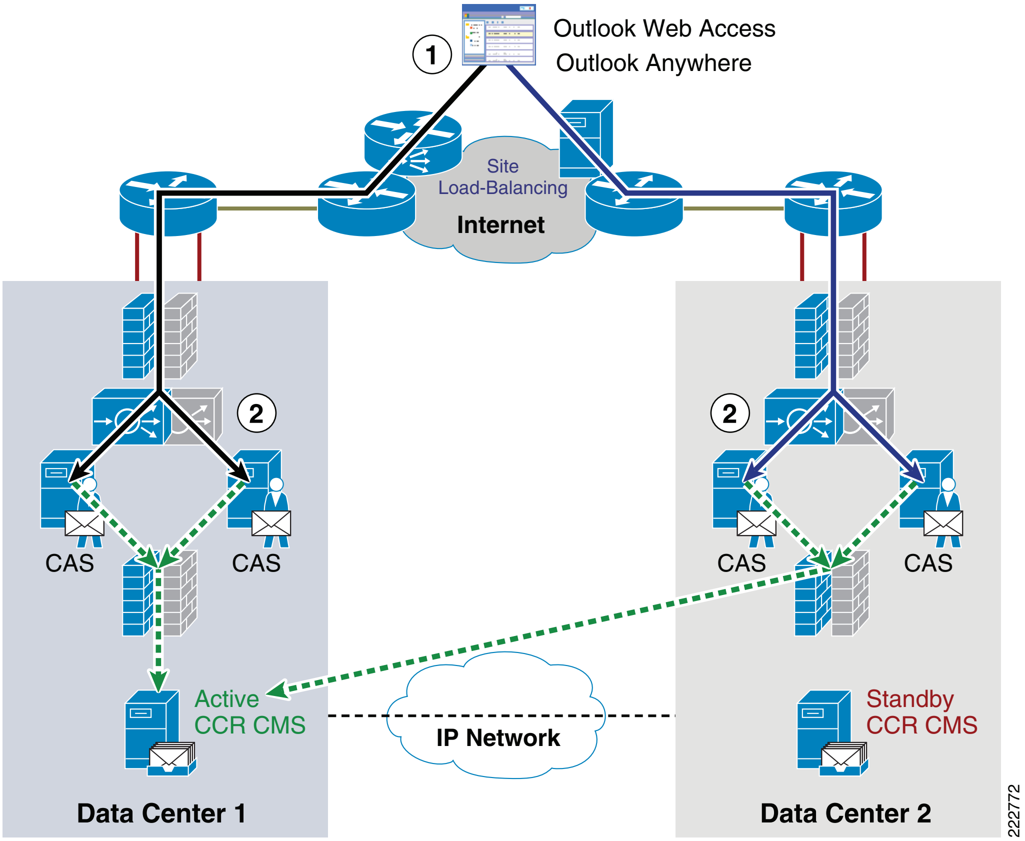

In Figure 8, the following occurs:

Step 1 The Cisco GSS is used to intelligently select the most appropriate site to load-balance the OWA traffic to. Alternatively, DNS round-robin can be used but DNS round-robin does not offer advanced features such as proximity-based routing, VIP tracking and service availability for each site and resource. The solid lines indicate possible paths the session may take depending on the load-balancing decision made for the CAS.

Step 2 Once the OWA traffic has been routed to a specific DC from Step 1, additional network services can be leveraged for security, performance and availability. Firewall services, intelligent server load balancing and SSL-offload can be performed on the OWA services at this step. After a specific server running the

CAS role has been selected and the above mentioned services performed, the OWA session is successfully connected to the CAS. The dotted lines indicated the paths from the CAS to the appropriate mailbox server that is associated with the user mailbox object.

In this example, all CAS-to-mailbox connections terminate on the active CCR CMS located in DC1. In the event of a CCR node failure at DC1, all CAS-to-mailbox connections automatically (or manually) terminate at the CCR CMS at DC2. This type of failover does not impact the steps for site and CAS load balancing, firewall services, or SSL-offload.

Additional information on the Microsoft Exchange 2007 CAS configuration options and Cisco-specific network configurations for this design will be discussed later in this document.

Hub Transport Server—Active/Active DC

The Hub Transport (HT) role has a built-in load balancing (round-robin) capability that allows each HT role to select another Microsoft Exchange 2007 role to route mail to. This is true of mail routing internal to the same AD site, a different AD site, and for external SMTP relay to ET roles (via EdgeSync subscriptions).

Other than providing firewall services for communication to/from HT roles in the active/active design, no other significant network services are being discussed in this document. Communications between HT roles and other roles is usually encrypted using Transport Layer Security (TLS) and also load-balanced within the HT role itself. External load balancing, SSL-offload and network optimization do not offer the same level of advantages for the HT role as with the CAS and ET roles.

One thing to note regarding the HT role and the single AD site design being discussed in this section is the communication between the HT and the mailbox server role. By default, mailbox servers use a list of known HT roles within their own AD site and will make a round-robin based decision on which HT role to send mail traffic to. Since the single AD site model spans multiple DC locations, the mailbox server may load-balance mail traffic to HT roles at a different DC location only to have the HT make its own load-balance decision and send the mail back to another Exchange server role in the original DC location.

Figure 9 illustrates a likely scenario where mail from a particular mailbox server needs to be routed to an external mail recipient. The mailbox server, by default, will load-balance to all HT roles in its own AD site (in this example, a single AD site spans both DC1 and DC2). In this example, a message needs to be routed from the mailbox server “ccr-mbx-01” to an external recipient and the round-robin choice this time is to “rtp2-hub-01” located in DC2. The message traverses the DC-to-DC interconnect link and arrives at “rtp2-hub-01”. As mentioned before, the HT roles also do round-robin load balancing on EdgeSync subscriptions within their own AD site and the round-robin choice this time is “rtp-edge-01” in DC1. Again, the message traverses back across the DC-to-DC interconnect and arrives at rtp-edge-01 and subsequently is relayed to the Internet.

Figure 9 Sub-optimal Active/Active Mail Flow Example

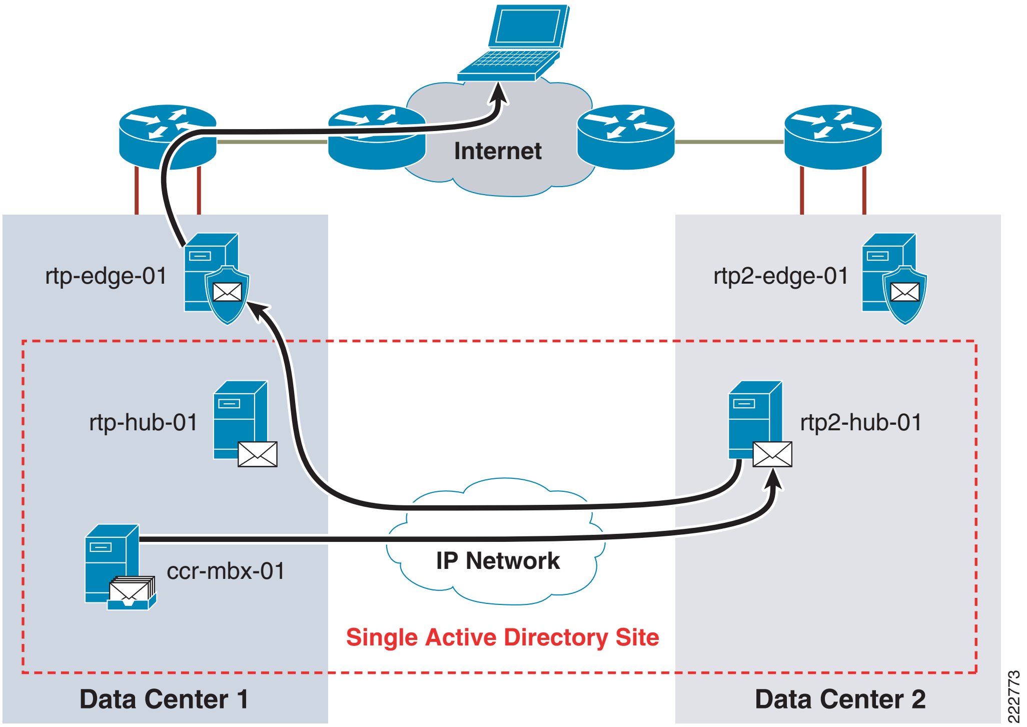

If this message routing is of concern to the network and Exchange administrator, it is possible to force the mailbox server to use only those HT roles that are physically located in the same DC as the mailbox server. This is done by modifying the **SubmissionServerOverrideList** for the mailbox server and list out only those HT roles located in the same DC as the mailbox server. For example, if the administrator wanted the mailbox cluster "ccr-mbx-01" to only use "rtp-hub-01" and rtp-hub-02" because they were physically located in the same DC, the administrator could run the following command in the Exchange Management Shell:

```
[PS] C:\>Set-MailboxServer -Id:ccr-mbx-01 -SubmissionServerOverrideList: rtp-hub-01,
rtp-hub-02
```

The setting is now altered in the mailbox server properties in Active Directory:

```
[PS] C:\>Get-MailboxServer -Id:ccr-mbx-01 | fl SubmissionServerOverrideList
SubmissionServerOverrideList : {RTP-HUB-02, RTP-HUB-01}
```

The mailbox server will now only use the configured HT roles listed. However, the administrator needs to ensure this setting is changed if there were a failure of one or both of the configured HT roles. An example of such a failure would be in the case of a total site failure in DC1. The mailbox server(s) now active in DC2 needs to be configured to use the HT roles located in that physical DC site. Without modifying the override list, the DC2 mailbox servers would only try to communicate with HT roles located in a site now unreachable (due to site failure).

Mailbox Server—Active/Active DC

In the single AD site with stretched CCR model, the most significant role that impacts the AD, Exchange and data center network design is the clustered mailbox server (CMS). The following list outlines some facts, requirements, and recommendations for CCR:

- CCR uses a two-node active/passive design that is built on top of the Microsoft Windows Server 2003 Cluster service

The cluster quorum design recommended by Microsoft is the Majority Node Set (MNS) with File Share Witness (FSW) – Quorum explanation and configuration is beyond the scope or purpose of this paper. Detailed information on the various quorum models, their configuration and best practices is found at:
http://technet.microsoft.com/en-us/library/e11b4f0b-53cb-4a38-a550-eba1cbbe31f8.aspx

- A minimum of two Network Interface Cards (NIC) is required in each server node and are divided into a “public” (client access) and “private” (cluster heartbeat) network.
- The nodes must communicate with each other over the same subnet (running the Mailbox role on Windows Server 2008 removes this requirement).
- Both CCR nodes must be in the same AD site.
- Microsoft recommends altering the tolerance for missed heartbeats on private and mixed interfaces to “10”. For more information, refer to: http://technet.microsoft.com/en-us/library/bb310759.aspx.
- Microsoft recommends a minimum of a gigabit link and latency no greater than 500 msec for the CCR link. For more information, refer to:
http://technet.microsoft.com/en-us/library/a26a07fc-f692-4d84-a645-4c7bda090a70.aspx.
- Messaging clients located in branch offices can be WAN optimized between the client and mailbox server under certain circumstances.
- Configure the transport dumpster feature on the HT roles to re-deliver messages after a lossy CCR failover has occurred. For more information, refer to
http://msexchangeteam.com/archive/2007/01/17/432237.aspx.

A stretched CCR implementation is one where each CCR node is in a separate physical DC and the communication between each node is “stretched” over the DC-to-DC interconnect. Based on this definition and the requirements listed above, the network and Exchange administrators can now implement the required components in each DC to connect the CCR nodes.

Figure 10 illustrates the high-level network layout between DC locations to support the stretched CCR design.

Figure 10 Stretched CCR – Network Layout

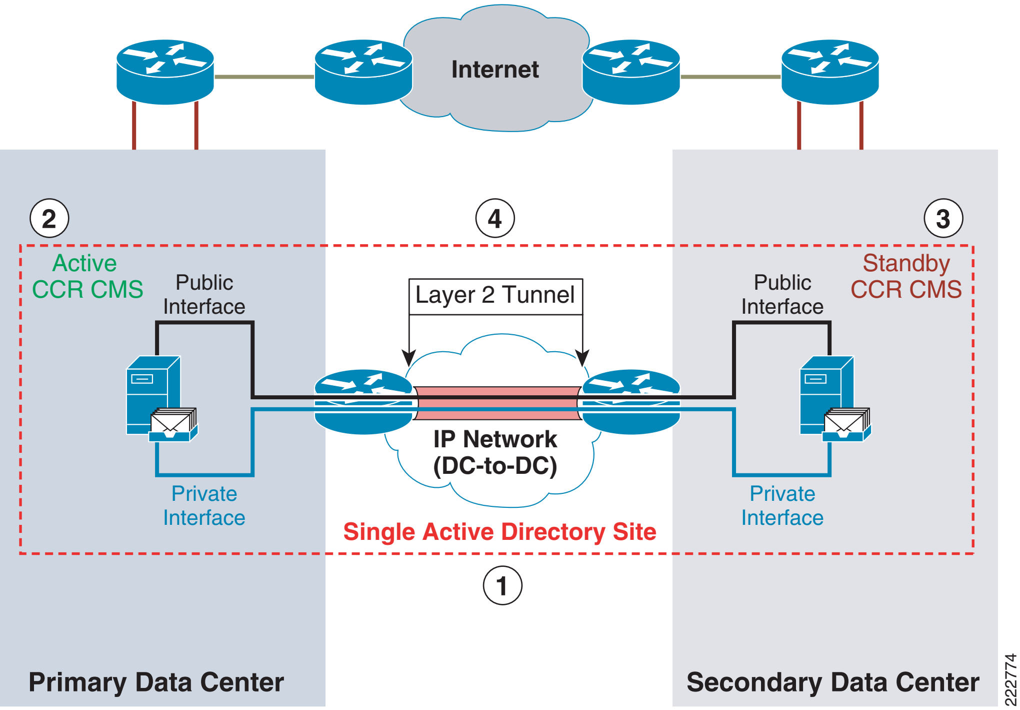

In Figure 10:

1. Both active and standby CCR nodes are in the same AD site that spans both physical DC locations.
2. The active CCR node is located in DC1 and is the only node to accept connections directly from MAPI clients, CAS roles and HT roles.
3. The standby CCR node pulls log files from the active CCR node using clustered network interface resources that now span the Layer 2 link between DC locations. This “pull” models is a significant improvement over past models where high latency between the active and standby nodes could adversely impact the active node and user experience. This is no longer the case in CCR.
4. The Layer 2 link is used due to the existing Microsoft Windows Server 2003 cluster requirements for same-subnet communications for both public and private interconnected links. The testing completed for this paper used an L2TPv3 configuration. Any transport/technology that satisfies the Layer 2 adjacency requirement for the cluster and also provides sufficient bandwidth and latency numbers can be used as an alternative to L2TPv3.

Log Shipping and Seeding

It is important to note that by default CCR uses the public interface to perform log shipping on all transaction logs and seeding. This must be accounted for in the planning and allocation of the capacity for the DC-to-DC interconnect links.

Exchange Server 2007 SP1 allows for the transaction log shipping and seeding/re-seeding to occur over “mixed-interfaces”, which are interfaces that support both client access and cluster heartbeat traffic. Enabling this functionality is important for many reasons, one being that if the public NIC connection between two CCR nodes goes down, logs can build up on the active CCR node. This not only causes a backup of logs that need to be shipped when the connection is restored, but in the event that the standby CCR node needs to be brought online as the active node, the logs will not be current. If support is enabled to allow for log shipping and seeding over mixed interfaces, automatic switchover of log shipping and seeding can occur in the event that the public NIC or DC-to-DC interconnect carrying the public VLAN or lambda fails.

The Exchange administrator can use the **`Enable-ContinuousReplicationHostName`** cmdlet in the Exchange Administrator Shell to enable mixed-interface support for log shipping and seeding. More information can be found at: http://technet.microsoft.com/en-us/library/bb690985.aspx.

In addition to enabling support for seeding and log shipping over mixed interfaces, it is very important to understand the network bandwidth and latency impact on CCR operation. Bandwidth is an important element to ensuring that CCR will perform seeding and log shipping properly, but bandwidth is not the most critical factor. Network latency between CCR nodes is the single most important network factor for ensuring that seeding and log shipping are completed in a timely manner. Latency that is too high will cause a backup of logs on the active CCR node and can prohibit access to the latest logs on the standby CCR node in the event of a failure. Ensuring the lowest possible network latency by selecting the proper transport type, distance between nodes, fine tuning Windows Server 2003 TCP/IP options and even the use of network optimizers such as the Cisco WAE are all proven ways to allow for timely seeding and shipping of logs and DB for CCR.

Hewlett Packard (HP) wrote the following document that discusses the impact of modifying the Windows Server 2003 TCP/IP options to realize better CCR seed and log shipping performance:

http://h71028.www7.hp.com/ERC/downloads/4AA1-4230ENW.pdf.

For more information on Windows Server 2003 TCP/IP tuning, refer to:
http://download.microsoft.com/download/2/8/0/2800a518-7ac6-4aac-bd85-74d2c52e1ec6/tuning.doc

Cluster Heartbeat Modifications

Microsoft has made recommendations to alter the default heartbeat tolerance for CCR environments. It is important to alter these values for the CCR nodes so that they are more tolerant of missed/lost heartbeats. This is especially important over long-distance DC-to-DC interconnections where latency or oversubscription may impact reliable reception of heartbeats. The cluster heartbeat values can be altered on each node (node should be in passive state) by using these commands:

```
C:\> cluster Clustername /priv HeartBeatLostInterfaceTicks=10:DWORD
C:\> cluster Clustername /priv HeartBeatLostNodeTicks=10:DWORD
```

To successfully complete this process, additional steps are required to stop and start the cluster service on each node. Refer to the instructions found at:
http://technet.microsoft.com/en-us/library/bb310759.aspx

Storage Requirements

CCR does not use shared storage. Unlike SCC which does use shared storage, CCR creates a separate copy of logs and database on the attached storage on both active and standby nodes. Log shipping is used to replicate logs created on the active node over to the standby node via a pull model (standby pulls logs from the active). Information on storage planning and requirements for CCR can be found here:
http://technet.microsoft.com/en-us/library/bb124518.aspx and
http://technet.microsoft.com/en-us/library/bb738147.aspx

Note While CCR does not use shared storage and therefore has no requirements for a SAN-based replication design between data centers, SCC does use shared storage. A future update to this document will include SCC considerations for both the front-end (IP network and intelligent network services) and back-end (SAN and SAN-based replication) in a multisite DC.

WAN Optimization

Messaging clients such as Microsoft Outlook 2007 can have the network connections between the client and mailbox server optimized by the Cisco WAAS solution. Similar to the CAS discussion with WAN optimization, the network and Exchange administrators must make a decision on whether or not client-to-server encryption for messaging clients located in a branch office is more or less important than the significant bandwidth savings and improved application response times are with WAN optimization. Much more detail on the messaging client requirements and the Cisco WAAS solution configuration can be found later in this document.

Edge Transport Server—Active/Active DC

Edge Transport (ET) roles should be deployed in a redundant way and are an ideal candidate for both role and site load balancing. Similar to the CAS role, deploying server and site load balancing for the ET role provides additional fault tolerance within the site and between sites as well as scalability of the Edge services.

The ET roles are deployed, as the name implies, at the edge of the network in a secured and load-balanced DMZ context. Some organizations will deploy their ET roles in a design that centralizes either inbound or outbound SMTP messages or perhaps both. One example of this is where an organization has multiple sites around the world and designs the Exchange environment to only allow SMTP messages into the organization via a site or sites within the United States, but permits outbound SMTP messages from any site internationally. Regardless of the reasons why or how the organization wants external SMTP messages to flow, scalability, security and availability are always required. In this document, ET roles are both site and role load-balanced and mail routing rules are equal between both DC sites. The Exchange configuration below shows two connectors, inbound and outbound, and they apply to all ET roles in both physical DC locations; mail is processed equally across the connectors between the ET and HT roles.

```
Identity                    AddressSpaces Enabled
--------                    ------------- -------
EdgeSync - ESE to Internet {smtp:*;100}  True
EdgeSync - Inbound to ESE  {smtp:--;100} True
```

Figure 11 shows a total of four Exchange servers running the ET role and they being deployed at the two DC locations. In this example, the ET role is deployed at the Internet data center edge in a DMZ context that is specifically configured for the ET role. As previously mentioned, the ET roles are not members of the internal AD forest and communicate with the internal HT roles via EdgeSync subscriptions. Later in the document configurations are shown on how the ET and CAS roles can be secured and/or load-balanced by the same Cisco products and features at the IDC, but in different contexts (for security and management purposes).

Figure 11 Edge Transport Deployment – Active/Active Data Center

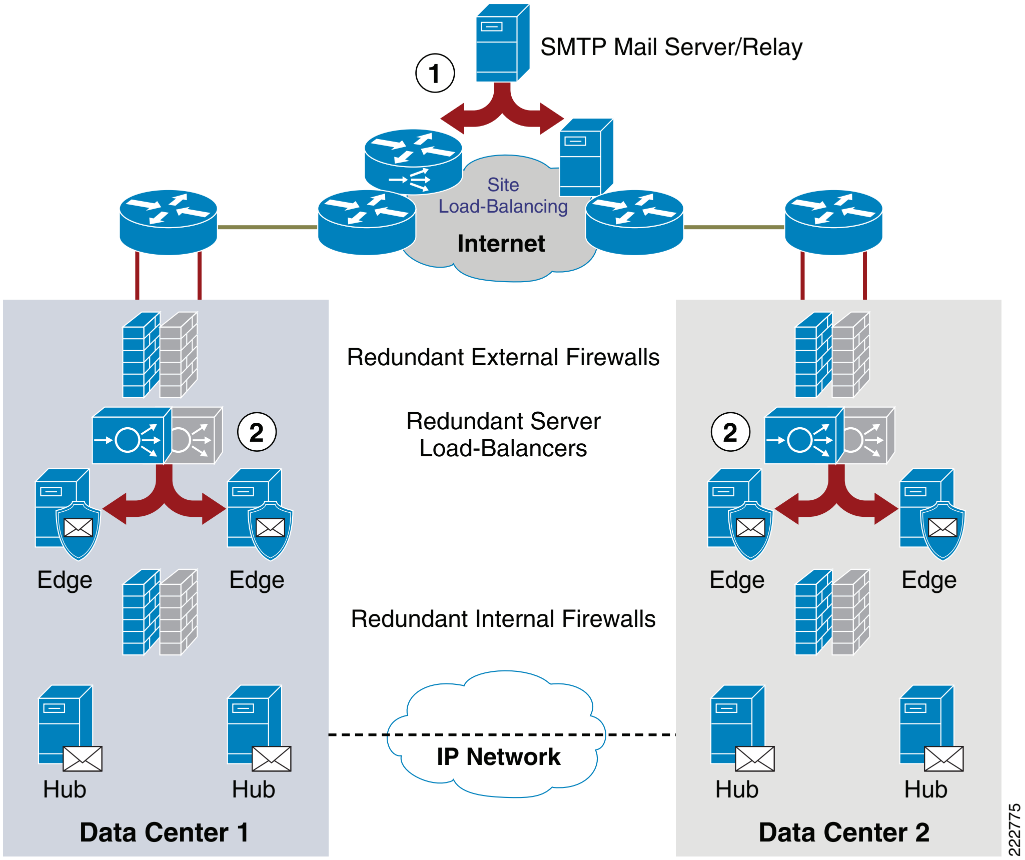

The numbered objects in Figure 11 correspond to the areas where the ET role can interoperate with networking services.

Step 1 Site selection and load balancing for SMTP mail relay via the Cisco GSS product or generic DNS round-robin.

Step 2 The Cisco ASA or FWSM can be used to provide firewall services. The Cisco ACE module for Layer 4 to Layer 7 load balancing can monitor the health of the ET services and intelligently balance traffic amongst multiple ET roles as well as report availability to the Cisco GSS.

Figure 12 shows an example traffic flow for the ET role when an external SMTP message is sent from the Internet to the organization.

Figure 12 Edge Transport Flow Example

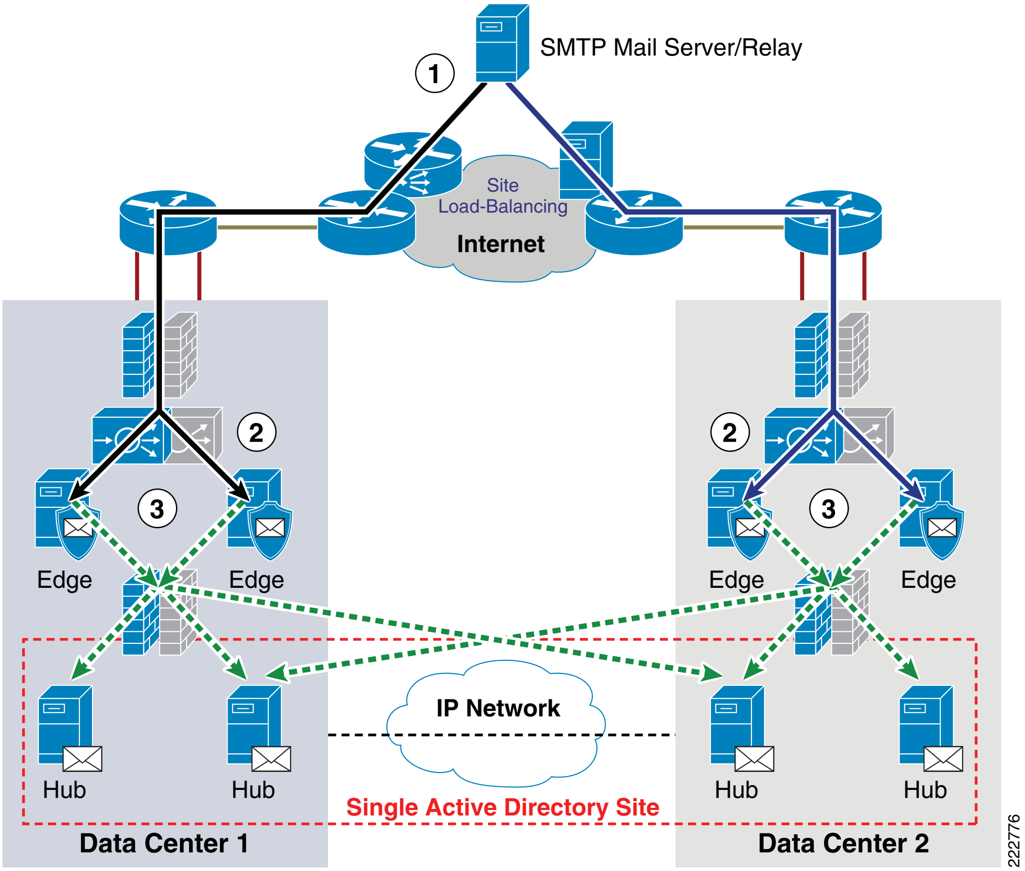

The traffic flow for Figure 12 is very similar to that of the CAS role:

Step 1 The Cisco GSS is used to intelligently select the most appropriate site to load-balance the SMTP traffic to. Alternatively, DNS round-robin can be used but DNS round-robin does not offer advanced features such as proximity based routing, VIP tracking, and service availability for each site and resource. The solid lines indicate possible paths the session may take depending on the load-balancing decision made.

Step 2 Once the SMTP traffic has been routed to a specific DC from Step 1, additional network services can be leveraged for security, performance and availability, such as firewall and server load balancing.

Step 3 After a specific server running the ET role has been selected and the above mentioned services performed, the SMTP messages are sent directly to an internal HT role via connectors configured during the EdgeSync subscription process. SMTP messages inbound and outbound between HT and ET roles are load-balanced using a round-robin method. The dotted lines indicate possible paths that may be taken from the ET to the HT roles based on the round-robin selection outcome.

The differentiating factor between this traffic flow and a flow with a multi-AD site model is that the ET roles does not send SMTP messages between DC locations to the other HT roles as they would be in a different AD site. EdgeSync subscriptions are on a per-AD site basis.

In the event of role or total site failure, nothing needs to take place in the configurations of the ET or HT roles. There is a single AD site for both DC locations and SMTP messages is routed to an alternative group of roles that are still available (such as those in the other DC) in the single AD site.

Additional information on Exchange Server 2007 ET role configuration and Cisco-specific network configurations for this design is discussed later in this document.

Multisite Active Directory—Local CCR + Remote SCR

As discussed earlier in this document, there are many technical and business drivers for having multiple DC locations active, but that comes at a cost. In the previous section, the design was focused on having every possible Exchange role active and to evenly distribute load between DC locations by using site selection. In this section, a topographically similar design is discussed, but using different options for Active Directory and Exchange Server 2007, specifically the mailbox server cluster options.

Figure 13 illustrates the high-level view of the design discussed in this section. The design discussed in this section uses the same two DC locations referenced before, but with only two AD sites instead of one. For simplicity sake, there is an AD site per DC-location as shown in Figure 13. Normal AD replication mechanisms is used to ensure AD is in sync between DC locations; this includes Exchange configuration and objects held in AD.

In this section, the DC design is active/standby for the entire application environment. The previous section discussed the components involved for active/active and it is important to discuss the active/standby design as an alternative for both network and Exchange administrators.

The HT roles uses the ET roles in their own AD site for relaying SMTP messages external to the organization. The mailbox server configuration can support a stretched CCR deployment that was discussed in the last section, but in this section a local CCR with a remote SCR design will be used. The DC design can be configured to support the Exchange Server 2007 CAS and ET roles in an active/active or active/standby configuration. As it relates to the Cisco part of the design, the main difference for these two roles in active/active or active/standby is how the Cisco GSS and/or DNS are configured for site preference. There are certainly other considerations to understand such as how CAS-CAS proxy works in this scenario as well as connectors and transport routing between ET and HT roles. These Exchange-specific topics are deep discussions for Exchange administrators and have little-to-nothing to do with the actual Cisco networking design; therefore the CAS/ET active/active component of this design will not be discussed.

As discussed later in this document, the CCR + SCR design easily allows for multiple backup data centers to exist in the event that the primary and even secondary DC become unavailable. A good example of this is a natural disaster such as a hurricane that impacts an entire region where the two main DC locations exist. The SCR capability in Exchange Server 2007 with SP1 allows for multiple copies of the mailbox server data (logs and DB) to be replicated to each DC regardless of location. This is an ideal configuration for geographically dispersed clusters.

Figure 13 Overview of Multiple Data Centers and AD sites with Local CCR + Remote SCR

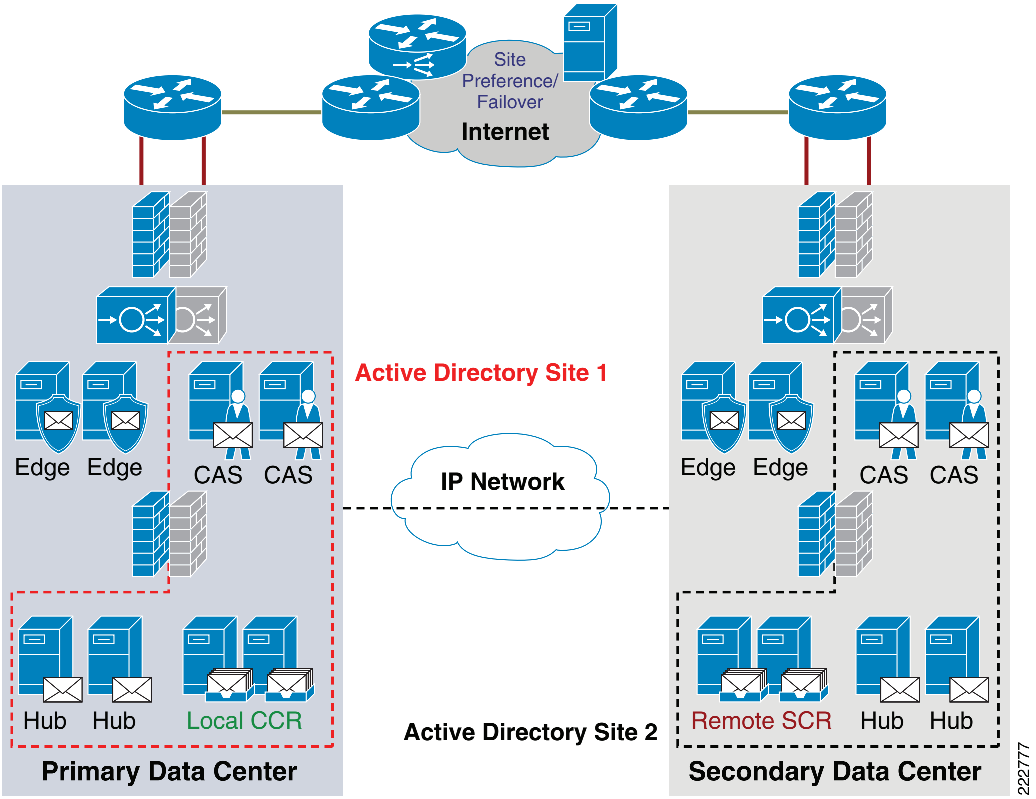

Message Flow

In an active/standby design, the message flow is much simpler and deterministic for messages coming into the DC location. In the active/active design, the Cisco GSS and DNS services evaluate which of the DC locations is most appropriate to send CAS (client access) or ET (SMTP) targeted traffic to. This can be as simple as round-robin or as complex as tracking the actual availability and load of the servers at one site in comparison to other sites. With active/standby all client access and messaging traffic goes go to the primary site. In the event of a site failure, the Cisco GSS and DNS can be configured to prefer the secondary DC (which is now active). The DC design requirement now changes from site load balancing to site preference and failover.

Figure 14 illustrates the traffic flow for the CAS and ET roles in an active/standby design.

Figure 14 Active/Standby CAS and ET Traffic Flow

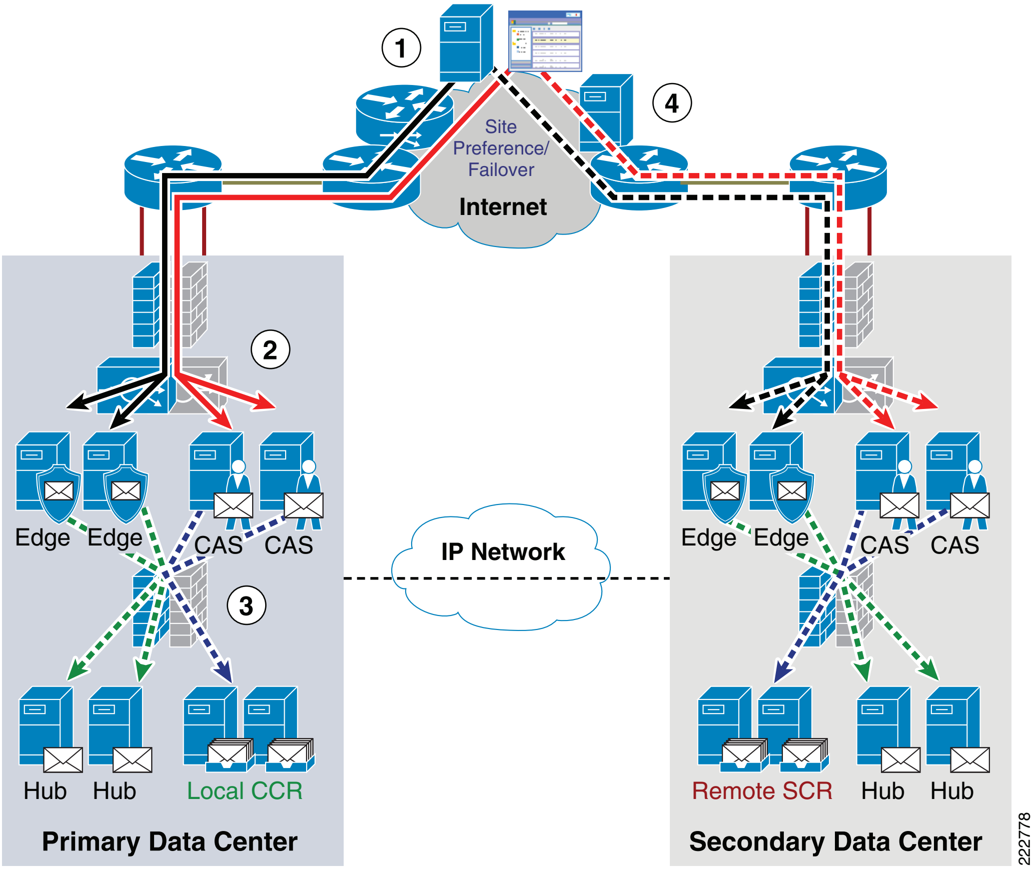

The following provide the steps for both CAS and ET targeted traffic flow:

Step 1 The CAS targeted traffic, such as OWA sessions (solid red line) or SMTP for ET (solid black line), are sent to the primary DC because the Cisco GSS and/or DNS are configured to prefer this site.

Step 2 The CAS and/or ET traffic arrives at the primary DC and has security, server load balancing, and SSL-offloading performed against the traffic as appropriate.

Step 3 The ET role relays the SMTP traffic to the primary DC HT roles based on round-robin selection. The CAS roles will connect to the appropriate active mailbox CCR node based on the mailbox association for the logged in user. The green and blue dotted lines indicate this flow.

Step 4 In the event of a site failure at the primary DC, the Cisco GSS and/or DNS entries for the CAS and ET roles can be modified to now prefer the secondary DC. The traffic (indicated by the dashed lines) flows to the Exchange server roles as they did in Step 2. The one change here is that the SCR node is activated as the active mailbox server for users and processes connections initiated from the CAS roles in the secondary DC.

Client Access Server—Active/Standby DC

The only network service for CAS that changes for the active/standby design is the site load-balancing service. In the event of a failure, the Cisco GSS and DNS server will be configured to prefer the Cisco ACE VIP entry for the CAS roles located at the secondary side. This must be a manual switchover because several things must occur before traffic is sent to the secondary site. There are several important steps that must be completed to ensure that client access requests can be successfully accepted at the secondary site. Some of these steps include:

Step 1 Active Directory access—Domain Controllers and Global Catalog servers must be online at the secondary site.

Step 2 Hub Transport roles—Message routing must be online in order to properly handle mail traffic for other AD sites and external SMTP traffic via the ET role.

Step 3 SCR Mailbox—There are several detailed steps required to ensure the SCR node is active with the correct logs and DB from the primary CCR cluster. This is not a trivial process and may take a significant amount of time depending on the environment and condition of the logs/DB on the SCR node(s).

Step 4 Edge Transport—In order for CAS attached users to send and receive external SMTP messages, the ET roles must be online.

Step 5 DNS Update—The Cisco GSS and DNS servers can now be updated to prefer the secondary DC.

Within the site the same requirements exist for firewall, server load balancing, network optimization, and SSL-offload as were required in the primary DC.

Hub Transport Server—Active/Standby DC

There are no additional requirements or considerations for the HT role in the active/standby design. The steps listed above in the CAS section include the requirement to ensure the HT role(s) are online in the secondary site to support mail routing for external AD sites and Internet source/destined SMTP messages.

Mailbox Server – Active/Standby DC

As was discussed in the single AD site with stretched CCR section, the most significant role that impacts a multisite DC design is the mailbox server role. In the local CCR + remote SCR design a local 2-node CCR cluster is deployed within the primary DC and either a single SCR node or clustered SCR deployment exists in the secondary site. The configuration and requirements discussed in the Mailbox Server—Active/Active DC, page 25, still have applicability in this design. However, the one difference is that the CCR nodes are within the same site and not stretched between DC locations. This significantly changes the requirements on the network and Exchange deployment.

Note CCR is not required in order to deploy SCR. SCR can be used as the target for a single standalone mailbox server as easily as it can with CCR.

In the stretched CCR design, the Windows Server 2003 cluster deployment requires that the public and private interfaces for the cluster be located on the same subnet. To support the requirement between DC locations, a Layer 2 tunnel was deployed so that the two VLANs for public and private were available at both DC locations. This is no longer a requirement when deploying SCR. The log shipping between the active CCR node (source) to the SCR node (target) at the secondary DC location can be done over a

Layer 3 link. This is a great change to the design because the log shipping between primary to secondary site can traverse an Layer 3 link on the DC-to-DC interconnect or alternatively, as a backup, traverse the front-end IP infrastructure such as a VPN. It is simply a routing decision to make this backup connection active.

Figure 15 illustrates the replication path between the local CCR nodes and the remote SCR node.

Figure 15 Local CCR Replication + SCR

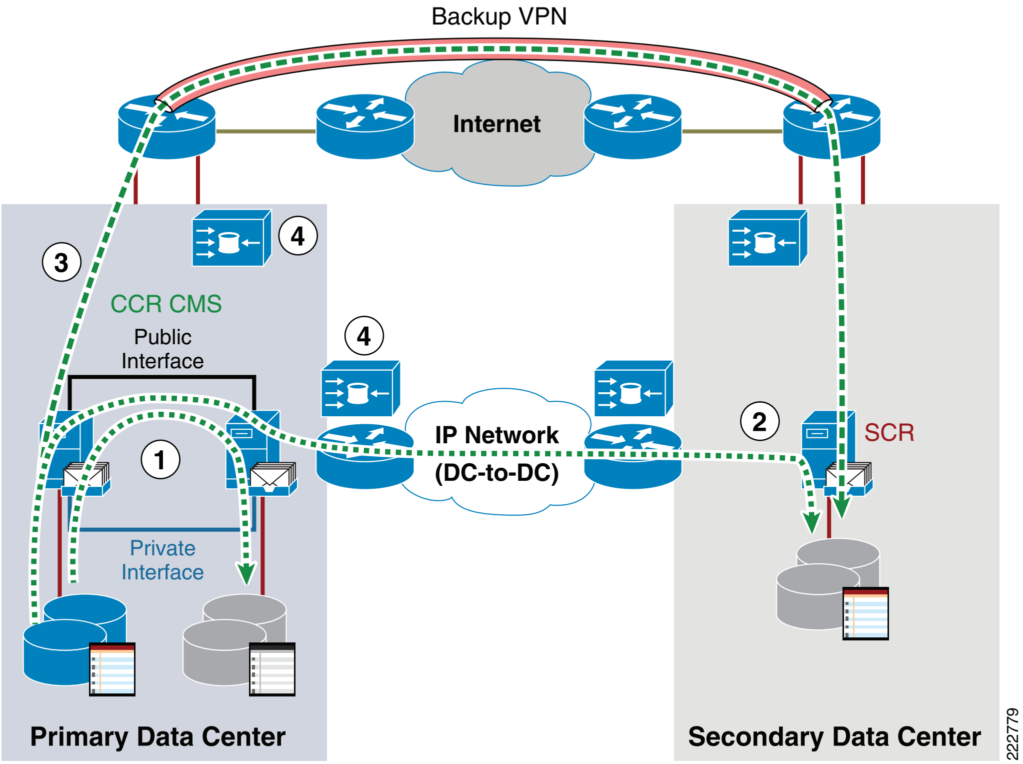

The process for Figure 15 is as follows:

Step 1 The active CCR node replicates the logs and DB for the attached storage group to the standby CCR node using a "pull" model. The replication is taking place over the public interface (default) between nodes, but may optionally replicate over a "mixed" interface. For more information, refer to Mailbox Server—Active/Active DC, page 25.

Step 2 The active CCR node is also replicating the logs and DB for the attached storage group to the SCR node in the secondary DC location. This replication is using the DC-to-DC interconnect link as the transport and the link can be Layer 3 or Layer 2.

Step 3 Optionally, as a backup to ensure continued replication between the active CCR node and the SCR node a second link between DC locations can be established using a technology such as site-to-site VPN or MPLS to route the replication traffic. This is a great fallback path in the event that the DC-to-DC interconnect link is lost for some reason.

Step 4 Note that server message block (SMB) is used to copy the logs between nodes. This allows for a WAN optimization solution such as the Cisco WAAS solution to optimize the bandwidth utilization and lower copy times between the two DC locations.

Edge Transport Server—Active/Standby DC

Similar to the CAS role, the ET role design and requirements do not change much in the active/standby design. Most network service modifications for the ET role are with the Cisco GSS and DNS, specifically the preference for the DNS MX records. Just like the CAS role, the ET role in the secondary DC can be automatically preferred in the event of a failure at the primary DC, but this should not be configured. For the same reasons as listed previously, the switch to using the ET roles at the secondary site must be made manually and with great care to ensure that traffic is not sent to the secondary site until the site is ready to successfully accept mail. Refer to the steps in the Client Access Server—Active/Standby DC, page 35.

The Cisco GSS, DNS, Cisco ACE and Firewall configurations are discussed for both the ET and CAS roles later in this document.

Optimization and Availability Support for Microsoft Exchange Server 2007 in a Cisco Multisite Data Center

It is important to give a brief overview of what roles within Exchange Server 2007 support various methods of load balancing (LB), fault tolerance (FT), network optimization and SSL-offload. Some of these methods are handled within the Exchange and/or server architecture, in standard ways such as with DNS or NIC-teaming or those methods that are provided by networking components.

In Table 1the Exchange 2007 server role is listed along with the method used to provide LB, FT, network optimization, and SSL-offload.

Table 1 Microsoft Exchange Server 2007 Role and LB, FT, HA Methods Supported

Microsoft Exchange 2007 Role	Site Load-Balancing	Server Load-Balancing	Fault-Tolerance	Network Optimization	SSL-Offloading
Client Access Server	Cisco Global Site Selector (GSS) and/or DNS round-robin	Cisco ACE, Microsoft Network Load-Balancing (NLB) or DNS round-robin	NIC-teaming, multiple CAS roles	Cisco WAE	Cisco ACE
Hub Transport Role	N/A	Handled internally by Microsoft Exchange	NIC-teaming, multiple Hub Transport roles	N/A	N/A
Mailbox Server	N/A	N/A	NIC-teaming, Clusters (LCR, CCR, SCR, SCC)	Cisco WAE	N/A
Edge Transport Role	Cisco Global Site Selector (GSS) and/or DNS round-robin	Cisco ACE, Microsoft NLB or DNS round-robin	NIC-teaming, multiple Edge Transport roles	N/A	N/A

Note The UM role was not tested and therefore not included in this table.

Enterprise Network Architecture

Data Center Network Components

The logical topology of the data center infrastructure can be divided into the front-end network and the back-end network, depending on their role:

- The front-end network provides the IP routing and switching environment, including client-to-server, server-to-server, and server-to-storage network connectivity.
- The back-end network supports the storage area network (SAN) fabric and connectivity between servers and other storage devices, such as storage arrays and tape drives.

Front-End Network

The front-end network contains three distinct functional layers:

- Core
- Aggregation
- Access

Figure 16 shows a multi-tier front-end network topology and a variety of services that are available at each of these layers.

Figure 16 Data Center Multi-Tier Model Topology

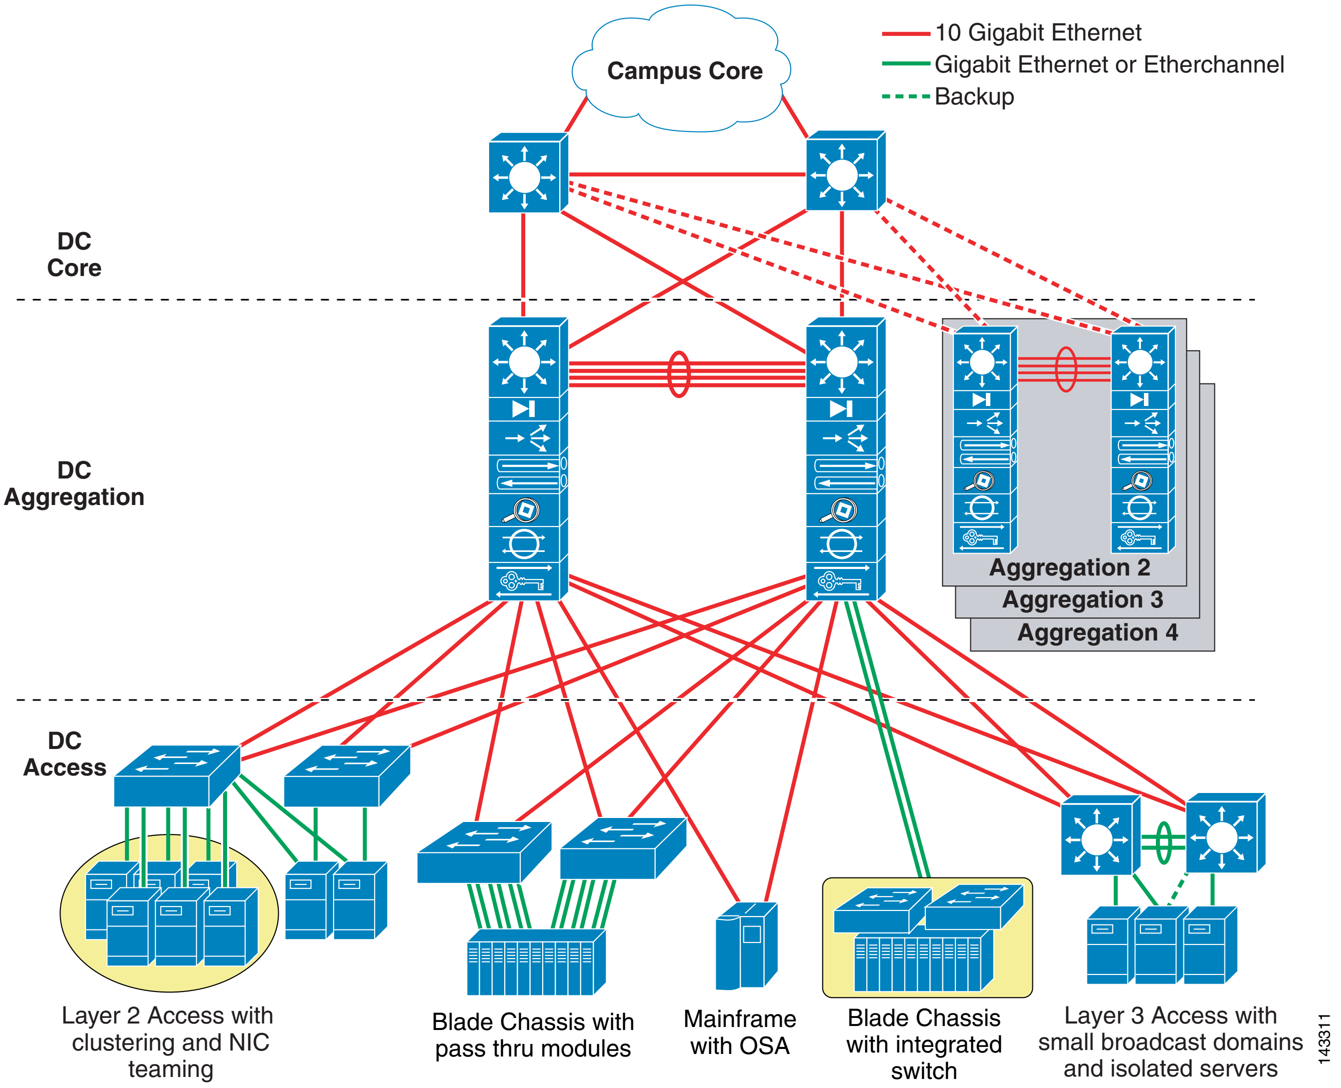
Campus Core
10 Gigabit Ethernet
Gigabit Ethernet or Etherchannel
Backup
DC Core
DC Aggregation
Aggregation 2
Aggregation 3
Aggregation 4
DC Access
Layer 2 Access with clustering and NIC teaming
Blade Chassis with pass thru modules
Mainframe with OSA
Blade Chassis with integrated switch
Layer 3 Access with small broadcast domains and isolated servers
143311

Core Layer

The core layer is a gateway that provides high-speed connectivity to external entities such as the WAN, intranet, and extranet of the campus. The data center core is a Layer 3 domain where efficient forwarding of packets is the fundamental objective. To this end, the data center core is built with high-bandwidth links (10 GE) and uses routing best practices to optimize traffic flows.

Aggregation Layer

The aggregation layer is a point of convergence for network traffic that provides connectivity between server farms at the access layer and the rest of the enterprise. The aggregation layer supports Layer 2 and Layer 3 functionality, and is an ideal location for deploying centralized application, security, and management services. These data center services are shared across the access layer server farms, and provide common services in a way that is efficient, scalable, predictable, and deterministic.

The aggregation layer provides a comprehensive set of features for the data center. The following devices support these features:

- Multilayer aggregation switches
- Load-balancing devices
- Firewalls
- Wide area application
- Acceleration
- Intrusion detection systems
- Content engines
- Secure Sockets Layer (SSL) offloaders
- Network analysis devices

Access Layer

The primary role of the access layer is to provide the server farms with the required port density. In addition, the access layer must be a flexible, efficient, and predictable environment to support client-to-server and server-to-server traffic. A Layer 2 domain meets these requirements by providing the following:

- Layer 2 adjacency between servers and service devices
- A deterministic, fast converging, loop-free topology

Layer 2 adjacency in the server farm lets you deploy servers or clusters that require the exchange of information at Layer 2 only. It also readily supports access to network services in the aggregation layer, such as load-balancers and firewalls. This enables an efficient use of shared, centralized network services by the server farms.

In contrast, if services are deployed at each access switch, the benefit of those services is limited to the servers directly attached to the switch. Through access at Layer 2, it is easier to insert new servers into the access layer. The aggregation layer is responsible for data center services, while the Layer 2 environment focuses on supporting scalable port density.

Layer 3 access designs are not widely deployed in current data centers. However, to minimize fault domains and provide rapid convergence, network administrators are seeking to leverage the benefits of Layer 3. Layer 3 designs do not exclude the introduction of network services, but the transparency of the service at the aggregation layer is more difficult to maintain. As with all access layer designs, the requirements of the application environments drive the decision for either model. The access layer must provide a deterministic environment to ensure a stable Layer 2 domain regardless of its size. A predictable access layer allows spanning tree to converge and recover quickly during failover and fallback.

Back-End Network

The back-end SAN consists of core and edge SAN storage layers to facilitate high-speed data transfers between hosts and storage devices. SAN designs (see Figure 17) 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.

Figure 17 SAN Topology

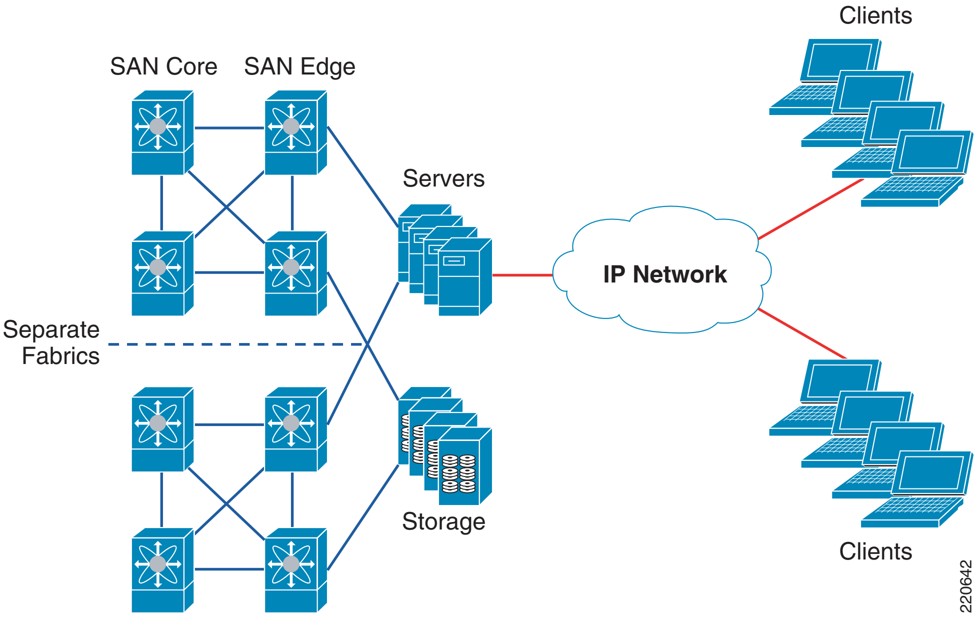

SAN Core Layer

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.

Note For more information on Cisco data center designs or other places in the network, see the following URL: http://www.cisco.com/go/srnd.

Branch Network Components

The enterprise branch provides remote users connectivity to corporate resources such as the centralized application services residing in the enterprise data center. The architectural design of the enterprise branch varies depending on the availability, scalability, security, and other service requirements of the organization.

The Cisco enterprise branch architecture framework defines the network infrastructure, network services, and application optimization capabilities of three typical branch deployment models. Figure 18 shows these three common branch solutions. Each of these profiles provides varying degrees of scalability and resiliency in addition to integrated network and application services.

Figure 18 Network Infrastructure Layer —Three Models

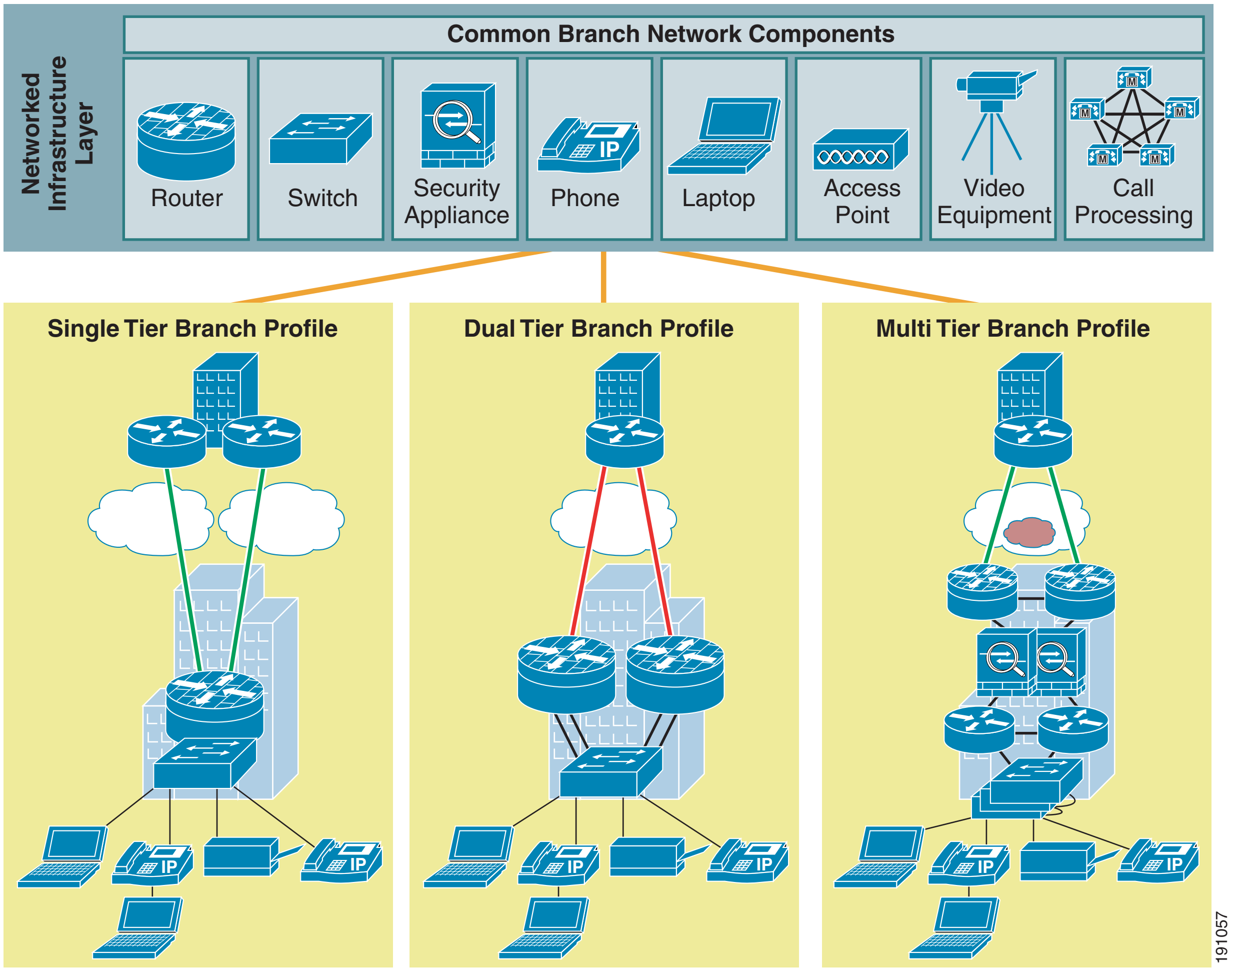

Note This document does not focus on enterprise branch design. For more information on Cisco data center designs or other places in the network, see the following URL: http://www.cisco.com/go/srnd.

Multisite Data Center Components

The Data Center Network Components section discussed the design of a single data center site consisting of front-end and back-end networking structures. This single-site design model delivers high availability, scalability, security and application optimizations within a solitary data center. Today, there is a larger requirement and growing trend to support the globalization of the enterprise, where multiple data center sites not only meet the local service level agreements and operational obligations of a single site, but also support the expansion of applications across locals. This extension of the enterprise creates a new set of conditions and possibilities for the network and application architects to address including:

- Site-to-site recovery
- Multisite load distribution

The introduction of multiple data centers into the enterprise provides site redundancy, removing the single point of failure a lone enterprise data center presents. Typically, one or more identical peer data centers are geographically dispersed to account for both natural and man-made incidents that may negatively influence business operations. The enterprise leverages these redundant resources for site failover and site recovery, it is a standby model. Applications normally contained within the primary data center element persist within the spare during adverse conditions. Site-to-site failover and recovery is a business continuance and disaster recovery strategy encompassing user to site, application to site and data to site rescue.

Multisite load distribution is an enterprise data center solution that builds upon the spare capacity introduced with multiple data centers poised to provide site-to-site recovery. To optimize the return on investment (ROI) and utilization of secondary data centers, the enterprise will seek to maintain applications across sites and distribute workloads efficiently. The enterprise architects must design a distributed network and application environment that meets the requirements of the user, the application, the data and most importantly the business.

Multisite enterprise data centers should address the following requirements:

- Local redundancy at each site for both network and application elements
- Application and data redundancy between sites
- Application availability via DNS and/or routing (IGP/BGP and RHI)
- Service provider redundancy (recommended)
- Site selection and stickiness technology
- Security (pervasive)

For more information on multisite data center, refer to the *Data Center Site Selection for Business Continuance* at www.cisco.com/go/srnd.

Design and Implementation Details

Design Goals

The enterprise network is a platform constructed to support a myriad of business functions; more specifically, applications. The traditional perception of the network relegates its role to one of data transport, providing a reliable fabric for the enterprise. This is a fundamental responsibility of the network infrastructure and should be enhanced rather than neglected. In addition to transport, the ubiquitous nature of the enterprise network fabric allows the introduction of intelligent network services to support business applications. This evolution of the network as an enterprise service platform is natural and supports the following application objectives:

- High availability
- Scalability
- Security
- Optimization

The Cisco data center architecture is a holistic approach that allows the network and the applications it supports to work together. The primary goals of this design are to increase the performance, availability, scalability, and manageability of enterprise applications within and between data centers, while simultaneously providing a secure environment. In addition, this design reduces the complexity and implementation time of enterprise applications in the data center using virtualization technologies and network design best practices. The remainder of this document focuses on each of these objectives when deploying a Microsoft Exchange application using the services of the Cisco data center infrastructure and Cisco empowered branch solutions.

Enterprise Data Center Design

Figure 19 represents the multisite test deployment of Microsoft Exchange 2007. Each data center site provides a highly available and robust network and network-based services for the local Exchange environment. The introduction of multiple data centers extends the n-tier Exchange application model between sites. Therefore, the network must address the state of the application tier at each site to provide user, application and data recovery services. To this end, the design leverages the following technologies:

- Site selection technologies combining DNS with high availability services such as KAL-AP available on the GSS.
- Route Health Injection (RHI) available on the ACE (also available on the CSM and CSS platforms).
- Layer 2 extension via L2TPv3 (pseudowire).

Figure 19 Multisite Data Center Testbed

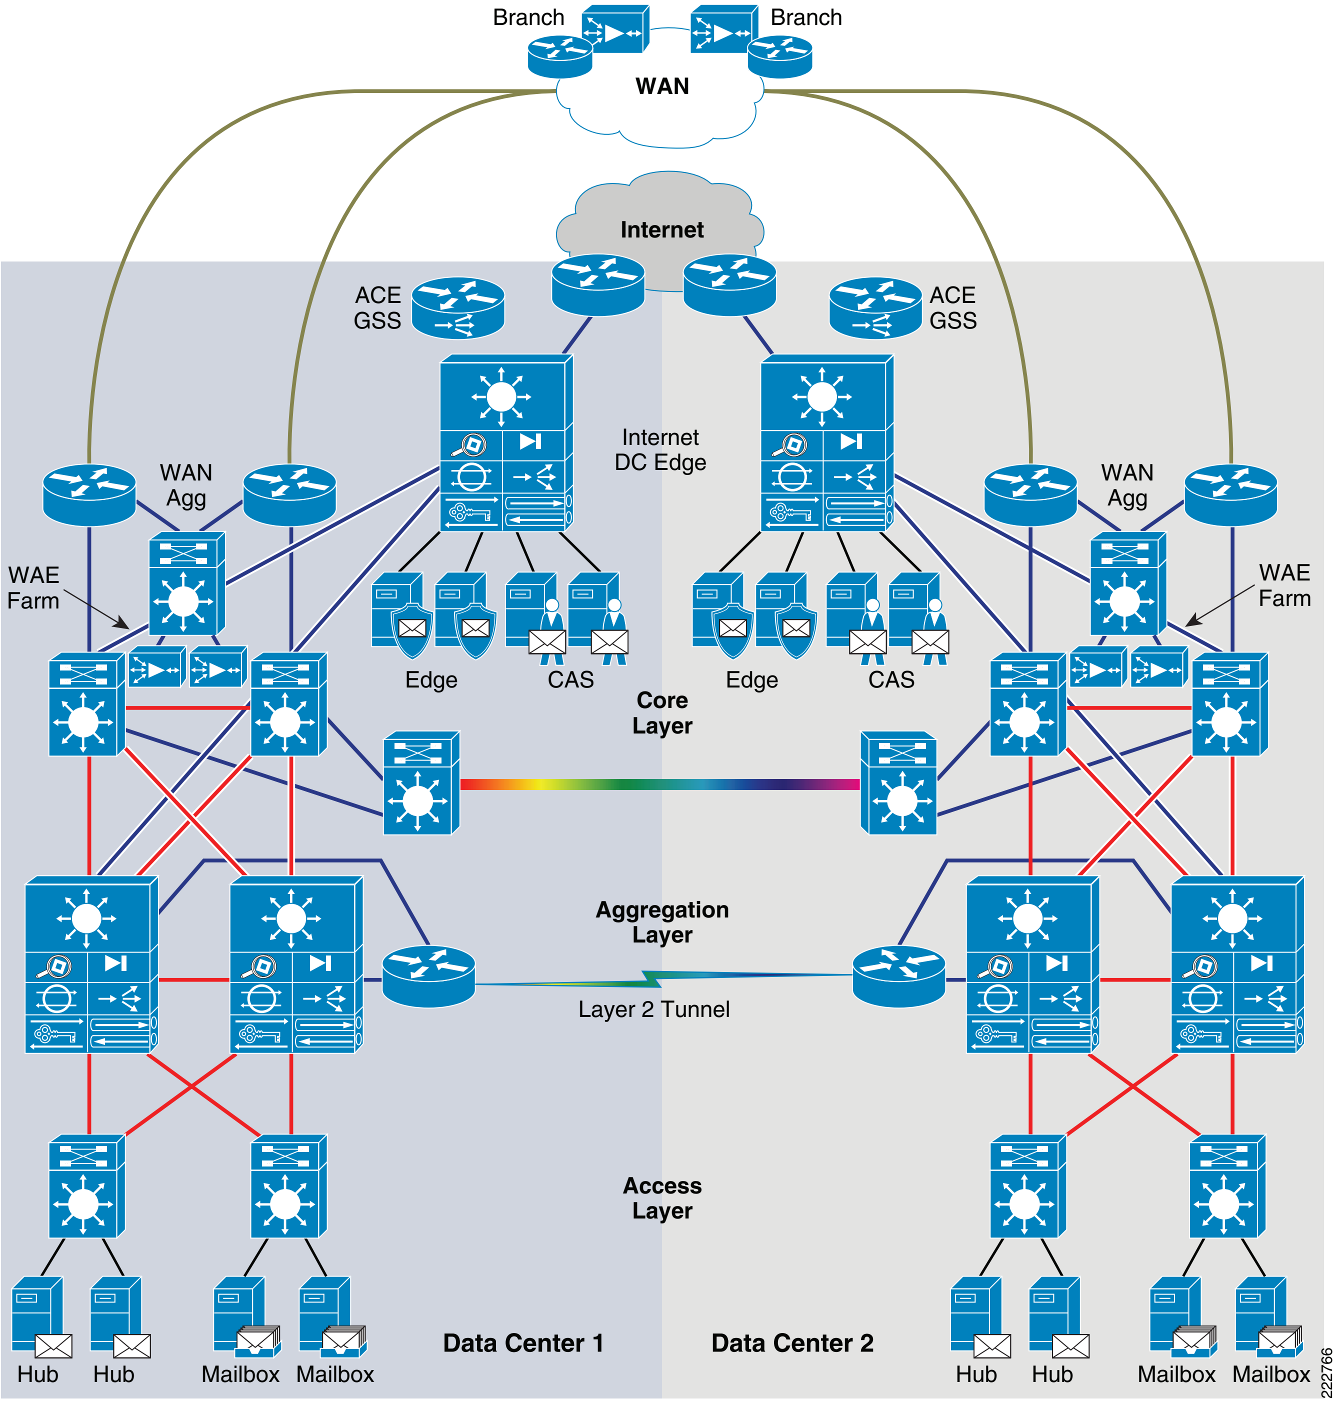

Site Selection

Site selection (or content routing) provides user recovery services associating a single user with an application instance. In Figure 19, the Cisco ACE Global Site Selector (GSS) provides this DNS-based service for the Exchange environment. The GSS appliances are the external DNS authoritative Name Servers for the enterprise providing A and MX records for the domain. The GSS appliances are deployed at the edge of the network as a cluster. Clustering the GSS boxes across multiple enterprise sites provides for a highly available and scalable DNS solution. It is important to note that each GSS houses an instance of Cisco's Network Registrar (CNR) that supplies the Mail Exchanger (MX) Resource Records to properly route Internet mail into the enterprise.

Typically, enterprise deployments leverage DNS-based round-robin and multiple MX records to distribute load across multiple Internet mail servers. This method provides redundancy and scalability but relies heavily on the client mail application to resolve DNS and MX records for mail delivery. Figure 20 is a snapshot of Microsoft's mail record configuration. Microsoft.com has three MX records referencing three different SMTP servers, mail a, b and c at Microsoft.com. Note that there is equal priority, in this case 10, given to each of these exchanges. MX preferences of equal value are the equivalent of round-robin load balancing. Lower MX preference values receive higher priority from the clients' mail applications. For example, an MX preference of 10 takes priority over an MX preference value of 20.

Figure 20 Example Enterprise MX Record Deployment

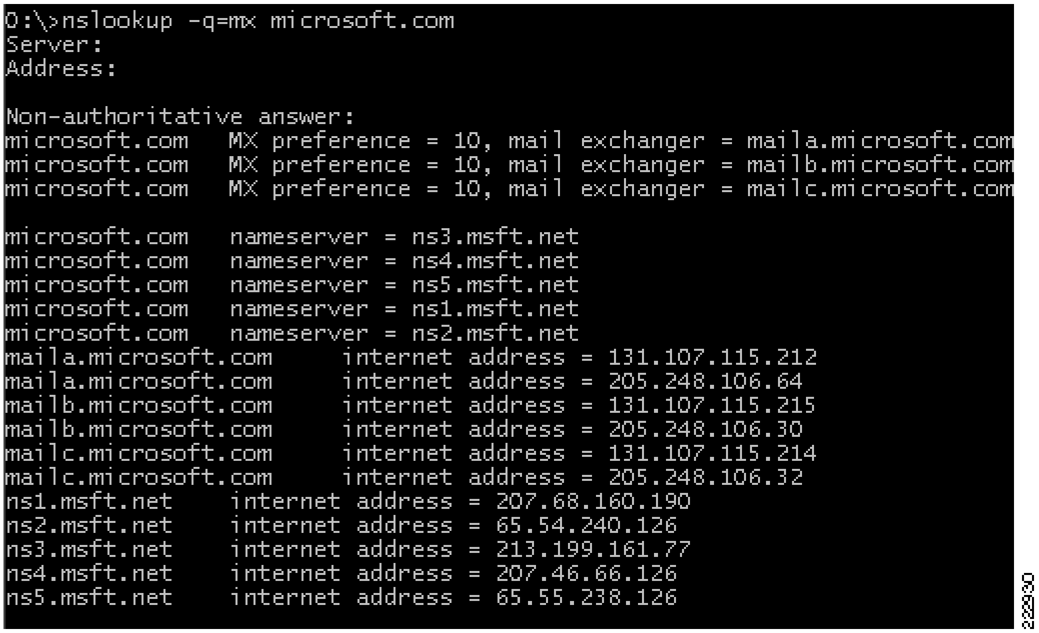

```
O:\>nslookup -q=mx microsoft.com
Server:
Address:

Non-authoritative answer:
microsoft.com   MX preference = 10, mail exchanger = maila.microsoft.com
microsoft.com   MX preference = 10, mail exchanger = mailb.microsoft.com
microsoft.com   MX preference = 10, mail exchanger = mailc.microsoft.com

microsoft.com   nameserver = ns3.msft.net
microsoft.com   nameserver = ns4.msft.net
microsoft.com   nameserver = ns5.msft.net
microsoft.com   nameserver = ns1.msft.net
microsoft.com   nameserver = ns2.msft.net
maila.microsoft.com     internet address = 131.107.115.212
maila.microsoft.com     internet address = 205.248.106.64
mailb.microsoft.com     internet address = 131.107.115.215
mailb.microsoft.com     internet address = 205.248.106.30
mailc.microsoft.com     internet address = 131.107.115.214
mailc.microsoft.com     internet address = 205.248.106.32
ns1.msft.net    internet address = 207.68.160.190
ns2.msft.net    internet address = 65.54.240.126
ns3.msft.net    internet address = 213.199.161.77
ns4.msft.net    internet address = 207.46.66.126
ns5.msft.net    internet address = 65.55.238.126
```

To provide a higher level of scalability and availability for inbound and outbound mail processing, administrators may choose to load-balance across multiple ET roles. A load-balancer, such as the Cisco ACE, provides high availability and scalability within a single data center and is able to communicate the state of the ET roles to the Cisco GSS that are globally aware. The GSS may probe the ACE VIP at each site using multiple probe types. Combining these probes allows the GSS to gain better insight into the state of the ET roles located at each data center, providing a more intelligent form of SMTP server (i.e., site selection).

Traffic Pattern

Figure 21 below illustrates the DNS traffic patterns and probing technology of the GSS. In this example, the GSS is actively probing the Cisco ACE service modules located at two distinct data centers in the enterprise. The Cisco GSS is aware of the ET application state. It reflects this knowledge in its resolution of DNS requests.

In this example, the GSS probes the ACE virtual context VIP hosting the ET roles using KAL-AP. As defined earlier, KAL-AP provides load and VIP state information for each VIP maintained within a context hosting a KAL-AP agent. To enable KAL-AP agent on the ACE, use the following commands:

```
kalap udp
  ip address <IP address> encryption md5 <password>
class-map type management match-any <MANAGEMENT>
  2 match protocol kalap-udp any
policy-map type management first-match <P-MANAGEMENT>
  class <MANAGEMENT>
    permit
interface vlan <VLAN>
  description ** Public Facing Interface **
  service-policy input <P-MANAGEMENT>
```

The IP address of the KAL-AP agent should match the address assigned to the ACE device or alias IP address. Apply a service policy to the GSS-facing, usually public-facing interface to permit the communication.

The GSS configuration requires one to define a shared keepalive of type KAL-AP and associate it with the KAL-AP agent defined on the ACE via the primary or secondary IP address fields, Figure 21 shows an example of this relationship. The 10.210.1.4 is an IP address defined on the bridged virtual interface of the virtual ACE context.

Figure 21 Example GSS Configuration of Shared KAL-AP

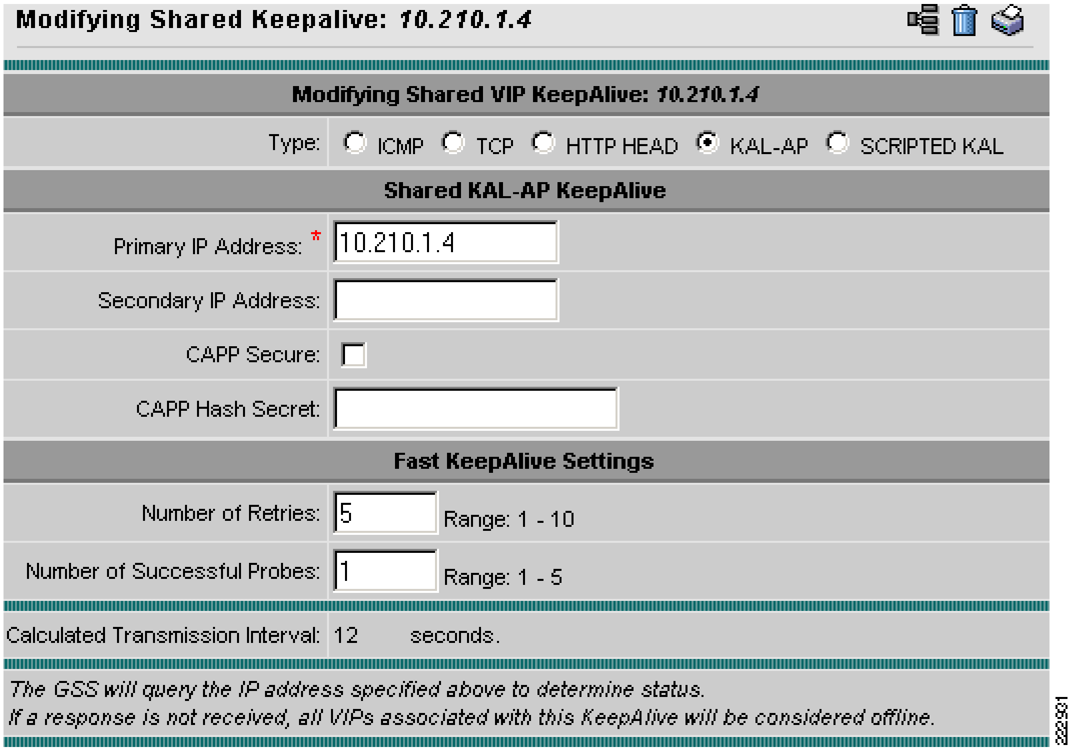

Modifying Shared Keepalive: *10.210.1.4*

Modifying Shared VIP KeepAlive: *10.210.1.4*

Type: ○ ICMP ○ TCP ○ HTTP HEAD ● KAL-AP ○ SCRIPTED KAL

Shared KAL-AP KeepAlive

Primary IP Address: *	10.210.1.4
Secondary IP Address:	
CAPP Secure:	□
CAPP Hash Secret:	

Fast KeepAlive Settings

Number of Retries:	5 Range: 1 - 10
Number of Successful Probes:	1 Range: 1 - 5
Calculated Transmission Interval:	12 seconds.

The GSS will query the IP address specified above to determine status.
If a response is not received, all VIPs associated with this KeepAlive will be considered offline.

KAL-AP VIP configuration screen (see Figure 22) allows one to set variables specific to the KAL-AP probe or one may leverage global parameters, timers and security settings, as shown in Figure 23. It is recommended to secure all communication between the GSS and ACE devices, to force this practice use the CAPP Hash Secret globally on the GSS.

Note The CAPP Hash Secret is equivalent to the *kalap udp* agent password on the ACE configuration.

Figure 22 Sample KAL-AP VIP KeepAlive Configuration

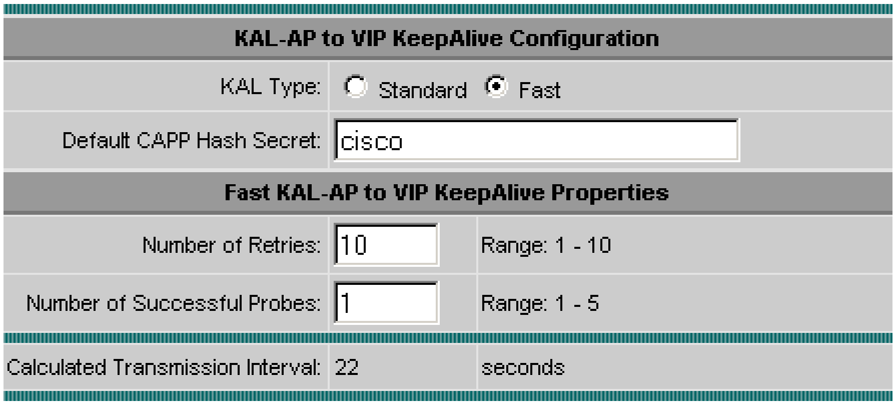

KAL-AP to VIP KeepAlive Configuration

KAL Type:	○ Standard ● Fast	
Default CAPP Hash Secret:	cisco	

Fast KAL-AP to VIP KeepAlive Properties

Number of Retries:	10	Range: 1 - 10
Number of Successful Probes:	1	Range: 1 - 5
Calculated Transmission Interval:	22	seconds

To allow the GSS to apply the KAL-AP knowledge to an associated A record representing the ACE VIP, assign the KAL-AP keepalive to the Answer (see Figure 23). Note that the VIP Address and KAL-AP keepalive reference IP addresses defined on the ACE context.

Figure 23 Example GSS Keepalive to Answer Configuration

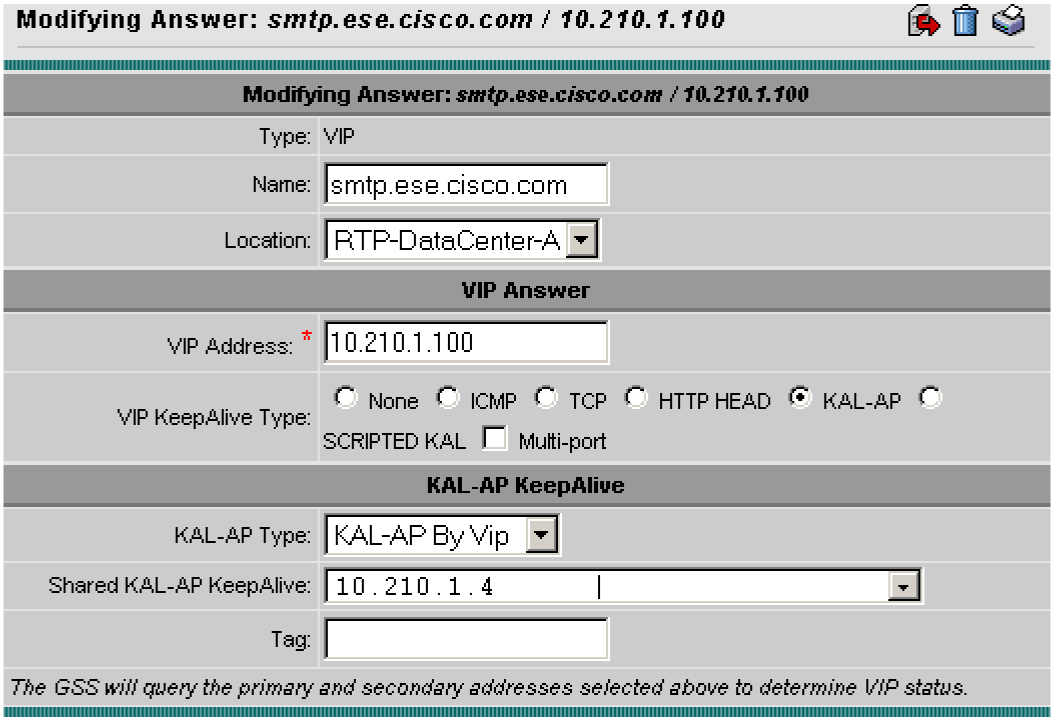

Modifying Answer: *smtp.ese.cisco.com / 10.210.1.100*

Modifying Answer: *smtp.ese.cisco.com / 10.210.1.100*

Type:	VIP
Name:	smtp.ese.cisco.com
Location:	RTP-DataCenter-A

VIP Answer

VIP Address: *	10.210.1.100
VIP KeepAlive Type:	○ None ○ ICMP ○ TCP ○ HTTP HEAD ● KAL-AP ○ SCRIPTED KAL □ Multi-port

KAL-AP KeepAlive

KAL-AP Type:	KAL-AP By Vip
Shared KAL-AP KeepAlive:	10 . 210 . 1 . 4
Tag:	

The GSS will query the primary and secondary addresses selected above to determine VIP status.

With the above GSS and CNR configurations completed, basic traffic flow can be tested. Figure 24 illustrates the DNS traffic patterns and probing technology of the GSS. In this example, the GSS is actively probing the Cisco ACE service modules located at two distinct data centers in the enterprise.

Figure 24 GSS MX Record Request

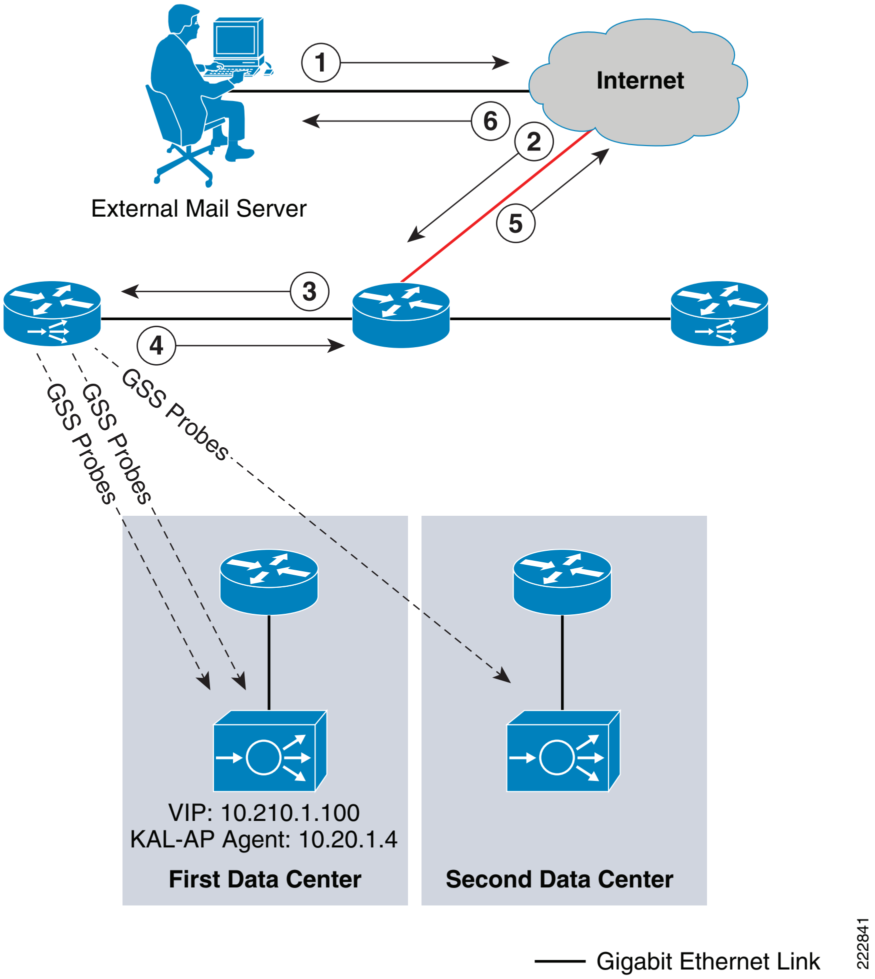

The following steps outline the external mail server requests for MX records:

Step 1 The external mail server issues a standard DNS query of type MX for domain ese.cisco.com (the test domain).

Step 2 The router receives the UDP packet and forwards to the GSS the authoritative name server for the domain.

Step 3 The GSS receives the DNS request. The CNR sub-process receives the MX request for the ese.cisco.com zone.

Step 4 The CNR responds with the resource record type of MX for ese.cisco.com. This resource record configuration directs all incoming mail for ese.cisco.com to the smtp.ese.cisco.com server. The smtp.ese.cisco.com server resolves in DNS to the ACE VIP. Figure 25 shows the configuration for the MX record within the CNR.

Step 5 The router forwards the packet.

Step 6 The external mail server receives the MX record from the Cisco GSS, the authoritative name server for the domain. See Figure 26.

Note The Cisco CNR can be configured to provide multiple MX records for redundancy.

Figure 25 CNR MX Record Configuration

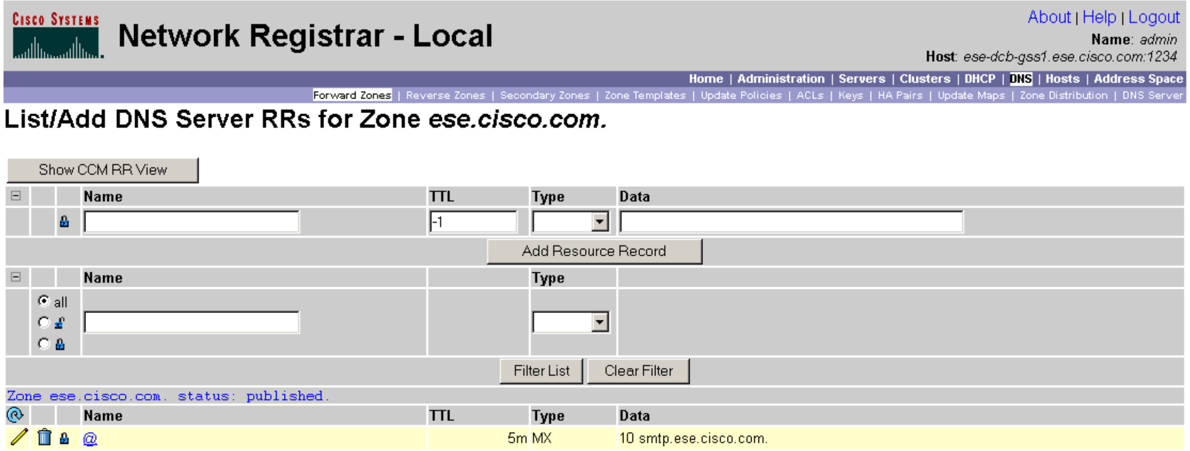
Cisco Systems
Network Registrar - Local
About | Help | Logout
Home | Administration | Servers | Clusters | DHCP | DNS | Hosts | Address Space
Forward Zones | Reverse Zones | Secondary Zones | Zone Templates | Update Policies | ACLs | Keys | HA Pairs | Update Maps | Zone Distribution | DNS Server
Show CCM RR View
Add Resource Record
all
Filter List
Clear Filter
222934

Name: admin
Host: ese-dcb-gss1.ese.cisco.com:1234

List/Add DNS Server RRs for Zone *ese.cisco.com.*

Name	TTL	Type	Data
	-1		

Zone ese.cisco.com. status: published.

Name	TTL	Type	Data
@	5m	MX	10 smtp.ese.cisco.com.

Figure 26 Example nslookup for ese.cisco.com

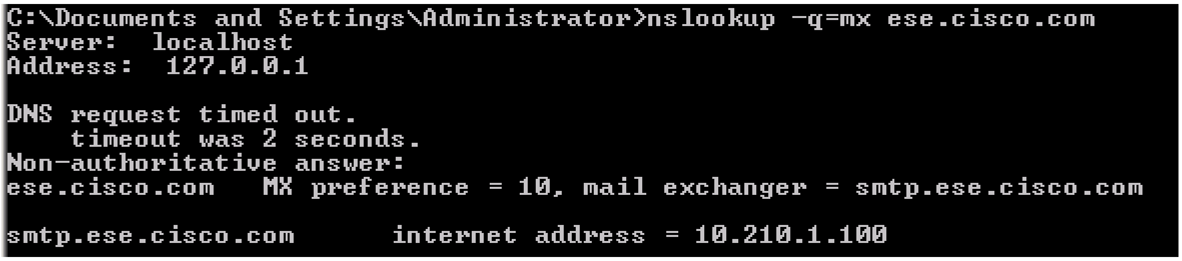

```
C:\Documents and Settings\Administrator>nslookup -q=mx ese.cisco.com
Server:  localhost
Address:  127.0.0.1

DNS request timed out.
    timeout was 2 seconds.
Non-authoritative answer:
ese.cisco.com   MX preference = 10, mail exchanger = smtp.ese.cisco.com

smtp.ese.cisco.com      internet address = 10.210.1.100
```

The MX record points to smtp.ese.cisco.com, which is a VIP on the ACE. The DNS name resolves to an IP address provided by the Cisco GSS as the authoritative name server for the enterprise. Per the example, the ACE VIP is returned by the Cisco GSS as an A record to the client. It is important to note the GSS only sends A records that are ‘active” according to its probes. The Cisco GSS may leverage local ACE application knowledge and its own keepalives to monitor the health of a site. This ensures the client will be directed to an active site from an application perspective and the DNS rules (availability, proximity, load, preference etc.) determined by the network administrator.

This DNS-based solution details how the Cisco GSS, CNR and ACE are able to provide a cohesive solution for Exchange ET site selection. Other Exchange servers such as the CAS that provide web-based access to remote users mailboxes may leverage the same type of site selection intelligence without the Resource Record features provided via the Cisco CNR.

Route Health Injection

Route Health Injection (RHI) allows the ACE to advertise host routes associated with any number of active virtual IP addresses hosted by the device. The injection of the host route to the remaining network offers Layer 3 availability and convergence capabilities to the application environment. In the Exchange test bed, the ACE advertises the VIP front-ending the CAS server farm to the Multilayer Switch Feature Card (MSFC) routing table. The MSFC distributes the route into the IGP, in this case OSPF, as a host route. At the edge of the Wide Area Network (WAN) OSPF is redistributed into BGP where the branch routers are made aware of the ACE VIP.

The following configuration is necessary on the ACE module to enable RHI:

```
!Define the VIP to be advertised
class-map match-all CAS_VIP

!10.211.1.100 is the VIP associated with the CAS server farm listening on port 80
  2 match virtual-address 10.211.1.100 tcp eq www

!This statement will enable RHI for the VIP defined under the CAS_VIP class
policy-map multi-match PM_CAS

  class CAS_VIP
    loadbalance vip inservice
    loadbalance vip advertise active
```

The host route is present on the MSFC, via the **show ip route | in 10.211.1.100** command:

```
S       10.211.1.100/32 [77/0] via 10.211.1.4, Vlan111
```

Route Health Injection deployments in the data center are described in Chapter 5 of the *Data Center—Site Selection for Business Continuance:*

http://www.cisco.com/en/US/netsol/ns656/networking_solutions_design_guidances_list.html#anchor3

Layer 2 Extension

As discussed in Tested Microsoft Exchange Server 2007 Deployment Models, page 20, the Exchange mailbox clusters have certain expectations of the network, namely Layer 2 adjacency for the public and private interfaces of the mailbox servers. In the tested Exchange 2007 scenario, L2TPv3, was the transport chosen to meet the Exchange 2007 application requirements as L2TPv3 allows for the transport of many Layer 2 protocols over an IP infrastructure. L2TPv3 is a Cisco solution for transporting Layer 2 packets over an IP network.

Figure 27 depicts the L2TPv3 deployment in the Exchange multisite testing. The Core layer consists of Catalyst 6509 with Supervisor 720 connected via Ten Gigabit Ethernet providing a high-speed fabric within and between data centers. At the Aggregation layer, the Catalyst 6509s provide Layer 2 and 3 services and uplink connectivity to the Access layer housing the Exchange 2007 Hub and Mailbox servers. To meet the requirements of Layer 2 adjacency for the Exchange mailbox cluster, a pair of Cisco 7200 routers reside in the Aggregation layer and are used to deploy L2TPv3. The Cisco 7200 platform has an NPE-G2 engine and is running the Cisco IOS version 12.4(15)T1.

Figure 27 L2TPv3 Deployment

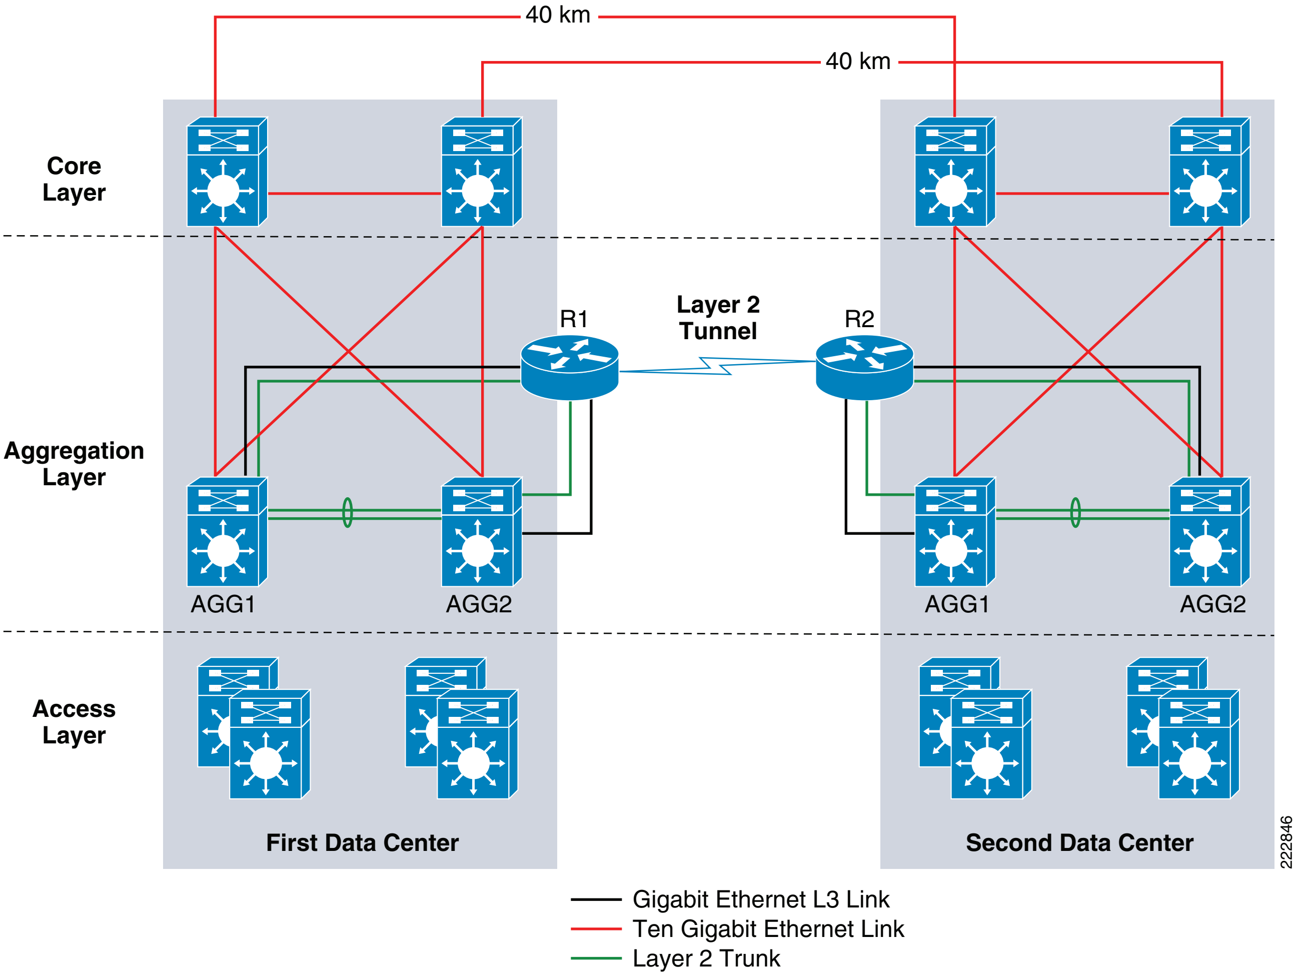

The L2TPv3 pseudowire is constructed using static per VLAN tunneling. This allowed for granular control of the traffic tunneled between the two sites and limited the domain creep occurring in the spanning tree environments of each data center. The only VLANs or subnets permitted across this tunnel were the public and private VLANs 74 and 76 respectively that are associated with the Exchange Mailbox cluster. Dynamic L2TPv3 sessions are also possible allowing for auto negotiation of the pseudowire between sites. The following configuration defines the control channel parameters manually, a static configuration:

```
l2tp-class ESE
 hello 10
 password 7 1511021F0725
 cookie size 8

pseudowire-class ESE1
 encapsulation l2tpv3
 protocol l2tpv3 ESE
 ip local interface Loopback0


interface Loopback0
 ip address 10.151.2.54 255.255.255.255
```

```
! Loopback referenced by remote pseudowire peer and local class
```

To enable this xconnect service, the Layer 2 Cisco 7200 sub-interfaces are configured as follows:

```
interface GigabitEthernet0/3.74
 encapsulation dot1Q 74
 xconnect 10.151.1.19 123 encapsulation l2tpv3 pw-class ESE1
!IP Address of remote peer and unique channel ID.  Pseudowire class <ESE1> defined
!previously.

interface GigabitEthernet0/3.76
 encapsulation dot1Q 76
 xconnect 10.151.1.19 456 encapsulation l2tpv3 pw-class ESE1
```

Note For more information including requirements, restrictions and configuration on L2TPv3 go to: http://cco/en/US/products/sw/iosswrel/ps1829/products_feature_guide09186a008016108b.html.

Enterprise Edge Design

This section focuses on the deployment of Edge and CAS roles at the enterprise edge of the data center. As previously mentioned, the Edge and CAS Exchange server roles are excellent candidates to take advantage of network-based services such as load balancing, SSL offload, firewall and other application optimization services. Figure 28 below highlights the flexibility and some basic design options afforded to the enterprise leveraging Cisco's virtual service technologies. These design options provide centralized load balancing, security, and optimization services for the application. In addition, the virtualization capabilities of both the FWSM and the ACE allow a single physical device to provide multiple logical devices to support a variety of application environments. System administrators can assign a single virtual device to a business unit or application to achieve application performance goals, requirements, or service-level agreements.

In Figure 28, the FWSM and ACE virtual contexts are logically positioned north and south of one another. Choice 'A' places the FWSM virtual context as a traditional firewall controlling access to the DMZ and network services cascading to the south. Choice 'B' leverages the firewall security services available in the Cisco ACE to provide protection and other application services. The FWSM context below the ACE context is another layer of protection into the internal enterprise. The final option depicted, choice 'C', relies on the security and application services of the ACE virtual context to protect and optimize the Edge of the enterprise network.

Figure 28 Enterprise Edge Design Choices

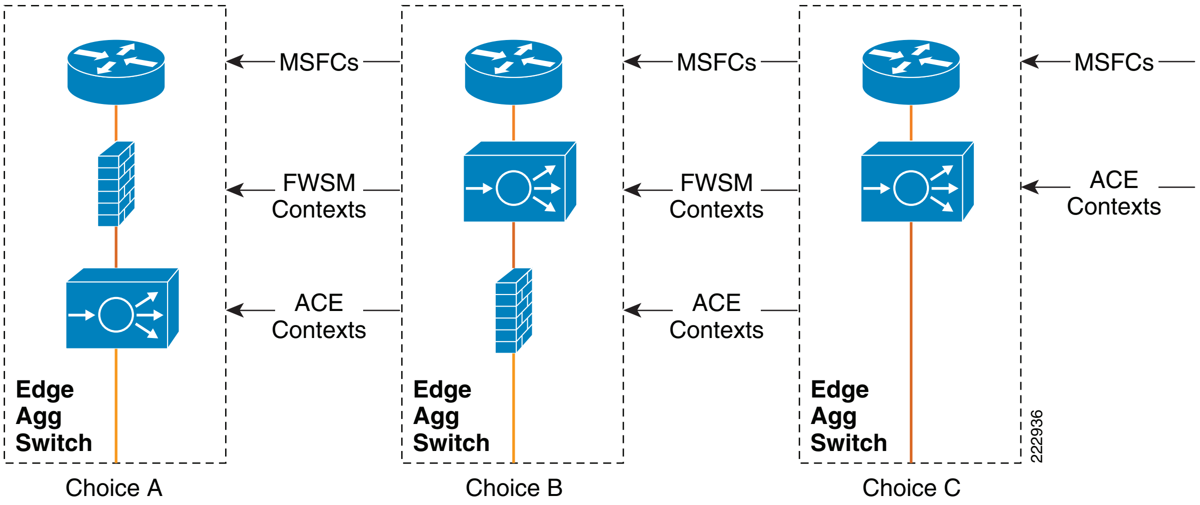

Figure 29 is a sample DMZ deployment for an Exchange 2007 application environment. In Figure 29, each Catalyst 6500 edge aggregation switch is an ingress point to the enterprise providing transport, application, and security services. The ACE virtual contexts residing in the Catalyst platform support firewall functionality, load balancing, and applications acceleration. The Microsoft Exchange 2007 CAS and Edge roles are positioned behind dedicated ACE virtual context explicitly configured to support each role. The following two sections detail the deployment of CAS and Edge roles behind their dedicated ACE virtual contexts.

Figure 29 Sample DMZ Deployment

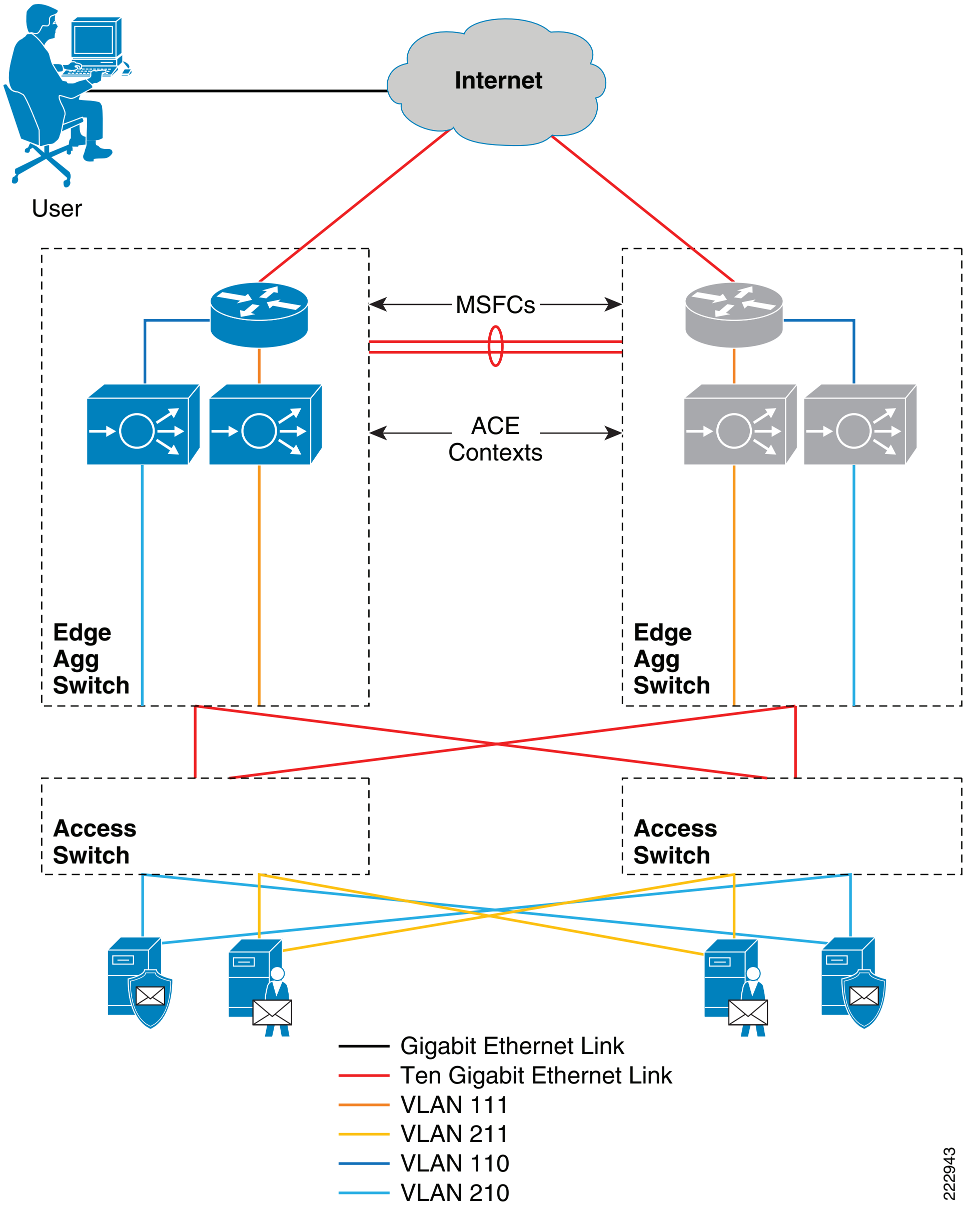

Client Access Server Role

The CAS role can benefit from the security, scalability and availability services of a Cisco ACE virtual context. In the tested topology, the ACE virtual contexts are deployed in transparent mode, meaning the ACE context is bridging traffic between VLANs. Figure 29 shows the public, or client facing VLAN as VLAN 111 and the internal or server facing VLAN as 211. The ACE creates a bridged virtual interface, BVI, to associate the two VLANs to one another. The CAS roles point to the MSFC as their default gateway. The ACE virtual context supports the CAS server farm by providing the following services:

- Security
- Load balancing

- Session persistence
- Offload services including SSL and TCP multiplexing
- Health monitoring
- Site selection

The design and configuration of these services in the ACE is dependent on the type of CAS traffic introduced.

WAN Optimization Traffic Patterns

As previously defined, the CAS role provides access to the Microsoft Exchange environment via all non-MAPI forms of communication. These different types of client applications and protocols involved are detailed in Table 2.

Table 2 Client Access Server Role Communication

Client Type	Communication Protocol
IMAP Client	IMAP4 / SMTP
POP Client	POP/ SMTP
Outlook Voice Access	RTP
Exchange ActiveSync	HTTP/HTTPS
Outlook Anywhere	MAPI over RPC over HTTP/HTTPS
Outlook Web Access (OWA)	HTTP/HTTPS
Outlook	MAPI over RPC

OWA allows users to have a robust or "thick-client" application experience within a web browser. Microsoft continues to introduce new features and functionality to this product that has standardized on the HTTP protocol for conducting remote transactions. Cisco technologies such as ACE and WAAS are well positioned to provide network-based services in an OWA environment because Microsoft leverages this standard form of communication.

Outlook Anywhere allows the Outlook client application to access its designated mailbox server using RPC over HTTP/HTTPS. This technology provides efficient access to users through enterprise firewall deployments. The CAS role supports this functionality, as do the Cisco ACE and WAAS technologies.

The following section details the flow of OWA or Outlook Anywhere traffic from an enterprise branch to the CAS roles residing in the enterprise edge of the data center. These traffic patterns include:

- Egress Client to Enterprise Edge
- Ingress Enterprise Edge to CAS Role
- Egress CAS Role to Enterprise Edge
- Egress Enterprise Edge to Client

This section details the transparent integration of WAAS technology.

Egress Client to Enterprise Edge

Clients in the branch using OWA or Outlook Anywhere services should consider using Cisco's application acceleration technologies, such as Cisco WAAS. Cisco WAAS requires a minimum of two WAE devices to auto-discover and deliver applicable application optimizations. To leverage these transparent optimizations across the WAN, deploy one or more WAEs at the remote branch and one or more WAEs at the enterprise data center, depending on availability and scalability requirements.

Note For more information on Cisco WAE branch deployments, see the *Enterprise Branch Wide Area Application Services Design Guide* at the following URL: http://www.cisco.com/application/pdf/en/us/guest/netsol/ns477/c649/ccmigration_09186a008081c7d5.pdf

Figure 30 Example Branch Traffic Pattern

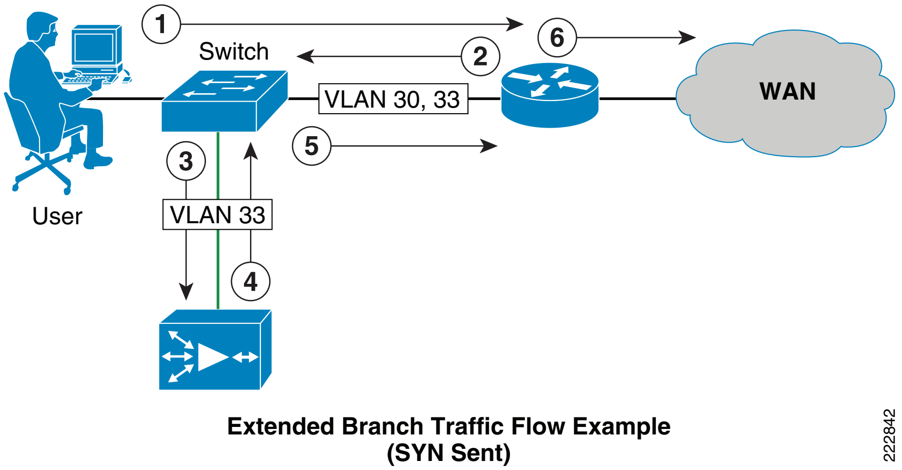

Figure 30 depicts the traffic pattern an OWA or Outlook 2007 client used to reach the Microsoft Exchange 2007 deployment in the data center, specifically the CAS role residing at the enterprise edge. The users must cross the WAN, a potential application bottleneck with its inherent latency and bandwidth challenges. To address this issue, the Cisco WAE is deployed at the remote branch. The following steps describe the initiation of communication between the Exchange client and the CAS role in the data center:

Step 1 The OWA or Outlook Anywhere client initiates a TCP connection via SYN packet to the ACE VIP front-ending the CAS server farm. The VIP is advertised via RHI from the ACE.

Step 2 The branch router transparently intercepts the TCP SYN using WCCPv2. WCCPv2 makes a load-balancing decision and the router Layer 2 redirects the flow to a specific WAE device in the service group. Service group 6 is defined with the following command and applied to the LAN interface of the router:

```
!WCCP service group will only intercept and redirect TCP flows
ip wccp 61
!Note the use of trunking on the LAN interface to segment and control traffic to the WAE
!versus traffic destined for the clients.
interface GigabitEthernet0/1.30
 description ** BRANCH DATA VLAN **
 encapsulation dot1Q 30
 ip address < IP address >
 ip wccp 61 redirect in

interface GigabitEthernet0/1.33
 description ** BRANCH WAE VLAN **
 encapsulation dot1Q 33
 ip address < IP address >
 ip wccp redirect exclude in
```

Note Use WCCPv2 ACLs to only redirect traffic destined for the WAN. Traffic confined to the branch LAN would not benefit from WAE services and would only introduce more load on the local WAE branch devices.

Step 3 The branch switch forwards the packet to the WAE device.

Step 4 The WAE device applies a new TCP option (0x21) to the packet if the application is identified for optimization by an application classifier. The WAE adds its device ID and application policy support to the new TCP option field. This option is examined and understood by other WAEs in the path as the ID and policy fields of the initial WAE device. The initial ID and policy fields are not altered by another WAE.

Note It is important to note that the Cisco WAE device has default optimizations enabled for HTTP/HTTPS or web traffic. Encrypted traffic will not benefit from all of the optimization techniques offered by the ACE. Compression and data redundancy elimination optimizations will be greatly reduced if not eliminated when encrypted traffic is processed by the WAE. However, the transport flow optimizations of the Cisco WAAS solution will continue to offer benefits including large initial TCP window sizes, scaling, and congestion handling techniques to reduce data loss.

Step 5 The branch switch forwards the packet to the branch router that is configured as the default gateway for the WAE devices. Note that the source and destination IP addresses of the initial request of the remote client remain unchanged.

Step 6 The router forwards the packet to the CAS VIP address across the WAN.

Ingress Enterprise Edge to CAS Role

Ingress traffic across the WAN destined for the Exchange CAS server farm will cross the WAN edge routers. Figure 31 clarifies the subsequent actions that occur for application optimization of the OWA or Outlook Anywhere clients.

Figure 31 Ingress Traffic from the WAN to the Data Center

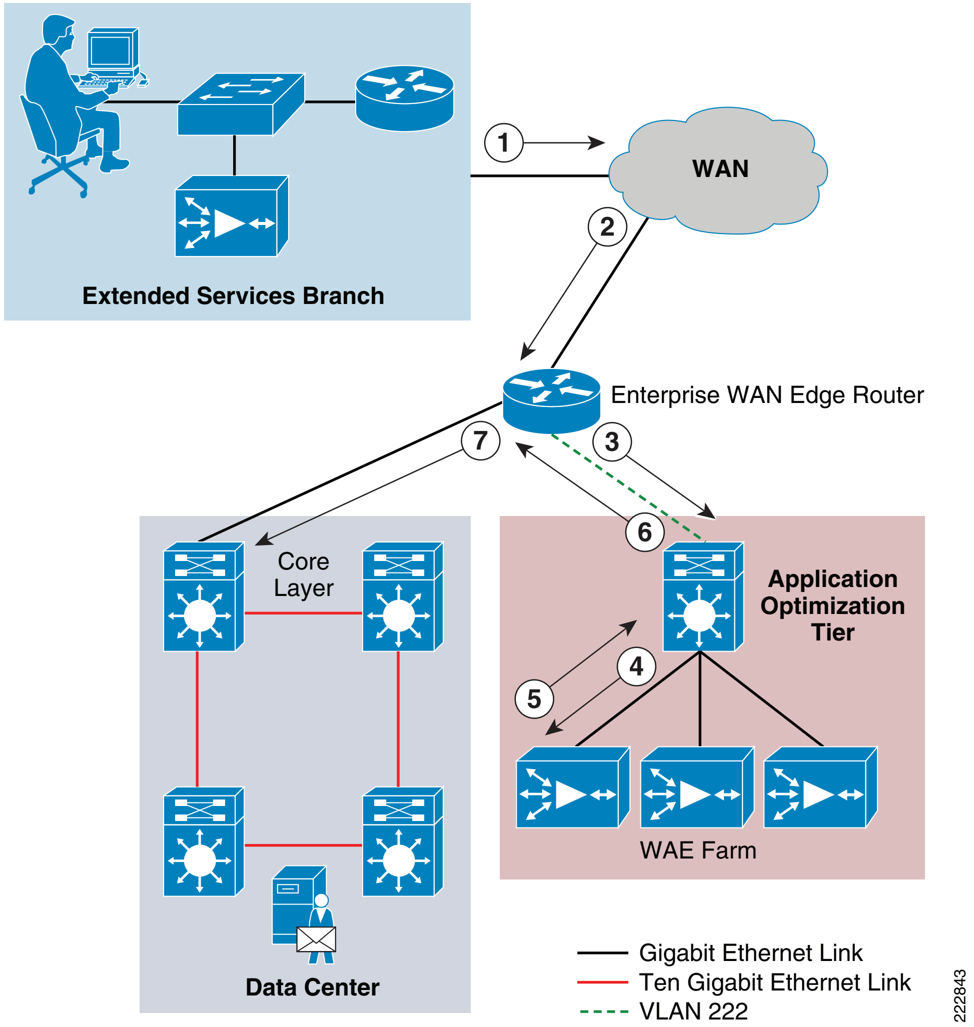

In Figure 31, the following takes place:

Step 1 Traffic exits the branch with the WAE TCP option (0x21) set.

Step 2 Traffic enters the WAN edge router with a destination IP address of the CAS VIP. The router forwards the packet to an application optimization-tier hosting a WAE farm via WCCPv2. The following commands are necessary on the WAN Edge Router to enable WCCPv2 redirection:

```
!Enable the WCCP service.
ip wccp 61
interface GigabitEthernet0/1
 description <<** Interface to WAN **>>
 ip address 192.168.22.1 255.255.255.0
 ip wccp 61 redirect in
!Enable HSRP or GLBP as a default gateway for the WAE Farm
interface GigabitEthernet0/3
 description <<** Interface to WAE Resources **>>
 ip address 10.222.1.3 255.255.255.0
 glbp 1 ip 10.222.1.1
 glbp 1 load-balancing host-dependent
```

The **ip wccp 61** service group command instructs the router to inspect traffic for any TCP packets. Matching packets should be load-balanced amongst service-group attached WAEs using assignment based on the source IP address of the packet. Inbound redirection combined with CEF is recommended to reduce the resource utilization of the router.

Step 3 The dedicated WAE VLAN carries the redirected traffic from the router to a specific WAE via MAC rewrite.

Step 4 The WAE receives the SYN containing the TCP option introduced by the branch WAE. The local data center WAE is aware of the branch WAE and what optimizations it supports. The WAE appends its ID and application policy support to the existing TCP option 0x21 field while storing the ID and application policy of the remote branch WAE. The WAE is registered to the WCCP service group in the router via the following commands:

```
!Register the WAE with the WCCP service groups in the WAN Edge Router
wccp router-list 1 <IP Address of WAN Edge Router>

!The routers in the list will use L2 redirection (MAC address rewrite) to send interesting
!traffic to the WAE
wccp tcp-promiscuous router-list-num 1 l2-redirect assign-method-strict
wccp version 2
```

Step 5 The WAE returns the packet to the router based on the egress method configured on the WAE, in this case via GRE. The WCCP protocol allows for the auto-negotiation of the GRE tunnel. The following command enables the GRE functionality:

```
egress-method negotiated-return intercept-method wccp
```

Note The default form of WAE egress traffic forwarding is IP-based. Typically, one uses IP forwarding when the WAE farms have a dedicated subnet, where as GRE based forwarding is leveraged when the WAE farm resides within the client and or server subnet. Use the **show egress-method** command on a WAE device to review the policy.

Step 6 The switch forwards the packet to the initiating WCCPv2 router.

Step 7 The router routes the packet to the ACE CAS VIP via Layer 3.

Transparent Service Integration, page 63 for ACE-to-CAS traffic flows.

Egress CAS Role to Enterprise Edge

One of the primary tenets to a well-designed network infrastructure is the ability to control traffic flow across the enterprise. From branch-to-data center or data center-to-data center, predictable traffic patterns are a hallmark to the stability and availability of an application. Traffic returning from the enterprise edge CAS roles with a symmetric path out of the data center will benefit from the application optimization offered at the WAN edge, providing acceleration to the HTTP/HTTPS transport of OWA and Outlook Anywhere clients.

Figure 32 defines the flow of traffic from the data center core to the branch in the test bed's Exchange 2007 environment. It is important to recognize that the traffic returning from the core is directed to the same WAE on egress as on ingress to enable WAN optimization.

Figure 32 Return CAS Traffic to the WAN Edge

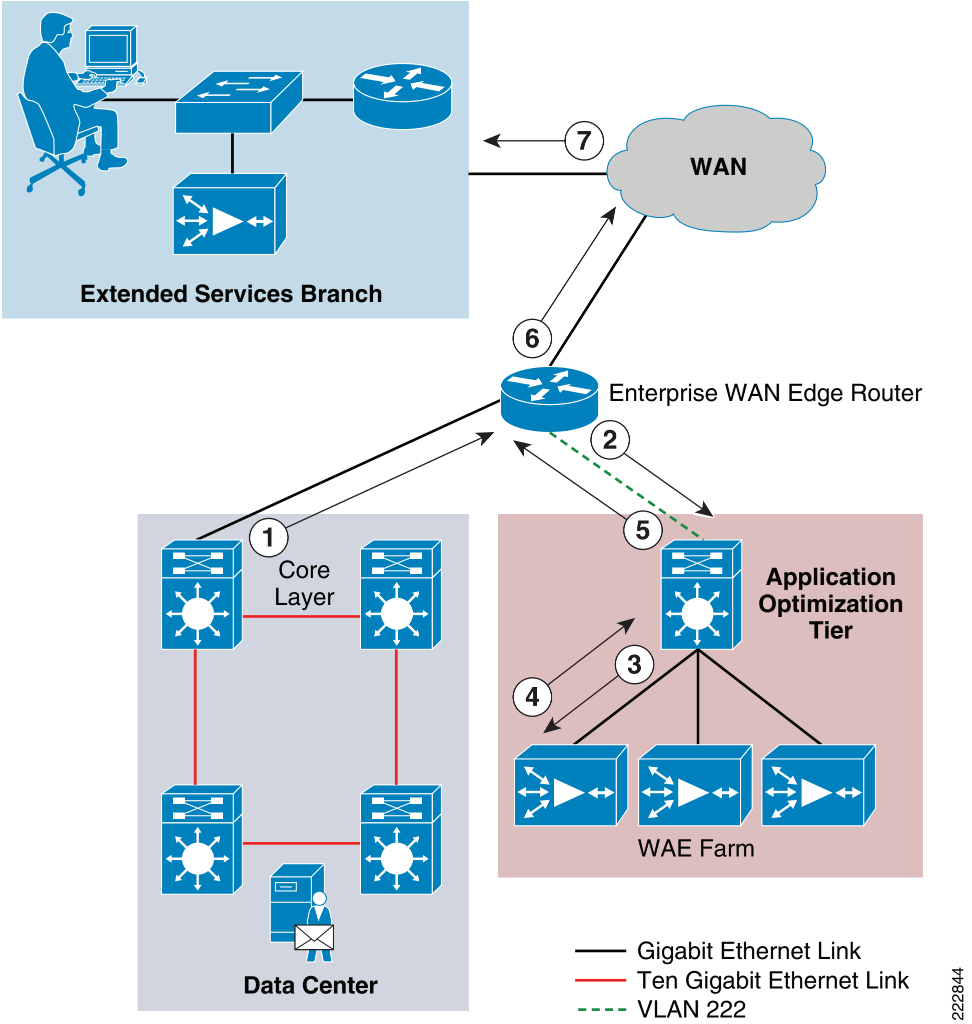

As shown in Figure 32, the following flow occurs from core to branch:

Step 1 SYN/ACK returned from the CAS serverfarm is routed to the WAN edge router.

Step 2 The router receives the packet on its internal interface with WCCPv2 enabled, service group 62. The following commands are necessary to enable this service on the 7200 router platform.

```
!Enable WCCP service goup
ip wccp 62
!Apply the service to the ingress interface from the data center core
interface GigabitEthernet0/2
 description <<** Interface to Data Center Core  **>>
 ip address 10.199.0.22 255.255.255.252
 ip wccp 61 redirect in
```

The **ip wccp 62** service group command instructs the router to inspect traffic for any TCP packets. Matching packets should be load-balanced amongst service-group attached WAEs using assignment based on the destination IP address of the packet.

Step 3 The packet returns to the same WAE, maintaining symmetric traffic flow and WAN optimization consistency.

Step 4 The WAE receives the SYN/ACK and adds TCP option 0x21 to the packet with its ID and application policy support. The packet is returned via the egress-method policy negotiated via WCCP to the intercepting router.

HTTP/HTTPS traffic is optimized by default on the Cisco WAE platform. OWA and Outlook Anywhere traffic fall in this category.

Step 5 The application optimization tier switch forwards the packet to the WAN edge router.

Step 6 The packet is routed to the branch.

Step 7 The branch router receives the packet.

Egress Enterprise Edge to Client

Figure 33 depicts the return traffic for TCP flows coming from the data center. The return of traffic to the WAE device located in the branch will complete the auto-negotiation and begin WAN optimization for the OWA and Outlook Anywhere clients.

Figure 33 Returning Traffic from the Data Center to Branch Client

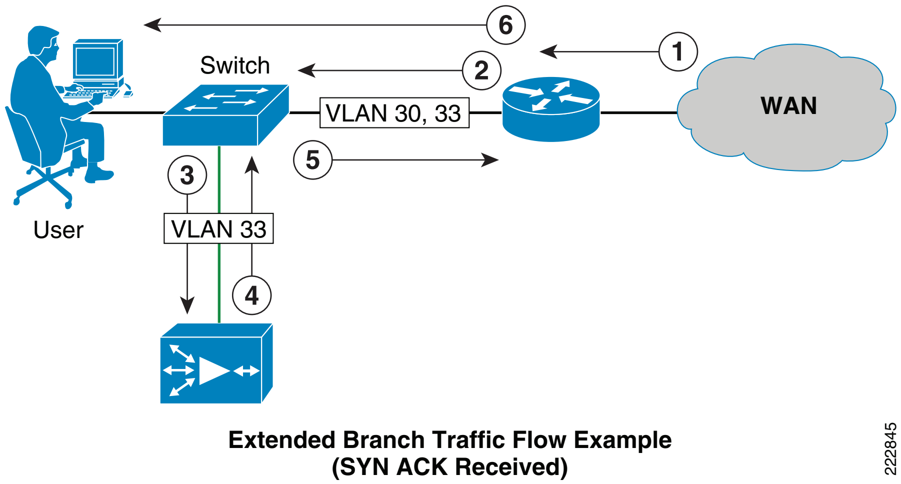

The flow in Figure 33 shows the following:

Step 1 The branch router receives the packet on the interface with a WCCP service group defined for inbound traffic interception.

```
interface GigabitEthernet0/0
 description ** WAN interface **
 ip wccp 62 redirect in
```

Step 2 Traffic is sent to the WAE VLAN on the LAN interface that is defined as a trunk.

Step 3 The WAE device receives the TCP flow. The WAE is aware of the WAE in the data center because the SYN/ACK TCP option 0x21 contains an ID and application policy. The auto-negotiation of the policy occurs as the branch WAE compares its application-specific policy to that of its remote peer defined in the TCP option. At this point, the data center WAE and branch WAE have determined the application optimizations to apply on this specific TCP flow.

The **show tfo connection** command on the branch or data center WAEs details the results of auto-negotiation for each TCP flow. Below is an example output for an SSL connection between an OWA SSL client and the CAS VIP. This view shows that only TCP optimizations are occurring for the HTTPS transaction , the four Ts indicate that TFO optimization is occurring at the local and remote WAE

```
Optimized Connection List
Policy summary order: Our's, Peer's, Negotiated, Applied
F: Full optimization, D: DRE only, L: LZ Compression, T: TCP Optimization
Local-IP:Port          Remote-IP:Port         ConId    PeerId            Policy
192.168.30.15:28073    10.211.1.100:443       17174    00:14:5e:a4:4f:66 T,T,T,T
```

Step 4 The traffic is sent to the router via a GRE tunnel.

Step 5 The switch forwards the traffic to the router.

Step 6 The router receives the traffic from the WAE VLAN and routes the packet to its final destination the client. The WAE VLAN is excluded from WCCP redirection via the **ip wccp redirect exclude in** command on the interface configuration.

Transparent Service Integration

This section discusses the integration of the Cisco ACE service module with the Microsoft CAS role. As previously mentioned, the CAS roles can leverage load balancing, health monitoring, and security features of the Cisco ACE. Figure 34 illustrates the tested deployment of the ACE and CAS roles. In this scenario, a virtual ACE context is deployed in transparent, or bridge mode, joining VLANs 111 and 211 logically. The VIP of the ACE abstracts the CAS roles located in the server farm.

Traffic Pattern

Figure 34 depicts the flow of OWA or Outlook Anywhere traffic in and out of the enterprise edge. It should be noted that the WAN optimization of Cisco WAAS technology is optional in this environment. The ACE context is unaware of these application services occurring across the WAN.

Figure 34 Transparent Service Integration with CAS

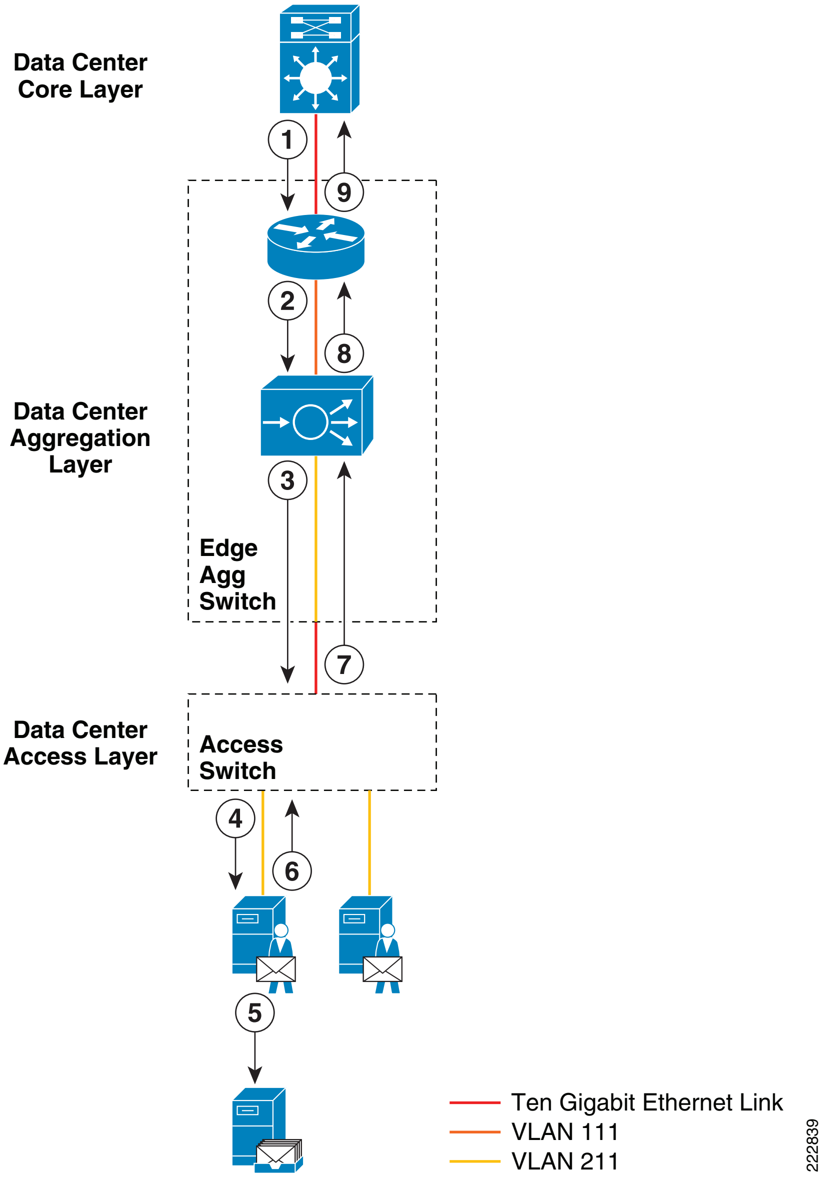

The integration of Microsoft Exchange's CAS roles and the ACE in transparent mode includes the following steps:

Step 1 The data center core routes incoming traffic to the ACE VIP.

Step 2 The MSFC of the aggregation layer receives the traffic and forwards it to the ACE VIP defined on VLAN 111 in this example. The ACE is in transparent mode and therefore leverages a bridged virtual interface to join VLANs 111 and 211. The following interface configuration is necessary:

```
interface vlan 111
  description ** Client side CAS **
!Joins the interface to the bridge group
  bridge-group 11
!This service policy defines the CAS VIP and the ACE contextual services it offers.
  service-policy input PM_CAS
  no shutdown

interface vlan 211
  description ** Server side CAS **
!Joins the interface to the bridge group
  bridge-group 11
  no shutdown

!This is the logical join between the two VLANs
interface bvi 11
  ip address 10.211.1.4 255.255.255.0
!The alias address is a virtual address shared between ACE contexts for redundancy
  alias 10.211.1.6 255.255.255.0
  description ** CAS Bridge Interface **
  no shutdown
```

Step 3 The ACE bridges the traffic to the CAS roles defined in its server farm after applying the associated load balancing, and security policies. The destination IP address becomes the selected CAS roles IP address, but the source address remains that of the originating client.

Note The policies defined within the ACE are discussed in ACE for OWA, page 65 and ACE for Outlook Anywhere, page 70.

Step 4 The CAS role receives the flow with the clients IP address.

Step 5 The CAS role communicates directly with the appropriate mailbox server for the user.

Step 6 The CAS sends flow back to its default gateway the MSFC.

Step 7 The ACE receives the packet on the VLAN 211 interface transparently bridging the communication from the CAS to the MSFC.

Step 8 The ACE forwards the flow to the MSFC, changing the source IP address from the CAS role to its VIP. The client is unaware of the “real” CAS role IP address.

Step 9 The MSFC routes the flow back to the client.

ACE for OWA

The CAS role supports browser based user connectivity to the Microsoft Exchange environment via a web browser. This thin client implementation has become increasingly popular and more functional for remote users. The CAS role services are available via Microsoft’s Internet Information Server (IIS) as a virtual directory named “owa”, therefore port 80 and 443 are the default points of access to the environment. In the tested configuration, the ACE module front-ends the environment providing the following services:

- SSL-offload
- Access control
- Load balancing
- Session Persistence
- Health Monitoring
- Server Redirect

To better understand the services provided via the ACE module in an OWA environment, it is best to look at an example. Figure 35 depicts an OWA logon session and will be used to describe the use of ACE services for OWA.

Note In this example, the client attempts to connect via HTTP but it is redirected to HTTPS as defined by the security policies of the enterprise. In addition, the Cisco ACE is positioned as an SSL proxy, offloading the SSL processing from the CAS role.

Figure 35 Example CAS OWA Logon

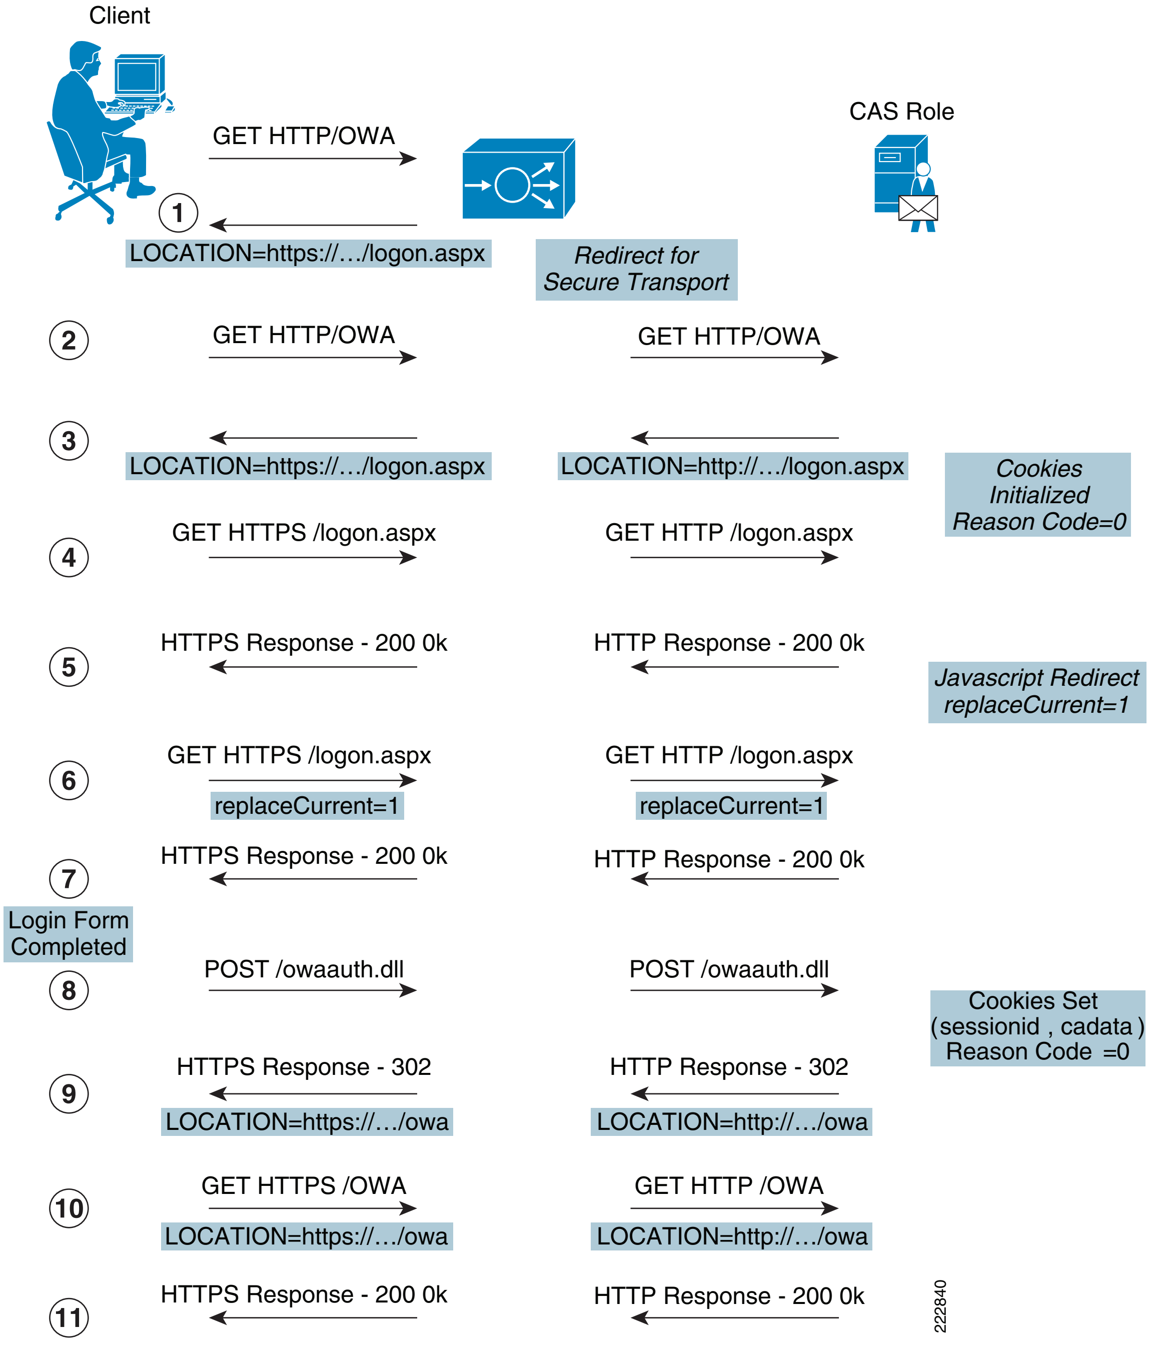

In Figure 35, the following flow occurs:

Step 1 The client browser attempts to access the OWA application at http://ese.cisco.com/owa in DNS this resolves to the ACE VIP defined in a class-map for ports 80 (HTTP) and 443 (HTTPS).

```
class-map match-all CAS-VIP
  2 match virtual-address 10.211.1.100 tcp eq www
```

```
class-map match-all CAS-VIP-443
  2 match virtual-address 10.211.1.100 tcp eq https
```

The policy map named PM_CAS leverages these virtual address definitions as class maps via the following commands:

```
policy-map multi-match PM-CAS
  class CAS-VIP
!Enable the virtual IP address
    loadbalance vip inservice
!Redirect load balance policy from port 80 to 443
loadbalance policy CAS-REDIRECT
    loadbalance vip icmp-reply active
!Enable RHI
    loadbalance vip advertise active
!Enable TCP reuse feature of ACE
    appl-parameter http advanced-options HTTP-PARAMS

  class CAS-VIP-443
!Enable the virtual IP address
    loadbalance vip inservice
!Session persistence and load balance policy
    loadbalance policy STICKYLB
    loadbalance vip icmp-reply active
!Enable RHI
    loadbalance vip advertise active
!Enable TCP reuse feature of ACE
    appl-parameter http advanced-options HTTP-PARAMS
!Enalbe ACE to perform SSL Services
    ssl-proxy server OWA
```

The client side interface, in this case VLAN 111, uses the PM-CAS policy map:

```
interface vlan 111
  description ** Client Facing Interface **
  bridge-group 11
  no normalization
!Apply policy map to interface
  service-policy input PM_CAS
  no shutdown
```

The policy map PM-CAS has a load-balance policy name CAS_REDIRECT under the class CAS_VIP. This class defines the VIP listening under port 80. The system administrator, in this example, wants all incoming traffic to leverage secure transport services, i.e. SSL. The ACE redirects the client via an HTTP 302 message to use the following LOCATION as defined in the HTTP header, https://owa.ese.cisco.com/owa. The CAS_REDIRECT policy uses a class of default allowing it to catch all traffic destined to the VIP on port 80.

```
policy-map type loadbalance first-match CAS-REDIRECT
!Catch all port 80 traffic to the CAS-VIP and send to OWA redirect serverfarm
  class class-default
    serverfarm OWA-REDIRECT-FARM

!A single redirect server is defined
serverfarm redirect OWA-REDIRECT-FARM
  rserver OWARedirectSSL
    inservice

!The redirect server defines the correct URL to access the OWA services of Exchange
!via a 302 HTTP header location statement
rserver redirect OWARedirectSSL
  webhost-redirection https://owa.ese.cisco.com/owa/ 302
  inservice
```

Note The redirect server may also leverage HTTP 301 messages. HTTP 301 messages are typically given in response to client POSTs and may be rejected by some browsers if received after a client GET message.

Step 2 The client issues a secure HTTP GET for the new location at https://owa.ese.cisco.com/owa. This ACE policy map PM-CAS supports this VIP defined under the CAS-VIP-443 class map. In addition, the ACE provides SSL-offload for the CAS roles with the SSL proxy server configuration. The ACE negotiates a secure session between itself and the client. The following commands define the SSL proxy service:

```
ssl-proxy service OWA
!Identifies the RSA key used for the OWA certificate
  key newKey
!Defines the certificate issued by the trusted provider
  cert OWA.cer
```

Note For more information on configuring the ACE to support secure transactions, see the appendix.

In addition to the SSL proxy service, the ACE provides load balancing via a load-balance policy named STICKLB. The STICKLB policy provides both a load balancing and a session persistence mechanism. The default load-balancing algorithm is round-robin. The session persistence method, in this case, leverages the CAS generated cookie named **sessionid** against the OWA server farm named CAS-FARM. The configuration is as follows:

```
!Define the OWA server farm
serverfarm host CAS-FARM
!Probe leveraged to monitor the server farm
  probe iProbe
!Return code statistics maintained by the ACE for the server farm
  retcode 200 500 check count
!Real CAS servers listening on port 80
  rserver RTP-CAS-01 80
    inservice
  rserver RTP-CAS-02 80
    inservice

!HTTP cookie persistence using the cookie named “sessionid”
sticky http-cookie sessionid STCKY-SESSIONID-GRP
!Set Cookie inactivity timeout (in minutes)
  timeout 20
!CAS server available for service
  serverfarm CAS-FARM
```

Note The default inactivity cookie timeout for OWA is 20 minutes. In this example, the ACE cookie inactivity timeout is also set to 20 minutes to remain consistent with the Microsoft CAS role deployment.

```
!This policy map defines the load balancing for all SSL traffic destined to the OWA farm
!with session persistence based on the OWA cookie “sessionid”
policy-map type loadbalance first-match STICKYLB
  class class-default
    sticky-serverfarm STCKY-SESSIONID-GRP
```

The ACE sends clear text HTTP to one of the OWA servers defined in the server farm.

Step 3 The CAS role receives the HTTP GET from the ACE and begins the logon procedure. This process consists of initializing all OWA related cookies and redirecting the client to the logon.aspx page. This redirection occurs via a HTTP LOCATION header. The ACE forwards the message, replacing the CAS role IP address with the CAS VIP address.

Note The CAS role is aware of the SSL-offload functionality of the ACE. To configure support for SSL-offloading on a CAS role, refer to: http://technet.microsoft.com/en-us/library/bb885060.aspx

Step 4 The client requests the logon page.

Step 5 To verify the client supports Javascript, a heavily leveraged feature in OWA, the CAS role requests the client perform a Javascript request to receive the logon page with a *replaceCurrent* value equal to 1. If the client successfully sends this Javascript redirect the OWA server believes the client meets the necessary browser functionality.

Step 6 The client requests a page via Javascript.

Step 7 The OWA server returns the logon form.

Step 8 The client posts the logon form to the “owaauth.dll”.

Step 9 The CAS role authenticates the user via Active Directory and sets the sessionid cookie. The sessionid cookie is returned to the client and learned by the ACE in a HTTP 302 message. The ACE operating at Layer 7 inspects the sessions for cookie information. The following output shows the output of a single sessionid cookie being learned by the ACE:

```
show sticky database
sticky group : STCKY-SESSIONID-GRP
type         : HTTP-COOKIE
timeout      : 20            timeout-activeconns : FALSE
  sticky-entry          rserver-instance                 time-to-expire flags
  ---------------------+--------------------------------+-------------+--------+
  13208929475664964043  rtp-cas-01:80                    1177           -
```

Note Microsoft recommends cookie or source IP-based session persistence. The ACE supports both in addition to cookie insert. Cookie insert allows the ACE to insert a cookie into the HTTP header of the client-server conversation.

Step 10 The client redirects to the new location, ironically the same URL as the original request.

Step 11 The OWA server responds with the client’s home page.

ACE for Outlook Anywhere

Outlook Anywhere allows the thick client application Outlook to connect to the CAS roles via HTTP/HTTPS. This communication technique was formerly known as RPC over HTTP(s). Outlook Anywhere essentially tunnels RPC traffic over HTTP or HTTPS. Outlook Anywhere allows firewalls without RPC fix-up capabilities to provide access to the Outlook client using the well-known ports of 80 and 443. In addition, Outlook Anywhere allows for MAPI connectivity over networks where latency is an issue. Traditionally, MAPI communication has issues where the network latency exceeds 250 milliseconds. Tunneling this traffic over HTTP(s) resolves this problem.

The ACE is more than capable of supporting HTTP and HTTPS traffic and is well positioned to provide services to Outlook Anywhere enabled clients. Outlook is not a web browser. Outlook does not support cookies. In this light, Microsoft recommends the use of source IP-based load balancing when leveraging a load-balancer with Outlook Anywhere clients. In addition, it is highly recommended to use secure transport of RPC, (i.e., SSL).

The configuration is identical to the OWA ACE configuration except for the type of session persistence. For Outlook Anywhere, deploy a source IP-based sticky group. To configure source IP-based sticky using the full IP address of the client, use the following ACE commands:

```
sticky ip-netmask 255.255.255.255 address source SRC-STCKY-GRP
  serverfarm CAS_FARM
```

The Outlook Anywhere client shows the connection state as “HTTPS” as illustrated in Figure 36.

Figure 36 Outlook Anywhere via ACE

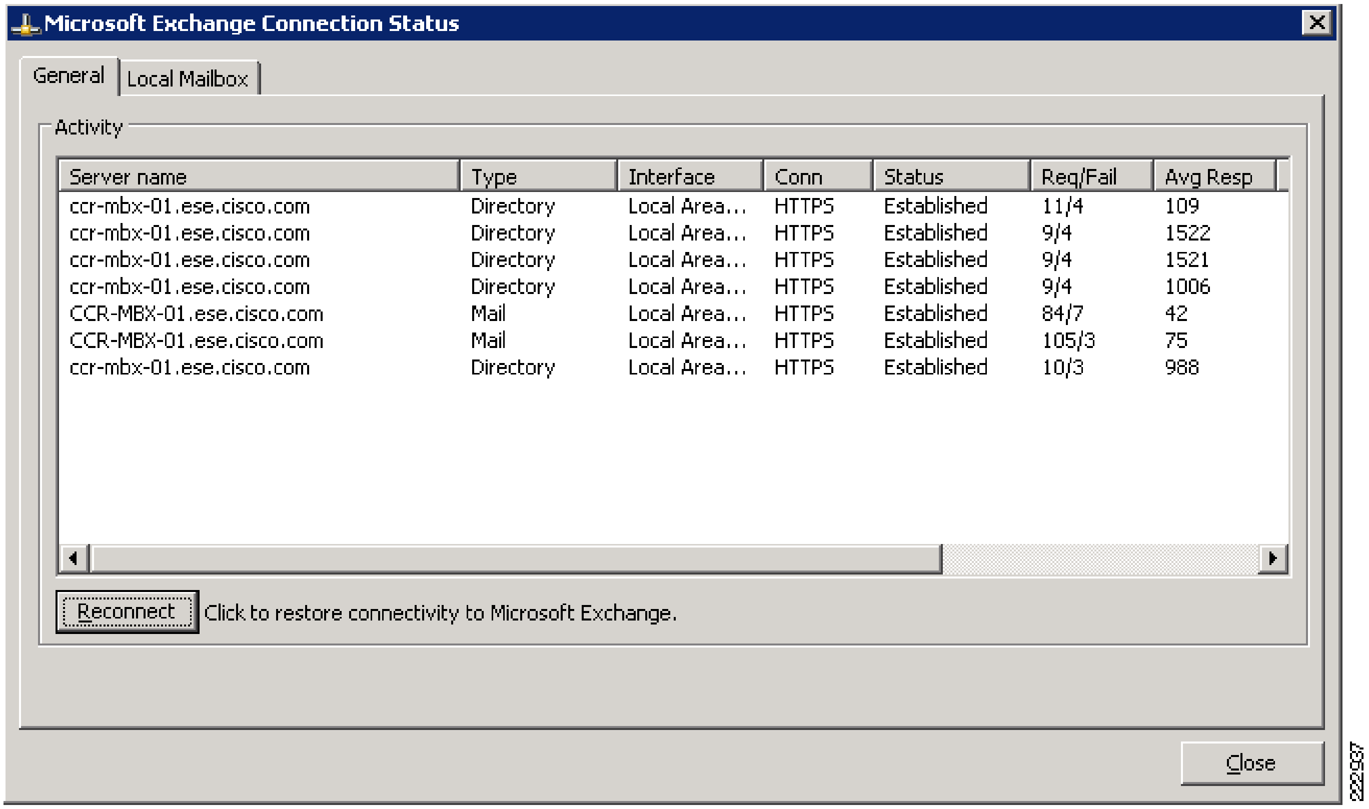

Server name	Type	Interface	Conn	Status	Req/Fail	Avg Resp
ccr-mbx-01.ese.cisco.com	Directory	Local Area...	HTTPS	Established	11/4	109
ccr-mbx-01.ese.cisco.com	Directory	Local Area...	HTTPS	Established	9/4	1522
ccr-mbx-01.ese.cisco.com	Directory	Local Area...	HTTPS	Established	9/4	1521
ccr-mbx-01.ese.cisco.com	Directory	Local Area...	HTTPS	Established	9/4	1006
CCR-MBX-01.ese.cisco.com	Mail	Local Area...	HTTPS	Established	84/7	42
CCR-MBX-01.ese.cisco.com	Mail	Local Area...	HTTPS	Established	105/3	75
ccr-mbx-01.ese.cisco.com	Directory	Local Area...	HTTPS	Established	10/3	988

Click to restore connectivity to Microsoft Exchange.

Security Considerations for the CAS Role

The CAS role is an access point for users to reach their mailbox. As such, the protection of the CAS role itself and it data paths is a primary consideration when planning an Exchange Server 2007 environment. Microsoft recommends the use of secure transport for all CAS communications: it is the default configuration of this server role. The ACE allows the network administrator to configure a central PKI infrastructure on the module and offload this client access service from the CAS roles. The ACE simplifies certificate management, preserves server resources and provides the level of transport security the Exchange environment demands.

From a traditional firewall perspective, the CAS role must be able to access the following ports, listed in the table below, to provide connectivity to all external clients and other Exchange entities. Configure the ACE module to provide the appropriate ACLs to support this configuration. Note that some of these features may be disabled by the server administrators and therefore will not require a “hole” in the firewall.

Table 3 CAS Communication Ports

Communication	Port(s)
Autodiscover service	80/TCP, 443/TCP (SSL)
Availability service	80/TCP, 443/TCP (SSL)
Outlook Web Access	80/TCP, 443/TCP (SSL)
POP3	110/TCP (TLS), 995/TCP (SSL)
IMAP4	143/TCP (TLS), 993/TCP (SSL)
Outlook Anywhere (formerly known as RPC over HTTP)	80/TCP, 443/TCP (SSL)
Exchange ActiveSync application	80/TCP, 443/TCP (SSL)
Client Access server to Unified Messaging server	5060/TCP, 5061/TCP, 5062/TCP, a dynamic port
Client Access server to a Mailbox server that is running an earlier version of Exchange Server	80/TCP, 443/TCP (SSL)
Client Access server to Exchange 2007 Mailbox server	RPC. (Dynamic Ports)
Client Access server to Client Access server (Exchange ActiveSync)	80/TCP, 443/TCP (SSL)
Client Access server to Client Access server (Outlook Web Access)	80/TCP, 443/TCP (SSL)
WebDAV	80/TCP, 443/TCP (SSL)

Note Microsoft strongly recommends the use of an application layer firewall to provide protection up the TCP stack.

Note For more information, refer to the *Data Path Security Reference* document at the following URL http://technet.microsoft.com/en-us/library/bb331973.aspx. This document is a comprehensive guide to all ports and data paths leveraged by the Microsoft Exchange 2007 server roles and clients.

Edge Server Role

The ET role is typically located at the edge of the enterprise. The Edge role primary role is to forward external traffic to the HT roles via SMTP, the Edge role acts as an SMTP proxy. In addition, the Edge roles are filters, removing spam, viruses and e-mail messages deemed inappropriate according to the enterprises security policy. The Edge roles are standalone, independent islands of defense in the DMZ of the enterprise. To make these roles highly available and scalable, enterprises typically deploy DNS-based round-robin or some form of load balancing.

For more information about DNS-based load balancing, refer to Site Selection, page 46.

Transparent Service Integration

This section discusses the integration of the Cisco ACE service module with the Microsoft ET role. As previously mentioned, the Edge roles can leverage the load balancing, health monitoring and security features of the ACE. Figure 37 illustrates the tested deployment of ACE and Edge roles. In this scenario, a virtual ACE context is deployed in transparent, or bridge mode, joining VLANs 110 and 210 logically. The VIP of the ACE abstracts the Edge roles located in the server farm.

Typically, the ET roles use two physical Network Interface Cards (NICs), a public or Internet-facing card and a private or internal-facing adapter. The Internet NIC should be positioned to handle SMTP traffic only when the private NIC supports SMTP and the Exchange-related control traffic.

Traffic Pattern

Figure 37 depicts the tested topology of a user sending an SMTP message to the enterprise. In this example, the MX records have already been resolved and the client is establishing a session to send an email.

Figure 37 Edge Role Traffic Pattern

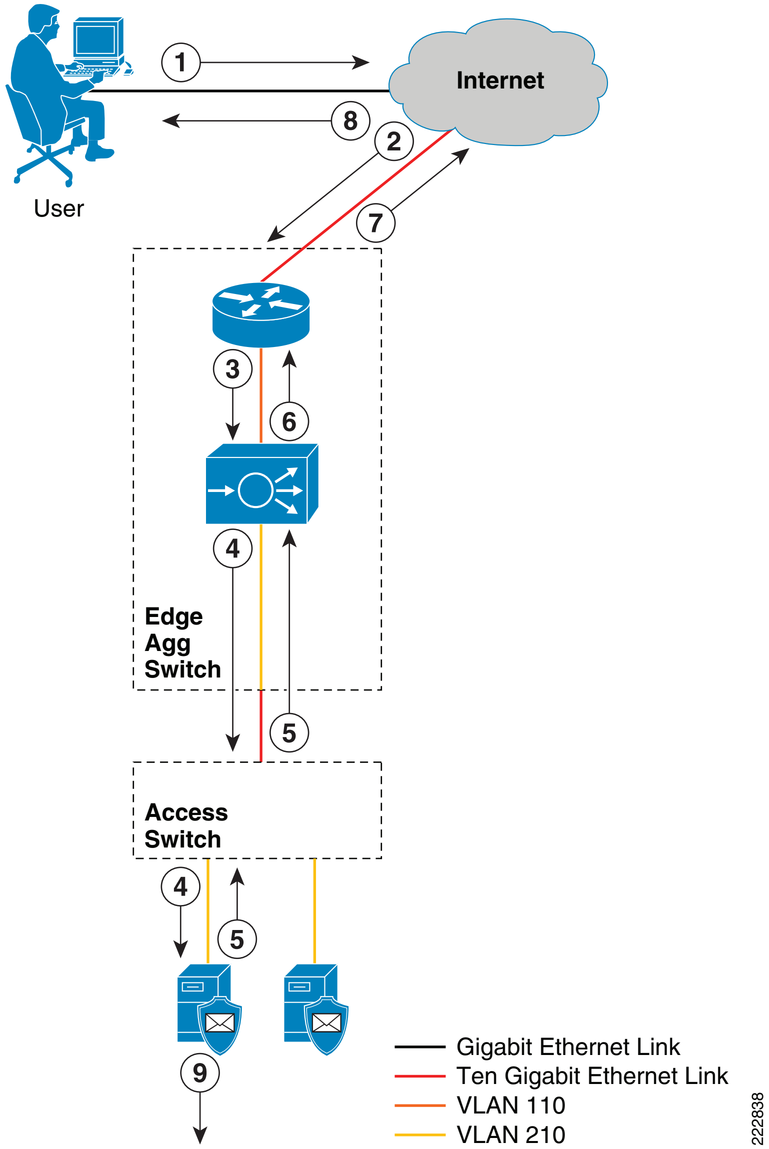

The following details the traffic flow for Figure 37:

Step 1 The user sends a SYN to the IP address associated with the DNS name of the MX record. This IP address is the VIP address of the ACE.

Step 2 The SYN is received by the router, in this case an MSFC, at the edge of the enterprise.

Step 3 The MSFC of the aggregation layer receives the traffic and forwards it to the ACE VIP defined on VLAN 110 in this example. The ACE is in transparent mode and therefore leverages a bridged virtual interface to join VLANs 110 and 210. The following interface configuration is necessary:

```
interface vlan 110
  description ** Public Interface **
!Places the interface in the bridge group
  bridge-group 10
!Service policy applied to the interface defines the ACE VIP and load balancing methods
  service-policy input PM-EDGE
  no shutdown

interface vlan 210
!Places the interface in the bridge group
  description ** Private Interface **
  bridge-group 10
  no shutdown
!Logical join of the two VLAN interfaces via the BVI
interface bvi 10
  ip address 10.210.1.4 255.255.255.0
  alias 10.210.1.6 255.255.255.0
  description ** Edge Bridge Interface **
  no shutdown
```

The ACE uses a service policy defined as PM-EDGE on the public VLAN 110 interface. The PM-EDGE policy-map definition is below:

```
!Policy Map leverages the EDGE-VIP class that defines the actual VIP
policy-map multi-match PM-EDGE
  class EDGE-VIP
    loadbalance vip inservice
!The Edge VIP is defined with this policy
    loadbalance policy PM-FM-EDGE-VIP
!Enable ping response on VIP
    loadbalance vip icmp-reply active

!Define the virtual IP address of the Edge server farm
class-map match-all EDGE-VIP
  2 match virtual-address 10.210.1.100 tcp eq smtp

!Create a catch all policy to send inbound SMTP traffic to the Edge server farm
policy-map type loadbalance first-match PM-FM-EDGE-VIP
  class class-default
    serverfarm EDGE-FARM
```

The Edge server farm configuration:

```
serverfarm host EDGE-FARM
!Health monitoring via an SMTP probe
  probe EXPROBE
!Define the real servers
  rserver RTP-EDGE-01
    inservice
  rserver RTP-EDGE-02
    inservice
```

The following SMTP probe is leveraged from the ACE:

```
probe smtp EXPROBE
  description Edge Server Probe
  interval 5
  passdetect interval 5
  passdetect count 2
  receive 1
  connection term forced
  open 1
!Successful SMTP message responses
  expect status 220 250
```

The **show probe** command displays the details of the SMTP health monitoring:

```
show probe EXPROBE

 probe       : EXPROBE
 type        : SMTP, state : ACTIVE
----------------------------------------------
   port      : 25      address     : 0.0.0.0         addr type  : -
   interval  : 5       pass intvl  : 5               pass count : 2
   fail count: 3       recv timeout: 1
                       --------------------- probe results --------------------
   probe association   probed-address  probes     failed     passed     health
   ------------------- ---------------+----------+----------+----------+-------
   serverfarm  : edge_farm
     real      : RTP-EDGE-01[0]
                       10.210.1.10     59101      1          59100      SUCCESS
     real      : RTP-EDGE-02[0]
                       10.210.1.11     59101      1          59100      SUCCESS
```

Define the real Edge roles and place them in-service:

```
rserver host RTP-EDGE-01
  ip address 10.210.1.10
  inservice
rserver host RTP-EDGE-02
  ip address 10.210.1.11
  inservice
```

Step 4 The ACE forwards the connection to the SMTP server based on the predictor configured under the server farm. In this instance, the ACE uses the default round-robin algorithm. It is important to note that the destination address is that of the real Edge role and not the ACE VIP.

Step 5 The Edge role responds with an SMTP 220 message indicating the “Mail Service is Ready” destined to the client’s IP address. The Edge role is configured to use the MSFC as a gateway to the internet.

Step 6 The response is forwarded to the MSFC after the Edge roles source IP address is replaced with that of the ACE VIP.

Step 7 The MSFC routes the packet to the client.

Step 8 The client receives the response from the Edge role and completes the SMTP transaction.

Step 9 The Edge role acts as a proxy and security filter for the enterprise. In this role, the Edge role inspects the SMTP message and determines if the message should be delivered based on the local security policies. In this example, the Edge role forwards the message to the HT role for final delivery. Figure 38 shows the original SMTP transaction with the client in frame 25. The client EHLO is received by the Edge role, the SMTP transaction completes. The Edge role begins to route the mail message to the local HT role, acting as a proxy for the remote client.

Figure 38 Edge Role is a SMTP Proxy Example

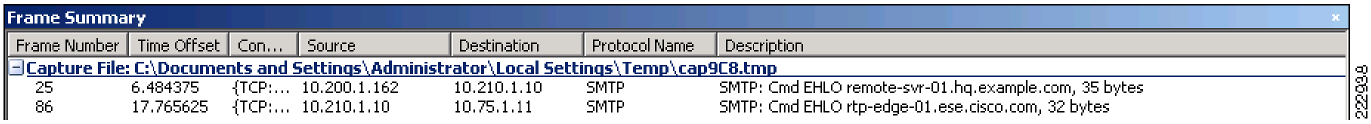

Frame Summary

Frame Number	Time Offset	Con...	Source	Destination	Protocol Name	Description
Capture File: C:\Documents and Settings\Administrator\Local Settings\Temp\cap9C8.tmp						
25	6.484375	{TCP:...	10.200.1.162	10.210.1.10	SMTP	SMTP: Cmd EHLO remote-svr-01.hq.example.com, 35 bytes
86	17.765625	{TCP:...	10.210.1.10	10.75.1.11	SMTP	SMTP: Cmd EHLO rtp-edge-01.ese.cisco.com, 32 bytes

Security Considerations for the Edge Role

The Edge role has very specific tasks, filtering incoming mail, and forwarding messages deemed safe. As an access point to the enterprise, the Edge role itself must be protected. To this end, Microsoft recommends the physical segmentation of traffic between private and public network adapters and leaving the Edge role disjoined from any domain. On the external interface, it is recommended to allow only SMTP and DNS traffic. The internal facing interface, must have allow access to HT and directory services. Table 4 details the ports necessary for the deployment of Edge roles in the DMZ.

Table 4 ET Role Interesting Ports

Communication	Port(s)
HT roles to ET roles	25/TCP (SSL)
ET role to HT role	25/TCP (SSL)
ET role to ET role	25/TCP (SSL), 389/TCP/UDP, and 80/TCP (certificate authentication)
Microsoft Exchange EdgeSync service	50636/TCP (SSL), 50389/TCP (No SSL)
Active Directory Application Mode (ADAM) directory service on ET role	50389/TCP (No SSL)

Note The ET role encrypts all communication by default except for the ADAM service. For information on the ADAM service, refer to Edge Transport Server, page 15.

Note For more information on securing the data paths of all the Microsoft Exchange 2007 servers go to “Data Path Security Reference” at http://technet.microsoft.com/en-us/library/bb331973.aspx.

Appendix

ACE SSL Proxy Configuration

Step 1 The configuration of the ACE module to support SSL requires several steps. The following details the configuration leveraged on the ACE to support the OWA and Outlook Anywhere clients. Note the default key bit size in Exchange 2007 is 2048:

Generate an RSA key:

```
crypto generate key 2048 newKey
```

Step 2 Create the certificate signing request parameters in the ACE:

```
crypto csr-params OWA
  country US
  state NC
  locality RTP
  organization-name RTP-ESE
  organization-unit DC
  common-name owa.ese.cisco.com
  serial-number 1
```

Step 3 Generate the Certificate Signing Request:

```
crypto generate csr OWA newKey
```

Step 4 Sample output from command:

```
crypto generate csr OWA newKey
-----BEGIN CERTIFICATE REQUEST-----
MIICqDCCAZACAQAwYzELMAkGA1UEBhMCVVMxCzAJBgNVBAgTAk5DMQwwCgYDVQQH
```

```
EwNSVFAxEDAOBgNVBAoTB1JUUC1FU0UxCzAJBgNVBAsTAkRDMRowGAYDVQQDExFv
d2EuZXNlLmNpc2NvLmNvbTCCASIwDQYJKoZIhvcNAQEBBQADggEPADCCAQoCggEB
AKzCTZEA4eVEzRo8M4VsOcuPn+6xeA5EaZBf8v5xcZEHuZgB8L3SrjakRbGSkH+i
Q3pGLEnq9xfCbXg4t3k7CHU46hD3KQEFV2s3Yo3x7qxlzoiwU6yzXCcMQQQb5ia7
Bd1mWZLOZiXa8JiTA3HhNDxitpsYHHLspntRlOBPTUyRMOvvTeYkofSOBvHbnZio
CJmprQfH06dUeVrwHMq3SNzi5Ewid/b8PVkp21g18fp3utqgVM+53sjL+VXfDuBt
JUfxlc99djrx9drtK65Cfd0xaRYoo6Oq1lzL84pN3qUB2Z8sOr2zkek6Xga6ecmz
jGKlGRLjgmCLnqsXJlO2m+0CAwEAAaAAMA0GCSqGSIb3DQEBBAUAA4IBAQA0Xt0n
q5cO29QVOgYOHkBGqkYNuP1Q9TR8SCImDISYHWNVJPXW/UrEEUGC648PAk8lKQ46
Thod1pytcypTvcyGG7JbnBaQ1V+V4rtI2nvsDvs/LLjH3/Cl0zRy1HXYJxjIeW/h
lOCRMxn6M6KtKVuBEnaKEZ4J9LfkDm6XxOoI8vxTo+39zWfN+3gPx0gLEmiAIk/U
sZhvhhQrbtxTCJGxnJRxyDoE7Y9Z6fWWzhIgldyixRUl02DoxgCVqwICnRmqi7VB
nSOSi62SvI7fqF8ppsG0a45NgKj3CMaSehWsVRFo8fgjo1qsPDVCKy2S4xt6TPZy
etni9/NkadCBV5Qq
-----END CERTIFICATE REQUEST-----
```

Step 5 The CSR must be exported to the Certificate Authority for a valid certificate to be issued. In the tested configuration, the Certificate Authority is a Microsoft Windows 2003 server running CA Version 5.2.3790.3959. The CA required the use of a certificate template in the CSR. To accommodate this requirement use the **certutils** command from a DoS prompt to associate a template attribute with the ACE CSR.

```
certreq –submit –attrib "WebServer" owa.req
```

A list of valid certificate template common names can be determined with the following command.

```
certutil -templates
```

Step 6 After the certificate is issued import the certificate into the ACE with the following command using one of the listed protocols:

```
crypto import ?
  ftp             Import a key/certificate from an ftp server
  non-exportable  Mark this key/certificate as non-exportable
  sftp            Import a key/certificate from an sftp server
  terminal        Accept a key/certificate from terminal
  tftp            Import a key/certificate from a tftp server
```

Step 7 Verify the successful import of the certificate:

```
crypto verify newKey OWA.cer
Keypair in newKey matches certificate in OWA.cer.
```

Step 8 The valid certificate should be associated with an SSL proxy service to provide client side offload. Note that the certificate and its associated key are defined in the proxy.

```
ssl-proxy service OWA
  key newKey
  cert OWA.cer
```

Step 9 Create an SSL parameter map to define the supported SSL ciphers and versions accepted by the ACE SSL proxy.

```
parameter-map type ssl sslparams
cipher ?
  RSA_EXPORT1024_WITH_DES_CBC_SHA  Accept RSA_EXPORT1024_WITH_DES_CBC_SHA cipher
  RSA_EXPORT1024_WITH_RC4_56_MD5   Accept RSA_EXPORT1024_WITH_RC4_56_MD5 cipher
  RSA_EXPORT1024_WITH_RC4_56_SHA   Accept RSA_EXPORT1024_WITH_RC4_56_SHA cipher
  RSA_EXPORT_WITH_DES40_CBC_SHA    Accept RSA_EXPORT_WITH_DES40_CBC_SHA cipher
  RSA_EXPORT_WITH_RC4_40_MD5       Accept RSA_EXPORT_WITH_RC4_40_MD5 cipher
  RSA_WITH_3DES_EDE_CBC_SHA        Accept RSA_WITH_3DES_EDE_CBC_SHA cipher
  RSA_WITH_AES_128_CBC_SHA         Accept RSA_WITH_AES_128_CBC_SHA cipher
  RSA_WITH_AES_256_CBC_SHA         Accept RSA_WITH_AES_256_CBC_SHA cipher
  RSA_WITH_DES_CBC_SHA             Accept RSA_WITH_DES_CBC_SHA cipher
  RSA_WITH_RC4_128_MD5             Accept RSA_WITH_RC4_128_MD5 cipher
  RSA_WITH_RC4_128_SHA             Accept RSA_WITH_RC4_128_SHA cipher
version ?
  all   All SSL versions
  SSL3  SSL Version 3
  TLS1  TLS Version 1
```

Step 10 Use the SSL parameter map by applying it to the SSL proxy service. For example:

```
 ssl-proxy service OWA
  key newKey
  cert OWA.cer
  ssl advanced-options sslparams
```

Outlook Anywhere Configuration

Deploying Outlook Anywhere requires both client and server modifications. This section describes the changes made to the CAS and Outlook test clients to verify ACE support of this feature.

Client Access Server (CAS)

Enabling Outlook Anywhere in Exchange is a two-step process involving:

- RPC over HTTP Windows Component
- CAS role configuration

To enable the RPC Windows Component, use the following steps:

Step 1 Click **Start**, point to **Settings**, click **Control Panel**, and then double-click **Add or Remove Programs**.

Step 2 Click **Add/Remove Windows Components**.

Step 3 On the Windows Components page, in the Components window, select **Networking Services**, and then click the **Details** button.

Step 4 On the Networking Services page, in the Subcomponents of Networking Services window, select the check box next to **RPC over HTTP Proxy**, and then click **OK**.

Step 5 On the Windows Components page, click **Next.**

Step 6 Click **Finish** to close the Windows Components Wizard.

From the Exchange Management Console, right click on the CAS role, where Outlook Anywhere support is to be deployed. The form in Figure 39 is shown.

Figure 39 Enable Outlook Anywhere Form

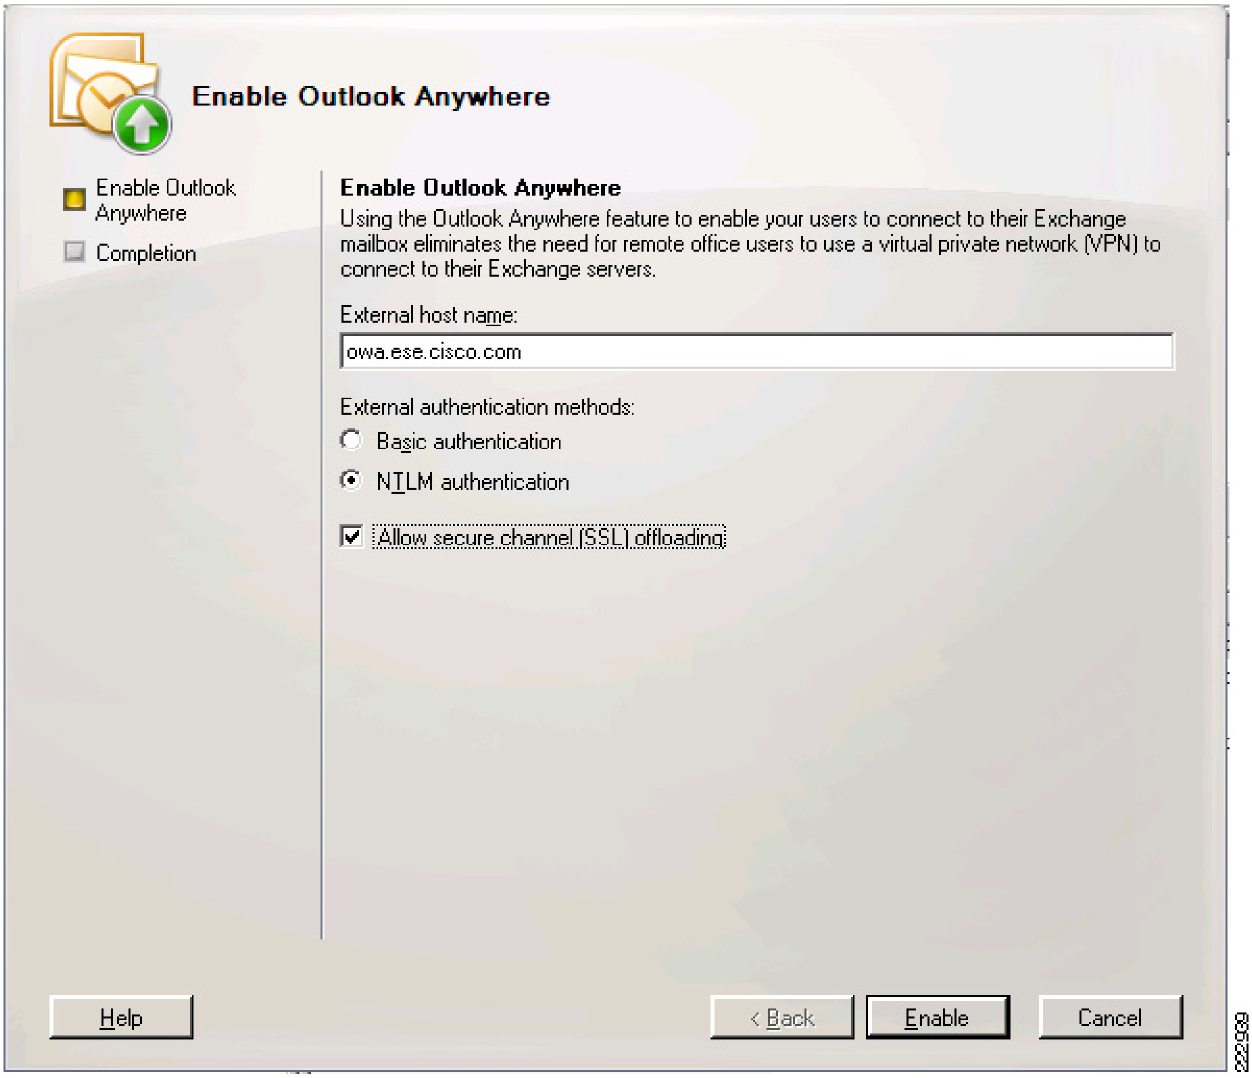

Enable Outlook Anywhere

Using the Outlook Anywhere feature to enable your users to connect to their Exchange mailbox eliminates the need for remote office users to use a virtual private network (VPN) to connect to their Exchange servers.

External host name:

owa.ese.cisco.com

External authentication methods:

○ Basic authentication

● NTLM authentication

☑ Allow secure channel (SSL) offloading

Step 7 External hostname field should be equivalent to the DNS name that resolves to the ACE VIP.

Step 8 Authentication is required, NTLM is the recommended method.

Step 9 Check the box to allow the ACE to offload the SSL services from the CAS role.

Step 10 Click **Enable** button. See Figure 40.

Step 11 Click **Finish** button.

Figure 40 Enable Outlook Anywhere Completion Screen Shot

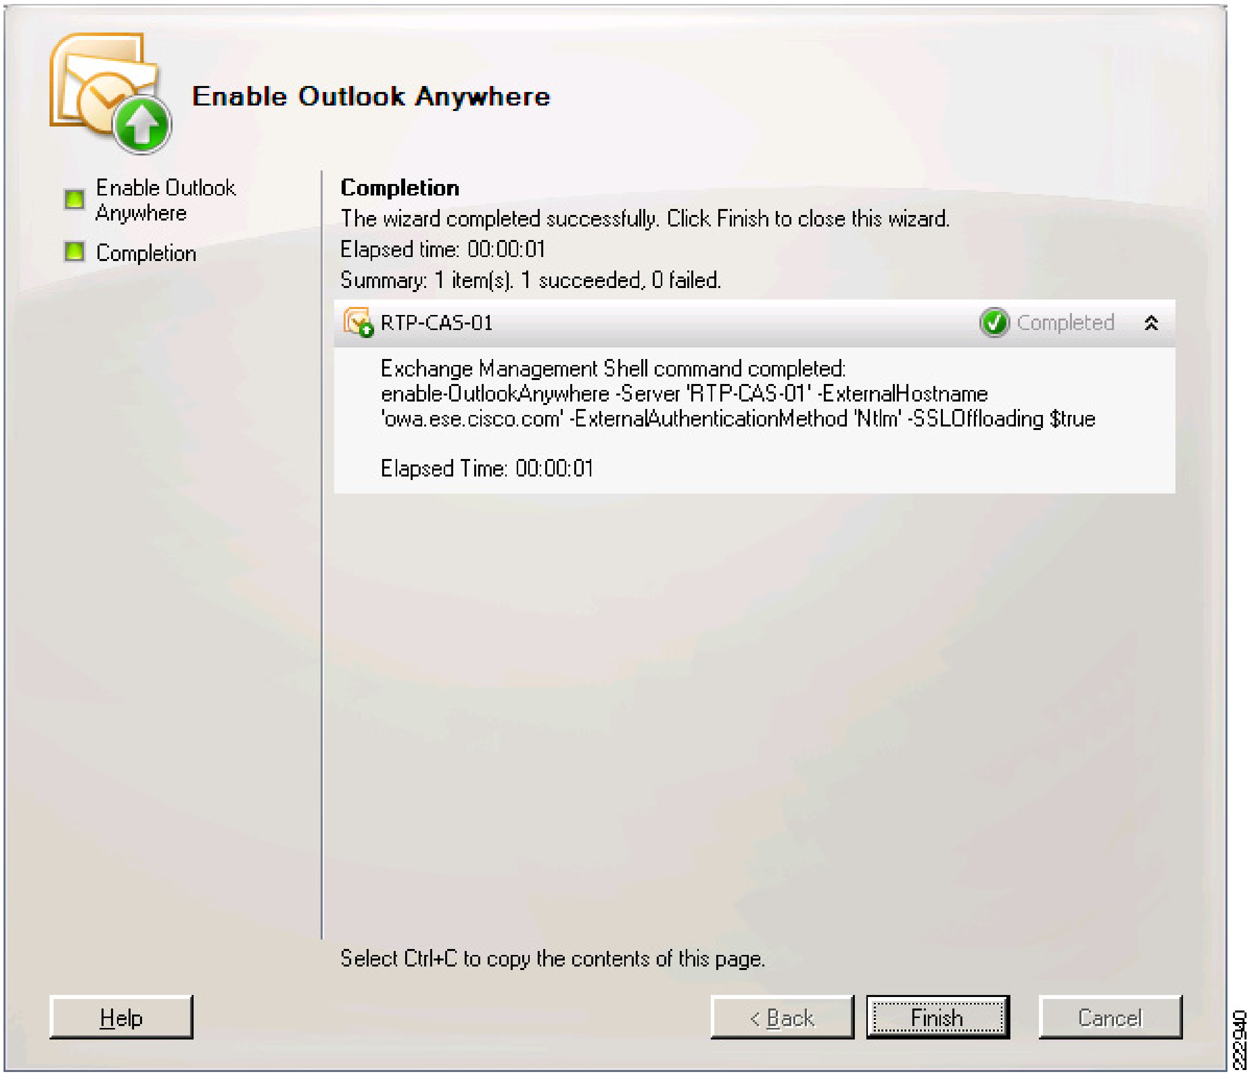

Completion

The wizard completed successfully. Click Finish to close this wizard.

Elapsed time: 00:00:01

Summary: 1 item(s). 1 succeeded, 0 failed.

Exchange Management Shell command completed:
enable-OutlookAnywhere -Server 'RTP-CAS-01' -ExternalHostname 'owa.ese.cisco.com' -ExternalAuthenticationMethod 'Ntlm' -SSLOffloading $true

Elapsed Time: 00:00:01

Select Ctrl+C to copy the contents of this page.

Note The external hostname and certificate common-name should match.

Outlook Client

To enable Outlook Anywhere on an Outlook 2007 client, perform the following steps:

Step 1 Open Outlook 2007 under *Tools -> Account Settings -> Change E-mail Account -> More Settings*

Step 2 Select the **Connection** tab.

Figure 41 Outlook 2007 - Enable Outlook Anywhere Connection

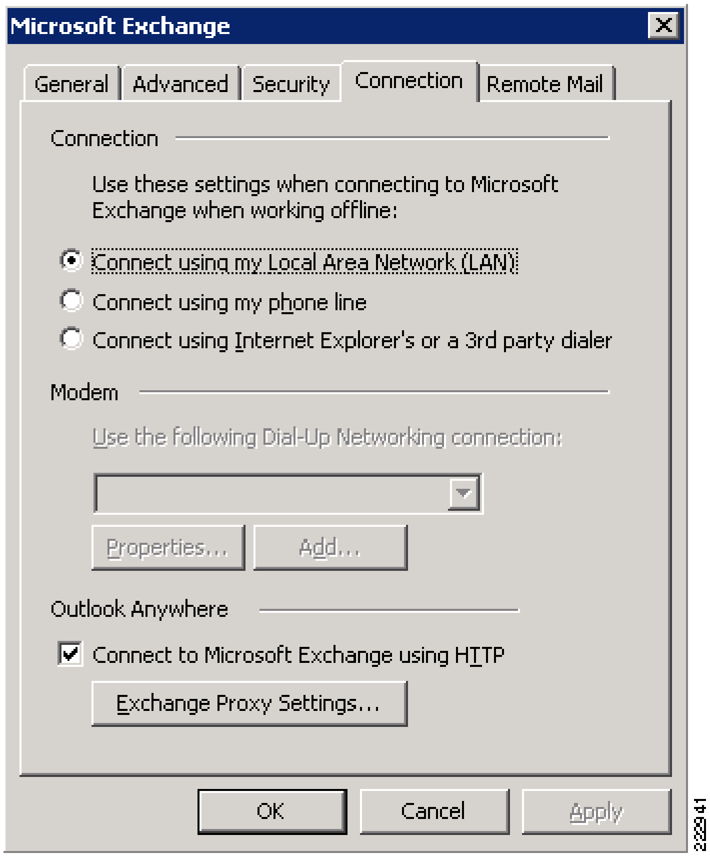

Connection

Use these settings when connecting to Microsoft Exchange when working offline:

● Connect using my Local Area Network (LAN)

○ Connect using my phone line

○ Connect using Internet Explorer's or a 3rd party dialer

Modem

Use the following Dial-Up Networking connection:

Outlook Anywhere

☑ Connect to Microsoft Exchange using HTTP

Step 3 Select **Connect to Microsoft Exchange using HTTP** checkbox and click on **Exchange proxy Settings**. See Figure 42.

Step 4 Complete the Proxy Settings form:

a. The URL is the FQDN of the ACE VIP.

b. Connect with SSL for secure transport.

c. Optional certificate parameter.

d. Force Outlook Anywhere on all connections via the “fast” and “slow” network check boxes.

e. Choose an authentication method.

Figure 42 Setting up Exchange Proxy Settings

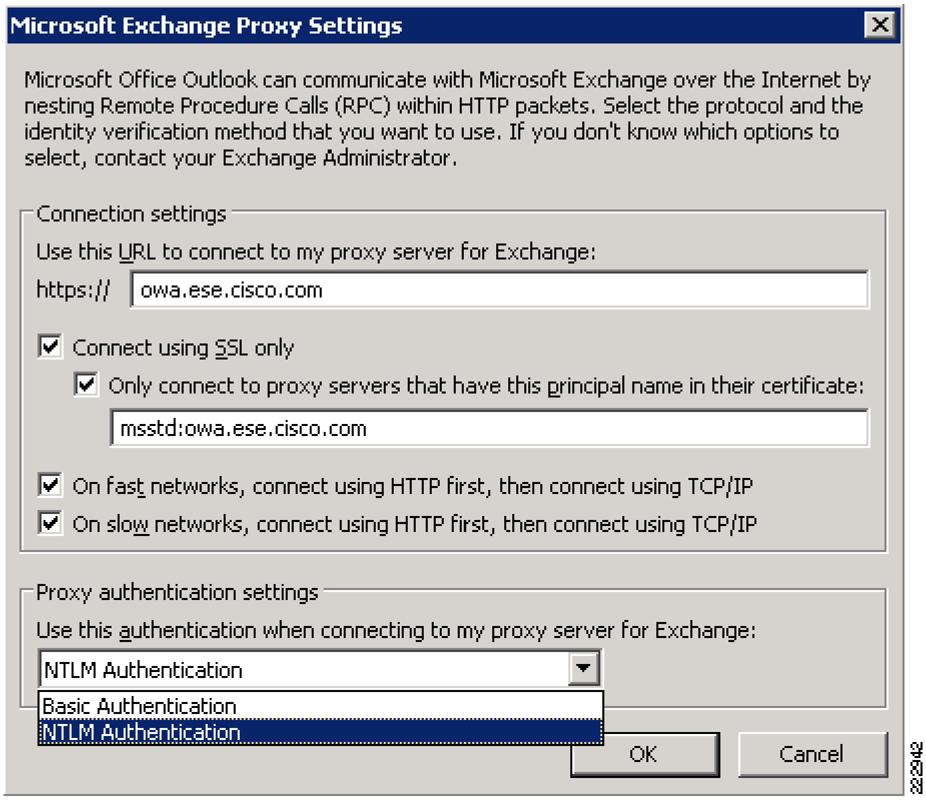

Microsoft Exchange Proxy Settings

Microsoft Office Outlook can communicate with Microsoft Exchange over the Internet by nesting Remote Procedure Calls (RPC) within HTTP packets. Select the protocol and the identity verification method that you want to use. If you don't know which options to select, contact your Exchange Administrator.

Connection settings

Use this URL to connect to my proxy server for Exchange:

https:// owa.ese.cisco.com

☑ Connect using SSL only

☑ Only connect to proxy servers that have this principal name in their certificate:

msstd:owa.ese.cisco.com

☑ On fast networks, connect using HTTP first, then connect using TCP/IP

☑ On slow networks, connect using HTTP first, then connect using TCP/IP

Proxy authentication settings

Use this authentication when connecting to my proxy server for Exchange:

NTLM Authentication

Cisco Validated Design

The Cisco Validated Design 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/validateddesigns.

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

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