Enterprise Data Center Wide Area Application Services (WAAS) Design Guide

This document offers guidelines and best practices for implementing Wide Area Application Services (WAAS) in enterprise data center architecture. Placement of the Cisco Wide Area Engine (WAE), high availability, and performance are discussed for enterprise data center architectures to form a baseline for considering a WAAS implementation.

Contents

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

Introduction

As enterprise businesses extend their size and reach to remote locations, guaranteeing application delivery to end users becomes increasingly important. In the past, remote locations contained their own application file servers and could provide LAN access to data and applications within the remote location or branch. Although this solution guarantees application performance and availability, it also means more devices to manage, increased total cost of ownership, regulatory compliance for data archival, and lack of anywhere, anytime application access. Placing application networking servers within a centralized data center where remote branches access applications across a WAN solves the management of devices and total cost of ownership issues. The benefits for consolidating application networking services in the data center include but are not limited to the following:

- Cost savings through branch services consolidation of application and printer services to a centralized data center
- Ease of manageability because less devices are employed in a consolidated data center
- Centralized storage and archival of data to meet regulatory compliance
- More efficient use of WAN link utilization through transport optimization, compression, and file caching mechanisms to improve overall user experience of application response

The trade-off with the consolidation of resources in the data center is the increase in delay for remote users to achieve the same performance of accessing applications at LAN-like speeds as when these servers resided at the local branches. Applications commonly built for LAN speeds are now traversing a WAN with less bandwidth and increased latency over the network. Potential bottlenecks that affect this type of performance include the following:

- Users at one branch now contend for the same centralized resources as other remote branches.
- Insufficient bandwidth or speed to service the additional centralized applications now contend for the same WAN resources.

- Network outage from remote branch to centralized data center resources cause “disconnected” events, severely impacting remote business operations.

The Cisco WAAS portfolio of technologies and products give enterprise branches LAN-like access to centrally-hosted applications, servers, storage, and multimedia with LAN-like performance. WAAS provides application delivery, acceleration, WAN optimization, and local service solutions for an enterprise branch to optimize performance of any TCP-based application in a WAN or MAN environment.

This document provides guidelines and best practices when implementing WAAS in enterprise architectures. This document gives an overview of WAAS technology and then explores how WAAS operates in data center architectures. Design considerations and complete tested topologies and configurations are provided.

Intended Audience

This design guide is targeted for network design engineers to aid their architecture, design, and deployment of WAAS in enterprise data center architectures.

Caveats and Limitations

The technical considerations in this document refer to WAAS version 4.0(3). The following features have not been tested in this initial phase and will be considered in future phases:

- Policy-based routing (PBR)
- Inline interception
- CIFS auto-discovery
- WAE interoperability with ASA firewalls

Although these features are not tested, their expected behavior may be discussed in this document.

Assumptions

This design guide has the following starting assumptions:

- System engineers and network engineers possess networking skills in data center architectures.
- Customers have already deployed Cisco-powered equipment in data center architectures. Interoperability of the WAE and non-Cisco equipment is not evaluated.
- Although the designs provide flexibility to accommodate various network scenarios, Cisco recommends following best design practices for the enterprise data center. This design guide is an overlay of WAAS into the existing network design. For detailed design recommendations, see the data center design guides at the following URL: http://www.cisco.com/en/US/netsol/ns743/networking_solutions_program_home.html.

Best Practices and Known Limitations

DC WAAS Best Practices

The following is a summary of best practices that are described in more detail in the subsequent sections:

- Install the WAE at the WAN edge to increase optimization coverage to all hosts in the network.
- Use Redirect ACL to limit campus traffic going through the WAEs for installation in the aggregation layer; optimization applies to selected subnets.
- Use Web Cache Communications Protocol version 2 (WCCPv2) instead of PBR; WCCPv2 provides more high availability and scalability features, and is also easier to configure.
- PBR is recommended where WCCP or inline interception cannot be used.
- Inbound redirection is preferred over outbound redirection because inbound redirection is less CPU-intensive on the router.
- Two Central Managers are recommended for redundancy.
- Use a standby interface to protect against network link and switch failure. Standby interface failover takes around five seconds.
- For Catalyst 6000/76xx deployments, use only inbound redirection to avoid using “redirection exclude in”, which is not understood by the switch hardware and must be processed in software.
- For Catalyst 6000/76xx deployments, use L2 redirection for near line-rate redirection.
- Use Multigroup Hot Standby Routing Protocol (mHSRP) to load balance outbound traffic.
- Install additional WAEs for capacity, availability, and increased system throughput; WAE can scale in near linear fashion in an N+1 design.

WAAS Known Limitations

- A separate WAAS subnet and tertiary/sub-interface are required for transparent operation because of preservation of the L3 headers. Traffic coming out of the WAE must not redirect back to the WAE. Inline interception does not need a separate WAAS subnet.
- IPv6 is not supported by WAAS 4.0; all IP addressing must be based on IPv4.
- WAE overloading such as the exhaustion of TCP connections results in pass-through traffic (non-optimized); WCCP does not know when a WAE is overloaded. WCCP continues to send traffic to the WAE based on the hashing/masking algorithm even if the WAE is at capacity. Install additional WAEs to increase capacity.

WAAS Technology Overview

To appreciate how WAAS provides WAN and application optimization benefits to the enterprise, first consider the basic types of centralized application messages that would be transmitted to and from remote branches. For simplicity, two basic types are identified:

- Bulk transfer applications—Focused more on the transfer of files and objects. Examples include FTP, HTTP, and IMAP. In these applications, the number of roundtrip messages may be few and may have large payloads with each packet. Some examples include web portal or lite 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 roundtrips of application protocol messages that may or may not have small payloads. Examples include Microsoft Office applications (Word, Excel, Powerpoint, and Project).

WAAS uses the following technologies to provide a number of application acceleration as well as remote file caching, print service, and DHCP features to benefit both types of applications:

- Advanced compression using DRE and Lempel-Ziv (LZ) compression

  DRE is an advanced form of network compression that allows Cisco WAAS to maintain an application-independent history of previously-seen data from TCP byte streams. 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 (TFO)

  Cisco WAAS TFO employs 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. 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 (CIFS) caching services

  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. WAAS provides a CIFS adapter that is able to 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

  WAAS can cache print drivers at the branch, so an extra file or print server is not required. By using WAAS for caching these services, client requests for downloading network printer drivers do not have to traverse the WAN.

- DHCP

  WAAS provides local DHCP services.

For more information on these enhanced services, see the *WAAS 4.0 Technical Overview* at the following URL:http://www.cisco.com/en/US/prod/collateral/contnetw/ps5680/ps6870/data_sheet_cisco_wide_area_application_services_WAAS_software_version_4_0.html.

Figure 1 shows the logical mechanisms that are used to achieve WAN and application optimization, particularly using WAAS.

Figure 1 Wide Area Application Services (WAAS) Mechanisms

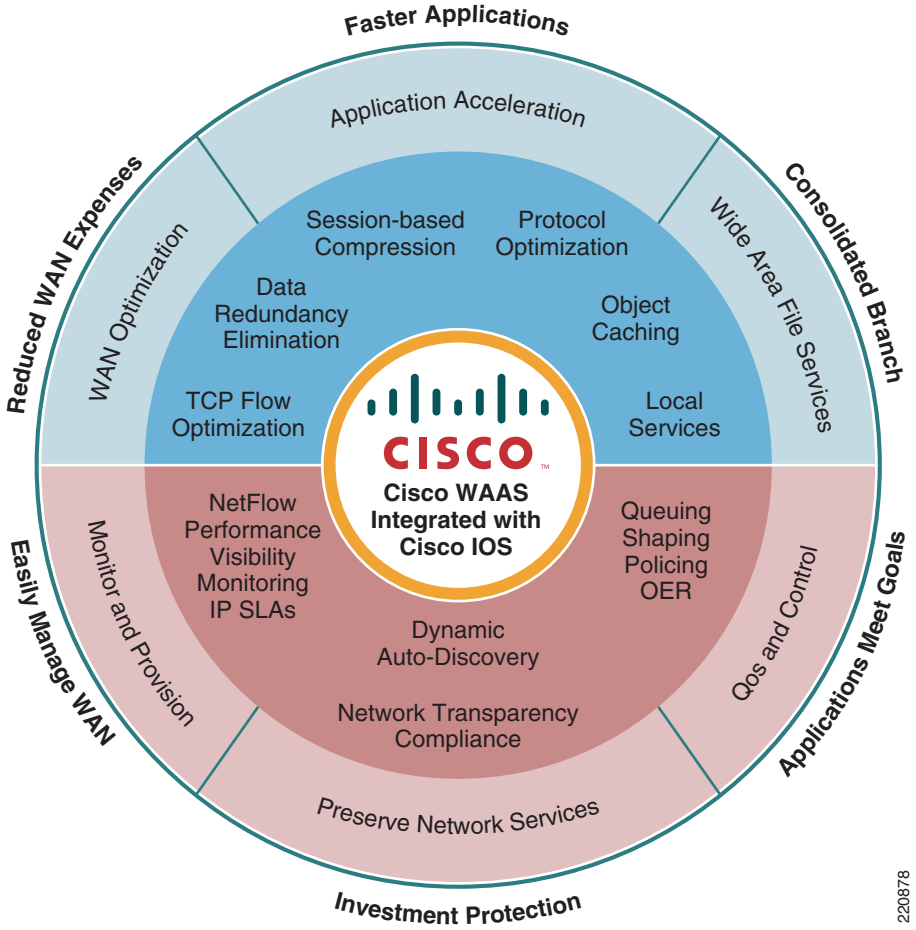

The WAAS features are not described in detail in this guide; the WAAS data sheets and software configuration guide explain them in more detail. This literature provides excellent feature and configuration information on a product level. Nevertheless, for contextual purposes, some of the WAAS basic components and features are reviewed in this document.

WAAS consists mainly of the following main hardware components:

- Application Accelerator Wide Area Engines (WAE) —The application accelerator resides within the campus/data center or the branch. If placed within the data center, the WAE is the TCP optimization and caching proxy for the origin servers. If placed at the branch, the WAE is the main TCP optimization and caching proxy for branch clients.
- WAAS Central Manager (CM)—Provides a unified management control over all the WAEs. The WAAS CM usually resides within the data center, although it can be physically placed anywhere provided that there is a communications path to all the managed WAEs.

For more details on each of these components, see the *WAAS 4.0.7 Software Configuration Guide* at the following URL:
http://www.cisco.com/en/US/docs/app_ntwk_services/waas/waas/v407/quick/guide/waasqcg.html.

The quantity and WAE hardware model selection varies with a number of factors (see Table 1). For the branch, variables include the number of estimated simultaneous TCP/CIFS connections, the estimated disk size for files to be cached, and the estimated WAN bandwidth. Cisco provides a WAAS sizing tool for guidance, which is available internally for Cisco sales representatives and partners. The NME-WAE is the WAE network module and deployed inside the branch integrated services router (ISR).

Table 1 WAE Hardware Sizing

Device	Max Optimized TCP Connections	Max CIFS Sessions	Single Drive Capacity [GB]	Max Drives	RAM [GB]	Max Recommended WAN Link [Mbps]	Max Optimized Throughput [Mbps]
NME-WAE-302	250	N/A	80	1	0.5	4	90
NME-WAE-502	500	500	120	1	1	4	150
WAE-512-1	750	750	250	2	1	8	100
WAE-512-2	1500	1500	250	2	2	20	150
WAE-612-2	2000	2000	300	2	2	45	250
WAE-612-4	6000	2500	300	2	4	90	350
WAE-7326	7500	2500	300	6	4	155	450

WAAS Optimization Path

Optimizations are performed between the core and edge WAE. The WAEs act as a TCP proxy for both clients and their origin servers within the data center. This is not to be confused with other WAN optimization solutions that create optimization tunnels. In those solutions, the TCP header is modified between the caching appliances. With WAAS, the TCP headers are fully preserved. Figure 2 shows three TCP connections.

Figure 2 WAAS Optimization Path

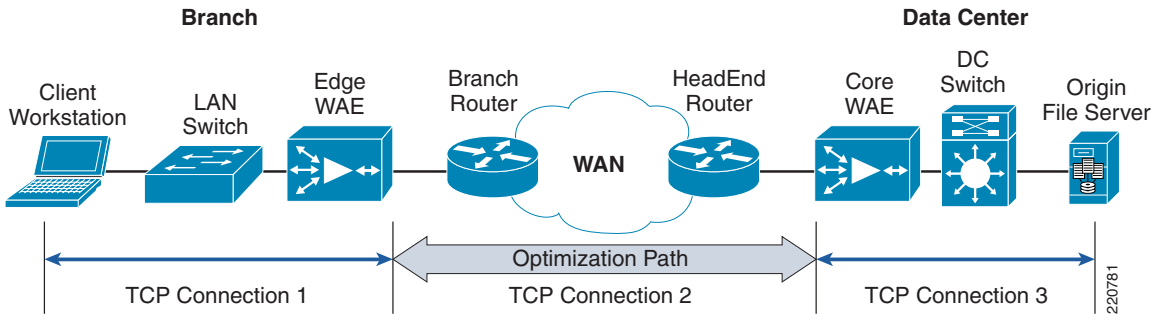

TCP connection #2 is the WAAS optimization path between two points over a WAN connection. Within this path, Cisco WAAS optimizes the transfer of data between these two points over the WAN connection, minimizing the data it sends or requests. Traffic in this path includes any of the WAAS optimization mechanisms such as the TFO, DRE, and LZ compression.

Identifying where the optimization paths are created among TFO peers is important because there are limitations on what IOS operations can be performed. Although WAAS preserves basic TCP header information, it modifies the TCP sequence number as part of its TCP proxy session. As a result, some

features dependent on inspecting the TCP sequence numbering, such as IOS firewall packet inspection or features that perform deep packet inspection on payload data, may not be interoperable within the application optimization path. More about this is discussed in Security, page 24.

The core WAE and thus the optimization path can extend to various points within the campus/data center. Various topologies for core WAE placement are possible, each with its advantages and disadvantages.

WAAS is part of a greater application and WAN optimization solution. It is complementary to all the other IOS features within the ISR and branch switches. Both WAAS and the IOS feature sets synergistically provide a more scalable, highly available, and secure application for remote branch office users.

As noted in the last section, because certain IOS interoperability features are limited based on where they are applied, it is important to be aware of the following two concepts:

- Direction of network interfaces
- IOS order of operations

For identification of network interfaces, a naming convention is used throughout this document (see Figure 3 and Table 2).

Figure 3 Network Interfaces Naming Convention for Edge WAEs

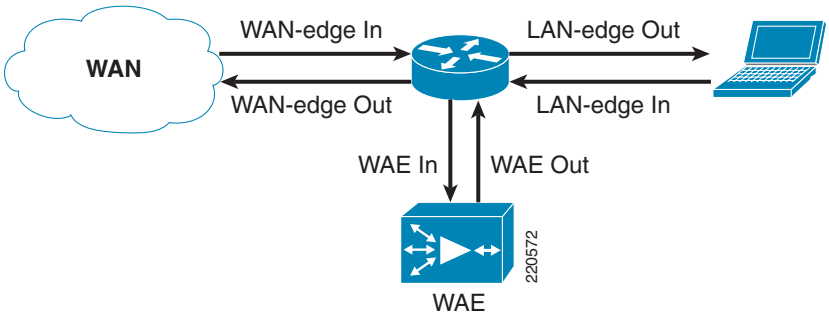

Table 2 Naming Conventions[1]

Interface	Description
LAN-edge in	Packets initiated by the data client sent into the switch or router
LAN-edge out	Packets processed by the router and sent outbound toward the clients
WAN-edge out	Packets processed by the router and sent directly to the WAN
WAN-edge in	Packets received directly from the WAN entering the router

Table 2 Naming Conventions[1]

Interface	Description
WAE-in	• From LAN-edge in—Packets redirected by WCCP or PBR from the client subnet to the WAE; unoptimized data • From WAN-edge in—Packets received from the core WAE; application optimizations are in effect
WAE- out	Packets already processed/optimized by the WAE and sent back towards the router: • To WAN-edge out—WAE optimizations in effect here • To LAN-edge out—no WAE optimizations

1. Source: http://www.cisco.com/en/US/tech/tk648/tk361/technologies_tech_note09186a0080133ddd.shtml

The order of IOS operations varies based on the IOS versions; however, Table 3 generally applies for the versions supported by WAAS. The bullet points in **bold** indicate that they are located inside the WAAS optimization path.

Table 3 Life of a Packet—IOS Basic Order of Operations[1]

Inside-to-Outside (LAN to WAN)	Outside-to-Inside (WAN to LAN)
• If IPsec, then check input access list • Decryption (if applicable) for IPsec • Check input access list • Check input rate limits • Input accounting • Policy routing • Routing • Redirect to web cache (WCCP or L2 redirect) • WAAS application optimization *(start/end of WAAS optimization path)* • **NAT inside to outside (local to global translation)** • **Crypto (check map and mark for encryption)** • **Check output access list** • **Inspect (Context-based Access Control (CBAC))** • **TCP intercept** • **Encryption** • **Queueing** • **MPLS VRF tunneling (if MPLS WAN deployed)**	• **MPLS tunneling (if MPLS WAN deployed)** • **Decryption (if applicable) for IPsec** • **Check input access list** • **Check input rate limits** • **Input accounting** • **NAT outside to inside (global to local translation)** • **Policy routing** • **Routing** • **Redirect to web cache (WCCP or L2 redirect)** • WAAS application optimization *(start/end of WAAS optimization path)* • Crypto (check map and mark for encryption) • Check output access list • Inspect (Context-based Access Control (CBAC)) • TCP intercept • Encryption • Queueing

1. Source: http://www.cisco.com/en/US/tech/tk648/tk361/technologies_tech_note09186a0080133ddd.shtml

The order of operations here may be important because these application and WAN optimizations, as well as certain IOS behaviors, may not behave as expected, depending on where they are applied. For example, consider the inside-to-outside path in Table 3.

Technology Overview

Deploying WAAS requires an understanding of the network from the data center to the WAN edge to the branch office. This design guide is focused on the data center. A general overview of the data center, WAN edge, and WAAS provides sufficient background for WAAS design and deployment.

Data Center Components

The devices in 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 4 shows a multi-tier front-end network topology and a variety of services that are available at each of these layers.

Figure 4 Data Center Multi-Tier Model Topology

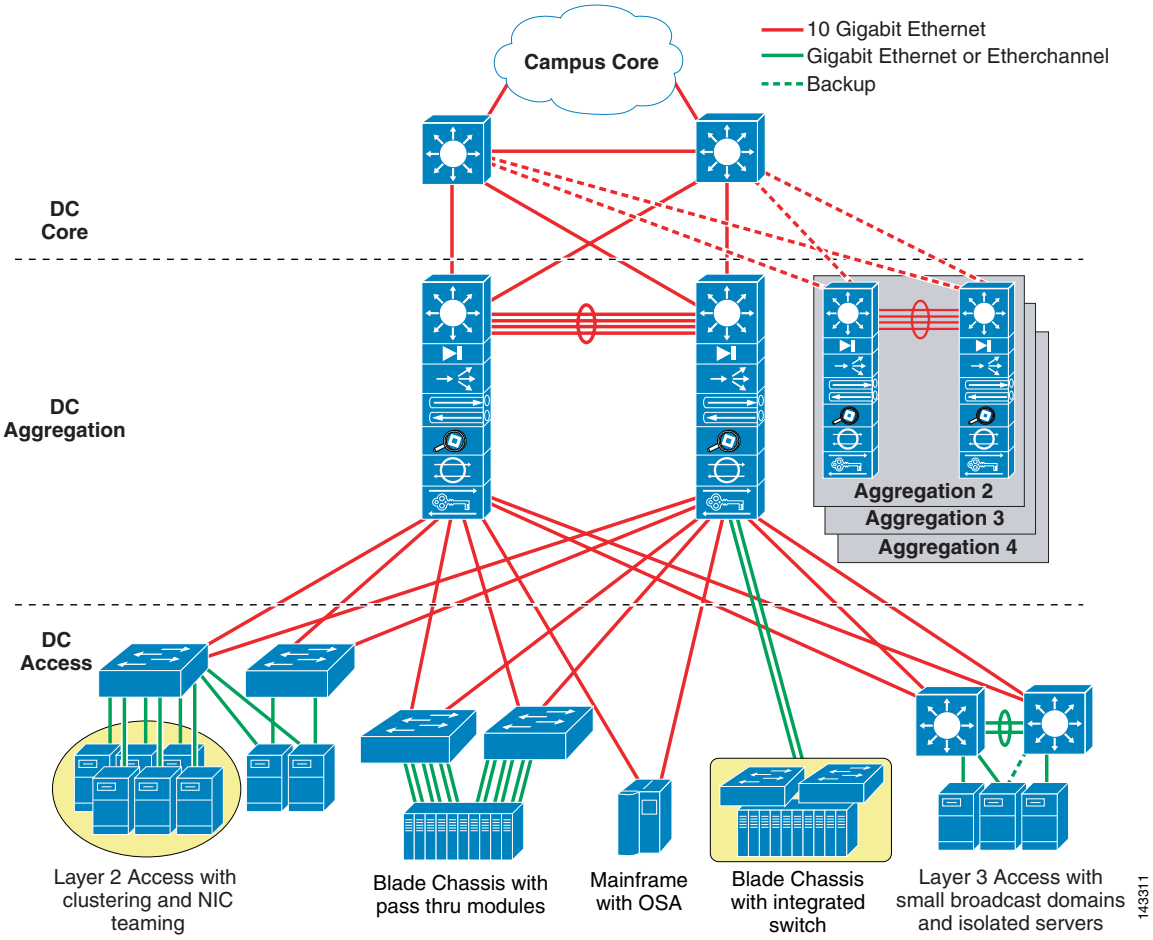

Core Layer

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 employs 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
- 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.

The access layer must provide a deterministic environment to ensure a stable Layer 2 domain. A predictable access layer allows spanning tree to converge and recover quickly during failover and fallback.

Note For more information, see *Integrating Oracle E-Business Suite 11i in the Cisco Data Center* at the following URL: http://www.cisco.com/en/US/docs/solutions/Enterprise/Data_Center/DC_App11i_DG.html.

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 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 5 shows the SAN topology.

Figure 5 SAN Topology

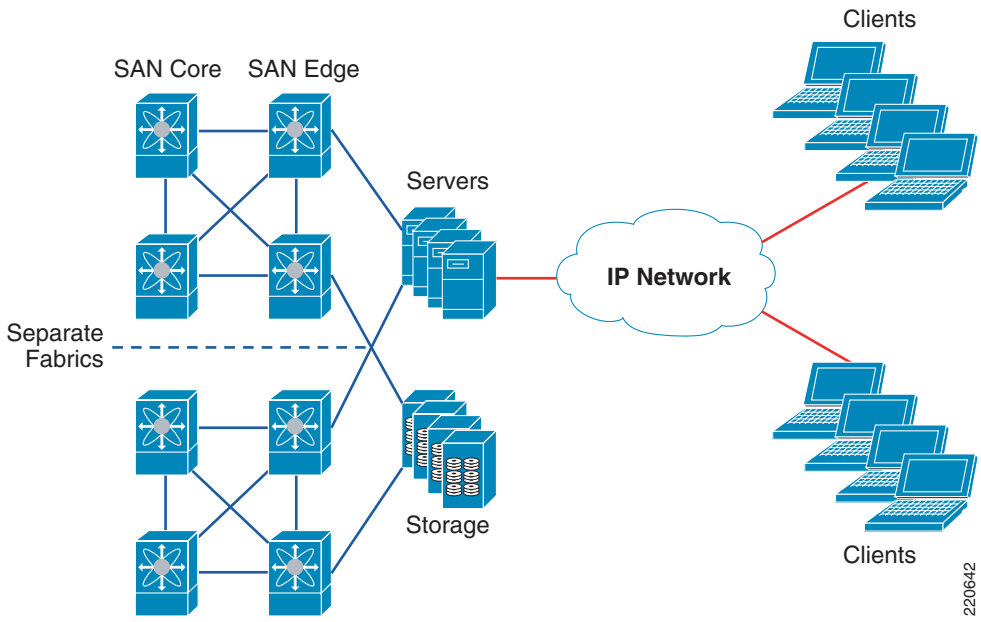

SAN Core Layer

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. Advanced storage functions such as virtualization, continuous data protection, and iSCSI are also found 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 not uncommon to segregate the storage devices to additional edge switches.

WAN Edge Component

The WAN edge component provides connectivity from the campus and data center to branch and remote offices. Connections are aggregated from the branch office to the WAN edge. At the same time, the WAN edge is the first line of defense against outside threats.

There are six components in the secured WAN edge architecture:

- Outer barrier of protection—Firewall or an access control list (ACL) permit only encrypted VPN tunnel traffic and deny all non-permitted traffic; they also protect against DoS attacks and unauthorized access.
- WAN aggregation—Link termination for all connections from branch routers through the private WAN.
- Crypto aggregation—Point-to-point (p2p), Generic Routing Encapsulation (GRE) over IPsec, Dynamic Virtual Tunnel Interface (DVTI), and Dynamic Multipoint VPN (DMVPN) provide IPsec encryption for the tunnels.
- Tunnel interface—GRE and multipoint GRE (mGRE) VTI interfaces are originated and terminated.
- Routing protocol function—Reverse Route Injection (RRI), EIGRP, OSPF, and BGP provide routing mechanisms to connect the branch to the campus and data center network.
- Inner barrier of protection—ASA, Firewall Services Module (FWSM), and PIX provide an inspection engine and rule set that can view unencrypted communication from the branch to the enterprise.

Figure 6 shows the WAN edge topology.

Figure 6 WAN Edge Topology

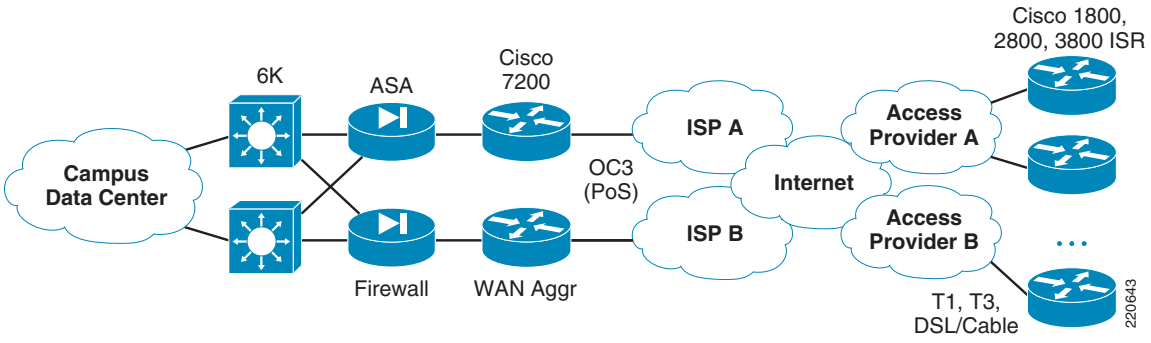

For more information on WAN edge designs, see the following URL: http://www.cisco.com/en/US/netsol/ns817/networking_solutions_program_home.html.

WAAS Design Overview

WAAS can be integrated anywhere in the network path. To achieve maximum benefits, optimum placement of the WAE devices between the origin server (source) and clients (destination) is essential. Incorrect configuration and placement of the WAEs can lead not only to poorly performing applications, but in some cases, network problems can potentially be caused by high CPU and network utilization on the WAEs and routers.

WAAS preserves Layer 4 to Layer 7 information. However, compatibility issues do arise, such as lack of IPv6 and VPN routing and forwarding (VRF) support. Interoperability with other Cisco devices is examined, such as the interactions with firewall modules and the Cisco Application Control Engine (ACE).

Design Requirements

Business productivity relies heavily on application performance and availability. Many current critical applications such as Oracle 11i, Seibel, SAP, and PeopleSoft run in many Fortune 500 company data centers. With the modern dispersed and mobile workforce, workers are scattered in various geographic areas. Regulatory requirements and globalization mandate data centers in multiple locations for disaster recovery purposes. Accessing critical applications and data in a timely and responsive manner is becoming more challenging. Customers accessing data outside their geographic proximity are less productive and more frustrated when application transactions take too long to complete.

WAAS solves the challenge of remote branch users accessing corporate data. WAAS not only reduces latency, but also reduces the amount of traffic carried over the WAN links. Typical customers have WAN links from 256 Kbps to 1.5 Mbps to their remote offices, with an average network delay of 80 milliseconds. These links are aggregated into the data center with redundant components.

The WAAS solution must provide high availability to existing network services. WAAS is also expected to scale from small remote sites to large data centers. Because the WAE can be located anywhere between the origin server and the client, designs must able to accommodate installation of the WAE at various places in the network, such as the data center or WAN edge.

Design Components

The data center is the focus of this document. The key components of any WAAS design consist of the following:

- Cisco high-end WAAS WAE appliance at the data center/WAN edge for aggregation of WAAS services
- Cisco high-end router/switch at the data center/WAN edge for WAAS packet interception
- Cisco NM-WAE or entry level WAAS WAE appliance for termination at the branch/remote sites
- Cisco ISR routers at the branch/remote office for WAAS packet interception

Core Site Architecture

The core site is where WAAS traffic aggregates into the data center, just like the WAN edge aggregates branch connections to the headquarters. However, unlike the WAN edge, WAEs can be placed anywhere between the client and servers. The following diagrams show two points in the network suitable for deploying WAAS core services.

WAE at the WAN Edge

Figure 7 shows WAAS design with WAAS WAE at the WAN edge.

Figure 7 WAAS WAE at the WAN Edge

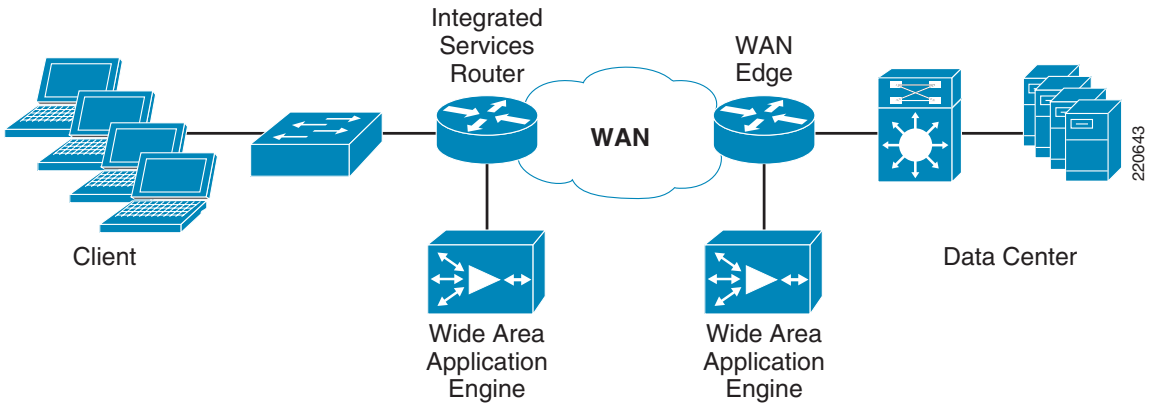

The WAN/branch router intercepts the packets from the client and data center servers. Both WAN edge and branch routers act as proxies for the clients and servers. Data is transferred between the clients and servers transparently, without knowing that the traffic flow is optimized through the WAEs.

WAE at the Aggregation Layer

Figure 8 shows the WAAS design with WAE at the aggregation layer.

Figure 8 WAAS WAE at the Aggregation Layer

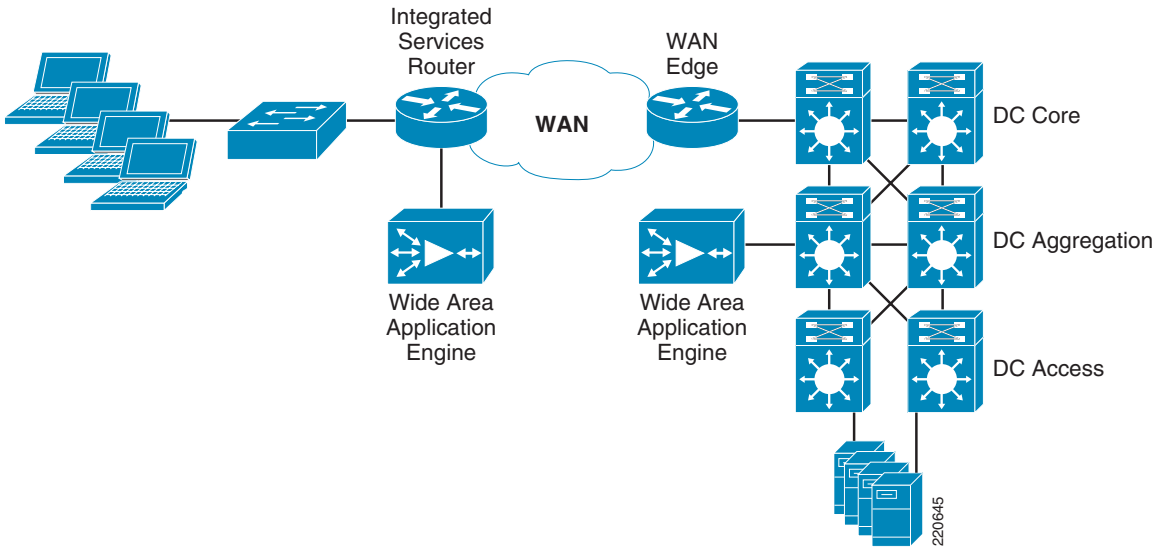

The aggregation switches intercept the packets and forward them to the WAE. The traffic flow is the same as the WAE at the WAN edge. However, much more traffic flows through the aggregation switches. ACLs must filter campus client traffic to prevent overloading the WAE cluster.

WAN Edge versus Data Center Aggregation Interception

WAAS traffic flow and operation is the same regardless of the interception placement. It is suitable to install the WAEs in two places in the network: the WAN edge and the aggregation layer. Each placement strategy has its benefits and drawbacks. The criteria for choosing the appropriate design are based on the following:

- Manageability of the ACLs
- Scalability of the WAEs
- Availability of the WAAS service
- Interoperability with other devices

Consider the following points when planning the WAE placement and configuration in the WAN edge or data center aggregation layer:

- Optimization breadth
  - WAN edge—Connections to any host in the data center/campus are optimized, even connectivity to another PC, unless ACLs are used selectively on optimized preferential servers.
  - Data center aggregation—Only servers connected to the aggregation/access switches are optimized. These hosts are in the data center and are already identified as critical servers.
- WAN topology
  - WAN edge—Complex WAN topologies such as asymmetric routing are supported by WAAS.
  - Data center aggregation—All traffic is directed to servers in the data center; asymmetric routing and complex WAN topologies are avoided in the aggregation layer.
- WCCP ACL configuration
  - WAN edge—ACL configuration is not required because only WAN traffic is optimized when the WAE device is placed at the WAN edge.
  - Data center aggregation—ACL configuration is required because only selected traffic (WAN) traversing the data center should be optimized. Campus and data center traffic must be excluded with ACLs to minimize unnecessary load on the WAEs.
- Physical WAE installation
  - WAN edge—The WAE is generally located in the telecom closet to co-locate with the rest of the WAN equipment.
  - Data center aggregation—The WAE is located in the actual data center facility with the added benefits of UPS, backup generators, and increased physical security.
- ACE integration
  - WAN edge—The ACE module works only on Cisco 7600 Series routers; deployment is limited to a specific hardware platform. Sites installed with Cisco 7200 Series routers are not able to take advantage of the ACE.
  - Data center aggregation—Most installations of aggregation switches are Catalyst 6500s, which do support the ACE module. The ACE is usually used for load balancing of server farms and other application-specific services in addition to the WAEs.
- Other services
  - WAN edge—By terminating the optimization path at the WAN edge, data center and campus traffic is not tampered with, preserving whole TCP packets.

- Data center aggregation—The optimization path extends to the data center aggregation layer. Other services such as deep packet inspection might be hindered because of compressed payload.

Design and Implementation Details

Design Goals

By providing reference architectures, network engineers can quickly access validated designs to incorporate in their own environment. The primary design goals are to accelerate the performance, scalability, and availability of applications in the enterprise network with the WAAS deployments. Consolidation of remote branch servers adds considerable savings to IT operational costs, while at the same time providing LAN-like application performance to remote users.

Design Considerations

Existing network topologies provide references for the WAAS design. Two of the profiles, WAE at the WAN edge and WAE at the WAN edge with firewall, are derivatives of the Cisco Enterprise Solutions Engineering (ESE) Next Generation (NG) WAN design. The core site is assumed to have OC-3 links. Higher bandwidth is achievable with other NGWAN designs. For more information, see the *Infrastructure Protection and Security Service Integration Design for the Next Generation WAN Edge v2.0* at the following URL: http://www.cisco.com/en/US/docs/solutions/Enterprise/WAN_and_MAN/IPSNGWAN.html.

High availability and resiliency are important features of the design. Adding WAAS should not introduce new points of failure to a network that already has many high availability features installed and enabled. Traffic flow can be intercepted with up to 32 routers in the WCCP service group, minimizing flow disruption. The design described is N+1, with WCCP or ACE interception.

For more details, see WAE at the WAN Edge, page 35 and WAE at Aggregation Layer, page 40.

Central Manager

Central Manager (CM) is the management component of WAAS. CM provides a GUI for configuration, monitoring, and management of multiple branch and data center WAEs. CM can scale to support thousands of WAE devices for large-scale deployments. The CM is necessary for making any configuration changes via the web interface. WAAS continues to function in the event of CM failure, but configuration changes via the CM are prohibited. Cisco recommends installing two CMs for WAAS deployment: a primary and a standby. It is preferable to deploy the two CMs in different subnets and different geographical locations if possible.

Centralized reporting can be obtained from the CM. Individually, the WAEs provide basic statistics via the CLI and local device GUI. System-wide application statistics can be generated from the CM GUI. Detailed reports such as total traffic reduction, application mix, and pass-through traffic are available.

The CM also acts as the designated repository for system information and logs. System-wide status is visible on all screens. Clicking the alert icon brings the administrator directly to the error messages.

Figure 9 shows the Central Manager screen with device information and status.

Figure 9 Central Manager Screen

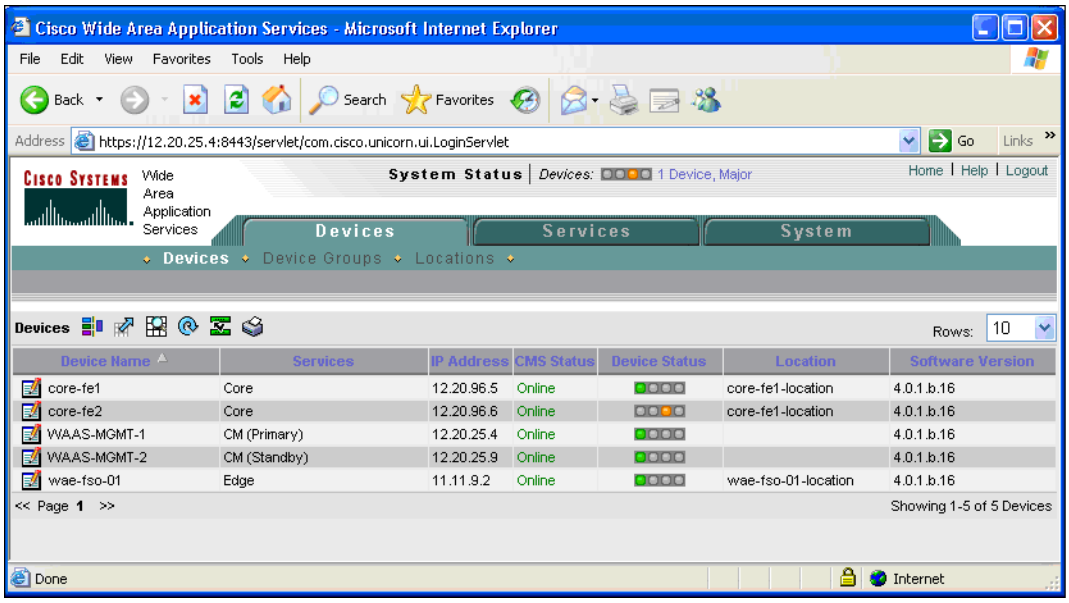

Central Manager can manage many devices at the same time via Device Groups.

CIFS Compatibility

CIFS is the native file sharing protocol for Microsoft products. All Microsoft Windows products use CIFS, from Windows 2003 Server to Windows XP. The Wide Area File Services (WAFS) adapter is the specific WAAS adapter for handling CIFS traffic. The WAFS adapter runs above the foundation layer of WAAS, such as DRE and TFO providing enhanced CIFS protocol optimization. CIFS optimization uses port 4050 between the WAEs. CIFS traffic is transparent to the clients.

Note The CIFS core requires a minimum of 2 GB RAM.

CIFS/DRE Cache

WAAS automatically allocates cache for CIFS. CIFS and DRE cache capacity varies among WAE models. High-end models can accommodate more disks, and therefore have more CIFS and DRE cache capacity. The DRE cache is configured as first in first out (FIFO). DRE contexts are WAE dependent. Unified cache management is not available in the current release.

For more information, see the *Cisco Wide Area Application Services Configuration Guide (Software Version 4.0.1)* at the following URL:
http://www.cisco.com/en/US/docs/app_ntwk_services/waas/waas/v401/quick/guide/wsqcg401.html.

Interception Methods

The ability for the WAE to “see” packets coming in and going out of the router is essential to WAAS optimization. The WAE is rendered useless when it loses this ability. There are four packet interception methods from the router to the WAE:

- PBR

- WCCPv2
- Service policy with ACE
- Inline hardware

Specifics of the interception methods as applied in various scenarios are discussed in detail in Implementation Details, page 35. As a reference, WCCPv2 is used in almost all configurations because of its high availability, scalability, and ease of use.

Table 4 shows the advantages and disadvantages of each interception method.

Table 4 Interception Method Comparison

	Pros	Cons
Policy-Based Routing	• No GRE overhead • Uses CEF for fast switching of packets • Provides failover if multiple next-hop addresses are defined	• Does not scale, cannot load balance among many WAEs • More difficult to configure than WCCPv2
WCCPv2	• Easier to configure than PBR • Uses CEF for fast switching of packets • Can be implemented on any IOS-capable routers (requires v2) • Load balancing and failover capabilities • L2 redirection available on newer CatOS or IOS products • Hardware GRE redirection is available on newer switching platforms	• More CPU intensive than PBR (with software GRE) • Requires additional subnet (tertiary or sub-interface)
Service policy (not tested)	• ACE-configurable load balancing • User-configurable server load balancing (SLB) and health probes • Provides excellent scalability and failover mechanisms	Works on ACE module only, requires Catalyst 6500/7600
Inline hardware (not tested)	• Easy configuration; no need for router configuration • Clear delineation between network and application optimization	Limited inline hardware chaining

Interception Interface

WCCP promiscuous mode uses the following:

- Service 61—Uses the source address to distribute traffic
- Service 62—Uses the destination address

Both these services can be configured on the ingress or egress interface.

Figure 10 shows two traffic flows; one from the client to server, and another from the server to client (blue lines are normal traffic, intercepted traffic are the dotted lines).

Figure 10 Interception Interfaces on the Router

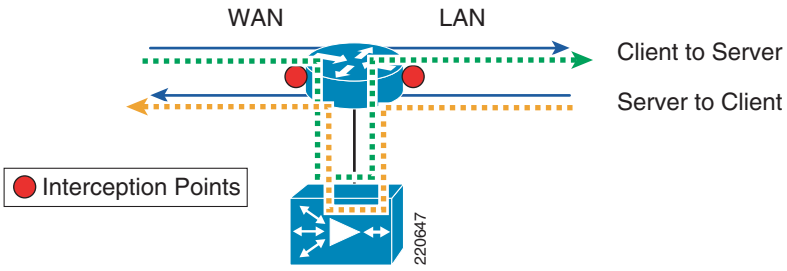

Both traffic flows need to be intercepted by the router and forwarded to the WAE. A number of interception permutations work. The rule is that Service 61 and Service 62 must be used, either on the ingress or egress interface. Both services can also be on the same interface; one for inbound, and another for outbound. The key is to capture both flows; one flow from the client to server, another flow from the server to the client. If an egress interface is used, the **redirect exclude in** command must be configured on the interface connecting to the WAE to avoid a routing loop.

For improved performance, use the **redirect in** command on both the WAN and LAN interfaces; for example, use **redirect in Service 61** on the LAN, and **redirect in Service 62** on the WAN, and vice versa. The packet is redirected to the WAE by the router before switching, saving CPU cycles. Aligning the same IP address on both flows for load distribution can potentially increase performance by using the same WAE for all flows going to the same server. Aligning the IP address based on the server increases DRE use. However, the WAE must be monitored closely for overloading because traffic destined for a particular server goes only to the selected WAE. The WCCP protocol has no way to redirect traffic to another WAE in the event of overloading. Overloaded traffic is forwarded by the WAE as un-optimized traffic.

Table 5 lists the Cisco WAAS and WCCPv2 service group redirection configuration scenarios.

Table 5 Cisco WAAS and WCCPv2 Service Group Redirection Configuration Scenarios

Scenario	Service Group 61	Service Group 62	Redirect Exclusion	Deployment Scenario
1	Inbound, LAN I/F	Inbound, WAN I/F	Not required	Most common branch office or data center deployment scenario
2	Inbound, WAN I/F	Inbound, LAN I/F	Not required	Functionally equivalent to scenario 1
3	Inbound, LAN I/F	Outbound, LAN I/F	Required	Common branch office or data center deployment scenario, used if WAN interface configuration not possible
4	Outbound, LAN I/F	Inbound, LAN I/F	Required	Functionally equivalent to scenario 3

Table 5 Cisco WAAS and WCCPv2 Service Group Redirection Configuration Scenarios

5	Inbound, WAN I/F	Outbound, WAN I/F	Required	Common branch office or data center deployment scenario where router has many LAN interfaces
6	Outbound, WAN I/F	Inbound, WAN I/F	Required	Functionally equivalent to scenario 5
7	Oubound, LAN I/F	Outbound, WAN I/F	Required	Works, but not recommended
8	Outbound, WAN I/F	Outbound, LAN I/F	Required	Works, but not recommended

GRE and L2 Redirection

Packet redirection is the process of forwarding packets from the router to the WAE. The router intercepts the packet and forwards it to the WAE for optimization. The two methods of redirecting packets are Generic Route Encapsulation (GRE) and L2 redirection. GRE is processed at Layer 3 while L2 is processed at Layer 2.

GRE

GRE is a protocol that carries other protocols as its payload, as shown in Figure 11.

Figure 11 GRE Packet

In this case, the payload is a packet from the router to the WAE. GRE works on routing and switching platforms. It allows the WCCP clients to be separate from the router via multiple hops. With WAAS, the WAEs need to be connected directly to a tertiary or sub-interface of the router. Because GRE is processed in software, router CPU utilization increases with GRE redirection. Hardware-assisted GRE redirection is available on the Catalyst 6500 with Sup720.

L2 Redirection

L2 redirection requires the WAE device to be in the same subnet as the router or switch (L2 adjacency). The switch rewrites the destination L2 MAC header with the WAE MAC address. The packet is forwarded without additional lookup. L2 redirection is done in hardware and is available on the Catalyst 6500/7600 platforms. CPU utilization is not impacted because L2 redirection is hardware-assisted; only the first packet is switched by the Multilayer Switch Feature Card (MSFC) with hashing. After the MSFC populates the NetFlow table, subsequent packets are switched in hardware. L2 redirection is preferred over GRE because of lower CPU utilization.

Figure 12 shows an L2 redirection packet.

Figure 12 L2 Redirection Packet

There are two methods to load balance WAEs with L2 redirection:

- Hashing
- Masking

Hashing

Hashing uses 256 buckets for load distribution. The buckets are divided among the WAEs. The designated WAE, which is the one with lowest IP address, populates the buckets with WAE addresses. The hash tables are uploaded to the routers. Redirection with hashing starts with the hash key computed from the packet and hashed to yield an entry in the redirection hash table. This entry indicates the WAE IP address. A NetFlow entry is generated by the MSFC for the first packet. Subsequent packets use the NetFlow entry and are forwarded in hardware.

Masking

Mask assignment can further enhance the performance of L2 redirection. The ternary content addressable memory (TCAM) can be programmed with a combined mask assignment table and redirect list. All redirected packets are switched in hardware, potentially at line rate. The current Catalyst platform supports a 7-bit mask, with default mask of 0x1741 on the source IP address. Fine tuning of the mask can yield better traffic distribution to the WAEs. For example, if a network uses only 191.x.x.x address space, the most significant bit can be re-used on the last 3 octets, such as 0x0751, because the leading octet (191) is always the same.

The following examples show output from **show ip wccp 61 detail** with a mask of 0x7. Notice that four WAEs are equally distributed from address 0 to 7.

```
wccp tcp-promiscuous mask src-ip-mask 0x0 dst-ip-mask 0x7

        Value SrcAddr    DstAddr    SrcPort DstPort CE-IP
        ----- -------    -------    ------- ------- -----
        0000: 0x00000000 0x00000000 0x0000  0x0000  0x0C141D05 (12.20.29.5)
        0001: 0x00000000 0x00000001 0x0000  0x0000  0x0C141D05 (12.20.29.5)
        0002: 0x00000000 0x00000002 0x0000  0x0000  0x0C141D06 (12.20.29.6)
        0003: 0x00000000 0x00000003 0x0000  0x0000  0x0C141D06 (12.20.29.6)
        0004: 0x00000000 0x00000004 0x0000  0x0000  0x0C141D08 (12.20.29.8)
        0005: 0x00000000 0x00000005 0x0000  0x0000  0x0C141D08 (12.20.29.8)
        0006: 0x00000000 0x00000006 0x0000  0x0000  0x0C141D07 (12.20.29.7)
        0007: 0x00000000 0x00000007 0x0000  0x0000  0x0C141D07 (12.20.29.7)
```

Following is the output from **show ip wccp 61 detail** with a mask of 0x13. Four WAEs are equally distributed across 16 addresses. If the IP address ranges are 1.1.1.0 to 1.1.1.7, the mask with 0x7 load balances better than the mask with 0x13, even though they have the same number of masking bits. Care should be taken when setting masking bits for balanced WAE distribution.

```
wccp tcp-promiscuous mask src-ip-mask 0x0 dst-ip-mask 0x13

        0000: 0x00000000 0x00000000 0x0000  0x0000  0x0C141D05 (12.20.29.5)
        0001: 0x00000000 0x00000001 0x0000  0x0000  0x0C141D05 (12.20.29.5)
        0002: 0x00000000 0x00000002 0x0000  0x0000  0x0C141D07 (12.20.29.7)
        0003: 0x00000000 0x00000003 0x0000  0x0000  0x0C141D07 (12.20.29.7)

        0004: 0x00000000 0x00000010 0x0000  0x0000  0x0C141D06 (12.20.29.6)
        0005: 0x00000000 0x00000011 0x0000  0x0000  0x0C141D06 (12.20.29.6)
        0006: 0x00000000 0x00000012 0x0000  0x0000  0x0C141D08 (12.20.29.8)
        0007: 0x00000000 0x00000013 0x0000  0x0000  0x0C141D08 (12.20.29.8)
```

Security

WCCP Security

Interactions between the WAE and router must be investigated to avoid security breaches. Packets are forwarded to the WCCP clients from the routers upon interception. Common clients include WAE and the Cisco Application and Content Networking System (ACNS) cache engine. A third-party device can

pose either as a router with an I_SEE_YOU, or a WCCP client with a HERE_I_AM message. If malicious devices pose as WCCP clients and join the WCCP group, they receive future redirection packets, leading to stolen or leaked data.

WCCP groups can be configured with MD5 password protection. WCCP ACLs reduce denial-of-service (DoS) attacks and passwords indicate authenticity. The group list permits only devices in the access list to join the WCCP group. After the device passes the WCCP ACL, it can be authenticated. Unless the password is known, the device is not able to join the WCCP group.

The following example is a password- and ACL-protected WCCP configuration.

```
ip wccp 61 redirect-list 121 group-list 29 password ese
ip wccp 62 redirect-list 120 group-list 29 password ese

access-list 29 permit 12.20.29.8
```

"Total Messages Denied to Group" shows the number of WCCP messages rejected by the switch that are not members of the ACL. "Authentication failure" shows the results of incorrect group passwords. In the following output, a device is trying to join the WCCP group but is rejected because of an ACL violation.

```
Agg1-6509#sh ip wccp 61
Global WCCP information:
    Router information:
        Router Identifier:                  12.20.1.1
        Protocol Version:                   2.0

    Service Identifier: 61
        Number of Cache Engines:            2
        Number of routers:                  2
        Total Packets Redirected:           0
        Redirect access-list:               121
        Total Packets Denied Redirect:      6
        Total Packets Unassigned:           0
        Group access-list:                  29
        Total Messages Denied to Group:     17991
        Total Authentication failures:      0
```

Service Module Integration

Service modules increase functionalities of the network without adding external appliances. Service modules are line cards that plug into the Catalyst 6500/7600 family. Service modules provide network services such as firewall, load balancing, and traffic monitoring and analysis. Within the layers of the data center network, service modules are commonly deployed in the aggregation layer. The aggregation layer provides a consolidated view of network devices, which makes it ideal for adding additional network services. The aggregation layer also serves as the default gateway in many of the access layer designs.

WAAS WAE placement in the network is discussed in earlier sections. With WAAS and services module integration, the role of service modules and WAEs have to be clearly identified. Service module and WAEs should complement each other and increase network functionality and services. A key consideration with WAAS and service module integration is network transparency. WAAS preserves Layer 3 and Layer 4 information, enabling it to effortlessly integrate with many of the network modules, including the ACE, Intrusion Detection System Module (IDSM), and others.

Application Control Engine

The Cisco Application Control Engine (ACE) is a service module that provides advanced load balancing and protocol control for data center applications. It scales up to 16 Gbps and four million concurrent TCP connections, making it ideal for large data center or service provider data center deployments. The

business benefits of ACE include maximizing application availability, consolidating and virtualizing server farms, increasing application performance, and securing critical business applications. ACE is available for the Catalyst 6500 and 7600 Series routers.

Table 6 shows ACE functionality and business benefits.

Table 6 ACE Functionality and Benefits

Functionality	Business benefit
Layer 3, 4–7 load balancing—High speed load balancing of server farms, firewalls, and other devices	Consolidation of server farms/application acceleration
SSL off-load—Initiates and terminates SSL connections on behalf of the servers, eliminates SSL processing on the server	Application acceleration
Hardware packet inspection—Inspecting traffic flow for protocol compliance, taking corrective action on out-of-compliance packets	Secured applications and data center
Virtual partitions—Multiple partitions (context) can be set up on the ACE, each with its own resources to allow the ACE to scale a large number of applications and server farms	Consolidation of server farms/secured applications

ACE/WAAS Integration Considerations

The following considerations are used in design and implementation of ACE with WAAS:

- Network interoperability

  WAAS and ACE are complementary technologies. They can integrate on various levels; one is simple network integration, another is WAAS with ACE load balancing.

- Network integration

  WAAS and ACE are devices connected to the network. There are no dependencies on either device. WAAS terminates the optimization path, and packets are forwarded to ACE for load balancing or packet inspection. This is a form of service chaining. This set up can be accomplished with WAE at the edge or WAE at the aggregation. A benefit of this approach is the segregation of network resources. ACE and WAAS resources are independent of each other, and can be managed separately, offering the network administrator operational flexibility. This is the preferred integration method for most deployments.

- WAAS with ACE load balancing

  This design increases the interaction between ACE and WAAS. ACE and WAAS now depend on each other, and should be viewed as an single service/entity. Rather than passing packets from WAAS to ACE, as in the above scenario, traffic comes into the ACE, ACE load balances traffic via the WAAS farms, and ACE passes traffic to the server farm. Because ACE load balancing scales higher than WCCP, this integrated approach enables WAAS to reach a higher number of connections. Using ACE to load balance WAAS is suggested for large scale enterprise or service provider data centers where networks traffic has scaled beyond WCCP capability, and where ACE is already deployed. Adding WAAS improves application performance for ACE load balanced server farms.

- Deep packet inspection/protocol compliance

ACE can perform deep packet inspection on HTTP, FTP, DNS, ICMP, and RTSP traffic. Inspections include port 80 misuse, RFC compliance, content/header/URL checksum, FTP reply spoofing, and many others. ACE can also analyze traffic for malformed packets and take corrective action. In a ACE load balanced WAAS context, ACE is in the optimization path, so these deep packet inspections cannot be performed. ACE contexts outside the optimization path can be configured with the deep packet inspection.

For more information on ACE deep packet inspection, see the following URL: http://www.cisco.com/en/US/docs/interfaces_modules/services_modules/ace/v3.00_A1/configuration/security/guide/securgd.html.

- Load balancing predictor

  The load balancing algorithms supported by ACE include the following:

  - Round-robin
  - Least connections
  - Hash address
  - Hash cookie
  - Hash header
  - Hash URL

  The ACE configuration guide recommends a hash address predictor for firewall load balancing because all packets belonging to the same connection must go through the same firewall. WAAS load balancing is similar with DRE caching. Each WAE device holds the cache for existing clients and servers connections. Subsequent network activities can potentially use previously cached data. Connectivity persistence should be preserved to reap the benefits of DRE. Hash addresses guarantee that the same source or destination IP uses the same WAE. Hash address should be used for WAAS load balancing.

Note Large enterprise or service provider might have proxies installed. All connections go through these proxies. Proxies can disrupt load balancing and mask network traffic information such as source and destination addresses.

  Round-robin and least connections can also be potentially used. Round-robin eventually populates all the WAE farm DRE cache with the same data, because all requests are evenly distributed to all WAEs. Each connection to the WAE farm cycles through different WAEs, resulting in duplicate DRE caches throughout the WAE farm. Round-robin is best used in high throughput deployments. Least connections assigns incoming connections to the WAE with the least number of connections. Again, it does not take into consideration the DRE caching. In the context of maximizing DRE cache usage, hash address is preferred over round-robin and least connections.

  L7 load balancing includes hash cookie, header, and URL. These load balancing techniques require payload inspection. ACE can perform L7 load balancing with unoptimized traffic. ACE cannot use L7 algorithms to load balance WAE farms because WAAS packets are compressed.

  ACE load balance on a per-connection basis. Incoming and outgoing traffic have to be on the same WAE for WAAS to work. Sticky-mac is required for ACE to forwarding traffic from and to the same WAE.

- WAE sharing

With WCCP, WAEs are load balanced and shared for all incoming connections. WCCP-intercepted traffic is forwarded to the WAE based on the bucket placement algorithm, hashing, or masking with IP addresses. For critical applications that cache data that cannot be on shared WAEs, such as service provider or financial institutions, WAEs can be segregated by ACE contexts. Each ACE context can have its own farm of WAEs. WAE DRE caches do not cross-contaminate.

- WAN edge or aggregation layer

  ACE/WAAS load balancing can be deployed in the WAN edge or in the aggregation layer. In the WAN edge, it functions similarly to a WAN aggregator, passing traffic between remote offices and the data center. Traffic is intercepted by ACE and forwarded to WAAS, which passes the traffic back to ACE. ACE then forwards the traffic to the data center. The series of steps are the same with ACE/WAAS at the aggregation layer.

  See WAE at the WAN Edge, page 35 and WAE at Aggregation Layer, page 40 for placement strategy.

This design focuses on WAAS load balancing with ACE in the aggregation layer. The ACE context is running as a Layer 4 server load balancer for WAAS. ACE functions such as SSL offload, Layer 7 load balancing, and protocol compliance are not necessary when ACE is load balancing WAAS. Other contexts or policies can continue to use full ACE functionality. Configurations in one context do not affect another context, with the exception of public IP addresses, which cannot be shared on multiple contexts.

Table 7 shows the features in WAAS load balanced context compared to the normal ACE context.

Table 7 WAAS Load Balancing versus Server Load Balancing

ACE feature	WAAS Load Balancing	Server Load Balancing
SSL offloading	No	Yes
L4 load balancing	Yes	Yes
L7 load balancing	Yes, on the server farm after WAAS, not WAAS farm	Yes
Protocol compliance	No	Yes

SSL offloading is not supported because ACE is now in the WAAS optimization path. The network connection terminates with the WAAS device, not ACE. Layer 7 load balancing methods such as URL and cookie-based load balancing are not used because of the lack of visibility of the payload, but Layer 7 load balancing can be done by the server farm after the WAAS farm. Protocol compliance is also not used for the same reason. ACE supports multiple contexts on the same line card: both WAAS and non-WAAS contexts at the same time.

ACE with WAAS Packet Flow

WAAS intercepts the packets at router endpoints; both the client and server. WAAS setup employs WCCP interception at the branch and data center. ACE load-balanced WAEs use ACE to intercept data center traffic.

Figure 13 shows traffic flow with ACE load balancing WAEs and server farm for the TCP handshake.

Figure 13 WAAS with ACE Load-Balancing Packet Flow

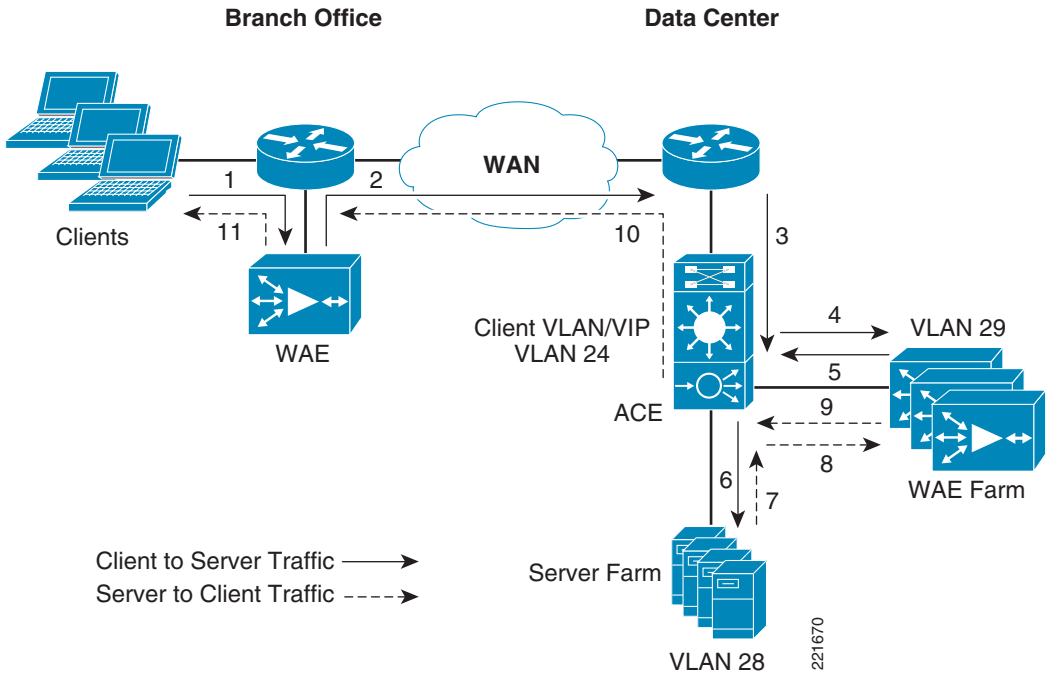

The following sequence takes place:

1. The client sends a packet to the server farm VIP address. A Syn packet is forwarded to the branch router, which intercepts the packet with WCCP. The packet is forwarded to the WAE.
2. The WAE marks the packet with TCP option 0x21 (First device ID and policy is marked), and forwards the packet out via the default gateway to the router. The router sends the packet to the WAN.
3. The packet arrives on the WAN edge router. Interception is not configured on the WAN edge router. The packet is forwarded to the switch and the ACE VIP.
4. The ACE checks the service policy on the client VLAN (vlan 24), and forwards the packets according to the service policy; in this case, to the WAE farm in vlan 29.
5. The WAE inspects the packet. It finds that the first device ID and policy is populated, and updates the last device ID field (first device ID and policy parameters are unchanged). The packet is forwarded back to the ACE via the default gateway.
6. Packets are routed and forwarded within the ACE to the server farm VLAN (vlan 28) by the appropriate service policy with TCP option 21 removed.
7. The server farm receives the packet and sends the Syn/Ack packet back to the client, with no TCP option. TCP options are usually ignored by the server, even if it is still in place.
8. Traffic from the server farm VLAN is matched and forwarded to the WAE farm on vlan 29. Sticky mac is enabled on the ACE. The ACE knows which WAE initiated the connection and sends the packet back to the originating WAE.
9. This is like Step 2, except for reverse traffic flow. The WAE marks the packet with TCP option 0x21 and forwards the packet back to the ACE via the default gateway.

10. Packets are sent to the client from the ACE. The branch router intercepts the packet and forwards it to the branch WAE. The branch WAE knows it initiated this connection (from the syn in step 1), and now it knows its first WAE in the path, itself. It also know the last WAE and optimization policy by examining the first device ID under option 21 on the syn/ack reply.
11. Branch WAE forwards the packet to the client.

The first and last WAE and optimization policy are now identified. TCP proxy for this connection on the WAEs start. Subsequent transfers on this connection from the client to server go through the WAE TCP proxy. The WAEs spoof client and server IP addresses, adding 2 GB to the sequence number of the WAE-to-WAE TCP connection. A big sequence number difference would prevent the client and/or server from the accidental use of the WAE-to-WAE TCP proxy connections.

WAE Network Connectivity

WAN Edge

In the WAN edge, the WAE can connect directly to the WAN router, which is not possible in many cases with multiple WAE deployments. Interfaces on the WAN router are scarce. A better alternative is to connect a switch to the WAN router, then attach WAEs to the switch. The switch not only expands connectivity capacity, it also provides better availability if properly configured. See WAE at the WAN Edge, page 35 for a sample topology.

Data Center Aggregation

In the data center, the WAE can connect to the aggregation or access switches. Because the interception is configured on the aggregation switch, connectivity to the aggregation switches results in faster traffic going in and leaving the WAE. Other services present in the aggregation switch include FWSM, ACE, and NAM. Aggregation switches also consolidate access switches to the core switches. Port availability on the aggregation switch should be considered.

Most of the WAE deployments with Catalyst 6000 switches use L2 redirection. The WAE can connect to access switches as long as it has L2 adjacency with the aggregation switch. Traffic has an extra hop to and from the access switch from the aggregation switch. This hop is insignificant in terms of the overall traffic path. In a highly available setup with standby interfaces, the same VLAN must be on both access and aggregation switches.

PortFast

PortFast reduces the time for spanning tree by immediately moving to forwarding state, bypassing block, listening, and learning states. The average time for switchport going into a forward state is approximately 30 seconds. Using PortFast reduces this time to approximately 5 seconds. The WAE is used as a host device, and does no routing or switching, but uses a default gateway. There is no need for WAE-connected ports to go through all the initialization states.

In the access layer, all host ports should be enabled with PortFast. Host ports in the aggregation layer are not as common. Because the aggregation switch has many switch connections, accidental connections from another switch to the WAE ports can occur. The local network administrator should able to provide guidelines for host ports in the aggregation switches.

Note As a caution, note that PortFast is used only with host ports; never connect any hubs, switches, or routers. Because PortFast skips some steps and moves directly to the forwarding state, it can cause spanning tree loops and possibly bring the network down.

For more information, see the following URL: http://www.cisco.com/en/US/products/hw/switches/ps700/products_tech_note09186a00800b1500.shtml.

Tertiary/Sub-interface

A tertiary or sub-interface and an additional routable subnet on the switch or router are necessary for transparent traffic flow between the client and server. When traffic is forwarded to the WAE from the router, the TCP headers are preserved. After the WAE processes the packet, it is sent back to the router with full header preservation, including the original source and destination IP address. For the router to identify traffic from the WAE, the subnet in which the WAE resides must be a distinct subnet to avoid the possibility of a routing loop. The subnet also needs to be routable because the WAEs keep communication with the Central Manager for system updates, status reporting, message logging, and configuration management. For WAAS deployment in the aggregation layer, a separate VLAN for WAEs is recommended for connecting multiple WAEs. Inline deployments do not require a tertiary or sub-interface.

High Availability

The WAAS service must be highly available in the data center. WAE does not incur downtime for clients; when the WAE is unavailable, the router removes the WAE from the WCCP list and forwards the packets normally. However, WAAS service interruptions can cause application delays (without optimization) for remote clients. In addition to the topics below, the WAE cluster should be configured with N+1 for high availability and scalability.

Device High Availability

The WAEs have many built-in high availability features. The disk subsystem is recommended to be configured with RAID 1 protection. RAID 1 is mandatory when two or more drives are installed in the WAE. With RAID 1, failure of the physical drive does not affect normal operations. Failed disks can be replaced during planned downtime. Multiple network interfaces are available. Standby interfaces can be configured for interface failover. A standby interface group guards against network interface failure on the WAE and switch. When connected to separate switches in active/standby mode, the standby interface protects the WAE from switch failure.

Loopback Interface

Loopback interface identifies the router to the WAEs. If the loopback interface is not defined, the highest available IP address is used. The WCCP protocol relies on the router ID to communicate to the service group. Router ID change leads to router view rebuilds. Flapping of the interface with router ID can cause lost connectivity to the service group. Although loopback interface is not mandatory, it is highly recommended, especially if high availability is a requirement.

The following log demonstrates **shut** and **noshut** interface loopback 0 that resulted in loss connectivity to the service group:

```
dc-7200-02(config-if)#shut
Mar  8 12:37:51.847 UTC: %LINK-3-UPDOWN: Interface Loopback0, changed state to up
dc-7200-02(config-if)#no shut
Mar  8 12:37:54.155 UTC: %LINK-5-CHANGED: Interface Loopback0, changed state to
administratively down
dc-7200-02(config-if)#shut
Mar  8 12:37:57.955 UTC: %LINK-3-UPDOWN: Interface Loopback0, changed state to up
Mar  8 12:37:58.955 UTC: %LINEPROTO-5-UPDOWN: Line protocol on Interface Loopback0,
changed state to upt
```

```
Mar  8 12:38:02.499 UTC: %LINK-5-CHANGED: Interface Loopback0, changed state to
administratively down
Mar  8 12:38:03.499 UTC: %LINEPROTO-5-UPDOWN: Line protocol on Interface Loopback0,
changed state to down
Mar  8 12:38:03.743 UTC: %WCCP-1-SERVICELOST: Service 61 lost on WCCP client 12.20.96.6
dc-7200-02(config-if)#no shut
Mar  8 12:38:20.295 UTC: %LINK-3-UPDOWN: Interface Loopback0, changed state to up
Mar  8 12:38:21.295 UTC: %LINEPROTO-5-UPDOWN: Line protocol on Interface Loopback0,
changed state to up
Mar  8 12:38:30.743 UTC: %WCCP-5-SERVICEFOUND: Service 61 acquired on WCCP client
12.20.96.6
```

Standby Interface Group

The WAE can be set up with a standby interface group. A standby interface is configured with the real IP address, while the physical interfaces are configured as part of the standby group. The physical network interface is connected to two different switches for redundancy. Although the physical interfaces are not configured with an IP address, they are in an UP state. The standby IP address is attached to the physical interface with the highest priority. In the event of an interface, link, or switch failure, the standby IP address attaches to the secondary physical interface. Failover time with the standby interface is approximately 5 seconds. Depending on the transaction, TCP session recovery is possible.

Standby interface supports GRE and L2 redirection with hashing. L2 redirection with masking is incompatible at this time.

WCCP High Availability

WAAS can be configured to be highly available with WCCP, PBR, inline, and the ACE module. This section describes WCCP high availability. The WCCP protocol can have up to 32 routers and 32 devices (WAEs) per service group. WCCP devices communicate with I_SEE_YOU and HERE_I_AM requests in ten-second intervals. In the event of a WAE failure and/or the WAE fails to respond within 25 seconds of the I_SEE_YOU request, the router sends a REMOVAL_QUERY to the WAE. If the WAE fails to respond within five seconds to the REMOVAL_QUERY message, the router removes the failed WAE and updates the WCCP client list. It can take up to 30 seconds for the router to detect failed WAEs. The message timers in WCCPv2 are fixed and are not tuneable. Existing connections are dropped in the event of a WAE failure. WAAS flow protection is supported when new WAEs are added to the service group.

One way to reduce failover time is use the standby interface. From observation, standby interface failover takes an average of 5 seconds, which is much less than the 30 seconds with WCCP. However, the standby interface does not protect against WAE device failure. The standby interface should be used in addition to WCCP.

For more information on WCCP, see the following URL:
http://www.wrec.org/Drafts/draft-wilson-wrec-wccp-v2-00.txt

HSRP/mHSRP/GLBP

Hot Standby Routing Protocol (HSRP) provides a fault-tolerant IP gateway. Two or more routers can form an HSRP group, comprising an active and standby router, and sharing a virtual MAC and virtual IP address. In the event of a failed active router, the standby router takes over the virtual MAC and IP address to continue forwarding packets. The HSRP protocol leaves redundant routers idle and under-used.

Multiple HSRP (mHSRP) groups can be configured with many virtual IP addresses and virtual MAC addresses. Multiple active gateways are set up on different routers within the same group, allowing load balancing for outgoing connections. mHSRP is HSRP with many groups on the same set of routers.

Gateway Load Balancing Protocol (GLBP) also provides a redundant router for IP hosts. The main difference between HSRP and GLBP is that GLBP allows all routers (active virtual forwarder) to participate in forwarding traffic. Unlike mHSRP, GLBP uses only one virtual IP address. A virtual MAC address is assigned to each active virtual forwarder (AVF). The active virtual gateway (AVG) responds to ARP requests with the virtual MAC address of the AVFs.

WCCP redirects traffic to the WAE cluster upon interception. Return traffic from the WAE cluster is forwarded back to the router via the default gateway. Multiple active gateways should be configured to load balance traffic leaving the WAE cluster. While GLBP can be used to load balance outgoing traffic, the AVG determines the load balancing method in a non-deterministic fashion. With mHSRP, manual assignment of the default gateway is more deterministic.

On the left of Figure 14, a single HSRP group uses one active router. When using multiple HSRP groups, traffic can be load balanced across many routers, as shown on the right of Figure 14.

Figure 14 Outbound Load Balancing with mHSRP

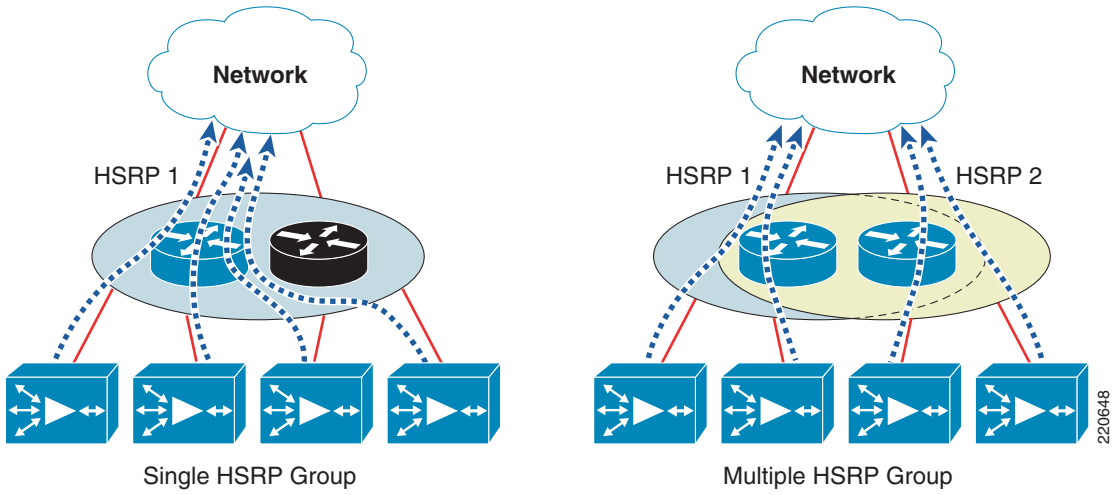

Scalability

Traffic in the data center can overwhelm any single device, so clustering of the core WAEs is recommended. Two WAEs are the minimum for a core WAE cluster. Additional WAEs can be added for N+1 configuration, up to a maximum of 32 WAEs with WCCP. WAAS service scales in a near-linear fashion with N+1 configuration. Number of connections, number of users, and traffic usage determines the WAE capacity required at the data center. NetFlow information, user sessions from the Windows server manager, and other network tools can assist in WAE planning. Table 8 provides current WAE family capacity and performance information.

Table 8 WAE Family Performance and Scalability

Device	Max Optimized TCP Connections	Max CIFS Sessions	Single Drive Capacity [GB]	Max Drives	RAM [GB]	Max Recommended WAN Link [Mbps]	Max Optimized Throughput [Mbps]	Max Core Fan-out [Peers]	CM Scalability [Devices]
NME-WAE-302	250	N/A	80	1	0.5	4	90		
NME-WAE-502	500	500	120	1	1	4	150		

Table 8 WAE Family Performance and Scalability

WAE-512-1	750	750	250	2	1	8	100	5	500
WAE-512-2	1500	1500	250	2	2	20	150	10	1000
WAE-612-2	2000	2000	300	2	2	45	250	30	2000
WAE-612-4	6000	2500	300	2	4	90	350	50	2500
WAE-7326	7500	2500	300	6	4	155	450	96	

Figure 15 shows N+1 WAE configuration.

Figure 15 N+1 WAE Configuration

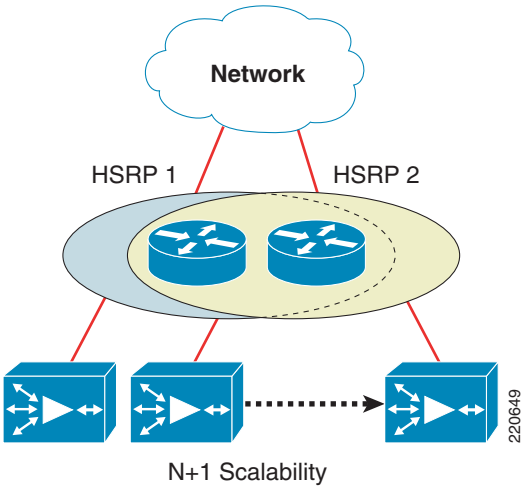

Table 9 shows the scalability for each of the interception methods. Of the four methods, WCCP and ACE integration are recommended in the data center. PBR and inline hardware are not recommended because of their limited scalability.

Table 9 Scalability of Interception Methods

Method	Remarks
PBR	• Minimal scalability • Cannot load balance among equal cost paths, provides failover only via next hop
WCCPv2	• High scalability • Up to 32 routers and WAEs in a service group • Load balancing with a hash algorithm or masking with appropriate hardware • Line rate redirection with Cat 6000 platform

Table 9 Scalability of Interception Methods

ACE	• High scalability • Handles up to 16,000 WAEs • Ability to run health probes on the WAE to determine best load balancing algorithm • Load balancing can be customized
Inline hardware	• Minimal scalability • Cannot have multiple equal load balancing WAEs because of physical design and limitations • Serial load balancing with spillover load balancing

Core Fan Out

Core WAE to edge WAE fan out determines the number of core WAEs required in relation to the number of edge WAEs. A minimum of two WAEs are recommended for high availability. The fan out ratio is determined by a number of factors, such as the number of WAEs in the branch offices, amount of network traffic, and number of TCP connections. A sizing tool is available internally that can help automate sizing decisions. (See Table 8.)

Implementation Details

This section details the integration of the WAAS WAE in the data center and WAN edge network. Although the WAE can be placed anywhere within the client and server, placements at the WAN edge and data center aggregation layer are preferred. The following section discusses in detail the two implementations and WAAS with ACE load balancing in the aggregation layer.

Central Manager

Central Manager is the main configuration tool for WAAS. As expected, WAAS continues to function even when the primary Central Manager is offline. However, no configuration changes can be made. Promotion of the standby Central Manager to the primary Central Manager is configured from the CLI of the standby Central Manager by changing the Central Manager role from standby to primary. When the failed primary Central Manager is back online, it must be configured for a standby CM role.

Central Manager needs to configured with proper device mode (from config mode):

```
device mode central-manager
```

Standby Central Manager must be configured as well:

```
central manager role standby
```

WAE at the WAN Edge

WAE at the WAN edge places the WAE at the WAN aggregation point (see Figure 16).

Figure 16 Network Topology of WAE at WAN Edge

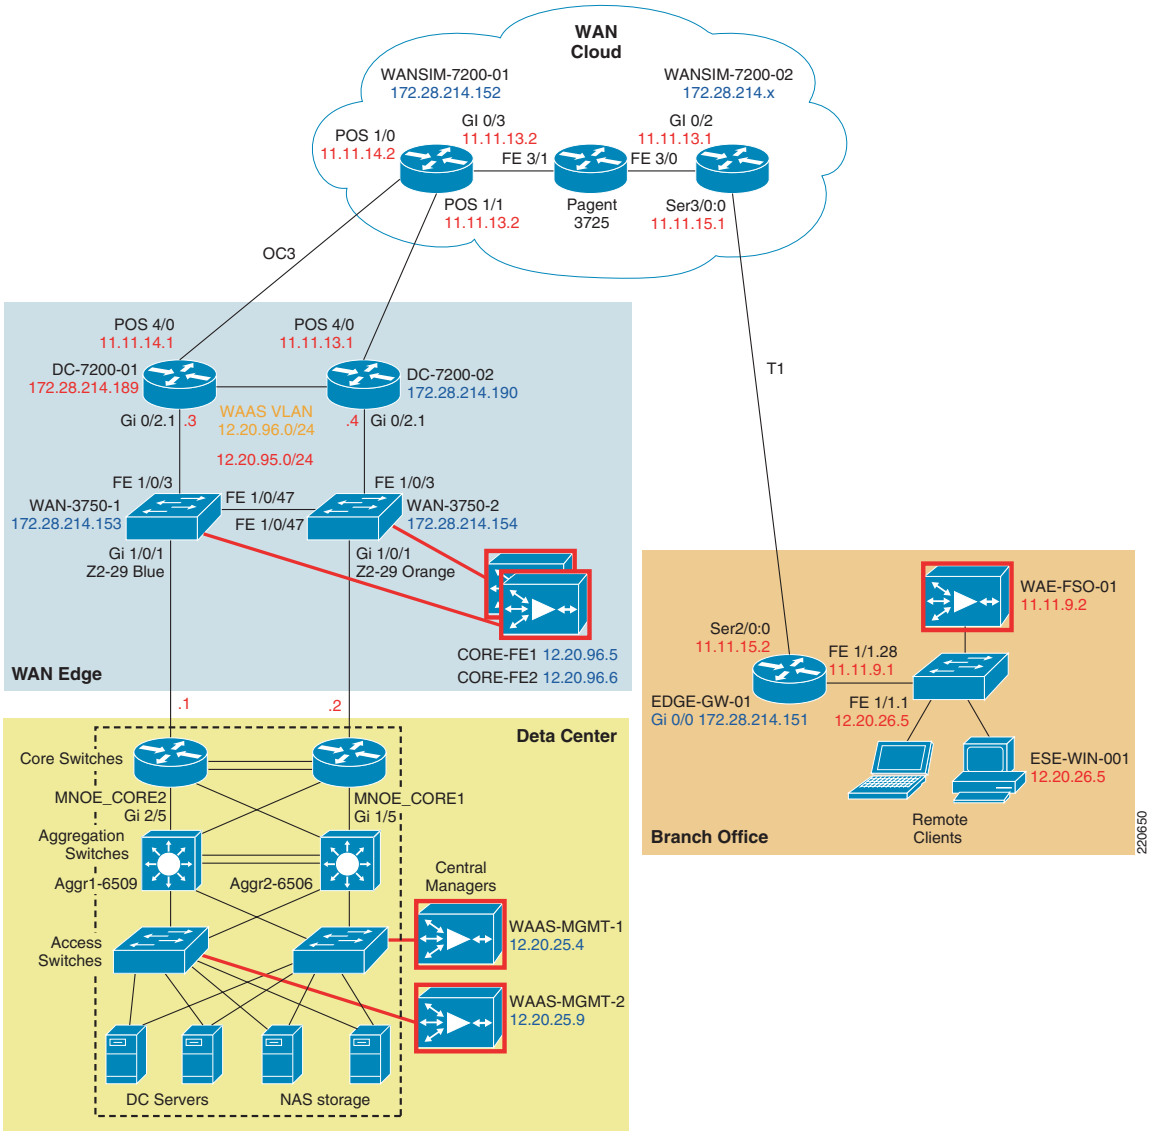

The data center and branch office are connected to a simulated WAN with OC3 connection at the WAN edge and T1 connection at the branch office. Network delay is set at 100 milliseconds. Core WAEs installed at the WAN edge connect to the two 7200 routers. The edge WAE is connected to the 3825 router at the branch office. Core traffic interception is placed at the WAN edge to optimize WAN-only traffic. Campus and data center traffic is not in the optimization path. The main benefit is that all WAN traffic is optimized without the burden of ACL management. The optimization path terminates at the WAN edge routers. Redundant routers and switches are in place to avoid any single point of failure. The switches share the same WAE VLAN and are interconnected. Failure of either the WAN edge router or the switch results in path re-convergence.

The WAN-facing interface of the router should have an outer barrier of protection such as ACLs or firewall configured. A WCCP redirect list is not necessary unless specific servers are targeted for optimization. WCCP provides inbound load balancing to the WAEs. The WAE forwards optimized traffic out via the default gateway. Traffic going out the WAEs is manually load balanced by two mHSRP groups with WAEs pointing to their respective mHSRP group.

The advantages and disadvantages of this architecture are discussed in WAAS Design Overview, page 16.

Following is a quick summary of benefits:

- No need to filter out WAN traffic
- Unencumbered by campus or data center traffic
- Optimization applies to all devices in the data center by default
- Extendable to include inner firewall
- Terminates optimization path at the WAN edge

Asymmetrical routing is supported as long as the routers are redirecting traffic to the same set of WAEs. For information on asymmetric routing, see the *Cisco Wide Area Application Services Configuration Guide (Software Version 4.0.1)* at the following URL: http://www.cisco.com/en/US/docs/app_ntwk_services/waas/waas/v401/quick/guide/wsqcg401.html.

For additional WAN edge information, see *Infrastructure Protection and Security Service Integration Design for the Next Generation WAN Edge v2.0* at the following URL: http://www.cisco.com/en/US/docs/solutions/Enterprise/WAN_and_MAN/IPSNGWAN.html.

Sub-Interface

The sub-interface is configured for the WAE at the WAN edge. On the WAN router, extra LAN interfaces are usually at a premium, so using the sub-interface reduces overall cost. This design is the one-arm approach, where the WAE connects to the router via one interface. A drawback of the one-arm approach is that traffic has to flow to the WAE and back to the router on the same interface. Because only WAN traffic passes through the WAN edge router, careful capacity planning can reduce the likelihood of over-saturating the sub-interface.

The following example show the configuration of the sub-interface on the branch router for the WAE VLAN:

```
interface GigabitEthernet0/2.96
 description WAE-VLAN
 encapsulation dot1Q 96
 ip address 12.20.96.1 255.255.255.0
```

An alternative is the two-arm mode, using two interfaces from the WAE and two interfaces from the router (see Figure 17). Traffic coming into the WAE and going out the router can be segregated. The two-arm approach offers better load distribution on the router with separate ingress and egress ports for WAE traffic. However, this configuration is not commonly used because of the additional interface required at the router.

Figure 17 Two-Arm Approach

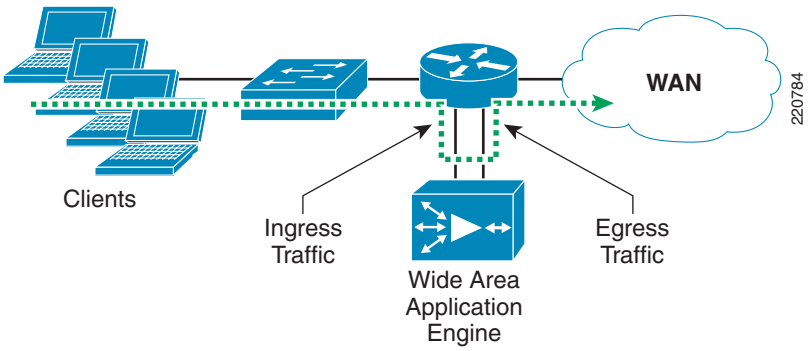

Interception Interface

Both service 61 and 62 interception is configured on the LAN interface of the router. Configuring both interceptions on a single interface requires redirect in and redirect out. Use of redirect out has a slight performance penalty because the packet is switched by the router before being forwarded to the WAE. The use of redirect out is justified. The WAN router is responsible for WAN aggregation, crypto, outer barrier ACL, and other functions. In some cases, the WAN-bound interface cannot be used for interception because of IPsec or VPN encryption. See the following section if CPU utilization is a concern with 7200 Series routers.

Both redirect in and redirect out are configured on the LAN-bound interface. This configuration is used on the 7200 Series routers. 6500s/7600s should use redirect in on two interfaces. For more information, see WAE at Aggregation Layer, page 40.

```
interface GigabitEthernet0/2.1
 encapsulation dot1Q 1 native
 ip address 12.20.95.4 255.255.255.0
 ip wccp 61 redirect in
 ip wccp 62 redirect out
```

GRE Redirection

GRE redirection is supported on most router platforms. The Catalyst 6500 and 7600 platforms also support L2 redirection. GRE Redirection is processed in software, which means redirection with GRE consumes CPU resources. Fortunately, the WAN routers do not have to handle as much traffic as do the data center switches. Although the 7200 platform can perform comfortably with WCCP and GRE redirection, the router should be monitored for CPU utilization because the WAN router also handles WAN aggregation, crypto aggregation, tunnel interfaces, and routing protocols. Upgrades to the 7600 platform should be considered if high utilization is observed with the 7200 platform. Segregating the crypto function to another router is also an alternative.

GRE redirection is configured by default if L2 or GRE redirection is not specified.

High Availability

Both WAN routers and WAEs are connected to the L2 switch set on the same VLAN. Interception is configured on both routers. The routers and the WAEs form a single WCCP service group. All components in this design are redundant; failure of any single component does not disrupt either the WAN or WAAS services.

Outbound Load Balancing

Intercepted traffic is forwarded from the routers to WAEs, which do not understand routing protocols. Optimized traffic is sent out via the default gateway, regardless of the source router. Two HSRP groups and two default gateways are configured for outbound traffic distribution. Traffic from the WAEs is manually load balanced between the routers. Additional WAEs can be added for increased throughput. The default gateway is manually configured on each WAE for appropriate out-bound traffic load balancing.

- DC-7200-01 mHSRP configuration:

```
interface GigabitEthernet0/2.96
 description WAE-VLAN
 encapsulation dot1Q 96
 ip address 12.20.96.3 255.255.255.0
 ip wccp redirect exclude in
 standby 1 ip 12.20.96.1
 standby 1 timers 1 3
 standby 1 priority 120
 standby 1 preempt delay minimum 1
 standby 2 ip 12.20.96.2
 standby 2 timers 1 3
 standby 2 priority 110
 standby 2 preempt delay minimum 1
```

- DC-7200-01 mHSRP configuration:

```
interface GigabitEthernet0/2.96
 encapsulation dot1Q 96
 ip address 12.20.96.4 255.255.255.0
 ip wccp redirect exclude in
 standby 1 ip 12.20.96.1
 standby 1 timers 1 3
 standby 1 priority 110
 standby 1 preempt delay minimum 1
 standby 2 ip 12.20.96.2
 standby 2 timers 1 3
 standby 2 priority 120
 standby 2 preempt delay minimum 1
```

Standby Interface

The standby interface is configured to protect from network interface, link, and switch failure. The default priority of the physical interface is 100. GigabitEthernet 2/0 is configured with a higher priority of 105. It is the active interface with the IP address when the WAE is online.

```
interface Standby 1
 description standby interface group
 ip address 12.20.96.5 255.255.255.0
 exit
!
interface GigabitEthernet 1/0
 description WAAS interface
 standby 1
 exit
interface GigabitEthernet 2/0
 description WAAS interface
 standby 1 priority 105
 exit
```

WAE at Aggregation Layer

Packet interception in the aggregation layer potentially exposes the WAE to all data center traffic. However, the aggregation layer is not burdened with high-speed switching like the core layer. Traditional services such as QoS, load balancing, and newer services such as the ACE and FWSM operate in the aggregation layer. (See Figure 18.)

Figure 18 Packet Interception at Aggregation

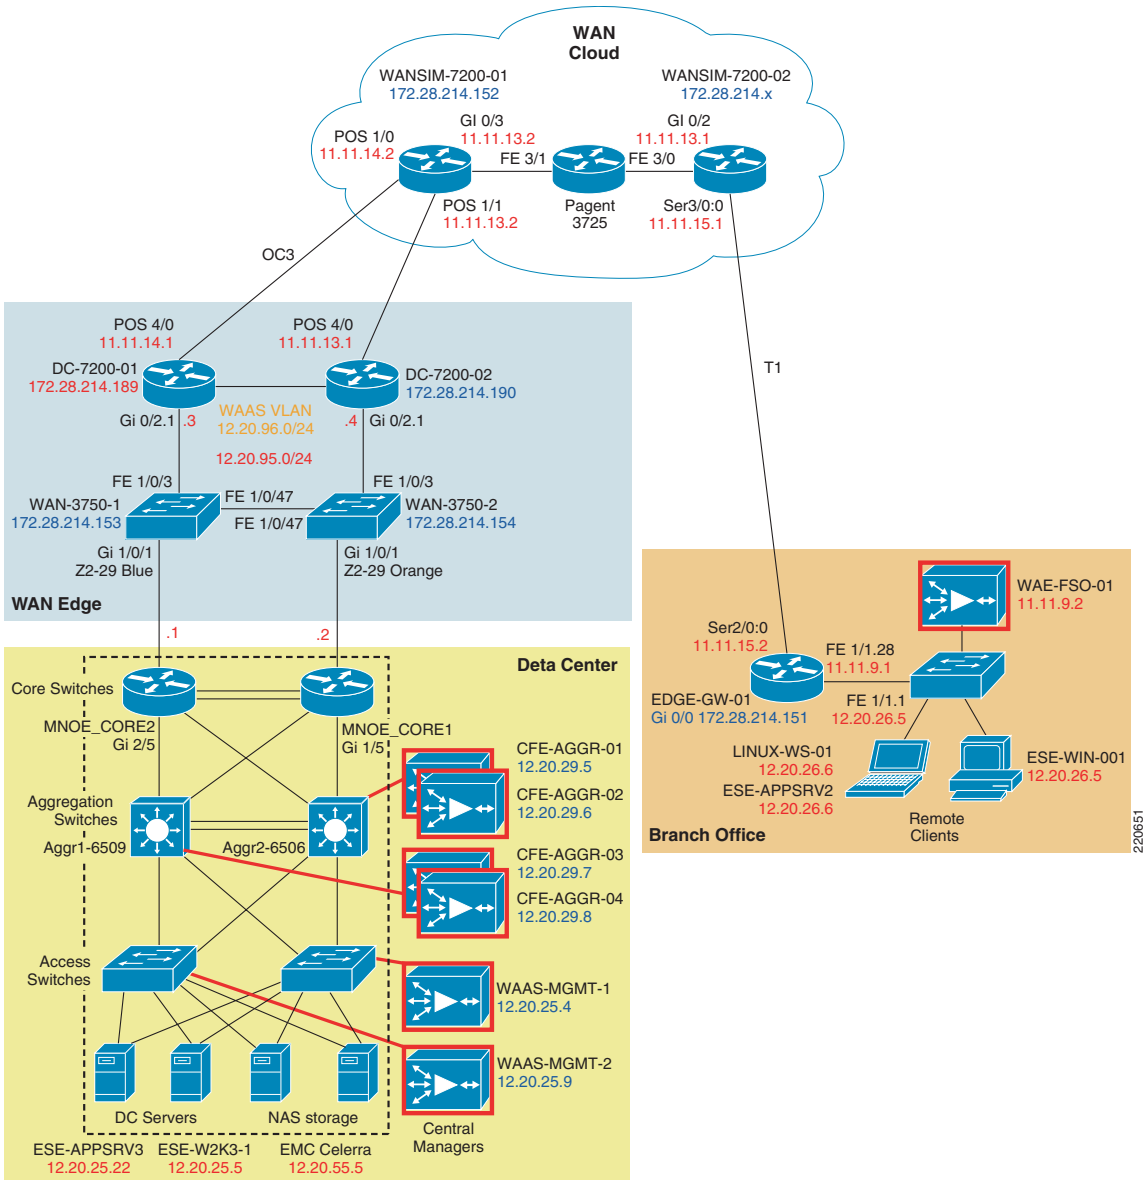

The WAE packet interception does not fit better at the core layer of the data center, because the data center core is reserved for high-speed switching (in some case, a shared core between campus and data center). Best practices suggest that end devices should not connect directly to the data center core. Unless ACLs are in place, every packet from the core is forwarded to the WAEs. Even if WAE devices can be placed in the core (which can be an almost impossible case to convince the network administrator), packet inspection by the WAEs leads to degraded network performance.

Interception at the access layer is not recommended either. Unless the data center is designed with a routed access layer, placing the WAE at the access layer simply does not work. WCCP and PBR work only on routed interfaces. Almost all the existing access layer designs are based on the L2 design. A benefit of the access placement is locality; it is located near the server farms. WAEs can be installed near each of the server farms for best performance. Packets do not have to traverse multiple switches from the WAEs to the server farms, unless they are outside the access switch with WAE installed. Managing WCCP or PBR interception at all access switches is challenging. There are many more access switches than aggregation switches. The task of mapping and managing WAEs to the corresponding access switches quickly outweighs the benefits.

The WAE cluster can be installed at the aggregation or access layer with VLAN extension, depending on the network administrator. Attaching the WAEs to the aggregation layer decreases the time for the packet to reach the WAEs. Traffic is coming and leaving the WAEs. If WAEs are connected to the access switches, capacity of the uplinks to the aggregation should be twice the total capacity of the WAEs.

WCCP is configured on the aggregation switches on uplinks to core switches and specific server VLAN interfaces. L2 redirection and separate in and out interfaces are configured for best performance.

Interception Interfaces and L2 Redirection

Interception needs to be configured on all uplink interfaces, and the access VLAN that needs optimization. Catalyst switches are often used in the aggregation layer. L2 redirection is preferred because it is hardware-based. L2 redirection must be configured on the WAEs, not on the Catalyst switches. The WAE and Catalyst switch negotiate which redirect and return method to use when the service group is formed. There can be many access VLANs on the aggregation switches. Redirection is configured on all VLANs that need optimization. Layer 2 switching ports, including trunk ports, are not eligible for redirection.

One suggested method of improving performance is by using the same server IP for all flows to map to the same WAE. DRE caching can be more efficient because the same information is not spread to other WAEs. Care must be taken to monitor the WAE for overloading, which can result in decreased performance with bypass traffic. The WAE does not forward overloaded traffic to another WAE for further optimization.

In the following example, TenGigabit Ethernet 4/1 is the WAN-facing interface configured to use the destination IP address (which is the server IP address) for load distribution. The VLAN 25 interface is configured to use the source address (which is also the server IP address) for load distribution. All flows to and from the server go to the same WAE for optimization, increasing DRE efficiency.

- Uplinks to core
  - interface TenGigabitEthernet4/1
  - ip address 12.20.20.2 255.255.255.0
  - ip wccp 62 redirect in
- Server VLAN
  - interface Vlan25
  - ip address 12.20.25.2 255.255.255.0
  - ip wccp 61 redirect in

Mask Assignments

GRE and L2 Redirection, page 23 explains the benefits of mask assignment. Additional mask assignment adjustment can help fine-tune the distribution of traffic to the WAE devices. Current 6500 hardware masking supports 7 bits. The default mask of 0x1741 assumes that all four octets in the IP address are significant. In some cases, the first octet is the same in a company with class A address space. The most significant masking bit is wasted with the default mask setting.

Assuming a company with 15.x.x.x/8 address space, a mask value of 0x1741 is not optimal because the first digit in the IP address is always 15. It does not make sense to use most significant bit, 0x1000 for masking.

Mask assignments are set up on the WAE. The first WAE that joins the WCCP service group determines the redirection method and masking value. Subsequent WAEs that join the group must have the same redirection and mask value setup; otherwise, they are not active participants in the WCCP group.

- CFE-AGGR-01

```
wccp tcp-promiscuous mask src-ip-mask 0x0 dst-ip-mask 0x1741
wccp tcp-promiscuous router-list-num 1 password **** l2-redirect mask-assign
```

WCCP Access Control Lists

WCCPv2 promiscuous interception forwards all packets from the router to the WAE. When the WAE is placed in the aggregation layer, every packet coming from the core to the aggregation layer is sent to the WAE. Most packets are pass-through traffic because most are coming from the data center or campus network. The WAE cluster can be overwhelmed by traffic flooding even with just pass-through packets. Additional latency is introduced by forwarding the packet to the WAE and sending it back to the switch. ACL-bounded WCCP interception at the ingress and egress interface is recommended. While the WAE supports ACLs in the device, the ACLs at the switch interfaces reduce needless packet forwarding and additional burden on the WAE.

The following is the packet control hierarchy from the switch to WAE:

1. Redirect list on the router
2. ACLs on the WAE
3. Application profile on the WAE

The following example defines two access lists: access-list 121 defining traffic going to 12.20.26.0/24 subnet, *and* access-list 120 for traffic coming from 12.20.26.0/24 subnet.

```
ip wccp 61 redirect-list 121 group-list 29 password ese
ip wccp 62 redirect-list 120 group-list 29 password ese

access-list 120 permit ip 12.20.26.0 0.0.0.255 any
access-list 121 permit ip any 12.20.26.0 0.0.0.255
```

Redirect exclude in

Speed is one of the benefits of using the aggregation layer to intercept traffic. The aggregation layer usually comprises a pair of high-end Catalyst switches. Combined with hardware L2 redirection, redirecting packets at the aggregation layer is almost at line rate. There is a caveat; the Catalyst hardware does not understand **redirect exclude in** commands. MSFC has to process packets with **redirect exclude in** on the interface. The only way to avoid using **redirect exclude in** is to not intercept on egress interfaces. There are performance penalties if **redirect exclude in** is used on the Catalyst platform.

By using only **redirect in** on all interceptions, **redirect exclude in** is avoided in the following example.

- Uplinks to core—Configure with **redirect in** on WCCP service 62

```
interface TenGigabitEthernet4/1
ip address 12.20.20.2 255.255.255.0
ip wccp 62 redirect in

interface TenGigabitEthernet4/2
ip address 12.20.41.2 255.255.255.0
ip wccp 62 redirect in
```

- Server VLAN—Configure with **redirect in** on WCCP service 61

```
interface Vlan25
ip address 12.20.25.2 255.255.255.0
ip wccp 61 redirect in
```

- WAE VLAN—Note that **redirect exclude in** is not necessary

```
interface Vlan29
description WAAS-WAE-VLAN
ip address 12.20.29.2 255.255.255.0
```

WCCP High Availability

High availability in the data center is configured with multiple WAEs, with load balancing via WCCP. The WAEs should be distributed equally among the switches: half on one switch, and half on another. Failed switches with WAEs attached affect only half the WAE cluster. Failure of the WAE does not significantly impact the WAAS service because the failed WAE is removed from the WCCP client pool. WCCP accommodates an N+1 design with up to 32 clients and 32 routers. A minimum of two WAEs are recommended. Additional WAEs can be added to scale, depending on the number of edge WAEs, traffic pattern, and number of TCP connections.

HSRP groups are configured like the WAN edge (see High Availability, page 31), allowing outbound load balancing of optimized traffic.

The standby interface is not configured in this scenario. L2 direction with masking is used and is incompatible with the standby interface. If L2 redirection with hashing is used, the standby interface should be configured.

WAAS with ACE Load Balancing

Figure 19 shows the WAAS with ACE load balancing topology.

Figure 19 WAAS with ACE Load Balancing Topology

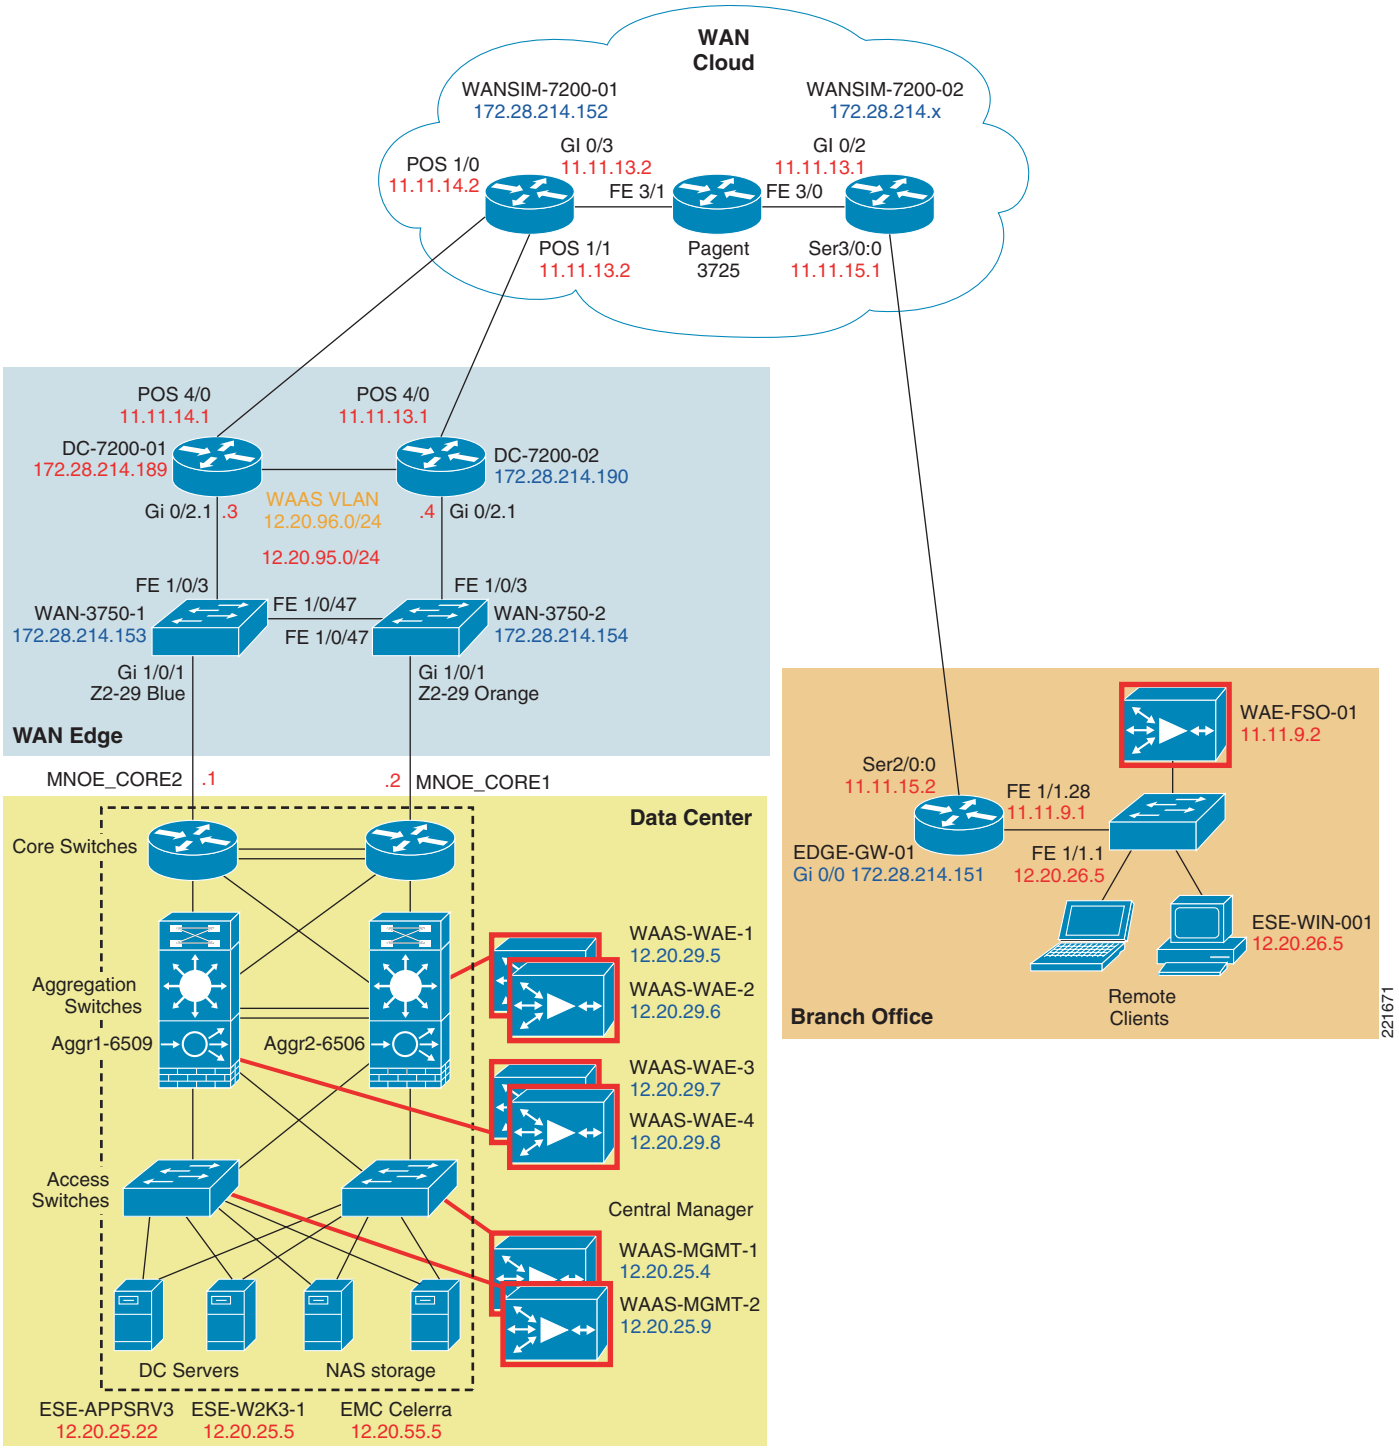

ACE operates in routed mode. The WAAS farms are load balanced by ACE. Traffic is then forwarded to the server farm. For traffic flow details, see ACE with WAAS Packet Flow, page 28.

Table 10 shows a comparison of WAAS farm configuration versus server farm configuration.

Table 10 WAAS Farm Configuration versus Server Farm Configuration

WAAS Farm Load Balancing Configuration	Server Farm Load Balancing Configuration

serverfarm host WAE-FARM
  transparent
  predictor hash address source
255.255.255.255
  probe SERVER_PING
  rserver WAAS_1
  rserver WAAS_2
  rserver WAAS_3
rserver WAAS_4
 	
class-map match-all L4_ANY_TCP
  2 match virtual-address 0.0.0.0 0.0.0.0
tcp any
 	
policy-map type loadbalance first-match
WAAS_POLICY
  class class-default
  serverfarm WAE-FARM backup SERVER_FARM
aggregate-state
policy-map multi-match L4_LB_WAAS_POLICY
  class L4_ANY_TCP
  loadbalance vip inservice
  loadbalance policy WAAS_POLICY
  loadbalance vip icmp-reply
  class L4_WEB_VIP_ADDRESS
  loadbalance policy WAAS_POLICY
 	
interface vlan 29
  description waas wae farm vlan
  ip address 12.20.29.2 255.255.255.0
  alias 12.20.29.1 255.255.255.0
  peer ip address 12.20.29.3 255.255.255.0
  no normalization
  mac-sticky enable
  no icmp-guard
  access-group input ALLOW_TRAFFIC
  service-policy input
L4_LB_VIP_WEB_POLICY
  service-policy input REMOTE-ACCESS
  no shutdown
	serverfarm host SERVER_FARM
	  probe SERVER_PING
	  rserver SERVER_1
	  inservice
	  rserver SERVER_2
	  inservice
	
	class-map match-all L4_WEB_VIP_ADDRESS
	2 match virtual-address 12.20.24.8 tcp any
	
	policy-map type loadbalance first-match
	WEB_POLICY
	  class class-default
	  serverfarm SERVER_FARM
	policy-map multi-match L4_LB_VIP_WEB_POLICY
	  class L4_WEB_VIP_ADDRESS
	  loadbalance vip inservice
	  loadbalance policy WEB_POLICY
	  loadbalance vip icmp-reply
	
	interface vlan 28
	  description server vlan
	  ip address 12.20.28.2 255.255.255.0
	  alias 12.20.28.1 255.255.255.0
	  peer ip address 12.20.28.3 255.255.255.0
	  no normalization
	  no icmp-guard
	  service-policy input L4_LB_WAAS_POLICY
	  no shutdown

Note that the WAAS farm has NAT disabled, with hash address predictor, match all policy, disabled normalization and backup server farm defined. Use Table 10 as a reference for the following topics.

WAAS Farm

WAAS farm is in the traffic path, intercepting and processing incoming and outgoing traffic. Unique features in configuring a WAAS farm include the following:

- Load balancing
- IP transparency
- Normalization
- Server farm redundancy

Mac-Sticky

A WAAS device must be in the traffic path, processing traffic bi-directionally. ACE directs incoming traffic to WAAS by using a load balancing predictor. Returning traffic from the server farms needs to follow the reverse traffic flow. This is accomplished with mac-sticky, which is reverse forwarding path with source MAC address. It redirects traffic from the server farm to the correct WAE that processed the incoming flow. Mac-sticky is defined with the **mac-sticky enable** command under the interface VLAN.

TCP Normalization

An ACE security feature provides ability to check for malformed IP and TCP packets. In this setup, ACE is used as a load balancing device. TCP normalization can potentially interfere with WAAS operations, such as sequence number changes and option 21 for auto-discovery. Cisco recommends that normalization be turned off for WAAS farms.

IP Transparency

NAT should not be used with WAAS. IP address information should be preserved for traffic going through WAAS devices. ACE should run in dispatch mode with load balancing WAAS devices, using native WAAS IP addresses. Dispatch mode can be enabled by the **transparent** keyword in the server farm definition.

Load Balancing Predictor

Various load balancing methods are described in ACE/WAAS Integration Considerations, page 26. WAAS can use any of the following methods:

- Load balancing methods
- Round-robin
- Least connection
- Hash address

Hash source address is used in this deployment for increased DRE efficiency.

Server Farm Redundancy

In the event of WAAS farm failure, the ideal scenario is to have a redundant WAAS farm. With no backup farm specified, the server farm is inaccessible if the WAAS farm is unavailable. Redundancy can achieved by sharing the WAAS farms with multiple application farms, a dedicated WAAS backup farm, or using the server farm as backup farm:

- Sharing multiple farms—WAAS farm A for server farm A, WAAS farm B for server farm B, and WAAS farm A is configured as backup for WAAS farms B, and vice versa. In this case, one WAAS farm is set up. Cost is shared among various applications.
- Dedicated WAAS backup farm—A dedicated backup WAAS farm is possible, but with increased cost.
- Using server farm as backup—The server farm itself can be specified as the backup farm. Even when the WAAS farm is down, the server farm is still available, although with no WAAS acceleration. Remember that sticky farms do not work as backup farms. There is no additional cost.

The backup server farm is defined in the server farm under the first match policy. The server farm itself is used as the backup farm in this design.

Traffic Matching

A server farm is normally accessed via the virtual IP (VIP) addresses. The load balancer advertises the VIP to the network. Incoming traffic is load balanced to the individual servers. The VIP address should be used by WAAS to the front-end server farms. Another way to match traffic is to match any traffic coming in to the ACE client VLAN. A match any policy can be used to front-end multiple servers or subnets, without defining each one in the service policy. However, a static route on the switch needs to forward relevant traffic to the client VLAN interface.

Both match any and VIP are used in this set up for maximum flexibility.

Server Farm Load Balancing

All predictors can be used to load balance the server farm. There are application-specific preferences for the load balancing algorithm. For simple HTTP-for-FTP services, simple round-robin can perform the task. Oracle, SAP, and other applications that require session persistence need more than Layer 3 load balancing. Cookies or URL-based predictors are better suited for these applications. Round-robin load balancing is the default predictor method.

ACE VLAN Policies

The WAAS VLAN is separate from the server farm and client VLAN. The WAAS VLAN is mac-sticky enabled (see Mac-Sticky, page 46). The service policy on the VLANs are also different. The client VLAN directs traffic to the WAAS farm with L4_LB_WAAS_POLICY. The WAAS VLAN sends traffic to the server farm VLAN via L4_LB_VIP_WEB_POLICY. Finally, the server farm VLAN forwards traffic back to the WAAS VLAN with L4_LB_WAAS_POLICY. The service policies define traffic flow from Step 3 to Step 9 in ACE with WAAS Packet Flow, page 28.

Central Manager Communication

WAEs communicate to the Central Manager for management traffic. Software updates, statistics collection, and management traffic are shared with the data path. Another option for the management interface is to use the secondary interface. An ACE load balanced server farm typically uses private VLANs and NAT. In the case with a WAAS farm, static routes and ACE VLAN permission must be enabled to facilitate communication to the Central Manager. Static routes are configured on the switch, forwarding WAE VLAN traffic to the ACE management interface:

```
Agg1-6509#
ip route 12.20.28.0 255.255.255.0 12.20.24.5
ip route 12.20.29.0 255.255.255.0 12.20.24.5
ACE Context
access-list ALLOW_TRAFFIC line 8 extended permit ip any any
access-list ALLOW_TRAFFIC line 9 extended permit icmp any any
interface vlan 29
  description waas wae farm vlan
    access-group input ALLOW_TRAFFIC
    …snip…
```

WAAS Standby Interface

The WAAS standby interface pings the default gateway every few seconds to verify connectivity. By default, the ACE VLAN interface is not pingable for enhanced security. With standby interface configured on the WAE, a large number of ping requests are observed. These ping requests result in performance degradation when the default gateway is not pingable. A quick fix is to make the default gateway pingable, which is accomplished by configuring the WAE VLAN interface as a management interface:

```
 interface vlan 29
  description waas wae farm vlan
service-policy input REMOTE-ACCESS
…snip…
```

Appendix A—Network Components

Network Component	Hardware Model	Software Version	Notes
Aggregation switches	CAT 6500 w/Sup 720	12.2 (18) SXF5	
ACE module	ACE10-6500-K9	Version 3.0(0)A1(4a)	Two ACE modules
Access switches	CAT 3750	12.2 (20) SE4	
Core WAE (WAN edge)	WAE-7326	4.0.1b16	Two WAEs
WAN router	7206 VXR w/NPE-G1		
Core WAE (aggregation)	WAE-511	4.0.1b16	Four WAEs
Branch router	3825 ISR		
Branch WAE	WAE-611	4.0.1b16	

Appendix B—Configurations

WAE at WAN Edge

DC-7200-01

```
!
! Last configuration change at 18:33:37 UTC Fri Jan 26 2007
! NVRAM config last updated at 18:34:20 UTC Fri Jan 26 2007
!
version 12.4
service tcp-keepalives-in
service timestamps debug datetime msec localtime show-timezone
service timestamps log datetime msec localtime show-timezone
service password-encryption
!
hostname dc-7200-01
!
boot-start-marker
boot system flash disk2:c7200-adventerprisek9-mz.124-9.T.bin
boot-end-marker
!
!card type command needed for slot 3
logging console warnings
enable secret 5 $1$KMm4$8XdMDkAnX5V53xVcMtHn2.
!
no aaa new-model
!
resource policy
```

```
!
ip wccp 61 group-list 96 password 7 045E1803
ip wccp 62 group-list 96 password 7 15171809
ip cef
!
ip domain name ese.cisco.com
ip ssh time-out 30
ip ssh logging events
ip ssh version 2
interface Loopback0
 ip address 11.11.18.1 255.255.255.255
!
interface GigabitEthernet0/1
 description connecting to dc1-7200-wan port Gi0/1 pagent
 no ip address
 duplex auto
 speed auto
 media-type rj45
 no negotiation auto
!
interface GigabitEthernet0/2
 description connecting to CAT-6500-AGG1 gi1/47
 no ip address
 duplex auto
 speed auto
 media-type rj45
 no negotiation auto
!
interface GigabitEthernet0/2.1
 description connecting  to CAT-6500-AGG1 gi1/47
 encapsulation dot1Q 1 native
 ip address 12.20.95.3 255.255.255.0
 ip wccp 61 redirect in
 ip wccp 62 redirect out
!
interface GigabitEthernet0/2.96
 description WAE-VLAN
 encapsulation dot1Q 96
 ip address 12.20.96.3 255.255.255.0
 ip wccp redirect exclude in
 standby 1 ip 12.20.96.1
 standby 1 timers 1 3
 standby 1 priority 120
 standby 1 preempt delay minimum 1
 standby 2 ip 12.20.96.2
 standby 2 timers 1 3
 standby 2 priority 110
 standby 2 preempt delay minimum 1
!
interface GigabitEthernet0/3
 ip address 172.28.214.189 255.255.255.0
 duplex auto
 speed auto
 media-type rj45
 no negotiation auto
!
!
interface POS4/0
 ip address 11.11.14.1 255.255.255.0
 no keepalive
!
router ospf 10
 log-adjacency-changes
 network 11.11.14.0 0.0.0.255 area 40
```

```
 network 12.20.95.0 0.0.0.255 area 40
 network 12.20.96.0 0.0.0.255 area 40
!
ip route 172.28.0.0 255.255.0.0 172.28.214.1
no ip http server
no ip http secure-server
!!
!
logging alarm informational
access-list 96 permit 12.20.96.5
access-list 96 permit 12.20.96.6
!
!
control-plane
!
!
gatekeeper
 shutdown
!
banner motd
Warning this is a private system.
Unauthorized access is prohibited.
Violators will be prosecuted.

!
line con 0
 stopbits 1
line aux 0
 stopbits 1
line vty 0 4
 exec-timeout 0 0
 password 7 045802150C2E
 login
 transport input telnet
!
!
end
```

DC-7200-02

```
!
! Last configuration change at 18:37:15 UTC Fri Jan 26 2007
! NVRAM config last updated at 18:38:10 UTC Fri Jan 26 2007
!
version 12.4
service tcp-keepalives-in
service timestamps debug datetime msec localtime show-timezone
service timestamps log datetime msec localtime show-timezone
service password-encryption
!
hostname dc-7200-02
!
boot-start-marker
boot system flash disk2:c7200-adventerprisek9-mz.124-9.T.bin
boot-end-marker
!
!card type command needed for slot 2
!card type command needed for slot 3
logging console warnings
enable secret 5 $1$FVS6$.eil6U/P/fsTlLbnngQjk1
!
no aaa new-model
```

```
!
resource policy
!
ip wccp 61 group-list 96 password 7 104B1A1C
ip wccp 62 group-list 96 password 7 15171809
ip cef
!
!
interface GigabitEthernet0/1
 description connecting to dc1-7200-wan port Gi0/2 pagent device
 no ip address
 duplex auto
 speed auto
 media-type rj45
 no negotiation auto
!
interface GigabitEthernet0/2
 no ip address
 duplex auto
 speed auto
 media-type rj45
 no negotiation auto
!
interface GigabitEthernet0/2.1
 encapsulation dot1Q 1 native
 ip address 12.20.95.4 255.255.255.0
 ip wccp 61 redirect in
 ip wccp 62 redirect out
!
interface GigabitEthernet0/2.96
 encapsulation dot1Q 96
 ip address 12.20.96.4 255.255.255.0
 ip wccp redirect exclude in
 standby 1 ip 12.20.96.1
 standby 1 timers 1 3
 standby 1 priority 110
 standby 1 preempt delay minimum 1
 standby 2 ip 12.20.96.2
 standby 2 timers 1 3
 standby 2 priority 120
 standby 2 preempt delay minimum 1
!
interface GigabitEthernet0/3
 ip address 172.28.214.190 255.255.255.0
 duplex auto
 speed auto
 media-type rj45
 no negotiation auto
!
interface POS4/0
 ip address 11.11.13.1 255.255.255.0
 no keepalive
!
router ospf 10
 log-adjacency-changes
 network 11.11.14.0 0.0.0.255 area 40
 network 12.20.95.0 0.0.0.255 area 40
 network 12.20.96.0 0.0.0.255 area 40
!
ip route 172.28.0.0 255.255.0.0 172.28.214.1
no ip http server
no ip http secure-server
!
logging alarm informational
```

```
access-list 96 permit 12.20.96.5
access-list 96 permit 12.20.96.6
!
!control-plane
!
gatekeeper
 shutdown
!
banner motd
Warning this is a private system.
Unauthorized access is prohibited.
Violators will be prosecuted.

!
line con 0
 stopbits 1
line aux 0
 stopbits 1
line vty 0 4
 exec-timeout 0 0
 password 7 045802150C2E
 login
!
!
End
```

CORE-FE1

```
device mode application-accelerator
!
hostname core-fe1
!
ip domain-name eselab-sj.com
!
primary-interface Standby 1
!
interface Standby 1
 description standby interface group
 ip address 12.20.96.5 255.255.255.0
 exit
!
interface GigabitEthernet 1/0
 description WAAS interface
 standby 1
 exit
interface GigabitEthernet 2/0
 description WAAS interface
 standby 1 priority 105
 exit
!
ip default-gateway 12.20.96.1
!
no auto-register enable
!
! ip path-mtu-discovery is disabled in WAAS by default
!
ip name-server 12.20.25.22
!
ntp server 12.20.25.4
!
wccp router-list 1 12.20.96.3 12.20.96.4
```

```
wccp tcp-promiscuous router-list-num 1 password ****
wccp version 2
! wccp slow-start is disabled in WAAS by default
!
username admin password 1 bVmDmMMmZAPjY
username admin privilege 15
username admin print-admin-password 1 29D5C31BFF3D8D25AAD3B435B51404EE 7D891AB40
2CAF2E89CCDD33ED54333AC
!
windows-domain netbios-name "CORE-FE1"
!
authentication login local enable primary
authentication configuration local enable primary
!
sshd allow-non-admin-users
sshd enable
!
!
central-manager address 12.20.25.4
cms enable
!
policy-engine application
   name WAFS
   name SQL
   name File-System
   name Systems-Management
...etc...
```

CORE-FE2

```
!
device mode application-accelerator
!
hostname core-fe2
!
ip domain-name eselab-sj.com
!
primary-interface Standby 1
!
interface Standby 1
 description standby interface group
 ip address 12.20.96.6 255.255.255.0
 exit
!
interface GigabitEthernet 1/0
 description WAAS interface
 standby 1
 exit
interface GigabitEthernet 2/0
 description WAAS interface
 standby 1 priority 105
 exit
!
ip default-gateway 12.20.96.2
!
no auto-register enable
!
! ip path-mtu-discovery is disabled in WAAS by default
!
ip name-server 12.20.25.22
!
ntp server 12.20.25.4
```

```
!
wccp router-list 1 12.20.96.3 12.20.96.4
wccp tcp-promiscuous router-list-num 1 password ****
wccp version 2
! wccp slow-start is disabled in WAAS by default
!
username admin password 1 bVmDmMMmZAPjY
username admin privilege 15
username admin print-admin-password 1 29D5C31BFF3D8D25AAD3B435B51404EE 7D891AB40
2CAF2E89CCDD33ED54333AC
!
windows-domain netbios-name "CORE-FE2"
!
authentication login local enable primary
authentication configuration local enable primary
!
sshd allow-non-admin-users
sshd enable
!
central-manager address 12.20.25.4
cms enable
!
policy-engine application
   name WAFS
   name SQL
   name File-System
... etc ...
```

EDGE-GW-01

```
!
Current configuration : 4623 bytes
!
version 12.4
service tcp-keepalives-in
service timestamps debug datetime msec localtime show-timezone
service timestamps log datetime msec localtime show-timezone
service password-encryption
!
hostname edge-gw-01
!
boot-start-marker
boot-end-marker
!
card type t1 2
logging buffered 51200 warnings
logging console warnings
enable secret 5 $1$z3gt$0OIoqNmp0oDyFj3b1C0Ff.
!
no aaa new-model
no network-clock-participate slot 1
ip wccp 61 redirect-list 9 password 7 00010003
ip wccp 62 redirect-list 9 password 7 13000417
ip cef
!
ip domain name ese.cisco.com
!
voice-card 0
 no dspfarm
!
!
crypto pki trustpoint TP-self-signed-775691316
```

```
 enrollment selfsigned
 subject-name cn=IOS-Self-Signed-Certificate-775691316
 revocation-check none
 rsakeypair TP-self-signed-775691316
!
!
crypto pki certificate chain TP-self-signed-775691316
 certificate self-signed 01
  3082024F 308201B8 A0030201 02020101 300D0609 2A864886 F70D0101 04050030
  30312E30 2C060355 04031325 494F532D 53656C66 2D536967 6E65642D 43657274
  69666963 6174652D 37373536 39313331 36301E17 0D303631 30313832 33313332
  335A170D 32303031 30313030 30303030 5A303031 2E302C06 03550403 1325494F
  532D5365 6C662D53 69676E65 642D4365 72746966 69636174 652D3737 35363931
  33313630 819F300D 06092A86 4886F70D 01010105 0003818D 00308189 02818100
  E4015874 59DCC15E 7E7DFBD3 871A407D A8C7068E A3BA1961 D7800717 AAF99F0F
  DC2DB065 9E425483 EE457ADB AB5B2132 4CB6301E 4F370A06 935CB6E7 9FD77214
  560CF9E1 1D431CF2 D355D2E8 1282E618 0172D080 B775CFF0 DCB5D5AA 10FA7488
  2B7D2888 E90C527A B7E562E9 2D807A21 317B20D0 E5D45DDA BFFE4AA3 DCE908E5
  02030100 01A37930 77300F06 03551D13 0101FF04 05300301 01FF3024 0603551D
  11041D30 1B821965 6467652D 67772D30 312E796F 7572646F 6D61696E 2E636F6D
  301F0603 551D2304 18301680 14504156 99A8C336 8C030E38 DD0D7F58 A3C3E5FC
  1A301D06 03551D0E 04160414 50415699 A8C3368C 030E38DD 0D7F58A3 C3E5FC1A
  300D0609 2A864886 F70D0101 04050003 81810046 432A489D 11E8AD58 0122A07C
  778C0A4A 18C7E165 4DD79106 E8A2FAB0 6741009D 732AAF77 09404C16 C543679D
  4FEBF8D0 94B9CF72 D8198874 8DBEFEC0 14A16D09 C1A097F3 B8A162DF FF427C8C
  3BC2AA7E 1A0F7523 C1AD094B AE5A173D 5AFE7F8D 235A2554 DF358EEB 365D28A4
  1D7A6D53 24B62906 E9905D03 BEE8B54F 5DADA5
  quit
username cisco privilege 15 secret 5 $1$uIDa$REVTO05wgTleEIGyRwSIW1
!
!
controller T1 2/0
 framing esf
 linecode b8zs
 cablelength short 110
 channel-group 0 timeslots 1-24 speed 64
!
controller T1 2/1
 framing esf
 linecode b8zs
 cablelength long 0db
 channel-group 0 timeslots 1-24
!
interface Loopback0
 description loopback interface
 ip address 11.11.18.3 255.255.255.0
!
interface GigabitEthernet0/0
 description $ETH-LAN$$ETH-SW-LAUNCH$$INTF-INFO-GE 0/0$
 ip address 172.28.214.151 255.255.255.0
 duplex auto
 speed auto
 media-type rj45
 no keepalive
!
interface GigabitEthernet0/1
 no ip address
 duplex auto
 speed auto
 media-type rj45
 no keepalive
!
interface GigabitEthernet0/1.1
 encapsulation dot1Q 1 native
```

```
 ip address 12.20.26.1 255.255.255.0
 ip wccp 61 redirect in
 ip wccp 62 redirect out
!
interface GigabitEthernet0/1.28
 encapsulation dot1Q 28
 ip address 11.11.9.1 255.255.255.0
 ip wccp redirect exclude in
!
interface Serial0/1/0
 no ip address
 shutdown
 clock rate 2000000
!
interface Serial0/1/1
 no ip address
 shutdown
 clock rate 2000000
!
interface FastEthernet1/0
 no ip address
 shutdown
 duplex auto
 speed auto
!
interface FastEthernet1/1
 no ip address
 duplex full
 speed 100
!
interface FastEthernet1/1.1
 encapsulation dot1Q 1 native
 ip wccp 61 redirect in
 ip wccp 62 redirect out
 shutdown
!
interface FastEthernet1/1.28
 encapsulation dot1Q 28
 ip wccp redirect exclude in
 shutdown
!
interface Serial2/0:0
 description WAN connection
 ip address 11.11.15.2 255.255.255.0
!
interface Serial2/1:0
 no ip address
!
router eigrp 1
 network 11.11.9.0 0.0.0.255
 network 11.11.15.0 0.0.0.255
 network 12.20.26.0 0.0.0.255
 no auto-summary
!
!
!
ip http server
ip http authentication local
ip http secure-server
ip http timeout-policy idle 5 life 86400 requests 10000
!
access-list 9 permit 11.11.9.2
!
control-plane
```

```
!
banner motd ^C
Warning this is a private system.
Unauthorized access is prohibited.
Violators will be proscecuted.
^C
!
line con 0
 login local
 stopbits 1
line aux 0
 stopbits 1
line vty 0 4
 exec-timeout 0 0
 privilege level 15
 password 7 094F471A1A0A
 login
 transport input telnet ssh
line vty 5 15
 exec-timeout 0 0
 privilege level 15
 password 7 094F471A1A0A
 login
 transport input telnet ssh
!
scheduler allocate 20000 1000
!
end
```

WAE-FSO-01

```
!
device mode application-accelerator
!
!
hostname wae-fso-01
!
ip domain-name eselab-sj.com
!
primary-interface GigabitEthernet 2/0
!
interface GigabitEthernet 1/0
 description public management interface
 ip address 172.28.214.192 255.255.255.0
 exit
interface GigabitEthernet 2/0
 description connection to edge-gw-01
 ip address 11.11.9.2 255.255.255.0
 exit
!
ip default-gateway 11.11.9.1
!
no auto-register enable
!
! ip path-mtu-discovery is disabled in WAAS by default
!
ip name-server 12.20.25.22
!
ntp server 12.20.25.4
!
wccp router-list 1 11.11.9.1
wccp tcp-promiscuous router-list-num 1 password ****
```

```
wccp version 2
! wccp slow-start is disabled in WAAS by default
!
username admin password 1 bVmDmMMmZAPjY
username admin privilege 15
username admin print-admin-password 1 29D5C31BFF3D8D25AAD3B435B51404EE 7D891AB40
2CAF2E89CCDD33ED54333AC
!
snmp-server community CLI_TRIGGER_4358 rw
!
windows-domain netbios-name "WAE-FSO-01"
!
authentication login local enable primary
authentication configuration local enable primary
!
sshd allow-non-admin-users
sshd enable
!
!
central-manager address 12.20.25.4
cms enable
!
policy-engine application
   name WAFS
   name SQL
   name File-System
   name Systems-Management
... etc ...
```

WAE at Aggregation Layer

This section contains WAAS-related output only.

AGGR1

```
ip wccp 61 redirect-list 121 group-list 29 password ese
ip wccp 62 redirect-list 120 group-list 29 password ese

vlan 25
 name WAAS-server-vlan
!
vlan 29
 name WAAS-WAE-vlan
!
vlan 55
 name NAS-vlan
!
interface TenGigabitEthernet4/1
 description to Core 1 - shut and copy to int vlan 18 if problems
 ip address 12.20.20.2 255.255.255.0
 no ip redirects
 no ip proxy-arp
 ip wccp 62 redirect in
 ip pim sparse-dense-mode
 ip ospf message-digest-key 1 md5 C1sC0!
 ip ospf network point-to-point
 ip ospf hello-interval 2
 ip ospf dead-interval 6
 logging event link-status
```

```
 spanning-tree guard loop
!
interface TenGigabitEthernet4/2
 description to Core2
 ip address 12.20.41.2 255.255.255.0
 no ip redirects
 no ip proxy-arp
 ip wccp 62 redirect in
 ip pim sparse-dense-mode
 ip ospf message-digest-key 1 md5 C1sC0!
 ip ospf network point-to-point
 ip ospf hello-interval 2
 ip ospf dead-interval 6
 logging event link-status
 load-interval 30
!
interface GigabitEthernet6/21
 description CFE-AGGR-03
 switchport
 switchport access vlan 29
 switchport mode access
 switchport port-security maximum 3
 no ip address
 no cdp enable
 spanning-tree portfast
 spanning-tree bpduguard enable
!
interface GigabitEthernet6/22
 description CFE-AGGR-04
 switchport
 switchport access vlan 29
 switchport mode access
 switchport port-security maximum 3
 no ip address
 no cdp enable
 spanning-tree portfast
 spanning-tree bpduguard enable
!
interface Vlan55
 description NAS-vlan
 ip address 12.20.55.2 255.255.255.0
 ip wccp 61 redirect in
 ip route-cache flow
 standby 1 ip 12.20.55.1
 standby 1 timers 1 3
 standby 1 priority 125
 standby 1 preempt delay minimum 1
 standby 1 name NAS-vlan
!
interface Vlan25
 description WAAS
 ip address 12.20.25.2 255.255.255.0
 ip wccp 61 redirect in
 ip route-cache flow
 standby 1 ip 12.20.25.1
 standby 1 timers 1 3
 standby 1 priority 110
 standby 1 preempt delay minimum 1
!
interface Vlan29
 description WAAS-WAE-vlan
 ip address 12.20.29.2 255.255.255.0
 ip route-cache flow
 standby 1 ip 12.20.29.1
```

```
 standby 1 timers 1 3
 standby 1 priority 110
 standby 1 preempt delay minimum 1
 standby 2 ip 12.20.29.4
 standby 2 timers 1 3
 standby 2 priority 120
 standby 2 preempt delay minimum 1
!
access-list 29 permit 12.20.29.5
access-list 29 permit 12.20.29.7
access-list 29 permit 12.20.29.6
access-list 29 permit 12.20.29.8
!
access-list 120 permit tcp 12.20.26.0 0.0.0.255 any
access-list 121 permit tcp any 12.20.26.0 0.0.0.255
!
```

AGGR2

```
ip wccp 61 redirect-list 121 group-list 29 password ese
ip wccp 62 redirect-list 120 group-list 29 password ese
!
vlan 25
 name WAAS-server-vlan
!
vlan 29
 name WAAS-WAE-vlan
!
!
vlan 55
 name NAS-vlan
!
interface TenGigabitEthernet4/1
 description to core 1
 ip address 12.20.30.2 255.255.255.0
 no ip proxy-arp
 ip wccp 62 redirect in
 ip pim sparse-dense-mode
 ip ospf message-digest-key 1 md5 C1sC0!
 ip ospf network point-to-point
 ip ospf hello-interval 2
 ip ospf dead-interval 6
 logging event link-status
!
interface TenGigabitEthernet4/2
 description to core 2
 ip address 12.20.50.2 255.255.255.0
 no ip redirects
 no ip proxy-arp
 ip wccp 62 redirect in
 ip pim sparse-dense-mode
 ip ospf message-digest-key 1 md5 C1sC0!
 ip ospf network point-to-point
 ip ospf hello-interval 2
 ip ospf dead-interval 6
 logging event link-status
!
interface GigabitEthernet5/19
 description CFE-AGGR-01
 switchport
 switchport access vlan 29
```

```
 switchport mode access
 switchport port-security maximum 3
 no ip address
 no cdp enable
 spanning-tree portfast
 spanning-tree bpduguard enable
!
interface GigabitEthernet5/20
 description CFE-AGGR-02
 switchport
 switchport access vlan 29
 switchport mode access
 switchport port-security maximum 3
 no ip address
 no cdp enable
 spanning-tree portfast
 spanning-tree bpduguard enable
!
access-list 29 permit 12.20.29.5
access-list 29 permit 12.20.29.7
access-list 29 permit 12.20.29.6
access-list 29 permit 12.20.29.8
!
access-list 120 permit tcp 12.20.26.0 0.0.0.255 any
access-list 121 permit tcp any 12.20.26.0 0.0.0.255
```

CFE-AGGR-01

```
! WAAS version 4.0.1 (build b16 Aug 28 2006)
!
device mode application-accelerator
!
hostname CFE-AGGR-01
!
ip domain-name eselab-sj.com
!
primary-interface GigabitEthernet 1/0
!
interface GigabitEthernet 1/0
 description connect to aggr switch
 ip address 12.20.29.5 255.255.255.0
 exit
interface GigabitEthernet 2/0
 ip address 172.28.214.151 255.255.255.0
 shutdown
 exit
!
ip default-gateway 12.20.29.1
!
no auto-register enable
!
! ip path-mtu-discovery is disabled in WAAS by default
!
ip name-server 12.20.25.22
!
!ntp server 12.20.25.4
!
wccp router-list 1 12.20.29.1 12.20.29.2 12.20.29.3
wccp tcp-promiscuous mask src-ip-mask 0x0 dst-ip-mask 0x1741
wccp tcp-promiscuous router-list-num 1 password **** l2-redirect mask-assign
wccp version 2
```

```
! wccp slow-start is disabled in WAAS by default
!
!
!
username admin password 1 bVmDmMMmZAPjY
username admin privilege 15
username admin print-admin-password 1 29D5C31BFF3D8D25AAD3B435B51404EE 7D891AB40
2CAF2E89CCDD33ED54333AC
!
authentication login local enable secondary
authentication configuration local enable primary
!
sshd allow-non-admin-users
sshd enable
!
!
central-manager address 12.20.25.4
cms enable
!
policy-engine application
   name SQL
   name File-System

... etc ...
```

CFE-AGGR-02

Configuration is same as CFE-AGGR-01 except for the following:

```
hostname CFE-AGGR-02
!
interface GigabitEthernet 1/0
 description connect to aggr switch
 ip address 12.20.29.6 255.255.255.0
 exit
interface GigabitEthernet 2/0
 shutdown
 exit
```

CFE-AGGR-03

```
! WAAS version 4.0.1 (build b16 Aug 28 2006)
!
device mode application-accelerator
!
!
hostname CFE-AGGR-03
!
!
!
!
ip domain-name eselab-sj.com
!
!
!
primary-interface GigabitEthernet 2/0
!
!
```

```
!
interface GigabitEthernet 1/0
 shutdown
 exit
interface GigabitEthernet 2/0
 ip address 12.20.29.7 255.255.255.0
 exit
!
ip default-gateway 12.20.29.4
!
no auto-register enable
!
! ip path-mtu-discovery is disabled in WAAS by default
!
ip name-server 12.20.25.22
!
!
!
ntp server 12.20.25.4
!
!
wccp router-list 1 12.20.29.1 12.20.29.2 12.20.29.3
wccp tcp-promiscuous mask src-ip-mask 0x0 dst-ip-mask 0x1741
wccp tcp-promiscuous router-list-num 1 password **** l2-redirect mask-assign
wccp version 2
! wccp slow-start is disabled in WAAS by default
!
!
!
username admin password 1 bVmDmMMmZAPjY
username admin privilege 15
username admin print-admin-password 1 29D5C31BFF3D8D25AAD3B435B51404EE
7D891AB402CAF2E89CCDD33ED54333AC
!
!
!
!
authentication login local enable primary
authentication configuration local enable primary
!
!
!
sshd allow-non-admin-users
sshd enable
!
!
central-manager address 12.20.25.4
cms enable
!
!
!
!
!
policy-engine application
   name SQL
   name File-System


... etc ...
```

CEF-AGGR-04

Configuration is the same as CFE-AGGR-03 except the following:

```
hostname CFE-AGGR-04
!
interface GigabitEthernet 1/0
 shutdown
 exit
interface GigabitEthernet 2/0
 description connect to aggr switch
 ip address 12.20.29.8 255.255.255.0
 exit
```

WAAS with ACE Load Balancing

CEF-AGGR-01 to 04

This is the same as the WAE at aggregation except that there are no WCCP statements. The following lines are removed:

```
wccp router-list 1 12.20.29.1 12.20.29.2 12.20.29.3
wccp tcp-promiscuous mask src-ip-mask 0x0 dst-ip-mask 0x1741
wccp tcp-promiscuous router-list-num 1 password **** l2-redirect mask-assign
wccp version 2
! wccp slow-start is disabled in WAAS by default
```

AGGR1 and AGGR2

This is the same as WAE at aggregation with vlan 24,28,29,30 deactivated. The following statements are appended to enable ACE module to take over these VLANs.

```
svclc multiple-vlan-interfaces
svclc module 1 vlan-group 1,2
svclc vlan-group 1  24,28-30,51,104,204,205,207-209,260,261,300-310,952-954  ← only vlan
28-30 is used in this config
svclc vlan-group 1  956-958,962-964,966-968
svclc vlan-group 2  82,105,107-109,160,161
```

ACE Module

This is the standard ACE configuration plus the following under the Admin context:

```
context WAAS-CONTEXT
  allocate-interface vlan 24
allocate-interface vlan 28-30

ft group 21
  peer 1
  priority 120
  peer priority 110
  associate-context WAAS-CONTEXT
  inservice

ACE-1/WAAS-CONTEXT# sh run← WAAS Context config
```

```
access-list ALLOW_TRAFFIC line 8 extended permit ip any any
access-list ALLOW_TRAFFIC line 9 extended permit icmp any any

probe icmp SERVER_PING
  interval 10
  passdetect interval 2
  passdetect count 2

rserver host SERVER_1
  description server 1
  ip address 12.20.28.5
  inservice
rserver host SERVER_2
  description server 2
  ip address 12.20.28.6
  inservice
rserver host WAAS_1
  description waas 1
  ip address 12.20.29.5
  inservice
rserver host WAAS_2
  description waas 2
  ip address 12.20.29.6
  inservice
rserver host WAAS_3
  description waas 3
  ip address 12.20.29.7
  inservice
rserver host WAAS_4
  description waas 4
  ip address 12.20.29.8
  inservice

serverfarm host SERVER_FARM
  probe SERVER_PING
  rserver SERVER_1
    inservice
  rserver SERVER_2
    inservice
serverfarm host WAE-FARM
  transparent
  predictor hash address source 255.255.255.255
  probe SERVER_PING
  rserver WAAS_1
  rserver WAAS_2
  rserver WAAS_3
  rserver WAAS_4

class-map match-all L4_ANY_TCP
  2 match virtual-address 0.0.0.0 0.0.0.0 tcp any
class-map match-all L4_WEB_VIP_ADDRESS
  2 match virtual-address 12.20.24.8 tcp any
class-map type management match-any REMOTE-MGMT
  2 match protocol ssh any
  3 match protocol telnet any
  4 match protocol icmp any
  5 match protocol http any
  6 match protocol https any

policy-map type management first-match REMOTE-ACCESS
  class REMOTE-MGMT
    permit
policy-map type loadbalance first-match WAAS_POLICY
  class class-default
```

```
    serverfarm WAE-FARM backup SERVER_FARM aggregate-state
policy-map type loadbalance first-match WEB_POLICY
  class class-default
    serverfarm SERVER_FARM
policy-map multi-match L4_LB_VIP_WEB_POLICY
  class L4_WEB_VIP_ADDRESS
    loadbalance vip inservice
    loadbalance policy WEB_POLICY
    loadbalance vip icmp-reply
policy-map multi-match L4_LB_WAAS_POLICY
  class L4_ANY_TCP
    loadbalance vip inservice
    loadbalance policy WAAS_POLICY
    loadbalance vip icmp-reply
  class L4_WEB_VIP_ADDRESS
    loadbalance policy WAAS_POLICY

interface vlan 24
  description ACE client vlan
  ip address 12.20.24.6 255.255.255.0
  alias 12.20.24.5 255.255.255.0
  peer ip address 12.20.24.7 255.255.255.0
  no normalization
  no icmp-guard
  access-group input ALLOW_TRAFFIC
  service-policy input REMOTE-ACCESS
  service-policy input L4_LB_WAAS_POLICY
  no shutdown
interface vlan 28
  description server vlan
  ip address 12.20.28.2 255.255.255.0
  alias 12.20.28.1 255.255.255.0
  peer ip address 12.20.28.3 255.255.255.0
  no normalization
  no icmp-guard
  service-policy input L4_LB_WAAS_POLICY
  no shutdown
interface vlan 29
  description waas wae farm vlan
  ip address 12.20.29.2 255.255.255.0
  alias 12.20.29.1 255.255.255.0
  peer ip address 12.20.29.3 255.255.255.0
  no normalization
  mac-sticky enable
  no icmp-guard
  access-group input ALLOW_TRAFFIC
  service-policy input L4_LB_VIP_WEB_POLICY
  service-policy input REMOTE-ACCESS
  no shutdown

ip route 0.0.0.0 0.0.0.0 12.20.24.1
```

Appendix C—References

- *Cisco Data Center Infrastructure 2.5*—http://www.cisco.com/en/US/docs/solutions/Enterprise/Data_Center/DC_Infra2_5/DCI_SRND_2_5_book.html

- *Cisco WAAS Configuration Guide*—
http://www.cisco.com/en/US/docs/app_ntwk_services/waas/waas/v4019/quick/guide/waasqcg.html

Note The configuration in this chapter deploys WAAS with ACE load balancing in the aggregation layer. Deployment depends on the location of the hosts. Only one deployment method should be used in any one ACE context. However, multiple ACE contexts can be created for different deployments options within the same chassis.